Summary of the Callisto Electrochemical Research Effort

Scope of effort

We made replication attempts of prior published electrochemical experiments by ostensibly credible researchers who claimed to have observed AHE calorimetrically. We added derivative variations as well as improved controls and diagnostics. All of the experiments were testing the original Pons-Fleischmann hypothesis that AHE results from a nuclear coalescence of deuterons in a highly-loaded Pd lattice. While there have been numerous claims of AHE involving other metals and also involving protium rather than deuterium as the nuclear feedstock, as well as much work on gas-phase systems, we felt that it was most important to focus on testing the original basis for claims of AHE.

Prior major replications

We acknowledge that we are far from being the first to conduct these sorts of replication attempts, but believe that the overall number and duration of our experiments far exceeded any other published work. Other than the flurry of replication attempts immediately after the original 1995 announcement, there were few experimenters willing to work further in the field, and those who did such work had very limited access to facilities and funds. We are aware of a few past significant replication efforts (i.e., likely in excess of $1M USD), but other than an effort by Google, they were not published in peer-reviewed journals, so may not have received the attention that they deserved. Earthtech conducted multiple different AHE replication attempts, and the work appears to have been of good quality with high-precision calorimetry, but their results (all negative as regards detecting AHE) are only documented on a website. Coolescence has published a YouTube video disclosing their failure to replicate AHE after a 13-year effort and has posted additional data online, but not compiled their overall effort into a publication. Guffey et al published an attempted replication of laser-stimulation experiments by Letts and Cravens, also with negative results. We understand that many scientists (and many journals) are not motivated to publish negative results, particularly in a field that is widely derided, but we nevertheless felt it would be a service to the technical community to go on record with our own results, including some of the ancillary discoveries.

Materials

The original Fleischmann-Pons electrolysis experiments employed a polycrystalline Pd cylindrical cathode and a helical Pt wire anode in 0.1 M LiOD:D_2O. This was the starting point for our AHE explorations, although we made many variations and perturbations to probe the parameter space. All experiments incorporated at least two types of controls that were run simultaneously whenever possible: an electrolyte formulated using

LiOH:H_2O instead of LiOD:D_2O, and/or a cathode that could not load with hydrogen, i.e., consisting of Pt rather than the 99.9% pure Pd normally used for "active" samples. Although some have claimed to observe AHE in light water systems, the nuclear physics of fusing protons is even less plausible than for deuterons. Moreover their nuclear properties (reaction cross sections, etc.) are so different that it is extremely unlikely that any AHE, if it existed, would have comparable orders of magnitude. Thus we believe light water to be a valid control, and suggest that those who have reported AHE from both isotopes were likely making faulty measurements in both cases.

Using Quartz Crystal Microbalances coated with thin films of Pd to directly measure mass loading of hydrogen during electrolysis, we found that protium uptake in Pd was highly preferential to deuterium uptake by more than 10x, e.g., a 3% H_2O contaminant of a D_2O electrolyte would result in 40% of the lattice sites in the Pd cathode being occupied with protons rather than deuterons. Accordingly, extreme care was taken to maintain the isotopic purity of the LiOD electrolyte (typically >99.8% as supplied by the vendor), through the use of an ultra-dry inert glovebox environment for storage and filling of sample vessels, and in most experiments the use of sealed calorimeter vessels containing a Pt recombiner.

Pd cathodes were typically >99.9% polycrystalline that in wire, foil, thin film, or bulk cylindrical forms. We also tested bars diced from single-crystal Pd wafers with (100), (110), and (111) surface orientations, as well as various Pd alloys that have been claimed to generate AHE. The alloying agents tested included B, Rh, Ag, and Ni in various concentrations as single constituents, as well as some complex multi-constituent alloys of our own choosing.

Loading measurements

Loading in bulk Pd cathodes was determined from 4-point measurements of bulk resistivity. Accuracy of this well-known method is highest for linear high aspect ratio cathodes (i.e., wires), which conveniently was our preferred sample geometry for most experiments. The resistance curve vs. loading has a maximum at around 75% loading, i.e., loading is a double-valued function of resistance, so it was important to monitor resistance continually throughout an experiment. The Pd must be uniformly loaded throughout its depth for this method to give accurate results. As loading in the bulk is a diffusion-limited process, small diameter wires were preferred (<1 mm) as they would stabilize in less than an hour.

Resistivity measurements are not useful for measuring loading of Pd thin films (thickness range of 0.1-2.0 μm in our work), as factors such as film thickness and grain size increase the resistivity by unpredictable amounts. As an alternative, the quartz crystal

microbalance (QCM) method provided us a direct measurement of mass gain, and hence loading (provided that isotopically pure electrolytes were used). The QCM method was usable up to about 80% loading, and correlated very well with coulometric stripping measurements. At higher loadings, hydrogen bubble evolution made QCM readings unstable, and problems with adhesion of the underlying electrodes to the quartz substrate were common, as the slightest micron-scale defect in adhesion could permit outgassing of high-pressure hydrogen, leading to runaway blistering.

As an alternative method for measuring thin film loadings, we developed optical reflectance measurements. Pd thin films of order 0.1-1.0 μm thickness were deposited onto a double-polished sapphire substrate along with an intervening ultrathin (1 nm) Ti layer to improve adhesion. Reflectance was measured from the back side of the sapphire, with only the Pd-coated front side was exposed to the electrolyte. A custom apparatus maintained focus and normal incidence alignment of a laser beam onto the sapphire. By comparing with X-ray diffraction measurements of film strain (a proxy for loading), we determined that a nearly linear relationship existed between strain and reflectance at a 785 nm wavelength, with a maximum reflectance change of slightly more than 22% approaching 100% loading. We later discovered that reflectance at a 1550 nm wavelength showed a non-monotonic relationship over the full 0-100% loading range, but with the very useful feature that the relationship was much more sensitive than at 785 nm when in the high (80-100%) loading range of interest to us, i.e., producing a reflectance change of ~5% at 80% loading increasing to ~20% at 100% loading. The thin-film reflectance method became our primary tool for ranking different additives for their effect on loading, which we then corroborated by testing the same additives on bulk Pd samples (wires) using the resistance method.

Loading enhancement

Beginning with Fleischmann and Pons' original work, it has often been alleged that a high deuterium loading fraction (D/Pd ratio) is an essential condition for creating AHE. Various threshold values have been claimed, generally in the range of 87.5-90%. This is a challenging level to achieve under electrolytic loading at ambient pressures and temperatures; in fact we believe it impossible to achieve using a truly pure LiOD electrolyte. Initial room-temperature loadings reach a peak at a particular current density, typically ~120 mA/cm^2 for most of our wire cathodes. Reduction in loading at higher current densities may be due to increased hydrogen gas evolution which leads to rapid outgassing from those portions of the cathode that lose contact with the electrolyte during bubble formation. This loading peak is generally under 80% (e.g., 73% for D and 77% for H, although this is temperature-dependent with higher temperatures resulting in lower

loadings, and conversely for lower temperatures). To the best of our knowledge, all claims of D loading above 80% have been achieved by modifying the surface of the Pd cathode from its initial nominally pure state, either naturally (as various trace ionic impurities in the LiOD electrolyte begin to plate out onto the surface of the cathode), or through the deliberate use of additives. A mere monolayer of electrodeposited impurities will alter the equilibrium loading, usually upward (e.g., by acting as a poison to the Tafel reaction wherein surface deuterons recombine to evolve gaseous D_2). It does not take long for this amount of surface contamination to accumulate despite using the highest purity LiOD salts available.

An experiment wherein loading is determined by random uncontrolled batch-dependent impurities that accumulate over time is obviously not a recipe for achieving reproducible results, and the large scatter in reported peak loadings by many prior experimenters makes this point. Since many of the published AHE claims that we were interested in replicating appeared to be due to naturally occurring contamination, we were aware that our most faithful attempts at replication would be imperfect, as there is no expectation that our LiOD feedstocks had the same types and quantities of impurities as those used by earlier experimenters. To create some determinism and explore the parameter space of possible relevant contaminants, we tested various additives in concentrations that we believed were sufficient to overwhelm the effects of trace impurities.

Many of the published additives alleged to enhance loading were found to be ineffectual in our experiments, e.g., P, Se, Te, Eu, Ce, S, SiO2, Cu, Al, Si, and B, although it is entirely possible that the chemical forms of our additives (oxidation state, anion if used, concentrations, etc.) were different, as prior publications generally did not disclose all of those details. We did find significant deuterium loading increases (at least 88%) from Bi, $K_2Sb_2(C_4H_2O_6)_2$, Tl, and Hg (although the latter gave highly inconsistent results). Our highest loadings were achieved with As(III) additives, either as $AsCl_3$ or as $NaAsO_2$; these routinely would give peak loadings of 95%, with a variation of +/-3%; this is contrary to a prior report that As(III) was ineffective as a loading enhancer. These were corroborated in different experiments by several different measurement methods (X-ray diffraction for foils, resistivity for wires, and reflectivity for thin films). However, we found that the As additives must be added promptly **after** the start of electrolysis but also **before** other impurities have had time to plate onto the cathode (e.g., within 20 minutes), and they also have the unfortunate property that the loading decays after a few hours, although this can be refreshed several times with added doses. Analysis of used cathodes revealed that a Pd-As intermetallic of substantial thickness (~1 um) forms on the surface of the cathode. Despite the relatively brief period of high loading, it provided enough time to run calorimetric tests for AHE, all of which gave negative results. We suggest that with further work it may be

possible to find other electrolyte chemistries that allow this high loading to be sustained. This could lead to interesting technological applications, such as an electrochemical generator of ultrahigh pressure hydrogen that would emit from any portions of the Pd that are protected from exposure to the electrolyte.

Calorimetry

The original Fleischmann-Pons work was an open cell, and relied primarily on isoperibolic calorimetry for their claims of excess heat. An open cell is especially problematic when performing calorimetry, as chemical energy in the form of oxyhydrogen gas is continually being vented, and heat is being continually lost through evaporation of the electrolyte and venting of the hot mixture of water vapor and oxyhydrogen. Without accurate measurements of the composition and temperature of the effluent gases (which were apparently not characterized), there is no way to determine these losses with any confidence. As for isoperibolic calorimetry, this method is not well suited to systems where there is a nonuniform internal temperature within the vessel. One can in principle compensate by using a multiplicity of temperature sensors and creating a multivariable mathematical model to calibrate the measurement, but this was not done.

In our work, we avoided open-cell calorimetry except in special situations (e.g., when the electrolyte chemistry produced non-recombinable gases such as N_2 or CO_2). Most of our calorimetry measurements of heat output used closed cells with internal Pt catalytic recombiners in the headspace. We used three different types of calorimeters as we progressed through our project. Initially we used a commercial Setaram MS80 differential calorimeter, which housed two pairs of 35-mm diameter cylindrical sample holders. Each pair would typically consist of one "target" cell and one control cell. Each target would contain a Pd cathode and LiOD electrolyte, and the controls would be nominally identical except that they would contain a LiOH electrolyte in one of the pairs, and a Pt cathode in the other pair. The MS80 is an excellent calorimeter, typically with ~2% accuracy under well controlled conditions. However it is a large unit and we had only one available for our use, which limited experimental throughput, especially considering that it would take approximately a day to thermally equilibrate prior to running a new experiment.

We soon decided that we would need access to multiple calorimeters to run more experiments in parallel. We implemented our own single-ended calorimeter design that we called the "HTCal" (for high-throughput); Fig. 2. The primary heat conduction path was through a thermal bar that was instrumented with 5 thermocouples and attached to a temperature-controlled cold plate. In steady state, approximately 87% of the sample's heat output transferred through the thermal bar; the remaining heat conducted mostly through foam insulation to the ambient enclosure that was temperature-controlled to

better than 0.1 K. Calibration of each cell was done using an internal heater resistor, both before and after an experiment to test for calorimeter drift. The HTCal had the benefit that the heat measurement did not rely on any temperature sensors internal to the cell whose readings might be corrupted by the presence of hot corrosive electrolyte. The multiple thermocouples also provided useful redundancy, as a faulty sensor could be recognized by software and excluded from the calculation. We did also have some internal temperature sensors (RTDs) which provided valuable corroboration of any heat signals that we observed, but they had a non-trivial failure rate in that environment despite careful encapsulation. Ultimately 26 HTCals were built, usually in clusters of four so that we could run four nominally identical cells that differed only in a single variable, such as a control cell. The HTCal's accuracy and stability in steady state was typically +/-2%, and any excess heat measurement of 5% or more was considered a potentially real AHE signal worth investigating, and 10% excess heat was considered definitive if it did not succumb to a classical explanation (which always turned out to the case). To facilitate comparison of measured heat output with expectations, we measured the thermal impulse response of each HTCal and convolved the input power (and hence the expected heat output in the absence of excess heat) with that impulse response.

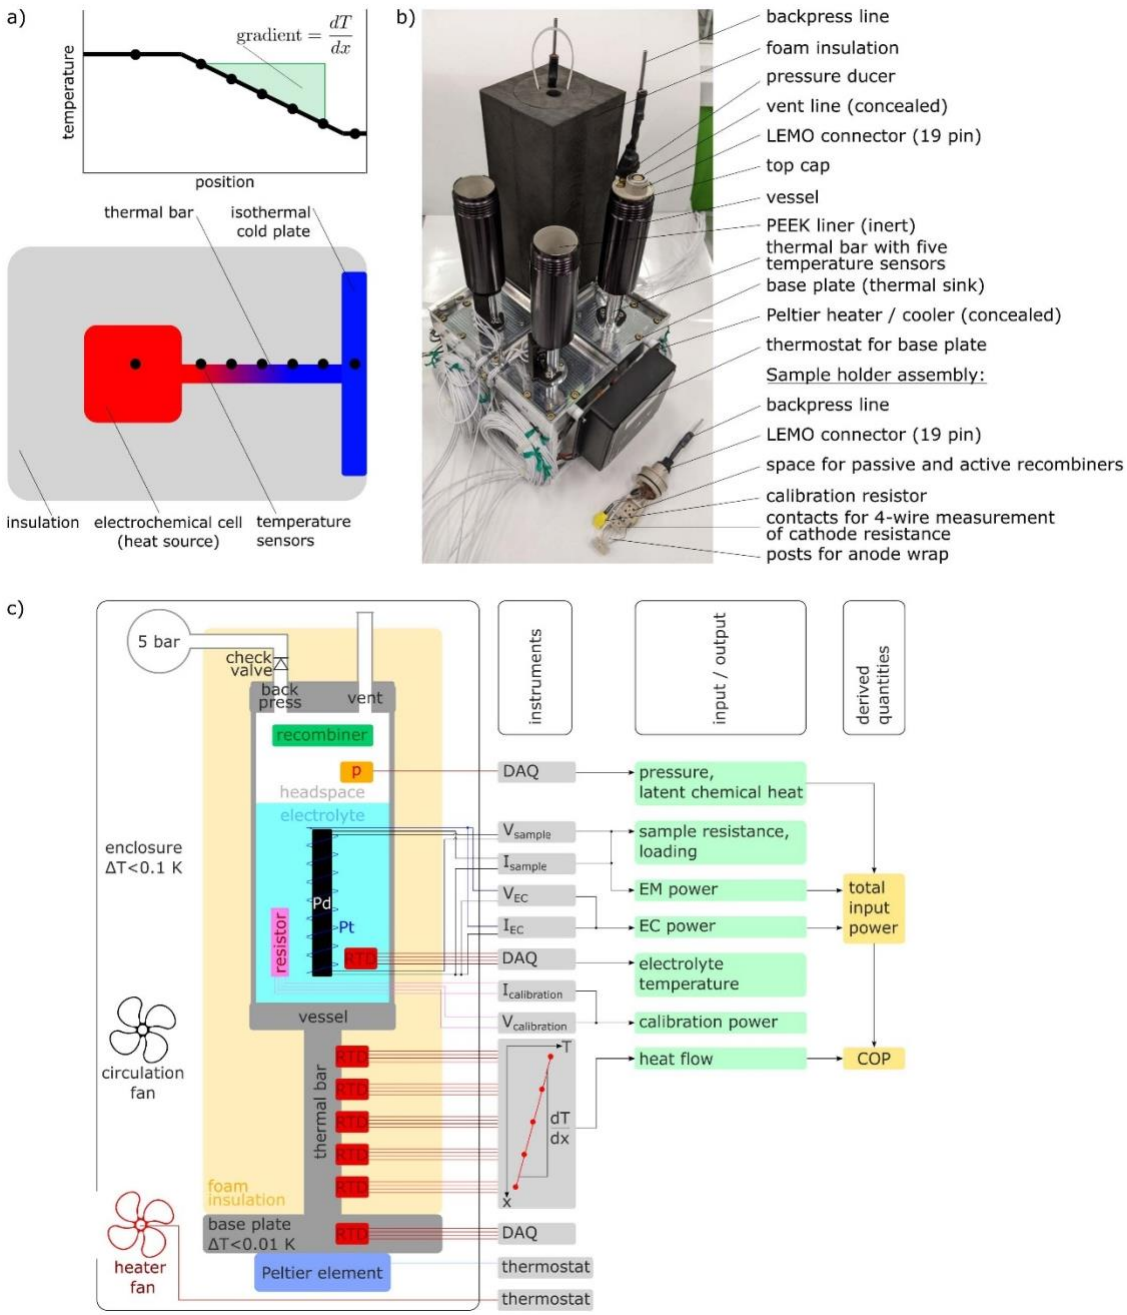

Fig. 2: High-throughput calorimeter (HTcal). a) Heat flow measurement principle of thermal bar (see eq. 2.10). b) Photograph of an HTcal unit with four cells. c) Electrical diagram of HTcal. Note that the vent port is closed during experiments.

Each HTCal contained a pressure sensor which was critically important, as it would indicate how well the Pt catalytic recombiner was functioning, and it also would allow for corrections to the calorimetric data (as it was a measure of stored chemical energy), as well as revealing if the vessel was leaking gas. A recombiner that gets partially wet from spray or contaminated with LiOD solids will become less effective, leading to a higher

partial pressure of oxyhydrogen in the headspace. We would occasionally observe heat bursts that correlated perfectly with a pressure drop that was proportional to the energy released, i.e., the recombiner suddenly became more effective and recombined most of the oxyhydrogen into water vapor, often as a detonation. We speculate that certain claims in the literature of anomalous heat bursts may have been simply this effect.

Towards the end of our program we deployed an ultrahigh throughput ("UHTP") 100-calorimeter array using the isoperibolic method, where we performed additional replications as well as tested various additive combinations. Two consecutive sets of experiments ran continuously for 2-3 months each to test for incubation effects, but with invariably negative results.

By the time our program ended, over 500 different types of experiments had been run, with a total of 200,000 hours of calorimeter time. Almost 90% of the samples and over 95% of the calorimeter hours were accumulated in closed-system calorimeters, which as discussed are significantly more accurate and less error prone than the open system calorimeters used in some of the AHE literature.

Key Replications

While we performed significant replication attempts on the work of nine different previous research groups, we will focus here on two in particular, as they were relatively recent, relatively well documented, and reported relatively large heat outputs (>>10% of input power, and >1 W). These were the works by Staker and by Barham et al.

Staker

McKubre and others have claimed that current flow in certain metal wires can stimulate AHE, and have suggested that the mechanism may be electromigration of deuterons (or protons, some have claimed) within the lattice. In 2019, Prof. Michael Staker published a particularly noteworthy paper, reporting excess heat from an open (i.e., continuously venting) 0.5 M LiOD electrolytic cell using a 0.5 mm diameter Pd wire cathode and Pt anode that had an applied axial current (typically 3 A) in addition to the usual cathodic bias. He chose as a control a second cell, run simultaneously in the same enclosure, which differed in two respects: it used a light water (LiOH) electrolyte and also an inert (i.e., non-loading) Pt cathode. He used the isoperibolic method of calorimetry (two thermocouples in each cell that were averaged, from which the outside room temperature was subtracted) and appears to have taken great care in calibrating the cells. The cells were continually replenished with D_2O or H_2O as appropriate to maintain a constant liquid level.

Staker made physical arguments suggesting that large concentration gradients of deuterons could be created by his electromigration current. To our knowledge, no one has previously experimentally tested the hypothesis that an axial cathode current can significantly affect hydrogen loading in an electrolytically loaded Pd, so we performed our own experiments by adding an extra voltage sense point at the midway point along the Pd wire cathode, allowing us to compare resistances of the two halves of the wire. Within our measurement error, we found no significant differences in the two resistances (hence we infer no measurable gradient in loading).

Staker reported a sustained modest (<10% of input power) amount of excess heat from the deuterated cell, based on his calibration curves. More spectacular was the report of a single cell that exhibited a step function in apparent heat after 42 days of operation in response to a very minor change in applied electrochemical current. This was originally reported as a 2.4 W output on 1.2 W power input (i.e., a COP of 2.0). However after one of us (DBT) made inquiries with Staker (who was consistently very gracious in responding) to get additional unpublished details, Staker published a correction stating that it should have read 4.43 W output on 2.66 W input (i.e., a COP of 1.67). This was still a significant result that we felt worth attempting to replicate.

We made many replication attempts of Staker's experiment, in both numbers of distinct experiments and also in their duration. Our first set of experiments used his isoperibolic calorimetry method, which employed a nested-test tube geometry to create a controlled thermal resistance between the electrolyte and the surrounding air. We found that the inferred heat output was very sensitive to the liquid level, and there are obvious heat-transfer considerations as to why this would be the case. Staker was well aware of this and his own calibration data shows that a level change of ~1 mm could lead to errors of multiple percentage points in the estimate of the excess heat fraction. However in our work and at the power levels that Staker reported, we found that the electrolyte could develop a foam layer on the surface with an onset at unpredictable times; this always led to an overestimate in heat output when it occurred. The overestimate was sometimes of order 10%, but in extreme cases of foaming it could be as high as 60%, which is not greatly different than Staker's step-function result. We used continuous video monitoring of our experiments so that we could correlate changes in apparent heat output with foaming. Despite our best efforts, and the use of thorough pre-cleaning of both glassware and all constituent materials, we were unable to suppress the onset of foaming throughout the duration of an experiment. In a discussion, Staker denied that any foaming occurred in his system, so perhaps there was some unknown factor that differed between our materials or our procedures, e.g. trace constituents in the LiOD electrolyte that affected surface tension.

As we could not validate the excess heat claims in our own replication of Staker's open-cell design due to the variability in heat transfer, we began using our closed-cell HTCal calorimeters (with a Pt recombiner in the headspace) to perform further long-duration experiments. The HTCal vessels were made of aluminum and had a thin 0.5 PEEK internal liner which is chemically compatible with the electrolyte. In contrast to Staker's apparatus, the HTCal design is insensitive to liquid level changes or foaming. Fig. 2 shows a typical 90-day experiment in which four cells were run simultaneously (three "active" cells and one inert control); all four indicated zero excess heat (COP=1.00+/-0.02) with very good stability. Power settings were periodically changed in an effort to recreate the step function observed by Staker, with no success. We ran multiple trials of this nature, including one cluster of four cells that ran continuously for a full year with no indications of excess heat from either the active cells or the control.

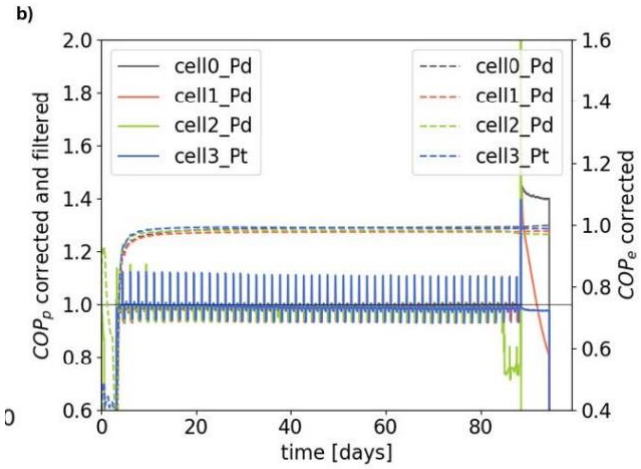

Fig. 2: Staker-like replication in a thermodynamically closed HTCal calorimeter. Cells 0 and 1: nominal Staker replication attempt (Pd with EM); Cell 2: Pd without EM; Cell 3: Pt with EM. COPp=Power COP; COPe=Energy COP. The periodic spikes are artifacts of power changes associated with the thermal time constant of the calorimeter, and not indicative of AHE.

We believe that Staker was a careful experimentalist who made the best of limited resources and accurately reported what he saw, but despite our best efforts were unable to validate the claim of sustained excess heat, let alone observing a step function change of the magnitude reported.

Barham et al (US Navy)

According to public reports, for many years, the US Navy has had an interest in exploring evidence of nuclear reactions (primarily emissions of small numbers of MeV-scale

energetic particles, although the prospect of a new source of nuclear energy had obvious relevance to the Navy's mission) occurring in electrolytic cells, specifically using "co-deposition" wherein Pd deuteride is electroplated onto an inert cathode during the experiment. According to public reports, this research began with work at their SPAWAR laboratory in San Diego, CA, and was the subject of a US patentand follow-up work continued at the NSWC laboratory in Indian Head, MD.

In November 2023, the latter group announced a remarkable claim in a reputable peer-reviewed journal that they had observed statistically significant AHE (5 Watts of excess heat above a 15 W input power, as measure in a Seebeck calorimeter that was open at the top) in relatively high-power electrolytic cells containing LiCl:D2O electrolyte (rather than the traditional LiOD) that was salted with $PdCl_2$ and sometimes $RhCl_3$ to induce co-deposition onto a Pt wire cathode. Two 0.27 T magnets were mounted on opposing sides of the calorimeter vessel to create a magnetic field that was preferably oriented parallel to the predominant electric field direction in the cell.

The November 2023 publication further claimed "Repeatable recipes are shared so that other groups may replicate and extend the present work", so we started our own replication attempts based on these recipes. Although the Navy claimed various electromagnetic emissions and also tracks in CR-39 that could suggest nuclear processes, these were not well correlated with the experiments claiming excess heat, and we chose not to test for those phenomena as they were deemed of minimal importance to our mission compared with finding evidence of excess heat. We built an open Seebeck calorimeter similar to the Navy's photographs. A closed calorimeter would have been a much more reliable instrument, but that would not have been a true replication, and moreover was not a near-term option as our recombiner technology was incompatible with chlorine gas emissions. Although our program was already winding down, we felt this was a particularly important replication, and had time to do over 50 trials, mostly derivative from the handful of runs which had been claimed to produce AHE, particularly the Navy's Experiment #19. We also ran light water controls to compare against the heavy water results (something that the Navy appears to have not done). All experiments were set to provide 15 W input power after reaching steady state, and all of our controls (which were used for calibration of the heat output) showed a maximum variation of +/-0.5 W about the nominal 15 W. None of the "target" (i.e., nominally nuclear active) experiments exceeded this margin of error, hence we concluded that no measurable AHE was occurring, and certainly not the 5 W excess heat that was claimed.

We also found no measurable change in heat output caused by the magnets, regardless of their orientation, and substituting identical degaussed magnets had no discernible effects.

We did however note that the magnets changed the distribution of temperatures within the cell (but not the overall heat output). Specifically, the cell had a homogenizing effect on the electrolyte temperature in the cell during electrolysis, effectively acting as a non-mechanical stirrer. This convective magnetohydrodynamic effect has been dubbed magnetoelectrolysis by some. If one is performing isoperibolic calorimetry (rather than Seebeck calorimetry) with just a single internal temperature probe, one could get the false impression that the magnetic field has affected the heat output. We suggest that prior experimenters who have made claims for beneficial effects on AHE from magnetic fields may have been misled, unaware of this classical explanation for why electrolyte temperatures may change in a magnetic field.

Finally, we point out two questions we have resulting from the Navy's publication. First, as acknowledged in their publication, experiment 25 which was nominally a control experiment (i.e., no $PdCl_2$ included) yielded an alleged excess heat energy, after making a subtraction of an alleged "6 sigma" threshold, of 400 kJ over 40 hours, vs. 45 kJ over 12 hours for their marquee result using $PdCl_2$. That is, the control experiment exhibited a time-averaged excess heat more than 2.5x that of the nominally active experiment. The Navy dismissed this discrepancy as due to "contamination". Second, a close examination of their temperature data showed that all of the experiments claiming excess heat exhibited an elevated temperature near the inert Pt anode, whereas they should have actually seen the elevated temperatures near the PdD-coated cathode, as that is the supposed nuclear-active material.

Other replications

In all, the Callisto Electrochemical team performed over 500 replication attempt experiments. Some of them were attempted to be replicated exactly, down to sourcing vintage components based on correspondence with the original scientists. In other cases, critical elements were replicated, but existing infrastructures (e.g., sample holders, calorimeters) were used. We found no signs of excess heat. In addition to the previously discussed work of Staker and Barham et al, we briefly mention here several other replications that we spent considerable time on.

Michael McKubre's work was noteworthy for its duration and documentation. He used various additives to achieve higher loading, incorporated electromigration, and claimed that long incubation times (>1 week) were often required. However, after reporting excess heat in a 1994 paper, in 1998 he allowed that the calorimetry was "seriously inaccurate" under the given electromigration conditions and concluded that "the presumption of repeatable excess heat production was premature". We agree, as none of our replication attempts showed compelling evidence of excess heat.

Prof. Graham Hubler published a report in 2022 of excess heat from 9 nm Pd nanoparticles in an electrochemical cell using a mass flow calorimeter. He reported excess power of 90 mW on 42 mW of input power, i.e., a COP of 1.9, occurring within 10 hours of the start of the experiment. However there were inconsistencies in his data, in particular a low cell temperature during periods of reported excess heat, suggesting that the cell was not the heat source. We nevertheless attempted our own replication. In view of the relatively low power levels, we used our most sensitive calorimeter (the Setaram MS80 differential calorimeter), and ran two different variations for 200 hours. No differential heat was observed vs. inert (Pt) control cells that exceeded the calorimeter's noise level of ~5 mW.

Melvin Miles of China Lake Naval Weapons center has published papers reporting excess heat of order 50-100 mW in open cells using isoperibolic calorimetry. The cathodes were Pd rods that had been alloyed with boron (0.50-0.75 wt%). We had our own alloy cathodes fabricated by a commercial vendor and performed multiple attempts at replication in our HTCal units, which are closed systems and hence less subject to errors. We found no evidence of excess heat.

Dennis Letts and collaborators have reported excess heat from laser stimulation of electrolytically loaded Pd:D foil cathodes. The experiment of particular interest to us was a paper with Peter Hagelstein that reported an experiment wherein two red diode lasers each of 30 mW output were tuned slightly differently to create a beat frequency modulation of the laser intensity incident on the cathode. Polarization was said to be important, as was the application of a magnetic field. Excess heat levels as high as 300 mW seemingly appeared at certain beat frequencies between 6 and 23 THz, which were believed to correlate with optical phonon band edges. Despite the fact that two prior researchers have reported being unable to replicate the effect, we found this experiment intriguing because it could provide insight into the underlying physics of the (alleged) AHE phenomenon. Moreover because electrolytic power and laser powers are all held steady during the experiment, a correlation between heat output and beat frequency would be detectable in our HTCal (modified to incorporate an opening in the foam and an optical window in the sample holder) even at levels as low as 10 mW, which is a small fraction of the claimed excess heat. Letts was very gracious in answering **our** inquiries and even provided two foils that had been partially processed in accordance with his recipes. Nevertheless, our own multiple attempts at replication, with typical run times of 1400 hours each, failed to show any detectable effect.

Storms has long made claims of excess heat, and had an excellent Seebeck calorimeter design. However most of his work used gas loading of samples at elevated temperatures,

which was outside of our focus on electrolytic systems. We made limited attempts to perform electrolytic experiments on nanoporous pressed Pd samples prepared in accordance with published materials and an online tutorial video, but found no excess heat. Storms later told us that key information existed outside of his publicly available materials, but by this time we were winding up our efforts so did not pursue the matter further.

Energetics was an Israeli company that in 2004 claimed to have measured 20 W of output heat on less than 1 W input power from an electrolytic Pd:D cell, purportedly through the use of multi-frequency stimulation (so-called "Superwaves"). McKubre et al has previously claimed a successful replication. Although we were skeptical of such claims, especially given that there has been no follow-up in the last 15 years, we nevertheless ran a series of experiments intending to create Superwaves (as best we could understand them from the limited documentation). We found no evidence of excess heat in our own experiments.

False positive signals

When we began our search for AHE, our attitude was that with episodic claims of positive results in the literature from multiple different experimenters, it seemed reasonable to suppose that there could be a real effect, but one that was difficult to replicate, especially as most of the experimenters had limited resources. As our own efforts proceeded, there were numerous instances where we observed what first appeared to be a positive AHE signal. However in each case, further diligent study revealed the signal to have been an artifact, either with certainty or with very high probability (the latter case would be, for example, an experimental artifact that could not be repeated on demand but matched the signature of a plausible effect that had been known to occur previously, such as an intermittent contact). There were a variety of different ways in which these false positives would manifest themselves, and each investigation would often require substantially more effort than the original experiment. We conjecture that such false positives likely occurred with many (if not all) of the prior reports of AHE, but due to limited resources or confirmation bias they were not recognized as such. Our conjecture is supported by the fact that most prior experimenters with credible scientific credentials reported only a few positive AHE signals of significance (i.e., producing more than a few % of excess heat, which in our view is the likely minimum error margin of most of the calorimetry methods used in the literature when one demands high statistical significance, our own methods included).

For the benefit of those who still wish to pursue experimental searches of AHE, we compile here a list of the various sources of false signals that we encountered:

1. **Incorrect accounting in COP calculations due to input power measurement error.** The late eminent physicist Richard Garwin, in a 2009 interview, cited this as a likely error in his review of the AHE literature. This can be due to contact resistance that was mischaracterized or developed over time, or electrical sneak paths that were not accounted for (e.g., from deterioration of wire insulation during long-duration tests). Reliance on source-measurement units to self-report voltage or current without independent corroborating measurements can also lead to errors. High-frequency inputs (AC or pulsed) are particularly prone to errors in power calculation. Another source of error is the measurement of input voltage at a point on a wire that is not at the thermal system boundary. This is especially a concern for wires carrying large current densities (e.g., in electromigration experiments). This can cause any level of COP error, in both thermodynamically open and closed calorimeters. When there is ohmic heating in the feed wires themselves, it becomes difficult to determine which fraction of the heat enters the calorimeter (and hence is accounted for), and which fraction conducts externally to the ambient through the wires. A multivariable calibration of the calorimeter (i.e., where both electrolytic power and electromigration current are varied independently) is required in such situations.

2. **Changes to the heat-transfer coefficient of the calorimeter (output heat measurement error).**

    - Changes to liquid level and electrolyte foaming: In certain thermodynamically open calorimeters, liquid level variations and electrolyte foaming both cause changes to the thermal transport properties of the calorimeter. We observed consequent misleading COPs as high as 1.6 in our Staker replication attempt experiments.
    - Environmental temperature instability: Good calorimetry requires temperature stability of the surrounding enclosure(s) on the order 100 mK for valid, high-precision measurements. When the surrounding environment temperature changes, the heat transfer of the calorimeter will be impacted.
    - Calibration drift: Heat flow measurements often rely on accurately and precisely measuring temperature, or some other proxy for heat flow. Whether using thermocouples, RTDs, or Seebeck sensors, all are susceptible to component failure, placement errors, and offsets which can drift over time.

- Insufficient understanding of reactor design: We characterized temperature heterogeneity in some replicated isoperibolic calorimeters. Certain designs can lead to misleading COPs on the order of 1.3.

3. **Unlocking stored energy via chemical reactions:** We have observed at least two types of chemical reactions in the headspace to date: Oxyhydrogen recombination (very common) and combustion of PEEK (rare but can occur with large-volume Pd samples that produce substantial orphaned oxygen). These reactions produce large near-instantaneous heat releases, which can present as a temporarily elevated COP. There were also indications that slower timescale corrosion reactions were releasing low levels of heat in some of our experiments.

4. **One-off experimental errors:** Equipment malfunctions, calculation errors, incorrect analysis code, etc. can lead to misleading COPs. This is especially true if sensors are non-redundant and experiments involve manual steps (such as manually changing currents or manually reading voltmeters).

Although not a part of the Callisto Electrochemical research effort, we are aware of the use of optical pyrometry as an alternative method for detecting excess heat. The parallel effort by the Callisto group used this method at times, and it can be subject to errors if the emissivity vs. wavelength $\varepsilon(\lambda)$ is not perfectly calibrated (or if it changes with time, as can occur during loading). We mention this because there are other currently active research groups who have reported AHE in high-temperature gas-loaded experiments where we believe that $\varepsilon(\lambda)$ was varying to some extent during the experiments (i.e., due to alloying, surface roughening, etc.). Moreover some of those groups also use a radiant lamp as a source of input power, and in such situations the absorbed heat (input power) to the sample will not be stable if $\varepsilon(\lambda)$ is time-varying, which can also lead to false conclusions. We have not seen this issue given the level of attention that we believe it deserves in the work by those groups.

Closing remarks

We recently passed the 30th anniversary of Fleischmann and Pons' original announcement of excess heat. We cannot prove that AHE does not exist, as it is always possible that there is some yet-undiscovered combination of variables that Fleischmann and Pons encountered by happenstance. However in the absence of any published recipe that produces replicable excess heat in the hands of qualified independent experimenters, as well as the lack of compelling physical theories, our opinion after years of diligent and well-resourced experimentation is that the original "discovery" was simply an error in energy

accounting, which is all too easy to make, and subsequent claims of positive results most likely have been similarly flawed.