Seismology

Executive Summary

In the Seismology approach, we are looking for the mechanical vibrations that we expect to be produced by anomalous heat events (AHE), rather than *directly* measuring temperature.

To date, we have not measured any isotopic difference from Seismology experiments, which we interpret as a negative result for AHE with 1 part in 1 trillion fractional site sensitivity. These experiments were unstimulated, and AHE measurements were not performed in our highest noise-performing experimental configuration before the project wound down.

Further progress was planned to expand into larger bulk samples including palladium slugs, to explore different stimulation methods, to monitor for much longer durations, and to further develop more nuanced machine learning techniques to interpret the data. However, many of these avenues were explored thoroughly across other projects in our global effort.

Concept

In the Seismology approach, we are looking for the mechanical vibrations that we expect to be produced by anomalous heat events (AHE), rather than *directly* measuring temperature. If an AHE event occurs anywhere in the sample, the material nearby will be rapidly heated. In order to measure temperature directly, we would need to choose a specific location on the sample to place temperature probes, and we would only be able to detect AHE occurring very close to the location of the probes. However, this rapid heating will also produce a *vibration* that will propagate throughout the *whole* material, as illustrated in Figure 1. This is analogous to vibrations from an earthquake (or underground nuclear test) propagating out over long distances from its local epicenter.

We are particularly interested in techniques that are sensitive over large regions of the material, which can detect small vibrations from rapidly induced changes in temperature. We can measure these small vibrations using piezoelectric detectors that generate electrical charge in response to mechanical stress. These detectors have a sensitive area on the order of millimeters, which is many orders of magnitude larger than the temperature sensors we use in other projects, where the detectable area is usually better measured in nanometers or microns. AHE occurring anywhere in a palladium film deposited on the surface of the sensor will cause a vibration through the whole material.

Furthermore, these sensors have a long-lasting response to a fast thermal impulse, so we can be confident that the effect of an anomalous heat event will last long enough to be measured. This can be demonstrated by studying how the sensor responds to a well-controlled thermal impulse from a nanosecond-pulsed laser. Figure 1 shows that the measurement's response to a 6 ns laser pulse lasts over 100 μs – over 4 orders of magnitude longer than the pulse. We expect the thermal impulse delivered by an AHE event to be much faster than 6 ns, but both can be effectively considered infinitely sharp impulses compared to the microsecond time scale relevant for vibrations, so we can argue that the vibrational response to an AHE event (or a short burst of AHE events) should last about as long as a nanosecond laser pulse of the same energy. We can also use these laser pulses to estimate the minimum detectable energy, which sets a limit on how many AHE events would need to occur within a burst in order to be detected. Crucially, these events would not need to occur at precisely the same moment as long they occur within a short enough window to be considered an impulse – potentially up to a few microseconds.

Our measurements are closely related to a widely used, non-destructive material analysis technique known as acoustic emission testing. In this technique, a material is put under a mechanical load, causing sudden, impulsive vibrations as the material begins to yield at an atomic scale. These impulses are referred to as 'acoustic emission hits', and they are detected and counted as an indicator that macroscopic yield is approaching. Acoustic emission testing can routinely detect elastic waves with very small kinetic energies, down to femtojoules, so detecting picojoule level thermal impulses due to AHE is within reason.

In our experiments, we monitored for acoustic emission hits while palladium was loaded in ultrapressure deuterium gas. We looked for presence versus absence of acoustic emission hits during corresponding parts of paired experiments using deuterium and protium. The presence of unexplained acoustic emission hits while a sample is highly loaded with deuterium, but not with protium, would be an interesting result, especially if the hits looked like sharp thermal impulses.

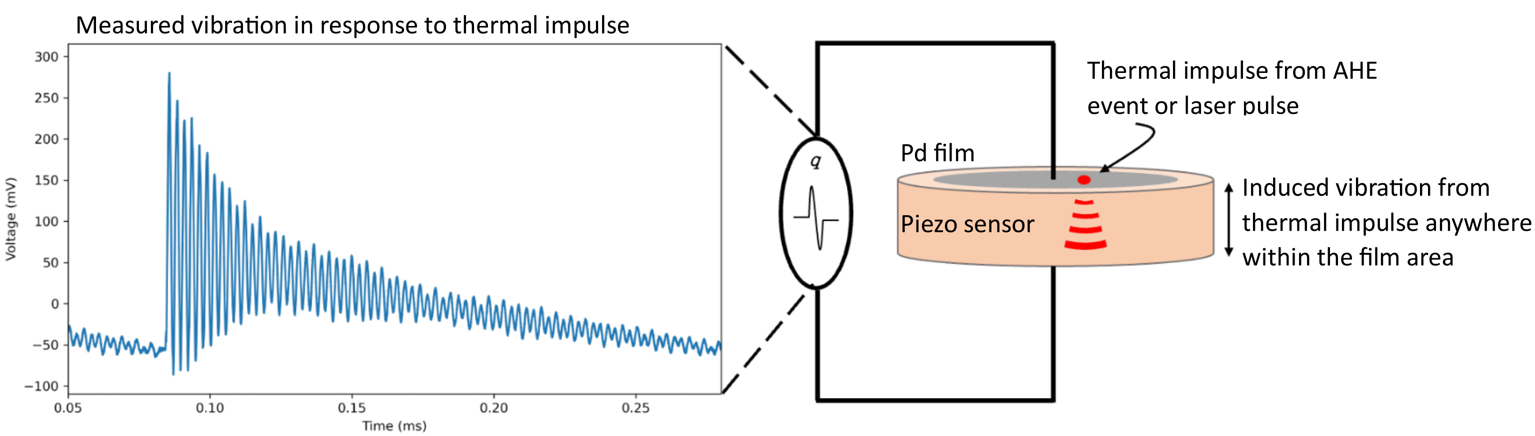

Figure 1: Seismology Concept

A fast thermal impulse occurring anywhere in a palladium film deposited on the surface of a piezoelectric sensor generates a vibration that can be measured on longer time-scales. The long-lasting response of this technique to AHE events occurring over a large area is a unique advantage to this approach.

Development of a Sample Based on a Piezoelectric Sensor:

Piezoelectric materials exhibit an electric dipole moment that changes as the lattice distorts. Practically, this means that a mechanical vibration results in an oscillating surface charge. When opposing surfaces of the material are connected through a sensitive amplifier capable of resolving the small charges that develop, piezoelectrics present an excellent method to measure small vibrations. In the field of non-destructive acoustic emission testing, sensors made from lead zirconate titanate (PZT) have become the industry standard due to their comparatively large dipole response to a given vibration. Depending on their precise geometry, these sensors can operate in a mechanical-electrical resonance, where the amplitude is a strong function of frequency, or in a broad-band mode where the response is similar across a wide frequency band.

In order to detect the vibration due to a thermal impulse, we selected a broad-band ultrasonic sensor manufactured by TDK (Part Number: B59050Z0206A031), since short impulses are inherently broad-band. Crucially, these sensors are available unpackaged and therefore can be easily integrated into a carrier assembly that is compatible with an ultrapressure hydrogen environment. It is worth noting that metallized PZT exhibits corrosion in hydrogen environments at elevated temperatures or long time scales, but this effect is minimal at the times and temperatures used in the discussed experiments[93]. We tested these sensors integrated with two different form factors of Palladium. Most reliably, we deposited a 100 nm thin film of Palladium directly onto the sensor electrode. We used a shadow mask to confine the film to the central 4 mm diameter of the sensor disk, and we deposited 5 nm Titanium adhesion layer before the 100 nm of Palladium, using 120W DC magnetron sputtering in a Kurt J. Lesker PVD 75 system.

Samples for our first proof-of-principle electrical measurements were very simple, consisting of wires soldered directly onto both sides of a PZT sensor inside of a small RF-shielded box (Pomona 3230) with BNC connections, in ambient atmosphere. These measurements demonstrated very low noise and set the benchmark performance level for our development efforts in an ultrapressure experimental environment. After these proof-of-principle measurements, we began implementing sensors into ultrapressure vessels, using the assembly shown in Figure 2.

[93] Mohammadabadi, A.S. Hydrogen-Induced Damage of Lead-Zirconate-Titanate (PZT), University of British Columbia, 2013. https://doi.org/10.14288/1.0073678.

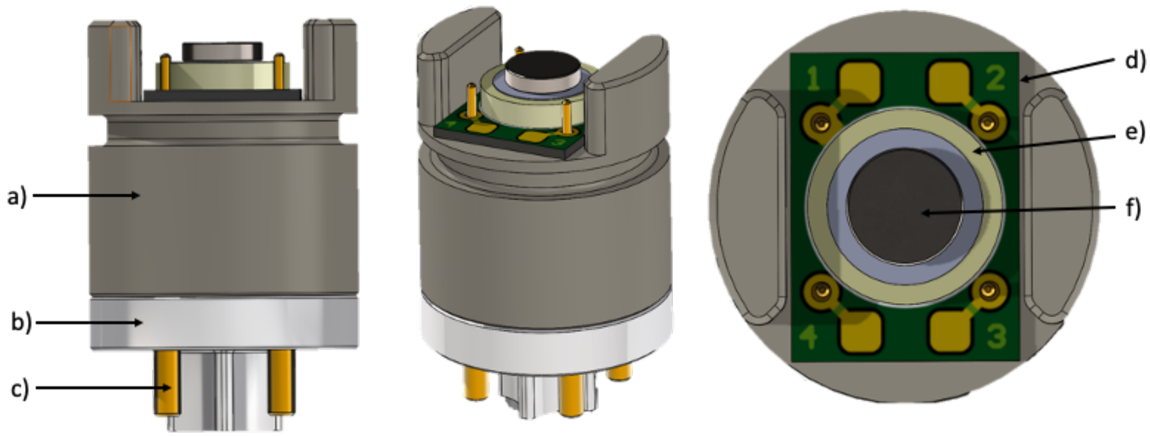

Figure 2: Seismology Sample Capsule Detail

In order to integrate piezoelectric samples into an ultrapressure environment while maintaining optical access to the surface, the sample must be held fast in the center of the cylindrical bore, and electrical contact must be made to both the top and bottom surface in a manner compatible with ultrapressure feedthrough technology. Here, the basic sample capsule is shown in detail, from which other shielded and lower-capacitance iterations were later developed. The first fully ultrapressure-compatible iteration consisted of palladium bulk (shown) or thin film adhered to the PZT sensor, which is mounted onto a ceramic carrier PCB using silver conductive epoxy. Contact to the top sensor electrode was made using a gold wire bond. Gold-coated copper pins were press fit into a PEEK capsule, secured with a fastener cap, and soldered into the carrier board using lead-tin solder. These pins fit into mating connector pieces which were brazed into sapphire to form an ultrapressure pin feedthrough.

a) PEEK capsule
b) Fastener cap
c) Gold-coasted copper pins
d) Ceramic (alumina) carrier PCB
e) PZT piezoelectric ultrasonic sensor disk (TDK Part Number: B59050Z0206A031). 1.02 mm thick by 5.00 mm diameter.
f) Palladium slug adhered to PZT sensor using Stycast 2850 epoxy. Alternatively, a 100 nm Pd thin film was deposited using DC magnetron sputtering at 120 W (Kurt J. Lesker PVD 75).

We quickly observed that noise coupling in a real experiment was far greater than in the well-isolated, benchtop proof-of-concept. In general, noise couples into an electronic measurement either capacitively or inductively.

To mitigate capacitive noise, the capacitance between the measurement leads and the noise source must be minimized. We identified that noise developing on the ultrapressure vessel itself was the primary culprit, so the capacitance between the leads and the vessel that contained the sample assembly needed to be minimized. We found that the capacitance was dominated by the ultrapressure feedthroughs. The capacitance between a brazed feedthrough pin and the vessel was measured to be around 7 pF. When epoxy feedthroughs were used instead, the capacitance was reduced to 3 pF. These epoxy feedthroughs consist of continuous, ceramic coated platinum wires fed through a narrow bore in a tungsten carbide plug and epoxied in using Stycast 2850 epoxy. The lower surface area of the wires compared to the brazed pins makes up for their closer proximity to the noise source and results in an overall reduction in capacitance. However, the performance could be improved by an order of magnitude if the feedthrough plug could be made from an insulator instead of conductive tungsten carbide, thereby pushing the noise source to a further distance from the confined wires. When sapphire was used instead of tungsten carbide, the system's overall capacitance was no longer dominated by the feedthrough. However, since sapphire is more susceptible to crack propagation, the modified Bridgman seal normally used with the tungsten carbide plugs needed to be replaced with a gold gasket seal, which required re-engineering of the ultrapressure vessel. The re-engineered vessel was successfully built and tested just before winding the project down, so we did not perform any AHE tests with the sapphire epoxy feedthrough plugs.

To minimize inductive noise, we employed several strategies together. First, we minimized the loop area formed by the measurement leads by using triaxial cables outside of the vessel. This involved designing custom breakout boards to wire the feedthrough wires into a triaxial connector mounted directly onto the vessel, as well as boards that broke the triaxial cable out into coaxial BNC connectors at the measurement amplifiers.

Second, we used radiative shielding at multiple scales. Materials with high magnetic permeability perform significantly better at shielding against electromagnetic radiation, especially at low frequency, so we chose shield materials with as high permeability as possible.

We enclosed the entire ultrapressure vessel, as well as the first pre-amplification stage, inside a Faraday cage to prevent radiative pickup from the external environment. Our first design for this Faraday cage used one-quarter-inch thick 410 stainless steel. However, sizeable holes in the enclosure were necessary to maintain free-space optical access and to feed through the beryllium copper ultrapressure plumbing.

To address these non-idealities, we modified an aluminum JRE 1522 RF Shielded Test Enclosure to add plates of high-permeability μ-metal (Magnetic Shield Corp) to all internal surfaces. The μ-metal plates were welded along their joined edges prior to bolting into the enclosure. We designed a custom feedthrough plate to pass the ultrapressure plumbing into the enclosure while maintaining the integrity of the shielding, and we delivered the nanosecond laser pulse using a fiber optic feedthrough rather than free space. This ultimately improved our external shielding substantially, but the project wound down before we could repeat AHE measurements with this improved noise floor.

In order to prevent radiative pickup in the internal measurement leads, we engineered small, cylindrical Faraday cage shields that were inserted along with the sample capsule into the inner bore of the pressure vessel. We machined these shields out of high-permeability Permalloy rods. Our highest performing sample assembly is shown in Figure 3, which integrates a sample capsule and a permalloy shield with an epoxy feedthrough.

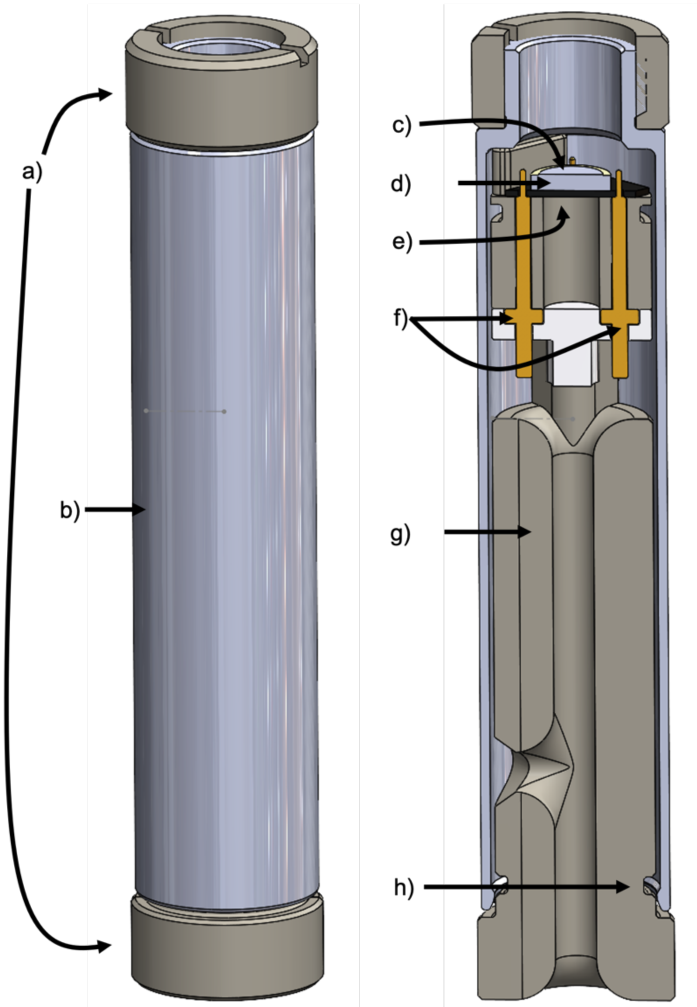

Figure 3: Seismology Sample Insert

The sample assemblies used in seismology experiments consist of a cylindrical, shielded capsule assembly that is inserted into the bore of an ultrapressure vessel. The area above the piezoelectric sample is clear for optical access through a sapphire window. This allows calibrated thermal impulses to be delivered to the sample surface using a nanosecond-pulsed infrared laser. Wires to measure the sensor are fed through an epoxy feedthrough in a tungsten carbide or sapphire piece to ambient pressure.

a) PEEK caps to prevent shorting of shield to vessel bore
b) Permalloy high-μ shield
c) 100 nm Palladium thin film. Shadow-mask patterned during DC magnetron sputtering at 120W (Kurt J. Lesker PVD75)
d) PZT piezoelectric ultrasonic sensor disk (TDK Part Number: B59050Z0206A031). 1.02 mm thick by 5.00 mm diameter.
e) Ceramic carrier PCB. Epo-Tek H20E silver conductive epoxy was used to adhere piezoelectric sensor to center contact on the PCB.
f) Gold-plated copper pins, soldered into the carrier PCB. Pins were soldered into the carrier PCB using lead-tin solder. Gold wire bond (not shown) connected the carrier PCB to the edge of the top contact on the piezoelectric disk.
g) PEEK capsule assembly, which protects the sensor and guides the feedthrough wires, which are laser welded to the pins during assembly.
h) Radial spring (not shown) sits in the PEEK groove and makes circumferential contact with the permalloy shield after it is snapped onto the capsule. Spring is laser welded to a feedthrough wire to provide optional ground to the shield.

Measurement Setup for Seismology Signal Acquisition:

Measurement Requirements:

Piezoelectric sensors provide rich data streams due to their broad frequency response and high sensitivity. To detect AHE, signals due to transient impulses are of the most interest. These impulses each last a few 10's to 100's of microseconds and may occur at any point during a several-hour long experiment, almost certainly without any degree of regularity or periodicity. As such, we need a measurement concept that is:

1) Sensitive preferentially to events that are abrupt and impulsive,
2) High enough bandwidth to resolve the character of 10-microsecond excitations,
3) Continuously monitoring over several hours.

The combination of these requirements presents a significant challenge for data acquisition and management. Collecting high-bandwidth, high-precision measurements continuously would result in terabytes of data for each experiment, which would be prohibitive for real-time transfer, storage, and analysis of the acquired data.

We addressed these issues by developing a hybrid acquisition scheme that combined high bandwidth and lower bandwidth data streams. The entire signal and data pipeline is described below and illustrated in Figure 4.

Signal Origination:

First, a thermal impulse in the palladium sample produces a localized, rapid expansion that propagates through the thickness of the piezoelectric sensor as an elastic wave – i.e. a vibration. This in turn causes a burst of electric charge at the surface electrodes of the sensor. This small sub-picocoulomb charge signal is conducted through gold wire bonds, down thin, ceramic-coated platinum wires and out of the ultrapressure vessel through a feedthrough connection. From there, the feedthrough wires are soldered into a vessel-mounted breakout board which connects them to the two inner conductors of a triaxial cable.

Signal Amplification and Filtering:

This triaxial cable connects to a charge pre-amplifier, which converts the small charge signal into a more easily measured voltage signal. We explored two approaches for this pre-amplification step, depending on how well the capacitive noise coupling was mitigated. See the below 'Note on Pre-Amplification' for further discussion.

After pre-amplification, we filter the signal to remove low frequencies that are not of interest. To accomplish this, we use the filtering capability of an SR560 low-noise voltage pre-amplifier. This also allows us to perform a differential subtraction of two inputs, and optionally to amplify yet again with some gain. Differential subtraction was essential when capacitive, common-mode noise dominated over our signal. However, when the dominant noise source is radiative / inductive pick-up, differential subtraction effectively doubles the noise, and better results are achieved using a single-ended measurement. We routinely applied a high pass filter with a 1 kHz cutoff frequency and 12 dB/octave roll-off.

Digital Sampling and Data Stream Management:

The output of the SR560 voltage amplifier is measured using a Zurich Instruments MFLI digital lock-in amplifier / oscilloscope. The signal up to this point has been entirely analog, with the bandwidth only limited by the response functions of the pre-amplifiers. However, now the analog signal must be digitally sampled and recorded. The MFLI input analog-to-digital converter (ADC) measures 16-bit samples at 60 MSa/sec, but the data could not be streamed to a data server and recorded continuously without losing large chunks of data[94]. However, we can take advantage of on-board processing to create two parallel data streams: one continuous, lower-bandwidth signal used to count impulsive events and to trigger the second data stream of individually triggered, short-duration, full-bandwidth oscilloscope traces showing the detail of each impulse event.

Collecting full-bandwidth triggered oscilloscope traces is straightforward once a few key parameters are known. Namely, the oscilloscope must be configured to acquire a trace long enough to capture the whole impulse response waveform. To characterize the impulse response of the sensor / measurement circuit, we delivered a controlled thermal impulse to the sample surface using a 1064 nm nanosecond pulsed laser (Trumpf TruPulse 2005 nano, selectable waveform 29, 6 ns pulse FWHM). We configure the scope to trigger on a rise in the signal input above a threshold set to roughly double the noise floor, and then fire a single laser pulse at moderate energy (~ 1 µJ). From the collected trace, we observed that the impulse response to a single pulse decays after about 100 µs. Thus, we configure the scope to acquire a 300 µs trace, at 60 MSa/sec bandwidth every time it is triggered. Scope traces can be acquired indefinitely so long as trigger events do not occur within 1 ms of each other. While this does allow for some missed triggers in principle, we are primarily interested in rare events. We could choose to trigger on the input signal directly, but that would be susceptible to triggering on false signals such as noise spikes.

Instead, we can demodulate the input signal (i.e. lock-in to a particular frequency) and use that as the trigger signal for the oscilloscope. The same impulse response we use to confirm the scope trace duration can be used to extract the primary carrier frequency of the sensor and measurement circuit, usually around 245 kHz. This is the characteristic frequency determined by the capacitances and inductances of the entire measurement circuit. In other words, it is not a mechanical resonance of the piezoelectric sensor. The second harmonic of this frequency can also be used effectively. When the input signal is demodulated at this frequency, with a fast, 10 µs time constant, the resulting processed signal is effectively the enveloping function of the impulsive events. It rises and falls once for each impulse, so it does not contain all the internal waveform information, but it is perfect for counting events and triggering the higher-resolution full-bandwidth oscilloscope traces.

Real-Time Monitoring and Data Storage:

The high-resolution scope traces and the continuous envelope function cannot both be streamed to the data server from the same instrument without data loss, so we instead configured the MFLI to output the envelope function to its auxiliary output, with some amplification, and we connected that output to the auxiliary input of a second MFLI instrument. This second instrument was used to stream the envelope function to the data server with 234 kSa/sec bandwidth, and to monitor the signal in real time during the experiment.

---

[94] Küng, B. (2015, July 7). Achieving High Data Transfer Rates with the UHFLI. Zurich Instruments. https://www.zhinst.com/americas/en/blogs/achieving-high-data-transfer-rates-uhfli

In summary, our measurement continuously monitors the piezoelectric sensor at a fast enough rate to detect rapid impulses, but not at such a high rate that the data acquisition pipeline would be overwhelmed. Instead, this continuous data stream is used to determine precisely when high-resolution data should be collected for a short time. We can use the continuously sampled envelope function to analyze how many impulsive events occurred, how densely, and in which parts of the experiment. We can use the high-resolution data to analyze the waveform of each impulse event individually.

Note on Pre-Amplification:

In closer-to-ideal configurations, where capacitive noise coupling was not a significant issue, we used a single-ended charge amplifier (Femto HQA-15M-10T). This amplifier has an extremely high gain of 10 V/pC, which makes it ideal for sensitive charge detection. We then passed the output through an SR560 low-noise voltage pre-amplifier with 1x gain, in single-channel source mode, purely to take advantage of the 12 dB/octave high-pass filter. The output of the SR560 was then measured using a Zurich Instruments MFLI lock-in.

However, the Femto HQA's single-ended geometry leaves the measurement susceptible to capacitive noise on the signal lead. Capacitance between the signal lead and any noise source (i.e. any conductor with a slightly fluctuating voltage) directly results in charge noise that can rail the amplifier output: $\delta q = C\delta V$. Indeed, our first measurements could not use the high-gain, single-ended configuration at all. The resulting noise railed the output completely, dominating even over a high-powered pulsed-laser calibration signal, as shown, for example, in Figure 5. It should also be noted that the PS-15-25-L switching power supply recommended for use with the Femto HQA-15M-10T, despite excellent performance compared to other switching supplies, still generates enough noise to be problematic for the extremely sensitive measurements we required. As such, we built a linear ±15 V power supply using two Acopian A15H1500 regulated linear power supplies, wired together in series with their midpoint held at ground. We also found it essential to provide a good ground connection to the preamplifier chassis through the output BNC shield.

When capacitive noise dominated over the signal, we used a differential charge amplifier that we built in-house. In this scheme, the charge from each surface of the piezoelectric sensor is individually measured using charge-to-voltage amplifiers built using development boards for the ADA4530-1 integrated circuit from Analog Devices. Then the output of each development board is connected to the two inputs of the SR560 low-noise voltage pre-amplifier, which is operated in differential source mode to subtract one input from the other. This amplifies the real signal while reducing the common-mode noise that couples into each measurement lead identically. In this case, each individual charge-to-voltage amplifier was lower gain than the corresponding Femto HQA, so that the output would not rail. We therefore used the SR560 for additional amplification, with 50 V/V gain, instead of purely for filtering.

Noise Characterization:

The sensitivity of our measurement also comes with a heightened vulnerability to electronic noise, made worse by the bustling experimental infrastructure of our ultrapressure lab space. Since every detail of the experiment had potential impact on noise performance, we needed a standardized way to characterize the noise performance of each specific configuration. Our standardized measurement captured two oscilloscope traces of the baseline input signal, connected to the sensor but without any external signal sources such as pressure ramps or laser shots. Thus, this baseline represented the inherent thermal and vibrational response of the sensor to its ambient environment as well as the electromagnetic interference present in the circuit. We routinely measured one oscilloscope trace at 60 MSa/sec sampling rate, for 80 ms, and one trace at 3.75 MSa/sec, for 1.2 seconds. In the time-domain, these traces enable qualitative observations of any periodic features, as well as quantitative metrics like the peak-to-peak and root-mean-square spread in the voltage signal. In the frequency domain, the Fourier transform of these oscilloscope traces offers standardized insight into the magnitude and effective 'fingerprint' of each noise source[95]. The frequency-domain resolution of a time-series is the reciprocal of its duration, and the maximum frequency component is half the sampling rate. Thus, the longer duration trace provides sub-hertz frequency resolution up to 1.875 MHz, and the shorter duration trace offers 12.5 Hz resolution for frequencies up to 30 MHz. Using only the longer duration spectrum would miss any noise feature above 1.875 MHz, but using only the shorter duration may fail to resolve sharp low-frequency peaks, for example at 60 Hz.

[95] Heinzel, G., Rüdiger, A., & Schilling, R. (2002). Spectrum and spectral density estimation by the Discrete Fourier transform (DFT), including a comprehensive list of window functions and some new at-top windows. https://holometer.fnal.gov/GH_FFT.pdf

Figure 4: Seismology Measurement Setup

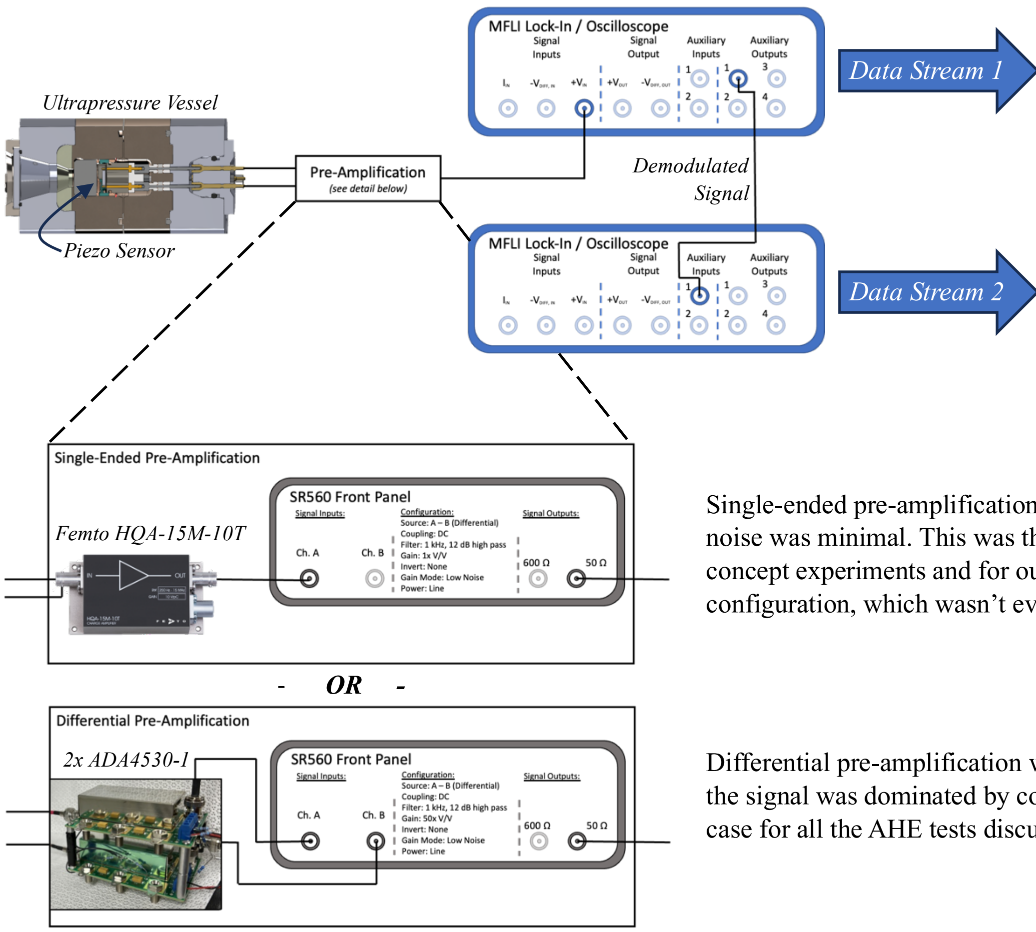

Triggered stream of high-resolution (60 MSa/sec) detail traces of individual events. Each is 300 µs in duration and must be separated by more than 1 ms from its neighbors to avoid loss.

Continuous stream (234 kSa/sec) of samples of the envelope function, which detects the shape of pulses occurring but not their internal detail. This is used to trigger Data Stream 1 and to put individual events in context.

Single-ended pre-amplification with very high gain was used when noise was minimal. This was the case for benchtop, proof of concept experiments and for our ultimate low-noise measurement configuration, which wasn’t ever used for AHE tests.

Differential pre-amplification with lower gain was used when the signal was dominated by common-mode noise. This was the case for all the AHE tests discussed in this report.

We traced most noise sources to various switching power supplies. If these supplies were part of our experimental setup, we could often remove or replace them. Some examples were switching supplies for a ThorLabs K-Cube motor controller, for the Femto HQA charge pre-amplifier, and for an internal transformer in the BRX series of programmable logic controllers. However, some noise features were present in our entire lab space, and we speculatively attributed them to larger equipment such as the facility air compressors, ventilation fans and HVAC equipment. Of note was a feature around 60 – 70 kHz that appeared to be dithered across a range of frequencies, which is a common method used to reach electromagnetic compatibility compliance in industrial equipment. These noise sources often contained both ground-coupled and radiative components, so we found it necessary to iterate over many different grounding and shielding configurations in order to optimize noise performance. The ultra-pressure plumbing and vessel acted both as a conductor that carried ground-coupled interference and as an antenna that picked up radiative interference. Ultimately, we found it necessary to minimize the capacitance between the measurement leads and the ultra-pressure vessel, as well as shield the vessel from radiative interference using a modified Faraday cage.

Figure 5: Mitigating Common-Mode Noise

a) b)

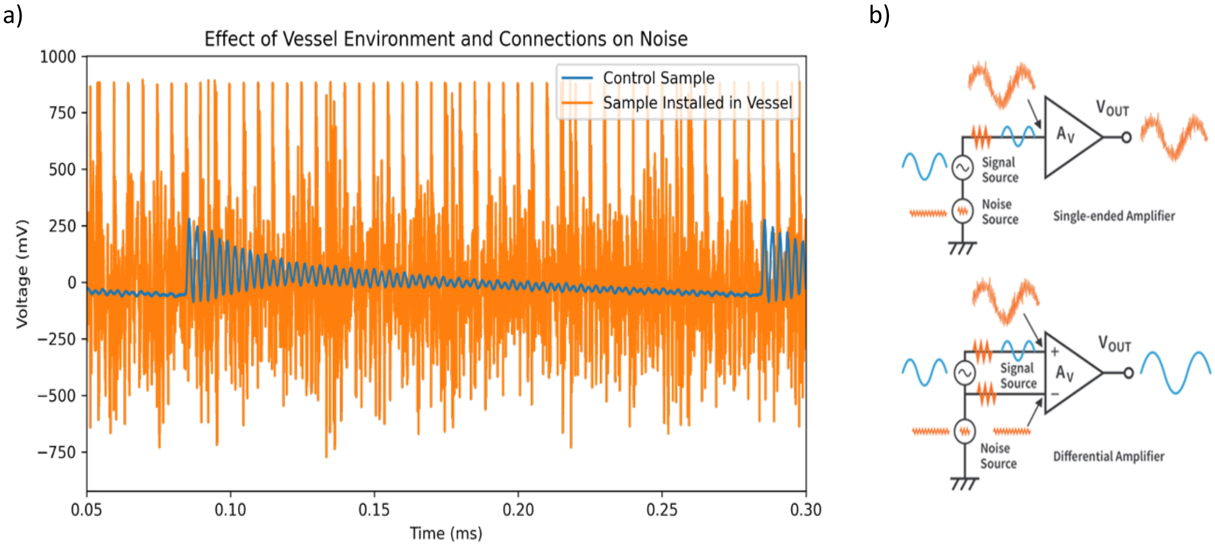

a) *Need for noise mitigation strategy*: Installing the sample in an ultrapressure vessel resulted in unacceptably high capacitive noise coupling and motivated an effort to mitigate common-mode noise.
b) *Single-ended vs. Differential Amplification*: Differential amplification is very effective at removing common-mode noise (Figure 6), which is present on both measurement leads, with the same phase. Noise from the same source that is injected capacitively onto each lead will inherently be common mode. However, differential amplification will make differential mode noise even worse. This type of noise occurs when noise from the same source couples onto each lead with an opposite phase, such as when the leads form a loop that acts as an inductive antenna. The optimal measurement strategy therefore depends on the noise environment.

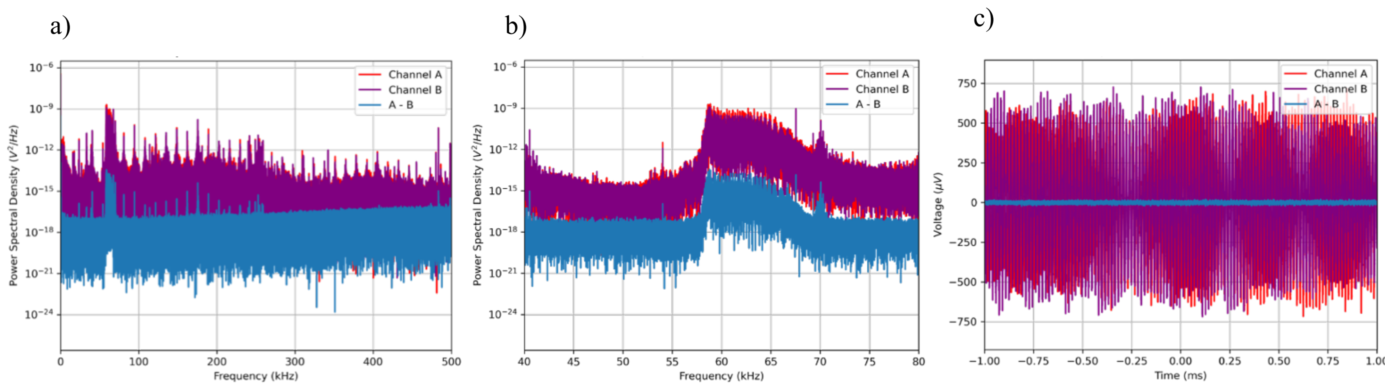

Figure 6: Differential Amplification Reduces Common-Mode Noise

a) Power spectral density of single-ended channels A and B vs. differential channel A – B, showing a reduction of up to 5 orders of magnitude in noise power.
b) Detail of a prominent noise feature at 60 – 70 kHz that often dominated the signal.
c) The time domain trace shows dramatic improvement.

Controlled Thermal Impulse Calibration:

Delivering a controlled thermal impulse to the sample surface using a pulsed laser has proven to be quite instructive, as already explained in a few instances above. Essentially, a nanosecond pulsed laser simulates the thermal impulse we expect from a burst of AHE events, but perfectly on demand and fully characterized.

The laser we used is a Trumpf TruPulse 2005 nano. Its wavelength is 1064 nm, in the near infrared, and the pulse waveform can be selected from a handful of options with different durations and pulse shapes. We used Waveform 29 as our AHE simulator, which is a gaussian pulse with a full width at half maximum of 6 ns, and 7 ns above 10% power. We also measured the response at the longest available pulse duration, using Waveform 0, which has a narrow pulse with a full width at half maximum of 35 ns with a slow trail off spending 250 ns above 10% power, and we didn't observe any difference in the signal for equivalent pulse energy. This supports our assertion that these pulses all occur in the impulsive limit, and are thereby reasonable approximations of an AHE event.

The laser also has a programmable repetition frequency, which we most often used during experiments to fire bursts of six shots spaced 1 ms apart from each other, or to deliver single, isolated shots. For characterization of the sensor response, we used continuous shots at 5 kHz so that several pulses would be observed in the same scope trace. However, repetition rates higher than 5 kHz would not allow the response to damp completely before the subsequent shot.

To control the pulse energy, we polarized the beam at its output, then adjusted the polarization using a half-waveplate, and split the resulting beam using a polarizing beam cube. To attenuate further than the extinction ratio allowed, we inserted an additional neutral density filter as needed. The pulse energy was measured in-situ.

For all of the discussed experiments, the beam was delivered to the sample using free-space optics. However, to achieve the best possible noise performance, we later rebuilt the optical system to use a fiber for final beam delivery to the ultrapressure vessel. This allowed us to maintain the integrity of our Faraday cage. See Figure 7 for details on the optical setup.

To determine our minimal detectable energy, we simply fired shots at decreasing pulse energies until we could not discern the pulse above the noise floor. The magnitude of the sensor response was consistent from sensor to sensor, for a given pre-amplifier configuration, so any variation in the sensitivity between experiments results completely from variations in the noise floor.

Figure 7: Delivering a Controlled Thermal Impulse Using a Pulsed Laser

a)

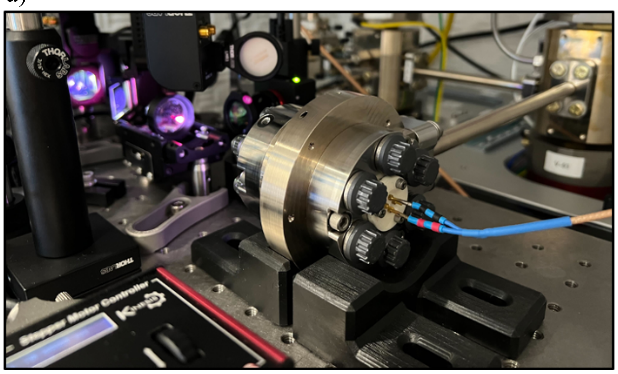

a) Free-space beam delivery to the sample through a sapphire window in the ultrapressure vessel.
b) Fiber beam delivery through an enclosed Faraday cage, which allows better noise performance. The beam still exits the fiber and passes through the sapphire window to reach the sample inside the vessel. Inset is a schematic of the energy attenuation and beam delivery scheme.

b)

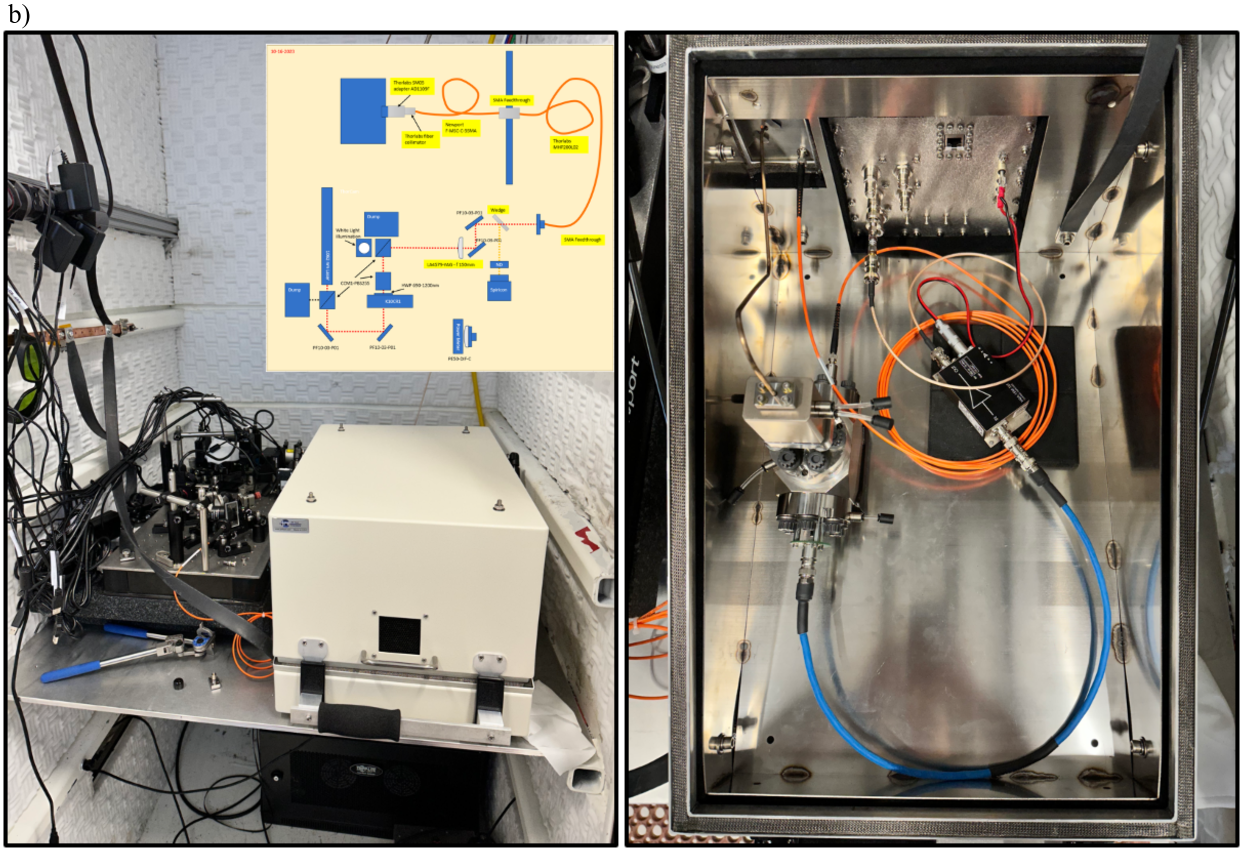

Figure 8 shows an example sensitivity test that can resolve a 30 nJ pulse. This equates to roughly 8000 AHE events. For a 100 nm thick, 4 mm diameter palladium thin film, that represents a fractional site sensitivity of 1 in 10 Trillion (1e-13). This itself is a conservative estimate, because we do not account for the loss due to imperfect absorption of palladium at 1064 nm. However, this was our best performing configuration, and our noise floor in the experiments discussed below was worse than this, with 300 nJ pulses easily detectable. That equates to a fractional site sensitivity of around 1 part in 1 Trillion (1e-12), which is already on par or slightly more sensitive than our most sensitive calorimeter. However, the calorimeter is sensitive to total energy regardless of event rate, and our measurement assumes the entire burst of events occurs quickly enough to be considered impulsive, which we estimate to be within a 10 microsecond window.

Figure 8: Measurement of Detectable Energy

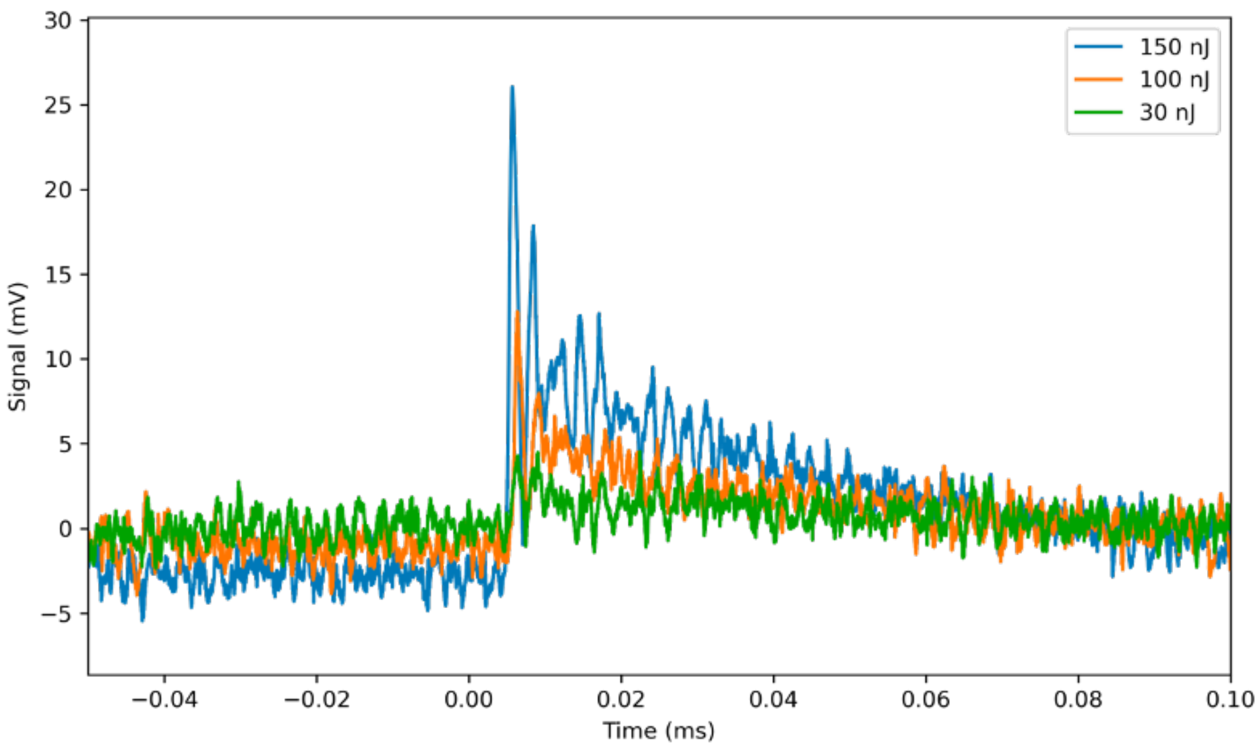

Measurement of minimum detectable pulse energy in experiment BE4191. Detecting 30 nJ in this sample geometry represents an equivalent to 8,000 AHE events, or a fractional site sensitivity of 1 part in 10 trillion. Typical experiments, including those discussed below, could detect 80,000 AHE events for a fractional site sensitivity of 1 part in 1 trillion, which is still on par with our best calorimetry.

Seismology Experiments Using Loaded PdD / PdH:

We performed Seismology measurements on palladium thin film samples in ultrapressure deuterium and protium up to 145 ksi. We incorporated laser shots into these experiments as verification of sensor performance, but these were not intended as a stimulation source, and the experiments are treated as passive (unstimulated). These experiments generate large data sets, even when the data streams are managed as described above, so we built analytical tools to aid our interpretations.

Our overall strategy is to progressively explain and rule out any impulse event that could be explained classically. We did not observe any significant isotopic differences between deuterium tests and protium tests, nor did we observe any meaningful number of unexplained impulse events at all in either case. Below, we discuss our experimental procedure and data analysis techniques.

Pre-Conditioning and the Kaiser Effect:
In non-destructive testing, acoustic emission hits occur when a material is placed under mechanical load. In our ultrapressure experiments, our sensor itself is placed under tremendous hydrostatic load, so it follows that many acoustic emission hits will develop from this mechanical stress alone. However, it is also well known that after a material is mechanically loaded and unloaded cyclically, acoustic emission hits are suppressed on subsequent cycles until the prior peak level has been surpassed. This is a phenomenological observation known as the Kaiser Effect.

We take advantage of this effect by subjecting our sensors to 5x pressure cycles in an inert gas (nitrogen), before ultimately loading them in either deuterium or protium. By the fifth nitrogen cycle, acoustic emission hits are almost completely suppressed, except during periods of active pressure ramping. Thus, we can focus our analysis on acoustic emission hits that occur during pressure holds without being overwhelmed by impulses occurring as the material responds to isostatic pressure.

It is essential to verify that cycling does not reduce the sensitivity of the sensor to a thermal impulse. We seek to suppress the origination of non-thermal impulse events – not to suppress our ability to detect events at all. To verify this, we measured the sensor's response to five laser pulses both before and after inert pressure cycling. We observed no decrease in sensitivity to thermal impulses, while impulses occurring due to pressure alone were greatly suppressed. See Figure 9.

Figure 9: Pressure Cycling Does Not Reduce Sensitivity

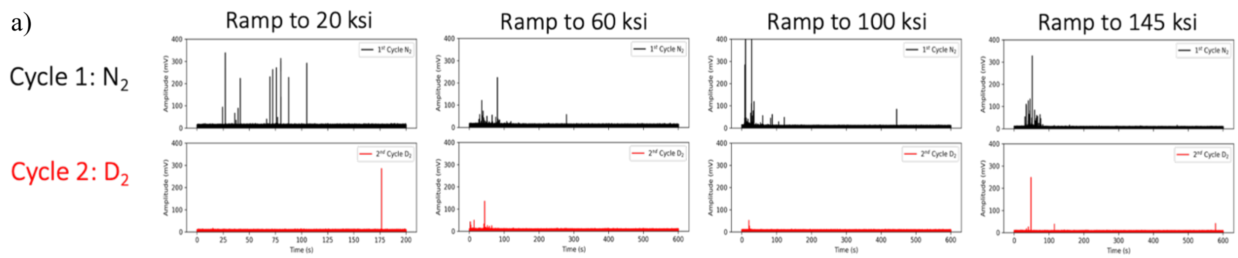

a) Demonstration of the effectiveness of one inert pressure cycle on suppressing acoustic emission hits due to the mechanical stress applied by pressure.

b) Demonstration that the sensitivity to thermal impulses is not affected by the pressure cycling.

It is worth noting that 5x pressure cycles suppress unwanted acoustic emission hits even better, but not strictly perfectly.

b) Response to thermal impulse from laser

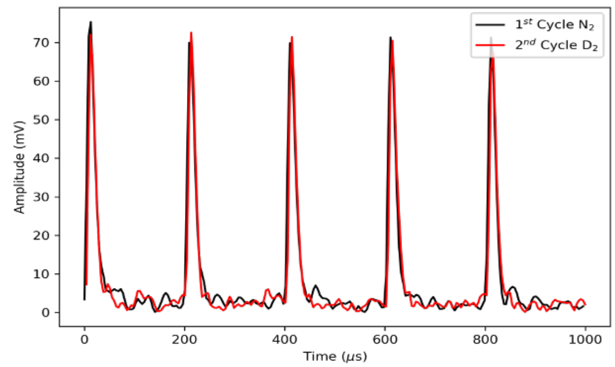

Experimental Procedure:

In these measurements, we seek to identify the presence of thermal impulses that are not accounted for by classical processes, especially when those impulses are present under deuterium loading and absent under protium loading.

First, we prepare the sample assembly as shown in Figure 2. Next we pressure cycle five times up to 145 ksi nitrogen to condition the sample. We collect both long-duration and triggered data streams during all these cycles, as described above. Next, we cycle to 145 ksi deuterium. By this point the number of triggered impulses per cycle has decreased dramatically. Next, we cycle up to 145 ksi protium, as an additional control.

Since we use the same thin film sample for both deuterium and protium loading, it is essential to run the deuterium cycle first. We have learned that thin films do not completely de-load after de-pressurization, so there is certainly still a substantial amount of deuterium loading during the protium cycle. If protium was introduced first, then we would never reach a state of high deuterium loading – because there would instead be a substantial fraction of protium in the sample. We argue that the protium-deuterium mixture still serves as a good control in which we do not expect AHE would be possible. However, the point is in fact moot because we ultimately did not observe any unexplained thermal impulses in the deuterium cycles in the first place – so there was nothing in particular to control for.

After the experiment, data was analyzed as described below. In each stage of analysis, we progressively classify and rule out acoustic emission hits until we are left only with events that have no clear explanation.

Figure 10: Long-Duration Data Gives Context to Each Event

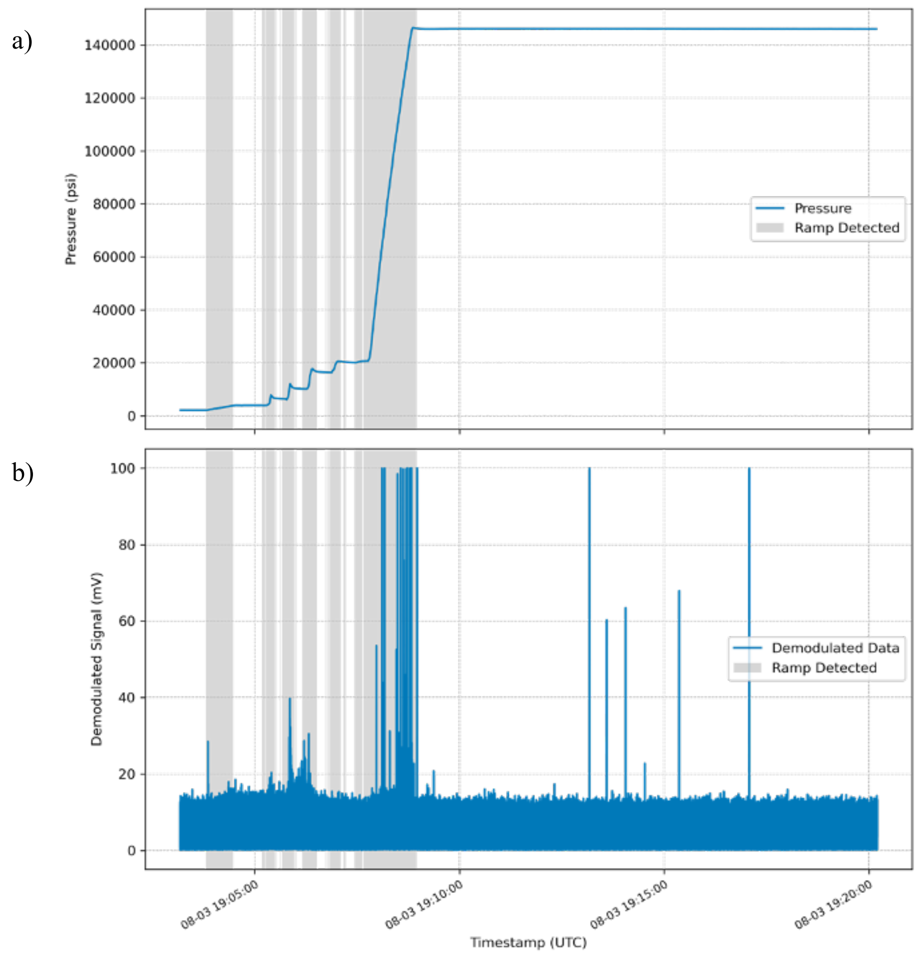

a) Pressure data is used to differentiate between times when the pressure is ramping and times when the pressure is holding.

b) Long-duration demodulated data shows that most events are clustered during pressure ramps, and assigns labels of interest to events observed during pressure holds, where we expect fewer classically explained events, for further scrutiny.

Analysis of Continuous Long-Duration Data Stream:

The primary advantage of the long-duration data stream is to place measured events in the context of what else was happening in the experiment at that moment. Along with the 234 kSa/sec demodulated envelope data, we also collect data from the pressure transducers at 1 Sa/sec. This data is used along with the trigger timestamp of each measured event to assign a pressure value, a gas identifier, and a label indicating whether the pressure was ramping or holding. These are all saved as columns in a data frame that summarizes each measured event. Most of our experiments record of order one thousand events.

Figure 10 shows an example of how this pressure data can be used to differentiate between events that occur during pressure ramps and events that occur during pressure holds.

Analysis of High-Resolution Triggered Data Stream:
We seek to use the high-resolution, triggered data to make interpretations about individual acoustic emission hits. We are interested only in events that appear to be sharp impulses, so to distinguish between these events and those of other origins, we employ several techniques from signal processing and machine learning.

First, we use a wavelet transformation of each data trace to obtain a fingerprint that shows how the frequency components of the signal evolve over time. Breaking a signal down into frequency components using Fourier analysis has been a staple of signal processing for centuries. However, transforming entirely to the frequency domain loses all information about the duration of the pulse in time, which is a crucial attribute to distinguish between impulsive and non-impulsive events. Instead, wavelet transformation produces a two-dimensional data set that includes both frequency information and time information[96]. Figure 11 shows an example of two measured acoustic emission hits, the demodulated envelope functions which triggered collection of the data, and the wavelet transforms using a convolution with the 'cgau3' complex gaussian wavelet.

Our experiments generate hundreds to thousands of detailed traces representing acoustic emission events, so we turned to unsupervised image clustering techniques in machine learning to group them based on how similar they are to each other.

First, we run a batch process to calculate the wavelet transforms of each event and save them as a png image. Next, we process the images through a neural network (ResNet-50) to assign an abstract, many-dimensional 'feature vector' to each image. Then we use principal component analysis to determine which hundred or so components of the feature vector are responsible for all of the variation between the images. Next we need to visualize which events are similar to each other, so we need to project the roughly 100-dimensional principal component space onto just two axes. To accomplish this, we use a t-stochastic nearest neighbor embedding (t-SNE) algorithm to optimize the plane we project onto so that it preserves most of the spread.

[96] Arts, L.P.A., van den Broek, E.L. The fast continuous wavelet transformation (fCWT) for real-time, high-quality, noise-resistant time-frequency analysis. *Nat Comput Sci* **2**, 47-58 (2022). https://doi.org/10.1038/s43588-021-00183-z

Figure 11: Wavelet Transform Analysis of High-Resolution Data

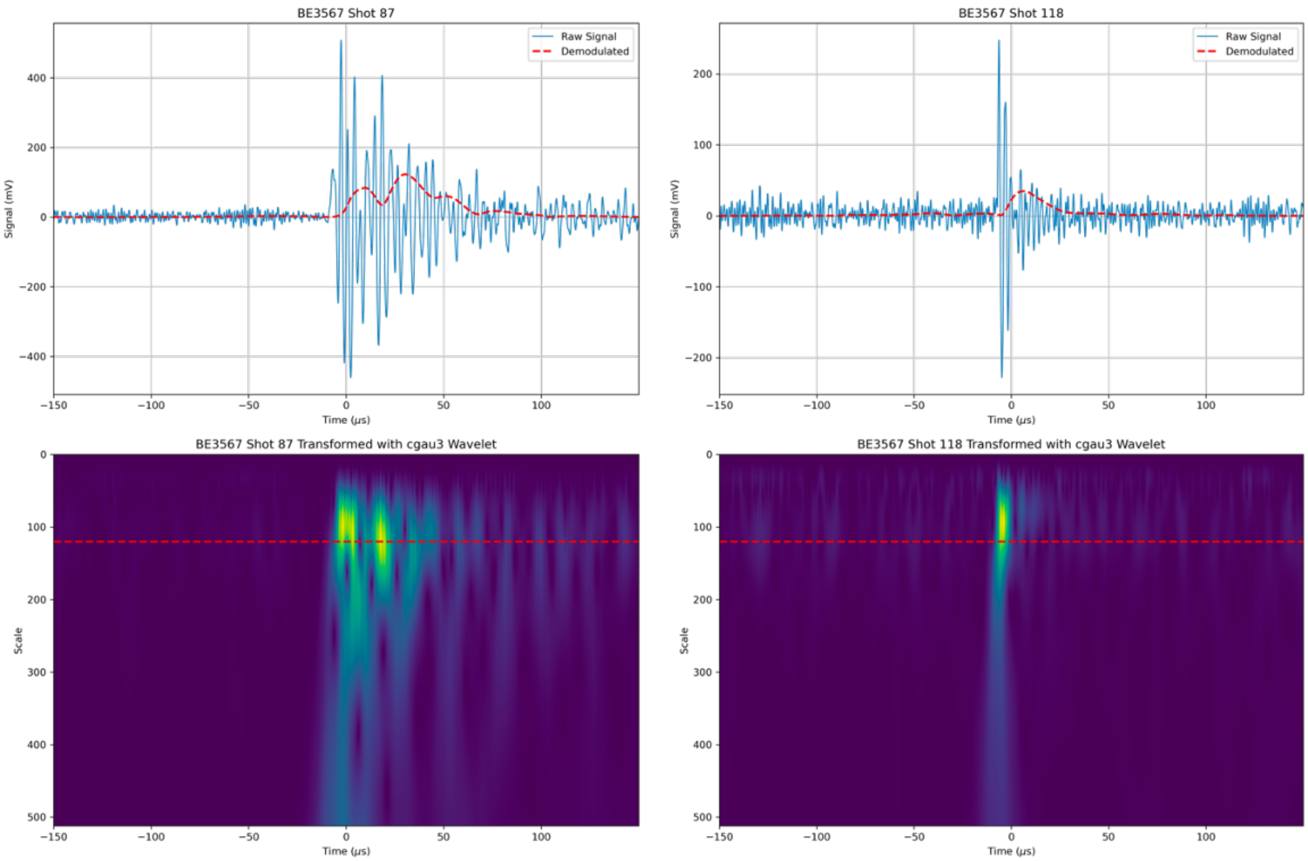

The above show acoustic emission hits that were recorded in high resolution, along with the demodulated data that triggered collection of the events. The wavelet transforms shown underneath the time traces represent ‘fingerprints’ that can be clustered to determine which events appear similar to thermal impulses. For example, thermal impulses are bursts of excited broad-band frequencies that decay quickly after their initial excitation, while other types of acoustic emission can show gradual increase or beating over time. The ‘scale’ axis shown on the y-axis in the wavelet transformations is proportional to frequency domain. The horizontal red dashed line in the wavelet transform images represents the center frequency of the demodulation.

a) An acoustic emission hit that occurred during pressure ramping.
b) An acoustic emission hit that resulted from a laser pulse.

The result is a plot like the one shown in Figure 12, where an abstract ‘similarity’ between two points is captured by the distance between them in the plane. The units of the axes are meaningless, because they ultimately originate from the abstract way the neural network encodes features within the images that are used to differentiate them from each other. We can quickly see that the clustering corresponds to real similarities because all laser shots group together. Furthermore, if we replace the points on the graph with image thumbnails, we can see that similar images are grouped together.

Figure 12: Machine-Learning Based Image Clustering

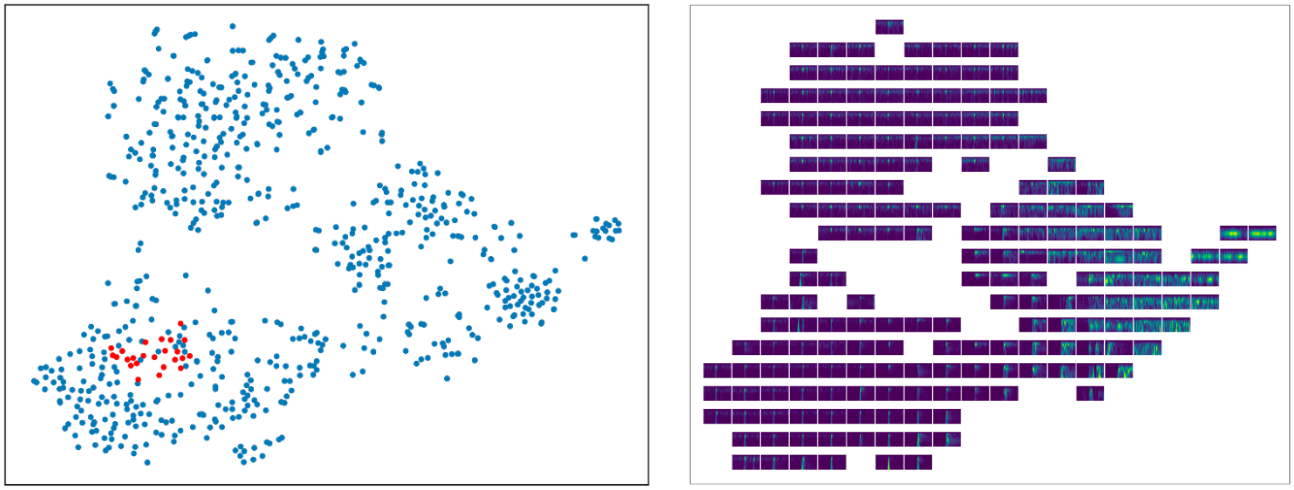

a) The 'latent space' representation of all the measured acoustic emission hits in a particular experiment. Points that are closer to each other on this plot are similar in an abstract way defined by the ResNet-50 neural network. Red dots indicate laser shots, which group together as they should.
b) Replacing each dot with a thumbnail image shows that indeed, acoustic emission hits with similar wavelet transforms are grouped together in this depiction.

Using the contextual data from the long-duration data stream, we can now progressively filter out events that occur during pressure ramps, and separate out events occurring under deuterium loading and protium loading.

Results:

To determine whether a given experiment contains impulse-like events that are unexplained, we start from the analysis described above and filter out unwanted events in steps.

For example, in BE3683 shown in Figure 13, we first separate events into those that occur under deuterium loading vs. protium loading, and then we further remove events that occurred during pressure ramps. In the experiment shown, some single-pulse laser shots were taken (red), and some multi-shot laser pulse trains that contain multiple shots in the same scope trace were taken (purple). After removing uninteresting events, we are left with only one possibly interesting impulse in the deuterium experiment, and only two in the protium experiment. This is not a meaningful isotopic difference, and the result is therefore interpreted as negative for AHE.

Figure 13: No Isotopic Difference Observed

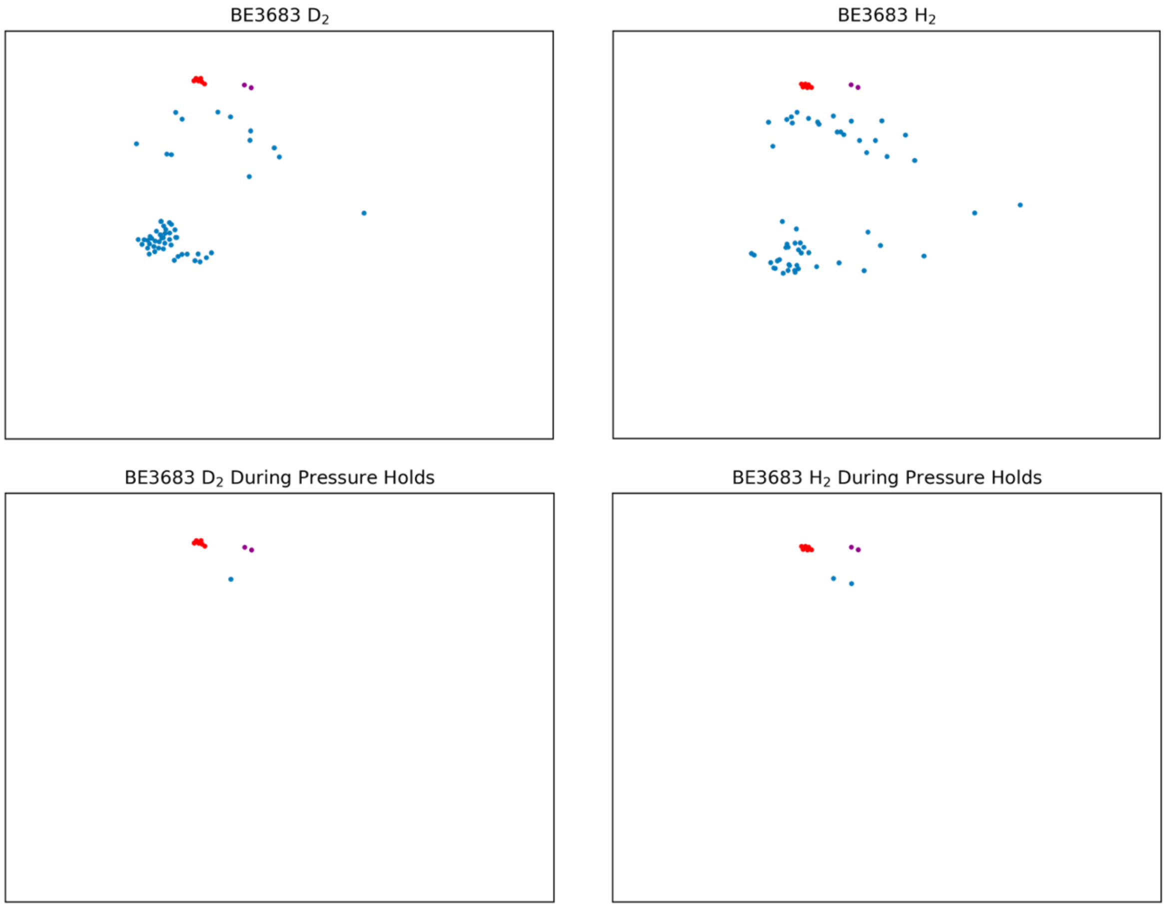

a) and b) Clustered events loaded under deuterium and protium. Red dots are single laser shots. Purple dots are multiple laser shots in the same scope trace. Blue dots are not laser shots.

c) and d) Remaining events after removing events that occurred during pressure ramps. No isotopic difference is observed. Furthermore, there are not a significant number of impulsive events that remain unexplained at all, indicating that our pre-conditioning and context-based filtering works well.

Conclusions

To date, we have not measured any isotopic difference from Seismology experiments, which we interpret as a negative result for AHE with 1 part in 1 Trillion fractional site sensitivity. These experiments were unstimulated, and AHE measurements were not performed in our highest noise-performing experimental configuration before the project wound down.

Further progress was planned to expand into larger bulk samples including palladium slugs, to explore different stimulation methods, to monitor for much longer durations, and to further develop more nuanced machine learning techniques to interpret the data. However, many of these avenues were explored thoroughly across other projects in our global effort.