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.

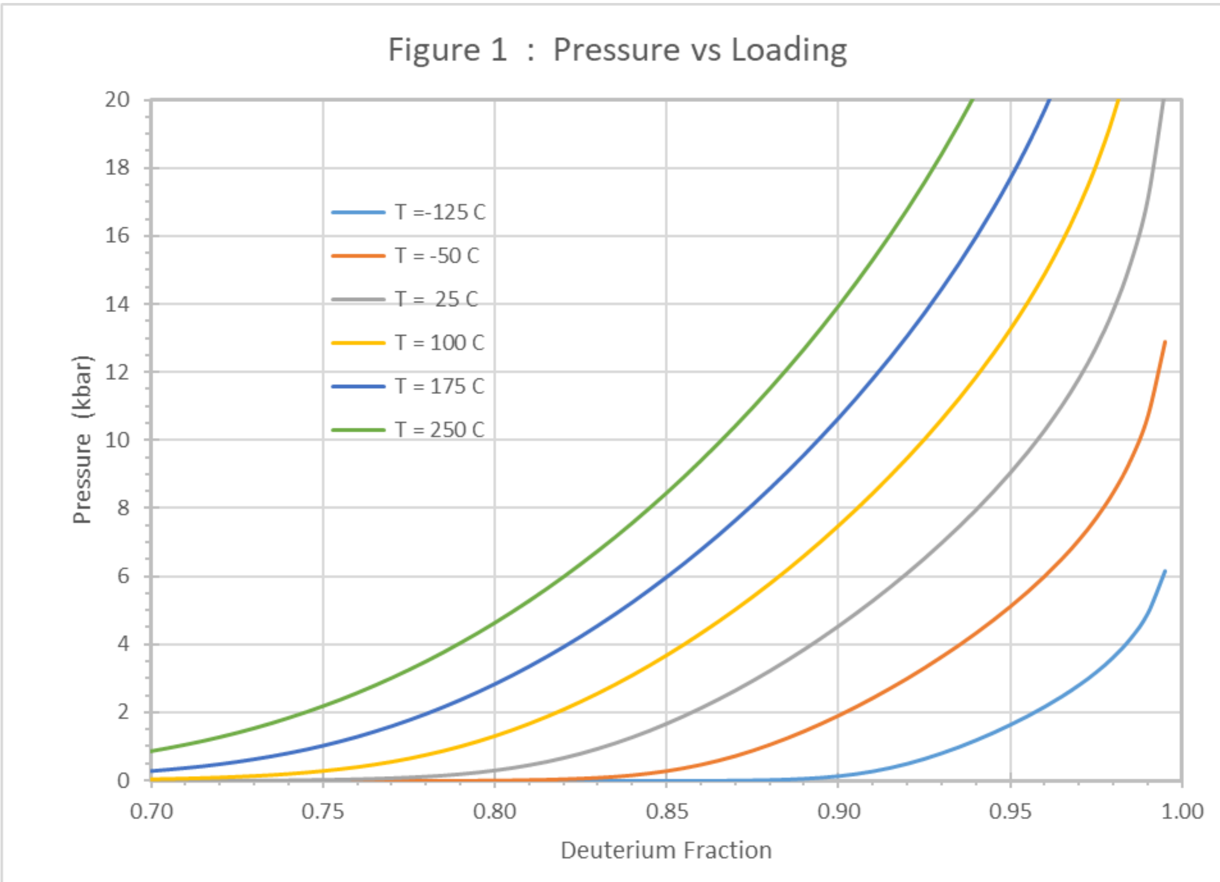

Later in this report, we will explain how we have been able to reach and sustain H_2 and D_2 gases at pressures above 100 ksi. For now, let's just point out that we learned to do so routinely (performing several hundred LENR experiments at these pressures) and with a perfect safety record (never any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