Callisto: High Pressure Loading

Final Report

February 11, 2024

Callisto team

Abstract

The Callisto project was dedicated towards searching for signs of low energy nuclear reactions (LENRs) occurring within deuterium-loaded palladium.

Our investigation was quite broad, both in attempting to induce LENR, and for trying to detect it:

> **Loading:** We have used very high pressure D_2 gas (up to 500,000 psi), as our primary way to load deuterium into palladium. In this way, we have reached stable, 100%, loading levels; i.e., stoichiometric PdD.
>
> **Stimulation:** We employed a wide range of methods, including heating loaded PdD up to its melting temperature, irradiating it with ultrashort-pulse laser beams, and applying intense electric fields, in an attempt to stimulate nuclear reactions within the PdD.
>
> **Material:** Many different forms of PdD have been investigated, ranging from single crystals and polycrystalline slugs, through wires, foils, and lithographically formed chips, to nanoparticles and large arrays of patterned nanostructures.
>
> **Detection:** In an effort to perform ever more sensitive searches for LENR energy, we have employed best-in-class commercial calorimeters, as well as a variety of high-temperature thermal sensors. This drive culminated in the development of detectors, using either physical or magnetic archaeology, capable of detecting the energy from even single LENR events.

As of late 2023, Callisto has not found any evidence of low energy nuclear reactions occurring in PdD.

We have concluded that such LENRs almost certainly do not exist.

Table of Contents

		Page
1	Background	1
2	Introduction to the Campaigns	13
3	[Omitted]	17
4	Loading	23
5	Thermal Ignition	73
6	Staring	90
7	Chips	124
8	Ultra-Short Pulse	165
9	Terahertz	196
10	High Pressure Electrolysis	205
11	Replications	219
12	Calorimetry	254
13	Physical Archaeology	286
14	Magnetic Archaeology	305
15	Seismology	339
16	Nuclear Detection	362
17	Conclusions	369
18	Recommendations	378
19	Appreciation	380

Background

There is a long history of efforts to identify and study LENRs, dating back to the pioneering 1989 work by Fleischmann and Pons. They claimed to observe nuclear fusion when electrolytically loading palladium (Pd) with deuterium (D) during electrolysis of heavy-water-based electrolytes. Because this work promised to produce fusion energy at near-ambient conditions, using only conventional electrolysis, it galvanized the scientific community into trying to replicate their results. However, most of these efforts were unsuccessful. Because of these failures, the lack of any plausible theoretical explanation as to how nuclear reactions could occur in such conditions, and (frankly) the fact that the claims came from electrochemists rather than physicists, the Fleischmann-Pons results were generally treated with suspicion. Hence, for about 30 years, LENR research was essentially blacklisted by the mainstream scientific community, with very little access to publication or funding. Nevertheless, research has continued, pursued by a fairly limited community of researchers. This community has continued to perform experiments, publish in a handful of journals and conference proceedings, and claim evidence of LENR, albeit without the clarity or consistency to convince the broader scientific community to reengage.

Several years ago, starting in 2015, Callisto was formed to investigate LENR. The reason was, of course, the enormous significance of LENR, if the effect exists, as an energy source. Our goal was to determine the scientific reality of LENR. If it could be demonstrated to exist, then follow-on efforts would determine how to turn scientific feasibility into engineering and economic reality. If, on the other hand, LENR could not be scientifically demonstrated, then this would also be an important, if far less desired, thing to know.

Unfortunately, as stated above, Callisto has not found any evidence of LENR. This report will lay out what our search has entailed, and why we believe it has established a strong case against LENR. Of course, it's virtually impossible to prove a negative, so it is possible that LENRs do indeed exist, and our effort has simply not been complete, or clever, enough to have found them. Accordingly, in addition to presenting the breadth of our search and the large portions of parameter space within which LENRs clearly do not occur, we'll also discuss some of the limitations of our search, and avenues which we do not believe have yet been closed off.

To begin, let's be explicit about what forms of LENR[1] we have been looking for, and which ones we have not been. We have been looking for deuterium-deuterium nuclear reactions occurring within deuterium loaded palladium or some of its alloys. We have not been looking for protium-based reactions, e.g., we have not pursued reports of LENR in nickel-protium systems. Furthermore, our emphasis has been upon a particular LENR hypothesis, namely a D-D reaction

[1] In this report, we use two different terms for the nuclear events we are looking for. LENR, standing for "low energy nuclear reaction" is often encountered in the outside literature. Internally, we often use AHE, "anomalous heat effect". While they are not formally identical, we'll use them interchangeably in this report.

in which the products are simply a He^4 nucleus and heat, rather than the high energy particles or radiation produced during conventional D-D reactions.

The most fundamental requirement for Pd-D LENRs, is that deuterium must be loaded into the palladium lattice. There are two major ways to do this. The one that Fleischmann and Pons used is electrolysis, using an electric potential to split deuterium from heavy water molecules, and force it into a palladium cathode. Because the technology needed to do this is inexpensive and readily available, this is the method that has been used in the vast majority of LENR experiments, both the purportedly successful ones and the unsuccessful ones. The Callisto electrochemical loading effort within our envelope was established to explore this avenue, allowing it to draw upon decades of preceding LENR research.

Callisto was tasked to explore an alternative route to loading deuterium into palladium, primarily by using very high pressure D_2 gas. Conceptually, this approach to loading is extremely simple (just apply pressurized D_2 gas to metallic Pd), even more so than electrolysis. In practice, however, it is not! The difficulty is the magnitude of the pressures required. Reaching loadings of about 65-70% (loading being defined as the atomic D-to-Pd ratio) is easy, just requiring a few bars of D_2 pressure. But a common belief in the LENR field, is that much higher loadings, e.g., 90% or more, are required. Unfortunately, the pressure cost to increasing loading to these levels is not linear, but exponential; reaching 90+% loadings can only be done with pressures in the 5,000 bar range[2]. It's because of this huge pressure requirement, that virtually no LENR researchers (with the notable exception of Baranowski's group in the very early days) has taken this route to achieving high loading.

Nevertheless, if one can actually reach and sustain such pressures, this approach does offer some major advantages over electrolytic loading. One, is that pressure loaded Pd-D is in equilibrium; the loading can be maintained indefinitely, and will last as long as the pressure does. Another major advantage is that gas loading is much chemically "cleaner" than electrolytic loading. Electrolysis is well known for gradually extracting even trace amounts of impurities from nominally pure electrolytes, causing complex (and virtually irreproducible) coatings on the surface where loading is taking place. Gas systems avoid these problems; as long as one can apply >5,000 bar pressures, they offer a clean, stable, and reproducible way to achieve highly loaded Pd-D. This is why, despite the technological challenges, the Callisto project was focused on high pressure loading.

A major question for LENR research is how it is possible for low energy nuclear reactions to occur, both how two deuterons can cross the coulomb barrier in order to react, as well as why the reaction energy appears only as lattice heating rather than as energetic nuclear particles. Both of these questions have, of course, been the focus of intense interest since the dawn of LENR research, and a large number of proposals have been posited to address one or both of the issues. However, none

[2] Throughout this report, we report pressure using a variety of different units. Using English units, "psi" means pounds per square inch, while "ksi" means 1000 psi; so 100,000 psi is 100 ksi. Using Metric units, "Pa" means Newtons per square meter, while "GPa" means 1 billion Pa. And one "bar" (nearly one atmosphere of pressure) is 100,000 Pa, while "kbar" means 1000 bars. So, 1 GPa is 10 kbar. To convert between English and Metric units; 145 ksi is 10 kbar, which is 1 GPa. Obviously, this can seem confusing, but our writeups, figures, and data have been generated and reported in this gamut of units.

of these explanations have been compelling enough to convince either the bulk of LENR community or, needless to say, virtually any of the wider population of physicists.

For our project, and for LENR research in general, this has meant that there is no real theoretical guidance as to how to generate LENRs. The field is purely experimental, with only historical lore as to what conditions have previously produced reports of LENRs, as well as a number of seemingly reasonable suppositions to guide an ongoing search for LENRs.

One of these, is that palladium does not actively participate in the nuclear reactions; they occur between two deuterons, not between a deuteron and a palladium nucleus. {It's hard enough to understand how the D-D coulomb barrier can be crossed; the D-Pd one is exponentially more difficult.} In this scenario, the role of the palladium is simply as a host medium, in which the deuterons exist and can react. The justification for focusing on palladium (as opposed to any other host material) is twofold. First, palladium is where the effect was first reported, and is the host where most later reports of LENR have come from; empirically there is little justification to look elsewhere. Secondly, palladium offers many advantages as a host; hydrogen atoms can be easily loaded (at least to ~ 70% concentration), while remaining mobile enough to approach each other (i.e., they are not chemically held apart).

Generating LENRs

Callisto's search for LENR was based on achieving a triad of conditions; the right loading, the right material, and the right stimulation. What "right" means, of course, is not known!

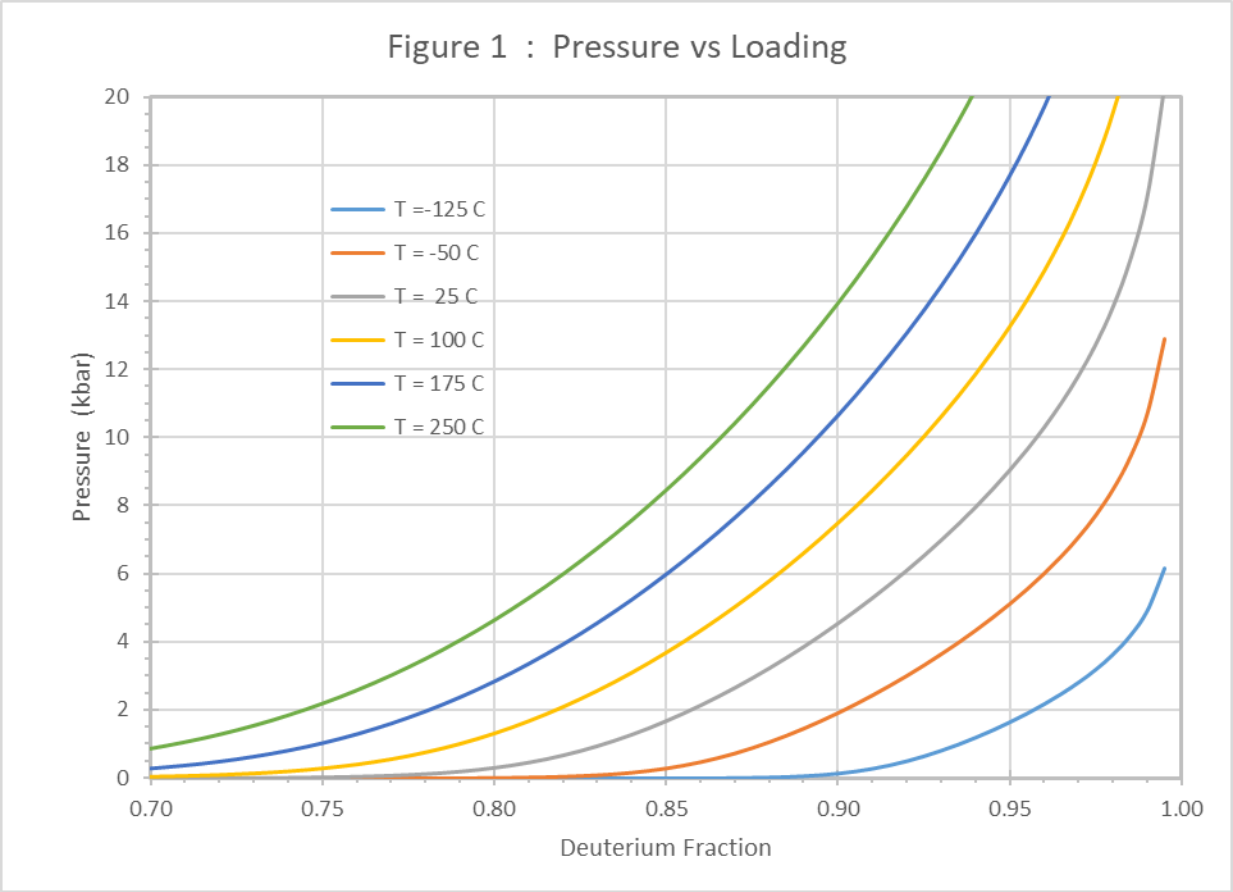
Generating LENRs : Loading

For loading, our goal has always been full, 100%, loading; one deuteron for each palladium atom. There is no proof that such perfection is necessary (or even desired); certainly, it is extremely unlikely that any of the previous LENR experiments (primarily electrolysis based) that reported seeing LENRs ever reached full loading. On the other hand, the belief in the field is that some threshold loading (e.g., 90% or more) is required, and that a prime weakness for most of the failed replication attempts has been insufficient loading. Hence, our thrust has always been to load as much as possible, with a goal of 100%.

We have achieved this goal, reaching 100% loading[3] via two different methods.

[3] Once loading is well above 99%, proving that it is actually 100.0% (as opposed to 99.8, 99.9, or 99.99%) is difficult; published correlations (e.g., to loading-metrics such as strain or resistivity) are not quantitatively established at these levels. Our belief that full 100% loading has been reached, is based on the fact that these loading-metrics, which normally keep changing as gas pressure is increased, stop changing once a high enough pressure is reached, despite further increases in pressure (other than changes from pure densification effects). This is discussed further in the Loading section of this report.

The first success came about at low temperatures, using the fact (evident in Fig. 1) that pressure requirements are lower for cold temperatures than for elevated ones. For most of our project's history, our pressure capability was restricted to about 10 kbar (i.e., 150 ksi, or 1 GPa). This is insufficient to reach full loading at room temperature, but has been successful at -100°C. While this was a nice achievement, it was also unsatisfying; not least because our ultimate reason for pursuing LENR research has always been powerplants, and these must operate at or above ambient temperatures. Accordingly, we have always desired even higher pressures, and have recently pushed these up to the 500 ksi level (over 30 kbar), achieving full, 100%, loading at room temperature.



Later in this report, we will explain how we have been able to reach and sustain H_2 and D_2 gases at pressures above 100 ksi. For now, let's just point out that we learned to do so routinely (performing several hundred LENR experiments at these pressures) and with a perfect safety record (never any injuries).

Although Callisto is primarily dedicated to gas loading, we have explored the use of our high-pressure capability to aid in electrolytic loading. This offers two potential advantages, performing electrolytic loading at temperatures well above 100°C, and using pressure to improve loading beyond that which electrolytic loading normally achieves. We have been able to perform electrolytic loading at pressures as high as 400,000 psi (i.e., 25,000 times ambient). With such

pressure-assist, we can reliably reach electrolytic-based loadings of 98%. But, since our D_2 loadings are all the way to 100%, these have remained our focus.

Generating LENRs : Material

For most of our effort, our "ideal" material has been pure palladium loaded with pure deuterium; "pure" meaning as close to 100% as possible, and in practice less than 10^{-4} or 10^{-5} impurities. However, the proper structural form for the PdD, has been much less obvious.

Our first goal was perfection, a single crystal of PdD. To that end, we have, for many years, purchased cm-sized single crystals of palladium, and loaded them with deuterium to values of over 95%. As is well known, the act of loading deuterium into palladium at ambient temperatures causes ~5.5% dimensional swelling in all directions, and induces an α–β phase transition; while the palladium may start out as a perfect crystal, it certainly isn't once loaded. We have solved this problem through "elastic loading"; using a pressure-temperature loading recipe which avoids the phase transition by doing the early portion of the loading at elevated temperatures. This has allowed us to achieve near-perfect crystals of highly loaded PdD. Unfortunately, this did not result in measurable LENRs.

While achieving "perfect" crystalline PdD is hard to do, at least it's a clear objective, and a singular one. Once one widens the material parameter space beyond this, there's a quandary; there are many more (effectively infinite) types of imperfect PdD than the single, perfect, one! Which of these many forms will lead to LENR? And, if we unsuccessfully try N possibilities, how do we know that is enough? Maybe the (N+1)st is the one that would have worked.

Obviously, there is no clear answer to this dilemma. There are a huge number of possible materials, with no practical way to search all of them.

So, here's what we have tried: Conceptually, split the space into two regimes; different material compositions, and different material structures.

For composition, we have focused on pure palladium, as well as its alloys with some other elements (Ag, Pt, Ni, Rh, B). Furthermore, in one scenario which will be described later, we have implicitly searched over most/all possible trace impurities. We have also taken a major detour, and done experiments with a number of rare-earth deuterides (e.g., Er and Yb).

For material structures, we have looked at a variety of different sizes, from nanoparticles, nanopillars, and nanofilms, up to macroscopic slugs. It has also been suggested that extremely distorted material, with many dislocations and cavities, is essential for LENR. As will be described later, our Thermal Ignition experiments naturally achieve such hyper-chaotic conditions.

Generating LENRs : Stimulation

The third condition we've attempted to reach, is applying the right stimulation to highly loaded PdD.

We had hoped that no stimulation would be needed; that simply achieving very high loading in PdD material would be enough to induce LENR. This possibility was aided by the fact that many of the reports of LENR from electrolytic experiments did not involve any explicit stimulation, beyond what might naturally be present during the process of electrolysis. Unfortunately, this did not prove to be the case. We have routinely loaded PdD to very high loadings in a wide variety of situations (perfect crystals, highly disorganized assemblages, nanostructures, at cryogenic-to-elevated temperatures, etc.) and then just watched for signs of LENRs; to no avail. Accordingly, we have come to believe that if LENRs are going to occur, the PdD will have to be stimulated in some fashion.

One potential reason for needing stimulation is to help deuterons get close enough to react. The deuterons in loaded PdD reside in octahedral sites within the FCC palladium lattice. There is one octahedral site for each palladium atom, and in fully loaded PdD, each deuteron is isolated, by itself, within its own octahedral site. These act as traps for the deuterons, isolating them from each other; the site-to-site separation is 2.9 angstroms, keeping the deuterons way too far apart to react with each other (at least in any conventional-type nuclear process). {Note that the two deuterons in a D_2 molecule are four times closer than this; and these certainly do not react.} Accordingly, some form of stimulation seems essential, driving deuterons out of their octahedral traps and giving them the opportunity to encounter one another.

Temperature may be enough; the natural thermal vibrations of the palladium lattice do jostle the deuterons inside their octahedral traps, and this leads them to occasionally escape. This turns out to happen fairly frequently; at room temperature, each deuteron will spill out of its trap in less than a nanosecond. This process is one which we had hoped would lead to LENR without the need for active stimulation. While that was not the case, the rate of thermal escapes grows rapidly with temperature. Hence, elevated temperatures, and the heating that causes it, are straightforward, and promising, forms of stimulation.

In practice, there are many different ways to heat a PdD target, hence there are potentially many different forms of thermal stimulation. However, viewed from the standpoint of a unit cell in the PdD lattice, these have almost identical effects. Deuterons vibrate within their octahedral sites with about a 100 femtosecond (fsec) period; heating pulses longer than this, from the deuteron's viewpoint, just cause very slow motion of the surrounding Pd atoms. The deuterons easily keep up with this motion, and stay centered within the traps, not picking up much differential back and forth vibration. For this "slow" heating, the details of the heat source, or its time profile, don't matter; all that matters to the deuteron, and its probability of escaping its octahedral trap, is the magnitude of the resultant temperature. This greatly reduces the parameter space of potential forms of thermal stimulations; most of the possible "differences" may matter a lot to the experimenter, but they don't to the deuteron. Accordingly, we have applied a variety of different thermal sources (e.g., lasers, resistive heating, thermal furnaces) but simply for their utility in a given experimental situation, not because of any fundamental physical difference.

There are, however, some features of thermal stimulation which can result in different types of stimulation. One occurs when the heating is applied on timescales comparable to, or even faster than, the deuteron oscillation time, i.e., 100 femtoseconds. Most heat sources are much slower than

this, so as mentioned above, are fundamentally indistinguishable. However, lasers with pulse lengths at or below 100 fsec are commercially available, and these can be used to deliver energy fast enough to force deuterons out of their traps. When the timescale of a pulse is much shorter than the deuteron's oscillation time (e.g., 50 fsec or shorter), the deuteron can be impulsively kicked out of its trap. Alternatively, the source can vary resonantly with the deuteron's oscillation (e.g., via a 10 THz laser), and efficiently "pump" its motion until it leaves the trap. The potential advantage of these types of short-pulse stimulation is not actually their effectiveness in helping deuterons escape from octahedral traps (remember, gradual thermal excitation does this quite well), but that they can drive each deuteron out simultaneously with all its neighbors. We have applied both of these types of stimulation; impulsive kicks with 35 fsec laser pulses, and resonant excitation with 10 THz radiation.

We have also exploited some other aspects of thermal heating, in order to achieve different types of stimulation. One of these exploits the fact that thermal diffusion is much faster than the diffusion of deuterons within the palladium lattice. This dichotomy has allowed us to stimulate very highly loaded PdD with temperatures as high as the melting limit of the lattice; it takes time for the deuterons to deload, and in the meantime, they experience the high temperature, and potentially undergo nuclear reactions. Another useful application is when we are interested in heating PdD near a surface, rather than in the bulk; by combining the short penetration depth of optical radiation with a pulse length shorter than that of thermal diffusion, the applied energy can be temporarily retained near the surface of a PdD target.

In practice, most forms of power applied by experimenters just result in one or another form of thermal stimulation. The different modalities may seem significant, but their actual effect is still thermal. And, since most of these methods are much slower than 100 fsec, their thermal effects are similar as well.

One form of input power which is potentially different is electrical fields. Light, of course, is one example of this. Usually, however, it is the absorbed power from the light that is more useful than the direct electric fields; i.e., light is just another (albeit, very convenient) way to deliver thermal stimulation. The actual electric fields from incident light are generally quite weak and do not penetrate far beyond a thin surface layer of PdD; furthermore, when using IR or visible wavelengths, the field oscillates too rapidly to have any cumulative effect on deuteron motion. In order to actually test electric field stimulation (not just its thermal effects), we have instead used capacitors to apply large electric fields; either by using PdD as one of the electrodes, or by using rare-earth deuterides as the dielectric.

There are certainly a number of other stimulation approaches that have been, or could be, used in LENR experiments. One, for instance, is fluxes of deuterons flowing within PdD; e.g., via electromigration or differential loading. We have designed such experiments, but not implemented them; one consideration being that, at the atomic scale, these cause much less deuteron motion than that due to natural thermal hopping. Another approach is to bombard a PdD target with incident particles, such as electrons or deuterons; the former is used in glow discharge experiments, and the latter to study D-D fusion enhancement by electron screening. We have not pursued either of these approaches.

In general, as with a search for the "right" material, there are potentially a great number of possible stimulation approaches that could be tried, searching for the "right" one. We have focused our search on the three methods discussed above; thermal, fast-pulse, and electric. There are two main reasons for not widening our search further. The first, is that a large proportion of other stimulation approaches are ultimately just different ways to apply thermal stimulation; not fundamentally different from what we've investigated. Another reason for not devising and pursuing more exotic, non-thermal, stimulation methods, is that none of the previous "successful" LENR experiments did so. Remember, the only reason for pursuing LENR research at all, is the experimental reports of nuclear energy production. Absent these, the LENR field would not exist; its conception was certainly not theory-driven. While, in principle, LENRs might be generated by exotic forms of stimulation, these were not needed to produce the results which the LENR field is based on. Hence, there is no reason to believe such stimulations are necessary, nor, of course, that they will work at all.

To summarize, we believe that creating LENR events in PdD depends on three factors; high enough deuterium loading, the proper material, and effective stimulation. We have reached perfect (100%) loading, have tested a range of material structures from perfect crystals through nanostructures and highly deformed material, and have tested a range of thermal stimulations from cryogenic to melting temperatures, as well as non-equilibrium ultrafast heating and electric fields. Unfortunately, none of these conditions have led us to detect LENRs.

Detecting LENRs

The preceding statement begs a question; how have we tried to detect LENRs? And, with what sensitivity; is it possible that we actually have generated LENRs, and simply missed detecting them?

For us, detectability has been a crucial pillar of our project. One element has been to drive our detection sensitivity down as low as possible; i.e., to be able to detect single LENR events, not simply the cumulative effect of billions of them. Another aspect has been the need to eliminate false positives; i.e., to reject signals that appear to be from LENR events, but aren't.

The nuclear reaction which we've been focused on detecting is a very different form of deuterium-deuterium reaction than the conventional one. Rather than creating energetic pairs of either a proton and triton or a neutron and a He^3 nucleus, the canonical LENR event is hypothesized to form only a stable He^4 nucleus, with the 23.8 MeV reaction energy appearing as heat within the PdD lattice.

Therefore, the detectable signals of LENR are expected to be helium and heat. Of course, the most clearcut sign of LENR would be to detect both of these together, in the expected ratio of 23.8 MeV per helium atom. Over the years, a number of experimenters have reported successfully doing so. Unfortunately, these experiments have not reached the levels of precision or reproducibility

necessary to convincingly demonstrate the reality of LENR. While correlated heat-helium levels would indeed be very strong evidence for LENRs, we have not chosen to pursue this approach.

The reason is that helium is a lot more difficult to measure, in the LENR context, than heat is. One difficulty is background suppression; helium is already present in the atmosphere, in the palladium and deuterium, and in all the other experimental equipment. So, one needs to distinguish LENR-generated helium, from pre-existing helium. Doing so, requires both sufficient amounts of LENR and robust controls (e.g., comparing helium measurements between PdH and PdD experiments).

Of course, heat measurements also face the need for background suppression, so also demand strong controls and above-threshold LENR levels. Helium poses more challenges for two reasons. First, the non-atmospheric background levels are rarely consistent; the helium levels in materials can differ from batch to batch, even for nominally identical materials; this makes it hard to perform reproducible experiments. The second challenge comes from the signal levels; for instance, 1 joule of heat corresponds to 0.5 picomoles of helium. The former is easy to detect, the latter is not. In particular, our goal has been to be able to detect individual LENR events; detecting 24 MeV of heat is doable, detecting a single helium atom is not. Accordingly, we have focused on heat measurements, not helium ones.

In principle, heat is easy to detect. The challenge, is proving that it comes from LENR, rather than the myriad of other sources present in the environment, as well as in any particular experiment.

To accomplish this, strong controls are essential. Our default control has been comparing experiments using PdD to ones using PdH. The advantage of isotopic controls is that both protium and deuterium load into palladium in the same way, and to very similar degrees. {Protium loads a bit easier than deuterium; so at a given pressure, PdH will have a slightly higher loading than PdD.} Because loading, and other chemical behaviors, are similar but not identical, PdH is more useful as a qualitative control for PdD than a tight quantitative one. Accordingly, if we see some unexplained morphological feature in a PdD experiment, we always repeat with PdH; if the same type of feature shows up there (even with greater or lesser frequency), we discount it as a sign of LENR. Likewise, if we observe some time-dependent temperature profile in a PdD experiment, but the same behavior happens in PdH, we discount it. However, because there are loading and chemical differences between protium and deuterium, we do not generally interpret quantitative differences between PdD and PdH experiments to be proof of LENR. The challenges to making quantitative PdH-vs-PdD comparisons will be illustrated later in the discussion of our Thermal Ignition campaign, and the extraordinary steps taken there to enable doing so.

The comparison between PdH and PdD experiments is only one of the controls we've used. Depending on the type of experiment being performed, other controls can be crucial. For instance, we often vary the stimulation used, comparing a test with stimulation to one without any, or testing with different amounts or varieties of stimulation. In other circumstances, we compare material effects; for instance, exposing Pd to equal pressures of D_2 versus He (one will load, the other won't), or placing either Pd or Pt into D_2 (again, one will load, the other won't). In all cases, we need to distinguish effects which might be due to LENR from ones due to more conventional (even

if unexplained) causes. The type of controls best suited to do so can vary from experiment to experiment, but their role is always essential.

The obvious way to look for heat from LENR is to do so directly. This might be done by simply measuring temperature; e.g., placing a temperature probe into the PdD or some other part of the experiment. If possible, however, it's good to perform the experiment within a calorimeter, so as to measure the overall heat production.

Calorimetric experiments have formed the basis of a great many previous LENR experiments, and many different types have been employed. One concern, however, is this lack of standardization. Usually, LENR researchers design and build their own calorimeters; e.g., to match the size and needs of the experiment. Unfortunately, despite the researcher's care in calibrating their instruments, outsiders generally have to take such measurements with a degree of faith; the apparatus is one-of-a-kind, with difficult-to-verify characteristics. To avoid these issues, we have concentrated on using commercial calorimeters. This choice carries some disadvantages (the experiment must adapt to the calorimeter, not vice versa), but also the key advantages of independence and traceability; the calorimeters come from specialized hands-off suppliers, have known verifiable characteristics, and are available to any other researchers.

Despite the advantages of calorimeters, many of our experiments did not use them, instead using other ways to measure the heat released from AHE events.

In many cases, temperature sensors simply provided more useful data than our calorimeters could. For instance, the calorimeters we were using had time constants of several minutes, so temperature probes were necessary to measure the details of faster thermal events. Another advantage came when we needed to measure higher temperatures than our calorimeters could handle; temperature sensors exist that can do so.

Probably the biggest limitation we ran into with calorimeters was due to size; they were often either too small, or too large, to be useful.

The "too small" problem is obvious; high pressure equipment is often large and complicated; some of our pressure vessels would fit within our calorimeters, while others could not. The "too large" problem arose because (to enhance sensitivity) we frequently wanted to shrink the size of our measurement region. In these cases, it was more valuable to us to know the temperature of a 10 μm region, than the heat produced by a cm-sized one. Or, even more significant, we wished to gather data from thousands or millions of individual small regions, not just a single large one; our calorimeters were simply too large and cumbersome a tool for that.

Accordingly, we developed and used a wide variety of non-calorimetric ways to measure thermal information in our experiments. Some were direct temperature probes (e.g., thermocouples and RTD[4]s); these were used both for large targets, and for small, chip-based ones. Other tools, such as IR cameras, gave us temperature information from thousands of separate sites across an image.

[4] RTD : Resistance Temperature Detector

We have also used two much less direct ways to attempt to detect LENR events by their energy release. One of these is seismology. Just as macroscopic events (e.g., earthquakes or nuclear tests) launch seismic waves through the earth, LENR events, being localized impulsive releases of energy, will also launch vibrations through the surrounding material. These persist for durations and distances well beyond the event itself. The vibration levels are small, and detection is complicated by background electrical noise, as well as by vibrations inherent in loading/deloading of PdD. Because of these challenges, we don't know how well this approach will ultimately work (development was still underway when Callisto ended).

The other indirect approach is archaeology. Its goal is to detect the permanent after-effects of LENR events (ideally of single individual events).

Single-event detection is a holy grail of LENR efforts, particularly in the current scientific verification phase, where we're trying to learn whether events occur at all. Traditional calorimeters, certainly including the commercial ones we use, are only sensitive enough to detect macroscopic levels of heat; however, this requires the occurrence of a very large number of LENR events (for instance, 1 mW of heat corresponds to 250 million LENR events per second). At our current levels of expertise, we may simply not be generating anywhere near this many LENR events. We need to know whether LENR is occurring at all; and for this, single event detection is the desired goal.

Detecting the heat from a LENR event, in real-time, is impractical. The thermal effects are most likely very local[5] (<100 nm), as well as quite fleeting (<100 psec). Monitoring a sample with either this spatial resolution or this time resolution, is quite doable; doing both simultaneously is not! Fortunately, however, we don't really care when an event occurs, just that it does. This is the attraction of looking at thermal history, aka archaeology, for trying to detect single LENR events.

We have used two different forms of archaeology:

One approach was magnetic; using the fact that raising the temperature of a magnetic material (e.g., above its Curie temperature) will permanently change its magnetic state. This transition occurs even if the temperature change is temporary and very quick (sub-psec). Other types of thermal history recorders are far too slow to be useful to detect LENR events. However, the speed of magnetic transitions enables the thermal pulse from a PdD LENR event to be detected as a permanent change in the local magnetic field of a nearby magnetic material.

Implementing this magnetothermal approach requires choosing the magnetic material, placing it very near the pressurized PdD target, and being able to magnetically image the material after the experiment. We developed two types of material-target systems. The initial version involved stripping overlayers from commercial hard-drive disks, and then depositing Pd films onto the disks. The subsequent approach was more sophisticated; we deposited arrays of small magnetic nanodots onto a Pd substrate. While we considered methods to actually image magnetic fields,

[5] I.e., to phonons of the nearby lattice, involving as many as tens of millions of atoms. As discussed in the Conclusions section of the report, if heat is spread globally (i.e., over micron ranges and billions of atoms), archaeological detection will not be effective.

magneto-optics and SEM-based, we ultimately fielded two scanning methods, one using magnetic force microscopy, and the other using TMR[6] read-heads salvaged from hard drives. As magnetic archaeology was a novel approach, it had just begun being operational in LENR experiments as Callisto concluded, and some aspects were still being developed and improved at the time.

The other approach was a more classical form of archaeology; looking for the physical scars which an LENR event will leave behind in the material. The energy from an LENR event is sufficient to melt a 100 nm bubble of PdD or to vaporize a 50 nm one. Damage on these size scales should be detectable, especially if it occurs near the surface of the material.

We have gone through three evolutions in learning how to look for LENR using this approach; all based upon SEM imaging in order to provide high resolution imagery over relatively large areas. The first, and simplest, effort, was to look for craters on the surface of a post-experiment Pd target. Craters proved to be hard to distinguish from the normal (generally deformed) surface, so we switched our archaeological efforts to nanostructured targets, where damage would be more visible. The first variety of these used arrays containing millions of Pd nanopillars, each one narrow enough (typically 60-70 nm) that an LENR event in a pillar should break it, and hence be visible in SEM images. Unfortunately, this approach gives ambiguous results; it's not possible to distinguish pillar breakage which might be due to LENR from that due to impacts with tiny debris particles endemic to experiments. Subsequently, we have fielded arrays of thin (20-30 nm thick, 10 μm wide) Pd nanofilms on inert substrates. These films are thin enough so that energy from an LENR event will blow a hole through the film, down to the substrate; the contrast between the exposed substrate within the hole and the Pd film surrounding it, makes such holes stand out in SEM imagery. Given this performance, and the ease of making thin film arrays in Pd and its alloys, the nanofilm approach became our standard archaeological tool.

[6] TMR : Tunnel Magneto-Resistance

Introduction to the Campaigns

Our search for LENR embodied a variety of campaigns, focused on trying to generate nuclear reactions (via combinations of loading, materials, and stimulation) and/or developing sensitive methods to detect them.

These will each be discussed, in detail, in the following sections of this report.

But first, here's a brief synopsis of what the campaigns were, and why we did them.

Let's first introduce two campaigns which did not directly try to make or detect LENR, but instead developed the technical infrastructure which was essential for the other campaigns to do so.

Synergy : Software and data-handling are ubiquitous parts of any modern scientific effort. But, in a fast-paced, parallel-track, program like ours, these usually become bottlenecks, retarding progress. A major reason is siloing; each experimenter, and often each of their key instruments, uses their own computers and codes, and retains its own data. Synergy was developed to address these issues. It provided a common platform to document all experiments, the materials used in them, and the data collected by them[7]. It was designed to make operation and interactivity as easy and seamless as possible; user-to-user, user-to-equipment, and user-to-the-past.

Loading : The most basic precondition necessary for PdD-based LENR, is loading deuterium into palladium. Callisto's primary route to do so, used high pressure D_2 gas. And, since we wanted loadings all the way to 100%, this meant we needed D_2 pressures as high as 500,000 psi. This is not an off-the-shelf capability, particularly since hydrogen embrittles almost all structural metals. This campaign had to develop the pumps to pressurize D_2, as well as vessels which were strong enough to host experiments at these pressures. In addition to developing the pressure hardware, we learned how to load deuterium into palladium without disrupting it, and how to measure near-stoichiometric loadings.

The next seven campaigns were focused on investigating different way to generate LENRs.

Thermal Ignition : After our early room-temperature experiments showed no signs of LENR, we searched for some form of successful stimulation. Temperature was the choice, but at a given pressure (e.g., 150 ksi) there is a strong tradeoff between loading and temperature, they can't both be high. This campaign was designed to give us both high loading and high temperature at the same time. We would load deuterium into Pd at room temperature (or even cryogenic), reaching 96-99% values. Then we would quickly heat the target with a laser. For a short time, 10 seconds or so, the loading would remain high; we looked for LENR during this high loading, high temperature period. Due to the nature of these experiments, we had to use temperature

[3] One guide to the reader; throughout the discussions there will be frequent references to BE#s or BD#s, e.g., BE4131. These are Synergy "pages" documenting our experiments (BE#) or devices (BD#).

measurements as the diagnostic. As a result, very careful controls (usually PdH vs PdD) became essential, and these evolved into our most complex LENR experiments.

Staring : This campaign was launched out of our desire to look for LENR using finer spacetime resolution than possible with full-target sensors (such as a calorimeter or a single temperature probe). Our approach was still temperature sensing, but done using a fast IR camera. Remote imaging allowed us to sense the temperature of thousands of small spatial pixels every millisecond or so. Our stimulation of choice was thermal, implemented by a programmable series of short laser pulses. Spatial control of the irradiation allowed us to heat different regions of the target separately. In the later stages of the project we used Staring's laser irradiation as a controllable, albeit fairly gentle, stimulation source for a variety of archaeology measurements

Chips : The most effective knob to turn, in order to detect ever fewer numbers of LENR events, is to reduce the size of the target being tested. The Chips campaign was focused on doing this, by using modern lithographic techniques to make small Pd targets, usually with integral stimulation sources and LENR detectors. We designed and fielded a variety of different systems. The canonical one involved a thin layer of Pd mounted above a Pt RTD-heater combination. In operation, the system would be placed in high pressure D_2 gas, thus loading the Pd; then current in the Pt would apply a short heat pulse to the PdD, and the RTD would record the resultant temperature history; this process would be repeated tens of thousands of times. Among the other Chips experiments were Estim ones, using micro-capacitors to apply electric field stimulation to Pd or rare-earth films; a fundamentally different type of stimulation than our normal thermal one. We also fielded Chips experiments using nanopillars and nanofilms for archaeological detection, as well as ones testing electric-field-concentrating Pd nanostructures such as bowties and THz metamaterials.

Ultra-Short Pulse : As we've discussed, most forms of thermal stimulation are fundamentally, on the atomic scale, very similar. However, once energy is applied on timescales comparable to, or much shorter than, the oscillation period of deuterons within the palladium lattice, things change; we can drive more intense deuteron motion, and can move neighboring deuterons coherently. The timescale needed is about 100 fsec; although this is short, there are many commercially available lasers with shorter pulses. We used both single pulses and resonant chains of pulses to attempt to stimulate LENRs. We used a variety of diagnostics for these experiments, including pump-probe sensing, thermal probes, and archaeology.

Terahertz : The outcome of thermal stimulation is predicted to be enhanced hopping of deuterons out of their octahedral traps, so that they can encounter each other and have the opportunity to undergo LENRs. Terahertz stimulation has the potential to enhance this hopping much more directly, by generating phonons tuned to resonantly drive deuteron oscillations. Since PdD is a metal, very little of arriving terahertz radiation can penetrate below the surface and cause hopping; hence intense terahertz sources are required. We tried to deliver strong enough terahertz power in two different ways, by nonlinear generation in plasma above the PdD surface, and by lithographically creating resonant terahertz metamaterials on the surface.

High Pressure Electrolysis : We have used our high-pressure capability to improve electrolytic loading. Doing so, required different (in some ways greater) hardware challenges than gas loading,

notably due to the corrosive nature of strong electrolytes. We were able to demonstrate 98% PdD loadings at pressures as high as 400 ksi. However, we found that the advantages of pressure leveled off; 150 ksi pressures worked as well as 400 ksi, and involve considerably more convenient hardware and electrolyte mixtures. With these systems, we performed high pressure electrolytic experiments, using our commercial calorimeters to look for LENR.

Replications : We are not the only LENR researchers to use D_2 gas in order to load Pd. Only a couple of efforts explored the very high pressure regime that we have, and they did not report seeing any sign of LENRs. There have, however, been a number of notable lower pressure experiments that have claimed to see LENRs, and so we have attempted to replicate two of these. One effort was carried out by Les Case in the early 2000s. He used commercial Pd-catalytic material (Pd nanoparticles on a carbon substrate), and cyclically loaded/deloaded this with D_2 gas at temperatures of around 200ºC; LENR detection comprised comparing temperatures reached with D_2 gas to those with H_2. Another was pioneered by Yoshiko Arata in the 1990s, and was extended after his death by other Japanese researchers. These experiments also use nanoparticles, typically PdNi combinations embedded in a ZrO_2 matrix, and exposed to elevated temperatures and pressures. In both of our replication attempts, we first tried using apparatus similar to theirs, and then migrated to our standard calorimeters; in neither case did we detect signs of LENR. In another, unrelated effort, we examined the 1989 LLNL electrolytic experiment. Here, we were not trying to look for LENR (Callisto electrochemical loading is performing replications experiments for that), but instead tried to identify possible non-LENR explanations for its results.

The final five campaigns concentrated on different ways of detecting LENRs.

Calorimetry : This is traditionally the gold-standard for thermal detection of LENR. Our main "wrinkle" was the exclusive use of commercial calorimeters, rather than custom-built ones. As previously discussed, this choice gave us sensitive, well-characterized calorimeters, while providing easier traceability and replicability of our measurements for other researchers. We used these calorimeters in a wide variety of high pressure LENR experiments. The most sensitive usage was in passive experiments, with no stimulation; these included bulk Pd targets, as well as Pd nanoparticles and aerogels. Since LENR likely requires some form of stimulation, we also performed experiments using resistive heating pulses. A final, quite different, use of our calorimeters, was to perform high pressure liquid-based electrolysis in them.

Physical Archaeology : This was our most sensitive LENR detector, and should be able to record single localized LENR events. It worked by using arrays of very small Pd nanostructures; each one thin enough so that energy from even a single LENR would permanently damage the nanostructure, in a way that would be visible in SEM images taken after the experiment. We fielded two different types of these structures; nanopillars, which were vertical Pd cylinders about 60-100 nm wide, and nanofilms, which were 20-30 nm thick films of Pd. Most of our nanopillar experiments occurred in the Staring and Chips campaigns. Later, because of their versatility and ease of fabrication, we concentrated on nanofilm detectors. In addition to use in Staring and Chips, we also used these in the Ultra High Pressure apparatus with both long and ultra-short laser stimulation, as well as in high-pressure electrolysis. Furthermore, we were able to field nanofilms using not just Pd, but also a wide variety of Pd alloys, and some other metals.

Magnetic Archaeology : This form of archaeology did not rely on permanently damaging a Pd target, but rather on the transient thermal spike from a LENR event changing the magnetic field in a nearby magnetic material. This is a novel, and very-fast-acting form of a thermal history sensor; it tells us what's most important (that a given temperature was reached) without having to determine exactly when that occurred (a fact that we care much less about). Implementing these detectors was technologically quite challenging; we learned to fabricate arrays of magnetic nanodots, and to image magnetic changes in them using either magnetic force microscopy or tunneling magnetoresistance read-heads.

Seismology : This was an attempt to measure the vibrations induced within materials by the sharp thermal impulse caused by a LENR event. Basically, it's an analog to how one detects earthquakes or underground nuclear tests; abrupt impulses launch waves which persist long after, and far away from, the event itself. We faced two key challenges: One, was detecting the weak electrical signals from an event's wave given the background electrical noise in our facility. The other, was the presence of non-LENR impulses as the Pd was interacting with the high pressure D_2 gas. This campaign was begun in the last year of our project; as a result, we were still working to resolve the two challenges, and have not determined the final sensitivity of this approach.

Nuclear Detection : Since the only products of conventional deuterium-deuterium reactions are energetic particles, we devoted some resources towards detecting nuclear products, postulating that they may at least constitute a side reaction to LENR's heat channel; and, crucially, that nuclear products would be a much less ambiguous sign of nuclear reactions than heat. We used two main approaches; using scintillators to look for γ-rays, neutrons, or charged particles escaping the sample during an experiment, and placing samples in gamma well detectors after experiments to look for delayed γ-rays resulting from nuclear transmutation.

[Section Omitted]

[Section Omitted]

[Section Omitted]

[Section Omitted]

[Section Omitted]

[Section Omitted]

Loading

Concept

Callisto has focused on using pressure (generally via D_2 gas) as the tool to achieve very high loadings of deuterium into palladium metal. This, of course, is not the D_2O electrolysis method used by most LENR researchers. Gas loading offers one extraordinary attraction; at a given pressure and temperature, the loading is governed by thermodynamic equilibrium. This makes loading stable, reliable, and reproducible. Experiments at the same pressure and temperature will reach the same loading, every time, and for every researcher. Furthermore, the loading will stay at this equilibrium value, for as long as the pressure and temperature are maintained.

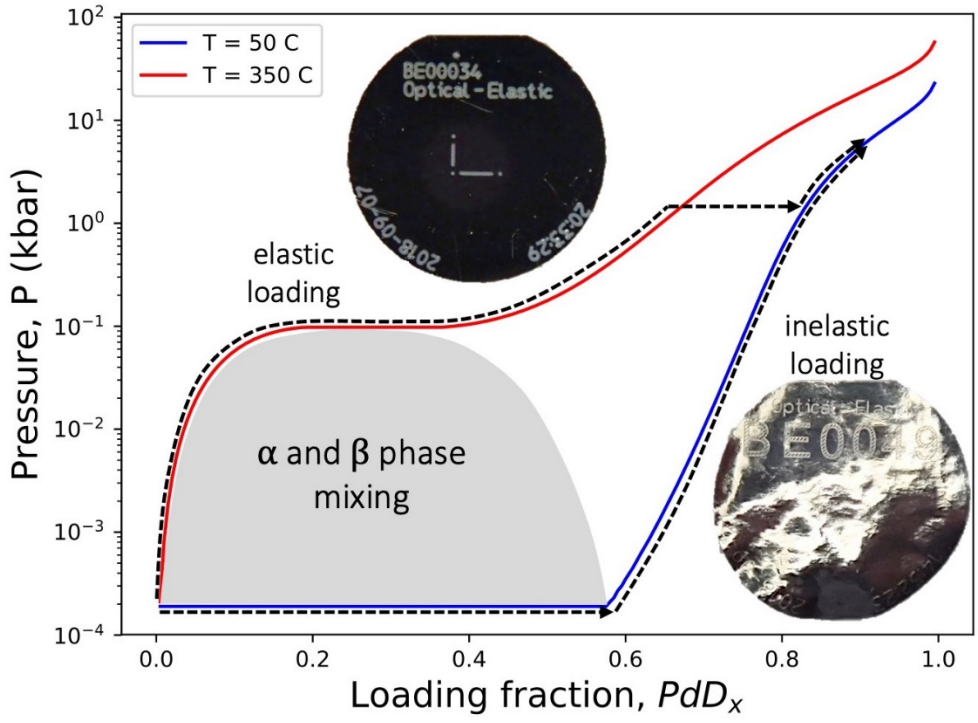
There is, however, one problem with gas loading; reaching high loadings requires extremely large pressures. For example, in order to reach a loading of 90% (often considered the threshold level for LENR), one needs a D_2 pressure of 70 ksi (70,000 psi), and then increasing this to 95% requires almost twice as high a pressure (135 ksi). These pressures are far beyond those commonly available to LENR researchers; while diamond anvil cells are indeed capable of pressures dwarfing these, the price is that their experimental volumes are only fractions of a millimeter in extent.
For Callisto to successfully operate gas loaded LENR experiments at loadings above 95%, we had to solve two practical problems. One was to develop the pressure equipment allowing us to field large (centimeter sized) LENR experiments at pressures of 150,000 psi and above. The other challenge was learning how to handle PdD under these conditions; having it survive the internal loading stresses and being able to accurately measure loadings as values approach 100%.

As protons/deuterons load into a metal lattice, material properties change (mechanical, optical, etc.). Furthermore, high fractional loadings (e.g., >90%) are theorized to be necessary for AHE. Hence, carefully characterizing a sample's loading level is essential to investigate AHE. A key challenge is that most loading measurements are proxies for loading i.e., they are indirect measurements that require calibration. However, gas loading has an advantage: since the loading is in thermodynamic equilibrium, it can be calibrated as a function of pressure and temperature. Hence, once we measure the equilibrium loading values with a true mass loading technique, we can calibrate indirect loading measurement techniques, which are more straightforward and enable in-situ, real-time measurements. We developed two true mass loading techniques: the UltraSorb (a form of a Sieverts apparatus) and the MicroSorb (a mass measurement using a microbalance). We leveraged three in-situ, indirect loading measurements: lateral strain, resistance, and reflectance. A nuance to loading (gas or electrolytic) is that the loading protocol matters. Loading deforms the lattice and leads to both nanoscale and macroscopic deformations that obscure loading and AHE diagnostics. Hence, we developed a gas loading protocol that we've termed "elastic loading", which uniformly loads a lattice and mitigates deformations.

Elastic Loading

Metal hydrides swell as they load protons/deuterons. For example, Pd swells up to ~17% volumetrically at full loading. The swelling is often uneven, which distorts the material on both nanoscales and macroscales and convolutes loading and AHE diagnostics. So, our objective is to uniformly load Pd without distortion during swelling. This preserves the surface morphology making it easier to detect AHE damage sites and enables optical loading measurements such as lateral strain and reflectivity. The key insight is to load (and deload) at high temperatures, which we've termed "elastic loading".

At high temperatures, protons/deuterons quickly diffuse through the Pd lattice, ensuring uniform loading throughout the material. At low temperatures, Pd loads more slowly which introduces loading gradients into lattice deforming the material. Pd's lattice constant is 3.89 Å while β-PdH_x and β-PdD_x lattice constants are 4.09 Å and 4.084 Å, respectively[8]. When loading gradients are present, the up to 0.2 Å lattice constant mismatch introduces significant stress into the lattice, which produces deformations. Room-temperature loading passes through the $\alpha - \beta$ phase mixing region, which makes distortions inevitable (**Fig. 1**). We've termed non-unform loading that passes through the $\alpha - \beta$ phase mixing region "inelastic loading".



[8] Schirber, J. E.; Morosin, B. "Lattice constants of β – PdH_x and β – PdD_x with x near 1.0." *Phys. Rev. B.* 1975, 12, 117-118.

Fig. 1: Inelastic loading is characterized by non-uniform loading and passage through the $\alpha - \beta$ phase mixing region, which produces nanoscale and macroscopic distortions. As shown, an inelastically loaded Pd foil visibly shows deformations. Elastic loading/deloading follows a high-temperature isotherm which ensures uniform loading and avoids the $\alpha - \beta$ phase mixing region, mitigating nanoscale and macroscopic distortions.

At temperatures above the phase threshold (e.g., 350°C or higher) the free energy no longer has a loading inflection, and so the $\alpha - \beta$ phase transition ceases to exist. The elastic loading protocol exploits this effect, by pre-heating the Pd to 350°C in an inert gas (or at high vacuum) and then slowly ramps the H_2/D_2 pressure up to >1100 psi; loading continues by cooling to room temperature and subsequently increases the pressure to the desired setpoint (e.g., for most AHE experiments, the setpoint was >100 ksi). To deload, the reverse process is followed. As an example, we tracked a Pd sample's surface roughness with a confocal microscope and measured the lateral strain between laser marked fiducials using optical images. We show the lateral strain during the elastic loading process and that the surface roughness remains constant (**Fig. 2**).

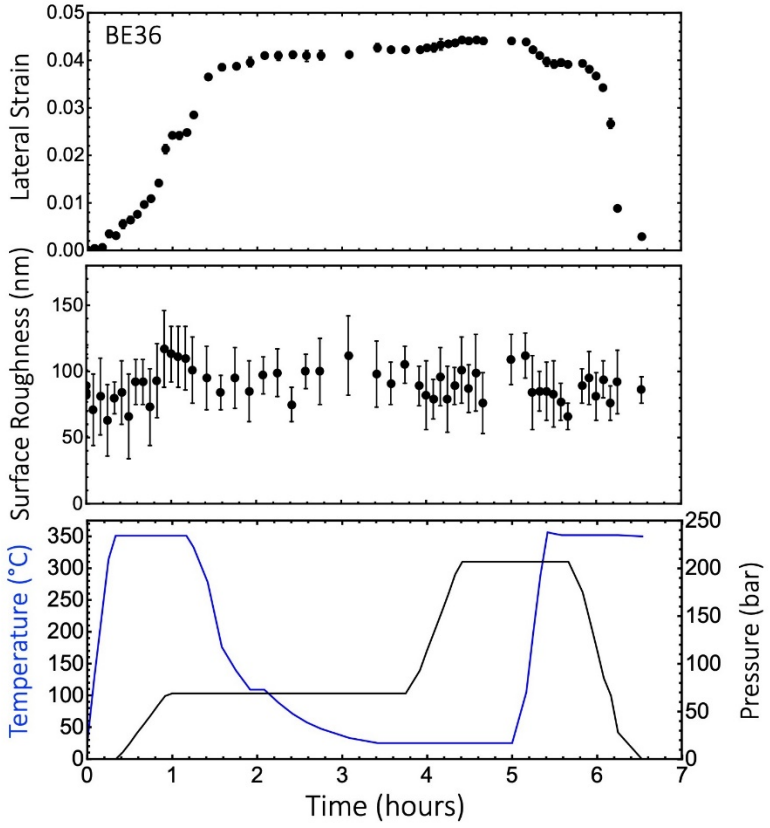


Fig. 2: The elastic loading protocol pre-heats the Pd to 350°C, slowly loads to ~75 bar (~1100 psi), cools to room temperature (~25°C) and then increases pressure the desired setpoint. The deload follows the reverse process. The Pd sample's lateral strain is proportional to the loading and shows the loading and deloading. The surface roughness remains constant during elastic loading and deloading.

While elastic loading mitigates macroscopic deformations, nanoscale distortions may still be present depending on the sample's crystallinity. We developed sample preparation processes (e.g., mechanical polishing, chemical cleaning, ion milling, etc.) that produce pristine Pd surfaces (**Fig. 3**). However, polycrystalline Pd sample surfaces are modified even after elastic loading. Grain boundaries contribute to most of the post-loading surface distortions: intergranular voids and cracks, and grain growth. Terracing is a feature of high loadings (e.g., >90%) and varies with grain orientation (**Fig. 3**). To mitigate grain boundary driven distortions, we moved to Pd single crystals: (100), (110), and (111) orientations. We've shown with XRD (X-Ray Diffraction) that the Pd is in fact single crystal (**Fig. 4**). After elastic loading and deloading, single crystal Pd sample surfaces remain pristine, successfully mitigating distortions driven by grain boundaries (**Fig. 5**).

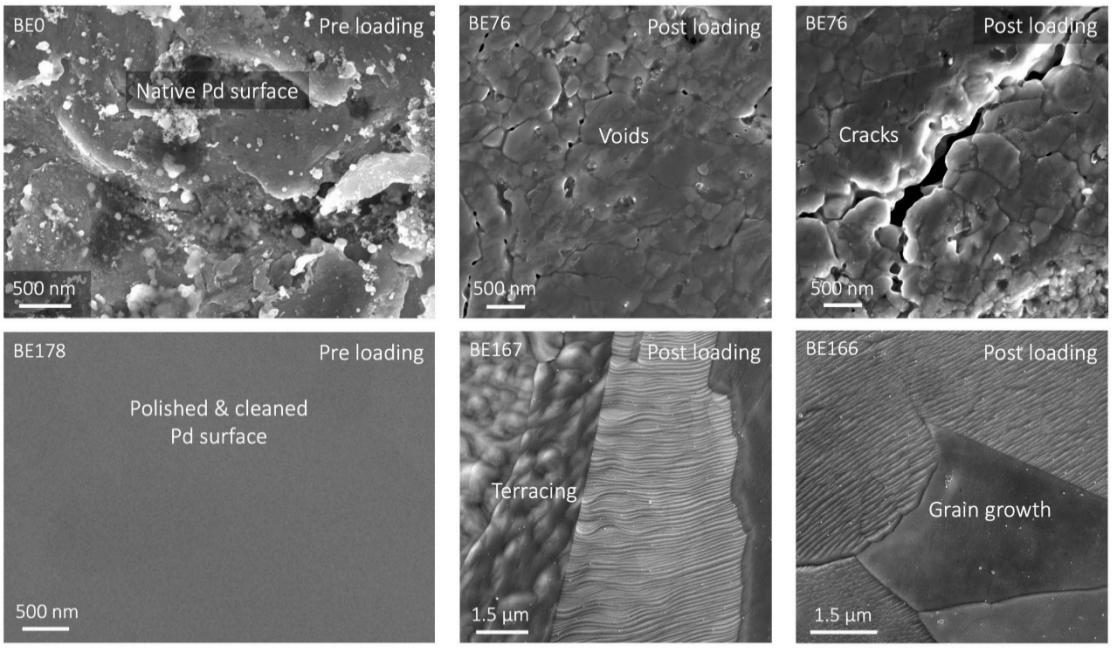


Fig. 3: Native polycrystalline Pd surfaces are rough and full of debris / contamination. After mechanically polishing and chemically cleaning, the Pd surfaces are pristine. Post elastic loading, polycrystalline surfaces are distorted (e.g., intergranular voids and cracks, grain growth, and terracing), primarily due to the presence of grain boundaries.

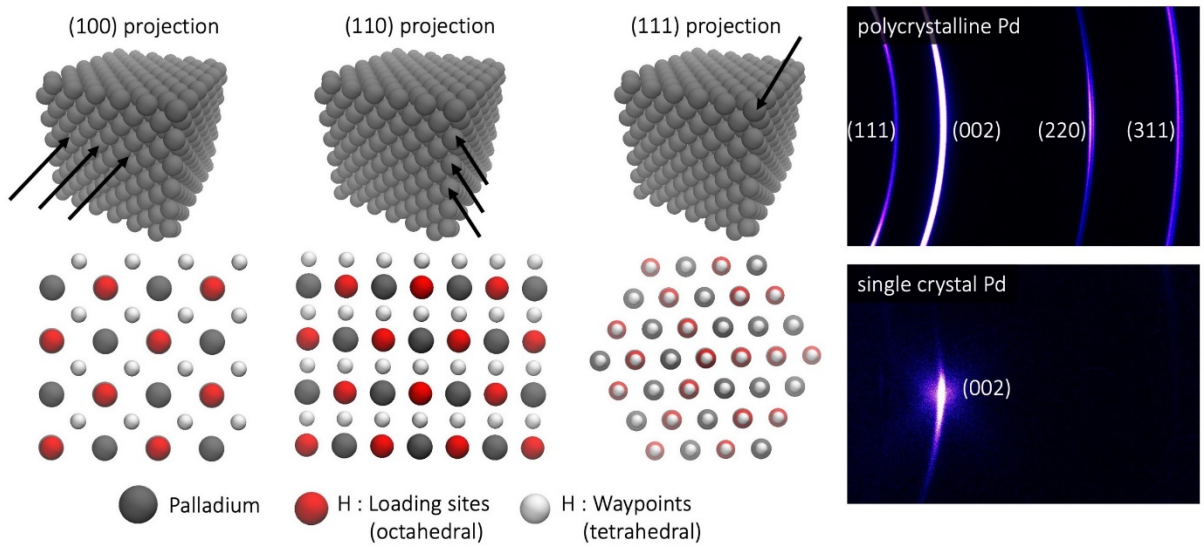


Fig. 4: Pd has an FCC crystal structure, and we show the (100), (110), and (111) orientations in addition to the octahedral and tetrahedral sites. Protons/deuterons load into octahedral sites. We confirm the Pd single crystal orientation using XRD.

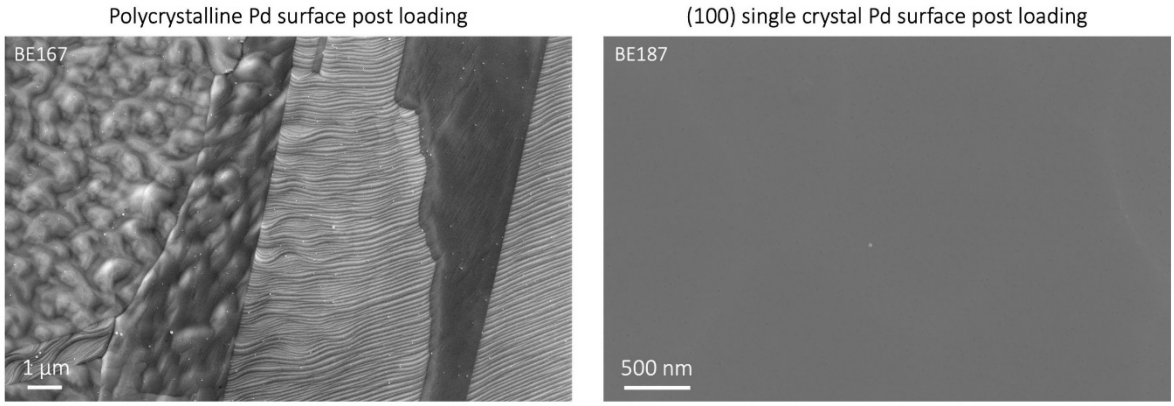


Fig. 5: Single crystal Pd surfaces remain pristine after elastic loading and deloading, successfully mitigating the distortions driven by grain boundaries in polycrystalline Pd.

Measuring Loading

To quantify loading values in metal hydrides, we developed two mass-loading measurement techniques (i.e., a direct measure of loading) and three indirect loading measurement techniques: UltraSorb & MicroSorb (direct measurements), and lateral strain, resistance, and reflectance (indirect measurements). For gas loading, our indirect loading measurements are calibrated as functions of gas pressure and temperature using our direct loading measurements. Once this is done, indirect loading measurements have proven to be very useful, both for convenience, and to monitor changes in loading. However, for determining absolute loading values, they are only as good as the underlying calibrations. So, for instance, indirect measurements based on calibrations

with Pd at room temperature, will not give accurate results for much different temperatures, or for Pd-alloys. In order to measure loading for general temperatures or materials, some form of direct loading measurement is essential.

Measuring Loading : Lateral Strain

A metal hydride swells as it loads protons/deuterons. Moreover, elastic loading of freestanding (i.e., unconstrained) samples ensures the strain is isotropic (i.e., the lateral and vertical strains are equivalent). With sufficient pixel resolution, the lateral strain (and thus loading) can be measured between fiducials on the sample surface using optical images.

Our objective is 1% loading resolution using lateral strain. For Pd, we estimate the volumetric strain Ω at full loading to be $\Omega = 0.17$; literature reports volumetric strains between 0.13-0.19[9]. We assume a linear relationship with loading. This is an approximation: strain is linear w.r.t to loading in α-phase and is linear w.r.t loading with a slightly different scaling in β-phase; however, using a single volumetric strain constant closely approximates the loading and is an accepted approximation in the technical literature. Note: we use a conservative volumetric strain constant (our measurements corroborate those literature which indicate it's closer to $\Omega = 0.16$).

If the sample is free standing (i.e., it's not constrained in any dimension), the strain is isotropic. Given this, we can approximate the loading (Γ) by $\Gamma = 3\varepsilon / \Omega$, where ε is the average measured lateral strain between tracked keypoints. Given our requirement of 1% loading resolution, we must be able to resolve strains of $\varepsilon_{min} = 0.000567$. By associating our strain resolution ε_{min} with our characteristic pixel resolution τ (i.e., $\tau = 1$ for a single pixel), we obtain the minimum required distance L_{min} between tracked keypoints by using the relation $\tau = L_{min} \times \varepsilon_{min}$. So, $L_{min} = 1765$ pixels at $\tau = 1$ and $\varepsilon_{min} = 0.000567$.

The required pixel size and field-of-view (FOV) depends on the sample geometry. As an example, a version of our ultra-pressure vessel allowed 4 mm diameter samples. To adequately accommodate the sample strain and tolerances on manufactured parts, the sample diameter was set to 3.65 mm, and the minimum distance between loading fiducials was set to 0.7 mm. This set the required pixel size at 0.4 μm/pixel. To provide an adequate buffer for the strain measurement, we preferred a FOV of at least 1.5 mm x 1.0 mm (**Fig. 6**).

[9] Peisl, H. "Lattice strains due to hydrogen in metals." *Hydrogen in metals I*. Springer Berlin Heidelberg, 1978. 53-74.

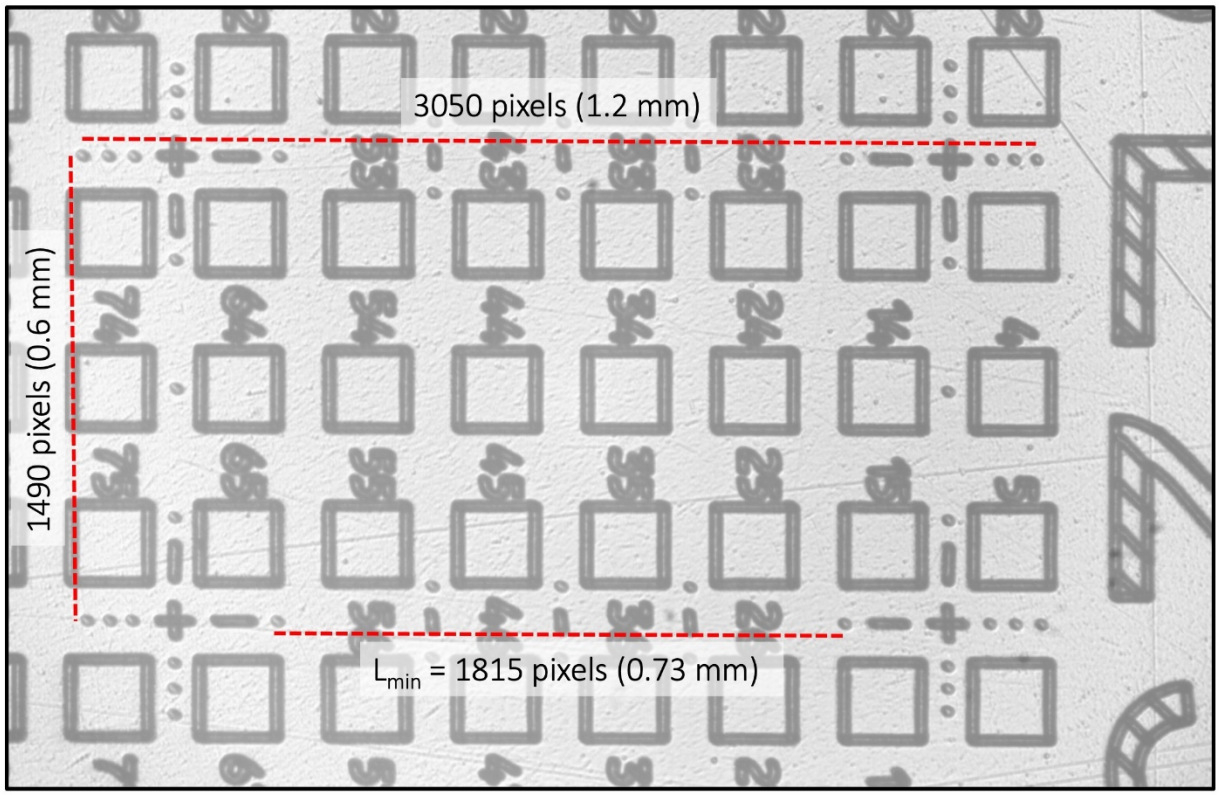


Fig. 6: Example image used for an in-situ strain measurement. The field-of-view is 3840x2484 pixels with 0.4 µm/pixel. The optical path contains three windows between the camera and the sample: camera → 0.5mm window → 0.5mm window → 7.5mm window → sample.

To enable lateral strain as a real-time loading diagnostic, we developed a computer vision software package using Python (**Fig. 7**). The measurement requires a source image (i.e., baseline image) of the unloaded Pd at a given pressure and temperature. Typically, the source image is set at 0 psi and the intended experiment temperature (often 25-50°C). At a set frame rate, target images are taken and compared to the source image. The sample surface has fiducials (typically laser-marked or lithographically patterned) that are detected in both the source and target images using the SIFT algorithm[10]. We match keypoints using a Brute Force Matcher[11] and KNN Matcher. We qualify the matches using a ratio test and RANSAC and filter the qualified matches by a minimum distance requirement (typically 700 pixels) and a minimum keypoints requirement (at least 4 valid matches are required, but we typically track and attempt to match 120 keypoints across each image). With a validated set of matched keypoints between the source and target images, *K,* we compute the Euclidian distance between each pair of keypoints in the source, $d_{src}(i,j)$, and target, $d_{dst}(i,j)$, images to measure the lateral strain:

[10] https://docs.opencv.org/4.x/d7/d60/classcv_1_1SIFT.html
[11] https://docs.opencv.org/4.x/d3/da1/classcv_1_1BFMatcher.html

$$\varepsilon = \frac{1}{N} \sum_{i,j \in K}^{N} \frac{d_{dst}(i,j) - d_{src}(i,j)}{d_{src}(i,j)}$$

The in-situ strain tracker is sensitive enough to measure both thermal strain and loading strain (**Fig. 7**).

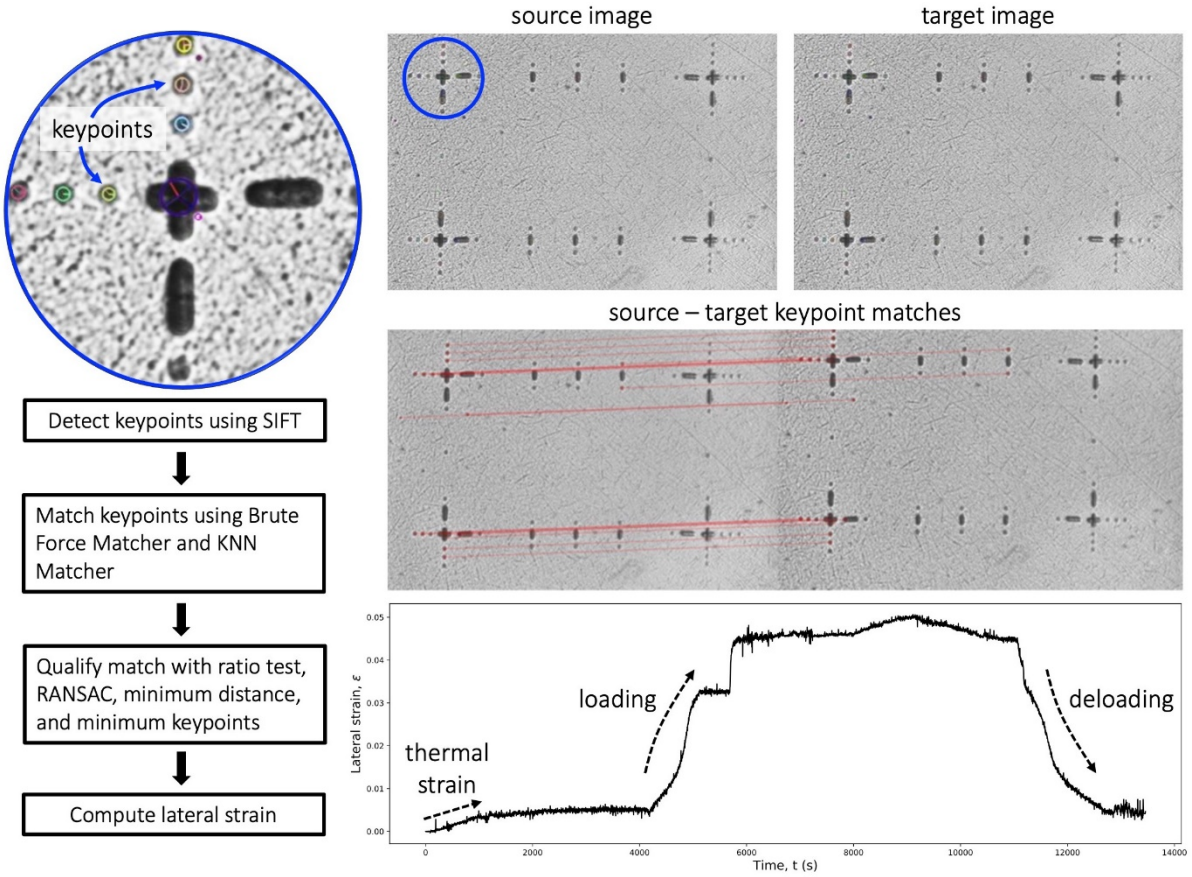


Fig. 7: To measure real-time lateral strain, keypoints are detected and matched between source and target images and strain is subsequently computed between the set of matched keypoints. Both thermal strain and loading strain are measured using this technique. Importantly, this technique relies on elastic loading/deloading to ensure the strain is isotropic.

A caveat to the lateral strain measurements is ensuring the sample surface remains in focus. When the focus drifts, it introduces error into the strain measurement. Both temperature and pressure changes require focus adjustments to avoid drift. As an example, we show a series of discontinuities in strain from focus drifts due to temperature changes (**Fig. 8**). When the focus is consistent between the source and target images, there is no strain error from drift. If the focus begins to drift, it can be re-adjusted, and the measurement will resume accurate strain measurements.

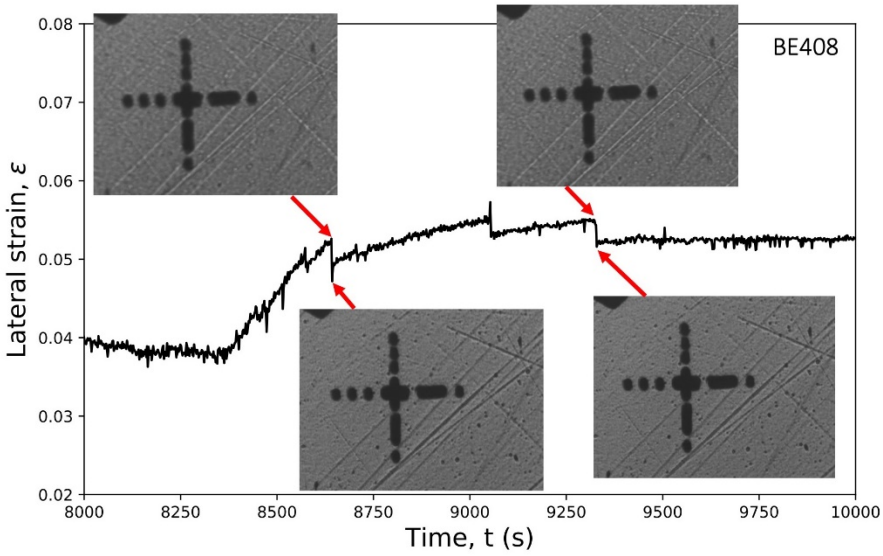


Fig. 8: Drifts in focus introduce error into the lateral strain measurement typically on the order of $\varepsilon_{err} = 0.002 - 0.003$. Full loading for Pd produces lateral strains of $\varepsilon = 0.057$ (assuming $\Delta V = 0.17$), indicating strain drifts of 0.002-0.003 leads to 3.5-5% loading errors.

We show an example strain measurement of a $Pd_{0.8}Rh_{0.2}$ alloy in **Fig. 9**. The lateral strain measures the thermal strain during the pre-heat to 350°C and subsequently measures the loading strain up to 5.7% at 100 ksi, 50°C. At a given pressure and temperature, the strain measurements are repeatable to < 0.0005 strain or < 1% loading (**Fig. 10**).

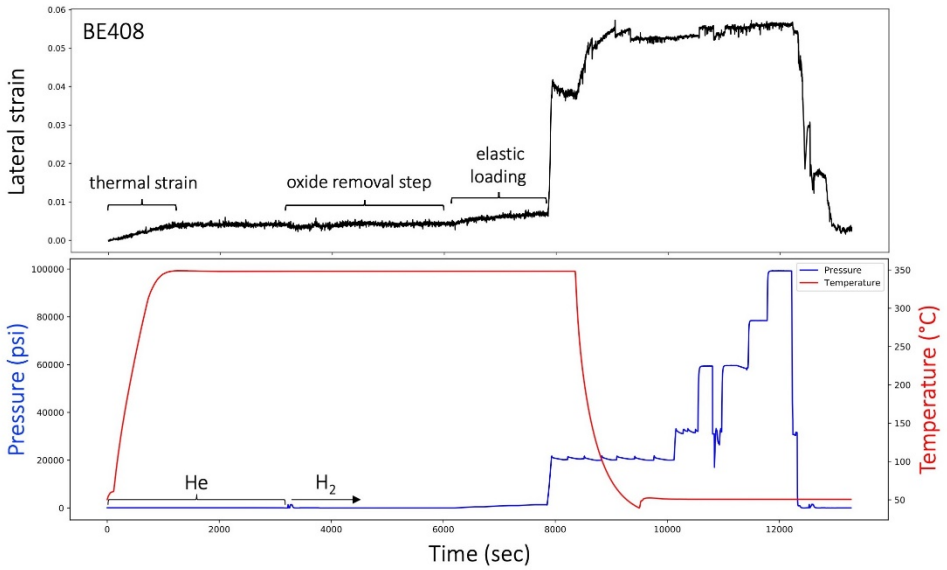


Fig. 9: A typical real-time strain measurement that senses both thermal and loading strains.

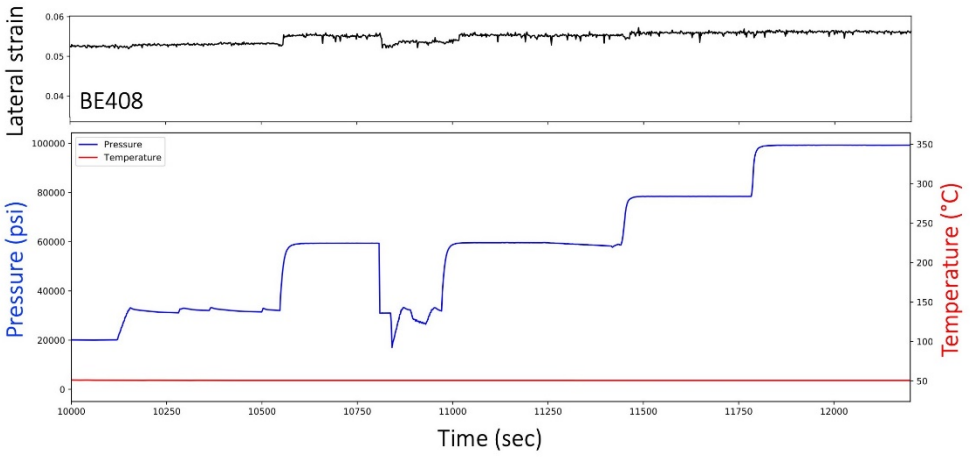


Fig. 10: Lateral strain measurements are repeatable to < 0.0005 strain (or < 1% loading) at a given pressure and temperature within a given experiment.

The strain measurement is highly reproducible sample-to-sample at a given pressure and temperature. We show several (P,T) points across three different samples in **Table 1**.

Experiment id	Pressure (psi)	Temperature (°C)	Lateral strain	Lateral strain std
BE34	3000	25	0.0442	0.000771
BE36	3000	25	0.0441	0.000306
BE50	3000	25	0.0443	0.000693
BE34	3000	350	0.0395	0.00105
BE36	3000	350	0.0393	0.000645
BE50	3000	350	0.0393	0.000533

Table 1: Lateral strain measurements are highly reproducible from sample-to-sample for a given pressure and temperature (i.e., to within 0.0002 strain or 0.35% loading).

Beyond tracking equilibrium loading, the lateral strain measurement is a useful diagnostic for kinetic behavior. For example, we wanted to confirm that cryogenic temperatures kinetically lock-in loading when out of thermodynamic equilibrium. To observe this, we loaded a Pd sample with 150 ksi, -100°C, D_2 gas and subsequently removed the D_2 gas and pulled high vacuum on the sample while remaining at -100°C. We observed no change in lateral strain (i.e., no deloading) for over 16 hours until the sample temperature was raised to 0°C, after which the sample began to deload (**Fig. 11**).

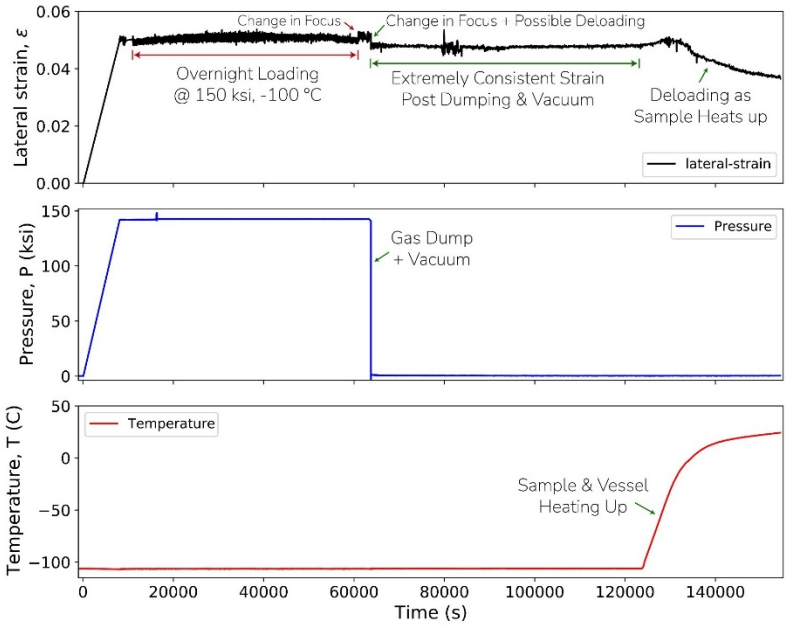


Fig. 11: The lateral strain measurement confirmed that cryogenic temperatures kinetically locks-in loading even when the gas is fully removed.

Measuring Loading : Resistive

The most common loading measurement approach in literature is the resistance ratio. The mapping between resistance ratio and loading is well documented[12] and our measurements generally agree with literature values (**Fig. 12**). However, most of these loading calibrations do not account for the effect pressure has on Pd/H(D). As we approach loadings nearing 1.0 where each 1% increase in loading requires ~100 ksi, the effect of pressure begins to make a difference. We find that the resistance of Pd decreases by ~2% from ambient pressure to 150 ksi, slightly changing our loading calibrations (~0.5% loading). We plot the resistance ratios for PdH_x and PdD_x wires at 25°C in **Fig. 13**. We've mainly used resistance ratio measurements with wire sample geometries.

[12] Zhang, Wu-Shou, Zhao-Fu Zhang, and Zhong-Liang Zhang. "Some problems on the resistance method in the in situ measurement of hydrogen content in palladium electrode." *Journal of Electroanalytical Chemistry* 528.1-2 (2002): 1-17.

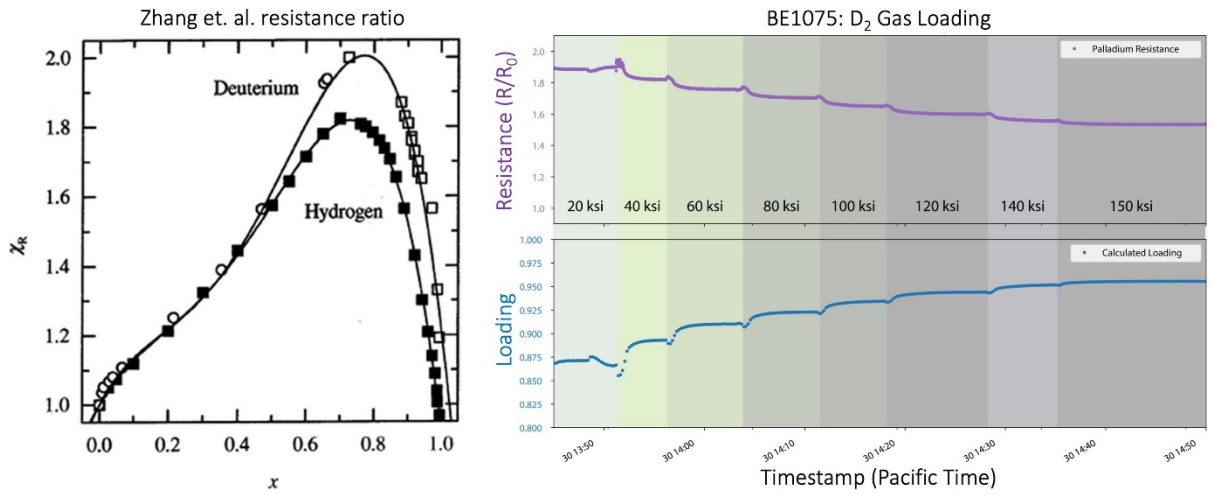


Fig. 12: The most accepted resistance ratio calibration curve from literature (Zhang) and a typical in-situ resistance ratio that we measured in a gas loading experiment with a Pd wire.

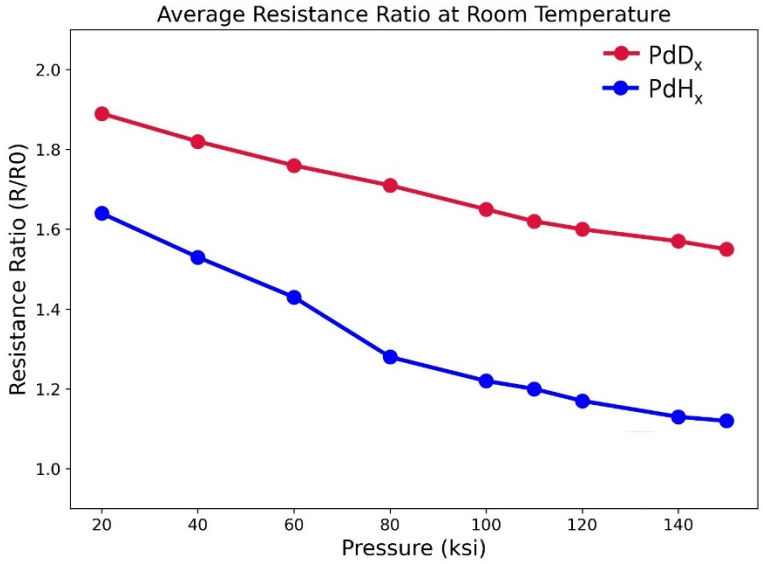


Fig. 13: The resistance ratios for PdH_x and PdD_x wires gas-loaded at 25°C up to 150 ksi.

Measuring Loading : Reflectance

As a Pd sample loads, its reflectance changes. We've found that 532nm wavelength is sensitive to loading: reflectivity changes by 15-20% when Pd loading changes from 0% to ~75% (**Fig. 14**). Importantly, live imaging allows continuous monitoring of non-loading factors that may alter reflectivity data such as beam position, focus, and sample surface. Due to the measurement's sensitivity to beam position, focus, and sample surface, reflectance measurements require thin-film sample geometries. We've attempted measurements in bulk samples, but the results were unreliable.

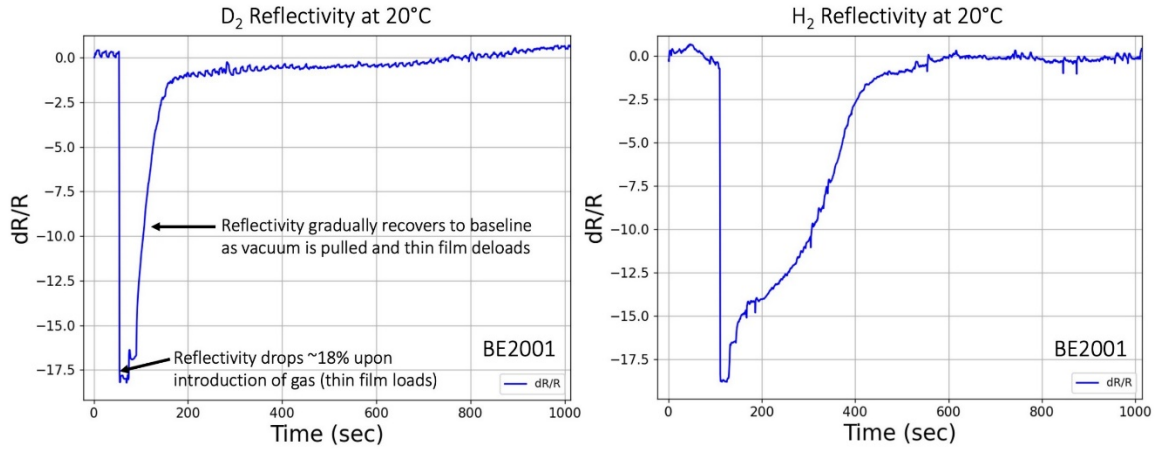


Fig. 14: Typical reflectivity data showing loading and deloading of thin-film Pd for H_2 and D_2 at 20°C. Both gases show loading times on the order of 1s (likely limited by gas flow rate of our apparatus, not intrinsic material properties). D_2 mostly deloads after 100-200s, while H_2 needs ~400s to deload to the same extent.

We've found reflectivity measurements to be a useful diagnostic to characterize deloading rates in thin films, but less useful as an absolute loading measurement. We found that deloading times increase ~1000x from 50°C to -20°C: from ~1 sec to >1000 sec (**Fig. 15**). At -80C, deloading is too slow to resolve at any measurable timescale (reflectivity did not measurably change after ~8 hours). From 20°C to 50°C, H_2 shows slower deloading times than D_2. Below 20°C, the data is less reproducible, but both isotopes show large increases in deloading times as temperature is reduced. Inconsistency could be related to inconsistent vacuum, mechanical shifts, or other factors.

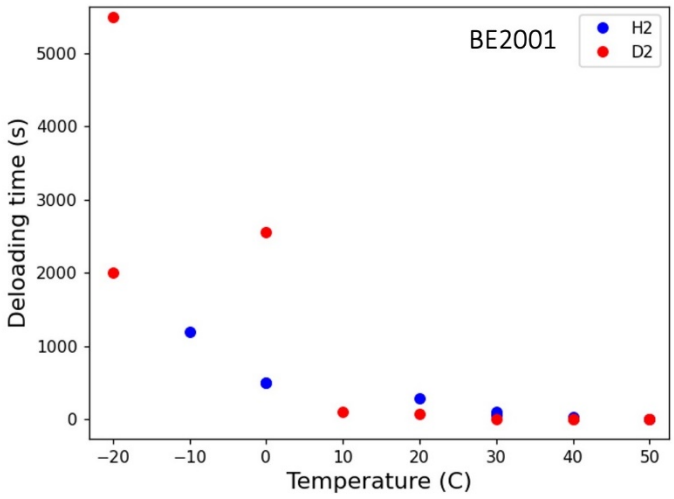


Fig. 15: Comparing D_2 and H_2 deloading kinetics of 20nm Pd thin films.

Beyond characterizing deloading kinetics, reflectivity measurements were a helpful diagnostic when developing purge sequences for our pressure systems. Reflectivity measurements showed that serial-dilution purge cycles at 50psi He were not effective; however, 100psi He cycles

completely deloaded the sample within a reasonable time i.e., we need high enough pressure to force gas through thin BeCu tubes on apparatus (**Fig. 16**). To ensure sufficient headroom, we use 5-7 purges at 1000psi to effectively exchange gas and deload the sample.

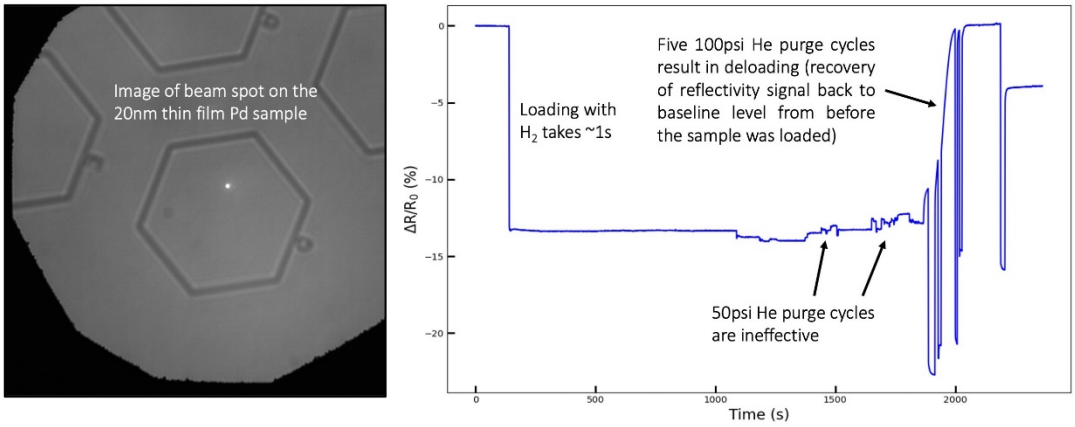


Fig. 16: Reflectivity measurements are used to test purge sequences with Protonpack and Chips panel.

Measuring Loading : Mass (UltraSorb)

To definitively measure mass loadings above 0.85 (i.e., a direct loading measurement with no calibration required), we developed a form of a Sieverts apparatus capable of handling pressures up to 10 kbar called the UltraSorb. The measurement involves a series of inert gas expansions between a pressurization-cell and an expansion-cell. These expansions enable precise volume calibration of the pressurization and expansion cells (**Fig. 17**).

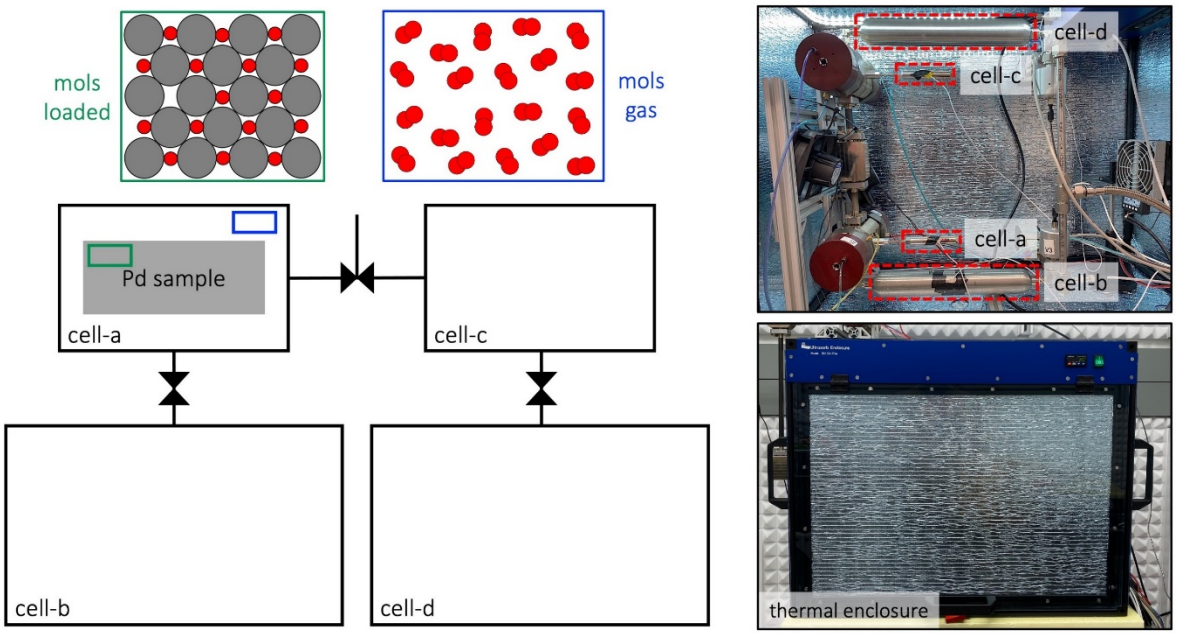


Fig. 17: The UltraSorb apparatus contains a pressurization-cell (cell-a) and an expansion-cell (cell-b) in addition to a gas density calibrator (cell-c & cell-d).

To measure loading at a given hydrogen pressure and temperature, a sample (Pd, Pd-alloy, etc.) is placed in the calibrated pressurization-cell and pressurized with hydrogen. Some hydrogen loads into the sample (mols-loaded) and the remaining hydrogen resides in the free volume around the sample (mols-gas); this sets the initial mols of hydrogen. The key challenge is differentiating the mols-loaded from the mols-gas. The hydrogen is then expanded into the calibrated expansion-cell. We know the expected mols of hydrogen from the free volume in the pressurization-cell, but the expansion-cell also measures excess mols of gas which come from the hydrogen loaded into the sample. Any residual mols of hydrogen in the sample (that did not deload during the initial expansion) can be measured as well by heating the sample and expanding into the calibrated expansion-cell. This overall excess is a direct loading measurement at the pressure and temperature in the pressurization-cell (**Fig. 18**).

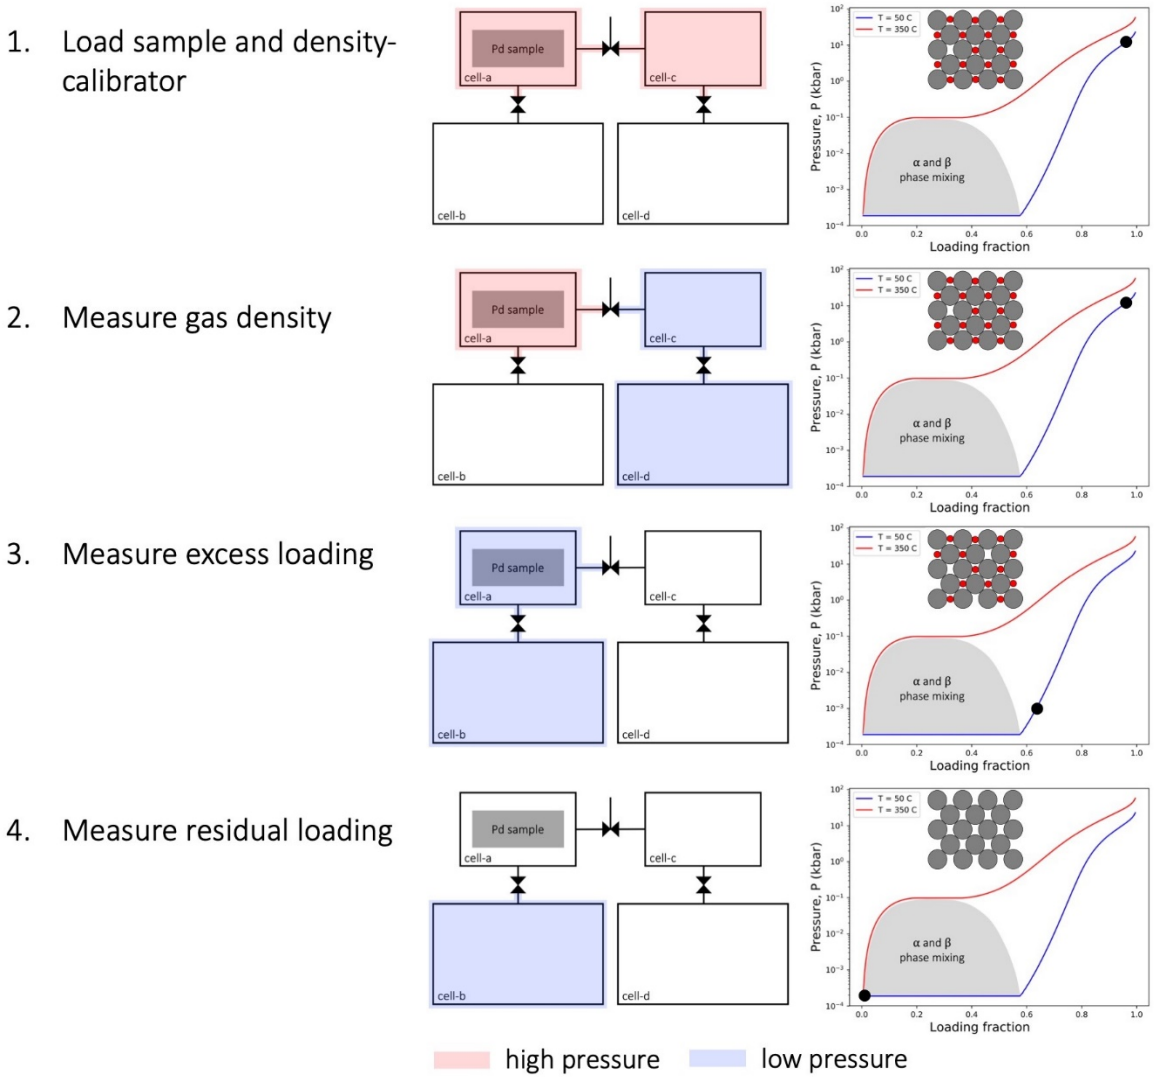


Fig. 18: The UltraSorb loading measurement requires calibrated cell volumes. The gas density is measured at the same pressure and temperature as the loaded sample. A high-pressure expansion measures the excess loading above 1 bar and subsequent low-pressure expansions measure the residual loading.

To track the mols of hydrogen in each cell, we must know the cell volume and gas density. Both cell volume and gas density require highly accurate pressure and temperature measurements as well as the gas's equation-of-state (EOS). Hydrogen's EOS is well known at low pressures, but not as well-known at ultra-pressures (e.g., 10 kbar); moreover, pressure transducers can have poor accuracy at ultra-pressures. To bypass the need to accurately measure ultra-pressures, we've designed a density calibrator into our apparatus, which directly measures the gas density at a particular pressure and temperature.

For the hydrogen and deuterium EOS, our codebase used the CoolProp library[13], which claims to reference the NIST Thermophysical Properties of Fluid Systems database[14]. Both CoolProp and NIST reference the Richardson et al. EOS[15] for deuterium. We confirmed the CoolProp library agrees with the NIST database (**Table 2**).

Pressure (psi)	Temperature (°C)	NIST – ρ (g/cm^3)	CoolProp - ρ (g/cm^3)
1000	25	0.010770	0.0107698
10000	25	0.078133	0.0781332
40000	25	0.17010	0.1701044
1000	100	0.0086552	0.0086551
10000	100	0.066204	0.0662041
50000	100	0.17226	0.1722616
1000	200	0.0068675	0.0068674
10000	200	0.055144	0.0551438
60000	200	0.17060	0.1705980

Table 2: The CoolProp library agrees well with the NIST database for the deuterium EOS.

Each loading point at a given pressure and temperature requires five different measurements: (1) expansion-cell volume calibration, (2) pressurization-cell volume calibration, (3) high pressure density calibration, (4) high pressure excess loading, and (5) low pressure residual loading. We define the nomenclature used in the measurements in **Table 3**.

UltraSorb Nomenclature	
V_p	Pressurization-cell volume
$V_{g(p)}$	Free gas volume in the pressurization-cell

[13] http://www.coolprop.org/index.html
[14] https://webbook.nist.gov/chemistry/fluid/
[15] Richardson, I.A.; Leachman, J.W.; Lemmon, E.W., "Fundamental Equation of State for Deuterium", J. Phys. Chem. Ref. Data, 2014, 43, 1, 013103.

V_e	Expansion-cell volume
$V_{g(e)}$	Free gas volume in the expansion-cell
V_g	Total free gas volume
V_{cal}	Calibration sample's volume (e.g., Au, Pt, etc.)
ρ_i	Gas density in state-*i*
V_s	The unloaded sample's volume (e.g., Pd)
e_L	The sample's strain due to loading
e_T	The sample's strain due to thermal expansion
n_g	Mols of gas
M_g	Molar mass of the gas
A	The number of atoms per gas molecule
n_{ex}	Excess mols of gas (molecular) deloaded from the sample in an expansion
$n_{ex(at)}$	Excess mols (atomic) deloaded from the sample in an expansion
$n_{tot(at)}$	Total excess mols (atomic) deloaded from the sample across all expansions
m_s	Unloaded sample mass (e.g., Pd)
M_s	Molar mass of the unloaded sample (e.g., Pd)
n_s	Mols of atoms in the unloaded sample (e.g., Pd)
Γ	The measured loading fraction

Table 3: Nomenclature used in the UltraSorb calculations.

Expansion-cell volume calibration measurement

The expansion-cell volume calibration requires the apparatus to have a small, high-pressure volume (the pressurization-cell) that expands into a much larger, low-pressure expansion volume (the expansion-cell) as well as a high purity noble metal sample (e.g., Au or Pt) that serves as a calibration volume (**Fig. 19**).

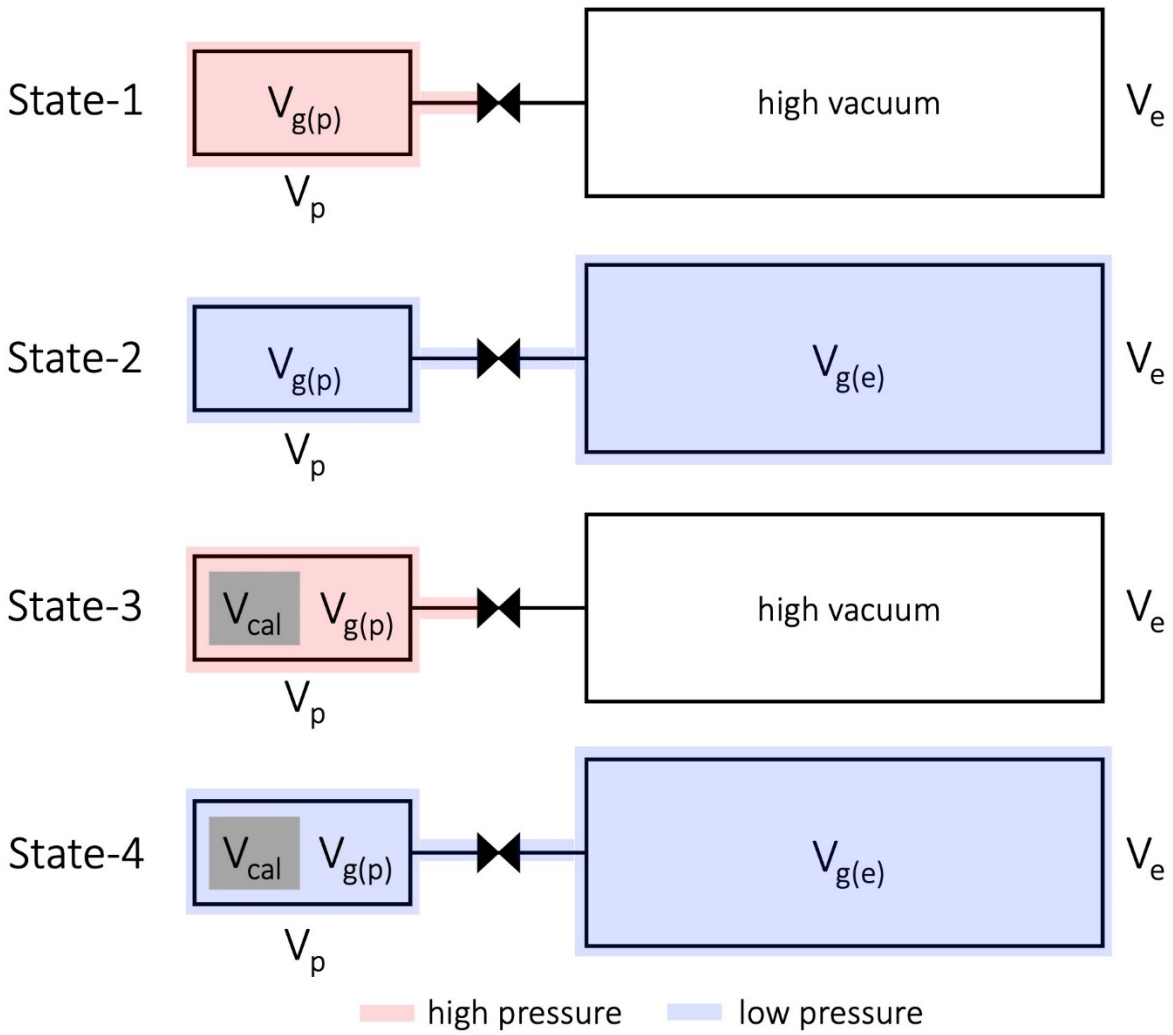


Fig. 19: Expansion-cell volume calibration measurement.

The measurement involves two expansions. The first expansion starts with empty pressurization and expansion cells (i.e., no calibration mass is included in the expansion) and involves two states:

- State-1: add inert gas to the pressurization-cell, while isolated from the expansion-cell at high vacuum.
- State-2: expand the inert gas in the pressurization-cell into the expansion-cell.

The second expansion add the calibration mass to the pressurization-cell and involves two states:

- State-3: add inert gas to the pressurization-cell, while isolated from the expansion cell at high vacuum.
- State-4: expand the inert gas in the pressurization-cell into the expansion-cell.

The expansion-cell volume calibration measurement has the following assumptions:

1. Zero leak requirement (i.e., the mols of inert gas are conserved between states).

2. The vessel strain of the pressurization-cell is negligible at low pressures (the measurements are done at lower pressures to take advantage of high precision and accuracy pressure transducers).
3. The calibration sample is high purity (i.e., > 4N); the sample's volume is the standard used to calibrate the cell volumes; we measure the sample mass to < 1 μg accuracy with a Mettler Toledo Microbalance, but the density is taken from NIST tables and assumes a high-purity sample.
4. The pressurization-cell volume remains constant between the two expansions (i.e., the vessel assembly with and without the calibration sample must be consistent).

Now, we derive the expansion-cell volume expression. Expanding state-1 into state-2 and using mass conservation, we get:

$$\rho_1 \times V_{g(p)} = \rho_2 \times \left(V_{g(p)} + V_{g(e)}\right)$$

Now, the free gas volumes consume the entire cell-volumes since no calibration sample is present: $V_{g(p)} = V_p$ and $V_{g(e)} = V_e$. Substituting the free gas volumes, we get:

$$\rho_1 \times V_p = \rho_2 \times \left(V_p + V_e\right)$$

Expanding state-3 into state-4 and using mass conservation, we get:

$$\rho_3 \times V_{g(p)} = \rho_4 \times \left(V_{g(p)} + V_{g(e)}\right)$$

Now, the free gas volume consumes the entire expansion-cell (i.e., it's empty), but only a subset of the pressurization-cell, which contains the calibration sample: $V_{g(p)} = V_p - V_{cal}$ and $V_{g(e)} = V_e$. Substituting the free gas volumes, we get:

$$\rho_3 \times \left(V_p - V_{cal}\right) = \rho_4 \times \left(V_p - V_{cal} + V_e\right)$$

Combining the state-1/state-2 and state-3/state-4 expansions, we get:

$$V_e = V_{cal} \times \frac{(\rho_1 - \rho_2)(\rho_3 - \rho_4)}{\rho_2 \times \rho_3 - \rho_1 \times \rho_4}$$

Pressurization-cell volume calibration measurement

The pressurization-cell volume measurement requires a known expansion-cell volume. This assumes an apparatus with a small, high-pressure volume (the pressurization-cell) that expands

into a larger, low-pressure expansion volume (the expansion-cell). The apparatus may or may not contain a non-loading sample volume that resides in the pressurization-cell (**Fig. 20**).

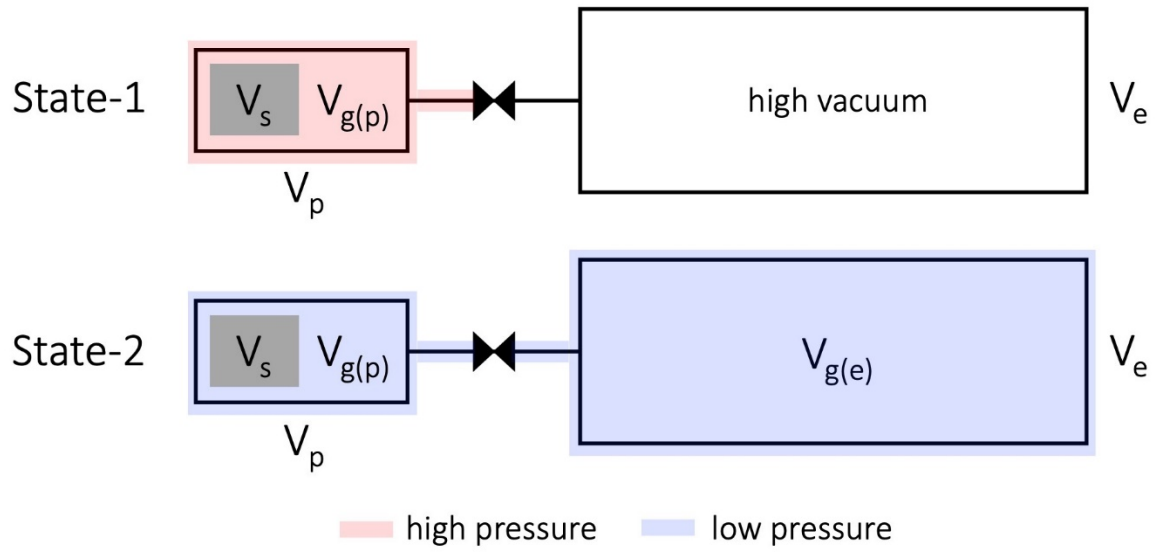


Fig. 20: The pressurization-cell volume calibration measurement.

The measurement has a single expansion:

- State-1: add inert gas to the pressurization-cell (at the target loading measurement pressure to ensure vessel-strain is accounted for), while isolated from the expansion-cell at high vacuum.
- State-2: expand the inert gas in the pressurization-cell into the expansion-cell.

The pressurization-cell volume calibration measurement has the following assumptions:

1. Zero leak requirement (i.e., the mols of inert gas are conserved between states).
2. The expansion-cell volume is known, i.e., it's been calibrated.
3. The sample is high purity such that the density taken from NIST tables is accurate. We measure the sample's mass to < 1 µg accuracy with a Mettler Toledo Microbalance.

Now, we derive the pressurization-cell volume expression. Expanding state-1 into state-2 and using mass conservation, we get:

$$\rho_1 \times V_{g(p)} = \rho_2 \times \left(V_{g(p)} + V_{g(e)}\right)$$

The free gas volume consumes the entire expansion-cell (i.e., it's empty) and a subset of the pressurization-cell if a sample is involved (the pressurization-cell may or may not contain the sample; so, we include V_s and set to it 0 if no sample is included): $V_{g(p)} = V_p - V_s$ and $V_{g(e)} = V_e$. Substituting for the free gas volumes, we get:

$$\rho_1 \times \left(V_p - V_s\right) = \rho_2 \times \left(V_p - V_s + V_e\right)$$

Solving for V_p we get:

$$V_p = V_s + V_e \times \frac{\rho_2}{\rho_1 - \rho_2}$$

High pressure density calibration measurement

At high pressures, pressure transducers are often inaccurate, and hydrogen's EOS is not well-known to our desired precision; so, we use our apparatus to directly measure gas density at high pressures. This measurement assumes an apparatus with a small, high-pressure volume (the pressurization-cell) that expands into a larger, low-pressure expansion volume (the expansion-cell). The apparatus may or may not contain a sample volume that resides in the pressurization-cell. Importantly, if the pressurization-cell contains a sample, it must not load hydrogen, e.g., it can be Au, Pt, etc. (**Fig. 21**).

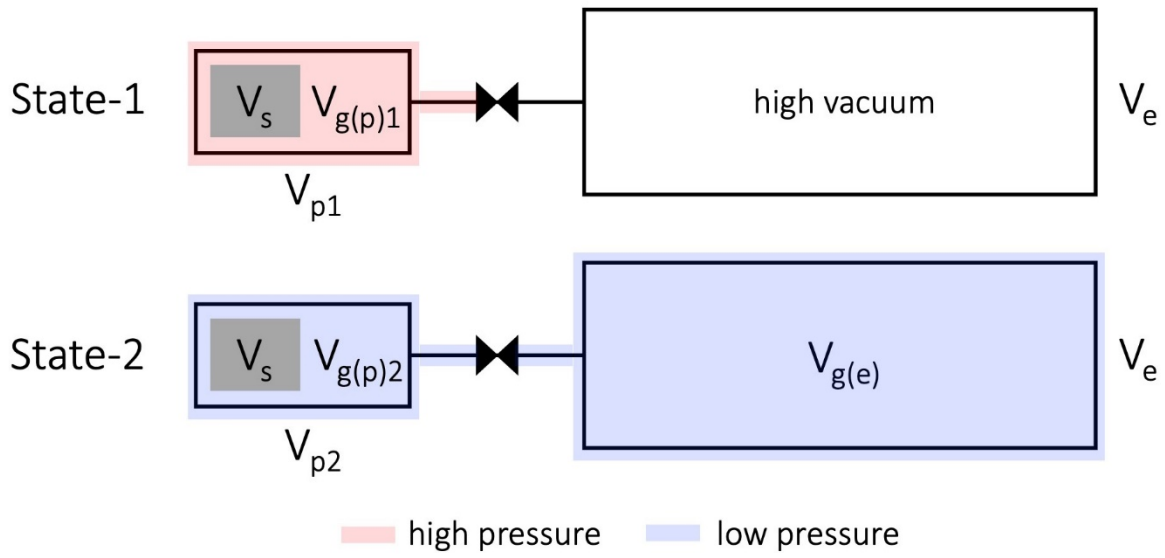


Fig. 21: The high-pressure, gas density calibration measurement.

The measurement has a single expansion:

- State-1: add the measurement gas (e.g., hydrogen) to the pressurization-cell (at matched pressure and temperature to the target loading measurement), while isolated from the expansion-cell at high vacuum.
- State-2: expand the gas in the pressurization-cell into the expansion-cell.

The high-pressure density calibration measurement has the following assumptions:

1. Zero leak requirement (i.e., the mols of gas are conserved between states).

2. The pressurization-cell volume is known, i.e., it's been calibrated.
3. The expansion-cell volume is known, i.e., it's been calibrated.
4. If an inert sample resides in the pressurization-cell, it's high purity such that the density taken from NIST tables is accurate. We measure the sample's mass to < 1 μg accuracy with a Mettler Toledo Microbalance.

Now, we derive the gas density expression. Expanding state-1 into state-2 and using mass conservation, we get:

$$\rho_1 \times V_{g(p)1} = \rho_2 \times \left(V_{g(p)2} + V_{g(e)}\right)$$

The pressurization-cell strains 1-2% at high pressures, so we must use a V_p that accounts for the vessel strain in state-1 to avoid a significant source of error. Since $V_e >> V_p$ (i.e., V_e is ~1000x larger than V_p), the expanded volume in state-1 ($V_{g(p)2} + V_{g(e)}$) is nearly all contained in V_e. So, we do not need to account for differences in the pressurization-cell vessel strain between the two states i.e., in state-2, $V_{g2} = V_{p2} - V_s + V_e \approx V_{p1} - V_s + V_e$. This calculation differentiates between V_{p1} and V_{p2}; however, it's acceptable to set $V_{p1} \approx V_{p2}$ since the error it introduces is well below the measurement resolution.

The free gas volume consumes the entire expansion-cell (i.e., it's empty) and a subset of the pressurization-cell: $V_{g(p)1} = V_{p1} - V_s$, $V_{g(p)2} = V_{p2} - V_s$, and $V_{g(e)} = V_e$. Substituting for the free gas volumes, we get:

$$\rho_1 \times \left(V_{p1} - V_s\right) = \rho_2 \times \left(V_{p2} - V_s + V_e\right)$$

Solving for the state-1 (i.e., high pressure) gas density, we get:

$$\rho_1 = \rho_2 \times \frac{V_{p2} - V_s + V_e}{V_{p1} - V_s}$$

<u>High pressure excess loading measurement</u>

The excess mols loaded into a sample at high pressures are measured with a single expansion. The measurement assumes an apparatus with a small, high-pressure volume (the pressurization-cell) that contains the sample of interest and a larger, low-pressure volume (the expansion-cell) (**Fig. 22**).

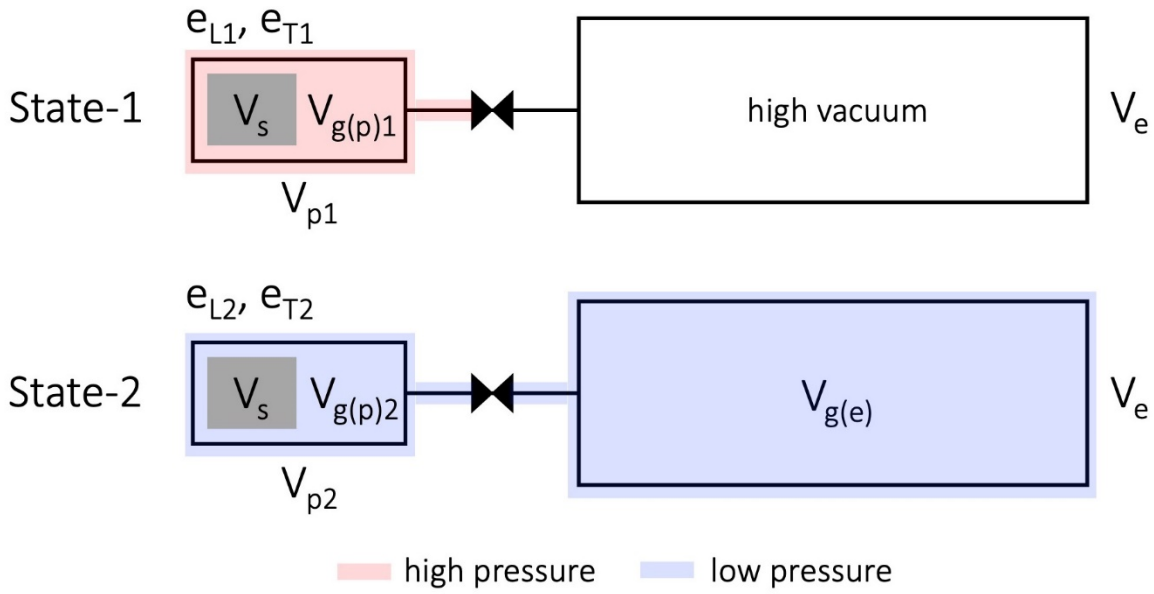


Fig. 22: The high-pressure, excess loading measurement.

The expansion consists of the following:

- State-1: add process gas (e.g., hydrogen) to the pressurization-cell that contains the sample of interest, which is isolated from the expansion-cell at high vacuum.
- State-2: expand the gas in the pressurization-cell into the expansion-cell and wait until the pressure plateaus in the expansion-cell.

The high-pressure excess loading measurement has the following assumptions:

1. Zero leak requirement i.e., the mols of gas are conserved between states.
2. The pressurization-cell volume is known, i.e., it's been calibrated.
3. The expansion-cell volume is known, i.e., it's been calibrated.
4. The gas density in state-1 is known, i.e., it's been calibrated; experimentally, the gas density calibration cell is simultaneously pressurized to the loading measurement cell to ensure the pressure is matched and the cells are held at the same temperature; this ensures the gas density is the same between the two cells.
5. The sample's lattice strain from loading is known, i.e., it's been calibrated.
6. The sample's thermal strain is known.
7. The sample is high purity such that the density taken from NIST tables is accurate. We measure the sample's mass to < 1 μg accuracy with a Mettler Toledo Microbalance.

Now, we derive the excess mols of atoms loaded into the sample. The free gas volume in state-2 is given by:

$$V_{g1} = V_{p1} - V_s \times (1 + e_{L1} + e_{T1})$$

The sample's lattice strain from loading, e_L, is calibrated from our lateral strain measurement technique (**Fig. 23**). We've found that the volumetric strain for PdH_x / PdD_x roughly scales as $e_L \approx 0.17 \times \Gamma$, where Γ is the loading fraction. For new materials, loading strain must be calibrated as a function of pressure and temperature as small changes in free gas volume in the pressurization-cell lead to large differences in measured loading. The sample's volumetric thermal strain, e_T, is approximated from the sample's coefficient of thermal expansion (CTE) taken from the NIST tables: $e_T \approx 3 \times CTE$. Thermal strain is only relevant for high-temperature measurements (e.g., >100 °C). It's important to note that the CTE will change as a function of loading, but it's a small effect in most materials (i.e., the error introduced is well below the resolution of the loading measurement).

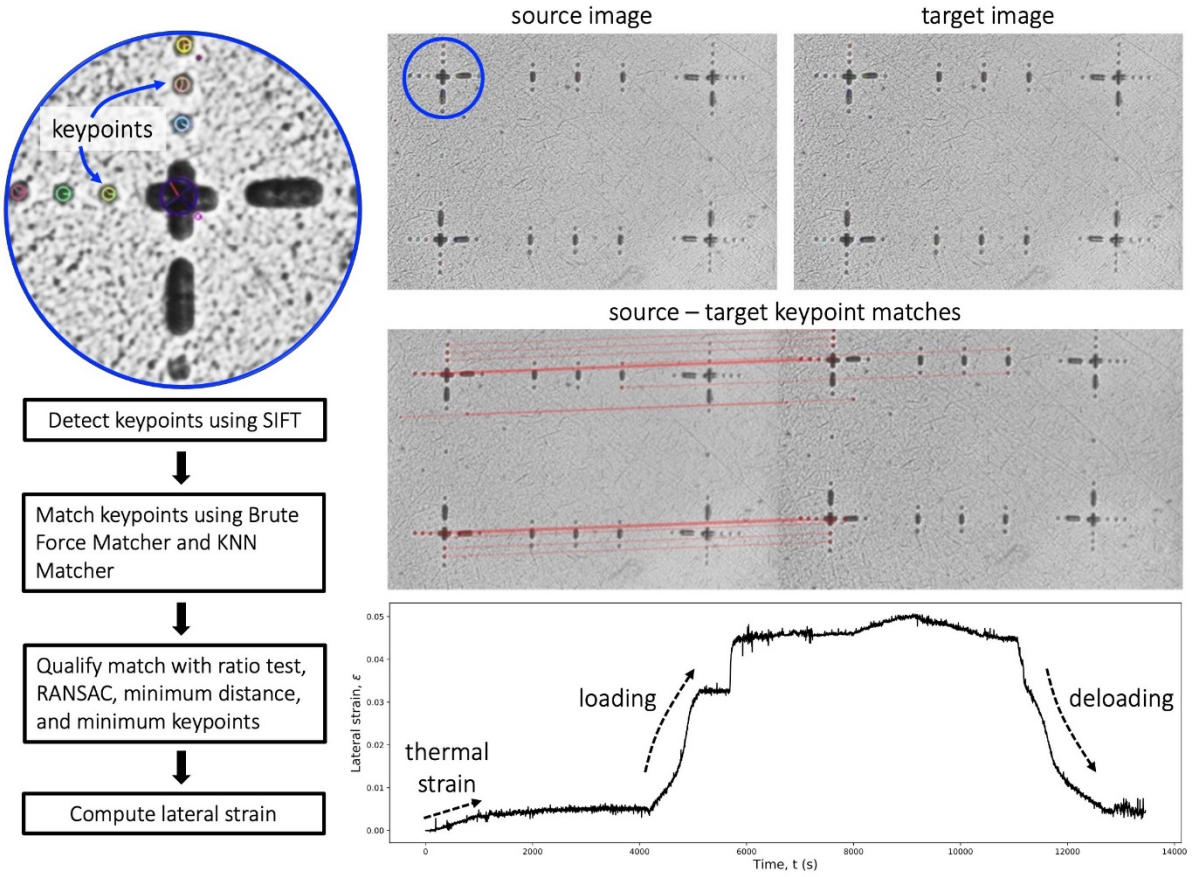


Fig. 23: Our in-situ lateral strain measurement calibrates the volumetric strain of a hydride as a function of pressure and temperature.

The mols of gas is given by $n = \rho V / M$. Substituting, we get:

$$n_{g1} = \frac{\rho_1 \times \left(V_{p1} - V_s \times (1 + e_{L1} + e_{T1})\right)}{M_g}$$

Similarly, we can write down the measured mols of gas in state-2 (after the gas from state-1 is expanded into state-2):

$$V_{g2} = V_{p2} - V_s \times (1 + e_{L2} + e_{T2}) + V_e$$

$$n_{g2} = \frac{\rho_2 \times V_{g2}}{M_g}$$

$$n_{g2} = \frac{\rho_2 \times \left(V_{p2} - V_s \times (1 + e_{L2} + e_{T2}) + V_e\right)}{M_g}$$

Now, the measured mols of gas in state-2 contains the mols of gas in state-1 (assuming mass conservation) as well as the excess mols of gas that were loaded in the sample in state-1. So, $n_{g2} = n_{g1} + n_{ex1}$. Solving for the excess mols of gas loaded in the sample in state-1, we get $n_{ex1} = n_{g2} - n_{g1}$. To get the excess mols of atoms loaded (versus mols of gas molecules), we multiply by the number of atoms per gas molecule: $n_{ex(at)} = A \times n_{ex}$. Substituting for n_{g1} and n_{g2}, we get:

$$n_{ex(at)} = A \times \left[\frac{\rho_2 \times \left(V_{p2} - V_s \times (1 + e_{L2} + e_{T2}) + V_e\right)}{M_g} - \frac{\rho_1 \times \left(V_{p1} - V_s \times (1 + e_{L1} + e_{T1})\right)}{M_g}\right]$$

Low pressure residual loading measurement

After expanding the high-pressure state and measuring the excess mols, the sample still has residual loading. For example, a high-pressure expansion generally reduces the pressure from 10 kbar to 1 bar. At 1 bar, 25 °C, PdH_x is still ~70% loaded. To measure the residual loading, the sample is heated (typically to ~200 °C) to deload the sample and is subsequently expanded into the expansion-cell, which is vented between each expansion and brought to high vacuum. The excess mols deloaded from the sample are calculated using the expression for $n_{ex(at)}$ established above. At these low pressures, hydrogen's EOS is accurate, so density-calibration measurements are not required. Furthermore, at low pressures, the pressurization-cell's vessel strain is not relevant. Expansions are continued until there is no observable increase in pressure in the expansion-cell i.e., the sample has fully deloaded. In typical measurement, 3 residual loading measurements are required to fully deload the sample. In summary, the total excess mols (atomic) is expressed by:

$$n_{tot(at)} = \sum_{i=1}^{N} n_{ex(at)i}$$

where N is the number of expansion cycles.

Loading calculation

To calculate the loading of the sample at the initial pressure and temperature, the total excess mols (atomic) deloaded from the sample is compared to the mols of lattice atoms: $\Gamma = n_{tot(at)}/n_s$. The mols of lattice atoms is given by:

$$n_s = \frac{m_s}{M_s}$$

So, the loading fraction can be expressed by:

$$\Gamma = \frac{M_s}{m_s} \times \sum_{i=1}^{N} n_{ex(at)i}$$

Example measurement

As an example measurement, we review BE2360, where we measured the PdD_x loading at 135 ksi, 50°C to be 92.6%. We used a 0.17757 cm^3 Pt sample and N_2 gas to calibrate the expansion-cell volumes and found them to be 501.65 ± 3.59 cm^3 (cell-b) and 493.67 ± 2.20 cm^3 (cell-d) (**Table 4**, **5**, **6**). Using the expansion-cell volumes, we calibrated the pressurization-cell volumes with N_2 gas and found them to be 0.51574 ± 0.002633 cm^3 (cell-a) and 0.49806 ± 0.000552 cm^3 (cell-c) (**Table 5, 6**). We simultaneously pressurized cell-a and cell-c with ~135 ksi, 50°C D_2 and measured the D_2 gas density to be 0.2337 g/cm^3 by expanding cell-c into cell-d (**Table 5, 7**). Finally, we expanded cell-a (containing the loaded PdD_x sample and ~135 ksi, 50°C of D_2) into cell-b to measure the excess loading above ~1 bar. We heated the PdD_x sample to ~200 °C and did three more expansions to measure the residual loading. It's important to note that we wait sufficient time for the sample to deload (i.e., when the pressure in the expansion-cell plateaus) for each expansion to ensure we measure the equilibrium value; depending on the sample thickness, each expansion can require 1-5 hours to equilibrate (**Fig. 24**). The sum of the loadings totaled 92.6% (**Table 5**, **8**).

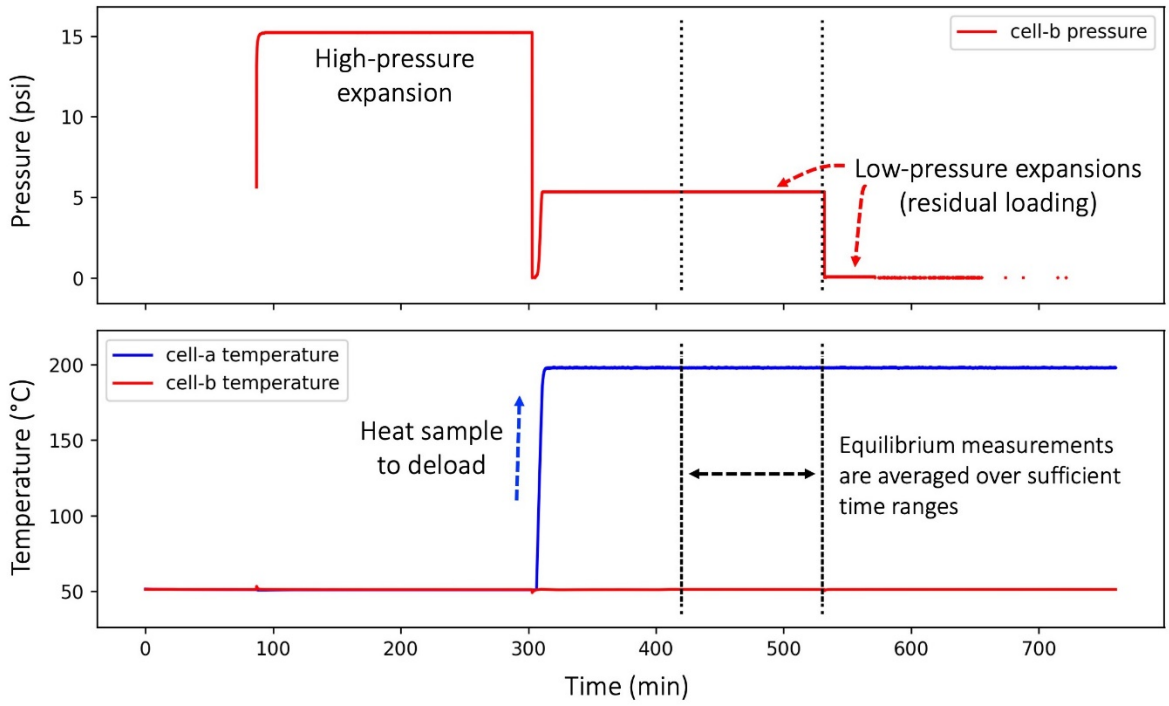


Fig. 24: A loading measurement includes a high-pressure expansion followed by low-pressure expansions while heating to fully deload the sample. Each expansion waits sufficient time for the sample to reach equilibrium and averages over a time range.

Sample id	Material	Mass (g)	Volume (cm^3)
BD1913	Pt	3.8089303	0.17757
BD3917	Pd	2.3122827	0.19237

Table 4: The samples used are a Pt sample for the expansion-cell volume calibrations and the Pd sample for the loading measurement.

Expansion id	Gas	Sample id	Pressurization-cell					Expansion-cell				
			Name	P (psi)	P std (psi)	T (°C)	T std (°C)	Name	P (psi)	P std (psi)	T (°C)	T std (°C)
BE2136-EX1	N_2	Empty	cell-a	1452.08	5.16	50.47	0.00148	cell-b	1.329	0.0566	50.10	0.00609
BE2136-EX2	N_2	BD1913	cell-c	1452.08	5.16	50.31	0.000963	cell-d	0.815	0.0611	50.62	0.00562
BE2136-EX3	N_2	Empty	cell-a	1460.95	5.11	50.43	0.00177	cell-b	1.338	0.0578	50.12	0.00608
BE2136-EX4	N_2	BD1913	cell-c	1460.95	5.11	50.25	0.000657	cell-d	0.819	0.0622	50.59	0.00425
BE2136-EX5	N_2	Empty	cell-a	1469.73	4.61	50.42	0.00272	cell-b	1.351	0.0573	50.12	0.006574
BE2136-EX6	N_2	BD1913	cell-c	1469.73	4.61	50.24	0.000824	cell-d	0.828	0.0619	50.56	0.00568
BE2136-EX7	N_2	BD1913	cell-a	1446.08	4.01	50.52	0.00139	cell-b	0.822	0.0572	50.19	0.00952
BE2136-EX8	N_2	Empty	cell-c	1446.08	4.01	50.52	0.000766	cell-d	1.321	0.0617	50.50	0.00995
BE2136-EX9	N_2	BD1913	cell-a	1467.93	4.58	50.42	0.00156	cell-b	0.839	0.0572	50.22	0.00890
BE2136-EX10	N_2	Empty	cell-c	1467.93	4.58	50.46	0.00148	cell-d	1.343	0.0619	50.49	0.00339
BE2136-EX11	N_2	BD1913	cell-a	1468.90	4.89	50.44	0.00178	cell-b	0.837	0.0557	50.23	0.00597
BE2136-EX12	N_2	Empty	cell-c	1468.90	4.89	50.45	0.000557	cell-d	1.342	0.0600	50.47	0.00265
BE2360-EX1	N_2	BD3917	cell-a	138613.8	172.17	51.91	0.00915	cell-b	8.497	0.00358	51.80	0.00237

BE2360-EX2	N_2	BD1913	cell-c	138613.8	172.17	51.71	0.00994	cell-d	8.509	0.00384	50.75	0.0173
BE2360-EX3	N_2	BD3917	cell-a	138530.1	70.46	51.89	0.00379	cell-b	8.506	0.00364	51.79	0.00267
BE2360-EX4	N_2	BD1913	cell-c	138048.6	15.66	51.28	0.00386	cell-d	8.484	0.00356	50.63	0.00399
BE2360-EX5	N_2	BD3917	cell-a	138046.9	32.84	51.88	0.00724	cell-b	8.391	0.00437	51.82	0.00599
BE2360-EX6	N_2	BD1913	cell-c	138046.9	32.84	51.09	0.0131	cell-d	8.492	0.00431	50.63	0.00335
BE2360-EX7	D_2	BD1913	cell-c	-	-	50.19	0.00529	cell-d	14.695	0.00436	50.35	0.00732
BE2360-EX8	D_2	BD3917	cell-a	-	-	51.78	0.00141	cell-b	15.765	0.00421	51.93	0.0113
BE2360-EX9	D_2	BD3917	cell-a	-	-	197.25	0.309	cell-b	5.347	0.00395	51.51	0.00231
BE2360-EX10	D_2	BD3917	cell-a	-	-	197.08	0.317	cell-b	0.0606	0.00399	51.49	0.00312
BE2360-EX11	D_2	BD3917	cell-a	-	-	197.10	0.323	cell-b	0.002	0.00399	51.39	0.00365

Table 5: The expansions table for the volume calibrations, gas density calibration, and loading measurements.

Calibration id	Name	V (cm^3)	V std (cm^3)	P (psi)	P std (psi)	T (°C)	T std (°C)	Expansions
VC8	cell-b	501.65	3.59	0.8326	0.00812	50.21	0.01240	BE2136-EX1 BE2136-EX3 BE2136-EX5 BE2136-EX7 BE2136-EX9 BE2136-EX11
VC9	cell-d	493.67	2.20	0.8205	0.00566	50.59	0.02621	BE2136-EX2 BE2136-EX4 BE2136-EX6 BE2136-EX8 BE2136-EX10 BE2136-EX12
VC24	cell-a	0.51574	0.002633	138396.9	306.02	51.89	0.01732	BE2360-EX1 BE2360-EX3 BE2360-EX5
VC25	cell-c	0.49806	0.000552	138236.4	326.81	51.36	0.31911	BE2360-EX2 BE2360-EX4 BE2360-EX6

Table 6: The volume calibrations table. The pressurization-cell/expansion-cell combinations are cell-a/cell-b and cell-c/cell-d).

Calibration id	Cell	Gas	ρ (g/cm^3)	T (°C)	Expansions	Volume Calibrations
GD11	cell-c	D_2	0.2337	50.19	BE2360-EX7	VC9, VC25

Table 7: The gas density calibration table.

Experiment id	Sample id	Hydride	P (ksi)	T (°C)	Loading (%)	Expansions	Volume Calibrations	Density Calibration
BE2360	BD3917	PdD_x	135.0	51.78	92.6	BE2360-EX8 BE2360-EX9 BE2360-EX10 BE2360-EX11	VC8, VC24	GD11

Table 8: The loading measurement results table.

Summary of key loading measurements

We primarily focused our UltraSorb loading measurements on PdD_x. Most of our AHE experiments were run between 100 – 150 ksi D_2 at 50°C. We show equilibrium loading values

ranging between 87 – 135 ksi D_2 at 50°C in **Table 9**. A subset of our experiments (specifically, our passive gas-loading Calorimetry campaign) ran at higher temperatures, so we've included a loading value at 108 ksi, 100°C D_2 (**Table 9**).

Experiment id	Hydride	Pressure (ksi)	Temperature (°C)	Loading %
BE2260	PdD_x	87.29	51.20	90.7
BE1971	PdD_x	101.71	50.87	91.2
BE2360	PdD_x	135.00	51.78	92.6
BE2192	PdD_x	108.68	99.88	89.4

Table 9: Summary of key UltraSorb experiments.

Measuring Loading: Mass (MicroSorb)

An alternative direct mass loading measurement is to simply weigh the loaded mass with a microbalance. To accomplish this, we developed a microcapsule that retains a sample of interest (generally a metal lattice) and has a weldable gas inlet (**Fig. 25**). After loading a metal lattice at a set pressure and temperature, the microcapsule is in-situ welded shut, which captures the thermodynamic state. The pressure and temperature surrounding the microcapsule can be brought to ambient conditions, enabling the microcapsule to be removed from the pressure system and have its mass measured with a microbalance. By measuring the mass before and after the loading experiment, the mass difference is a direct measurement of the mols of hydrogen/deuterium. Some of the hydrogen mass is from the free gas internal to the capsule, i.e., not loaded into the sample, so it must be properly accounted for. We've measured the internal volume of the microcapsules using X-ray tomography with a voxel size of 7µm.

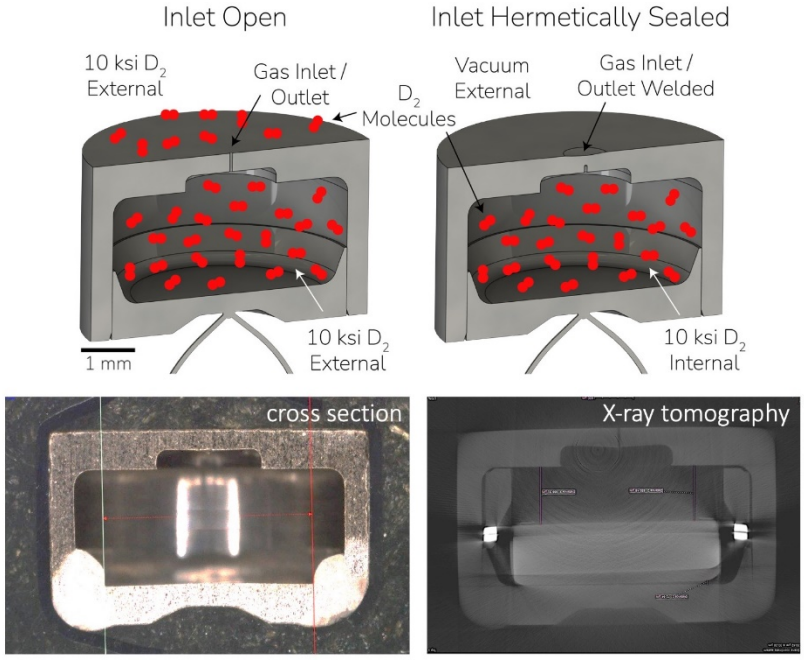


Fig. 25: A microcapsule can be used for a direct-mass loading measurement by incorporating a weldable gas inlet and weighing pre/post the experiment on a microbalance.

Now, let's derive the analytical expressions for the MicroSorb measurement (the nomenclature is given in **Table 10**). The microcapsule contains a sample, retention-spring, and free-gas (**Fig. 26**).

MicroSorb Nomenclature	
V_s	Sample volume
V_r	Retention-spring volume
V_i	Internal capsule volume
V_g	Free gas volume
V_c	Capsule volume (internal volume and wall volume)
ρ_s	Sample density
ρ_r	Retention-spring material's density
ρ_c	Capsule material's density
ρ_g	Gas density
ρ_{air}	Air density
e_s	Sample strain due to loading
m_s	Sample mass
m_r	Retention-spring mass
m_c	Empty capsule mass
m_g	Mass of the free gas
m_{air}	Mass of displaced air
m'	Buoyancy corrected mass
m_{pre}	Mass of entire μ-capsule assembly prior to the experiment
m_{post}	Mass of entire μ-capsule assembly after the experiment
m_{ex}	Excess mass from residual gas and loading
m_l	Mass loaded in the sample
M_l	Molar mass of the atomic species loaded into the sample (atomic species)
M_s	Molar mass of the sample
Γ	Loading fraction

Table 10: MicroSorb nomenclature used in the calculations.

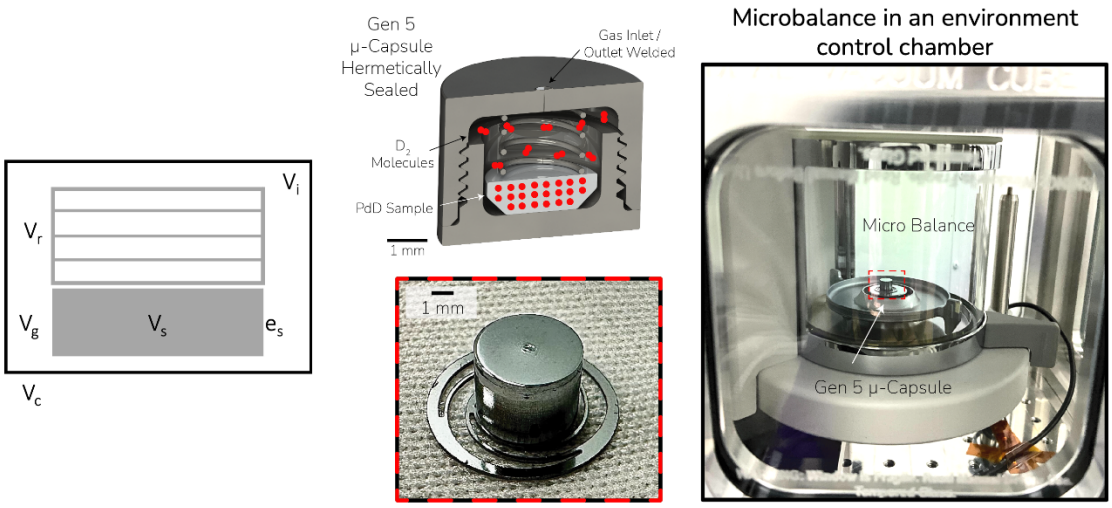


Fig. 26: The MicroSorb capsule contains a sample, retention-spring, and free gas. The capsules are weighed before and after the experiment to determine the loading.

The sample volume is obtained by measuring the unloaded sample mass and referencing the NIST density for the material: $V_s = m_s/\rho_s$. Similarly, the retention-spring's volume is obtained by measuring the spring's mass and referencing the spring's material density: $V_r = m_r/\rho_r$. The internal volume of the microcapsule, V_i, can be measured with X-ray tomography and destructively confirmed by cross sectioning (**Fig. 25**). We've found that the capsules drawing-specification is close to the measured volumes. The sample strains when loaded, e_s, which we've calibrated with lateral-strain measurements. The free gas volume is then given by:

$$V_g = V_i - V_r - V_s \times (1 + e_s)$$

And the mass of the free gas is simply $m_g = \rho_g \times V_g$, where the gas density is obtained from gas's equation-of-state at the measured pressure and temperature. We weigh the full capsule assembly before the experiment, m_{pre}, and can obtain the empty capsule mass:

$$m_c = m_{pre} - m_s - m_r$$

During the experiment, the microcapsules are in-situ hermetically sealed, so we must use buoyancy corrected masses for the sample, spring-retention, and empty microcapsule:

$$m_s' = \frac{m_s}{1-\rho_{air}/\rho_s},\ m_r' = \frac{m_r}{1-\rho_{air}/\rho_r},\ m_c' = \frac{m_c}{1-\rho_{air}/\rho_c}$$

Now, the capsule volume is the sum of the internal volume and wall volume (we get the wall volume from the buoyancy-corrected empty capsule mass and density of the capsule material):

$$V_c = V_i + \frac{m_c'}{\rho_c}$$

The mass of displaced air is given by:

$$m_{air} = \rho_{air} \times V_c$$

We weigh the entire microcapsule assembly after the experiment, m_{post}, and account for the mass of displaced air:

$$m_{post}' = m_{post} + m_{air}$$

Now, we can calculate the excess mass, which is the difference between the post mass and the buoyancy corrected sample, spring-retention, and empty capsule:

$$m_{ex} = m_{post}' - m_s' - m_r' - m_c'$$

The excess mass is a combination of the loaded mass and residual free mass. So, we calculate the loaded mass by:

$$m_l = m_{ex} - m_g$$

The loading is then computed by comparing the mols of loaded species ($n_l = m_l/M_l$) to the mols of lattice atoms ($n_s = m_s'/M_s$):

$$\Gamma = \frac{n_l}{n_s}$$

As an example measurement, we review BE2983. A Pd sample was placed in a microcapsule and held at (145 ksi, 150°C) D_2 for 20 minutes, (145 ksi, 25°C) for 1 hour, and (145 ksi, -100°C) for 1 hour (note: the sample was 500μm thick, so 1 hour at -100°C is not nearly enough time to hit equilibrium loading; nevertheless, we expect to measure a loading >94%, which is the estimated loading at 145 ksi, 25°C). Holding the temperature at -100°C, the pressure was reduced to 20 psi and the microcapsule was in-situ welded shut. The microcapsule was removed from the apparatus and weighed at 20°C. The material densities are given in **Table 11**, volumes in **Table 12**, and masses in **Table 13**. We approximate the volumetric strain from 100% loading to be $e_s = 0.17$. Using the molecular weights of Pd (106.42 g/mol) and D (2.014 g/mol), we converted the masses to mols and found the loading to be $\Gamma = 0.968$.

Component	Density (g/cm^3)
Pd sample	12.023
Retention-spring (WRe3)	19.5
Capsule (MoRe 47.5/52.5)	13.5
Air	0.001025
D_2 (20 psi, 20°C)	0.0002279

Table 11: Material densities used in the BE2983 MicroSorb measurement.

Component	Volume (cm^3)
Pd sample (V_s)	0.004362
Retention-spring (V_r)	0.000419
Internal capsule (V_i)	0.018567
Free gas (V_g)	0.013044
Capsule (V_c)	0.081720

Table 12: Volumes used in the BE2983 MicroSorb measurement.

Component	Mass (g)	Buoyancy-corrected Mass (g)
Pd sample (m_s)	0.052438	0.052443
Retention-spring (m_r)	0.008179	0.008180
Empty capsule (m_c)	0.852493	0.852569
Free gas (m_g)	2.9734e-6	-
Displaced air (m_{air})	9.8473e-5	-
Full assembly - pre (m_{pre})	0.913110	-
Full assembly - post (m_{post})	0.914057	0.914155
Excess mass (m_{ex})	0.00096341	-
Loaded mass (m_l)	0.00096044	-

Table 13: Masses used in the BE2983 MicroSorb measurement.

Stoichiometric Loading

To fully exhaust the loading-hypothesis for AHE, we needed to confirm stoichiometric (i.e., 100%) loading. We developed two paths to stoichiometric loadings: cryogenic temperatures and high-pressure gas.

Stoichiometric Loading: Cryogenic

Loading increases with decreasing temperature (i.e., it lowers the required pressures to reach high loadings), but the kinetics exponentially decrease with low temperatures. The challenge in reaching stoichiometric loadings at cryogenic temperatures is the slow loading kinetics, i.e., it could take days-weeks to load depending on the sample geometry (surface area and diffusion length). To reduce the time to reach stoichiometry, we ran cryogenic experiments at high-pressure (and added most of the loading at higher temperatures before cooling) and modified the sample geometries (e.g., thin-films, nanoparticles, etc.). We ran typical cryogenic experiments at ~150 ksi,

-100°C, which should be stoichiometric in PdD and PdH (**Fig. 27**). As an example, we show a cryogenic calorimetry experiment run at 130 ksi, -100°C D_2 (PdD_x is in Ch2), where we observed no evidence of AHE (**Fig. 28**).

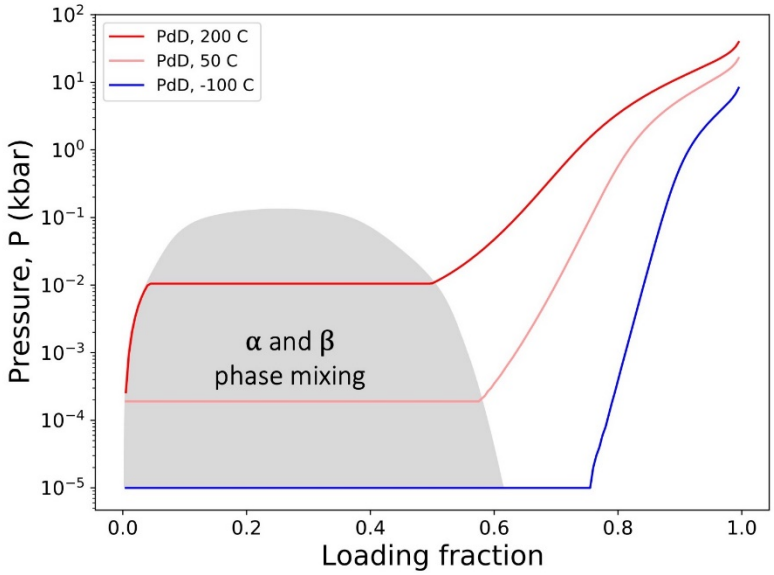


Fig. 27: Stoichiometric loadings can be achieved at 10 kbar, -100°C.

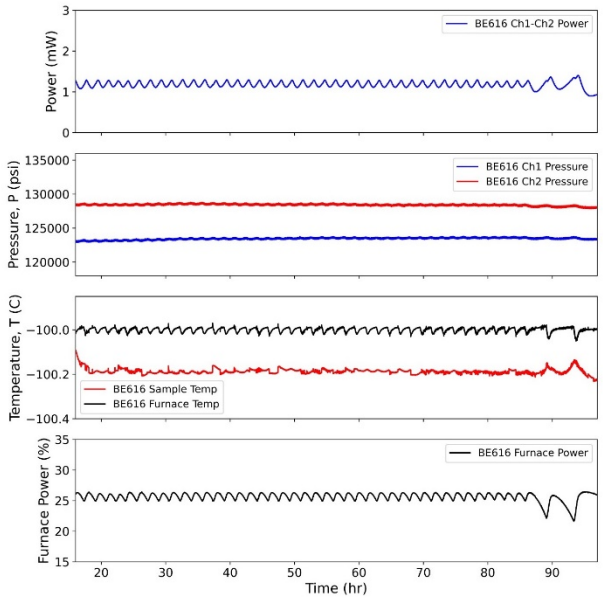


Fig. 28: A passive differential calorimetry experiment ran at 130 ksi, -100°C.

To confirm stoichiometric PdD at 150 ksi, -100°C, we ran a MicroSorb experiment. To reduce the diffusion path and increase surface area, we in-situ modified the Pd sample. Cavities and surface blisters form in metal hydrides when rapidly heated and pushed too far out of equilibrium (i.e., thermal diffusion >> mass diffusion). The result is the creation of new surfaces in the bulk of the

metal hydride (**Fig. 29**). In BE3270, the in-situ modification decreased the diffusion path of deuterons from ~440 µm to ~30 µm and increased the surface area. The MicroSorb measurement showed a loading of 99-100%.

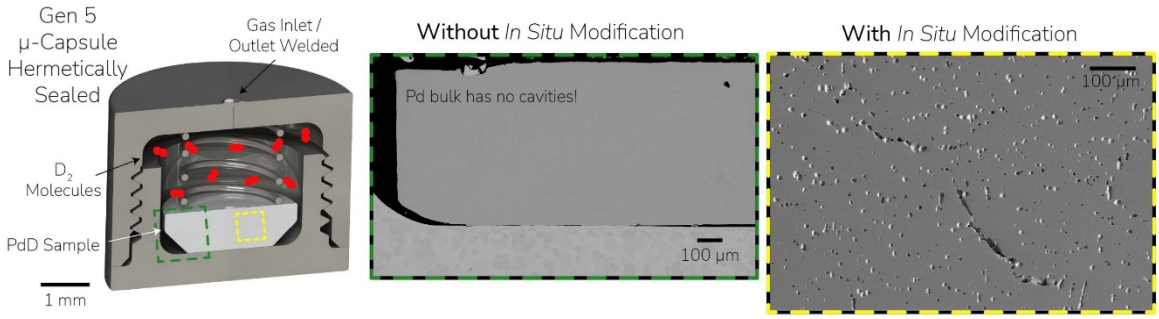


Fig. 29: In BE3270, a MicroSorb experiment in-situ modified a Pd sample to decrease the diffusion path and increase surface area while loading the Pd at 150 ksi, -100°C. The loading was measured at 99-100%.

Stoichiometric Loading: High Pressure Gas

High pressure is a direct method to achieving stoichiometric Pd/D(H) loadings. However, the pressure (at 25°C) above which stoichiometric loadings are guaranteed is not well established. Diamond anvil cells have been used in conjunction with synchrotron x-ray diffraction measurements, but this approach has large error bounds in both loading (~5%) and pressure (~100 ksi) (**Fig. 30**).

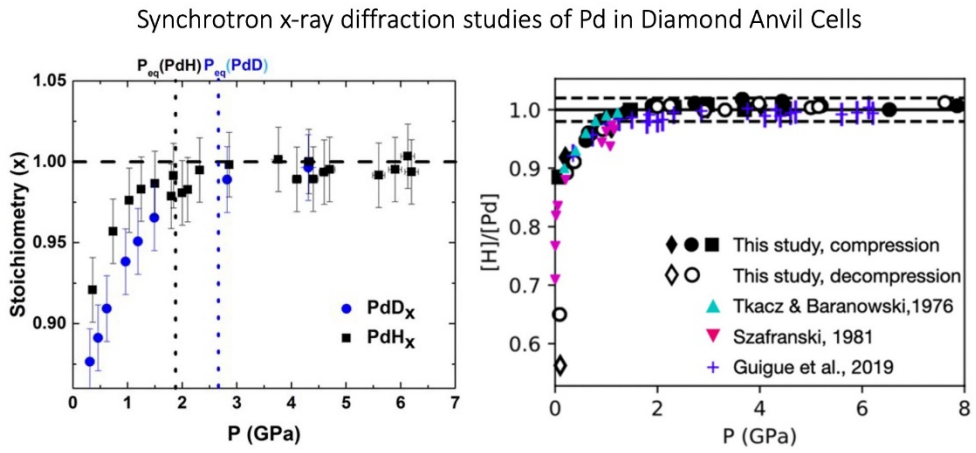


Fig. 30: Synchrotron x-ray diffraction studies of Pd hydrides in Diamond Anvil Cells[16,17].

[16] Z. Geballe, et. al, Phys. Rev. B 103, 024515 (2021)
[17] B. Guigue, et. al, J. of Appl. Phys. 127, 075901 (2020)

Electrical resistivity of PdH_x and PdD_x is the most common metric used to quantify loadings in Pd, but the ability to differentiate between 99% and 100% loading is difficult. While the mapping between resistance ratio and loading is well documented and regularly referenced, there is some spread in the resistivity values reported in literature in the last 1-2% of loading, and there are no reports of resistivity measurements past 435 ksi in D_2 (**Fig. 31**). We completed resistive gas measurements up to 550 ksi and have run several controls to decouple the effects of loading and pressure to better determine the pressures required for 100% D/Pd.

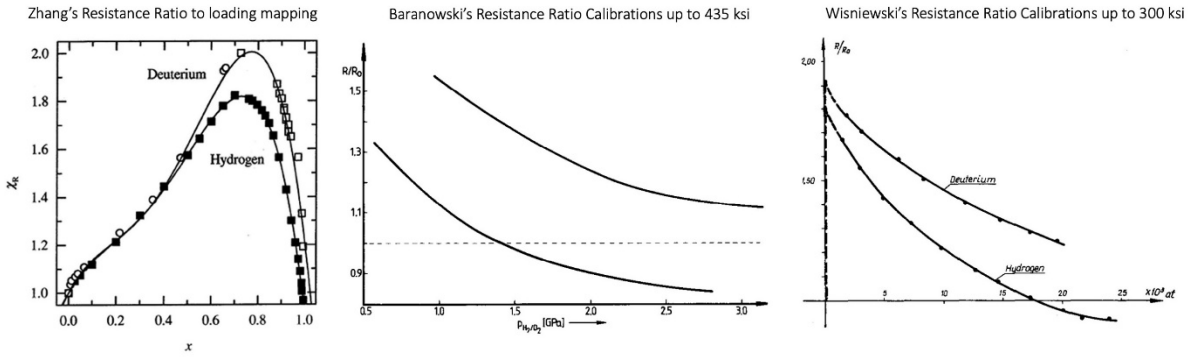


Fig. 31: Resistance ratio calibrations from literature shows spread in the last 1-2% of loading[18,19,20].

To confirm stoichiometric PdD/H loadings via resistivity, we look for the point at which the resistance ratio stops changing with increasing pressure. Diamond anvil literature shows that the lattice strain does not change from 2 GPa to 8 GPa, indicating that once the Pd octahedral sites are fully occupied, the lattice does not load further, i.e., 100% loading is the max limit for Pd. However, for PdH we continue to observe decreases in resistivity above 300 ksi; similarly, for PdD we continue to observe decreases in resistivity above 400 ksi and measure R/R_0 values as low as 1.1, 5% lower than the lowest numbers reported in literature (**Fig. 32**).

[18] Z. Zhang, et. al, J. of Elec. Analytical. Chem. 528 (2002)
[19] R. Wisniewski, Review of Scientific Instruments, 41, 3 (1969)
[20] B. Baranowski, et. al., J. of the Less-Common Metals, 158 (1990)

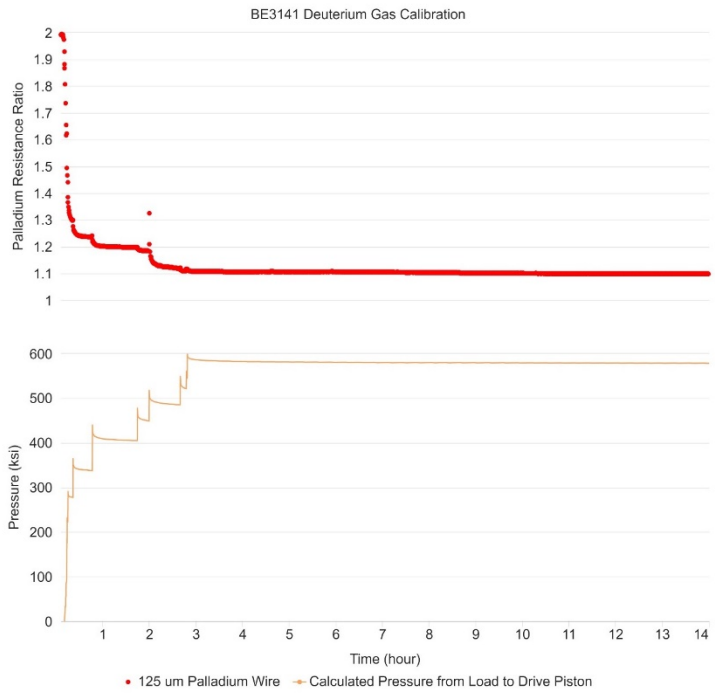


Fig. 32: In BE3141 we ran a resistance ratio calibration of a 125 µm diameter Pd wire up to >550 ksi D_2 at 25°C. The resistance continues to drop with increasing pressure.

To build confidence that we are in-fact reaching stoichiometry with PdH and PdD, we need to decouple the effects of pressure and loading on the resistivity of Pd at >300 ksi pressures. Hence, we performed measurements with only N_2 and with H_2/D_2:N_2 gas mixtures. In N_2, we confirmed pressure effects on Pd resistivity are non-negligible above 150 ksi (**Fig. 33**). To establish stoichiometry, we find the point where R/R_0 converges for the pure gases (H_2/D_2) and gas mixtures (H_2/D_2:N_2). For protium, we targeted an 84:16 H_2:N_2 mixture to reach ~380 ksi of H_2 partial pressure at 450 ksi of total pressure (~70 ksi of N_2 partial pressure). For deuterium, we targeted an 82:18 D_2:N_2 mixture to reach ~410 ksi of D_2 partial pressure at 500 ksi of total pressure (~90 ksi of N_2 partial pressure).

For PdH we find nearly identical final R/R_0 values in H_2:N_2 mixtures as we did at ~400 ksi H_2. This suggests that we are at PdH stoichiometry >330 ksi and that the effect pressure has on resistance is what continues to drive resistivity lower. Similarly, for PdD we find the R/R_0 values in D_2:N_2 mixtures converge at 500 ksi, indicating that any further decrease in resistance is due to pressure alone, not loading.

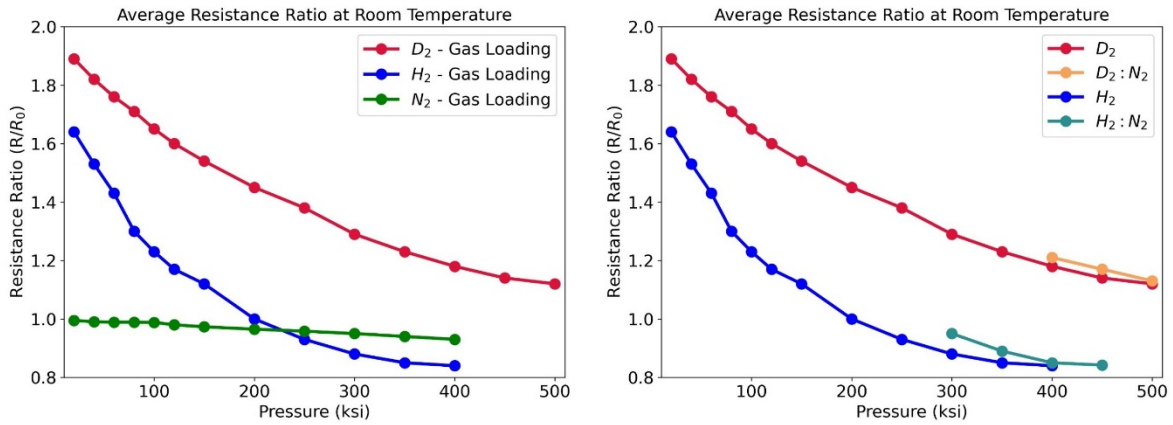


Fig. 33: Pressure effects on resistance are non-negligible at pressures >150 ksi as shown the by N_2 control. To decouple the effects of pressure and loading on the resistivity of Pd, we performed measurements with H_2/D_2:N_2 gas mixtures.

Through a combination of hydrogen and deuterium gas mixtures with inert gases, as well as access to sufficiently high pressures, we determined the pressures required to reach stoichiometry. For PdH, stoichiometry is around 330 ksi, higher than the ~270 ksi values proposed in literature. For PdD, we find that ~400 ksi is not enough, and that pressures >450 ksi are needed. Our gas-mixture calibrations are the first demonstrations to decouple resistance changes from pressure and loading. While the pressures required to reach stoichiometric H/Pd and D/Pd are higher than we anticipated, we are well suited to operate at pressures beyond what is required for stoichiometry to ensure we are studying Pd under unity loading conditions at ambient temperatures.

Hardware

A fundamental hypothesis to exploring the existence/non-existence of AHE is the necessity of high fractional loadings (e.g., 90-100%). From a gas-loading perspective, this requires high-pressure, hydrogen-compatible systems (e.g., >10 kbar H_2/D_2). Commercial systems, which can handle H_2/D_2 at these pressures, do not exist and very few components (pumps, valves, vessels, etc.) are commercially available. So, we put significant effort into building zero-leak high-pressure systems and a facility to support them.

Hardware: Safety

High-pressure hydrogen introduces safety risks: combustible gas, embrittlement (i.e., strong, ductile materials become brittle), shrapnel from pressure failures, etc. To mitigate these risks, we built armored chambers that housed our pressure sources and experimental apparatuses (**Fig. 34**). All pressure failures, combustions, etc. were confined inside the chamber, mitigating any personnel risks; all high-pressure experiments were performed remotely, without humans in the chambers. Nevertheless, we still added isolated air exchange systems and fire suppression systems on each chamber to mitigate the risk of hydrogen fires inside the chambers. It's important to note that our compressed-hydrogen volumes were very small (e.g., < 10 cm^3 at high pressures), so the total stored energy was low, and the risk of run-away hydrogen fires was nonexistent. In total, we built a "main" chamber that housed pressure sources and fed pressure into 8 experimental chambers: 5x chambers that fit a full-sized optics table and 3x mini chambers that housed compact experimental apparatuses.

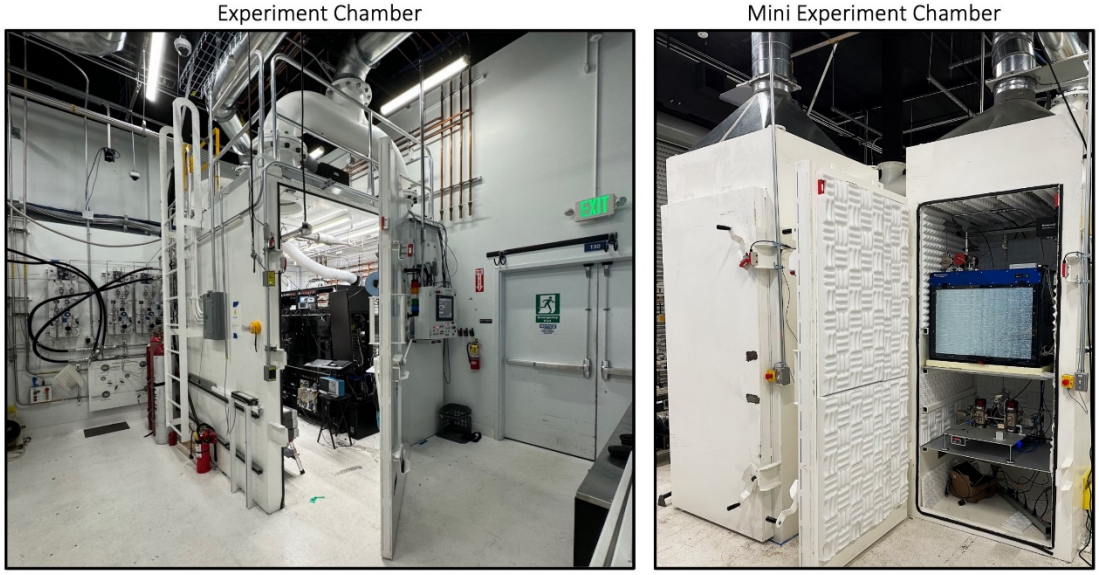


Fig. 34: Armored and vented chambers housed our high-pressure hydrogen sources and experimental apparatuses.

Hardware: Pressure Sources

We developed 4 different pressure sources (**Fig. 35, Table 14**). To start, Powerpack 1 was designed to use vendor-supplied components (e.g. valves and pumps) and had a maximum pressure capability of 100 ksi. We quickly ran into bottlenecks with vendor supplied components as high-pressure requires constant maintenance (e.g., seal stacks wear quickly, sub-components break, etc.) and we had to frequently send parts to vendors for refurbishment. This led to significant down-time. Moreover, vendor-supplied valves and pumps were leaky and Powerpack 1 could only hold high pressure for < 1 min. So, we designed in-house components (valves, pumps, fittings) to increase reliability and enable rapid maintenance, and over the years designed and built 3 different zero-leak systems: Powerpack 2, Screwjack, and Protonpack. Most high-pressure components were made with Toughmet® due to its strength and hydrogen resistance.

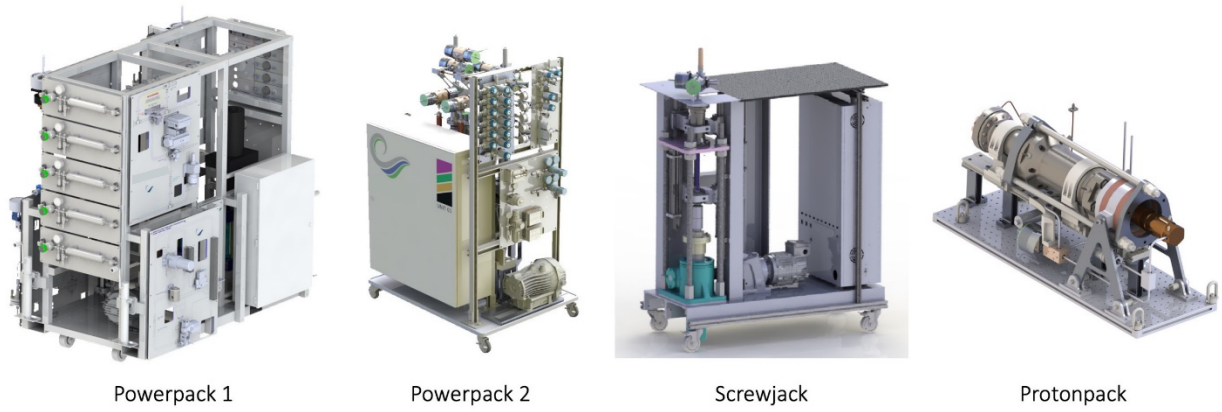


Fig. 35: Pressure sources developed at Callisto.

Source	Pressure Capability	Drive	Media
Powerpack 1	100 ksi	Hydraulic	Gas
Powerpack 2	150 ksi	Hydraulic	Gas
Screwjack	150 ksi	Electric	Gas / Liquid
Protonpack	150 ksi	Pneumatic	Gas

Table 14: High-level specs of the pressure sources developed at Callisto.

Powerpack 2 fully integrated in-house components (discussed more below), was our first zero-leak system, and increased our pressure capability from 100 ksi to 150 ksi. To increase control over a pump's piston position (which is important for liquids as they are highly incompressible i.e., small changes in piston position lead to large pressure changes), we moved to a Screwjack design. The Screwjack utilized our valves and pumps but replaced hydraulic actuation with a screw-jack driven by a motor. The Screwjacks were our most stable and reliable systems. Due to limited facility space and the need for additional experimental apparatuses, we opted to build mini chambers. To decouple the mini chambers from sharing pressure with other chambers, we built a small pressure source, the Protonpack, that fit inside a mini chamber along with an experimental apparatus. It was a pneumatically actuated system that was the most compact, but also the most hazardous due to a lack of control over the piston during leaks.

Our 150 ksi zero-leak valves were essential to all pressure systems and experimental apparatuses (**Fig. 36**). Our valves reduced the number of possible leak paths from 6 (in vendor-supplied valves) to 2 and updated the seal materials to metal rings and Teflon (vendor-supplied valves used metal, elastomers, and even leather). The updated seals achieved zero-leak and increased the cleanliness of our systems, which is important for debris-free AHE experiments. Seals that shed material (such as leather) would often contaminate our samples.

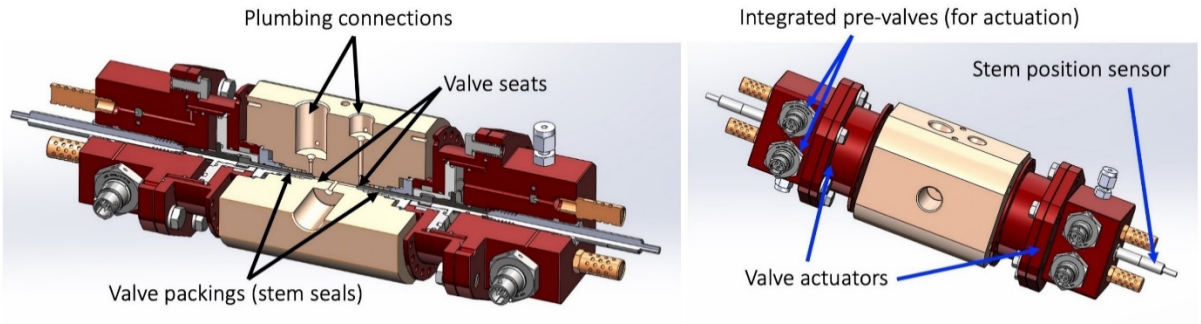


Fig. 36: 150 ksi compatible manifold valve.

We designed and fabricated 150 ksi zero-leak pumps (**Fig. 37**). The pump's seal is a small Teflon ring which is compressed ~10% by a carbide packing ram to a far higher internal stress than the surrounding gas. We found this seal approach much more effective than pressure-energized designs, which require a certain minimum pressure to seal. The linkage to the actuation cylinder is fixed in one degree of freedom but unconstrained in 5 degrees of freedom. This assures alignment of the rod with the seal pack. Pumping is accomplished by a tungsten carbide rod displacing volume inside the Toughmet bore.

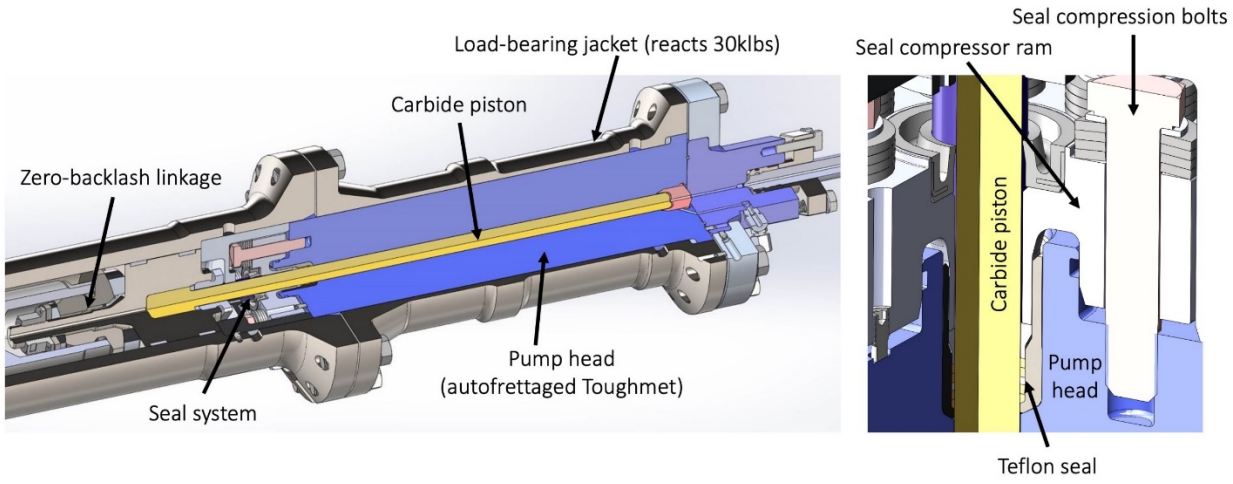


Fig. 37: 150 ksi intensifier (i.e., pump) that was designed and fabricated in-house.

In addition to our 150 ksi pumps, we developed a 20 ksi pump to enable finer control over the pressure from 0-20 ksi (**Fig. 38**). Initially, we used commercial boost pumps which do not give fine pressure control and require an additional pressure regulator. Our 20 ksi pumps use commercial off-the-shelf seals and implement an actuator/pump linkage similar to the 150 ksi pumps.

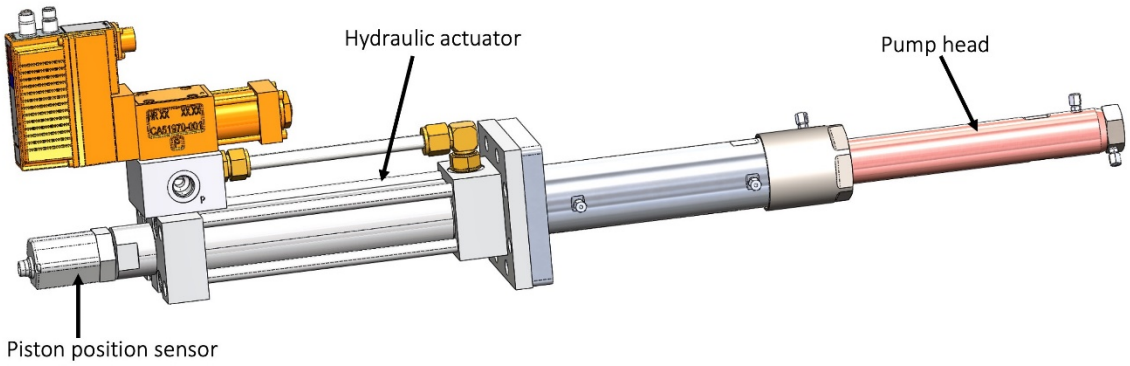


Fig. 38: 20 ksi intensifier (i.e., pump) that was designed and fabricated in-house.

Another highly useful component we developed was a 150 ksi tube fitting (**Fig. 39**). We experienced tube failures (i.e., ruptures) due to torsional stress from commercial fittings (the seat compression is generated by a single gland nut, transmitting torque into the tube). Our fittings removed the torsional stress with a bolted flange design and never led to a tube failure.

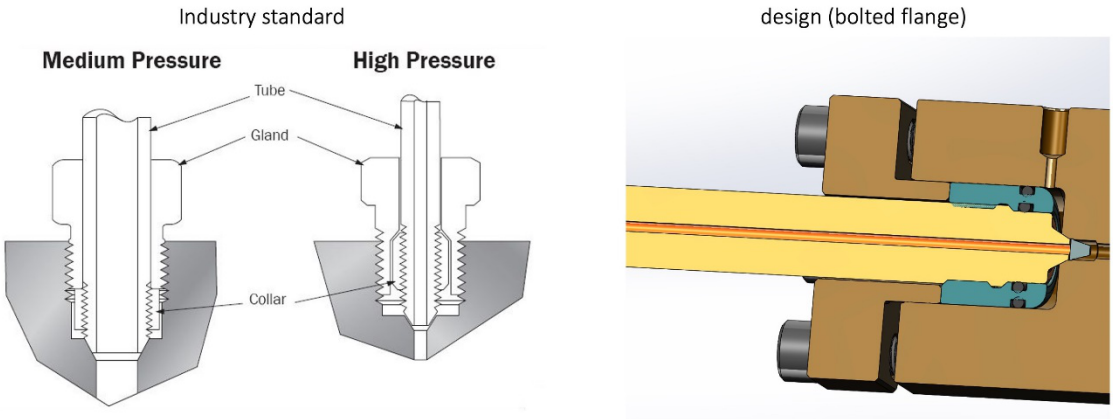


Fig. 39: Commercial fittings generated seat compression with a single gland nut, which transmits torque into the tube. Callisto fittings generate seat compression with bolts and no torque is transmitted into the tube.

Hardware: Pressure Vessels

Pressure vessels are central to high pressure experiments: they house the sample, control the pressure and temperature environment, and enable access for stimulation and in-situ diagnostics. While we developed an array of vessels, the main workhorse was the "Hammerhead" vessel, which evolved into a versatile and modular testbed (**Fig. 40**, **Table 15**). Depending on the experiment,

various volume, vacuum-level, thermal, optical, and electrical options could be configured to achieve the desired environmental and stimulation conditions.

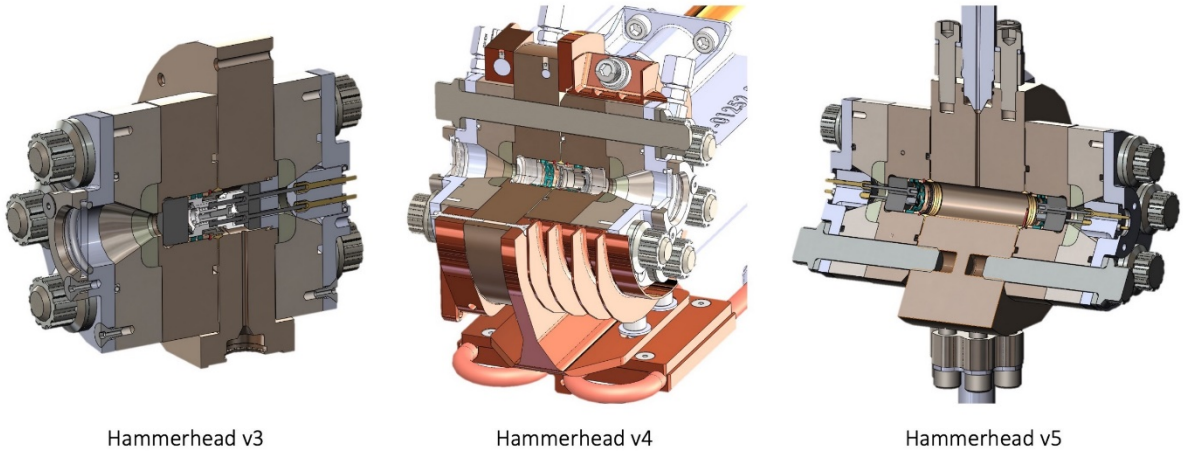


Fig. 40: The Hammerhead vessel served as a versatile testbed, enabling a wide range of pressure and temperatures in addition to providing access to stimulation sources and diagnostics.

Hammerhead	Pressure Range	Temperature Range	Features
v1	1mbar to 10 kbar	25°C to 50°C	• Electrical feedthroughs • Optical port
v2	1mbar to 10 kbar	-120°C to 50°C	• Cryo • Electrical feedthroughs • Optical port
v3	1nbar to 10 kbar	-120°C to 50°C	• Cryo • High vacuum • Electrical feedthroughs • Optical port
v4	1mbar to 10 kbar	-120°C to 350°C	• Cryo + High temperature • Two optical ports
v5	1mbar to 10 kbar	25°C to 50°C	• Corrosion resistant • Large internal volume • Dual electrical feedthroughs

Table 15: The Hammerhead vessel served as a versatile testbed, enabling a wide range of pressure and temperatures in addition to providing access to stimulation sources and diagnostics.

Developing zero-leak 150ksi electrical feedthroughs was a challenge. We initially started with brazing pins to single crystal sapphire windows to take advantage of our optical hardware. However, the brazed sapphire approach was unreliable and often cracked. So, we developed feedthrough pins with a self-energizing seal (**Fig. 41**). Pressure from the test cell pushes against the head of the pin, which transfers force to a ceramic traveling anvil. This type of design improves as sealing pressure increases. Manufacturing the seals requires micromachining of both metallic and ceramic parts, but the pins are easy to assemble and do not require specialized equipment.

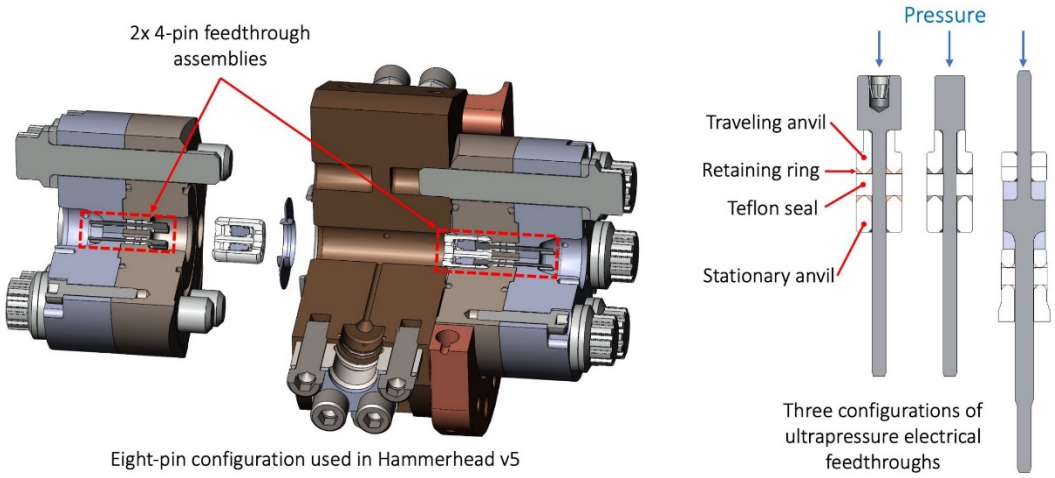


Fig. 41: Zero-leak, 150ksi electrical "pin" feedthroughs take advantage of a self-energizing seal that improves sealing as pressure increases.

Hardware: >300 ksi

While our 150 ksi hardware was effective at studying AHE above 90% loadings, we need much high pressures to hit 100% (stoichiometric loading) at room temperatures. To bypass designing >150 ksi valves, tubing, and fittings, we pursued a combined intensifier-vessel design. Sealing >300 ksi hydrogen gas is notoriously difficult and has only been historically demonstrated by 2 academic groups outside of the diamond anvil cell literature with little-to-no documentation. Hydrogen compatibility significantly reduces the materials we can work with, particularly high strength steels, which were successful in our >300 ksi liquid apparatuses. While Toughmet was our workhorse hydrogen-compatible material at lower pressures, >300 ksi is well beyond Toughmet's ~140 ksi yield strength. So, we moved to multi-layer vessels to constrain the materials (**Fig. 42**).

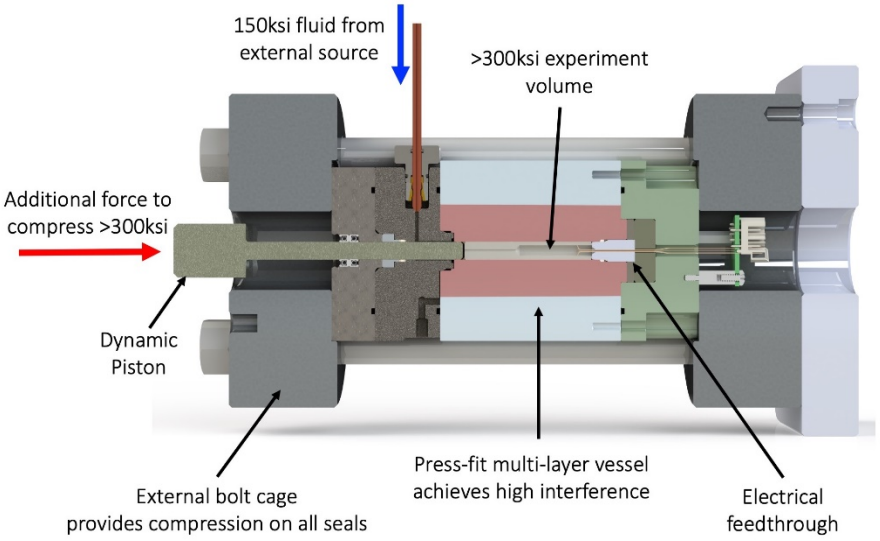


Fig. 42: Multi-layer vessel design for >300 ksi.

By placing a liner under compression with a jacket, the stress of the inside wall can be placed under compression during assembly to decrease the tensile stress during pressurization (**Fig. 43**). The higher the interference between the two layers, the greater the compression. The limit of how much interference one can achieve is determined by the yield strength of the jacket and inner liner. The interference pressure between an internal layer and external layer is given by:

$$p_{int} = \frac{\delta_R}{R\left(\frac{R^2(\nu_i - 1)}{E_i(R^2 - r_i^2)} - \frac{R^2(\nu_o - 1)}{E_o(R^2 - r_o^2)} + \frac{r_i^2(\nu_i + 1)}{E_i(-R^2 + r_i^2)} + \frac{r_o^2(\nu_o + 1)}{E_o(R^2 - r_o^2)}\right)}$$

where E_i is the modulus of elasticity of the internal layer, E_o is the modulus of elasticity of the external layer, ν_i is the Poisson's ratio of the internal layer, ν_o is the Poisson's ratio of the external layer, δ_R is the radial interference, R is the final radius of the interface, r_i is the radius of the internal layer, and r_o is the radius of the external layer.

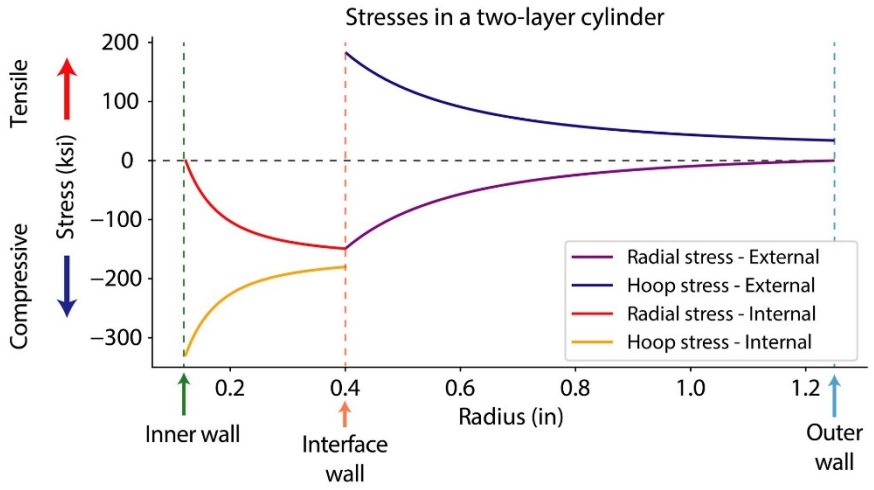


Fig. 43: Stresses in a two-layer cylinder.

To reach pressures >300 ksi, we need to keep the bore from dilating. We began to employ extremely high interference, multi-layer vessels to keep the internal bore under compression. We accomplished this with tapered fits (instead of thermal fits) because of the ability to achieve significantly larger interferences. While extremely stiff materials like Tungsten Carbide make ideal jacket materials, they are extremely poor under tension. We employ 3-layer designs to keep the carbide contained as an intermediate layer. A Toughmet inner sleeve (which mitigates hydrogen embrittlement) is pressed into the Tungsten-Carbide/Steel composite jacket, achieving theoretical interference pressures >250 ksi (**Fig. 44**). This is beyond the elastic regime, collapsing the inner liner. Aside from the ability to use carbides and ceramics as the vessel material or a significantly larger footprint, this is as radially constrained as we can get. The vessel is honed to the correct diameter, but the material properties are ambiguous in this plastic regime. With the

bore under large amounts of compression from the high interference jackets, we can now seal directly along the bore of the vessel with a simple anti-extrusion ring and PTFE liners (**Fig. 45**). The seal engages with the piston only after filling with >100 ksi gas, removing most of the compressibility of the gas. The seal translates along the bore of the vessel as the piston is driven deeper to achieve higher pressures. This is only possible with minimal bore dilation.

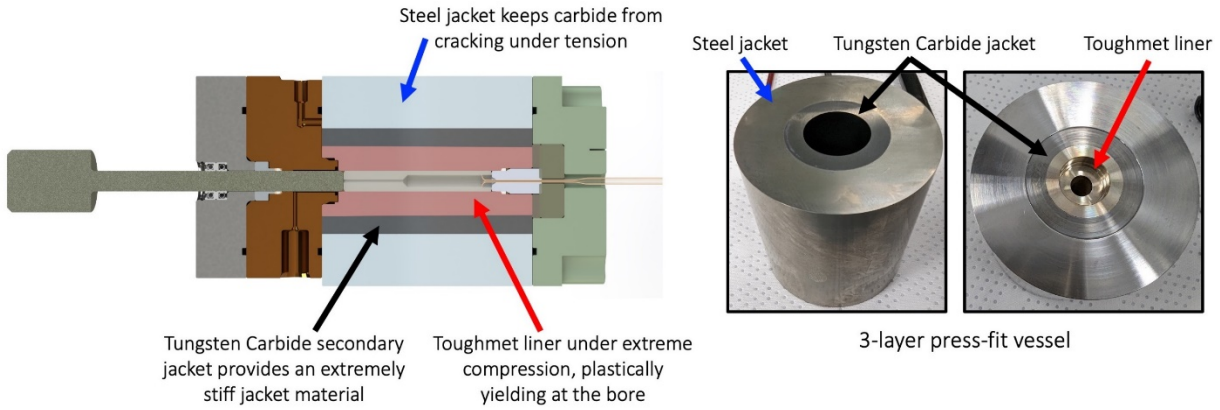


Fig. 44: A 3-layer press-fit vessel utilizes a steel jacket to keep a Tungsten Carbide jacket from cracking, which in turn keeps a Toughmet liner under extreme compression.

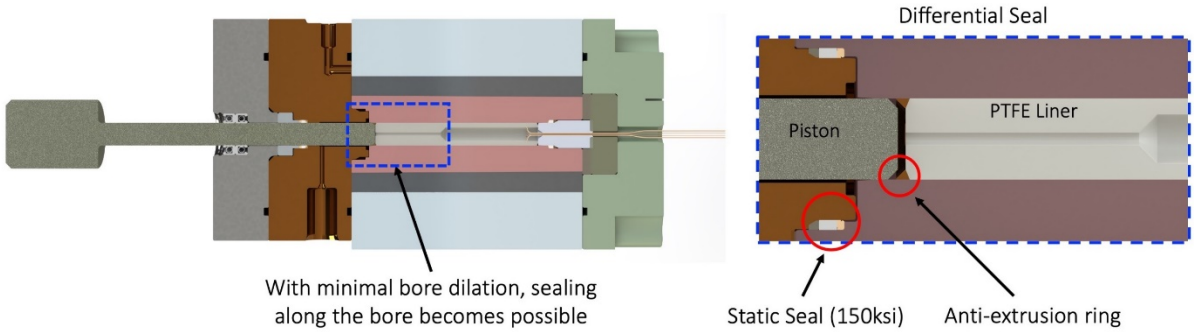


Fig. 45: With the bore under compression, the seal is directly along the bore of the vessel with an anti-extrusion ring and PTFE liner.

With 3-layer vessels, we could consistently reach ~350-400 ksi pressures. At pressures >400 ksi, the innermost liners significantly dilate, and the seals extrude into the gap between the piston and liner. To address this, we began to sequentially autofrettage the vessels to higher pressures. This got us to 450 ksi. Eventually, we ran into limits of radial compression and autofrettage. At >450 ksi, we consistently observed that liners grow axially. With axial compression >200 ksi, we constrain the material in all dimensions and our limitation now becomes the strength of our piston. Using the highest strength grade carbides for pistons, we could access pressures near 600 ksi, operating well past the elastic regime. These are the highest pressures reached with hydrogen outside of diamond anvil cells, pushing past Baranowski's limits.

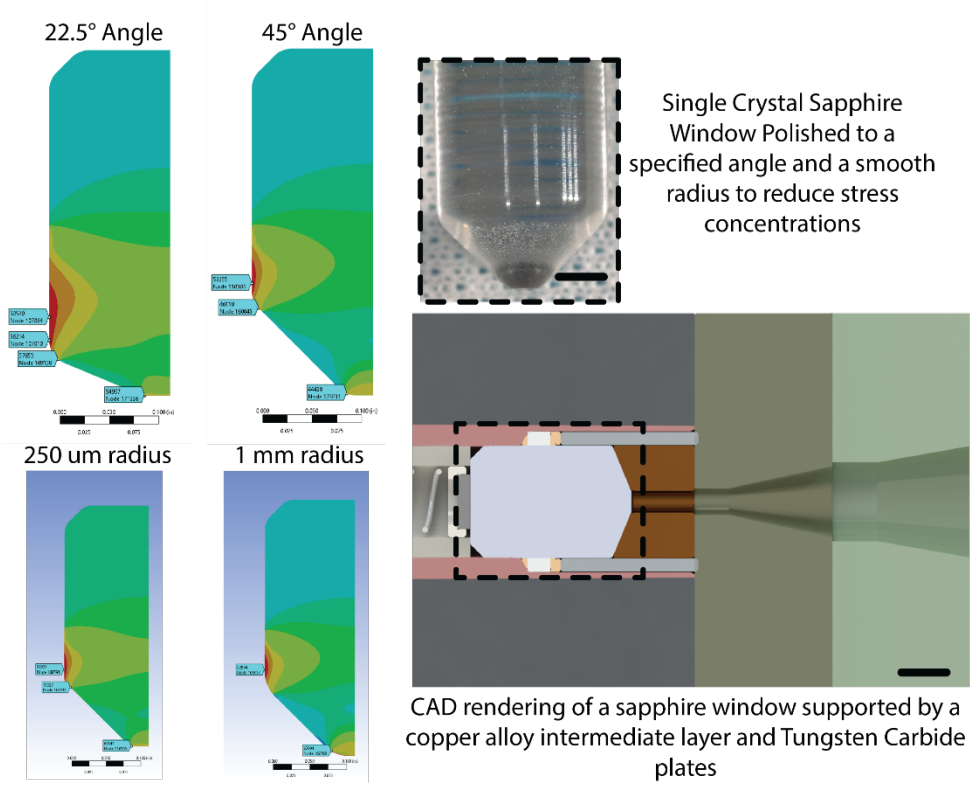


Fig 46: Single crystal sapphire window designed to take advantage of a cone, distributing the load over a larger area. Black scale bar shown is 2 mm. We used FEA modeling to determine the optical angle and radius of curvature to reduce stress concentrations at the edges. The images shown compare the simulated stresses at 450 ksi of gas pressure

To contain gas pressures >300 ksi, several strategies are employed to keep any material from experiencing stresses above its limit. For example, by pre-stressing the innermost layers through interference fits, we decrease the partial pressure that layer is exposed to. However, the electrical and optical ports into the vessel are exposed to the full pressure differential, as they provide access from atmospheric conditions to the pressure inside the vessel. These are typically the weak links in pressure vessels and where material choice is most limited.

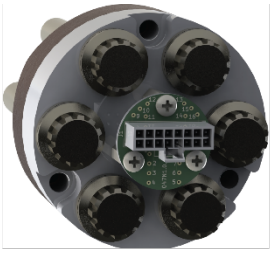
Optical Access > 300 ksi:

For optical access into the pressure vessels, we utilized single-crystal sapphire and silicon carbide windows. Aside from diamond, these are the best material candidates that are optically transparent. However, even single-crystal sapphire and silicon carbide contain minor defects that serve as crack initiation points, particularly at edges, leading to a significant reduction in the operable pressure. To overcome this limitation, we investigated a variety of geometries to find which ones are most resilient to pre-existing cracks.

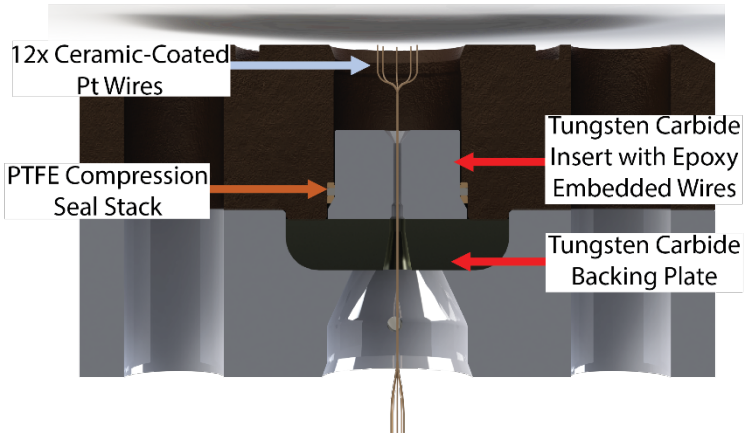
Our most successful design was to use a 45-degree cone on the low-pressure side of the optical window. A grinding process was developed to achieve a desired geometry, and diamond polishing

was used to remove any defects that were introduced in the grinding step. The process of adding an angle to the optical window removed the defects already present in the window as well as distributed the load over more surface area. A copper alloy intermediate layer matches the angle on the sapphire and serves as an intermediate layer between the tungsten carbide backing layer and the brittle sapphire windows. Behind the sapphire window, we found that a softer material, like copper, served a critical role in ensuring that the load distributed as evenly as possible to the sapphire. Stiff (e.g. tungsten carbide) intermediate layers were tested with angles matching the sapphire windows, but typically ended in crack propagation at pressures >300 ksi, while softer materials (e.g. copper alloys) could consistently operate at pressures > 450 ksi. The aperture of the copper intermediate layer and tungsten carbide backing flange was limited to ~2 mm for 7 mm sapphire windows for operation at 450 ksi. A low ratio of unsupported / supported area is needed to prevent the sapphire window from experiencing extreme tensile stresses.

Electrical Access > 300 ksi:



Feedthrough PCB Assembly CAD rendering



Similar to the optical window port, electrical feedthroughs experience the full pressure differential (>300 ksi to atmosphere). Tungsten carbide (WC) inserts were used with a PTFE-compression seal stack around the periphery. In short, a thin PTFE ring is compressed between the WC insert and the vessel wall, and opposing copper-alloy rings are used to prevent the extrusion of the PTFE at elevated pressures. An 800 µm hole is drilled into the WC insert, where 12 ceramic-coated Pt microwires (125 µm diameter) are embedded with an epoxy and aluminum oxide encapsulant (50 nm Al_2O_3 nanopowder mixed with Stycast 2850FT). These electrical feedthroughs were stable up the highest pressures we've tested, providing electrical access at pressures >500 ksi.

Measuring hydrogen pressures >300 ksi is non-trivial. Intuitively, one can calculate the pressure p by $p = F/A$, where F is the applied force and A is the area the force is applied to. However, this method includes the frictional forces from sealing. We attempted to measure the vessel strain at the outer diameter; however, due to operating in the plastic regime, it's difficult to measure accurate results >300 ksi. Resistive manganin sensors are often utilized to infer pressure at different temperatures, given its resistivity is temperature-insensitive and pressure-sensitive. However, it's well documented that the pressure response of manganin resistive sensors becomes highly non-linear past 400 ksi[21,22]. Through a combination of these approaches, we can determine the minimum bound of pressure and find that the most accurate pressure data is calculated from the force to drive the piston into the high-pressure area minus the friction of the seals along the bore (**Fig. 47**). To bound the impact of seal friction, we move the piston forward and backward equal amounts around our pressure setpoint. When the piston moves forward/backwards, the friction acts in the positive/negative direction, respectively.

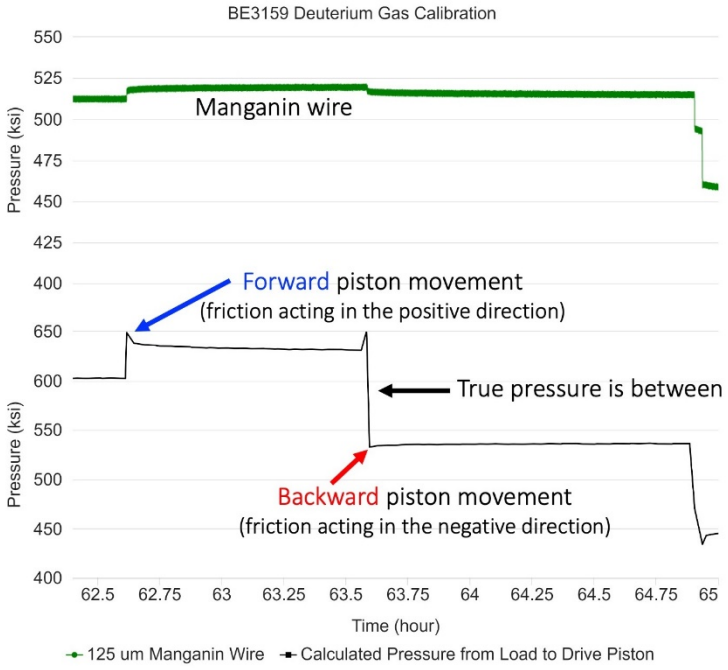


Fig. 47: Pressure measurements >300 ksi using a Manganin wire and load to drive the piston; forward & backward piston movements around the pressure setpoint bound the impact of seal friction on the load.

Hardware: Machine Shop

Due to our hardware quantity, exotic materials, and tight tolerances (e.g., 0.001 inch), fully outsourcing our parts to external machine shops became prohibitive in time. For example, vendors

[21] L. Xiang, et. al., Rev. of Sci. Instr. 91, 095103 (2020)
[22] R. Zeto et. al., J. of Appl. Phys. 50, 5 (1968)

often quote 6-8 week lead times for a single, tight-tolerance part with non-standard materials. Given the number of iterations required to achieve our record-setting hydrogen pressures (for large volume), we built an in-house machine shop to accelerate our cycle-times (**Fig. 48**). As an example, during our >300 ksi hardware project, our machine shop enabled >15 unique design iterations, made thousands of match-fit parts, and performed hundreds of part modifications over the course of about 4 months. If completely outsourced, the >300 ksi hardware project would have taken 5-10 years assuming a 4-8 week cycle time.



Fig. 48: Our in-house machine shop significantly accelerated our design-build-test-learn cycle-time.

Thermal Ignition

Introduction

One of the leading hypotheses put forward by the LENR community for potential AHE observations is that the process is driven by a *phonon-mediated* mechanism; in other words, vibrations of atoms within a material could drive deuterons to close enough proximity to undergo fusion. The main experimental "knob" for tuning these atomic vibrations is simply temperature: The amplitude and frequency of these vibrations and motion of the atomic lattice is known to scale exponentially with increasing temperatures[23]. While, in theory, this is a straight-forward process (there are many ways to heat a material to extreme temperatures), doing so in a controlled manner that allows for studying and looking for potential AHE events is significantly more difficult.

Our Thermal Ignition campaign focuses on the development of systems that can deliver *rapid and intense* temperature rises to a loaded Pd sample while maintaining the ability to have extreme sensitivity to potential AHE events. In the following text, we will discuss the evolution of this experimental effort as well as a detailed perspective of the final effort, TI-12, that provides unambiguous data supporting that **AHE is not occurring via a thermal stimulation process.** This is shown at temperatures up to and in excess of melting for $PdD_{1.0}$ and $PdH_{1.0}$ (unity loading).

Synopsis of earlier TI campaigns

Our prior experimental campaigns, TI1 – TI11, were focused on the development of an accurate temperature reading under ultra-pressure conditions. In the earliest renditions of TI, we monitored the sample temperature optically by way of pyrometry; this method uses the light emitted from the heated PdD_x sample as a metric of temperature (i.e., consider a "red hot" metal; the amount of light emitted and its color are directly related to the temperature of the metal). Shortly after the start of this campaign, we observed a form of isotopic contrast: Optical pyrometry suggested that the temperature of PdD_x was greatly exceeding the temperature of PdH_x under identical stimulation conditions. Ultimately, we determined that this difference in measured temperature is an artifact unique to ultrapressure environments. Generally, symmetric molecules such as H_2 and D_2 do not absorb light due to a lack of dipole moment. However, at sufficiently high densities/pressures, these molecules collide frequently enough to *temporarily* gain a dipole moment; this process, called collision-induced absorption, is highly sensitive to the mass of the atoms composing the molecule. In other words, the 2x difference in mass between H_2 and D_2 led to different amounts of

[23] The distribution of these atomic vibrations follows Bose-Einstein statistics, where the number of atoms at an energy level, E, scales as $\propto \exp(-E/kT)$.

light absorption at the wavelengths that are used for optical pyrometry, thus causing the observed isotopic contrast and obfuscating our ability to resolve potential AHE events.

To bypass this limitation, we developed ultrapressure electrical feedthroughs to allow for in-situ thermometers. This allowed us to accurately quantify the sample temperature by way of a thermocouple which is known to be insensitive to gaseous environments and negligibly change readout due to pressure. This major hardware upgrade for TI ultimately led to several other campaigns that became greatly improved with the option of electrical feedthroughs for our ultrapressure vessels.

With thermocouples for temperature readouts, we again observed isotopic contrast between measurements under 150 ksi D_2 and H_2. Specifically, we found that PdD_x was heated to a significantly higher temperature than PdH_x under identical stimulation conditions. To confirm whether this effect was due to AHE events or chemical in origin, we performed identical experiments using a material that does not load – Platinum – and is therefore not AHE active. With an identical sample geometry, we confirmed that the peak temperature difference that we initially observed with loaded Pd was repeatable in our Pt control. This observed isotopic contrast was attributable to the difference in cooling power between the two gases; under identical pressures, H_2 cools more effectively than D_2, thus explaining the reduced temperature in PdH_x relative to PdD_x.

In our TI 11 campaign we established the reliability of our *in situ* ultrapressure temperature measurements. This campaign notably deviated from the transient heating protocol we established in all other TI campaigns, instead we investigated high temperature equilibrium loading conditions in PdD and PdH at high pressures. So rather than exploiting a transient state where we are heating faster than we can deload, like in our prior campaigns (and TI12), in TI 11 we sustained an equilibrium state up to the melting point of the material. We were interested at looking at this uncharted territory of keeping a sample at higher temperatures with the resulting lower loadings and exploring the result of that tradeoff to see if we observe anything anomalous and interesting. Figure 1 is a schematic of the experimental setup.

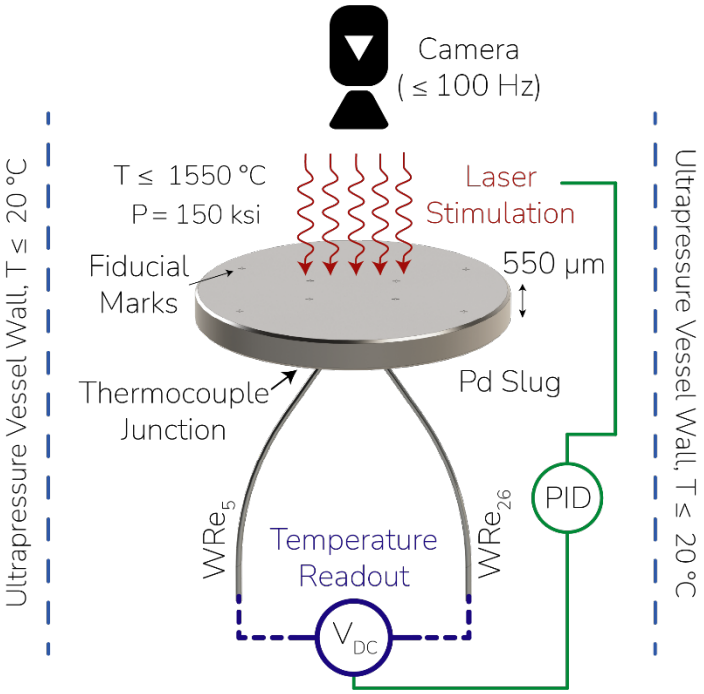


Figure 1: TI-11 Experimental schematic, showing that we are heating by way of a high power laser, in a feedback loop with our thermocouple temperature measurement.

Our TI 11 campaign was the first to melt PdD and PdH under equilibrium conditions, and maintain a melted pool of the PdD/PdH atop the material slug. Initially in our TI 11 campaign we measured a notable difference in the melting point between PdH and PdD. This was later corrected by re-engineering the electrical feedthroughs to remove the inaccurate cold junction compensation offset. Figure 2 below shows the summary of these results, across 13 experiments, showing no significant spread in the measured melting point of PdD and PdH.

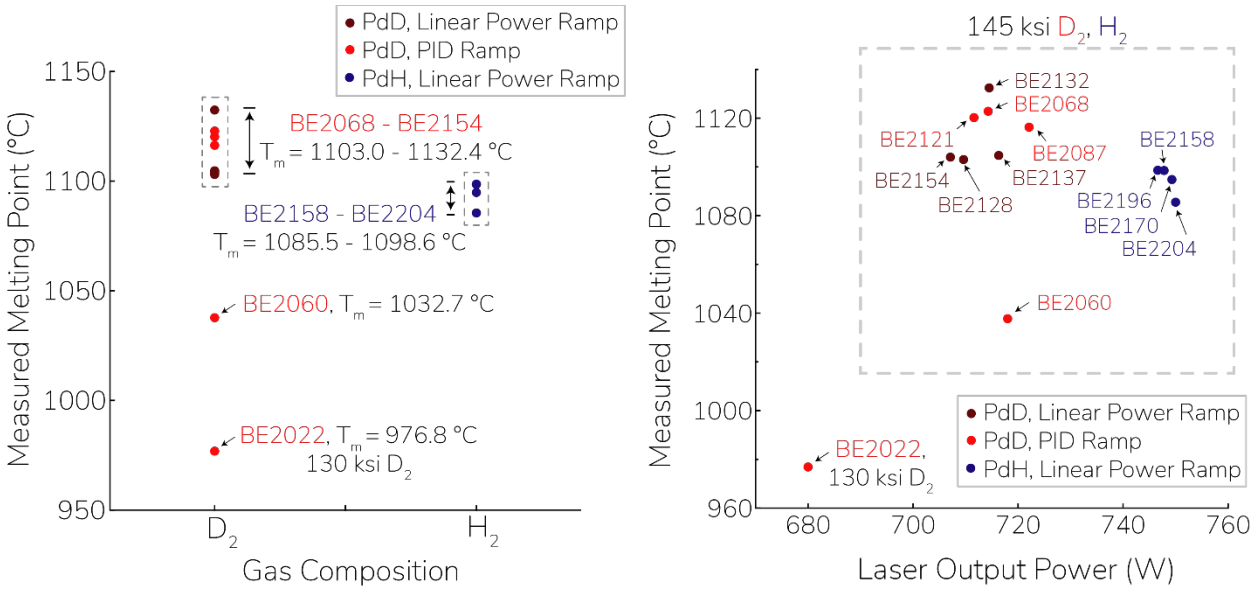


Figure 2: TI-11 campaign summary, showing that we largely eliminated the difference in measured melting point between PdD and PdH.

Unfortunately, no anomalous results were observed in our TI 11 campaign despite holding just below or above the melting temperature for several minutes to hours under ultra pressure conditions. Our last campaign, TI 12, focused on maximizing our resolution, eliminating the need to overcome the rapid cooling of the ultra-pressure deuterium gas to observe an anomalous result.

Thermal Ignition 12 (TI12)

TI12: Advantages and Overview

With experimental data verifying the need for controlled atmospheres (i.e. between the two isotopes), we started our TI12 campaign. This campaign is defined by 3 major aspects relative to prior iterations.

First, TI12 continued to use thermocouples rather than optical measurements to accurately monitor sample temperature. However, in TI11, we observed a significant need for proper cold-junction compensation; in other words, we modified our thermocouples to provide highly accurate temperature measurements, displaying only a few degrees error in the absolute temperature based on known material melting points.

The second major change in TI12 was the introduction of cryogenically cooled samples. This cryogenic cooling allows us to load PdD_x to 100%, stoichiometric conditions even at pressures of only 150 ksi; this was confirmed through multiple mass measurements. Additionally, by cryogenically cooling, we can "lock-in" our loading conditions – to allow for stimulation in

identical environments, TI12 relies on loading the sample under high pressure gas, cooling to <-80 C, and then evacuating the vessel and subjecting the sample to vacuum. Due to the ultra-slow kinetics at these low temperatures, the sample is found to *not* deload from stoichiometry even under these vacuum conditions.

TI12's third, and likely most important, modification was the design and implementation of hermetically sealed microcapsules. Although the environment is void of gas/cooling power differences between isotopes, once the sample is rapidly heated with a laser, it will undergo rapid deloading of the stored gas. This gas expulsion is akin to evaporative cooling: the sample temperature immediately drops. More importantly, this temperature drop occurs at different rates for the two isotopes, causing differences in thermal profiles and peak temperatures. By sealing the 100% loaded PdD_x/PdH_x in a hermetic capsule, we can deliver rapid thermal stimulation and contain the deloaded gas within the capsule. The PdD_x/PdH_x samples will thus cool differently, but the capsule as a whole will not. By monitoring the capsule temperature, there is no net energy loss; we have effectively created an adiabatic (no heat exchange / no mass loss to the environment) system. In addition to minimizing heat loss, these microcapsules are also ideal for measurements of isotopic contrast between H_2 and D_2 gases; we have observed PdH_x and PdD_x to have different optical properties, and the use of a microcapsule allows *equal* absorption of laser power between the two, eliminating this anomalous isotopic contrast. This combination allows for unprecedented resolution of potential AHE events.

Due to the critical importance of each, we discuss the experimental validation of these aspects in more detail below. This is followed by our primary experimental results regarding potential AHE events in this experimental campaign.

TI12: Experiments at cryogenic vacuum conditions

One of the major issues with prior TI campaigns was the gaseous environment. Particularly at ultra-pressure (150 ksi) conditions, H_2 and D_2 have different cooling powers, producing a non-AHE-related isotopic difference between the two. To bypass this, *removing* the environment by surrounding the sample with vacuum is simply the most ideal scenario: The sample is now thermally isolated. To highlight the importance and need for vacuum conditions, consider Fig. 3 below. As shown, not only does high pressure gas cool samples at extremely fast rates (and differing rates for different gases!), but the gas leads to a reduction in peak temperatures at early times. This reduced temperature is indicative of "lost" heat at these timescales. One can easily envision that, should an AHE event occur here, that the AHE-related heat could be lost from this convective cooling.

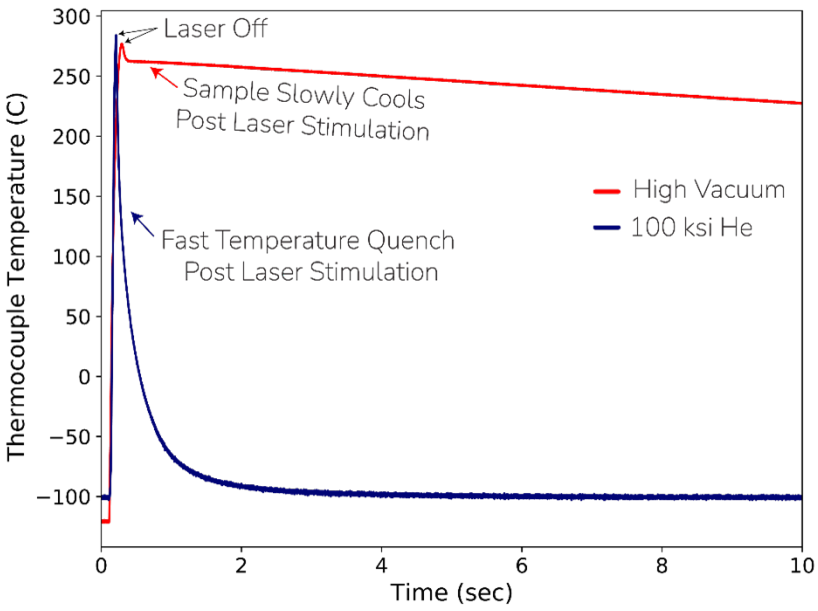


Figure 3: Comparison of temperature profiles following TI laser stimulation of a sample in vacuum (red line) and 100 ksi of He gas (blue line).

While vacuum conditions are thus ideal from an AHE-detection perspective, they are non-trivial to obtain for our samples. Palladium loads (or deloads) hydrogen atoms based on the external pressure: Higher pressures lead to higher loadings. Thus, when *loaded* palladium is exposed to vacuum, the sample will *immediately* begin to deload. As one of the major hypotheses regarding AHE is a need for high loadings, this limitation bottlenecks the scope of our experimental campaign.

To address this, we introduced cryogenic cooling to our sample. By loading the sample at 150 ksi and then cooling to temperatures lower than -80 C, we gain two major benefits in these TI12 experiments. First, at these temperatures, the loaded hydrogen atoms become "locked in" to the Pd lattice; in other words, deloading does not occur at – 80 C. Quantitatively, the loading/deloading kinetics of PdD_x/PdH_x follow the Arrhenius equation, where the rate of deloading, is given by $Rate = A \exp\left(-\frac{E}{RT}\right)$, where A is a constant, E is the activation energy for loading, R is the universal gas constant, and T is the sample temperature. Using the activation energy for Pd-H from literature, Fig. 4 shows the rate of deloading as a function of temperature; one can see that at our cryogenic conditions of approximately 200 K, the kinetics for this deloading go to zero, whereas after stimulation (>1000 K) the rate increases by orders of magnitude.

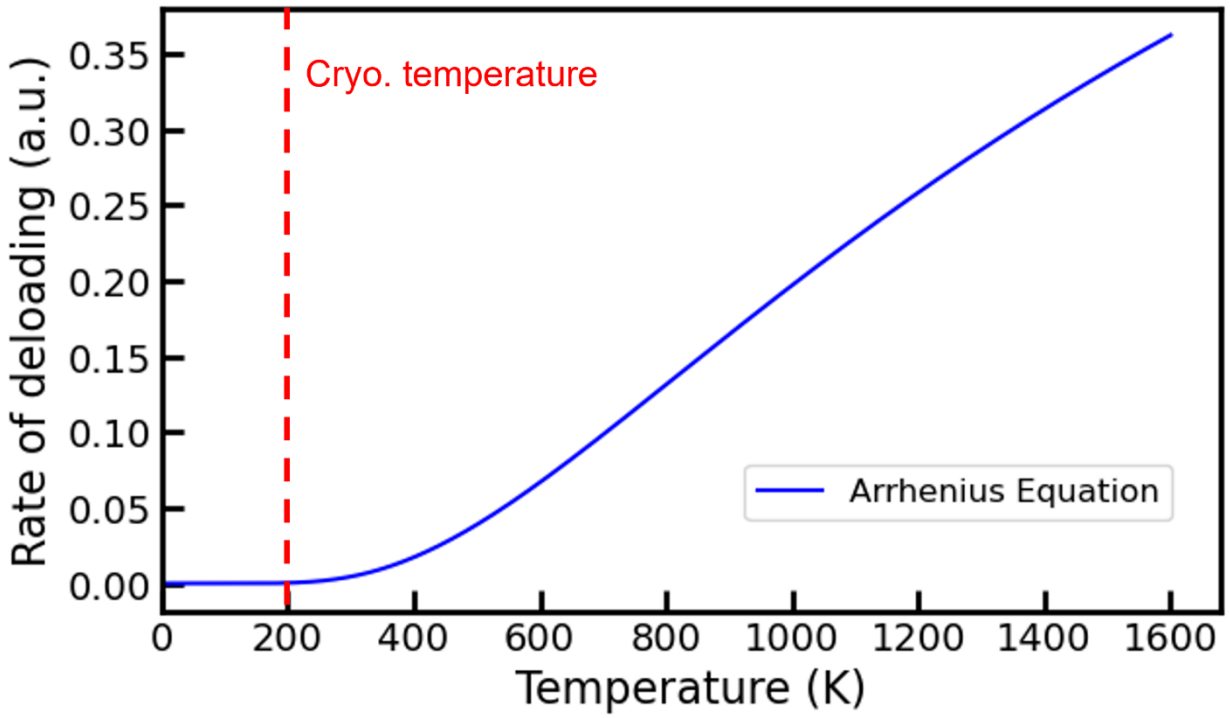


Figure 4: Arrhenius equation for Pd-H highlights the need for cryogenic temperatures to "lock in" loading by reducing the rate of deloading to near-zero at ~200 K.

To further support this, we experimentally verified this critical aspect of our experiment by performing time-resolved loading measurements. As discussed in other sections, the strain, or expansion of the lattice, is directly related to the amount of hydrogen stored within Pd. Thus, we performed experiments in which we monitored the strain during ultrapressure loading, followed by cryogenic cooling and exposure to vacuum. These results are highlighted in Fig. 5 and Fig. 6.

As shown in Fig. 5, we perform a control of these strain measurements by exposing the Pd sample to hydrogen at 150ksi and -100 C for over 24 hours: This provides a direct measure of the error associated with a sample in equilibrium conditions. One can see in the early times of these data that the strain increases upon pressurization and remains constant during this exposure period. Then, after multiple hours and the sample maintained at -100 C, we remove gas from the system and expose the sample to vacuum. As can be seen, for nearly 17 hours, the strain remains unchanged in these cryogenic vacuum conditions. Then, upon heating, deloading occurs and the strain begins to drop.

Figure 6 provides a direct comparison of these temperature-dependent deloading kinetics by comparing the time-resolved strain measurements at -100 C vs. 50 C. At 50 C we observe a gradual reduction in strain while the -100 C sample remains constant. These kinetic rates are in reasonable agreement with the expected Arrhenius-like deloading rates calculated in Fig. 4. Further, we provide a comparison to a Pt control under identical pressure/temperature conditions to highlight

the degree of strain induced by the conditions in the absence of loading; this effect is observed to be negligible relative to the loading-dependent strain in Pd.

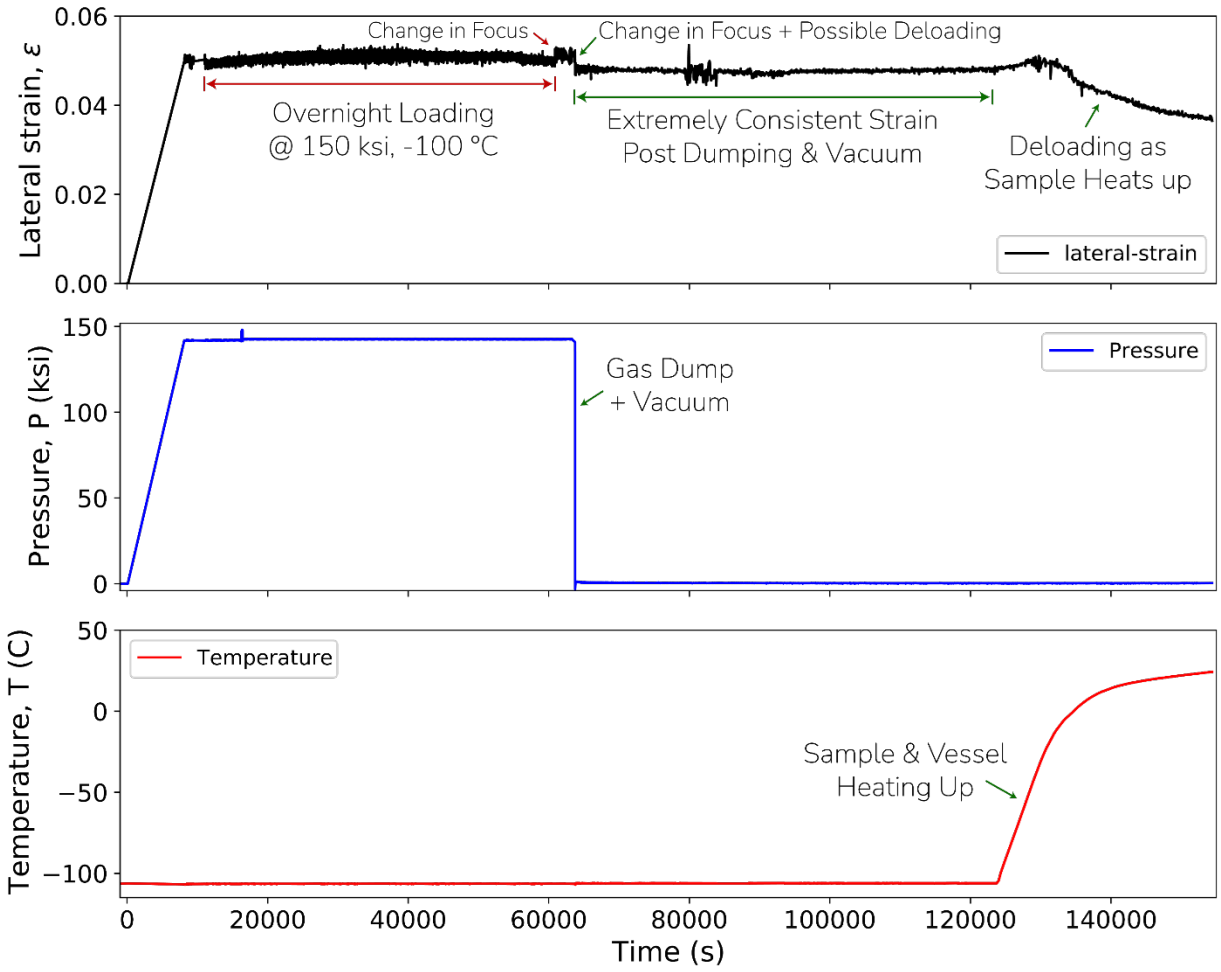


Figure 5: (Top) In-situ strain measurements under varying conditions of gas exposure and temperature. (Middle) Pressure during these strain measurements (Bottom) Temperature during these strain measurements.

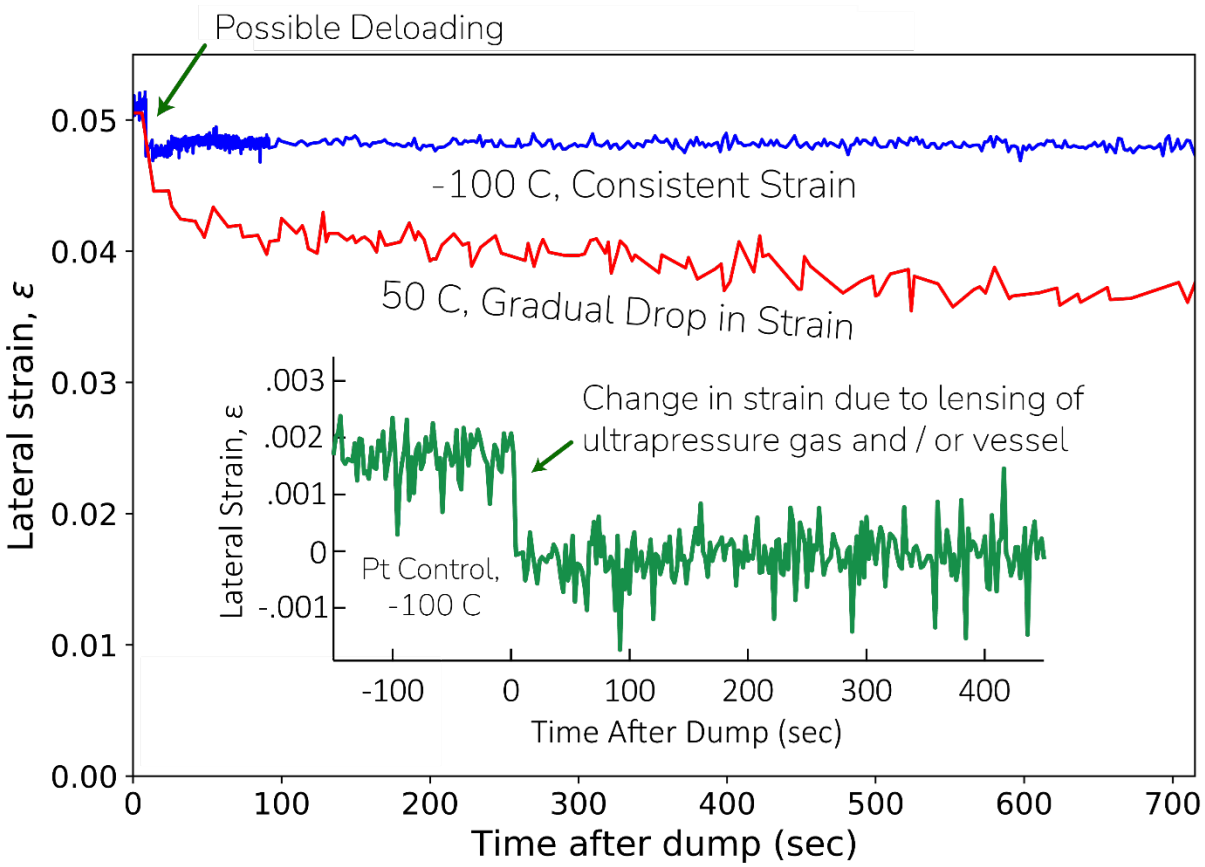


Figure 6: Sample strains for Pd-H exposed to 150ksi hydrogen at -100 C (blue line) and 50 C (red line). A Pt control (green) under identical conditions is shown to highlight the degree of strain induced by pressure alone in the absence of loading.

TI12: Adiabatic microcapsules and thermal isolation

The second major revision introduced for TI12 was the use of microcapsules within our ultra-pressure vessel. These hermetically sealed capsules prevent gas from leaving the system under study and provide adiabatic conditions (no heat loss/mass loss). In doing so, "evaporative" cooling is eliminated from our analysis and thus eliminates anomalous isotopic differences associated with such cooling.

The necessity for *true* hermetic sealing of these capsules, which is an extremely non-trivial task to achieve, can be observed in Fig. 7 below. As shown in Fig. 7, in earlier experiments, we attempted to solve this cooling issue by "capping" the sample in thin, hydrogen-impermeable membranes, namely silicon nitride films. While this approach *significantly* reduced the amount of gas that escapes from the PdDx upon thermal stimulation, we continued to observe a large degree of excessive cooling due to the small amounts of gas molecules escaping. In other words, we experimentally determined that *any* gas escaping from the measurement volume leads to errors in

temperature measurements. Thus, hermetic sealing, or sealing in which no gas escapes, is a *requirement* to test AHE to the highest resolution feasible.

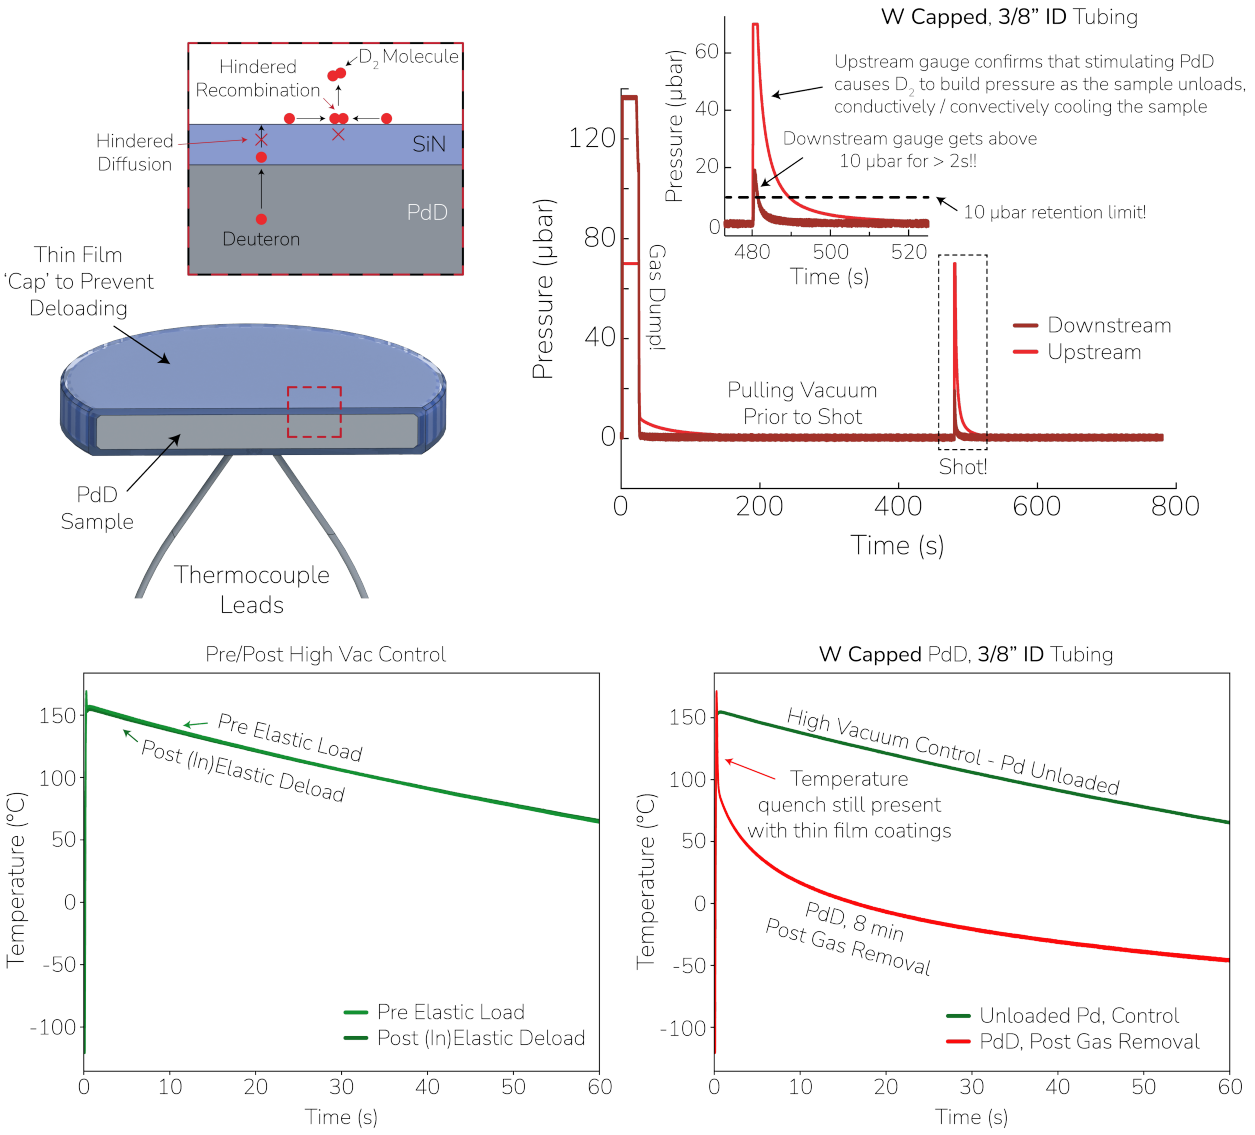


Figure 7: (Top left) Schematic of our SiN-capped Pd samples to attempt to avoid gas deloading. (Top right) Measurements of the gas pressure with the SiN cap relative to the amount of gas lost upon pulling vacuum. (Bottom left) Vacuum controls of unloaded Pd. (Bottom right) Comparison of these vacuum controls to an identically stimulated PdD_x slug. As shown, the SiN cap reduces *quenching, but does not solve the problem.*

The iterations of these microcapsules are shown below in Fig. 8. After multiple experimental cycles, our final design, Gen. 5, provides multiple modifications from earlier iterations that are worth highlighting. First, the capsule is designed in such a way to *maximize* heat transfer between the capsule and the loaded Pd; in TI12, we stimulate the top surface of these capsules and rely on the heat transfer to the underlying Pd to reach the temperatures of interest. By using capsules made

of a MoRe alloy ($MoRe_{0.47}$ wt%), this heat transfer is greatly improved by the high thermal conductivity of MoRe. Notably, MoRe also has a melting temperature of nearly 2500 C, allowing us to stimulate to temperatures in excess of Pd's melting point; MoRe also does not undergo hydrogen loading, thus preventing convective cooling due to its subsequent deloading upon heating. Second, we added and continued improving on various thread designs as a "first-defense" for mechanical stability once the internal pressure spikes following thermal stimulation. This is then supported and made hermetic through an in-house method of laser welding the two components after threading; we found that significant care must be taken in this laser welding process due to the formation of microscale voids that can occur and ultimately lead to premature failure. Notably, when *any* voids were observed following welding, the capsules leaked upon thermal stimulation, despite passing helium leak checks prior to stimulation; in other words, these voids "rupture" due to the excessive heat load supplied by our laser.

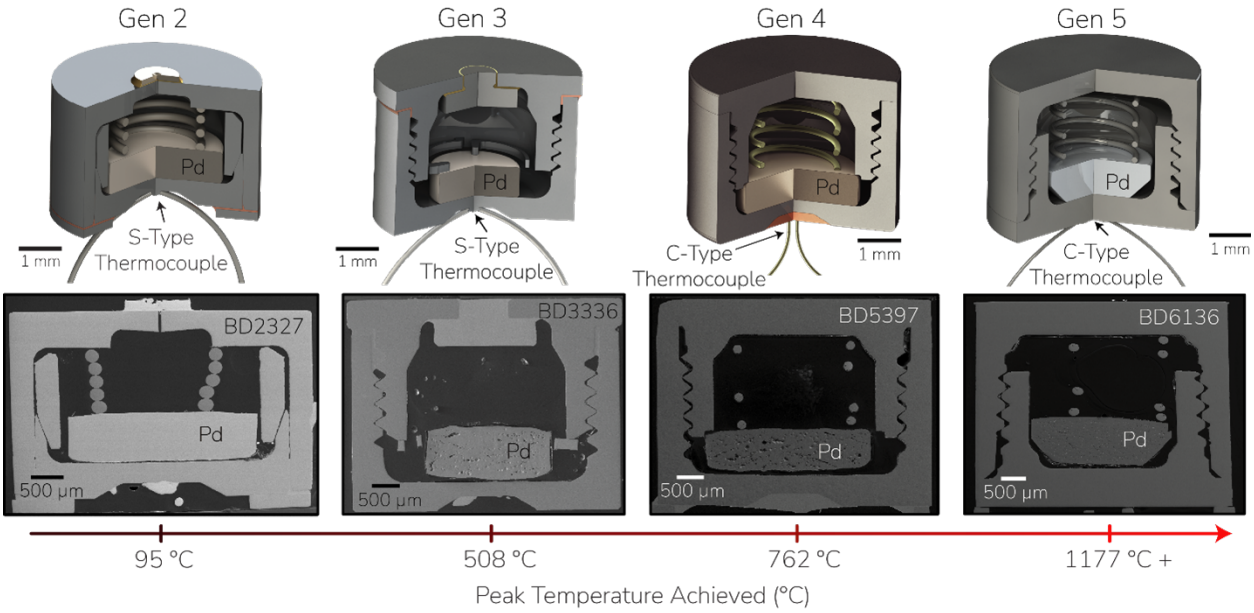


Figure 8: Experimental iterations and improvements on our microcapsules over multiple generations.

In addition to the microcapsule design itself, significant effort was made to ensure true thermal isolation between this capsule and the surrounding pressure vessel. As there is no way to levitate a capsule within a vacuum environment, a mechanical retention is required. This retention must be designed to minimize heat flow between the microcapsule and the pressure vessel itself. To understand the importance of this retention, consider Fig. 9, which shows the iterations of various retention designs and their effect on the temperature distribution. As shown, our final Archimedean spiral provides near the same isolation as expected by true, "perfect" thermal isolation, with only a few percent loss of heat (rather than orders of magnitudes with the prior design!).

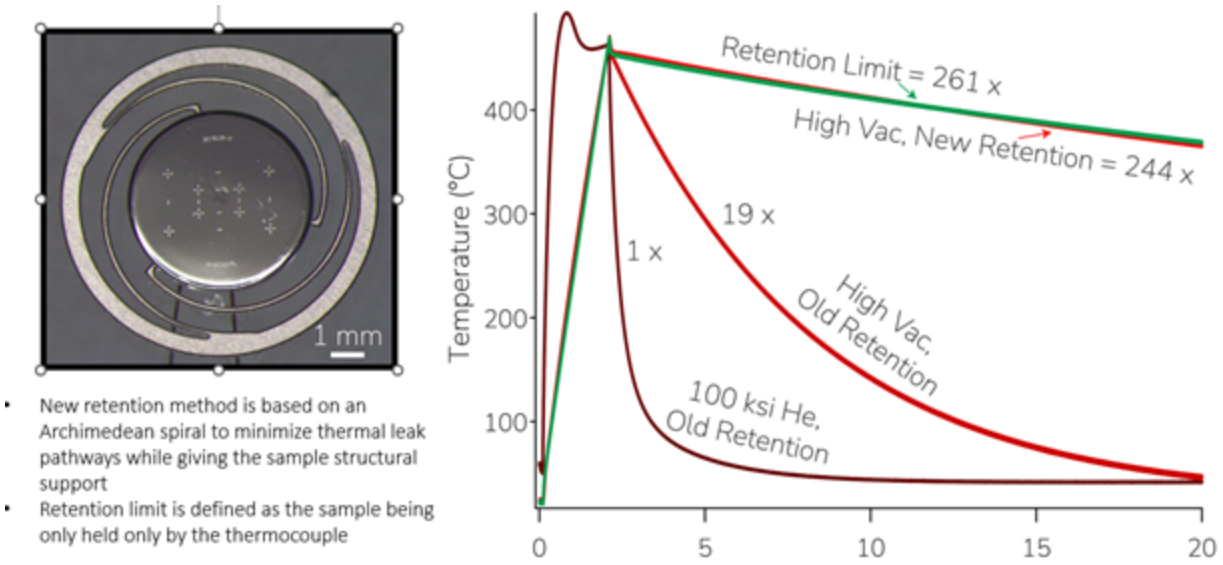


Figure 9: (Left) Final microcapsule retention design for TI12 experiments. (Right) Comparisons of temperature profiles with varying retention designs and "perfect" isolation for comparison. This Archimedean spiral design is within a few percent of the achievable limit, in comparison to our prior retentions that were "leaking" orders of magnitude more heat.

TI12: Experimental protocol and results

With these designs in mind, consider now our experimental protocol: Palladium is placed within a microcapsule and the entire microcapsule assembly is contained within a 150ksi pressure vessel via the Archimedean spiral retention. At this stage, the capsule is entirely sealed (determined by He leak testing) except for a small, 100 μm diameter hole on the front surface; this hole allows gas to enter the capsule for loading.

Prior to loading the Pd, we subject the sample to vacuum, reaching approximately 10^{-7} Torr. With an unloaded Pd sample, we stimulate the sample with a given laser profile (this is varied through our experiments, but generally is a 10 ms pulse that reaches temperatures in excess of Pd melting) and measure the resultant temperature profile. This temperature profile acts as our control. This control is performed for *all* D_2 and H_2 experiments; small changes in microcapsule or Pd mass can lead to large changes in the observed temperature. Thus, our unloaded sample acts as the ideal comparison.

With the experimental control now obtained for a given sample, we load the Pd sample by exposing it for long durations (~hours) to 150ksi D_2/H_2 to ensure homogeneous loading throughout the bulk of the material. Once the Pd is homogeneously loaded, we *slowly* cool the vessel from 25 C to -100 C; by moderating the rate of which we cool the sample, there is sufficient time for *increased* loading to occur at these cryogenic temperatures. These cooling rates are confirmed to be valid for achieving stoichiometric, 100% loading by performing mass measurements of the microcapsule

before and after cooling. This is done by mimicking our TI12 protocol except for stimulation: The capsule is simply removed, and its mass compared to that without loading.

With the vessel now at -100 C, deloading of the sample is inhibited, allowing us to introduce vacuum to the vessel. This first step removes the 150ksi gas and replaces the environment with ~200 psi of D_2. In this relatively low-pressure environment, we perform *in situ* laser welding of the 100- µm hole to finalize the hermetic sealing of the capsule. The slight pressurization, rather than true vacuum, eliminates debris formation from evaporation of the melt pool during this weld step; we observed that true vacuum leads to our laser window becoming obscured and causing anomalous changes in temperature profiles due to reduced light reaching the microcapsule.

With the capsule now sealed, we remove the external 200 psi of D_2/H_2 and again pull true vacuum on the system, reaching identical vacuum to our original control stimulation conditions. At this point, the Pd is 100% loaded with a hydrogen isotope and under identical environmental conditions to our pre-experiment control. Then, we apply same the laser stimulation as our control and compare the two temperature profiles.

A schematic of our final microcapsule design along with representative data of this experimental procedure is shown in Fig. 10 below. In this case, the two stimulation shots are observed to be within 50 C of each other; this small reduction for PdD is also observed for PdH and is attributed to the change in heat capacity due to the increased mass (loaded hydrogen) within the microcapsule volume relative to the control shot. One major takeaway from these data is that the control and loaded PdD cool at *identical* rates, a key benefit of the hermetically-sealed microcapsules that we have designed.

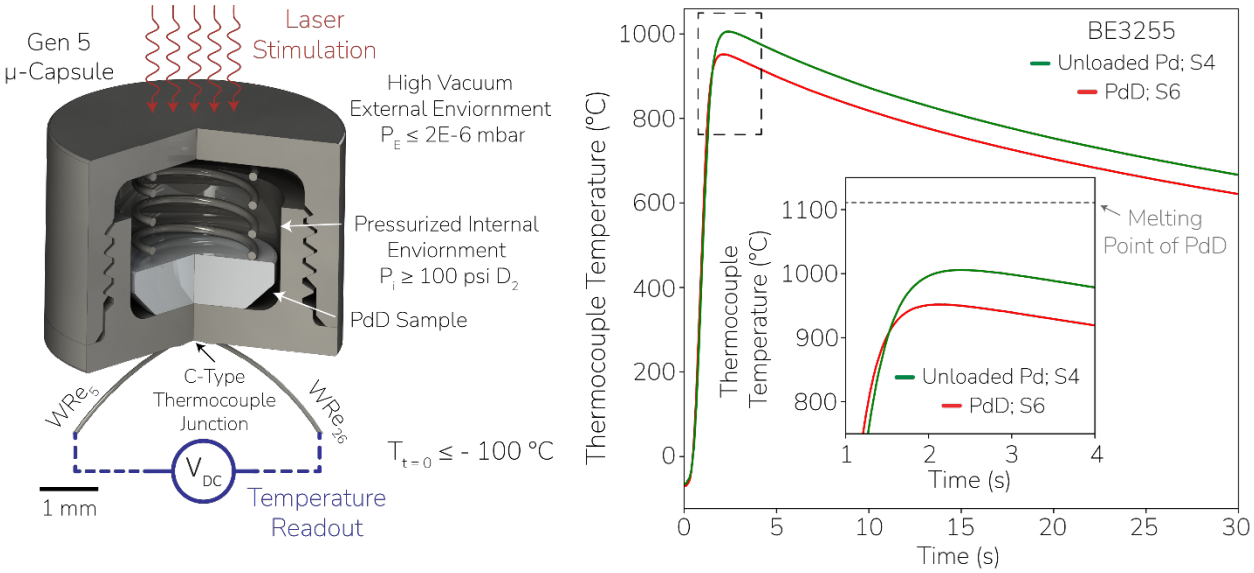


Figure 10: (Left) Schematic of our TI12 microcapsule with laser stimulation and thermocouple readouts. (Right) Representative temperature profiles following stimulation for (green) unloaded Pd, which acts as our control, and (red) stoichiometric PdD after our loading and capsule sealing protocol.

To further extend the stimulation conditions to the "maximum limit," we continued to increase our stimulation conditions until we exceeded the melting temperature of PdD. In Fig. 11 below, we show both the pressure profiles, representative of the degree of evaporative cooling that is experienced by the microcapsule, along with the temperature profiles following stimulation. As observed with lower stimulation conditions, loaded PdD continues to show *lower* temperature excursions than our unloaded Pd control, suggesting an absence of AHE heat input.

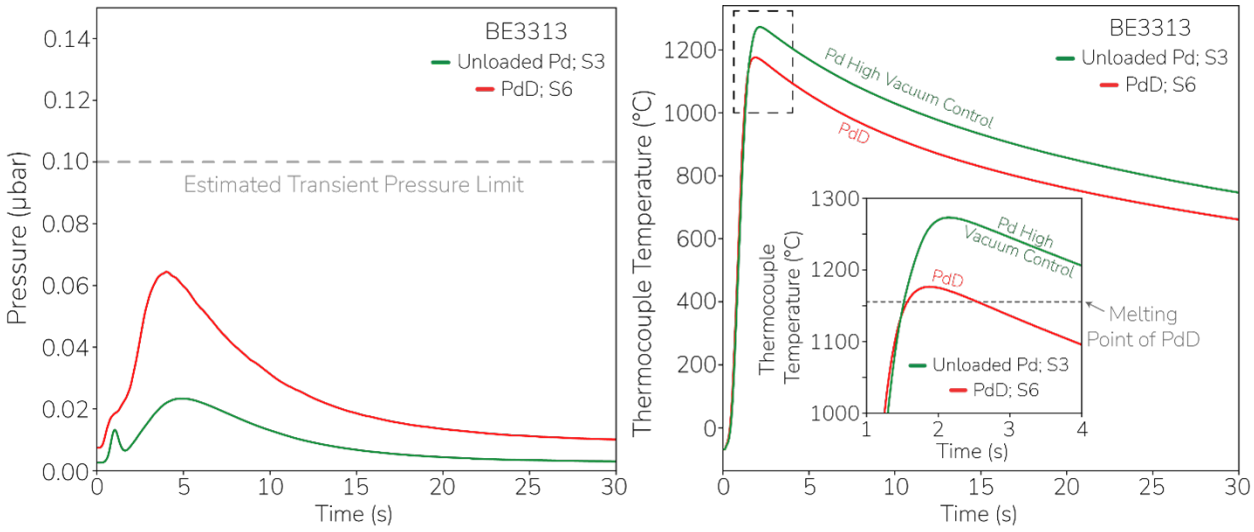


Figure 11: (Left) Transient pressure data following stimulation for (red) PdD and (green) our control. The slight increase in PdD is due to adsorbed gas species as the sample is held at cryogenic conditions. (Right) Temperature profiles following stimulation; both profiles are observed to be in excess of the melting point of PdD. Again, we observe PdD having a lower *temperature rise than the control shot, suggesting a lack of excess heat production.*

To verify melting was achieved, we cross-section the microcapsule after the experiment, allowing us to observe the Pd structure after stimulation was performed. An example of this SEM data is shown below in Fig. 12. The signature of melting here is the large deformation of the Pd itself: The palladium is found to have melted and conformed to the surrounding microcapsule, rather than retaining its original shape. In addition to this clear observation that melting occurred, we note the internal microstructure of the Pd contains large (~25 μm) voids due to the rapid deloading of hydrogen gas during stimulation. These features are consistent in both PdH and PdD experiments.

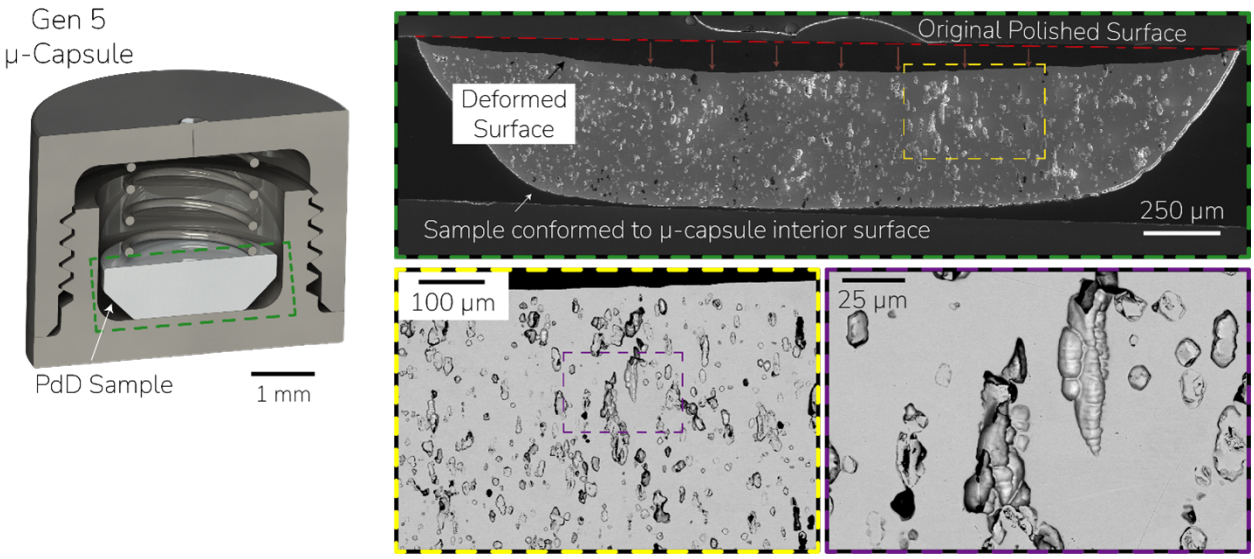


Figure 12: (Left) Microcapsule with the green dashed line representing the cross-sectioned area. (Right, top) SEM micrograph of this dashed region. The Pd is observed to have conformed to the surrounding capsule and collapsed inwards, highlighting that melting occurred during laser stimulation. (Right, bottom) Micrographs of the voids that are formed due to the rapid release of hydrogen gas. These microstructural features are observed for both PdD and PdH.

TI12: Overview and conclusions

Experiment ID (BE####)	**Material (at%)**	**Peak Temperature**	**Experimental Resolution (Ratio of AHE active sites vs available sites)**	**AHE?**
BE3323	$PdD_{1.0}$	1204 °C	2.9E-8	No
BE3861	$Pd_{75}Rh_{25}D_x$	1186 °C	9.0E-8	No
BE3801	$Pd_{75}Ni_{25}D_x$	901 °C	7.0E-8	No
BE3822	$Pd_{75}Ag_{25}D_x$	879 °C	6.2E-8	No
BE3829	$Pd_{75}Pt_{25}D_x$	1204 °C	4.7E-8	No
BE3845	$Pd_{92.5}B_{7.5}D_x$	945 °C	6.1E-8	No

BE3922	NiD_x	1086 °C	6.8E-8	No

Table 1 : Key TI12 experiments

Across our TI12 campaign, we repeated the above protocol and analysis with several materials, including many Pd alloys, to investigate the potential role of alternate atomic species. Table 1 shows this data, along with the experimental peak temperatures (determined based on the known melting temperatures: all samples are stimulated in excess of melting). In each case, we determine the experimental resolution in terms of how many AHE events would need to happen to be observable. **In the worst-case scenario, 1 in 100 million potential AHE sites would have to undergo AHE to be observable**. In other words, our TI12 campaign places a lower limit and suggests that *if* AHE were to be occurring, it is happening in less than 1 in 100 million potential cases.

To further support the accuracy of our TI12 apparatus and the above claim regarding the lack of AHE, we confirmed our resolution by using a known, thermally activated exothermic reaction (i.e., a non-nuclear means of generating "excess" heat). To confirm our resolution (i.e.., adiabatic assumption), and experimentally generate error bounds, we performed a TI12 experiment using a 50 / 50 at% mix of Pd and Al powder in place of PdD_x, where the creation of the intermetallic alloy is an exothermic reaction; this is akin to thermite without relying on oxygen transfer and occurs under high vacuum. This control experiment reacts the $Pd_{50}Al_{50}$ alloy to completion, with an expected exothermic reaction of 175 kJ/mol. Based on the known mass added, we anticipate 19.5 J of generated additional heat from the exothermic reaction. Indeed, we measure 21.9 J as "excess" heat in our experiment, or approximately 12% error between the measured and expected value.

Across all TI campaigns, we have rigorously tested nearly 12 orders of magnitude in pressure conditions, ranging from ultrahigh vacuum to 1 GPa. We've exceeded the melting temperature of PdD and PdH, directly measuring 100% loading for both isotopes. We have also expanded our search to other materials and alloys. We have done all of the above while maximizing our system resolution, requiring as little as 10^{-6} per-cent of the AHE sites to be active for us to unambiguously detect AHE, exceeding the resolution of the best available calorimeters available. This parameter space is highlighted below in Fig. 13, but concludes that, under these conditions, **we unfortunately have not seen any discernible evidence of AHE in our Thermal Ignition experiments**.

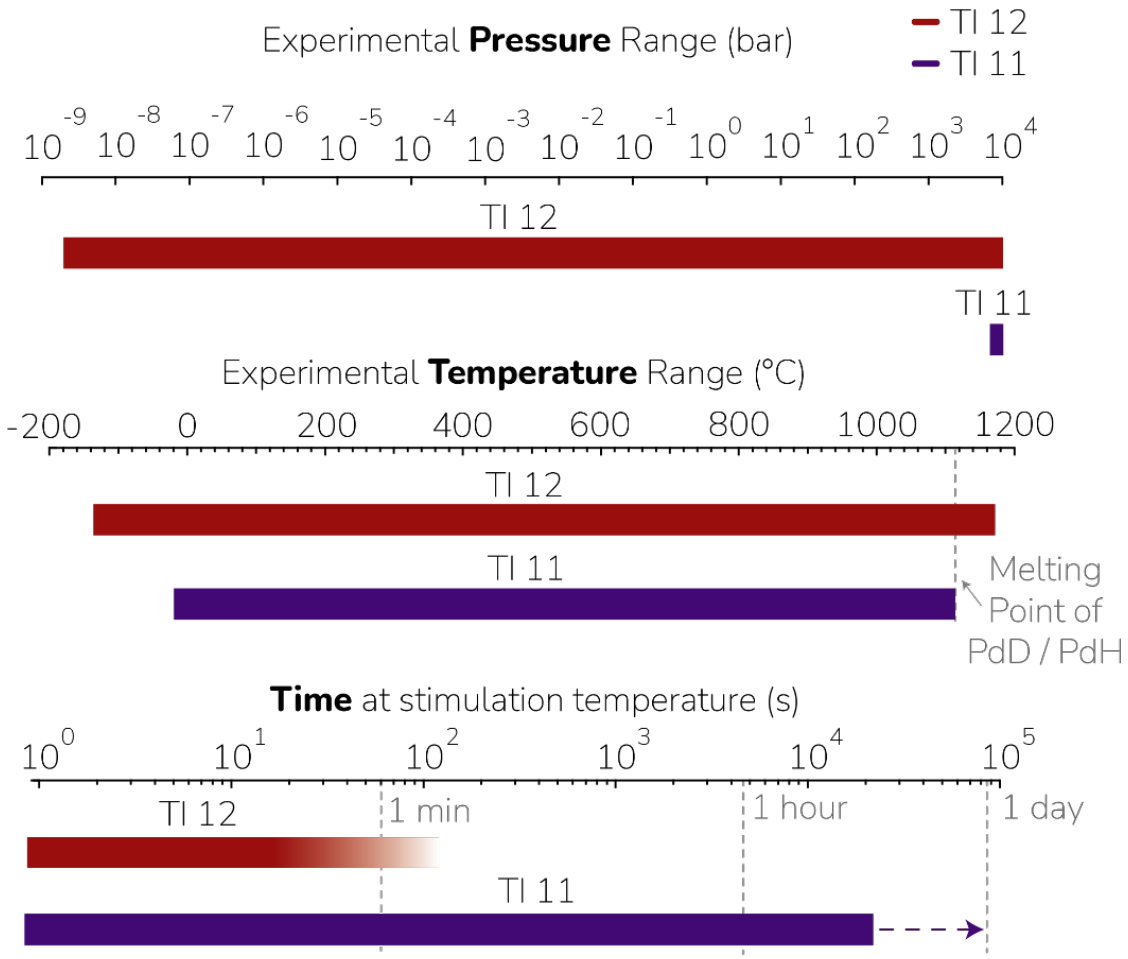


Figure 13: Parameter space explored across TI campaigns. Multiple orders of magnitudes on various experimental conditions were varied, in addition to the material list shown in the above table.

Staring

Introduction

The Staring project had two main types of experiments. The first type was long-duration, continuous thermal imaging of loaded samples to look for evidence of LENR in the form of hot spots on the surface. The second type was archaeology experiments where post-facto, ex-situ imaging of a sample could look for evidence that LENR had occurred with single-event sensitivity. For both types of experiments, we had the ability to thermally stimulate the sample with a laser with variable powers and pulse durations as a way to potentially induce LENR.

Following Planck's law for blackbody thermal emission, any object at a given temperature will emit thermal radiation whose intensity and wavelength distribution depend on the object's temperature and emissivity. Within our temperature range of interest, roughly room temperature to 1000C, the best wavelength band to observe hot spots is the mid-IR band, spanning wavelengths from around 1.5-5 μm. At shorter wavelengths, such as with near-IR and visible light, the total amount of emitted light for a surface within our target temperature range is reduced compared to the mid-IR, which would limit the amount of signal we'd be able to capture. This would make it more difficult to resolve signal from noise, and make it impractical to satisfactorily resolve temperature differences within our desired temperature range. For longer wavelengths (e.g., far-IR), the total amount of thermally emitted light would not be a concern, but there would already be a substantial background signal at room temperature and the fractional increase in emitted photons vs temperature would be smaller, which would limit the contrast between a surface at room temperature and an elevated temperature. Balancing the desire for a strong enough signal and an ability to resolve temperature differences within our desired range, the mid-IR wavelength band provides the most ideal solution. For these experiments, we used mid-IR cameras with the ability to detect light in the 3-5 μm or 1.5-5 μm bands. Capturing the signal over the full band the camera could detect maximized our total signal and the probability that we'd be able to observe a thermal event.

Our long-duration thermal imaging experiments provided some unique advantages over other types of detection methods and allowed us to test a few different hypotheses about the triggering and detection of LENR. First, performing imaging as opposed to some form of detection integrated over the volume of the sample (as with calorimetry, for example) allowed us to have spatial resolution over the sample and have equal sensitivity to LENR events over every part of the sample surface. This would allow us to directly observe if LENR occurred in small localized areas of the larger sample, and possibly shed light on the conditions necessary for LENR by comparing areas where we observed hot spots to areas where we did not. In addition, the fast frame rates of our camera (we used the fastest IR cameras available commercially) would allow us to see thermal dynamics on a few-ms time scale. Together, these properties would allow us to detect LENR events

smaller in magnitude and spatial extent, and shorter in duration than could be seen in a spatially-integrated and time-integrated detection scheme such as calorimetry. Lastly, the ability to perform automated, real-time filtering of images to isolate those with possible thermal signals allowed us to perform continuous imaging for hours at a time without compromising our spatial or temporal resolution. This allowed us to test the hypothesis that a sample needs to remain loaded for some extended period of time (hours to days) before LENR typically occurs. Any image with a potentially interesting thermal signal would be automatically flagged by our live filter and saved for additional post-facto analysis.

For our archaeology experiments, we tested a variety of different sample types over the lifetime of the staring project. The main assumption we operated under was that a single LENR event would release 24 MeV of thermal energy instantaneously into the surrounding sample volume, which would alter the sample within an area localized around the event. Each sample type was designed with the criteria that small numbers of LENR events would cause clear, permanent changes to the sample that would be observable in some form of post-facto scanned imaging.

There were two main classes of samples we measured for archaeology. The first were samples where an LENR event would cause observable damage or destruction. The primary types of samples that fell under this category were nanopillars and patterned nanofilms. For these types of samples, we performed SEM (scanning electron microscopy) scans over the entire active area of the sample pre-experiment and post-experiment to look for changes that occurred as a result of the loading and laser stimulation. Imaging the surfaces pre-experiment is necessary for us to eliminate potential false-positives where some damage was pre-existing prior to loading and laser stimulation of the sample.

The second class of archaeology samples were those where the heat from an LENR event would alter the polarization of magnetic domains in a layer of magnetic material above or below a Pd layer. For these magnetic archaeology experiments, we scanned samples pre and post-experiment with MFM (magnetic force microscopy) to image the polarization of the magnetic layer, and again looked for changes that we observed in the post-experiment images compared to the pre-experiment images. These experiments are discussed in detail in their own section of this report and will not be covered in detail here.

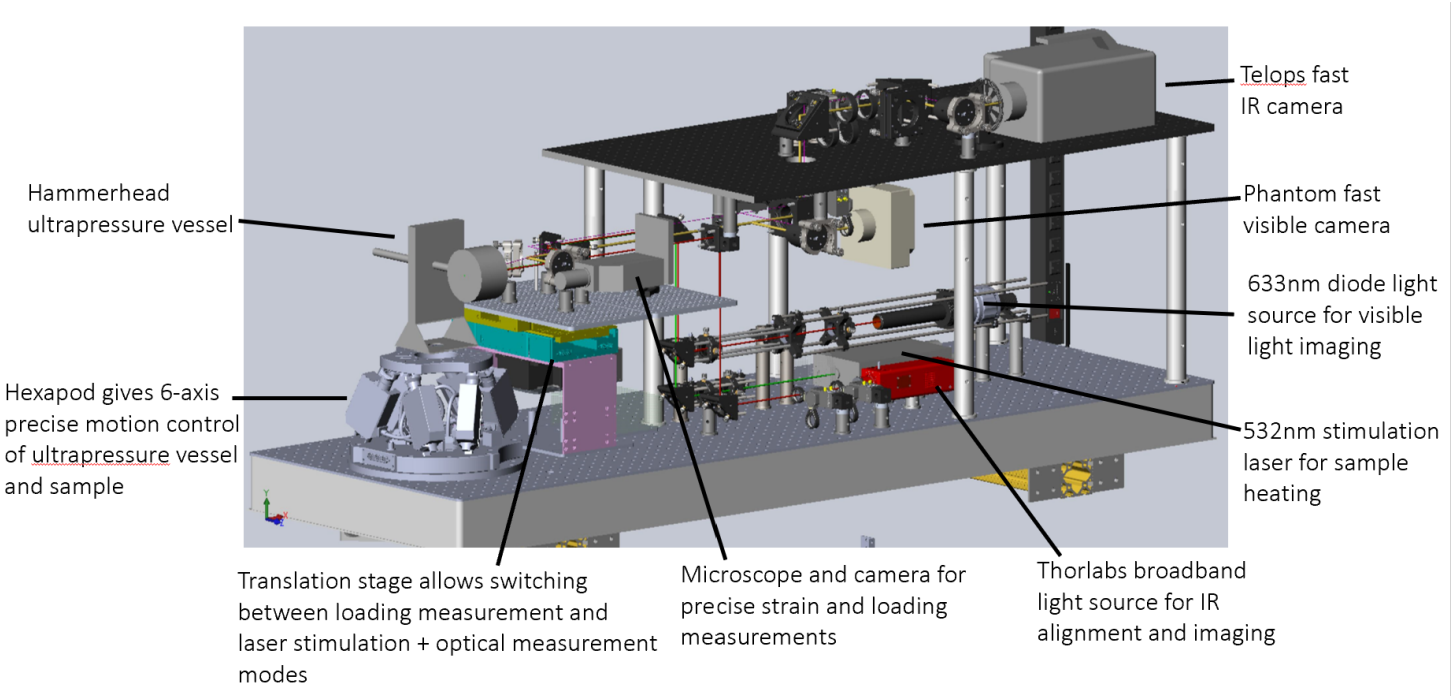
For all types of experiments, we could load a sample with either deuterium to attempt to detect LENR, or with hydrogen as a control. Any potential evidence of LENR would need to be able to be consistently replicated in deuterium but not hydrogen.

Hardware and Apparatus

Figure 1 shows a diagram of the optical setup used for staring experiments with some of the key components indicated.

Our hammerhead ultrapressure vessel contained the sample and allowed pressurization with a variety of gases (H_2, D_2, He, N_2, Ar) from vacuum to 145ksi, and was able to hold temperatures from -80C (used when cryogenically cooling a sample) to 380C (used when elastically loading a bulk sample). The ultrapressure vessel is mounted on a 6-axis hexapod, which allows precision control of the sample position over the 3 principal spatial axes and 3 principal tilt directions. This ensures that we can simultaneously align a stimulation laser onto the desired region of the sample, illuminate the sample with visible light for visible imaging, and collect the thermally emitted IR light for thermal imaging.

Figure 1 - A diagram of the Staring experiment optical apparatus



The hammerhead vessel could be configured with an ultrapressure sapphire window on one side for optical access to the sample, and electrical feedthroughs on the other side. Our primary use for the electrical feedthroughs was for RTDs on the surface of magnetic archaeology calibration samples, which were used to correlate laser pulse conditions to temperature achieved at the surface of the sample. Other archaeology samples such as patterned nanofilms and nanopillars, and samples typically used for long duration thermal imaging, did not require the use of the electrical feedthroughs.

A precision translation stage allowed us to move a microscope body containing an objective lens and compact Pixelink camera into or out of place in front of the hammerhead vessel. When in place, this microscope could image the sample with better than 5 μm spatial resolution. This higher-resolution imaging mode was used to perform strain measurements on bulk samples to characterize their loading, and was frequently used to closely investigate the surface of the sample and optical windows for any contamination or damage before, during, and after the experiment.

When the microscope was moved out of place, we could perform laser stimulation simultaneously with visible and IR imaging. A gold OAP (off-axis parabolic mirror) was moved into place in front of the hammerhead vessel to collect reflected and emitted light from the sample. By using a gold reflective optic instead of a lens to collect light from the sample, we were able to collect the entire range of wavelengths from visible to mid-IR without chromatic aberrations, allowing simultaneous visible and mid-IR imaging without large chromatic distortions. After the OAP, a germanium beamsplitter separated the visible light and IR light so that each wavelength range could be directed to the appropriate camera.

We used a Phantom S640 visible-light camera for live monitoring of the sample during experiments and for alignment of the stimulation laser onto the desired area of the sample surface. Our imaging system with this camera had a spatial resolution of around 15 μm. Continuous monitoring during experiments was required to ensure the sample remained in position and to check for contamination or any changes to the sample surface that might occur as a result of loading or laser stimulation, such as cracking, roughening, darkening, or blistering. A 633nm diode light source served as the illumination for this camera. The imaging optics used with this camera produced an image with 1X magnification. The clear aperture of the hammerhead vessel's ultrapressure window was around 5mm in diameter, and the camera sensor was significantly larger than this, so we could easily resolve the entire sample area that was possible to view through the ultrapressure window of the hammerhead vessel.

A Thorlabs SLS202L broadband light source allowed us to illuminate the sample with mid-IR light, which was used to align the IR sample image into the IR camera. This ensured that if the sample became hot enough to emit IR photons, we would capture the emitted light in the IR camera.

The fast-framerate IR cameras we used for thermal imaging were the Telops Fast-M2K and Fast M3K models, which provided IR imaging in the 3-5 μm or 1.5-5 μm wavelength bands, respectively. We alternated between the two models over the lifetime of the project because one camera would occasionally need to be repaired, and swapping in the other camera avoided prolonged periods of experimental downtime. Our imaging system with these cameras had a spatial resolution of around 60 μm when used in a 1X magnification configuration, and a spatial resolution of around 15 μm when used in a 4X magnification configuration. These cameras could operate at a maximum framerate of 2000-3000fps. For some long-duration acquisitions, we operated at 2000fps near the camera's operation limit, but running the camera at high framerates for long periods of time could cause heating of the sensor and camera body that would increase noise, so for some experiments we operated at a more conservative 200fps. We typically operated this camera with an active sensor area of 5.76mm x 5.76mm. With a 1X magnification system, we could resolve the full area of the sample in the sensor, since the hammerhead clear aperture was around 5mm. When we used a 4X magnification system, the field of view of the sensor was correspondingly reduced to around 1.44mm x 1.44mm, so we could not image the entire sample area at once.

For most of the lifetime of the project, we used a 532nm, 100W stimulation laser from IPG. This laser was used for all of our long-duration IR imaging experiments and most of our archaeology

measurements. This laser had a tunable average power from 1W to 100W and output pulses at a fixed pulse duration of 4ns and fixed frequency of 40MHz. The laser can be triggered on and off electronically, but the switching times are on the order of a few ms for the laser to fully ramp on or switch off, meaning that the minimum time we could have the laser on would be between 1 and 10ms. It would not have a constant power during this time as it ramped up to turn on and ramped down to switch off.

Since the deloading and reloading times of Pd thin films and nanostructures like nanopillars was likely shorter than a few ms, we wanted to achieve faster on/off modulation of the laser. To attain this, we used a Pockels cell, a device commonly used in ultrafast lasers for fast modulation. The Pockels cell contains an optical crystal with a large electro-optic coefficient, meaning that the index of refraction for light polarized along a particular crystal axis depends on the applied voltage across the crystal. This property allows the crystal to be used as a voltage-dependent waveplate. For example, if incoming light is polarized properly relative to the crystal axes, then when no voltage is applied across the crystal, the polarization of light will remain unchanged, but if a sufficiently large voltage is applied, the polarization of the light will be rotated 90 degrees. Placing a polarizer after the Pockels cell, we can achieve near-zero transmission when the voltage is switched off, and near-100% transmission when the voltage is switched on. With this scheme, the speed of on/off modulation of the laser is only limited by the speed at which the voltage across the crystal can be switched on and off. We used a Pockels cell with a BBO crystal designed for our laser wavelength of 532nm, which enabled us to gate the laser on and off with durations ranging from 50ns to 5ms.

In the final few months of the project, we used a 1070nm, 100W laser for stimulation while the 532nm laser was being repaired. This laser was used for some of our archaeology experiments on patterned nanofilms and thermal history samples. This laser could output pulses with a tunable pulse duration and frequency. When using this laser, we used a pulse duration of 4ns and a frequency of 3.39MHz, or a pulse duration of 50ns and a frequency of 320kHz (since the total energy output of the laser can't exceed 100W, the pulse frequency drops as the pulse duration increases). We modulated this laser on and off with an arbitrary function generator with time windows ranging from 10us to 3ms.

Since the laser essentially provides thermal stimulation, the particular wavelength used should not have made a meaningful difference for the triggering of LENR.

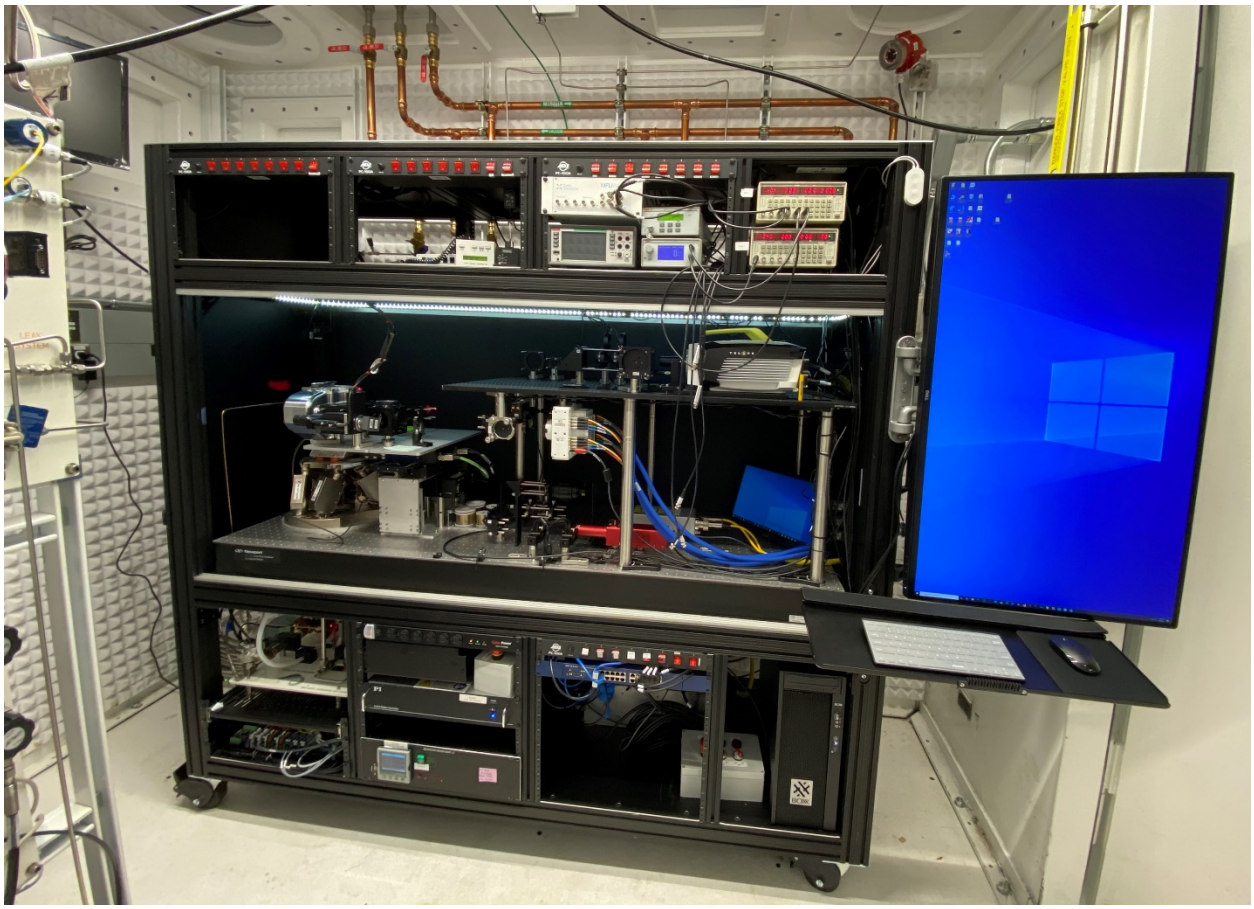


Figure 2 - A photo of the full Staring experimental apparatus.

Figure 2 shows the full Staring experimental apparatus in one of our ultra-pressure chambers. The powerpack that served as the ultra-pressure source is out of the field of view of the image. The optical setup sits within a compact, mobile enclosure designed and constructed by our optics team that leaves room for a powerpack in the same ultra-pressure chamber as the apparatus. The enclosure has integrated bays above and below the optical setup for measurement electronics, ultra-pressure valves and controls, various experimental controls, and our PC that needs to be located close to the cameras for fast data transfer.

Long Duration Thermal Imaging Experiments (LDTIE)

These experiments used long-duration, continuous thermal imaging of loaded samples to look for evidence of LENR in the form of hot spots on the surface.

LDTIE: Experimental Methods, Data Acquisition and Live Image Filtering

Our long duration thermal imaging experiments went through several different iterations over the lifetime of the project as we improved our hardware and software capabilities and investigated different types of samples. In this section, we will focus on the results from the most mature versions of the project (part of the ST3 campaign), where we got the cleanest data and were able to draw the most solid conclusions.

For all of our thermal imaging experiments, we used a 532nm, 100W IPG laser as the stimulation source. As mentioned above, this laser outputs pulses with a tunable average power from 1W to 100W, and a fixed pulse duration of 4ns and a fixed frequency of 40MHz. We used the Pockels cell to modulate the laser on and off with time windows ranging from 50ns to 5ms, and we used laser stimulation powers of around 5W – 80W.

Our IR camera would capture the most signal from a thermal LENR event when the frame exposure time was matched to the duration of the thermal event. Since we did not know a-priori how long a typical LENR thermal event would be, we wanted to make the camera frame exposure time as short as possible to be more sensitive to short thermal events. However, when the exposure time becomes too short and close to the hardware limit of the camera, the variation in thermal counts frame to frame due to noise becomes large, making it hard to distinguish frames with real thermal signals from variations due to noise. Balancing these tradeoffs, we determined that using an exposure time of 5us would be best for most acquisitions – it was the shortest exposure time we could use while remaining sufficiently able to distinguish real thermal signal from random variations frame to frame. Longer exposure times could be used when we wanted to cover a larger time window and possibly catch thermal events that we would have missed with a short exposure time.

We hypothesized that LENR was most likely to occur either when the stimulation laser was on, or at some characteristic time delay after the laser stimulation pulse. We set the Pockels cell pulse frequency equal to the camera frame acquisition frequency by using a common trigger for both, so there would be one frame captured for every Pockels cell laser pulse. We also had the ability to set the relative time delay between the laser pulse and the camera frame to test if we saw thermal events at some time after the laser pulse.

We could also have a longer time scale duty cycle for the laser, where the Pockels cell would typically be allowing laser pulses through for 5s out of every 20s. The camera would still be acquiring frames during this time, ensuring we captured many images with the laser off as well as with it on. Duty cycling the laser in this way prevented the laser from causing gradual long-duration heating of the sample, hammerhead vessel, or optical setup.

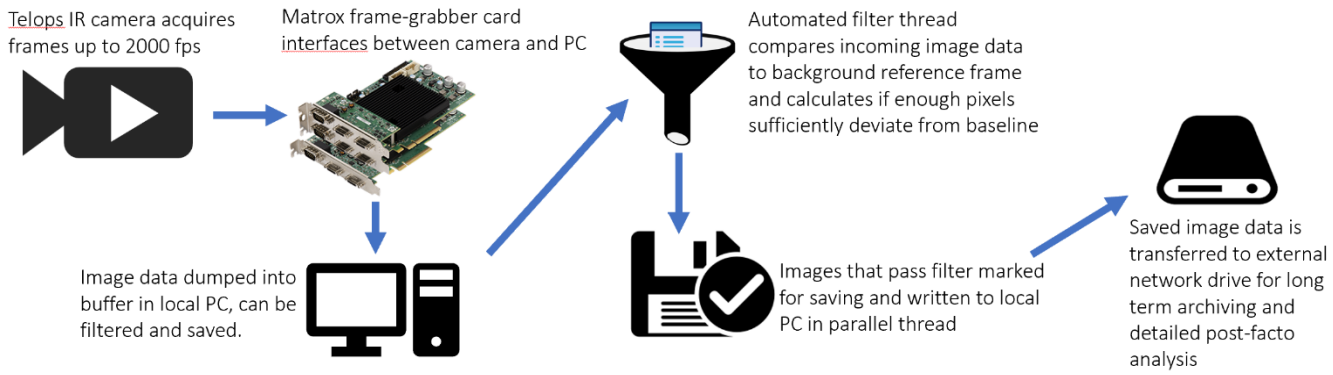


Figure 3 – A schematic of the long-duration thermal imaging data acquisition and processing pipeline.

In order to perform long-duration experiments, we needed the capability to perform real-time filtering of images as they were acquired to isolate those that may have interesting thermal signals and save only those images for future analysis. With the camera operating at 2000fps, it would generate around 150MB of images per second, which would be too high of a data rate for all the images to be saved to a disk for post-facto analysis. This data rate comes out to over 500GB per hour, which means that a single experiment could generate many TB of data. Storing this large volume of data and processing it post-facto would be very cumbersome. By filtering images in real time, we greatly reduced the saving rate to a rate that a single high-performance machine could handle, and greatly cut back the volume of data that would need to be stored long-term and analyzed post-facto.

Figure 3 shows a schematic of our data acquisition and processing pipeline. First, the camera acquires frames at a frame rate typically between 200fps and 2000fps. The data is transferred from the camera to a Matrox frame grabber card via high bandwidth Camera Link cables. The frame grabber card is installed in a high-performance PC that needs to be located near the camera to avoid long Camera Link cable lengths. Once the frame grabber card receives the raw image data, it can deposit the images in a buffer in the local PC. The local PC compares each incoming image placed into the buffer to a background reference frame, and calculates the difference (in either percentage or in absolute raw counts) between each pixel of the incoming image and the reference image. If enough pixels have sufficiently high increases in counts relative to the reference, then that image will be flagged for saving. A parallel thread saves the images that have been marked for saving, along with metadata for each image. Some image metadata is saved for every frame the camera acquires, but the full raw image data is only saved for frames that pass the filter. After the acquisition run ends, the data that was saved to the local PC is transferred to an external network drive for long-term storage where it can be accessed for more detailed post-facto analysis.

During live image acquisition, the slowest part of the process is the saving of images that pass the filter, which was one of the factors ultimately limiting the acquisition speed. We found that saving a frame was roughly 10 times as long as filtering a frame, so the saving could keep up with the filtering if less than 10% or so of the frames were flagged for saving. If more frames than this were

marked for saving over a period of a few seconds, then they would not be able to be saved before being deleted from the buffer to make room for new incoming frames. In this case, we would lose frames that had potential thermal signals. Our maximum save speed was around 200fps, and our maximum acquisition speed due to filtering was around 2000fps. Some acquisitions were conducted at this limit of 2000fps. However, for some of our thermal imaging experiments, running the camera for long durations near its maximum frame rate would heat up the sensor, causing the sensor cooling mechanism to operate at maximum capacity continuously and introduce noise, so we ran some acquisitions at 200fps when we observed this behavior starting to occur. Running the camera at this slower framerate ensured that the sensor never became noisy after running for hours continuously.

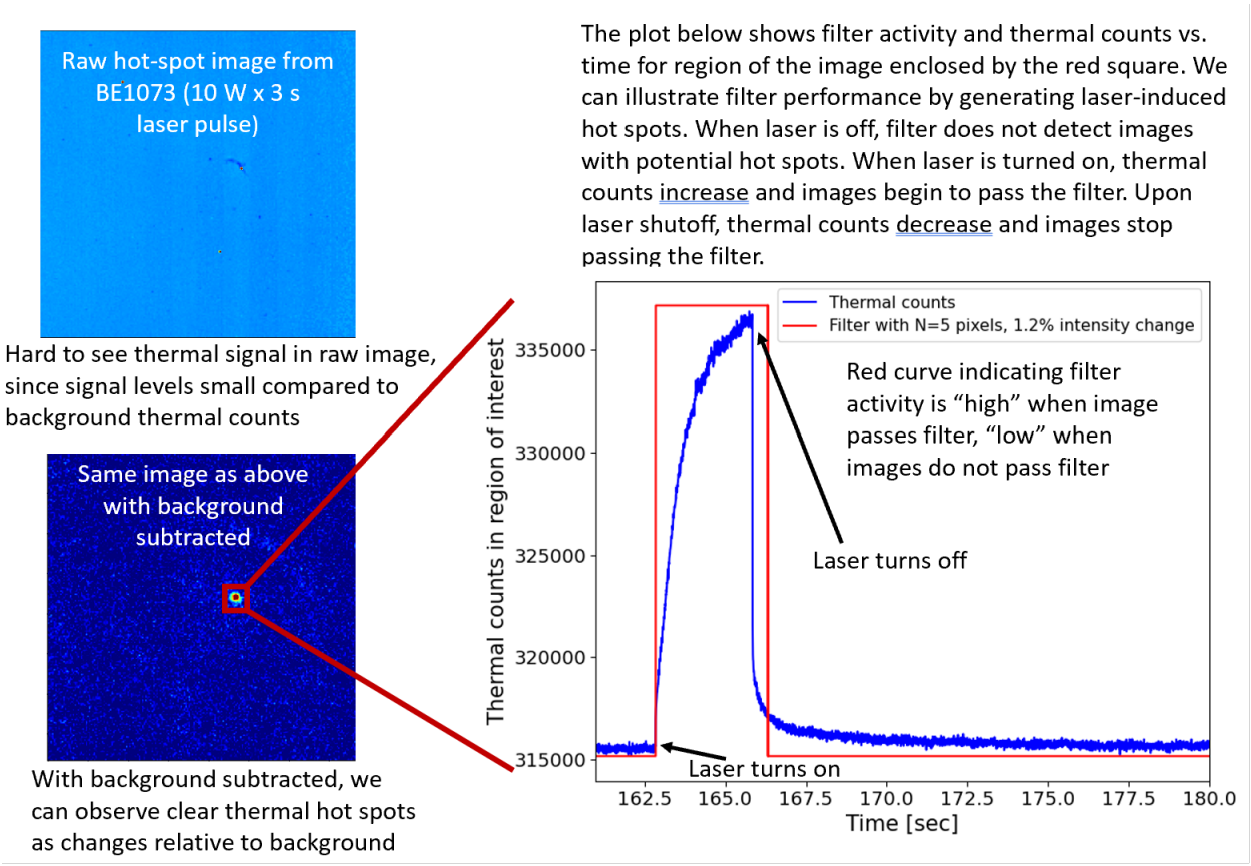


Figure 4 – An illustration of background subtraction and real-time filter behavior.

At the beginning of each acquisition, a background frame is acquired when the laser is off. This can either be a single frame or can be an average of several frames. The background frame is stored as a reference to which subsequent frames will be compared. Since the thermal signals we were looking for might be small compared to the background levels of thermal counts from sensor noise, and since the background levels of thermal counts could drift over time by amounts larger than the signals we'd want to detect, filtering images by looking for fractional changes relative to a background or baseline rather than looking at raw unprocessed images made detection of potential small signals more achievable. This is illustrated in the left half of Figure 4, which shows a raw

image from the camera and the same image with the background subtracted. After background subtraction, the contrast between the thermal hot spot and the rest of the image is enhanced.

Once the background frame is established, all subsequent frames are compared to the background frame on a pixel-by-pixel basis. A time interval was set between background frames. After the time interval had expired, a new background frame would be acquired and incoming frames would then be compared to the new background frame. Since the background level of thermal counts in the sensor could drift over time, we found it essential to re-acquire the background frame frequently. If enough time passed without doing this, then all the incoming frames would be sufficiently different from the background frame to pass the filter, and we would quickly overwhelm the image saving thread with too many frames marked for saving that did not actually contain a true thermal signal. The speed of drift could depend on many factors, such as the ambient temperature, camera sensor temperature, or air flows near the camera. Durations of 30-60s between background frames sometimes resulted in too much background drift, so we settled on using 4s between background frames the standard, which was a short enough interval to have drift never be an issue.

We assume that there is no LENR event happening during the background frame, since the laser was off during that time. The background frames were also a small fraction of the total frames (0.5% to 0.05%), so the chance of a spontaneous, non-laser induced thermal event happening during background acquisition was low. If such an event did somehow occur, the background would become valid again when the next background frame was acquired a short time later.

We employed two schemes for filtering, which could be operated independently of each other (so both could be active simultaneously or only one of the two schemes could be active). The first scheme was a multi-pixel filter, which we developed first and used for every long-duration thermal imaging experiment. The second scheme was a single-pixel filter, which we only implemented and had operational for the final few experiments of the long duration imaging project, where it was used for some of our experiments on peened Pd slugs.

In the multi-pixel filter, we select a number of pixels N and a percent change value. After subtracting the background frame from the frame being analyzed, the filter calculates the percent change for each pixel of the frame relative to the background. If N or more pixels show a percent change greater than the set value, then the frame is flagged for saving. This filter is best for detecting thermal spots that span at least N pixels, but have small signals. Typical values for N were 4-8 pixels. We tended to set the percent change value slightly above the range of most noise fluctuations, so that we could be as sensitive to true signals as possible without saving lots of frames that would pass the filter from noise or drift. This filter is at risk of missing a frame where a small number of pixels (less than N) show a very large signal.

In the single-pixel filter, the filter finds the single brightest pixel for every frame, and compares it to the brightest pixels seen in earlier frames. If the maximum pixel for the frame being considered has one of the top M values seen so far (typical values chosen for M were between 100 and 1000), the frame is flagged for saving. This filter ensures that if a frame has even one pixel that is an outlier in brightness, we will save it even if all the other pixels are close to baseline.

The right half of Figure 4 shows an illustration of the real-time multi-pixel filter behavior (5 pixels with a 1.2% change relative to background and a 50us camera exposure time), establishing that when we artificially induce a hot spot with the laser, the image is flagged by the filter for saving, and when the sample has cooled sufficiently after the laser pulse, the filter stops flagging images as possibly containing a thermal signal.

LDTIE : Sample Types Measured

We used several different sample types for our long-duration thermal imaging experiments. One sample type was Pd thin films, typically 20nm in thickness, on a substrate consisting of 1.5 μm of SiO_2 on top of 500 μm of Si. SiO_2 is a poor thermal conductor, so any heat generated in the Pd thin film from LENR would keep the Pd hotter for a longer duration than if the Pd were on bare Si, making detection more likely. Thin films also were a robust sample type, able to withstand multiple pressure (0ksi to 140ksi and back) and temperature (30C to -85C and back) cycles without any cracking, degradation, discoloration, or other changes. This made them a good candidate for a long-duration optical measurement, where we could remain confident the sample's absorption and emissivity would remain mostly constant over the course of an hours or days long hold at some pressure and temperature. For experiments with these samples, the IR optical system was set to form images at the IR camera with 1X magnification, where the image of the sample on the sensor would match the physical sample size.

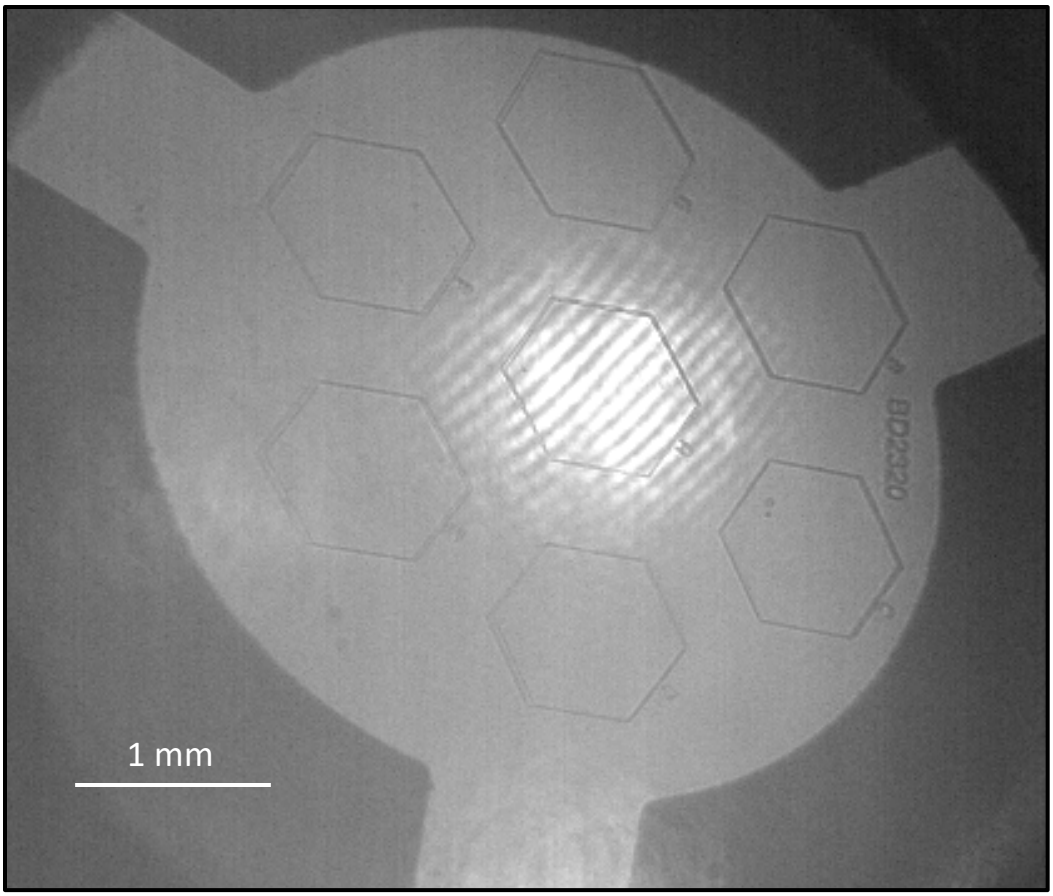


Figure 5 – A visible-light camera image of sample BD2320 (20nm Pd thin film) loaded under 145ksi D_2. The stimulation laser is aligned onto the center hexagon.

Our thin films were typically patterned with 7 hexagons as seen in Figure 5, which is an in-situ image taken with the Photron visible camera of sample BD2320 at 145ksi D_2 during experiment BE2131. This patterning allowed us to test different laser stimulation conditions on different demarcated areas of the sample. We typically used a 1.4mm diameter stimulation laser spot on these samples as seen in the image, which was a spot size large enough to cover one hexagon, but small enough to not stimulate neighboring hexagons.

We ran one long-duration thermal imaging experiment (BE2404) on an erbium thin film sample to test if a rare-earth material loaded with D_2 would show any signs of LENR. This experiment was very similar to BE2131, except with an erbium thin film instead of Pd.

We also measured deformed bulk Pd samples, to test if sample deformation or defects would cause LENR to occur. We used bulk Pd samples for these experiments, since deformations were too difficult to introduce in a controlled, reproducible way to a thin film without causing widespread cracks and holes. For experiments with these samples, the IR optical system was set to form images at the IR camera with 4X magnification, to give us increased spatial resolution compared to the 1X system (~15 μm vs ~60 μm).

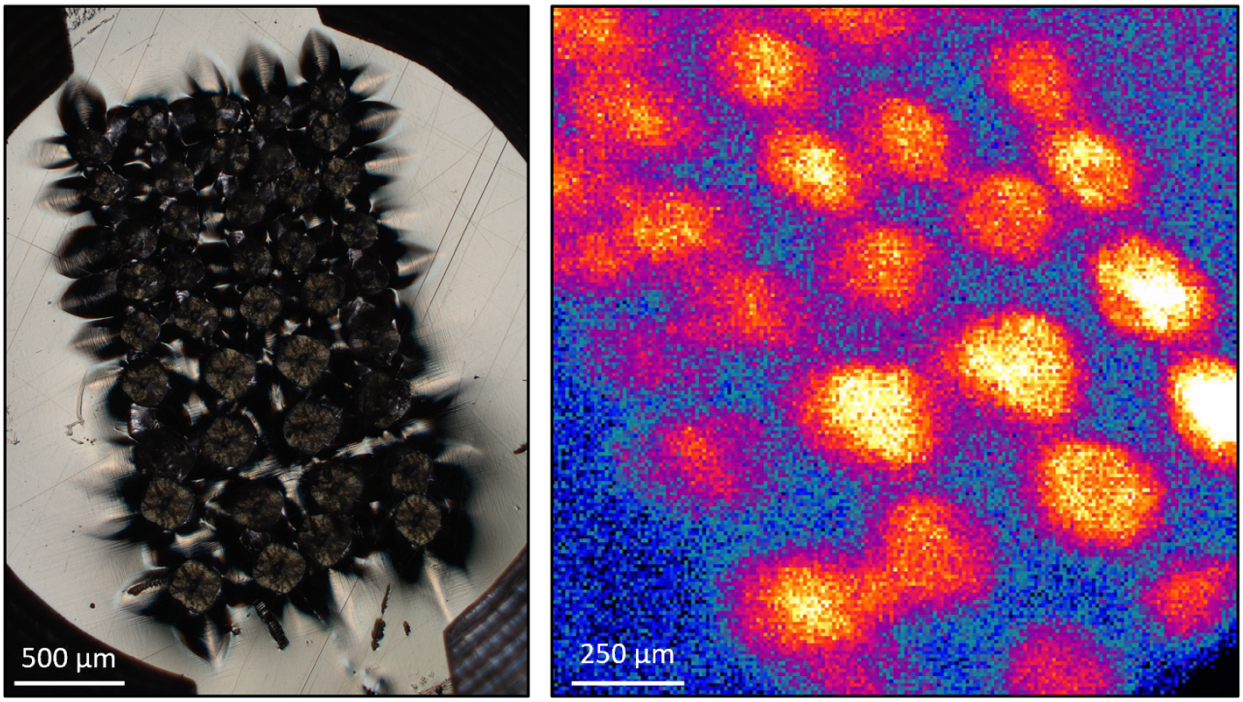


Figure 6 – Ex-situ visible-light camera image (left) and in-situ Telops IR thermal image (right) of peened Pd slug BD1263 used in BE3181 (sample under 100psi argon). Indentations from peening have a higher emissivity than the undeformed, flat crystal surface and thus appear bright in the thermal image.

The first type of deformed samples we tested were peened samples, where an electromagnetically accelerated stylus creates indentations on the sample surface. With our peening setup, we could controllably form indentations of a desired depth and diameter on the surface by controlling the stylus height above the sample and its impact speed. We typically used indentations with depths of 50-100 μm and diameters of 100-200 μm. Figure 6 shows visible-light and IR images of one of our peened Pd slugs, sample BD1263 used in BE3181. In the figure, the peened indentations are darker in the visible-light image since the sample surface is not flat in those areas and does not reflect light cleanly. In the IR image, the sample is being laser-heated under 100psi argon, a low-pressure state where the entire sample surface can reach a fairly uniform temperature (since gas cooling is not very strong). The peened indentations are brighter in the IR image since the deformed areas have a higher emissivity than the surrounding areas of the surface. The sample has two regions of indentations, one with smaller indentations of around 90 μm depth and one with larger indentations of around 130 μm depth (measured ex-situ with optical profilometry with our

Keyence confocal microscope). We also used a 1.4mm laser spot with this sample type so we could selectively stimulate one or the other peened regions.

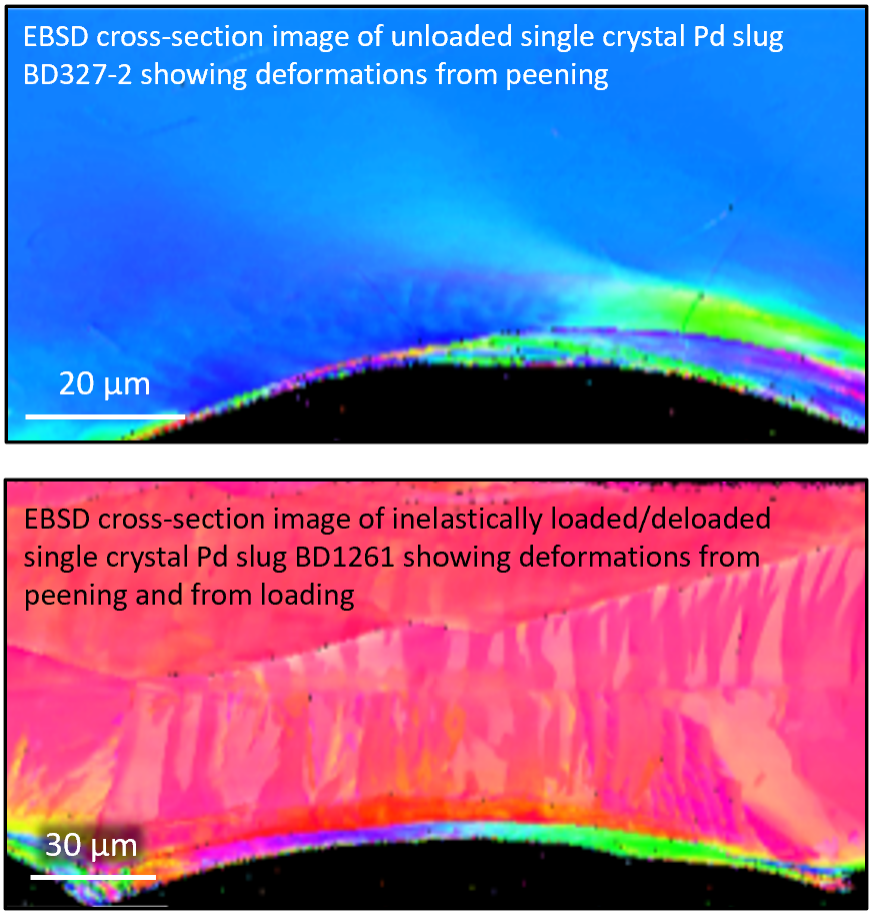


Figure 7 – EBSD on cross sections of single crystal Pd slugs with and without inelastic loading.

To verify that the deformations introduced by peening ex-situ persist after a loading cycle, we performed EBSD (electron backscatter diffraction) on cross-sections of single-crystal Pd slugs with peened indentations, as seen in figure 7. EBSD measures the rotation of crystal domains. The single-crystal sample had only one crystal domain pre-peening (hence the single-crystal designation), but as seen in the top image of the figure, peening caused deformations near the surface where the indentation occurred. Alterations to crystal structure from peening are mostly concentrated within 5-10 µm of the sample surface, with smaller deformations extending up to ~40 µm into the bulk of the sample. The bottom image shows a sample after inelastic loading and deloading, showing that the deformations from peening persist after even an inelastic loading and deloading cycle.

The other type of deformed samples we tested were axially strained samples. For these samples, a rod of Pd was pulled axially until the yield point to induce strain, and then slugs were cut out from the strained rod such that the strain direction was normal to the slug surface. We only ran one clean

long-duration thermal imaging experiment (BE3560) with this type of sample, again using a 1.4mm diameter laser spot for stimulation.

LDTIE : Results and Discussion

Here, we will focus on the results of long-duration thermal imaging experiments that were carried out with live-filtering capabilities in place, and with individual stimulation laser pulses from the Pockels cell synced or controllably delayed relative to individual IR camera frames.

In the earlier stages of the project, we conducted some experiments with short-duration shots, where a single laser pulse was followed by a period (<1 min) of thermal imaging, with images stored on the onboard buffer of the Telops IR camera and downloaded to a computer for analysis. Overall, the behavior we observed in these experiments showed no evidence of LENR and was consistent with conventional behavior. In these experiments, we would observe the sample heat up from the laser stimulation when the laser power was sufficiently high, and cool immediately after the laser shut off. If the laser stimulation was too weak to directly heat the sample enough to see in the thermal camera, we would observe no anomalous thermal signals. The results from these experiments will be skipped over in this report due to higher quality experiments being conducted later on.

We also conducted some long-duration experiments with live filtering, but individual IR camera frames not synced with Pockels cell pulses. For any given individual frame, we did not know the relative timing between that frame and the laser stimulation pulses. These experiments also showed no anomalous thermal signals that might have been evidence of LENR. The results from these experiments will also be skipped over in this report due to higher quality experiments being conducted later on.

Our thermal imaging experiments with live-filtering and synced IR camera frames/Pockels cell laser pulses were the newest, highest quality set of experiments that had the highest chance of detecting LENR, so we will go over some characteristic results in this section. These experiments are all from the ST3 campaign. A table of these experiments is included in Table 1. These experiments all had samples loaded under 140ksi D_2 pressure. We used a steady-state sample temperature of 30C for all but one of these experiments to ensure that the reloading kinetics (which are greatly slowed at cryogenic temperatures) would be fast enough for the sample to reload in between laser pulses, and to increase our sensitivity to thermal events by holding the sample at a temperature not too far below our apparatus' single-pixel thermal detection threshold of ~70C (discussed in more detail below). One experiment (BE1760) was conducted at -80C to test if holding a sample at cryogenic temperatures and attaining the slight increase in loading that occurred at that temperature was necessary to trigger an LENR event. Overall, these experiments also did not show any anomalous thermal signals that might indicate LENR. Controls in H_2 were not very useful or necessary in this campaign, since we did not see any potentially interesting thermal signals in D_2 to require comparison against a sample loaded with H_2. Some of these experiments were conducted with just the multi-pixel filter implemented, but the final run of experiments was conducted with both the single and multi-pixel filters active.

Calibrations using a Pd thin film sample with an RTD on the surface allowed us to match IR camera thermal counts to sample temperature measured from the RTD, and established that we could see a discernible signal above the noise threshold for a single pixel when the sample surface had heated to around 70C. BE2663 using sample BD4861 (50nm Pd thin film with RTD) was the experiment used to establish this limit. Thus, we would be able to resolve a real thermal signal if a small region of the sample (enough to cover 1 or more pixels) got up to around 70C or more for a duration roughly equal to or greater than our camera's exposure time (assuming the frame exposure was synchronized with the thermal event). As mentioned above, 5us was our most commonly used exposure time.

At 30C, there was no resolvable thermal signal in the camera. Thus, it made sense to keep the sample at a steady-state temperature of around 30C, so that we would have no background thermal signal, but a relatively small temperature increase (40C) would show up in the camera, allowing us to be as sensitive as possible to low energy thermal events. If the sample started at a lower temperature, it would still have to heat up to 70C for a signal to be detectable, which would be a larger temperature increase requiring a thermal event with more energy. In order to heat up a (60 μm)^3 volume of Pd from 30C to 70C instantaneously, a thermal event would need to release around 25 μJ of energy, so our thermal imaging experiments were sensitive to spontaneous thermal events of roughly this energy level or greater.

Experiment ID	Sample temp	Sample Type	Camera frame rate	Hours of image data acquired	Live image filters used
BE1760	-80C	Pd thin film	2000fps	15 hours	Multi-pixel
BE1859	30C	Pd thin film	2000fps	21 hours	Multi-pixel
BE2131	30C	Pd thin film	2000fps	58 hours	Multi-pixel
BE2404	30C	Er thin film	2000fps	38 hours	Multi-pixel
BE2906	30C	Peened Pd slug	2000fps	33 hours	Multi-pixel
BE3093	30C	Peened Pd slug	200fps	18 hours	Multi-pixel
BE3181	30C	Peened Pd slug	200fps	50 hours	Single and multi-pixel
BE3355	30C	Peened Pd slug	200fps	34 hours	Single and multi-pixel
BE3560	30C	Strained Pd slug	200fps	36 hours	Single and multi-pixel

Table 1 – A list of the best long duration thermal imaging experiments in the ST3 experimental campaign conducted at 140ksi D_2, 30C, with the 532nm stimulation laser and synced or controllably delayed Pockels cell pulses and IR camera frames.

If the thermal event covered multiple pixels, allowing us to average the thermal counts over multiple pixels, then the noise would be reduced and we would be able to detect a signal at a lower temperature, but 70C is the threshold for a single pixel. If a measured sample had an emissivity different from the calibration sample, then the detection limit would vary accordingly. Higher

emissivities would allow us to detect signals at lower temperatures, and vice-versa. Since we ran the calibration on a Pd thin film, it should be accurate for Pd thin film samples. Samples with surface deformations such as peened samples would have a higher emissivity, meaning that the actual temperature threshold for detection would be below 70C.

BE2131 is one of our higher-quality experiments that had both live filtering and synced Pockels cell/camera frames. This experiment used a 20nm Pd thin film sample, BD2320, seen above in Figure 5. The sample was loaded at 30C in 140ksi D_2. We used a 2kHz Pockels cell pulse frequency, matching the IR camera running at 2000fps. Pockels cell laser pulses were 2us in duration, while IR camera exposure times ranged from 5us to 20us. Laser powers incident on the sample were between 20 and 30W. For this experiment, we captured images with the stimulation laser pulses synchronized with the camera frames, or with the camera frames delayed relative to the laser pulses by values ranging from 5us to 20us. This experiment used the multi-pixel filter only. The IR imaging system magnification was 1X.

After background subtraction and live filtering, we analyzed the IR images that were flagged as potentially containing a thermal signal and saved. Figure 8 shows some representative IR images. Three main features were found in the saved images, all of which have conventional explanations and are discussed in more detail below. None of the saved images appeared to have a thermal signal that might serve as evidence of LENR. There was no observable effect on the images from delaying camera frames relative to laser pulses by some time interval compared to having the camera frames synced with laser pulses, other than the expected behavior where a strong enough laser pulse synced with the camera frame would result in a hot spot in the camera frame at the location of the laser from direct laser heating of the sample.

The first feature that was commonly seen in saved images was small number of randomly scattered bright pixels both on and off the sample area (see images 1-9 in Figure 8). Since these pixels occur both on and off the sample, they can't correspond to a thermal signal on the sample. This behavior is consistent with random noise occasionally causing some frames to have enough bright pixels to pass the filter – we intentionally set the filter level so that almost all normal images will be filtered out, but occasionally random noise will cause some to get through. This ensures that the filter is as sensitive as possible to small signals that are just above the noise level.

The second feature seen in saved images was a hot spot at the laser spot position that appeared when the IR frame was synced with the laser pulse and the laser pulse was strong enough to cause sufficient direct heating of the sample that it was visible in the thermal camera (see images 4, 10-14, 16). This hot spot from the laser disappears when the laser power is reduced or when the IR frame is delayed relative to the laser pulse, allowing the sample enough time to cool from the laser stimulation before the frame was acquired.

The third feature causing an image to be flagged and saved was the appearance of an irregular spatial pattern (see images 10-16), where the areas of bright pixels do not correspond to the sample area. This spatial pattern is due the sensor heating up from long operation at near maximum frame rates and the sensor cooling mechanism (a rotary Stirling cooler, which had to operate at maximum

capacity due to the sensor heating up) causing spatial variation in the pixel responsivity. This effect tended to disappear in later experiments when lower framerates were used.

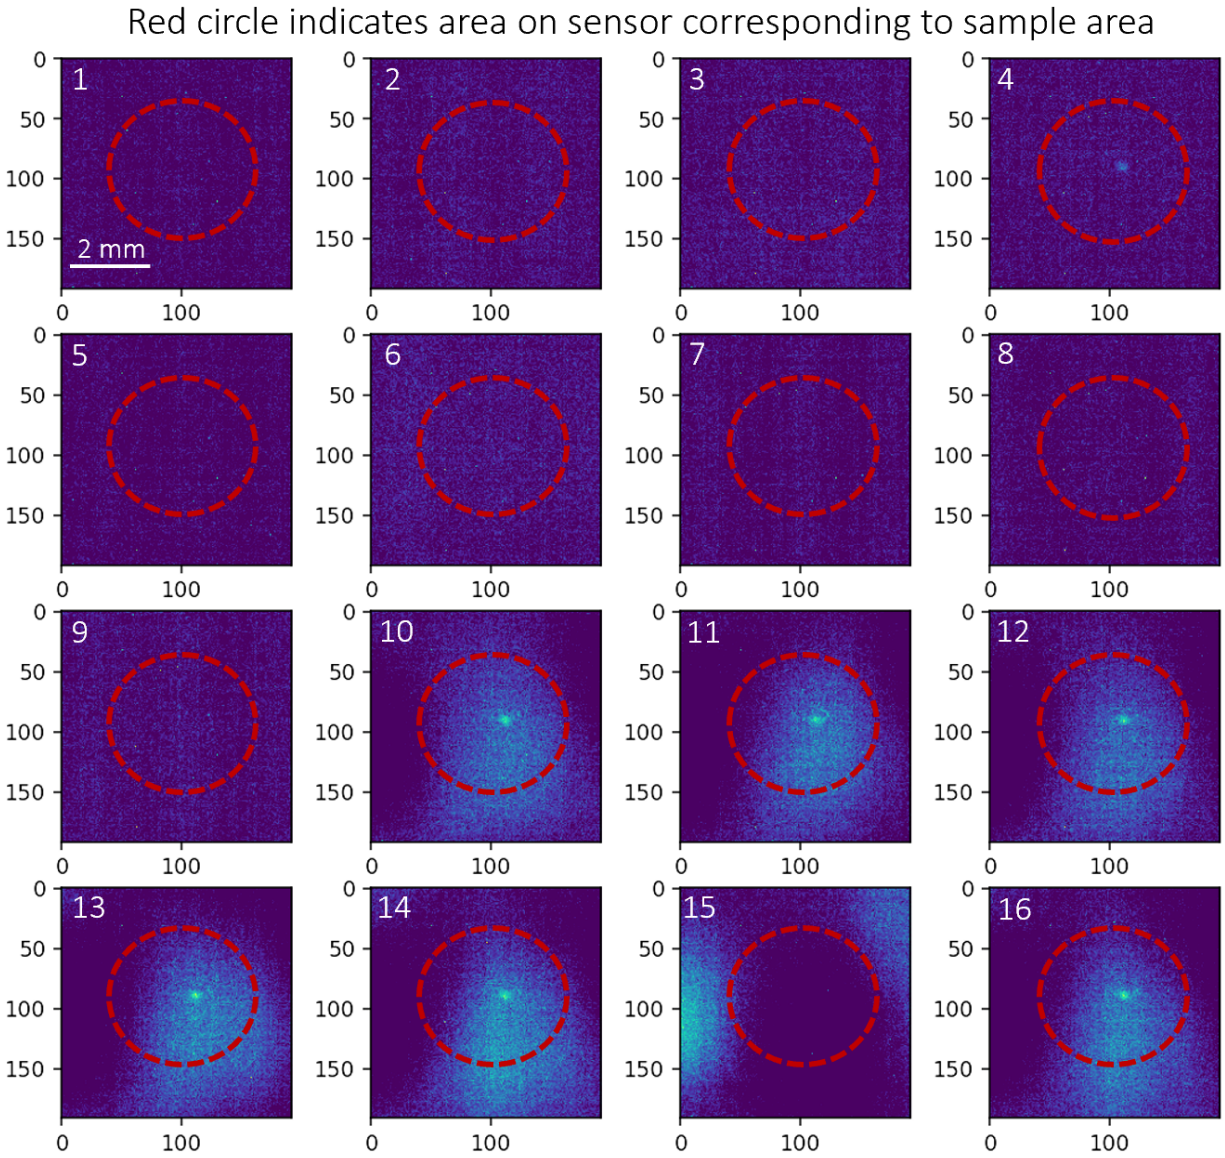


Figure 8 – Representative IR images that passed filtering in BE2131 and were flagged for saving.

For our long-duration thermal imaging experiments on peened samples, we implemented the single-pixel filter discussed above in addition to the multi-pixel filter to improve our sensitivity to images containing just a single pixel that was an outlier in brightness. BE3181, using peened sample BD1263 (shown above in Figure 6), was one of our best experiments with this sample type. In this experiment, we ran the IR camera and Pockels cell at 200Hz, which allowed the camera sensor to remain cool throughout the experiment, avoiding the irregular spatial pattern discussed

above that was due to sensor heating/cooling. The Pockels cell pulse duration was again 2us for this experiment. We also used low enough laser powers, ranging from 8W to 36W incident on the sample, where the direct heating from the laser did not result in thermal hot spots detectable by the IR camera. Note that although we used a higher maximum laser power on the peened Pd slug than the Pd thin film discussed above (36W vs 30W), we did not see direct hot spots from laser heating on the peened slug but did with the thin film. This is because the bulk Pd slug can conduct absorbed heat from the laser at the surface into the bulk of the material, while the thin film is on top of an insulating SiO2 layer that impedes it from conducting heat away from the surface, so less laser power can heat the thin film by a greater amount.

Figure 9 shows some representative data from BE3181. In this experiment, the only type of saved frames that passed the live filter were those with randomly bright pixels due to noise, as seen in the top panel of the figure. There were no spatial patterns or features to the bright pixels, which were always randomly distributed throughout the images. In addition, the single-pixel filter gave us additional insight by recording the value of the single brightest pixel for each frame (regardless of whether it passed the filter and was saved), and saving the set of frames with the highest values for the brightest pixel. Plotting the value of the single brightest pixel for each frame vs time, we saw that there were no pixels that were outliers in brightness compared to the distribution – the brightest pixels for each frame fell within a fairly tight distribution of values, showing that it was random fluctuations that determined which frames had the brightest single-pixels vs others. In the bottom panel of Figure 9, the values of the brightest individual pixels for each frame are plotted over a 1-hour acquisition period, showing that there are no outliers.

There is also a small periodic oscillation with a period of around 1100-1200 seconds that can be seen in the plot. The magnitude of the oscillation is 1-2% of the pixel counts value (10-20 counts oscillation magnitude for a roughly 1000 thermal counts pixel value). The frame-to-frame noise is roughly 50 counts. This is probably due to some cycling of the environmental temperature/air flow or the camera sensor temperature that is causing a regular, though small, modulation of the sensor dark counts background. The appearance of this behavior in many acquisitions independent of the laser conditions, presence or absence of an installed sample, gas used to pressurize, etc. establishes this solidly as a systematic apparatus/camera artifact and not related to a real thermal signal originating from the sample.

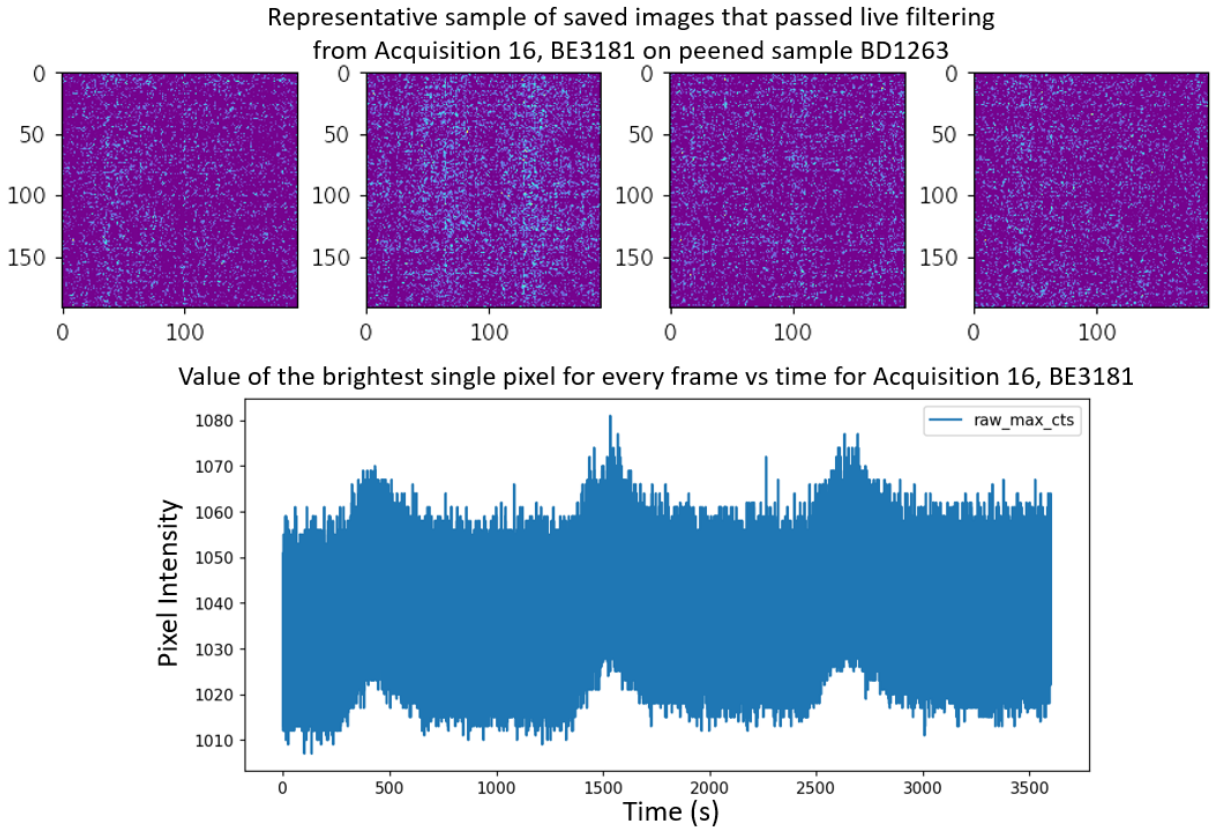


Figure 9 – Some representative data from Acquisition 16 of B3181. The top panels show some representative saved images that passed live filtering, and the bottom panel shows the values vs time for the single brightest pixel for each individual frame over the 1-hour time period of the acquisition.

We also ran one clean experiment on an axially strained Pd slug, BE3560 which used sample BD6832. We used the single-pixel and multi-pixel filters, and operated at 200Hz, so overall this experiment was very similar to BE3181, with only the sample type changed. The results were very similar, with no frames saved with any spatial clusters of hot pixels, and no bright pixel values that were outliers relative to the distribution and thus potentially due to a real thermal signal rather than noise.

Staring Archaeology (SA)

These were archaeology experiments where post-facto, ex-situ imaging of a sample could look for evidence that LENR had occurred with single-event sensitivity.

Staring Archaeology : Experimental Methods and Sample Types Measured

With our archaeology experiments, we loaded and performed laser stimulation of nanostructured samples with the staring apparatus, but we did not have in-situ detection or diagnostics to look for LENR. Instead, we relied on a comparison of pre-experiment and post-experiment scanned imaging over the active area of the surface to look for changes that occurred during the experiment. The samples were designed such that a single LENR event, assumed to release 24MeV of thermal energy instantaneously into the local crystal lattice, should cause irreversible changes that would be detectable in the post-experiment imaging. Interesting results would need to occur in samples loaded in deuterium but not hydrogen, and show up in the post-experiment images but not the pre-experiment images. We did observe differences in post-experiment images compared to pre-experiment images with our various sample types, but all of such results were ultimately attributable to conventional explanations and did not provide evidence of LENR.

One main class of archaeology experiments were those using magnetic archaeology samples, where the heat from a single LENR event would alter the polarization of magnetic domains in a layer of magnetic material below a Pd layer. For these experiments, we used MFM (magnetic force microscopy) as the pre and post experiment scanned imaging method. With these samples, we first performed calibration experiments to characterize the laser powers and pulse durations at which the laser heating was sufficient for magnetic polarizations to start to flip. Then, we would perform laser stimulation below these levels, with the idea that the extra heat from induced LENR would induce flipping of the magnetic domains. Most of these experiments were conducted with the sample loaded at 30C (ST4 campaign), but some were conducted with the sample loaded at -60C (ST4-2 campaign). These experiments are discussed in detail in their own section, so will be skipped over here.

The other main class of archaeology experiments were those using samples that were designed to exhibit damage from a single LENR event, which comprised the ST5 campaign. The two types of samples used in this campaign were Pd nanopillars and patterned nanofilms of pure Pd and various Pd alloys. With these types of samples, we used SEM (scanning electron microscopy) for the pre and post experiment scanned imaging. The details of fabrication of the Pd nanopillars and patterned nanofilms will be covered in detail in their own section of this report.

Pd nanopillars were the first sample type we measured intensively. Samples fabricated using electroplating gave the highest quality samples in terms of cleanliness and uniformity, so our discussion here will focus on samples fabricated using that method. For all electroplated nanopillar samples, experiments were conducted at 30C with our 532nm stimulation laser, using a beam spot size of 300 μm and using the Pockels cell to modulate the laser pulse duration.

After iterating on nanopillars through a few different versions of the samples, experimental protocols, and analysis tools, we moved on to primarily using patterned nanofilms, which proved to be a simpler and more reliable platform for these measurements. The patterned nanofilms were much less susceptible to incidental damage from conventional causes, such as impacts with debris particles, that might appear similar to LENR. This made the nanofilm experiments more robust

and made the analysis of results simpler, without compromising on single-event sensitivity. All our patterned nanofilm experiments were carried out at 30C. Some used the 532nm stimulation laser with the Pockels cell and a beam spot size of 300 μm. After our 100W, 532nm laser broke, our last run of experiments on Pd alloys used our 100W, 1070nm laser with a 500 μm spot size.

Staring Archaeology : Data Analysis, Results, and Discussion

Our initial attraction to Pd nanopillars was that we could fabricate large, uniform arrays of millions of pillars with diameters less than 100nm, which were thin enough that a single 24MeV LENR event should cause a pillar to break. Samples were fabricated in a 3x3 array of square nanopillar pads, with each of the 9 pads a 250 μm side-length square containing 400k-450k individual nanopillars. Across the full sample of 9 pads, there were nearly 4 million nanopillars. The stimulation laser spot size was set to a 300 μm diameter, so it could fully cover one of the 9 pads at a time without heating up any of the neighboring pads.

The sample was divided into different pads so that we could test up to 9 different laser stimulation conditions (power, pulse duration, and number of pulses) with each ultrapressure run. Usually, one of the 9 pads in each experiment would be left un-stimulated as a further type of control. Using the ATLAS mapping capability of our Zeiss SEM, we were able to scan the entire active sample area over a time period of less than 24 hours with a spatial resolution of 10nm/pixel. The leftmost image of Figure 10 in the upper row shows an optical microscope image of one of our nanopillar samples, showing the 9 pads.

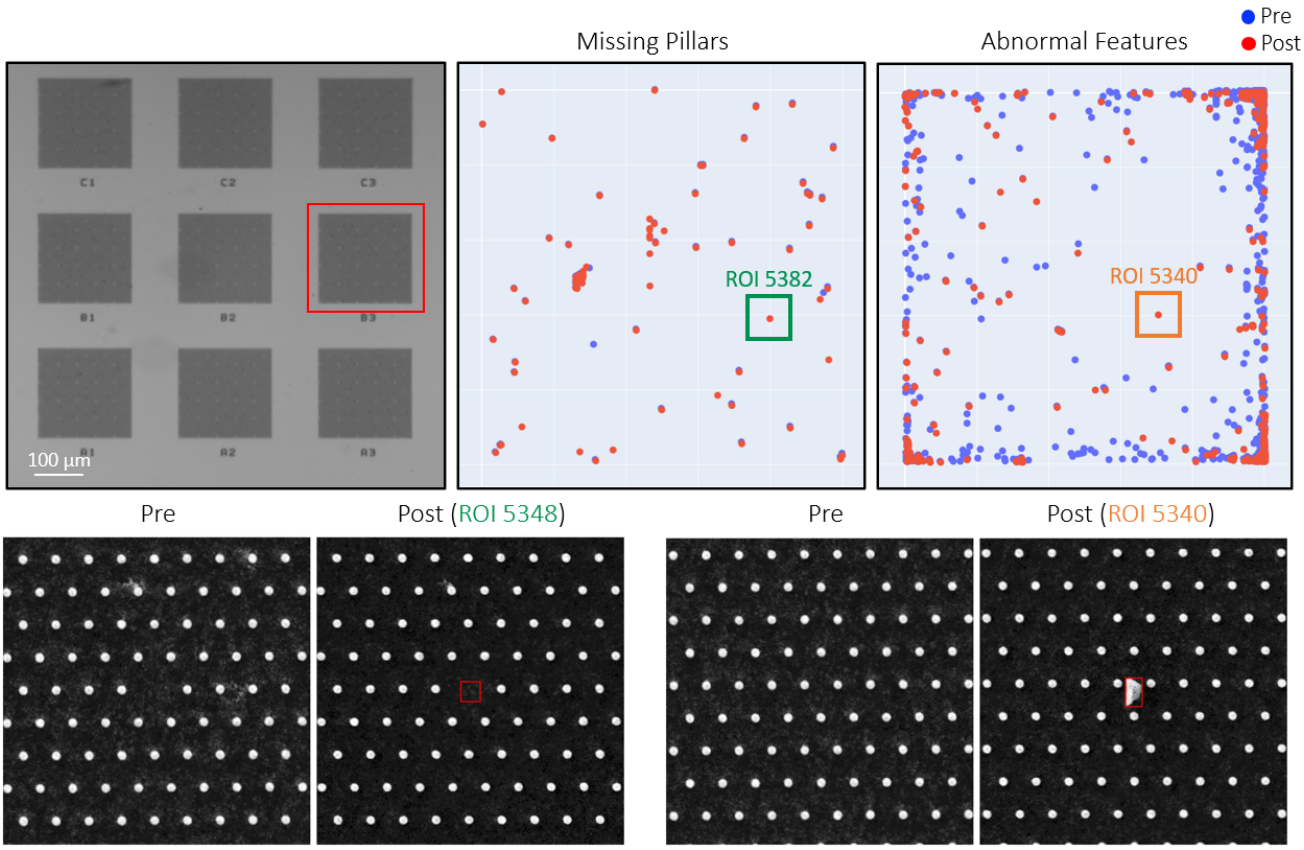


Figure 10 – Upper left: an optical microscope image of a nanopillar sample showing the 3x3 array of nanopillar pads. Upper center and right: Examples of 2D maps for a given pad for missing pillars and abnormal features, showing the spatial positions of those features in the pre and post experiment ATLAS images. Bottom row: One pair of pre/post ATLAS images showing a missing pillar feature, and one pair of pre/post ATLAS images showing a particle that was detected as an abnormal feature.

The pillars were robust with regard to pressure and loading, but extremely fragile with respect to handling mishaps or debris impacts. In every experiment, even with our highest quality samples and careful handling, there would always be some broken pillars that existed pre-experiment and some pillars that would be knocked over during handling or from debris particles during the ultrapressure run.

To try to efficiently catalog the differences in pillars pre and post experiment and identify potentially interesting differences, we developed software tools to automatically detect and flag two main types of features from the large dataset of ATLAS images that mapped the sample: missing pillars, where we would observe an empty spot where we expected a pillar to be located, and “abnormal features”, which were features in the nanopillar array that deviated in size and shape from a typical pillar. This was intended to help us detect pillars that had melted or been knocked over, as well as debris particles. The software tools would assign each feature an ID number, and could display a 2D map of each pad with the positions and ID numbers of each feature, and whether it was present in pre or post ATLAS images. The second and third images in the upper row of

Figure 10 show an example of the 2D maps for a given pad, one showing missing pillars and one showing abnormal features.

We found that there would always be a large number of missing pillars and abnormal features for every experiment, for samples loaded both in D_2 and H_2, so we could not use the simple presence or absence of these features as evidence of LENR. The numbers of each type of feature could also vary greatly from factors such as handling and cleanliness of each sample, so statistics and total numbers of features at the pad-level or sample-level also turned out to not be useful for discerning if evidence of LENR was present. Thus, it was necessary to look at individual features to evaluate if each one might serve as evidence of LENR or if it had a conventional explanation, and to compare features across samples loaded in D_2 and H_2 to see if there were systematic isotopic differences.

In our software tools, we had the ability to display the pre and post experiment ATLAS images for each feature detected so we could analyze them individually. The bottom row of figure 10 shows one example of a pre/post image pair for a missing pillar and for an abnormal feature. In this illustrative example, the missing pillar is present in both the pre and post image, while the abnormal feature is a debris particle, i.e., not a nanopillar, that appears in the post image but not the pre image.

In our experiments on electroplated nanopillars, we tried several different experimental protocols to look for isotopic differences between pillars loaded in H_2 vs D_2. One protocol was to take two separate samples that were as identical as possible and perform the same laser stimulation conditions on each. One sample would be loaded with D_2 and the other with H_2. Another protocol was to perform sequential H_2 and D_2 experiments on the same nanopillar sample, taking ATLAS images after each ultrapressure run. Some samples were tested with H_2 first, others with D_2 first. For both of those protocols, we used laser stimulation energies below the damage threshold. Since we were looking for damaged pillars as evidence of LENR, we wanted to avoid false-positives by causing direct damage from the laser.

As a third protocol, we did perform laser stimulation up to and above the damage threshold on some nanopillar samples to test if there were any systematic isotopic differences (e.g., if the damage was considerably and consistently more extreme in a sample loaded with D_2 than one loaded with H_2). This way, we could be potentially sensitive to LENR even if it occurred for laser stimulation near or above the damage threshold.

In all of those types of experiments, we observed no systematic differences between samples loaded with H_2 vs D_2. For the below-damage-threshold experiments, both isotopes exhibited missing pillars and pillars knocked over from scratches or debris, but no damage that looked to be from LENR. For the above-damage-threshold experiments, both isotopes exhibited similar behavior in the manifestation of the damage from laser heating.

Table 2 shows a list of the experiments we performed on electroplated Pd nanopillars in the Staring apparatus. Matched pairs where sequential ultra-pressure runs were performed on the same nanopillar sample are listed in the same row in the table.

Experiment ID	Sample temp	Gas(es) and pressure(s) used for loading	Laser energy above damage threshold?
BE1634	-90C	145ksi D_2	no
BE1677	-80C	145ksi D_2	no
BE1835	30C	145ksi D_2	no
BE2043	30C	145ksi D_2	yes
BE2111	30C	145ksi D_2	no
BE2229	30C	145ksi H_2	no
BE2316	30C	145ksi D_2	no
BE2342	30C	145ksi H_2	no
BE2544/BE2569	30C	145ksi H_2 then 145ksi D_2	no
BE2588/BE2647	30C	145ksi H_2 then 145ksi D_2	no
BE2802/BE2827	30C	145ksi D_2 then 145ksi H_2	yes

Table 2 – A list of experiments performed in the staring apparatus on electroplated nanopillar samples, using the 532nm laser with the Pockels cell for pulse modulation.

One of our cleanest nanopillar matched pairs using one sample loaded in H_2 and another in D_2 was BE1835 (D_2 run) and BE2229 (H_2 run). Figure 11 shows SEM images of each of these samples, establishing that the pillars for each are nearly identical in height and diameter.

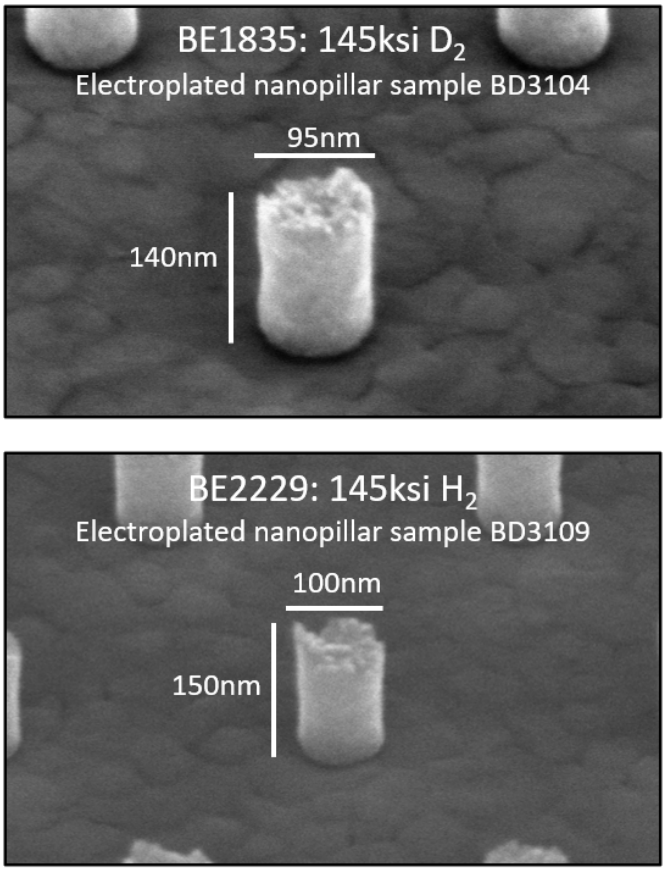


Figure 11 – Two nearly identical nanopillar samples used in a matched pair of experiments – one sample loaded in D_2, the other in H_2.

In this matched pair of experiments, one pad of each sample was not shot with the laser to serve as a control pad. The other 8 pads allowed us to test 8 different pulse duration and laser power combinations. We used two pulse durations of 200ns and 1us. For each pulse duration, we used four different laser powers of 14, 26, 48, and 78W.

We found that each sample had generally similar numbers of missing pillars and abnormal features. There was no correlation between missing pillars or abnormal features and the laser energy, showing that all laser energies were comfortably below the damage threshold – direct laser heating did not affect the pillars on any of the pads.

Figure 12 shows a comparison of representative missing pillar features detected in each sample in this matched pair of experiments. Almost all missing pillars detected in each sample existed in both the pre and post images. However, there were a small number of missing pillars in the post-images in the D_2 sample (BD3104) that were not in the pre-images, as shown in the figure. For the H_2 sample, all the missing pillars identified were present in the pre and post images. Although this may appear to be an isotopic contrast, the missing pillars in the D_2 sample generally occurred in clusters with other missing pillars, and often near scratch marks or debris particles, indicating they

were knocked over mechanically from scratching or debris. Repetitions of similar experiments established that scratching and debris could commonly knock over or dislodge pillars. The presence of large numbers of dislodged pillars and common conventional explanations made it difficult or impossible to claim any given missing pillar was evidence of LENR. {This ambiguity and difficulty with identifying true LENR evidence from conventional damage was a large part of the motivation for eventually switching to thin films instead of nanopillars as our physical archaeology platform.}

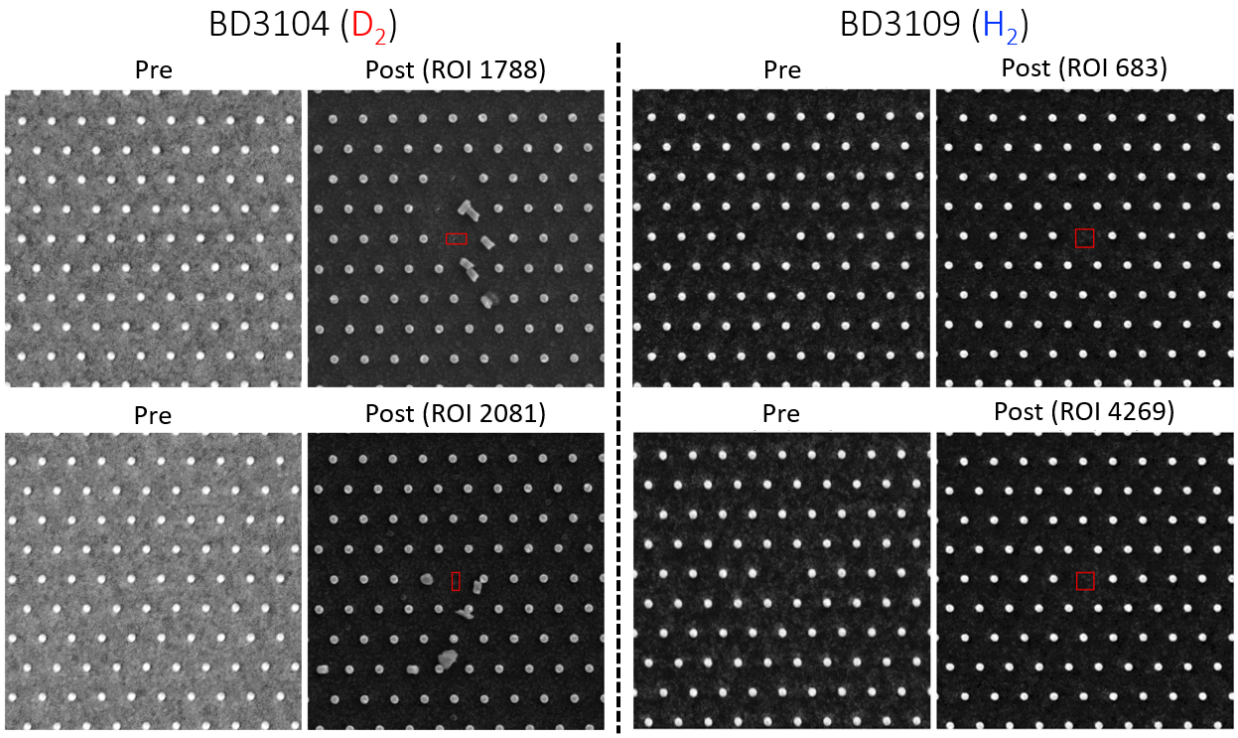


Figure 12 – Examples of missing pillars in BD3104 used in BE1835 (D_2), and BD3109 used in BE2229 (H_2).

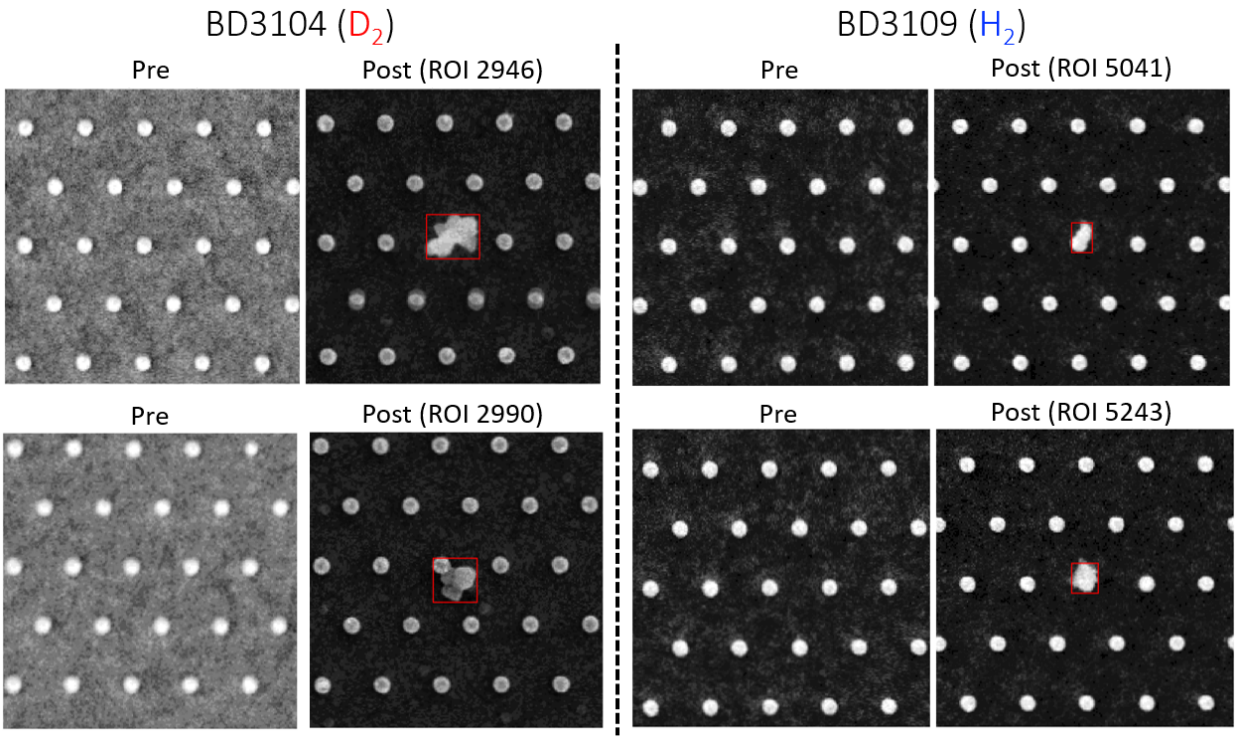


Figure 13 – Examples of abnormal features in BD3104 used in BE1835 (D_2), and BD3109 used in BE2229 (H_2).

Figure 13 shows a comparison of representative abnormal features detected in each sample in this matched pair of experiments. Both samples showed similar behavior, with relatively similar concentrations of debris particles showing up as abnormal features in both samples. We did not detect any misshapen or melted pillars that might serve as evidence of LENR.

We also tried repeated laser stimulation experiments on the same nanopillar sample loaded first under H_2, then D_2, or vice-versa. Figure 14 shows images of two samples that were each loaded with both gases successively, but in opposite orders. In these types of experiments, we observed no isotopic contrasts. All missing pillars either existed pre-experiment or appeared to be from scratches or debris particles, and both H_2 and D_2 runs showed similar but small amounts of debris that accumulated on the surface during an experiment. Figure 15 shows examples of missing pillars and abnormal features on the same sample for an H_2 run and a D_2 run.

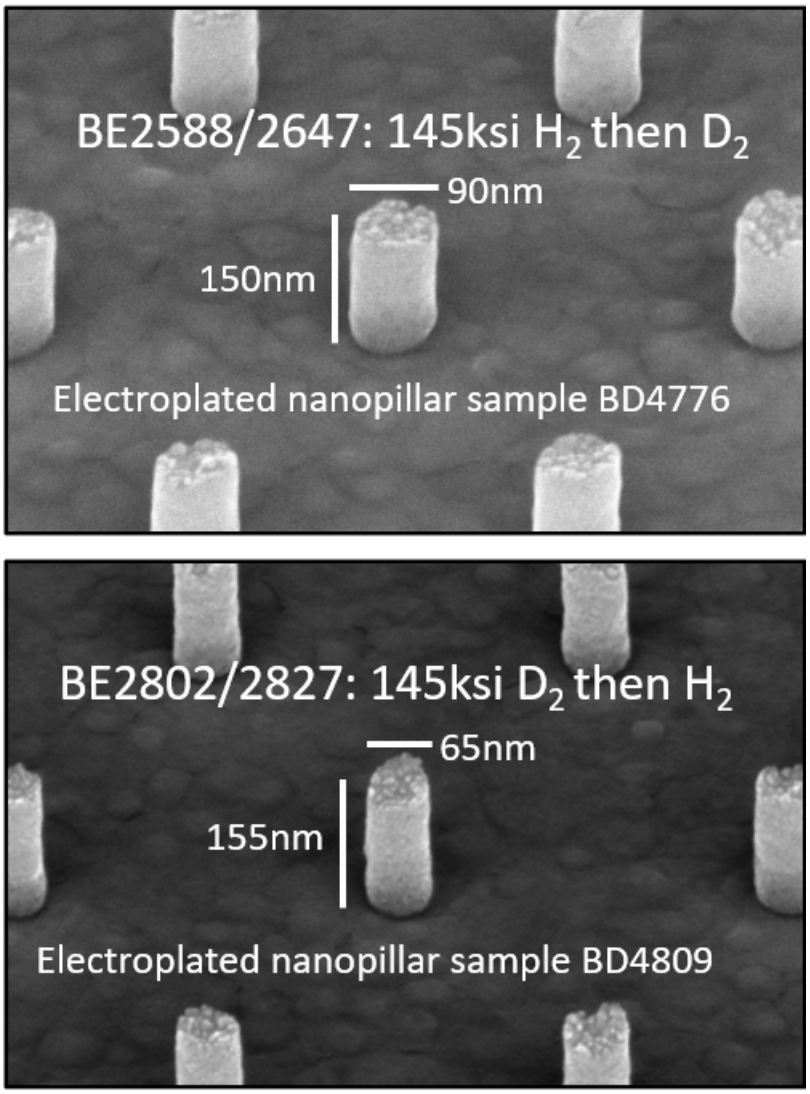


Figure 14 – Two nanopillar samples, one loaded with H_2 then D_2, the other with D_2 then H_2.

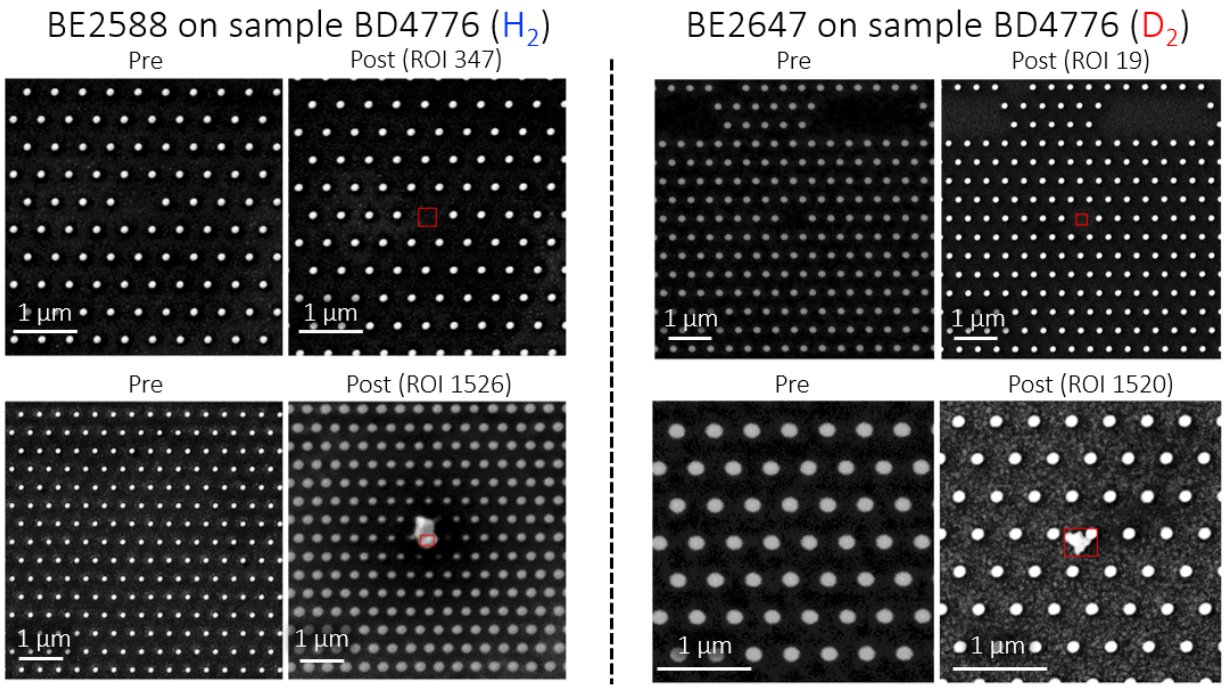


Figure 15 – Examples of missing pillars and abnormal features for H_2 and D_2 runs on the same nanopillar sample.

For nanopillar samples measured with laser stimulation energies above the damage threshold, we also saw no isotopic contrast that might indicate LENR. Figure 16 shows some representative images for BE2827 and BE2802, a matched pair of experiments on the same nanopillar sample with intentional laser damage. The sample was loaded first with D_2, then with H_2. One pad was damaged in D_2, and a different pad was damaged under H_2. Both damaged pads show melted or ablated pillars close to the area of highest laser intensity (left side of the images). As distance from the area of highest laser intensity increases (moving left to right in the images), the pillars gradually become less deformed, as expected.

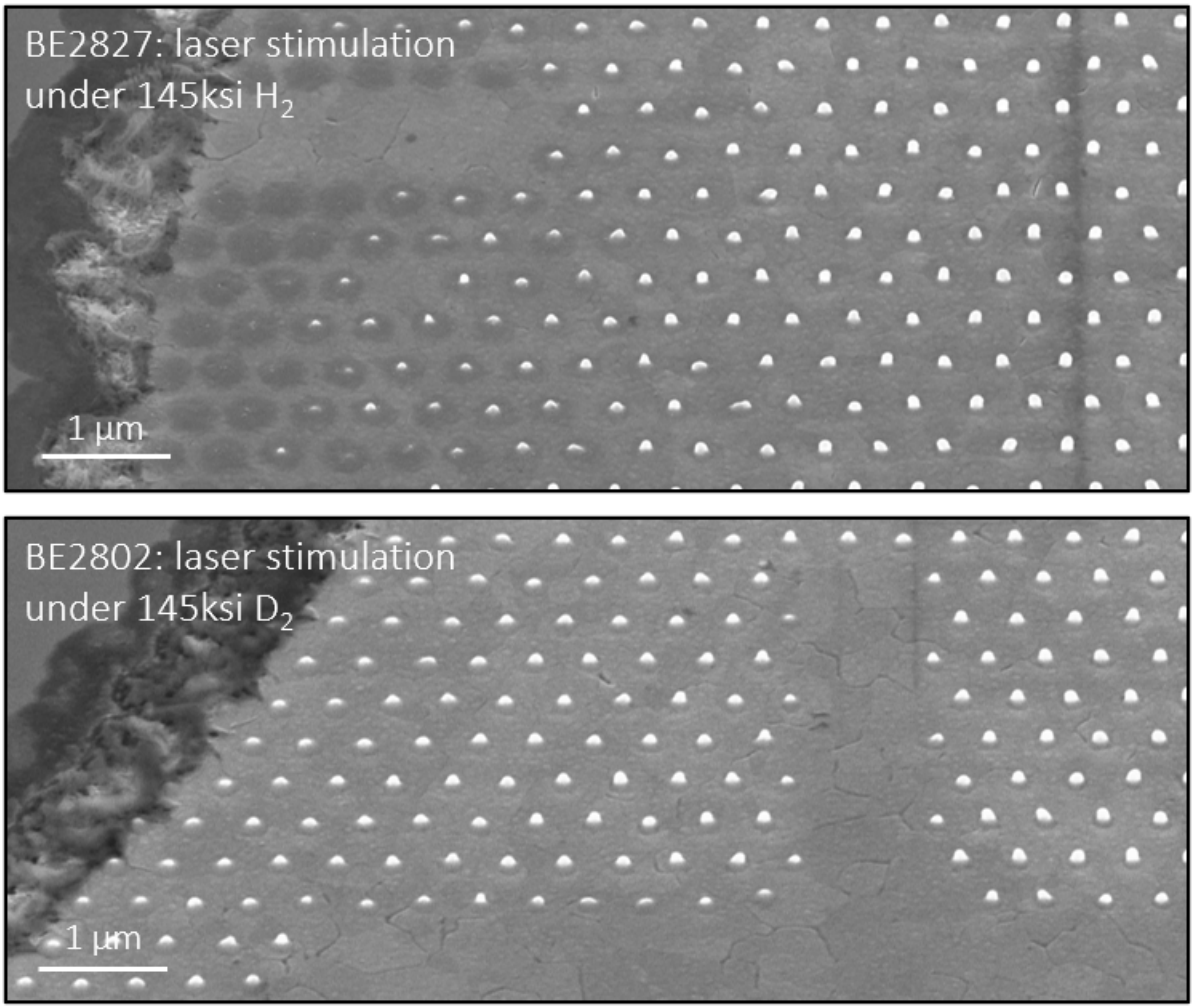


Figure 16 – Examples of pillars intentionally damaged from direct laser heating for a sample loaded in D_2 followed by H_2. The damage in each gas shows similar behavior.

The lack of isotopic differences in our electroplated nanopillar experiments and the unavoidable presence of many missing pillars and debris particles motivated a transition to a simpler and more robust sample type that could still provide single-LENR-event sensitivity.

Patterned Pd nanofilms proved to be a robust, sensitive sample type for physical archaeology. 20nm thick films could be fabricated with very few pre-existing defects or holes (less than 10 over an entire sample), which was 10-100X fewer defects compared to our best nanopillar samples. This thickness of 20nm ensured that thermal energy from an LENR event would cause a hole to appear in the film. These samples were also much more robust to routine handling or particles, and could undergo pre and/or post experiment cleaning to remove particles and residue without damaging the thin film. Using these types of samples, the quantity of features we needed to investigate on each sample was greatly reduced, allowing us to manually look over the ATLAS

images rather than rely on computer detection of features and allowing us to draw stronger, less ambiguous conclusions about the presence or absence of LENR.

After a few iterations with slightly different fabrication methods for the samples, we settled on fabricating our nanofilm samples on bare Si substrates so that the Pd or alloy thin film would be in direct contact with an electrically conducting layer. This aided in avoiding charge buildup during SEM imaging, which causes blurry images and degraded spatial resolution. Samples consisted of an array of 3x3, 4x4, or 5x5 pads, each pad containing an 8x8 or 10x10 array of nanofilms that were squares with a side length of 10 μm. As with our nanopillar samples, the laser spot size (300 μm for 532nm laser, 500 μm for 1070nm laser) was large enough to cover one pad without touching other neighboring pads, so we could test various laser powers and pulse durations with each ultra-pressure experiment.

The patterning of each larger pad into an array of nanofilms allowed us to have regions of high contrast with sharp boundaries between the Pd or alloy nanofilm areas and the areas of bare substrate, which was essential to focus the SEM properly during ATLAS imaging. A bare, unbroken, unpatterned thin film would be much more difficult to bring into focus and result in lower quality ATLAS images.

Figure 17 shows an optical microscope image of one of our pure Pd patterned nanofilm samples, as well as some ATLAS images of individual nanofilms showing representative pre and post experiment results from BE3003. We were able to identify contamination as opposed to damage to the thin film.

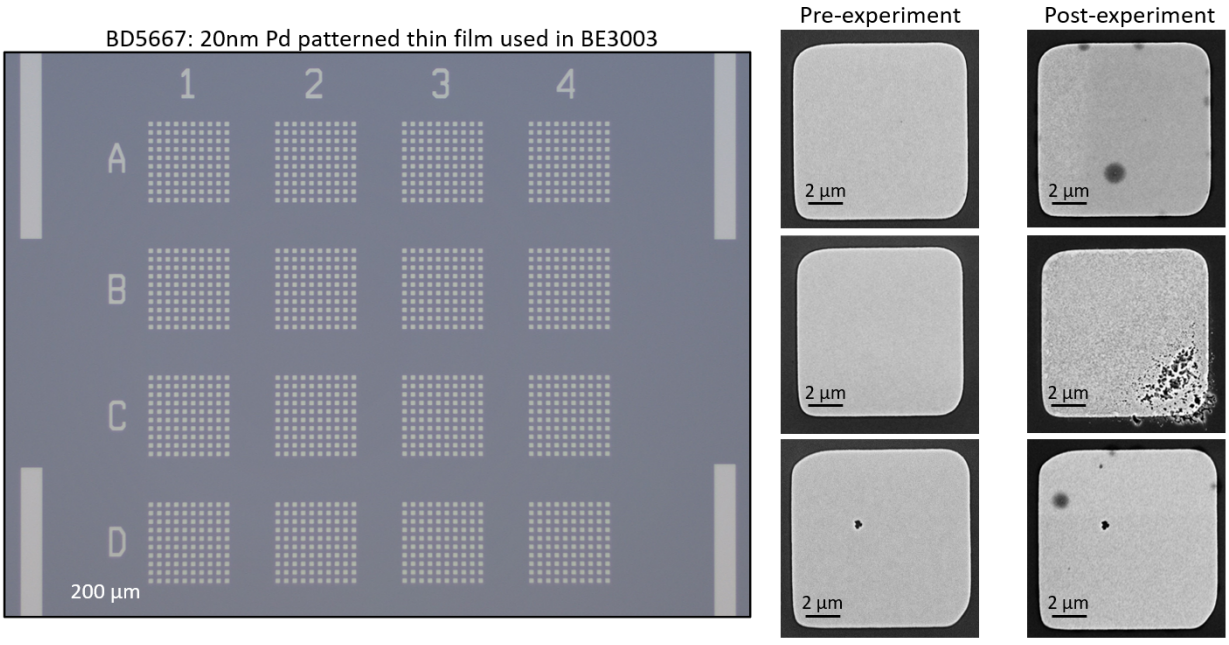


Figure 17 – Optical microscope image and some representative pre/post ATLAS image pairs of individual nanofilms for BD5667 (pure Pd nanofilm sample) used in experiment BE3003.

The top row of pre and post ATLAS images in Figure 17 shows some contamination spots on a nanofilm that are in the post image but not the pre image. The middle row shows damage to a nanofilm in the post image from ablation from laser heating, and the bottom row shows a small damage spot that was present in both the pre and post images, so was pre-existing prior to the experiment. There is also some contamination visible in the bottom pair of pre/post images.

Experiment ID	Sample material	Gas and pressure used for loading	Laser used	Laser energy above damage threshold?
BE3003	Pure Pd	145ksi D_2	532nm	yes
BE3155	Pure Pd	145ksi D_2	532nm	no
BE3321	Pure Pd	145ksi D_2	532nm	no
BE3439	Pure Pd	145ksi H_2	532nm	no
BE3517	Pure Pd	145ksi D_2	532nm	no
BE3725	PdRh alloy	145ksi D_2	532nm	no
BE3764	Pure Pd	145ksi D_2	532nm	no
BE3979	PdAg alloy	145ksi D_2	1070nm	no
BE3997	PdB alloy	145ksi D_2	1070nm	no
BE4005	PdPt alloy	145ksi D_2	1070nm	yes
BE4007	PdNi alloy	145ksi D_2	1070nm	no
BE4049	Pure Pd	145ksi D_2	1070nm	no
BE4065	PdNi alloy	145ksi D_2	1070nm	yes
BE4077	PdB alloy	145ksi D_2	1070nm	yes
BE4177	Pure Pd	145ksi D_2	1070nm	yes

Table 3 – A list of experiments performed in the Staring apparatus on patterned nanofilm samples. All experiments used a sample temperature of 30C.

Table 3 shows a list of our patterned nanofilm archaeology experiments. In all of our patterned nanofilm experiments, of which BE3003 is a representative example, all damage was the result of laser heating above the damage threshold, or was pre-existing prior to the experiment, or appeared to be the result of mechanical scratching from particles or handling.

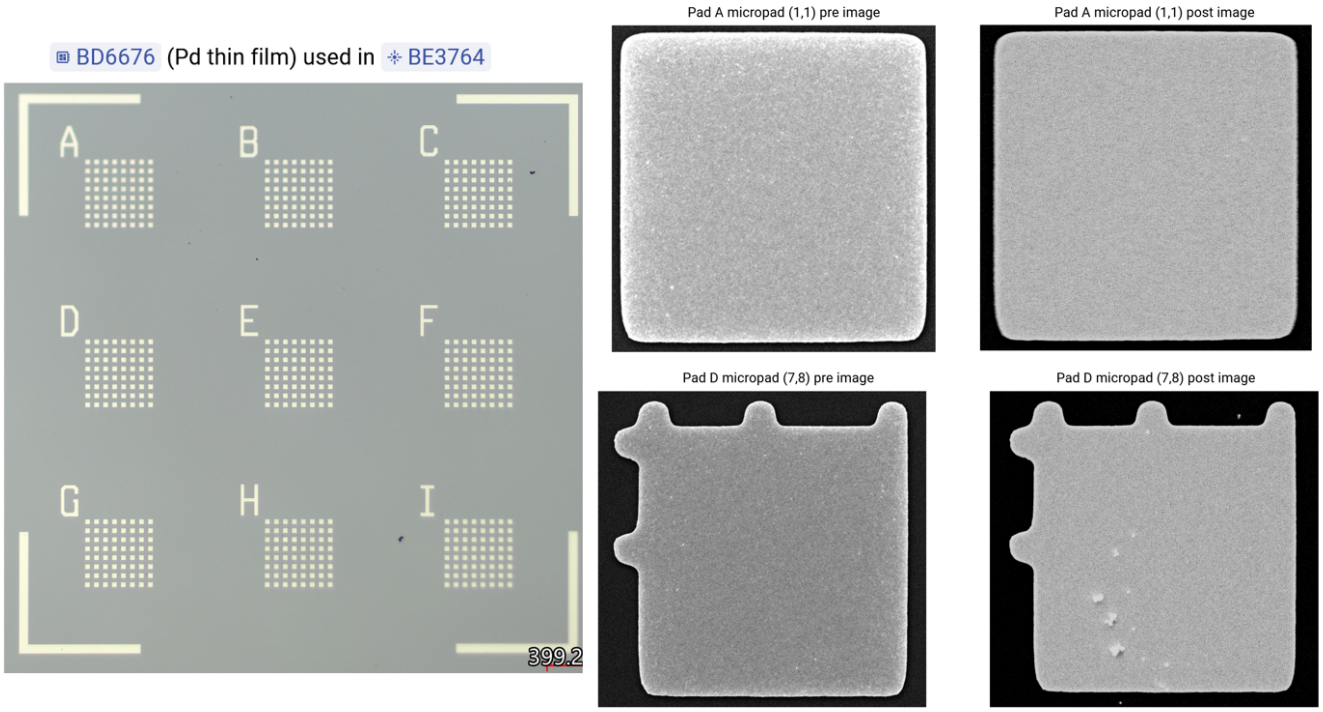


Figure 18 : An optical microscope image and some representative ATLAS images of pre and post experiment nanofilms from BD6676, a pure Pd thin film sample used in experiment BE3764.

After iterating on our experimental protocols and attempting several repeats, we achieved a few very clean experiments with pure Pd and the alloys PdB, PdNi, PdPt, PdRh, and PdAg. Our best experiment was BE3764 on a pure Pd thin film (BD6676). The ATLAS images were very high quality, and across the entire sample, there was no damage to the thin film in the pre or post experiment images. There were a very small amount of debris particles present that did not cause damage. Figure 18 shows an optical microscope image and some representative ATLAS images of pre and post experiment nanofilms from this experiment. Most nanofilms showed no discernible change between the pre and post images (e.g., top row), but a few nanofilms showed a small amount of debris in the post images that did not cause damage to the thin film (bottom row). Given the single-event sensitivity, the large number of lattice sites over the entire sample, and the complete lack of any damage, this experiment strongly establishes that LENR (specifically, the D-D to heat modality) does not occur in a Pd thin film loaded and stimulated under the conditions used in that experiment. Our other patterned nanofilm experiments, though not quite as high quality, collectively serve as a strong body of evidence that LENR did not occur in the pure Pd and Pd alloy thin film samples we measured.

Chips

The thrust of this campaign was to use the tools of modern lithography in order to reduce the size of palladium targets, with an eye towards both making novel structures to help generate LENRs as well as increasing our sensitivity for detecting them.

Nanostructures

The goal of the Nanostructures effort was to use our cleanroom facilities and nanofabrication resources to quickly prototype custom nanostructures that were used to support mainline experimental campaigns. The primary focus was to make structures that enabled stimulation and detection of AHE. The Nanostructures campaign primarily focused on structures for physical archaeology (Fig 1): plasmonic structures using pyramids and bowties, as well as nanopillars, nanofilms, and split ring resonators. These structures were then tested in mainline campaigns.

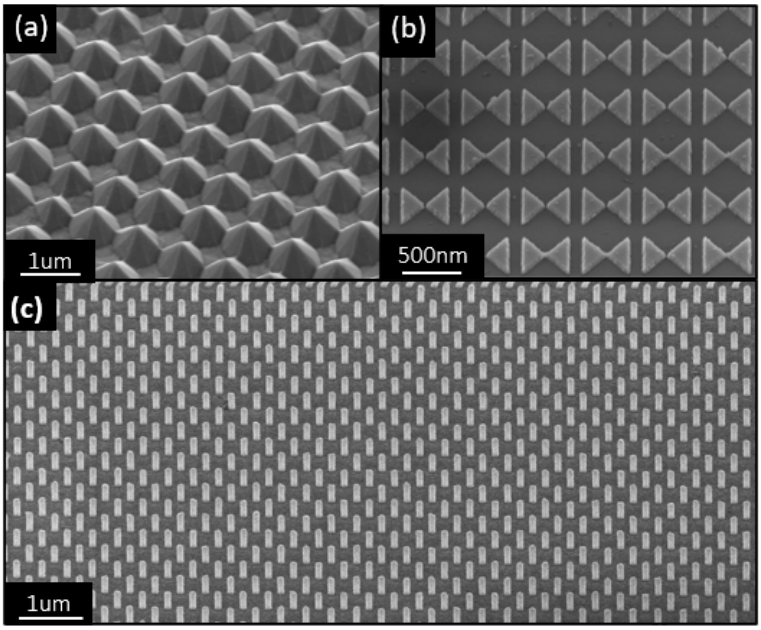


Figure 1: The Nanostructures campaign focused on developing several types of structures: (a) nanopyramids, (b) bowties, and (c) nanopillars {discussed in the Physical Archaeology section of this report}.

Nanostructures : Nanopyramids

Introduction: Plasmonics refers to strong light-matter interactions that result in collectively oscillating free electrons at metal surfaces, and generally involving high electric fields. A variety of efforts were made to design plasmonic structures which enhance the field and may trigger AHE. For example, simulations showed that metal pyramids have high enhancement at edges and sides

of the pyramids. The challenge, however, is using the optically lossy material of choice, palladium. Our group discovered a method to anisotropically etch single-crystal palladium and used this to fabricate plasmonic nanopyramids. The accompanying goal was to try to decrease dimensions so as to increase field enhancement. This was difficult and required using ebeam lithography, finer control of surface polishing and crystallinity, and improved etching handling procedures to reduce the sizescale of inverted pyramids down to 200 nm. Furthermore, in the process of reducing the inverted pyramid size, we found that we could tune this etching method to create a variety of pyramid-like nanostructures with potential for plasmonic enhancement, shown in Figure 2.

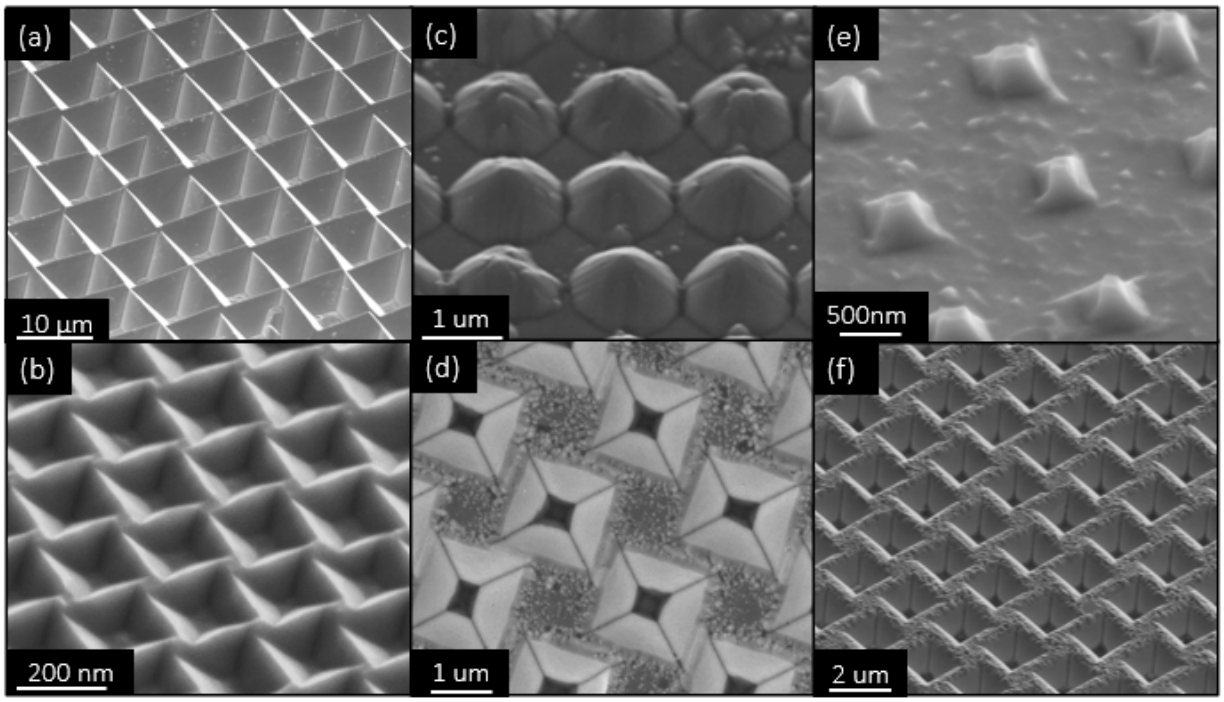


Figure 2: Finetuning the fabrication process enabled fabrication of a variety of pyramid geometries, including (a) larger inverted pyramids fabricated using photolithography, (b) inverted pyramids down to 200 nm pitch using ebeam lithography. Rotating the orientation of the pattern in relation to the Pd crystal planes resulted in different etch geometries (c,e), and etching in two steps (d,f) resulted in nano-texturing on top of the formed structures.

Experimental Methods: Anisotropic etching of palladium was integral to our approach to forming nanopyramids. To our knowledge, we have not seen explicit mention of anisotropic etches of palladium in literature, though this likely results from the significant expense of obtaining adequate quantities of single crystal palladium for such etching studies. We found that both of the chemical etchants of palladium we tried, Transene Palladium etchant (25-30% $FeCl_3$, 3-4% HCl, >60% H_20) and piranha (4:1 sulfuric acid: hydrogen peroxide) do in fact preferentially etch along the (100) plane, with a much slower etch rate along the (111) plane resulting in pyramidal type structures. This is exactly the behavior exhibited by the KOH etchant commonly used in fabrication processes for Si, which like Pd, also exhibits a fcc crystal structure. We suspect that any chemical etchant of palladium will likely result in similar etching behavior.

Additionally, the flatness of the 5 mm diameter palladium slug substrates was also crucial to the success of forming reliable structures. Extensive effort was focused on refining the polishing technique to ensure flatness and subsequent etching uniformity, including efforts such as optimizing the etching time to be brief enough to prevent vibrations from resulting in waviness across the surface, yet long enough to ensure the removal of particles. Well-polished palladium slugs greatly benefited the lithography, in terms of maximizing usable patterned area by minimizing the resist edge bead, in addition to maintaining the focal plane across the sample during exposure.

Significant effort was spent on refining the lithography and etching parameters. We did initially try fabrication of alignment marks to enable orienting the lithography pattern with respect to the crystal orientation; however, we found that rotational alignment of the pattern using our direct write e-beam tool (Durham Magneto-Optics MicroWriter ML3) resulted in the mirrors of the digital micromirror device (DMD) array misaligning to the pattern and limiting the achievable resolution of the tool. Therefore, our samples patterned using photolithography were not pre-aligned to the crystal rotation. Our final photolithography process was as follows: we first applied hexamethyldisilazane (HMDS) to the polished Pd slugs to promote adhesion to the resist (SPR 3612), which we then spin on at 6 krpm for 45 s, achieving ~900 nm resist thickness, followed by a resist prebake at 90 °C on a hotplate for 1 minute. We exposed the pattern at 80 mJ with varying rotation angles to try to ensure that one of the patterns aligned well to the crystal orientation desired, then performed the post-exposure bake at 115 °C for 1 minute, developed in MF319 for 1 minute, and then hardened the resist at 115 °C for 1 min. Etching was done using the Transene ferric chloride etchant at 60 °C, followed by DI water rinse. We targeted etching times so that the nanopyramids in our ideal patterns etched into V-shaped trenches. We removed resist in acetone, isopropanol, and dried the samples with a nitrogen gun.

The ebeam lithography process was performed using a JEOL JBX 6300 at Stanford Nano-Shared Facilities. First, to promote adhesion to the sample surface, the palladium slugs underwent a UV ozone clean at 100 °C for 10min followed by a dehydration bake at 180 °C for 5min. Then, PMMA 950K A4 resist was spin-coated at ~2500 rpm for 1min, resulting in resist thickness of ~180-190 nm. The resist was then baked at 180 °C for 2min, followed by exposure in the JEOL with conditions dependent on pattern. Then, the sample was developed in 1:3 MIBK:IPA for 30 s. Etching of the much smaller patterns was done in the Transene ferric chloride etchant at room temperature due to the slower etch rates at room temperature, and resist was then removed in acetone, isopropanol, and dried with a nitrogen gun.

Results: We targeted structures for use in the Staring apparatus which had a 532 nm laser with a tunable spot size from 300 μm - 1.4 mm. Since the structures were now extremely tiny relative to these spot sizes (1000x smaller at 200 nm), we were not limited in space and could fabricate a multitude of pyramid types in each hexagon target region (Fig 3). In each hexagon, we fabricated ‘inverted pyramids’ and ‘Aztec pyramids’ at four different orientations with sizes ranging from 200 nm to 1000 nm. In total, this is 72 different types of pyramids per hexagon. Each hexagon was identical and allowed for a different laser-stimulation condition in Staring, allowing for 72 ‘experiments’ to be run per high-pressure test. Our rapid development allowed us to progress from just six different structure types to 72 structure types per experiment over a 3 month period,

enabling high throughput in completing the testing of many structure types. However, post-imaging of these structures following laser stimulation in the Staring apparatus did not reveal any meaningful differences with the pre-imaging which would indicate anomalous activity such as LENRs.

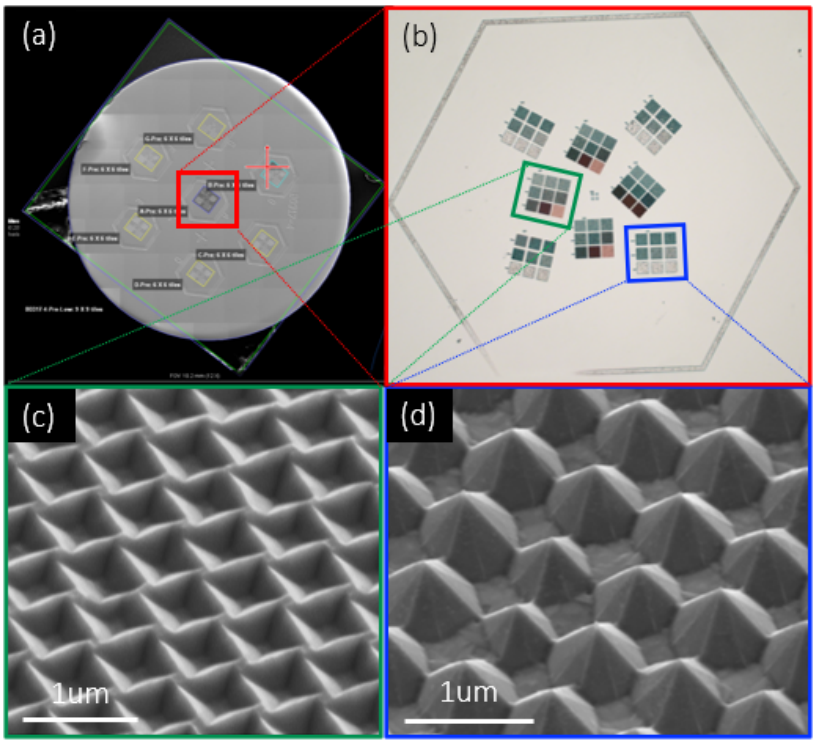


Figure 3: The nanopyramid chips (a) were composed of seven identical stimulation hexagons (b), each with two types of pads: four pads of inverted nanopyramids rotated at four different orientations (c), and 4 pads of nano Aztec pyramids rotated at four different orientations (d). Each pad contains arrays with lithographically defined sides ranging from 200 nm to 1000 nm. Each chip contained a total of 72 different pyramid types per shot condition.

Nanostructures : Bowties

Introduction: In addition to the pyramid structures in bulk palladium, arrays of optical antennas were patterned on-chip and optimized for strong interaction with our stimulation lasers. Nano bowtie antennas are a type of optical antenna that are commonly used for their ability to confine and enhance electromagnetic fields at the nanoscale. In contrast to the majority of our other stimulation techniques which are fundamentally thermal, plasmonic structures induce extreme transient electric fields in the skin depth of the palladium. As with other physical archaeology sample types which confine the palladium volume, the center ‘bridge’ of the bowties where the field enhancement is the strongest are 20-40 nm in width and 50nm thick making them the target of the archaeology analysis.

Simulation and Optimization: The primary goal of simulations was to optimize the peak field enhancement of incident 800 nm laser pulses on palladium. The Lumerical code was used to evaluate the peak field enhancement while exploring various bowtie geometries, structure layouts, and material stacks. It was found that much higher field enhancement and a thermal isolation of the bowties could be achieved using a metal-insulator-metal material stack. A gold reflector layer was deposited on a silicon wafer followed by a CVD silicon dioxide film. The thickness of this layer was optimized for each sample batch according to the wavelength of the stimulation laser.

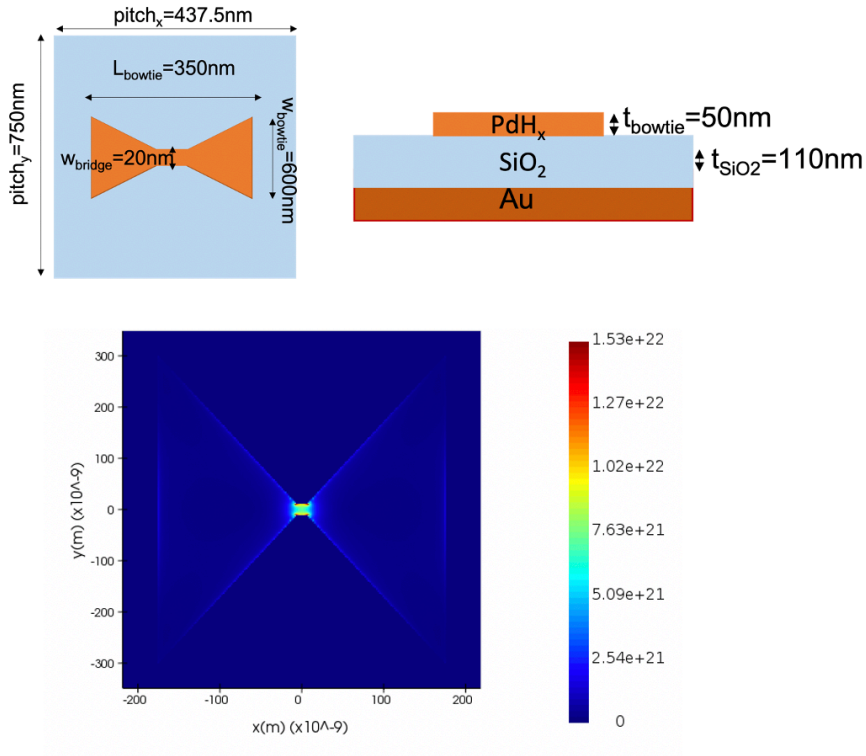


Figure 4. Simulation results and design dimensions for 800 nm wavelength optimized metal-insulator-metal bowtie antenna. Simulations suggest a peak field enhancement factor of 355 in the palladium.

Bowties were tiled with each adjacent bowtie rotated 90 degrees to minimize dependance on polarization alignment of the stimulation laser.

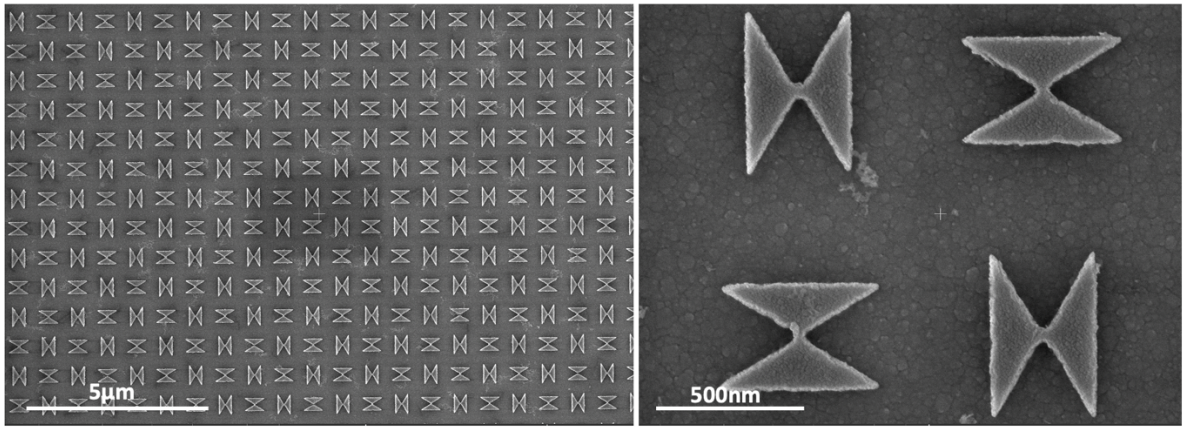


Figure 5. SEM images of bowtie antenna arrays optimized for 1064 nm wavelength.

In short-pulse stimulation experiments when nearing the damage threshold for the devices, bowties were observed to eject material from the center bridge. This effect was produced in both deuterium and protium and determined to be of classical origin.

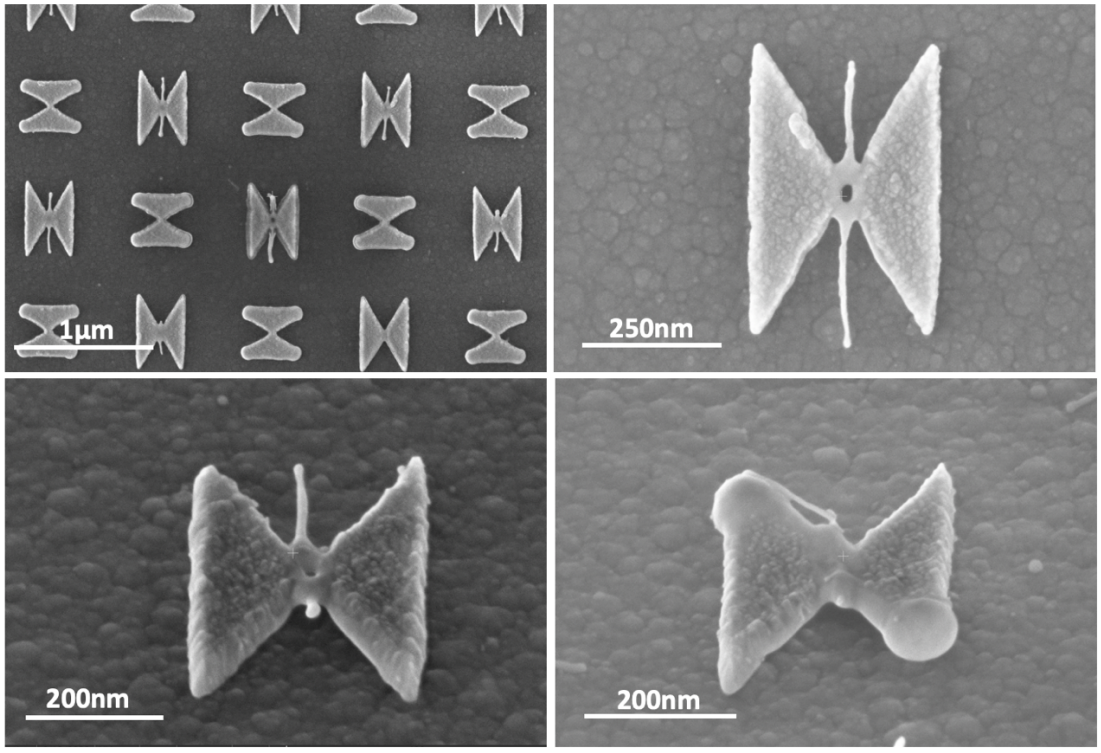


Figure 6. Callisto experiment BE1438. SEM post-analysis of short pulse laser stimulated, loaded palladium bowtie antennas.

Chips

Traditionally, work in the LENR field looks for a macroscopic (millijoules or more) energy signature, due to a cumulative reaction of billions or more of LENR events. In trying to pursue how and when an anomalous event occurs, the Chips project umbrella worked toward looking for less numerous events. This was made possible using modern lithography and nanofabrication techniques to make the target material mass much smaller, enabling access to smaller length scales and timescales. Smaller length scales also enabled devices which combine sensing and stimulation on-chip, negating the need for supporting equipment such as optics. The small footprint also enabled fitting multiple devices onto a single chip in ultra-pressure vessels, enabling faster experimental throughput.

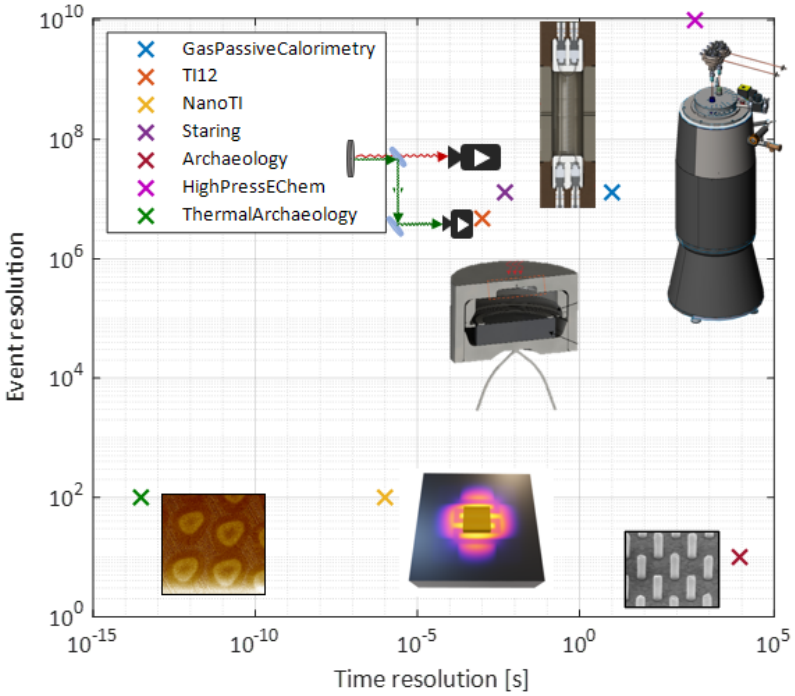


Figure 7: Estimations of the sensitivity of various Callisto campaigns.

Combining sensing and stimulation on-chip into smaller form factors also enabled new stimulation and sensing modalities. One track explored in Chips was resistive heating, which is based on manipulation of electrons, which are inherently easier to localize than optical photons. Lithographically defined resistive heaters can create spatially localized thermal profiles and enable precise command of current and temperature swings; however, resistive heating does not approach the short timescales of femtosecond laser pulses which can probe beneath lattice-vibration timescales. Another track explored in Chips was electrical stimulation using nanofabrication techniques to localize electric fields for non-thermal stimulation. This includes both plasmonic stimulation, not an easy task with such a lossy material as palladium, as well as more classical field enhancement using narrowed tips or capacitors. Nanofabrication also enabled creation of microstructures designed for sensing. Chips has explored resistive sensing for both determining

loading and temperature anomalies as well as various geometries for 'physical archaeology'-type sensing, which enabled looking for single events.

Chips : NanoTI

Introduction: Thermal stimulation has been extensively explored across our organization and by others working in the field, and for good reason: on a small enough scale, nearly all stimulation ultimately results in manipulation of lattice vibrations, i.e., phonons. Likewise, thermal sensing has also been extensively investigated, using different methods such as calorimetry, or detection by thermocouples or resistive temperature detectors (RTDs). NanoTI brought thermal sensing and stimulation onto the chip itself by using lithographic techniques to delineate the target material and the resistive heater/sensors. Similar to bulk thermal ignition experiments, NanoTI used rapid heating to push the palladium to a non-equilibrium state where it is both hot and highly loaded. In contrast to bulk experiments, the use of NanoTI chips allowed us to heat very small Pd samples over short timescales and repeat such pulses millions of times during a single test. Since so little material is involved, few-event sensitivity should be easier to attain, since by the specific heat equation less energy, and hence fewer LENRs, are required. The NanoTI project sought evidence of 24 MeV events by searching for small variations in the temperature data as well as examining post-SEM morphology.

NanoTI pushed the limits of thermal stimulation to the melting and degradation temperature limits of materials in the system (~1900C for SiN) and pursued targeting different stimulation regimes. This campaign also operated at the extremes of pressure, reaching the highest pressures reported to date (aside from diamond anvil literature), enabling access to unity loading of the palladium. Unity loading has been thought to be a crucial link toward achieving low-energy nuclear reactions; however, we did not observe any anomalous events under such conditions. With the NanoTI approach, we were able to run tests with as little as 600 femtograms of Pd. Our system achieved temperature sensitivity of 7-10 °C, which equates to resolution of >500 events over 1 ns (Fig 8). We collected data from more than 448 devices for more than 1.5 million pulses. We used SEM imaging for archaeology-type observation of anomalous differences in pre vs post images.

Under the extreme testing conditions of NanoTI, many material properties and behaviors are poorly understood. Our work toward testing under such conditions enabled us to gain understanding of such phenomena as palladium loading kinetics under high thermal gradients and high loading fractions, and behavior of thin film meandering RTDs at high temperatures. Additionally, the testing conditions of NanoTI required much hardware development and engineering ingenuity to make testing more convenient and reliable.

Design and Simulation: The NanoTI device had much smaller Pd mass than many of the other campaigns, which by the specific heat equation, implies less energy is needed to result in a temperature increase (Fig 8b). However, less mass does equate to fewer lattice sites for AHE to occur, and assuming LENR events are probabilistic, this equates to lower probability of such events. Due to the system design which enables very short stimulation pulses, repetitive stimulation of the Pd for a multitude of pulses is relatively quick, enabling us to theoretically increase the probability of events. For the design (Fig 8a), we placed a layer of Pd on top of a meandering platinum resistor which acts as both the stimulation source and the thermal sensor.

The SiN passivation was intended to help mitigate gas cooling effects and prevent deloading during stimulation, in order to maximize the time that the Pd stays hot and highly loaded. Applying a stimulation voltage across the meandering resistor results in a thermal gradient across the palladium (Fig 8c). Such a thermal gradient was intended to stimulate an LENR event, which would result in heat quickly conducting out of the metal to be detected by the RTD. Using FEM simulations (COMSOL) to check our estimates, a suite of 1000 AHE events of 3.84 pJ energy would create a temperature spike in the Pd (Fig 8d blue line) and would also result in a spike in temperature on the Pt RTD at a point directly underneath the AHE locations (Fig 8d green line). However, experimental measurements of the RTD sensor are measuring all the way across the platinum meander, essentially averaging over the entire RTD, which results in a much lower temperature rise (Fig 8d red line); this illustrates the need for a fast temporal response and a highly sensitive thermal sensor.

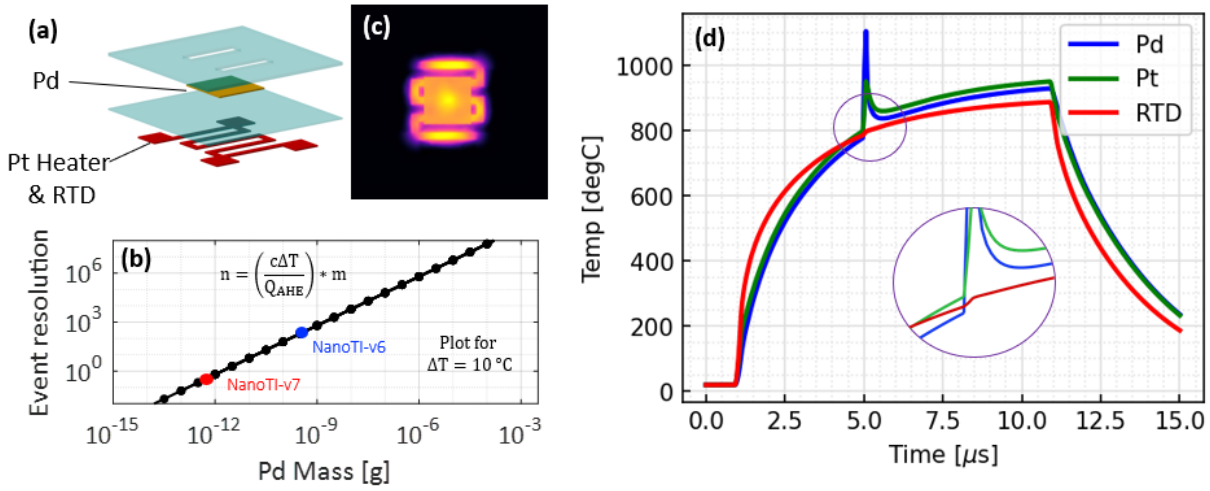


Figure 8: Design of NanoTI. (a) Diagram showing palladium target material on top of platinum meandering resistor which acts as both a heater and RTD. Silicon nitride passivation prevents palladium from shorting the platinum resistor, and silicon nitride also caps the palladium to ensure it stays hot and highly loaded. (b) Palladium mass is directly proportion to event resolution. (c) COMSOL simulations illustrate the effect of an event in the palladium as the heat spreads throughout the device. (d) An AHE event (blue) in the Pd, assuming 1000 3.84 pJ events over 100ns, superimposed on a 10us thermal pulse. This results in a spike in temperature in the Pt (green) and a much smaller spike averaged over the RTD (red), illustrating the need for a fast and sensitive thermal sensor.

Experimental methods:

(a) *Device fabrication:* Optimization of nanofabrication was essential to ensure repeatability and reproducibility of our experiments; we worked to improve our fabrication processes to minimize defects using a nanofabrication network that encompassed our own facilities, in addition to several other facilities in the Bay Area. With each device generation, we improved fabrication processes, enabling us to increase RTD sensitivity and decrease Pd target size. Fabrication of the NanoTI device involved starting with thermally oxidized silicon wafers; initial simulations showed that 1.5 μm of silicon dioxide on silicon was a good balance between oxidation time and heat capacity. We then defined the meandering platinum resistors. For the larger resistors down to 1.25 μm linewidths, bilayer photolithography is necessary to ensure clean lines (we coated the wafers with LOL resist, baked at 185 C to ensure minimal undercut for the narrow line gaps, coated with Shipley positive photoresist 3612, baked, exposed the pattern in a Heidelberg direct write tool,

developed in MF-26A, and descummed the resist in an oxygen plasma). For the smaller resistors with linewidth down to 150nm, we used ebeam lithography. Evaporation of the titanium adhesion layer and the platinum ensured that we could then cleanly perform liftoff in NMP, followed by acetone, IPA and N_2. We then used PECVD to deposit silicon nitride as an electrical insulation layer. Then, we repeated the same lithography and liftoff process with evaporation of palladium to define the target material. We again used PECVD to deposit the capping silicon nitride layer, and we used monolayer photolithography with RIE etching to define the capping region. This is followed by removing the resist in an oxygen plasma. Optical and SEM imaging confirmed the desired device morphology (Fig 9).

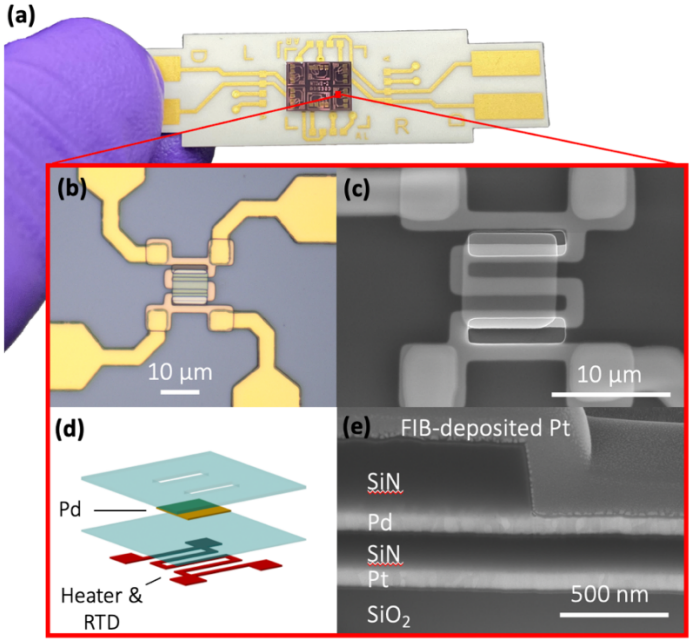
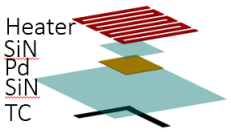
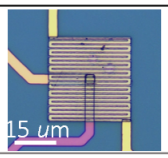
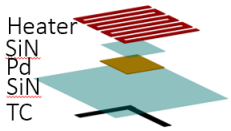
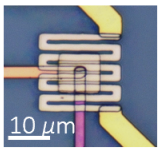
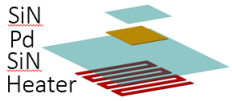
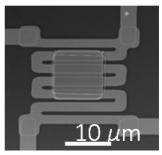
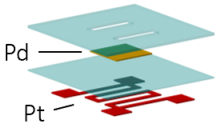
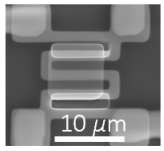
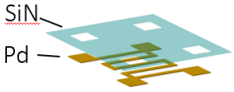
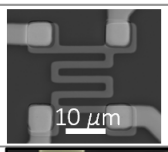
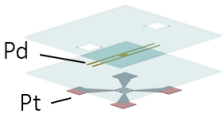
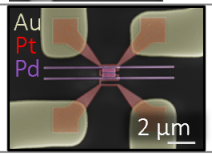


Figure 9. We designed the device for increased sensitivity by keeping the RTD size small relative to our Pd target, near the minimum of lithographical constraints. (b,c) Optical and SEM images of top-down view of device, respectively. (d,e) Diagram of device material stack and FIB cross-section, respectively. Following nanofabrication, we mounted the device on an alumina PCB (a) and then into the pressure testing vessel.

Table 1. Summary of the design iterations of nanoTI

Device	Description	Stack	Image
NanoTIv1	Pt heater / Pd target / Thermocouple *Contains Pd sizescales from 5-50 µm.*		
NanoTIv2	Pt heater / Pd target / Thermocouple *Contains Pd sizescales from 5 to 20 µm. Includes a variety of schemes, including a heating pad, a meandering heater, and a heater-RTD desgin*		
NanoTIv3 NanoTIv4 NanoTIv5	Pt-heater RTD / Pd target *Focuses on heater/RTD design. These device iterations focused on process improvements*		
NanoTIv6-1	Pt heater-RTD / Pd target *Contains Pd sizescales of 10µm, 30 µm. Also contains Pt control chips*		
NanoTIv6-2	Pd heater-RTD *Contains Pd sizescale of 15µm with varying geometries of SiN capping*		
NanoTIv7	Pt heater-RTD / Pd target *Contains Pd sizescales of 350nm, 1µm*		

(b) *Temperature coefficient of resistance (TCR) calibration:* Resistive sensing enables direct and rapid measurements of the state of the device, but resistance measurements are just a proxy for temperature and so an RTD requires calibration to accurately estimate temperature. Temperature coefficient of resistance (TCR or α) is described by the equation $\frac{R}{R0} = 1 + \alpha\Delta T$. For temperatures up to 200 °C, we used a Fluke 1944 with the test chip in an enclosed positive pressure nitrogen environment to prevent oxidation of materials during calibration heating. For temperatures above 200 °C, we heated the chip in a rapid thermal annealer (RTA) under positive pressure argon or in vacuum. We found that above 500 °C, the devices exhibited melting-type failures, due to void formation and agglomerations nucleating at defect sites and grain boundaries[24,25]. We applied temperature ramps to either the Fluke or the RTA and measured reference temperatures from a reference thermocouple in the RTA (Fig 10a,b). We measured device resistance using a Keithley

[24] Kane, W. M., Spratt, J. P., & Hershinger, L. W. (1966). Effects of Annealing on Thin Gold Films. Journal of Applied Physics, 37(5), 2085–2089. doi:10.1063/1.1708676

[25] Firebaugh, Samara L., Klavs F. Jensen, and Martin A. Schmidt. "Investigation of high-temperature degradation of platinum thin films with an in situ resistance measurement apparatus." Journal of Microelectromechanical systems 7.1 (1998): 128-135

DAQ6510 and plotted the resistance ratio versus temperature in order to approximate the TCR using a linear fit of data points taken after temperature equilibrated on the device (Fig 10c). TCR calibration enables estimation of temperature reached, but several factors add uncertainty to the estimation. First, TCR calibrations were not performed up to the highest temperatures likely reached in our experiments; they were performed up to $\frac{R}{R0}$ of ~1.8 while devices may reach $\frac{R}{R0}$ of ~2.9, which may result in underestimating temperature reached by ~100 °C. Secondly, stimulation itself results in material changes (Fig 11); higher stimulation energy results in higher temperature uncertainty, up to ~300 °C uncertainty. Temperature measurements near device melting are inherently unreliable. Thirdly, when performing TCR measurements on NanoTI devices using Pd RTDs (described below) as opposed to Pt RTDs, we have to consider that experiments are performed at ultra pressure on loaded Pd as opposed to TCR calibration near atmospheric pressure on unloaded Pd; temperature increase results in more electron scattering thereby increasing resistance while adding H or D to Pd also increases baseline resistance, so estimation of this effect on TCR is difficult. Pre-annealing the sample may reduce material changes due to stimulation input energy and loading[26]; none of the samples described in this report were pre-annealed, since our ultimate diagnostics were either a change in measured resistance or morphological changes.

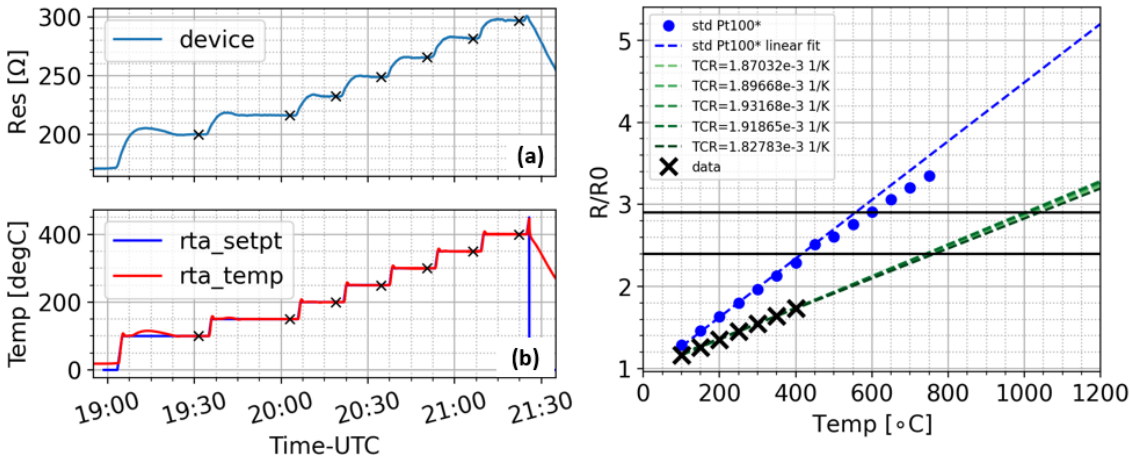


Figure 10. Example TCR calibration of NanoTIv7 performed in an RTA. To calculate TCR, we measure the device resistance (a) and the RTA temperature (b). We choose data at which temperature has settled, marked with an x, and plot the two data sets against each other (c, green lines). Is linear extrapolation a good fit for data in that temperature range? We see from data comparing raw data from a standard Pt100 resistor to its linear fit (blue dots and blue dashed line), that using the linear extrapolation of R/R0 can underestimate the actual temperature. For NanoTIv7 (green dashed lines), the relevant R/R0 range is ~2.4-2.9 which equates to estimated temperatures of ~800-1000 °C, based on linear extrapolation of three datapoints from the dataset.

[26] Frazier, Gary A., and R. Glosser. "Characterization of thin films of the palladium-hydrogen system." Journal of the Less Common Metals 74.1 (1980): 89-96.

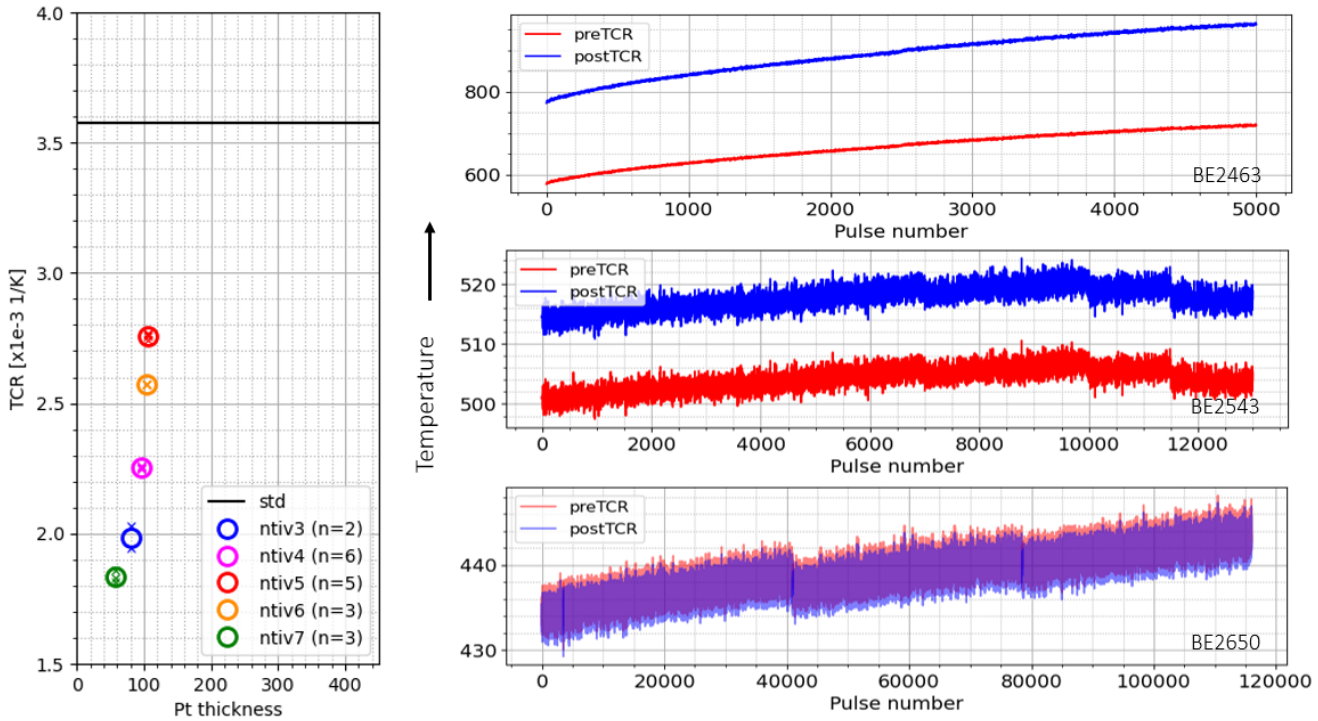


Figure 11. (Left) TCR roughly increases with film thickness. Higher TCR results in more sensitivity to temperature; for a given R/R0 range, the estimated temperature range is smaller. However, the fabrication process constrains the maximum film thickness. (Right) Effect of stimulation on TCR calibrations. Stimulating at higher temperatures resulted in a greater deviation between the pre and post TCR, resulting in larger uncertainty of the estimated temperature reached. These graphs show the estimated temperature during successive stimulation pulses based on the pre versus the post TCR. This data suggests that measured temperatures below 500 °C have uncertainty of 5°C due to stimulation alone, whereas temperatures below 900 °C have uncertainty near 300 °C.

(c) *Testing instrumentation and procedure:* Measurement speed and sensitivity were critical to our experiments. Off-the-shelf resistance measuring equipment is slow due to the time needed to calculate resistance, so we built a circuit to directly measure voltage, and we post-processed to find resistance ourselves, enabling us to measure at ≥1 MHz. We used a PXIe-5413 function generator for applying voltage to the device circuit, which consisted of the device in series with a current-sensing resistor. We measured voltages across the device and across the current-sensing resistor using a PXIe-DAQ6510, and post-processed using custom Python code to find the resistance and estimate temperature reached. In a typical experiment, we applied pulse trains of 100 µs period and 5 µs or 10 µs pulsewidth with voltages targeting a desired temperature, and recorded the data from those pulses (Fig 12a). As we continued our pressure experiments, we greatly increased our number of thermal stimulation pulses to increase the amount of material preconditioning and to increase event probability, as AHE literature frequently mentions the necessity for multiple loading/deloading cycles. For these experiments having a thousand or more thermal cycles, it becomes impossible to look through all pulses for anomalies. Therefore, what we do is plot a single datapoint from each pulse, for instance at 7us into the pulse, and look for anomalies (Fig 12b) over the collection of all pulses. This process is duplicated for 40 separate timepoints within a standard 10 msec pulse, and for each one, a plot is made showing the temperature at the given timepoint for each of the thousands of pulses; the collection of these 40 plots for one experiment is shown in Figure 12c. In the device BD4103-30umA, we saw that after 15000 pulses, we observed an exciting rise in temperature followed by loss of data as our sensor

broke (Fig 12b). However, subsequent experiments showed that this effect also occurs in both D_2 and H_2.

We performed data analysis to look for small temperature deviations indicative of AHE (Fig 12b,c). Higher stimulation energy may increase event probability. However, continuous escalation in this energy invariably approaches the material melting thresholds. Melting the sensor affects accuracy of the temperature diagnostic, so analysis is performed on data prior to melting. After the experiment, we use archaeology-type imaging to look for isotopic differences.

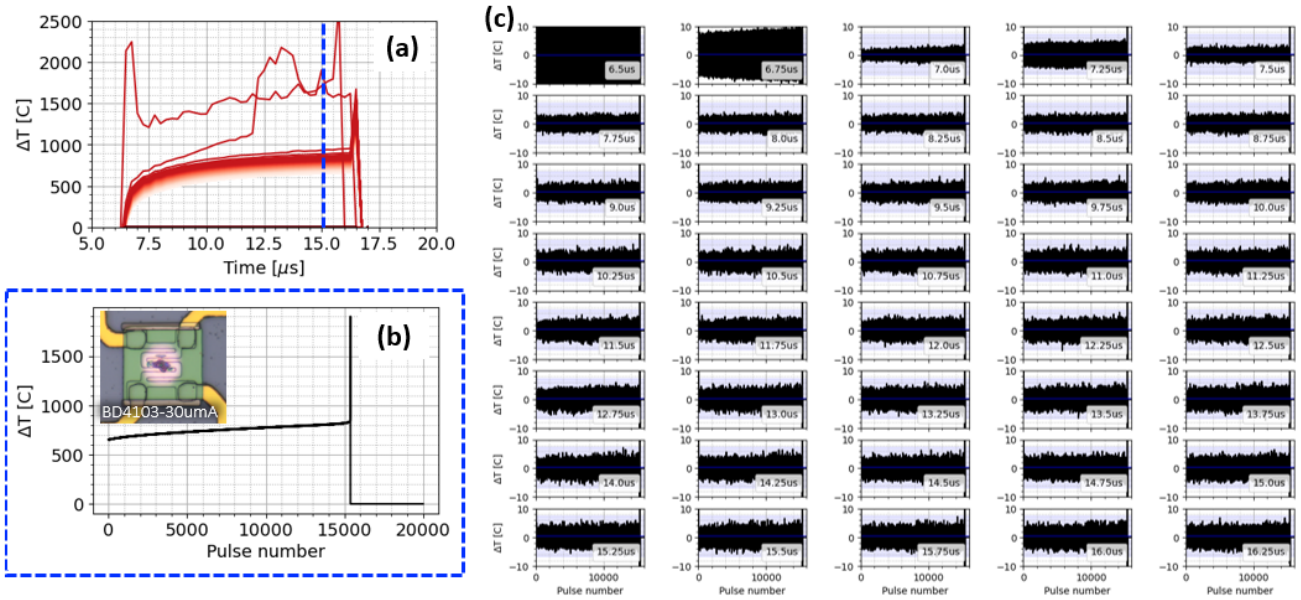


Figure 12. (a) Temperature data from a pulse train of 10 μs pulses, overlaid. The first pulse is in light red and the last pulse is in dark red. The temperature increases and then deviates from the mean prior to melting. (b) We examine the temperature at a specific timepoint into the pulse for all pulses in the pulse train, showing the clear temperature increase prior to device failure. The inset shows morphology of the melted device. (c) We examine this trend of temperature versus pulse number for all timepoints in the pulse prior to melting.

Loading kinetics: A main assumption behind NanoTI is that stimulation with high-temperature pulses on a short-enough timescale should prevent deloading. Another assumption is that capping the Pd film should help slow down deloading. We can test these hypotheses using Pd resistors. Additionally, investigations of thin film Pd measurements under gaseous conditions hold interest for hydrogen transport and storage applications. While Pd is a common hydrogenation catalyst for applications such as hydrogen storage, its kinetics under thermal or electrical stimulation are largely unexamined. Many investigations into Pd-H or Pd-D loading systems are carried out by electrolytic methods using Pd wire electrodes and thin film properties are relatively unexplored. Here, we present our key findings on thin film properties of Pd under stimulation.

We showed that we can disentangle the loading measurement from temperature and measure resistive loading curves under pressures of up to 500,000 psi; we found that thin film loading curves highly depend on intrinsic Pd film variances. We verified Pd loading qualitatively by observing reflectance decrease during loading and more quantitively with resistance measurements. One critique of measuring loading using resistance is that fundamentally, resistance

is a function of both loading and temperature. Measurement of resistance involves applying a voltage which results in some localized heating, especially when the device is designed to increase heating efficiency. Typical multimeters apply 0.7 μA – 1 mA of current to measure a 100 Ω resistor[27]. Figure 13a,b shows that resistance measurements of the Pd resistor device do not appreciably change within noise when applying less than 5 mA of current to the device, showing that the device does not measurably heat up. This factor allows us to apply a measurement offset to the stimulation pulse, enabling us to effectively measure loading (Fig 13c) as well as thermal decay (Fig 14a). We measured resistance during pressurization by applying test shots. We can then plot resistance versus pressure; assuming Zhang's curve[28] to be applicable for these thin films, we can estimate the loading fraction for each pressure point and therefore plot resistance versus loading fraction. First, we estimated loading based on resistance curves (Fig 13c). We overlaid nine curves from nine different devices and two device batches and found that the resistance versus loading curve highly depends on batch. This is more than likely due to Pd film thickness differences in the different batches.

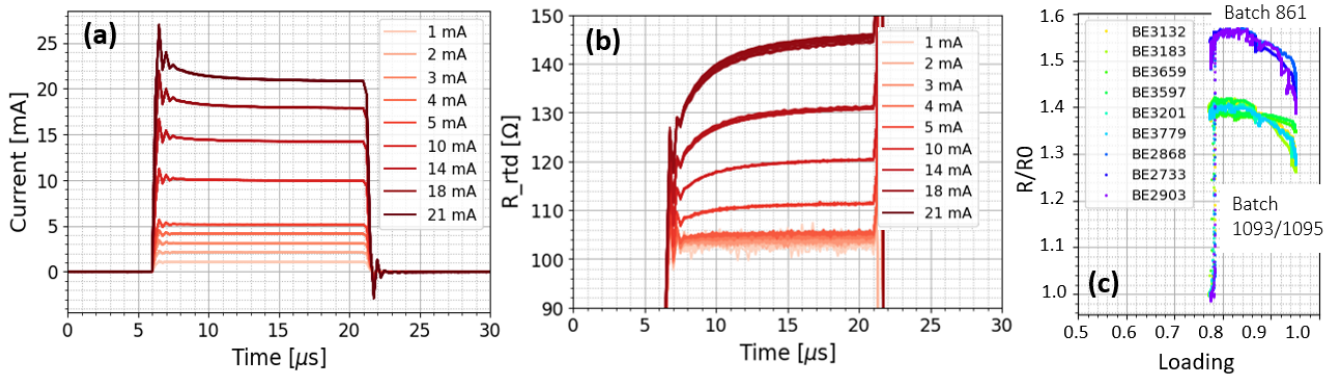


Figure 13: (a) Estimation of loading as a function from R/R0 data from two fabrication batches. (b) Current measured across device for increasing applied voltage, with the legend showing average current. (c) Calculated resistance from the data in (b), showing that below 5 mA, resistance and therefore temperature does not change to a large degree. This enables us to use <5 mA to minimize temperature-induced resistance changes and probe purely for loading.

[27] "Model DAQ6510: Data Acquisition and Multimeter System Specifications." Keithley: A Tektronix Company, April 2018, [https://download.tek.com/document/SPEC-DAQ6510A_April_2018.pdf]

[28] Zhang, Wu-Shou, Zhao-Fu Zhang, and Zhong-Liang Zhang. "Some problems on the resistance method in the in situ measurement of hydrogen content in palladium electrode." *Journal of Electroanalytical Chemistry* 528.1-2 (2002): 1-17.

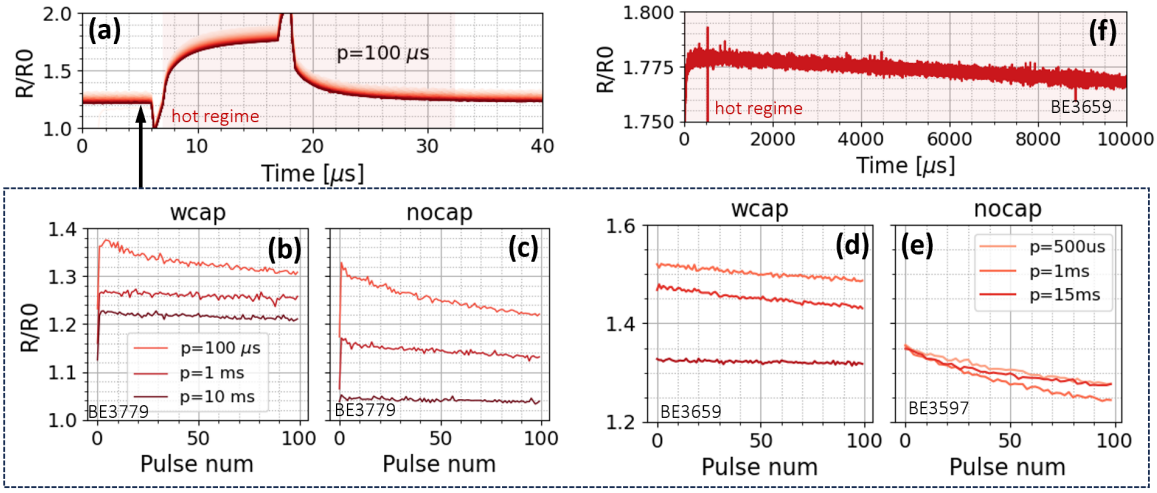


Figure 14: (a) For these loading kinetics tests, we measure prior to the hot regime where temperature should not be convoluted with loading and R/R0 values should be indicative of loading. (b,c) Under 5V stimulation (~450 C), the sample significantly deloads, even with increased pulse period for both the capped and uncapped devices. Capping does not appear to prevent reloading. (d,e) Under 3.2V stimulation (~150 C), the sample reloads for the uncapped sample but not the capped sample. (f) Loading a sample and then stimulating in a N_2 environment devoid of H_2 allows probing deloading timescales. This graph shows temperature data in the hot regime of a long 10 ms pulse, which allows viewing the transition from increasing resistance as temperature increases to decreasing resistance as deloading occurs; this transition occurs around 1 ms, so the timescales for deloading are in the regime of ~1 ms.

We investigated sample loading and deloading kinetics with varying parameters, including pulse period, pulse width, hydrogen loading state, and the effect of nitride capping (Fig 14). For these tests, we measured the sample shortly before the hot regime to avoid convolution of loading with R/R0. We performed testing in two types of conditions: in the first, we loaded samples with deuterium and then switched to a 145 ksi nitrogen environment devoid of hydrogen (Fig 14a,f), estimating that the sample is ~60-65% loaded during experimentation. In the second, we tested at 145 ksi D_2, which is equivalent to~96% loading (Fig 14b-e). NanoTI pulses are kept short at <20 µs to avoid deloading, however we found that pulse widths up to 100 µs do not appear to result in further deloading. Deloading happens around 1 ms for ~150C stimulation (Fig 14f), indicating that short pulses may not be necessary purely from a deloading perspective. Under short pulse stimulation, the R/R0 vs pulse number curve has different expected behavior depending on loading regime. At low loadings, R/R0 is before the loading hump (loading curve depicted in Fig 13c), so deloading results in decreasing R/R0. At high loadings past the loading hump, deloading results in an increase followed by a decrease in R/R0. For all loading cases, capping did not necessarily help prevent deloading, but it did prevent reloading (Fig 14b-e). This informed the design of nanoTIv7 which did not have a SiN capping layer. Higher pulse periods do result in more time to reload between pulses, resulting in a less sloped curve (Fig 14b-e), though more quantitative comparison of the effect of pulse period between test conditions is difficult since the starting loading values of different test conditions are not the same. At low loadings, all samples return to baseline R/R0. We suspect that every high temperature stimulation pulse (Fig 14b,c) likely deloads the sample, but the sample reloads under a high-pressure hydrogen environment.

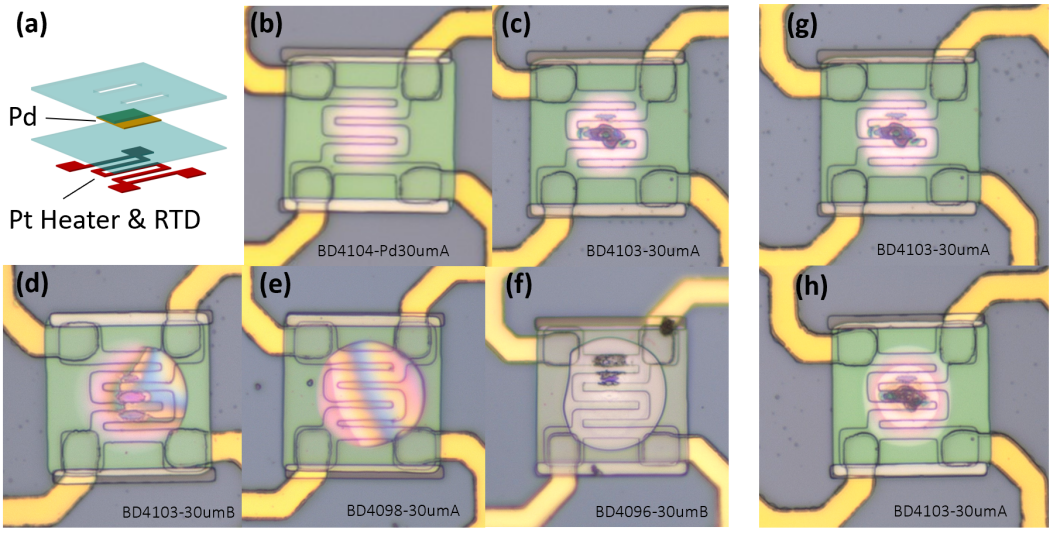


Figure 15. (a) Diagram of Pt heater-type design of NanoTI. (b-f) Evolution of nitride morphology change. The refractive index changes (b-c) due to changes in the H_2 content in the PECVD nitride film. Temperature changes and CTE (coefficient of thermal expansion) mismatch result in the nitride beginning to delaminate (d-e), before fully delaminating off the device (f). The same sample after stimulation under hydrogen environment (g) and after reloading under hydrogen environment without stimulation (h) exhibits an annular coloration change, indicating that the color change does not depend upon thermal stimulation and is likely solely due to gaining hydrogen.

We investigated the effect of capping the samples for the purpose of preventing deloading and found clear coloration differences in optical images of the SiN capping layer during experiments. SiN is a great passivation material for a variety of reasons -well-characterized, easy-to-develop, high breakdown strength. SiN deposited by PECVD acts as a very efficient hydrogen barrier because of its high hydrogen content[29].The observed optical changes in the film (Fig 15) lead us to a few speculations. Either the H has diffused out under thermal stimulation or more H has loaded in under thermal stimulation. If the hydrogen has diffused out, it would diffuse into the underlying palladium and this would be evident in resistance data. However, testing under nitrogen shows the same behavior in both a sample capped with SiN and a sample without the SiN capping, leading us to infer that that the color change in nitride is due to gaining hydrogen.

Results: The bulk of our experiments were performed using platinum resistors which acted as both the stimulation source and the sensor. Extensive experimentation enabled us to successfully target various temperature regimes. We first targeted a high thermal stimulation regime which results in material melting of both the Pd as well as the Pt heater-RTD sensor; this of course results in changes to the temperature data near the melting point of the Pt as the resistor deforms. Temperature data deviations near the melting point resulting from morphological changes of melting cannot be easily distinguished from anomalous nuclear events. For these devices we analyzed temperature data prior to the data deviations for small spikes which may be indicative of events (as described in the methods section) and found nothing of significance above the noise floor. Post-imaging confirmed melting morphology of both the Pd target and Pt heater-sensor. Repeated tests in both D_2 and H_2 reveal no morphological differences. After determining the

[29] Nemanič, Vincenc, et al. "Hydrogen permeation through silicon nitride films." *Journal of alloys and compounds* 539 (2012): 184-189.

required input energy for melting, we backed off and targeted melting of only the Pd; this enabled us to analyze the entire temperature data set.

We repeated these tests using ultra high-pressure hardware targeting >450 ksi, enabling testing at full stoichiometric loading. Figure 16a-b shows the first case of targeting Pd melting only, and representative experiments in D_2 and H_2. The morphology shows that the Pd melted, and we do see that the nitride has completely delaminated in the D_2 case. Again, we see no qualitative evidence in the temperature data or morphology data to indicate that we definitively have AHE. Then, we tried the second case of targeting both Pd and Pt melting; example experiments in D_2 and H_2 are shown in Figure 16c-d. Melting both the Pd and Pt results in changes to the temperature data; however again, we see no qualitative evidence in the temperature data or morphology data to indicate we have AHE.

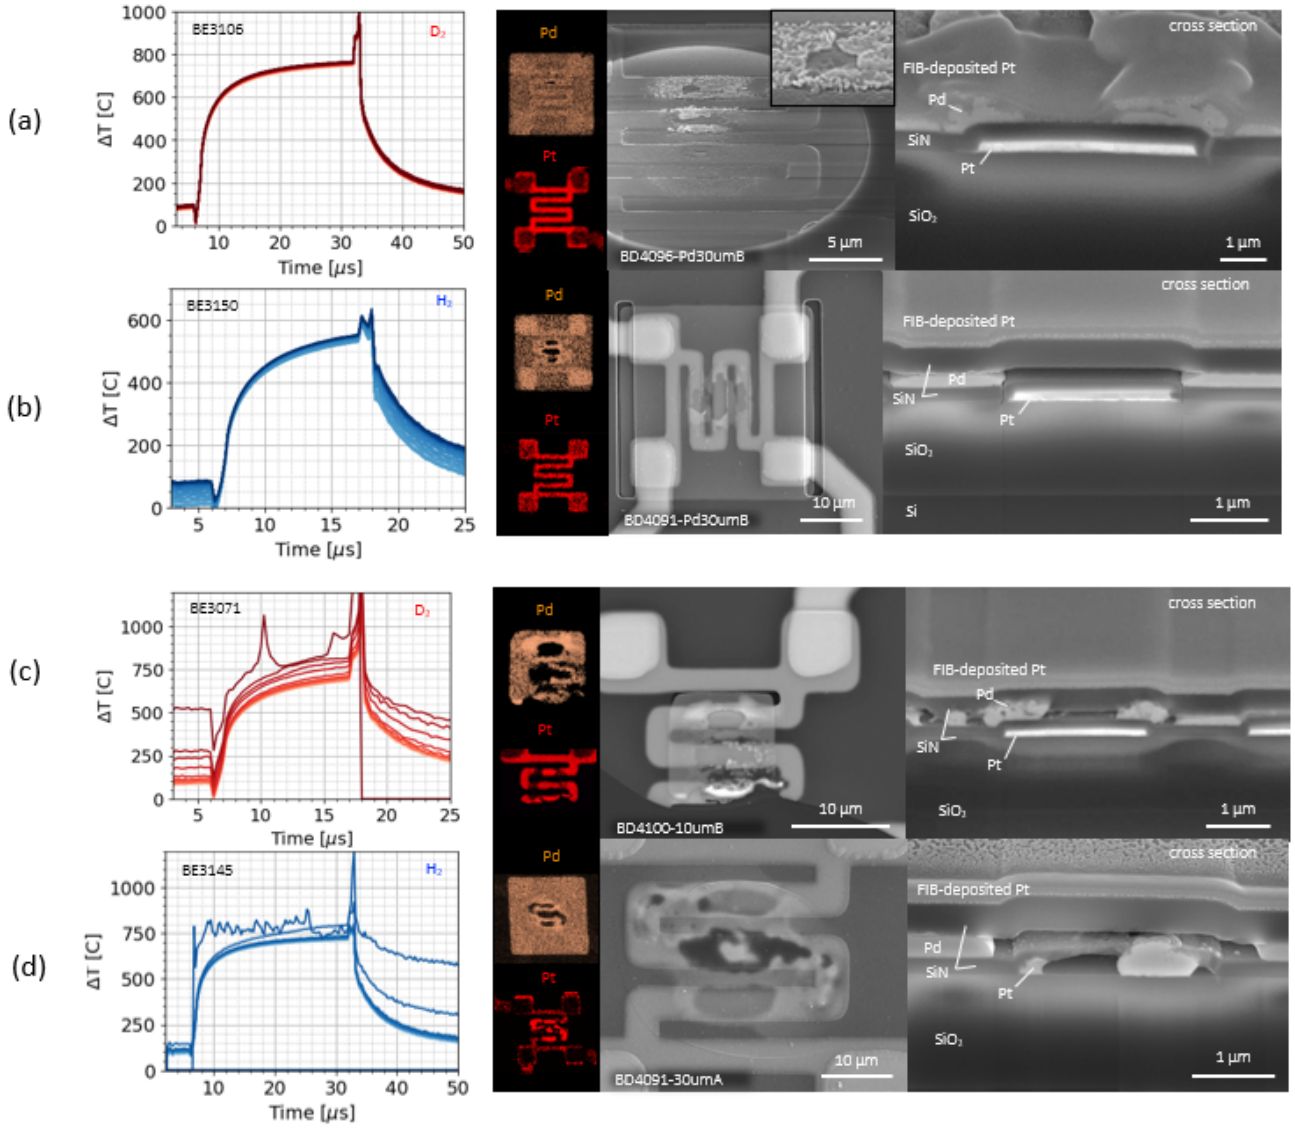


Figure 16: Experimental results from testing Pt resistors. Examples of temperature and archaeology data from searching for AHE at >450 ksi (stoichiometry) for testing below melting (a-b) and up to melting (c-d) under both deuterium (a,c) and hydrogen (b,d). For all experiments, we also performed analysis of temperature vs pulse number (see Fig 12). All experiments showed similar trends in the R/R0 and the temperature characteristic. Post-imaging shows similar characteristics in the hydrogen controls. The behavior is identical for 145 ksi.

In the Pt resistors, we were conducting heat to and from the Pd, but another approach was to directly stimulate the Pd. This has the advantage of directly probing AHE as well as directly measuring loading. The drawback is that we use resistance to measure both temperature and loading, and these may be convoluted when the chip is stimulated (as described in Experimental Methods). When running an experiment at pressure, we use our ability to verify that resistance stops changing to confirm that we are fully loaded at that isobar prior to stimulation. We are also able to monitor the resistance before and after each stimulation shot to verify loading.

For these Pd resistor devices, we targeted a stimulation regime which enables us to test for many cycles prior to device melting. As for all NanoTI-type devices, we tested at 145 ksi and at near

unity loading. In all cases, we saw no isotopic difference in resistance data, and we saw a more isolated melt morphology that persists at higher pressures in hydrogen as well as deuterium. For these Pd resistor chips as well, we have not yet observed any anomalies indicative of AHE. None of the NanoTI Pt heaters or Pd heaters have shown clear evidence of AHE across >448 devices.

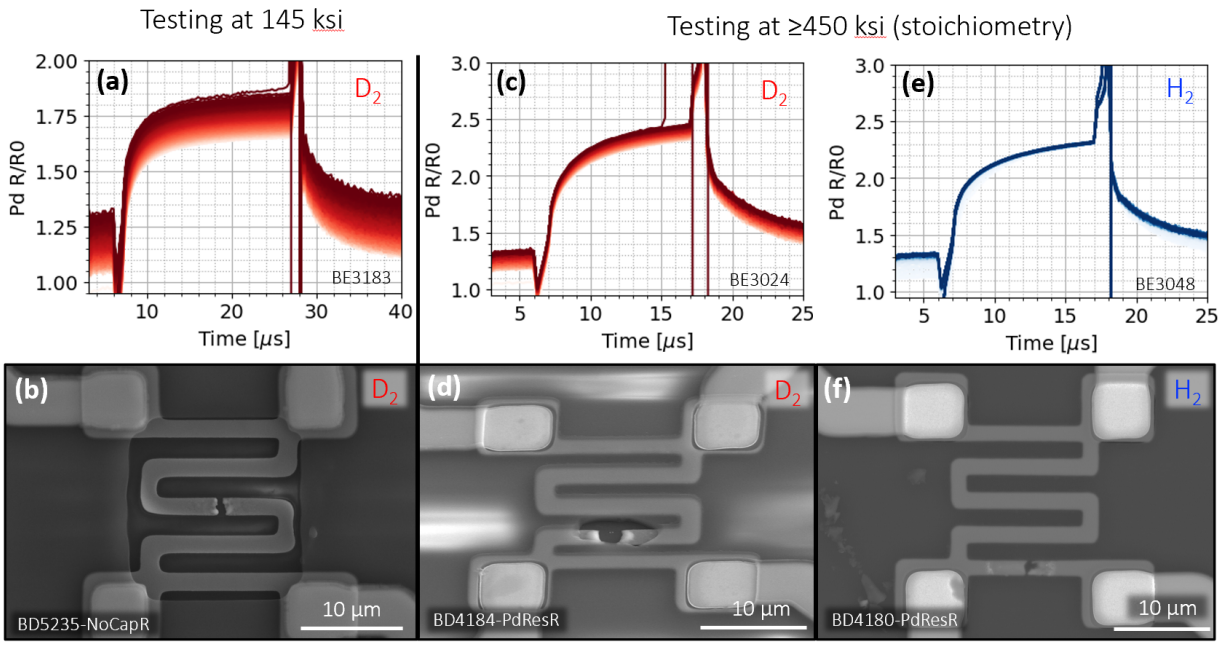


Figure 17: Experimental results from testing Pd resistors. Examples of temperature and archaeology data from searching for AHE at 145ksi (a-b) and at >450 ksi (stoichiometry) (c-f). For all experiments, we observed similar trends in the R/R0 or temperature characteristics. Post-imaging shows that the isolated melt morphology persists at stoichiometry and in the hydrogen control.

Table 2. Selected experiments are listed here to highlight representative test conditions

Experiment / Sample	Description	Gas	Pressure	Test condition
BE3106 / BD4096-Pd30umB	NanoTIv6 – Pt heater/RTD	D_2	≥450 ksi	Below melting
BE3150 / BD4091-Pd30umB	NanoTIv6 – Pt heater/RTD	H_2	≥450 ksi	Below melting
BE3071 / BD4100-10umB	NanoTIv6 – Pt heater/RTD	D_2	≥450 ksi	Above melting
BE3145 / BD4091-30umA	NanoTIv6 – Pt heater/RTD	H_2	≥450 ksi	Above melting
BE3183 / BD5235-NoCapR	NanoTIv6 – Pd resistor	D_2	145 ksi	Above melting
BE3024 / BD4184-PdResR	NanoTIv6 – Pd resistor	D_2	≥450 ksi	Above melting
BE3048 / BD4180-PdResR	NanoTIv6 – Pd resistor	H_2	≥450 ksi	Above melting
BE3533 / BD6048-1umC	NanoTIv7 – Pt heater/RTD	D_2	145 ksi	Below melting
BE3537 / BD6049-1umC	NanoTIv7 – Pt heater/RTD	H_2	145 ksi	Below melting
BE3438 / BD6052-1umB	NanoTIv7 – Pt heater/RTD	D_2	145 ksi	Above melting
BE3502 / BD6050-1umB	NanoTIv7 – Pt heater/RTD	H_2	145 ksi	Above melting

Chips : Pillars on Pads

Introduction: Pillars on Pads combines the rapid, on-chip thermal cycling approach of NanoTI with the palladium nano-pillar arrays to enable single event detection. This technique utilizes a thin film platinum resistor to quickly heat palladium nano-pillars loaded with deuterium from high pressure D_2 gas, thereby stimulating phonon generation and deuteron flux. The design, shown in Figure 18, features the palladium sample structured as an array of nano-pillars. The pillars are designed to have minimal diameter and thermal mass so that a single AHE event (24MeV) depositing its energy at any lattice site should lead to melting or ablation of the host pillar, effectively turning each nano-pillar into a single event detector. Calculations suggest pillar diameters of less than 100nm are needed for a single 24MeV event, if the energy is entirely thermalized to the pillar in order to produce visible damage. Post experimental analysis uses scanning electron microscopy (SEM) to identify any missing or altered pillars, and comparing the pre and post-analysis determines whether any AHE events occurred throughout the experiment.

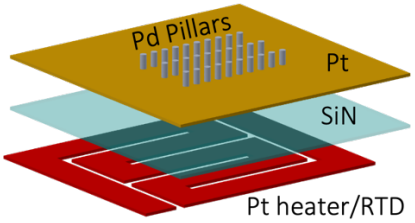


Figure 18. Diagram showing the Pillars on Pads material stack and heater/RTD geometry.

The stimulation of these pillars is achieved through the application of 10µs duration voltage pulses along the serpentine platinum resistor. These pulses induce joule heating in the resistor which conductively heats the pillars via the chip's material stack. To measure the relative temperature changes in the platinum, the voltage drop across the resistor was measured during stimulation pulses and the calculated resistance was correlated to previously established resistance-to-temperature calibrations. The combination of the chip's small thermal mass in the active region and the ultrapressure gas environment's strong cooling power allow the device temperature to return to ambient within tens of microseconds, as shown in the simulation in Figure 19. Another change which occurs with each pulse, is a partial deloading of deuterium as temperature rises, followed by reloading in the interval between pulses; both happening quickly due to the small size of the pillars. This rapid return to baseline, for both temperature and loading, enables the device to undergo thousands of temperature cycles per second.

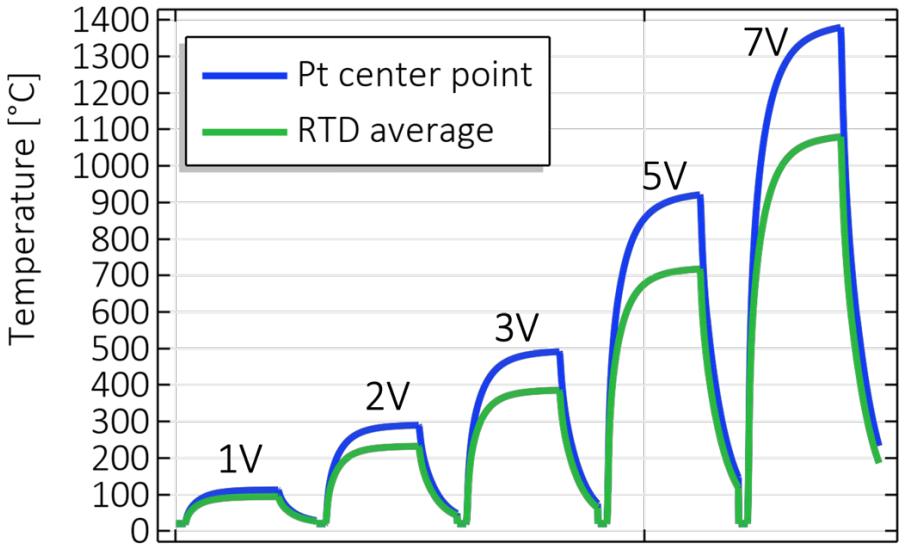


Figure 19. Simulations showing the expected temperature rise for different pulse amplitudes in a 145ksi H_2 environment.

This approach has enabled precise control of the peak stimulation temperature all the way to the melting point of palladium while delivering single event resolution.

Design and Fabrication: Leveraging insights from our NanoTI initiatives and previous nano-pillar experiments, we have refined the Pillars on Pads chip design and fabrication process. A re-engineered platinum heater/Resistance Temperature Detector (RTD) allows for a higher density of nano-pillars on each chip while streamlining the electroplating workflow. Utilizing electron beam lithography with a dual-layer resist, we achieved high-resolution patterning for the platinum heater/RTD, ensuring a smooth surface topology. We then deposited four gold contacts for precise current delivery and monitoring. A silicon nitride layer was deposited for electrical insulation, followed by a gold seed layer to facilitate palladium electroplating. Subsequently, electron beam lithography defined an array of holes in the e-beam resist atop the platinum RTD/heater arms, through which palladium was electroplated to construct the nano-pillars as shown in Figure 20. Through much experimentation, we finetuned the e-beam lithography and electroplating process to achieve a consistent pillar profile. Higher e-beam doses during lithography and longer development time lead to more consistent wetting of the gold seed layer during electroplating but also result in larger pillar diameters. Conversely, using too low dose or development time leads to missing pillars due to residual resist at the base of the hole. Our process was optimized to result in pillar diameters of approximately 70 nm with a pillar yield of >99%. Additionally, a non-lithographic technique was developed for reducing pillar diameters as low as 25 nm using Ar^+ ion milling at grazing incidence angles such that the pillar wall is preferentially sputtered away. This modification enhances the pillars sensitivity, potentially allowing for more obvious damage to the pillar.

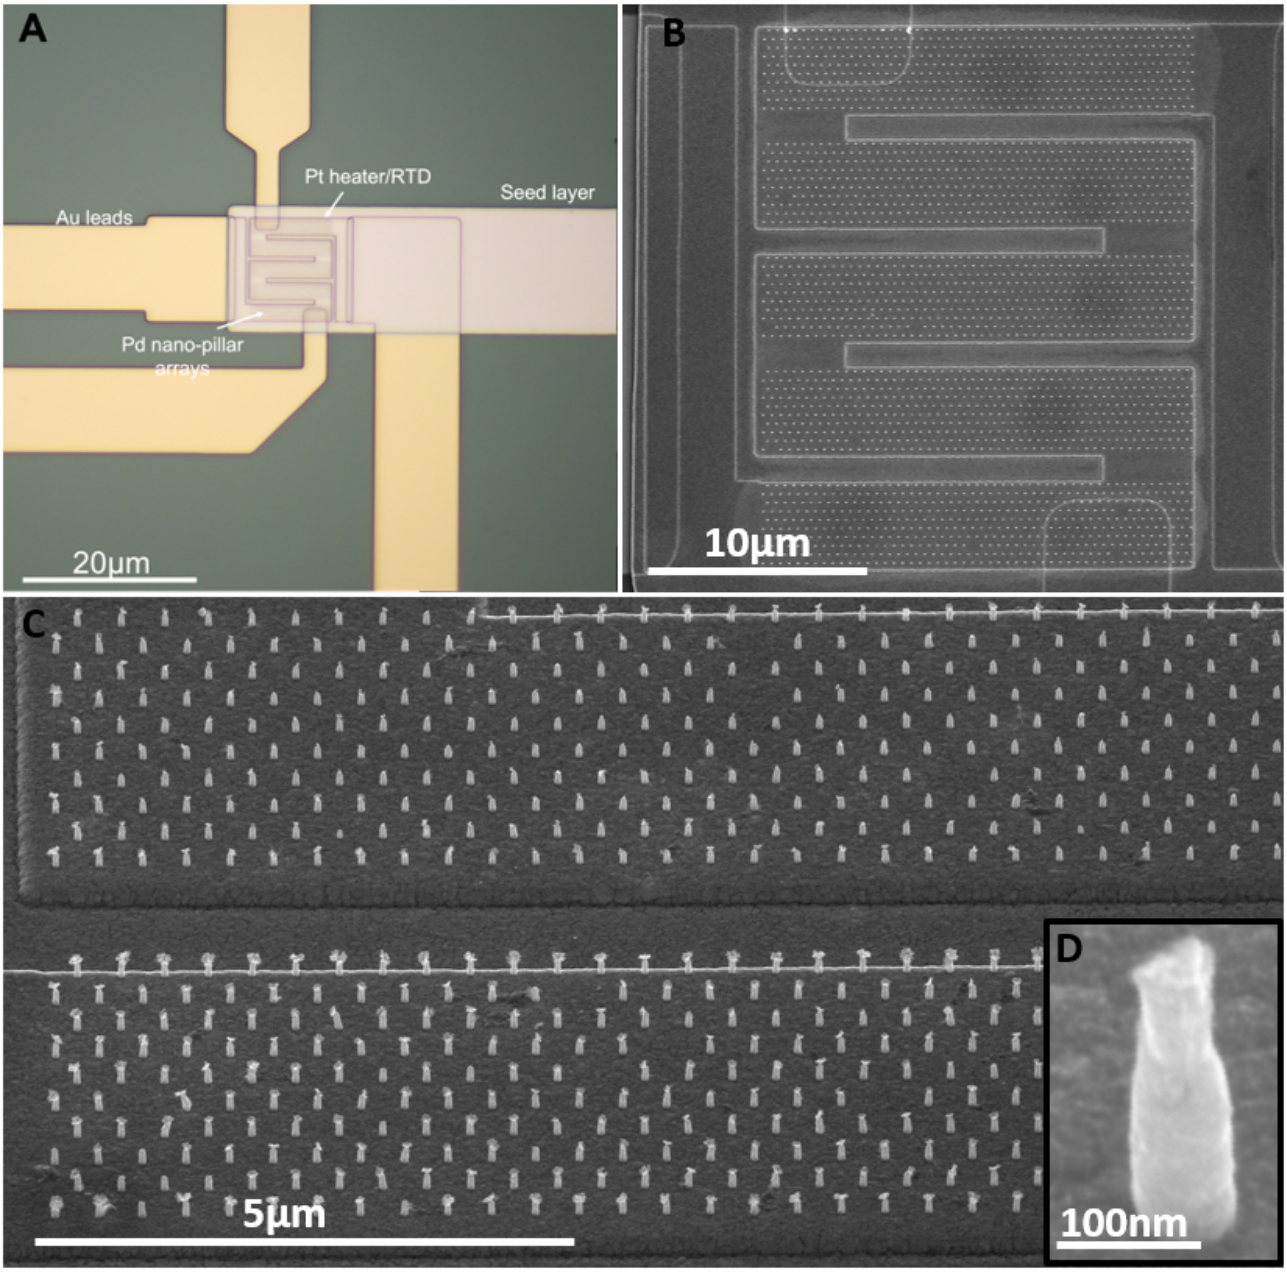


Figure 20. A) Confocal image of Pillars on Pads device. B) Top down SEM micrograph of a Pillars on Pads device. C) SEM micrograph of Pd nano-pillars on the device at a 52degree tilt. D) SEM zoom in for an individual pillar viewed at 52degree tilt.

Experimental Methods: In our experimental setup, palladium nano-pillars were loaded in ultra-pressure hydrogen or deuterium gas, with testing pressures between 100,000 and 500,000 psi. Thermal stimulation of the pillars was achieved by delivering electrical pulses through the platinum heater. These pulses, 10 µs in duration and spaced by 100 µs periods, were meticulously calibrated. The voltage of each pulse train was systematically varied from 1 to 8 volts. Initial

calibration involved a progression of pulses with incrementally increasing voltages to precisely establish the peak temperature profile before the main stimulation pulse train.

As discussed in the NanoTI section, to ensure accurate temperature measurements, we assessed the Temperature Coefficient of Resistance (TCR) for each device within a precision temperature-controlled thermowell, under an inert atmosphere, from 50°C to 250°C. This pre-experimental calibration was critical because the intense pressure of the gases could potentially alter the resistance values and, consequently, the resistance-temperature relationship of the platinum heater/RTD. Despite platinum's TCR generally exhibiting linearity within the -100°C to 250°C range, our observations confirmed that this relationship becomes non-linear at the higher temperatures induced during the experiments.

A drift in resistance in the platinum heater/RTD was observed during the high-temperature cycling, which was not entirely unexpected given the extreme conditions. The as-deposited platinum has very small grain size and becomes annealed during the stimulation pulse trains. However, this drift and the altered post-stimulation TCR values presented too much uncertainty to report an accurate peak temperature. One other concern was simply that the RTD occupied the full pad area, so could only measure an average temperature of the pad, not the peak temperature in the center of the pad. Given the challenges in calibrating TCR at temperatures exceeding 900°C within an ultra-high pressure environment, and the spatial variation in the temperature across the pad, we adapted our methodology to report “temperatures” in terms of resistance ratios, denoted as R/R_0. These ratios were then empirically linked to melting points, recognizing the variations due to the nonlinear TCR, the inherent differences in baseline resistance, and the differences in cooling effects between the different pressures and process gases used.

Over the course of our experiments, one evolution of the device design, from using a gold to a platinum seed layer beneath the Pd pillars, was instrumental. The gold layer, while initially serving as a suitable contact for electroplating, resulted in surface roughening that interfered with the SEM post-experiment analysis when we began exploring higher stimulation temperatures. This roughening, identified as partial gold layer melting when the R/R_0 exceeded approximately 1.6, was eliminated by transitioning from a gold seed layer to a platinum seed layer. This change facilitated higher stimulation temperatures without inflicting thermal damage to the chip's structure.

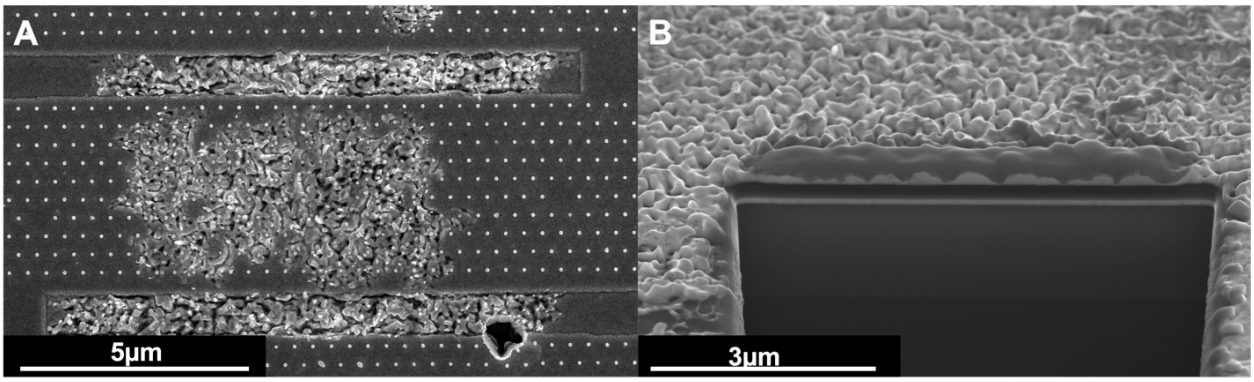


Figure 21. A) SEM micrograph of thermally damaged chip following high temperature stimulation. B) FIB cross section confirming the melting is in the Au seed layer.

With the enhanced chip design, we proceeded to identify the damage threshold for the pillars. This was achieved through a series of pulse trains, each consisting of 10,000 thermal cycles, with progressively increasing voltage up to 8 V under a pressure of 145 ksi with D_2. A peak R/R_0 ratio of approximately 2.3 was the critical point at which pillars in the central, hottest zone of the chip were observed to be missing. Notably, when the pulse amplitude was slightly reduced under identical cycling conditions, the pillars remained intact, as confirmed by post-experiment characterization. Replicating this process using 145ksi H_2 yielded comparable results, suggesting that thermally stimulated palladium does not damage more easily when loaded with deuterium compared to protium.

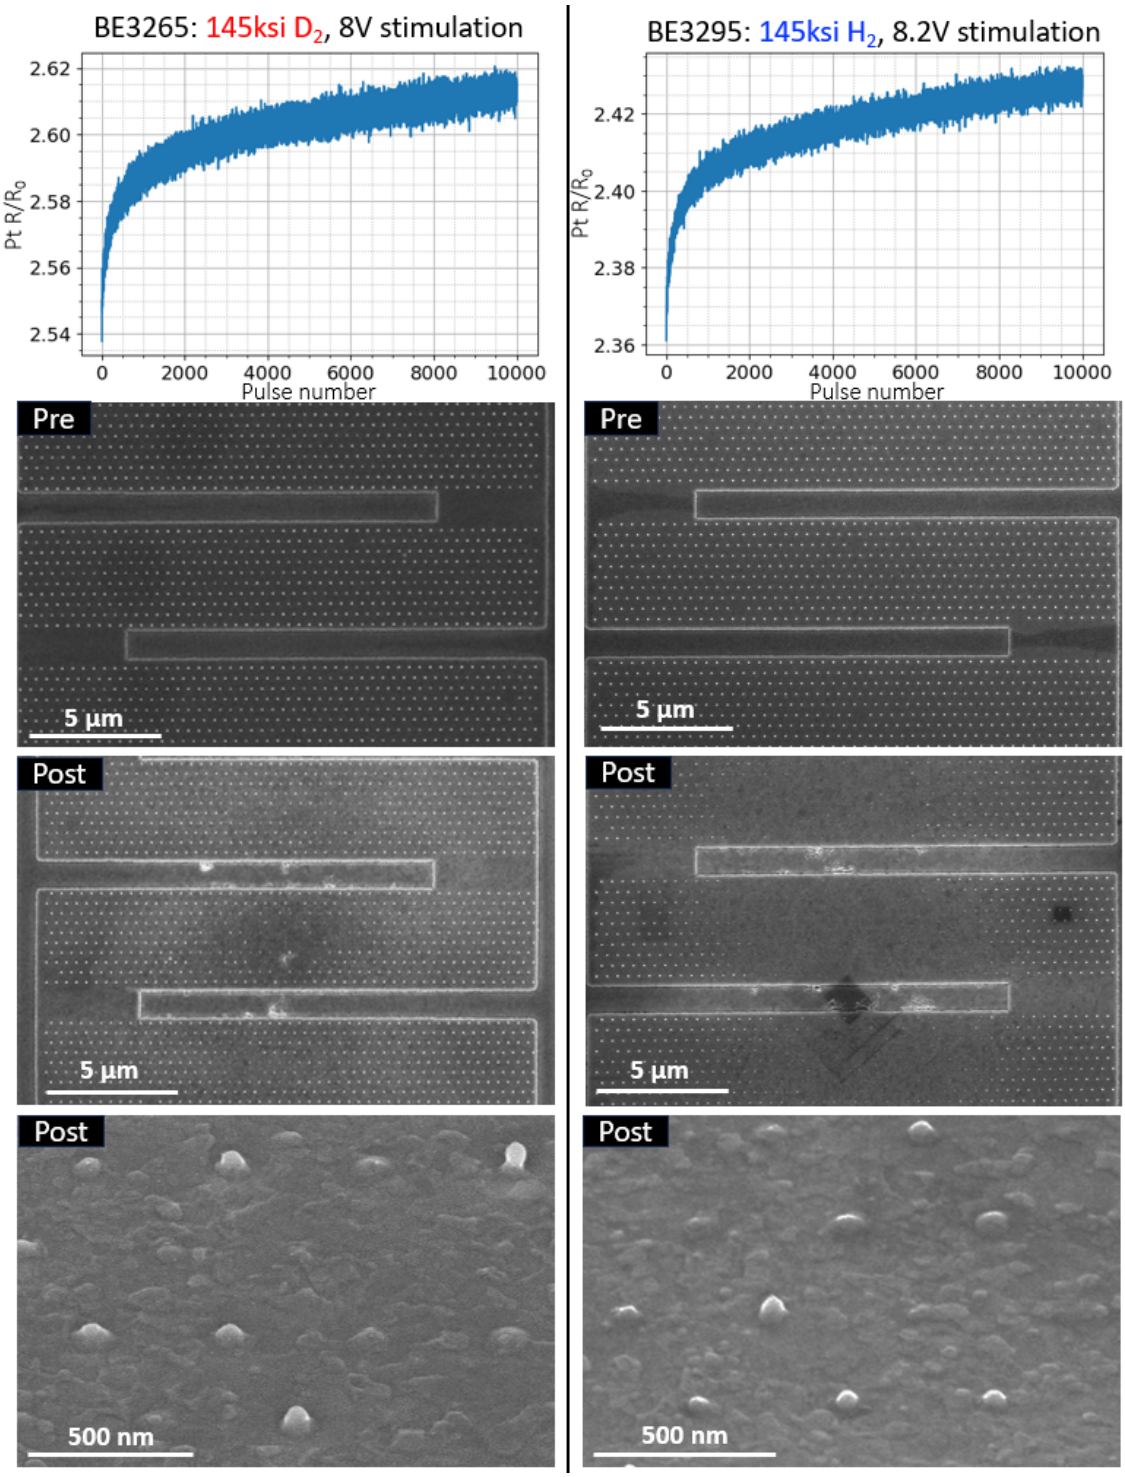


Figure 22. Experimental comparison of pillars stimulated above damage threshold under protium vs deuterium. 10,000 thermal cycles were used to stimulate loaded palladium nano-pillar arrays. Due to the lower molecular weight, H_2 gas has a higher cooling power for the same pressure, requiring more power to reach the same peak stimulation temperature. The plots (D_2 on the left, and H_2 on the right) show peak R/R_0 for each pulse in the pulse train. Pre and post SEM analysis show pillars ablated during stimulation at lower temperatures in H_2 compared to D_2, suggesting no AHE or self-heating occurs in D_2.

An additional experimental campaign was then run using these devices at higher pressures where the palladium nano-pillars are expected to be at unity loading, i.e., where 100% of the octahedral sites in the palladium lattice are filled with a proton or deuteron. These experiments again showed no indication of excess heat generated in deuterium loaded Pd compared to protium loading.

Table 3. Pillars-on-Pads experiments

Experiment	Gas	Pressure (ksi)	Thermal Cycles
BE3265	D_2	145	10,000
BE3295	H_2	145	10,000
BE2940	D_2	300	20,000
BE3008	D_2	380	50,000
BE3126	H_2	300	50,000
BE3200	D_2	460	10,000
BE3335	D_2	450	10,000
BE3390	D_2	470	10,000

We have theorized that high fractional loadings are necessary, but not sufficient for AHE. When earlier experiments did not show signs of anomalous events at >90% loading, we theorized that the vacancy fraction may be important i.e., the difference between 90% and 95% loading is ~5%, but the difference between the corresponding 10% and 5% vacancies is a factor of 2. Chips projects including NanoTI and Pillars-on-Pads have put this hypothesis into question by not observing qualitative differences between low-90% loading and unity PdD loadings during full AHE experiments (stimulation & detection).

Chips : Heat Pad Archaeology

Introduction: In addition to using nano-pillars for single event detection or 'archaeology', palladium films of sufficiently small thickness should be unable to dissipate 24 MeV of energy impulsively dumped into the lattice, without undergoing local vaporization of the palladium. This vaporization will create a hole in the thin palladium film, exposing the underlying substrate. The contrast between the exposed substrate within the hole and the palladium film surrounding it will then be visible during SEM imaging. While both nano-pillar and thin film array targets were fabricated and deployed to other stimulation campaigns, standalone, on-chip stimulation/detection devices were also tested.

Design and Experimental Methods: The first tests of on-chip thin film archaeology devices utilized the basic chip designs which had proved reliable for Pillars on Pads. These chips were modified by replacing the seed layer and electroplated palladium nano-pillars with a

lithographically patterned square of 25 nm thick palladium on top of the silicon nitride passivation above the serpentine platinum heater/RTD.

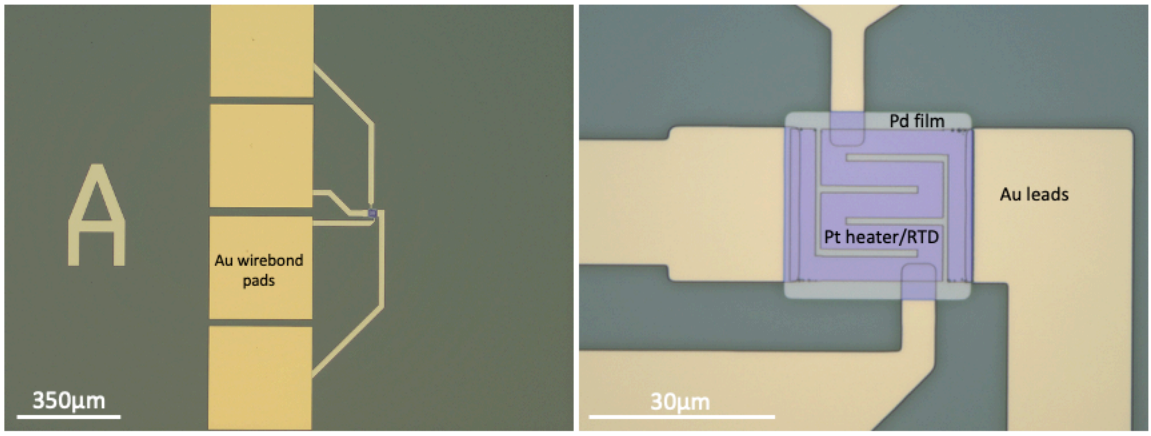


Figure 23. Annotated confocal microscope images of Heat Pad Archaeology V1 devices.

Before each experiment, the palladium thin film was extensively pre-imaged in an SEM. Because this chip platform was extensively tested in both the 150 ksi and 500 ksi H_2/D_2 systems in Pillars on Pads experiments, minimal testing was performed in the 150 ksi system. As with Pillars on Pads testing, stimulation was achieved by electrically pulsing the platinum heater/RTD with 10 µs pulses separated by 100 µs. Typically, 10,000 pulses were used in each stimulation experiment. The figures below (Fig 24, 25) show a hydrogen and deuterium experiment at unity loading (i.e., stoichiometric PdH or PdD) below the thermal damage threshold of the chips.

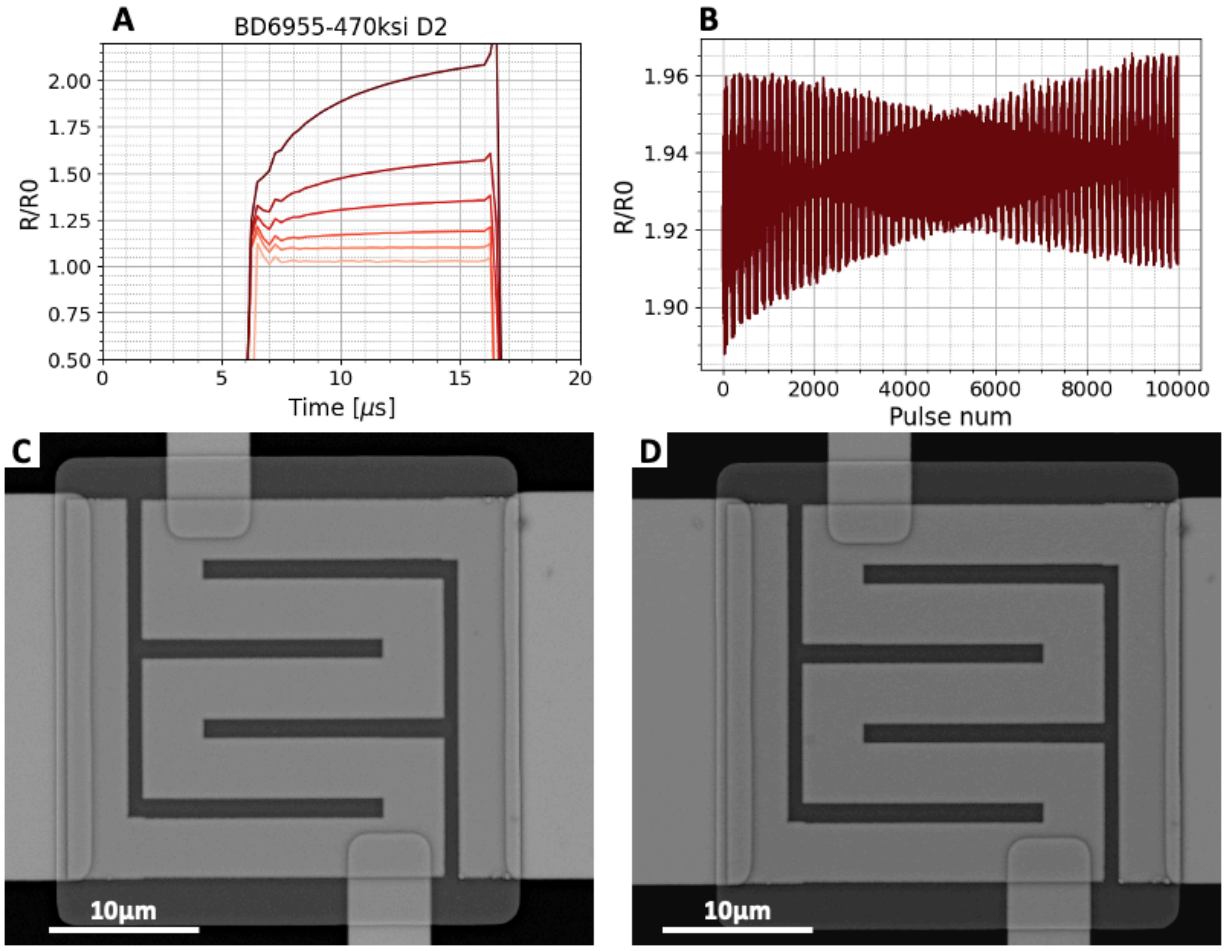


Figure 24. Callisto experiment BE3612. The chip was loaded with deuterium gas at 470ksi. A) Calibration shots. A voltage sweep is used to determine the pulse amplitude to be used in the pulse train. B) The R/R_0 for each shot in the stimulation pulse train taken at t=16μs. The apparent electrical noise present in this data was an anomaly not seen in other experiments. C) Pre-experiment SEM of the palladium film. D) Post-experiment SEM of the palladium film showing no evidence of film damage. Many higher resolution images were taken across the sample confirming the film was not visibly changed through the experiment.

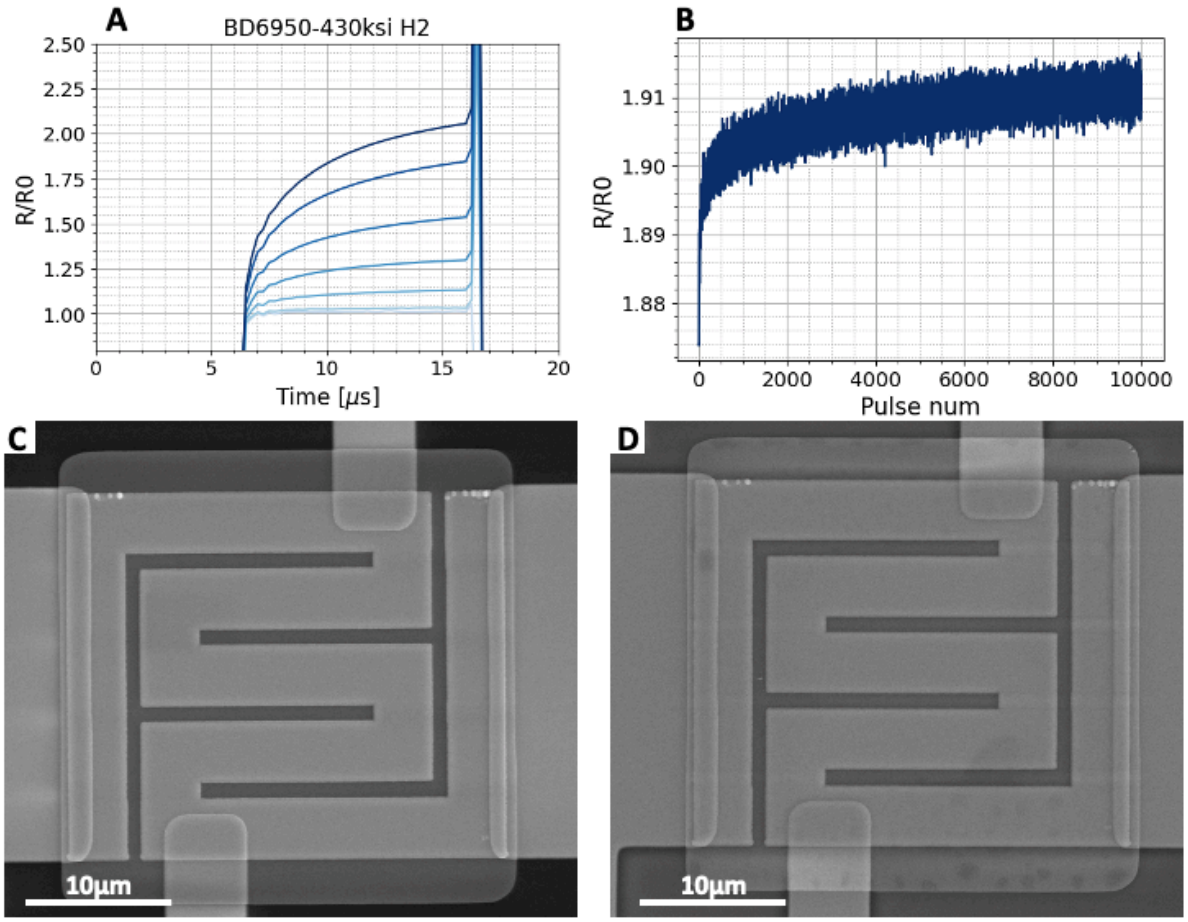


Figure 25. Callisto experiment BE3615. The chip was loaded with protium gas at 430ksi as a control test. A) Calibration shots. A voltage sweep is used to determine the pulse amplitude to be used in the pulse train. B) The R/R_0 for each shot in the stimulation pulse train taken at t=16μs. C) Pre-experiment SEM of the palladium film. D) Post-experiment SEM of the palladium film showing no evidence of film damage. Many higher resolution images were taken across the sample confirming the film was not visibly changed through the experiment.

Each subsequent experiment targeted a higher temperature until changes to the film were observed.

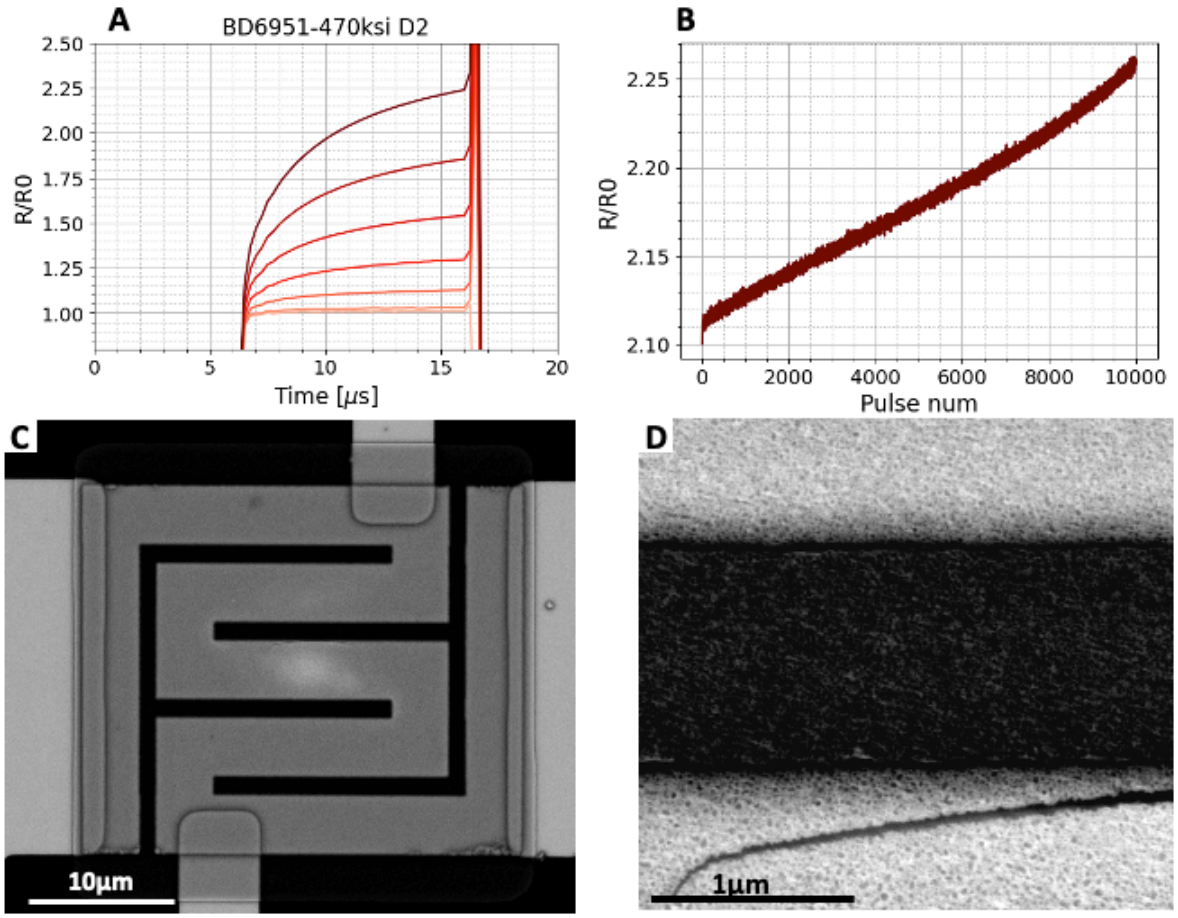


Figure 26. Callisto experiment BE3623. The chip was loaded with deuterium at 470ksi. A) Calibration shots. B) The R/R_0 for each shot in the stimulation pulse train taken at t=16μs. The positive slope in this plot is indicative of changes occurring in the platinum heater/RTD due to the high temperatures and is typically observed just before heater failure. C) Post-experiment SEM of the palladium film. A lighter region is observed in the region of the device which gets the hottest during stimulation. D) Higher magnification SEM of the bright region. Nano-pore are observed to have formed during stimulation.

Experiment BE3623, shown above in Figure 26, pushed stimulation to the very limit of the chip's capabilities. Approaching the melting point of palladium is expected to cause agglomeration of the film due to surface tension and on short time scales this can yield a nano porous film, however, activation of AHE in these films is expected to present as nano-holes and so their appearance was a result of interest. The question, of course, was whether they were due to near-melting stimulation, or due to AHE. Many attempts were made to replicate this experiment in protium but every following chip failed during the pulse train after matching calibration shots (Figures 27, 28). Eventually, attempts were made to replicate this result in deuterium with slightly lower temperatures but nano pores were never observed at lower stimulation temperatures. Later experiments, with our more advanced V3-4.3 chip designs, did develop nano-pores for both H_2 and D_2 loading (Figures 31, 32).

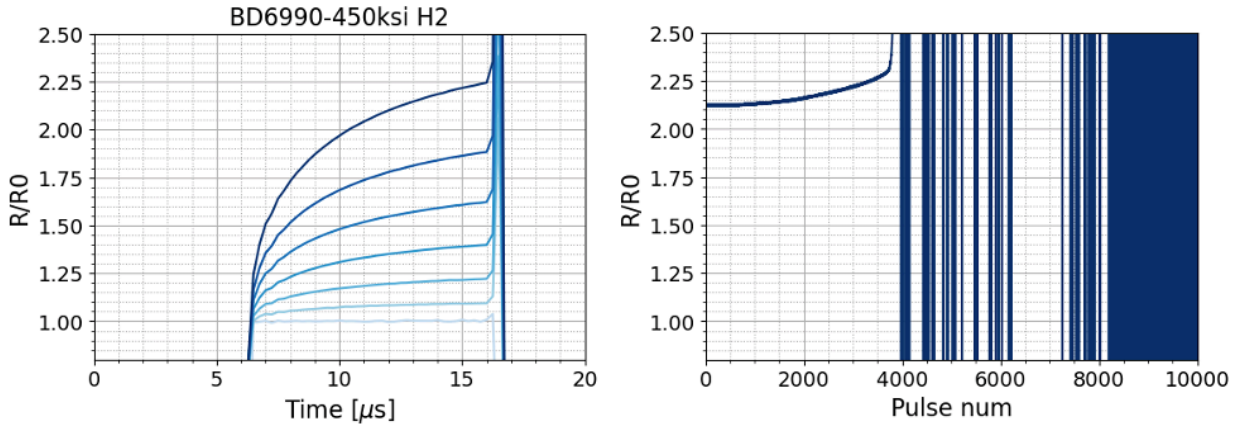


Figure 27. Callisto experiment BE3634. Calibration shots matched temperature closely to BE3623 but device failure occurred during stimulation.

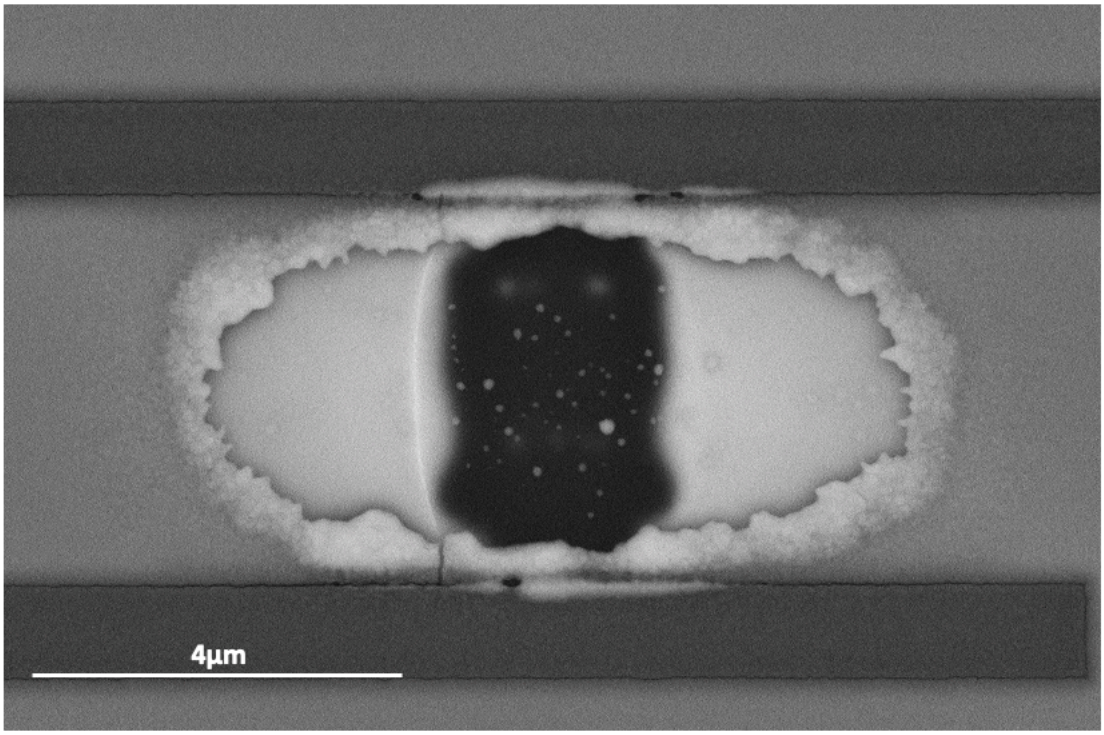


Figure 28. Callisto experiment BE3634. SEM of platinum heater/RTD following failure.

These results made it clear that this platform was too limited in palladium area, palladium visibility, and temperature capabilities. New versions of the films on heaters devices were designed and fabricated utilizing the nanofilm arrays described in the Staring and Physical Archaeology

sections of this report. The Atlas SEM mapping of large areas and computer vision could then be used to evaluate samples for the formation of holes following experiments over much larger areas of palladium. In order to heat an area containing a 3x3 grid of pads, each comprising a 64 nanofilm array, significantly more energy is required compared to the small heated regions in the nano-TI and Pillars on Pads chips. In order to rapidly heat the large area via joule heating, high current densities are required. This necessitates a much higher ratio between resistances of the resistive heater and the current source leads which can be accomplished by material selection and geometry. In order to have a smaller and more predictable thermal gradient within the heated region, a square 'heat pad' was chosen for the resistive heater geometry with an area of 1mm x 1mm which can easily accommodate the desired nanofilm arrays. NiCr was chosen as the heat pad material for its excellent thermal stability and high resistance. Extensive testing and many iterations of this device yielded a device capable of heating full nanofilm array samples to temperatures exceeding that of previous devices. Several amperes were needed with pulse widths of 150-250µs. To prevent significant heat accumulation, pulses were spaced apart by 50ms.

Devices, shown in Figure 29, were fabricated using photolithography and evaporation for all layers except for the silicon nitride passivation which was deposited via a chemical vapor deposition process and etched back using reactive ion etching. The device stack was 600µm sapphire wafer substrate, 100nm nickel/chromium 80/20 (wt%) heat pad, 250nm gold leads, 150nm silicon nitride passivation, 100nm platinum RTDs, 200nm gold leads, 25nm palladium nanofilm arrays.

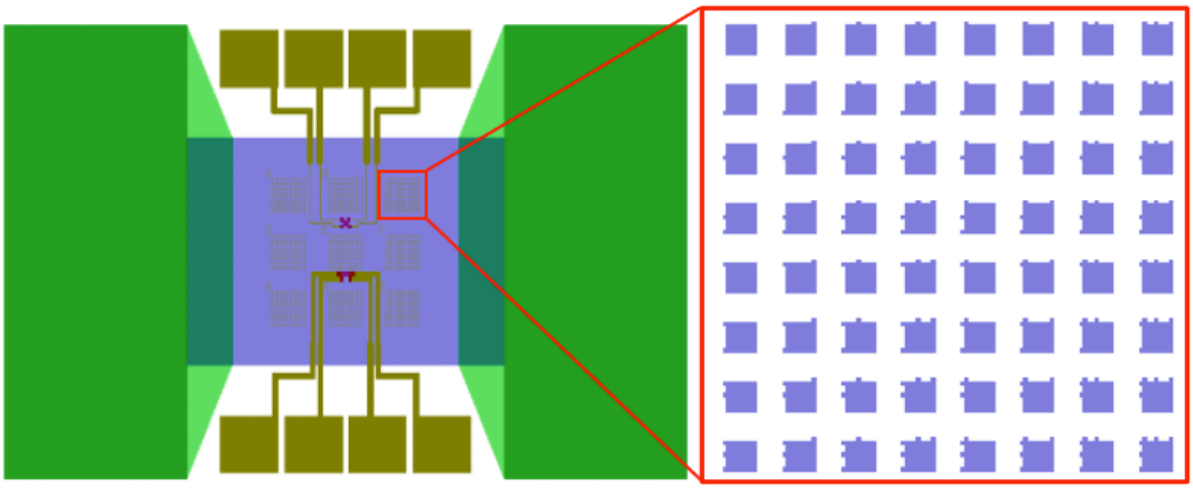


Figure 29. Designs for FoP_V3-4 device showing the whole chip and the micro pad arrays containing 10x10µm palladium pads with binary labeling for ease of detection when analyzing SEM images.

These devices were tested in the 150 ksi system, again, increasing stimulation temperature with each test until changes to the palladium were detected, data for one test is shown in Figure 30

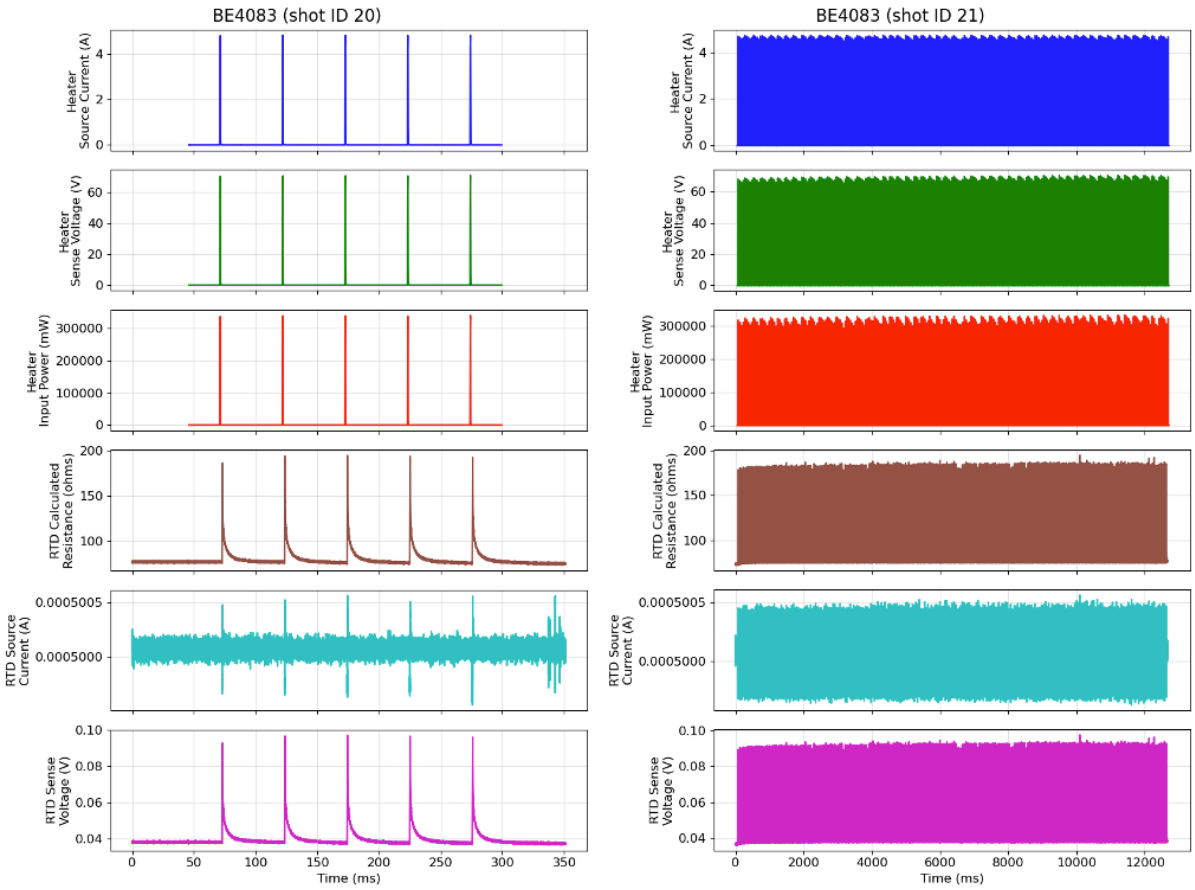


Figure 30. Callisto experiment BE4083. Chip loaded under 145ksi deuterium gas. The source meter and RTD data for calibration shots are shown in the plots on the left for 5x pulses of 5A amplitude (triggered), 250µs width, and 50ms period. The plots on the right show source meter and RTD data for a stimulation pulse train of 250 out of 500 total pulses.

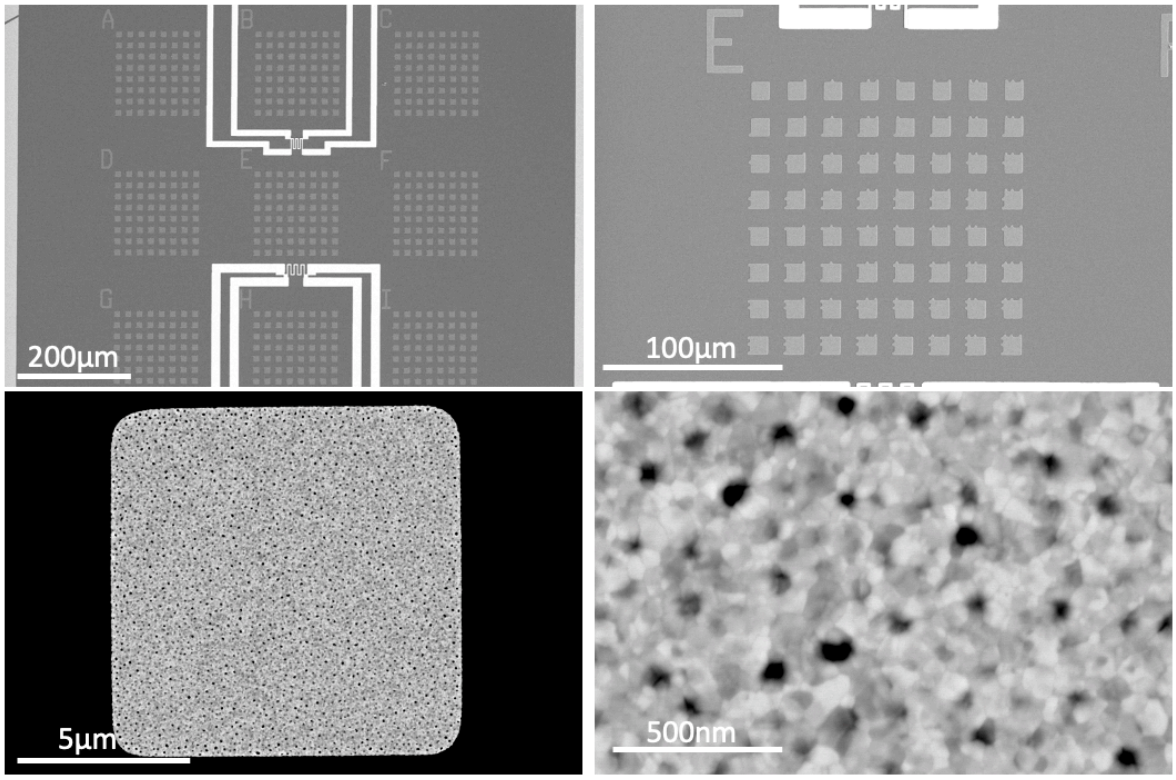


Figure 31. SEM post analysis of Callisto experiment BE4083. Nano pores were observed throughout all palladium regions on the device.

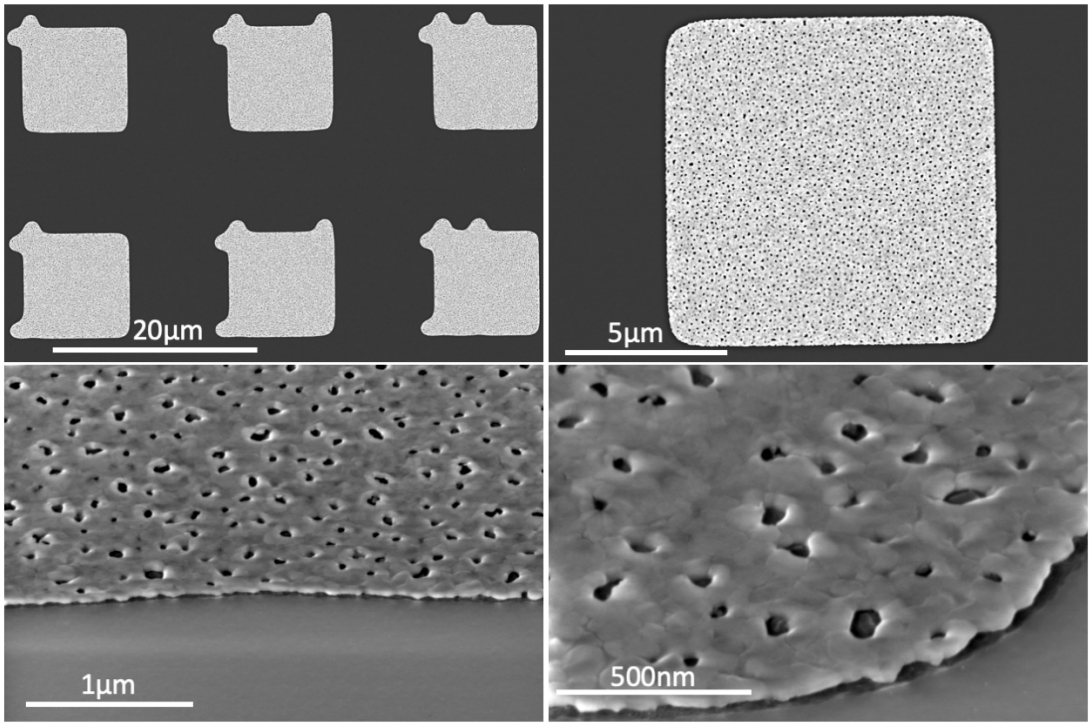


Figure 32. SEM post analysis of Callisto experiment BE4103. Device was loaded under 145ksi hydrogen and stimulated using conditions matching BE4083.

The presence of the nanopores shown above in Figure 31, raised the same question as discussed for the previous experimental design; were they due to near-melting stimulation, or due to AHE? In this case, we were able to repeat the experiment, using protium gas rather than deuterium. With the nano pore effect replicated in protium, it was ruled out as an indicator of AHE (Figure 32). The temperature limit where this effect began to occur was set as the thermal damage threshold for the device. Extensive testing was then carried out below the damage threshold but no other anomalies were detected.

Table 4. Heat-Pad experiments

Experiment	Chip version	Gas	Pressure (ksi)	RTD R/R_0	Thermal cycles
BE3542-B	V1	D_2	145	1.52	10,000
BE3542-E	V1	D_2	145	1.92	10,000
BE3563-E	V1	D_2	145	2.25	10,000
BE3563-B	V1	D_2	145	2.34	10,000
BE3612	V1	D_2	470	1.94	10,000
BE3615	V1	H_2	430	2.01	10,000
BE3634	V1	H_2	450	2.25	10,000
BE3623	V1	D_2	470	2.26	10,000
BE3664	V1	H_2	450	2.11	10,000
BE3675	V1	D_2	480	2.15	10,000
BE3898	V3-4.3	D_2	145	2.34	1,000
BE3926	V3-4.3	D_2	145	2.43	1,000
BE3972	V3-4.3	D_2	145	2.42	500
BE4083	V3-4.3	D_2	145	2.74	500
BE4095	V3-4.5	H_2	145	2.78	500
BE4103	V3-4.5	H_2	145	2.74	500
BE4141	V3-4.5	D_2	145	2.84	500

Chips : Electrical Stimulation

Thermal stimulation is not a new type of AHE stimulation mechanism -our mainline campaigns generally focus on thermal and optical stimulation –but electrical stimulation is relatively unexplored as a stimulation mechanism. Here, the emphasis is the actual electric fields and/or the flow of electrons, not simply resistive heating or electrolysis. This project campaign aimed for both forms of on-chip electrical stimulation, and we approached this with a variety of structures. One of our initial aims was to target extremely energetic electrons and ionization of deuterium at the palladium surface; at extremely high pressures and high electric fields, we expected dielectric breakdown to ionize the gas. Another initial aim was to use neighboring pillars as electrodes and attract deuterium atoms to induce AHE on one of the pillars as ‘Lightning Rods’. A third aim was

to use high frequency stimulation near the breakdown strength of the dielectric to induce AHE; the structure was essentially a metal-insulator-metal capacitor in a 'Capacitive Estim chip', operating at high field strengths at high frequencies, and the top thin palladium film acted as single event detectors.

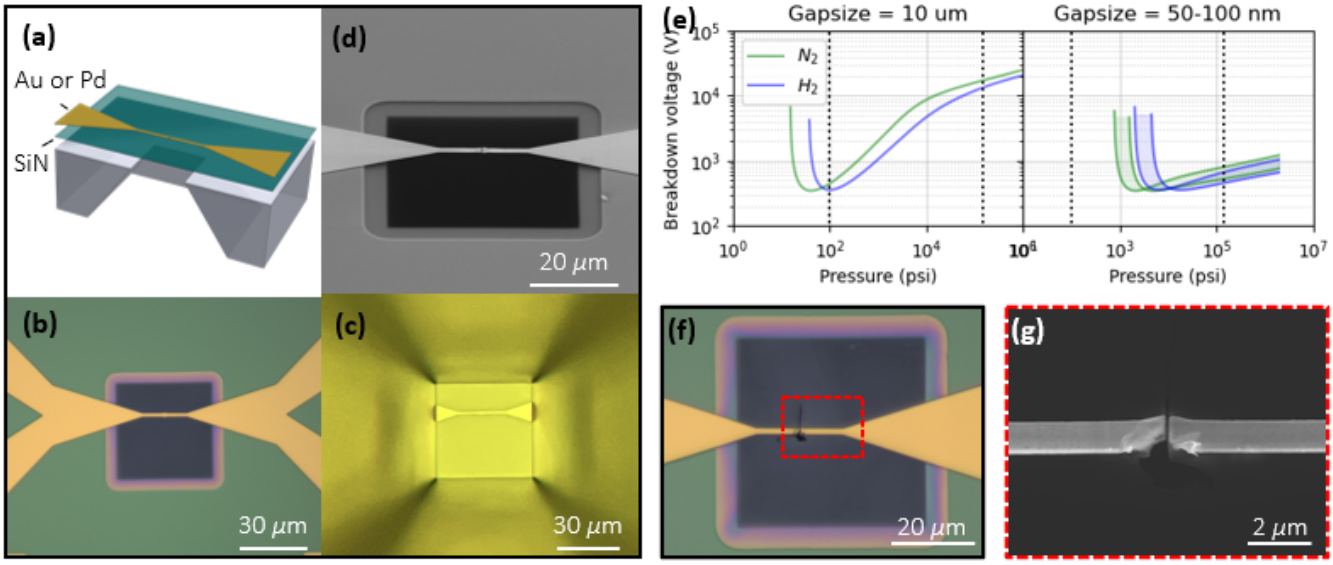


Figure 33. Suspended electrical stimulation chip. The chip design (a) consists of the gold or Pd electrode patterned on a SiN membrane suspended over an isotropically etched hole in the thermally insulated Si substrate. Optical images show the top-down morphology (b) as well as the morphology from the back of the device (c). The tilted SEM image (d) shows the suspension of the electrode over the hole in the Si substrate. The modified Paschen curve (e) provided guidance for experimental testing. We demonstrated breakdown across this device, as shown in the optical image (f) and SEM image (g).

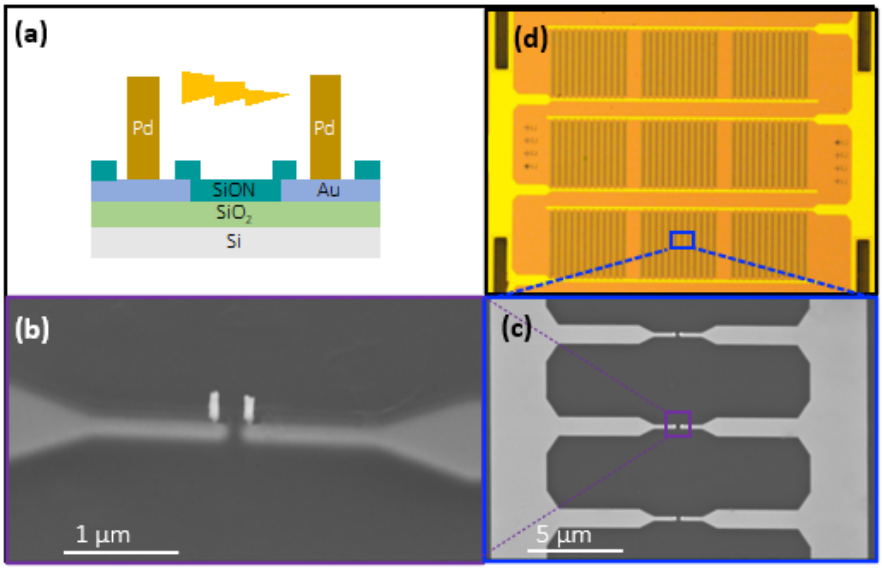
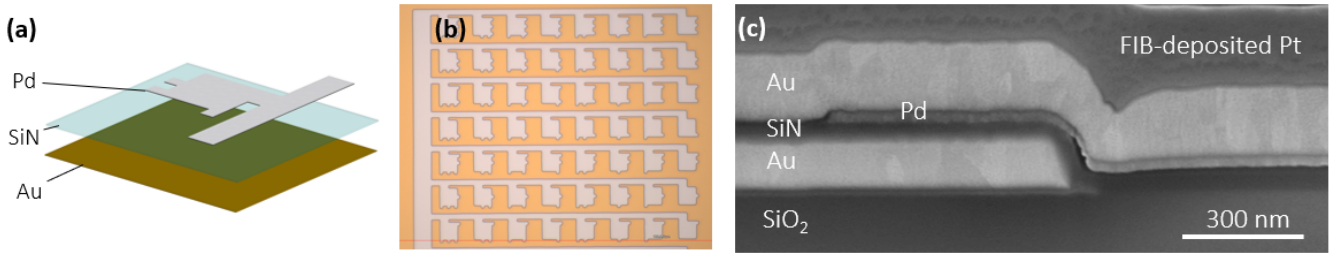


Figure 34. Lightning Rod chip. This chip design (a) consisted of electroplated Pd nanopillars perched on electrodes defined by e-beam lithography. SEM images (b-c) confirmed the expected morphology. We varied the distances between nanopillars from 60 nm to 150 nm, and defined an array of nanopillars shown in the optical image (d) to parallelize testing and maximize the likelihood of an event.

Stimulation at Pd(s) -D_x(g) interface: In the suspended chip design (Fig 33), we aimed to induce ionization of D_2 at the PdD_x surface (as a means of stimulation) by aiming for high electric fields between neighboring PdD_x electrodes. Likewise, the lightning rod chip design (Fig 34) also aimed to induce ionization of D_2 at the nanopillar PdD_x surface by aiming for high electric fields between neighboring Pd nanopillars. In the lightning rod design, we electroplate nanopillars on top of an electrode array defined by e-beam lithography and found that the design had substrate breakdown issues. In the suspended design, we developed MEMS-based fabrication processes to pattern electrodes on a floating membrane to circumvent substrate breakdown. The Paschen curve provided guidance for designing experiments for electrical breakdown at high pressures. We demonstrated breakdown across the device and tested in the presence of nuclear scintillation-based detectors, which showed no increase in counts relative to the background.

Fabrication consisted of starting with double-sided polished Si wafers, thermally oxidizing them with 1.5 μm SiO_2 to ensure electrical isolation and depositing 80 nm LPCVD SiN. We used bilayer photolithography to pattern 10 nm Cr adhesion and 100 nm Pd, deposited by e-beam evaporation. We then use bilayer photolithography with metal evaporation again to define the 10 nm Cr adhesion and 100 nm Au bondpads. We then used photolithography with backside alignment to pattern the SiN window on the backside of the wafer, followed by KOH and BOE etching to etch through the Si and SiO_2, defining the window.

Capacitive electrical stimulation: In this chip design, we aimed to use high fields either statically or at high frequencies to target near-surface electrons in the Pd. We fabricated devices in a metal-insulator-metal (MIM) structure with thin Pd on top to be sensitive to AHE events, as shown in Figure 35. We stimulated by applying bias to either the gold or the palladium and applying either a DC or AC-type pulse. We then used archaeology to look for anomalies. We iterated over a few fabrication batches to increase the breakdown field, reaching breakdown fields of ~6.7 MV/cm, and tested below and above breakdown. We tried fields below and above breakdown up to 7.7 MV/cm and frequencies up to 1 MHz for both polarities (i.e., electric field either into or out of the Pd) but did not observe isotopic differences.



(d) Test conditions	145ksi D_2 gas		150 ksi heavy water
Waveform:	DC-type	AC-type	DC
Results:	≤1 MHz Both polarities	≤1 MHz Both polarities	Current density ~ 250 mA/cm^2 Both polarities
	Non-isotopic holes observed	Non-isotopic holes observed	No obvious morphological anomalies

Figure 35. Capacitive electrical stimulation chip. The chip design (a) consists of a bottom gold electrode, a thin SiN passivation layer and dielectric, and the thin Pd device layer patterned into a pad array to make post-imaging analysis easier. The array of 64 pads is shown in the top-down optical image (b). The focused ion-beam scanning electron microscopy (FIB-SEM) image shows the cross-section at the border of the bottom gold electrode; at this interface, the Pd has to cross over the bottom gold electrode, necessitating the top gold to cover the Pd to avoid that edge as a breakdown point. This cross-section also confirms expected thicknesses of the layers, enabling calculation of the breakdown field.

Fabrication consisted of using substrates of thermally oxidized Si to electrically isolate the device. We used bilayer photolithography to pattern the evaporated metal electrodes of 5 nm titanium adhesion and 50 nm gold, which also act as wirebonding contact pads. This is followed by passivation with PECVD 50 nm silicon nitride, patterned, and etched to reveal the contact pads. We then used bilayer photolithography to pattern a thin 30 nm evaporated Pd device region, and then defined a thicker 5 nm adhesion and 150 nm gold region. Because we wanted the electric field to directly reach the Pd surface, there was no adhesion layer placed between the Pd and SiN. This thicker gold layer is necessary to ensure the step shown in Figure 35c does not limit the electric field at which breakdown occurs.

To measure breakdown voltages and stimulate at below or above the breakdown condition, we used an NI PXI system waveform generator PXIe-5413 to control a Trek Model2220 high-voltage amplifier. We tested under varying frequency conditions using both a DC-type pulse wherein we applied oscillated +/- 1 V around a given voltage, and a true AC-type pulse wherein we oscillated from 0V up to the given voltage. We also tested both polarities, wherein either the positive bias was applied to the Pd or the negative bias was applied to the Pd. The thin Pd film allowed us to use archaeology to look for holes, indicative of AHE. Capacitive electrical stimulation above breakdown voltages results in hole morphology (Fig 36). However, these holes were also apparent in hydrogen and were indicative of breakdown effects. Capacitive electrical stimulation below breakdown revealed no isotopic differences in the comparison between the pre and post-images for D_2 experiments which were not also apparent in H_2 gas (Fig 37a,b).

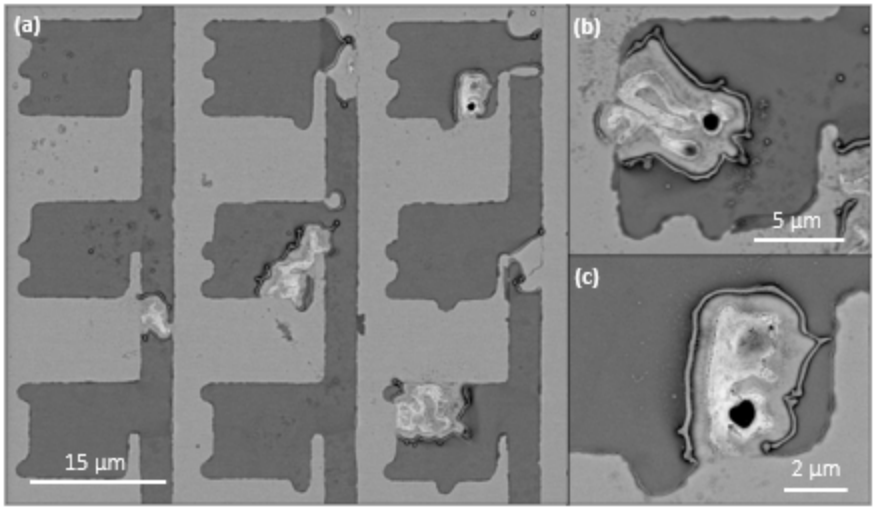


Figure 36. Stimulation above breakdown fields on the pads resulted in randomized hole morphology that persisted across both H_2 and D_2 gas and both field polarities. Images (a,b,c) involve different magnifications.

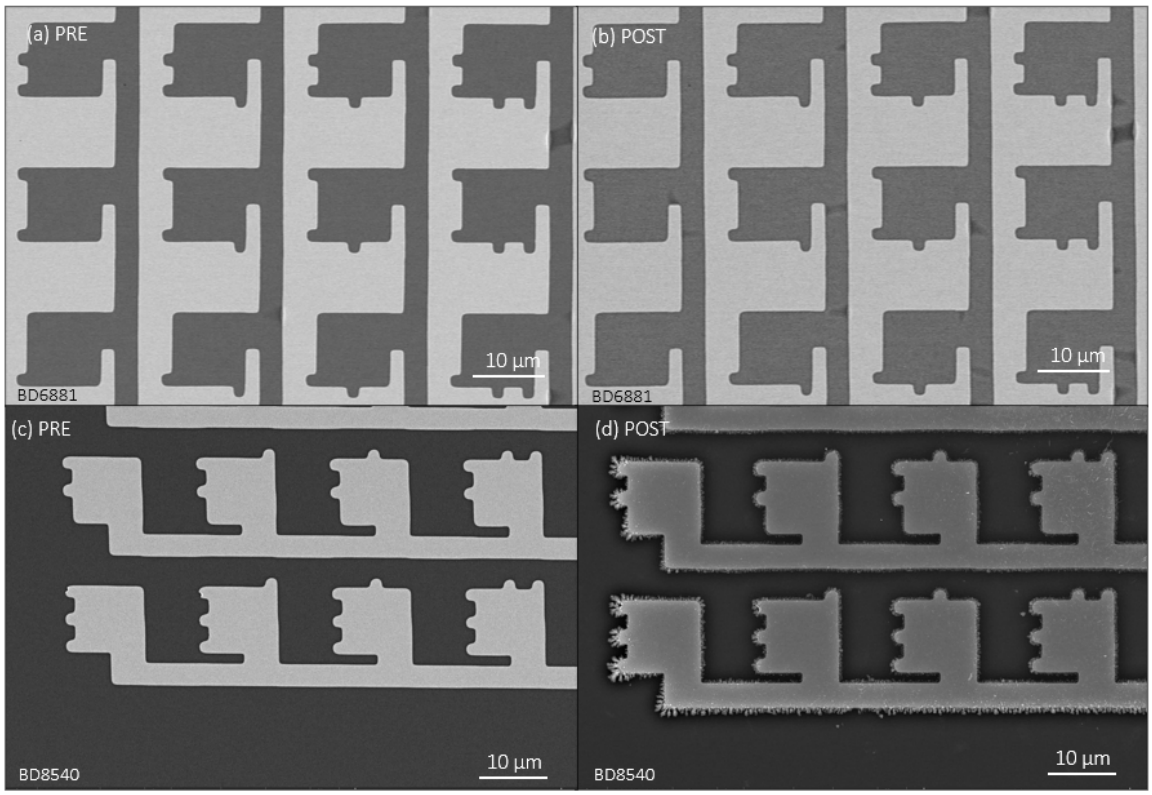


Figure 37. Stimulation below breakdown in conditions of 145 ksi D_2 gas pre (a) and post (b), and conditions of 150 ksi heavy water pre (c) and post (d). The gas testing showed no holes which would be indicative of anomalous events, though some delamination of the Pd film is evident, as no adhesion layer was used. The heavy water testing also shows no anomalous holes though some plating debris on the film edges is evident.

In heavy water (D_2O), we redesigned the samples to include a thin 5 nm Ti adhesion layer. Without the adhesion layer, the samples completely delaminate in the Pd region, likely as a result of a combination of the high pressure and the presence of liquid. We tested 4 devices at 150 ksi in

heavy water at current densities near 250 mA/cm2 (Fig 37c,d) in order to employ nano-film detection in high pressure electrolysis experiments . None of the samples exhibited indications of delamination though the samples do exhibit some minor electroplating.

Ultra-Short Pulse

Executive Summary

Here, we investigate the use of femtosecond laser pulses, or ultrashort pulses (USPs), as a means of stimulating and detecting potential AHE events. The primary difference for this stimulation method is that it is *non-thermal* in nature: excitation occurs at timescales faster than atomic motion. To rigorously test the stimulation conditions, we vary the wavelength from UV through THz, spanning various material processes ranging from deep electron excitation to direct phonon interaction and vary the pulse intensity from well-below damage threshold to intensities in which high-energy particles are observed to be emitted from the material system. These stimulation methods are aided by the use of Transient Reflectivity (TR) as a means of femtosecond-resolution probing of temperature and thus potential AHE events, as well as Archaeology-type samples to allow for single event sensitivity. While PdD_x and PdH_x demonstrate interesting behavior relative to their unloaded counterpart, these studies provide clear evidence that AHE is not occurring due to USP stimulation.

Ultra-short pulses: Stimulation and detection of AHE

A plethora of hypotheses on the "mechanism" for inducing nuclear reactions within a solid-state material exist within LENR literature. However, these can nearly all be summarized within two broad categories: Electron-driven events or lattice-driven events. In the former case, one example theory is that the electrons in the material can "screen" the native repulsion that occurs between two deuterons (notably, this is a well-understood process in both solid-state physics as well as high-energy physics), with the general idea being that two positively charged deuterons will be brought together through their equal attraction to a centering electron and thus reach a proximity in which fusion occurs. In the latter case for lattice-driven reactions, one may hypothesize examples such as the potential for deuterons to simultaneously "hop" into an identical position within a lattice. As this lattice position is stable and is a strong local minimum within the energy landscape, the deuterons may be capable of reaching proximities for fusion to occur. However, direct manipulation of either parameter is a non-trivial task.

Few methods exist of tuning electronic states or the local lattice environment without corresponding changes to the host material. Generally, meaningful changes to electronic states require dopants, which, in turn, also affects the surrounding lattice. Similarly, to induce something akin to the hopping-rate or diffusivity of atoms within a material, one must distort the lattice itself. In other words, it is generally difficult to modify the conditions surrounding a deuteron within a lattice without significant lattice distortion in itself. This nuanced complication can mean that our

assumptions of the *initial* lattice state, and hopping within that initial state, are incorrect: During hopping, the "lattice state" may be modified, leading to difficult interpretations.

One notable method that can achieve such changes to either the electronic or lattice-based states is via optical stimulation. In fact, this is the *primary* method of inducing changes to either state independently of the other: Electrons can shift energy levels simply by absorption and emission of light, and the hopping rate is dictated by vibrational states (i.e., phonons) which, at least in some cases, can be IR-active and thus "stimulated" by light. However, *meaningful* changes will not be expected for low light intensities; thermal radiation persists at nearly all infrared wavelengths continuously at room temperature and yet we have not observed lattices undergoing constant fusion (in fact, the power emitted by thermal radiation from a human body at room temperature [~1,000 Watts] is significantly higher than most reports of anomalous heat!). A similar argument can be made for electronic states: Both inter- and intraband transitions can occur continuously in a metal that is irradiated by even standard room lights, with both transitions having non-zero probabilities (in most materials) at room temperature simply due to Fermi-Dirac statistics. In other words, to expect *meaningful* modification of the states of interest, one likely requires a source of *intense* light.

Ultrashort pulses (USPs) of light can overcome these limitations and provide a unique method of both stimulation and detection. From a stimulation perspective, one can envision multiple scenarios that may lead to induced nuclear reactions. For example, USPs can deliver significant amounts of energy on timescales faster than atomic motion can occur (GW power densities in a single pulse). For context, the femtosecond lasers at Callisto Research can deliver mJ energy levels per-pulse in under 50 femtoseconds: providing a lower limit on the *instantaneous* power of 1 mJ/50 femtoseconds ~ 20 GW. This power is equivalent to *twenty* nuclear power plants all delivered within a single pulse. In other words, a single laser pulse can deliver as much power as the world's largest power plant but all delivered into a single material, thus providing access to highly energetic and strongly non-equilibrium lattice conditions.

From a detection perspective, these USPs have multiple advantages. First, optical probing has the benefit of being a "stand-off" technique – difficult micro- and nanofabrication methods to integrate sensors to the material of choice can be avoided. This benefit provides a means to interrogate the state of materials in extreme environments without concern of the sensor being modified; as we'll discuss later, it turns out that ultra-pressure environments lead to some difficulties for optical probing which are ill-discussed in literature. Second, the ultra-short timescales of these pulses dictate the measurement's temporal resolution. When monitoring behavior of a material with continuous-wave lasers, the temporal resolution is generally dictated by the measuring electronics. For example, if the absorbance or reflectance of a material changes due to temperature, the observed change in this behavior is limited by the rate of which the photodetector and oscilloscope can detect such changes. In contrast, for USPs, the detectors/electronics observe an impulse, delta-function-like response – the amplitude of this pulse is directly proportional to the reflectance. In such a case, this single value of reflectance has a temporal resolution equivalent to the *laser* pulse width: Any "smearing" of this delta function is known to be the result of the electronics. Generally, for the works in this section of the document, our lasers have pulse durations less than 100

femtoseconds. This provides the ability to observe events with a time resolution *faster* than typical fusion rates (~1 ps).

Section overview

Our investigations into the use of these USPs for the stimulation and detection of AHE events are described in the sections below. These sections are as follows:

- USP1. Stimulation methods
- USP1a. Single pulse stimulation
- USP1b. Pulse train stimulation
- USP2. Detection methods
- USP2a. Transient Reflectivity
- USP2b. Archaeology

USP1: Stimulation Methods

As discussed above, USPs provide a unique method of stimulating materials: These pulses can deliver massive amounts of energy on timescales *faster* than atomic motion and electronic equilibration. Thus, these pulses provide one of the only methods to induce **non-thermal stimulation** of materials. Generally, Callisto Research's efforts investigated various flavors of *thermal* stimulation; different ways of delivering heat and monitoring for anomalous heat production via nuclear events. In our USP experiments, however, the pulses' ability to deliver energy at timescales *faster* than thermalization of electrons and phonons (vibrations of the lattice/atoms) provides our only look into non-thermal stimulation.

In the next two sections, we will look at our work on UV, visible, and near-infrared wavelength laser pulses (250 – 1200 nm) for such non-thermal stimulation. These efforts all surround the notion of non-thermal processes via direct stimulation of *electrons* in the material. In a later section, Terahertz, we manipulate these USPs to produce intense terahertz pulses, thus providing a look into direct stimulation of *phonons* in the material.

USP1a: Single pulse stimulation

In most of our experiments, we look at the effect of a *single* laser pulse that has a duration of less than 100 femtoseconds, or an array of pulses that are sufficiently spaced in time that they are independent and do not lead to accumulative behavior. To understand the effect of these pulses on

a PdD_x sample and the corresponding "non-thermal" behavior, we can turn to the two-temperature model (TTM)[30].

The TTM describes the material response when subjected to USPs with relatively high accuracy and is widely used throughout literature to understand this behavior in metals. The premise of the TTM is that when a pulse of light interacts with a metal, the high electron density of the material screens the pulse; this screening is what causes high reflectivity of metals. This screening works by the electrons oscillating in-phase with the pulse, causing the pulse to reflect. However, the oscillatory motion of the electrons causes them to gain some amount of energy; we can describe this energy gained by an electron temperature. In other words, when a USP interacts with a metal, it causes the electrons to heat up. In contrast, during this interaction, the lattice remains cold as the pulse was "screened" by the electrons; thus, the atoms themselves have not gained any energy and the lattice temperature is unchanged. During this time, the material is in a non-equilibrium, non-thermal state (i.e., the material isn't hot, only the electrons are!). After a finite amount of time, the electrons start to lose energy via collisions with atoms, causing the lattice to heat up. This behavior is shown in Fig. 1 below; the base code for the TTM can be found in **BE1209** and is used frequently in later experiments.

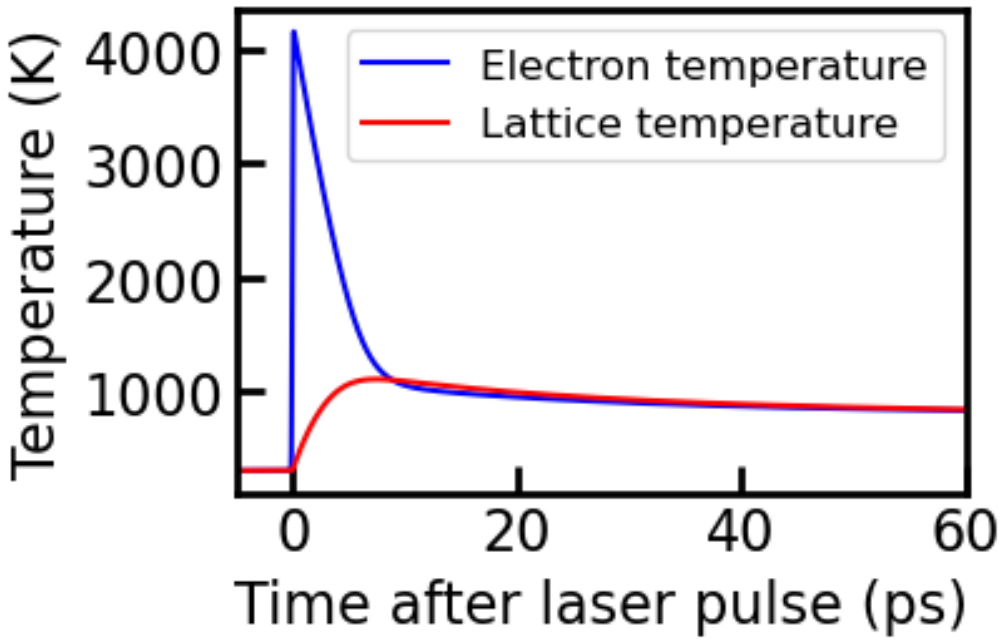


Figure 1: Electron and lattice temperatures following a 100-femtosecond laser pulse stimulating a Pd sample as calculated with the two-temperature model.

[30] Chen, J.K. *et al., International Journal of Heat and Mass Transfer* **49**, 307-316 (2006).

While this general TTM description works extremely well and is highly validated through experimental research in USPs, it is also extremely nuanced and only operates well in certain regimes or under certain assumptions. For example, the time that it takes for electrons to equilibrate with the lattice is proportional to the electron-phonon coupling factor, a generally intrinsic property of materials. However, knowledge of this material constant is determined by fitting the TTM to experimental data. As shown later, we determine this constant for PdD_x in our experiments. Another complication is that the *type* of USP that interacts with the sample can drastically change the behavior of this TTM process. For example, **by changing the wavelength of light, we intrinsically modify the way in which light interacts with the metal**.

Above, we stated that electrons screen the incident light pulse in metals. This is absolutely true, but only for certain wavelengths of light. At shorter wavelengths (higher photon energies), the electrons in metals can undergo *interband transitions*. This occurs by the wavelength of light causing an electron to jump from the valence band to the conduction band of the metal; this is the cause of non-silver colors of metals such as Au and Cu. In complex metals, this transition from oscillatory behavior to electronic transitions can be ill-defined, where many transition states exist through visible wavelengths[31]; this is the case for Pd and PdD_x. This complication, paired with the lack of theory surrounding AHE, suggests that *all* electronic transitions need to be considered.

While much of our work centers around either 800 nm or 400 nm stimulation/detection wavelengths, we also pursued these electronic transitions through the use of an OCPCA system (Optical Parametric Chirped Pulse Amplification). At a high-level, this is effectively a *tunable-wavelength* femtosecond laser system, capable of producing <100 femtosecond laser pulses (generally closer to ~35 fs), with tunability from 260 to 950 nm. This range allowed us to stimulate and investigate across nearly 5 eV of transition states. The relative energy levels and wavelength spectra of this system are outlined in **BE1235**.

USP1b. Pulse train stimulation

In our single pulse studies, the focus was on non-thermal stimulation via excitation of electrons. In parallel, we developed a method to investigate phonon-resonant stimulation using visible-wavelength USPs (again, later, we investigated the use of resonant THz laser pulses for an alternative approach to this hypothesis). The logic in this experimental thrust is as follows:

One hypothesis for AHE is that the deuterons within PdD_x oscillate "into" each other and undergo fusion. The exact details of this are, of course, vague due to a lack of sufficient mathematical theory. However, it is worth noting that atoms within a lattice are *always* oscillating; these collective oscillations are referred to as phonons. For any two-atom crystal basis (e.g., Pd + D), the material can support *optical* phonons, which are out-of-phase atomic vibrations. At a basic level, this arises from the mass mismatch between the two different atoms. These are referred to as optical phonons, as the frequency for these vibrations typically lies within optically accessible

[31] Lin, Z. *et al., Physical Review B* **77**, 075133 (2008).

frequency ranges; in many cases, these vibrations can induce direct absorption of light. For PdD_x, it has been established by both experimental neutron scattering and first-principal calculations that this phonon frequency lies around 10-14 THz. The hypothesis put forth toward AHE was specifically that *stimulation* of this phonon may cause sufficient oscillatory motion to cause deuterons to reach close enough proximities to undergo fusion.

Figure 2 below provides a visual contrast between acoustic phonons, which are the primary *heat* carriers in materials and are the modes that are stimulated in *thermal* stimulation methods, and these optical phonons investigated here, which are only weakly stimulated through thermal processes.

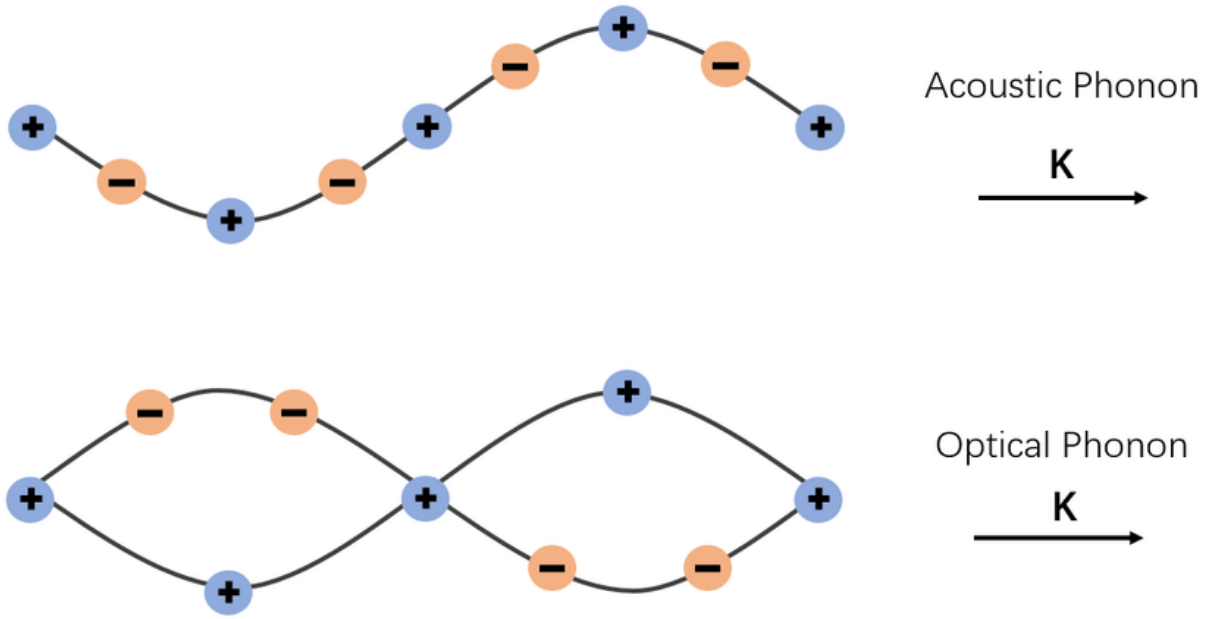


Figure 2: 1D visualization of (top) acoustic phonons and (bottom) optical phonons. One can envision the negative ions as Pd atoms and the positive ions as deuterons within the lattice. In both our "pulse train" and Terahertz stimulation efforts, the objective is to directly stimulate these out-of-phase optical modes in PdD_x, as acoustic phonons are sufficiently stimulated through thermal stimulation approaches.

To access this ~14 THz optical mode via *visible,* non-resonant USPs, we developed what we refer to as "pulse train" stimulation, where a train of ~35 fs optical pulses are separated by a time of $(14\ \text{THz})^{-1} \sim 70$ fs. This provides energy delivered to the lattice in-resonance with the motion of this optical mode, and resonant stimulation is a known method of increasing the amplitude of any oscillatory motion; does this increase in optical mode amplitude induce AHE?

These pulse trains are produced by leveraging the large frequency bandwidth of a USP. Due to the uncertainty principle, **the shorter a laser pulse, the more wavelengths/frequencies** (larger

bandwidth) it must contain[32]. In other words, when one refers to an 800 nm, 35 fs laser pulse, they really mean a ~770-830 nm laser pulse! The temporal shape of the USP is simply the Fourier transform of the sum of the frequencies composing the laser pulse; for an ideal ("Fourier transform limited") laser pulse, the sum of these different frequencies equates to a Gaussian time-envelope and is what we use for single pulse stimulation studies. **By manipulating the different wavelengths composing the pulse, we can manipulate the temporal shape of the laser pulse**.

This manipulation is done by modulating the *phase* of each frequency component. While this can be achieved experimentally through a number of approaches (most common example being phase-shift masks[33]), we induce phase modulation through spatial-light modulation (specifically, we use a 4f spatial light modulation geometry, shown in Fig. 3 below). To do this, we "split" the various wavelengths spatially using a diffraction grating; this is akin to sending light through a prism, where each wavelength exits at a different position. The wavelengths are then reflected from a liquid crystal array that is programmed with a specific spatial pattern; similar to phase-shift masks, the liquid crystal causes varying degrees of interference and thus imposes different phases to each wavelength component. After reflection from the liquid-crystal array, the wavelengths are recombined on the diffraction grating; the new pulse now has the imposed frequency modulation and thus temporal modulation.

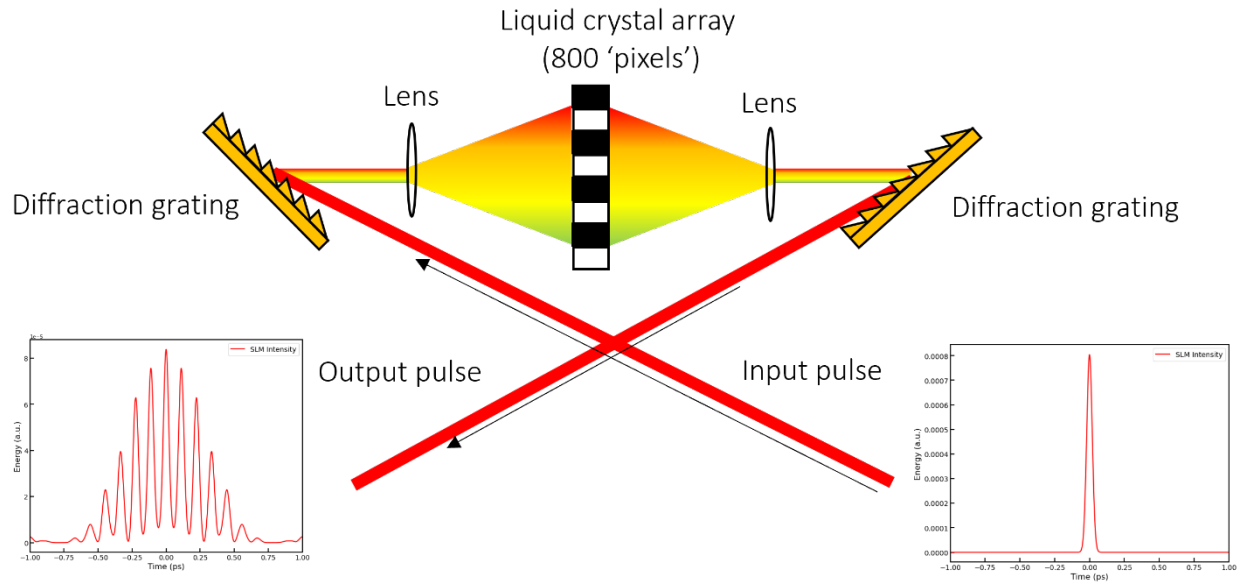


Figure 3: 4f spatial light modulation geometry (illustrated in a transmission layout) used to convert 35 fs laser pulses into "trains" of pulses at a given frequency.

Note, to generate "sharp" pulses in the time domain, the frequency domain requires sharp autocorrelations. This implies that we require *pseudo-random sequences of* phases in the frequency

[32] $\Delta E \Delta t > h$, where Δt is the pulse duration and $\Delta E = h\Delta f$, where f is the optical frequency; this optical frequency is inversely proportional to wavelength ($f = \frac{c}{\lambda}$).

[33] https://en.wikipedia.org/wiki/Phase-shift_mask

domain; non-random arrays will not produce sharp autocorrelations and thus not produce “clean” pulses in the time-domain. To do this, we introduced a maximum length sequence algorithm that generates the required “randomness” for phase modulation. These algorithms and the related development of this pulse train development are outlined in BE1231.

To observe the effects of these pulse trains on the material response, we implement our two-temperature model using a heat source equivalent to the calculated pulse array via our algorithm. An example of such is shown in Fig. 4 below, where the temperature can be shown to accumulate at all frequencies, but sharp transitions in temperature are observed at the high- and low-end of frequencies; these transitions correspond to different diffusion timescales within the material stack. These calculations allow us to separate *known* thermal accumulation effects from potential AHE effects. For example, one should certainly expect *more* damage at high pulse-train frequencies simply due to heat accumulation; this extent of additional damage can be pre-determined in the absence of AHE through such modeling approaches.

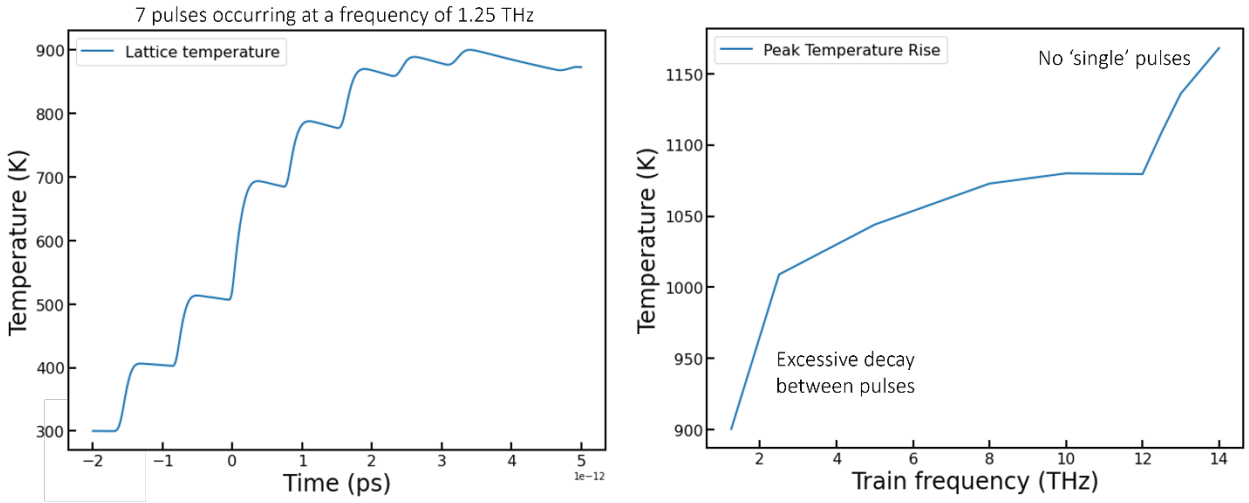


Figure 4: (Left) Example of TTM calculation using a 1.25 THz pulse train. (Right) The calculated peak temperature rise as a function of pulse-train frequency.

USP2: Detection methods

Due to the short timescale associated with USPs, most traditional thermometers do not have a time-response adequate for detection of the temperature rise associated with the short pulse (and thus not a sufficient temperature rise for AHE event detection!). To understand this aspect better, consider laser machining: Modern laser machining relies on ultrashort pulses for material removal because the pulse is capable of ablating material away faster than heat is deposited to the surrounding region. Thus, if a thermometer were placed nearby, it would observe negligible temperature rise; any anomalous heat events that occurred on a similar timescale would have

similar issues. Thus, choice in an adequate thermal sensor is critical to detecting our developed stimulation methods.

In our USP efforts, we rely on two primary methods for AHE detection: Transient reflectivity and Archaeology. In transient reflectivity (TR) we leverage a *second* USP that acts as an ultrafast thermometer, providing temporal resolution better than 100 fs. The second method, physical archaeology, primarily relies on observation of material damage due to a given laser pulse. In both cases, controls with protium are critical to separating laser-induced damage from AHE-induced damage. Importantly, the **ultrafast aspect of USPs eliminates the need for removing convective differences between the two isotopes of hydrogen**: The USP delivers and dissipates heat faster than convective timescales (~femtoseconds vs. microseconds). This greatly simplifies our analysis for both detection methods.

USP2a: Transient Reflectivity

Transient Reflectivity (TR) is an optical "pump-probe" technique that allows for thermal detection at ~mK temperature excursions without the need for micro- and nanofabrication methods or direct contact with the sample of interest. Generally, at least for metals such as Pd and PdD_x, the sample of interest can be directly stimulated/heated and directly measured using TR. This non-contact approach thus avoids the effects of thermal resistances associated with the thermometer itself or the thermometer/sample interface. To put this thermal resistance problem in perspective and highlight the benefits of TR, we can consider Fig. 5, which highlights a range of *interface* resistances between two dissimilar materials.

As shown, one can find that the thermal resistance associated with the interface alone can be equivalent to the addition of up to 100 nm of SiO_2 to the interface. Notably, when developing in-contact microscale thermometers, one will generally aim to minimize insulator thicknesses specifically to avoid this issue of mismatched temperatures between the thermometer and the sample of interest. In the case where one wants to interrogate the temperature of a metal such as Pd or PdD_x, these intermediate insulators are a requirement, and thus introduce not only their own intrinsic resistances, but *two* additional interfaces that impose a fundamental limit on the accuracy of the thermometry technique. Thus, the ability to eliminate these additional thermal resistances and *directly* interrogate the temperature of a sample of interest is a unique benefit to optical thermometry methods such as TR.

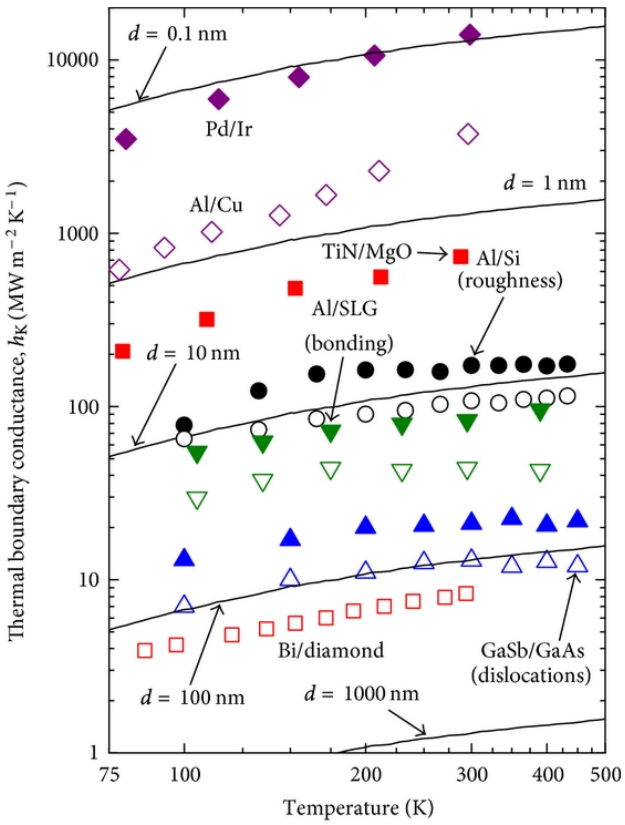


Figure 5: Thermal boundary conductance (~heat flow at material interfaces) for different material systems. The solid lines highlight the equivalent conductance if varying thicknesses of SiO_2 were added between the two materials, highlighting the ideation behind non-contact, optical thermometry.

TR, and other pump-probe methods, rely on using a laser pulse to stimulate/generate heat on a target surface while using a second laser pulse to measure the resultant temperature rise. In the case of our experiment, this is achieved by optically *splitting* a single laser pulse into two components; one portion of the beam is now the pump/stimulation pulse while the second beam acts as a probe pulse. To determine the time-evolution of this temperature rise induced by the stimulation pulse, we introduce a varying offset between the arrival time of the two pulses (i.e., the time-delay between the two). This is done using a retroreflecting mirror mounted to a motorized stage; as the mirror is mechanically moved to varying distances, the probe must travel the additional distance, d, before reaching the mirror, thus creating an additional time delay corresponding to $t_{delay} = \frac{2d}{c}$. The factor of two is produced from our experimental geometry in which the probe passes twice down the retroreflector. In our experiment, the pump pulse is frequency doubled to 400 nm while the probe pulse remains at the fundamental of 800 nm, allowing for spectral filtering of the two beams. A schematic representation of Transient Reflectivity is shown in Fig. 6 below.

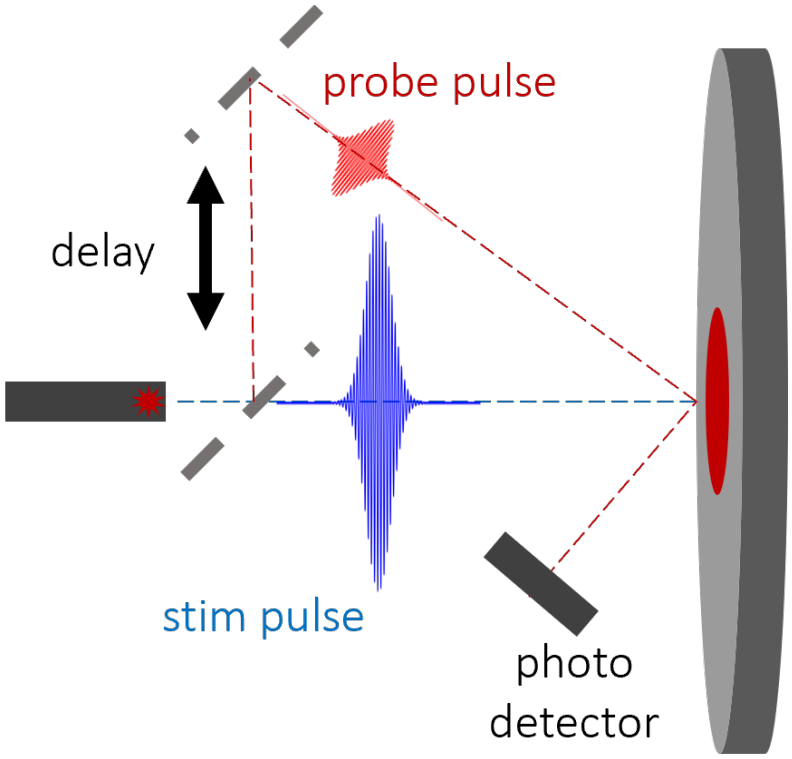


Figure 6: Schematic of Transient Reflectivity experiments. A single laser pulse is split into two components: A stimulation pulse and a probe pulse. The stimulation pulse creates a temperature rise on the sample surface while the probe pulse measures the resulting change in reflectivity, and thus temperature, of the sample surface. The temperature evolution in time is determined by repeating this process with varying path lengths of the probe pulse.

The objective in these TR experiments is, at a high-level, to look for AHE events through the observation of increased probe signal (i.e., temperature) following a stimulation pulse. This, in theory, may appear as an increased signal at all time-delays (which would be the case where AHE happens *during* the stimulation pulse and simply increases the average temperature/energy deposited) or a second event that occurs sometime after the stimulation pulse (i.e., if AHE happens 10 picoseconds following stimulation, then it will appear as though a second laser stimulated the sample).

During early investigations of AHE with Transient Reflectivity (USP2-1 through USP2-7 campaigns) several sample geometries and test conditions were investigated, including varying thicknesses of foils, cryogenic temperatures, loading conditions, and variations on the measurement hardware itself. The general objective in these experiments was to observe isotopic contrast – differences in the TR signal between PdH_x and PdD_x. However, multiple experimental concerns obfuscated the ability to reliably "look" for AHE events in these early investigations.

First, most of our early studies relied on the use of bulk sample geometries. However, materials with thicknesses exceeding a critical length scale tend to produce non-traditional TR signal due to

the non-equilibrium conditions associated with laser stimulation. A demonstration of this is shown below in Fig. 7, where we compare experimental TR measurements of a 100 nm vs. 20 nm PdD_x film under identical loading conditions (145 ksi D_2). As shown, when thicker films are used, there can be conditions in which the "hot electron peak" is not observed. This peak is due to non-equilibrium electrons that are initially excited by the pump pulse. After some time, generally a few picoseconds, these hot electrons lose energy to the lattice and the sample follows traditional heat conduction. However, in sufficiently thick films, these electrons can be ballistic and traverse into the depth of the thick layer, whereby they undergo thermalization away from the measurement surface. Similarly, the thermal decay of energy from the surface is dependent on the thickness of the film due to the aforementioned effects of thermal resistances at interfaces: In thick films of Pd, the material is highly thermally conductive and causes rapid decay of the induced temperature rise. However, in sufficiently thin films, not only are ballistic effects removed (they remain within the measurement volume), but heat conduction is limited by interfacial thermal resistances, causing the layer to stay hot for longer durations. This critical thickness is primarily dictated by the electron-phonon coupling factor of the metal.

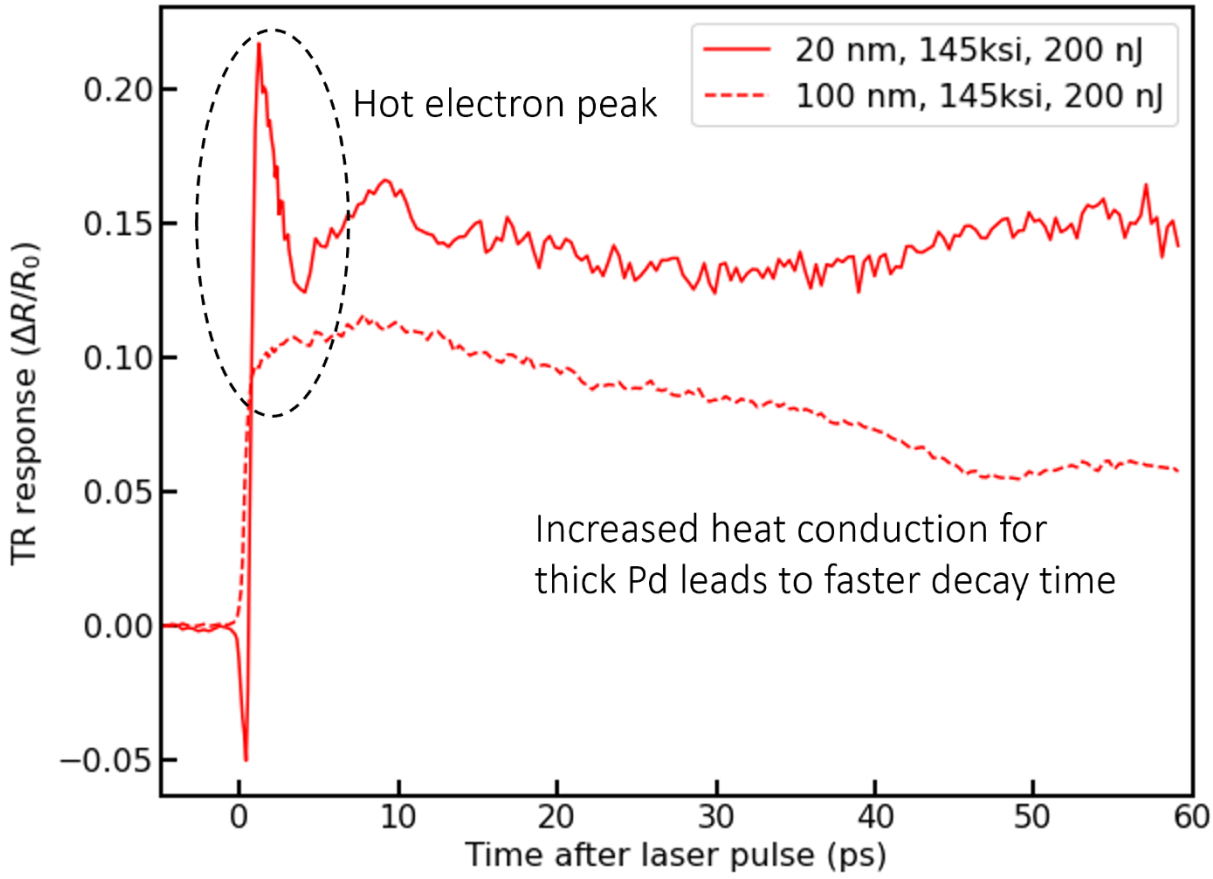


Figure 7: Comparison of experimental TR data for a 20 nm (solid line) and 100 nm (dashed line) PdD_x film at 145ksi pressure of D_2 gas. Contrast between the curves can be observed at nearly all time delays due to the non-equilibrium heat deposition from USP stimulation.

A second issue, as also highlighted in our Physical Archaeology section, is supercontinuum generation. At sufficiently high intensities, when USPs interact with dielectric media, the pulse can generate an array of non-linear optical processes. One of these, termed supercontinuum generation, leads to the production of white light. In other words, our 400 nm or 800 nm pulse can be spectrally broadened through the visible spectrum, taking a new spectral width that exceeds the filters. The issue with this process is that our TR experiment relies on spectral filtering of the two pulses; if one of them becomes broadband, the pulse will transmit through these spectral filters and produce anomalous signals in the photodetector. Second, this supercontinuum generation greatly distorts the spatial distribution of energy, causing local "hot spots" within the pulse, which can be interpreted as potential AHE events if not avoided. In our ultrapressure hardware, the sample is positioned behind a thick sapphire window – in our early experiments, the sample was simply too close to the window, thus requiring the USP to be highly focused when transmitting through the window. This leads to locally high intensities within the sapphire, creating supercontinuum, and leading to anomalous signals. This effect is shown below in Fig. 8.

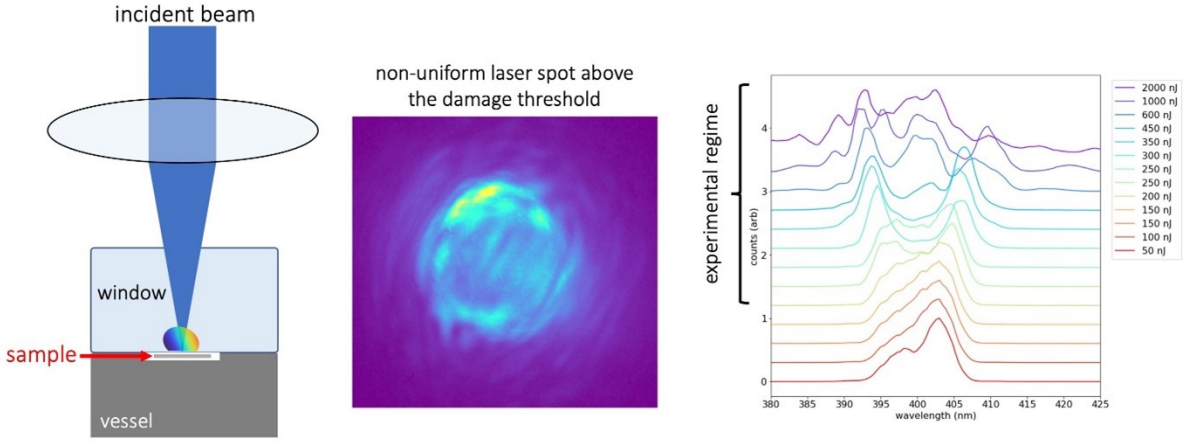


Figure 8: Supercontinuum generation from USP interactions with our high-pressure sapphire windows. (Left) Schematic of the high focusing conditions when the sample is placed too close to the window. (Middle) Image of an inhomogeneous laser spot due to this supercontinuum generation. (Right) shows the pump-pulse spectrum when supercontinuum occurs.

These two primary issues were resolved in experiments from USP2-8 onwards; more details on these types of artifacts that can produce "AHE-like" signals are discussed in USP-Appendix. Notably, upon resolving these issues, our TR experiments were highly repeatable and produced data consistent with two-temperature model calculations. Shortly thereafter, we observed highly reproducible isotopic contrast between PdD_x and PdH_x. An example of such contrast is shown in Fig. 9 below. As shown, one can observe that the data prior to a time-delay of zero (i.e., before the sample is stimulated), the two isotopes are identical. However, following stimulation with a pump pulse, PdD_x was observed to have *consistently* greater TR signals than its PdH_x counterpart.

Interestingly, this contrast was greater at intermediate loading conditions (<90%). As also shown in Fig. 9, increased pressures cause the two isotopes to converge, with only a small increase in TR signal for PdD_x.

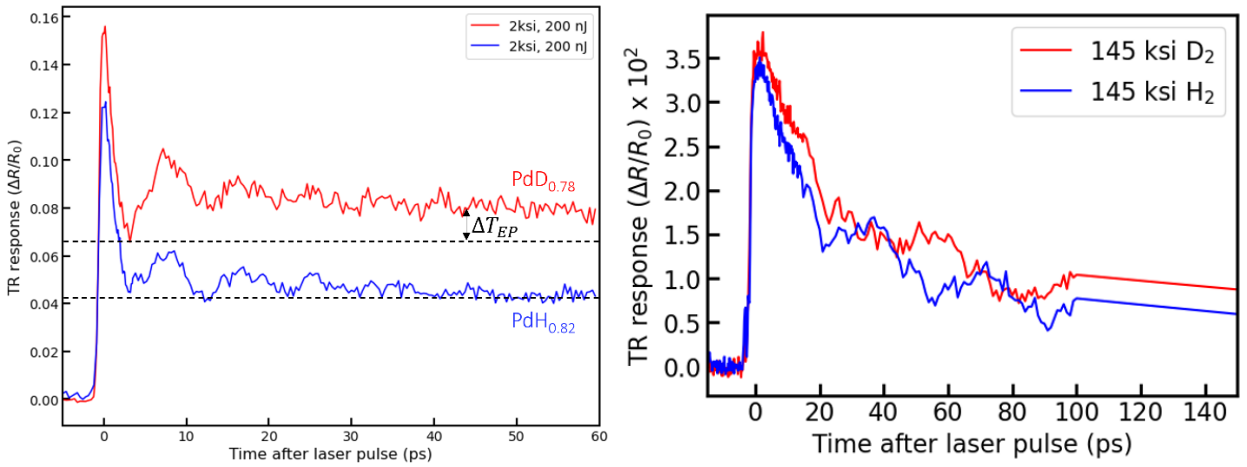


Figure 9: (Left) TR data for PdD_x (red) and PdH_x (blue) at pressures of 2 ksi and nominally 75% loading conditions; PdD_x was consistently observed to have a greater TR signal than its PdH_x counterpart. (Right) The same data acquired at 145ksi (~95% loading).

This isotopic contrast suggests one of two phenomena. First, the increased signal in PdD_x could suggest that the temperature is greater due to additional heat being generated within the sample; this could be the result of increased pump absorption in the metal or AHE events. The second possibility is that PdD_x produces a stronger change in reflectivity for a given increase in temperature. Different materials are known to produce varying thermoreflectance coefficients, and it is possible that PdD_x and PdH_x are simply "different materials" from a band structure perspective. This is certainly plausible given the known contrasts in electrical properties between the two isotopes.

Our first attempt at testing these hypotheses is shown in Fig. 10 below, where we performed Au controls in the two gas conditions. If similar behavior was observed here, it would indicate that the difference between deuterium and protium gas environments is responsible for the observed effect; notably, in the past, we had observed such isotopic differences due to collision-induced absorption in the mid-infrared. However, as shown in Fig. 10, we observed the *opposite* response, where H_2 gas produced a greater signal than D_2 gas. While certainly an isotopic difference, it suggests that the differences between PdD_x and PdH_x are *greater* than observed in our TR data!

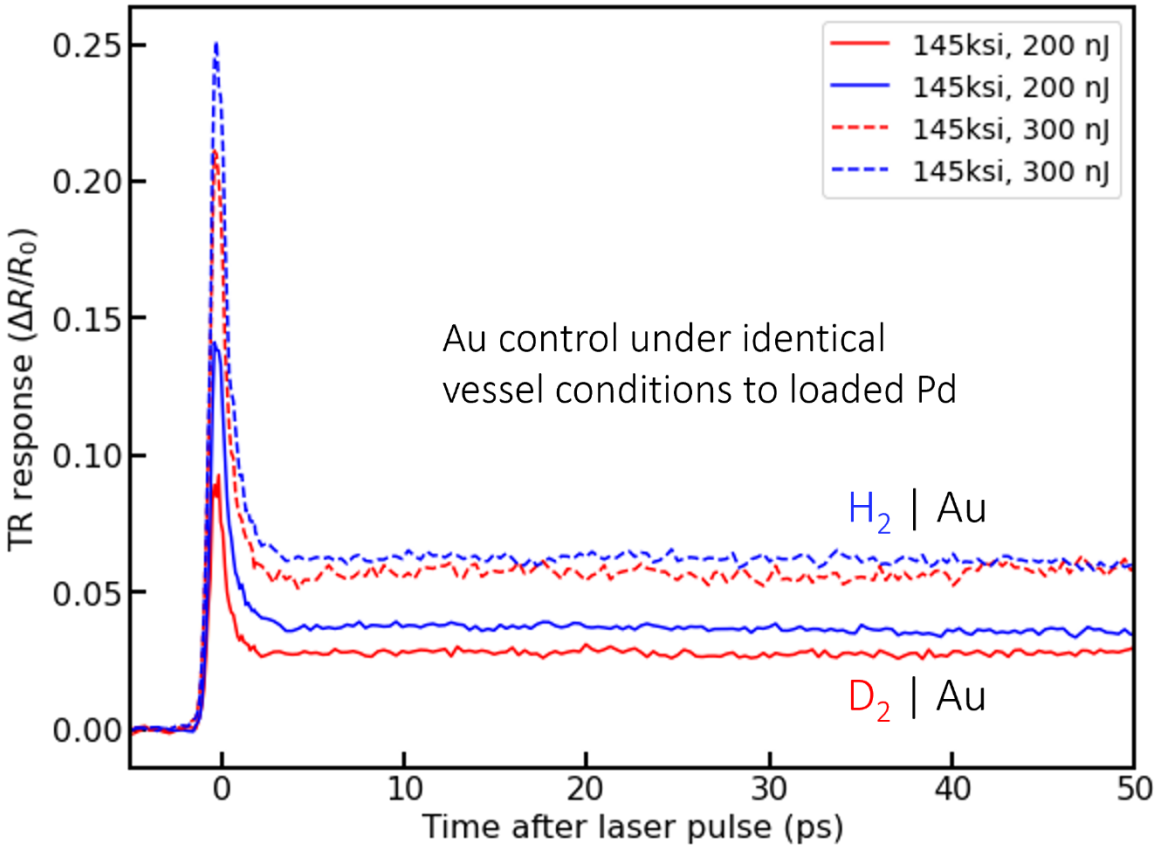


Figure 10: First attempt at resolving the observed isotopic contrast for PdDx vs. PdHx. In these experiments, we performed TR measurements on Au controls to isolate the role of optical differences between high pressure H_2 vs. high pressure D_2. However, we observe opposing behavior relative to our Pd experiments, where in the controls, H_2 is observed to produce a larger signal than D_2.

To further test these hypotheses, we rebuilt our TR experiment to allow for a new pump-probe geometry. Two major changes were made: First, we designed the setup to allow for measurements at much greater pump-probe time delays. This allows us to extend the datasets to timescales associated with heat diffusion and determine whether this observed isotopic contrast is simply the result of non-equilibrium behavior; it also allows us to look for potential AHE events at times much later than the pump stimulation. The second change was to allow us to pump and probe different sides of the material system. In doing so, by using an appropriate material stack, we can stimulate Pd, while measuring the temperature in an adjacent Au layer. Since we have determined the response of Au at elevated pressures and know Au does not produce AHE (Au doesn't load), this acts as an ideal optical thermometer. The concept of this is shown in Fig. 11 below.

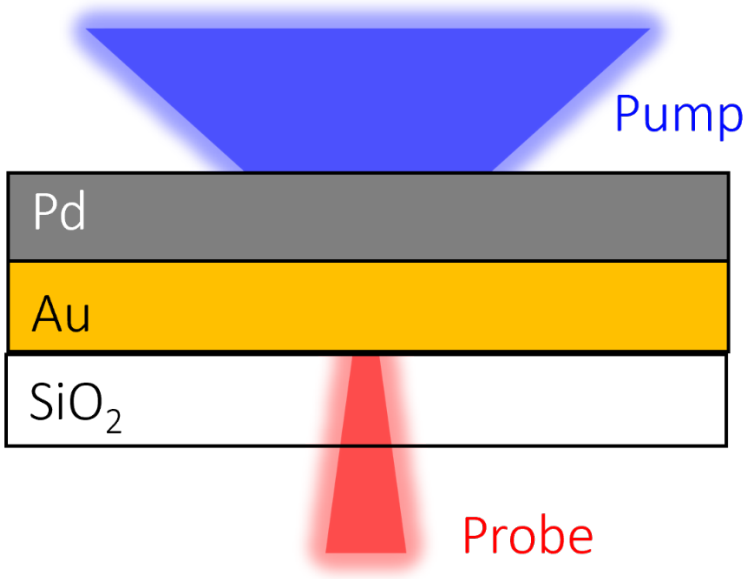


Figure 11: Schematic of revised TR setup to allow for stimulation of Pd while using Au as a constant optical thermometer.

While many experiments were conducted to verify the accuracy of this approach and eliminate any potential spurious signals not associated with AHE, the general conclusion is that the isotopic contrast *inverts* relative to pumping and probing the Pd alone. In other words, this experimental setup verifies that the observed isotopic contrast in which $PdD_x > PdH_x$ is solely due to differences in thermo-optical responses between the two material systems and is **not** associated with anomalous heat events. An example of these data is shown below in Fig. 12, where it can be observed that upon the probe interacting with Au following PdD_x/PdH_x stimulation, we no longer observe an increased signal for the deuterated metal. Two salient differences in these data from our prior TR experiments are due to the change in geometry. First, the slow rise in signal is because we are no longer stimulating the measured surface – heat must flow through Pd into Au before a measurable signal occurs; this slow rise is simply heat diffusion from the two layers. Second, the oscillations in the signal are primarily the result of Brillouin scattering within the SiO2 layer. When a rapid thermal excursion occurs, the thermally expanding metal launches an acoustic wave into the materials below; in this case, the Pd launches a wave into the Au/SiO2. This acoustic wave creates a local density gradient that is optically detectable on the probe beam; when the acoustic wave is distanced from the Au surface at integer multiples of the wavelength of light, interference will occur, causing local minima and maxima in the optical signal. In other words, the period of oscillation is found to match classical Brillouin scattering through $f_{osc} = \frac{2nv}{\lambda_{probe}}$, where n is the refractive index of the material (SiO_2), v is the sound velocity of the material, and λ_{probe} is the wavelength of the probe beam.

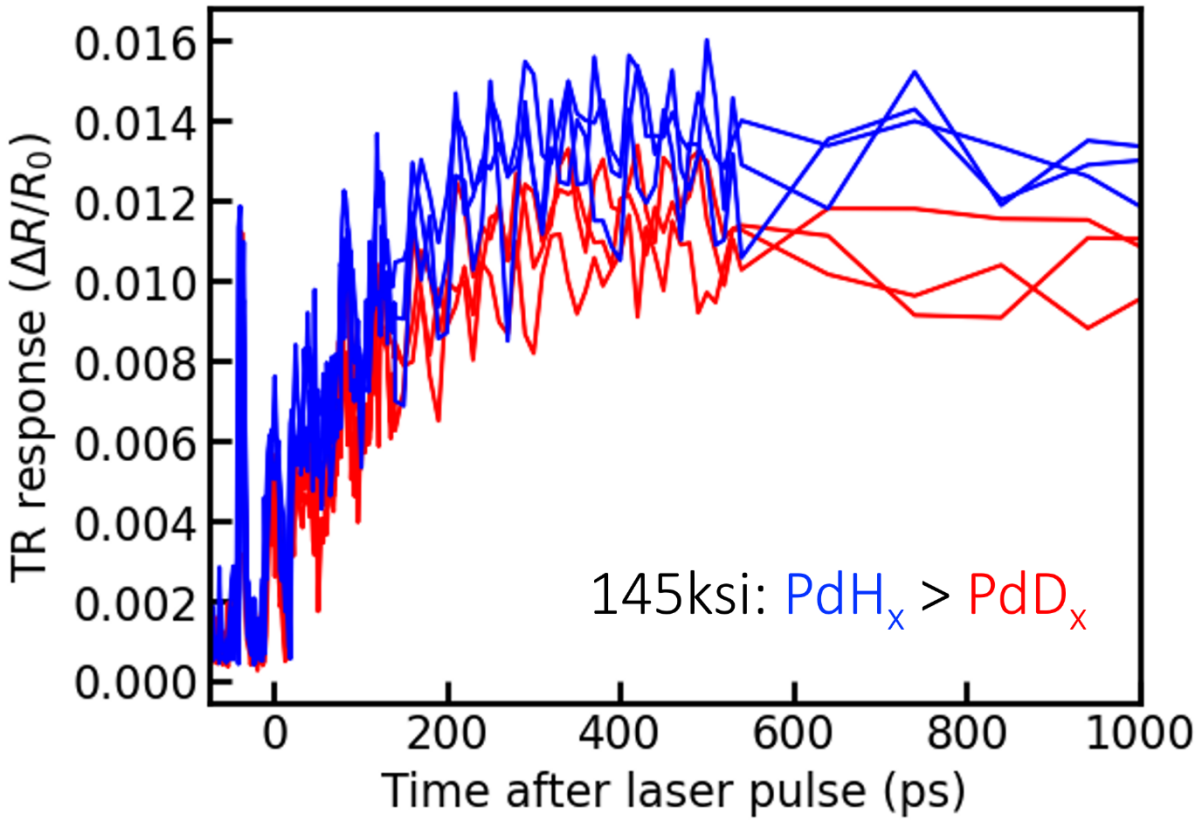


Figure 12: Experimental data from our revised TR setup in which the pump beam stimulates the Pd surface while the probe measures the thermal response of an Au film in direct contact with the Pd. The observed isotopic contrast is now consistent with control experiments on Au, in which the signal in PdH_x exceeds that of PdD_x. This experiment thus verifies that the original isotopic contrast was solely the result of differences in the material response functions.

Not only does this revised experiment verify a lack of AHE in PdD_x via laser stimulation, but it also suggests that optical probing is not a reliable method for AHE investigations when isotopic contrast is the sole method of potential AHE verification – the material response functions between PdH_x and PdD_x are simply too different for valid interpretation. Further, these experiments are repeated on rare-earth elements, such as Erbium, that allow for loadings near 300%. In no cases are AHE events observed.

USP2b. Archaeology

Our second method of looking for potential AHE events is through Archaeology. Opposed to TR, this method provides no direct *in situ* measurement of temperature rises. Rather, Archaeology detection relies on *ex situ*, *postmortem* investigations of material morphologies following

excitation. In other words, we compare the relative damage profiles between PdH_x and PdD_x to determine whether AHE occurred. The logic behind this is that the laser may induce some degree of damage to PdH_x, but if AHE events occurred, additional damage will have arisen due to the increased temperature excursion. In some cases, the laser is kept below the damage threshold, and the material is designed in such a way that the temperature rise associated with AHE should be sufficient to damage the material without any laser input. This structure size can be determined by a quick calculation on the temperature rise associated with a single AHE event; the specific heat of Pd (C_s) is 0.24 J g-1 K-1, the density (ρ) is 12.02 g cm-3, and we anticipate AHE to produce 24 MeV of energy internally. Assuming the material melts at the Pd melting temperature of 1830 K (the melting point is suppressed when loaded, so this *over*estimates the size threshold; sizes below this will melt once loaded if AHE occurs) the nanostructure volume to melt from a single AHE event is given by $V_{nanostructure} = \frac{24\ \text{MeV}}{1830\ \text{K} * C_s \rho} = 7.3 * 10^{-16}\ \text{cm}^3$. For a ~100 nm diameter pillar, this puts an upper length of 900 nm. In other words, for structures smaller than this critical volume, we can anticipate a single AHE event causing material melting to occur. Similarly, in the case where we use 20 nm Pd films, we can expect a single AHE event to induce a hole of at least 100 nm in the film, readily observable with electron microscopy methods.

At a first pass, one can use thin-film archaeology as a means of investigating potential AHE events. Generally, observing the *onset* energy of damage is the parameter of interest: If damage to a film is observed at lower intensities in PdD_x than PdH_x, it would be a potential indication that additional anomalous heat may have been generated within the deuteride. However, *discrete* measurements of this threshold are difficult and time-consuming: To accurately determine the potential addition of 24 MeV, one must vary the laser energy in increments of 24 MeV and find the onset of damage (i.e., a *massive* number of data points must be obtained). Further, the onset of damage is typically not observable: Damage takes the form of local melting and resolidification at low laser energies, which, at the nanoscale, is not observable until sufficient material modification has been performed (i.e., *observable* damage occurs much later than the *onset* of damage). In thin films, however, this energy threshold can be determined through a few (<10) laser shots on a single target through simple analytical modeling. Conveniently, for a Gaussian laser pulse, the threshold fluence (F_{th}, energy/unit area) is given by $D^2 = \omega_0^2 \ln\left(\frac{F}{F_{th}}\right)$, where D is the damage radius at a given fluence, F, and ω_0 is the beam waist. In other words, by plotting the *measured* damage area vs. incident laser energy, one can determine the threshold fluence via extrapolation.

We perform these thin film studies in a variety of shot conditions, including wavelengths spanning from 280 – 900 nm excitation, THz-tailored pulse trains, and excitation energies from well-below threshold to well-above threshold. An example of this data stream is shown in Fig. 13 below for pulse trains taken at 145 ksi of H_2 and D_2. In no cases do we observe statistically significant isotopic contrast with this approach.

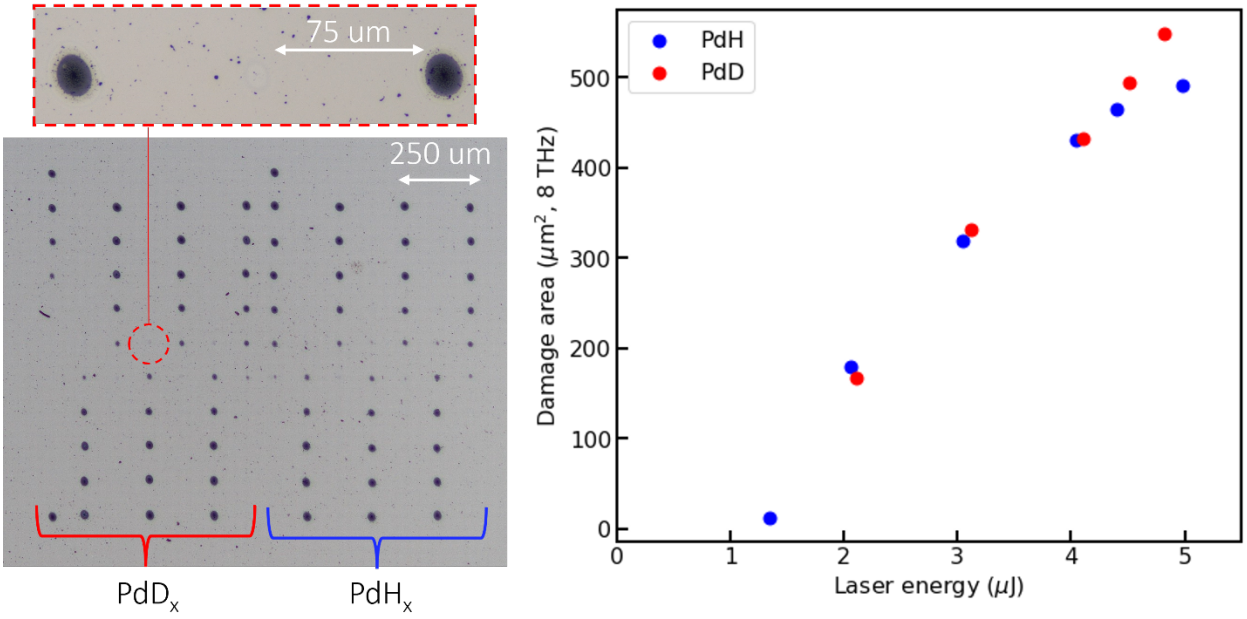


Figure 13: (Top left) Zoom in of damage to a 20 nm Pd film with a pulse train, (Left) Arrays of pulses for PdD_x and PdH_x with varying energies taken at different pulse train frequencies spanning from 0 to 12 THz and (Right) Example of the measured damage area vs. incident laser energy for a pulse train at 8 THz on PdD_x and PdH_x.

To more rigorously look for possible AHE events, we extend these thin film studies by performing post-mortem SEM imaging of nanofilm arrays. To test this USP stimulation at 100% loading conditions and consider the role of various materials/Pd alloys, we introduced a <100 fs pulse into our 400 ksi pressure vessel. As an additional comparison, we repeat this experiment using a long-duration thermal hold just below the melting temperature of the material of choice; these holds are achieved by introducing a continuous-wave (CW) laser. In this campaign, we thus intentionally remain *below* the damage threshold of the material due to the limited number of positions on a sample that can be obtained (i.e., our 400 ksi pressure vessel has a significantly reduced aperture relative to 150 ksi). Here, we look for >100 nm holes that form due to potential AHE events, as the laser is intentionally kept below the damage threshold. These films are 20 nm thick and span a range of material systems. This method allows for ***single-AHE event sensitivity.*** Table 1 highlights the results of this campaign; however, we note that no anomalous damage features of significance were observed or were not observed to be in isotopic contrast.

Table 1: Experiments performed at >400ksi with CW and USP stimulation with single event sensitivity.

Experiment ID	Sample ID	Material	Laser Source
BE3722	BD6828	Pd	3 kW, CW
BE3727	BD6829	Pd	3 kW, CW
BE3737	BD7171	$Pd_{0.75}Rh_{0.25}$	3 kW, CW

BE3739	BD7172	$Pd_{0.75}Rh_{0.25}$	3 kW, CW
BE3746	BD7318	$Pd_{0.75}Rh_{0.25}$	3 kW, CW
BE3750	BD7331	$Pd_{.75}Ni_{0.25}$	3 kW, CW
BE3752	BD7333	$Pd_{.75}Ni_{0.25}$	3 kW, CW
BE3752	BD7332	$Pd_{.75}Ni_{0.25}$	3 kW, CW
BE3756	BD7334	$Pd_{.75}Ni_{0.25}$	3 kW, CW
BE3757	BD7416	$Pd_{.75}B_{0.25}$	3 kW, CW
BE3761	BD7167	$Pd_{0.75}Rh_{0.25}$	3 kW, CW
BE3762	BD7418	$Pd_{.75}B_{0.25}$	3 kW, CW
BE3766	BD7417	$Pd_{.75}B_{0.25}$	3 kW, CW
BE3767	BD7418	$Pd_{.75}B_{0.25}$	3 kW, CW
BE3773	BD7635	$Pd_{.75}Pt_{0.25}$	3 kW, CW
BE3776	BD7636	$Pd_{.75}Pt_{0.25}$	3 kW, CW
BE3777	BD7637	$Pd_{.75}Pt_{0.25}$	3 kW, CW
BE3780	BD7508	$Pd_{.75}Ag_{0.25}$	3 kW, CW
BE3781	BD7509	$Pd_{.75}Ag_{0.25}$	3 kW, CW
BE3783	BD7510	$Pd_{.75}Ag_{0.25}$	3 kW, CW
BE3786	BD7639	$Pd_{.75}Pt_{0.25}$	3 kW, CW
BE3790	BD7511	$Pd_{.75}Ag_{0.25}$	Ekspla, USP
BE3794	BD7319	$Pd_{0.75}Rh_{0.25}$	Ekspla, USP
BE3795	BD7640	$Pd_{.75}Pt_{0.25}$	Ekspla, USP
BE3800	BD7512	$Pd_{.75}Ag_{0.25}$	Ekspla, USP
BE3803	BD7321	$Pd_{0.75}Rh_{0.25}$	Ekspla, USP
BE3805	BD7513	$Pd_{.75}Ag_{0.25}$	Ekspla, USP
BE3808	BD7322	$Pd_{0.75}Rh_{0.25}$	Ekspla, USP
BE3812	BD7356	$Pd_{.75}Pt_{0.25}$	Ekspla, USP
BE3815	BD7357	$Pd_{.75}Pt_{0.25}$	Ekspla, USP
BE3819	BD7335	$Pd_{.75}Ni_{0.25}$	Ekspla, USP
BE3821	BD7423	$Pd_{.75}B_{0.25}$	Ekspla, USP
BE3823	BD7422	$Pd_{.75}B_{0.25}$	Ekspla, USP
BE3828	BD7337	$Pd_{.75}Ni_{0.25}$	Ekspla, USP
BE3830	BD7426	$Pd_{.75}B_{0.25}$	Ekspla, USP
BE3831	BD7338	$Pd_{.75}Ni_{0.25}$	Ekspla, USP
BE3850	BD7324	$Pd_{0.75}Rh_{0.25}$	Ekspla, USP

Additionally, we performed experiments using nano-features designed to enhance light-matter interactions. Further, these were performed at conditions sufficient to achieve 100%, stoichiometric loading (>400 ksi). Finally, we also performed experiments on rare-earth deuterides/hydrides to achieve >300% loading conditions.

One of our original test structures used 8-sided pyramidal structures to allow for plasmonic excitation of the Pd. An example of these structures and their effect on the optical field is shown below in Fig. 14, where pyramidal structures with varying pitch (**p**) are shown; the color is a

representation of the electric-field enhancement due to strong light-matter coupling (i.e., plasmonic enhancement). This type of optical modeling played a large role in our choices of nanostructures for USP stimulation and, as described later, THz stimulation.

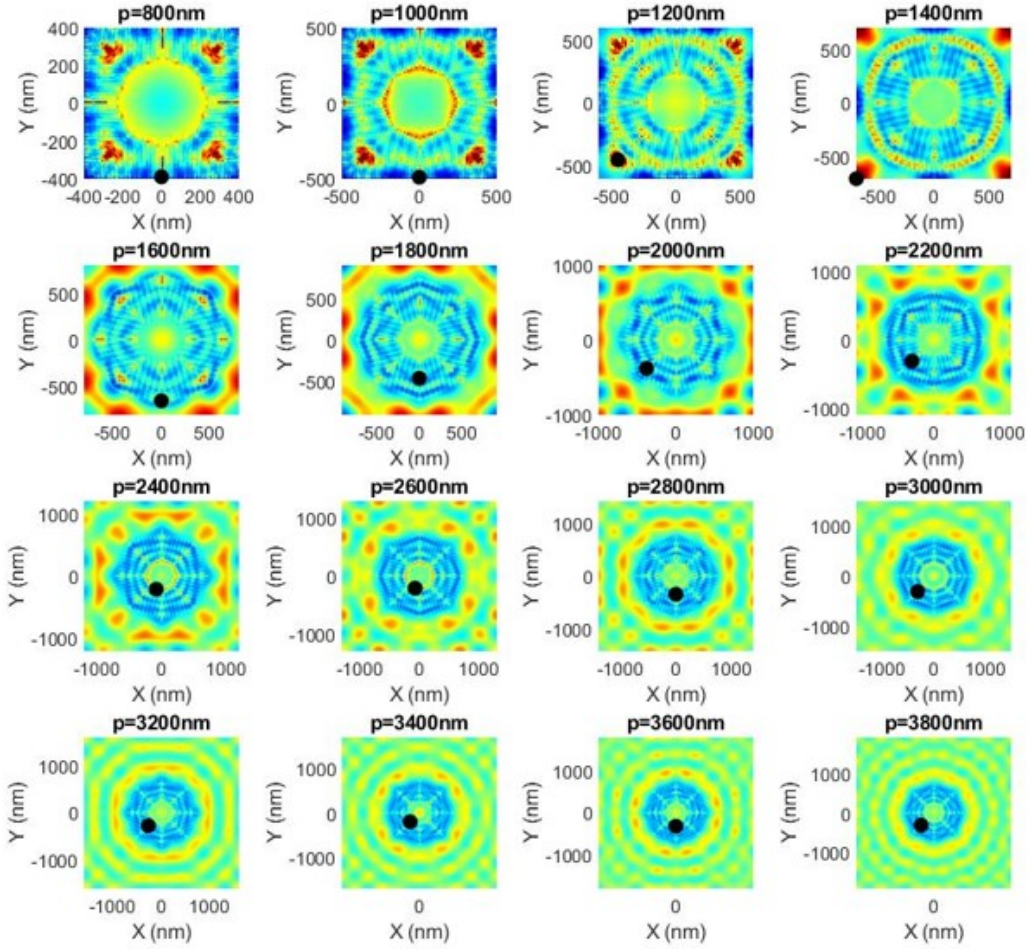


Figure 14: Electric field enhancement (color axis) for 8-sided pyramidal structures with varying pitch. The design wavelength was 400 nm. Calculations are performed using Lumerical.

Overall, there were a number of interesting phenomena/observations throughout this experimental campaign. One example, shown in Fig. 15, highlights nanoscale "jetting" that occurred due to USP stimulation of "bowties" (designed for internal plasmonic enhancement) at 145ksi. Interestingly, only bowties with this plasmonic enhancement displayed such features. To the best of our knowledge, no such process has been observed in literature and is truly unique. However, in no cases, including experiments performed at >400 ksi and achieving stoichiometric loading, nor geometries that allow for strong-field enhancement, do we observe isotopic contrast suggestive of

AHE. Including the bowties in Fig. 15, even these phenomena unique to our loaded Pd studies are found to occur in both deuteride and hydride alike.

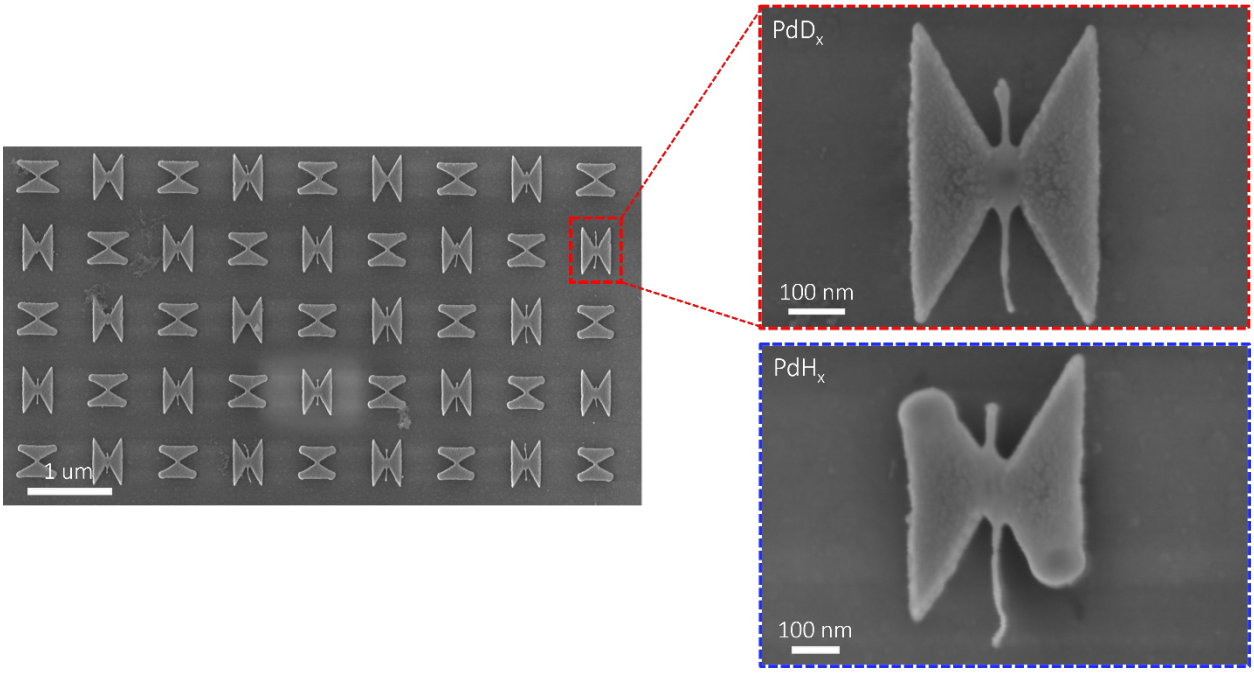
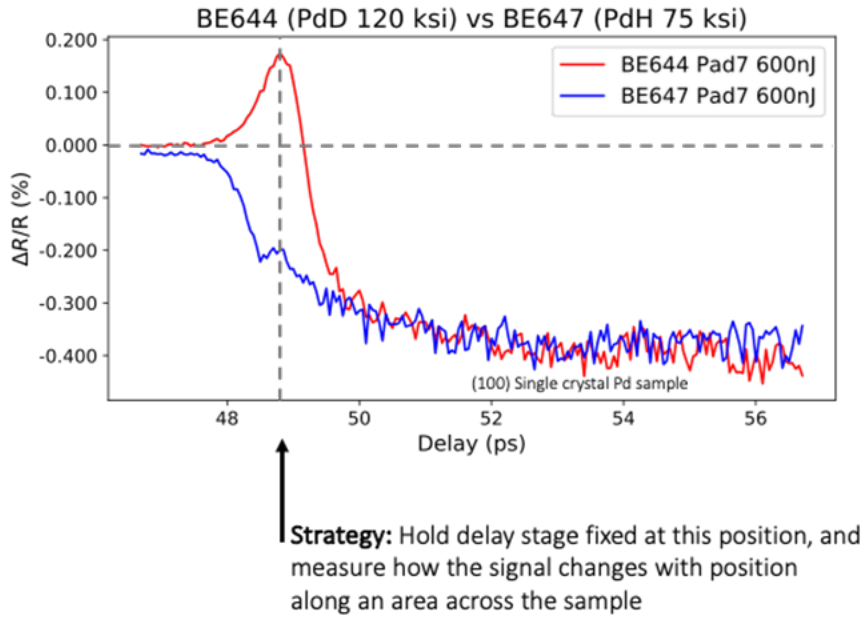


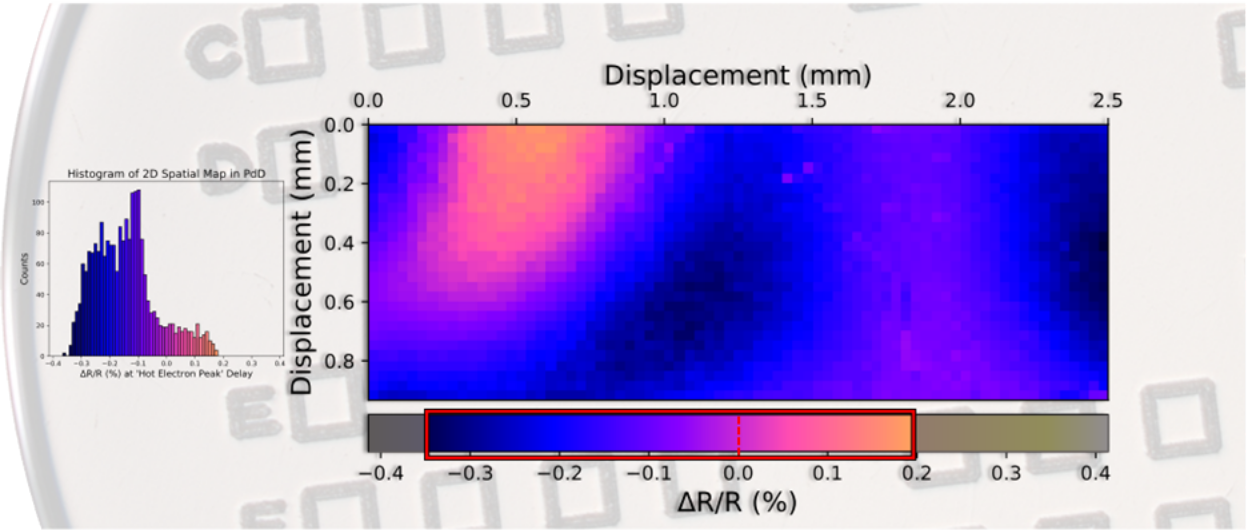
Figure 15: Example of nanostructures used for USP Archaeology studies. (Left) Overview image of "bowties" designed for plasmonic enhancement at the central point of two triangular structures. (Right) Shows unique "jetting" of melted material from the location of plasmonic enhancement; interestingly, this only occurs for orientations that undergo such enhancement. However, as shown, both deuteride (red) and hydride (blue) show these features.

USP-Appendix: Artifacts in TR measurements

TR "Imaging" for spatial variations due to mechanical stress on optical windows:

Figure 16: Concept and Example Result for Fixed-Delay 2-D Scan of Short Delay Feature



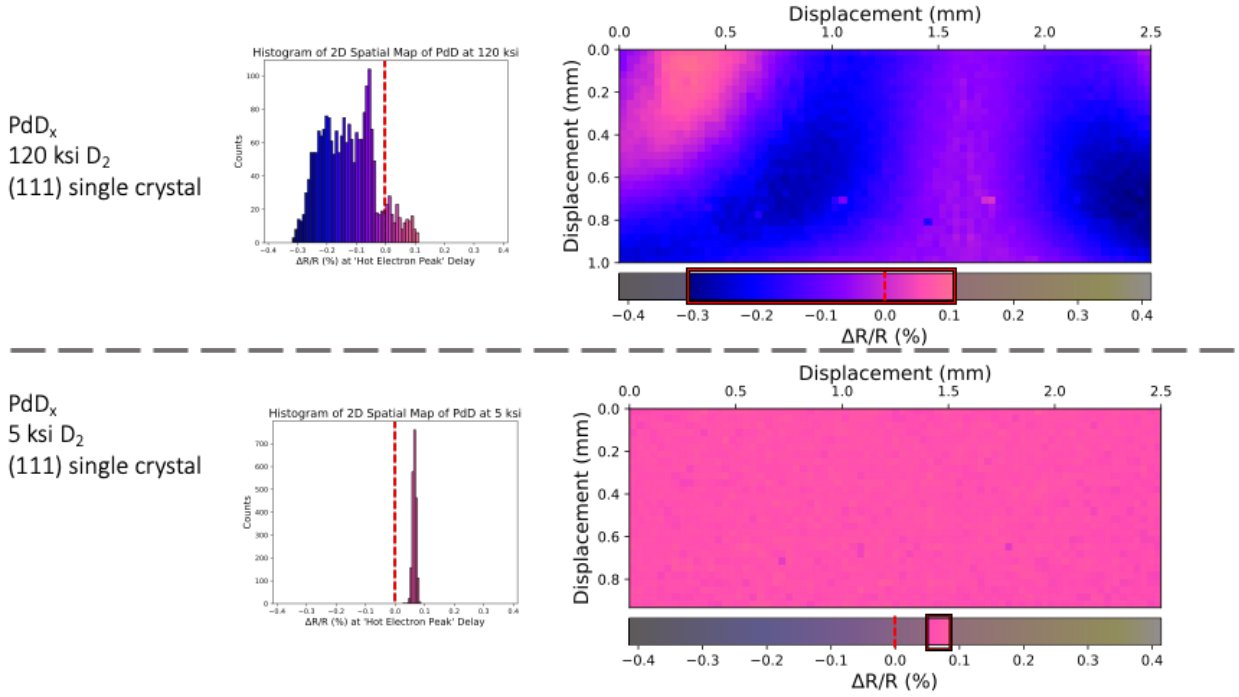


Concept for fixed-delay spatial scan. Delay stage is held at the position of interest, and the sample is scanned in two dimensions while the magnitude of the TR signal is recorded. Example map is shown overlain with an optical image of the single-crystal palladium sample. Histogram to the left indicates the spread of the feature of interest over the mapped area, and the range is highlighted in the color bar.

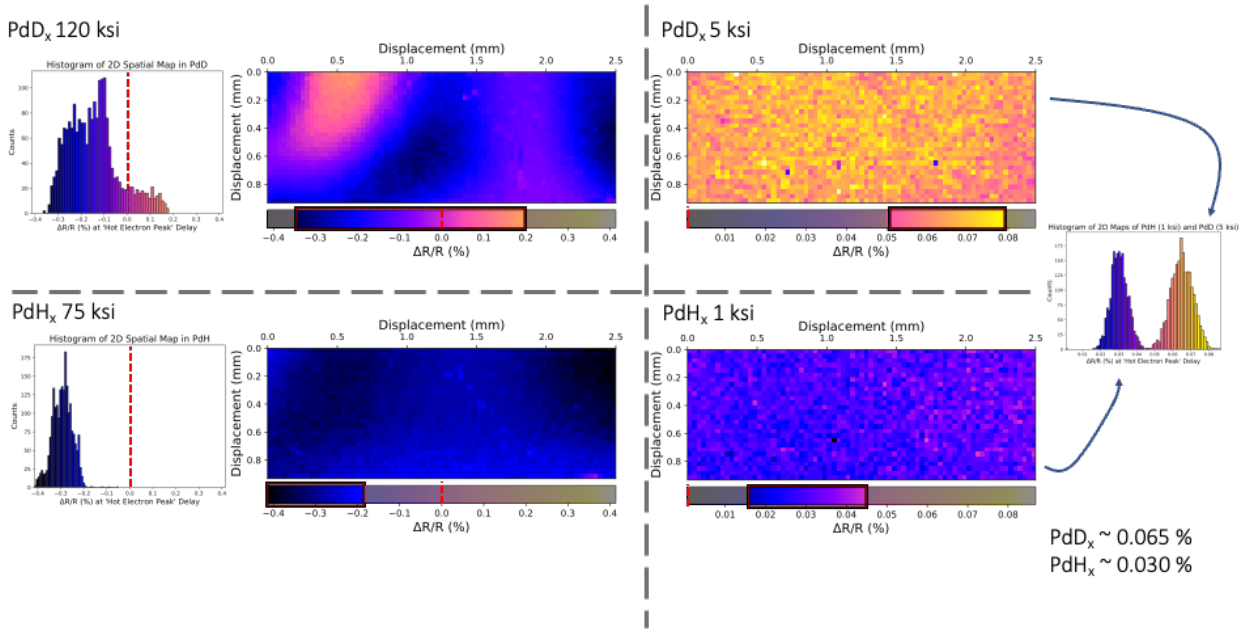
Transient reflectance data for the most part remained consistent between isotopes. However, it is worth noting that in early TR experiments, a sharp feature *occasionally* appeared at short pump-probe time delays of up to 2 picoseconds; this feature initially displayed as a form of isotopic contrast. To understand these differences, we performed fixed-delay two-dimensional scans to map this feature over the sample area in both isotopes. Figure 16 above details this approach and shows an example result.

In our investigations of this behavior, we identified several key observations:

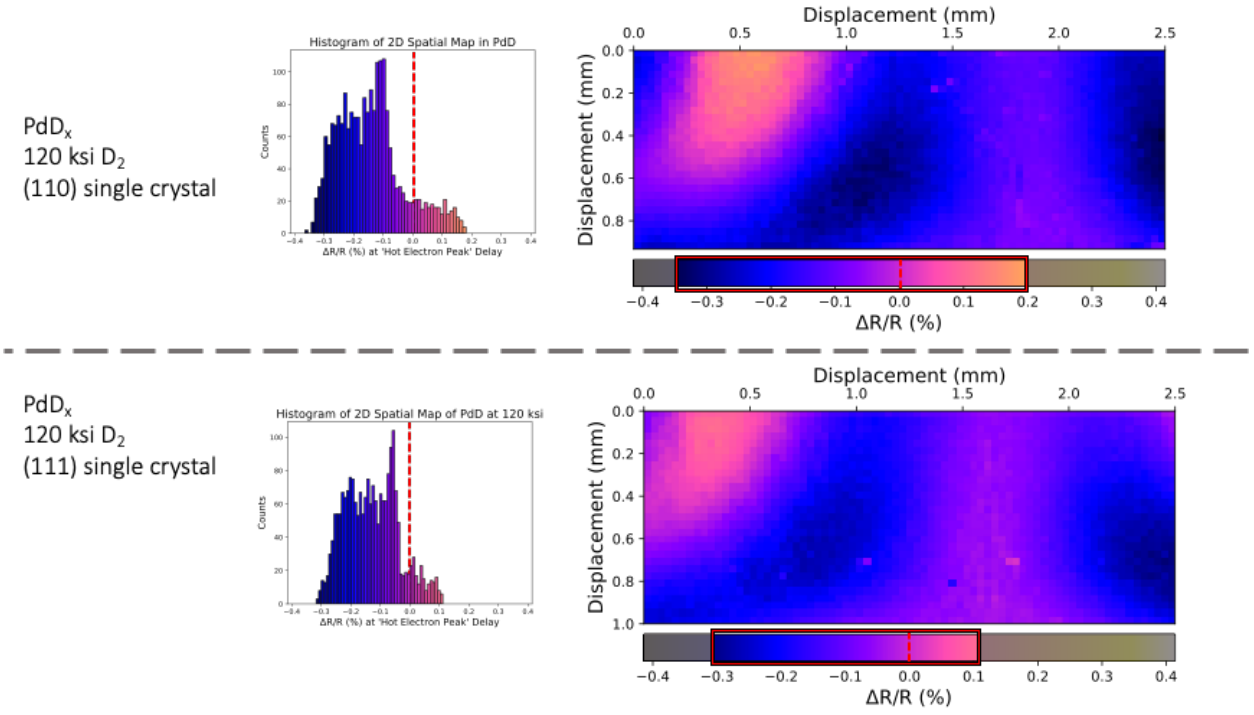
1) The spatial variation only appears at high pressures.



2) There appear to be isotopic differences in the short-delay feature, with larger positive magnitudes appearing in deuterium. This trend persists even at lower pressures when the distributions are narrower. Note: here we've matched *loading* rather than pressure.



3) Nearly identical spatial patterns appear for different samples, which **indicates this is an apparatus effect rather than an interesting sample effect**.



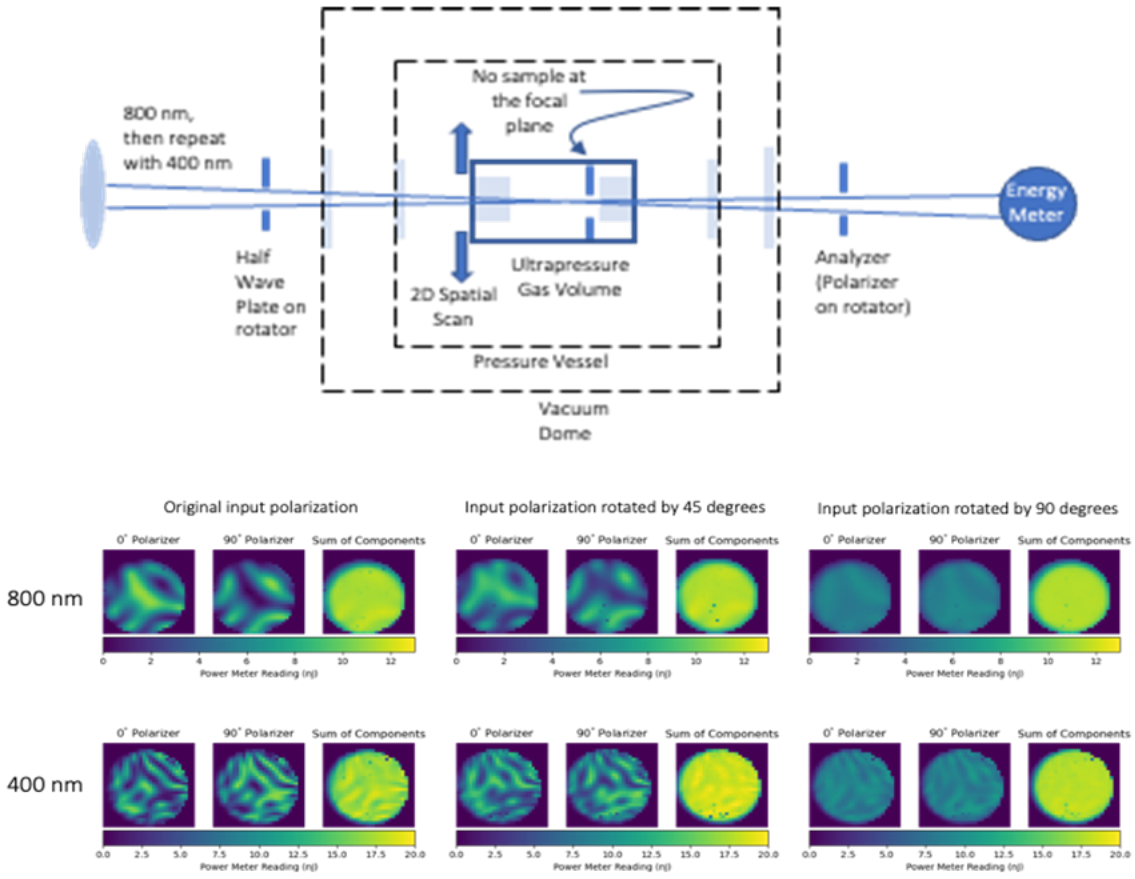
4) We also performed extensive measurements on gold samples, which do not load, and observed similar patterns. Our gold results are summarized below:
    - The TR signal at the peak delay shows spatial variation
        - Pattern ***depends on the apparatus***, not the sample
        - Pattern appears ***only at the peak delay***, not a few ps later
        - Pattern occurs ***only at high pressure, in H/D*** but not He
    - The Baseline 800 nm reflectance also shows spatial variation
        - Pattern ***depends on the apparatus***, not the sample
        - Pattern is ***not delay dependent***
        - Pattern occurs ***only at high pressure, in all gases***

As we demonstrate below, these observations are consistent with two different polarization shifts. Because transmission and reflection in different optics depends on polarization, shifts in the polarization couple directly into magnitude changes in our measured data. Our results indicate two different polarization shift mechanisms. One mechanism involves a stress dependent birefringence in the sapphire pressure windows, which shifts the baseline reflectance and results in a characteristic pattern based on the window geometry alone. The other mechanism is a polarization shift induced by the pump, in ultrapressure diatomic gases, which therefore occurs only shortly after the pump and probe occur at the same time. This explains the short-delay feature that is delay-dependent and that does not normalize away with the baseline reflectance.

Transmission Measurements in Ultrapressure Gas Reveal Polarization Effect

After determining that the spatial variation in the short-delay feature followed very similar patterns even after changing the sample, it became clear that the origin of the variation rested in the design of the apparatus. We performed a series of transmission measurements to explain the origin of the observed effect. As shown in Figure 17, we measured the energy transmission of each polarization component independently. We found that the total energy transmitted did not vary strongly with position, but the polarization did vary. Furthermore, the shape of the revealed pattern is consistent with the shape of the spatial variation in the short delay features in our TR data. This suggests that the sapphire window exhibits stress-dependent birefringence, leading to a spatially varying polarization shift in the transmitted light.

Figure 17: Polarization Varies Spatially but Transmitted Energy is Consistent



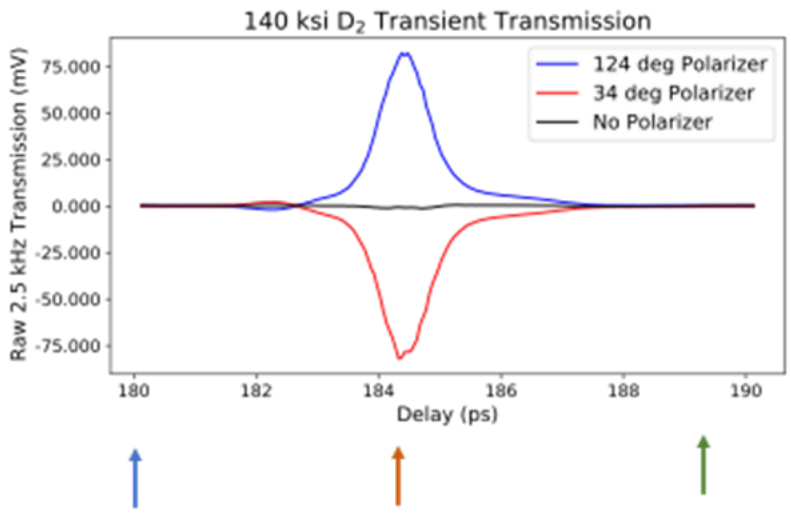
We measured the polarization of the transmitted probe and pump beams as we scanned over the whole ultrapressure bore. One scan was taken with the analyzer polarizer at 0 degrees, and then a second scan was taken with the analyzer at 90 degrees. These measurements were summed to demonstrate that the total energy remained consistent across the spatial scan, even though the polarization changed considerably. The measurement was repeated for three different input polarization angles.

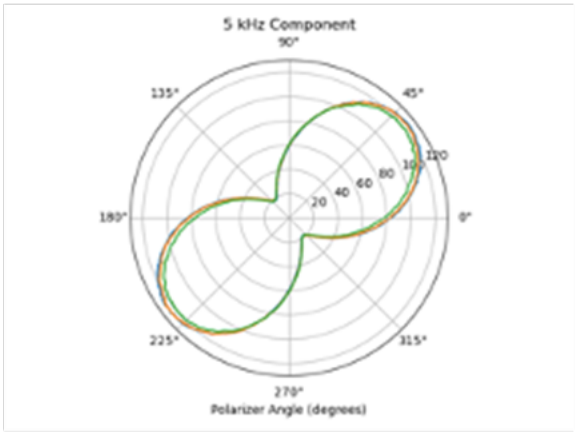
Transient Transmittance Measurements Explain Short-Delay Feature as Gas Effect

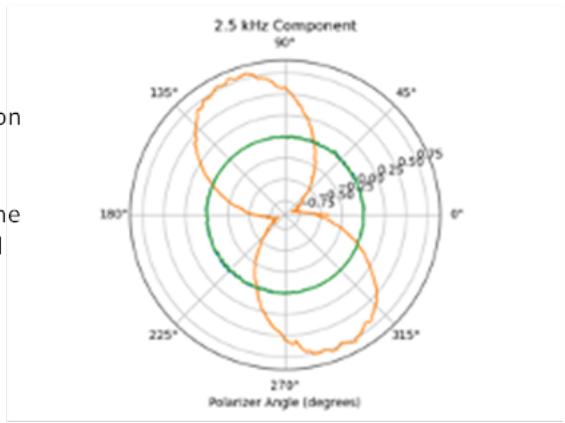
The polarization shift due to the window alone is not enough to account for the variation we observe at short delays. Namely, the effect should normalize away in $\Delta R / R$ if the pump does not induce additional polarization change in the probe. To measure any induced pump-induced polarization change, we performed transient transmittance spectroscopy using an analyzer polarizer.

We found that indeed, with no sample present in the vessel, there is a polarization shift induced in the probe beam when it overlaps with the pump beam, as shown in Figure 18. We suggest this is some form of a Kerr effect in the ultrapressure gas, which is consistent with its appearance in diatomic gases but not in helium. This entirely explains our observed short-delay feature in transient reflectance data as a pump-induced polarization artifact.

Figure 18: Transient Transmittance Showing Pump-Induced Polarization Shift







a) Transient transmittance data showing pump-induced polarization shift. This explains the variable short-delay feature in our transient reflectance data.
b) Bowtie plot showing the variation in the measured 5 kHz baseline transmitted probe light at three delay stage positions (indicated by the blue, orange, and green arrows on the TT plot) as the analyzer polarizer is rotated.
c) Bowtie plot showing the variation in the measured 2.5 kHz component in the transmitted probe light (i.e. the modulation due to interaction with the pump) as the analyzer polarizer is rotated. At the delay stage positions indicated in blue and green, there is no pump-related signal. However, at the delay stage position indicated in orange, the signal shows strong polarization dependence that is not the same as the baseline transmittance.

Isotopic Difference at Longer Delays Explained as an Interference Effect:

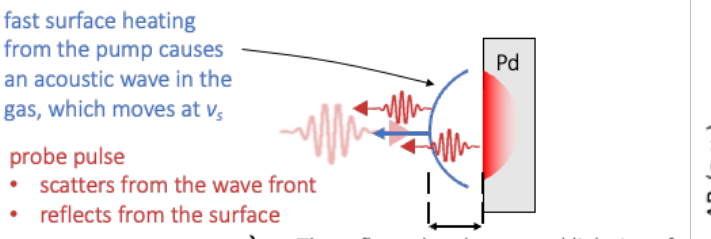
At pump-probe delays from 10 ps – 300 ps, we observed an isotopically different waveform in transient reflectance data, shown in Figure 19, which we ultimately explained as an interference

effect. The pump pulse superheats the surface and causes a shockwave to propagate out into the gas at the speed of sound. Part of the probe pulse then scatters off this wavefront and interferes with the reflected pulse. The interference is constructive when the wavefront has moved half a wavelength during the pump-probe delay time. The speed of sound in the dense gas depends both on pressure and on isotope, so the period of the observed waveform is different between deuterium and protium.

In the experiment shown, BE544, a palladium slug sample was elastically loaded up to 100 ksi in deuterium. Transient reflectance data was acquired at three different stimulation pump energies, at two different locations on the sample (labeled points A and B). The 10 ps – 300 ps waveform showed a magnitude that increased with pump energy, with a consistent period of ~100 ps. The waveform was consistent at both points measured. The sample was then elastically de-loaded, and then elastically loaded up to 100 ksi in protium. The measurements were repeated at two new locations (labeled Points C and D). The observed period in 100 ksi protium was ~80 ps.

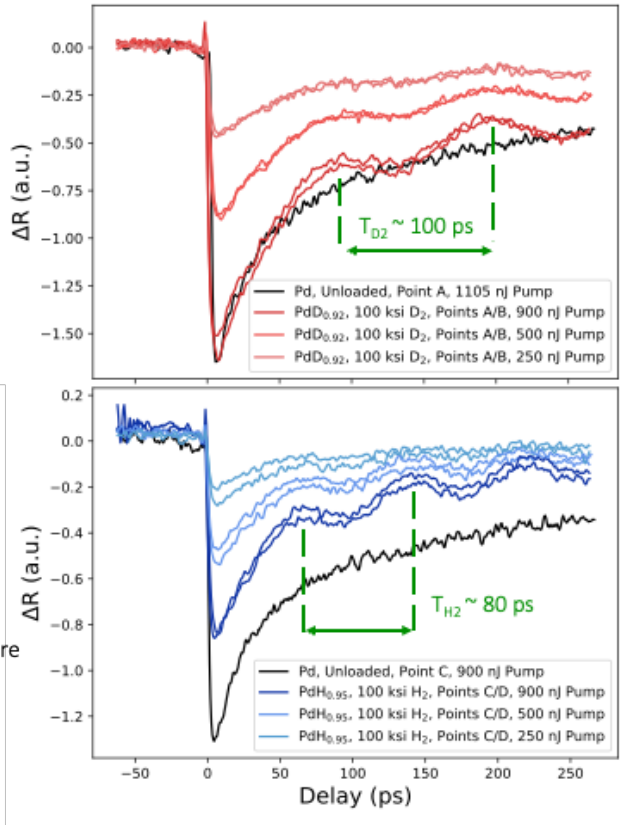
Figure 19: Longer-Delay Oscillations Explained as an Interference Effect

The pump pulse superheats the surface of the palladium, which causes an acoustic shockwave to propagate into the gas. Probe light that scatters off of the wavefront can interfere with the reflected light when the right conditions are met. These conditions depend on the speed of sound, so the effect appears different between isotopes.



$$v_s \mathrm{T} = \frac{\lambda}{2n}$$

The reflected and scattered light interfere constructively when the wave front has moved a half wavelength during the delay, T, between pump and probe



Argument for Focusing on Long Delay Features

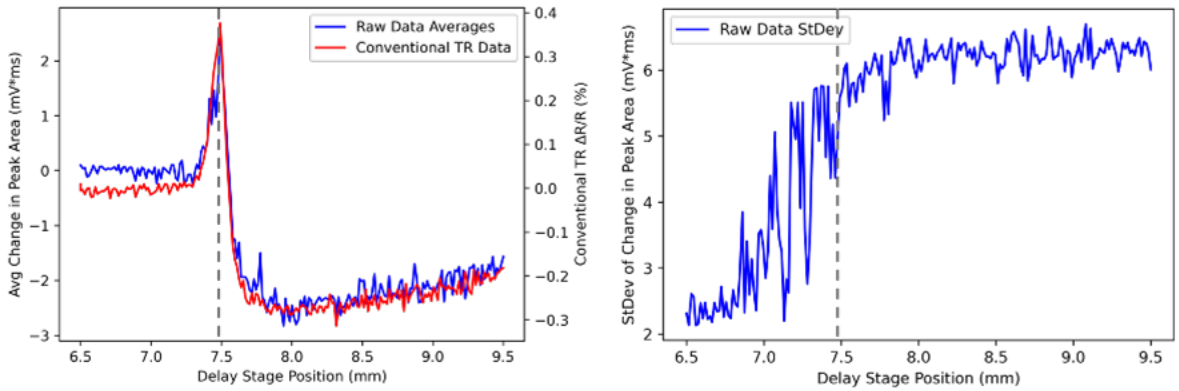
Given the observed variation in short-delay features, we suggest that long-delays are far better indicators of the presence or absence of AHE. Admittedly, AHE is not expected to occur in precisely the same way for every stimulated pulse. However, any AHE that does occur would uniformly raise the temperature of the sample, crucially *for all times after the AHE event before complete temperature relaxation*. This means that even in an averaged measurement, some isotopic difference should be apparent at longer delays, before complete relaxation, unless AHE does not happen at all in the overwhelming majority of pulses; the Statistical Transient Reflectance technique discussed below was an attempt to address this possibility. Thus, later measurements focused on much longer delays than those presented above.

Statistical Transient Reflectance Measurement:

In conventional transient reflectance measurements, lock-in demodulation is used to characterize an average effect over thousands of pulses. However, because AHE is expected to be a rare event, it is possible or even likely that most individual pulses will not trigger AHE, and the averaged signal will not be measurably affected by the rare events. To address this concern, we developed a technique that collects raw, time-domain oscilloscope measurements of the photodiode response to each individual pulse. This still allows us to compute the average, for example by retroactively applying a digital lock-in filter, or by identifying a consistent quantity such as the peak height or area under the curve for each pulse. Additionally, however, we can quantify the distribution and compute quantities like the standard deviation that can give insight to the prevalence of rare outlier events.

We performed proof of concept experiments using this technique, but due to higher priorities this was not implemented in operational AHE experiments.

Figure 20: Proof of Concept Data for Statistical Transient Reflectance



Conventional TR data is shown in red (right value axis) and is shown to be consistent with data extracted from computing the area under each pulse in a time domain oscilloscope trace (left value axis). Additionally, computing the area under each pulse in a time domain trace enables measurement of distributional quantities such as the standard deviation. Here, we see that the standard deviation increases when the probe arrives after the pump, which we suggest is the expected effect given the intensity noise in the pump pulses.

Terahertz

Executive Summary

Here, we extend our use of ultrashort laser pulses as a means of generating intense THz-wavelength pulses as a means of potential stimulation of AHE. These intense pulses are further enhanced by the implementation of PdD_x phonon-resonant split-ring structures that act as microscale antennas. Ultimately, we find that the electron-driven optical properties of metallic PdD_x causes screening that well exceeds the potential interaction of light with the phonon mode of interest.

Stimulating optical phonons for AHE

A non-negligible fraction of LENR/AHE literature discusses the importance of phonons, or vibrations of atoms, within palladium deuteride toward generating anomalous heat events. It is worth noting that such literature is relatively qualitative in nature and does not provide mathematical rigor to such "theories," but, nonetheless, we seek to leave no stone unturned. The difficulty with the lack of rigor in these discussions is that there is negligible work in the LENR community that is definitively stimulating these atomic vibrations of interest within PdD_x. In other words, many authors cite the importance of these phonons, but their methods do not stimulate the phonon modes!

This limitation is likely due to the fact that there are very few ways of "coupling" to phonons in a material system, particularly metals. As described in our USP section, we tested this theory by "beating" the sample via repetitive short pulse heating at the resonant frequency associated with the Pd-D optical phonon mode centered at ~14 THz. While this qualitatively may be expected to increase the amplitude of the phonon, it certainly is not a "known" method in literature for phonon stimulation.

In fact, the only known method of directly interacting with and manipulating phonons at this frequency of interest is through optical stimulation, where the frequency of the *photons* are resonant with the frequency of the *phonon*/lattice vibration. This is fundamentally different than our above resonant heating approach: In the resonant heating approach in USP, we simply "hit" the material repeatedly at 14 THz. However, this mechanical driving force is simply not equivalent to a direct interaction with atoms themselves. In contrast, when light has a frequency or wavelength resonant with the atomic vibration, direct coupling can be directly observed; this is the basis for a number of spectroscopy methods such as FTIR (Fourier Transform Infrared

Spectroscopy) and has been shown in multiple studies to provide a manipulation of phonons, including an increase in their amplitude[34].

In the following sections, we discuss our attempts at using resonant-frequency optical pulses (THz pulses) to stimulate AHE events. This includes the development of intense THz sources, measurements on the coupling between these sources and loaded Pd, as well as the development of resonating structures to enhance the electric field intensity experienced by the PdD_x.

Developing intense THz sources

The primary method of generating *intense* THz sources is via non-linear processes that occur when an intense, ultrashort pulse from a laser interacts with a gas or liquid. Consider the electric field amplitude of a laser pulse is given by $E = \sqrt{P * 2\eta}$, where P is the power flux of the laser source and η is the resistivity of free space. Thus, to generate a sufficient electric field, E, for nonlinear processes (nonlinear processes typically have susceptibilities 6-8 orders of magnitude lower than linear susceptibilities), the required power increases with increasing field intensity (i.e., more power is needed to drive small changes in E when E is large than when E is small – there are diminishing returns!). This limitation suggests that extremely large powers are required for large electric fields. For this reason, ultrashort pulse lasers become an ideal "driving" function for the electric fields of interest; ultrashort pulse lasers deliver their total energy in extremely short periods of time. Since the power follows $P = \frac{Energy}{\tau_p}$, where τ_p is the duration of the pulse, the *shorter* the pulse duration, the greater the power, even for the same net energy input. There are no other methods of generating equivalent power densities; for example, a factor of $\sim 10^{12}$ is gained by relying on a picosecond pulse rather than a 1 second pulse duration. The relative intensity of the pulses used in this work produces *peak* powers of nearly 100 GW, the equivalent of nearly fifty large-scale nuclear reactors.

Qualitatively, this process of THz generation via ultrashort pulse lasers functions by the extremely high electric field of the pulse driving electrons out of the molecules in a gas/liquid; an example image of this plasma-forming process is shown below in Fig. 1. Notably, in the absence of the graphical depiction of the laser and corresponding THz emission, both beams are undetectable by eye or with standard cameras. The ionized gas/plasma is visible due to the large number of wavelengths emitted from this ionization process; these wavelengths span into the THz region under appropriate conditions.

[34] Woerner, M., *et al., Physical Review Letters* **122**, 107402 (2019).

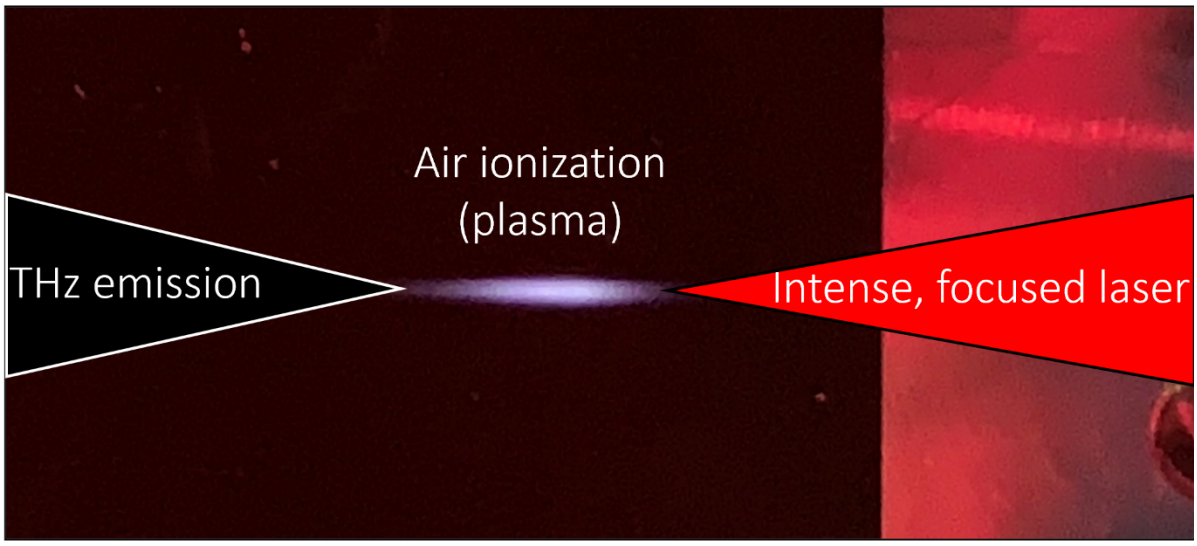


Figure 1: Experimental image of laser-induced ionization for THz pulse emission.

Due to the large wavelength of THz pulses, there are few methods available for *detecting* the generation of these pulses. As mentioned above, standard cameras do not detect these wavelengths of light; in fact, even specialty electronics specifically designed for "far wavelength" operation do not generally reach the frequencies of interest. The primary method for detecting infrared light of these wavelengths is via interferometry, where the infrared light is split into two paths and the constructive/destructive interference is measured as a function of path length (this is the principle of FTIR discussed later). While this works sufficiently for *continuous* waves of light, the method struggles at the low duty cycles for intense, quasi-single cycle THz pulses of interest here. Rather, we rely on a "pump-probe" method of detecting the generated THz pulse.

Our method of detecting these THz pulses is a four-wave mixing process referred to as Air-Biased Coherent Detection (ABCD)[35]. As we use an 800 nm pulse to *generate* the THz wave, we can also use the pulse to *detect* the THz wave. This is done by splitting off a small, ~10% fraction of the 800 nm pulse prior to generating the air-plasma. Once the air-plasma is generated, we use Si optics to remove the remnant 800 nm beam within the THz beam path; this isolates our THz "pump" beam to wavelengths more than ~1200 nm, thus eliminating the visible components of the air-plasma as observed in Fig. 1. From this, we recombine the 10% "probe" beam at 800 nm and the THz "pump" beam such that they interact with each other spatially and temporally; by varying the temporal overlap of the two beams, we can extract the temporal profile of the THz pulse by reading-out the second harmonic of the 800 nm probe as a function of time-delay.

Quantitatively, this ABCD process is dictated by the following equation,

$$E_{2\omega}(t) \propto \chi^{(3)} E_{THz}(t) E_{\omega}(t) E_{\omega}(t),$$

[35] *X. Lu and X.-C. Zhang, Applied Physics Letters* **98**, 151111 (2011).

where $\chi^{(3)}$ is the third-order susceptibility of the detection medium (here, the medium is air), E_ω is the electric field of the fundamental probe wavelength (here, ω = c/800 nm), E_{THz} is the generated THz pulse, and $E_{2\omega}(t)$ is the generated second harmonic via four-wave mixing of the two electric fields (here, 2ω = c/400 nm). Notably, this detection is simply the inverse of the generation process which follows,

$$E_{THz}(t) \propto \chi^{(3)} E_{2\omega}(t) E_\omega(t) E_\omega(t),$$

as we also rely on the second harmonic generation of our 800 nm generation beam to create a sufficiently intense THz pulse.

An example of this time-domain detection method is shown below in Fig. 2. As shown, the second harmonic signal displays various oscillatory features within the electric field associated with the THz pulse. While this is our primary information of interest, it is inconvenient for understanding the resultant pulse. Rather, we can perform a Fourier-transform on this time-domain signal to extract the frequencies contained within the pulse; this is shown on the bottom of Fig. 2. This displays the wide range of frequencies contained within the THz pulse; our goal is to generate *intense* peaks as close as possible to 10-14 THz. In the example shown here, we can see that the highest nearby frequency peak is around 8 THz and corresponds to the most intense peak within the upper waveform.

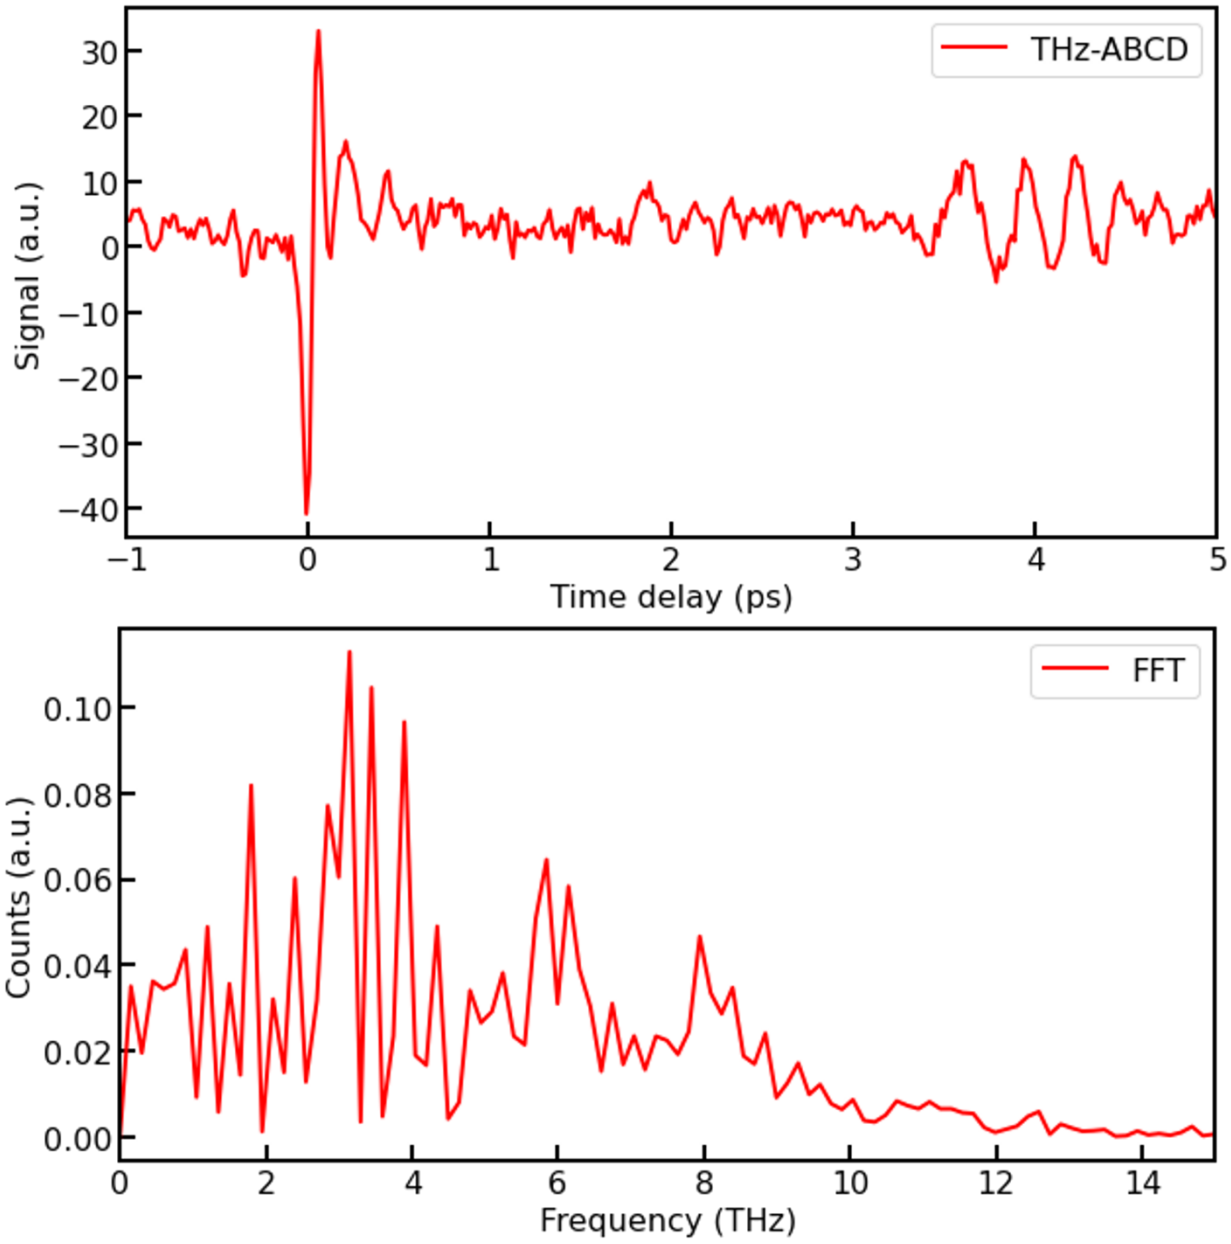


Figure 2: (Top) Representative ABCD signal produced via four-wave mixing of the generated THz pulse with an additional 800 nm "probe" beam. (Bottom) Fourier-Transform of the upper ABCD waveform, providing a clearer representation of the produced frequencies.

We estimate the peak electric field based on several experimentally measured properties of the pulse. The peak electric field for a THz pulse is given by,

$$\left|E_{peak}\right|^2 \approx \frac{2W\eta_0}{\tau A},$$

where W is the pulse energy, η_0 is the impedance of free space, τ is the pulse width, and A is the beam area. The pulse width can be extracted through the ABCD waveform, e.g., as shown in Fig. 2. The beam size is determined using a THz Beam Profiler (WinCamD-THz); we measure the pulse width at the beam waist. The pulse energy is given by bolometry measurements of the produced beam. These measurements, when performed in a dry atmosphere (water molecules have strong absorption bands in this frequency range, greatly reducing the beam intensity), suggest field intensities on the order of 500 – 1000 kV cm^{-1}.

Resonators

One approach to further increase the electric field experienced by a material is via beam interactions with a metamaterial. For example, radio waves rely on antennas for enhancement of the electric field and optical waves can interact with plasmon waves within a metal for enhancement of the electric field. In the former case, antennas are generally developed toa maximum frequency of about 0.1 THz, beyond which resonant structures become significantly more difficult to create. On the other end of the spectrum, metal plasmons and other polaritonic modes are generally supported from 100s of THz down to about 30 THz, where, again, a gap is observed. This range, 0.1 – 30 THz, is referred to as the "THz gap," as useful power generation and receivers (antennas) are unfeasible.

While a large amount of research has been performed to reduce this gap, there remains difficulties in resonant enhancement of the electric field. To address this and further increase the electric field intensity of our THz wave, we investigated the production of split-ring resonators[36] that specifically enhanced the frequency range of interest (10-14 THz).

To optimize the design of these structures, finite element calculations were performed on a large variety of metastructures; these calculations were guided by electromagnetic constant measurements of Pd/PdD_x and other materials of interest (e.g., the dielectric substrate, gap materials, etc.). Further, we performed measurements post-fabrication to verify these calculations. An example of these finite element calculations is shown in Fig. 3 below, where the color-axis displays the normalized electric field (electric field / incident electric field) at the resonant frequency. Additionally, an example of the idealized design for 10-14 THz enhancement is shown in Fig. 3, where we determined that *double* split-ring resonators (dSRRs) led to even greater enhancements in the electric field. Ultimately, our simulations suggest that the electric field within the gap of these dSRRs is of order 100-1000x the incident electric field intensity; these resonators are made from Pd to utilize it as a potential stimulation material.

[36] Driscoll *et al., Applied Physics Letters* **91**, 062511 (2007).

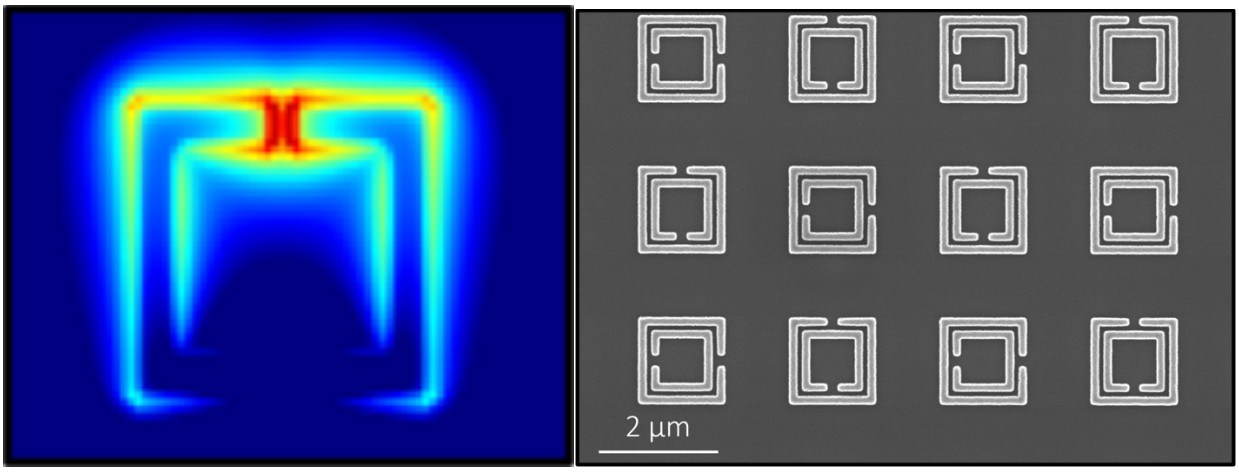


Figure 3: (Left) Finite element calculations of a single split ring resonator. The color axis is the normalized electric field, where red denotes enhancements in the electric field and blue denotes no enhancement. Most of the field enhancement, as expected, occurs within the "split" of these structures. (Right) Double split-ring resonators designed and built by us for improved enhancement in the 10-14 THz wavelength range.

To further highlight the benefit of these structures, we show below, in Fig. 4, a direct comparison between various resonator structures, including no resonator, a single split ring, and split rings with offset gaps (i.e., not at the midpoint of a side). We experimentally measure the *transmission* of light through these structures to determine the resonant enhancement: The absorption of light is directly proportional to the "coupling" between the electromagnetic wave and the structure of interest. Coupling in excess of a "non-resonant" structure is direct verification of electric field enhancements. Further, our simulations provide not only enhancement, but an expected transmission function: the simulated transmission can be directly compared to our experimental measurements to allow for improved quantification of this enhancement.

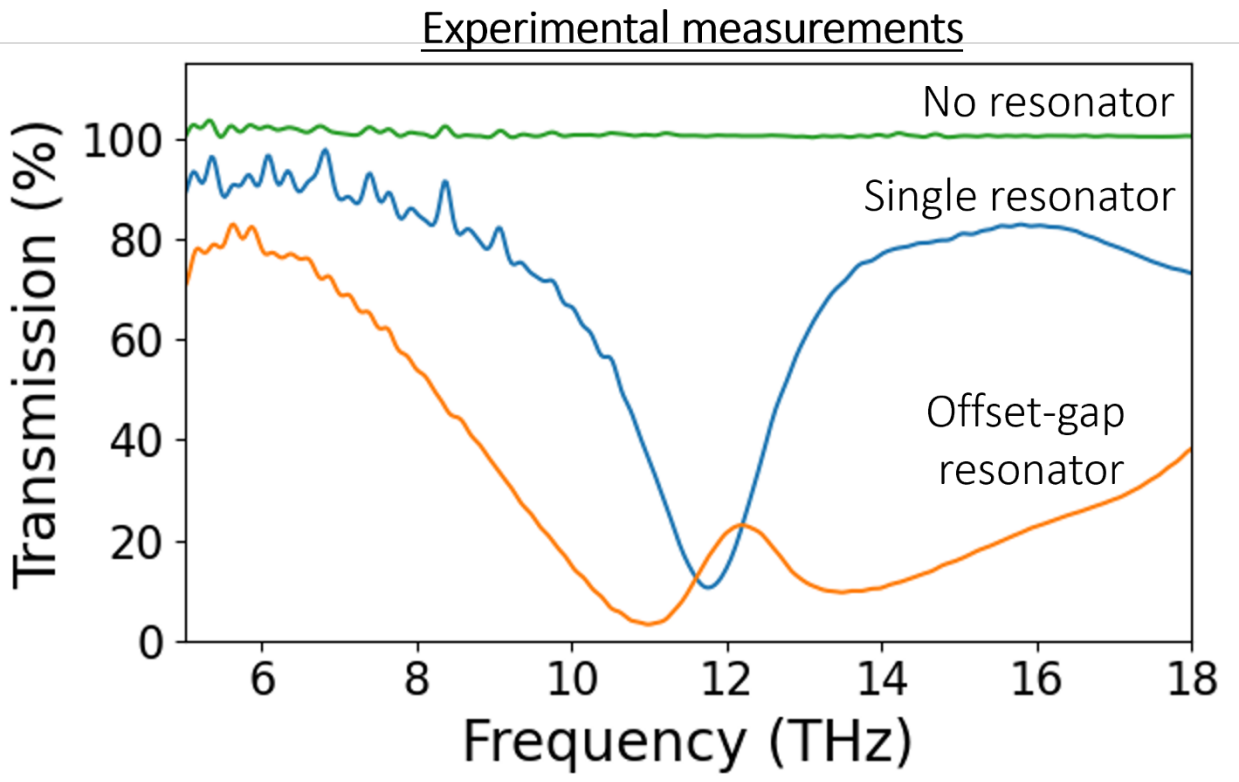


Figure 4: Experimental FTIR measurements on (green) Si with no resonator, (blue) Si substrate with a single split ring made of Pd, and (orange) an offset-gap designed Pd resonator on a Si substrate.

Notably, we performed THz pulse stimulation on structures with various resonant frequencies/materials to look for AHE events via the production of "hole" events within the metamaterial. These experiments were performed by sweeping/rastering the THz beam over a wafer containing a large number of these dSRR structures, as well as moving the beam through the focal plane (i.e., performing this experiment at repeated distances from the expected beam focus) to ensure that the beam reached maximum intensity. Ultimately, **we did not observe laser-induced nor AHE-induced hole/void formation within these dSRRs using the intense THz pulses described above**. We believe that there is simply insufficient coupling between the light and metal of interest (PdD_x); this notion is well-supported by literature and internal measurements on PdD_x and discussed in more depth below.

Electromagnetic coupling to phonons in metals

While THz and mid-infrared coupling of light into *dielectric* materials is well-known, predictable, and commonly studied, there is a large lack of literature on the coupling of mid-infrared light to

phonons within *metals* such as Pd. Generally, this is not surprising as the dielectric function that describes light-matter interactions in metals follows the Drude model. While the Drude model can support resonant interactions with, in theory, an oscillatory mode, one would not expect this to occur at energies *below* the plasma frequency; in fact, this limitation is the reason why metals do not support strong coupling in the mid-infrared and researchers have turned to dielectrics!

The Drude model is derived from kinetic theory (quasi-classical calculations) that assumes atoms within a material are "rigid" and simply act as stationary scattering sites for the electrons that cause electrical conduction (and reflectivity) within metals. Notably, this simple model provides *extremely* good agreement with the optical properties of metals at wavelengths far from *electronic* transitions. The optical properties of a metal follows the Drude dielectric function, ϵ, given by,

$$\epsilon(\omega) = \epsilon_0 \left(1 - \frac{\omega_p^2}{\omega^2 + i2\omega\gamma}\right),$$

where ϵ_0 is a constant representing the dielectric function at infinite frequencies, ω_p is the plasma frequency (the frequency in which electrons stop responding to light: this is generally at deep UV wavelengths for metals), and γ is the scattering rate of electrons due to the rigid lattice. This function is notoriously smooth and leads to the extremely high reflectances of metals in infrared and greater wavelengths (i.e., THz wavelengths).

The key point, is simply that the optical properties of metals such as Pd/PdD_x are well-described by a model that invokes *no* consideration of the phonons other than their role on free-electron scattering; the interaction of light with metals at THz wavelengths is dictated by electrons, *not* phonons of interest for our AHE studies. We confirmed this by a number of FTIR experiments on monitoring the absorption of THz/mid-infrared light with PdD_x and PdH_x – in no cases did we observe absorption due to phonon modes of interest. This makes electromagnetic stimulation a non-ideal method for testing the phonon hypothesis of driving AHE; we explore alternative approaches to this hypothesis in our USP effort.

High-Pressure Electrochemistry

Executive Summary

This report presents a comprehensive analysis of experimental research conducted in the field of high-pressure electrochemistry, focusing on achieving high loadings at elevated temperatures. The study spans various experiments, from explorations in supercritical water environments to high-pressure, ambient-temperature conditions, and includes substantial hardware developments to accommodate these unique experimental settings. The primary objective of this research is to facilitate higher loadings at elevated temperatures than what could be achieved with gas-loading alone.

The initial hypothesis centered on using supercritical water as a medium for electrochemical loading, aiming to maintain a uniform, bubble-free electrochemical loading environment at high temperatures. However, achieving measurable H/Pd loading levels at temperatures >200°C proved to be challenging, and we could not measure stable loading in a supercritical water environment. We then focused on exploring the potential symbiotic relationship between high-pressure and electrochemical loading driving forces and found that we could achieve higher loadings than what could be achieved with gas pressure alone under equivalent conditions. We also found that high-pressure electrochemical (HPE) loading could maintain near-unity loading conditions indefinitely but found no evidence of LENR. We found that the highly non-equilibrium nature of HPE would produce blisters within the Pd lattice, as well as heat-producing chemical reactions, but these observations were non-isotopic and chemical in nature.

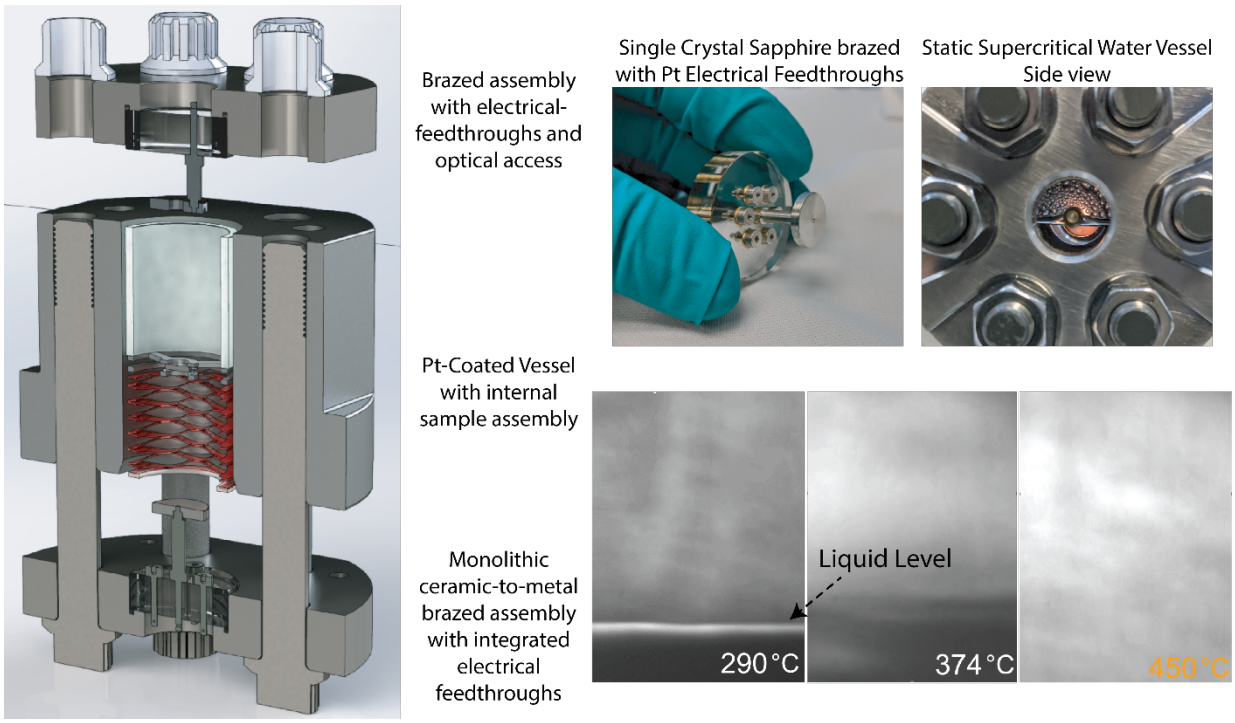


Figure 1: A high pressure, corrosion resistant vessel was used to test in supercritical water conditions. Because of the need to operate at pressures above 10 ksi and 374°C with corrosive liquids, material choice was limited. Seals and electrical and optical access were custom designed to meet these requirements. The three panels in the lower right are optical images of the water/electrolyte in the vessel at 3 different temperatures (i.e., below, at, and above the critical point).

Experimental Concept

The goal with high pressure electrochemical loading is to achieve high Pd loadings at elevated temperatures to test if this was a possible route to LENR. Since electrolytic loading doesn't work well at high temperatures, we aimed to explore if high-pressure may offer a way around this. To accomplish this goal, we proposed electrochemical loading under supercritical conditions. The supercritical regime is defined as beyond the critical point (> 374 °C and > 3200 psi for water) and is essentially a form of gas, that, above 374 °C, will never liquify. From electrolytic loading, the fugacity comes from an electrochemical potential, not ultra-pressure, and in a supercritical electrolyte.

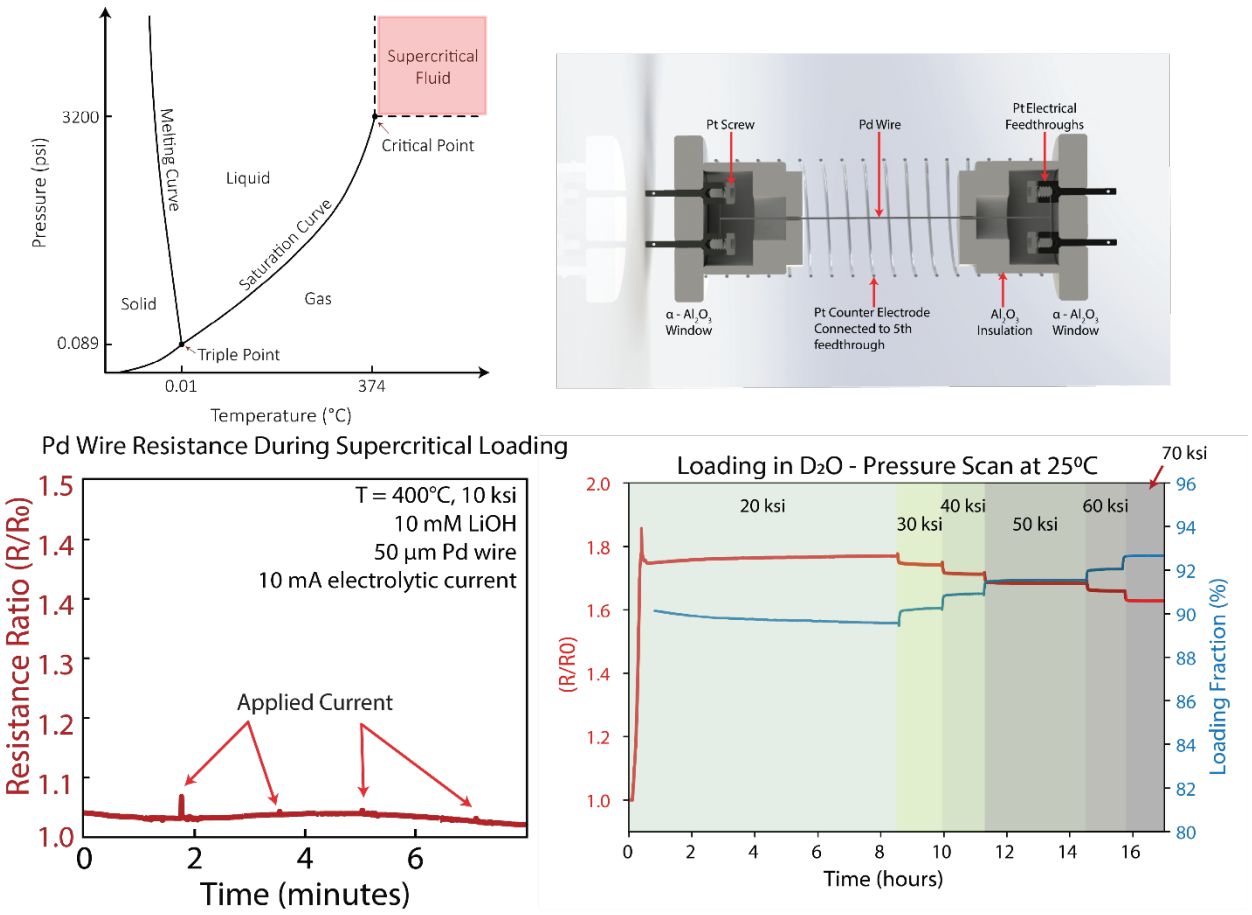


Figure 2: To operate in the supercritical regime with water-based electrolytes, temperatures >374°C and pressures >3200 psi are held in a corrosion-resistant pressure vessel. Using 4-wire resistivity measurements, we could not detect stable resistance changes that would suggest we were loading Pd. In contrast, when electrochemically loading Pd at increasing pressures, higher and higher Pd loadings are observed. These loading levels are higher than gas loadings at equivalent pressures.

Experiments were conducted in various settings, including high pressure vessels with water/electrolyte mixtures and intensifiers for generating pressure, similar to ones described in the Hardware section. These setups allowed for the exploration of ultra-high pressures and temperatures and their effects on electrochemical loading. To work in these high temperature, high pressure, corrosive environments, we had to make a few key hardware improvements, including creating high pressure seals compatible with high temperatures, preventing oxidation, and most importantly, corrosion mitigation. We found that supercritical water, especially when combined with electrolytes, became an extremely corrosive environment that required a lot of engineering to keep in control.

Figure 1 shows the high pressure, corrosion resistant vessel that was designed to operate in supercritical water conditions. Inconel 718 was the primary material used for the vessel body because of its high temperature stability and corrosion resistance, and all parts were coated with 50 nm Cr and 1000 nm of Pt to further reduce corrosion. Seals and electrical and optical access

were custom designed to meet these requirements. We brazed Pt electrical feedthroughs into single crystal sapphire windows with Au to maintain both optical and electrical access into the vessels. The use of Au, Pt and sapphire was critical to mitigate corrosion and survive the high temperature environment. Post-experiment analysis showed no significant contaminants on Pd sample surfaces.

The high-pressure vessel was filled with a water/electrolyte mixture in a N_2 glovebox (to remove the presence of oxygen) and sealed shut. The vessel is heated up to 450°C and the water level is monitored with a camera on the optical window. Sample images are shown in Figure 1, where the water line is clearly visible at temperatures <374°C and disappears as the temperature is increased above this point. Once temperature is dropped below the critical point, the water reliquefies, and the water line is optically visible.

In these environments, we tested 50-250 μm diameter Pd wires, spot welded to 4 Pt wires for resistivity measurements and in between a Pt counter electrode. Current densities ~ -50 - 150 mA/cm^2 between the Pd and Pt electrodes were applied, and we found that we could not observe stable resistance changes when loading in supercritical water. 0.1M H_2SO_4 and 0.1M LiOH/LiOD electrolytes were tested with identical results (several dozen tests done with varying concentrations and electrolytes, but these were the best candidates). A plot of an example measurement, wherein we measure the resistance of a palladium wire as we apply an electrochemical bias, is shown in Figure 2. We found that while we may observe small transients in the resistance of the Pd wire, possibly indicating small amounts of loading, they weren't stable and were within the noise of temperature fluctuations of a few degrees.

We confirmed that this was not an artifact of being in a supercritical environment, and we observed the same behavior in the liquid phase at high temperatures, between 250 to 350°C, confirming that this was a challenge with electrochemical loading at high temperatures. This left us to ask the question, at what temperatures can we sustain high loadings. Our hope was to find that with the assistance of pressure, we can overcome some of the difficulty in maintaining high electrolytic loadings at elevated temperatures.

HPE in the Liquid Phase

HPE : Up to 110 ksi with water electrolytes

This section is focused on understanding the effect pressure has on electrolytic loading at room temperature conditions. When performing these experiments, we quickly found that while we had significant experience with high-pressure gas systems, corrosion from electrolytes presented some new challenges. Primarily, corrosion from electrolytes would damage our pressure generating hardware. After a few uses, valves are rendered useless due to corrosion from electrolytes, and our high-pressure seals would wear out at a significantly faster rate. In addition, the hydrogen

generated from high-pressure electrolysis made hydrogen embrittlement a key limitation in material choice. Figure 3 is an image of an Inconel 718 high-pressure tube that ruptured due to hydrogen embrittlement. While we took advantage of copper superalloys (e.g., Toughmet

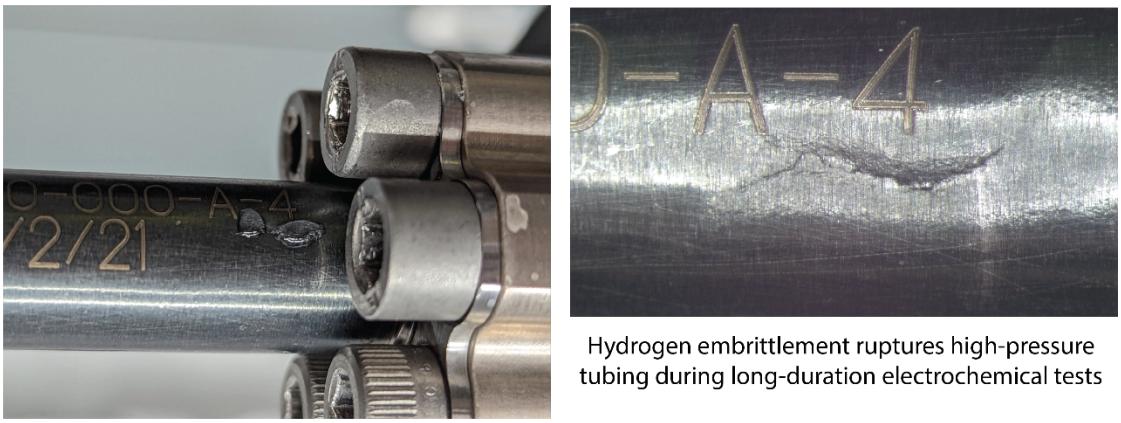


Figure 3: Inconel 718 is the material of choice for high pressure vessels and tubing to withstand corrosive liquids, but the hydrogen generated from long-duration electrochemical loading leads to hydrogen embrittlement.

TS160U) in our gas-based hardware to avoid hydrogen embrittlement, in electrolytic usage copper contamination would build up on Pd surfaces and inhibit Pd loading. We explored a variety of coatings and found that sputtered Pt helped mitigate corrosion, but no coating was robust enough to withstand these conditions for more than 2-3 tests.

To limit the effects of corrosion, we developed a 1-to-1 ratio fluid separator capable of translating pressure from our existing pressure-generating equipment without exposing our valves and intensifiers to corrosive electrolytes (Figure 4). Using spring-energized seals, a fluid separator piston is housed in a smooth (Ra < 4 μ–inch), Pt-coated Inconel 718 bore. As one side is pressurized, the spring-energized seals translate the piston, compressing the electrolyte volume and translating the pressure between the inert and electrolyte fluids without any exchange of liquid. This significantly decreases the surface area exposed to corrosive electrolytes, leading to cleaner electrochemical systems, and further abstracting away the difficulty with working with corrosive fluids. The fluid separator assisted in focusing our corrosion mitigation measures to just the high-pressure vessel and not any of the pressure generation hardware. Sputter coating a 1 μm thick Pt layer inside the high-pressure vessel proved to be sufficient to limit corrosion in HPE tests and resulted in clean working electrode surfaces following electrochemical loading, even after week-long experiments. There are no signs of Cu, Fe, or Zn which are usually common contaminants on electrolytically loaded samples, as seen in Figure 5.

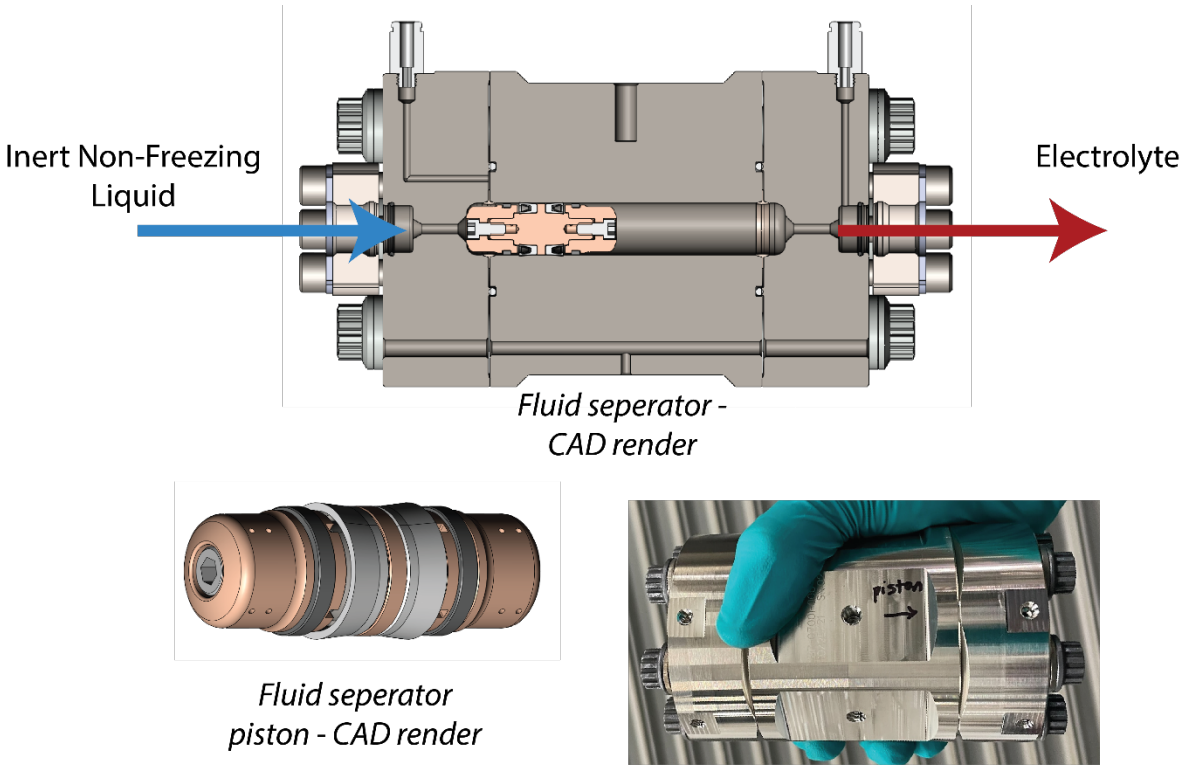


Figure 4: An ultra-high pressure fluid separator was designed to protect ultra-high pressure generation equipment from corrosive electrolytes, minimizing the amount of hardware that needed to meet strict corrosion-resistance requirements. A fluid separator piston is housed in a smooth (Ra < 4), Pt-coated Inconel 718 bore. As one side is pressurized, the spring-energized seals translate the piston, compressing the electrolyte volume and translating the pressure between the inert and electrolyte fluids without any exchange of liquid.

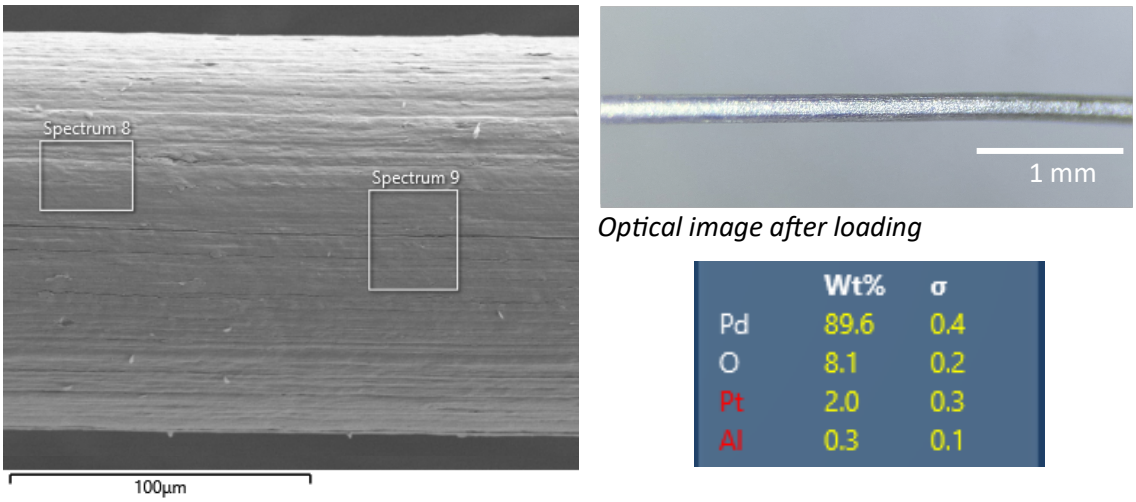


Figure 5: SEM/EDS analysis of 125 µm Pd wire loaded in high-pressure e-chem apparatus for over a week.

With the hardware improvements we made, we were able to test electrolytic loading at pressures

up to 110 ksi. Testing was done on bulk Pd wires and foils with no additives and 0.1 – 0.5M LiOD in D_2O electrolytes. Inside the pressure vessel, the resistance of Pd is measured through 4-wire resistance measurements as the proxy for loading, and a thin Platinum wire coiled around the Palladium cathode provides a measurement of the average temperature of the electrolyte near the Palladium cathode. External to the vessel, the temperature of all the key components is monitored up and downstream of the test vessel. In Figure 6 is a sample experimental HPE data stream from a light water test. We find that as we apply the electrochemical potential, as seen in the 2nd and 3rd plots, Pd begins to load as seen in the first subplot. After about an hour, Pd reaches a loading of 98-99%, and this loading remains extremely stable during the duration of the test. On the bottom two subplots, the temperature of a few components in the system is shown along with pressure data.

At pressures up to 110 ksi, we found that we could reach loadings of 99% and 97% with light and heavy water, respectively, without the use of loading enhancers. These results are fundamentally new - there are no loading enhancers used, and we reached loadings that are extremely rare in the electrochemical loading field, which previously has been nearly impossible without loading enhancers. If we turn off the electrochemical potential, the Pd sample immediately begins to deload, demonstrating that this is fundamentally different than typical gas phase experiments. Pressure alone is not leading to loading. More notably, these loadings are extremely stable and could be sustained for as long as the experiments are run, largely due to some of the work to control for contamination in the test systems, such as platinum coating all wetted materials past the fluid separator.

As a reference point, the loadings achieved at 110 ksi with HPE loading were equivalent to what is measured with 200 ksi of gas pressure loading. We propose two primary explanations for why an increase in electrolytic loading of bare Pd with increasing pressures is observed. The first is the reduction of bubbles with pressure. In normal electrolysis, bubbles are present at the electrode surfaces and are known to act as local deloading sites and create large concentration gradients. With high pressures, bubble size is significantly reduced, and the solubility of gas is increased, which further mitigates bubble formation. In addition, the local availability of hydrogen/deuterium at the Pd surface determines electrolytic loading. It's possible that the increased partial pressure of dissolved H_2/D_2 in the electrolyte has an additive effect on loading, or affects the relative reaction rates of Volmer-Tafel/Volmer-Heyrovsky processes.

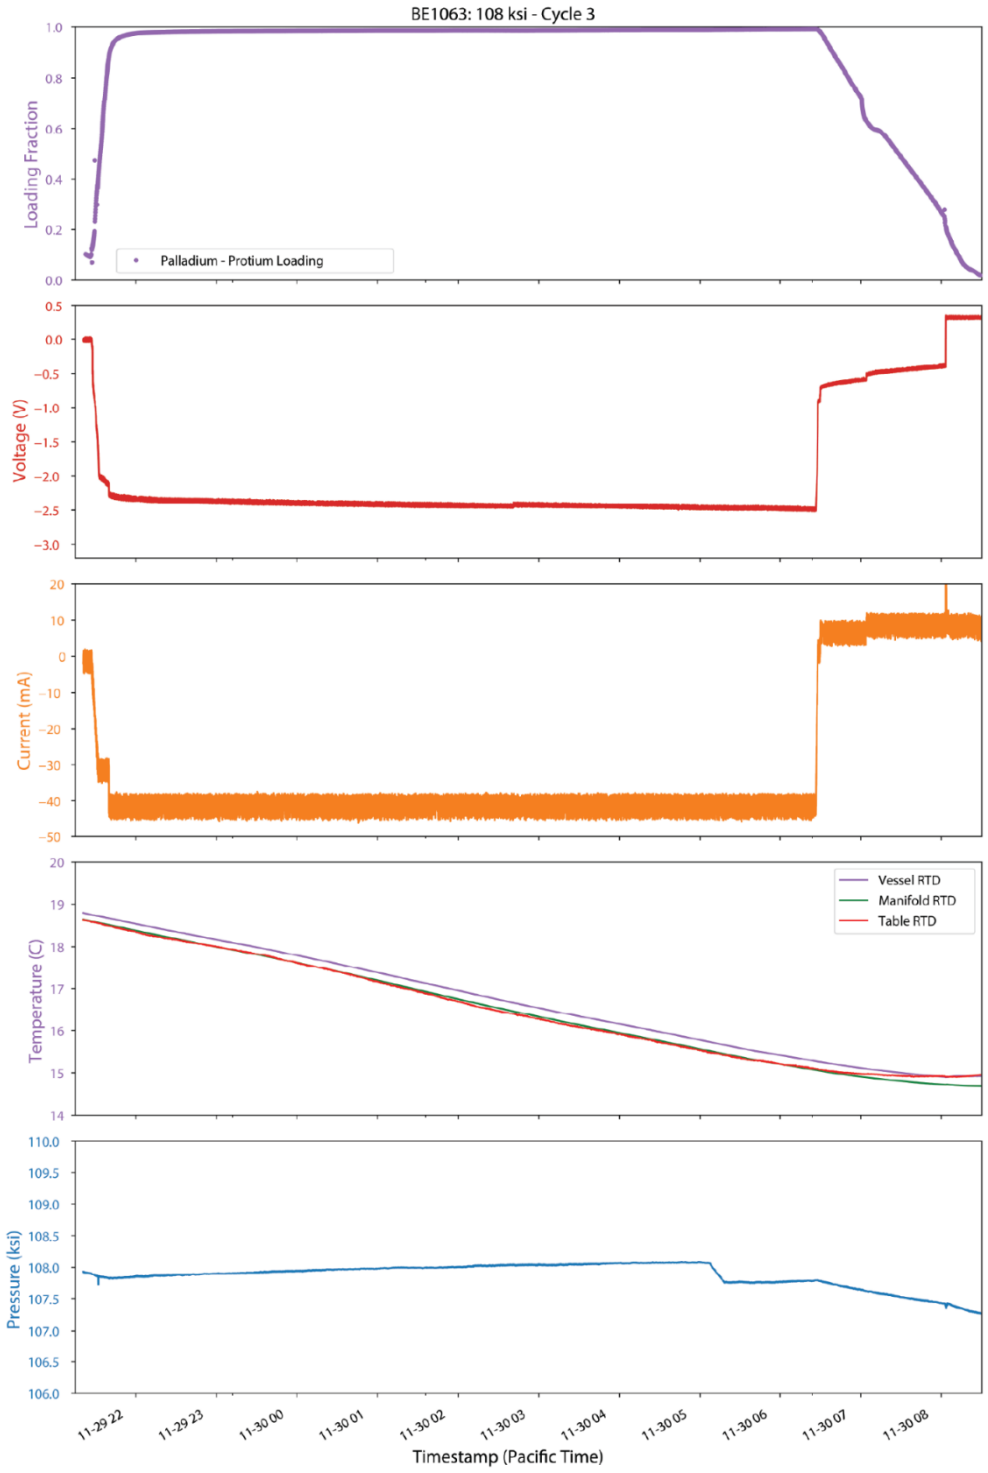


Figure 6 Data from a representative H/Pd electrochemical loading test. Loading is measured through 4-wire resistivity, and the electrochemical voltage and current are shown in the second and third subplots. Temperature is recorded in various points across the high-pressure electrochemical setup, along with pressure (fourth and fifth subplots). This is a sample loading cycle of a 250 µm Pd wire at >100 ksi. At high pressures, high loadings are sustained nearly indefinitely (days) with no measurable fluctuations.

HPE : Above 110 ksi with alternative electrolytes:

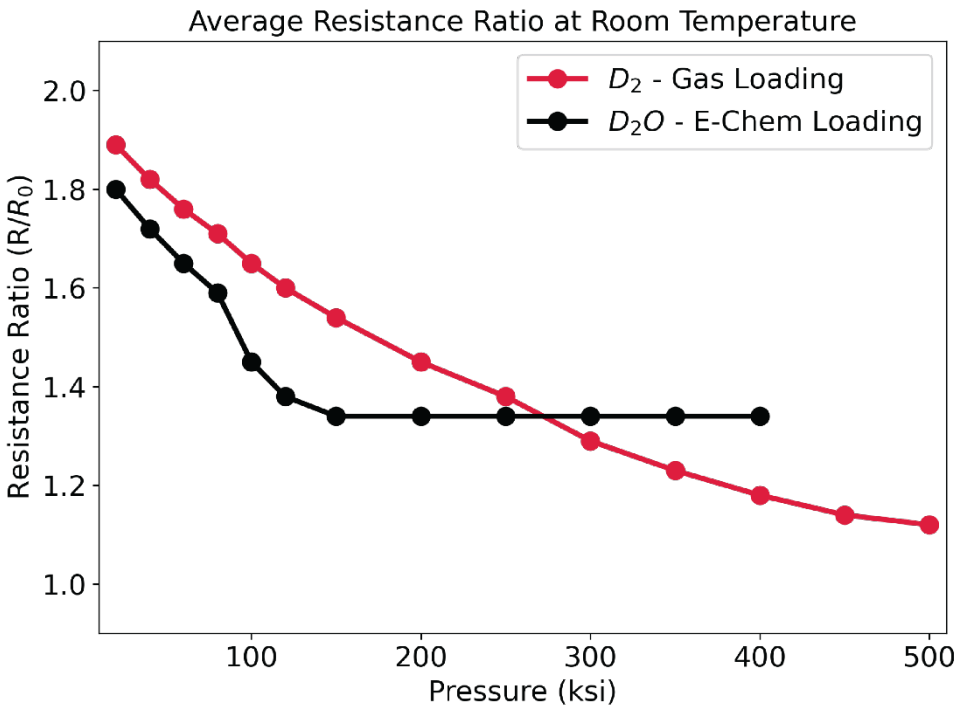


Figure 7 High pressure electrochemical loading summary. The average resistance ratio of all of our experiments in gas and HPE setups is plotted as a function of the maximum pressure tested. Resistance ratio is the primary metric used to estimate Pd loading fraction in these experiments.

As we pushed to higher pressures, we began to run into the high pressure freezing of water-based electrolytes at around 110 ksi. Water-based electrolytes undergo solidification/clathrate formation above 110 ksi at room temperature. Essentially, once we passed the freezing point of the water-based electrolyte, the loading would become extremely unstable, likely due to changes in the electrical properties of electrolytes during freezing. This effect takes attempting to push the pressure knob further with solid electrolytes alone off-the-table, and we needed to figure out a way to remain liquid. Pressure-induced freezing limits how much we can continue to push pressure using water-based electrolytes at room temperature. Initially, we tested increasing the temperature of LiOD in D_2O electrolytes up to 120°C to test if this would bypass any phase changes from occurring, but found that high-pressure solidification still occurred at temperatures up to these temperatures. This meant that temperature alone wasn't going to get us around high-pressure solidification.

Dilute methanol in water has been reported in literature to bypass freezing up to 300 ksi, and we were able to empirically confirm that this was viable up to 220 ksi. Past 220 ksi, 5-20% deuterated methanol (CD_3OD) in LiOD/D_2O electrolytes also solidified, likely due to phase separation of the water from the methanol and so acidic electrolytes had to be used. We tested with acidic electrolytes which are typically used for cryogenic loading, primarily, 1M deuterated hydrochloric acid (DCl) in deuterated methanol (CD_3OD) and 10% deuterated sulfuric acid (D_2SO_4, 96-98 wt. % in D_2O) in deuterated ethanol (CD_3CD_2OD). We found that these electrolytes were stable up to the maximum pressures we tested with HPE, 400 ksi.

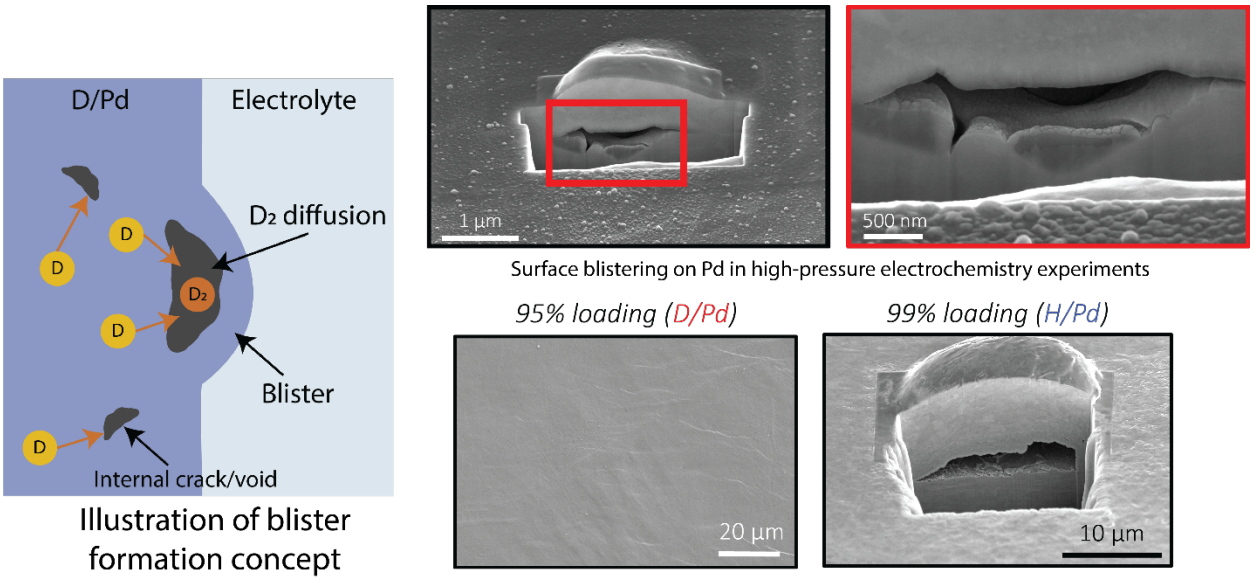


Figure 8 Blistering is observed on electrochemically loaded targets only when loadings >97% are reached. While their origin is unknown, blisters 1-30 µm in size are observed in both D/Pd and H/Pd, and no isotopic contrast is observed. We believe that the highly non-equilibrium nature of electrochemically loaded Pd is what leads to these significant blisters.

Figure 7 is a meta-summary of all HPE loading experiments in heavy water, compared to gas loading with D_2. Here, the measured resistance ratio as a function of the pressure is directly compared between gas and electrochemical loading experiments at identical pressures. We have not been able to find that HPE loaded bulk samples can get above 98% D/Pd loading with increasing pressures. Simply testing electrolytes that don't immediately solidify at 400 ksi didn't yield higher loadings than what was observed at 110-150 ksi. We believe that pressure-assisted electrochemical loading reaches a point of diminishing returns near 98% D/Pd loading. HPE can be very beneficial when pressure capabilities are limited, but at high enough pressures, there is little to no benefit to loading with HPE than with gas. For D/Pd, this cross-over happens at ~275 ksi.

HPE : Surface blistering in highly loaded HPE samples

One of the interesting features we consistently observed with HPE loaded samples is blister formation, particularly in experiments where Pd reached loadings above 97%. Post-experiment characterization of the palladium surfaces from our near-stoichiometric loading experiments revealed that the Pd surface would form microscopic blisters throughout the sample. These blisters vary in size, between 300 nm to 10 µm. Figure 8 is a SEM image of palladium foils tested in HPE. A slice of these blisters can be taken using focused-ion beam milling, where a thin protective platinum film is locally deposited on the area of interest, and local milling cuts the area in the middle of one of these blisters. On the top two images, we can see that these blisters are several hundred nanometers deep into the bulk, sometimes up to a micron, and we can confirm that there's a cavity behind where the blister is on the surface.

As a reference point, on the bottom right is an image of a palladium surface from a light water test that reached near-unity loadings and formed blisters, in comparison to an experiment that only reached 95% loadings. Through extensive cross-sectioning, we can isolate that these features are in close-proximity to some form of lattice distortion, many times from an impurity. There was a large amount of stochasticity in the presence of blistering and the density and size of these blistered surfaces, particularly with palladium hydride. We performed focused tests with our ultra-sensitive calorimeters to search for any meaningful evidence of heat production. Those results are discussed in more detail in the Calorimetry section. In short, no evidence of heat production was found for experiments that formed blistered Pd surfaces in comparison to those that didn't.

Table 1 is a summary of the experiments we ran with samples of different morphologies across both isotopes. In short, we observe blister formation in bulk-Pd of varying thicknesses and surface treatments, both polished and roughened. There are several reports in literature of blisters forming in metals from hydrogen and deuterium gas. These blisters are typically found in metals such as Be, Cu, Al when they're exposed to large fluxes of H_2 or D_2 plasma. It's proposed that these blisters form when protons or deuterons are introduced into the metal, in that case by irradiation with hydrogen or deuterium plasma, and segregate into defect sites where they recombine and form a blister, instead of diffusing to the surface. While HPE conditions are different than these reports, they both share similarities in that the Pd lattice is in a highly non-equilibrium state.

Table 1 Summary of completed HPE experiments at > 150 ksi.

Sample Type	Isotope	Number of Experiments	Peak Loading %	Blistering?
Wire	D_2O	1	98	1/1
Wire	H_2O	2	98/99	0/2
Polished Foil	D_2O	19	97/98	17/19
Polished Foil	H_2O	10	99	0/10
Roughened Foil	D_2O	1	98	1/1
Roughened Foil	H_2O	5	99	4/5
Single Crystal Slug	D_2O	2	R/R_0 is not reliable	0/2

HPE Loaded Thin Films

So far, we've discussed work done on bulk Pd. However, most of our AHE-diagnostics in Callisto are reliant on thin-film Pd to achieve the highest levels of sensitivity. While the resistivity of H/Pd and D/Pd as a function of loading is well-understood in bulk, thin-film measurements have never been done in literature to the best of our knowledge. In gas loading, we found that thin-film resistivity matched up with bulk Pd very well. Figure 9 plots the resistivity of thin films in gas as a function of pressure in comparison to bulk. While the exact magnitudes of resistivity are offset, they both follow the same trend, and a clear point of diminishing returns with increasing pressures is observed. We attribute minor variations in resistivity to the changes in geometry, as well as the role that constrained films might have, but overall, the trends are the same between bulk and thin films.

However, in HPE loaded thin films, we found that thin films of varying thicknesses become significantly more conductive as they load, reaching resistivity values well below their initial starting point. Additionally, the low resistivity values are reached at relatively moderate pressures as low as 70 ksi. To the right of Figure 9 is an example HPE experiment with a 20 nm Pd thin film, where pressure is increased from 70 ksi to 220 ksi and the resistivity is relatively unaffected. While this suggests that additional pressure has little effect on the loading of thin-films, it's unclear what loading the Pd thin-film is at. It's also worth noting that at ambient pressures, this does not happen. Resistivity stays above its starting value throughout the loading cycle. Table 2 is a summary of the lowest resistivity values of thin-films and at what pressures those are reached, where it's clear that increasing pressures don't enhance loading for HPE loaded thin films past 50 ksi. Regardless of what loading thin films are at, we don't observe any evidence of damage to the thin films that would suggest any relevance for LENR.

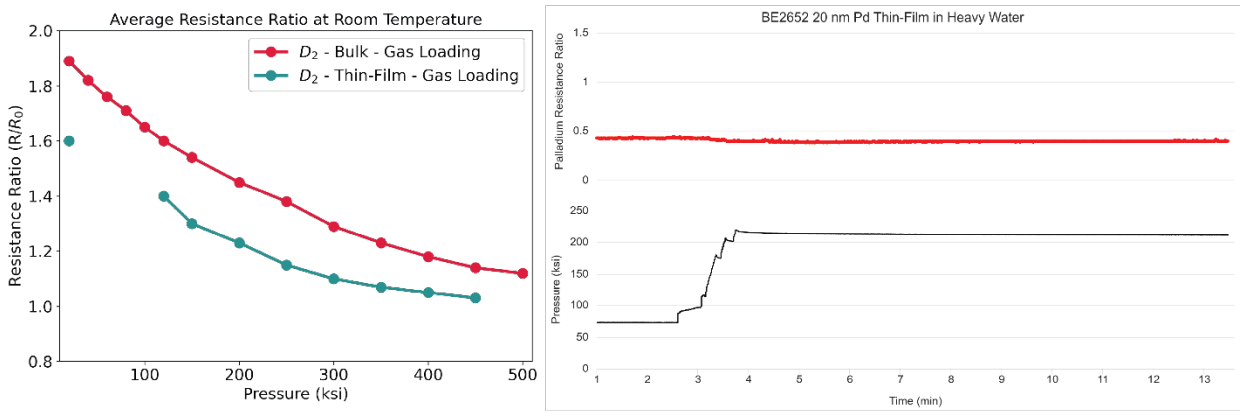


Figure 9 The resistivity of Pd thin films between 20-200 nm thick were measured with both gas and electrochemical loading as a function of pressure. Gas-loaded thin-films exhibited resistivities different in magnitude than gas-loaded bulk Pd, but the trends were consistent. HPE loaded thin films, however, displayed resistivity values significantly lower than bulk Pd. In addition, HPE loaded thin-films showed no signs of continued decreases in resistance past 80 ksi, while gas-loaded thin-films continue to show decreasing resistance at pressures up to 400 ksi.

Table 2 The range of resistivity ratios of thin films between 20-200 nm thickness are shown below under various testing conditions.

Resistivity Ratios of Thin-Films			
	Isotope	Min R/R_0	Pressures tested
Gas	H_2	0.88	400 ksi
HPE	H_2O	0.4	50-220 ksi
Gas	D_2	1.03	450 ksi
HPE	D_2O	0.4	50-400ksi

Conclusion

As far as we know, we're the only ones that have ever done electrochemistry at 10-20 kbar pressures, even outside of the context of loading Palladium. The occurrence of high energy events (discussed in more detail in the Calorimetry section) is sporadic, and we believe them to be chemical in origin, but we don't know mechanistically how they emerge. However, these are electrochemical phenomena and do not suggest any relevance for LENR. The potential symbiotic relationship between electrolytic and pressure loading warrants deeper exploration. Understanding this relationship could lead to more effective methods for achieving high loadings at elevated temperatures. It's unclear if there remains any benefit to performing bulk high pressure electrochemical loading if there is already access to pressures high enough to reach stoichiometry with gas, but electrochemistry may be its own form of stimulation. With electrochemically loaded thin films, we find that they become significantly more conductive than they began. This has never been reported on before and is potentially interesting, but we do not have reason to believe that this behavior is relevant for LENR.

Replications

Since Martin Fleischmann and Stanley Pons first reported anomalous heat generation in deuterated palladium electrolysis, most AHE experiments have built upon historical reports. While Callisto has sought to probe the less explored gas loading systems and utilize novel stimulation, we have attempted to replicate some of the previously published reports of AHE, primarily those for gas-loaded systems. Experiments were prioritized to be replicated if the work had commonalities with our D-D thesis regarding AHE, the publishing scientist was a particularly credible source, and the work had been reported to have been successfully reproduced by others.

Arata

Japan has a long history of research in the LENR field, particularly with a focus on gas phase loading and nano-structured palladium samples. Funding in Japan has come from both government and major industry and has attracted some reputable scientists. Yoshiaki Arata was known as an established physicist and leader in nuclear physics in Japan. His early work in LENR involved palladium nanoparticles and palladium nanocomposites in a thick-walled palladium vessel used as a cathode in electrolytic experiments which he termed the 'double structure cathode' or DS-cell. Later, this same approach was used in gas loading where he reported elevated temperatures on the interior of the DS-cell when the process gas was deuterium, but no excess heat was detected under protium[37]. In his final work he stopped using the DS-cell and focused on nano-composite samples in a custom apparatus. Throughout his entire body of LENR work, Arata et al. reported the detection of excess heat.

We elected to replicate two of Arata's experimental campaigns. The first used the DS-cell and a nano-composite sample composed of ZrO_2 and Pd atomic clusters (nano particles having D=<5nm). This experiment was intriguing due to our findings that Pd nanoparticles tend to agglomerate when loaded (Figure 1), and possibly lose the potentially unique properties of nanoparticles. However, when anchored to a support matrix of ZrO_2, they are prevented from agglomerating allowing the testing of loaded and isolated nanoparticles. The publications regarding these experiments provided details of numerous control experiments showing the effect only occurred with a specific sample and gas type.

[37] Arata, Y., & Zhang, Y.-C. (2005). Development of "DS-Reactor" as a practical reactor of cold fusion based on the "DS-cell" with "DS-Cathode". In Proceedings of the 12th International Conference on Cold Fusion (ICCF-12).

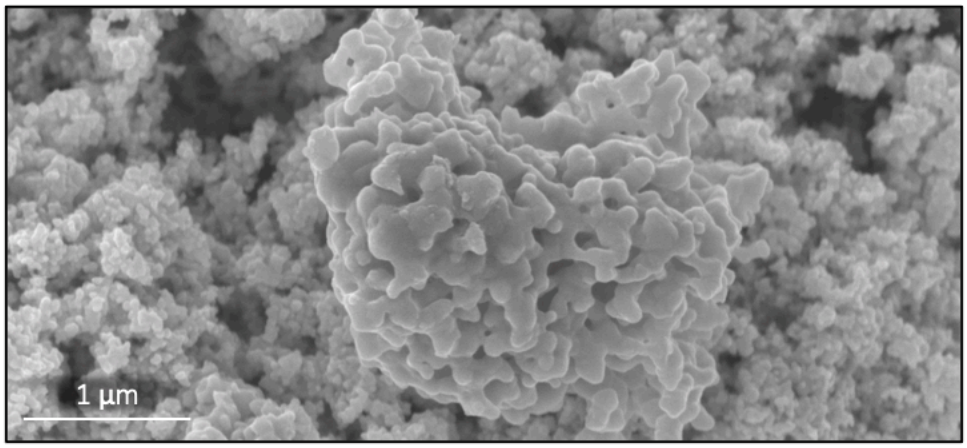


Fig 1. Callisto SEM image shows the agglomeration of palladium nanoparticles following exposure to H_2 gas.

The second of Arata's experiments replicated by Callisto focused on alternative nano-composite compositions in a more sophisticated heat detecting apparatus. We were able to successfully replicate some of the reported samples and match experimental conditions in state-of-the-art Setaram calorimeters.

Arata : Material Replication

The nanocomposites that Arata claimed were a key component of his success were comprised of a zirconium dioxide matrix with interspersed palladium atomic clusters, or sub 5nm nanoparticles. As outlined in his ICCF12 presentation and references therein, the nanocomposite was synthesized by first melt-spinning an alloy of 65at% Zr, 35at% Pd into thin, amorphous ribbons. The glassy alloy was then annealed in air to oxidize the zirconium and allow the palladium to precipitate into atomic clusters. For our replication (Figure 2), the alloy was first made by arc melting zirconium and palladium in proper ratios under an Ar atmosphere. The resulting alloy pellets were strip cast into amorphous ribbons at the Royce Institute at the University of Sheffield. This process involves inductively melting the alloy and injecting the molten alloy onto a liquid nitrogen cooled, spinning copper wheel. This process can quench the liquid at a rate of >1,000,000C/s preventing the formation of any crystalline structure during freezing. Due to the highly reactive nature of zirconium, we used custom quartz crucibles for this process which proved to maintain purity of the amorphous alloy. The calcination of the amorphous ribbons using the conditions provided by Arata et al. did not yield results matching his characterization, so many different annealing conditions were tested until characterization matched. Annealing for 10 hours at 500C in atmosphere yielded the closest results to the literature reports (Figures 3,4). {However, all of the other synthesized materials were also tested for AHE activity.} The resulting brittle ribbons were ground into a fine powder using a planetary ball mill at 600RPM for 30 minutes. XRD analysis

with phase identification and crystallite size was the primary method which we used for characterizing each batch of material. In the XRD analysis, the desired result would show complete segregation of the palladium and zirconium (no intermetallic phases identified), full oxidation of the zirconium, no oxidation of the palladium, and small average crystallite size for pure palladium phases. Annealing for 10 hours at 500C fulfilled nearly all of these requirements but did show oxidation of the palladium. The nano-composite was then reduced in a tube furnace for 30 minutes at 400C under flowing forming gas (5%H_2, balance Ar).

Synthesis of nanoPd·ZrO_2 composite

1. $Zr_{65}Pd_{35}$ is melted and rapidly quenched resulting in an amorphous alloy
2. The amorphous alloy is annealed in air allowing the Pd and Zr to phase-segregate. The Zr reacts with air to form ZrO_2 while the Pd precipitates into atomic clusters
3. The material in further annealed under reducing conditions to remove any oxide from the Pd

X-Ray Diffraction Analysis

Step	Identified phases	Avg crystallite size (nm)
1	Amorphous	N/A
2	ZrO_2 PdO	9.5 4.9
3	ZrO_2 Pd	7.6 7.8

Fig 2. Steps in synthesis process and XRD characterization after each step

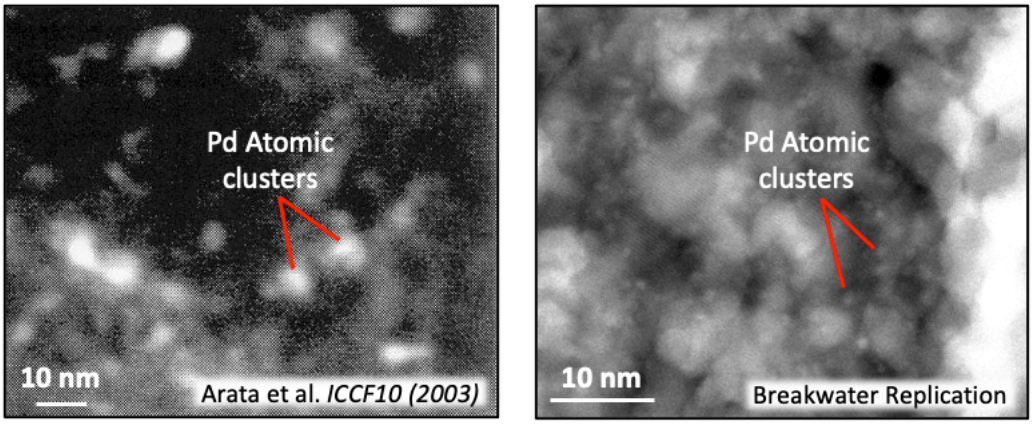


Fig 3. Comparison of transmission electron micrographs published by Arata et al[38]. and Callisto's replicated material.

[38] Arata, Y., & Zhang, Y. C. (2003). Development of compact nuclear fusion reactor using solid Pycnodeuterium as nuclear fuel. Paper presented at the 10th International Conference on Cold Fusion.

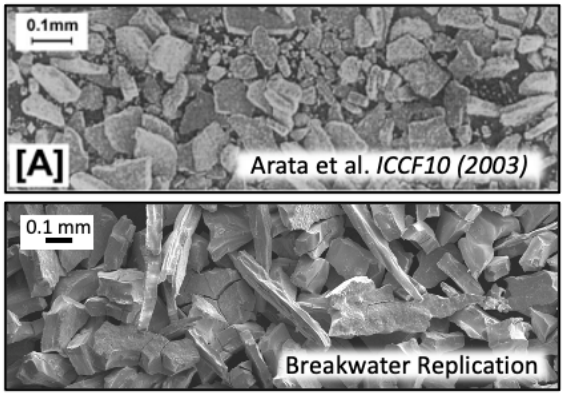


Fig 4. Comparison of scanning electron micrographs published by Arata et al. and Callisto's replicated material.

In addition to claiming nano-palladium facilitated AHE, Arata also reported that loading measurements on the nanocomposite showed 250% D/Pd loading at only 3atm and up to 400% D/Pd loadings near 100atm. Using a Quantachrome iSorb HP1, we measured the uptake of our replicated nanocomposites. When H_2 or D_2 uptake was normalized to the palladium content of the sample, the corresponding loading fraction very closely matched the expected loading for bulk palladium (Figure 5). It did not match the much higher loadings reported by Arata.

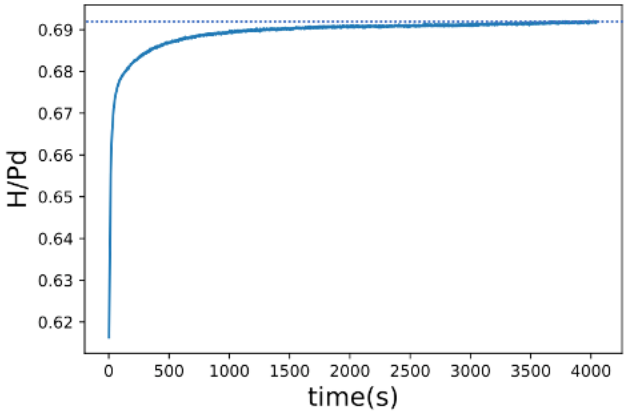


Fig 5. A hydrogen uptake measurement of replicated nanocomposite with <10nm palladium particles taken at 40C and 20bar. The uptake was normalized to the palladium content of the sample. Expected fractional loading for bulk palladium under these conditions is 0.71 H/Pd.

Arata : Experimental

Arata's DS-reactor was comprised of a 3mm wall thickness palladium pressure vessel containing sample material nested inside a stainless steel pressure vessel. Each vessel had isolated pressure

and electrical access. A heating mantle surrounded the outer pressure vessel and was used to heat the entire apparatus. The DS-cell had access to vent/vacuum, a pressure gauge(P_{in}), and a temperature sensor embedded in the sample material(T_{in}). The outer vessel was equipped with vent/vacuum, a pressure gauge(P_{out}), a temperature sensor(T_{out}), and a high pressure process gas inlet.

For Arata's experiments, the initial conditions were ambient temperature and vacuum in both the inner and outer vessels. The outer vessel was then pressurized to P_{out}=~40-80atm of process gas (H_2/D_2) and the heating jacket turned on targeting a temperature T_{out}=~140C. The process gas was then able to diffuse through the palladium DS-cell walls and into the sample. Because the heater was wrapped around the outer vessel and heat must conductively diffuse from the outer wall to the sample, it is to be expected that T_{out}> T_{in}. Arata reported this to be true for control experiments using H_2 process gas and for bulk palladium when using D_2 gas. However, when using palladium black or zirconia supported palladium nano particles with D_2, it was reported that the sample temperature was hotter than the temperature outside of the DS-cell, i.e., T_{in}> T_{out} (Figures 6,7).

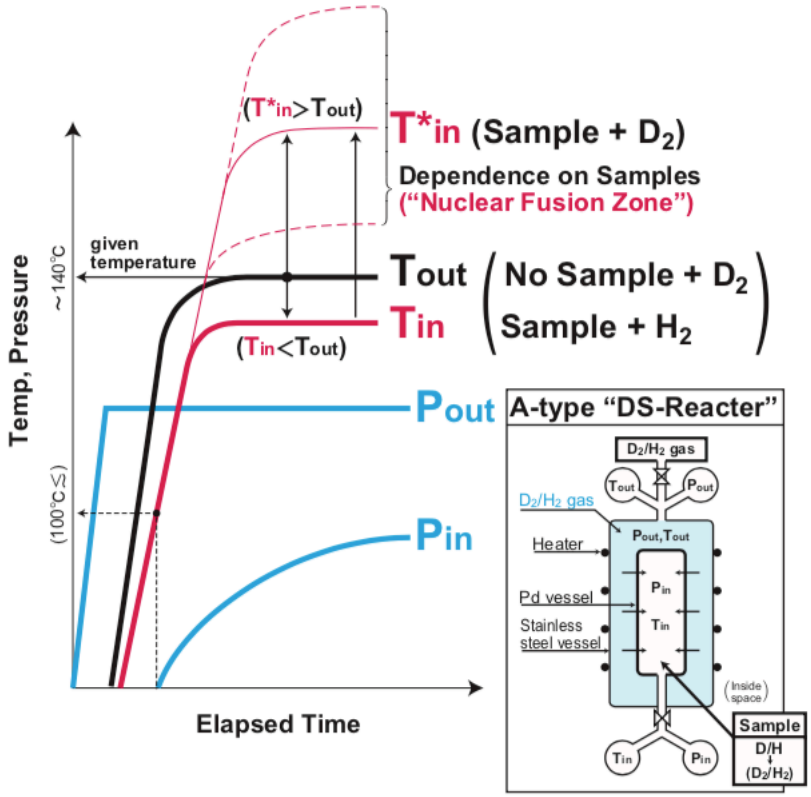


Fig 6. Apparatus and experiment description from "ICCF12: Yoshiaki Arata, M.J.A. Yue-Chang Zhang Center for Advanced Science and Innovation, Osaka University."

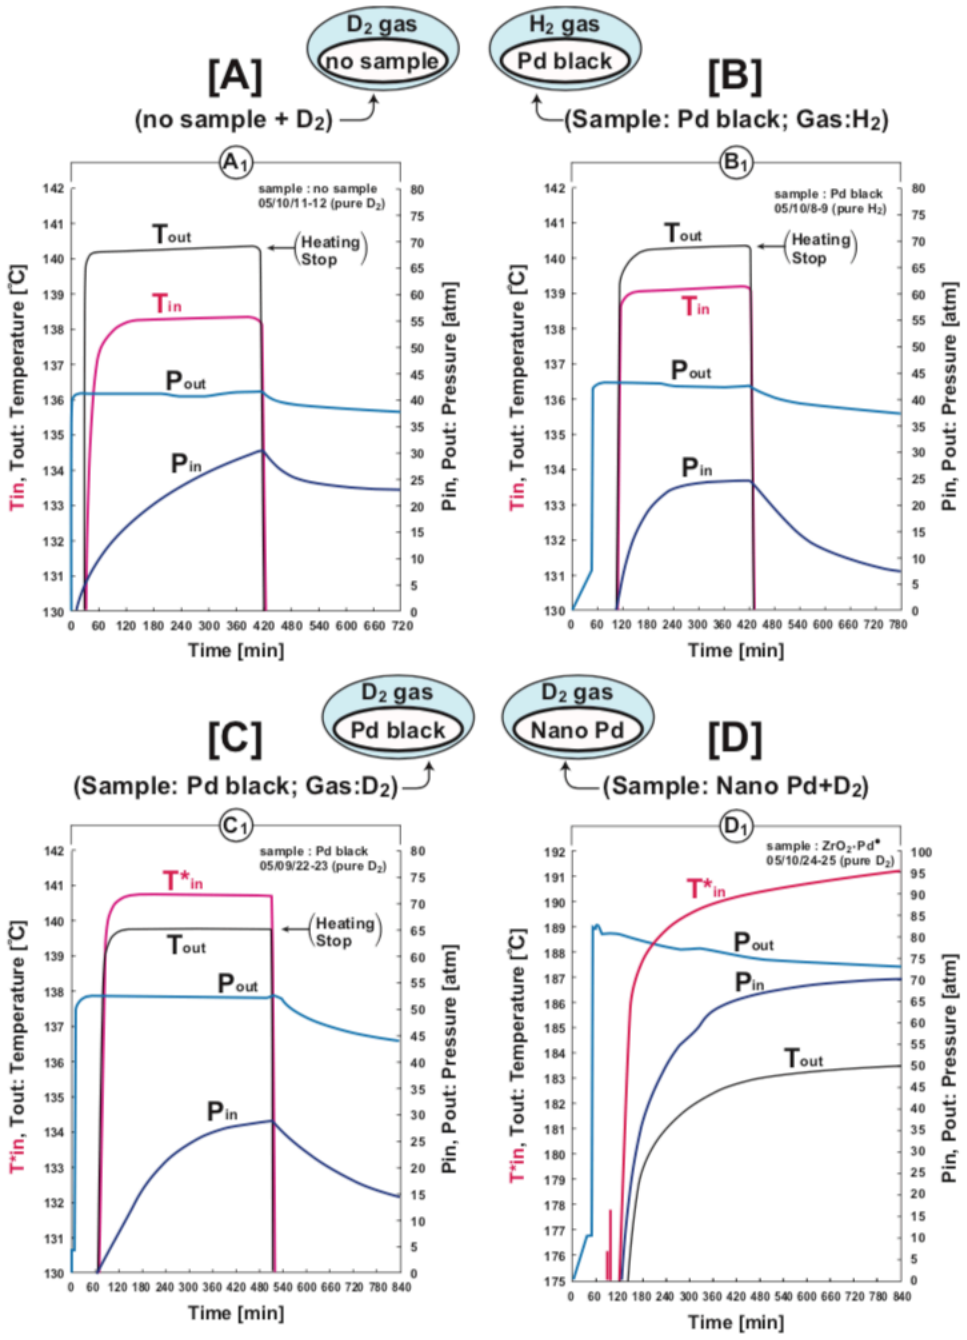


Fig 7. Experiment summaries from "ICCF12: Yoshiaki Arata, M.J.A. Yue-Chang Zhang Center for Advanced Science and Innovation, Osaka University."

Our replication apparatus (Figure 8) connected the outer stainless steel vessel to a manifolded pressure system with vent/vacuum, hydrogen, deuterium, and helium process gasses. A pressure transducer and electrical passthroughs were installed between the manifold and vessel using 125µm diameter insulated platinum wires connected to an RTD element in contact with the outer wall of the DS-cell to represent T_{out} and P_{out}. The DS-cell was machined from a palladium rod with

14mm outer diameter, 50mm length, to reach an 8mm diameter bore, with 47mm bore depth. The DS-cell was sealed by brazing the palladium to an Inconel cap/tube stub using Nioro (82Au 18Ni). The cap was connected to a 1/8in tube which was fed through the 1/4in inlet line to the stainless steel vessel. The inner tube contained electrical passthroughs connected to an RTD which was immersed into the sample material installed in the DS-cell.

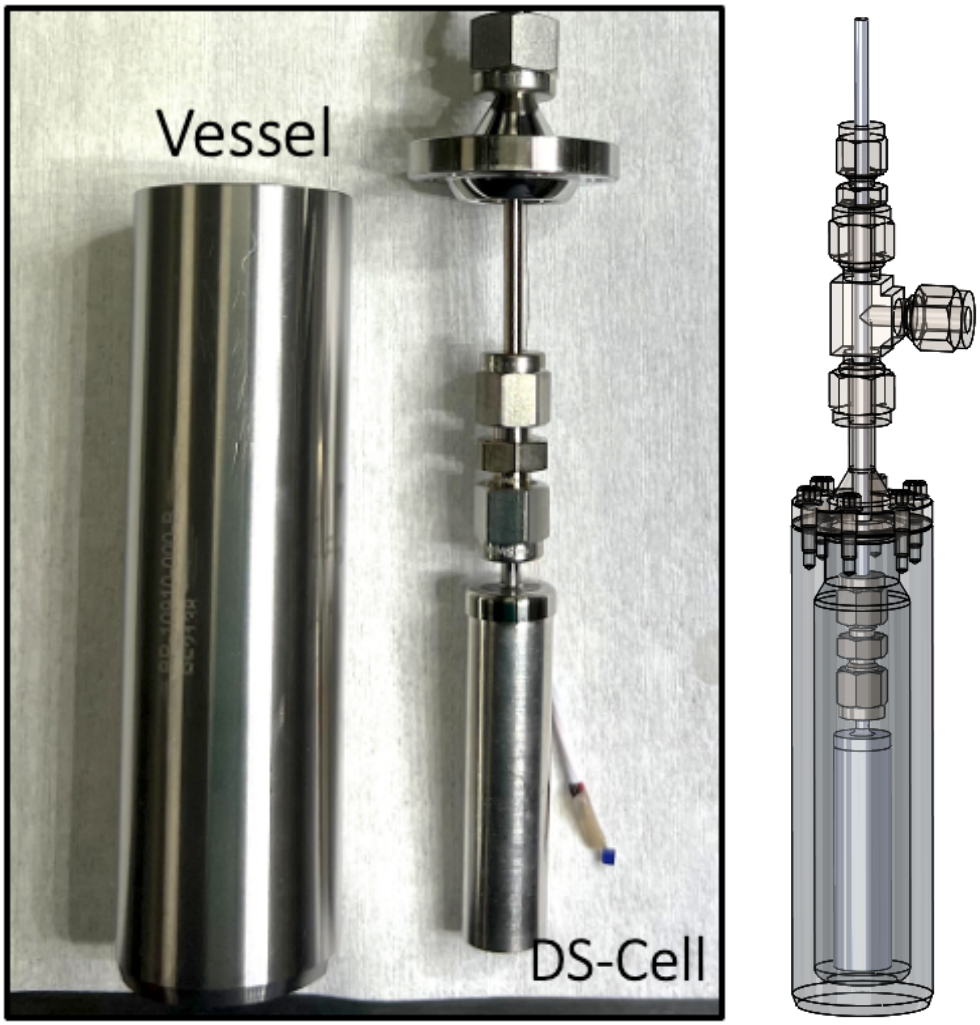


Fig 8. Callisto replication vessel and DS-cell

This apparatus was used to screen multiple batches of replicated Arata nano-composite samples and multiple variations of the procedures. In nearly all experiments, $T_{in} < T_{out}$, indicating a negative AHE result. In the set of experiments where $T_{in} > T_{out}$, it was determined to be caused by irreversible changes to the RTD upon pressurization as the offset persisted after the heater was turned off and the system returned to ambient conditions (Figure 9).

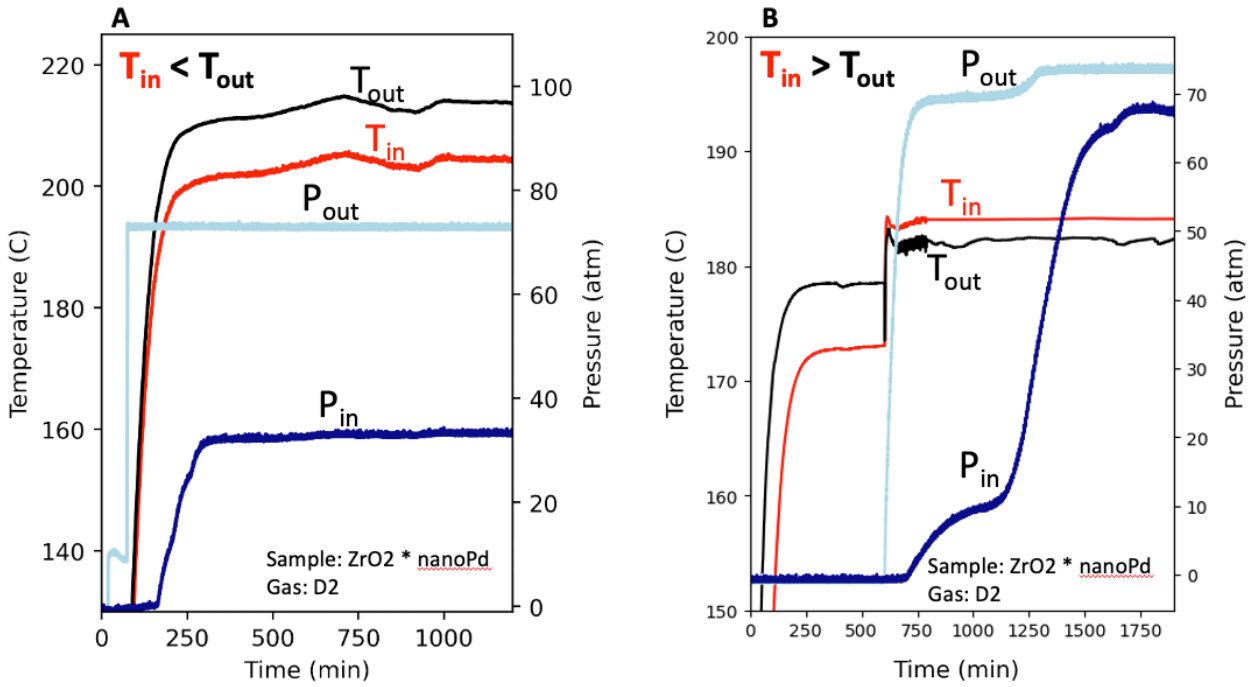


Fig 9. A) Callisto Experiment BE2606. The outer vessel was pressurized while simultaneously turning on a constant power heater. A leak was suspected but it was unclear why P_{in} did not equilibrate with P_{out}, however, many of Arata's published experiments show similar characteristics. As expected $T_{in}<T_{out}$. B) Callisto Experiment BE2758. The heater was commanded to a setpoint of 200C and both RTDs were allowed to equilibrate before pressure was introduced to the outer vessel. Upon pressurization, the reading of both RTDs increased with $T_{in}>T_{out}$. While this was unexpected, D_2 gas had not diffused into the sample at the time of the inversion and so the cause was not suspected to be AHE. Further testing revealed the pressure caused irreversible changes to the RTD.

While the temperature probes in different locations could theoretically indicate the generation of heat within the sample material, it is not an accurate or sensitive metric.

A more sophisticated approach is to use differential calorimetry where heat flow is directly measured between two nearly identical systems with one exposed to D_2 and the other exposed to H_2. This methodology, employing a C80 Setaram calorimeter, is further discussed in the Calorimetry section of this report.

For the Arata-replication calorimetry experiments, custom sample vessels were designed to maximize sample volume and minimize thermal mass. Early designs showed standard hardware was inconsistent at sealing at 175C and eventually, a mono-block vessel was designed which had all welded joints and VCR fittings (Figure 10).

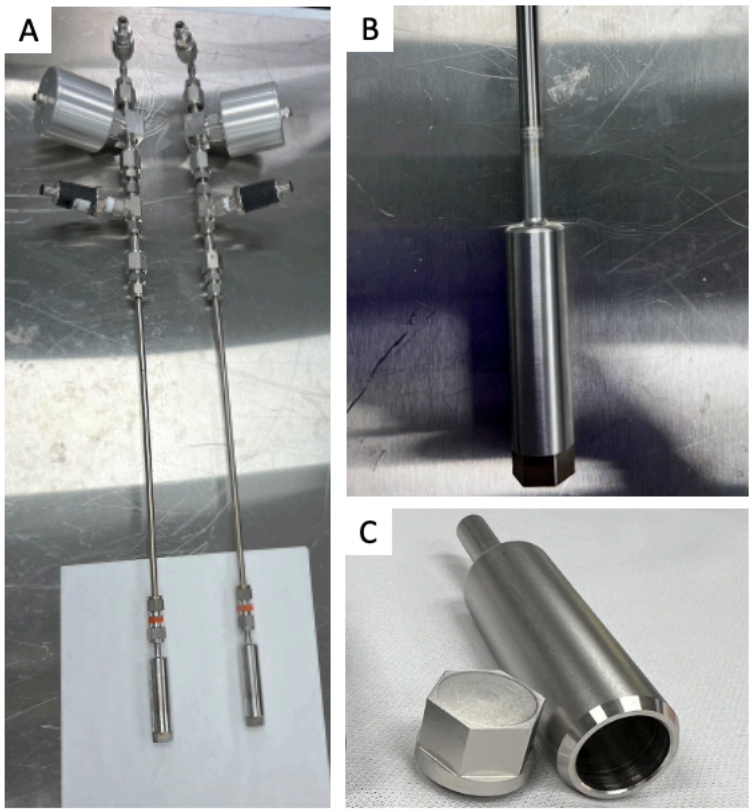


Fig 10. A) Early downtube assembly showing pressure inlet, isolation valves, pressure transducers, and filtered fittings to the vessels. B) Leaks were prevented by orbital tube welding the mono-block vessel to the downtube assembly. C) In order to achieve leak free mono-block vessels with large sample volume and small thermal mass, the vessels were cut open, precision hollowed out, and welded back together.

In these experiments, sample masses were matched within 1%, typically between 5-10g each, and exposed to 1000psi at 175C. Each experiment exposed the samples to multiple pressure/temperature cycles lasting 4-100 hours. During the active period of each cycle, pressures and temperatures were allowed to equilibrate before an isolation valve was closed on each channel. Between cycles, the samples were allowed to degas at 175C under vacuum. Figure 11 shows the results of a representative experiment, displaying no measurable excess heat.

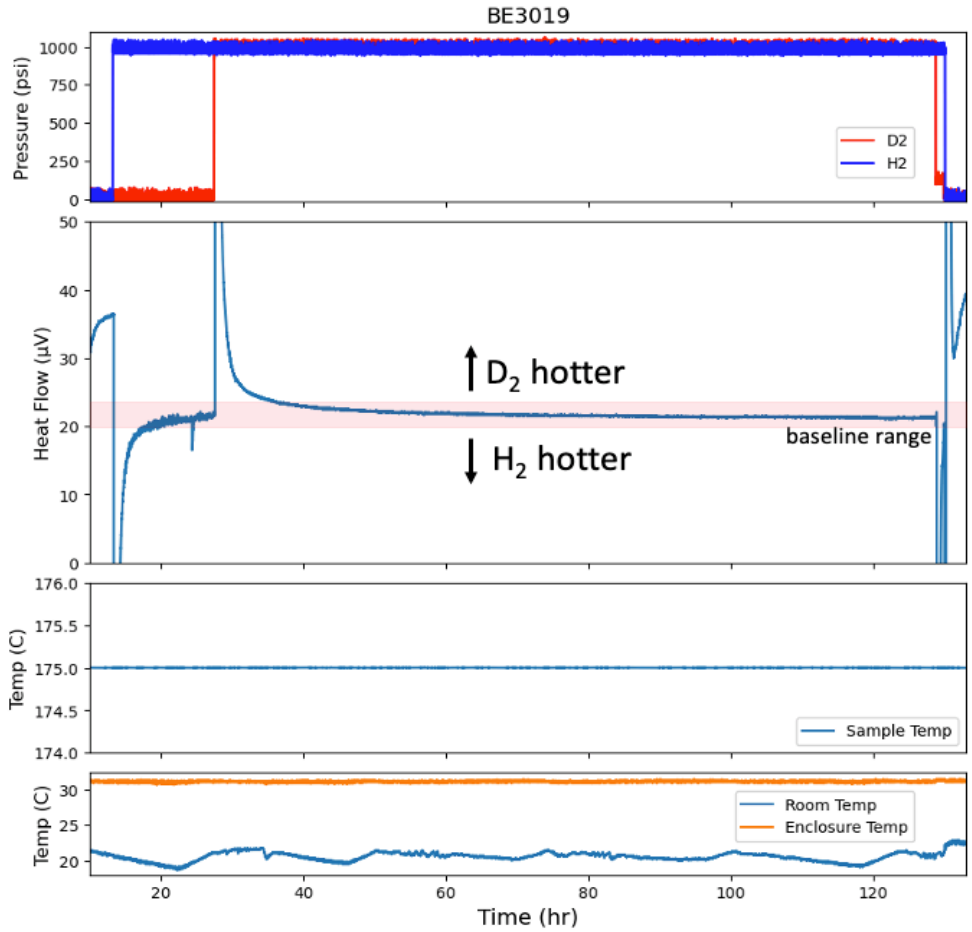


Fig 11. Callisto Experiment BE3019. Replicated ZrO_2 nano-composite is tested via differential calorimetry at 1000psi H_2/D_2 and 175C. Heat flow data does not show evidence of heat generation in either channel. The conversion from µV to milliwatts is not fully characterized but expected to be near 20-30µV/mW. The bottom plots show highly stable temperature of the calorimeter and in the calorimeter enclosure.

Following extensive testing of the Arata work, post-Arata work of the same type was explored in Japan. These experiments utilized similar nano-composite materials but replaced some of the palladium with varying ratios of nickel. The materials were exposed to H_2/D_2 pressures up to 1MPa and temperatures up to 300C in an improved apparatus which did not use the DS-cell. The updated apparatus contained six thermocouples and four RTD's which were calibrated to applied-heat using a heater in the sample chamber, and a mass flow controller to measure gas uptake by the sample. This body of work reported as much as 24watts of excess heat. The experimental apparatus was, in effect, a calorimeter and only required the development of samples for our replication effort. Of the fifteen samples listed, we elected to replicate the one with the highest reported excess heat. This nanocomposite was fabricated using the same basic procedure used in Arata's nano-palladium samples with a change to the starting alloy composition. Denoted 'PNZ7', the starting alloy contained 65at% zirconium, 30.625at% nickel, and 4.375at% palladium. The alloy was again made via arc melting and melt-spun into amorphous ribbons for us using the same procedure as

previously described at the University of Sheffield. XRD analysis of the material following the specified anneal (400C, 60hrs) revealed all nickel and palladium constituents were bound to Zr; analysis results are listed below (**Fig. 12**) from BE3834.

Sample ID	Phases Identified	Average Crystallite Size (nm)
BE3834	Zr_2Ni – Nickel Zirconium Tetragonal, S.G.: I4/mcm (140) PDF# [04-003-7047]	19.7 ± 3.7
	ZrO_2 – Baddeleyite Monoclinic, S.G.: P21/c (14) PDF# [04-002-5426]	16.1 ± 1.7
	Zr_3Pd_4 – Palladium Zirconium Hexagonal, S.G.: R-3 (148) PDF# [04-003-1186]	9.8
	ZrO_2 – Tazheranite Cubic, S.G.: Fm-3m (225) PDF# [04-002-8314]	6.6

Fig 12. XRD analysis of PNZ7-1. Amorphous ribbons were annealed in air at 450C for 60 hours. The analysis failed to identify any phases of Ni or Pd which were not bound to Zr. Zr was not fully oxidized.

The temperature and calcination time were adjusted for a series of subsequent batches prepared with the intention of fully oxidizing and segregating the zirconium from the palladium and nickel (Figures 13,14).

Sample ID	Phases Identified	Concentration (± 5 wt%)	Average Crystallite Size (nm)
BE4066	ZrO_2 – Baddeleyite Monoclinic, S.G.: P21/c (14) PDF# [04-002-5420]	65.2%	14.5 +/- 0.2
	NiO – Bunsenite Cubic, S.G.: Fm-3m (225) PDF# [04-006-6925]	13.6%	10.8 +/- 0.2
	Zr_2Ni – Nickel Zirconium Tetragonal, S.G.: I4/mcm (140) PDF# [04-003-7047]	10.1%	29.6 +/- 0.5
	ZrO_2 – Tazheranite Cubic, S.G.: Fm-3m (225) PDF# [04-002-8314]	9.9%	9.2 +/- 0.2
	Zr_2Pd – Palladium Zirconium **??** Tetragonal, S.G.: I4/mmm (139) PDF# [04-007-7762]	1.3%	8.0 +/- 0.5

Fig 13. XRD semi-quant analysis of PNZ7-3. Amorphous ribbons were annealed in air at 600C for 15 hours. A phase of nickel was identified not bound to Zr. Annealing parameters still appear to be insufficient for full Zr oxidation and segregation.

The loading of each material batch was measured at 40C and 300C under 20bar H_2 (Figures 15,16). The iSorb HP1 carefully calibrates the void volume of the sample vessel and measures the pressure lost when the vessel is dosed with H_2 gas to report moles of gas lost to sample absorption. Using the assumption that no gas is absorbed by the Zr/ZrO_2, the loading fraction is calculated based on the nickel + palladium content. Because the degree of oxidation within the sample is not well known, XRF measurements were used to measure the mole fraction of nickel and palladium in the nanocomposite sample (Figure 17).

Element	Conc.		Element	Conc.
O	17.4		Ni	17.4
Al	0.10		Cu	0.037
Si	0.12		Zn	0.007
Ca	0.027		Zr	58.5
Cr	0.085		Pd	8.12
Mn	0.024		W	1.05
Fe	0.24			

Fig 14. XRF analysis of PNZ7-2. Amorphous ribbons were calcined in air at 500C for 10 hours. This annealing condition was taken from the process used to match Arata's nanocomposite. Concentrations are reported as wt%.

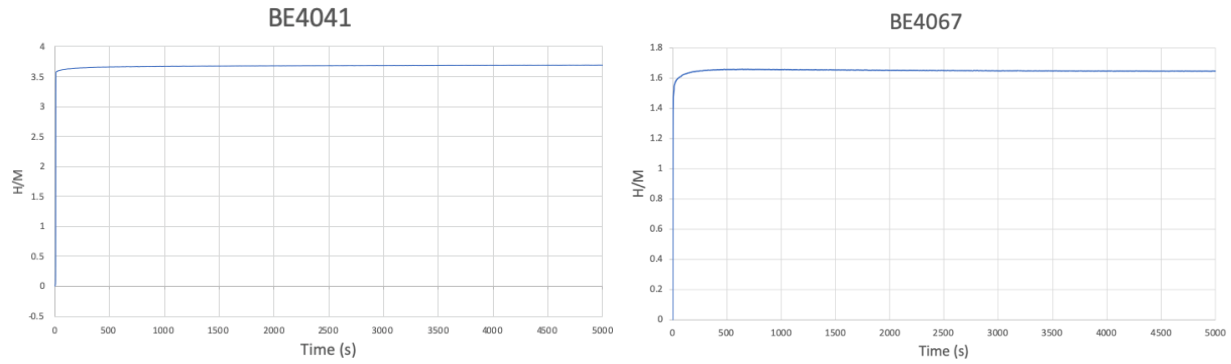


Figure 15. BE4041) 40C 20bar H_2 loading measurement of PNZ7-2. Measured uptake of 3.69H/M, M=Ni+Pd. Kitamura et al. reported 3.4H/M for room temperature D_2 loading at <1MPa. BE4067) 300C 20bar H_2 loading measurement of PNZ7 synthesis trial 2 shows measured uptake of 1.65 H/M compared to reported 1.3 D/M at '>250C' and <1MPa.

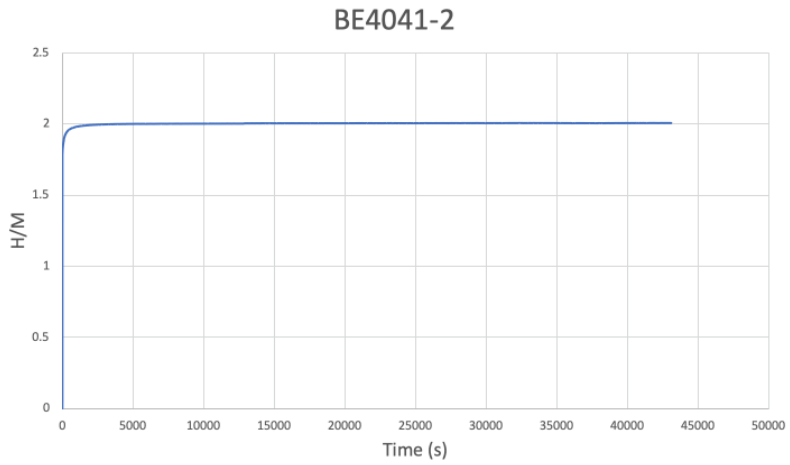


Fig 16. BE4041-2) A second 40C 20bar H_2 loading measurement of PNZ7--2 was performed following BE4041 and BE4067. This result showed a substantial decrease in measured loading indicating a possible irreversible mechanism for pressure loss during measurement such as reduction of NiO species. H_2 absorption by Pd and Ni is expected to be reversible and should not vary with cycle number.

Sample ID	Phases Identified	Concentration (± 5 wt%)	Average Crystallite Size (nm)
BE3908	ZrO_2 – Baddeleyite Monoclinic, S.G.: P21/c (14) PDF# [04-002-5420]	47.5	17.6 ± 7.0
	Zr_2Ni – Nickel Zirconium Tetragonal, S.G.: I4/mcm (140) PDF# [04-003-7047]	45.0	11.8 ± 5.6
	NiO – Bunsenite** Cubic, S.G.: Fm-3m (225) PDF# [04-006-6925]	6.9	11.6
	ZrO_2 – Tazheranite** Cubic, S.G.: Fm-3m (225) PDF# [04-002-8314]	0.5	12.5
	Zr_3Pd_4 – Palladium Zirconium** Hexagonal, S.G.: R-3 (148) PDF# [04-003-1186]	0.1	6.6

Fig 17. XRD analysis of PNZ7-2 shows the presence of reduceable species in significant concentrations.

AHE tests were performed in a Setaram C80 calorimeter by loading each calorimeter vessel with identical samples and exposing them to ~100psi H_2/D_2 at 275-300C. Shorter 'conditioning'

pressure cycles were used as described by Kitamura et al.[39] followed by long duration isolated holds under static pressure/temperature where heat flow was carefully observed at equilibrium conditions. Figure 18 shows an example of such an AHE test.

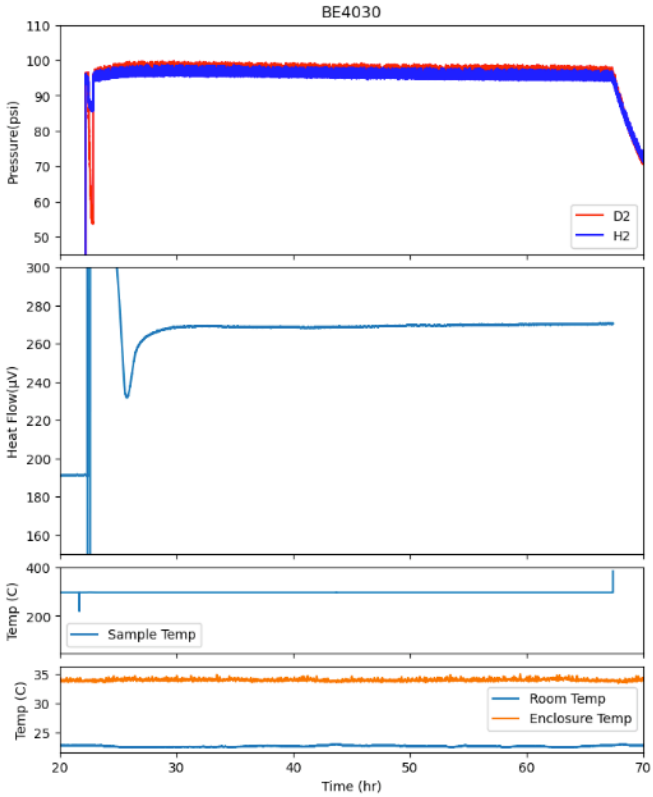


Fig 18. Callisto calorimetry experiment BE4030. PNZ7 synthesis trial 2 material was tested at 300C and 100psi H_2/D_2 via differential calorimetry for AHE activity. No indications of AHE were detected over the long hold. Note that baseline heat flow is unique for each system build and offsets from 0μV are not an indication of heat generation in either channel.

Arata : Conclusion

In total, 15 AHE tests were carried out replicating the Arata and post-Arata body of work (Table 1). While some of these experiments included slight pressure leaks, temperature sensor drift, or other apparatus related anomalies, none of our testing showed any indication of AHE.

[39] Kitamura, A., Takahashi, A., Takahashi, K., Seto, R., Hatano, T., Iwamura, Y., Itoh, T., Kasagi, J., Nakamura, M., Uchimura, M., Takahashi, H., Sumitomo, S., Hioki, T., Motohiro, T., Furuyama, Y., Kishida, M., & Matsune, H. (2018). Excess heat evolution from nanocomposite samples under exposure to hydrogen isotope gases. International Journal of Hydrogen Energy, 43, 16187–16200. https://doi.org/10.1016/j.ijhydene.2018.06.187

Table 1: Arata replication experiments

Exp ID	Apparatus	Process pressure (psi)	Process temperature (degC)	Sample type	AHE Detected?
BE1578	**Case style apparatus**	**50**	**240**	**Arata nanocomposite**	**No**
BE1639	**Case style apparatus**	**930**	**240**	**Arata nanocomposite**	**No**
BE1673	**Case style apparatus**	**930**	**170**	**Arata nanocomposite**	**No**
BE1766	**Case style apparatus**	**1085**	**200**	**Arata nanocomposite**	**No**
BE1878	**C80**	**1085**	**180**	**Arata nanocomposite**	**No**
BE1928	**C80**	**1000**	**175**	**Arata nanocomposite**	**No**
BE2023	**Arata DS-cell**	**72**	**230**	**Arata nanocomposite**	**No**
BE2173	**C80**	**850**	**175**	**Arata nanocomposite**	**No**
BE2261	**Arata DS-cell**	**1085**	**210**	**Arata nanocomposite**	**No**
BE2606	**Arata DS-cell**	**1085**	**180**	**Arata nanocomposite**	**No**
BE2758	**Arata DS-cell**	**1085**	**230**	**Arata nanocomposite**	**No**
BE3019	**C80**	**1000**	**175**	**Arata nanocomposite**	**No**
BE3957	**C80**	**100**	**300**	**PNZ7-1**	**No**
BE4030	**C80**	**100**	**300**	**PNZ7-2**	**No**
BE4126	**C80**	**100**	**275**	**PNZ7-3**	**No**

Case

Dr. Les Case was an MIT trained chemical engineer prior to pursuing LENR research. His work focused on very simple methods to achieve the anomalous heat effect using commercially available hydrogenation catalysts. In addition to the straightforward activation and detection of the effect, his experiments were reportedly successfully reproduced at Stanford Research Institute by Mike

McKubre and Russ George[40,41] giving sufficient credibility to the work to warrant testing. The Case experiments exposed hydrogenation catalyst (typically palladium particles on activated carbon) to D_2 gas at elevated temperatures and moderate pressures using constant power heating. The equilibrium sample temperature using D_2 was reportedly higher by 5-20°C relative to that involving similar H_2 exposure. A common theme between this work and the experiments by Arata et al. were the use of nano particles and a support structure to prevent agglomeration upon loading. Both Les Case and Russ George presented on this work citing a $D + D \rightarrow He^4 + 24$ MeV reaction.

In our replication of the Case experiments, we employed a number of commercial catalyst materials matching the descriptions stated by Case and exposed them to the same processes. This included moderate H_2/D_2 pressures and high temperatures. In addition to a setup designed to mimic the apparatus used by Case where temperature sensors are used to monitor the sample temperature throughout testing, the sample materials were each subjected to the same conditions and processes in a high sensitivity differential calorimeter.

Case : Materials

Case reported at ICCF7[42] that AHE was detected in his apparatus. He investigated using many different platinum group metal catalysts supported on activated carbon including about 0.4-1% palladium, platinum, iridium, and rhodium, with palladium being by far the most active and other substrates appearing inactive. While extensive characterization of these materials was not presented, the preferred material used by Case was G-75E (0.5%Pd on activated carbon) and G-75D (0.4%Pd on activated carbon) sourced from United Catalysts, Inc. Unfortunately, at the time of this replication work, this company was no longer operating. Instead, 5 candidate materials were selected with varying forms of activated carbon and palladium content (Figures 19,20).

[40] Mallove, G. (1999, July). Progress in Les Case's Catalytic Fusion. *Infinite Energy Magazine*, 4(23).

[41] George, R. (1999, March 26). Production of 4He from deuterium during contact with nano-particle palladium on carbon at 200° C and 3 atmosphere deuterium pressure. Paper presented at the American Physical Society Centennial Conference.

[42] Case, L.C. (1998). Catalytic Fusion of Deuterium into Helium-4. In *Proceedings of The Seventh International Conference on Cold Fusion*. Vancouver, Canada: ENECO Inc.

Callisto Material	Palladium Content	Activated Carbon Source	Vendor
BM524	1%	Egg shells, unreduced	Fisher Scientific
BM536	5%	Unreduced	Fisher Scientific
BM537	1%	Coconut shells, reduced	Cocoworks
BM538	0.5%	Coconut shells, reduced	Cocoworks
BM891	5%	Unreduced	Acros Organics

Figure 19. Case replication materials list.

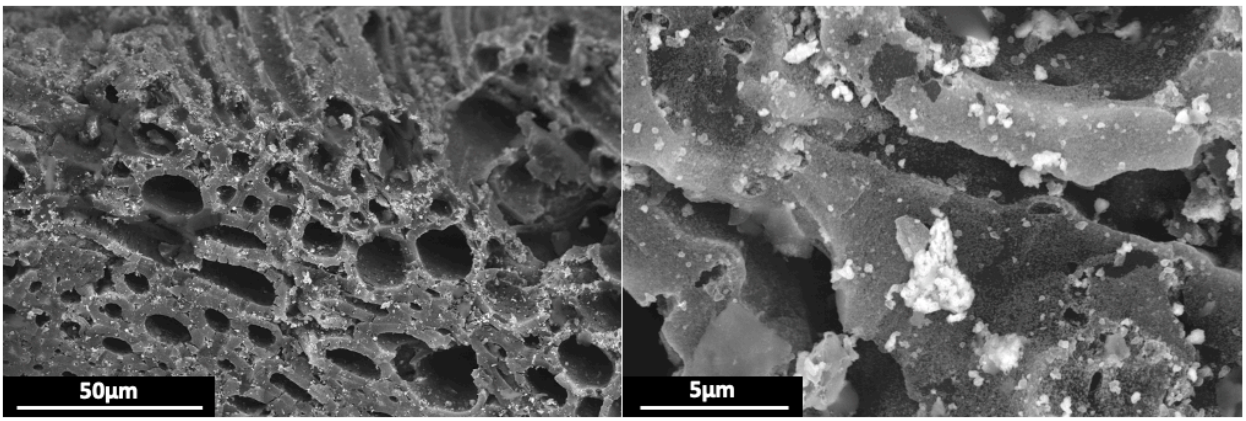


Figure 20. SEM images of BM537.

Case : Experimental

As outlined in the writings of Case from ICCF-7, Russ George via Catalyst Helium online, and interviews of Les Case by Gene Mallove published in Infinite Energy magazine, the protocol for generating excess heat was as follows; first, the candidate sample material is loaded into an oval stainless steel pressure reactor (1.6L) which is nested in a constant power heating jacket. The vessel is pumped down to vacuum and hydrogen gas is introduced to approximately 50psi. The heating jacket is then turned on targeting an equilibrium temperature of approximately 150-200C and the

thermocouple on the inside of the reactor is monitored until a stable temperature is reached. Once this temperature is recorded, the heater is turned off and the gas vented. This process is repeated until a consistent temperature is measured between cycles. These 'conditioning' cycles were reportedly to take any chemical reactions to completion and establish a baseline temperature. The following cycle would introduce deuterium gas in place of hydrogen and reportedly, for AHE active materials, the resulting equilibrium temperature would be higher than the baseline temperature by 5-20C. Case also performed mass spectroscopy on the resultant gas from experiments where AHE was detected which found above ambient levels of helium. Helium detection as a diagnostic for AHE, however, is difficult to perform properly, and the presence of background helium can easily lead to false positives.

In Callisto's replication experiments, a stainless-steel vessel of similar geometry and volume was fitted with pressure and temperature sensors and installed into a heating jacket connected to a variac for adjustable constant power output (Figure 21).

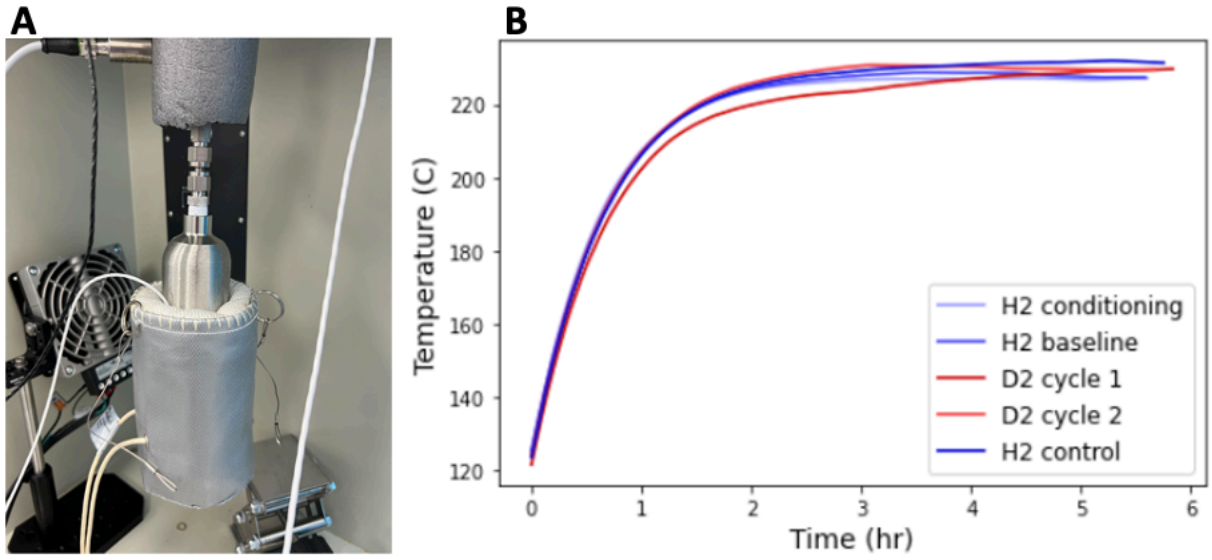


Figure 21. A) Case replication apparatus containing pressure and temperature sensors, a temperature-controlled enclosure, and constant power heating jacket. B) Callisto replication experiment BE1530. 42g of 1% palladium on activated carbon is exposed to repeated cycles of H_2 and D_2 gas. After pressurizing, the heater is turned on and the system is allowed to reach an equilibrium temperature. The plot shows the temperature profiles for each cycle. No significant difference between the H_2 and D_2 cycles was observed.

As with the Arata experiments, the experiments were ultimately moved from the replication apparatus to the Setaram C80 calorimeter to directly test for the generation of heat. Representative experimental results are shown in Figures 22, 23.

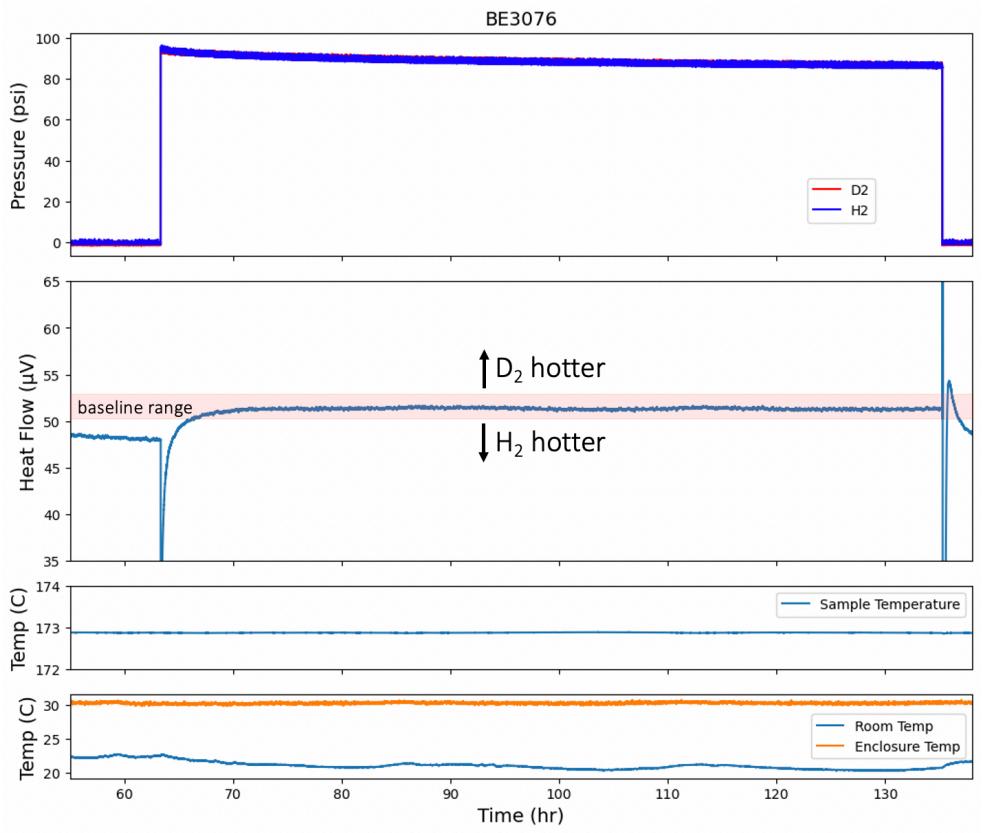


Fig 22. Callisto calorimetry experiment BE3076. Samples of BM538 with mass 2.727g (H_2 channel) and 2.737g (D_2 channel) were subjected to 2x conditioning cycles and 2x long holds at 175C set point. Samples were left under vacuum for 12 hours between pressure cycles to degas. Each channel used only one isotope through all cycles. No indication of excess heat was detected.

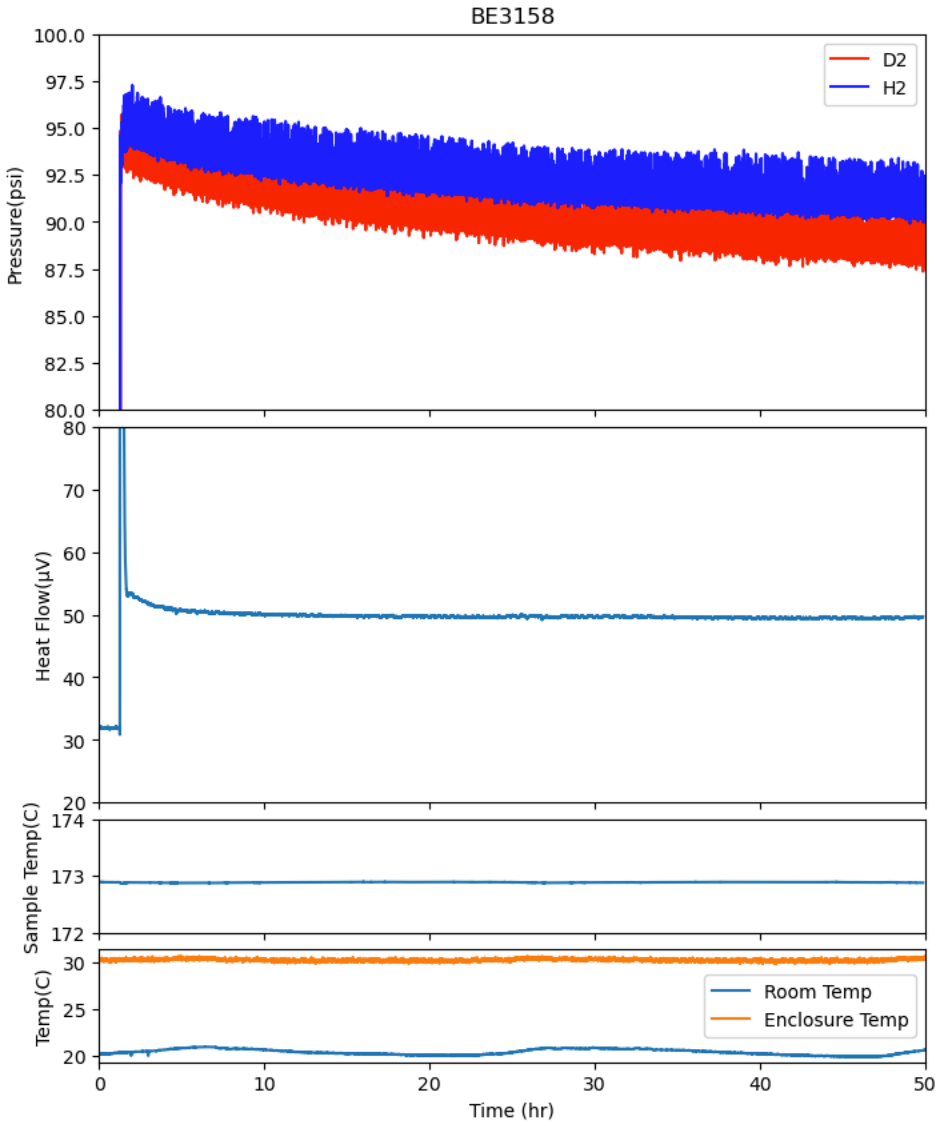


Figure 23. Callisto calorimetry experiment BE3158. Samples of BM536 with mass 2.51g (H_2 channel) and 5.21g (D_2 channel) were loaded into calorimeter and exposed to 6 pressure cycles at setpoint T=175C. The heat flow data for the final long duration pressure hold shown here shows no evidence of heat generation from either channel.

Case : Conclusion

Sixteen AHE experiments were carried out between the Case replication apparatus and the Setaram C80 calorimeter. Many of these experiments showed slight pressure decay within the vessels when isolated, which was suspected to be adsorption or chemical reactions. In the Case replication apparatus, no experiment showed significant temperature discrepancies between pressure/temperature cycles of different isotopes using equal heating power and pressure. In

differential calorimetry experiments, no significant heat flow was detected between H_2 and D_2 channels for any experiment. In conclusion, none of the evidence for AHE or isotopic contrasts reported by Case were detected in Callisto replication experiments.

Table 2: Case replication experiments

Experiment	Apparatus	Material	AHE Detected?
BE1399	Case style apparatus	BM524	No
BE1437	Case style apparatus	BM537	No
BE1477	Case style apparatus	BM538	No
BE1530	Case style apparatus	BM537	No
BE1721	Case style apparatus	BM891	No
BE2300	C80-calorimeter	BM537	No
BE2381	C80-calorimeter	BM537	No
BE2439	C80-calorimeter	BM538	No
BE2513	C80-calorimeter	BM538	No
BE2591	C80-calorimeter	BM538	No
BE2690	C80-calorimeter	BM537	No
BE2822	C80-calorimeter	BM538	No
BE3076	C80-calorimeter	BM538	No
BE3158	C80-calorimeter	BM536	No
BE3264	C80-calorimeter	BM524	No
BE3327	C80-calorimeter	BM524	No

LLNL : 1989 Experiment

Nearly all AHE claims involve small amounts of excess heat relative to input power (i.e., 5-50%) and so can potentially be explained by measurement errors, lack of controls, and conventional mechanisms. However, a few historical AHE claims observed exotic thermal runaways, well above such small levels.

A specific observation that strongly motivated our investigation into AHE was Lowell Wood's 1989 observation at LLNL (Lawrence Livermore National Laboratory). After Callisto electrochemical loading spent 3 years running LLNL replications without (yet) observing AHE, Callisto investigated potential classical (non-nuclear) explanations and found plasma electrolysis to provide a reproducible explanation that qualitatively matches Wood's observations.

Since no recorded data exists from the 1989 experiment, we will summarize Wood's recalled observations (**Fig. 24**).

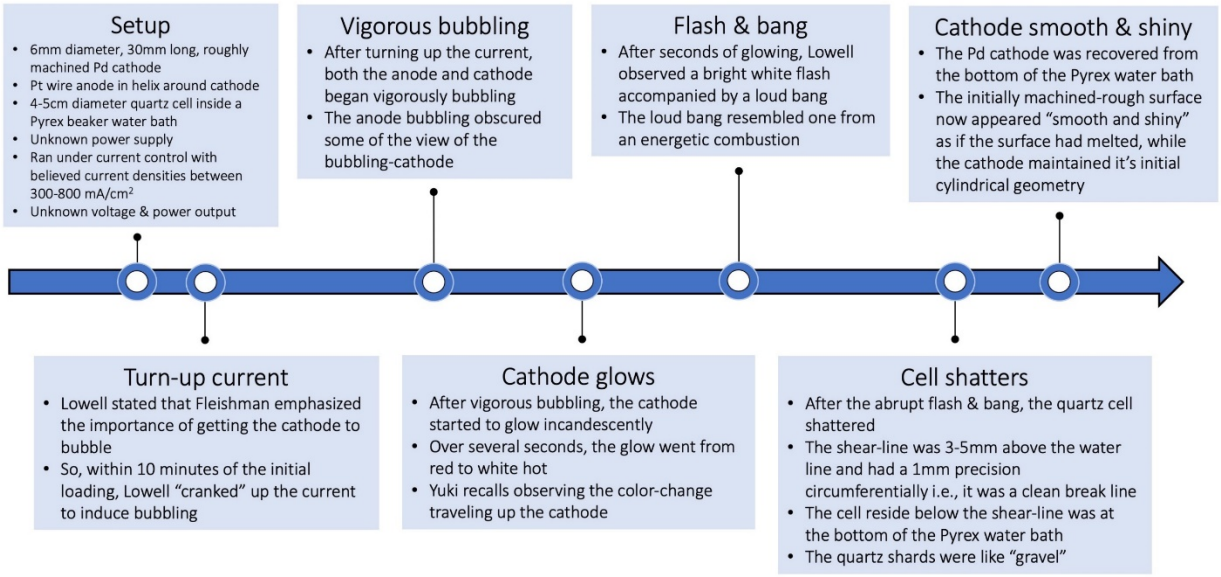


Fig. 24: Wood's recalled observations from the 1989 LLNL thermal runway experiment.

The apparatus involved a 6mm diameter, 30mm long, roughly machined Pd cathode. Prior to the experiment, the Pd cathode had visible surface deformations from the rough machining. A Pt wire was used as the anode and wrapped in a helix around the Pd cathode. The Pd cathode and Pt anode were placed in a 2-5cm diameter quartz cell with D_2O (heavy water) and ~0.5M ^{6}LiOD (Lithium-6 deuteroxide), which resided in a larger Pyrex-beaker water bath. Wood did recall at one point that the experimental cell was capped (similar to 1989-1990 Fleischmann cells), but this detail was fuzzy and not consistently recalled. The experimental cell and water bath had comparable water-levels, started at room temperature, and were believed (there was no temperature diagnostic) to warm up to 70-80°C throughout the experiment. **The exact power supply is unknown** (it was believed to be a readily available power supply from Wood's lab). The power supply was run in current control at current densities between 300-800 mA/cm^2; the voltage was unknown and ultimately depends on the exact power supply. Wood specifically remembers discussing with Fleischmann the importance of running at current densities above 300 mA/cm^2; so, Wood used whatever voltage was necessary to achieve >300 mA/cm^2.

During the experiment, Wood stood 1-2 feet away from the apparatus. Fleischmann emphasized to Wood the importance of getting the cathode to bubble; so, approximately 10 minutes into the initial loading of the Pd cathode, Wood "cranked up" the current to induce bubbling. It's important to note that ~10 minutes of electrolytic loading of a 6mm diameter Pd cathode means that while the surface may have been well loaded, the overall sample was only ~1% loaded[43]. After increasing

[43] After loading 100's of Pd wires, we've empirically observed a loading rate of ~100 μm/hr at room temperatures (e.g., it takes ~1 hour to load a 250 μm diameter Pd wire at 25°C). This abstracts the kinetic components (i.e., water splitting and H^+ diffusion), but typical timescales indicate loading is surface limited at room temperatures. Assuming

the current, both the Pt anode and Pd cathode started vigorously bubbling; the bubbling from the Pt anode obscured some of the view of the Pd cathode. Some unknown time after the vigorous bubbling started (it was likely within minutes), the Pd cathode started to incandescently glow. Over several seconds, the glow went from red to white hot. Another participant, Muriel Ishikawa, recalls observing the color change travel up the Pd cathode. After seconds of glowing, Wood observed a bright white flash, which was accompanied by a loud bang (one that "resembled an energetic combustion"). Roughly coincident with the abrupt flash-&-bang, the experimental cell shattered. The quartz cell had a shear-line 3-5mm above the water line with ~1mm circumferential precision and everything below was in shards (described like "gravel") at the bottom of the water bath. The Pd cathode was recovered from the bottom of the water bath. The initially rough surface now appeared "smooth and shiny" as if the surface had melted, while maintaining its cylindrical geometry.

Wood attempted "sporadic" replications in the following weeks but was largely preoccupied with programmatic responsibilities and was unable to replicate the runaway. The replication attempts were believed to use the same equipment (power supply, Pd cathode, etc.). Wood does recall that they *did not* precisely reproduce the anode/cathode geometry.

At Callisto, we investigated 3 possible classical (i.e., non-nuclear) explanations to the '89 LLNL runway: intermetallic alloying, H_2/O_2 recombination, and plasma electrolysis.

LLNL : Intermetallic Alloying

It's likely that many "excess heat" reports are nothing more than calorimetry errors and less-than-rigorous scientific methods. However, there are a few interesting sidenotes that indicate traditional chemical reactions could produce excess heat, particularly from authors reporting frequent production of LENR.

In Storms' well-known 1996 paper[44], he added Aluminum metal powder to his electrolyte: "… Aluminum metal (2-20 ppm) added to the electrolyte after the palladium has achieved its maximum deuterium content is sometimes useful in initiating excess heat production." The enthalpy of formation for a Pd-Al alloy is 95 kJ/mol: this mixture is commonly used in pyrotechnics and as a reactive nanolaminate; it's akin to thermite without oxygen transfer.

Similarly, McKubre found that Al and Si additives enhances loading[45] (and thus AHE): "… the presence of approximately 200 ppm of dissolved aluminum (or silicon) in the electrolyte was found to facilitate the reproducible attainment and maintenance of high loadings." Palladium silicide can be readily formed and has a (stoichiometry-dependent) enthalpy of formation of ~60 kJ/mol.

---

a constant loading rate of 100 μm/hr, the 6mm diameter Pd cathode only loaded ~16 μm of its surface or ~1% of the total volume.

[44] E. Storms, *J. Fusion Technol.* **29**, 261 (1996).

[45] M. McKubre et al., *Frontiers of Cold Fusion*, Universal Academy Press, Inc., Japan (1992).

Finally, Letts added an unknown rare earth to his electrolyte[46]: "The cathode was a Gold substrate 5mm x 8mm x 0.75mm with Palladium plated onto the substrate with a rare earth additive provided by Cravens. The cathode didn't respond well to the laser stimulation until the rare earth solution was added to the electrolyte. With the additive, the cathode became responsive to a laser operating at 657nm with a radiant output of 30 mW." One can only guess what was added; nevertheless, this is another example of a potentially highly reactive alloying additive being added to the electrolyte.

The authors in the above experiments do not discuss whether the amount of their additives are sufficient to account for the reported amounts of excess heat.

To test whether intermetallic alloying reactions could qualitatively explain the '89 LLNL observations, we intertwined 200 µm Pd and 200 µm Al wires and drove power across them while in both air and water (**Fig. 25**). As expected, the Pd-Al reaction energetically proceeds in air but cannot propagate in water i.e., the 95 kJ/mol from the Pd/Al reaction cannot overcome the strong cooling power of water. Thus, intermetallic alloying reactions produce heat in gaseous environments over short durations but not enough to produce the '89 LLNL major thermal runaway in water.

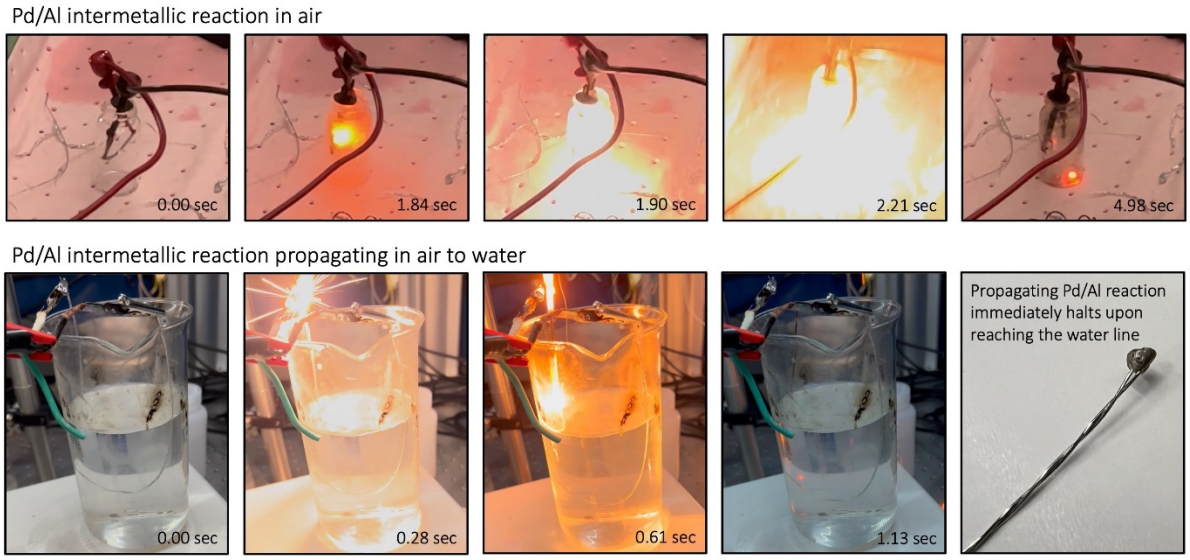


Fig. 25: Pd-Al intermetallic alloying produces 95 kJ/mol but cannot propagate in water.

LLNL : H_2/O_2 Recombination

[46] Letts et al., *Condensed Matter Nuclear Science*, pp. 159-170 (2005).

In thermodynamic equilibrium, loading decreases with temperature at a fixed pressure. In other words, a loaded lattice can be deloaded through heating. Furthermore, the deloading rate increases with temperature and the enthalpy of formation of water is ~286 kJ/mol. So, as a sample deloads H/D, the H/D can recombine with O_2 to form water, which heats the sample surface. As the surface heats up, the deloading rate increases, which in turn further heats up the sample i.e., it can create a thermal runaway scenario. An ideal sample geometry has a high surface-area to volume ratio (e.g., a foil), which enables more surface sites for recombination and reduces the deloading distance.

To test this hypothesis, we loaded a 50mm x 50mm x 2mm Pd foil in 1 bar H_2 and subsequently removed it from the vessel and placed it in air. As the sample deloads, water forms on the surface and reduces the available surface sites for recombination. So, we heated the foil to 100°C to evaporate the surface water and subsequently removed the heating source. As anticipated, the foil underwent a thermal runaway, eventually reaching >400°C (based on an IR camera measurement) and incandescently glowing red hot (**Fig. 26**).

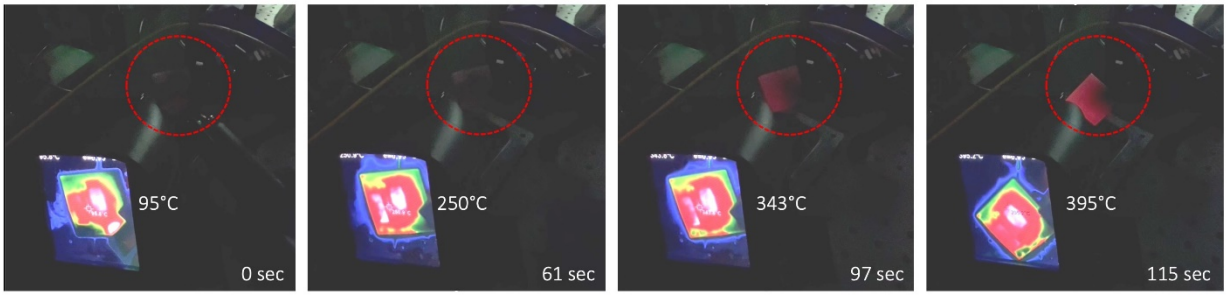


Fig. 26: Thermal runaway of a Pd foil from H_2/O_2 recombination in air.

To better quantify the temperature rise and timescale, we used our Thermal Ignition apparatus to run a highly controlled experiment. We elastically loaded a Pd slug with protium and subsequently reduced the temperature to -100°C to lock-in the loading. Prior to loading the Pd slug, we took control measurements with 20 psi N_2 at -100°C and laser heated the Pd to 137°C over 10 seconds. This calibrated our required laser input to heat to >100°C and showed the expected heating/cooling rates without excess heat produced. With the PdH slug loaded at -100°C, we fully purged our system and switched gases to 20 psi O_2. We then laser heated the PdH to 137°C over 10 seconds and shutoff the laser. The PdH slowly self-heated to ~200°C over 18 seconds and eventually ran away to 965°C over 3 seconds (**Fig. 27**).

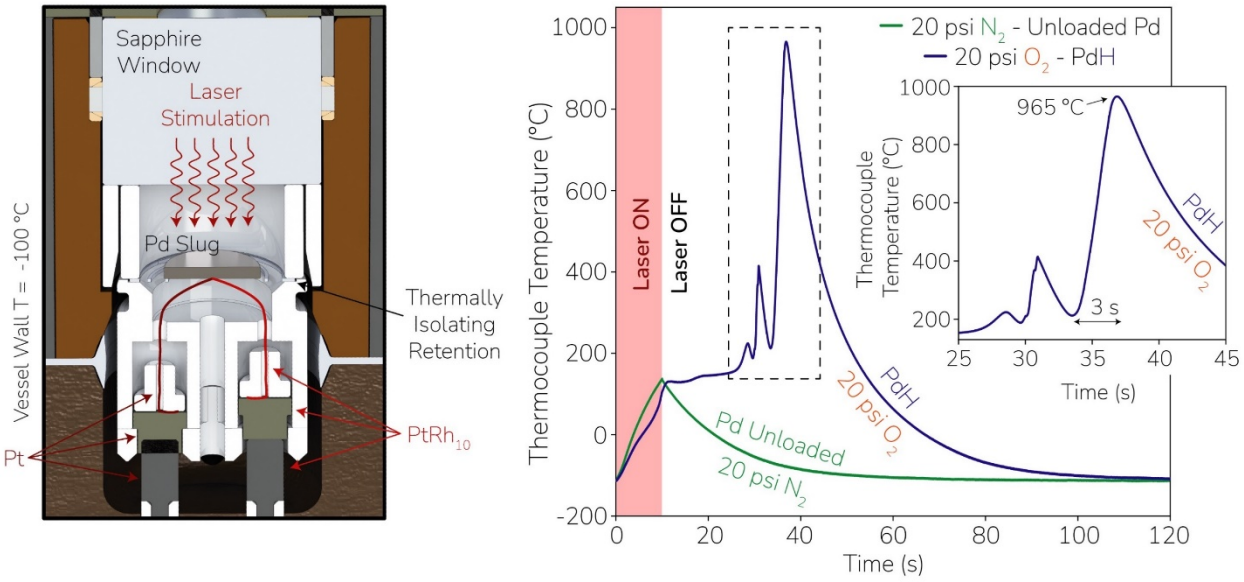


Fig. 27: Thermal runaway of a PdH slug in 20 psi O_2 from H_2/O_2 recombination.

After clearly establishing thermal runaways from H_2/O_2 recombination in gaseous environments, we developed a split cell to empirically determine if the ~286 kJ/mol from H_2/O_2 recombination could overcome the strong cooling power of water. Our split cells allow for fluid/gas, gas/gas, and fluid/fluid experiments with electrical and optical access. In a pure liquid environment, it's a stretch to explain how enough O_2 could contact the PdH surface to enable significant levels of H_2/O_2 recombination. In electrolytic cells, O_2 is generated at the anode during water splitting and would need to contact the PdH surface at sufficient quantities and rates. To test the best-case scenario, we used our split cell in gas/liquid configurations (i.e., one side of the slug is in contact with gas, while the other side is in contact with liquid). First, we ran a N_2/H_2O control, where we laser heated a Pd slug on the N_2 side with the backside in contact in H_2O. As expected, we observed a plateau in temperature around 100°C (the measured ΔT was ~70°C with a T_0 of ~30°C) i.e., the Pd|H_2O system does not exceed the H_2O boiling point (**Fig. 28**). Out of due diligence, we ran an experiment with a loaded PdH slug and O_2/H_2O in the split cell. When laser heating in O_2, we observed a minor increase in peak temperature (still around 100°C) and did not thermal runaway (**Fig. 29**). This proves H_2/O_2 recombination is not responsible for the '89 LLNL thermal runaway.

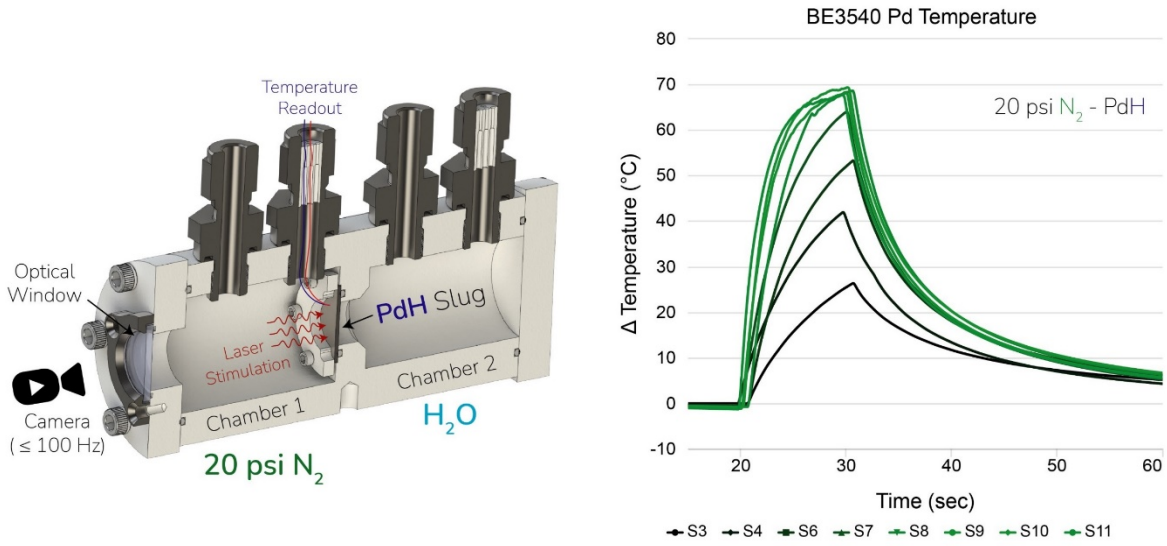


Fig. 28: Using a split-cell, we show that a Pd slug in contact with H_2O does not exceed 100°C (the ΔT plateaus near 70°C with a T_0 of ~30°C).

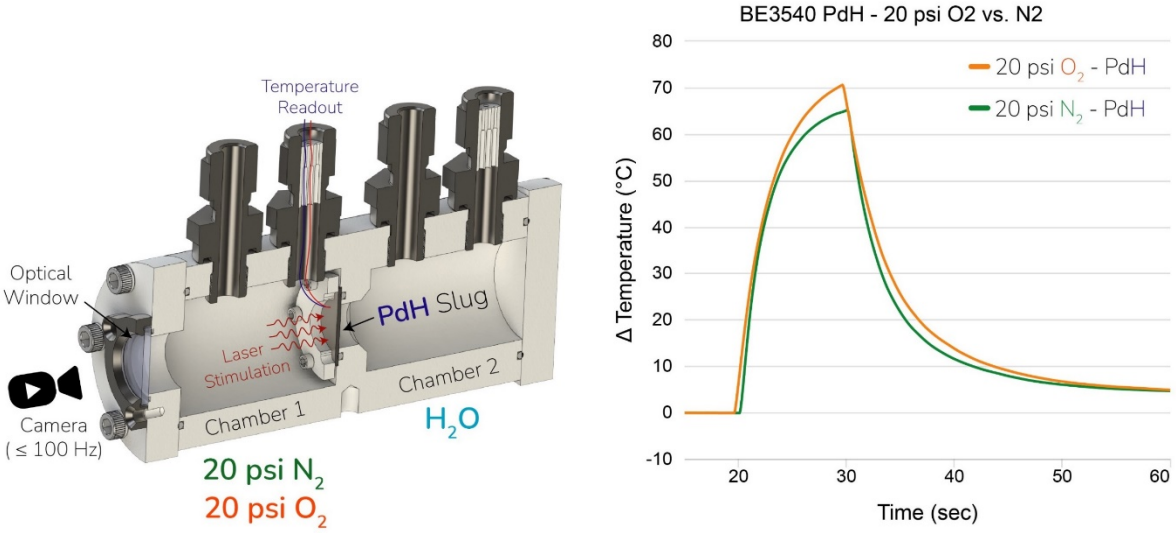


Fig. 29: Using a split-cell, we show that O_2 (which undergoes H_2/O_2 recombination) produces slightly higher temperatures than an N_2 control, but the peak temperature remains around 100°C (ΔT = 70°C with a T_0 of ~30°C) and no thermal runaway is observed.

Given our control experiments, specifically those with laser-heating Pd in contact with water, it became challenging to understand how classical mechanisms could enable a Pd cathode to incandescently glow in water, reaching temperatures >1000°C: the power output would need to be significant. To put rough numbers to the required power output, we calculated the energy density

required to heat the surface of a 6mm cylinder in water to 1000°C (i.e., glowing "white hot" as observed in the '89 LLNL experiment) using COMSOL. The heating is assumed to occur in a 15 μm annulus at the surface. We found minimum energy density to glow white-hot is 2.9 kJ/cm^3 (i.e., instantaneous) and 407 kJ/cm^3 is required for glow for 1 second (**Table 3**). A 6mm diameter, 30mm long Pd cathode has volume of 0.85 cm^3 and 8.2x10^{15} surface atoms, which is ~10^{-8} mol. If we assume chemical reactions are occurring only with surface atoms and take an upper bound of 1000 kJ/mol for a reaction[47], we could instantaneously produce 10^{-5} kJ, which is orders of magnitude below the 2.5 kJ required for the cathode to instantaneously glow. If not just the immediate surface, but the aforementioned annulus of 15 μm, reacts, then 10^{-3} moles are involved, producing 1 kJ, near the required value.

Duration (sec)	Required Energy Density (kJ/cm^3)
0.000	2.9
0.001	35.6
0.010	104
0.100	267
1.000	407

Table 3: The energy density required to heat a 15 μm annulus at the surface of a 6mm cylinder in water to 1000°C.

While chemical reactions are highly unlikely to produce a glowing cathode in water, a nuclear explanation presents challenges as well. There is certainly plenty of energy available; highly loaded PdD can (if all deuterons react) produce over 10^8 kJ/cm^3; so only a tiny fraction of the available deuterons must react.

However, even if we assume a nuclear explanation, certain observations from the '89 LLNL experiment seem to be at odds: the cathode glowed for seconds *and* the beaker shattered *and* only the cathode surface melted. Thermal diffusivity in the Pd cathode is significantly faster than in water so, for seconds-long heating durations, the cathode is nearly in thermal-equilibrium. So, if the cathode surface melted, and glowed for seconds, much of the bulk should have as well[48]. Moreover, for the beaker to shatter from a shock wave, there must have been rapid heating (e.g., < millisecond); but the cathode was observed to glow for seconds.

Two observations (that only the surface melted, and that the beaker shattered) seemingly require very short, sub-millisecond, heating durations, and are at odds with the observation of seconds-long glowing. It's possible that the glowing didn't actually last for seconds, but was simply perceived to do so due to human physiology. But, while a sub-millisecond heat pulse might explain

[47] The highly exothermic thermite reaction of Fe_2O_3 and Al produces ~850 kJ/mol.
[48] The timescales of atomic diffusion matter, but it's hard to envision surface reformation without the bulk failing to some extent.

both the surface-deep heating and the beaker shattering, it leaves us having to understand how nuclear events across a large surface region are time-synchronized so tightly.

Perhaps a glowing cathode in direct contact with water is too simple of a picture. What if we had a scenario where the liquid surrounding the cathode locally boils, creating a vapor environment and reducing the power required to heat?

LLNL : Plasma Electrolysis

Plasma electrolysis[49] is a nontypical electrochemical process that uses electric current to produce a plasma from an electrolyte solution. At high enough voltages, electrolyte solution around the cathode can heat up and vaporize, forming a vapor envelope. Within this vapor layer, the high temperatures and electric field can cause the molecules to ionize, creating a plasma, which is typically visible as a glowing region around the cathode (**Fig. 30**). The specific conditions under which plasma forms depends on a range of factors including the composition of the electrolyte solution and electrodes, the anode/cathode geometry, and the applied current and voltage.

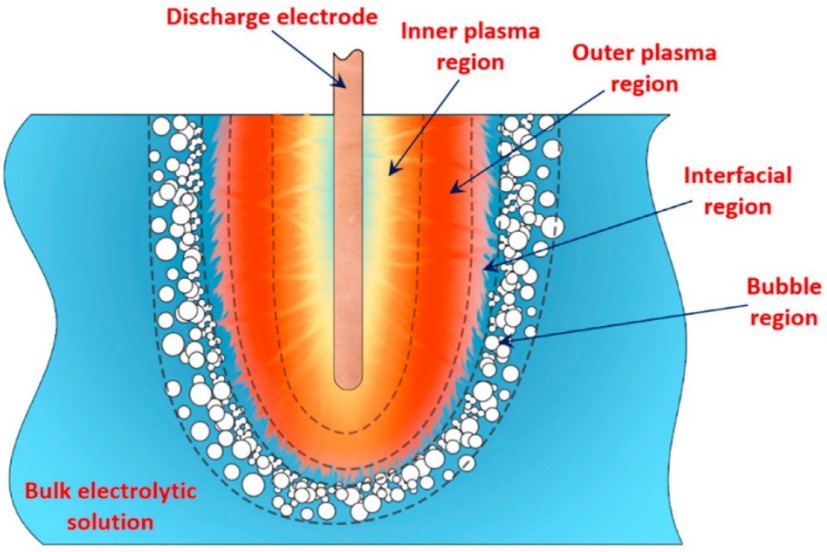


Fig. 30: Plasma electrolysis can occur in electrochemical cells when high enough voltages are applied. Figure taken from S. Bespalko[50].

[49] Plasma electrolysis is also referred to as plasma-driven solution electrolysis (PDSE) and contact glow-discharge electrolysis (CGDE).
[50] Bespalko, S., Mizeraczyk J., *Energies*, **15**, 7508 (2022).

Fundamental to generating plasma electrolysis is a sufficient power supply. We developed two, independent apparatuses with different power supplies (**Fig. 31**). One setup used an 80 V, 62 A, 5 kW power supply (Magna Power Electronics TSA80-62), while the other used a higher voltage 230 V, 20 A, 1.5 kW power supply (NI RMX-4126). Importantly, these power supplies are common and it's plausible that units with similar specifications were in the '89 lab at LLNL. We replicated the '89 LLNL setup (e.g., cathode geometry, current density range, quartz cell, etc.) and followed Wood's protocol: cranking up the current to induce vigorous bubbling. Doing so, we observed plasma electrolysis as low as 30-40 V with initiation current densities < 3 A/cm^2 and much lower sustainer currents. At voltages < 80 V, current densities < 3 A/cm^2, and 0.5M LiOH in H_2O, we observed both sustained glow discharge (**Fig. 32**) and intermittent glow discharge (**Fig. 33**). At a set voltage and current density, we found that the cathode/anode geometry was important to reproducibly observe plasma electrolysis (this matches Wood's recollection, where the cathode/anode geometry was inconsistent in his replications and no additional runaways were observed).



Fig. 31: Power supplies used in '89 LLNL replication experiments.

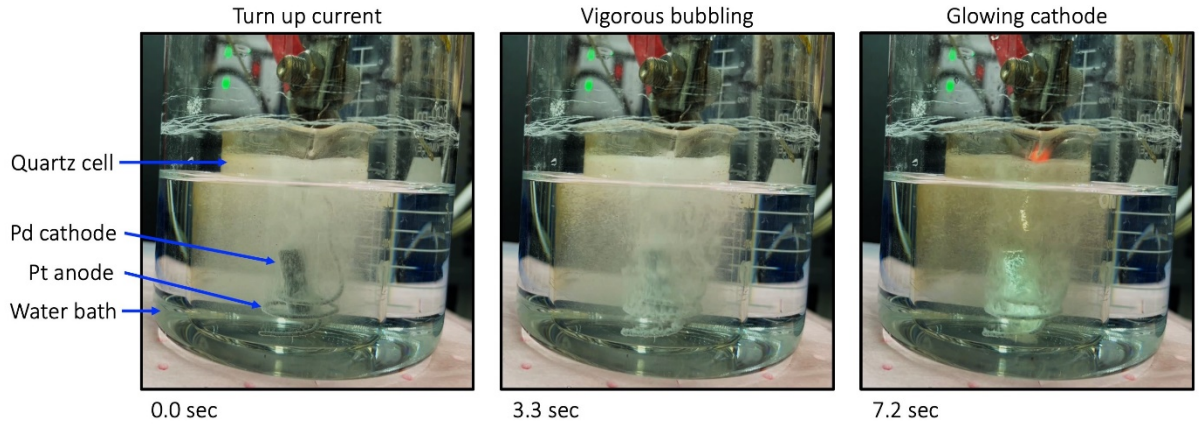


Fig. 32: Sustained glow discharge is observed as low as 30-40 V and < 3 A/cm^2 in 0.5M LiOH in H_2O.

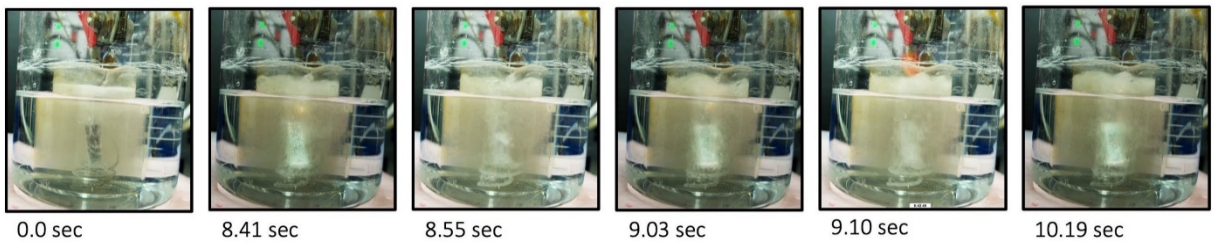


Fig. 33: At times, intermittent glow discharge is observed as low as 30-40 V and < 3 A/cm^2 in 0.5M LiOH in H_2O.

After observing plasma electrolysis at lower voltages (e.g., 30-40 V), we used a higher voltage supply to probe up to 230 V (using 0.25M LiOH/LiOD in H_2O/D_2O, respectively). At higher voltages, the plasma electrolysis is exceedingly violent, quickly increasing the water temperature to nearly 100°C and heating the cathode to a glowing white-hot state (**Fig. 34**). Importantly, during plasma electrolysis, the current density is < 800 mA/cm^2, which matches the conditions from the '89 LLNL experiment (**Fig. 35**). Furthermore, our Pd cathodes often had melted surfaces, again matching the '89 LLNL description. Moreover, the plasma electrolysis is coupled with loud "flashes and bangs", which we attribute to H_2 / O_2 combustion. The Faradaic efficiency in plasma electrolysis is dozens of times higher than normal electrolysis, producing orders of magnitude more hydrogen[51]. When the H_2 and O_2 produced from the electrolysis are in proximity to the plasma, sparks form and combust the H_2 / O_2. To prove plasma electrolysis is unrelated to loaded Pd cathodes, we ran Pt/Pt controls and observed qualitatively the same plasma electrolysis and H_2 / O_2 combustions.

[51] Susanta K. Sen Gupta, *Plasma Sources Sci. Technol.*, **24** 063001 (2015).

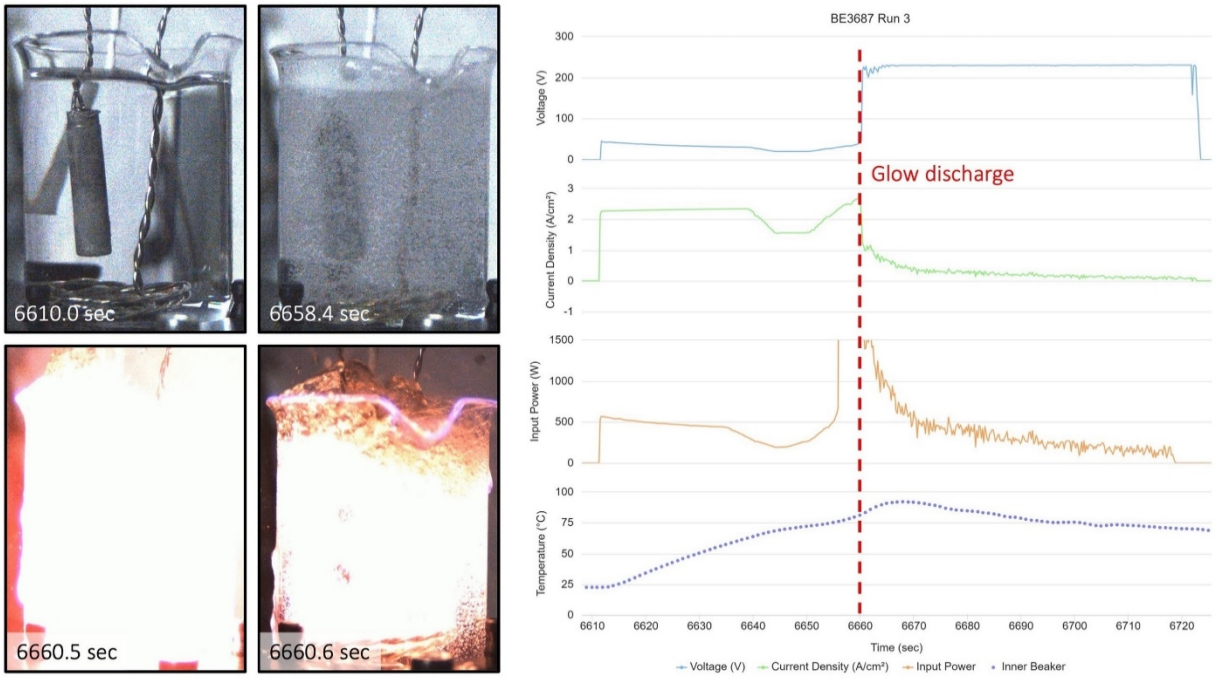


Fig. 34: In both 0.25M LiOH and LiOD, increasing current density produces vigorous bubbling after which violent plasma electrolysis is observed at 230 V, < 3 mA/cm^2.

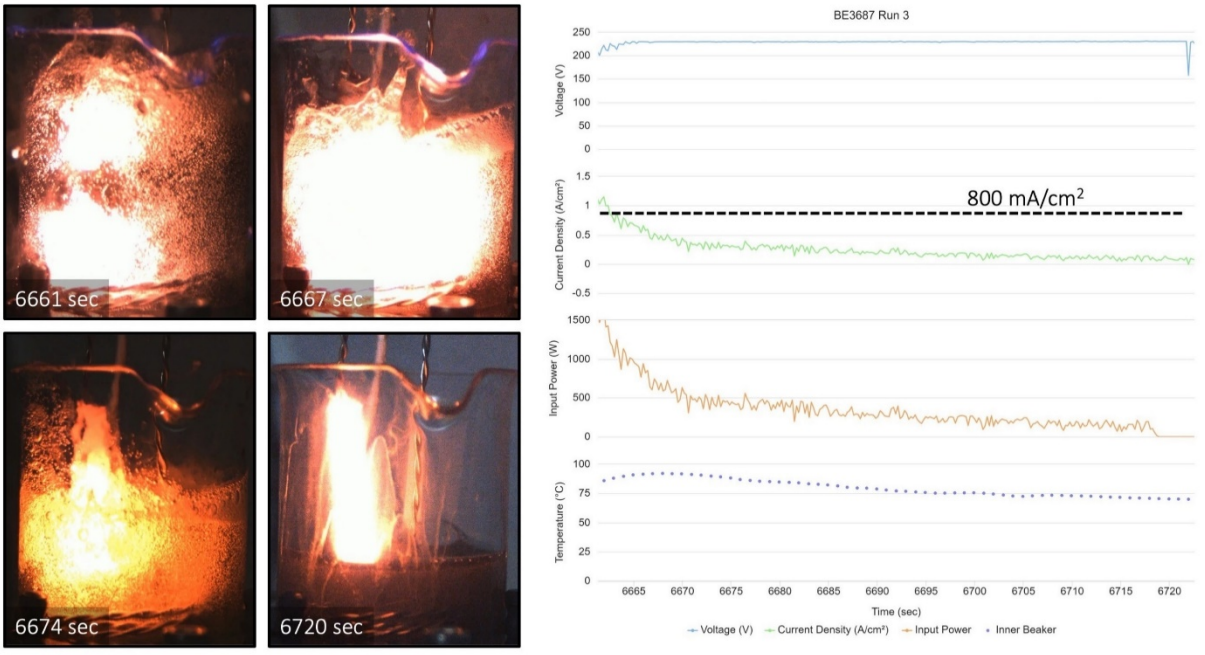


Fig. 35: In both 0.25M LiOH and LiOD, violent electrolysis occurs < 800 mA/cm^2, matching the observations in the '89 LLNL experiment. The plasma electrolysis is accompanied by loud H_2 /

O_2 combustions (as seen at 6667 sec). The water level is reduced due to splashing, revealing a white-hot Pd cathode even after the voltage and current are turned off at 6720 sec.

While the H_2 / O_2 combustions are most likely responsible for the shattered cell in the '89 LLNL experiment, it's important to highlight the shattered-cell observation is not scientifically testable. Without knowing the exact state of the glass used in the '89 experiment, we cannot meaningfully test theories for glass failures. The strength of glass is highly dependent on external factors (e.g., crack propagation occurs at significantly lower stresses with the presence of water and under heating[52]), pre-existing conditions (e.g., repeated, cyclic heating leads to increased defects[53]), and composition[54]. At a photovoltaic module reliability workshop at NREL[55] (National Renewable Energy Laboratory), C. Barry discussed "why glass sometimes breaks." Their data on glass failure showed a >3x spread in failure stress and was summarized well: "Will my glass break? Nobody knows for sure. You can't tell how strong it is, until it was."

Keeping the non-testability of the '89 shattered cell in mind, we built an apparatus to test H_2 / O_2 combustions in water in *unused* (i.e., as delivered from the vendor) quartz cells (**Fig. 36**). The apparatus was remotely controlled and mixed O_2 + $2H_2$ in a small, isolated volume. To ignite the gas mixture, we used a modified spark plug, which applies a high voltage transient and breaks down the gas between two electrodes. We tested as low as 2 mL of an O_2 + $2H_2$ mixture in direct contact with 50 mL quartz beakers (1.5mm average wall thickness, with 1.5-2.3mm thickness at the base) and observed catastrophic glass failures in all cases. Depending on the volume of O_2 + $2H_2$, the shard size distributions varied, ranging from a few large shards to dust. To simulate a scenario of a H_2 / O_2 combustion in the water (i.e., it's not in contact with the glass), we mixed O_2 + $2H_2$ in a thin latex balloon suspended in water. We found that up to 3 mL of an O_2 + $2H_2$ mixture was unable to break the quartz. Moving the thin latex balloon in-contact with the quartz wall (while still suspended in water) did catastrophically break the quarts (**Fig. 37**). So, we concluded that a combustible < 3mL O_2 + $2H_2$ mixture must be in contact with glass to induce fracture.

It's important to highlight that Mizuno ran plasma electrolysis experiments between 2000 – 2005 and observed *one* catastrophic explosion[56]; however, this had nothing to do with loaded Pd cathodes as Mizuno used a Pt anode and W cathode, neither of which load with deuterium.

[52] S.J. Grutzik et. al, *Journal of Non-Crystalline Solids: X*, 16, 100134 (2022).
[53] W. Feng et. al, *Corrosion and Materials Degradation*, 2(3), 412-446 (2021).
[54] G. Rene, *Journal of Non-Crystalline Solids*, 316(1), 1-11 (2003).
[55] C. Barry, *Photovoltaic Specialists Conference* (2009).
[56] T. Mizuno, *The 12th International Conference on Condensed Matter Nuclear Science* (2005).

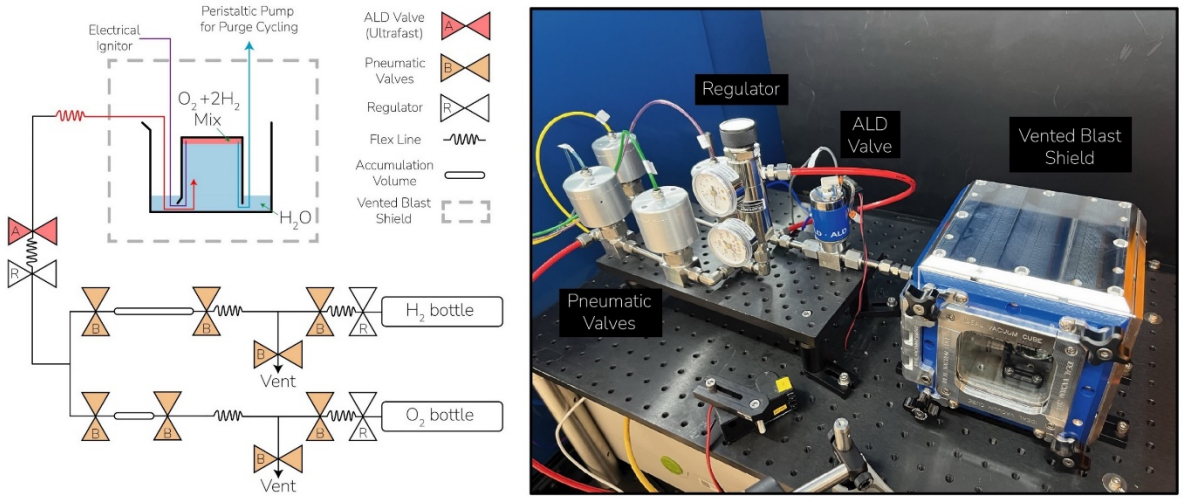


Fig. 36: Apparatus used to study H_2 / O_2 combustions in water.

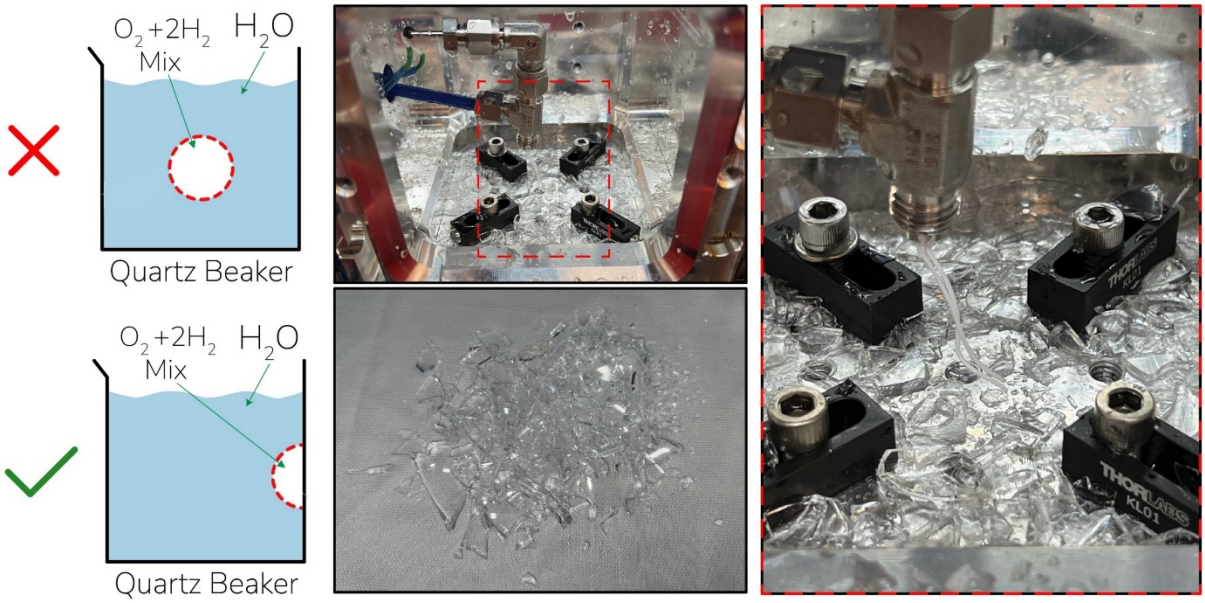


Fig. 37: As long as a combustible O_2 + 2H_2 mixture is in-contact with the quartz cell wall, the cell catastrophically fails.

In summary, we've reproducibly explained (in qualitative terms) the key observations of the '89 LLNL experiments: "cranking up" the current density in current control leads to "vigorous bubbling" and subsequently a "glowing cathode" that undergoes a color change from "red to white

hot" after which a "flash and bang" occurs, shattering the cell[57] and leaving a melted Pd cathode surface (**Fig. 38**). We argue that the voltage of the "unknown power supply" is the key to the '89 LLNL experiment: if capable of >30 V at the 800 mA/cm^2 current density believed to have been used, then plasma electrolysis and H_2/O_2 combustions can occur and explain the observations.

Hopefully, the ongoing Callisto electrochemical loading replications of the '89 LLNL experiment will demonstrate LENR. In that event, then it is reasonable to assume that the '89 LLNL experiments were also due to LENR. However, if they do not, then the one-off nature of the '89 LLNL experiment (and the lack of any quantitative data records from it), means that we will never actually know the cause of the LLNL results. LENR is one possible explanation of those results. But, given the Callisto experiments, we believe that LENR is no longer the only possible explanation. Hence, the existence of the '89 LLNL results is not sufficient to establish the existence of LENR "beyond a reasonable doubt".

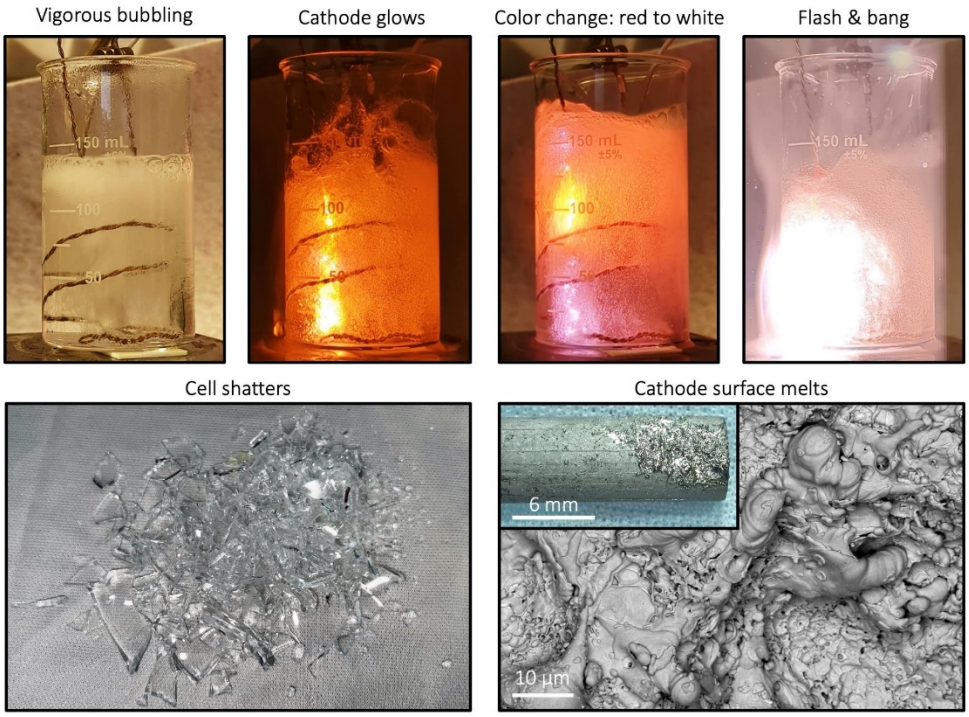


Fig. 38: The key observations from the '89 LLNL experiment have been qualitatively reproduced and explained by plasma electrolysis and H_2/O_2 combustions.

[57] All cell shattering experiments were done in the controlled H_2 / O_2 combustion apparatus.

Calorimetry

Executive Summary

Using a series of commercially purchased differential calorimeters to quantitatively measure the amount of energy released or absorbed within an isolated Pd-based system, we have not observed any observable quantities of heat that can be attributed to anomalous heat events (AHE) or low energy nuclear reactions (LENRs). Calorimetry is one of the few experimental campaigns carried out in our AHE search that is directly measuring heat flow from an isolated system. Methods used were focused on bulk AHE detection and not in very low mass/thin film type geometries since our calorimeters require substantial heat activity to detect AHE and we need this effect to be substantial for it to be useful and applicable to real world clean energy technologies.

A multitude of sample geometries, temperature regimes, thermal stimulations, and loading pathways were employed inside our calorimeters in efforts to induce and capture AHE. Pd nanoparticles to bulk foils were loaded in environments where deuterium pressures reached 150 ksi at temperatures ranging from -150 °C to 200 °C. Loadings up to unity were achieved, and the calorimeter system was monitored in an equilibrium state between 3-14 days for each experiment. Time sensitivity was based off thermal time constants of each calorimeter which ranged from 3-14 min; however, 5mW pulsed power as short as 100 ms could be detected. We used a variety of differential calorimeters manufactured by the Setaram company; differing in characteristics such as temperature range, available sample dimensions, and sensitivity. Overall sensitivity for each calorimeter type can be found in the table below, but in general each calorimeter required hundreds of billions of AHE events to occur over a few minutes to be observable through calorimetry. This level of activity translates to having >1 AHE event for every 100 million loaded lattice sites to be detected by our calorimeters.

Table 1- AHE Sensitivity

Calorimeter	Sensitivity	Fractional site sensitivity
MS80	50 µW	1 in 10^{11} sites
BT215	200 µW	1 in 10^{10} sites
C600	2 mW	1 in 10^{9} sites

Background

Calorimetry is the main experimental method to obtain thermodynamic data in an isolated, well controlled, and defined environment. The basic concept in calorimetry is to measure, quantifiably, how much energy/mass is inserted and/or removed from a system. The system, with one channel, is what is being monitored and anything that crosses the system boundary to the surroundings is considered an input or output from the system. In the calorimeters we used, a sensor called a thermopile[58] is attached to the outside of the channel while remaining within the system boundaries, and this is what measures which direction and to what magnitude heat is transferring around the channel.

The specific type of calorimetry employed for this project is called differential calorimetry which adds an additional control system to improve the sensitivity and resolution of our calorimeters. Instead of having only one channel inside the calorimeter there are two. They are identical channels in mass, location from calorimeter center, materials, and environmental inputs. In most of our experiments the control channel has protium gas or light water instead of deuterium or heavy water. This control will undergo the same loading ramps and be introduced to identical peak pressure conditions but we do not expect 24 MeV AHE production to form. Therefore, since the sample material, pressure environment, and temperature are the same between the 2 channels, we can subtract out the control thermopile signal from the thermopile attached to the active channel, taking a differential. This method removes many classical reactions that emit or absorb heat during Pd loading including: gas compression and enthalpy of loading summed up in Equation 1. With these more classical phenomena accounted for in the differential signal, ideally, the only signal that would be added to produce excess heat in the PdD channel is AHE. In a control test the delta between the channels would be zero in a fully closed system. The excess heat would be the measured power over a given time period. The general method of accounting for input energy in a system is the coefficient of performance energy (COP_{energy}) (Equation 2). If the COP or efficiency of the system is significantly above 1 it would indicate one of the following: energy accounting is erroneous causing false positive, classical reaction is occurring, or AHE.

$$Eqn\ 1 \quad \Delta E_{in} = \left[m_{control} \int c_{v_{PdH}} dT + \Delta H_{loading}\right] - \left[m_{active} \int c_{v_{PdD}} dT + \Delta H_{loading}\right]$$

$$Eqn\ 2 \quad COP_{energy} = \frac{E_{measd}}{E_{in}} > 1\ ,\ if\ E_{measd} > E_{in}$$

[58] Thermopile: a circuit consisting of a series of thermocouples, used for measuring minute changes in temperature (indicated as a voltage) or for generating thermoelectric current. In this case the unit is μV.

Background : Passive Calorimetry

No additional energy was introduced into the system outside of (1) loading enthalpy which self-heats, and (2) protium or deuterium isochoric gas compression (above 20 ksi up to 150 ksi) during experimental set-up prior to sample/calorimeter isolation for the remainder of the experiment. The absence of externally supplied stimulation energy makes these our most sensitive calorimetric AHE experiments. Controlling energy and mass input points and closing them off during the stable phase of an experiment makes this a closed-type calorimeter. Gas was applied using ultrapressure downtubes in each calorimeter channel. To minimize any thermal gradients up these downtubes, and thereby minimize heat flow that was not being captured radially around the channel's thermopile, a heater collar [59]was clamped to the top of each down tube using a PID loop to maintain the same set temperature at the top of the downtube as the internal calorimeter (Figure 1). Once the system is isolated there is no pathway for mass to escape. Ideally, any energy transfer would have to occur through conduction radially outward from the vessel wall into the surrounding thermopile sensor.

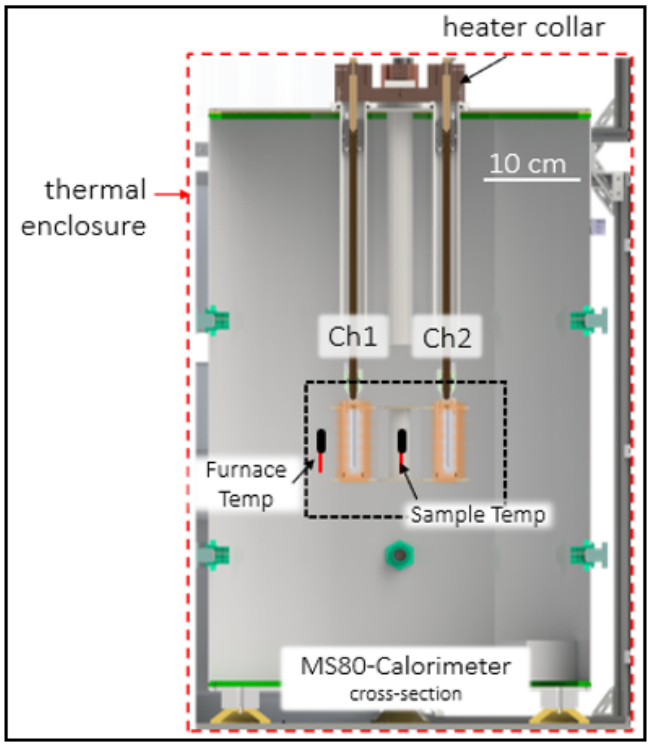


Figure 1: Heater collar location in experimental set-up.

The motivation for running these closed-type calorimetry experiments is to maximize our AHE signal sensitivity by not introducing additional noise to the system such as mass loss, electrical heaters, convection currents, etc. Removing these signals reduces uncertainty in our heat flow measurement as well, and since AHE should only be occurring in the PdD channel we should see a clear divergence in heat flow potential indicating the PdD is heating up.

Without any additional forms of stimulation, however, we are implicitly demanding that AHE occurs through loading Pd with deuterium alone. We are asking: do the kinetics of Pd loading alone cause an AHE run away reaction to a high enough intensity or flux for the thermopile

[59] Heater Collar: Bulk copper heater clamped at the top of the calorimeter and channel downtubes and maintained at the same set temperature as the calorimeter.

to measure a detectable amount of released heat from the PdD lattice? Does a loaded Pd lattice in isolation form local AHE spontaneously?

Background : Active/Stimulated Calorimetry

The active calorimetry conducted in these studies mainly focuses around adding additional electronics that were internal to each sample channel and therefore contributed to an input power over time that must be accounted for in our heat flow analysis for AHE when calculating COP_{energy}. All stimulation in gas calorimetry tests was focused on thermal stimulation either through Joule heating of the Pd sample itself or through conduction and convection heat transfer paths between a heater wire and the Pd sample. These experiments are mainly asking the question, do loading gradients caused through temperature gradients in the Pd cause deuteron hopping to a large enough extent to where we increase the probability of 2 deuterons meeting and fusing in an interstitial site?

Using active calorimetry, as opposed to passive calorimetry, involves a tradeoff; does the advantage of stimulation in helping/allowing AHE offset the lower achievable sensitivity? Fortunately, of course, this is not an either/or proposition; we were able to perform both types of experiments.

Background : Setaram Calorimeter Suite

The differential calorimeters used throughout this project were sourced from Setaram Inc. The company has extensive knowledge of the design and production of these calorimeters and is world class in terms of quality, resolution, sensitivity, and reproducibility in both academic and industrial research settings. Sourcing all calorimeters from this reputable supplier made sense in both time savings and data quality, instead of designing and troubleshooting our own custom differential calorimeters from scratch. It also, of course, will aid any other researchers interested in replicating our experiments; they will face no need to attempt to resurrect (or trust) a custom, home-built calorimeter, Four main calorimeters were used throughout the duration of the project, a basic overview of each calorimeter's specifications can be found in Table 2.

MS80 – The largest calorimeter by size containing 4 sample channels, each 35 mm in diameter, this calorimeter allowed for larger mass samples (up to 15 g) in each channel (Figure 2). Channels 1 and 2 are a differential pair while channels 3 and 4 were a second pair within the same calorimeter. The second pair would allow a second experiment to run in parallel or run as another control pair to account for any environmental changes to the system over the course of the experiment. The MS80 is the most sensitive out of all the calorimeters employed but most limited in set temperature range of 25-200 °C. It is also the least sensitive in time with a thermal time

constant [60]of 12 min. The largest heat flux range that could be reliably measured was between +/- 500 mW between channels. If heat flow from one channel to another exceeds 500 mW in this differential measurement the system saturates [61]the sensors, and we no longer would have an accurate heat flow signal until the heat flux between channels drops below 500 mW once more. Both gas and liquid type loading/AHE experiments were conducted inside this system.

C600 – 2 channel differential calorimeter optimized for high set temperature experiments (25-600 °C) (Figure 2: Cross sections of a) thermopile sensor, b) MS80, c) C600, d) BT215 with liquid N2 jacket around channels 2). The highest temperature explored was 200 °C before the channels were damaged due to a pressurized sample vessel catastrophically rupturing inside the calorimeter damaging the thermopile beyond repair. Though this instrument was damaged, the higher temperature parameter space was well explored through other campaigns including Thermal Ignition and NanoTI.

C80 – 2 channel differential calorimeter optimized for high set temperature experiments (25-300 °C). The unit dimensions are similar to the C600 unit with a slightly longer bore accessing the bottom of each channel where the thermopile sensitive volume is located. This unit was mainly utilized in initial low-pressure tests (<10 ksi) to confirm Pd loading, as well as replication experiments that are explained in detail in their respective sections of this overall report.

[60] Time constant (τ): time it takes for the heat flow signal to decay to e^{-1} from its peak maxima. It is related to calorimeter heat capacity and conductive path between the sample vessel and constant temp block, controlled by the calorimeter construction. Expressed as the time for the temperature at the element of a thermistor, with no load applied, to change to 63.2%, during a sudden change in temperature.

[61] Heat flow saturation: The thermopile sensor has a limit on sensitivity of the electrical circuit measured in uV, if the potential exceeds a limit of 20000 uV (converted to 500 mW) it is outside the linear range of the circuit and therefore inaccurate to heat flow measurement.

BT215 – 2 channel differential calorimeter optimized for cryogenic set temperatures down to -190 °C (Figure 2). This instrument was mainly used to run experiments both passive and stimulated while at unity loading conditions.

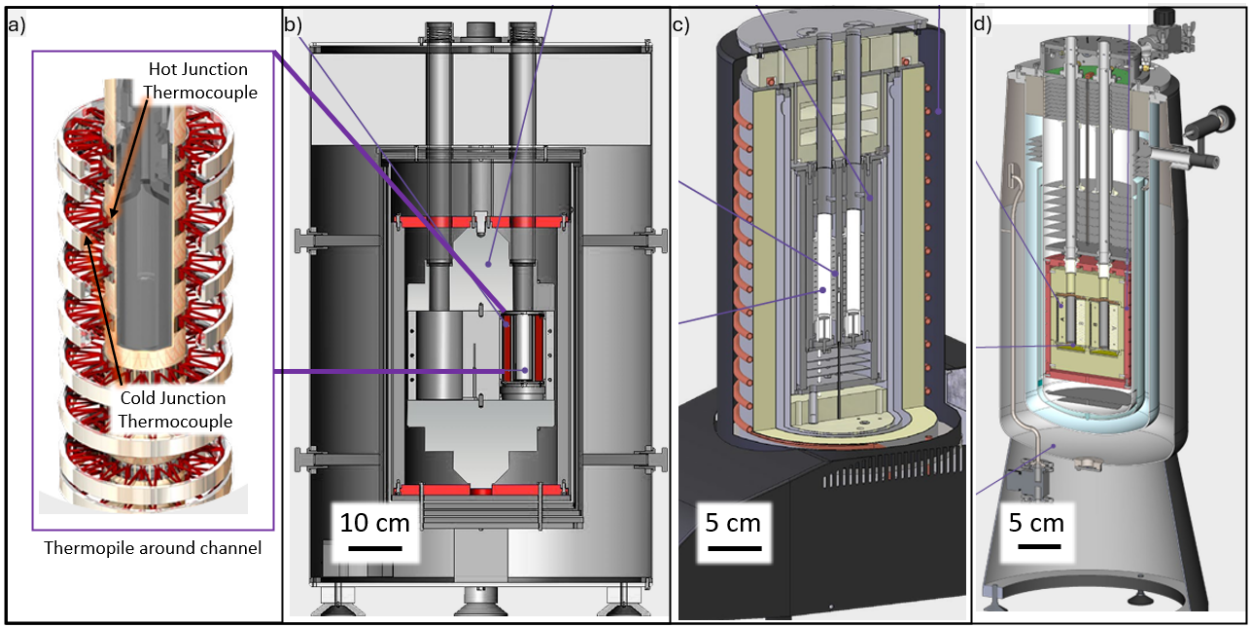


Figure 2: Cross sections of a) thermopile sensor, b) MS80, c) C600, d) BT215 with liquid N2 jacket around channels

Table 2: Calorimeter Specifications

Calorimeter Unit	Bore Diameter	Bore Length	# of Channels	Temperature Range	Sensitivity
MS80	35 mm	700 mm	4	30 - 200 °C	50 µW
C600	17 mm	350 mm	2	30-600 °C	2 mW
C80	17 mm	450 mm	2	30-300 °C	200 µW
BT215	17 mm	600 mm	2	-190-20 °C	200 µW

Methods

Methods : Samples

All samples used in calorimetry were focused on pure Pd in its various form factors that we could obtain or produce multiple grams of material. Below are the sample selection requirements we used in grading which samples were most feasible in our differential calorimeters.

- Quantity of available material to purchase should be > 5 g to allow for multiple experiments to be conducted.
- Surface area of material form factor should be maximized to reduce overall H/D loading times during experimental set up to channel isolation
- No mass should be able to escape the system throughout the duration of the experiment.
- Pd purity must be at least 99.95 at. % and ideally ≥99.99 at % pure.

The requirement for large mass negated the use of any of our very thin film geometries since the thermal mass of films is <<1 mg which would require large heat fluxes to form in a given Pd film, i.e., for >0.1 % of the Pd interstitial sites to form AHE events over a few seconds.

Foils and Wires

Pd foils of ≤250 μm were used for large sample mass experiments where we could fit >5 g into a single channel (Figure 4). High purity wires ranging from 20-250 μm in diameter were also used for some experimental runs to further increase surface area of the Pd (Figure 5). Bulk quantities of foils and wires were acquired from various material suppliers including Goodfellow, ACI Alloys, and California Fine Wire.

Nanoparticles

Significantly reducing the length scale of our Pd to where the ratio of surface area to volume was much higher was an important parameter to turn since there have been claims in LENR literature that nanoscale structures of Pd were important for loading and AHE. The size scale and geometry produced loading enhancements at the surface of the nanoparticle down to roughly 5 nm below which deterioration of absorption occurs[62,63]. Pd nanoparticles (NPs) were produced in-lab using Pd hydrogel to aerogel sublimation chemical processing. This utilized various precursor Pd salts and catalysts to induce Pd precipitation into a fine Pd nanostructure[64,65,66]. We were also able to obtain larger quantities of Pd nanoparticles with similar reaction chemistry from Goodfellow allowing us to install much larger masses >1 g of nanoparticles in each channel. The nominal diameter for nanoparticles within a chain-like structure were <20 nm in diameter with 99.99 at %

[62] Eduardo A. Crespo, Margarita Ruda, Susana Ramos de Debiaggi, Eduardo M. Bringa, Fabián U. Braschi, Graciela Bertolino, Hydrogen absorption in Pd nanoparticles of different shapes,
International Journal of Hydrogen Energy, Volume 37, Issue 19, 2012, Pages 14831-14837, ISSN 0360-3199

[63] Sachs, C, Pundt, A, Kirchheim, R, Winter, M, Reetz, M T, and Fritsch, D. "Solubility of hydrogen in single-sized palladium clusters". 2001

[64] Burpo FJ, Nagelli EA, Morris LA, McClure JP, Ryu MY, Palmer JL. A Rapid Synthesis Method for Au, Pd, and Pt Aerogels Via Direct Solution-Based Reduction. J Vis Exp. 2018 Jun 18;(136):57875. doi: 10.3791/57875. PMID: 29985323; PMCID: PMC6101767.

[65] Burpo, F., Nagelli, E., Morris, L., McClure, J., Ryu, M., & Palmer, J. (2017). Direct solution-based reduction synthesis of Au, Pd, and Pt aerogels. Journal of Materials Research, 32(22), 4153-4165. doi:10.1557/jmr.2017.412

[66] Du R, Joswig JO, Hübner R, Zhou L, Wei W, Hu Y, Eychmüller A. Freeze-Thaw-Promoted Fabrication of Clean and Hierarchically Structured Noble-Metal Aerogels for Electrocatalysis and Photoelectrocatalysis. Angew Chem Int Ed Engl. 2020 May 18;59(21):8293-8300. doi: 10.1002/anie.201916484. Epub 2020 Apr 6. PMID: 32187791; PMCID: PMC7317422.

purity claimed. 99.9% purity was verified with in-house resources including energy dispersive x-ray spectroscopy.

Methods : Sample Retention for each Campaign

Ultrapressure Vessel Design

The hardware component that made all calorimeter experiments possible at these extreme pressure conditions was the cylindrical pressure vessel. The vessel was optimized to ensure containment of hydrogen-based gases, and therefore the chosen material, Toughmet, was selected to avoid embrittlement from exposure to hydrogen. A double compression seal was designed for both the small and large bore vessel types. The smaller vessels used a "football" type Toughmet insert that sealed against both the vessel inlet and a BeCu downtube that provided gas access all the way down the calorimeter bore length (Figure 3a). Once at ultrapressure conditions each channel used ultrapressure valves at the top of each channel on the outside of the calorimeter to isolate the channel from the larger pressure source to ensure that energy and mass transfer is minimized along that pathway. The larger pressure vessel used what was called internally as, a UFO seal, sealing the vessel cap to the main vessel body. The vessel cap then used a similar sealing methodology as smaller vessels to connect to the BeCu downtube gas inlet (Figure 3b). For both sets of hardware all sealing surfaces were captured in a leak detection system that would alert the user that a given sealing surface was leaking. This is crucial because a small leak is manifested in the calorimeter

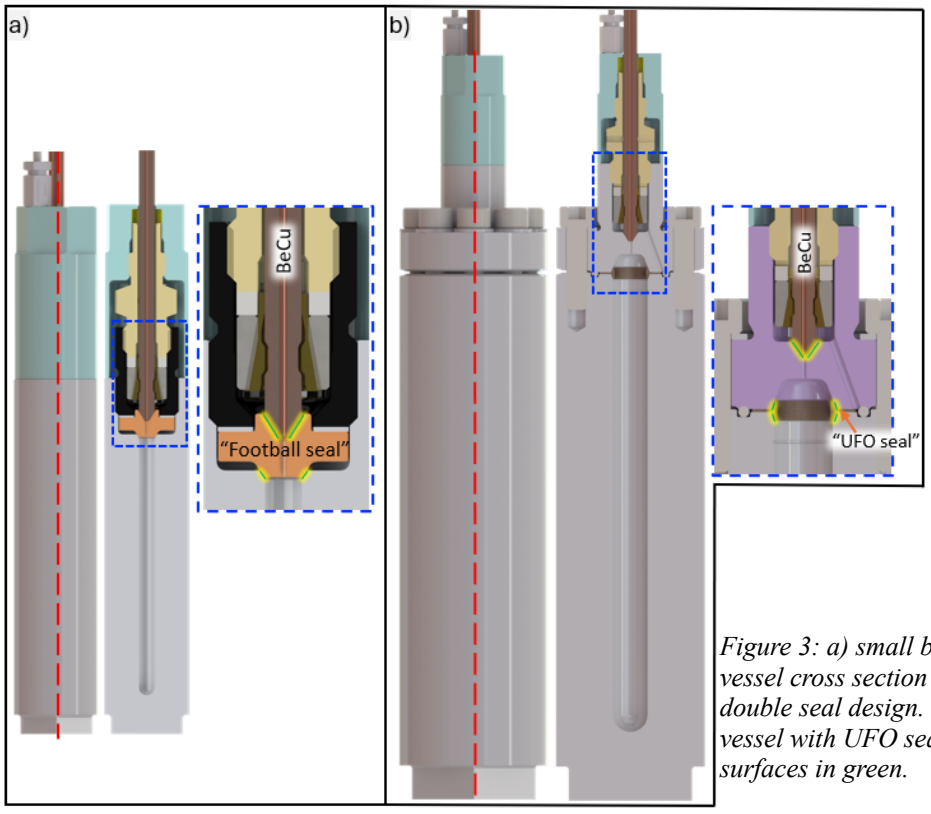


Figure 3: a) small bore pressure vessel cross section showing double seal design. b) large bore vessel with UFO seal. Sealing surfaces in green.

as heat flow, and so the auxiliary leak detection system was able to quickly confirm or deny if a transient heat flow signal was pressure based. All channels were monitored by ultrapressure transducers during the entire course of the experiment.

Passive Foil/Wire

The foils and wires were held loosely within a thick-walled pressure vessel that was then inserted into each channel of the calorimeter. Foils were rolled up tightly with a gold mesh sheet separating foil surfaces from one another to allow gas access and not slow down protium/deuterium loading kinetics at each surface (Figure 4). In other experiments, microwires of 20 μm diameter were bundled together in packs of 2000 to maximize the mass that could be placed inside a pressure vessel (Figure 5).

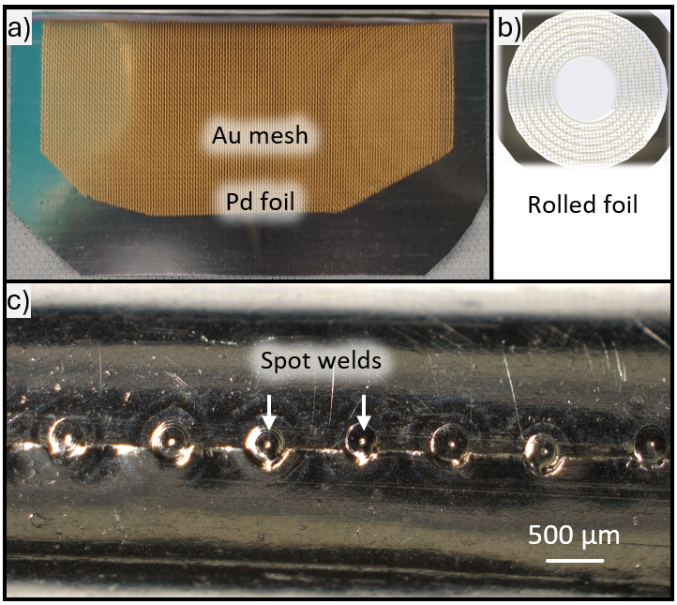


Figure 4: a) Rolling process with mesh interlayer, b) axial view of rolled sample, c) spot welds of foil seem to prevent unraveling.

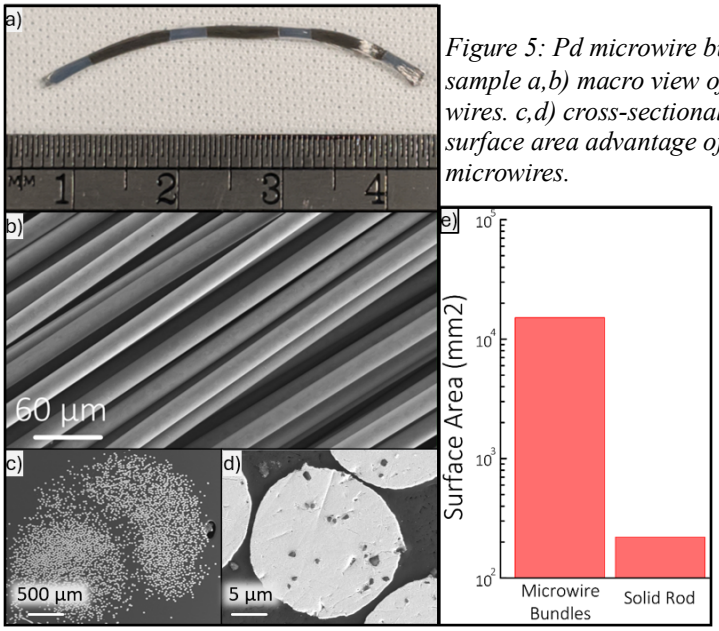


Figure 5: Pd microwire bundle sample a,b) macro view of 20 µm wires. c,d) cross-sectional views. e) surface area advantage of microwires.

Passive Nanoparticles retention

Due to nanoparticles being inherently small, they are prone to becoming airborne which can be dangerous for users in proximity to the experiment and could easily contaminate many instruments, including the calorimeter, increasing uncertainty in future experiments. We are also inserting nanoparticles into a high-pressure environment, that could leak catastrophically, worsening this contamination and health issue. The retention designed for passively monitoring nanoparticles in our MS80 calorimeter can be seen in Figure 6. This stainless-steel tube was

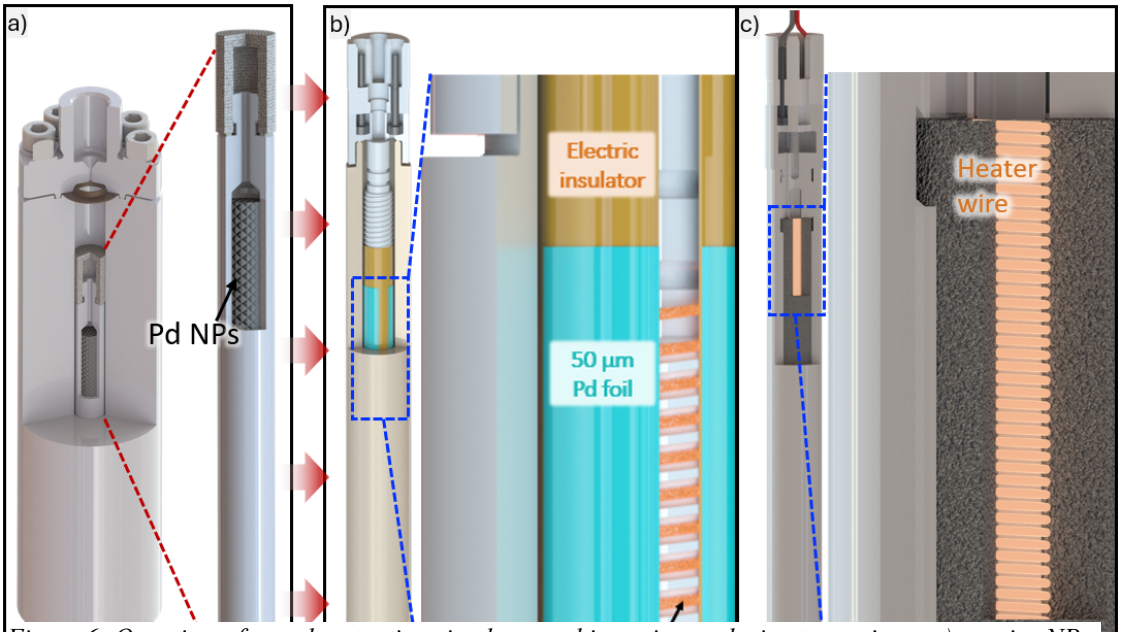


Figure 6: Overview of sample retentions implemented in various calorimeter projects. a) passive NP. b) foil/wire stimulation. c) NP stimulation.

closed on one end with a compression seal and had a porous stainless steel gas filter threaded on the other. The average pore size of this filter was 0.5 µm with a 1 mm thick cross-section making escapement of particles through the filter difficult without becoming entrapped. Protium and deuterium flow into the retention was still possible allowing nanoparticle loading and to prevent inadvertent formation of uncontrolled pressurized volume within the vessel. This minimizes risk for personnel when handling material and any damage to instruments near or internal to the calorimeter.

Stimulated Foils/Nanoparticles

Thermally stimulating foils and NPs required the same type of heater wire composed of 150 µm diameter NiCr80. For the foil retention, bare NiCr wire was spooled onto a 2-start ACME thread design to maximize the length of wire used and ensure wires entered and exited the retention on both sides of the vessel (Figure 6). Kapton film (25 µm thick) or bare fiberglass fibers were inserted in between the Pd foil and the heater wire to prevent shorting.

Fiberglass insulated NiCr wire was used in the NP retention to maximize surface contact of the NPs to the heater spool in the center (Figure 6c). Similar to the passive NP retention, only one end of the retention was open to gas and also contained a porous titanium sheet with pore size of 1 µm to minimize NP escapement throughout the duration of the experiment.

High Pressure electrochemistry

High pressure electrochemical loading inside the differential calorimeter required improvements to electrical access to each pressure vessel. To account for all energy inputs going into the vessel for each channel a total of eight isolated Pt wires were used in ultrapressure electrical feedthroughs. These feedthroughs are identical in design concept to epoxy feedthroughs discussed in the hardware section of this report, with the additional constraint of sealing around a 5/16” inner diameter instead of a ½” or ¼” diameter used in later iterations of 500 ksi HPE experiments that

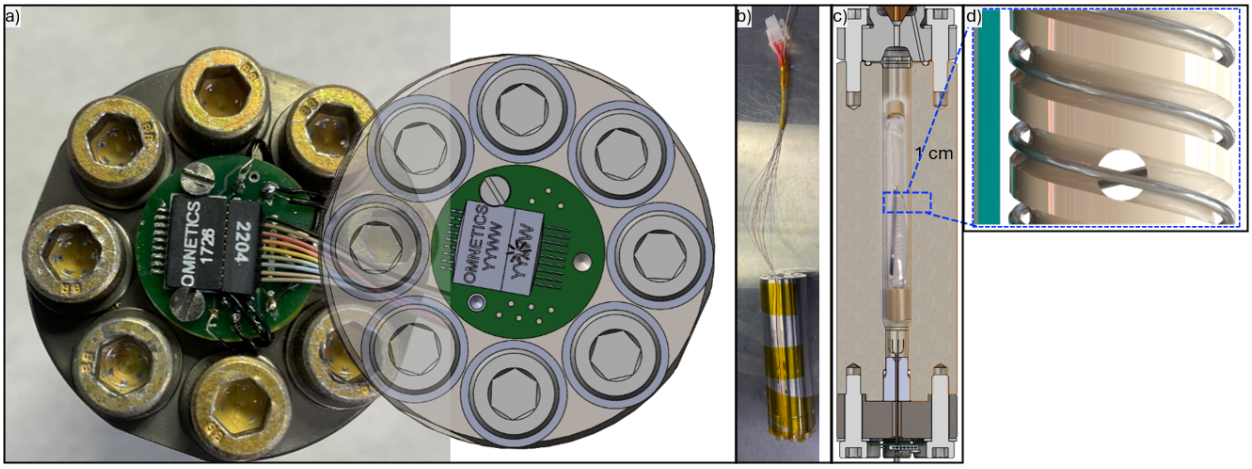


Figure 7: a) Looking at the bottom of the HPE calorimeter vessel with custom PCB connector and wires that were fed along a track on the side of the vessel. b) Overview image of wires fed along outside of vessel. c) vessel cross section with detail d) showing retention of electrodes (Pd is the foil, Pt the spiral wire).

were not carried out in a closed calorimeter enironment. These feedthroughs carried the constant current voltage across the Pd cathode and the Pt anode, *in-situ* 4-wire RTD in contact with electrolyte, *in-situ* 4-wire resistance Pd circuit for H/D loading montoring, and voltmeter circuit. All circuits were terminated at the bottom of the vessel as seen in Figure 7 where a custom PCB connector was used to maintain low profile and maximize vessel wall contact with calorimeter channel wall. The entire wired vessel assembly is seen in Figure 7b,c. A teflon mesh was placed between the Pd and Pt electrodes to ensure there were no short circuits during the experiment.

Results

Results : Passive calorimetry : At 50ºC to 150°C

Experiments completed where we kept the pressure and temperature constant, once the calorimeter was pressurized, were split up according to sample form factor and set temperature. Experiments completed successfully always maintained a 100 ksi or 150 ksi gas pressure internal to each channel. Table 3 summarizes the various experimental conditions we conducted on Pd samples with each experiment, on average, remaining in an isolated, stable state for ≥4 days. All tests were completed with a protium control in 1 of the 2 channels.

Table 3: Passive Calorimetry Parameter Space

		Temperature (°C)					
		-100	-50	50	100	150	200*
Pd Form factor	Foil						
	Wire						
	3µm powder						
	NPs (15 nm)						

* No other samples were tested at 200 °C due to damage of the high temperature calorimeter (C600)

Interpretation of the data focused on any excess-heat measurements that would then need to be thoroughly examined through instrument, calibration, user error, or classical explanations, before postulating an AHE explanation. The data interpretation relies heavily on establishing a 'baseline', with deviations from this baseline indicating that one channel of the calorimeter is hotter or colder than the other. The detection resolution of the experiments is determined by the magnitude of these baseline fluctuations. The resolution for most passive experiments was defined by upper and lower limits defined by +/- 2σ from the baseline average. For heat flow data reported in this section the

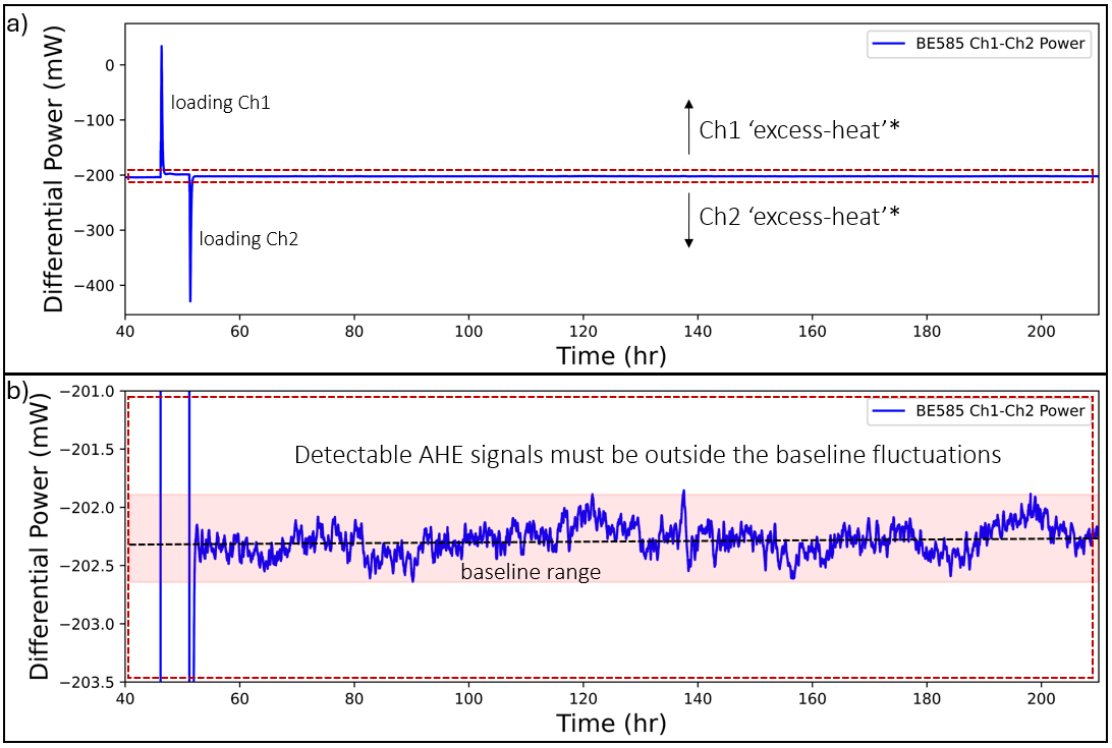


Figure 8: a) Overview of calorimeter heat flow in a typical experiment. Loading itself produces heat flow from H/D loading into the Pd lattice and is ignored for analysis. b) Heat flow axis is expanded to only the baseline range where the calorimeter stabilized after gas isolation. No abnormal signals detected.

units are in mW (converted from thermopile raw signal collected in μV).

Most experiments did not show anomalous behavior beyond the baseline fluctuations and therefore did not warrant further analysis and investigation of spurious heat flow signals. For calorimeter experiments using foils and wires between 50 and 100 °C the heat flow data was extremely stable (Figure 8). Fluctuations within baseline are small temperature gradients across the large calorimeter thermal mass which are <0.25 °C and occur over 24-hour timescales, much longer than any short-term releases of heat within a given channel. Anomalous signals observed throughout all passive calorimetry fell within the categories listed in Table 4.

Table 4: False Signal Sources in Passive Calorimetry

Cause of False signal	Resulting Heat Flow Profile	Controls
Pressure loss (<1000 psi) in 1 or more channels	Gradual but consistent diversion from heat flow baseline, 'cooling' the channel that is leaking through mass loss	• Monitored by pressure transducers
Catastrophic pressure loss (>1000 psi)	Sharp diversion from heat flow baseline, 'cooling' the channel that is leaking through convection and mass loss	• Monitored by pressure transducers • Internal channel RTDs
Heater collar failure	Baseline shift over 3-4 hours then stabilizing	• >4 RTD probes on heater collar looking for temperature gradients across the collar
Calorimeter enclosure heater failure	Baseline range 2-5x larger and fluctuates cyclically over 24 hour period	• >4 RTD probes at various locations on cart external to calorimeter
Calorimeter heater power shifts	Baseline shift over 3-10 hours then stabilizing	• Recording input power levels from calorimeter heater with Setaram software
Liquid nitrogen source is empty	Heat flow drifts uncontrollably until stabilizes with room temperature	• Liquid nitrogen reservoir liquid level probe • RTDs in reservoir • Facility tank fill % recorded every minute

Prior to operating matured calorimetry experiments that maintained exceedingly good thermal stability and isolation from external variables outside the calorimeter system, one experiment produced a heat flow signal indicating the PdD was releasing heat. This test was conducted at 100 ksi deuterium and at 50 °C with no other thermal stimulation, a completely passive test. The differential heat flow is plotted in Figure 9Figure 9: a) Heat flow data showing numerous spikes in the direction that indicates PdD emitting energy (baseline noise subtracted). b) Spread in peak power amplitude reached for a given heat flow spike, not all peaks are of same magnitude. The spikes in heat flow were sporadic throughout the duration of the test, not occurring at specific times of the day or in line with variations in pressure change or external RTD variation (see Appendix for further data streams).

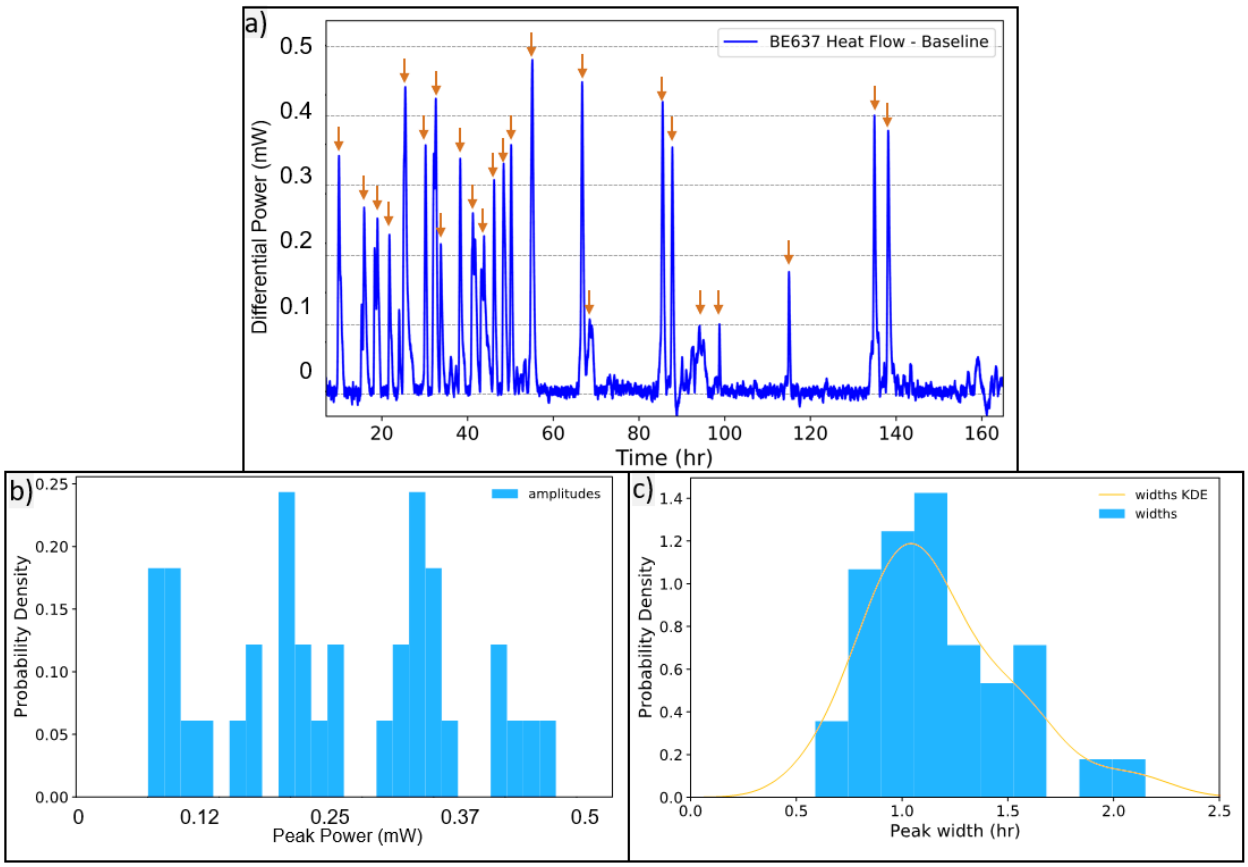


Figure 9: a) Heat flow data showing numerous spikes in the direction that indicates PdD emitting energy (baseline noise subtracted). b) Spread in peak power amplitude reached for a given heat flow spike, not all peaks are of same magnitude. c) Spread in how long the peak forms and decays over time showing that there is variable duration of power for each event.

This signal never formed again within any of our calorimeter systems including the same MS80 that this data was collected.

At the time this signal was observed we did not have a fully enclosed thermal enclosure or minor forms of vibration dampening to the calorimeter which could lead to extraneous influence on the heat flow data that was not detected by auxiliary sensors. Extensive post characterization on the Pd sample was completed, with no abnormal surface morphologies or compositional variation measured compared to the source material. We did find evidence that minor movements of the channel downtube would produce large transient behavior over short time scales of 10-30 min in thermopile voltage signal, but we could never definitively locate the source of this heat flow signal. And after iterating and reproducing upwards of dozens of these passive calorimetry experiments with stable heat flow results, we concluded this signal was not relevant to observing AHE or characterizing its behavior because it could not be replicated. Further examples of other false signals can be found on Synergy under the Calorimetry Project.

Results : Passive calorimetry : At -150ºC to -50°C

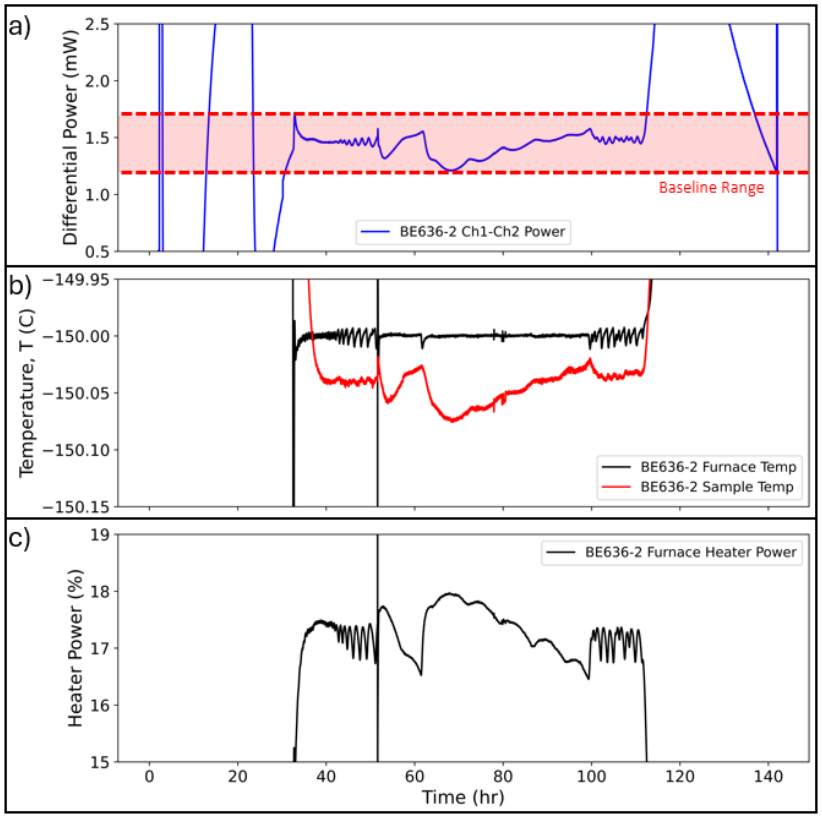


Figure 10:a) Heat flow axis showing small fluctuations over 140 hours related to calorimeter heater and LN2 cooling instabilities. b) RTDs of samples and furnace (heater) to maintain -150 °C. c) Power input from BT215 converted to % of total constant current DC power.

Performing differential calorimetry at these temperature regimes in both bulk and NP forms helped ensure we were maximizing our loading in the Pd lattice, and (because of very low surface kinetics and deuterium diffusion) locked in the Pd loading once isolated and stabilized at these temperatures[67]. The Pd was loaded at room temperature and 150 ksi protium/deuterium prior to cooling down to final set temperature to ensure the kinetics timescales of the loading process would not exponentially decrease and hence not fully load when <-50 °C. Figure 10 shows a typical heat flow result from a cryogenic calorimetry run at -150 °C where we see no significant deviations from the baseline readout between the control channel and the active PdD channel. The short-term deviations over 10-20 min timescales are from the BT215 heater fluctuation power to maintain -150 °C, and was confirmed by noting that the heater power changed 20-30 seconds before we start to see the heat flow response clearly indicating this heat flow change is not caused by internal channel sample reactions. In Appendix Figure 19, another calorimeter experiment is shown where we kept baseline fluctuations to <0.2 mW and still observed no unexplained anomalies in the readout. NPs in passive experiments showed similar results to foils and wires. Images and additional details on NPs based experiments can be found in the Appendix.

Results : Passive calorimetry : At 200°C

When elevating the set temperature, the equilibrium loading process results in a reduced peak loading of the Pd lattice; however, the deuterons will have a higher average kinetic energy for

[67] F. Scaramuzzi, Gas loading of deuterium in palladium at low temperature, Journal of Alloys and Compounds, Volume 385, Issues 1–2, 2004, Pages 19-27, ISSN 0925-8388.

interstitial site hopping, testing the hypothesis that increased lattice site hopping increases probability for 2 deuterons to meet and fuse at a single site. The highest temperature that we could achieve while at 100 ksi and inside a closed calorimeter was 200 °C inside the C600 unit. Out of the 2-3 tests in which we collected data at 150-200 °C we did not observe any unexplained heat flow fluctuations in our data. Unfortunately, we could not run further tests at these levels safely because we found that the pressure vessel material, holding the sample and gas inside the calorimeter channel, fails catastrophically from high temperature embrittlement of the Toughmet material at elevated temperatures (Figure 11). Through repeat vessel failure experiments, external to a calorimeter, we found these failures occur regardless of gas type, i.e., also occurred in He. Since this failure effectively can bisect a vessel and form projectiles, we discontinued our efforts in probing high temperature holds over long time periods with these vessels.

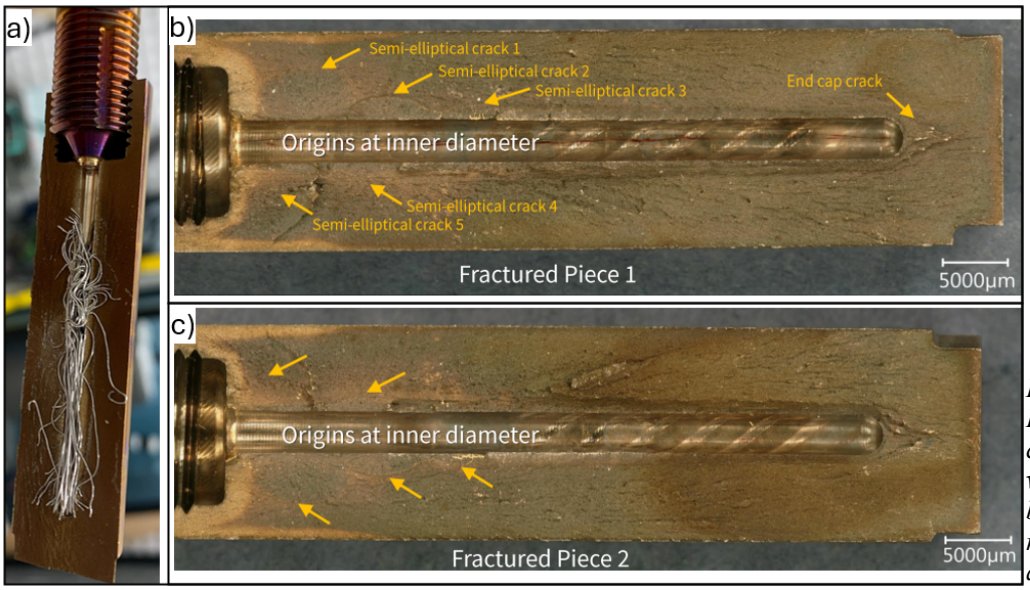


Figure 11: a) Ruptured C600 channel vessel with Pd wires inside. b) and c) both halves of a vessel rupturing axially with a brittle fracture.

Results : Stimulated calorimetry : Foils and Nanoparticles at 50ºC to 100°C

Up to this point the results we've described have only been looking at loading and set temperature as variables that could induce AHE to occur spontaneously. In this and remaining results sections we are actively adding energy to the Pd system through conduction heating means in order to attempt to knock the system out of equilibrium. Adding energy to the system over a few seconds timescale is testing if inducing loading fluxes/deloading to the sample could increase short term deuteron site hopping enough to cause D-D fusion.

For us to understand if we are producing more energy than we are putting in we are utilizing the COP_{energy} calculation discussed in the concept section, and we are adding an additional variable to the energy balance of the system seen in Equation 3. As a reminder, the heating is from a NiCr wire so the DC input power can be calculated using current (I) and voltage (V). Heating these wires in close proximity to the sample, we were able to heat our Pd samples to at least 350°C and above 450 °C at the surface of Pd in contact with wires based on COMSOL simulations (Figure 13). Equations 3 and 4 are only for stimulating a single channel in the calorimeter, so do not treat the differential signal; however, for us to heat the sample up to >350 °C we would saturate the heat flow signal to the upper or lower range of the thermopile sensor. Therefore, we performed a separate series of experiments where we applied the same input power to both channels and ideally subtract out the input power signal since it is a differential calorimeter (Equation 5 and 6). This is not an ideal system, so we had to account for about 2 % uncertainty in the COP_{energy} calculation.

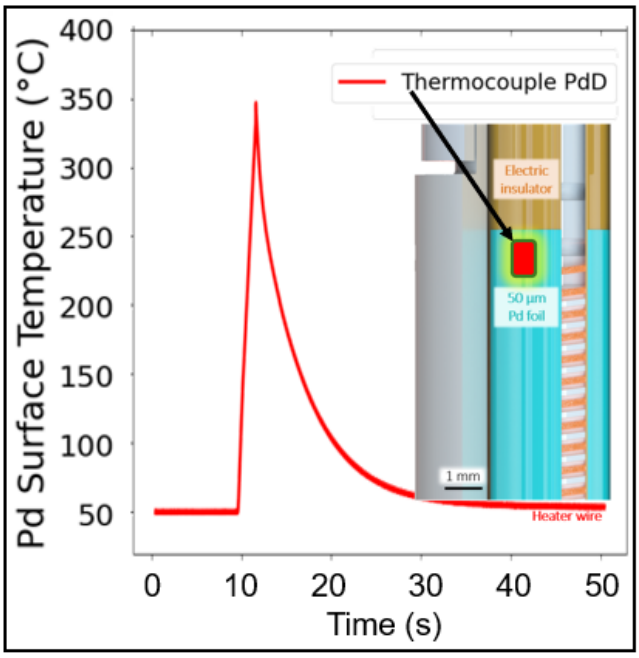


Figure 13: In-situ thermocouple readout during thermal pulse of wire heater. CAD model in insert used in COMSOL heat transfer simulation.

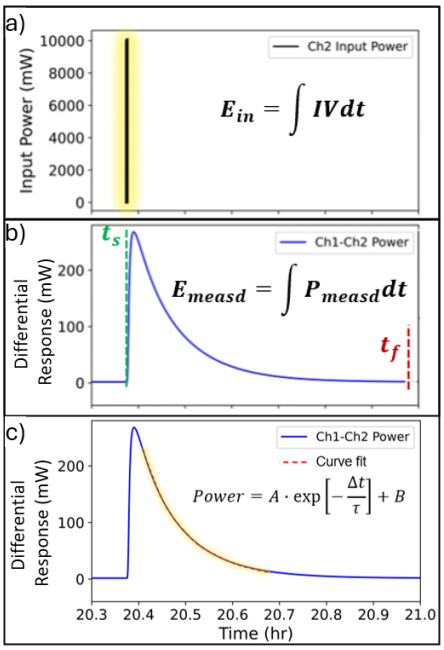


Figure 12: a) Electrical pulse sent through heater wire. b) Calorimeter heat flow response. c) Exponential decay fit of heat flow.

Eqn 3

$$\Delta E_{1ch} = \left[m_{control} \int c_{v_{PdH}} dT + \Delta H_{loading} + \int IV_{PdH} dt\right] - \left[m_{active} \int c_{v_{PdD}} dT + \Delta H_{loading}\right]$$

Eqn 4

$$E_{measd} = \int P_{measd} dt \ , \quad COP_{energy} = \frac{E_{measd}}{E_{1ch}}$$

Eqn 5

$$\Delta E_{2ch} = \left[m_{control} \int c_{v_{PdH}} dT + \Delta H_{loading} + \int IV_{PdH} dt\right] - \left[m_{active} \int c_{v_{PdD}} dT + \Delta H_{loading} + \int IV_{PdD} dt\right]$$

Eqn 6

$$COP_{energy} = \frac{E_{measd}}{E_{2ch}}$$

The experimental results presented are analyzing the result's COP_{energy} value and determining if it is significantly above the value of 1 indicating excess heat in one channel, and calculating the thermal time constant from the calorimeter response in heat flow for each thermal stimulation pulse. Changes in the thermal time constant can indicate changes to the conductive path between the differential channels in the calorimeter. Monitoring if the conductive path changes during a test is important to ensure our thermal calibrations are still applicable. Figure 14 shows the COP and time constant results for Pd Foils stimulated in one channel at a time. These results give no indication of increasing thermal time constant, and COP remains below 1, indicating we are not observing excess heat from this stimulation method while at 150 ksi. We are not able to capture

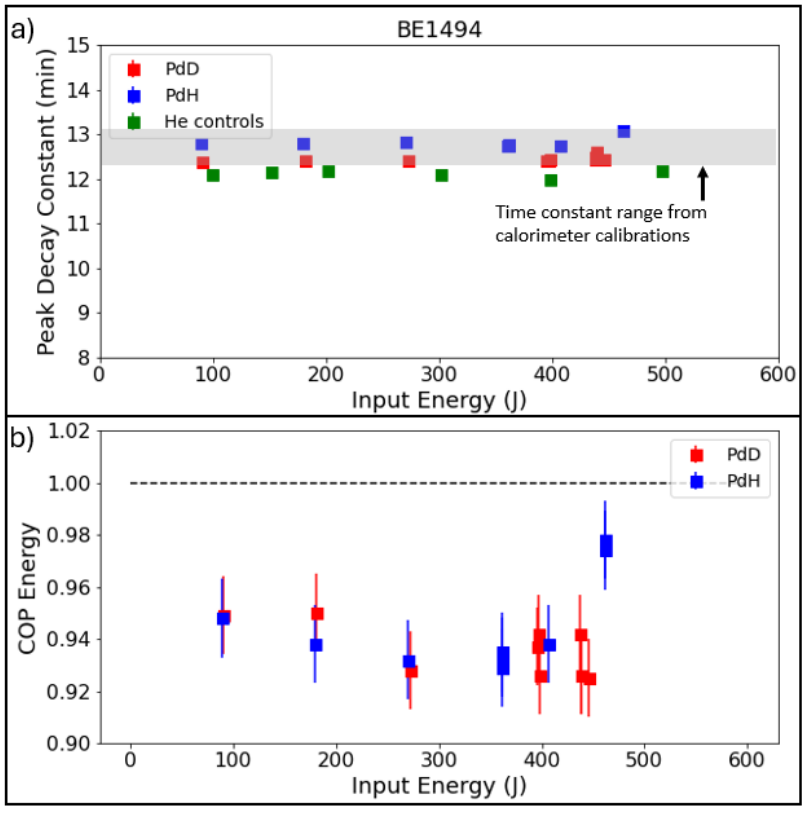


Figure 14: a) Calculated time constants from various levels of thermal stimulation energies. b) COP energy for each stimulation condition with a dashed line threshold set at 1.0.

100% of the input energy from the heater wire with the thermopiles, which is why we are below a COP of 1. One major reason for this is that we had over 6 feet of fine gauge copper leads running down the channel bores inside the calorimeter that are not within the thermopile sensor area, and therefore any minor joule heating from these wires was not captured inside the closed system.

Thermally stimulating NPs in both protium and deuterium was conducted similarly to foils with the main difference being a separate retention to safely handle and test the NPs. The results in Figure 15 show a similar story for both 3 µm Pd particles/ grains and 15 nm NPs where the time constant and COP are stable across the range of thermal input energies we used. All COPs in all experimental runs remained below 1 and the results between the PdH control and PdD were with experimental uncertainty of the measurement. No AHE signals detected.

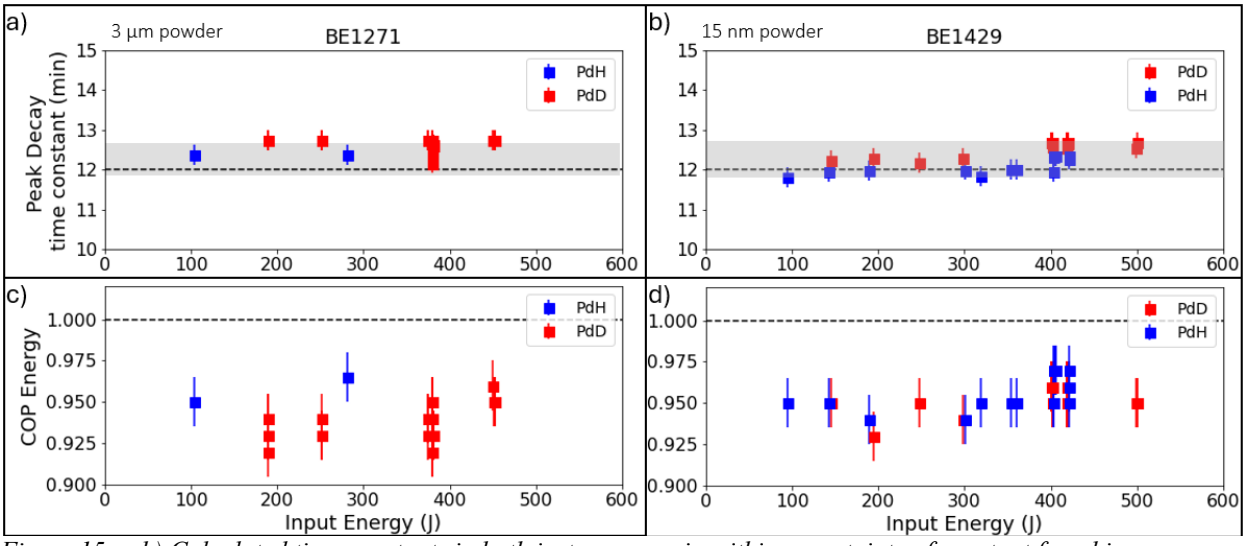


Figure 15: a,b) Calculated time constants in both isotopes remain within uncertainty of constant found in calibrations. c,d) COPs for various stim conditions, none above 1.0.

Results : Stimulated calorimetry : Wires and Nanoparticles at -100ºC to -50°C

Cryogenic campaigns with identical forms for thermal stimulation tested the same hypothesis as stated in the previous section, i.e., that loading fluxes across a sample lattice might induce D-D fusion. Below are COP results collected at -100 and -50 °C calorimeter set temperatures, points at which the initial loading condition of the Pd wire and or NPs should be at or near unity in protium and deuterium (Figure 16). Results show we did not observe excess heat at these cryogenic conditions and under these thermal impulse parameters.

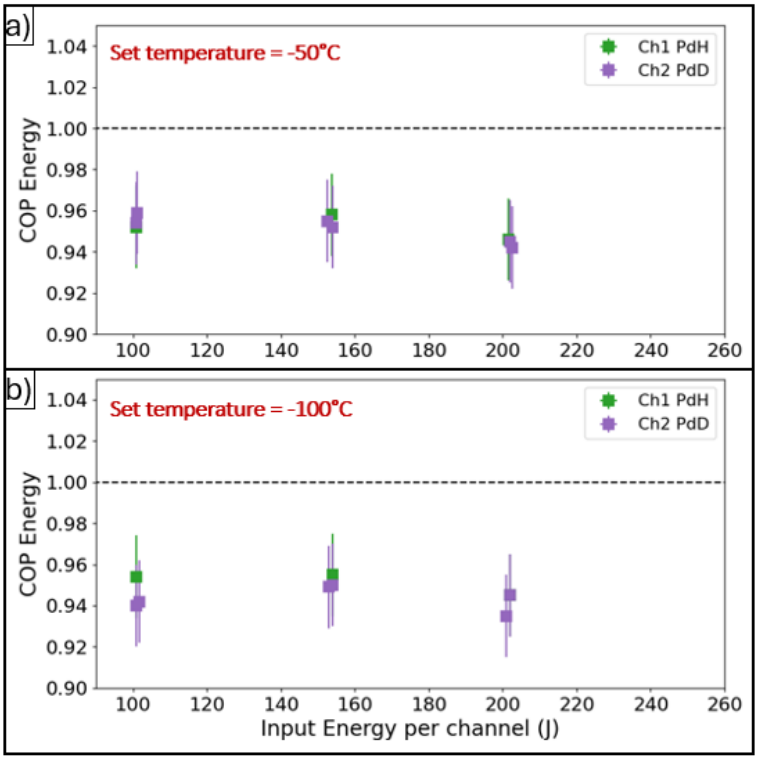


Figure 16: a,b) COP results at -50 and -100 °C set temperatures where no stimulation conditions result in a COP >1.

Results : HPE calorimetry :– At 35 to 100 °C

The materials and electrolytes used for this set of calorimeter experiments are identical to ones used in the general high pressure electrochemical experiments conducted between 150-500 ksi. For more detail on electronics used and current densities used please refer to the HPE section of this report. Both heavy and light water HPE experiments were carried out inside the MS80 calorimeter. The main salt used employed in the high purity water was LiOH/LiOD depending on which isotope was being tested in a given channel vessel. As shown in figures in the sample retention section the input power and other diagnostic read out circuits gained access to the ultrapressure liquid environment through the bottom of the vessel inside each channel of the MS80. The downtube supplied the pressure to the system utilizing a screwjack pressure system external to the thermally isolated environment around the calorimeter. For each test the calorimeter channels were first electrochemically loaded with an initial constant current set point between 5-15 mA. This was completed with full submersion of the electrodes in the electrolyte, and allowed for faster initial loading of the Pd. This section of heat flow data is not analyzed since the sample is not at pressure and the calorimeter is in a transient state making analysis in terms of measured power highly inaccurate during this time period. Once the sample reaches near the peak R/R0 value (i.e., 2.0 for PdD) the pressure ramp is started for both channels at the same time and rate. The dual channel press-up is only possible when we are using Pt as the control and are able to use a single water isotope for a given experiment. The screwjack is hooked up to each channel in parallel,

effectively tying both channel volumes together. It should be noted the heat flow loss through the path between the 2 channel downtubes is negligible since the volumes are only connected to each other by a 750 μm diameter BeCu tube of over 2 m in total length, making the conduction path exceedingly small in both cross-sectional area and overall electrolyte fluid conductance over that length scale. The only section of heat flow data analyzed was the portion after a total pressure between 150-200 ksi was reached and isolated. This pressure is typically held in this equilibrium state for 20-40 hours for a given experiment.

To avoid isotopic contamination a separate screwjack system was dedicated to a given water isotope, heavy or light water. The control sample utilized for each calorimeter had a cathode made of Pt with similar surface area, mass, and surface finish to the Pd foil used in the active channel. Making the control a Pt-Pt system allowed us to monitor heat flow formation that may occur due to small variations in Pd loading over the course of the equilibrium period of the experiment. Aside from loading and measured power, other data streams collected to account for all input power signals include: variation in lead resistance and source/sense circuits for the in-situ RTD, all of which have measurable levels of input power to the electrochemical system inside the thermopile area or are counted as power loss on the outside of the calorimeter.

A clean set of HPE calorimeter experiments can be seen in Figure 17 where one is using heavy water and the second light water. Both experiments are using close to the same current density set points with an input current of -30 mA. Accounting for all measurable sources of input power and power losses the measured power (red or blue) is slightly below the input power (green) over the majority of the experimental time window displayed. Points where we see input power drop below measured power levels translate to a COP_{power} <1.05. This COP_{power} level is below the threshold to be statistically significant and auxiliary data streams including abrupt voltage changes, pressure variation, and external RTD variation occur seconds before these increases in heat flow indicating this is a system response to small changes in equilibrium between system and surroundings. No measurable levels of excess heat at either 35 °C or as high as 100°C were observed during these experimental runs. More detail can be found on the Synergy experimental pages and Appendix for this section.

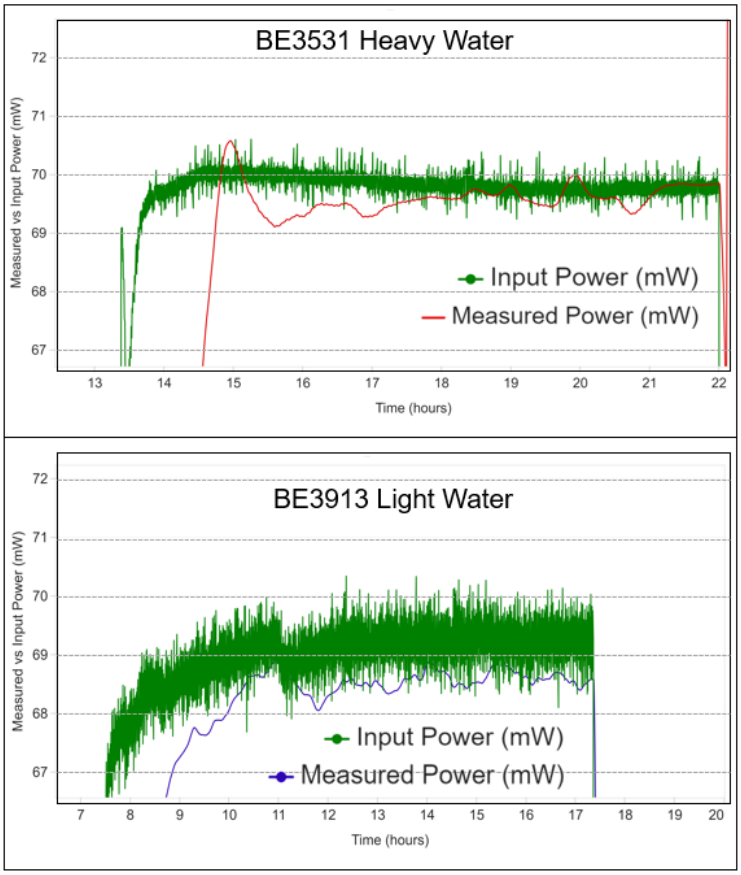


Figure 17: Light and heavy water electrochemical heat flow data overlayed with the total input power to the system. The input power remains above measured power due to the calorimeter being a real system with unaccounted for heat leaks.

During the course of operational HPE with these calorimeters we did observe heat flow signals that turned out to be false positives. These signals formed when there was an abrupt voltage change/instability, usually a drop of a few volts at constant current, where then the heat flow settled back to equilibrium after 3-4 time constants measuring power that was significantly above input power levels being supplied to the system (Figure 18). This effect was reproducible from experiment to experiment and more importantly could be replicated in light water systems. This inversion between input and measured power could also be produced just using Pt electrodes without Pd altogether making it clear this has nothing to do with lattice mediated fusion. At the time the experimental results were collected, we were still utilizing an electrolyte solution that contained heavy or light forms of methanol. This low molarity of methanol was used to stop the electrolyte from freezing at ultrapressure conditions >100 ksi while at 35 °C.

After troubleshooting possible errors in input power accounting to explain this result it was decided to increase the set temperature to 50 °C or above and remove the methanol additive altogether to remove it as a possible variable. With the methanol removed this phenomenon, increased measured power after a voltage instability, disappeared from the heat flow data in both light and heavy water experiments at ultrapressure conditions (Figure 17). The hypothesis at the time of this report is that

the methanol was possibly reducing in some fashion, forming a side exothermic reaction that was only occurring in the electrolyte. This would be a more classic chemical reaction that was reaching its own equilibrium reaction rate explaining the sustained increase in measured power over time.

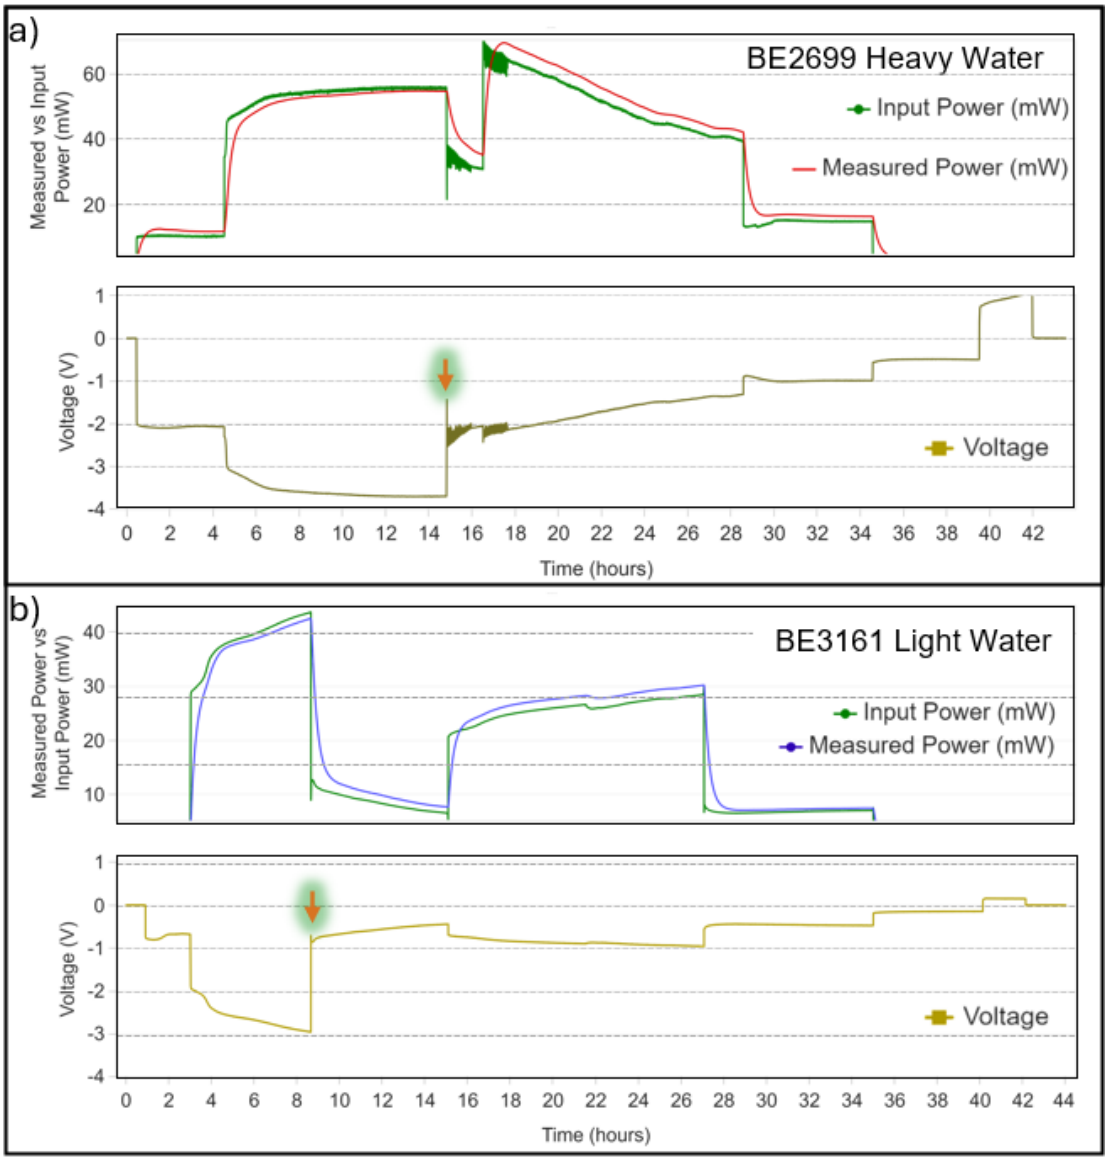


Figure 18: a) Heavy water HPE experiment showing a voltage instability form between 14-15 hours after which measured power remains above the calculated input power. b) Same results but in Light water and with only Pt electrodes.

Conclusion

While maintaining exceptional thermal and electrical isolation of the differential calorimeter channels, the parameter space explored in these calorimeter studies has not produced any definitive signal of LENR based AHE phenomena. Through passive calorimetry experiments probing set temperatures from -150-200 °C at gas pressures up to 150 ksi, spontaneous lattice enhanced D-D fusion (i.e., using unstimulated PdD) could not be detected through the bulk direct heat flow detection mechanism used in calorimetry. Stimulated Pd samples, regardless of sample form factor were heated up to 350 °C in near unity loading conditions with no excess heat measured by any of the Setaram calorimeters employed. Under any of the experimental conditions completed during this campaign, no environment was produced in either equilibrium or highly transient states that delivered quantifiable levels of AHE.

Calorimetry Appendix

This decision to directly interpret the voltage potential read out was made for the following reasons. Firstly, the system was a closed-type calorimeter with no additional energy or mass inputs once the channels were isolated at 100 or 150 ksi pressure with ultrapressure valves located on the outside of the calorimeter. As stated in the concept section the absence of additional inputs to the system we are monitoring minimizes the noise measured by each thermopile, maximizing the signal to noise ratio for actual sample-based reactions or perturbations. The raw potential signal, based off a known, stable baseline signal, can be interpreted directly to indicate if either channel's heat flux, control or active, is increasing or decreasing over time. Secondly, converting the potential signal to power using the thermal calibration coefficient [68]only provides a scaling factor to the already existing profile in the heat flux (μV) data since we used constants based on linear calibrations. The conversion to power becomes truly important when we do see a signal and we want to understand how much energy over time was released or absorbed into one of the sample channels.

[68] Thermal calibration coefficient (*K):* ($\mu V.mW^{-1}$) is used to convert electrical signal from the heat flux sensors (thermopiles) to the differential power inside the system. This study used a liner coefficient which was sufficient to assume the calibration as linear across the full range of the thermopile (-20000 to 20000 μV).

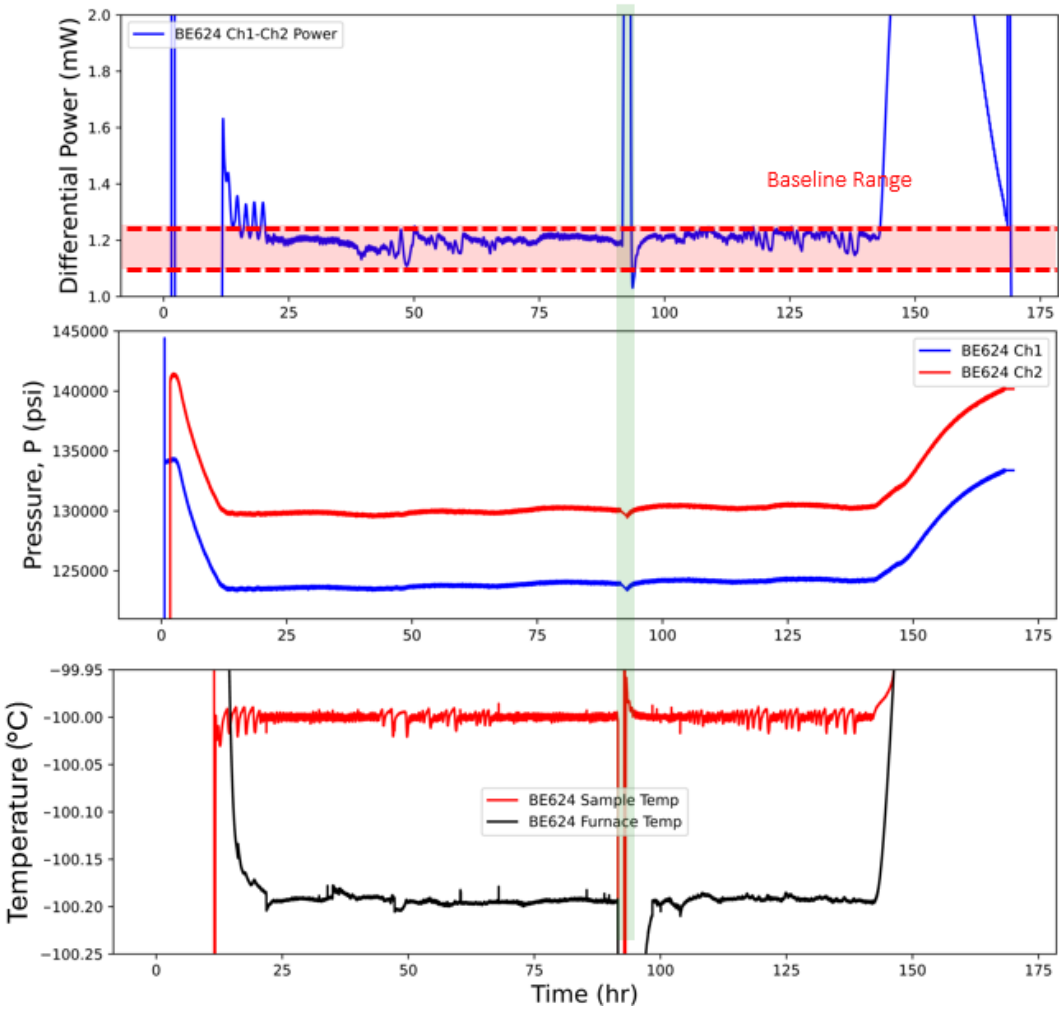


Figure 19: Heat flow results from another cryo-calorimtery test at 130 ksi in both H and D gases. Pressure decreases during cooling, explaining the drop from 140 to 130 ksi. At 90 hours the calorimeter heater shut off leading to a temperature dip to -150 °C, this was caught and corrected within 5 hours of the event. Sensitivity in this passive test was <0.5 mW.

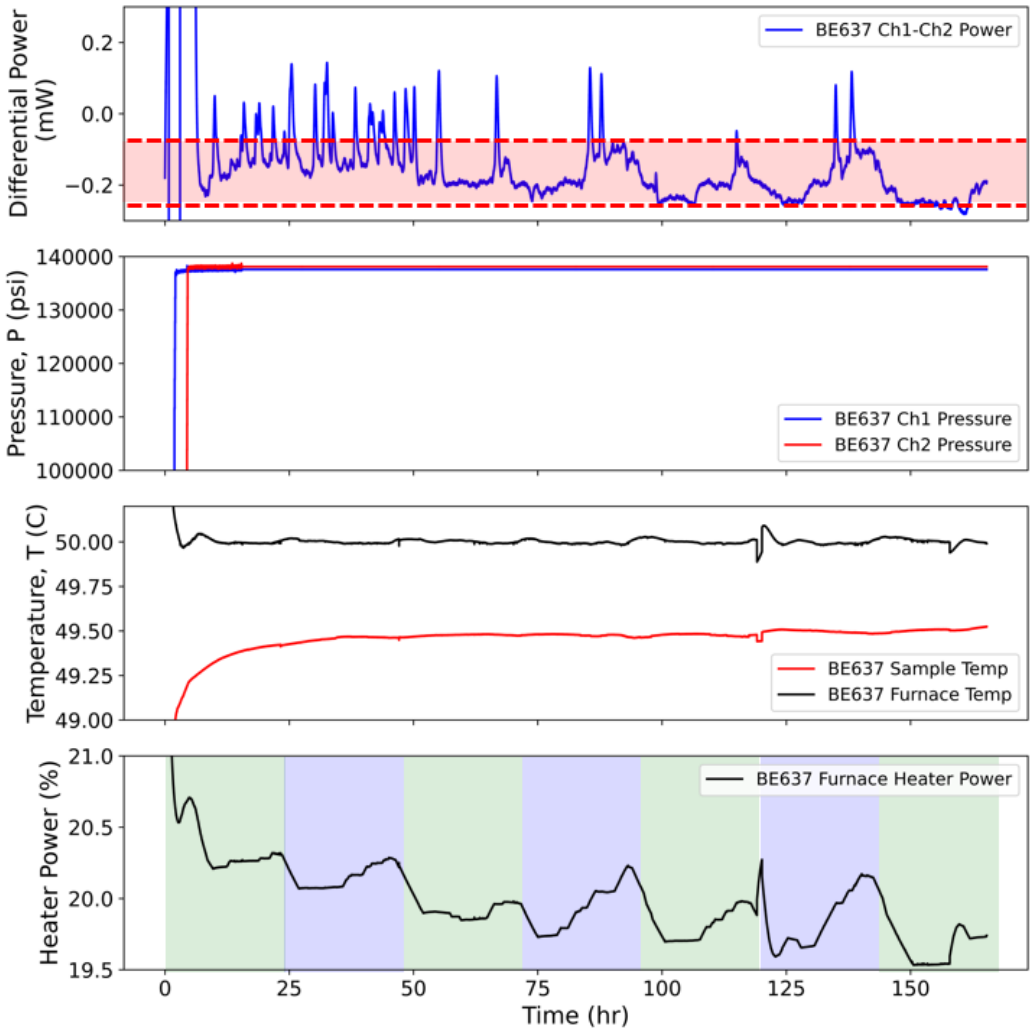


Figure 20: Auxiliary data included for BE637 that showed these odd heat spike events. Pressure and temperature trends show nothing correlating with these spikes. The heater power from the calorimeter itself is slowly decreasing over time and showing cyclical behavior over 24 hours due to small fluctuations in the temperature of the room the calorimeter enclosure was in. Blue and green zones designate 24 hour increments for the eye to follow easily.

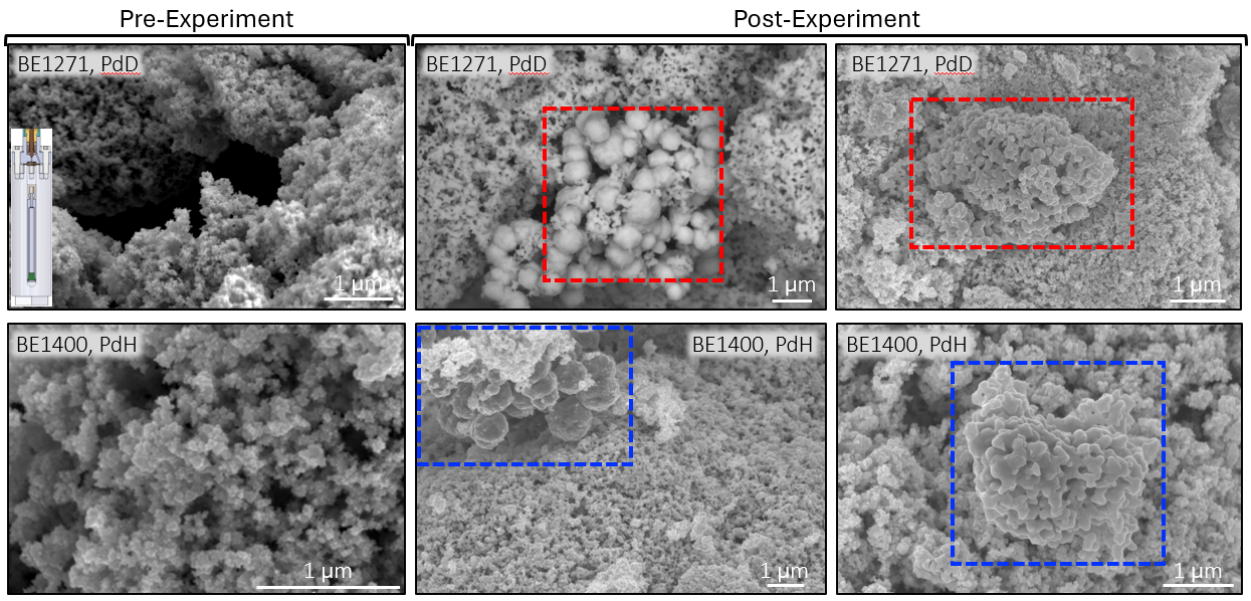


Figure 21: Pre and post SEM images of nanopowder used in passive calorimetry experiments for both protium and deuterium. Agglomerations of nanoparticles can be seen in post images independent of the gas isotope. We were also able to produce these agglomerations from pressurizing alone with nitrogen indicating this is a possible mechanical compression effect or these agglomerations always existed and were not caught in pre imaging procedures.

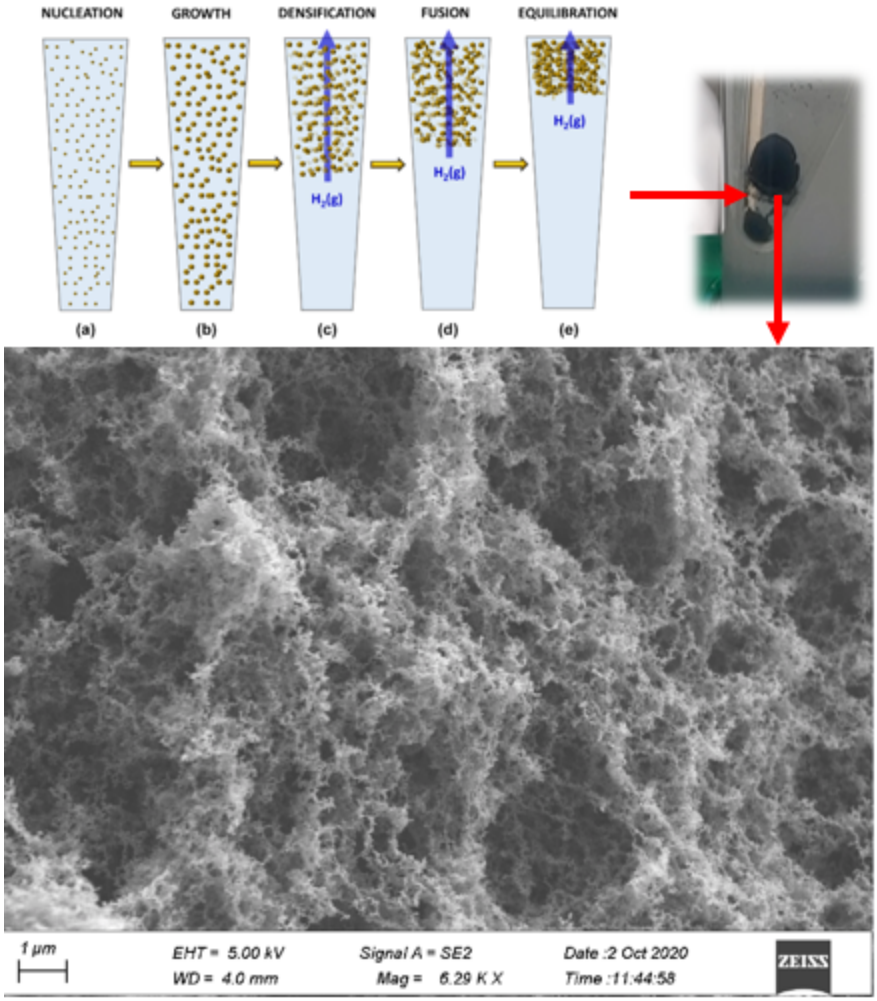


Figure 22: Simplified diagram of the process of how the aerogel and nanopowders were made inhouse and through external vendor. The nucleation rate, Growth, and densification were controlled through initial molarity, solvent temperature, and reaction time to obtain reproducible results. These fused nanostructures dehydrated from their solvent through sublimation freeze drying in high vacuum conditions over 24 hours. (Note: We could not confirm external vendor processing protocol but based off identifiable elemental contamination found in the vendors product we are confident they utilized the same precursor salts and solvents to obtain these nanostructures.)

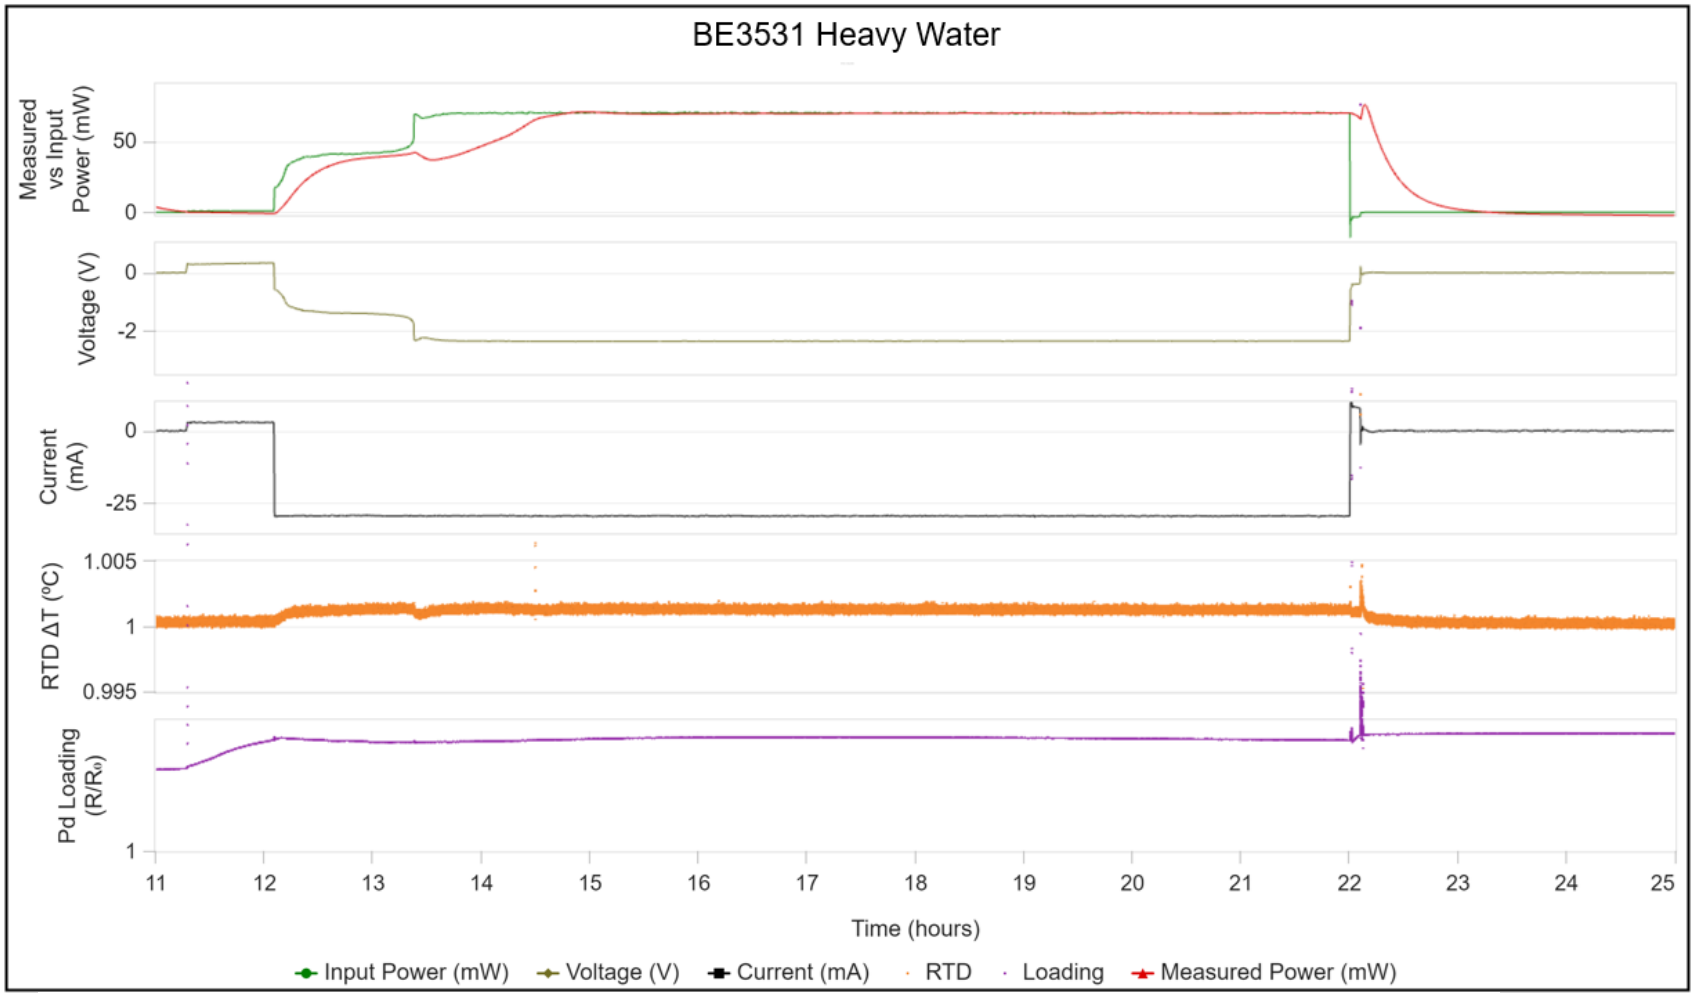


Figure 23: Auxiliary data for HPE calorimetry test in heavy water. Here we are showing how loading (R/R0), solvent temperature, and measured power change over time due to changes in voltage, current and voltage on the electrochemical circuit.

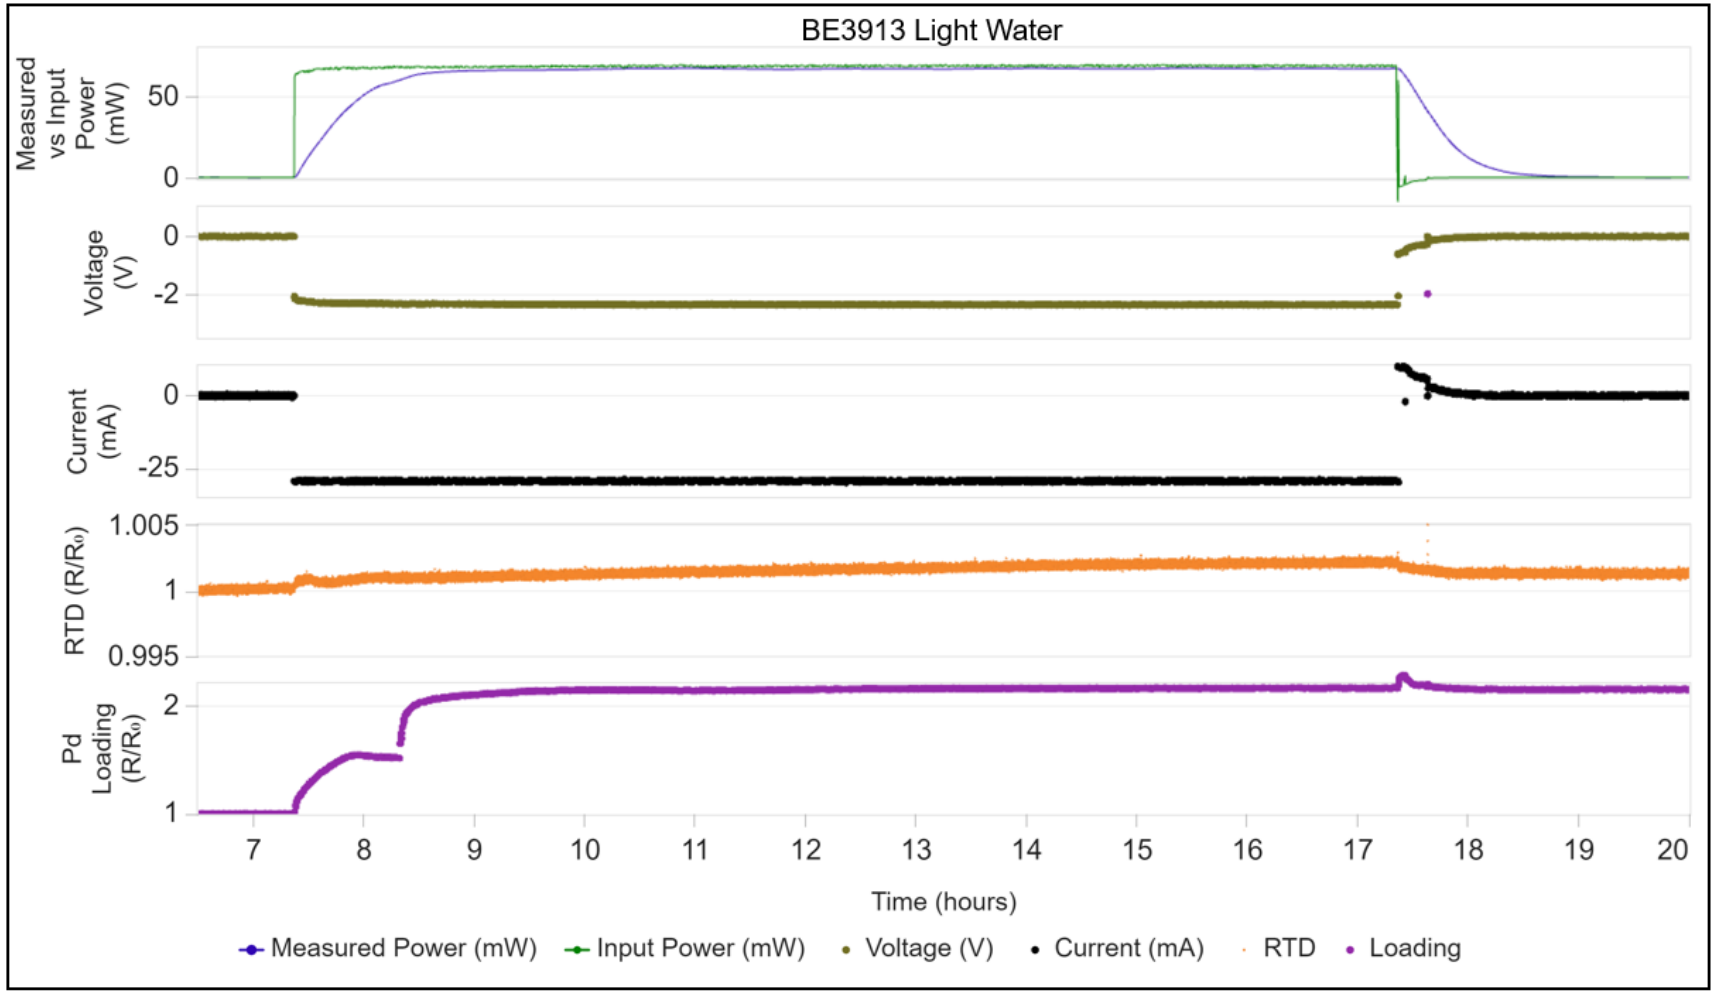


Figure 24:Auxiliary data for HPE calorimetry test in light water. Here we are showing how loading (R/R0), solvent temperature, and measured power change over time due to changes in voltage, current and voltage on the electrochemical circuit.

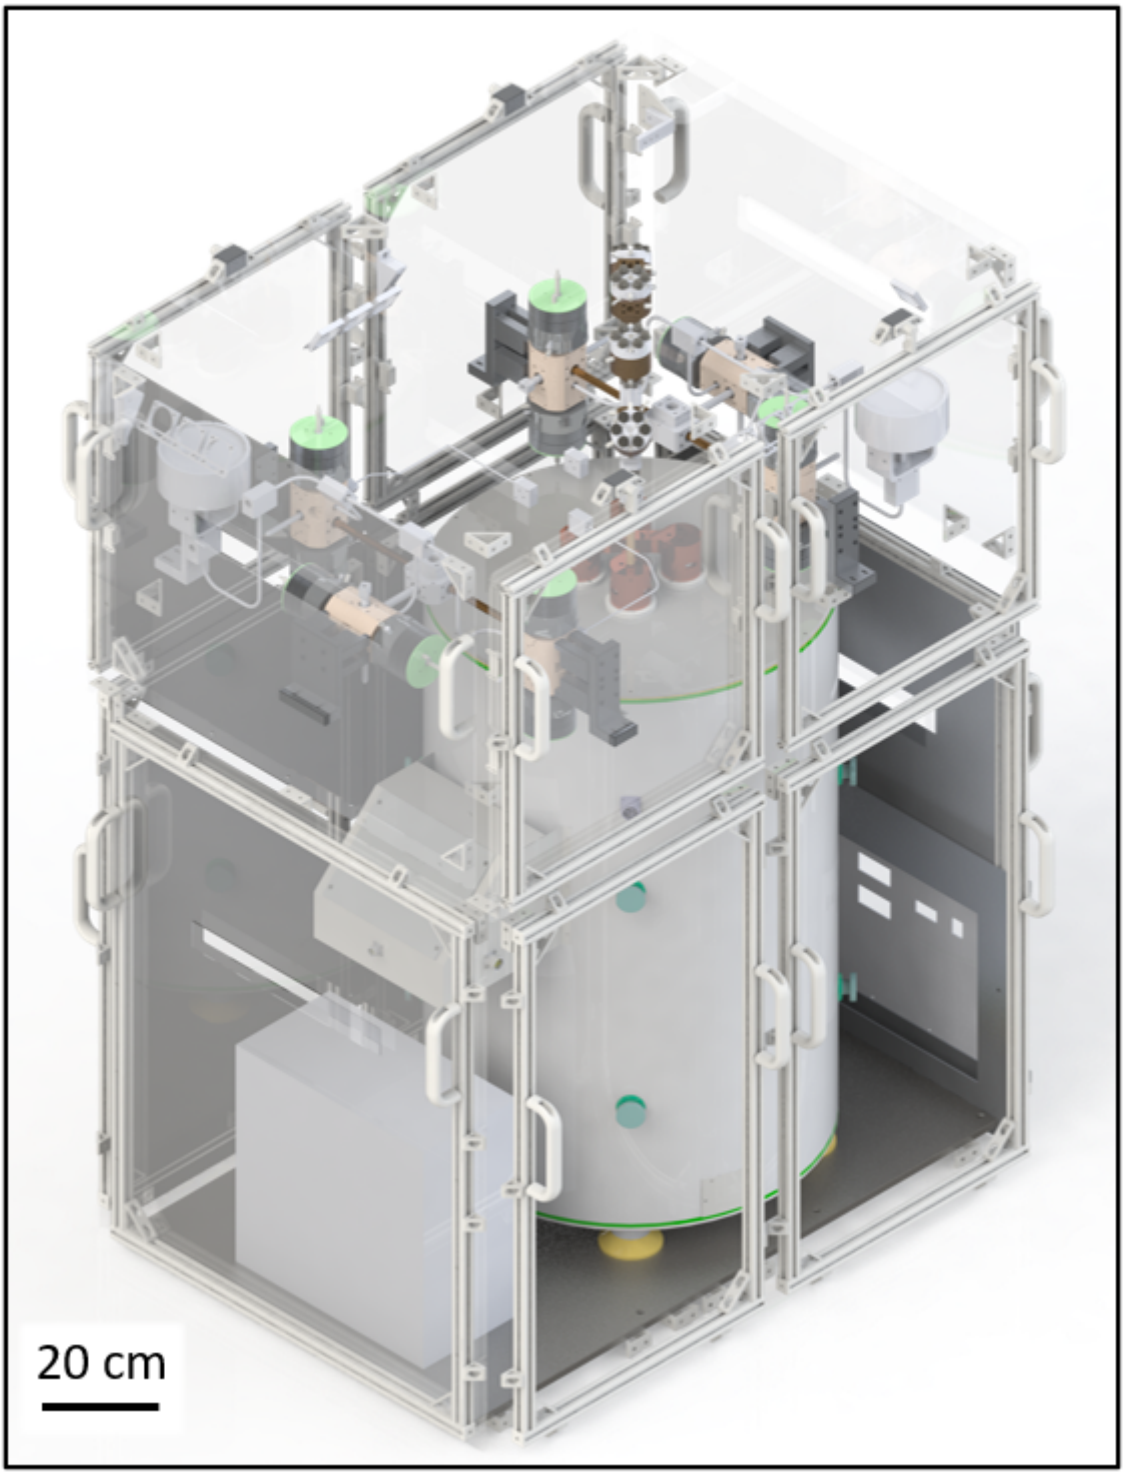


Figure 25: MS80 Thermal Enclosure, isometric view. The cart includes the MS80 calorimeter, 12 ultrapressure gas valves, and temperature and humidity PID controls to always maintain a constant 30 +/- 0.5 °C around the calorimeter shell.

Physical Archaeology

The energy from an LENR event is sufficient to melt a 100 nm bubble of PdD or vaporize a 50 nm one. Damage on these size scales should be detectable, especially if it occurs near the surface of the material, where it can be imaged. Physical archaeology refers to looking for these physical scars, which energy from an LENR event will leave behind in the material.

We have gone through three evolutions in learning how to look for LENR using this approach; all based upon SEM imaging in order to provide high resolution imagery over relatively large areas. Here we compare SEM images of the sample before and after the experiment, looking for damage caused by the experiment.

Initial Approach

In our initial approach, we focused on free standing, bulk sample geometries (e.g., Pd foils and slugs). Native Pd sample surfaces (i.e., as received) are too noisy to make useful interpretations. So, we developed mechanical polishing and chemical cleaning techniques to prepare pristine Pd surfaces prior to AHE experiments (**Fig 1**). However, loading significantly distorts the sample surface, convoluting SEM image data. We developed the elastic loading technique to uniformly load, but the presence of grain boundaries in polycrystalline Pd still led to distortions such as intergranular voids and cracks, grain growth and terracing. So, we switched to single crystal Pd samples, which successfully mitigated the distortions driven by grain boundaries and preserved the pristine surface after elastic loading and deloading (**Fig. 2**).

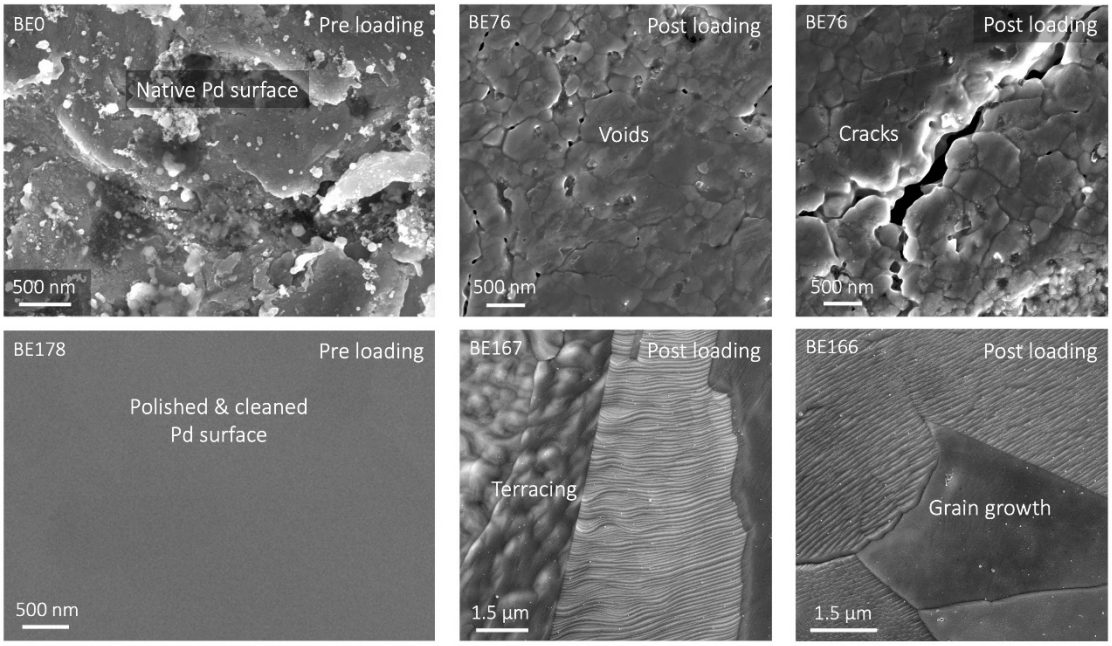


Fig. 1: (Top left) Native polycrystalline Pd surfaces are rough and full of debris / contamination. (Top center and right) Even polished samples become distorted by loading/deloading. (Bottom left) After mechanically polishing and chemically cleaning, the Pd surfaces are pristine. (Bottom center and right) However, post elastic loading, polycrystalline surfaces are distorted (e.g., intergranular voids and cracks, grain growth, and terracing), primarily due to the presence of grain boundaries.

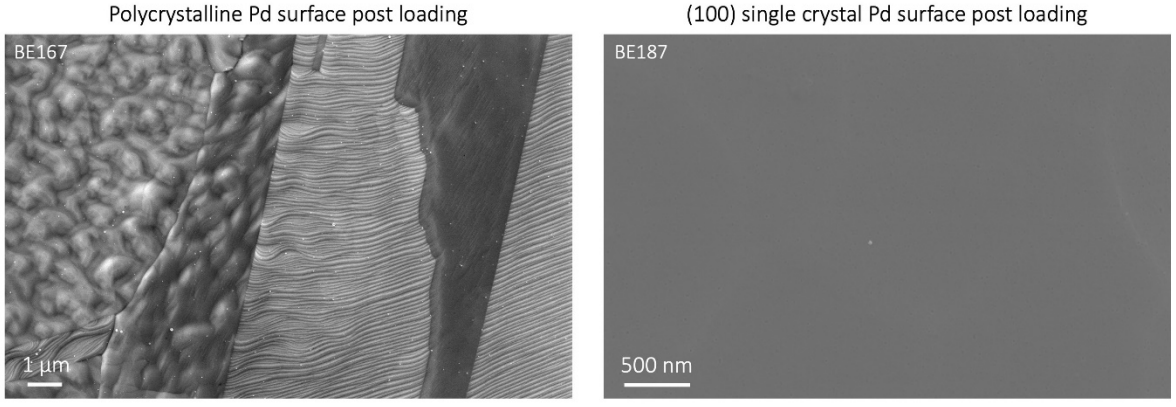


Fig. 2: Single crystal Pd surfaces remain pristine after elastic loading and deloading, successfully mitigating the distortions driven by grain boundaries in polycrystalline Pd.

Beyond developing sample preparation and loading techniques to preserve the sample surface, cleanliness is essential for valid archaeology data i.e., debris / contamination can obscure SEM images. As an example, we observed “linear features” (LFs) in early experiments. The LFs were

observed in both protium and deuterium with varying dimensions and concentrations (**Fig. 3**). LFs tended to appear in areas with nanoparticle contamination (aluminum oxide, Al_2O_3, nanoparticles were used to mechanically polish the samples prior to the experiments and were not entirely removed). Initial EDS (energy dispersive x-ray spectroscopy) analysis was unable to identify the chemical composition of the LFs due to the strong background signal of the Pd sample (i.e., when the LFs were on the Pd surface, the EDS only identified Pd). So, we used an in-situ manipulator to pick a LF off the Pd surface and perform EDS in isolation. This clearly showed the LFs were AlF_3 and most likely generated from a reaction of the Al_2O_3 nanoparticles and F from the lubricants in the pressure source (**Fig. 4**). Thus, archaeology relies on clean, contaminate-free surfaces to yield useful insights and avoid needless rabbit holes.

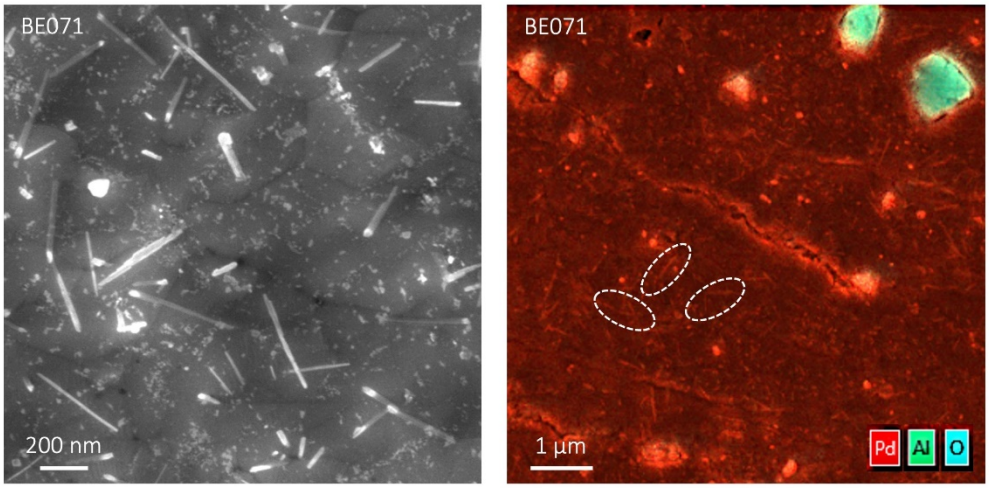


Fig. 3: Linear features (LFs) were identified on both protium and deuterium loaded Pd surfaces. EDS was unable to chemically distinguish LFs from the background Pd signal.

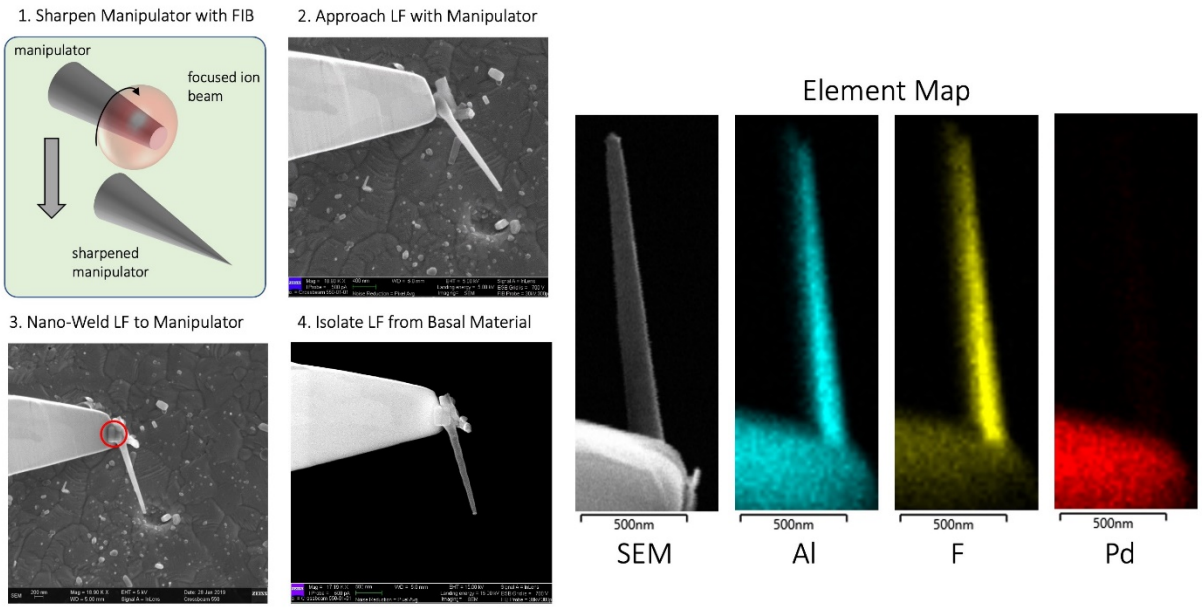


Fig. 4: Using a micromanipulator, a LF was picked off the Pd surface and EDS imaged in isolation, clearly revealing the LFs composition: AlF_3.

A final learning from the initial physical archaeology campaign was the importance of avoiding supercontinuum generation in our short-pulsed laser stimulation experiments. Using a Coherent Astrella short-pulsed laser with < 50 fs pulse durations at wavelengths of 400nm and 800nm, we performed stimulated AHE experiments at loadings >0.9 in Pd/D(H) and used pre/post SEM imaging as our AHE diagnostic. Valid stimulation experiments require fluences below the Pd/D(H) damage threshold, so as not to convolute laser damage with AHE damage sites (we expected 24 MeV AHE events to leave ~100nm voids/melt-spots in the Pd surface). We did observe ~100nm melt spots on some PdD_x samples (**Fig. 5**), but it was inconsistent (i.e., using identical loading and pulse conditions, we did not always observe the melt spots). We confirmed with cross sections that the melt spots were accompanied with sub-surface < 100nm voids (**Fig. 6**). We found that our samples were too close to our pressure windows and our focused, short-pulse beam was undergoing a non-linear optical effect that non-uniformly broadened the spectrum of incident light called supercontinuum generation (**Fig. 7**). This created illumination conditions that inconsistently damaged our samples. By increasing the sample-to-window distance, we avoided supercontinuum generation and confirmed that the melt-spots were in-fact laser damage (**Fig. 8**). So, clean archaeology experiments with short-pulse laser stimulation must avoid supercontinuum generation and stimulate below damage threshold.

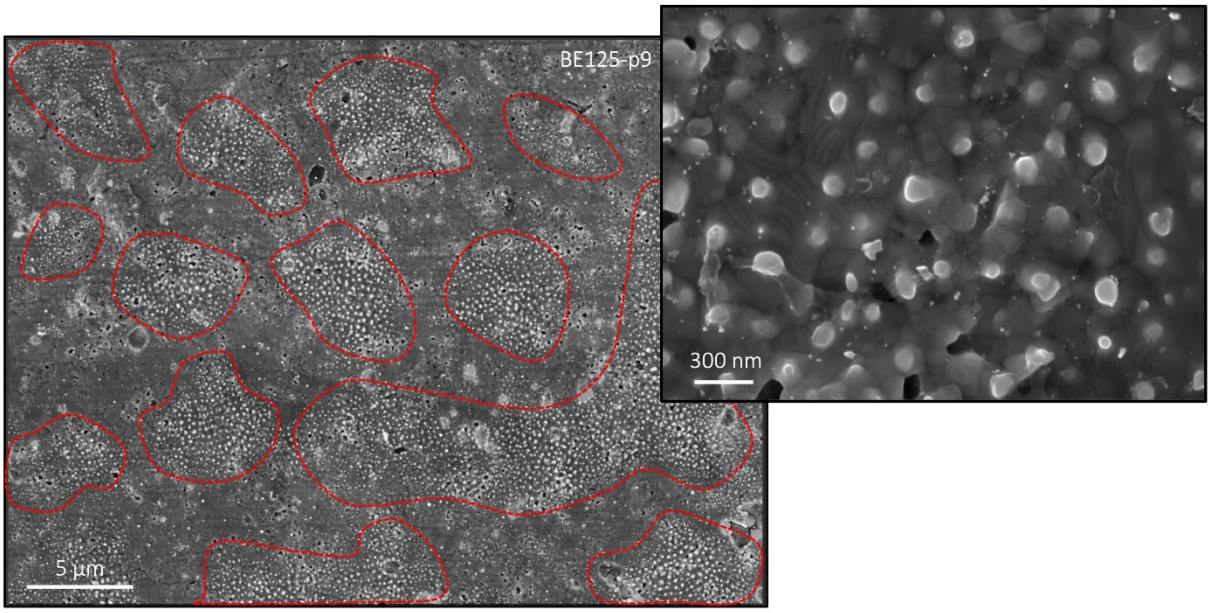


Fig. 5: Some PdD_x samples displayed patches of ~100nm melt spots after stimulation, but it was inconsistent.

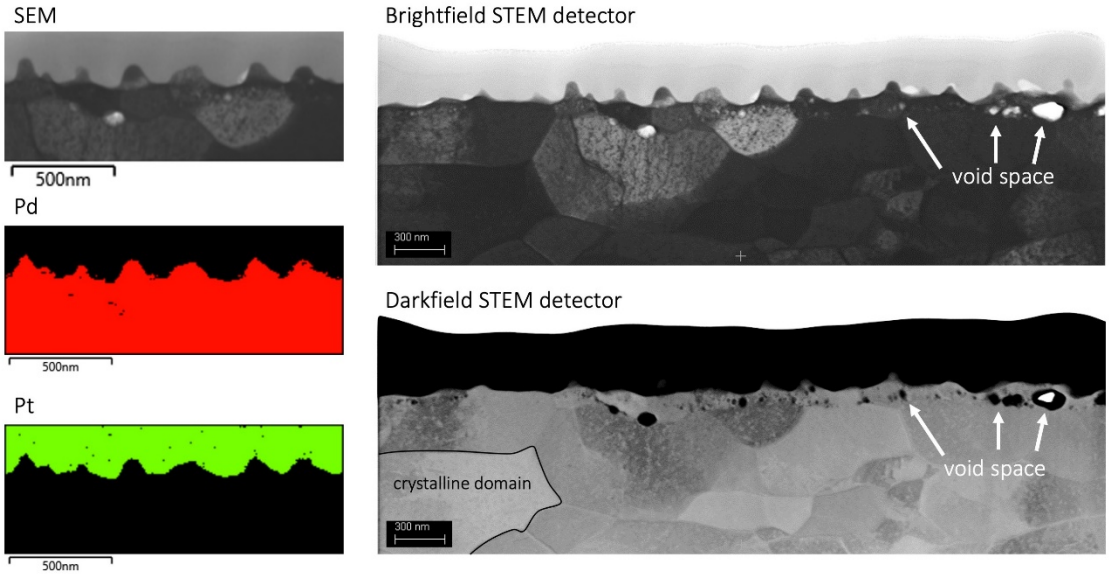


Fig. 6: The melt spots in PdD_x samples were accompanied by sub-surface <100nm voids.

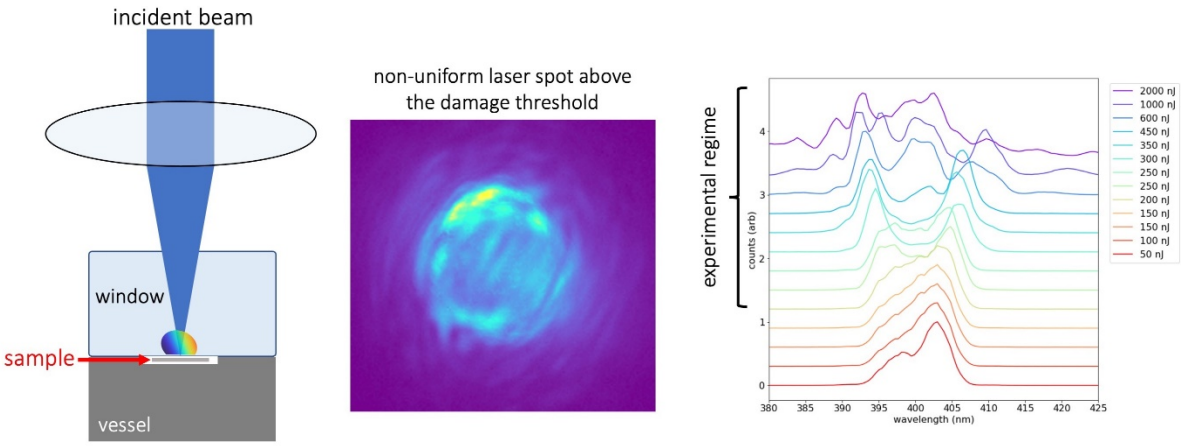


Fig. 7: With samples in close proximity to thick pressure windows, a focused, short-pulse beam can undergo supercontinuum generation and create inconsistent illumination conditions.

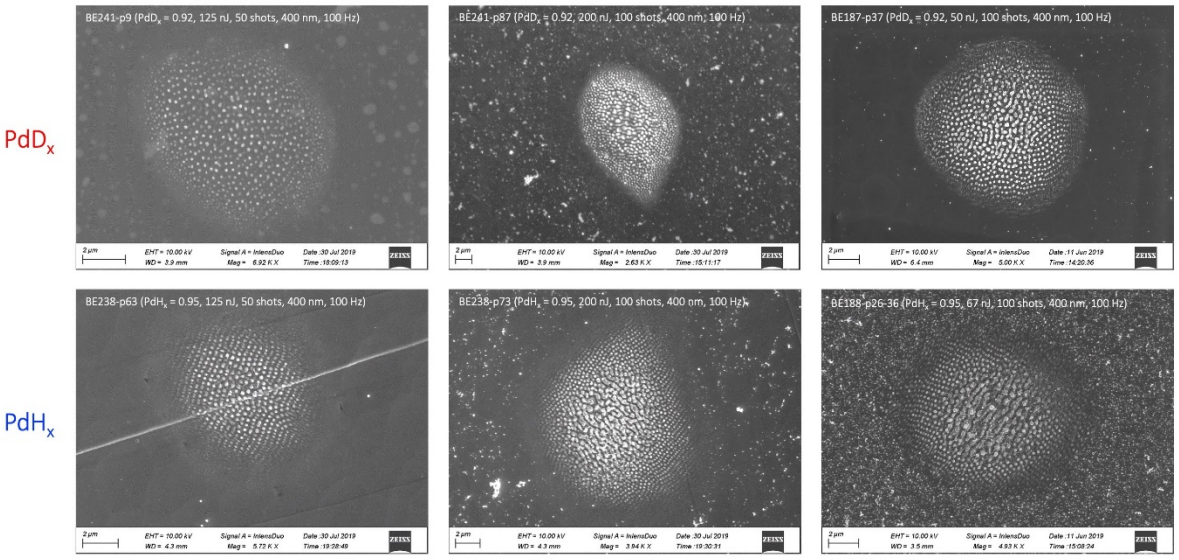


Fig. 8: Increasing the sample-to-window distance avoids supercontinuum generation and yields reproducible damage spots across PdD_x and PdH_x.

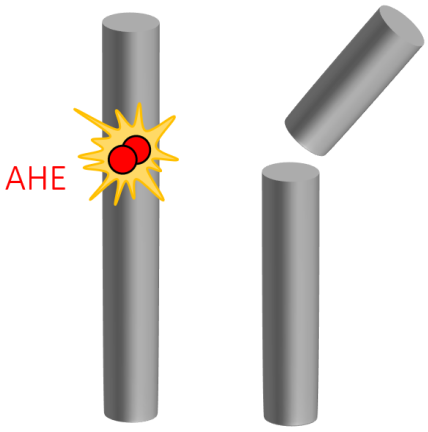
Nanostructure Archaeology

Following our initial archaeological approach, we switched our efforts to nanostructured targets, where damage would be more apparent. One key advantage is achieved by fielding nanostructures small enough so that the heat from an AHE event will damage or destroy them. This is done by

keeping their size (in at least one direction) well below 100 nm. A second way to make damage more visible, was to use arrays of discrete nanostructures. Not only is damage often more apparent in a small nanostructure, but they provide the damage sites with clearcut addresses; researchers can agree on locations for additional commentary or inspection.

Nanostructure Archaeology : Nanopillar Arrays

Our first, and most technologically most sophisticated, type of nanostructures developed for archaeology were Pd nanopillars. Each of these was a small cylindrical pillar, typically 60 - 100 nm wide and about 200 nm tall. The concept behind them is shown below; energy from a single 24 MeV AHE event will be spatially confined by their small diameter, and severely damage, if not destroy, the pillar.



Two other characteristics of Pd nanopillars are worth noting. Because of their small dimensions, they load (and deload) with deuterium very quickly. And, since they are only anchored at their base, most of the swelling caused by deuterium loading is unrestrained, so does not cause internal stresses or disruptions to the Pd lattice.

Obviously, a problem with nanopillars is that they are small. So, in order to test a significant amount of PdD material, we need to field a very large number of nanopillars. We did so by creating large 2D arrays of nanopillars on a substrate, typically containing millions of individual pillars.

Nanopillar arrays were first formed by sculpture, carving them out of bulk Pd samples by using focused ion beam (FIB) milling (**Fig. 9, Fig. 10**). First, 0.5mm thick, 5.5mm diameter 'slugs' were fabricated from mono-crystalline rods of palladium and their surface was polished to nanometer or sub-nanometer scale surface roughness via chemical mechanical polishing and grazing incident Ar^{+} ion milling. The pristine surfaces were laser marked with fiducials to measure sample strain during loading and target boxes which could be targeted by stimulation lasers in experiments.

Finally, a gallium FIB column in an SEM was programed to mill recessed arrays of nano pillars into the sample surface within the target boxes.

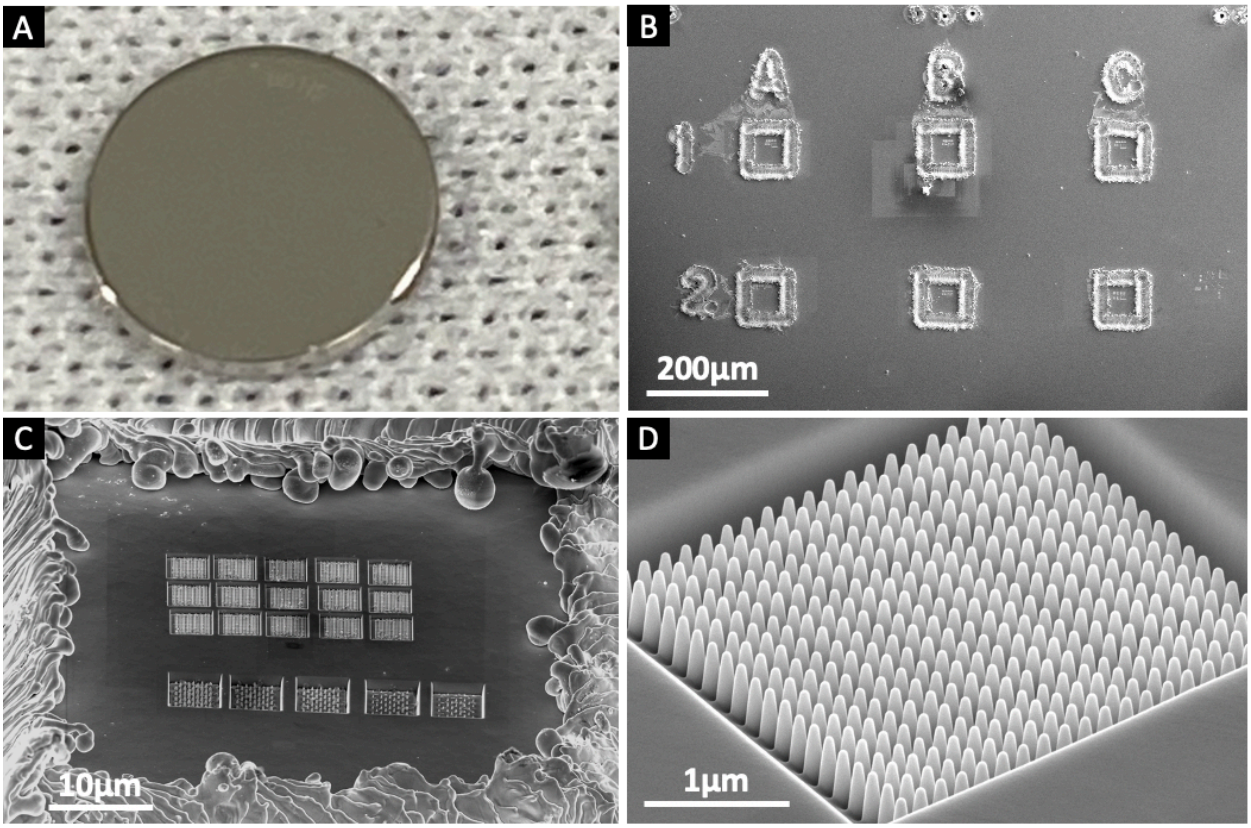


Figure 9. A) A single crystal palladium slug polished to sub-nanometer RA surface roughness. B) Laser marked target pads shown on the surface of the palladium slug. C) A target pad containing 20 FIB pillar arrays. Each array utilized slightly different milling conditions or programmed milling pattern to generate an assortment of pillar geometries and array layouts on each pad. D) A FIB pillar array containing 400 individual palladium nano pillars.

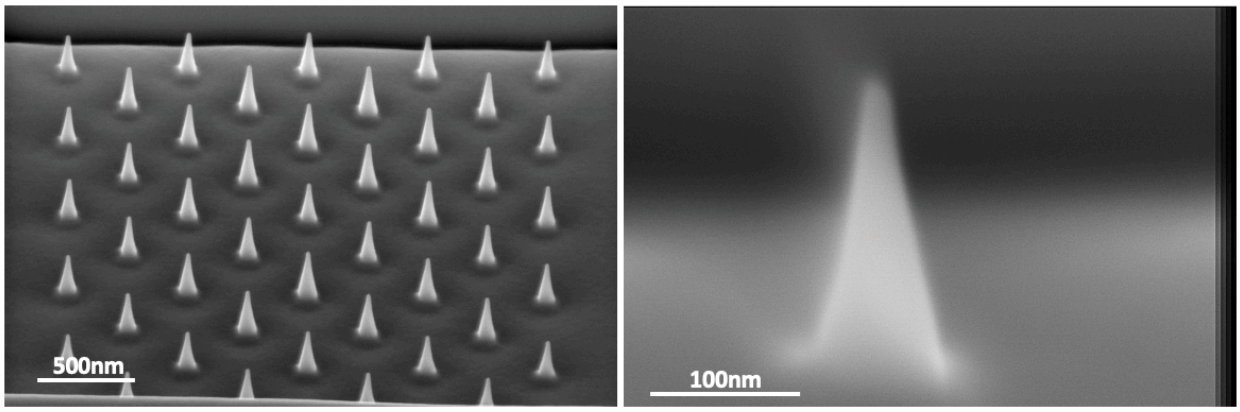


Figure 10. The size of FIB pillars could be decreased by further exposure to the ion beam.

During short pulse stimulation of deuterium loaded targets containing FIB pillars, archaeology found evidence of altered pillars (**Fig. 11**). However, this result was eventually replicated in protium and was therefore determined to be an effect of classical origin, not a result of LENRs.

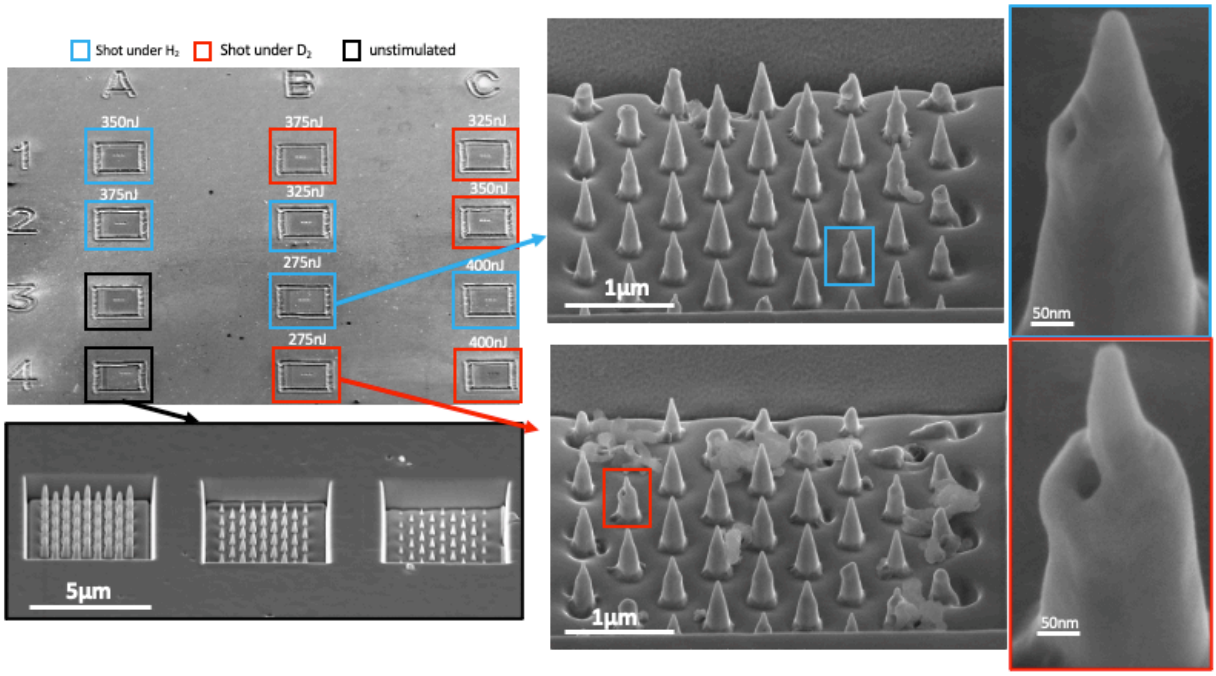
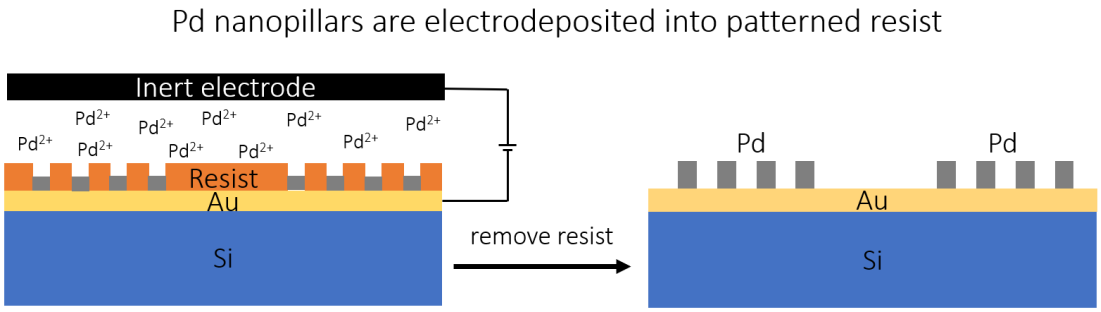


Figure 11. Callisto experiment BE827. FIB nanopillars on single crystal Pd are loaded at ultrapressure (H_2/D_2) and stimulated using a high fluence, ultra-short pulse laser. Pillars illuminated by 275nJ pulse energy exhibited cavities and other structural changes in the archaeology analysis. The effect was produced under both hydrogen and deuterium and was found to be not isotopically dependent.

Milling pillars using FIB is a time intensive process which limits the number of pillars on each sample. A more scalable approach was developed using electron beam lithography to pattern arrays of holes in a resist layer on gold coated silicon wafers (**Fig. 12**). Palladium was then electroplated onto the wafer to form palladium nanopillars in the holes. Extensive effort was required to optimize the lithography process for hole diameter and uniformity of the resulting arrays. Lower e-beam doses and development time lead to smaller diameter holes/pillars but risk residual resist at the base of the hole preventing wetting during electroplating, leading to missing pillars within the array. Through extensive testing, a process for wetting and plating the proper palladium thickness was achieved.



- Electrodeposition is commonly used to fill vias in the semiconductor industry; Plating vias on the nanoscale is more challenging due to the need for e-beam lithography
- We developed a method to plate nanoscale vias with near-perfect uniformity and reproducibility down to 60nm

Figure 12. Pd nanopillar fabrication procedure.

The nanopillars arrays (**Fig. 13**) formed hexagonal patterns, with 400 nm spacing between nanopillars. Most of these arrays were intended for use with laser stimulation, so were patterned into a 3x3 array of separate nanopillar-filled pads, spaced apart so that each pad could be irradiated with different stimulation conditions. Each pad was 250 microns on a side, and contained 450,000 nanopillars; so there were over 4,000,000 pillars in a typical experiment. Substrates were prepared by evaporating 10nm titanium followed by 150nm gold onto Si wafers. The wafers were then coated in 300nm thick PMMA 950k and large cross marks near the corner of the wafer were exposed using DUV and a physical mask. The large cross marks had a combined area of 1cm^2 and accounted for the overwhelming majority of the exposed gold, making the current density in the holes much more stable and controllable during electroplating. The virtual mask for e-beam lithography contained 30nm wide hexagons to define each pillar and was exposed onto the PMMA resist with a dose typically between 3,000 and 5,000$\mu C/cm^2$. The pattern was then developed and descummed using 15s O_2 plasma following development, and again before electroplating. Wafers were pre-wet by submerging in $PdCl_2$ based electroplating solution in a vacuum chamber for 5 minutes. Electroplating was typically done using constant current at 10mA with plating times ranging from 30 seconds to 2 minutes. After electroplating, wafers were diced using a micromachining laser and cleaned with a series of solvents including acetone to remove the PMMA resist.

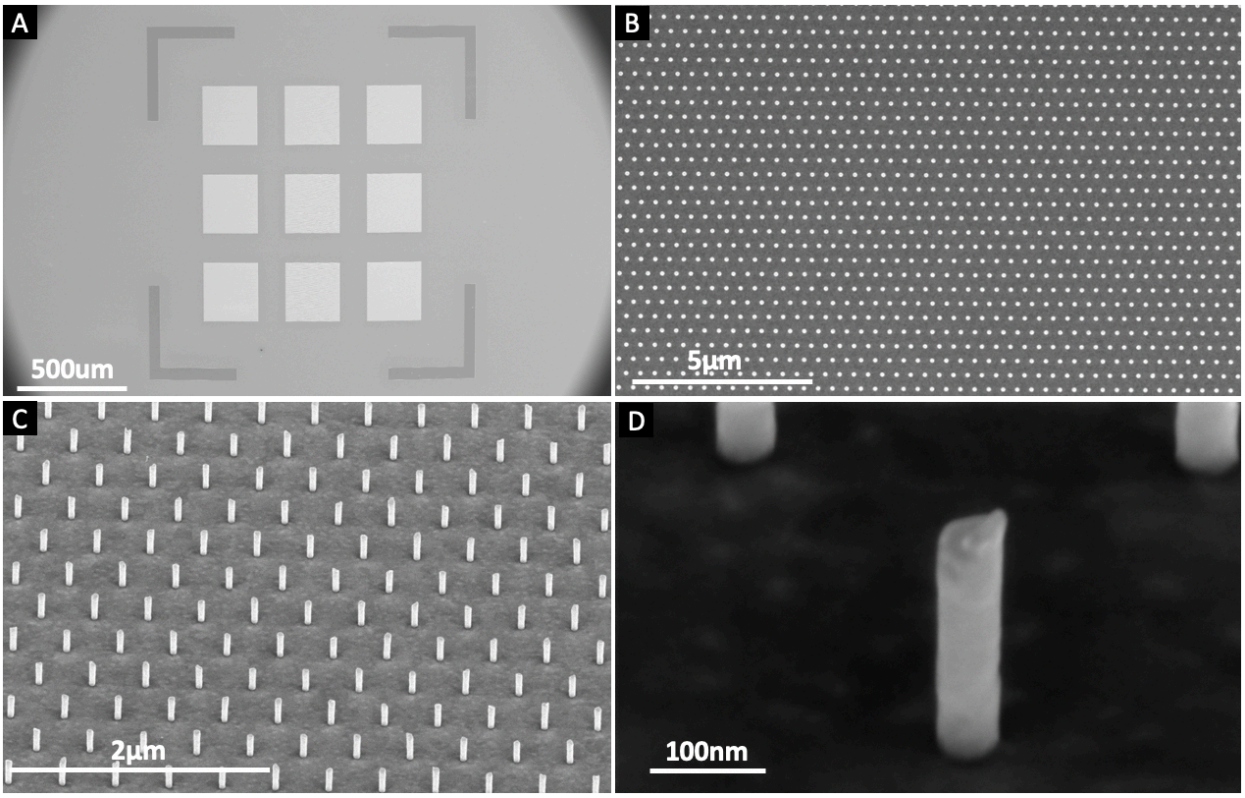


Figure 13. Nanopillar array samples. Each sample contains a 3x3 grid of nanopillar filled pads, each pad 250x250µm and containing nearly 450,000 pillars.

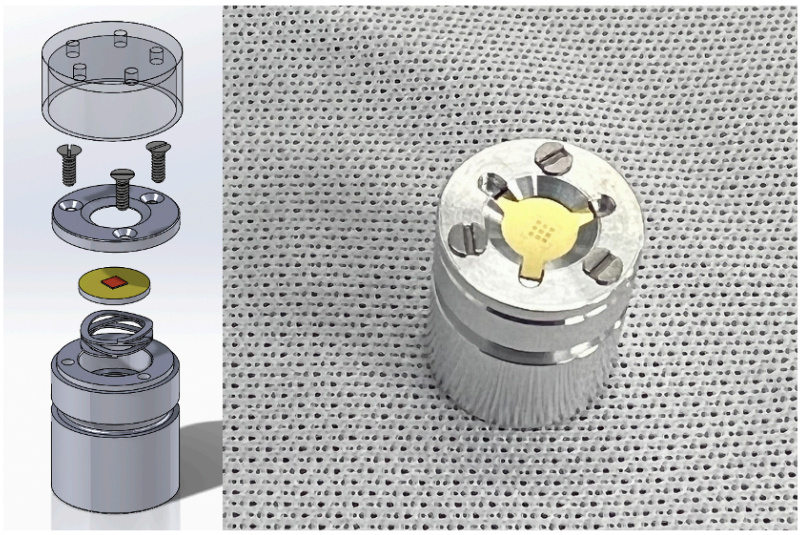


Figure 14. A capsule was developed to hold archaeology type samples for easy handling and installation in ultrapressure vessels with optical access.

Following the palladium electrodeposition on the wafer, a micromachining laser was used to cut out the samples into 5.5mm diameter disks. These disks are then installed in a custom capsule to protect the delicate samples during handling and allow easy installation of the sample into ultrapressure systems (**Fig. 14**). Other than their susceptibility to damage during handling, the nanopillars proved extremely robust during their exposure to 145 ksi gas pressures, and to being loaded and deloaded with protium or deuterium.

These large arrays were never perfect; errors in lithography or electrodeposition would leave occasional gaps in the array. In practice, this was not a serious problem; the errors were rare (typically less than 1 per 1000 pillars). Even more fundamentally, we would record pre-experiment SEM images, allowing us to ignore pre-existing damage when analyzing the experiment. The one form of experimental damage which proved most problematic was due to small debris particles. As a practical matter, most ultrapressure experiments inadvertently included small numbers of "debris" particles (e.g., around a micron in size, and made from silica or hydrocarbons). Gas flow would drive these into the nanopillar arrays; each impact breaking one or more nanopillars. On a global basis, this predation was a small effect, but it was functionally more serious than fabrication errors, because it occurred during the experiment, and could not be discounted by pre-experiment SEM images. A pillar broken or destroyed by an AHE event may be impossible to distinguish from one hurt by a debris particle.

Most of the AHE nanopillar experiments were performed in our Staring apparatus; this is described in detail in the Staring section of the report. To briefly summarize, this apparatus allowed nanopillars to be loaded with either H_2 or D_2 gas at 145 ksi pressures, achieving loadings of 95% or above. The pillars were organized into a grid (usually 3x3) of separated pads, each one a 250 μm square, containing about 450,000 nanopillars. A full experiment therefore, would contain over 4,000,000 nanopillars; each of which acting as a single-event LENR detector.

Obviously, it is impractical for experimenters to personally inspect each pillar in an experiment containing 4,000,000 of them. To deal with this, we developed automated image-processing algorithms (discussed further in the Staring section) to inspect SEM images (acquired both before and after each experiment) and flag "interesting" features for human attention. Two types of filtering were used, one looking for abnormal features (ones that didn't look like normal pillars) and the other for missing pillars.

Usually, the cause of the flagged features was easy to identify. For instance, most missing pillars were already missing in the SEM pre-images; e.g., were due to fabrication errors. Likewise, many of the abnormal features were clearly (due to their size, shape, or position) debris particles, not damaged pillars. Pillars which were found to be missing or damaged in post-experimental SEMs, but not in the pre-images, were always more interesting. However, the cause of these could almost always be identified by one or more "smoking guns". For instance, since an LENR event should mostly affect its host pillar, groups of damaged or missing pillars were indications of handling errors, debris impacts, or (sometimes) too much stimulation energy. And sometimes, assigning

blame to debris particles was clearcut, e.g., the particle was still adjacent to the damage, or several broken pillars lay along a straight line.

After excluding all of these false-positive features, there were usually only a very small number (1 or 2 per experiment) with no obvious cause. However, the fact that they were so rare, were not systematically reproducible, and did sometimes occur with PdH, made it impossible to prove they were caused by LENR. Some could have been, but in each case, they could have been caused by a debris particle which was later carried away with the gas.

This ambiguity meant that nanopillar archaeology was not a definitive tool for either proving or disproving LENR.

Nanostructure Archaeology : Nanofilm Arrays

While nanopillars were too damage prone to give clearcut evidence for or against LENR, the attraction of being able to detect single LENR events remained compelling. So, rather than abandoning archaeology, we simply found a more suitable nanostructure, nanofilms.

These are basically just thin films of Pd, deposited on the surface of an inert substrate. The film is made thin enough (20 – 30 nm) so that energy from an LENR event will blow a hole through the film, down to the substrate; the contrast between the exposed substrate within the hole and the Pd film surrounding it, makes such holes stand out in broad-area SEM imagery.

For convenience in image analysis, we subdivided the Pd film into an array of separate regions, 10 μm squares called nanofilms, with bare substrate in between. Not only did the discrete nature of nanofilms make cataloging locations much simpler and unambiguous, the contrast between the Pd nanofilms and the substrate between them, made it possible to maintain a tight and consistent SEM focus throughout each image sequence.

Nanofilms offer two major advantages over nanopillars; robustness and ease of fabrication.

Thin films are much less vulnerable to damage from handling or debris impacts than are nanopillars. And, when damage does occur, gouging a hole down to the substrate, the shape of the hole hints at its cause, i.e., irregular for debris, smooth and circular for LENR. This, of course, is far different than simply being faced with a missing nanopillar. These facts, that nanofilms experience less damage, and that damage which does occur is informative, mean that nanofilms are a much less ambiguous diagnostic than were nanopillars.

Nanofilms are also much easier to make and can be formed from virtually any metal. Deposition of thin metal films is a very standard and repeatable fabrication process. And, since the individual nanofilms were 10 μm in lateral size, not nanometric, the arrays could be patterned with simple photolithography, not requiring the e-beam lithography used for making nanopillars. Because thin film deposition is a pretty straightforward process, we were able to extend our nanofilm fabrication from Pd films to a wide range of other metals, including many Pd alloys, Ni, and even several rare-earths. In principle, this would have been possible for the electroplating used to make nanopillars,

but in practice, each different composition would have required its own development effort. With nanofilms, switching materials turned out to be quick and straightforward.

In fabrication, a substrate is coated with 1μm thick SPR3612 photoresist over a 200 nm thick liftoff layer. The virtual mask is exposed, and the pattern is developed. Five nm of titanium adhesion layer followed by 25 nm of the desired test material is deposited by e-beam induced evaporation, followed by liftoff of the excess metal. In one case, the nanofilm capacitors used for electrical stimulation in the Chips campaign, we did not use any adhesion layer, since the intent was to apply electric field directly to the palladium surface.

Various substrates were used with these devices depending on the application. In some experimental setups, sapphire wafer substrates were needed due to frequent fracturing of silicon based devices. Sapphire based samples could also employ a 5 nm thick gold layer over the substrate prior to photolithography to aid in charge dissipation during SEM analysis. Thermally oxidized silicon could be used to increase the thermal isolation of the micro pads, but made high resolution SEM imaging challenging and was deemed unnecessary.

For use in LENR experiments, we usually fielded nanofilms in 8x8 arrays; each array forming a separate pad, spaced apart from neighboring pads (**Fig. 15**). Each pad had 64 individual nanofilms, and occupied a 150 μm by 150 μm square. The number of pads, and the separation between them varied depending on the experiment in which they were used. In some cases, e.g., in the Staring campaign, we could tailor our laser beam to match the size of a pad; hence by spacing pads far enough apart, each one could be irradiated with different stimulation conditions. In other experiments, the entire target area received the same stimulation, and so there was no drive to increase pad-to-pad separation.

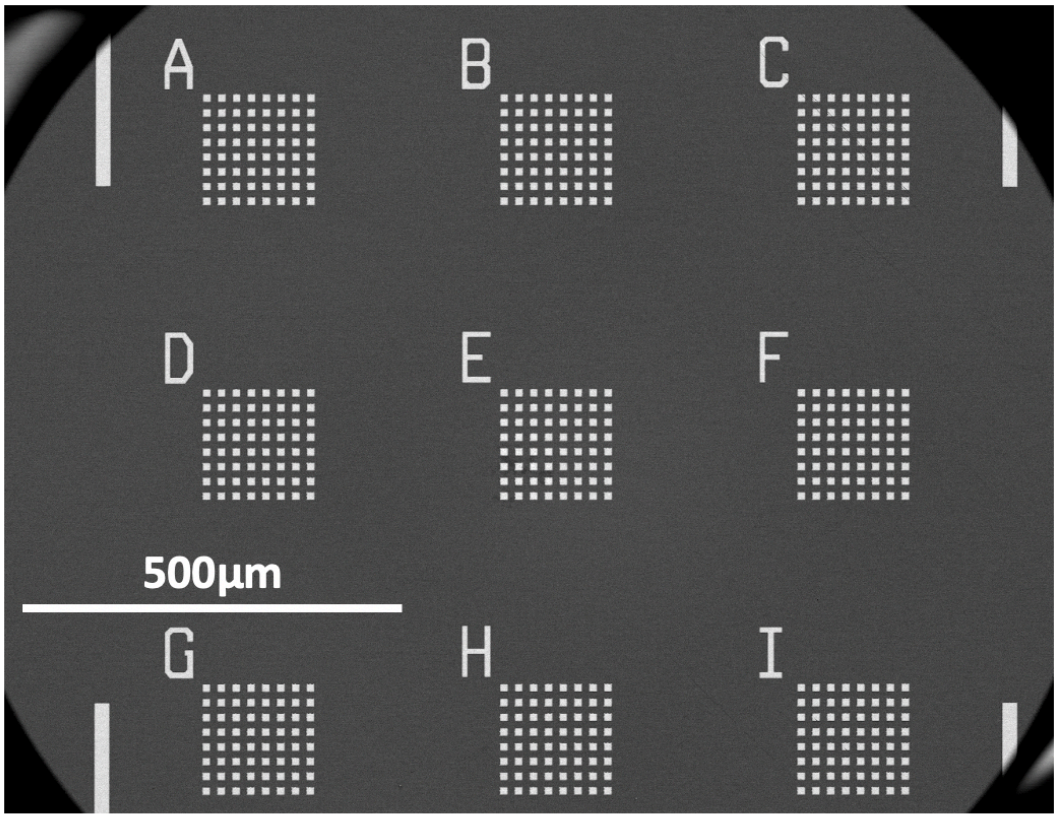


Figure 15. SEM of micro pad array sample.

As mentioned above, a prime reason for splitting large Pd films into arrays of smaller nanofilms, was the ability to identify the locations of interesting features; i.e., in being in a particular nanofilm. The 8x8 array format was chosen to allow us to lithographically mark each individual nanofilm with binary row-column marks, uniquely identifying its location within the array (**Fig. 16**).

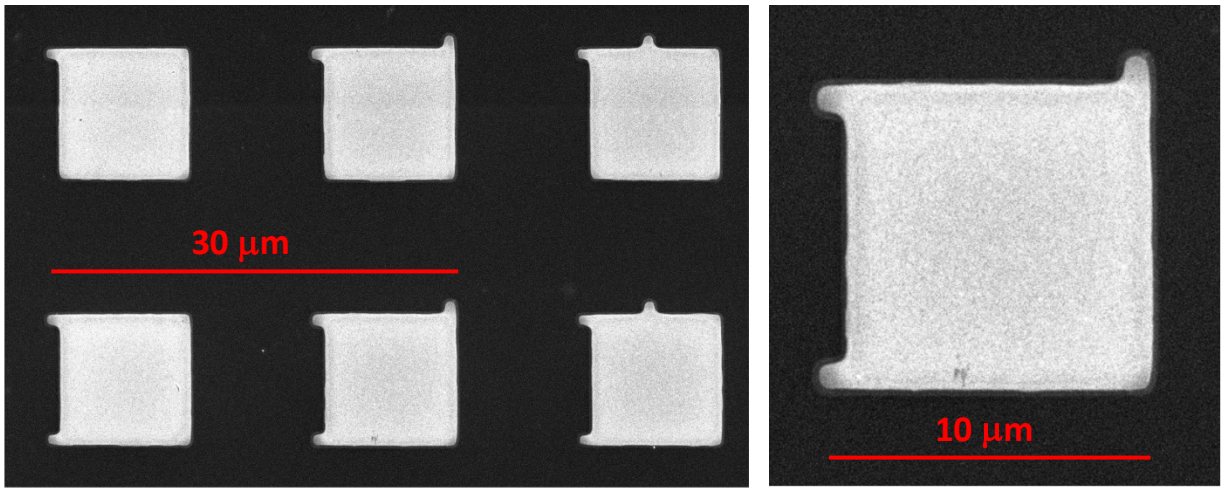


Figure 16. Left: Six nanofilms from rows 4-5 and columns 0-2 of array. Right: Nanofilm (5,1) at row 5, column 1 site of array.

A typical nanofilm experiment involved 576 nanofilms, each 10μm x 10μm in size. We anticipated that LENR events, if present, would be manifested as holes of roughly 0.1 μm in size. These would show up in SEM images of the nanofilm as dark spots against the lighter background of the Pd surface. While we did begin development of a machine-based image analysis tool to search for such holes, in practice it was not essential; human inspection proved effective enough.

The first inspection stage was to simply look at the post-experiment SEM image of each nanofilm, noting the presence of any dark spots. There were typically two causes for these; surface features (debris or contamination), and holes down to the substrate. We used contrast to help distinguish these; holes were black, while surface features were usually not (ranging from white to dark grey). In some cases, black features were clearly irrelevant (e.g., due to scratches or gouges). But, when such causes were not certain, we would refer back to the pre-experiment SEM image of the same nanofilm. Features which were already present before the experiment (**Fig. 17**) were obviously not caused by LENR, while those that only appeared in the post-experimental images potentially were.

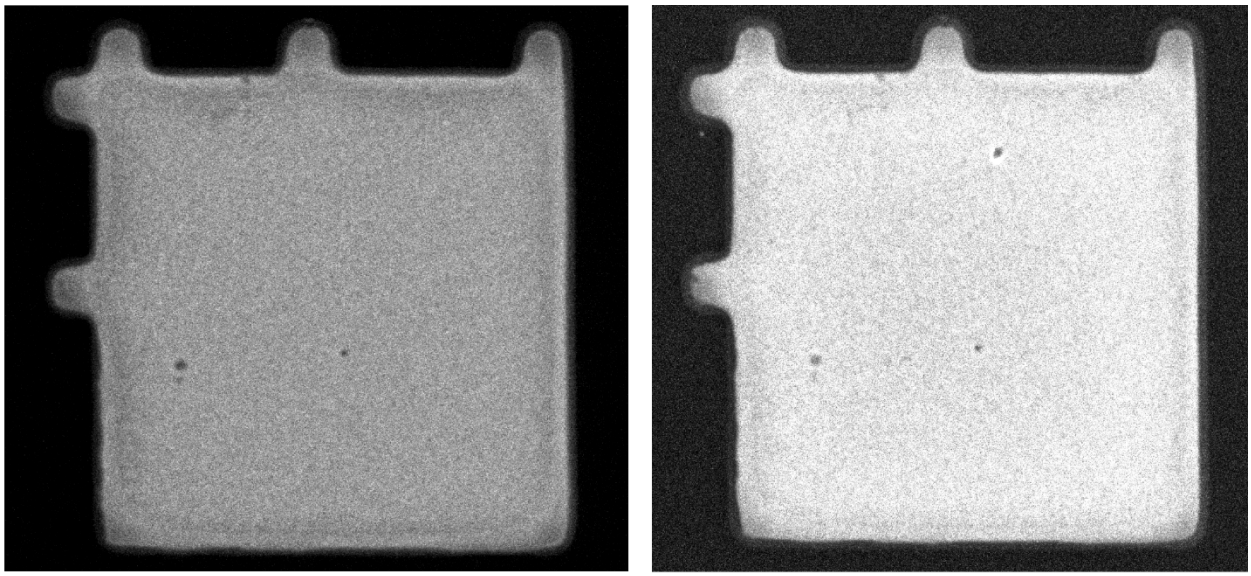
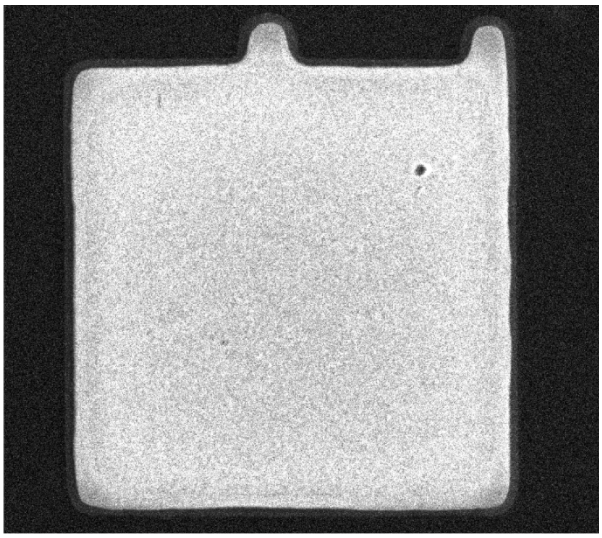
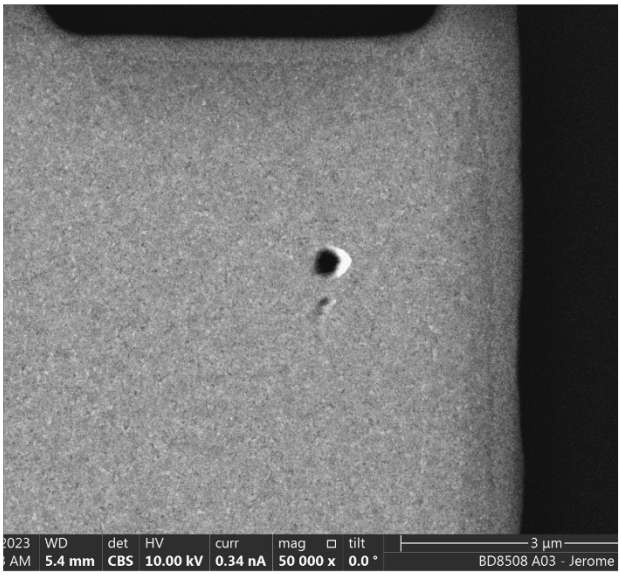


Figure 17. Two SEM images of nanofilm A67 from target BD8508 used in experiment BE4178. On the left is the pre-experimental image, and on the right is the post-experimental one. The lower two spots are present in both images, so are irrelevant to LENR. While the top spot in the post-SEM image is potentially interesting, the white halo around it indicates a debris origin.

After each experiment, there were often image features which could not be explained just from the first-stage inspection. These were black objects (high-contrast compared to the metal film) and only present in the post-experimental SEM view, not the pre-experimental images. We subjected such objects to a second-stage inspection, using spot images with higher magnification and a pair of different SEM detectors (**Fig. 18**), one based upon secondary electron emission, and the other upon electron backscatter. Such images not only addressed the specific features being re-imaged, but also helped us to interpret types of features which could be characterized in the first-stage inspection without needing re-imaging.

Figure 18. Example of second-stage image analysis from BD8508. Top-right: The original post-experimental SEM image of nanofilm A03, with unexplained black spot. Bottom-left: Higher magnification SEM image using backscatter detector. Bottom-right: Higher magnification SEM image using secondary-electron detector. These spot images show that this feature (and similar white-bordered ones) are impact features, not LENR holes.







In practice, the time needed to perform nanofilm experiments was dominated not by the actual operational time, but rather by the time required for SEM imaging. Producing a suite of high-resolution SEM images for 9 pads of 8x8 nanofilm arrays (i.e., 576 individual nanofilms) would typically require a full day of SEM time; and, of course, mean forgoing other usage for that period. Performing a fully documented and controlled experiment required 4 such imaging periods; having each PdD experiment matched with a corresponding PdH one, and with all nanofilms SEM-imaged both before and after the experiment. Because of this large resource commitment, we sometimes used a shortcut procedure to increase throughput four-fold, in which only post-imaged PdD versions were performed. The rationale in these cases, was that in the event that PdD nanofilms were not showing interesting results in post-experiment SEM images, there was little need to expend time on taking pre-experimental images, or to running PdH control versions.

Nanostructure Archaeology : Nanofilm Experiments

We used nanofilm arrays as single-event detectors in a wide range of different LENR experiments.

The first three cases, listed below, were our most extensive applications of nanofilm archaeology. All of them tested not just pure Pd, but also five different Pd alloys (Pd with B, Ni, Ag, Rh, Pt).

Staring: In this case, the Staring laser was used as a stimulation source, providing thermal stimulation at 145 ksi with two different laser wavelengths, and a variety of intensities and durations. The results, and usage of nanofilms, are described in much further detail in the Staring section of this report.

Ultra-Short-Pulse: We used nanofilm arrays to test for LENR in the extreme environment of both ultra-high pressure and non-thermal stimulation (using the Ekspla femtosecond laser). A tabulation of these experiments is given in Table 1 of the Ultra-Short-Pulse section of the report.

High Temperature Equilibrium: We used nanofilm arrays to test for LENR in an extreme form of pressure-temperature equilibrium; ultra-high pressure and near-melting temperature. A tabulation of these experiments is (as above) given in Table 1 of the Ultra-Short-Pulse section of the report.

The three cases above involved many experiments, each containing 576 nanofilms. Most of the nanofilms showed no hole-like features. When features were observed, they were individually inspected and categorized, as described earlier. Most of the pure Pd targets were pre-imaged, so many "suspicious" features in post-images could definitively explained. Many of the alloy experiments were not pre-imaged, which meant that this sieve method was not available for some of the observed features. However, in the course of our image analysis, we learned a lot about the appearance of surface debris and impact damage, and how to use image shape and contrast to distinguish them from AHE-type holes. In the course of these three types of experiments we saw no features which appeared to be caused by LENR events.

The following four cases were less extensively pursued, and used only pure Pd targets, not any of its alloys.

Heater Pad Stimulation: This set of experiments placed arrays of nanofilms onto resistively heated pads in a chip format. Functionally, this is similar to the Staring experiments discussed above; still thermal stimulation, but applied resistively rather than by a laser. These tests were discussed in the Chip section of this report.

Cryogenic: In this experiment (BE4131) we used no explicit stimulation; the goal was to look for LENR in a cryogenic situation. The sample was loaded in 145 ksi D_2 and held at -160ºC for 12 hours, so it formed stoichiometric PdD, having very little thermal stimulation, but a greater chance of Bose-Einstein effects than in other cases.

Electric Field Stimulation: In some experiments, we used the electrical capacitor arrays, described in the Chips section of this report, for nanofilm archaeology. These looked for single

LENR events caused by intense electric fields (both static and oscillatory). Operationally, these archaeological tests were similar to our normal capacitive ones, the only difference being the introduction of high resolution SEM imaging.

High Pressure Electrochemistry: After developing nanofilm archaeology, we naturally wanted to use it to look, with single-event sensitivity, for LENR caused by electrolytic loading, not just the gas-loading that all our other nanofilm experiments used. We employed the same type of linked-together Pd nanofilm arrays developed for the Chip capacitor tests as the cathodes for high-pressure D_2O-based electrochemistry.

As these arrays are different from those used for gas-loaded experiments, they're shown below (**Fig. 19**), along with one of the nanofilms as it appeared after the experiment.

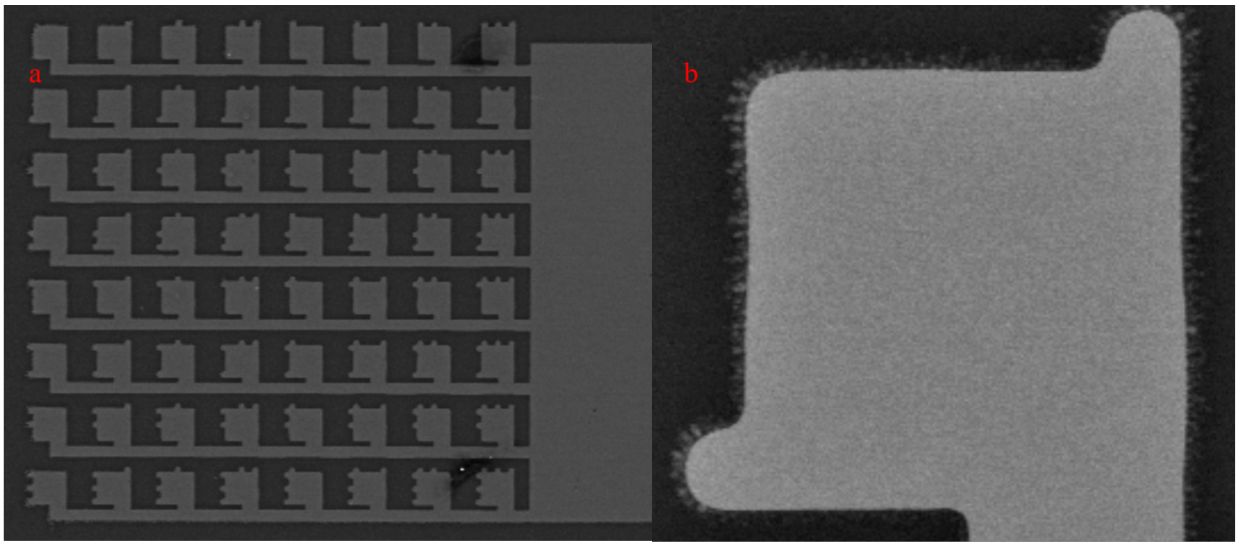


Figure 19. Post-experimental SEMs for BE4199. (a) The 64 linked-together nanofilms of pad C. (b) Nanofilm C11. Note the filamentary features extending from the edges of the nanofilm.

While not as extensive as the three upper cases of nanofilm archaeology experiments, none of these four types of experiments detected any signs of LENR events.

Magnetic Archaeology

Executive Summary

Individual LENRs are believed to deposit energy locally and quickly, only affecting a small volume (~ 100 nm in size) for a very short time (~ 1 psec) before the energy dissipates into the surrounding material. Detecting temperature rises, at unknown space-time locations, with either this temporal or this spatial resolution is challenging enough; doing both simultaneously is well beyond our current capabilities.

In principle, thermal history recorders offer us an attractive way to reduce this dimensionality; they'll tell us whether or not an event has occurred at a given site, but not when it did so. This is a tradeoff well worth taking. However, there is one caveat; to be useful for LENR detection, the "recorder" must be quick enough to respond on psec timescales; otherwise LENR events will be washed out and missed. Conventional history recorders, such as thermal paints or phase change materials rely on chemical reactions or atomic motion, and are far too slow to detect psec-scale events. However, magnetic effects, being based on electron spins, occur much faster, and will respond on sub-psec timescales. Accordingly, we have developed magnetically-based thermal history recorders, and used them in our magnetic archaeology campaign to try to detect LENR.

Single domain, nanoscale magnets have been used for decades in the data storage industry as a method for storing information in a non-volatile medium. The information is saved to memory by flipping or reorienting the natural magnetic polarity of a local metallic film structure, domain, that can later be read back effectively as a 0 or 1. The magnetic fields are normally very stable, and can persist for years. However, the nanomagnets can become destabilized from 1 to 0 if locally heated to a high enough temperature. This thermal-destabilizing effect, responsive to local heat input, is what made this method of detection truly attractive for detecting LENR events. At these length scales and the temperature requirements for a magnet to flip, we were able to form nano sized dots (≤200 nm) that could detect a single 24 MeV D-D event within a few 100 nm of its center. With Pd films adhered in close proximity to these sensors we have not observed any observable quantities of heat that can be attributed to anomalous heat events (AHE) or low energy nuclear reactions occurring in the Pd lattice. Methods implemented here were focused on magnetic sensor design and read-out of magnetic signals from these sensors making sure to maximize signal to noise in both magnetic polarity resolution and spatial resolution. FePt alloy films were developed for nanodot sensors, and magnetic force microscopy (MFM) as well as scanning tunnel magneto resistance (TMR) were used as sensor readouts.

Background

The fundamental mechanism magnetic archaeology is based around, is that the alignment of magnet dipole moments becomes disordered at elevated temperatures, at which point they are susceptible to outside magnetic field influence. A hard magnetic material is generally understood to possess high remanent magnetization (M_R), meaning in the absence of large externally applied fields (H) it retains a relatively large density of magnetic dipole moments that are oriented along a single vector, compared to the density at full saturation (M_S). Hard magnets also require larger coercive H fields to flip the orientation of these dipoles from vector $\hat{\imath}$ to $-\hat{\imath}$. This is depicted through its coercivity, H_C. At increased temperatures both M_R and H_C typically decrease, and at a critical temperature, referred to as the Curie Temperature (T_{Curie}), these values are reduced to 0 where there

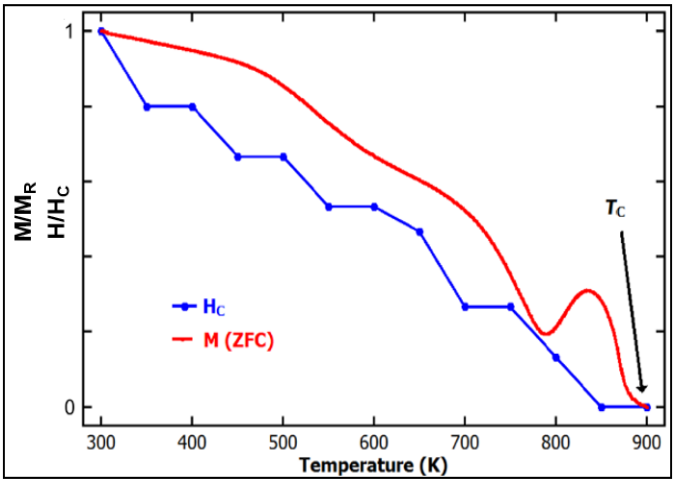


Figure 1: Thermo-magnetic plot of the coercive field strength (blue), normalized to the room temperature value, indicating switching field requirements. Zero-field-cooled (ZFC) thermos-magnetization plot (red; $\mu_0 H = 0.01$ T), also normalized to its room temperature value, indicating a Curie temperature of TC < 900 K.[1]

is no measurable long range magnetic order. This loss of magnetic order persists while the material remains above T_{Curie}, but upon cooling, the polarity of each atom can still be reoriented if exposed to an external magnetic field.

Understanding and applying this basic principle above, people have been able to store information through magnetic polarity vectors pointing up or down on a 2D matrix of nanodomain structures. As mentioned in the summary, this phenomenon has been extensively applied and optimized for non-volatile memory storage in computer hard disk drives (HDDs). A hard magnetic thin film that contains trillions of nanodomains (bits) on the surface are polarized up or down (0 or 1) using a magnetic write head that scans across the surface applying brief but high intensity magnetic fields localized to <100 nm. Standard hard drives utilize perpendicular magnetic recording (PMR) where the vectors point out of the platen plane. State-of-the-art hard drives use heat-assisted magnetic recording (HAMR) methods to (1) decrease writing field needed to flip a domain, (2) increase

domain or bit stability at room temperature, (3) and further decrease size of the domain size (<10 nm) increasing memory density >1Tb/in^2). The stability of a domain with respect to temperature is inversely proportional to the domain's coercivity and magnetization (Figure 1[69]).

The complexity of the film stack in HDD substrates or platens is more complicated than just having a single hard magnetic layer on top of a rigid glass substrate. In Figure 2 it can be understood that there are ≥7 layers that each contribute to the writability, bit stability, and scratch/corrosion protection of the hard magnetic film. For the purposes of this report please refer to more detailed literature that explains form and function of PMR HDD technology[70,71].

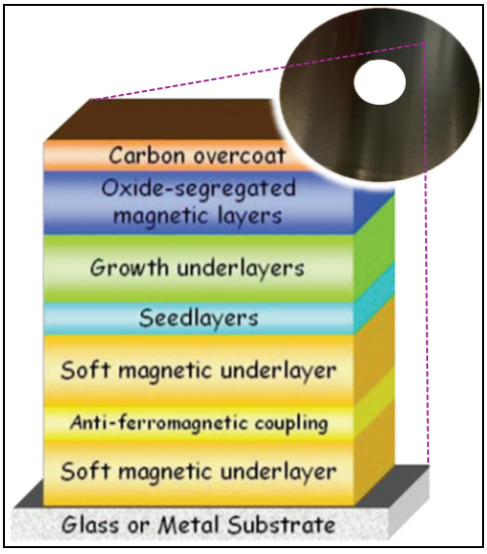


Figure 2: Traditional film stack to support reading and writing patterned magnetic media. The layers in green, lightt blue and yellow are to assist in bit stability, grain orientation, and signal enhancement for read-out. The dark blue layer is the film that actualy stores information in a hard magnetic material.

The ability for local heat sources to be able to distort and destabilize the magnetic structure of a magnet makes these material systems very useful for detecting AHE in Pd thin films, nanoparticles, and other Pd form factors that have at least one dimension reduced to the micro/nanoscale. It can effectively record a thermal event with high spatial accuracy that we can then read back using techniques that detect magnetic field, fluxes, or gradients. Literature has shown magnetic structure can be changed at picosecond and femtosecond timescales when using lasers as the pump and/or probe[72,73,74]. Since the thermal effects from D-D fusion type nuclear events may persist for

[69] Scheunert, G.; Cohen, S. R.; Kullock, R.; McCarron, R.; Rechev, K.; Kaplan-Ashiri, I.; Bitton, O.; Dawson, P.; Hecht, B.; Oron, D. Beilstein J. Nanotechnol. 2017.
[70] Kane, Margaret M. "Fabrication and characterization of perpendicular magnetic anisotropy thin-film CoCrPt grown on a Ti underlayer". 2015.
[71] Dong, Y., "Micromagnetic Study of Perpendicular Magnetic Recording Media", PhDT, 2012.
[72] Proenca, M.P., Muñoz, M., Villaverde, I. *et al.* Deterministic and time resolved thermo-magnetic switching in a nickel nanowire. *Sci Rep* **9**, 17339 (2019).
[73] Vahaplar, K et al. "Ultrafast path for optical magnetization reversal via a strongly nonequilibrium state." *Physical review letters* vol. 103,11 (2009).
[74] Hirofumi Suto, Akira Kikitsu, Yoshinori Kotani, Tomoyuki Maeda, Kentaro Toyoki, Hitoshi Osawa, Nobuaki Kikuchi, Satoshi Okamoto, Tetsuya Nakamura; Time-resolved imaging of an operating hard-disk-drive write head using nano-beam x-ray magnetic circular dichroism. *J. Appl. Phys.* 7 October 2020.

picosecond or longer timescales, these materials as magnetic sensors should have more than enough time resolution to capture an event that occurs in close proximity. Furthermore, since the read-out method only needs to measure changes in magnetic field direction and/or amplitude, this method of detection does not require permanent film damage (Pd or sensor) for the user to detect if a thermal event/reaction occurred! Changes in magnetic structure do not necessitate physical changes to the lattice structure. This is crucial in data reliability, repeatability, and reproducibility between sample to sample and experiment to experiment. We do not need to be as careful with surface cleanliness and local debris prior to an experiment since most debris will not have a magnetic signature and remain invisible.

The sequence of events of how a magnetic domain would record an AHE event is depicted in Figure 3. We first magnetize all domains in a single direction, e.g., "North" (pointing out of the surface), using a magnetizer set to pulse a 4 Tesla field through the film. Once the Pd film is loaded under ultrapressure conditions the stimulation pulse, either through Joule heating or laser absorption, would stimulate the Pd up to a threshold temperature that is still below the switching temperature (T_S) for the magnetic domains for a given external field (H_{ext}) amplitude. If an LENR-related thermal event occurs this would further increase the temperature of the Pd, and by conduction the magnetic sensor, above T_S, switching local domains in the direction of H_{ext}. After

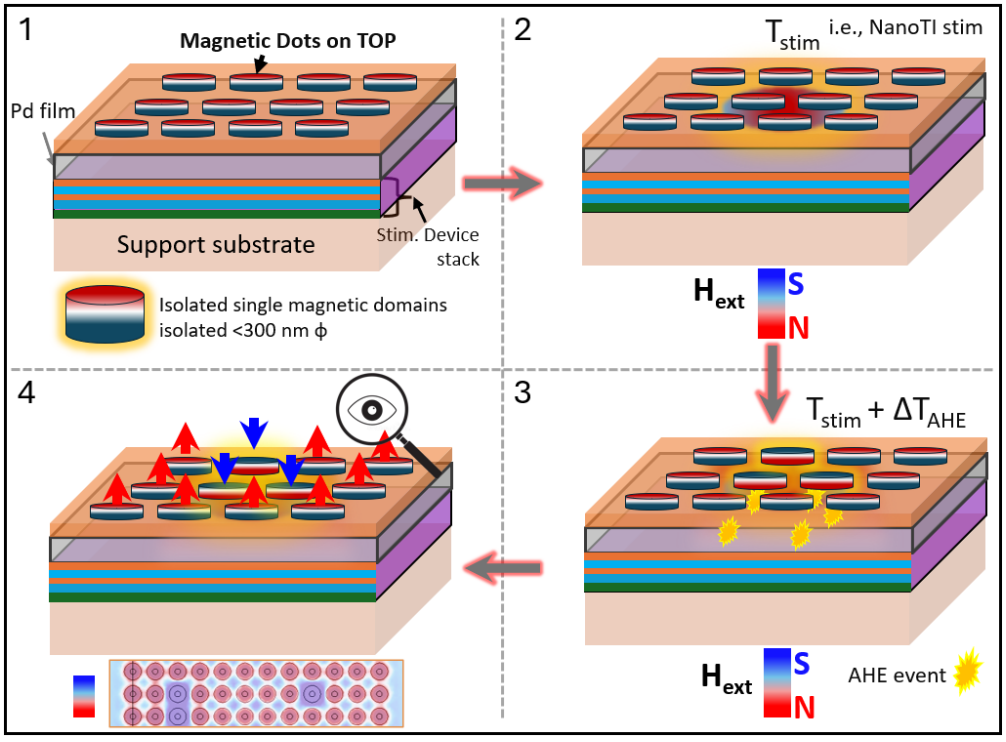


Figure 3: a)All spatially isolated single domains are mangetized in one direction in a preconditoned state. b) The Pd is thermally stimulated from external source to below a switching threshold temperature. c) If AHE forms locally the additonal input energy will increase the dots temperature enough to flip the domain orientation to the direction of the applied H field. d) post inspection of themagnetic signal for each domain is carried out to see if domains were altered during the course of the experiment.

the system cools back down to room temperature (within μs→ ms) the switched domains now persist within the film structure, and their magnetic fields can be read out in post experiment characterization to confirm existence of local LENR.

The main drawbacks to this detection method are that we are currently limited to thin films, and we cannot know, in real time, at what point in an experiment a thermal event occurs, we can only know if it did or did not occur *ex-situ* with high spatial sensitivity.

The methods used to read or scan the surface of each sample, magnetic force microscopy and tunneling magnetoresistance, will be discussed further in the methods section. However, the reasoning for selecting these two methods of magnetic detection came down to spatial resolution, ease of implementation, and ability to scan quickly over 100s of microns while maintaining high fidelity. The magnetic force microscopy (MFM) technique is a proven tool for qualitatively measuring changes in emitted fields from the surface of a material, in this case a planar film. This technique allows for scanning at high spatial resolution, but at modestly small scanning arrays <50x50 μm^2. Tunneling magnetoresistance (TMR) microscopy allows for a quantitative readout of resistance changes in a sensor's circuit from changes in local fields making it more robust. This technique, which is used by read heads in commercial hard disk drives (HDDs), required significant design and prototype work though, since high spatial resolution TMR sensors are optimized for steady, non-contact, uni-directional scanning inside an HDD, not the bi-directional, surface-contact, raster scanning used in our application.

Methods

Methods : Detecting magnetic polarity

Given the thin film morphology for these magnetic materials, the domains that are polarized up or down have exceedingly small domain volumes with emitted field (*B*) decaying as a function of inverse distance cubed from the center of the domain (*Equation 1*). First principle magnetic simulations at the micro and nanoscale indicate that >99 % of the magnetic field intensity is lost at ≥1 µm gap away from the thin film surface to the read-out probe that is measuring the field (Figure 4). This necessitates the requirement that our read-out method for AHE magnetic sensors must be able to probe the emitted fields <<1 µm away from the film surface to maintain a high signal to noise ratio.

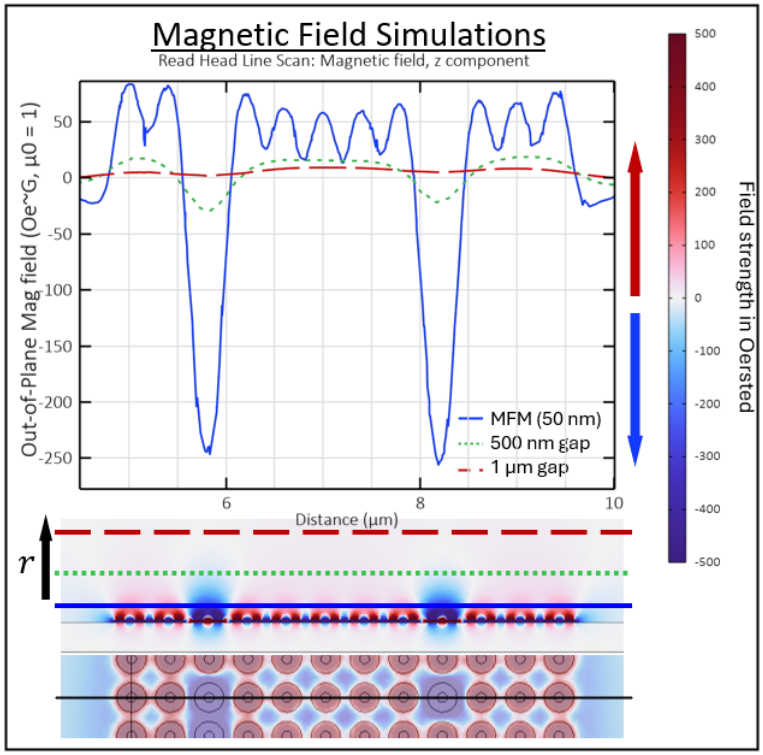


Figure 4: COMSOL permanent magnet simulation. Blue, green, and red scans plotted are line scans from a probe at a set distance ***r*** *from the magnetic domain surface. Y axis indicates the magnitude and sign of the emitted field strength. Domains are polarized up or down from the film surface (50 nm- 1 µm).*

Equation 1

$$B = \frac{\mu_0 I A \sqrt{1+\cos^2(\theta)}}{4\pi r^3}$$

1: Polar Magneto Optical Kerr Effect (MOKE) Microscopy

This technique takes advantage of changes in photon incident and reflected angles off a magnetic material surface. Using a polarized ultraviolet light source, the Kerr rotation of the polarized light can be captured over the entire focal plane making it an ideal technique for obtaining magnetic data over large surface areas in a short period of time[75,76]. The contrast can be enhanced taking the differential between 2 images of positive and negative angle polarizer rotation, increasing contrast between domains that have opposite magnetic polarity[77]. The Kerr rotation is directly related to the sign and amplitude of the local magnetic field the light interacts with, as its irradiance drops exponentially under the surface. The great advantage for this technique is that it is able to quickly obtain magnetic polarity or moment maps of the sample surface over 100s µm in a single image. An example of a polar MOKE image is in Figure 5 from a credit card magnetic strip. However, this technique was ultimately not implemented in our experimental approach for magnetic characterization and analysis due to fundamental optical limitations in lateral resolution based off the Rayleigh criterion for available 450 nm wavelength sources and 50x objective lens. This means we would not be able to reliably resolve <500 nm magnetic features with high confidence.

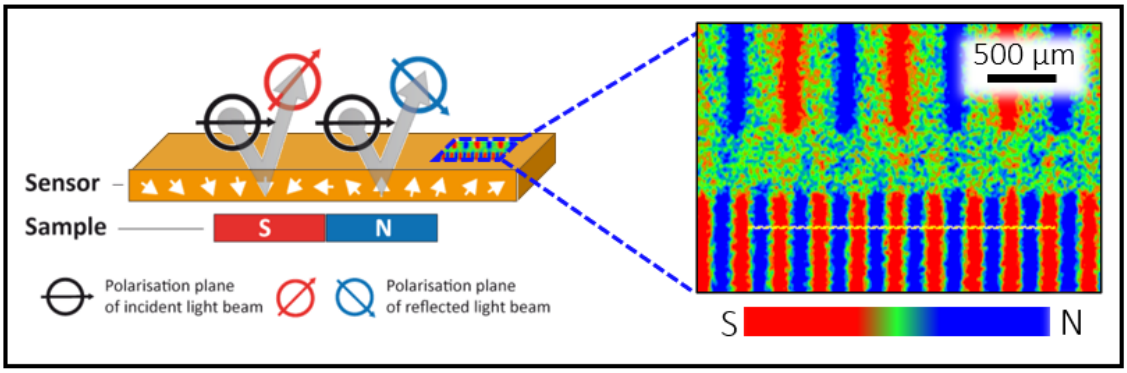


Figure 5: When the MOKE sensor is in close proximity to the sample surface the soft magnet in the MOKE sensor reorients to align with the moment polarity in the sample. A polarized light is reflected off the back of the sensor film being rotated slightly in the local orientation of the magnetic field upon reflection (left). The incident and reflected signals are differentiated to produce the blue to red color gradient image (right). Image is a magnetic strip.

[75] Jeffrey McCord 2015 *J. Phys. D: Appl. Phys.* **48** 333001.
[76] M. Takezawa, K. Maruko, N. Tani, Y. Morimoto, J. Yamasaki, T. Nishiuchi, S. Hirosawa; Magnetic domain observation of hydrogenation disproportionation desorption recombination processed Nd–Fe–B powder with a high-resolution Kerr microscope using ultraviolet light. *J. Appl. Phys.* 1 May 2010; 107 (9): 09A724.
[77] H. Z. Yang, Y. J. Chen, S. H. Leong, C. W. An, K. D. Ye, M. J. Yin, J. F. Hu; A multi-functional testing instrument for heat assisted magnetic recording media. *J. Appl. Phys.* 7 May 2014; 115 (17): 17B726.

2: Magnetic Force Microscopy (MFM)

MFM is a technique that operates identical to a traditional AFM that scans across the surface of a sample, usually moving in a serpentine path. The atomically sharp tip (≤80 nm radius) is coated with a thin layer of CoCr which is traditionally defined to have medium/hard magnetic properties. The tip is polarized parallel to the tip axis, so the CoCr domains have a polarity pointing up or down when the cantilever is mounted into the AFM chassis and therefore is most sensitive to out-of-plane stray magnetic fields. When the tip is near a local magnetic field with high magnetization it will deflect and change the phase of the resonant frequency of the cantilever depending on the polarity of the stray field (Figure 6). MFM typically requires 2 scans across the same line path. For this project, the first pass scans the surface topography in AC tapping mode, the second (also called a nap scan) lifts the tip above the surface so the amplitude and phase of the resonant frequency are only influenced by external magnetic fields or local charges, not topography. The charge effects are mitigating by grounding the tip and sample to a common ground. This allows us to obtain both topography and magnetic contrast over the same sample surface area at <10 nm resolution (Figure 7 and Figure 8)[78].

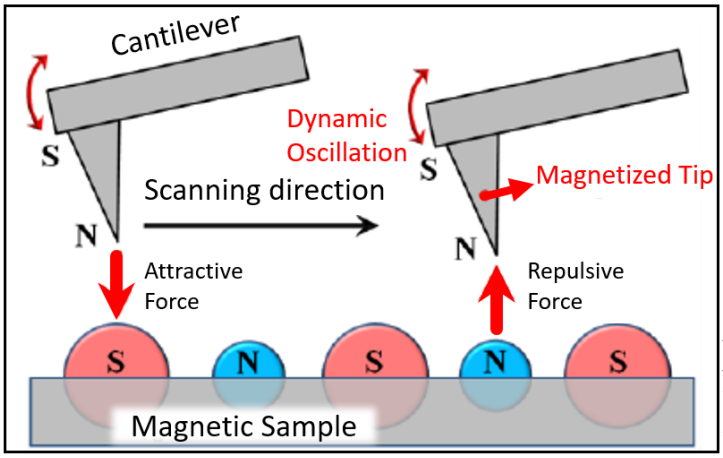


Figure 6: Pictorial diagram of how the AFM-MFM tip interacts with a magnetic surface with local variation in North or South magnetic polarization.

For this research all MFM scans are set-up so that an increase in phase angle indicates a sample domain is polarized north (field out of the surface) and a decrease in phase angle means the domain is polarized south (field into the surface). This method is inherently slow, being limited to small scan areas to maintain exceptional spatial fidelity, and without calibration is a qualitative measurement for magnetic polarity only, north or south. In detecting AHE, we only need the north or south polarity information and not necessarily the amplitude. This technique was the main

[78] T. Bublat, D. Goll; Influence of dot size and annealing on the magnetic properties of large-area L1-FePt nanopatterns. *J. Appl. Phys.* 1 October 2011.

diagnostic used for measuring magnetic polarity in all sample types through the course of magnetic archaeology projects.

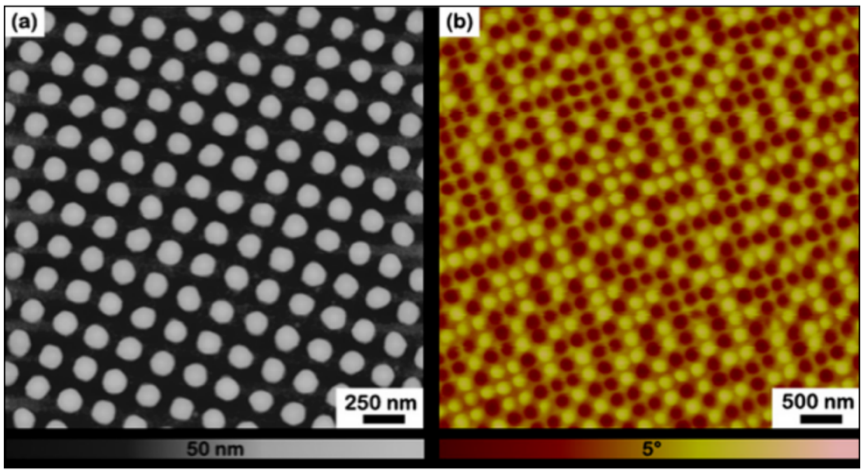


Figure 7: a) Topography image from first scan, AC tapping mode of magnetic dot features. b) MFM phase contrast on the nap scan showing each dot contains a single domain oriented up or down with repect to the sample surface.

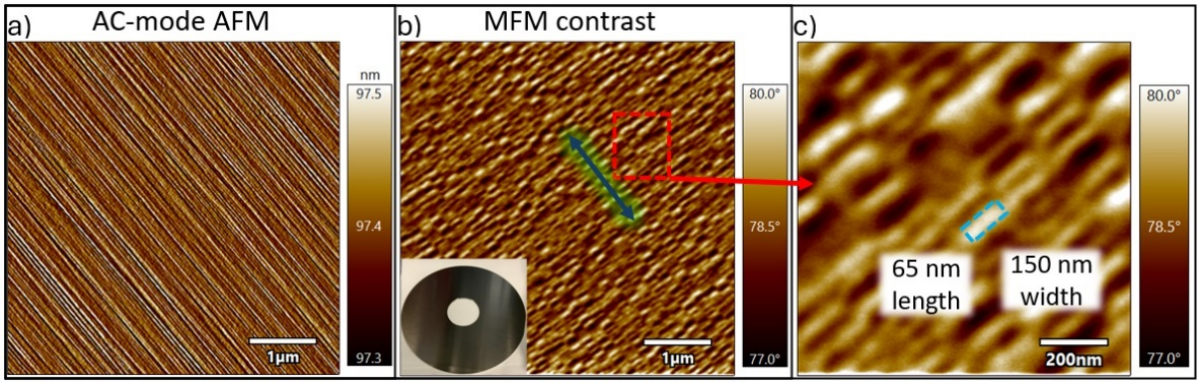


Figure 8: AFM scan of HDD platen used in our chips as starter substrates. a) topography scan showing low roughness striations from the read/write head arm. b,c) MFM contrast scan showing "racetrack" of magnetic bits along the write head path.

Image Analysis Pipeline

To speed up the analysis process for comparing pre and post experiment MFM magnetic polarity scans for each individual dot/magnetic domain on samples with >10,000 dots, an image filter software pipeline was developed to automate this visual process. The program located each dot within the MFM scan as seen in Figure 9a by first looking at the topography image collected at the same time and location the MFM contrast data is collected. An edge detection module is employed to locate the edges of all dots in the image based off sharp gradient changes in pixel intensity. With the location of each dot or region of interest (ROI) logged, ROIs are then mapped directly onto the MFM scan, Figure 9c. Depending on the threshold values for pixel area and average pixel intensity, set by the user, the ROI for each dot is filtered as having either North (red) or South (blue) polarity.

We also accounted for physically missing dots in the array and negated the polarity information in the MFM scan for that area, leaving a white space in the final area map.

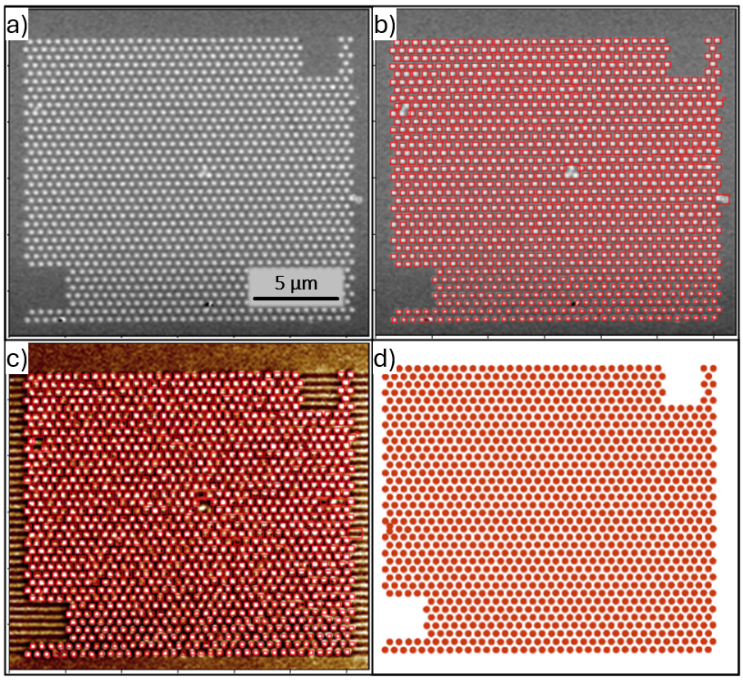


Figure 9: a) Topography scan used for initial map for located physical coordinates for each dot feature in the array. b) Resulting ROI boxes outline boundaries for each dot. c) ROIs from b) mapped to the MFM scan where the pixel intensity within each ROI are averaged and normalized to background pixel intensity (a dark brown shade). d) based on pixel intensity thresholds each dot is measured to be either oriented North (red) or South (blue). In this case this array is in a preconditioned state of all North polarized domains.

3: Tunneling Magnetoresistance (TMR) Microscopy

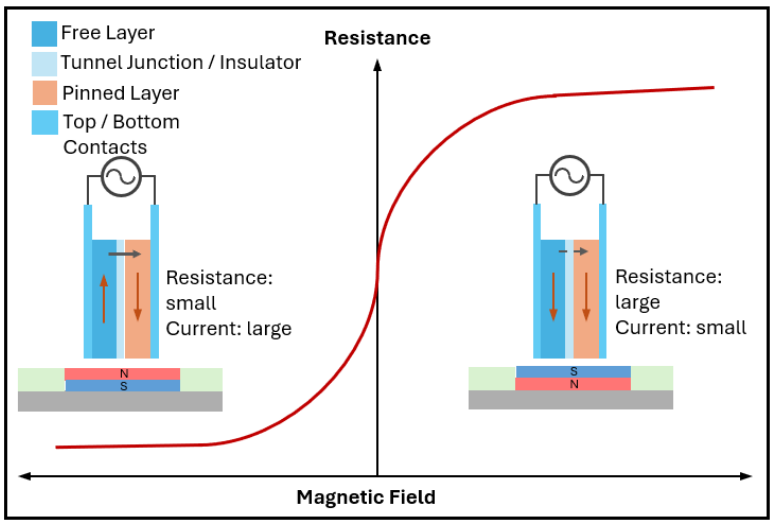


Figure 10: Resistance change of a TMR circuit as a function of local magnetic field change from north to south until the tunnel junction is saturated.

The last method of read-out used for measuring our samples is TMR which exploits a magnetoresistive effect that occurs in a magnetic tunnel junction with an ultrathin insulating

layer[79]. One side of the junction is pinned to have a magnetic orientation that does not change while the free layer can be flipped in orientation to be parallel or antiparallel by local external fields of an HDD or magnetic sample in general[80,81]. Measured sensor resistance is dependent on the local magnetic field of a sample surface and is a quantitative measurement of the local magnetic vector down to a spatial resolution of roughly 50 nm and dependent on the physical dimensions of the TMR junction (Figure 10). This technology is used in read heads of hard drives to achieve high storage densities and has been optimized for flying over an HDD platen spinning at >5000 rpm on what is called a slider; billions of dollars have been gainfully invested in this development. However, the samples for AHE experiments could not be easily scanned by high-speed rotation and therefore we needed to design our own in-house TMR microscope that could scan selective areas of a sample with a resolution requirement of ≤100 nm.

To solve this problem, TMR read heads were salvaged from existing HDD units and operated using an AC voltage driven source. Data read out was done by monitoring variation in current across the tunnel junction and calculating the change of resistance. Two off-axis optical microscopes were used to locate the sample alignment marks while nanometer resolution, single axis x, y, and z piezo motors were used to drive the sample position underneath the TMR read head with sub 5 nm precision (Figure 11). A more detailed set of SEM images are in Figure 12 showing the actual tunneling junction. An example TMR scan can be seen in Figure 13 for a sample containing 300

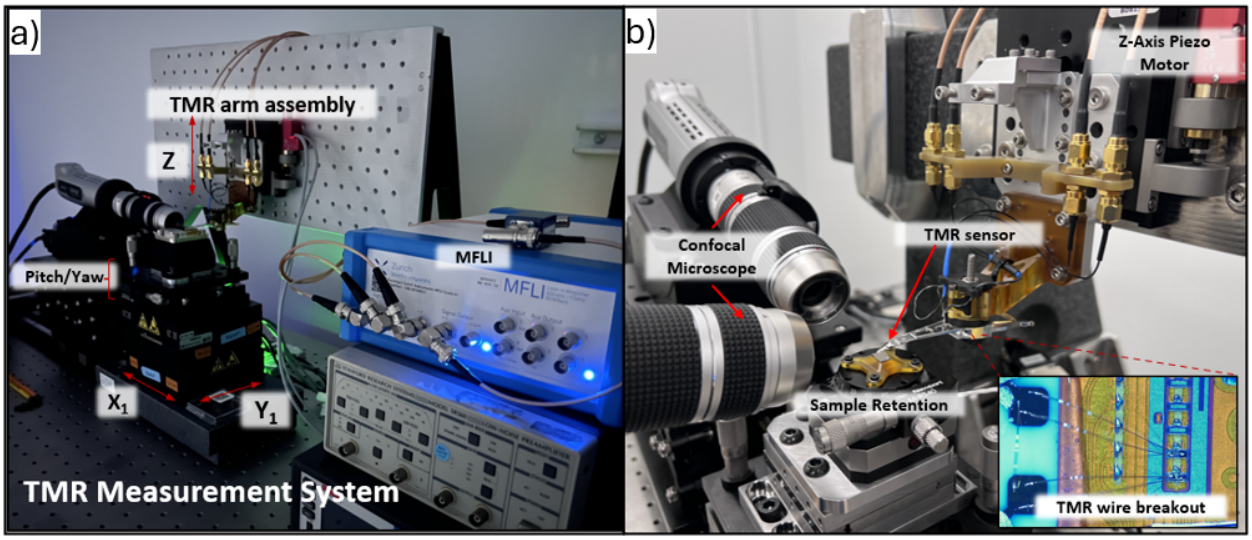


Figure 11: a) TMR apparatus overview. TMR assembly contains the read head details in the next figure. b) close-up image of off-axis confocal optical microscopes pointed at sample surface to align sample fiducials. Detail view shows the custom wire bonding required to interface with and control the read head sensor that was harvested from the HDD.

[79] Zhang, X., Butler, W. (2016). Theory of Giant Magnetoresistance and Tunneling Magnetoresistance. In: Xu, Y., Awschalom, D., Nitta, J. (eds) Handbook of Spintronics. Springer, Dordrecht.
[80] D. Mitin, M. Grobis, M. Albrecht; Scanning magnetoresistive microscopy: An advanced characterization tool for magnetic nanosystems. *Rev. Sci. Instrum.* 1 February 2016; 87 (2): 023703.
[81] Nakatani T, Gao Z, Hono K. Read sensor technology for ultrahigh density magnetic recording. MRS Bulletin. 2018;43(2):106-111.

nm FePt dots having all their polarities oriented in the same direction. The dots appear elongated due to slight drift in the motor positions during the time it takes to complete one scan. Unfortunately, this measurement technique was never deployed for measuring actual AHE samples due to the project being wrapped up prior to full TMR readiness.

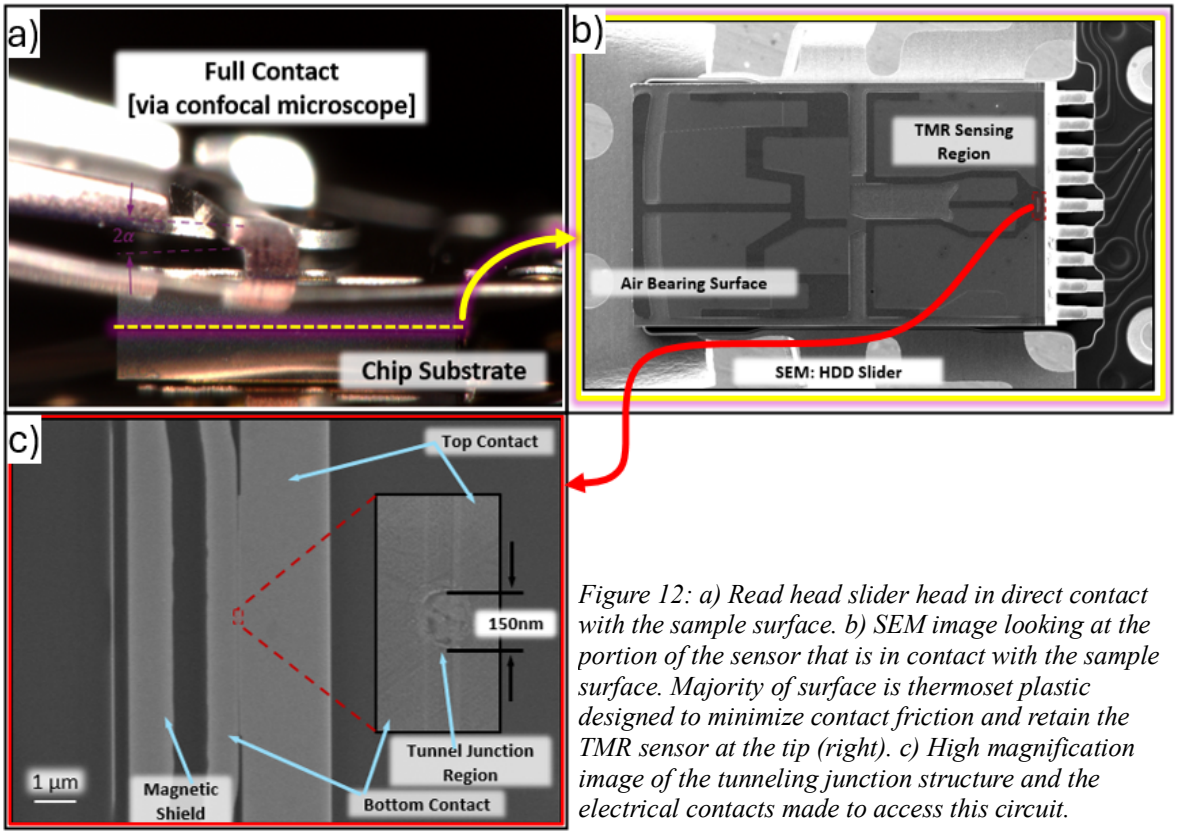


Figure 12: a) Read head slider head in direct contact with the sample surface. b) SEM image looking at the portion of the sensor that is in contact with the sample surface. Majority of surface is thermoset plastic designed to minimize contact friction and retain the TMR sensor at the tip (right). c) High magnification image of the tunneling junction structure and the electrical contacts made to access this circuit.

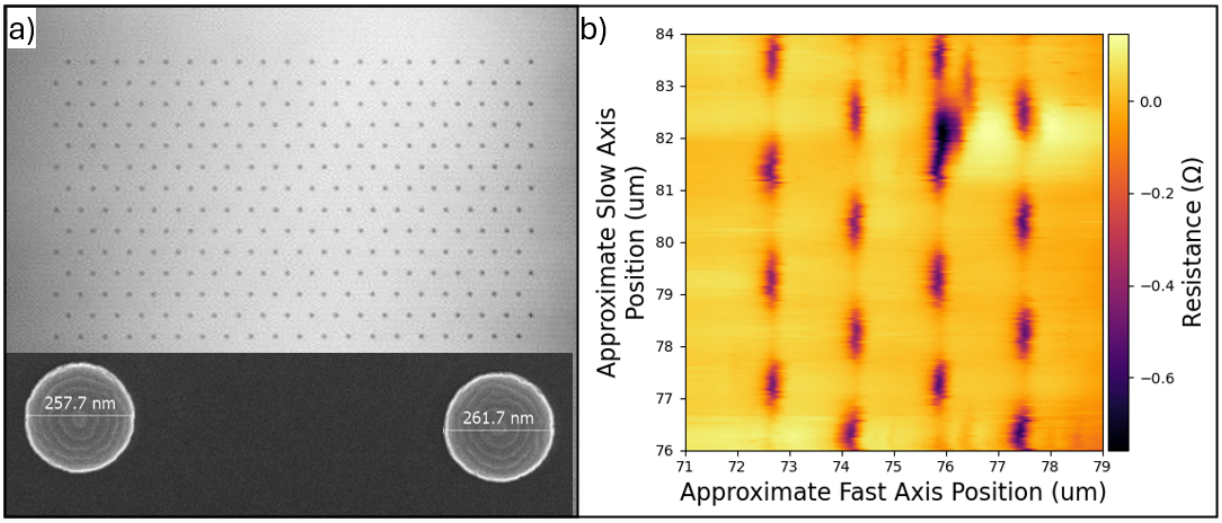


Figure 13: a) SEM image and high magnification image of 250-300 nm dots scanned by TMR system. b) TMR resistance-change read out over an 8x8 µm area. The dots were polarized in a Northward polarity indicated by a decrease in resistance. The dots are elongated in the slow axis due to small drift in the piezo motor position over the course of the scan.

Methods : Fabricating Samples

Materials selection for the magnetic AHE sensor was based off multiple requirements that were chosen to maximize phase stability (temperature and chemical), minimize susceptibility to sensor degradation from handling and extraneous fields, and minimize complexity of film architecture. Selection criteria are listed below.

1. High coercivity (H_C) at room temperature (>5 kOe) up to 200 °C
2. Curie temperature (T_C) variable between 350-900°C
3. Single domain size with high magnetocrystalline and shape anisotropy down to 50 nm
4. Out-of-plane magnetization with magnetic easy axis aligned perpendicular to substrate surface
5. Minimal reactivity to hydrogen gas
6. High remanent magnetization (M_R) flux density for bit stability both *in situ* and *ex situ*
7. Can be lithographically patterned, etched, and coated while maintaining compositional and phase integrity

The two material systems chosen were CoPtCr and FePt. Both have extensive literature in memory applications, and therefore have a good track record of fulfilling criterion 1-4. Internal tests were completed to confirm or deny the last 3 requirements essential for these magnetic archaeology tests.

1: CoPtCr- Standard for PMR in HDD platens

Instead of ourselves synthesizing the ternary phase film structure of magnetic CoPtCr using complex sputtering processes and exotic metallic seed layers, initial magnetic archaeology magnetic sensors utilized the existing film architecture in HDDs that can be purchased off the shelf. The HDDs contain between 2-6 platens that are glass substrates coated with CoPtCr films ready to be used in storage applications. The general film stack is seen in Figure 14 where you can see the complexity of this technology which has already been optimized for this specific industry. Purchasing allowed us to cheaply obtain many HDD platens with effectively <100 nm domain sensors within days instead of months. Subjecting these platens to ultrapressure hydrogen conditions sooner rather than later in the project helped in affirming our hypothesis that these films were robust in these harsh environments at an early stage of development. It is important to clarify here, that ultimately for the final iterations of sample development in this project, we moved away from using HDD platens as the material for magnetic sensing. This was due to issues with repeated delamination, increased complexity of reading out the magnetic signature of this film, low Curie temperature (~150 °C), and low amplitude of M_R signal making it a noisy dataset to filter and analyze. However, as an initial, rapidly implementable, approach it was a useful step to take.

Experimental chips were fabricated using these platens as substrates. Chips for both NanoTI and Staring type AHE experiments were designed, fabricated, and tested with the CoPtCr sensor being buried underneath a layer of gas-loadable Pd. Due to the bulk service temperature limit of these platens roughly set at 150 °C all fabrication processes needed to minimize substrate heating to below 140 °C so as to not degrade the complex layer structure. The HDD platens firstly need to have the diamond like carbon (DLC) coating removed using a diamond/oxide etch clean, the DLC coating made the surface hydrophilic making it difficult to spin coat positive photoresist for further pattern development.

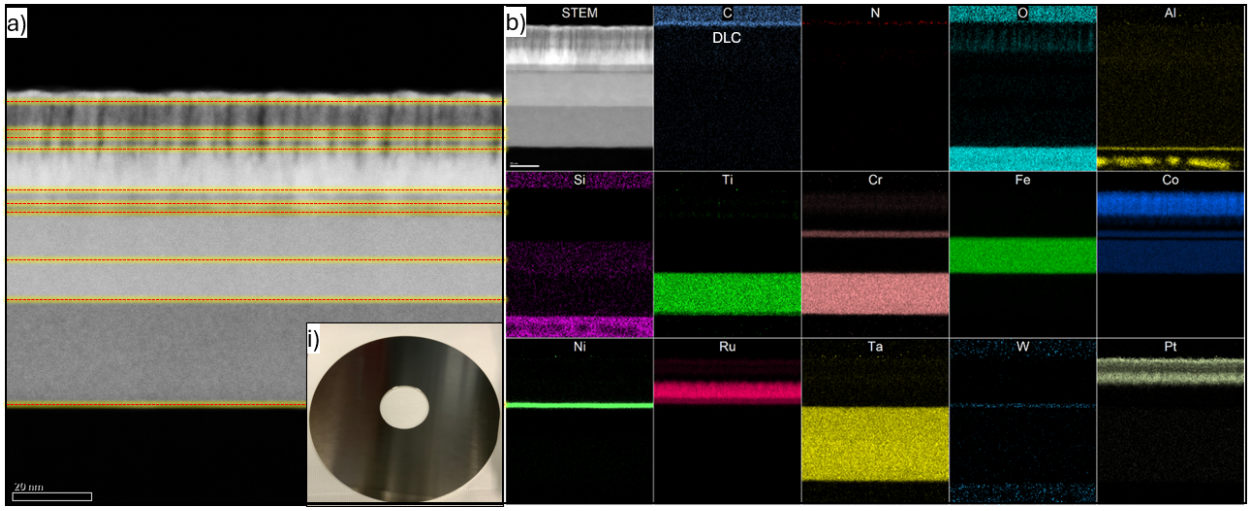


Figure 14: a) FIB lamella cross section imaged with STEM showing individual layers of the top 75 nm from the surface of an HDD platen. Below these layers is a SiO_2 based substrate up to 500 µm thick. There are >10 layers in this architecture. B) STEM-EDS maps of the same cross-section area in a). Each color indicating a specific element called out for each sub-image. The applications for each layer are similar to the general overview structure displayed in the concept section.

2: **CoPtCr – Staring Chips**

After DLC removal, 50 nm Pt markers were deposited using e-beam evaporation with a 5 nm Cr adhesion layer. 50 nm of SiO_2 acted as a passivation layer between the CoPtCr and Pd sample that is placed on top to prevent unwanted alloying and possible interference with the CoPtCr crystalline structure. It also slowed diffusion of protium or deuterium reaching the magnetic layer during a test. Plasma enhanced chemical vapor deposition (PECVD) was used for all dielectrics, including

SiO_2, for these chips. Finally, 50 nm Pd was patterned on top with 5 nm Cr adhesion layer allowing for unimpeded gas loading at the Pd surface (Figure 15).

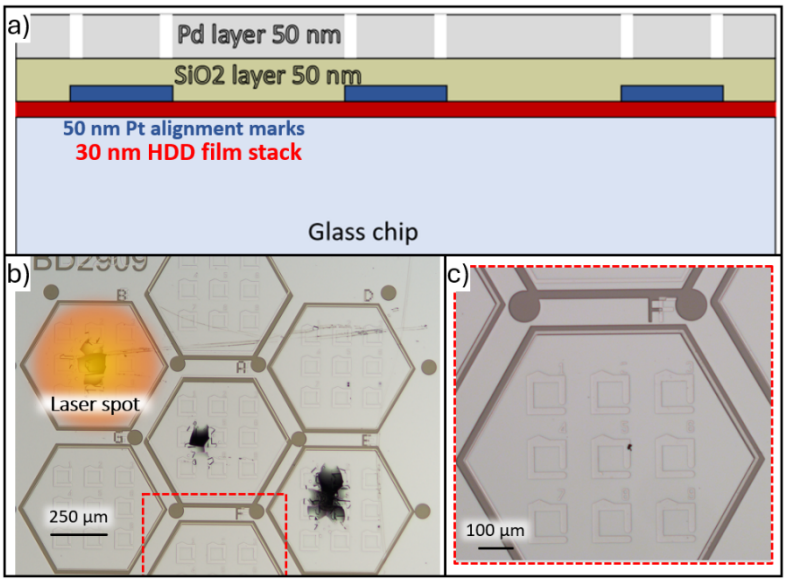


Figure 15: a) Chip cross-section in L-edit software showing multi-layer film structure. Pd is on top to ensure loading and direction absorption of laser power from Staring laser. b) Example top-down view of the HDD staring chip with stim pads A,B, and F damaged from laser shots. c) Higher magnification of individual laser stim pad with 9 sub-boxes (100x100 µm). These boxes are used to locate the MFM and or TMR sensors over the ROI prior to starting a pre or post scan for magnetic read-out.

3: **CoPtCr – NanoTI Chips**

The additional film stack for these chips used in the NanoTI experimental campaign are similar to the architecture used in the non-magnetic-archaeology chips used in NanoTI. For more detail on the methodology of how the chip operated for thermal coefficient calibrations and AHE experiments please refer to the Chips campaign section of this report. Chip fabrication for magnetic archaeology followed the film stack in Table 1. Up to eight 20 x 20 µm Pt serpentine heater devices, to provide resistive stimulation during the experiments, were patterned onto a single chip to account for lower device yield and shorter experimental turnaround times (Figure 16). Various sizes of Pd targets were made between 5-30 µm square to vary the peak temperature and thermal gradients reached for each target, stimulated at the same current and voltage parameters. Pulse widths for stimulation varied between 5-20 µs with 5-25% duty cycle.

Table 1: HDD -Thermal Arch NanoTI Fabrication Steps

Step #	Material	Method	Feature
1	5 nm Cr/ 50 nm Pt	E-beam Evap	Alignment marks
2	20 nm SiO_2	PECVD	Passivation
3	5 nm Cr/100 nm Pd	E-beam Evap	Pd target
4	100 nm SiO_2	Oxide Etcher (ICP)	Passivation/loading windows for Pd
5	5 nm Cr/ 100 nm NiCr	E-beam Evap	Pt heater/RTD
6	10 nm Cr/200 nm Au	E-beam Evap or Sputter	Interconnects/Contacts

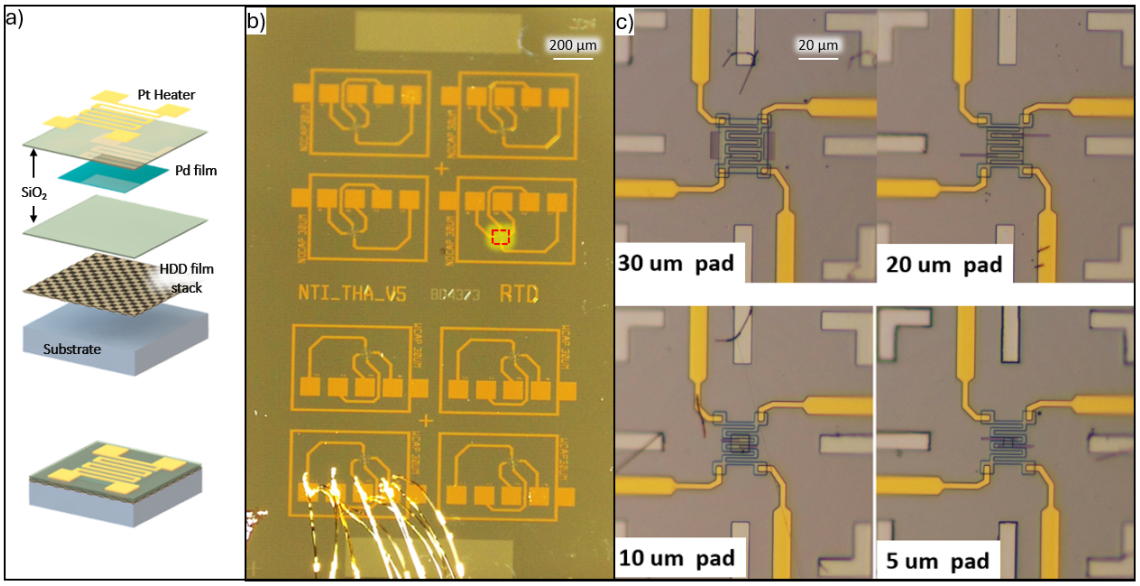


Figure 16: a) Exploded view of film stack used in HDD NanoTI type chip architecture. b) Top-down view of chip with 8 stimulation circuits with one wire-bonded to an alumina PCB ready for vessel installation/calibration. c) Pd pad circuits ranging from 5-30 µm in side-length with Pt heater on top.

4: FePt

The material system selected for producing our own magnetic sensors was the binary FePt system. Like CoPtCr, FePt is heavily studied in the memory industry, specifically in the HAMR field since it has a much higher coercivity at room temperature making it difficult to rely on applied field alone to "write" data. The magnetic phase of FePt is $L1_0$ fcc phase with an easy axis along the [001] direction which is grown to have the (001) plane parallel to the sample surface. The plane orientation is formed in two main processes in literature 1) using seed layers like MgO (200) single crystal substrates so the lattice parameters between the 2 materials match up in the preferred [100] orientation, (2) growing FePt on Si/SiO_2 and use rapid thermal annealing (RTA) heating the material up to >550 °C to cause thermal expansion mismatch between FePt and SiO_2 which is minimized along the FePt (100) plane[82,83,84,85] (Figure 17). To keep most traditional CMOS and silicon microfabrication processes easily transferable, the later method of achieving the $L1_0$ FePt phase was selected.

[82] Yingguo Peng, Jian-Gang Zhu, David E. Laughlin; perpendicular thin film deposited by alternating sputtering at elevated temperature. J. Appl. Phys. 15 April 2006; 99 (8): 08F907.
[83] Jae-Song Kim, Yang-Mo Koo, Byeong-Joo Lee, Seong-Rae Lee; The origin of (001) texture evolution in FePt thin films on amorphous substrates. J. Appl. Phys. 1 March 2006; 99 (5): 053906.
[84] Albrecht, M. and Brombacher, C. (2013), Rapid thermal annealing of FePt thin films. Phys. Status Solidi A, 210: 1272-1281.
[85] Liang-Wei Wang, Wen-Chieh Shih, Yun-Chung Wu, Chih-Huang Lai; Promotion of [001]-oriented L1-FePt by rapid thermal annealing with light absorption layer. Appl. Phys. Lett. 17 December 2012; 101 (25): 252403.

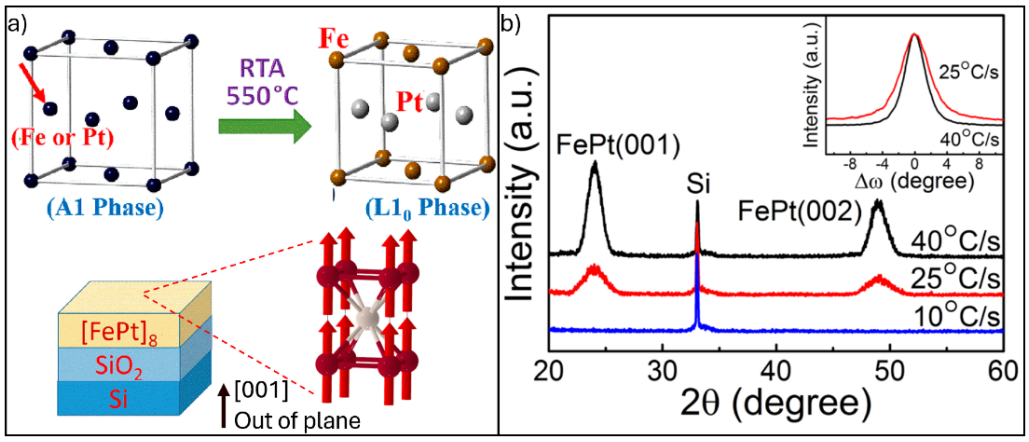


Figure 17: a) Overview of phase transition of as-deposited FePt A1 disordered phase film to L1$_0$ by thermal annealing. During this transiiton if strained properly the film will become textured to (001). b) Resulting XRD spectra for various heating rates indicating that high thermal mismatch is beneficial for FePt(001) texturing.

Having a hard magnetic material that we can deposit at any point in the fabrication process made patterning these films with either metal lift-off or reactive ion etching much easier than their CoPtCr HDD counterpart, and gave flexibility in experimental design to put the FePt sensor either on top or underneath the Pd film target.

Our standard process for producing FePt thin films with desired magnetic properties was to first use a DC-sputter target of FePt (52 at. % Fe) at room temperature in 5 mtorr Argon gas. Target power during sputtering was set to 75 W to deposit 20 nm thick films on Si (100) wafers with 3000 Å thermal SiO_2. The wafer is then thermally annealed in a RTA furnace at a ramp rate of 100 °C/s to a peak temperature of 750 °C and a hold time of 1 min in $\leq 1.0 \times 10^{-5}$ Torr vacuum. Film quality and structure was characterized using a MoK_α XRD source, AFM, EDS, and vibrating sample magnetometry (VSM). Film composition was 53 at.% Fe, 47 at.% Pt possessing similar XRD spectra (Figure 17). If annealing time is held longer than 2 min the L1$_0$ (111) peak forms indicating misorientation of the magnetic easy axis, decreasing magnetic anisotropy and therefore degrading out-of-plane properties of M_R and H_C. Surface roughness from AFM and SEM indicates average grain size to be 120 nm. A 5-10 nm thick layer of SiO_2 is coated on top of the FePt film using atomic layer deposition to protect the magnetic layer during further processing and handling.

All chip designs used this recipe for FePt synthesis and were deployed into 3 distinct experimental campaigns including NanoTI, Staring, and 400 ksi HPE femto-second laser stimulation. To enhance the coercivity for FePt, the films were etched to form isolated nanodots. This etch process formed our nanodot arrays. The physical isolation of domains from their neighbors has been shown to enhance the stability of individual single domains since the external stray fields emitted from individual neighboring domains is much lower even just a few 100 nm laterally from each dot center (Figure 18). This decreased the possibility of false positives from dot crosstalk due to strong stray fields close to the sample surface.

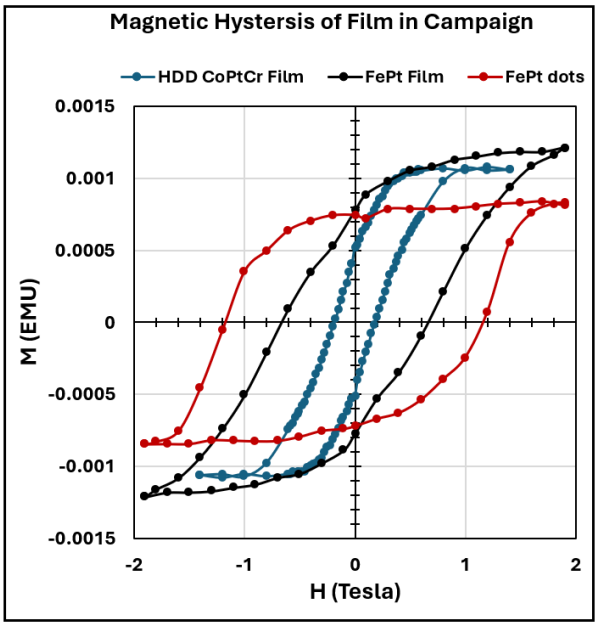


Figure 18: Magnetic Hysteresis loops collected at room temperature in nitrogen. Coercivity is enhanced for FePt after patterning to isolated nanodot length scales. Magnetization levels of FePt are slightly lower than HDD CoPtCr alloy. Film volumes kept constant with sample dimension and thickness. Dots likely have a higher M than shown.

5: **NanoTI Magnetic Archaeology Chips**

The fabrication steps for FePt sensors on NanoTI type chips are in Table 2 below. The starting substrate is Si (100) with 3000 Å SiO_2. Notice for these chips we were able to deposit the FePt layer on top of the NanoTI device (with its Pt heater and RTD) and the Pd film. The FePt dot features for this chip iteration form nano-ring-like structures which performed satisfactorily during AHE tests (Figure 19). Final chip dimensions were 5 x 7 mm.

Table 2: FePt NanoTI Fabrication Steps

Step #	Material	Method	Feature
1	5 nm Cr/ 100 nm Pt	E-beam Evap	Alignment marks
2	10 nm Cr/150 nm Au	E-beam Evap or Sputter	Interconnects/Contacts
3	100 nm SiO_2	PECVD	Passivation
4	5 nm Cr/ 50 nm Pd	E-beam Evap	Pd pads
5	5 nm Cr/ 100 nm NiCr	E-beam Evap	Pt heater/RTD
6	100 nm SiO_2	PECVD/Oxide Etcher (ICP)	Passivation/loading windows for Pd
7	20 nm FePt	DC Sputter/Positive E-beam lithography/Reactive Ion Etch	FePt 300 nm dots

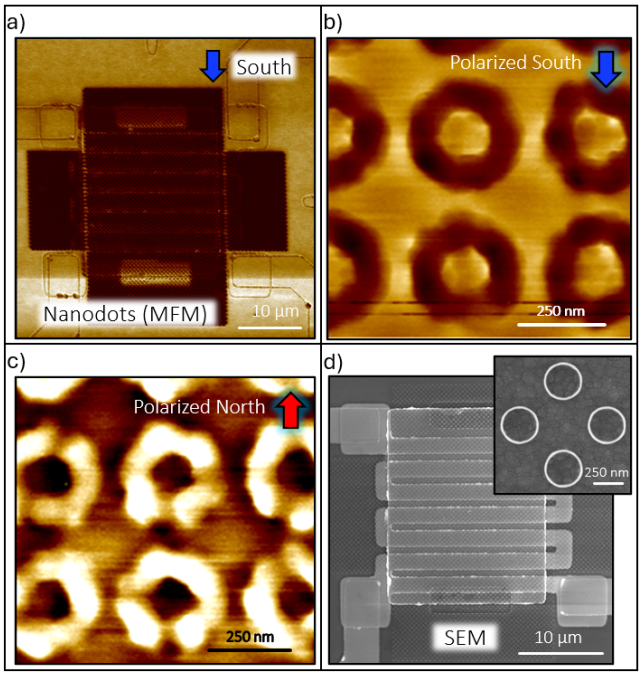


Figure 19: a) Nanodot array in preconditioned state all oriented down into the surface of the sample. b) Small area scan (1x1 µm) for dots polarized South. c) polarized North. d) SEM image of the same nanodot array, the dots are thin enough to be slightly transparent to the 10keV electron beam used to image the Pt heater. Note: these dots were processed using positive E-beam lithography leading to nanoring like strictures instead of dots, but the magnetic isolation and enhancement persisted in this morphology.

6: Staring Chips – Magnetic Archaeology

The fabrication steps for FePt sensors on Staring type chips are listed in Table 3 below. The starting substrate is Si (100) with 3000 Å SiO_2. The FePt dot sensor array is underneath the Pd for these samples. The final chip dimensions were 2x2 mm in order to fit into the 400 ksi vessel which had a bore of <0.25". With the FePt layer being annealed on a well characterized 3000 Å SiO_2 layer the remanent magnetization of these dots was more uniform than NanoTI chips allowing us to space the dots closer together and ultimately form sub 100 nm, single domain magnets for even greater spatial resolution (Figure 20). For most magnetic archaeology AHE experiments we utilized the 250 nm nanodot arrays instead of 100 nm due to having a larger stock of 250 nm generation chips (Figure 21). Aside from Staring, these chips were also used in 400 ksi experiments that stimulated the sample with a femtosecond laser.

Table 3: FePt Staring Fabrication Steps

Step #	Material	Method	Feature
1	20 nm FePt	DC Sputter/RTA 750 °C	FePt film
2	5 nm SiO2	ALD	Passivation/Resist adhesion
3	FePt	Negative E-beam lithography/Reactive Ion Etch	FePt 250 nm dots
4	FePt	RTA Anneal	Repair $L1_0$ phase
5	10 nm SiO2	PECVD	Passivation
6	5 nm Cr/ 30 nm Pd	E-beam Evap	Pd target

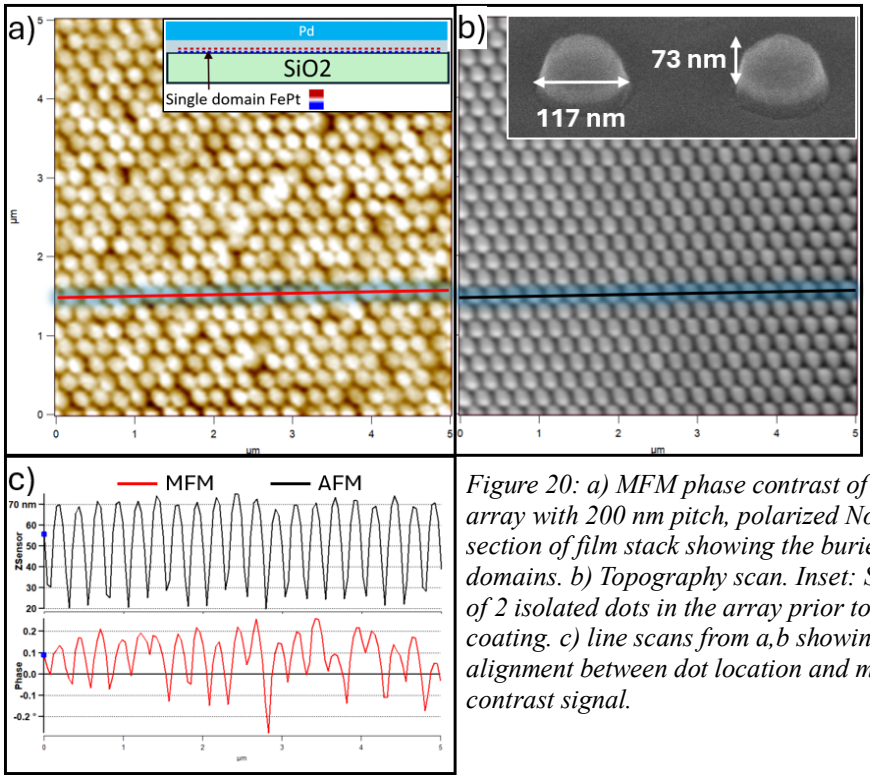


Figure 20: a) MFM phase contrast of 100 nm dot array with 200 nm pitch, polarized North. Inset: cross section of film stack showing the buried single domains. b) Topography scan. Inset: SEM tilted view of 2 isolated dots in the array prior to SiO_2 and Pd coating. c) line scans from a,b showing clear alignment between dot location and magnetic phase contrast signal.

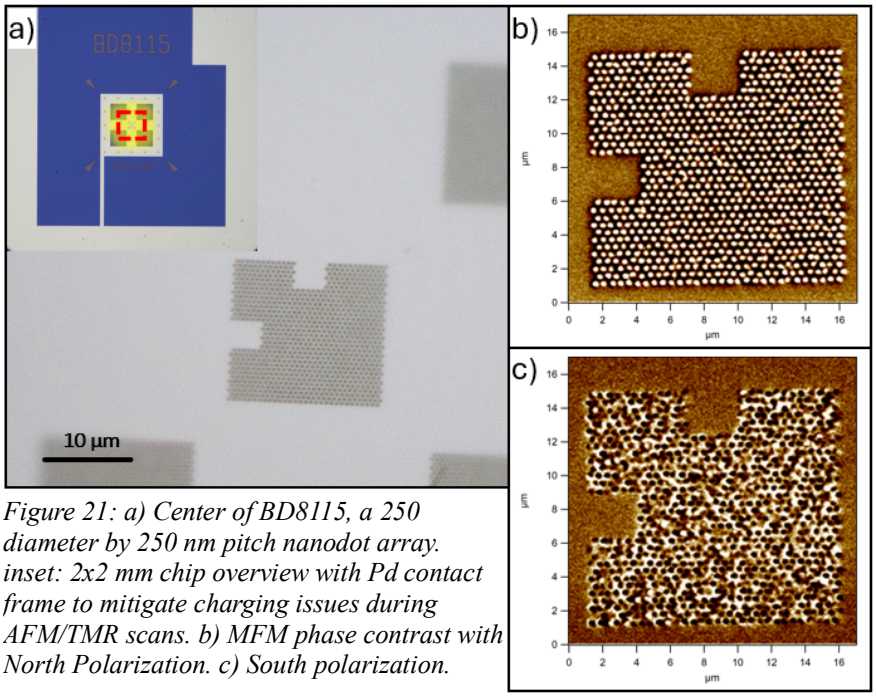


Figure 21: a) Center of BD8115, a 250 diameter by 250 nm pitch nanodot array. inset: 2x2 mm chip overview with Pd contact frame to mitigate charging issues during AFM/TMR scans. b) MFM phase contrast with North Polarization. c) South polarization.

1: NanoTI-Magnetic Archaeology

The electronics used for sending and receiving electrical signals to the device are the same used in the main Chips campaign. Further details on the electronics set-up and pulse sequence is detailed in the Chips section. On top of the existing retention requirements for NanoTI chips including electrical contacts through vessel and wire-bonding, magnetic archaeology requires an applied field using a permanent magnet. This applied field was implemented using a 4 mm diameter by 2 mm thick SmCo magnet with gold plating to minimize hydrogen embrittlement of the brittle compressed SmCo powder matrix. The location of this magnet within the vessel can be seen in Figure 22. The ceramic PCB used for holding the chip in the center of the vessel is the same component used in other chip gas experiments.

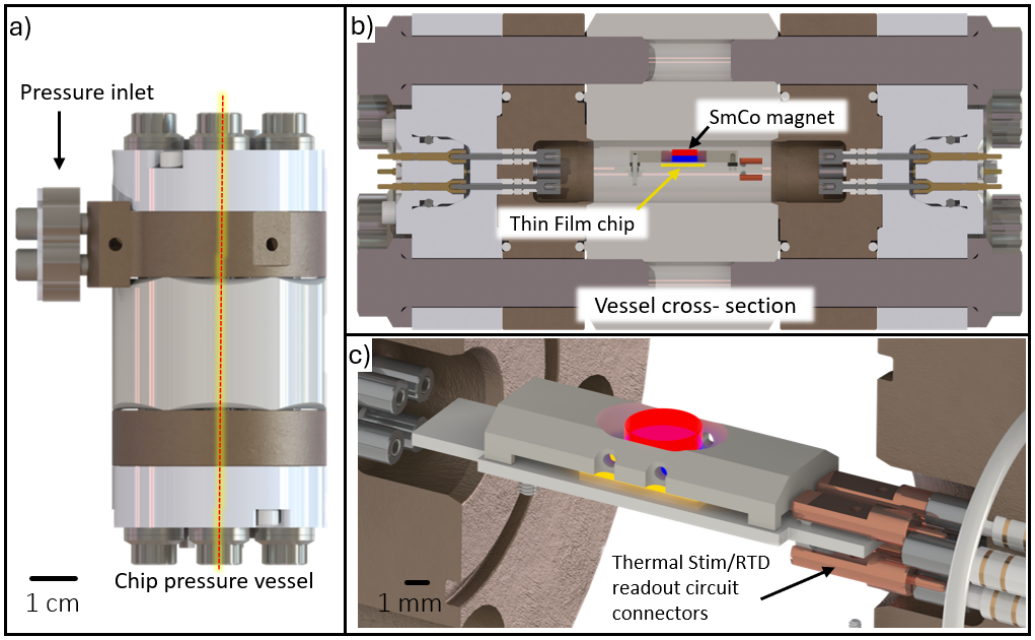


Figure 22: a) Pressure vessel overview. b) cross-section of vessel displaying locating of magnetic archaeology NanoTI PCB/chip. 2 sets of 4- pin ultrapressure feedthroughs for electrical access. C) Isometric view of chip and SmCo magnet retention with the magnet being held off the chip surface about 300 µm.

2: Staring – Magnetic Archaeology

Similar to NanoTI magnetic archaeology requirements, the sample had to be exposed to a low amplitude applied magnetic field so that when the nanodots were heated to a high enough temperature their polarity would be reoriented to the direction of the applied field prior to cooling back to room temperature. Correlating the surface temperature of the sample with a given laser profile and input power further required electrical contacts inside the pressure vessel to access the

on-chip 4-wire RTD. The retention displayed in Figure 23 fulfills these retention requirements with the SmCo permanent magnet positioned behind the sample, allowing for laser access to the sample area on the front. The strength of the magnetic field at the set distance between the FePt film and the permanent magnet along the vessel axis was measured externally using a Teslameter. The field

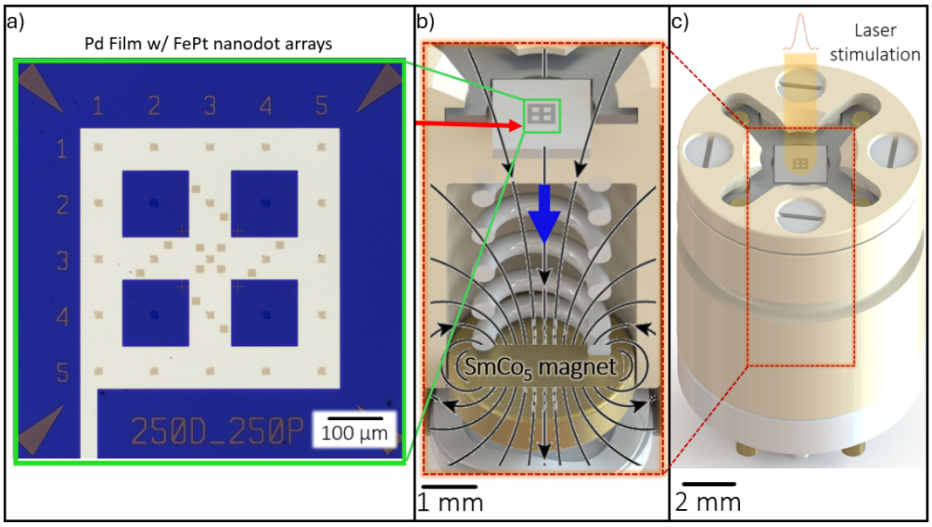


Figure 23: a) Pd window-like structure on top of >25 nanodot arrays each with 1000 dot sensors in each. b) Cross-section of sample retention for staring. The SmCo magnet is offset using a spring to reduce the field intensity at the surface of the sample to -100 mT. c) Staring sample retention with 1070 nm wavelength laser incoming from above.

amplitude was typically set to -100 mT. Further detail on the laser set-up for this and similar experiments can be found in the Staring campaign section.

3: HPE 400 ksi – Magnetic Archaeology

The 400-500 ksi pressure vessel added a tighter constraint for retaining the sample inside the vessel while keeping it properly oriented with optical access window for femtosecond laser stimulation experiments. The vessel diameter was reduced from 0.5” in 150ksi vessels to ≤ 0.25”. The solution for holding the sample and permanent magnet in place during the duration of the experiment, without damage to the sample during install or removal is detailed in USP and HPE sections of this report. In short, a set of countersunk Toughmet retaining rings hold the 2x2 mm chip loosely along the optical axis. A compression spring presses all surfaces of the Toughmet ring, sapphire window, and permanent magnet gently into one another so the retention remains in place prior to press-up to 400 ksi and beyond.

Results

Utilizing these submicron magnetic sensors as heat detectors produced experimental results that were reproducible between various campaigns leading the same general result of observing no indication of local thermal events within the Pd film. We can also have high confidence that the magnetic sensors deployed for this AHE search should be sensitive to a single 24 MeV event.

These results are covered in further detail in the subsections below moving from magnetic material stability testing for both HDD and FePt systems to full pressure stimulation experiments in both protium and deuterium for NanoTI, Staring, and 400 ksi HPE vessel campaigns.

Results : Ultrapressure gas degradation tests

HDD films, FePt films, and FePt nanodots were subjected to deuterium gas up to 400 ksi. All films were coated with 10 nm PECVD SiO_2 as a protective layer. The films were left under this deuterium and protium environment for 1 hour prior to removal and post characterization with MFM. No obvious changes to magnetic polarity could be measured using MFM, and magnetic hysteresis loops of films exposed to hydrogen showed no significant changes in coercivity and magnetization (Figure 24).

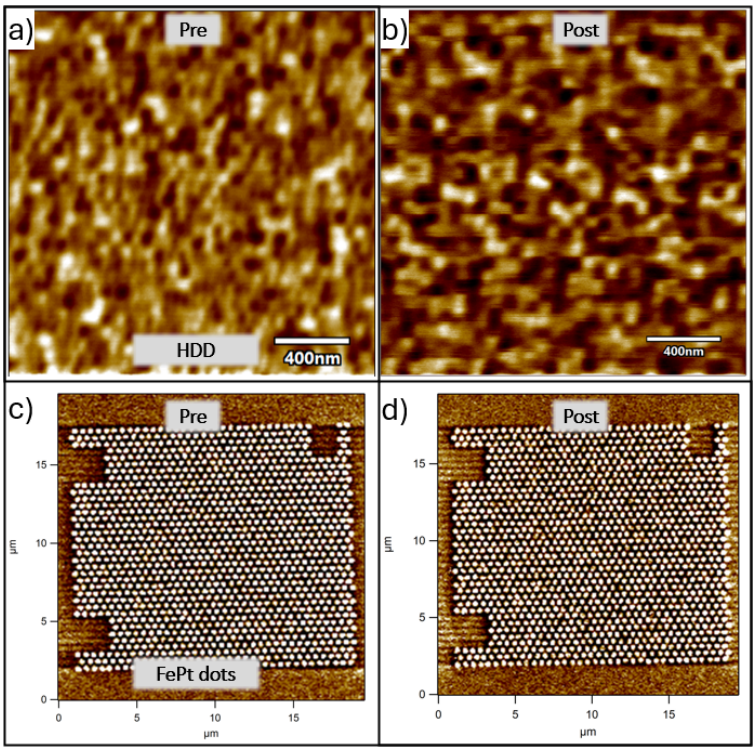


Figure 24: a, b) HDD, CoPtCr media before exposure to ultrapressure hydrogen and after showing the magnetic domain structure is retained. c, d) MFM scan of FePt nanodot array showing the magnet domain strucutre is retained after ultrapressure testing.

Results : Confirmation of ultrafast switching of domains

To confirm that both HDD and FePt type magnetic sensors could record nuclear events that may occur as fast as the picosecond timescale, a series of experiments were completed that used a femtosecond 450 nm laser to locally pump energy into the Staring and NanoTI type film stacks

within an area of 10-30 µm in diameter. The pulse width for all tests was on average 40 fs with a rep-rate of 1-5 Hz. The excitation fluence for each pulse was increased between 0.01- 10 µJ/cm^2, spanning the range from below domain reversal threshold to above film damage threshold. An external applied field with an amplitude of 150 mT was applied perpendicular to the sample surface to reverse domains that have become destabilized from local heating within the laser spot. For HDD devices the fluence level for domain reversal was roughly 0.60 µJ/cm^2, and FePt devices 1.1 µJ/cm^2 (Figure 25). These results confirm the magnetic materials used can record a local thermal stimulus, over fsec-psec timescales, in the magnetic domain polarity of the film. This is fundamental in HAMR and ultrafast optical switching; however, it was imperative to confirm this effect with the films being used as AHE detectors[86,87].

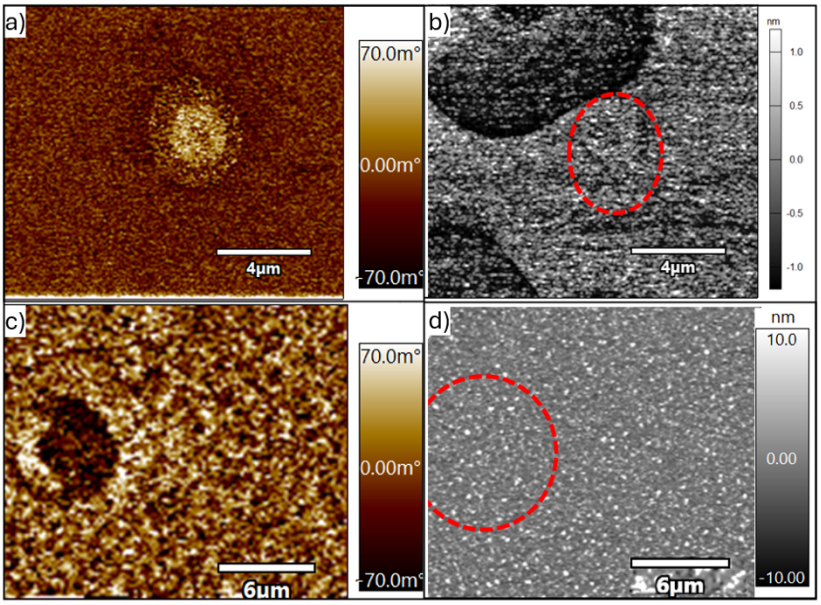


Figure 25: a) MFM phase contrast of CoPtCr magnetic HDD film showing clear domain structure change where the femtosecond laser impinged on the surface. b) Same scan area as a) showing topography, no physical damaged manifested within laser spot (red). c) MFM phase contrast of FePt film showing clear domain structure change (in presence of a South oriented H_{ext}) where the femtosecond laser impinged on the surface. d) No physical damage observed in the same location shown in c).

Results : Single Event Detection Simulation

Based off published values for FePt, hydrogen gas, and CMOS based chip architectures, a thermal simulation using COMSOL Multiphysics was performed to help quantify our AHE sensitivity

[86] Kelvin Yao, Felix Steinbach, Martin Borchert, Daniel Schick, Dieter Engel, Filippo Bencivenga, Riccardo Mincigrucci, Laura Foglia, Emanuele Pedersoli, Dario De Angelis, Matteo Pancaldi, Björn Wehinger, Flavio Capotondi, Claudio Masciovecchio, Stefan Eisebitt, and Clemens von Korff Schmising. *Nano Letters* **2022** *22* (11), 4452-4458.

[87] Xie Z, Cai Y, Tang M, Zhou J, Liu J, Peng J, Jiang T, Shi Z, Chen Z. Fluence and Temperature Dependences of Laser-Induced Ultrafast Demagnetization and Recovery Dynamics in $L1_0$-FePt Thin Film. *Materials*. 2023; 16(14):5086.

using this magnetic detection method[88,89]. Accounting for convection and conduction coefficients and using a 24 MeV heat source as the thermal input to the system, we defined the radial sensitivity of a single 250 nm FePt dot from its center. To reduce model complexity this simulation only models heat transfer in 3D with thin film approximation. The change in magnetic properties based off the thermal input was not modeled, and instead a threshold temperature of 300 °C was defined, above which it is assumed the nanodot single domain would destabilize and flip in the applied field direction. Other assumptions made for this model are 1) a 20 ps pulse width over which the 24 MeV event would be released and 2) the event sphere has a radius of 2 nm over which the energy will be emitted from. Figure 26 shows the 3D representation of this film stack and the change in

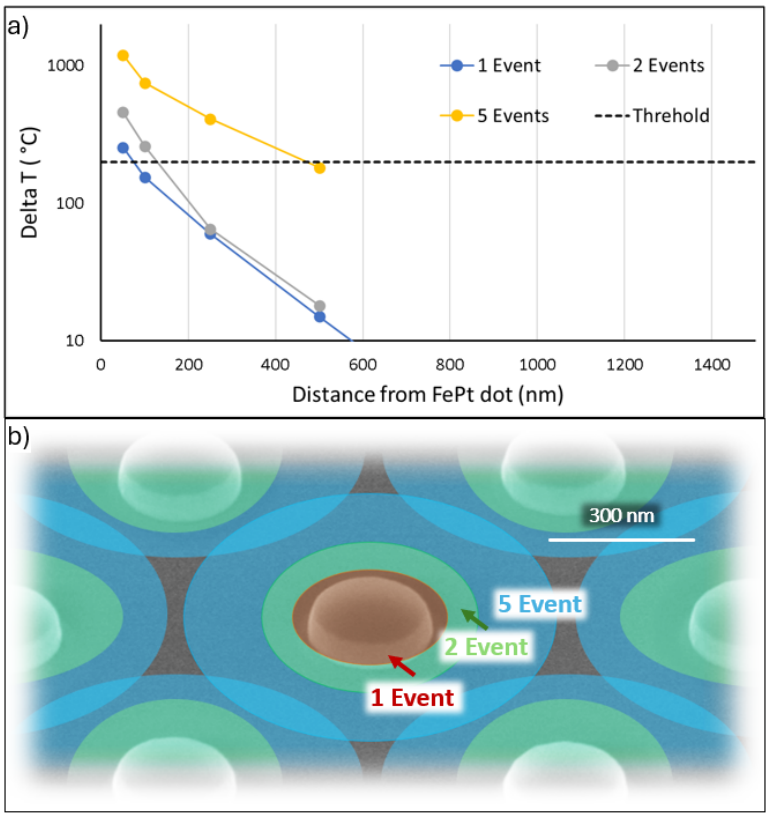


Figure 26: a) Temperature change vs. lateral distance from dot center as a function of 24 MeV AHE events. b) Tilted SEM view of FePt nanodot array with overlayed sensitivity color map.

temperature as a function of radial distance from dot center and number of events occurring within 20 ps. Through this basic heat transfer model, we can have a level of confidence that we have single event sensitivity within 100-200 nm from the center of a 250 nm nanodot. Within the 250 nm nanodot array with 500 nm pitch from center to center, 1-5 AHE events should be detectable using this magnetic recording approach.

[88] S. E. Russek and W. E. Bailey, "Magnetic domain structure and imaging of Co-Pt multilayer thin-film nanostructures," in *IEEE Transactions on Magnetics*, vol. 36, no. 5, pp. 2990-2992, Sept 2000.
[89] H. Z. Yang, Y. J. Chen, S. H. Leong, C. W. An, K. D. Ye, M. J. Yin, J. F. Hu; A multi-functional testing instrument for heat assisted magnetic recording media. *J. Appl. Phys.* 7 May 2014.

Results : Magnetic Archaeology – NanoTI FePt dots (300 nm)

The thermal stimulation applied for these experiments was the same used in all other NanoTI chip campaigns. Further detail on experimental design of the electrical circuit can be found in the Chips section. With a stimulation pulse width of 10 μs and varying the voltage between 1-6 V the heater circuit was able to heat up to 350 °C without thermally degrading the circuit (Figure 27). Running 3-4 samples in deuterium at 145 ksi we did not observe domains reversal from the preconditioned state of south when heating below 325 °C. The MFM contrast clearly indicates there is no individual sensor being flipped to the applied field direction of north (Figure 28). With these stimulation conditions no observable AHE events could be confirmed.

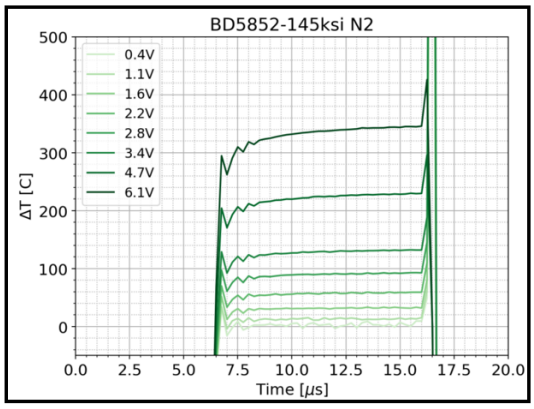


Figure 27:Resulting temperature increase of the Pt heater circuit as a function of input voltage with constant current. We can stimulate above and below any Curie Temperature the films may possess.

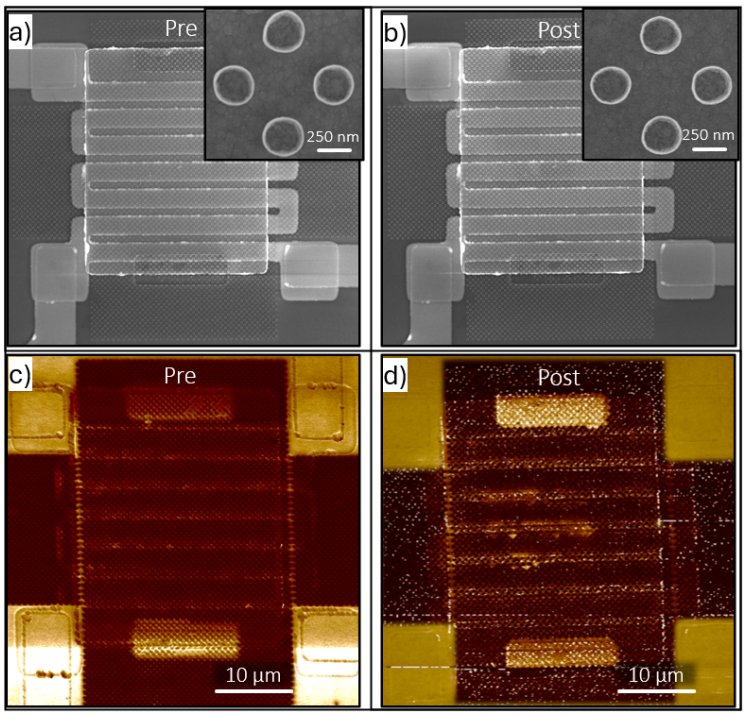


Figure 28: a) Pre SEM image of NanoTI with FePt nanodot sensors on top. b) post SEM showing no obvious signs of damage. c,d) Pre/Post MFM showing no signs of domain reversal from south to North.

These chips proved to be difficult to reproduce between wafer batches due to variations in the PECVD SiO_2 layer that supported the FePt dots. Furthermore, with the Pt heater underneath the FePt, this Pt heater needed to be processed during fabrication through rapid temperature changes up to 750 °C, which unfortunately can change the resistance of the Pt heater circuit leading to variation in our thermal calibrations as well. This variation leads to large uncertainties (50-100 °C) in how hot we are heating our Pd and FePt sensors. With this in mind, the remainder of magnetic archaeology stimulation was focused on laser-based heating using the 1070 nm and femtosecond 450 nm wavelengths.

1: HDD/CoPtCr Sensors

To not damage our CoPtCr sensors, temperature calibrations were completed using an on chip RTD to measure the peak average surface temperature of a 750 µm diameter 1070 nm laser spot (Figure 29a). This calibration located the damage threshold and switching threshold powers for the CoPtCr film. The damage threshold measured to be >18 W, while the domain switching threshold, regardless of the applied field, was between 16-17 W (Figure 29b,c). AHE experiments were carried out in both 145 ksi protium and deuterium with input power set to 16 W. Neither condition produced reliable domain reversal in the stimulated laser spot region. Ultimately the fact these HDD films were continuous in nature with little space between domains, made it exceedingly difficult to interpret the MFM contrast scans between pre and post. Neighboring domains would

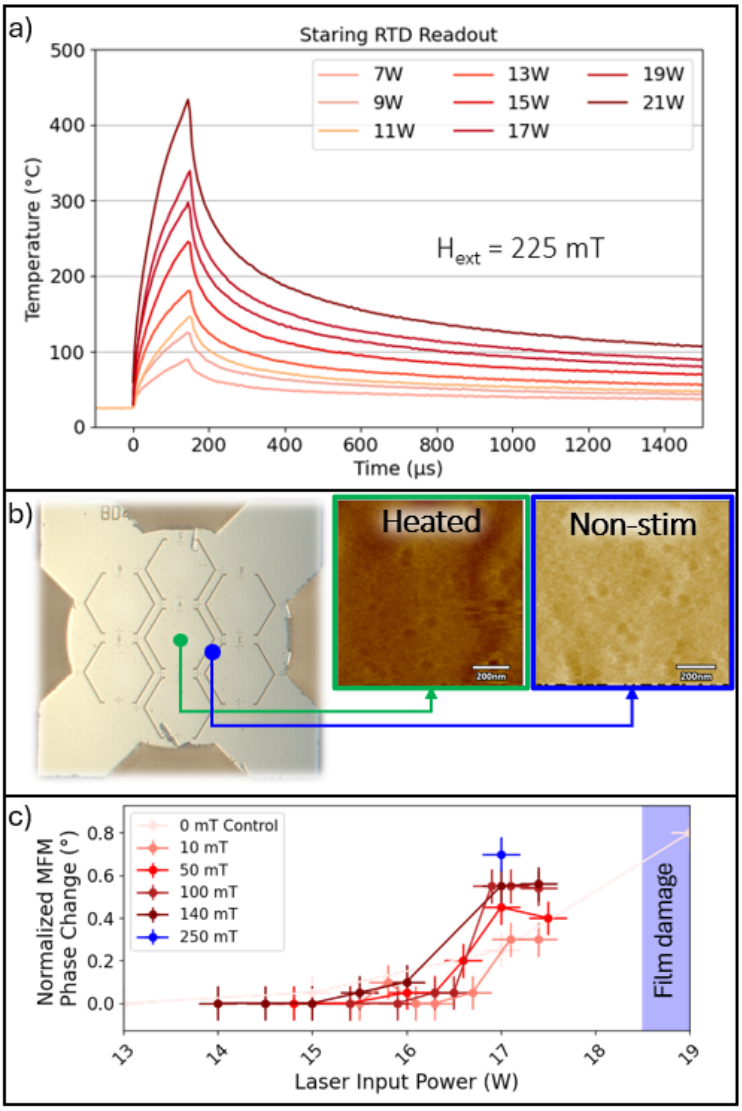


Figure 29: a) RTD calibration temperature profiles for each laser input condition. b) Method used to measure changes in MFM phase contrast, subtracting out background (blue) from the stimulated region (green). c) Phase contrast normalized to a complete domain reversal (=1). There is a slight trend of decreased power required for increasing Hext, but not significant enough to be controllable under these extreme environments.

flip others stochastically. For this reason, we moved quicky to FePt based sensors where we obtained much better stability and reproducibility.

2: FePt dots (250 nm)

The FePt nanodot with Pd films were exposed to 150 ksi deuterium and protium gases while being stimulated with a 1070 nm wavelength laser at various input power conditions and pulse lengths. The beam diameter was 1 mm, and a calibration RTD chip with the same film stack as our actual magnetic sensor samples was used to calculate how hot the surface of the Pd was during laser heating. Knowing the peak temperature of the film is critical because there is a maximum service temperature that FePt films can operate at without completely destabilizing and that is the Curie Temperature (T_C). This calibration ensures the laser parameters on our Pd/FePt samples will keep the peak temperature well below T_C to ensure we do not get false positives from laser stimulation alone. The calibration results can be found in Figure 30 where we heated the calibration chip up to 600 °C. For all Staring tests we aimed to keep the film temperature at or below 300°C so we remain well below the average T_C of 450°C and variability around this temperature up to 50 °C in bulk films based off possible local variability in phase structure, composition, finite size effects[90,91,92].

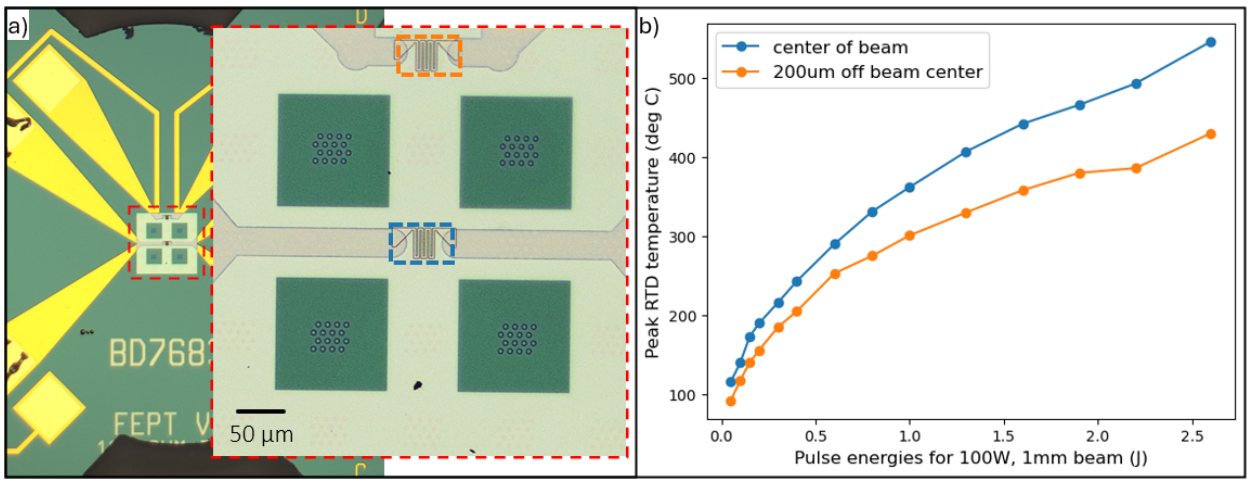


Figure 30: a) Calibration RTD chip with 2 4-wire RTD circuits with 1 mm diameter beam, inset: center of chip with serpentine Pt RTDs. b) Peak temperature values for each pulse of 100 W at varying pulse lengths between 1-70 ms.

A total of 10 individual Pd/FePt samples were loaded and thermally stimulated in the Staring apparatus, 5 deuterium and 5 protium. All samples that were heated up to a peak temperature of 300 °C or below showed no indication of AHE. Clean PdH_x and PdD_x samples showed no clear

[90] L Rohman *et al* 2020 *J. Phys.: Conf. Ser.* **1491** 01 2013.
[91] Nguyen Thanh Binh, Sergiu Ruta, Ondrej Hovorka, Richard F. L. Evans, and Roy W. Chantrell, Phys. Rev. B 106, 054421 – Published 17 August 2022
[92] Bublat, T. & Goll, D.. (2010). Temperature dependence of the magnetic properties of L10-FePt nanostructures and films. Journal of Applied Physics. 108.

changes in magnetic polarity when scanned using MFM after stimulation. All domains remained oriented in the northward direction out-of-plane (Figure 31). With this result there were no indications that the Pd film heated up locally due to AHE activity.

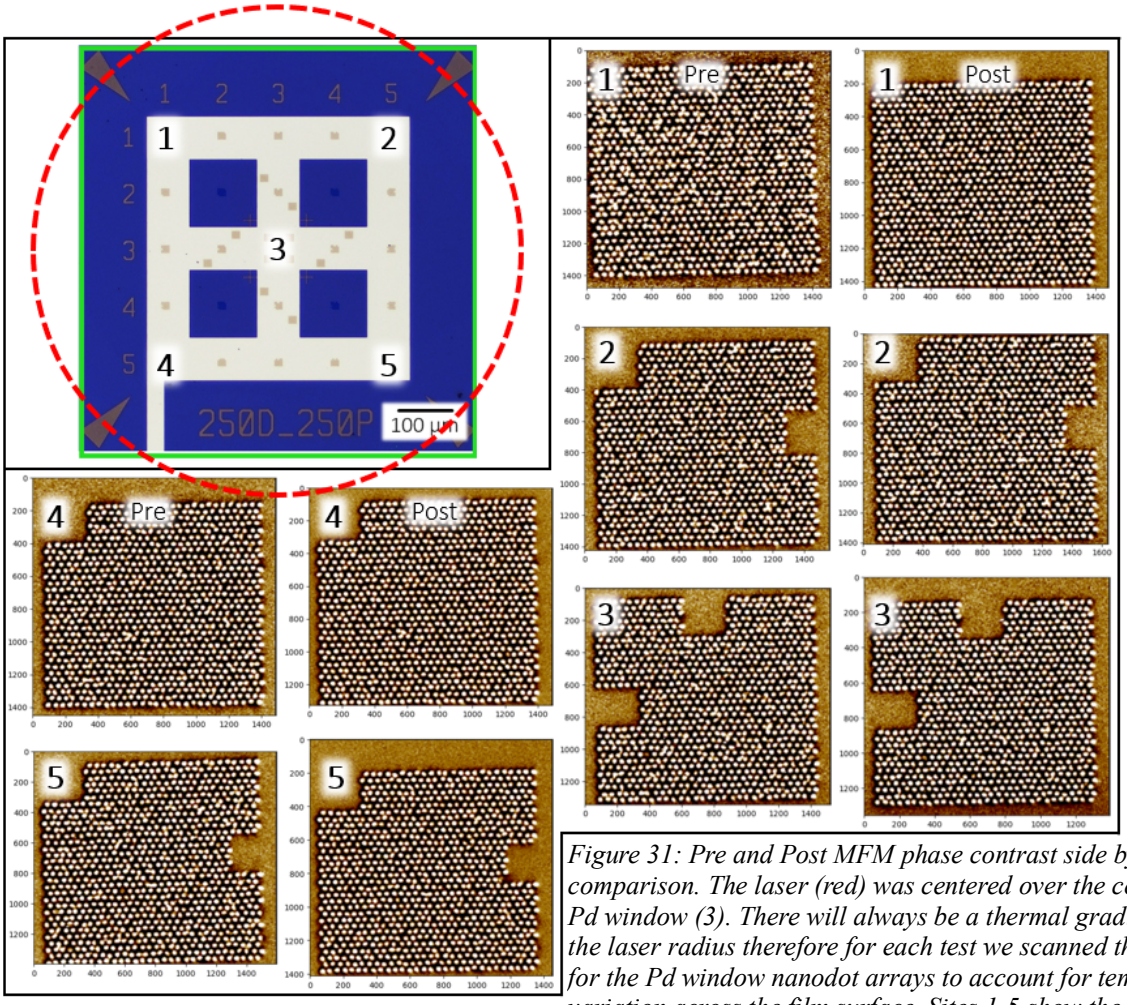


Figure 31: Pre and Post MFM phase contrast side by side comparison. The laser (red) was centered over the center of the Pd window (3). There will always be a thermal gradient along the laser radius therefore for each test we scanned the 4 corners for the Pd window nanodot arrays to account for temperature variation across the film surface. Sites 1-5 show the raw MFM phase scans where all domains remain polarized in the preconditioned orientation, North.

There were sources of contamination that could cause a false positive in our MFM readout. This contamination mainly formed through initial poor sample handling practice and powerpack inlet line particulate adhering to the surface of the sample. If a carbon-based contamination or particle with increased roughness is locally covering some of the FePt nanodots, there is increased absorbance in those locations leading to higher peak temperatures locally within a nanodot array. These false positives are caught by collecting overview SEM images of each nanodot array after a pressure test is completed (Figure 32). This contamination effect occurred in the femtosecond laser experiments as well and was accounted for in final conclusions for this study.

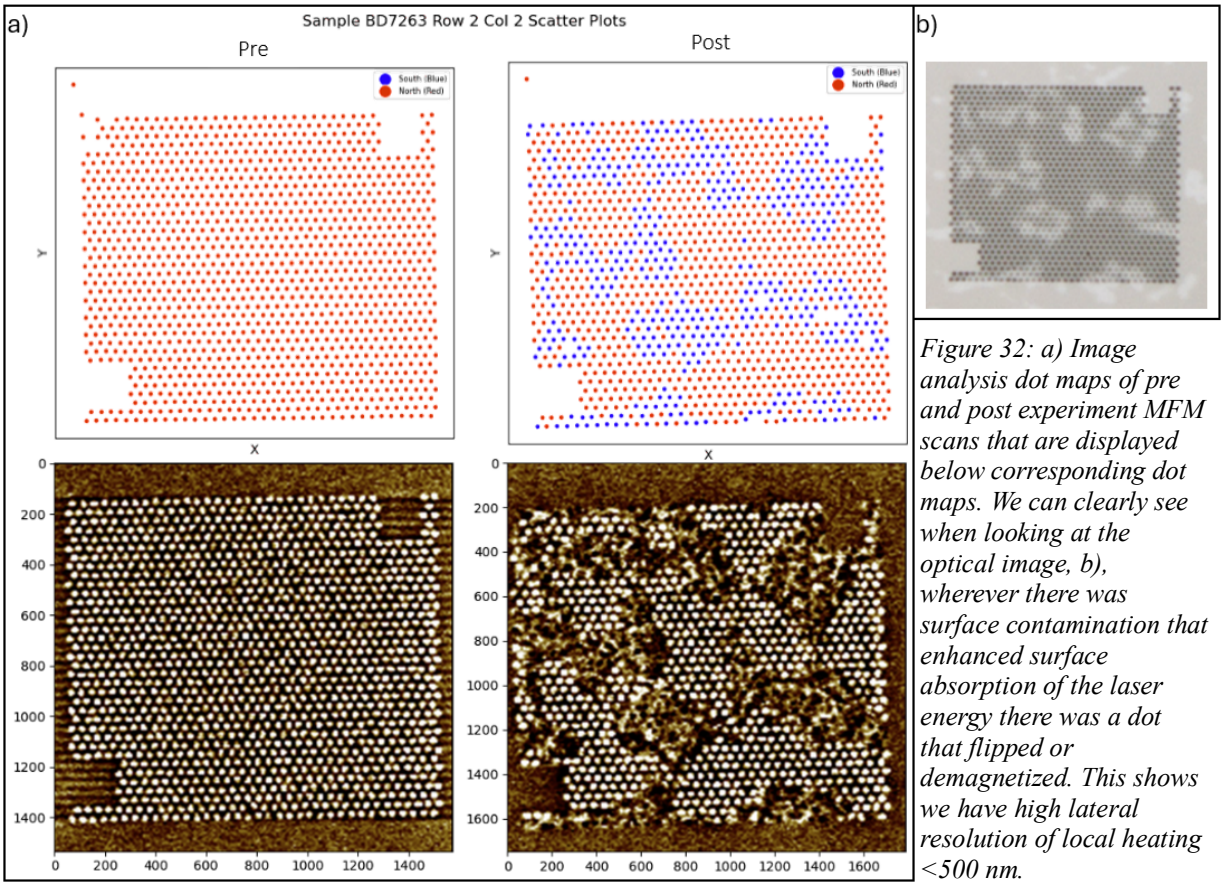


Figure 32: a) Image analysis dot maps of pre and post experiment MFM scans that are displayed below corresponding dot maps. We can clearly see when looking at the optical image, b), wherever there was surface contamination that enhanced surface absorption of the laser energy there was a dot that flipped or demagnetized. This shows we have high lateral resolution of local heating <500 nm.

Results : Magnetic Archaeology : 400 ksi, femtosecond laser

Exposing the Pd/FePt nanodot array chips to 400 ksi ensured maximum deuteron and proton loading prior to stimulation using a 450 nm femtosecond laser. For these experiments laser parameters were selected based off damage threshold and domain reversal threshold calibration studies where the beam diameter was focused to 150 μm and pulse width <100 fs. If we saw no domains flip in either protium or deuterium at a given input energy level, the temperature reached locally around each nanodot is known to be below the domain reversal temperature (T_S). Directly measuring the film temperature in a 400 ksi gas environment and at these ultrafast timescales was not feasible at the time of this study. A total of 8 thermal stimulation tests were completed with laser parameters that stimulated the sample below T_S, 6 in deuterium and 2 in protium as controls. Figure 33 shows multiple nanodot arrays at various locations around the center of the sample. All five locations show no clear changes in domain polarity from north indicating no hot spots were detected in those locations before, during, or after laser stimulation. Table 4 lists the entirety of all AHE experiments the various magnetic sensors were utilized for heat detection.

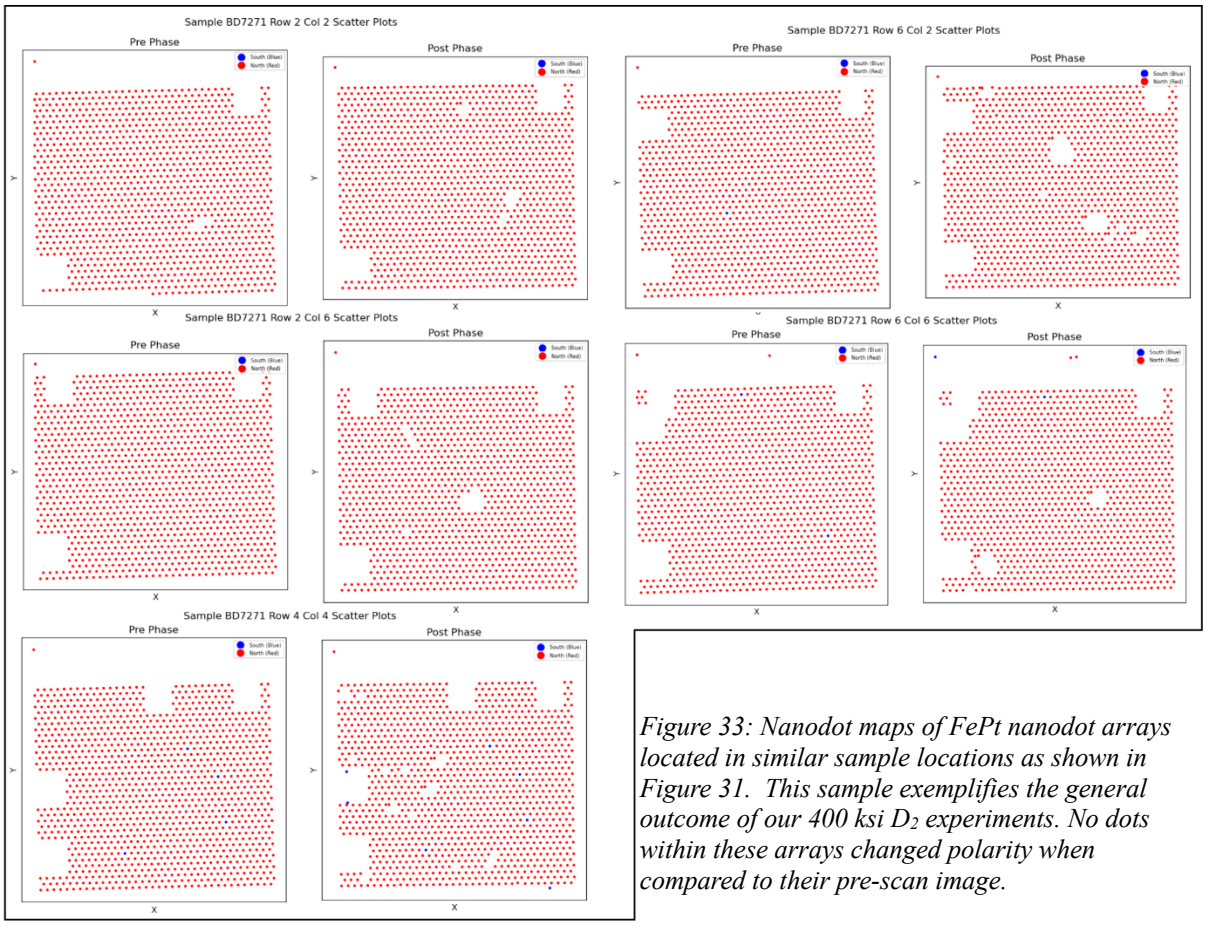


Figure 33: Nanodot maps of FePt nanodot arrays located in similar sample locations as shown in Figure 31. This sample exemplifies the general outcome of our 400 ksi D_2 experiments. No dots within these arrays changed polarity when compared to their pre-scan image.

*Table 4: AHE Experiments using Magnetic Sensor**

NanoTI				
BE	Sample type	Gas	Pressure (ksi)	Peak Temperature (°C)
2351	HDD/CoPtCr	N2	120	350
2442	HDD/CoPtCr	D2	140	450
2518	HDD/CoPtCr	H2	140	450
2621	HDD/CoPtCr	H2	140	450
3297	HDD/CoPtCr	D2	100	300
3318	FePt 300 nm dots	D2	100	300
Staring				
BE	Sample type	Gas	Pressure (ksi)	Peak Power (W) / Temperature
1845	HDD/CoPtCr	D2	145	19.2
1936	HDD/CoPtCr	D2	145	19.2
2059	HDD/CoPtCr	D2	145	23
2290	HDD/CoPtCr	D2	145	17
2519	HDD/CoPtCr	H2	120	17

2532	HDD/CoPtCr	H2	145	17
2710	HDD/CoPtCr	D2	140	15
2722	HDD/CoPtCr	D2	145	15
2732	HDD/CoPtCr	D2	145	15
2739	HDD/CoPtCr	D2	140	15
3611	FePt 300 nm dots	D2	150	11
3715	FePt 300 nm dots	D2	140	13/300 °C
3738	FePt 300 nm dots	D2	140	12/250 °C
3778	FePt 300 nm dots	D2	140	12/250 °C
3824	FePt 300 nm dots	D2	145	No Stim
3990	FePt 300 nm dots	D2	145	300 °C
4013	FePt 300 nm dots	D2	145	250 °C
4045	FePt 250 nm dots	D2	140	250 °C
4058	FePt 250 nm dots	H2	145	250 °C
4063	FePt 250 nm dots	H2	145	250 °C
4093	FePt 250 nm dots	D2	145	250 °C
4097	FePt 250 nm dots	D2	145	200 °C
4101	FePt 250 nm dots	H2	145	200 °C
4121	FePt 250 nm dots	D2	145	250 °C
4123	FePt 250 nm dots	H2	140	250 °C
4144	FePt 250 nm dots	H2	140	50 °C
4189	FePt 250 nm dots	D2	140	50 °C
4190	FePt 250 nm dots	D2	140	250 °C
400 ksi femtosecond laser				
BE	Sample type	Gas	Pressure (ksi)	Peak Power (W)**
3899	FePt 250 nm dots	D2	400	No Stim
3906	FePt 250 nm dots	H2	400	75 %
3941	FePt 250 nm dots	H2	400	60 %
3964	FePt 250 nm dots	H2	400	30 %
3978	FePt 250 nm dots	H2	400	15 %
4022	FePt 250 nm dots	H2	400	7.5 %
4057	FePt 250 nm dots	D2	400	7.5 %

*The BEs in this table are only experiments that were clean AHE tests, these are not including calibrations, optimization experiments, pressure test, offline stim testing, and most inert gas tests.
**Peak power value is based of % of full power of femtosecond laser upstream from pressure vessel, refer to USP for values and further details

Conclusion

Reflecting back, LENR is believed to produce local hot spots that form and dissipate quickly, only affecting a small volume (~ 100 nm in size) for a very short time (~ 1 psec). The extreme difficulty of detecting temperature pulses with both this temporal and this spatial resolution was the reason for developing the various magnetic heat sensors discussed in this section. Throughout our suite of experiments, we have confirmed that the sensors fabricated can respond to ultrafast energy

pulses at the local <1 μm length scales in proximity to a PdD_x lattice. With the single event sensitivity achieved for Pd thin film-based samples, we have not observed AHE production at the various pressure, temperature, and loading conditions probed in this study.

Seismology

Executive Summary

In the Seismology approach, we are looking for the mechanical vibrations that we expect to be produced by anomalous heat events (AHE), rather than *directly* measuring temperature.

To date, we have not measured any isotopic difference from Seismology experiments, which we interpret as a negative result for AHE with 1 part in 1 trillion fractional site sensitivity. These experiments were unstimulated, and AHE measurements were not performed in our highest noise-performing experimental configuration before the project wound down.

Further progress was planned to expand into larger bulk samples including palladium slugs, to explore different stimulation methods, to monitor for much longer durations, and to further develop more nuanced machine learning techniques to interpret the data. However, many of these avenues were explored thoroughly across other projects in our global effort.

Concept

In the Seismology approach, we are looking for the mechanical vibrations that we expect to be produced by anomalous heat events (AHE), rather than *directly* measuring temperature. If an AHE event occurs anywhere in the sample, the material nearby will be rapidly heated. In order to measure temperature directly, we would need to choose a specific location on the sample to place temperature probes, and we would only be able to detect AHE occurring very close to the location of the probes. However, this rapid heating will also produce a *vibration* that will propagate throughout the *whole* material, as illustrated in Figure 1. This is analogous to vibrations from an earthquake (or underground nuclear test) propagating out over long distances from its local epicenter.

We are particularly interested in techniques that are sensitive over large regions of the material, which can detect small vibrations from rapidly induced changes in temperature. We can measure these small vibrations using piezoelectric detectors that generate electrical charge in response to mechanical stress. These detectors have a sensitive area on the order of millimeters, which is many orders of magnitude larger than the temperature sensors we use in other projects, where the detectable area is usually better measured in nanometers or microns. AHE occurring anywhere in a palladium film deposited on the surface of the sensor will cause a vibration through the whole material.

Furthermore, these sensors have a long-lasting response to a fast thermal impulse, so we can be confident that the effect of an anomalous heat event will last long enough to be measured. This can be demonstrated by studying how the sensor responds to a well-controlled thermal impulse from a nanosecond-pulsed laser. Figure 1 shows that the measurement's response to a 6 ns laser

pulse lasts over 100 µs – over 4 orders of magnitude longer than the pulse. We expect the thermal impulse delivered by an AHE event to be much faster than 6 ns, but both can be effectively considered infinitely sharp impulses compared to the microsecond time scale relevant for vibrations, so we can argue that the vibrational response to an AHE event (or a short burst of AHE events) should last about as long as a nanosecond laser pulse of the same energy. We can also use these laser pulses to estimate the minimum detectable energy, which sets a limit on how many AHE events would need to occur within a burst in order to be detected. Crucially, these events would not need to occur at precisely the same moment as long they occur within a short enough window to be considered an impulse – potentially up to a few microseconds.

Our measurements are closely related to a widely used, non-destructive material analysis technique known as acoustic emission testing. In this technique, a material is put under a mechanical load, causing sudden, impulsive vibrations as the material begins to yield at an atomic scale. These impulses are referred to as 'acoustic emission hits', and they are detected and counted as an indicator that macroscopic yield is approaching. Acoustic emission testing can routinely detect elastic waves with very small kinetic energies, down to femtojoules, so detecting picojoule level thermal impulses due to AHE is within reason.

In our experiments, we monitored for acoustic emission hits while palladium was loaded in ultrapressure deuterium gas. We looked for presence versus absence of acoustic emission hits during corresponding parts of paired experiments using deuterium and protium. The presence of unexplained acoustic emission hits while a sample is highly loaded with deuterium, but not with protium, would be an interesting result, especially if the hits looked like sharp thermal impulses.

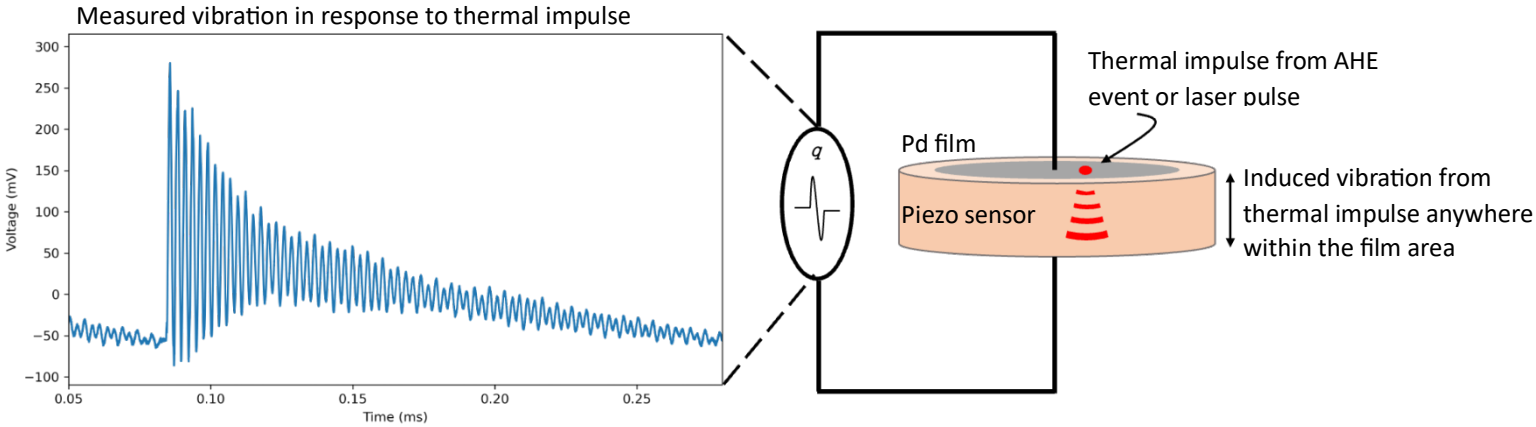


Figure 1: Seismology Concept

A fast thermal impulse occurring anywhere in a palladium film deposited on the surface of a piezoelectric sensor generates a vibration that can be measured on longer time-scales. The long-lasting response of this technique to AHE events occurring over a large area is a unique advantage to this approach.

Development of a Sample Based on a Piezoelectric Sensor:

Piezoelectric materials exhibit an electric dipole moment that changes as the lattice distorts. Practically, this means that a mechanical vibration results in an oscillating surface charge. When opposing surfaces of the material are connected through a sensitive amplifier capable of resolving the small charges that develop, piezoelectrics present an excellent method to measure small vibrations. In the field of non-destructive acoustic emission testing, sensors made from lead zirconate titanate (PZT) have become the industry standard due to their comparatively large dipole response to a given vibration. Depending on their precise geometry, these sensors can operate in a mechanical-electrical resonance, where the amplitude is a strong function of frequency, or in a broad-band mode where the response is similar across a wide frequency band.

In order to detect the vibration due to a thermal impulse, we selected a broad-band ultrasonic sensor manufactured by TDK (Part Number: B59050Z0206A031), since short impulses are inherently broad-band. Crucially, these sensors are available unpackaged and therefore can be easily integrated into a carrier assembly that is compatible with an ultrapressure hydrogen environment. It is worth noting that metallized PZT exhibits corrosion in hydrogen environments at elevated temperatures or long time scales, but this effect is minimal at the times and temperatures used in the discussed experiments[93]. We tested these sensors integrated with two different form factors of Palladium. Most reliably, we deposited a 100 nm thin film of Palladium directly onto the sensor electrode. We used a shadow mask to confine the film to the central 4 mm diameter of the sensor disk, and we deposited 5 nm Titanium adhesion layer before the 100 nm of Palladium, using 120W DC magnetron sputtering in a Kurt J. Lesker PVD 75 system.

Samples for our first proof-of-principle electrical measurements were very simple, consisting of wires soldered directly onto both sides of a PZT sensor inside of a small RF-shielded box (Pomona 3230) with BNC connections, in ambient atmosphere. These measurements demonstrated very low noise and set the benchmark performance level for our development efforts in an ultrapressure experimental environment. After these proof-of-principle measurements, we began implementing sensors into ultrapressure vessels, using the assembly shown in Figure 2.

[93] Mohammadabadi, A.S. Hydrogen-Induced Damage of Lead-Zirconate-Titanate (PZT), University of British Columbia, 2013. https://doi.org/10.14288/1.0073678.

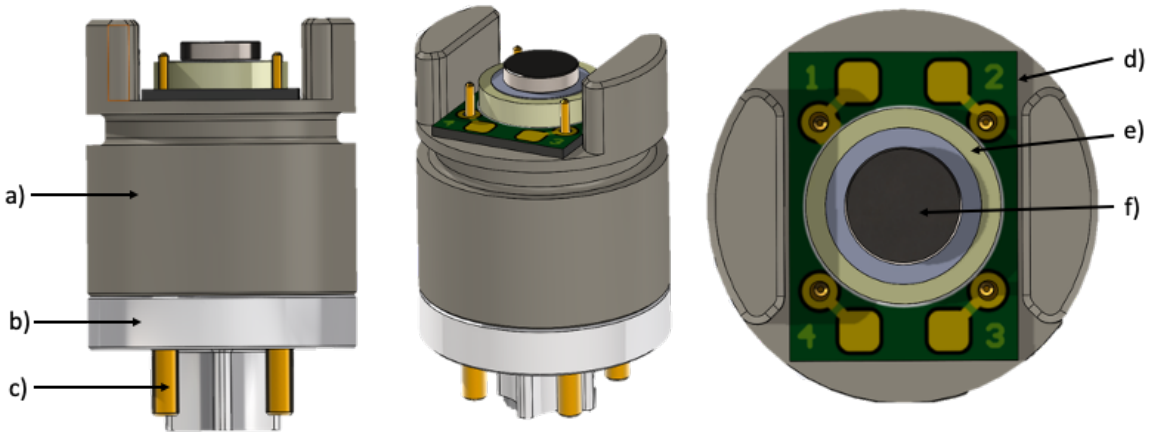


Figure 2: Seismology Sample Capsule Detail

In order to integrate piezoelectric samples into an ultrapressure environment while maintaining optical access to the surface, the sample must be held fast in the center of the cylindrical bore, and electrical contact must be made to both the top and bottom surface in a manner compatible with ultrapressure feedthrough technology. Here, the basic sample capsule is shown in detail, from which other shielded and lower-capacitance iterations were later developed. The first fully ultrapressure-compatible iteration consisted of palladium bulk (shown) or thin film adhered to the PZT sensor, which is mounted onto a ceramic carrier PCB using silver conductive epoxy. Contact to the top sensor electrode was made using a gold wire bond. Gold-coated copper pins were press fit into a PEEK capsule, secured with a fastener cap, and soldered into the carrier board using lead-tin solder. These pins fit into mating connector pieces which were brazed into sapphire to form an ultrapressure pin feedthrough.

a) PEEK capsule
b) Fastener cap
c) Gold-coasted copper pins
d) Ceramic (alumina) carrier PCB
e) PZT piezoelectric ultrasonic sensor disk (TDK Part Number: B59050Z0206A031). 1.02 mm thick by 5.00 mm diameter.
f) Palladium slug adhered to PZT sensor using Stycast 2850 epoxy. Alternatively, a 100 nm Pd thin film was deposited using DC magnetron sputtering at 120 W (Kurt J. Lesker PVD 75).

We quickly observed that noise coupling in a real experiment was far greater than in the well-isolated, benchtop proof-of-concept. In general, noise couples into an electronic measurement either capacitively or inductively.

To mitigate capacitive noise, the capacitance between the measurement leads and the noise source must be minimized. We identified that noise developing on the ultrapressure vessel itself was the primary culprit, so the capacitance between the leads and the vessel that contained the sample assembly needed to be minimized. We found that the capacitance was dominated by the ultrapressure feedthroughs. The capacitance between a brazed feedthrough pin and the vessel was measured to be around 7 pF. When epoxy feedthroughs were used instead, the capacitance was reduced to 3 pF. These epoxy feedthroughs consist of continuous, ceramic coated platinum wires

fed through a narrow bore in a tungsten carbide plug and epoxied in using Stycast 2850 epoxy. The lower surface area of the wires compared to the brazed pins makes up for their closer proximity to the noise source and results in an overall reduction in capacitance. However, the performance could be improved by an order of magnitude if the feedthrough plug could be made from an insulator instead of conductive tungsten carbide, thereby pushing the noise source to a further distance from the confined wires. When sapphire was used instead of tungsten carbide, the system's overall capacitance was no longer dominated by the feedthrough. However, since sapphire is more susceptible to crack propagation, the modified Bridgman seal normally used with the tungsten carbide plugs needed to be replaced with a gold gasket seal, which required re-engineering of the ultrapressure vessel. The re-engineered vessel was successfully built and tested just before winding the project down, so we did not perform any AHE tests with the sapphire epoxy feedthrough plugs.

To minimize inductive noise, we employed several strategies together. First, we minimized the loop area formed by the measurement leads by using triaxial cables outside of the vessel. This involved designing custom breakout boards to wire the feedthrough wires into a triaxial connector mounted directly onto the vessel, as well as boards that broke the triaxial cable out into coaxial BNC connectors at the measurement amplifiers.

Second, we used radiative shielding at multiple scales. Materials with high magnetic permeability perform significantly better at shielding against electromagnetic radiation, especially at low frequency, so we chose shield materials with as high permeability as possible.

We enclosed the entire ultrapressure vessel, as well as the first pre-amplification stage, inside a Faraday cage to prevent radiative pickup from the external environment. Our first design for this Faraday cage used one-quarter-inch thick 410 stainless steel. However, sizeable holes in the enclosure were necessary to maintain free-space optical access and to feed through the beryllium copper ultrapressure plumbing.

To address these non-idealities, we modified an aluminum JRE 1522 RF Shielded Test Enclosure to add plates of high-permeability μ-metal (Magnetic Shield Corp) to all internal surfaces. The μ-metal plates were welded along their joined edges prior to bolting into the enclosure. We designed a custom feedthrough plate to pass the ultrapressure plumbing into the enclosure while maintaining the integrity of the shielding, and we delivered the nanosecond laser pulse using a fiber optic feedthrough rather than free space. This ultimately improved our external shielding substantially, but the project wound down before we could repeat AHE measurements with this improved noise floor.

In order to prevent radiative pickup in the internal measurement leads, we engineered small, cylindrical Faraday cage shields that were inserted along with the sample capsule into the inner bore of the pressure vessel. We machined these shields out of high-permeability Permalloy rods. Our highest performing sample assembly is shown in Figure 3, which integrates a sample capsule and a permalloy shield with an epoxy feedthrough.

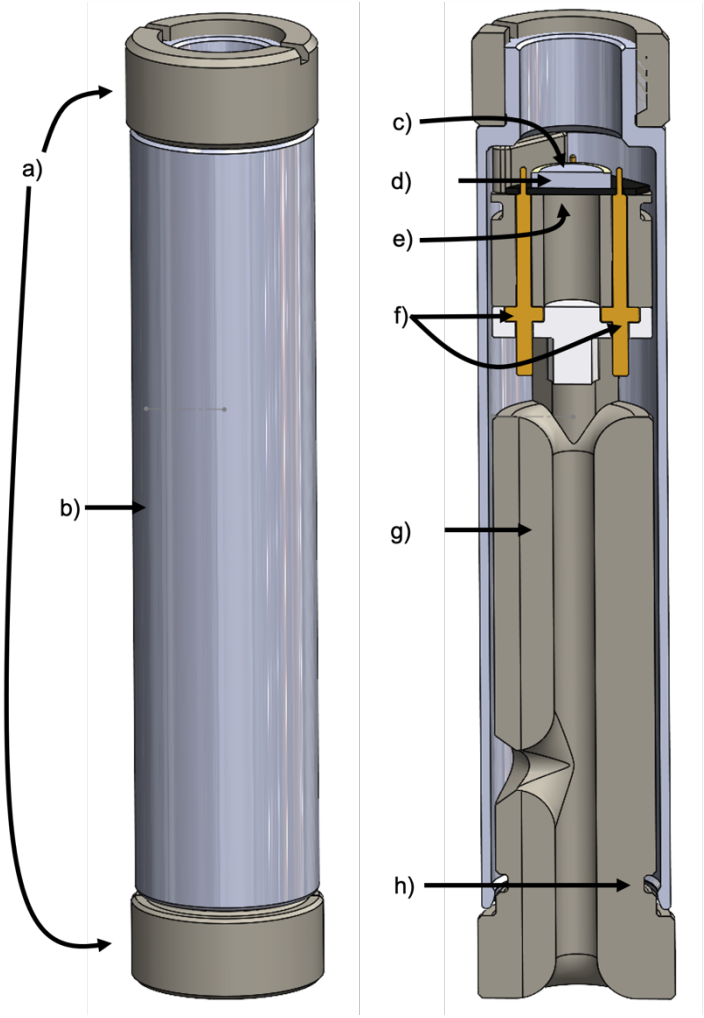


Figure 3: Seismology Sample Insert

The sample assemblies used in seismology experiments consist of a cylindrical, shielded capsule assembly that is inserted into the bore of an ultrapressure vessel. The area above the piezoelectric sample is clear for optical access through a sapphire window. This allows calibrated thermal impulses to be delivered to the sample surface using a nanosecond-pulsed infrared laser. Wires to measure the sensor are fed through an epoxy feedthrough in a tungsten carbide or sapphire piece to ambient pressure.

a) PEEK caps to prevent shorting of shield to vessel bore
b) Permalloy high-μ shield
c) 100 nm Palladium thin film. Shadow-mask patterned during DC magnetron sputtering at 120W (Kurt J. Lesker PVD75)
d) PZT piezoelectric ultrasonic sensor disk (TDK Part Number: B59050Z0206A031). 1.02 mm thick by 5.00 mm diameter.
e) Ceramic carrier PCB. Epo-Tek H20E silver conductive epoxy was used to adhere piezoelectric sensor to center contact on the PCB.
f) Gold-plated copper pins, soldered into the carrier PCB. Pins were soldered into the carrier PCB using lead-tin solder. Gold wire bond (not shown) connected the carrier PCB to the edge of the top contact on the piezoelectric disk.
g) PEEK capsule assembly, which protects the sensor and guides the feedthrough wires, which are laser welded to the pins during assembly.
h) Radial spring (not shown) sits in the PEEK groove and makes circumferential contact with the permalloy shield after it is snapped onto the capsule. Spring is laser welded to a feedthrough wire to provide optional ground to the shield.

Measurement Setup for Seismology Signal Acquisition:

Measurement Requirements:

Piezoelectric sensors provide rich data streams due to their broad frequency response and high sensitivity. To detect AHE, signals due to transient impulses are of the most interest. These impulses each last a few 10's to 100's of microseconds and may occur at any point during a several-hour long experiment, almost certainly without any degree of regularity or periodicity. As such, we need a measurement concept that is:

1) Sensitive preferentially to events that are abrupt and impulsive,
2) High enough bandwidth to resolve the character of 10-microsecond excitations,
3) Continuously monitoring over several hours.

The combination of these requirements presents a significant challenge for data acquisition and management. Collecting high-bandwidth, high-precision measurements continuously would result in terabytes of data for each experiment, which would be prohibitive for real-time transfer, storage, and analysis of the acquired data.

We addressed these issues by developing a hybrid acquisition scheme that combined high bandwidth and lower bandwidth data streams. The entire signal and data pipeline is described below and illustrated in Figure 4.

Signal Origination:

First, a thermal impulse in the palladium sample produces a localized, rapid expansion that propagates through the thickness of the piezoelectric sensor as an elastic wave – i.e. a vibration. This in turn causes a burst of electric charge at the surface electrodes of the sensor. This small sub-picocoulomb charge signal is conducted through gold wire bonds, down thin, ceramic-coated platinum wires and out of the ultrapressure vessel through a feedthrough connection. From there, the feedthrough wires are soldered into a vessel-mounted breakout board which connects them to the two inner conductors of a triaxial cable.

Signal Amplification and Filtering:

This triaxial cable connects to a charge pre-amplifier, which converts the small charge signal into a more easily measured voltage signal. We explored two approaches for this pre-amplification step, depending on how well the capacitive noise coupling was mitigated. See the below 'Note on Pre-Amplification' for further discussion.

After pre-amplification, we filter the signal to remove low frequencies that are not of interest. To accomplish this, we use the filtering capability of an SR560 low-noise voltage pre-amplifier. This also allows us to perform a differential subtraction of two inputs, and optionally to amplify yet again with some gain. Differential subtraction was essential when capacitive, common-mode noise dominated over our signal. However, when the dominant noise source is radiative / inductive pick-up, differential subtraction effectively doubles the noise, and better results are achieved using a single-ended measurement. We routinely applied a high pass filter with a 1 kHz cutoff frequency and 12 dB/octave roll-off.

Digital Sampling and Data Stream Management:

The output of the SR560 voltage amplifier is measured using a Zurich Instruments MFLI digital lock-in amplifier / oscilloscope. The signal up to this point has been entirely analog, with the bandwidth only limited by the response functions of the pre-amplifiers. However, now the analog signal must be digitally sampled and recorded. The MFLI input analog-to-digital converter (ADC)

measures 16-bit samples at 60 MSa/sec, but the data could not be streamed to a data server and recorded continuously without losing large chunks of data[94]. However, we can take advantage of on-board processing to create two parallel data streams: one continuous, lower-bandwidth signal used to count impulsive events and to trigger the second data stream of individually triggered, short-duration, full-bandwidth oscilloscope traces showing the detail of each impulse event.

Collecting full-bandwidth triggered oscilloscope traces is straightforward once a few key parameters are known. Namely, the oscilloscope must be configured to acquire a trace long enough to capture the whole impulse response waveform. To characterize the impulse response of the sensor / measurement circuit, we delivered a controlled thermal impulse to the sample surface using a 1064 nm nanosecond pulsed laser (Trumpf TruPulse 2005 nano, selectable waveform 29, 6 ns pulse FWHM). We configure the scope to trigger on a rise in the signal input above a threshold set to roughly double the noise floor, and then fire a single laser pulse at moderate energy (~ 1 μJ). From the collected trace, we observed that the impulse response to a single pulse decays after about 100 μs. Thus, we configure the scope to acquire a 300 μs trace, at 60 MSa/sec bandwidth every time it is triggered. Scope traces can be acquired indefinitely so long as trigger events do not occur within 1 ms of each other. While this does allow for some missed triggers in principle, we are primarily interested in rare events. We could choose to trigger on the input signal directly, but that would be susceptible to triggering on false signals such as noise spikes.

Instead, we can demodulate the input signal (i.e. lock-in to a particular frequency) and use that as the trigger signal for the oscilloscope. The same impulse response we use to confirm the scope trace duration can be used to extract the primary carrier frequency of the sensor and measurement circuit, usually around 245 kHz. This is the characteristic frequency determined by the capacitances and inductances of the entire measurement circuit. In other words, it is not a mechanical resonance of the piezoelectric sensor. The second harmonic of this frequency can also be used effectively. When the input signal is demodulated at this frequency, with a fast, 10 μs time constant, the resulting processed signal is effectively the enveloping function of the impulsive events. It rises and falls once for each impulse, so it does not contain all the internal waveform information, but it is perfect for counting events and triggering the higher-resolution full-bandwidth oscilloscope traces.

Real-Time Monitoring and Data Storage:

The high-resolution scope traces and the continuous envelope function cannot both be streamed to the data server from the same instrument without data loss, so we instead configured the MFLI to output the envelope function to its auxiliary output, with some amplification, and we connected that output to the auxiliary input of a second MFLI instrument. This second instrument was used to stream the envelope function to the data server with 234 kSa/sec bandwidth, and to monitor the signal in real time during the experiment.

[94] Küng, B. (2015, July 7). Achieving High Data Transfer Rates with the UHFLI. Zurich Instruments. https://www.zhinst.com/americas/en/blogs/achieving-high-data-transfer-rates-uhfli

In summary, our measurement continuously monitors the piezoelectric sensor at a fast enough rate to detect rapid impulses, but not at such a high rate that the data acquisition pipeline would be overwhelmed. Instead, this continuous data stream is used to determine precisely when high-resolution data should be collected for a short time. We can use the continuously sampled envelope function to analyze how many impulsive events occurred, how densely, and in which parts of the experiment. We can use the high-resolution data to analyze the waveform of each impulse event individually.

Note on Pre-Amplification:

In closer-to-ideal configurations, where capacitive noise coupling was not a significant issue, we used a single-ended charge amplifier (Femto HQA-15M-10T). This amplifier has an extremely high gain of 10 V/pC, which makes it ideal for sensitive charge detection. We then passed the output through an SR560 low-noise voltage pre-amplifier with 1x gain, in single-channel source mode, purely to take advantage of the 12 dB/octave high-pass filter. The output of the SR560 was then measured using a Zurich Instruments MFLI lock-in.

However, the Femto HQA's single-ended geometry leaves the measurement susceptible to capacitive noise on the signal lead. Capacitance between the signal lead and any noise source (i.e. any conductor with a slightly fluctuating voltage) directly results in charge noise that can rail the amplifier output: $\delta q = C\delta V$. Indeed, our first measurements could not use the high-gain, single-ended configuration at all. The resulting noise railed the output completely, dominating even over a high-powered pulsed-laser calibration signal, as shown, for example, in Figure 5. It should also be noted that the PS-15-25-L switching power supply recommended for use with the Femto HQA-15M-10T, despite excellent performance compared to other switching supplies, still generates enough noise to be problematic for the extremely sensitive measurements we required. As such, we built a linear ±15 V power supply using two Acopian A15H1500 regulated linear power supplies, wired together in series with their midpoint held at ground. We also found it essential to provide a good ground connection to the preamplifier chassis through the output BNC shield.

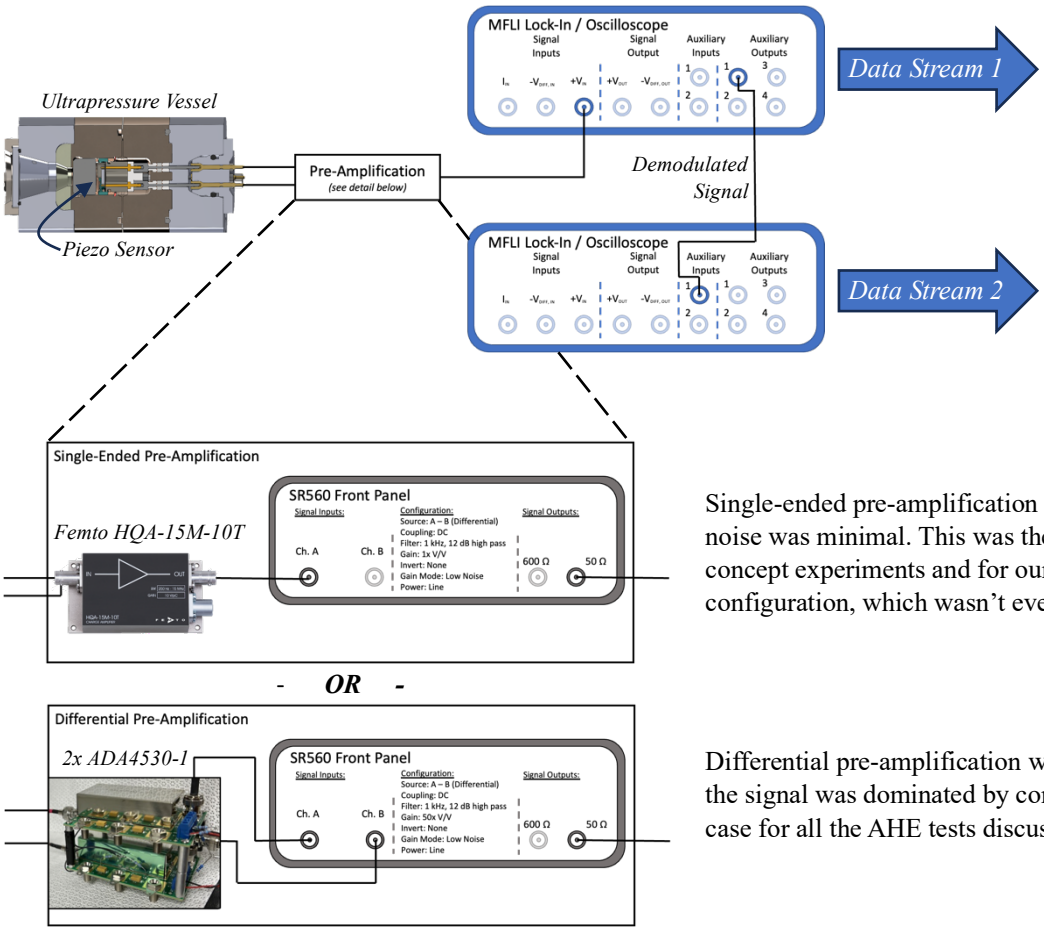
When capacitive noise dominated over the signal, we used a differential charge amplifier that we built in-house. In this scheme, the charge from each surface of the piezoelectric sensor is individually measured using charge-to-voltage amplifiers built using development boards for the ADA4530-1 integrated circuit from Analog Devices. Then the output of each development board is connected to the two inputs of the SR560 low-noise voltage pre-amplifier, which is operated in differential source mode to subtract one input from the other. This amplifies the real signal while reducing the common-mode noise that couples into each measurement lead identically. In this case, each individual charge-to-voltage amplifier was lower gain than the corresponding Femto HQA, so that the output would not rail. We therefore used the SR560 for additional amplification, with 50 V/V gain, instead of purely for filtering.

Noise Characterization:

The sensitivity of our measurement also comes with a heightened vulnerability to electronic noise, made worse by the bustling experimental infrastructure of our ultrapressure lab space. Since every detail of the experiment had potential impact on noise performance, we needed a standardized way to characterize the noise performance of each specific configuration. Our standardized measurement captured two oscilloscope traces of the baseline input signal, connected to the sensor but without any external signal sources such as pressure ramps or laser shots. Thus, this baseline represented the inherent thermal and vibrational response of the sensor to its ambient environment as well as the electromagnetic interference present in the circuit. We routinely measured one oscilloscope trace at 60 MSa/sec sampling rate, for 80 ms, and one trace at 3.75 MSa/sec, for 1.2 seconds. In the time-domain, these traces enable qualitative observations of any periodic features, as well as quantitative metrics like the peak-to-peak and root-mean-square spread in the voltage signal. In the frequency domain, the Fourier transform of these oscilloscope traces offers standardized insight into the magnitude and effective 'fingerprint' of each noise source[95]. The frequency-domain resolution of a time-series is the reciprocal of its duration, and the maximum frequency component is half the sampling rate. Thus, the longer duration trace provides sub-hertz frequency resolution up to 1.875 MHz, and the shorter duration trace offers 12.5 Hz resolution for frequencies up to 30 MHz. Using only the longer duration spectrum would miss any noise feature above 1.875 MHz, but using only the shorter duration may fail to resolve sharp low-frequency peaks, for example at 60 Hz.

[95] Heinzel, G., Rüdiger, A., & Schilling, R. (2002). Spectrum and spectral density estimation by the Discrete Fourier transform (DFT), including a comprehensive list of window functions and some new at-top windows. https://holometer.fnal.gov/GH_FFT.pdf

Figure 4: Seismology Measurement Setup



Triggered stream of high-resolution (60 MSa/sec) detail traces of individual events. Each is 300 µs in duration and must be separated by more than 1 ms from its neighbors to avoid loss.

Continuous stream (234 kSa/sec) of samples of the envelope function, which detects the shape of pulses occurring but not their internal detail. This is used to trigger Data Stream 1 and to put individual events in context.

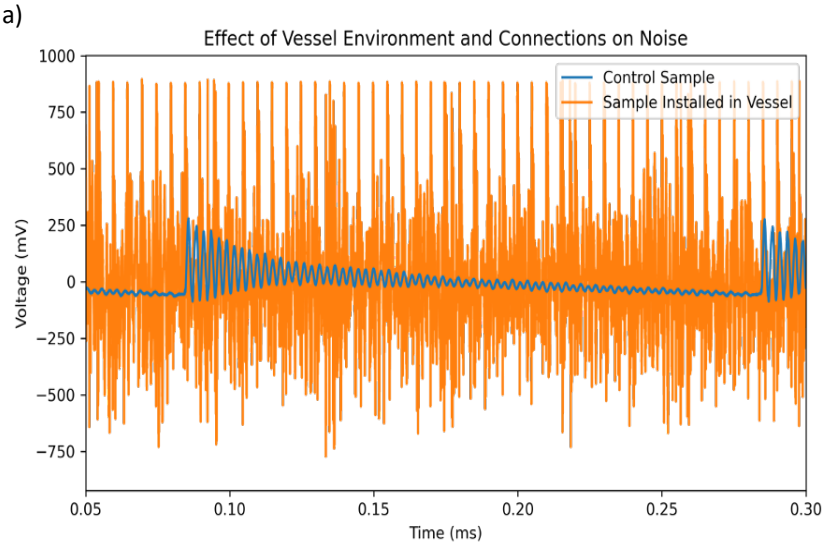
Single-ended pre-amplification with very high gain was used when noise was minimal. This was the case for benchtop, proof of concept experiments and for our ultimate low-noise measurement configuration, which wasn't ever used for AHE tests.

Differential pre-amplification with lower gain was used when the signal was dominated by common-mode noise. This was the case for all the AHE tests discussed in this report.

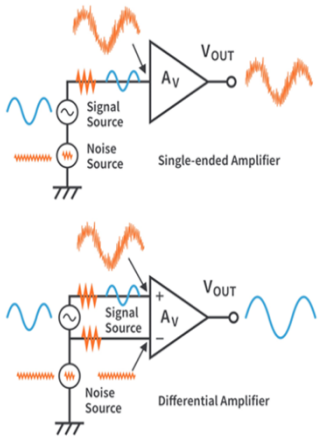
We traced most noise sources to various switching power supplies. If these supplies were part of our experimental setup, we could often remove or replace them. Some examples were switching supplies for a ThorLabs K-Cube motor controller, for the Femto HQA charge pre-amplifier, and for an internal transformer in the BRX series of programmable logic controllers. However, some noise features were present in our entire lab space, and we speculatively attributed them to larger equipment such as the facility air compressors, ventilation fans and HVAC equipment. Of note was a feature around 60 – 70 kHz that appeared to be dithered across a range of frequencies, which is a common method used to reach electromagnetic compatibility compliance in industrial equipment. These noise sources often contained both ground-coupled and radiative components, so we found it necessary to iterate over many different grounding and shielding configurations in order to optimize noise performance. The ultra-pressure plumbing and vessel acted both as a conductor that carried ground-coupled interference and as an antenna that picked up radiative interference. Ultimately, we found it necessary to minimize the capacitance between the measurement leads and the ultra-pressure vessel, as well as shield the vessel from radiative interference using a modified Faraday cage.

Figure 5: Mitigating Common-Mode Noise

a)



b)



a) *Need for noise mitigation strategy*: Installing the sample in an ultrapressure vessel resulted in unacceptably high capacitive noise coupling and motivated an effort to mitigate common-mode noise.

b) *Single-ended vs. Differential Amplification*: Differential amplification is very effective at removing common-mode noise (Figure 6), which is present on both measurement leads, with the same phase. Noise from the same source that is injected capacitively onto each lead will inherently be common mode. However, differential amplification will make differential mode noise even worse. This type of noise occurs when noise from the same source couples onto each lead with an opposite phase, such as when the leads form a loop that acts as an inductive antenna. The optimal measurement strategy therefore depends on the noise environment.

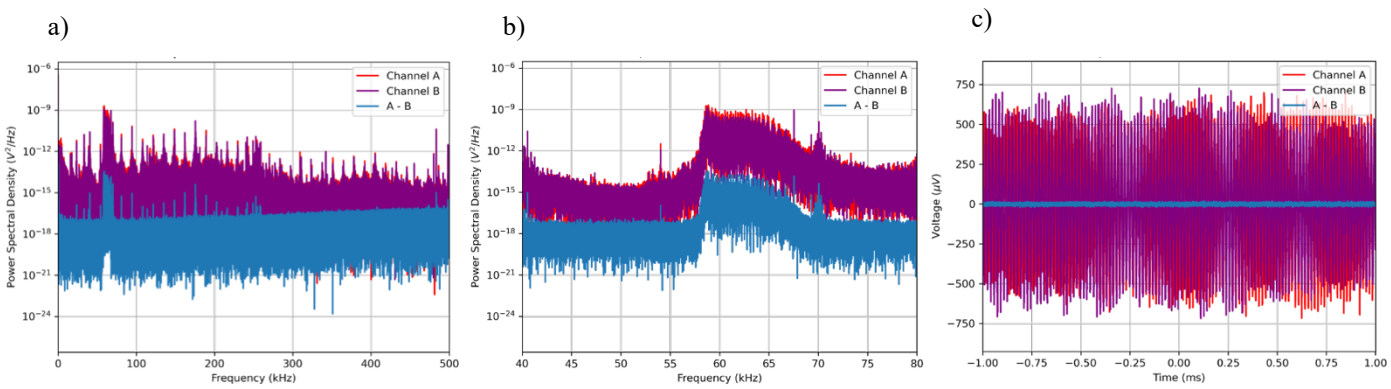


Figure 6: Differential Amplification Reduces Common-Mode Noise

a) Power spectral density of single-ended channels A and B vs. differential channel A – B, showing a reduction of up to 5 orders of magnitude in noise power.
b) Detail of a prominent noise feature at 60 – 70 kHz that often dominated the signal.
c) The time domain trace shows dramatic improvement.

Controlled Thermal Impulse Calibration:

Delivering a controlled thermal impulse to the sample surface using a pulsed laser has proven to be quite instructive, as already explained in a few instances above. Essentially, a nanosecond pulsed laser simulates the thermal impulse we expect from a burst of AHE events, but perfectly on demand and fully characterized.

The laser we used is a Trumpf TruPulse 2005 nano. Its wavelength is 1064 nm, in the near infrared, and the pulse waveform can be selected from a handful of options with different durations and pulse shapes. We used Waveform 29 as our AHE simulator, which is a gaussian pulse with a full width at half maximum of 6 ns, and 7 ns above 10% power. We also measured the response at the longest available pulse duration, using Waveform 0, which has a narrow pulse with a full width at half maximum of 35 ns with a slow trail off spending 250 ns above 10% power, and we didn't observe any difference in the signal for equivalent pulse energy. This supports our assertion that these pulses all occur in the impulsive limit, and are thereby reasonable approximations of an AHE event.

The laser also has a programmable repetition frequency, which we most often used during experiments to fire bursts of six shots spaced 1 ms apart from each other, or to deliver single, isolated shots. For characterization of the sensor response, we used continuous shots at 5 kHz so that several pulses would be observed in the same scope trace. However, repetition rates higher than 5 kHz would not allow the response to damp completely before the subsequent shot.

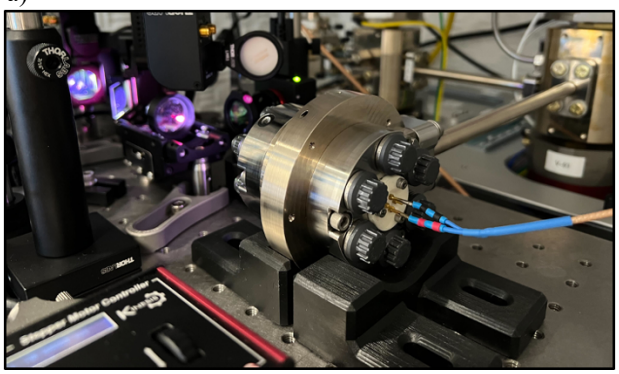
To control the pulse energy, we polarized the beam at its output, then adjusted the polarization using a half-waveplate, and split the resulting beam using a polarizing beam cube. To attenuate further than the extinction ratio allowed, we inserted an additional neutral density filter as needed. The pulse energy was measured in-situ.

For all of the discussed experiments, the beam was delivered to the sample using free-space optics. However, to achieve the best possible noise performance, we later rebuilt the optical system to use a fiber for final beam delivery to the ultrapressure vessel. This allowed us to maintain the integrity of our Faraday cage. See Figure 7 for details on the optical setup.

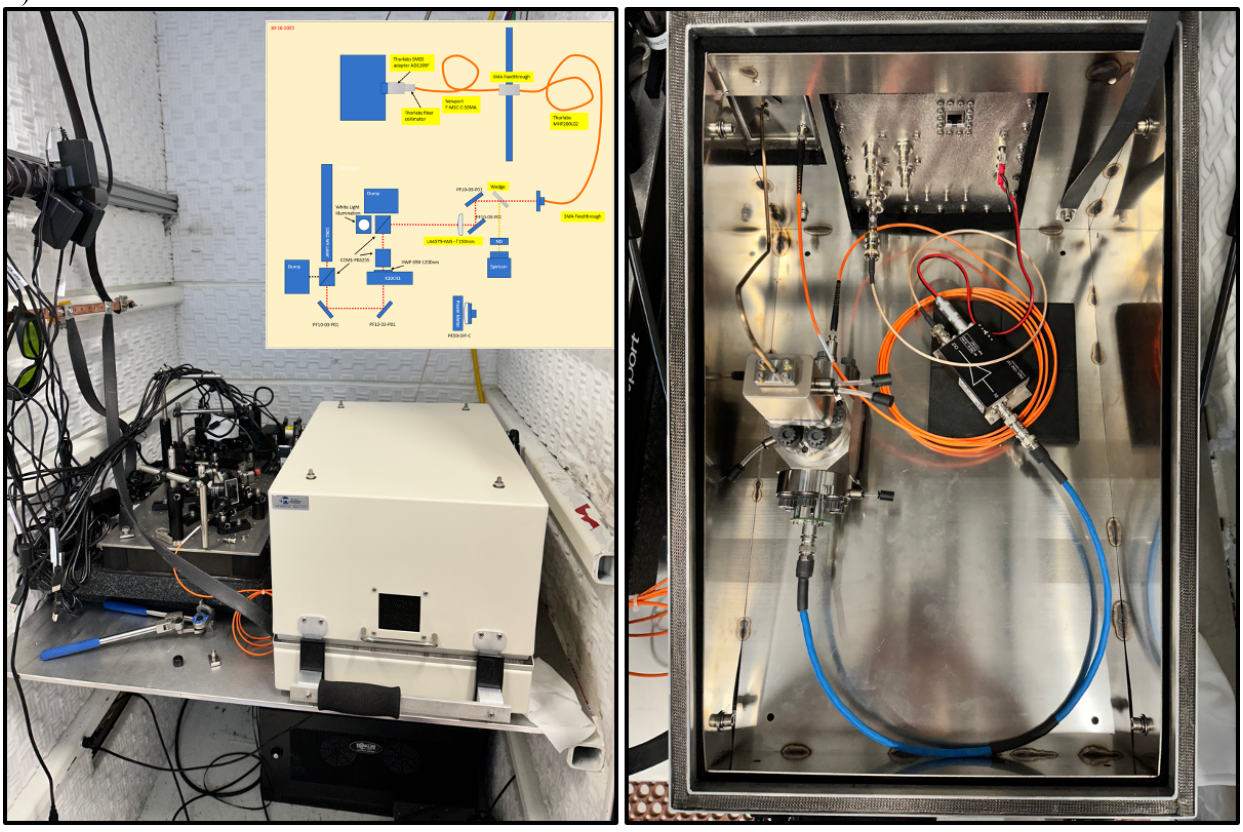
To determine our minimal detectable energy, we simply fired shots at decreasing pulse energies until we could not discern the pulse above the noise floor. The magnitude of the sensor response was consistent from sensor to sensor, for a given pre-amplifier configuration, so any variation in the sensitivity between experiments results completely from variations in the noise floor.

Figure 7: Delivering a Controlled Thermal Impulse Using a Pulsed Laser

a)



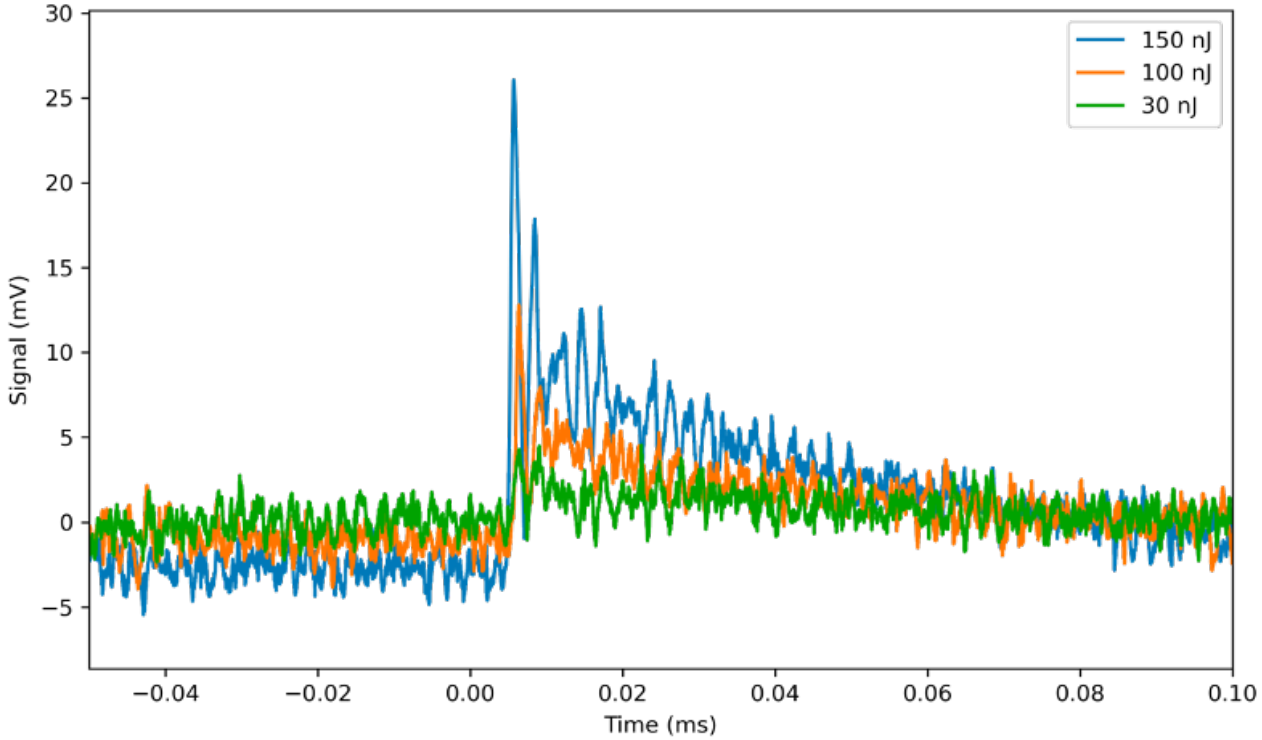
b)



a) Free-space beam delivery to the sample through a sapphire window in the ultrapressure vessel.
b) Fiber beam delivery through an enclosed Faraday cage, which allows better noise performance. The beam still exits the fiber and passes through the sapphire window to reach the sample inside the vessel. Inset is a schematic of the energy attenuation and beam delivery scheme.

Figure 8 shows an example sensitivity test that can resolve a 30 nJ pulse. This equates to roughly 8000 AHE events. For a 100 nm thick, 4 mm diameter palladium thin film, that represents a fractional site sensitivity of 1 in 10 Trillion (1e-13). This itself is a conservative estimate, because we do not account for the loss due to imperfect absorption of palladium at 1064 nm. However, this was our best performing configuration, and our noise floor in the experiments discussed below was worse than this, with 300 nJ pulses easily detectable. That equates to a fractional site sensitivity of around 1 part in 1 Trillion (1e-12), which is already on par or slightly more sensitive than our most sensitive calorimeter. However, the calorimeter is sensitive to total energy regardless of event rate, and our measurement assumes the entire burst of events occurs quickly enough to be considered impulsive, which we estimate to be within a 10 microsecond window.

Figure 8: Measurement of Detectable Energy



Measurement of minimum detectable pulse energy in experiment BE4191. Detecting 30 nJ in this sample geometry represents an equivalent to 8,000 AHE events, or a fractional site sensitivity of 1 part in 10 trillion. Typical experiments, including those discussed below, could detect 80,000 AHE events for a fractional site sensitivity of 1 part in 1 trillion, which is still on par with our best calorimetry.

Seismology Experiments Using Loaded PdD / PdH:

We performed Seismology measurements on palladium thin film samples in ultrapressure deuterium and protium up to 145 ksi. We incorporated laser shots into these experiments as verification of sensor performance, but these were not intended as a stimulation source, and the experiments are treated as passive (unstimulated). These experiments generate large data sets, even when the data streams are managed as described above, so we built analytical tools to aid our interpretations.

Our overall strategy is to progressively explain and rule out any impulse event that could be explained classically. We did not observe any significant isotopic differences between deuterium tests and protium tests, nor did we observe any meaningful number of unexplained impulse events at all in either case. Below, we discuss our experimental procedure and data analysis techniques.

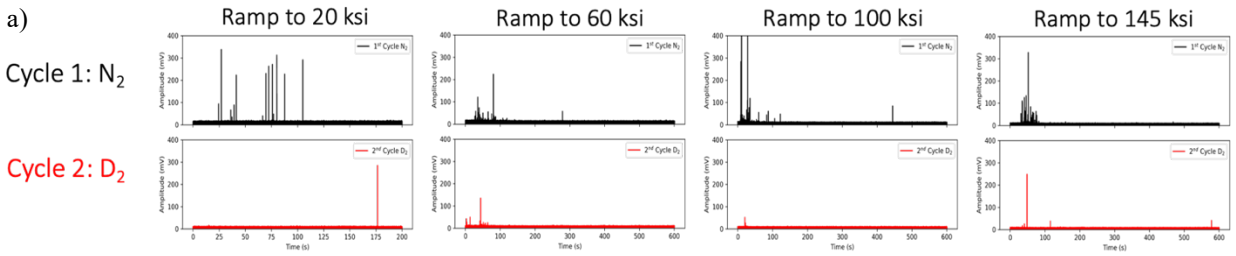
Pre-Conditioning and the Kaiser Effect:

In non-destructive testing, acoustic emission hits occur when a material is placed under mechanical load. In our ultrapressure experiments, our sensor itself is placed under tremendous hydrostatic load, so it follows that many acoustic emission hits will develop from this mechanical stress alone. However, it is also well known that after a material is mechanically loaded and unloaded cyclically, acoustic emission hits are suppressed on subsequent cycles until the prior peak level has been surpassed. This is a phenomenological observation known as the Kaiser Effect.

We take advantage of this effect by subjecting our sensors to 5x pressure cycles in an inert gas (nitrogen), before ultimately loading them in either deuterium or protium. By the fifth nitrogen cycle, acoustic emission hits are almost completely suppressed, except during periods of active pressure ramping. Thus, we can focus our analysis on acoustic emission hits that occur during pressure holds without being overwhelmed by impulses occurring as the material responds to isostatic pressure.

It is essential to verify that cycling does not reduce the sensitivity of the sensor to a thermal impulse. We seek to suppress the origination of non-thermal impulse events – not to suppress our ability to detect events at all. To verify this, we measured the sensor's response to five laser pulses both before and after inert pressure cycling. We observed no decrease in sensitivity to thermal impulses, while impulses occurring due to pressure alone were greatly suppressed. See Figure 9.

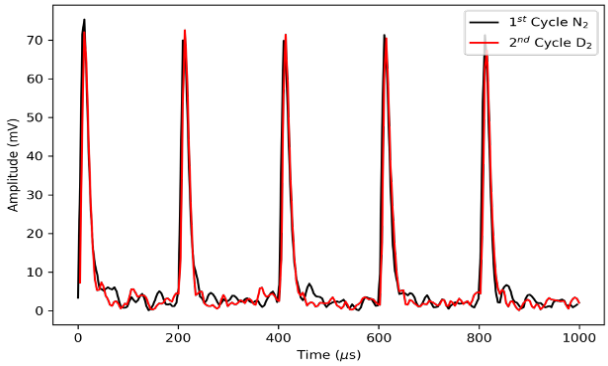
Figure 9: Pressure Cycling Does Not Reduce Sensitivity



a) Demonstration of the effectiveness of one inert pressure cycle on suppressing acoustic emission hits due to the mechanical stress applied by pressure.
b) Demonstration that the sensitivity to thermal impulses is not affected by the pressure cycling.

It is worth noting that 5x pressure cycles suppress unwanted acoustic emission hits even better, but not strictly perfectly.

b) Response to thermal impulse from laser



Experimental Procedure:

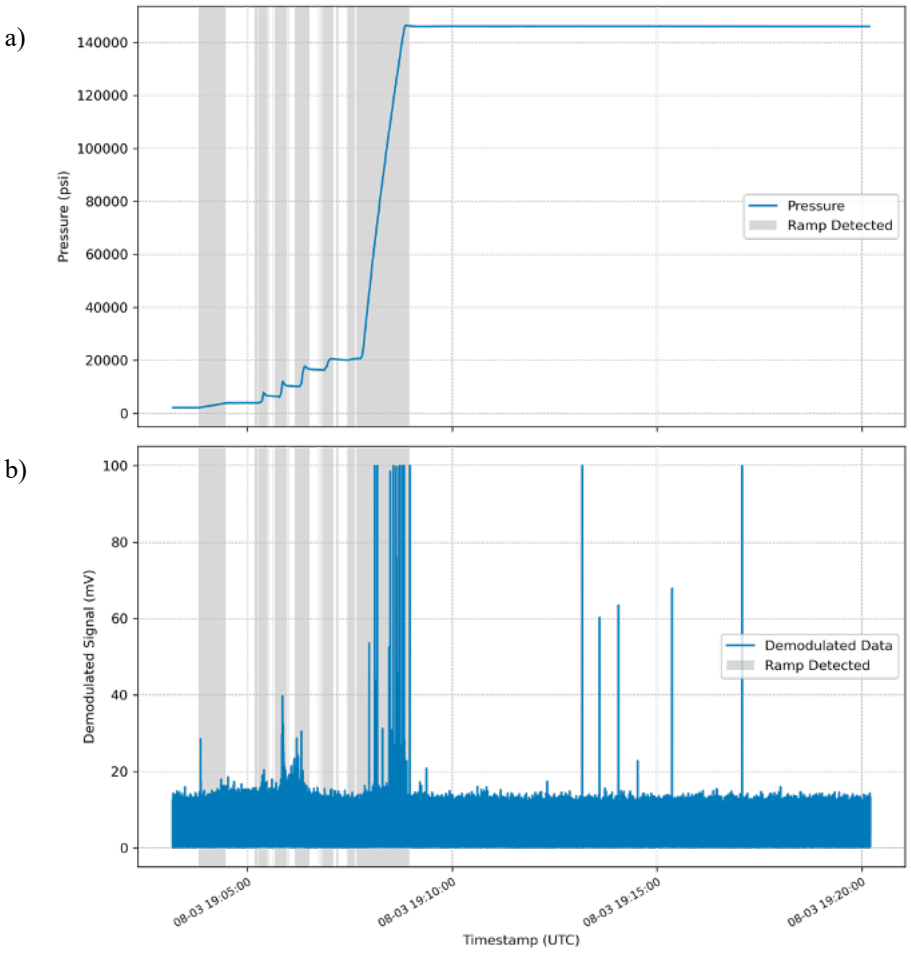
In these measurements, we seek to identify the presence of thermal impulses that are not accounted for by classical processes, especially when those impulses are present under deuterium loading and absent under protium loading.

First, we prepare the sample assembly as shown in Figure 2. Next we pressure cycle five times up to 145 ksi nitrogen to condition the sample. We collect both long-duration and triggered data streams during all these cycles, as described above. Next, we cycle to 145 ksi deuterium. By this point the number of triggered impulses per cycle has decreased dramatically. Next, we cycle up to 145 ksi protium, as an additional control.

Since we use the same thin film sample for both deuterium and protium loading, it is essential to run the deuterium cycle first. We have learned that thin films do not completely de-load after de-pressurization, so there is certainly still a substantial amount of deuterium loading during the protium cycle. If protium was introduced first, then we would never reach a state of high deuterium loading – because there would instead be a substantial fraction of protium in the sample. We argue that the protium-deuterium mixture still serves as a good control in which we do not expect AHE would be possible. However, the point is in fact moot because we ultimately did not observe any unexplained thermal impulses in the deuterium cycles in the first place – so there was nothing in particular to control for.

After the experiment, data was analyzed as described below. In each stage of analysis, we progressively classify and rule out acoustic emission hits until we are left only with events that have no clear explanation.

Figure 10: Long-Duration Data Gives Context to Each Event



a) Pressure data is used to differentiate between times when the pressure is ramping and times when the pressure is holding.
b) Long-duration demodulated data shows that most events are clustered during pressure ramps, and assigns labels of interest to events observed during pressure holds, where we expect fewer classically explained events, for further scrutiny.

Analysis of Continuous Long-Duration Data Stream:

The primary advantage of the long-duration data stream is to place measured events in the context of what else was happening in the experiment at that moment. Along with the 234 kSa/sec demodulated envelope data, we also collect data from the pressure transducers at 1 Sa/sec. This data is used along with the trigger timestamp of each measured event to assign a pressure value, a gas identifier, and a label indicating whether the pressure was ramping or holding. These are all saved as columns in a data frame that summarizes each measured event. Most of our experiments record of order one thousand events.

Figure 10 shows an example of how this pressure data can be used to differentiate between events that occur during pressure ramps and events that occur during pressure holds.

Analysis of High-Resolution Triggered Data Stream:
We seek to use the high-resolution, triggered data to make interpretations about individual acoustic emission hits. We are interested only in events that appear to be sharp impulses, so to distinguish between these events and those of other origins, we employ several techniques from signal processing and machine learning.

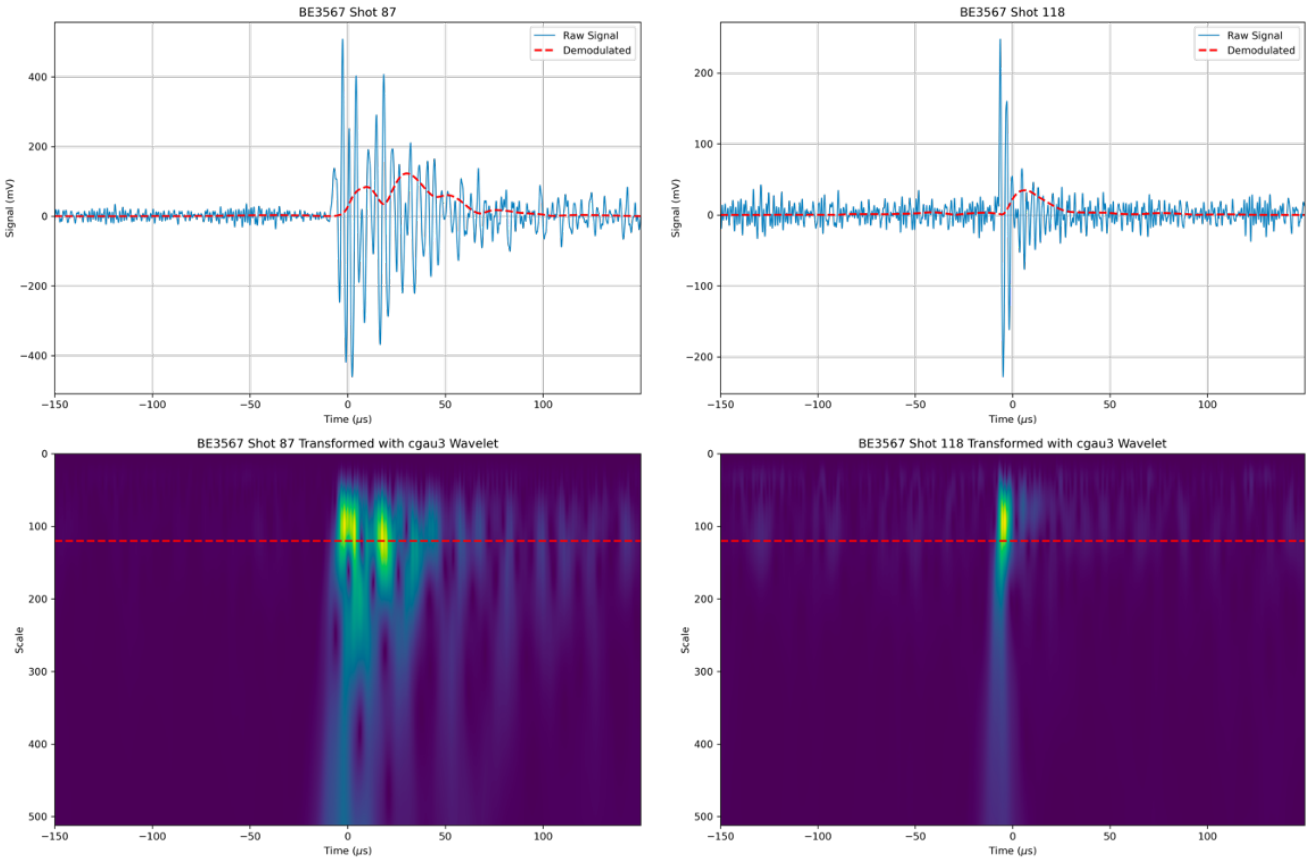
First, we use a wavelet transformation of each data trace to obtain a fingerprint that shows how the frequency components of the signal evolve over time. Breaking a signal down into frequency components using Fourier analysis has been a staple of signal processing for centuries. However, transforming entirely to the frequency domain loses all information about the duration of the pulse in time, which is a crucial attribute to distinguish between impulsive and non-impulsive events. Instead, wavelet transformation produces a two-dimensional data set that includes both frequency information and time information[96]. Figure 11 shows an example of two measured acoustic emission hits, the demodulated envelope functions which triggered collection of the data, and the wavelet transforms using a convolution with the 'cgau3' complex gaussian wavelet.

Our experiments generate hundreds to thousands of detailed traces representing acoustic emission events, so we turned to unsupervised image clustering techniques in machine learning to group them based on how similar they are to each other.

First, we run a batch process to calculate the wavelet transforms of each event and save them as a png image. Next, we process the images through a neural network (ResNet-50) to assign an abstract, many-dimensional 'feature vector' to each image. Then we use principal component analysis to determine which hundred or so components of the feature vector are responsible for all of the variation between the images. Next we need to visualize which events are similar to each other, so we need to project the roughly 100-dimensional principal component space onto just two axes. To accomplish this, we use a t-stochastic nearest neighbor embedding (t-SNE) algorithm to optimize the plane we project onto so that it preserves most of the spread.

[96] Arts, L.P.A., van den Broek, E.L. The fast continuous wavelet transformation (fCWT) for real-time, high-quality, noise-resistant time-frequency analysis. *Nat Comput Sci* **2**, 47-58 (2022). https://doi.org/10.1038/s43588-021-00183-z

Figure 11: Wavelet Transform Analysis of High-Resolution Data

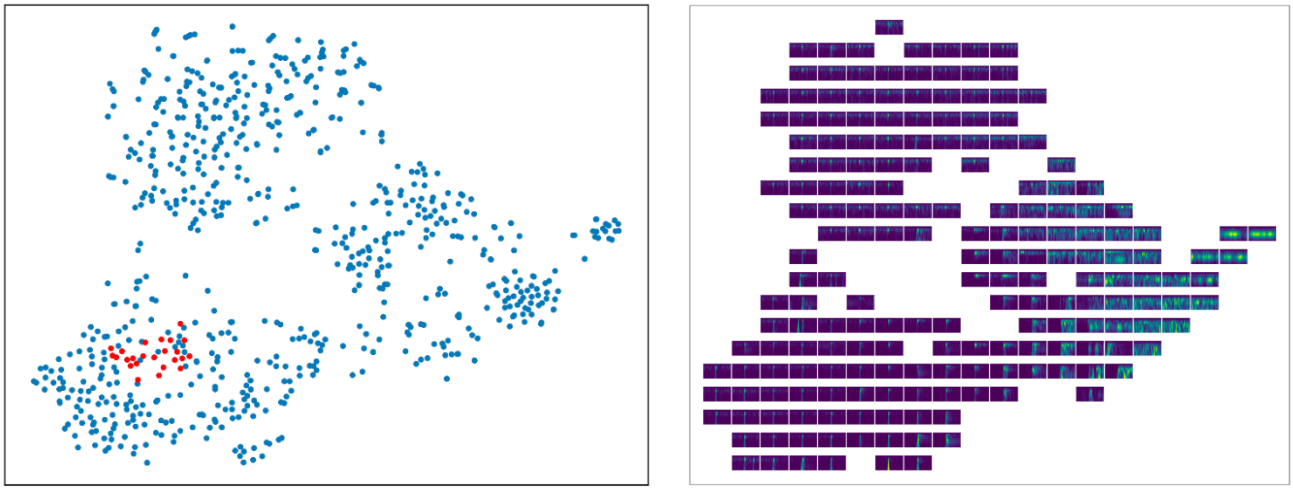


The above show acoustic emission hits that were recorded in high resolution, along with the demodulated data that triggered collection of the events. The wavelet transforms shown underneath the time traces represent 'fingerprints' that can be clustered to determine which events appear similar to thermal impulses. For example, thermal impulses are bursts of excited broad-band frequencies that decay quickly after their initial excitation, while other types of acoustic emission can show gradual increase or beating over time. The 'scale' axis shown on the y-axis in the wavelet transformations is proportional to frequency domain. The horizontal red dashed line in the wavelet transform images represents the center frequency of the demodulation.

a) An acoustic emission hit that occurred during pressure ramping.
b) An acoustic emission hit that resulted from a laser pulse.

The result is a plot like the one shown in Figure 12, where an abstract 'similarity' between two points is captured by the distance between them in the plane. The units of the axes are meaningless, because they ultimately originate from the abstract way the neural network encodes features within the images that are used to differentiate them from each other. We can quickly see that the clustering corresponds to real similarities because all laser shots group together. Furthermore, if we replace the points on the graph with image thumbnails, we can see that similar images are grouped together.

Figure 12: Machine-Learning Based Image Clustering



a) The 'latent space' representation of all the measured acoustic emission hits in a particular experiment. Points that are closer to each other on this plot are similar in an abstract way defined by the ResNet-50 neural network. Red dots indicate laser shots, which group together as they should.
b) Replacing each dot with a thumbnail image shows that indeed, acoustic emission hits with similar wavelet transforms are grouped together in this depiction.

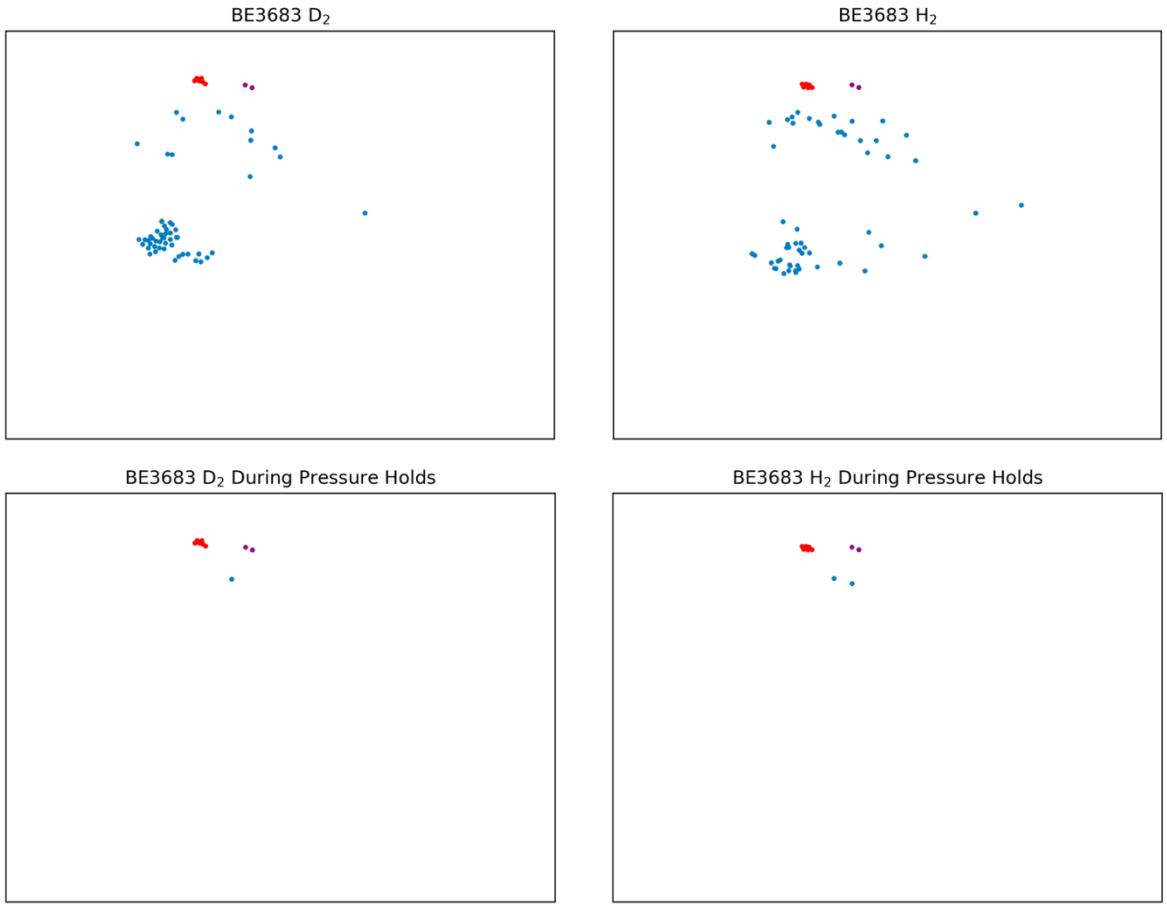
Using the contextual data from the long-duration data stream, we can now progressively filter out events that occur during pressure ramps, and separate out events occurring under deuterium loading and protium loading.

Results:

To determine whether a given experiment contains impulse-like events that are unexplained, we start from the analysis described above and filter out unwanted events in steps.

For example, in BE3683 shown in Figure 13, we first separate events into those that occur under deuterium loading vs. protium loading, and then we further remove events that occurred during pressure ramps. In the experiment shown, some single-pulse laser shots were taken (red), and some multi-shot laser pulse trains that contain multiple shots in the same scope trace were taken (purple). After removing uninteresting events, we are left with only one possibly interesting impulse in the deuterium experiment, and only two in the protium experiment. This is not a meaningful isotopic difference, and the result is therefore interpreted as negative for AHE.

Figure 13: No Isotopic Difference Observed



a) and b) Clustered events loaded under deuterium and protium. Red dots are single laser shots. Purple dots are multiple laser shots in the same scope trace. Blue dots are not laser shots.

c) and d) Remaining events after removing events that occurred during pressure ramps. No isotopic difference is observed. Furthermore, there are not a significant number of impulsive events that remain unexplained at all, indicating that our pre-conditioning and context-based filtering works well.

Conclusions

To date, we have not measured any isotopic difference from Seismology experiments, which we interpret as a negative result for AHE with 1 part in 1 Trillion fractional site sensitivity. These experiments were unstimulated, and AHE measurements were not performed in our highest noise-performing experimental configuration before the project wound down.

Further progress was planned to expand into larger bulk samples including palladium slugs, to explore different stimulation methods, to monitor for much longer durations, and to further develop more nuanced machine learning techniques to interpret the data. However, many of these avenues were explored thoroughly across other projects in our global effort.

Nuclear Detection

Executive Summary

In addition to our search for anomalous heat events, we also considered the possibility that D-D fusion may occur within a lattice and emit traditional nuclear byproducts. We performed experiments ranging from post-mortem "staring" to look for delayed emission of particles to in-situ measurements during ionization of PdD_x with high intensity laser pulses. In no cases did we observe nuclear products suggestive of deuteron fusion.

Background

The general premise of the LENR field is the observation of anomalous heat events; these events are suggested to be the coupling of nuclear byproducts to a material lattice, namely a 24 MeV photon, that is emitted during D-D fusion events to phonons within the material. One difficulty in this assumption is the mechanism by which the 24 MeV couples to the lattice: Generally, a high-energy *phonon* (lattice vibration) has an energy on the order of 10s of meV. This large mismatch between the nuclear and phonon energies (9 orders of magnitude difference in energy) is a major stumbling block in theoretical attempts to understand LENR.

Rather than assuming an unknown scattering/coupling process arises for the emitted byproduct, we also considered the possibility that D-D fusion is occurring within the lattice and simply emitting the expected uncharged nuclear radiation – a 24 MeV photon or a 2.45 MeV neutron. Further, we consider the possibility that a reaction occurs between Pd ions and the deuterons within the lattice, leading to transmutation of the Pd atoms; this transmutation would lead to the emission of a characteristic x-ray/gamma ray that is easily identifiable with isotope libraries. These forms of nuclear detection are implemented in several of our experiments: These nuclear detection experiments range from us performing a standard "stimulation" experiment on PdD_x followed by placing the material in an ultra-sensitive "well detector" that eliminates background counts, to active stimulation of materials to conditions well above their ablation/ionization threshold with arrays of various detectors actively looking for particle emission. As detailed in this section, we ultimately **observe no evidence of nuclear transmutations or fusion reactions (i.e., LENR) within a solid-state lattice via nuclear detection.**

Detectors and analysis

Our experiments on nuclear detection rely on three primary detector types: Active gamma detectors, gamma ray well detectors, and both plastic and liquid scintillators. In this section, we provide a brief background on each detector type, their advantages, and their use for LENR investigations.

One of the primary claims in LENR/AHE research is that D-D fusion within a solid-state lattice results in the emission of 24 MeV that couples to the lattice to generate heat. However, one channel of traditional D-D fusion would simply emit this as a 24 MeV photon without heat generation and be detectable by gamma ray detectors. Thus, we implement two forms of gamma ray detection. The first is with an in-situ detector that allows for gamma detection during stimulation experiments; in other words, the detector can be placed next to our sample/pressure vessel during any experiment performed. This looks for "prompt" gamma rays that are generated during stimulation. The second form of gamma detection is performed using a gamma ray well detector; in this instrument, a sample is placed within a deep well that is highly shielded by highly dense, thick walls to almost entirely remove radiation that is *not* emanating from the sample of interest. While this well detection method significantly improves signal to noise ratios, the geometry of the well limits the use to only post-mortem measurements of samples that have been stimulated. In other words, we use a gamma ray well detector to investigate potential *delayed* gamma emission – gamma rays that have been emitted sometime after stimulation occurs. These may take the form of delayed D-D reactions or signatures of transmutation as the isotope decays. Our other method of detection is via liquid or plastic scintillators. These detectors provide the advantage of not only detecting gamma rays, but detecting *any* ionizing radiation that reaches the detector. This allows us to broaden our search to include potential 2.45 MeV neutrons which are likely to occur from a branching ratio perspective.

Prompt gamma rays from D-D fusion or transmutation of elements

To investigate potential prompt gamma rays, we use an ORTEC high purity Germanium (HPGe) solid state detector. For these experiments, we place the detector near our pressure vessel; based on the geometry used, our detector's solid angle corresponds to a 5-10% detection probability[97] assuming isotropic emission. This "prompt gamma" detector was used in nearly every experimental campaign, including Calorimetry, USP, TI, HPE, and CHIPS.

[97] Our "standoff" distance, the distance between detector and sample, varies from 6" to 12" in most experiments, corresponding to ~5 and ~10% solid angle coverage, respectively.

In this example, shown in Fig. 1, we can see example spectra for both isotopes of Pd as well as a control experiment with no Pd in the calorimeter. As shown, there is no discernible differences between the plots, suggesting no additional nuclear activity in any given case. However, it is worth highlighting the background features from having an exposed detector (i.e., a non-well detector with sufficient shielding to external radiation); a large amount of background counts is due to the elements that exist within normal construction materials, for example, the peak at 511 keV is due to Na-22; the large signal at ~0 keV is simply background noise from thermal excitations; etc. Overall, **the fact that there is no isotopic variation observed in *any* prompt gamma measurements suggests a lack of nuclear-related processes within these experimental campaigns.**

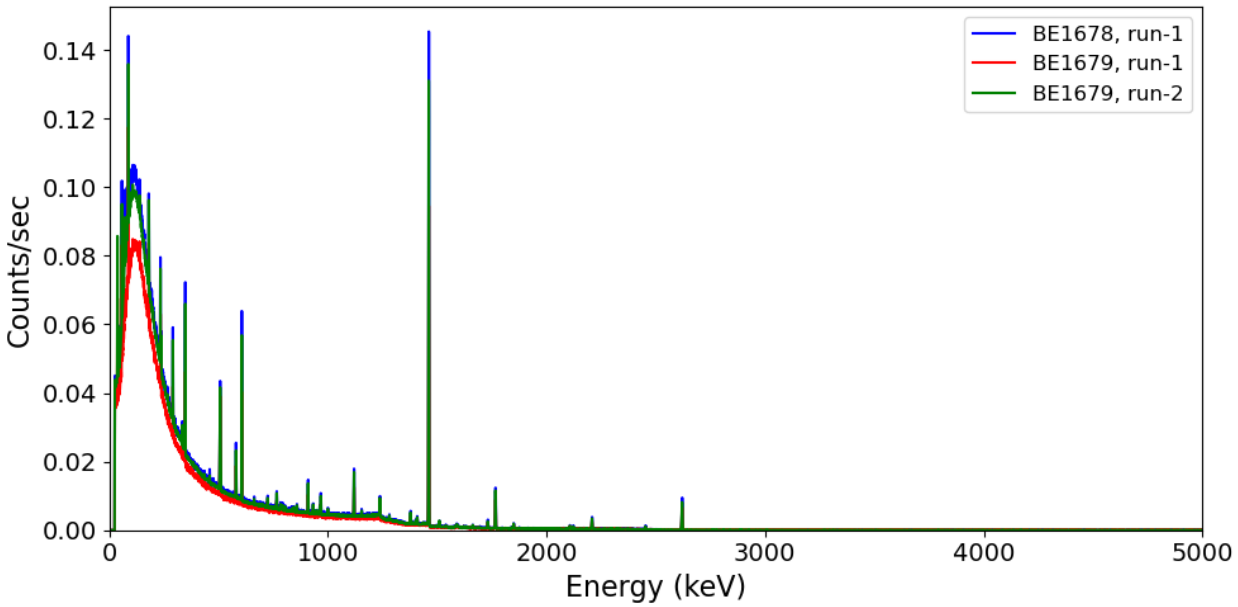


Figure 1: Prompt gamma emission detected during a loading experiment with (red) PdD_x, (blue) PdH_x, and (green) no sample.

Delayed gamma rays from D-D fusion or transmutation of elements

One alternative approach to detection of prompt emission is placing an already-stimulated sample of PdD_x within a gamma well detector that reduces these background counts by orders of magnitude. This approach would allow for a greatly increased probability of detecting transmutations of Pd that may be due to nuclear events. The decay chains following neutron capture by isotopes Pd^{102}, Pd^{108}, or Pd^{110} involve steps slow enough so that emitted gamma rays can be detected by promptly placing the Pd into a gamma well detector after its use in an experiment. Notably, this form of detection was used in multiple studies by NASA, in which deuterated materials such as Ti and Er were exposed to high energy gamma rays for prolonged durations; afterwards, the samples were placed in a low-background environment and delayed

gamma events were measured. In these studies, the authors observe non-negligible increases in gamma rays suggestive of transmutation; while this work was indicative of solid-state driven nuclear reactions, it is critical to recall their use of >1 MeV stimulation methods, which is energy well-in-excess of the energy required to drive traditional nuclear processes. In other words, this method of delayed gamma emission has been utilized to identify when *lattice-based* nuclear reactions occur *at high energies*; comparable evidence at low energies, which is our primary interest, has not been observed.

An example of a delayed-gamma experiment is shown below in Fig. 2; this data was obtained following a calorimetry experiment using Pd nano-powder. One immediate contrast relative to our prompt data in Fig. 1 is the significant reduction in additional peaks: While the Na-22 peak at 511 keV is still present, for example, a large number of other isotopes associated with building materials have been removed. Further, one can note the total counts/sec is reduced by over an order of magnitude when comparing Fig. 2 to Fig. 1.

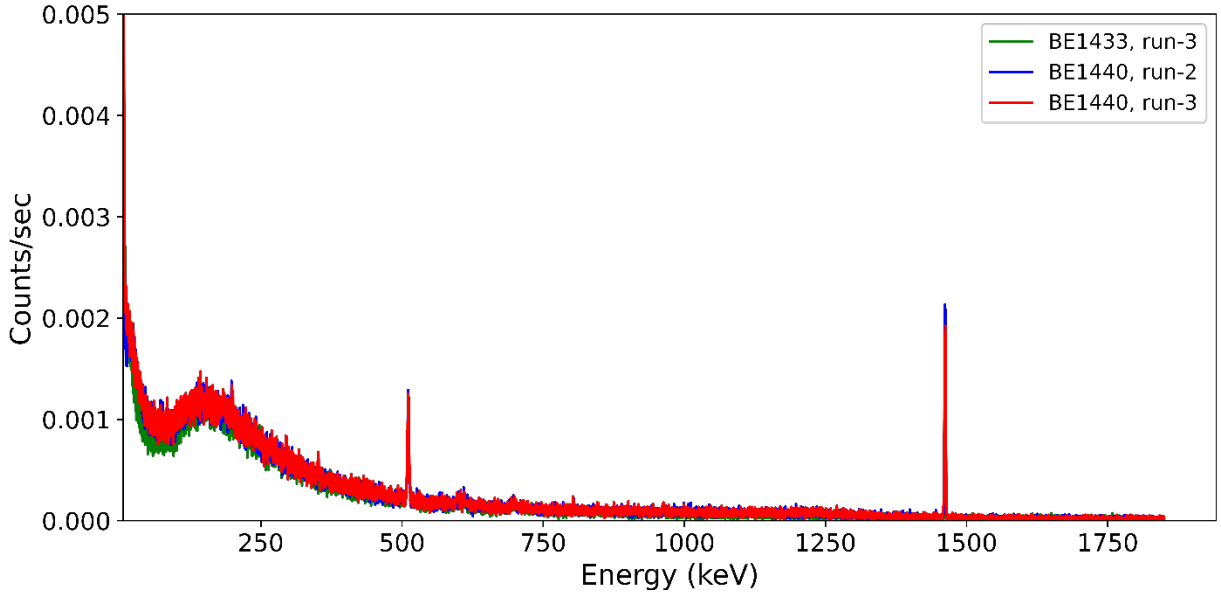


Figure 2: Delayed gamma emission detected during a calorimetry experiment with (red) PdD_x, (blue) PdH_x, and (green) with no sample.

Notably, we were able to perform *in-situ* electrochemical experiments with PdD_x and PdH_x using this gamma well detector in BE4056, thus greatly reducing background noise while providing *prompt* gamma detection. Despite the order of magnitude reduction in background counts, no form of isotopic contrast nor transmutation elements were observed. Further, we note that time-localized spikes in gamma-ray counts *were* observed in a few instances; for example, during a case in which one of our pressure-vessels failed, there was a corresponding increase in the total counts per second observed in the detector. While one may argue that the failure *could* be indicative of AHE driven during the depressurization, these results were able to be mimicked by us through other seismic

events and are observed during earthquakes or intentional motion of the detector itself. Again, the lack of isotopic contrast in *delayed* gamma experiments is suggestive of a lack of deuterium LENR.

Prompt gammas, neutrons, and electrons

In contrast to our HPGe detectors, the use of modern liquid scintillators allows for detection of *any* ionizing radiation above some energy threshold. The intensity of the scintillation signal provides a measure of the radiation's energy, while the shape of the scintillation signal provides a measure of the type of particle. In other words, to differentiate different particle types, one can rely on pulse-shape discrimination (PSD), where the signals are binned according to both their pulse shape as well as their energy. An example of this is shown below in Fig. 3; the contour provides a measure of the number of counts that are measured by the detector at a given intensity (energy) and pulse shape (PSD), allowing differentiation, in this example, of neutron and gamma rays over a wide energy spectrum.

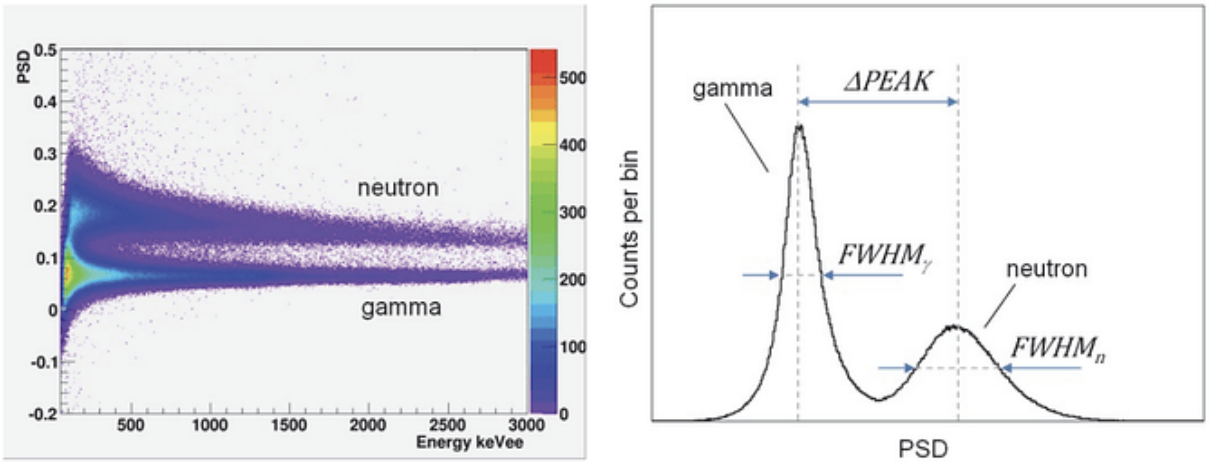


Figure 3: (Left) Contour of measured counts (color axis) for a given pulse shape (y-axis, PSD) and varying intensities/energy levels (x-axis). This allows for simultaneous measurements of different types of ionizing radiation. (Right) Comparison of a gamma-ray signal in contrast to a neutron-induced signal.

In one instance, we relied on these liquid scintillators to look for potential LENR events in our HPE calorimetry studies, where large voltage spikes were observed during electrochemical loading. An overview of this experiment is shown below in Fig. 4 for PdD_x. Despite large variations in the electrochemical characteristics over time, the liquid scintillators measure a constant number of events per second. To understand the sensitivity of these detectors to LENR reactions, we note that four detectors are surrounding the HPE vessel; each detector has a *conservative* estimate of 3.5% areal coverage. Given the Pd mass (50 mg) and near-unity loading, there are an estimated $3*10^{20}$ deuterons ($1.5*10^{20}$ potential reaction sites since 2 deuterons are

involved in the reaction). At our background counts per second of <50 cps, we can use these values to set a probability on the fusion rate for us to observe an event. If we assume that a signal of >100 cps is easily resolved, this suggests that we need 100 cps / $(0.035 * 1.5 * 10^{20}$ reaction sites) = $2 * 10^{-17}$ events *per deuteron pair* each second. In other words, these measurements would suggest that in ~10^{17} seconds we may observe events exceeding the background. This analysis procedure includes *any* ionizing radiation and is not discriminating based on pulse shape nor intensity.

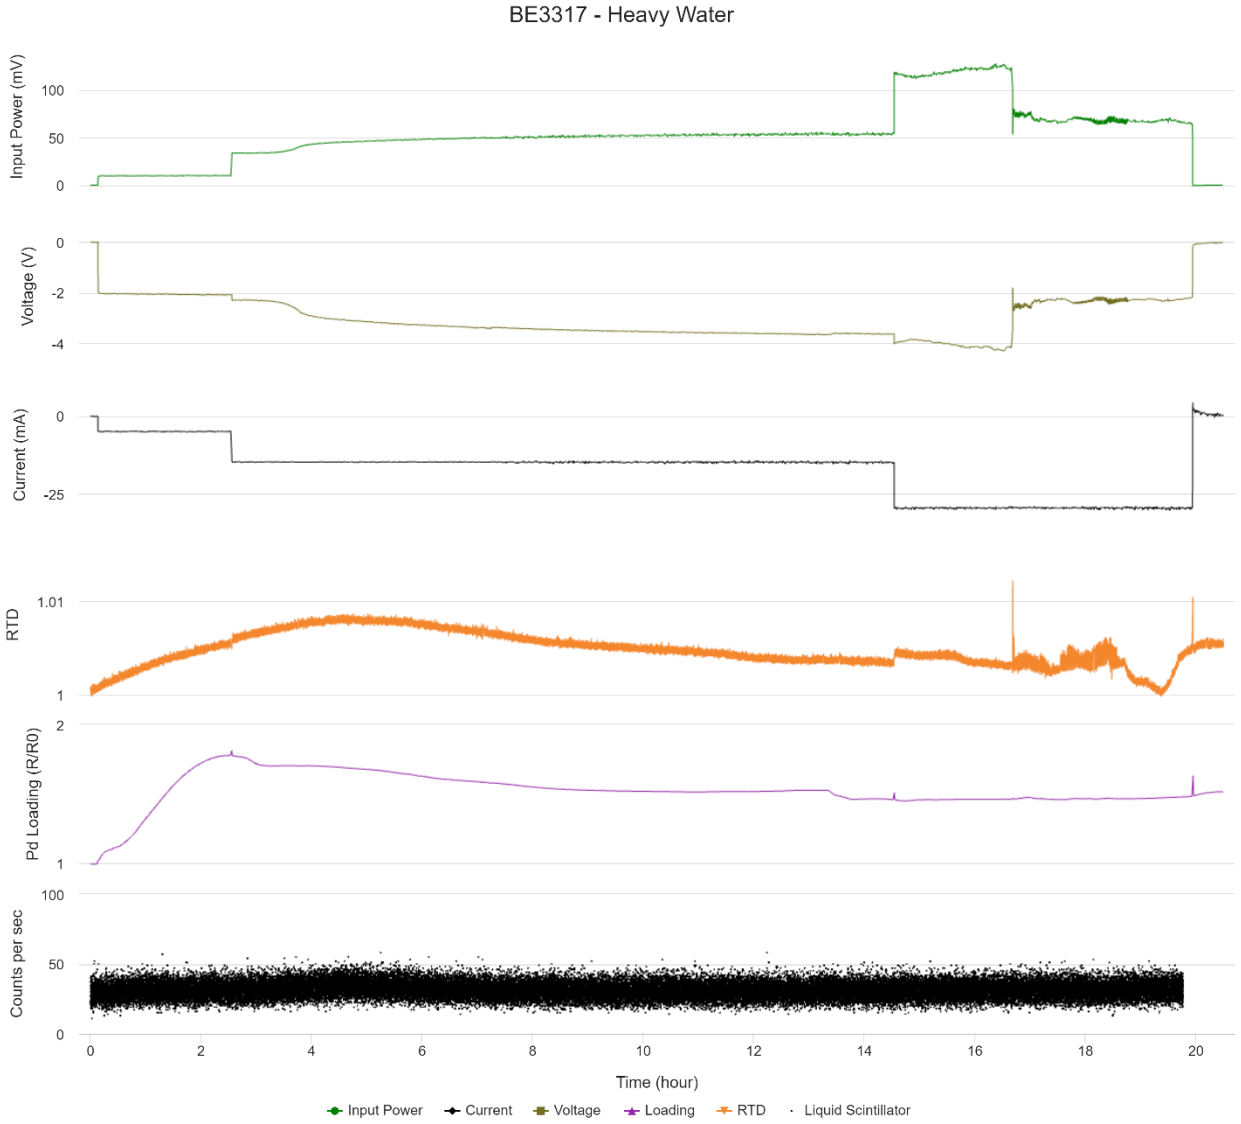


Figure 4: Overview of data from BE3317 in which 4x liquid scintillators were placed around an HPE vessel during various voltage instabilities/anomalies. The bottom plot provides the scintillator data of counts per second.

In other instances, we implemented arrays of these detectors in our USP campaign. In particular, we used multiple forms of nuclear detection to investigate potential D-D fusion events when stimulated with an ultrashort laser pulse at intensities in excess of 10^{16} W cm^{-2} – at these power densities, the material is effectively vaporized. These experiments are of interest due to evidence suggesting that D-D fusion can occur *in the absence* of a lattice with comparable laser intensities in molecular deuteron jets/clusters[98,99]. While we observed significant amounts of ionizing radiation being produced in PdD_x, increased counts were observed in PdH_x. Further, a number of alternate, non-loading materials exhibited even greater production of ionizing radiation at these laser intensities. Figure 5 displays the comparison of radiation produced across several materials in a vacuum environment, including "pre-processed" Pd that was stimulated in a gaseous environment before the radiation measurements[100]. This combination of effects clearly rules out AHE or D-D related reaction processes; additional nuclear diagnostics suggest this is a form of inverse Bremsstrahlung, where electrons are accelerated by the high intensity laser.

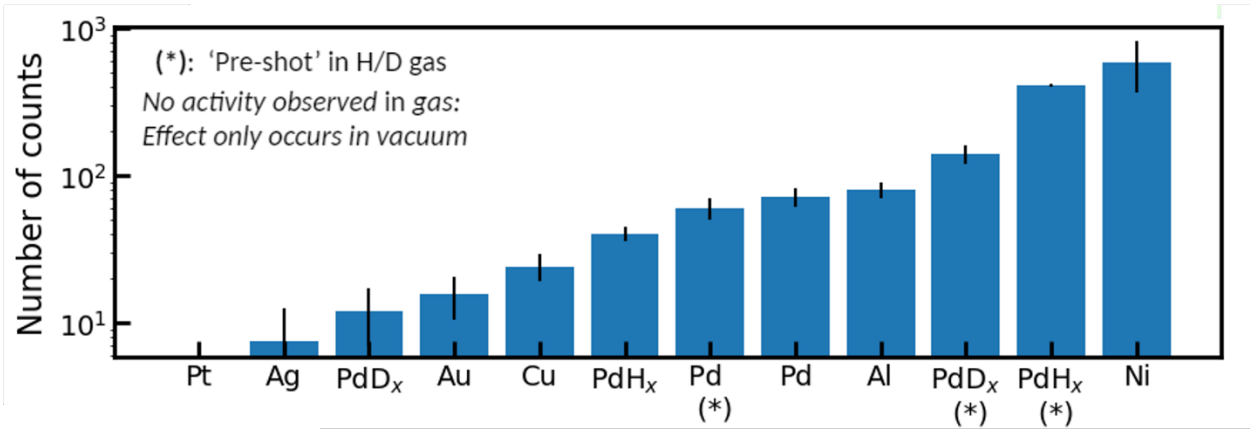


Figure 5: Relative amount of produced ionizing radiation with high intensity laser stimulation across different materials, both loaded and unloaded.

Summary

In summary, across multiple forms of nuclear detectors and implementation into a variety of stimulation and "staring" based experiments, we observe no evidence of energetic radiation from LENR processes or AHE-like effects. These observations provide lower limits on the probability of nuclear processes occurring and producing conventional byproducts; this lower limit suggests the likelihood of observing nuclear reactions on a meaningful timescale is effectively null.

[98] Ditmire *et al., Nature* **398**, 491 (1999).
[99] Grillon *et al., Physical Review Letters* **89**, 065005 (2002).
[100] We observe *no* produced radiation with a gaseous atmosphere; vacuum is required to observe this effect. This is the result of nonlinear optical effects that occur due to the high intensities; the beam is effectively absorbed *by* the gas, reducing the intensity on-target.

Conclusions

I : We did not see any evidence of LENR.

To be specific, we performed a suite of aggressive, yet very different, LENR campaigns trying to detect AHE events. None of them succeeded in doing so.

Loading: Our hardware operated at D_2 pressures up to 500 ksi pressures, allowing us to reach 100% PdD loading in many situations. We used this capability in a wide range of LENR experiments; with targets ranging from slugs to films to chips, stimulations from thermal to electric-field to ultra-short laser pulses, and detection methods from calorimeters to RTDs to magnetic and physical archaeology. No LENR was detected.

Thermal Ignition: We tested highly loaded PdD and PdH targets in two very intense thermal stimulation situations; at temperatures as high as 1200°C (where the palladium hydride melted). In one case, TI-11, the targets were in thermal equilibrium, sacrificing some loading in exchange for long duration. The other case, TI-12 was more dynamic; loadings would start at well over 95% (despite these extreme temperatures) but then rapidly decay over time as the material deloaded. From a material standpoint, these were macroscopic bulk targets (not thin films), and experienced a huge range of lattice deformations, from near perfect crystals to being literally ripped apart. No LENR was detected.

Ultra-Short Pulse: This was used to test a unique, non-thermal, form of stimulation. We loaded Pd targets with D_2 at pressures up to 450 ksi, and irradiated them with very short (~ 35 fsec) laser pulses. This was done both with isolated single pulses, sequences of temporally separated pulses, and with correlated bursts of pulses. We used both pump-probe reflectivity measurements and nanofilm archaeology to analyze results of the stimulation. No LENR was detected.

Staring: This was a two-pronged campaign using laser irradiation as a thermal form of stimulation for 150 ksi-loaded targets. In one effort, we literally stared at the target for long times, determining the temperature evolution of thousands of separate spatial pixels by recording thousands of IR images of the target. In the other effort, we used the laser as a stimulation source, irradiating both Pd and Pd-alloy targets, and using both magnetic and nanopillar/nanofilm archaeology to determine results. No LENR was detected.

Chips: The purpose of the Chips campaign was to test the response of small, lithographically made, targets to electrical stimulation; applied either by resistive heating or by intense electric fields. The smallness of the targets decreased the mass which would have to be heated-up by LENR energy, thereby increasing the detection sensitivity. In some experiments, we used small on-chip RTD thermometers to measure the response of each pulse in multi-thousand-pulse sequences. In other experiments, the chips contained arrays

of nanopillars or nanofilms for single-event archaeological detection. No LENR was detected.

High Pressure Electrolysis: We used our high pressure systems in concert with electrolysis to achieve higher and more stable loadings (98%) than achievable in usual electrolytic LENR experiments. With this capability, we ran a number of LENR experiments both in and out of our commercial calorimeters. No LENR was detected.

Replications: We attempted to replicate two of the most influential gas-loaded experiments; those of Arata and Case. We performed these studies both in similar apparatus to that which they had used, as well as in our commercial calorimeters. No LENR was detected.

Calorimetry: Here we used high-sensitivity commercial calorimeters to look for LENR-induced energy from a variety of PdH and PdD targets and situations. We tested nanoparticle and aerogel targets as well as macroscopic foil and wire ones. In most experiments, the targets were gas loaded, but we did perform ones based on high-pressure electrolytic loading. Experiments were run at cryogenic temperatures, ambient ones, and temperatures up to 200ºC. Our most sensitive experiments were passive, with no stimulation energy to account for; but, given the probable importance of stimulation, most experiments used monitored amounts of resistive heating. No LENR was detected.

Physical Archaeology: We developed two types of nanostructures, arrays of nanopillars and nanofilms, to serve as LENR sensors with single-event sensitivity. Nanofilms were used as detectors in many different experimental campaigns under a wide range of loading conditions (150-450 ksi D_2 gas, D_2O electrolysis, cryogenic loading) and stimulations (lasers, ultra-short pulse lasers, resistance heating, electric fields), and with materials using Pd as well as several Pd alloys. No LENR was detected.

Magnetic Archaeology: We developed a new type of ultra-fast thermal history detector based on the permanent change, induced in a magnetic material, by the fast thermal pulse from a nearby AHE event. We developed arrays of nanoscale ferromagnetic dots to detect thermal changes in D_2-loaded Pd films. The pre and post experimental state of the targets was measured with nanometric resolution by magnetic force microscopy. No LENR was detected.

Nuclear Detection: We searched for nuclear products in two ways, looking for energetic products emitted during experiments, and for delayed γ-rays generated after them. No LENR was detected.

II : We have concluded that LENRs[101] do not exist

Obviously, this is a belief, not a proven fact. It's very hard to prove a negative, and we haven't done so.

Nonetheless, it's a strong belief, and a conclusion which we're confident in.

Here are the reasons why:

1. **The breadth and quality of our search.**

   Our search has thoroughly covered the three main conditions which seem crucial to enable LENR; loading, stimulation, and material. That search, much of which was performed with the ability to detect even single LENR events, did not find any LENRs.

   Loading: High loading is believed by most of the LENR community to be an essential precondition to allow LENRs to occur. It is also a major challenge to LENR researchers, and continues to be despite 35 years of effort; very few can routinely achieve loadings above 90%.

   In contrast, we can always, *every single time*, reach loadings over 95%. And, using two different methods (cryogenic temperatures, or ultra-high pressure gas) we can reliably reach full, 100%, loading.

   Stimulation: It's not clear whether active stimulation is essential for LENR. Certainly, it is simplest to perform experiments without applying any. We have done so, with 95-100% loading, using both perfect and imperfect PdD lattices, and while using very sensitive detection (passive calorimetry and single-event archaeology). None of these experiments detected LENR.

   Because of this, we have considered active stimulation to be important.

   The simplest form of stimulation, and the most commonly used in the field, is thermal. This presents a plausible mechanism both to commence LENR (increased hopping of deuterons out of their individual octahedral sites) and to propagate the occurrence of LENRs across a target. We have run experiments heated up to the melting temperature; which is the maximum amount of thermal stimulation possible in a PdD lattice. We have also done so, holding at these maximum values, over a wide range of time durations, from nanoseconds to seconds. None of these experiments detected LENR.

[101] To be specific, the flavor of LENR we're talking about here, is a nuclear reaction between two deuterons in a (largely) Palladium lattice, producing virtually all of its 24 MeV reaction energy as heat in the nearby lattice, not as the energetic particles seen in conventional D-D fusion.

We have also used a variety of fundamentally different forms of stimulation; from ultrashort-pulse lasers (producing nonthermal excitation), to plasmon excitation, to resonant terahertz radiation (producing optical phonons), and to intense electrical fields. None of these experiments detected LENR.

Material: It is reasonable (and often suggested) that the proper lattice material might be essential for LENR; either due to its composition or its structure. The extreme case here is perfection, both perfect composition, i.e., pure Pd (or at least 99.999%) and perfect structure (single crystal). We learned how to maintain such perfection even with high loading, but saw no LENR.

While there is only one form of a perfect material, there are near-infinite varieties of imperfect ones. We obviously have not tried them all.

We have performed LENR experiments with Pd nanomaterials, both aerogels and nanoparticles, as well as more structured forms such as nanopillars and nanofilms. There is a belief by some researchers that large numbers of dislocations, or other forms of highly strained material are conducive to LENR. Our Thermal Ignition samples literally tear themselves apart after transient heating; in the process, these undergo a full gamut of intense and chaotic deformations. None of this suite of structural variations resulted in detectable LENR.

There are also a huge number of ways to field impure palladium, depending on the type of impurities and their concentrations. We have only explicitly experimented with a few of these, having made and tested nanofilms using several different Pd alloys. Despite the use of nanofilm archaeology, none of these demonstrated LENRs. It should also be pointed out that the single-event sensitivity of nanofilm archaeology provides a strong implicit test of impurities. Each nanofilm contains over 10^{11} palladium atoms and a 9 pad experiment has 10^{14}; no matter how pure the palladium is, it is virtually guaranteed that every possible impurity exists, and that many of the dual impurity combinations do as well. If the presence of impurity atoms within the palladium lattice causes even a single nearby deuteron-deuteron LENR, the nanofilm will reveal it; none were seen.

Detection: Combining the proper loading, stimulation, and material to manage to generate LENRs appears to be quite difficult; doing so, yet failing to detect them would be extremely regrettable.

Accordingly, we have always emphasized high sensitivity detection. This began with commercial differential calorimeters, and then led to high temperature thermocouples, lithographic RTDs, femtosecond pump-probe detectors, and piezoelectric seismology. This effort culminated with our development of archaeological single-event detectors. None of these detected LENR.

Our inability to detect LENR, despite such an extensive and sensitive search over the LENR landscape, strongly contributes to our belief that LENRs simply do not exist.

2. We strongly doubt that other researchers have actually detected LENR

- If other researchers did see LENRs, we would have also.

  Over the past 35 years, many researchers have reported detecting LENRs. Most of these experiments were based on electrolytic loading, not our high-pressure gas loading approach.

  We have not tried to replicate such electrolytic experiments (in multiple cases, the Callisto electrochemical loading project *has* attempted to), and can certainly not *prove* that they failed, nor offer informed commentary as to why they might have failed.

  Nonetheless, we believe they did fail; for the reason given above: If they had seen LENR, we should have also.

  The reason for this, is that their work takes place in the same physical environment (specified by loading, stimulation, material) that we have explored. And, for the most part, they are in "vanilla" portions of this parameter space; ones that our search has *very* thoroughly covered.

  Loading: While many LENR reports claim that high loading is important, none of them involve loadings as high as we reach!

  Stimulation: Many LENR reports were from experiments with no active stimulation; just that inherent at normal temperatures (< 100^{o}C), or in the course of electrolysis. Of the active stimulations that have led to reports of LENR, most were one or another form of heating. However, once well above picosecond timescales, all forms of thermal stimulation are fundamentally the same. We have thoroughly covered this landscape, and to the maximum intensity possible!

  There are some reports of LENR in which specialized types of stimulation were used (e.g., terahertz beat frequencies, charged particle impacts, deuterium fluxes) but none of these were involved in the pioneering and primary reports which the LENR field is based on.

  Material: The primary LENR reports use the same basic material that we have, high purity palladium. In terms of structure, there is no consensus as to what form is needed, reports of what "works" are all over the map. However, the given recipes are all global ones (e.g., grain sizes, annealing conditions, surface roughness, nanometric crevices); the local structures are uncontrolled and stochastic. While details of local structures are bound to be unique, all of these types of structural deformations have occurred in our experiments.

LENR, if it exists, is not "hiding". The parameter spaces in which it has been reported to exist are well within those which we have covered. The fact that we have not seen LENR there, is another reason casting doubt on its existence.

- Another reason for doubt, is the failure of the direct replication experiments that we have performed.

  Rather than merely searching in the same basic parameter space as other researchers, we have, in a few cases, tried to search with the identical conditions. This has not been our primary focus, both because relatively few of the LENR reports have utilized gas loading, and because none of these "successful" experiments have been in our high-pressure, high-loading regime. But, as discussed previously in this report, we have tried to replicate two well-known results; low-pressure gas loading experiments by Case and by Arata. In neither case did we see signs of LENR, as they had reported.

  We also attempted to understand the high-energy event seen in the 1989 LLNL experiment. This was an ambient pressure electrolytic experiment, and hence Callisto electrochemical loading has been leading efforts to detect LENR during replication attempts. Our focus was different; we were interested to learn if the experiment could be "mistreated" in ways so as to produce results similar to those originally seen. We found that there were indeed excitation conditions which could produce qualitatively similar behavior (due to its preliminary nature, no quantitative data was obtained in 1989). Whether or not the 1989 behavior resulted from similar mistreatment is, unfortunately, probably unknowable.

- There are no reliable, reproducible, LENR experiments.

  This is clear; if anyone could make LENR succeed, on-demand, the "debate" would be over!

  The fact that, after almost 35 years of research, the sporadic reports of successful LENR are only that, and have not led to reproducible results, is hard to ignore.

3. We don't theoretically understand how LENRs can occur

On the face of it, two aspects of LENRs are implausible; how two deuterons can cross the coulomb barrier between them and get close enough to undergo a nuclear reaction, and, if they do, why the reaction products are just heat and He^4, not the energetic particles produced from conventional deuterium fusion.

These are not new concerns; they have been major issues since the initial announcement of LENR in 1989. The seeming impossibility of either or both of these results has always been one of the main reasons why LENR research was shunned by the mainstream physics community.

Obviously, we (both our sponsor and the Callisto researchers) were well aware of this problem when we commenced work on Callisto. However, there were (and always have been) a variety of theoretical suggestions as to how these things could take place. And, more fundamentally, there's the fact that in science, experiments trump theory!

However, as time has gone by, our experiments have shown no indication of LENR. Nor have the various approaches proposed to explain LENRs advanced beyond the conceptual stage to become fully developed theories, let alone compelling ones.

Accordingly, this lack of an understandable mechanism continues to nag, and remains a significant reason for us to doubt the existence of LENRs.

Why might our conclusion rejecting LENR be wrong?

We've explained why we do not believe that LENRs exist. However, we've certainly not *proven* this to be the case; our belief might be wrong. Here are some possible reasons why:

Other researchers do claim to have seen LENR: If any of them are correct, our conclusion is clearly wrong!

The burden of proof required for them to demonstrate LENR is not particularly onerous, just deliver convincing and *reproducible* experimental results! To date, 35 years in, no one has managed to do so.

Our search is incomplete: We've searched a very wide swath of the loading, stimulation, and material parameter space underlying LENR. However, we haven't searched it all. In particular, while we've reached a wide range of individual loading, stimulation, and material conditions, we haven't satisfied all the simultaneous combinations of these. So, there are portions of parameter space which we haven't explored; it is certainly possible that LENRs only exist in one of these regions.

We consider our search over loading and over stimulation conditions to be definitive.

No higher loading has been, or can be, achieved; our search here is complete.

There are certainly different types of stimulation that can be tried: But, the most common type, thermal stimulation has been fully explored; different implementation details may be tried, but these won't change the fundamental thermal effects. There are still types of nonthermal stimulation which we have not explored, but we don't consider those we know of to be significant; some we don't believe to be physically effective, while others aren't relevant to LENR.

It's our search over materials that is the least complete. While we have experimented with perfect materials, and many types of imperfect ones, there are certainly many other instances of imperfect materials which we have not investigated.

We doubt that these will lead to LENR, but certainly cannot prove this.

LENR may require a "magic" material: One explanation for why, despite our broad search over parameter space, we might not have seen LENR, is that a "magic" material (in composition and/or structure) is required.

This could explain why some researchers are actually correct in claiming to have seen LENR. Given imperfect knowledge of, and control over, microscopic material properties (a particular issue in electrolytic experiments) this could also explain the irreproducibility of LENR results, even in experiments performed by the same researcher and conditions. I.e., sometimes, the proper material exists, and LENR occurs, but usually it does not.

However, why a special form of PdD would be essential, and how it would lead to LENR, is, to be diplomatic, unknown.

That one is required is certainly a possibility, and impossible to discount. But, meaningful support for this proposition will only come if someone actually proffers up such a material.

Electrolysis might be special: Most of Callisto's research has been based on loading palladium via high pressure D_2 gas. In contrast, most of the outside LENR research (and particularly the early pioneering work that initiated the field), uses D_2O electrolysis to load deuterium into palladium. It's possible that there is something fundamentally different in these two processes; LENR may only be possible for electrolysis, and never for gas loading.

If LENRs occur within the interior of the PdD lattice, the difference between gas and electrolytic loading shouldn't matter; the local deuterium ions only "know" their local environment. However, it is possible that LENRs are a purely surface phenomena, occurring only within the first few layers of the lattice. If this is so, then electrolytic loading is fundamentally different from gas loading, and this may well explain why none of our gas-based experiments have detected LENR.

A note, however; Callisto has performed a series of high-pressure electrolytically loaded LENR experiments. These have achieved very high and stable loadings (~ 98%), yet have failed to detect LENR. (And, of course, our sister program, Callisto electrochemical loading, is exclusively focused upon electrolytic loading; they will report on these separately.)

Nanofilms may not work: While nanofilm archaeology is just one of our approaches to detection of LENR, it is our most sensitive one. The fact that it provides the ability to detect even single LENR events, and can be utilized in a broad range of different LENR experiments, has naturally made the failure to detect LENR in any of its applications a very significant component of our belief that LENR does not occur.

This naturally leads to a question; is there any systematic flaw in nanofilm archaeology which makes this an invalid diagnostic tool? Here are two possibilities which could:

We have focused on "heat – helium" LENRs, defining them as deuterium-deuterium reactions in which 24 MeV of reaction energy is transferred as heat to the nearby palladium lattice. However, note that if LENR energy appeared as energetic particles (as in a conventional deuterium-deuterium reaction), these would escape from the thin nanofilms, depositing only a trivial amount of energy in them; nanopillars would not break, and nanofilms would not develop holes. Alternatively, if the reaction energy does indeed appear only as lattice vibrations, but is inherently spread over huge ranges (e.g., thousands of lattice spacings) then it would be too diffuse to damage nanofilms or nanopillars; archaeology would not work.

Another concern would be if the lattice of a nanopillar or nanofilm simply can't support LENRs. For instance, if LENR involves some form of collective behavior, nanostructures may be too thin or have too few atoms to allow LENRs to occur. More conventionally, the palladium lattices of these nanostructures may be too deformed or disordered to host LENRs. Since no one actually knows how LENRs work, neither of these concerns can be dismissed out of hand.

We may have been looking for the wrong type of LENRs: We have focused our search on a specific form of LENRs; the reaction of two deuterons in a palladium lattice, producing heat and He^4. If LENRs instead involve a different reaction, then our search has been misdirected, and may well miss seeing them.

Several possible scenarios illustrate the problems. Obviously, if neither palladium or deuterium are involved, our search would fail. But, involvement of protium as a reactant would also be problematic, since we often use H_2 experiments as a null control.

Even if the LENR reaction involves two deuterons in palladium, it might work fundamentally differently than what we've been assuming. We've mentioned, for instance, the fact that reaction energy will not be retained in thin structures if the energy is delivered as energetic particles; this would affect not just nanofilms and nanopillars, but also magnetothermal detectors or chip-based systems.

Another issue, would be if LENR generation required some form of collective quantum interaction (e.g., Bose-Einstein condensation). Such suggestions have certainly been made in the past, and if correct, would probably require considerably different experimental conditions than we've been using. However, beyond qualitative suggestions, there is no developed theory or experimental evidence supporting this possibility; certainly, none of our cryogenic experiments detected LENRs.

Recommendations

- In October 2023, Callisto recommended that its LENR search be terminated.

  This was done, effective November 1, 2023

- Given our work, we recommend against future LENR investigations.
  Nevertheless, if future efforts are desired, we propose monitoring the field, waiting for *strong* evidence for LENR to appear before deciding to resume efforts.

  This evidence might constitute a compelling theoretical explanation of LENR. In this case new work should be scientific, to experimentally demonstrate LENR.

  This evidence might involve convincing experimental demonstrations, by others, of LENR. In this case the first focus would be replication. But afterwards, the primary effort should be the engineering exploitation of this advance.

Appreciation

We wish to express our enormous appreciation for the extraordinary treasure and personal attention which the sponsor has devoted to the Callisto effort.

We greatly regret that we were not able to successfully detect LENR.