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

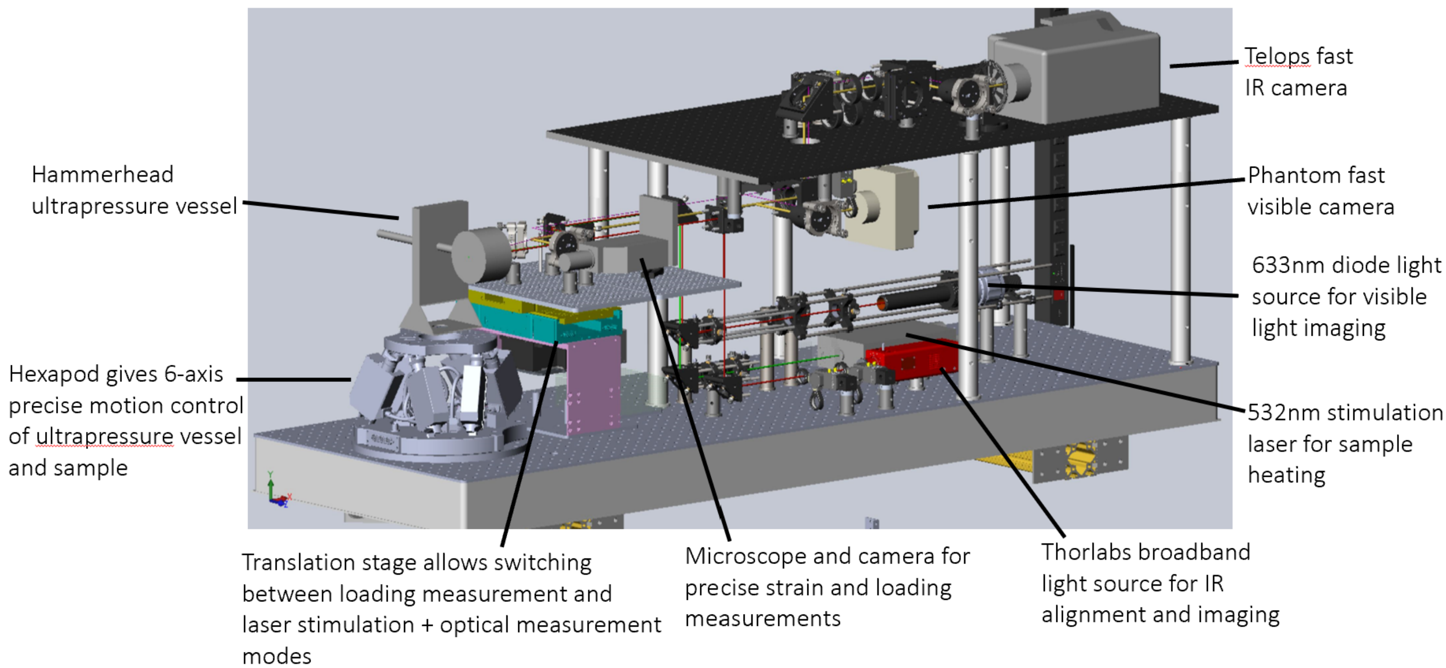

The hammerhead vessel could be configured with an ultrapressure sapphire window on one side for optical access to the sample, and electrical feedthroughs on the other side. Our primary use for the electrical feedthroughs was for RTDs on the surface of magnetic archaeology calibration samples, which were 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.

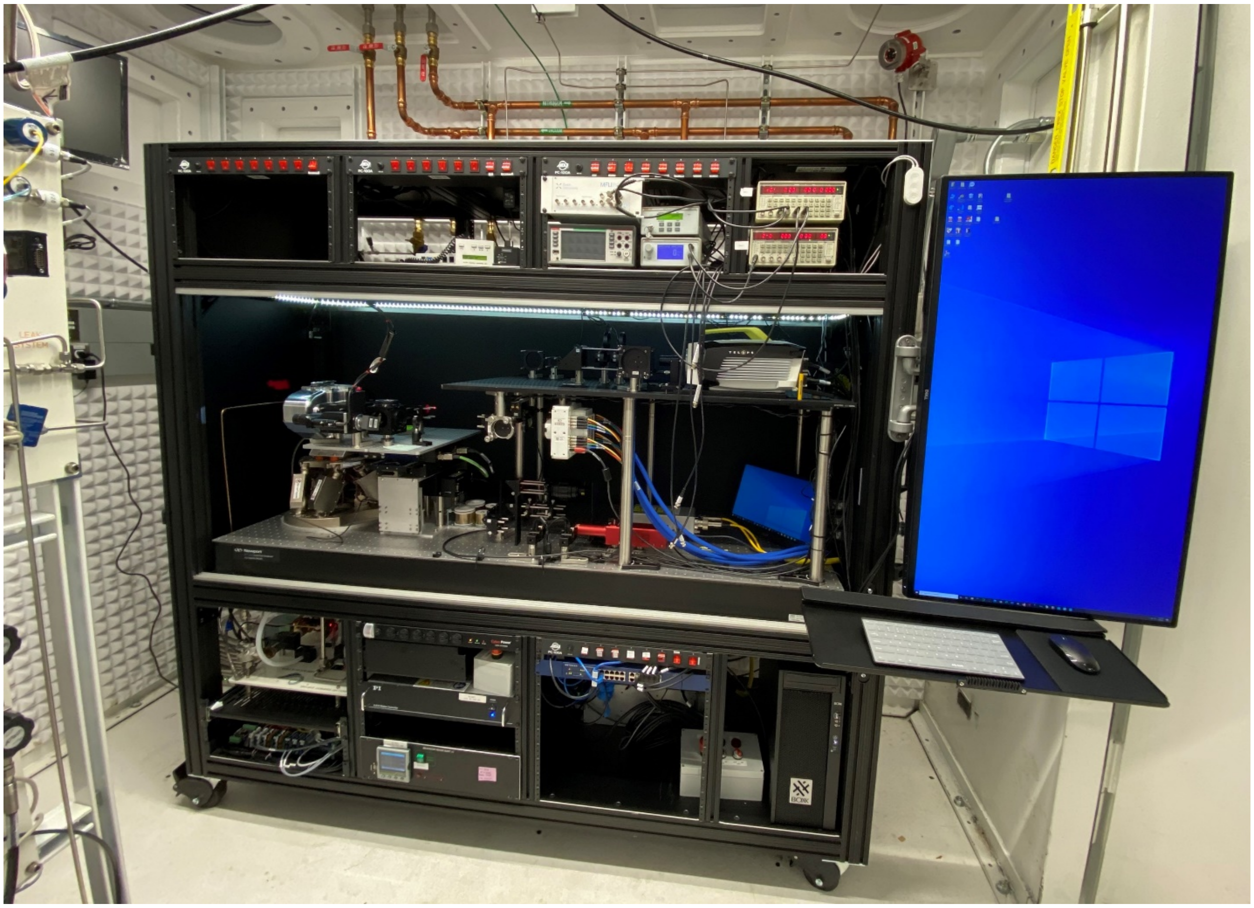

Figure 2 - A photo of the full Staring experimental apparatus.

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.

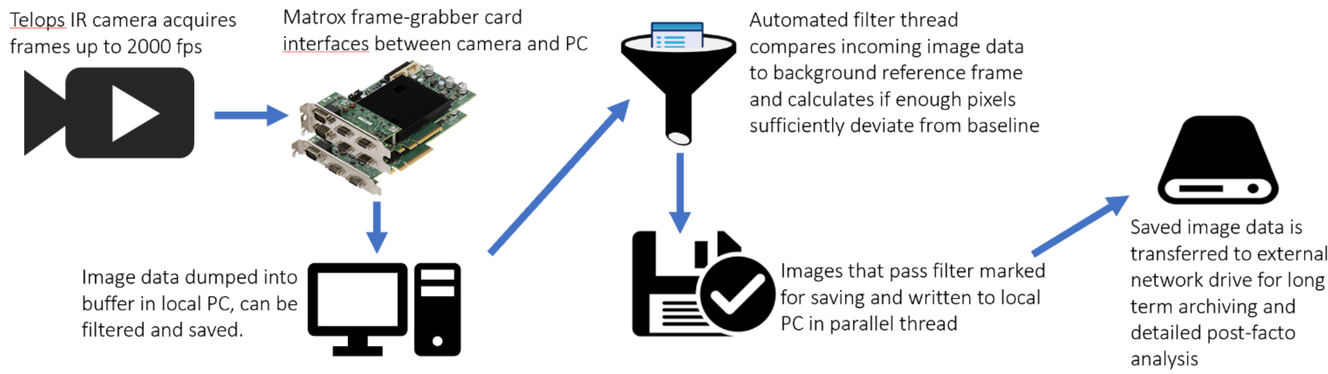

Figure 3 – A schematic of the long-duration thermal imaging data acquisition and processing pipeline.

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.

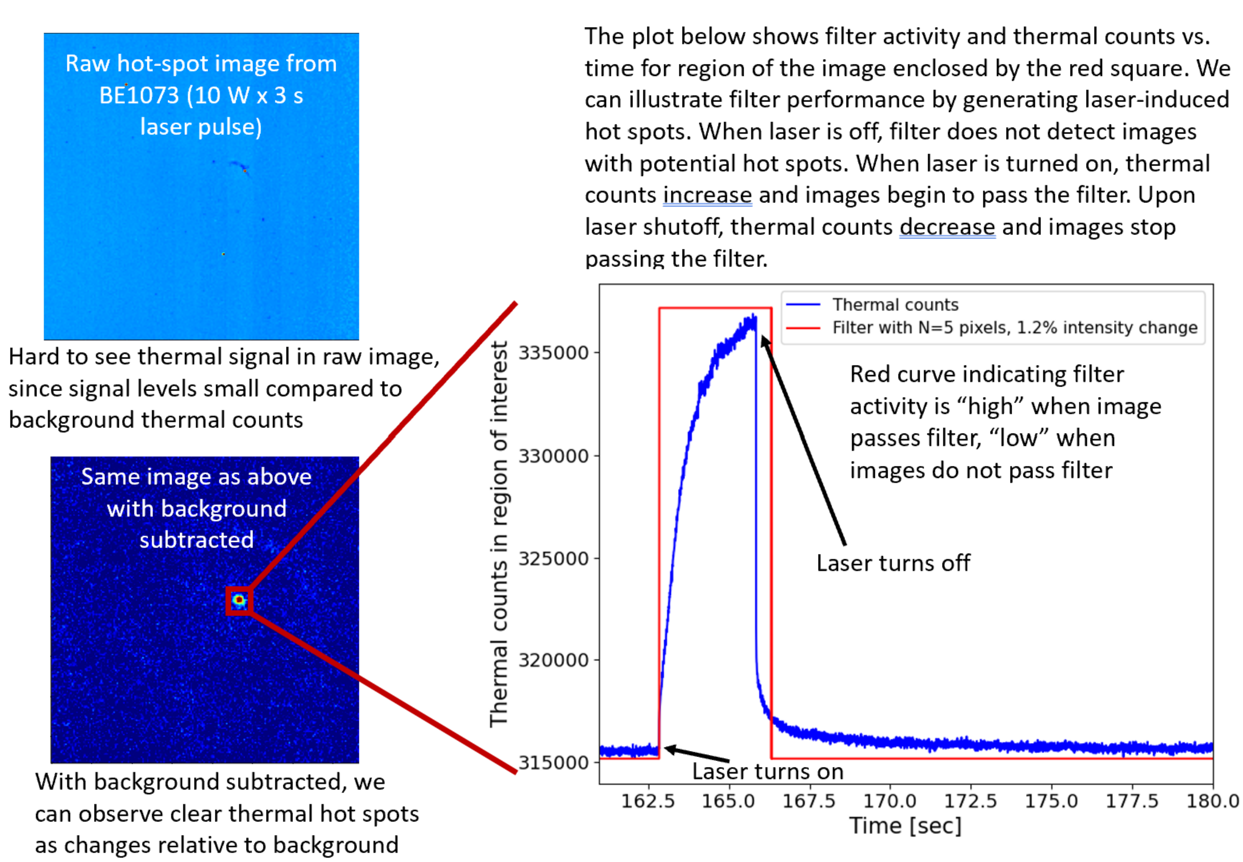

Figure 4 – An illustration of background subtraction and real-time filter behavior.

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 SiO2 on top of 500 μm of Si. SiO2 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.

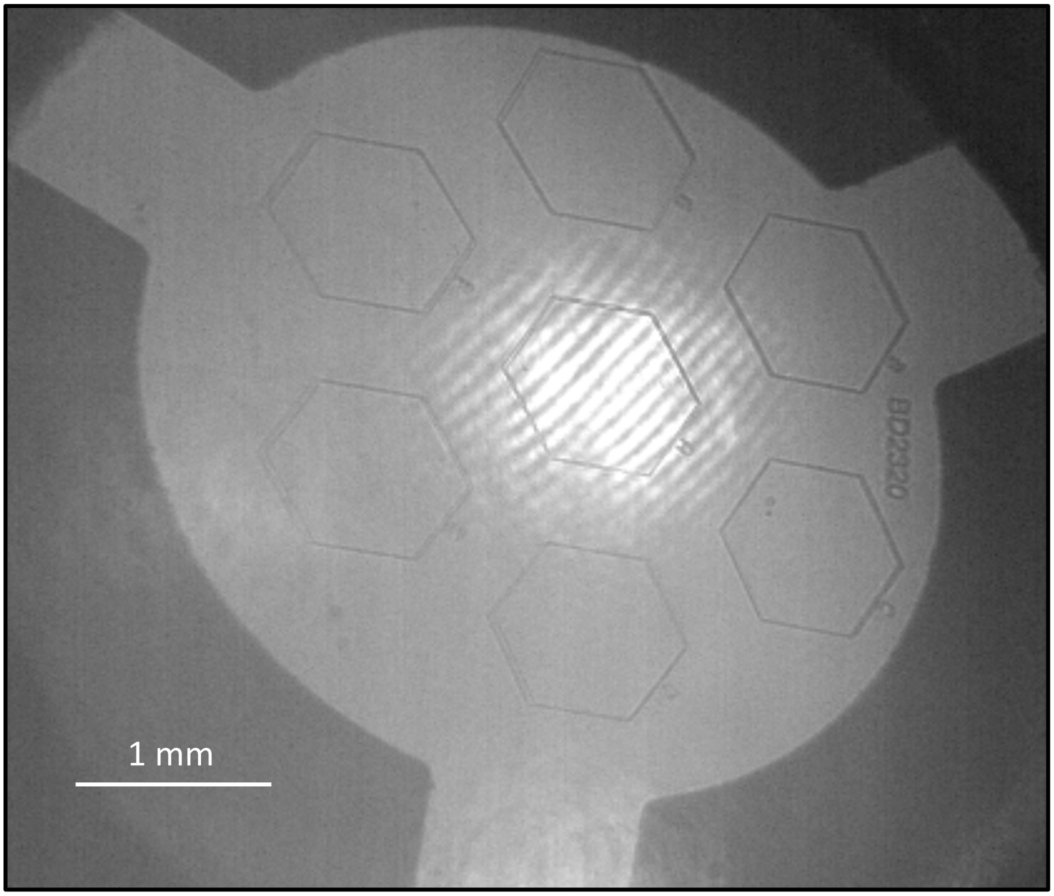

Figure 5 – A visible-light camera image of sample BD2320 (20nm Pd thin film) loaded under 145ksi D_2. The stimulation laser is aligned onto the center hexagon.

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).

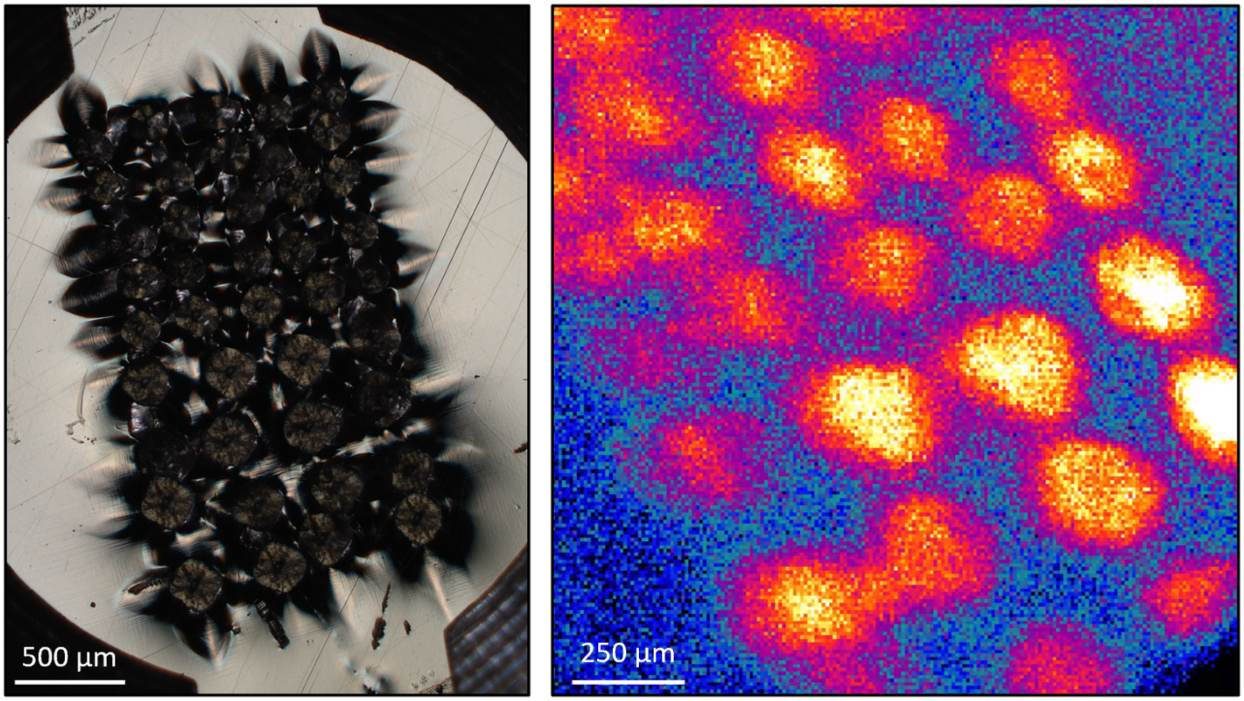

Figure 6 – Ex-situ visible-light camera image (left) and in-situ Telops IR thermal image (right) of peened Pd slug BD1263 used in BE3181 (sample under 100psi argon). Indentations from peening have a higher emissivity than the undeformed, flat crystal surface and thus appear bright in the thermal image.

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.

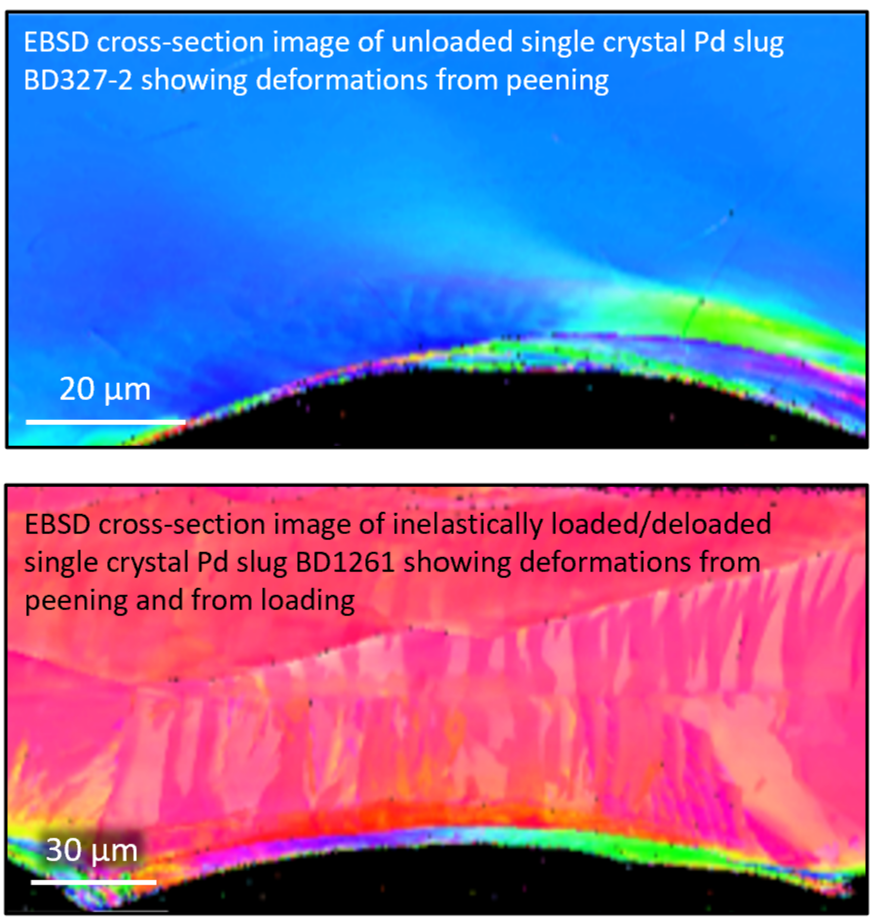

Figure 7 – EBSD on cross sections of single crystal Pd slugs with and without inelastic loading.

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.

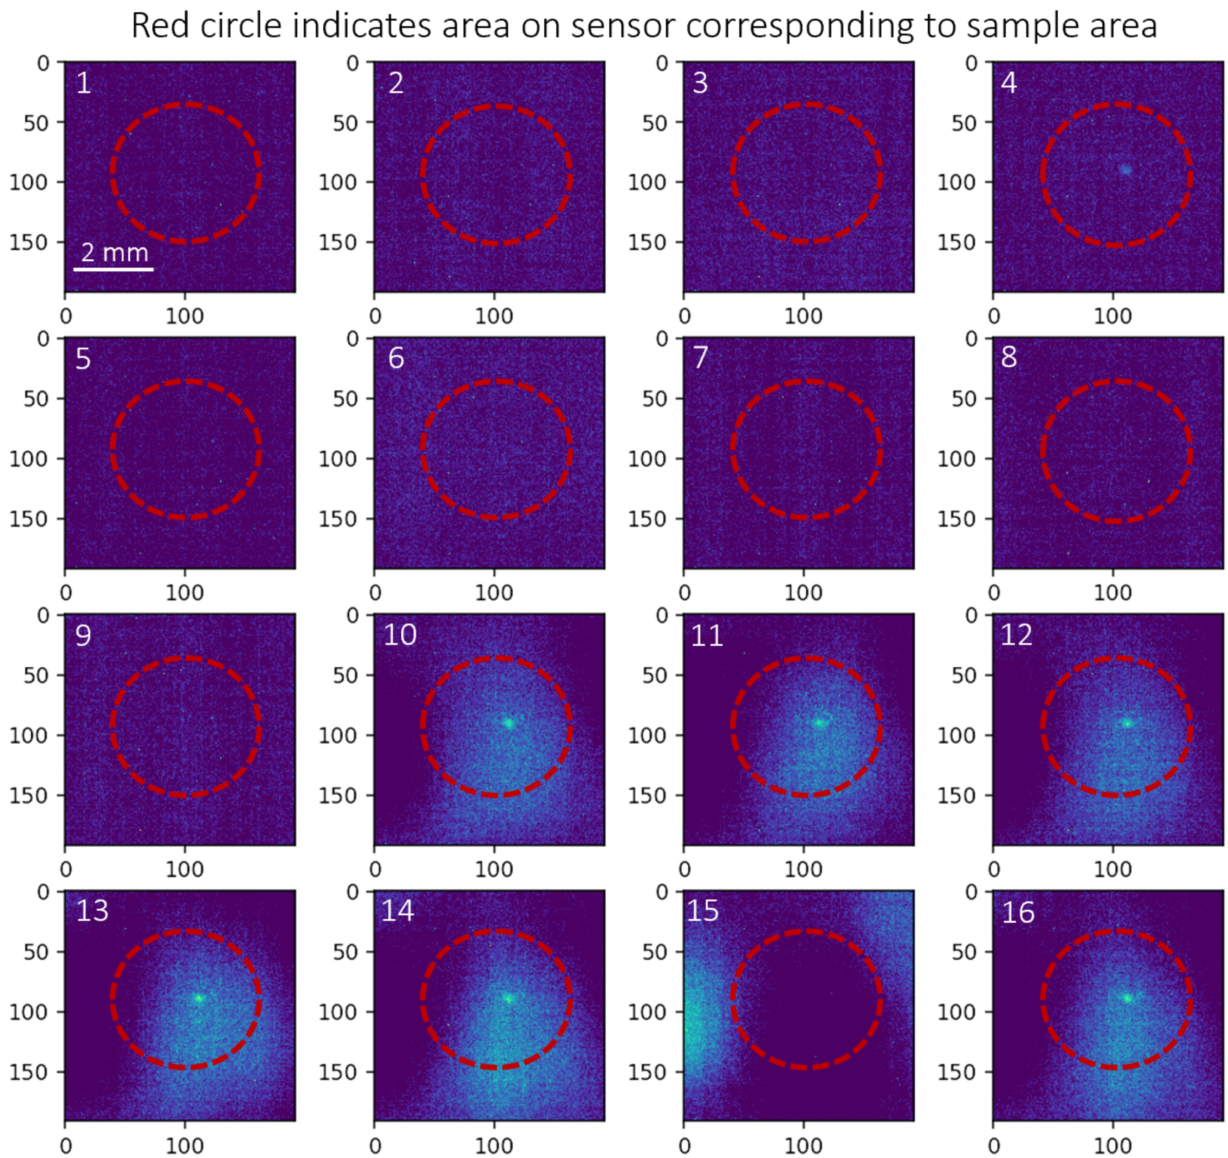

Figure 8 – Representative IR images that passed filtering in BE2131 and were flagged for saving.

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 SiO_2 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.

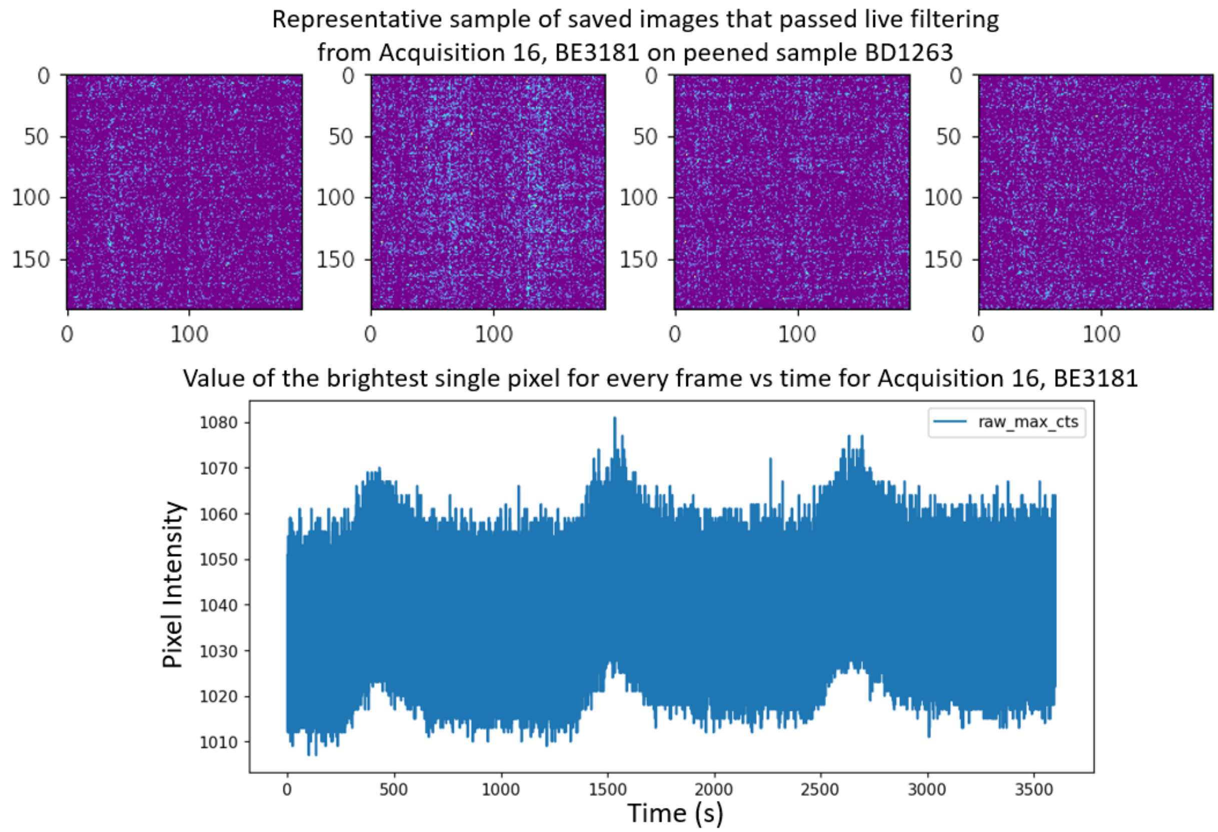

Figure 9 – Some representative data from Acquisition 16 of B3181. The top panels show some representative saved images that passed live filtering, and the bottom panel shows the values vs time for the single brightest pixel for each individual frame over the 1-hour time period of the acquisition.

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.

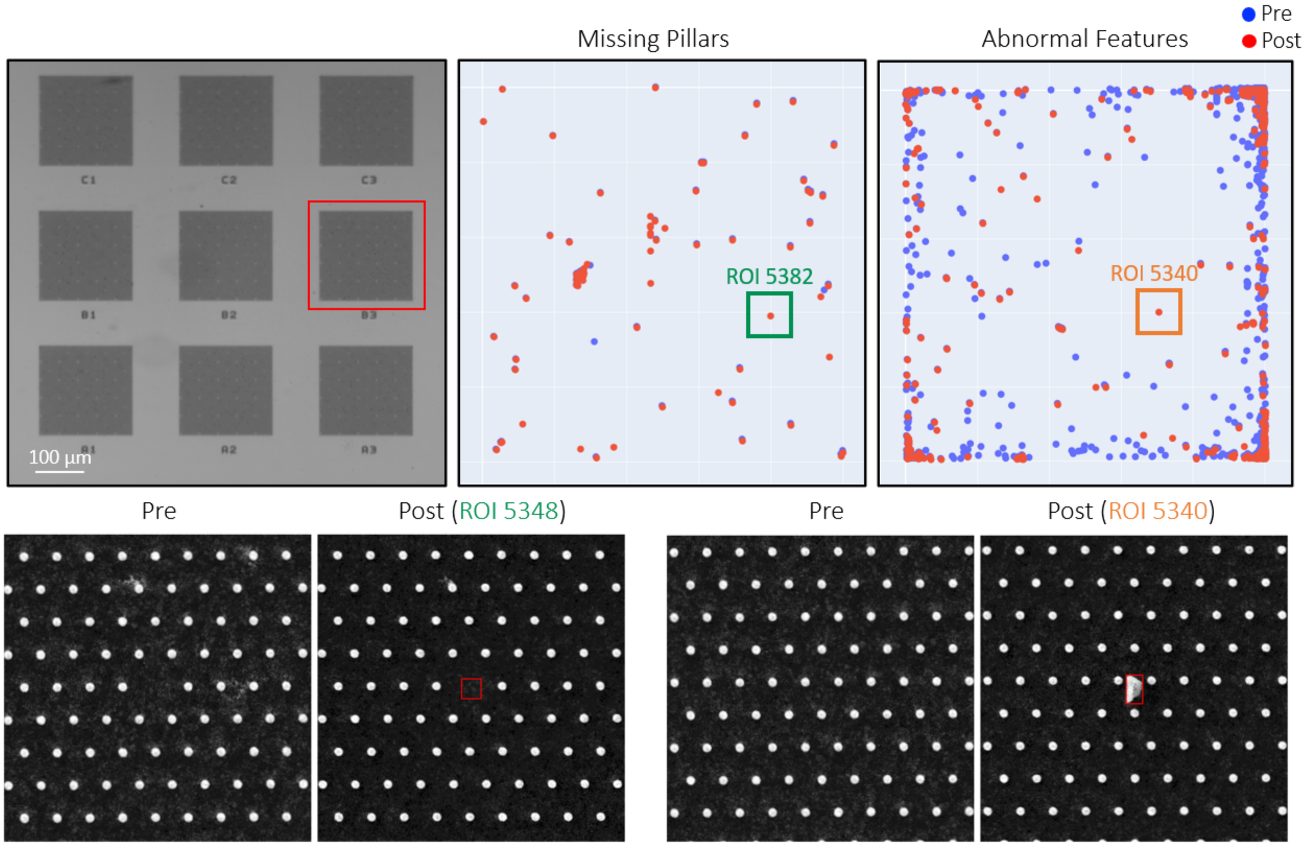

Figure 10 – Upper left: an optical microscope image of a nanopillar sample showing the 3x3 array of nanopillar pads. Upper center and right: Examples of 2D maps for a given pad for missing pillars and abnormal features, showing the spatial positions of those features in the pre and post experiment ATLAS 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.

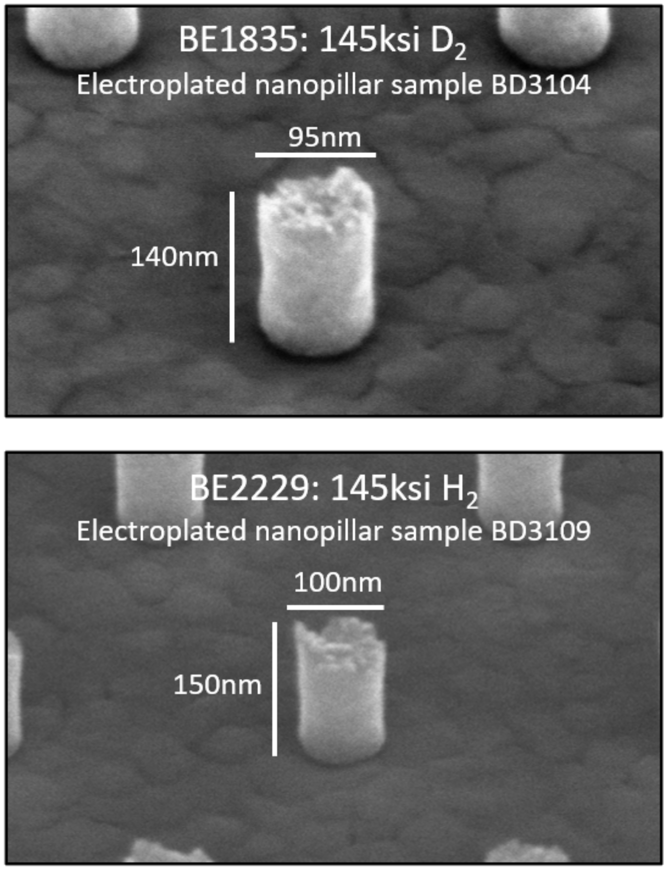

Figure 11 – Two nearly identical nanopillar samples used in a matched pair of experiments – one sample loaded in D_2, the other in H_2.

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.}

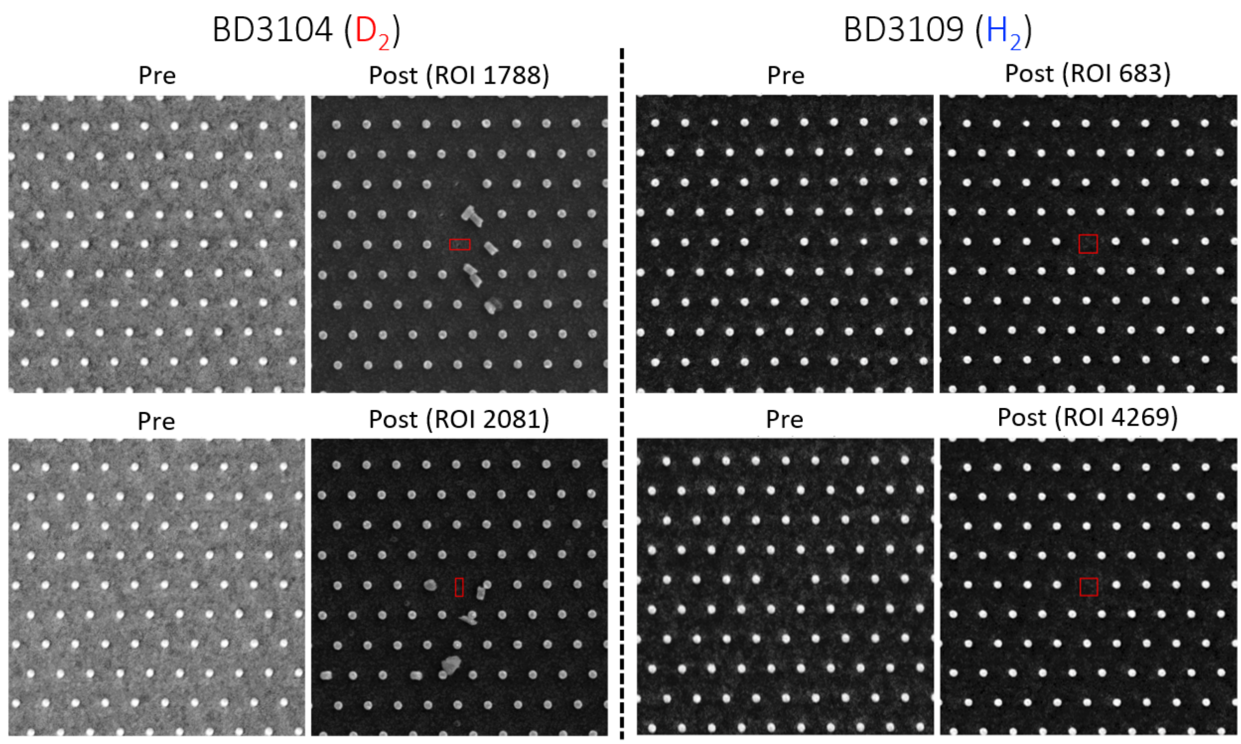

Figure 12 – Examples of missing pillars in BD3104 used in BE1835 (D_2), and BD3109 used in BE2229 (H_2).

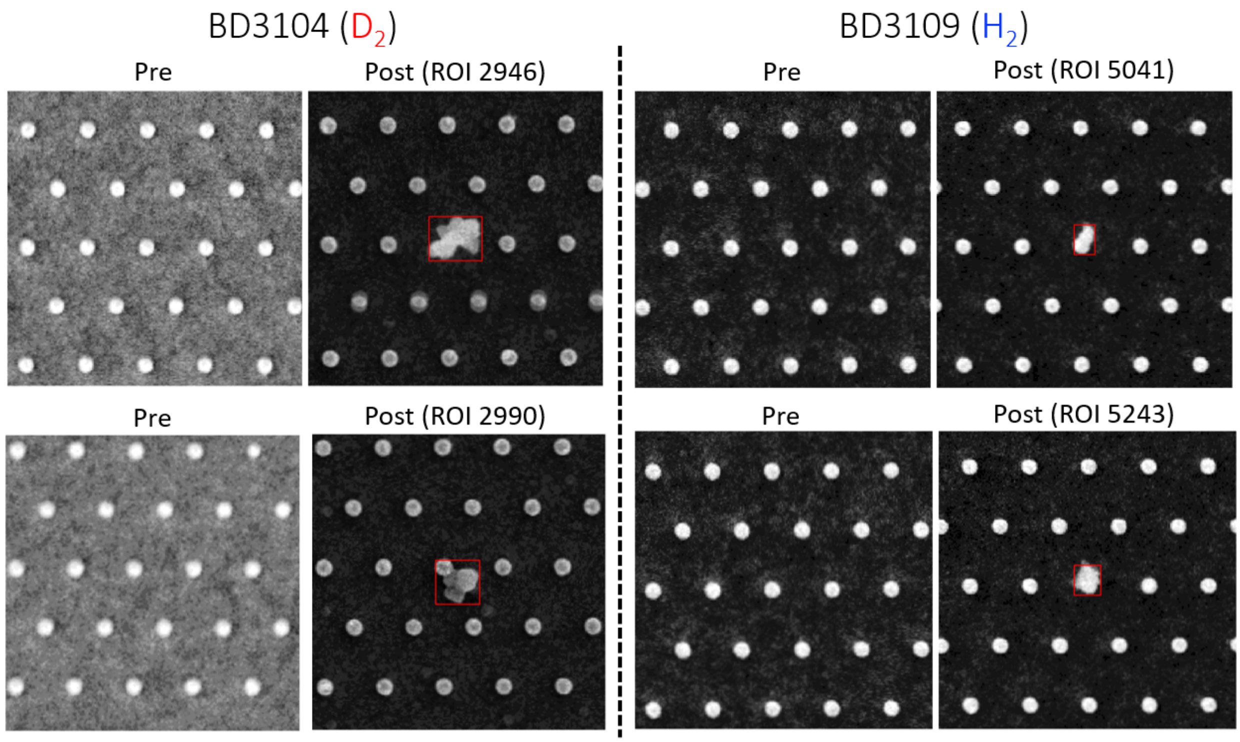

Figure 13 – Examples of abnormal features in BD3104 used in BE1835 (D_2), and BD3109 used in BE2229 (H_2).

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.

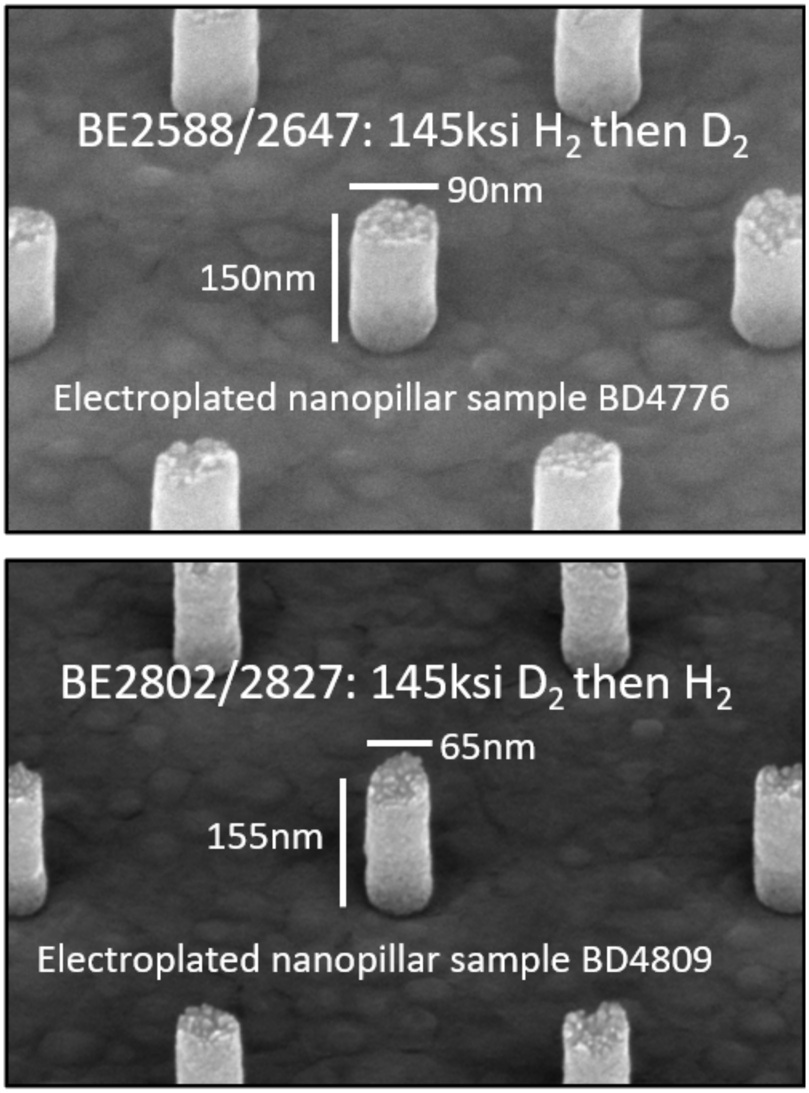

Figure 14 – Two nanopillar samples, one loaded with H_2 then D_2, the other with D_2 then H_2.

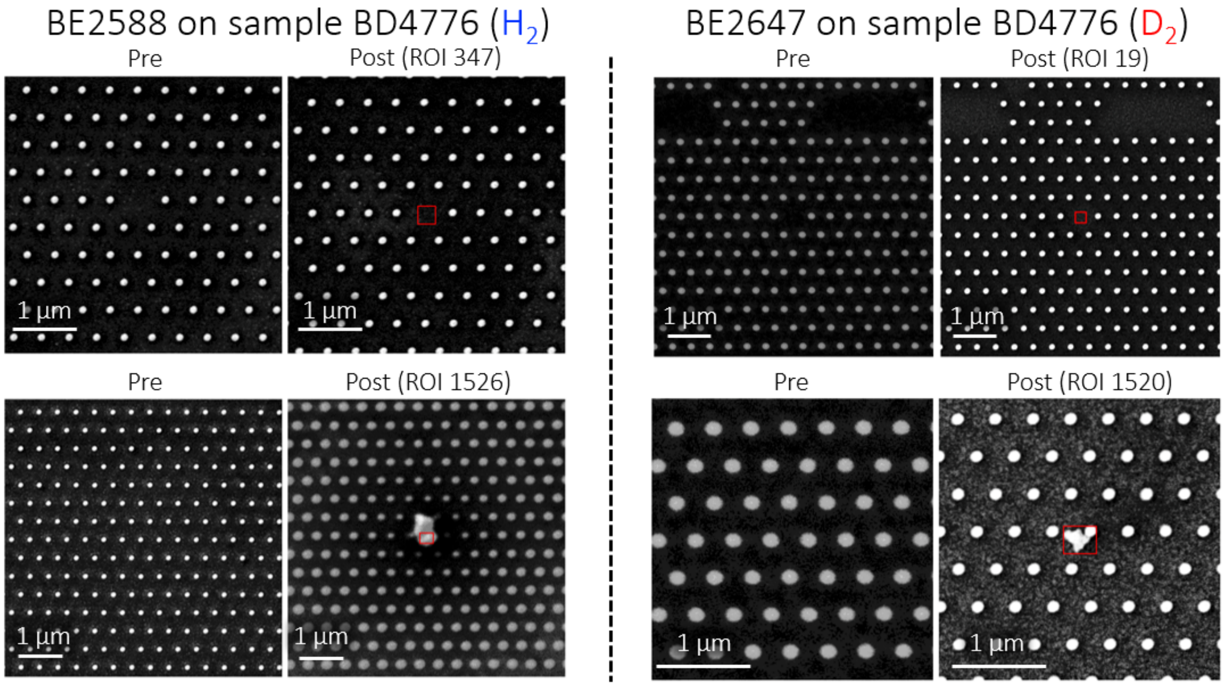

Figure 15 – Examples of missing pillars and abnormal features for H_2 and D_2 runs on the same nanopillar sample.

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.

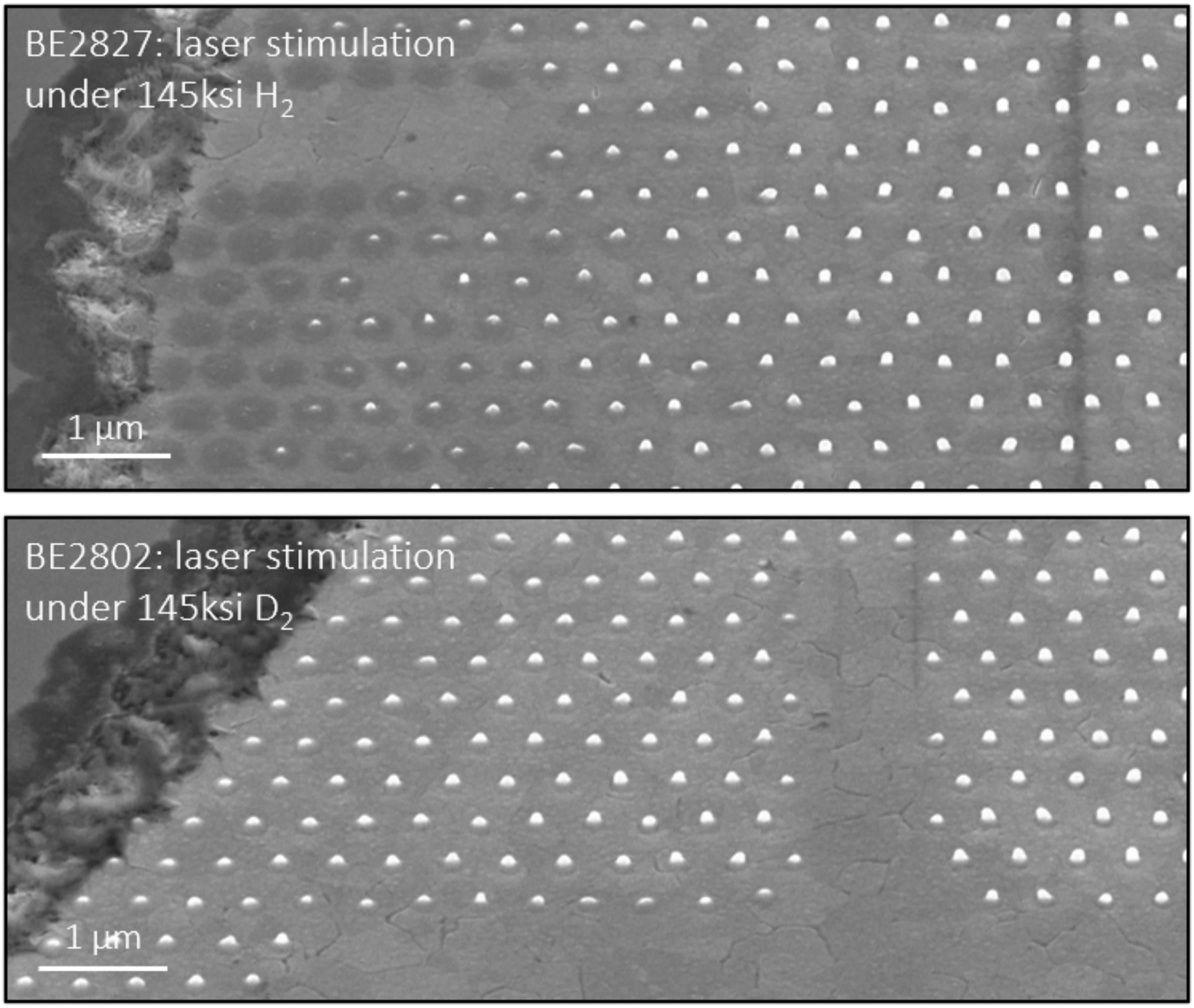

Figure 16 – Examples of pillars intentionally damaged from direct laser heating for a sample loaded in D_2 followed by H_2. The damage in each gas shows similar behavior.

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.

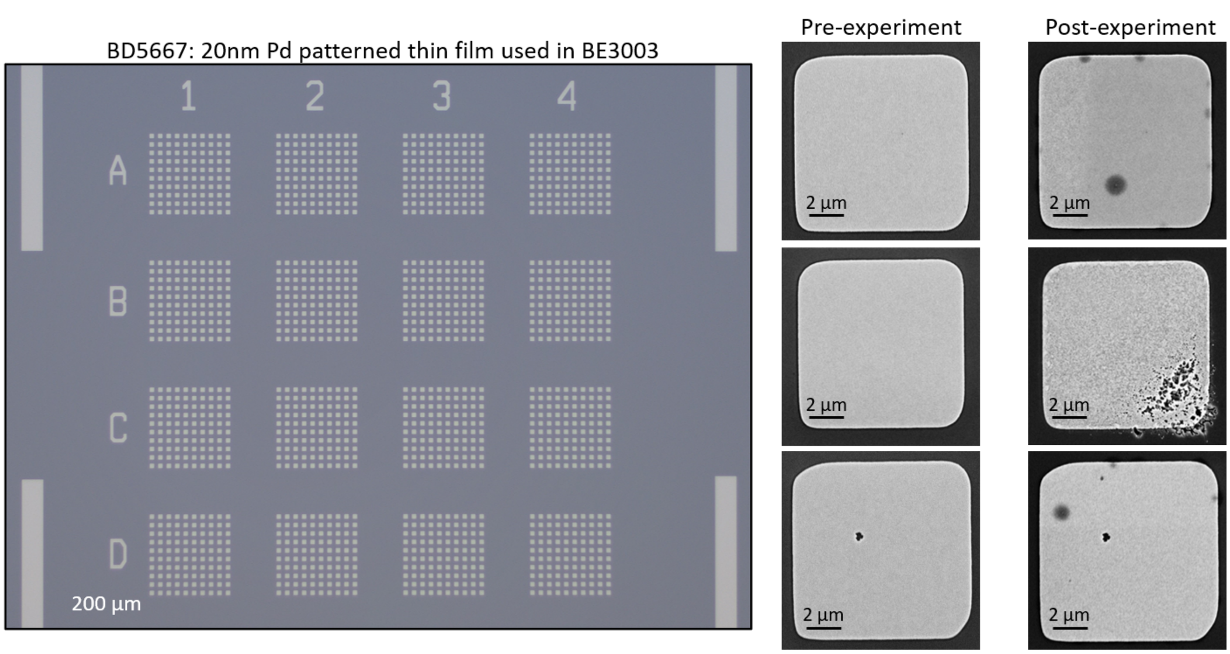

Figure 17 – Optical microscope image and some representative pre/post ATLAS image pairs of individual nanofilms for BD5667 (pure Pd nanofilm sample) used in experiment BE3003.

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.

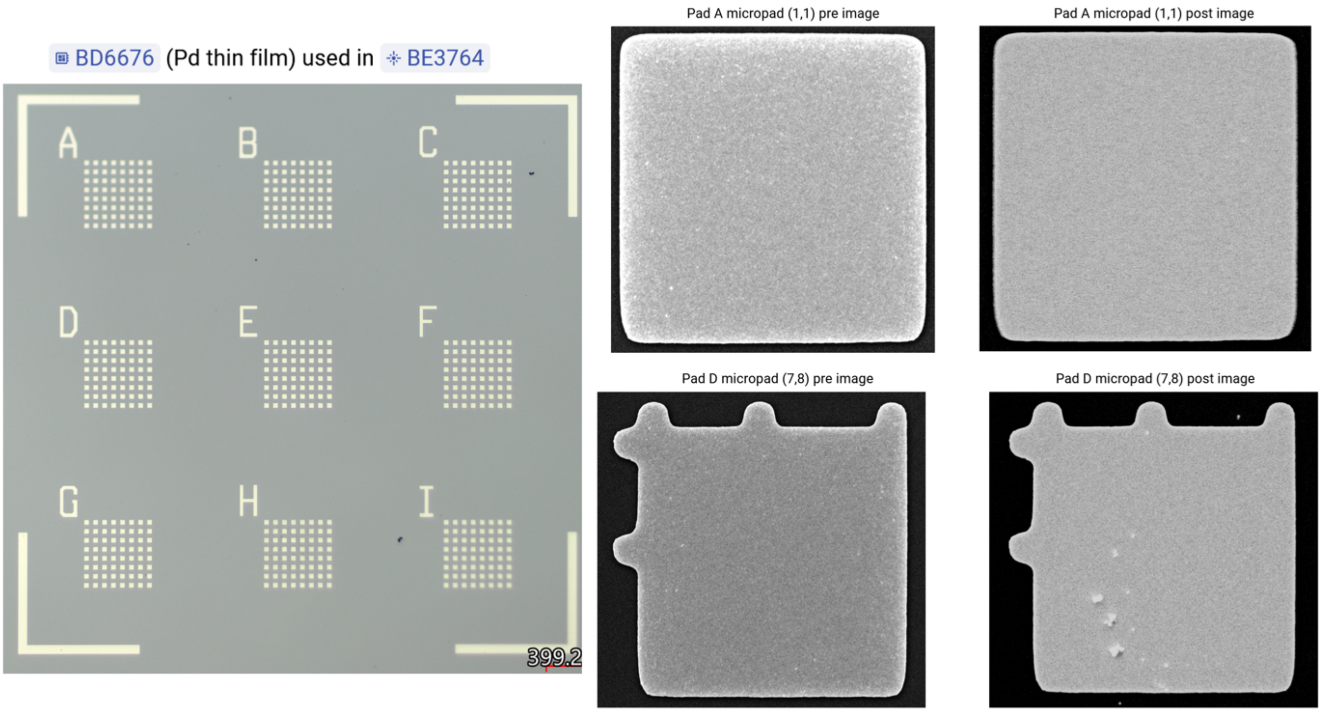

Figure 18 : An optical microscope image and some representative ATLAS images of pre and post experiment nanofilms from BD6676, a pure Pd thin film sample used in experiment BE3764.

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.