Adventures in EM Side-channel Attacks

NCC Group - Research and Internal Projects

October 24, 2025

Prepared By
Aaron Kondziela

1 Adventures in EM Side-channel Attacks

Introduction

The Hardware and Embedded Systems practice of NCC Group had cause to delve into some particularly crafty EM side-channel work in late 2024, thanks to a request to duplicate the Eucleak vulnerability. The original research (https://ninjalab.io/eucleak/) performed by Thomas Roche and others at NinjaLab in Quebec, demonstrated a long-standing subtle information leak in the Infineon secure element used in the YubiKey 5 Series products, and many other devices using that chip. By exploiting it, they demonstrated it was possible to extract the secret ECDSA key, which would enable cloning of the hardware token, and entirely destroy that root of trust. Thankfully, Infineon quickly patched the vulnerability, and YubiKeys with firmware 5.7+ are no longer vulnerable. Our task was to replicate this attack.

Originally, the Eucleak work was a two-year research project, with many dead-ends, ultimately exploiting extremely subtle timing signal leaks within a very complex system. We had two weeks to pull it off. Furthermore, the original researchers had access to commercially-available micro-H-field probes; but given our timeline, these were not an option for our work. However, H-field probes are fundamentally simple things, and we chose to tackle the challenge regardless, and fabricate our own probing solution. This is the story of that adventure.

H-Field Probes

The probes involved in this research use the most basic principle of electromagnetism - where there is an electric current, there will also be a magnetic field. And, when the electric current changes, so does the magnetic field. Likewise, this changing magnetic field will induce an electric current in a nearby conductor.

H-field probes exploit this effect using a simple loop of wire. Typically, they are constructed using a single turn of wire at the end of a bit of coaxial cable, or the equivalent embodied in a PCB. The center conductor forms the loop, and the outer shield is extended around the loop for shielding against electric fields, having a break in the middle to prevent it from becoming, effectively, another loop of wire, but one which is shorted out and would only serve to dampen the magnetic field reception.

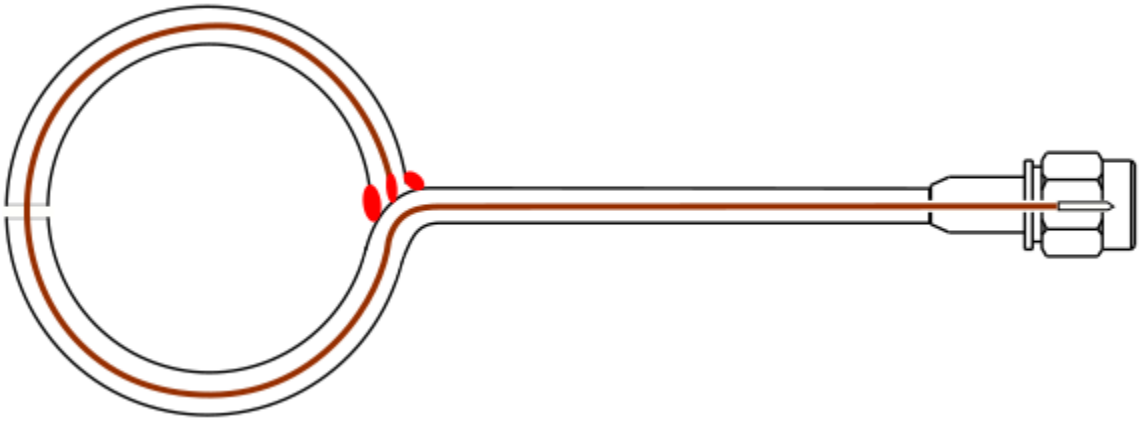

Figure 1: H-field probe illustration

The commercial micro-probes used in the Eucleak research were from Langer, one of the few commercial offerings of micro-H-field probes. They are well-made, incorporating front-end amplifiers and well-controlled frequency response. In particular, a 500-micron (0.5mm) probe was used for the Eucleak research. While this is quite small, it does not require unattainable equipment to produce, so we chose to create a probe of similar characteristics.

Probe Fabrication

We began with a length of 1mm diameter 50-ohm coaxial cable, the size everyone has lying around.

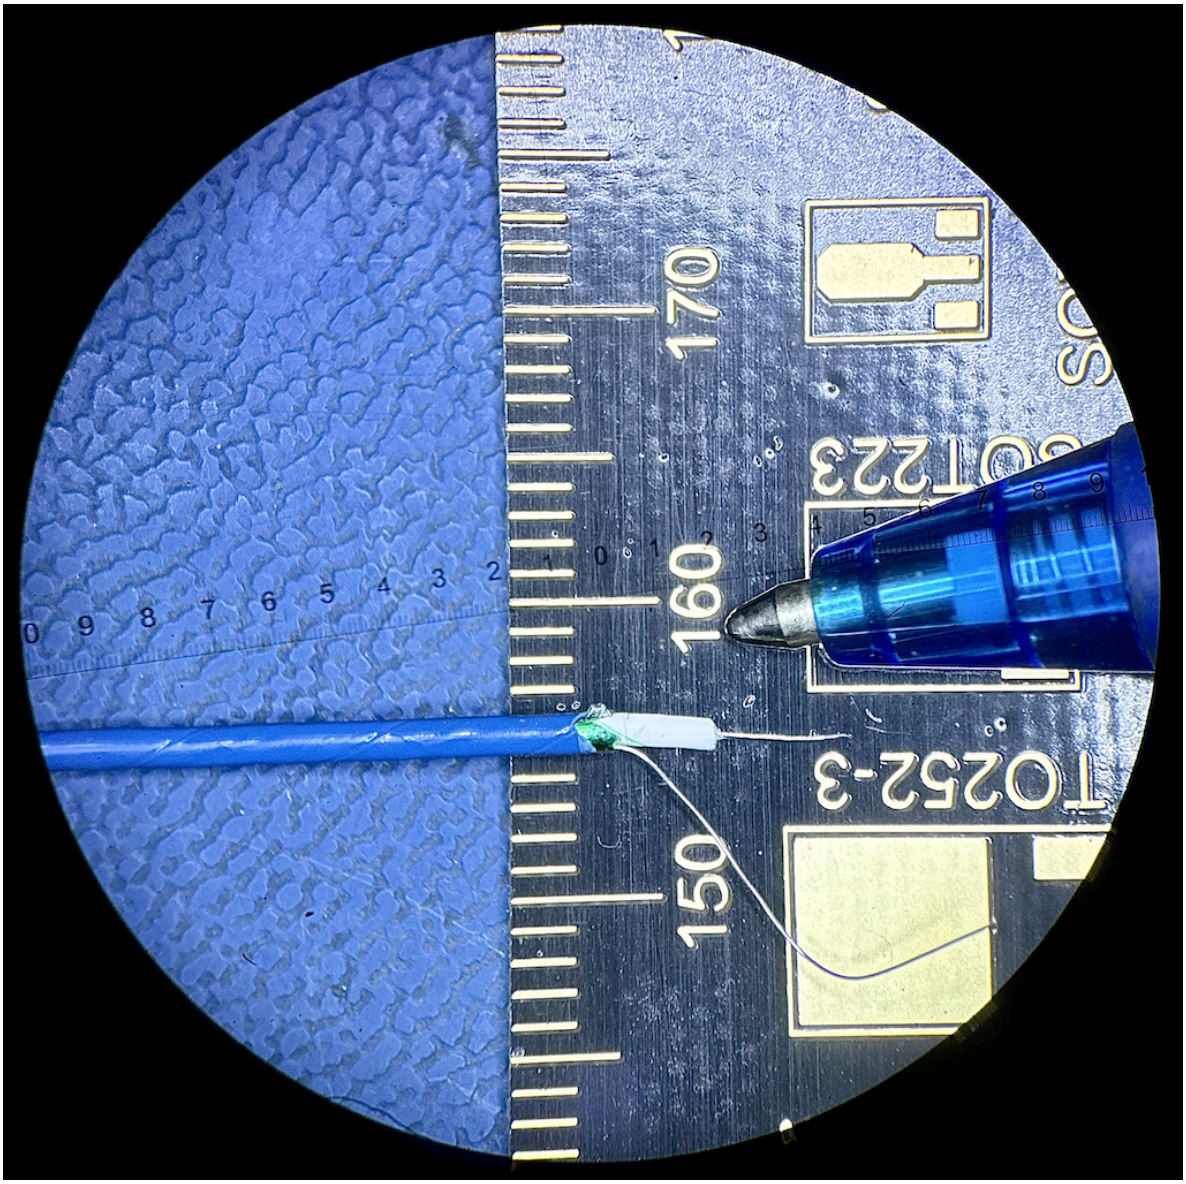

Figure 2: 1mm micro coaxial cable

A loop was fabricated from the center conductor, using a 500-micron PCB via drill as a mandrel. This was micro-soldered to the outer shield of the coax, using a microscope and a steady hand.

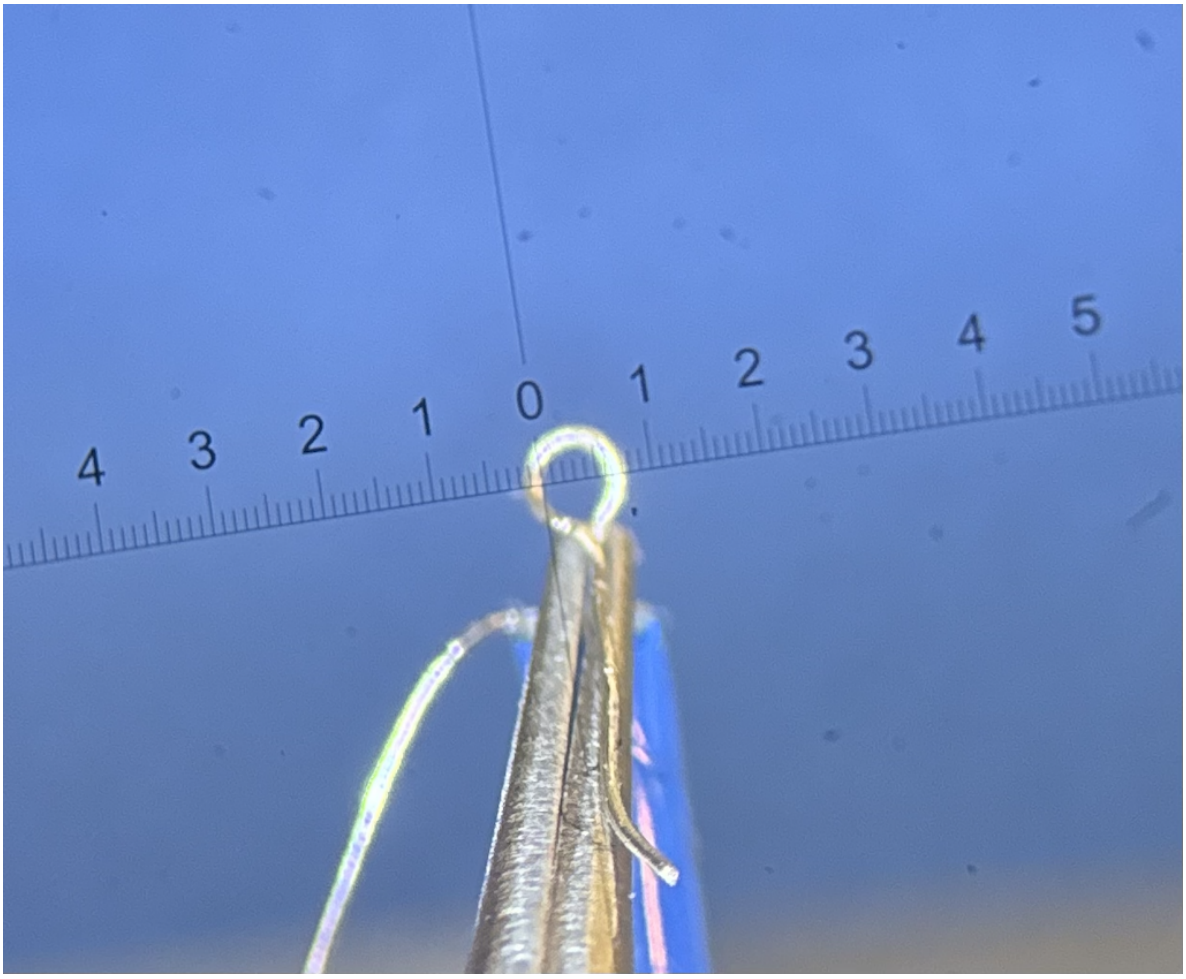

Figure 3: Forming the loop

After a bit of massaging with tweezers and surgical blades, a suitable 500-micron loop was formed at the end of the coax. We decided not to bother with E-field shielding at this scale, as the given application should not bear too strongly on that source of noise. If we needed to, however, plans were in place to insulate the loop with a specific electrical varnish, and coat that with silver-conductive paint as a shield, then laser-ablate a notch therein, to avoid the shorted-loop concern.

Figure 4: Fingerprint for scale; 1 millimeter reticle

Chassis Fabrication

The probe tip itself is, understandably, quite delicate, and requires physical protection. It also requires electrical protection – The coax probe must feed into an LNA (low-noise amplifier), and the entire assembly, having significant gain, is susceptible to interference, and needs stout shielding. For shielding and physical support, the 1mm coaxial cable was fed through a length of very small copper tube, salvaged from a discarded HVAC thermocouple assembly. This tube was attached to an alloy chassis, wherein the first-stage LNA and power distribution would be housed. The appropriate machining work was performed, and components mounted, resulting in a functional front-end to the system.

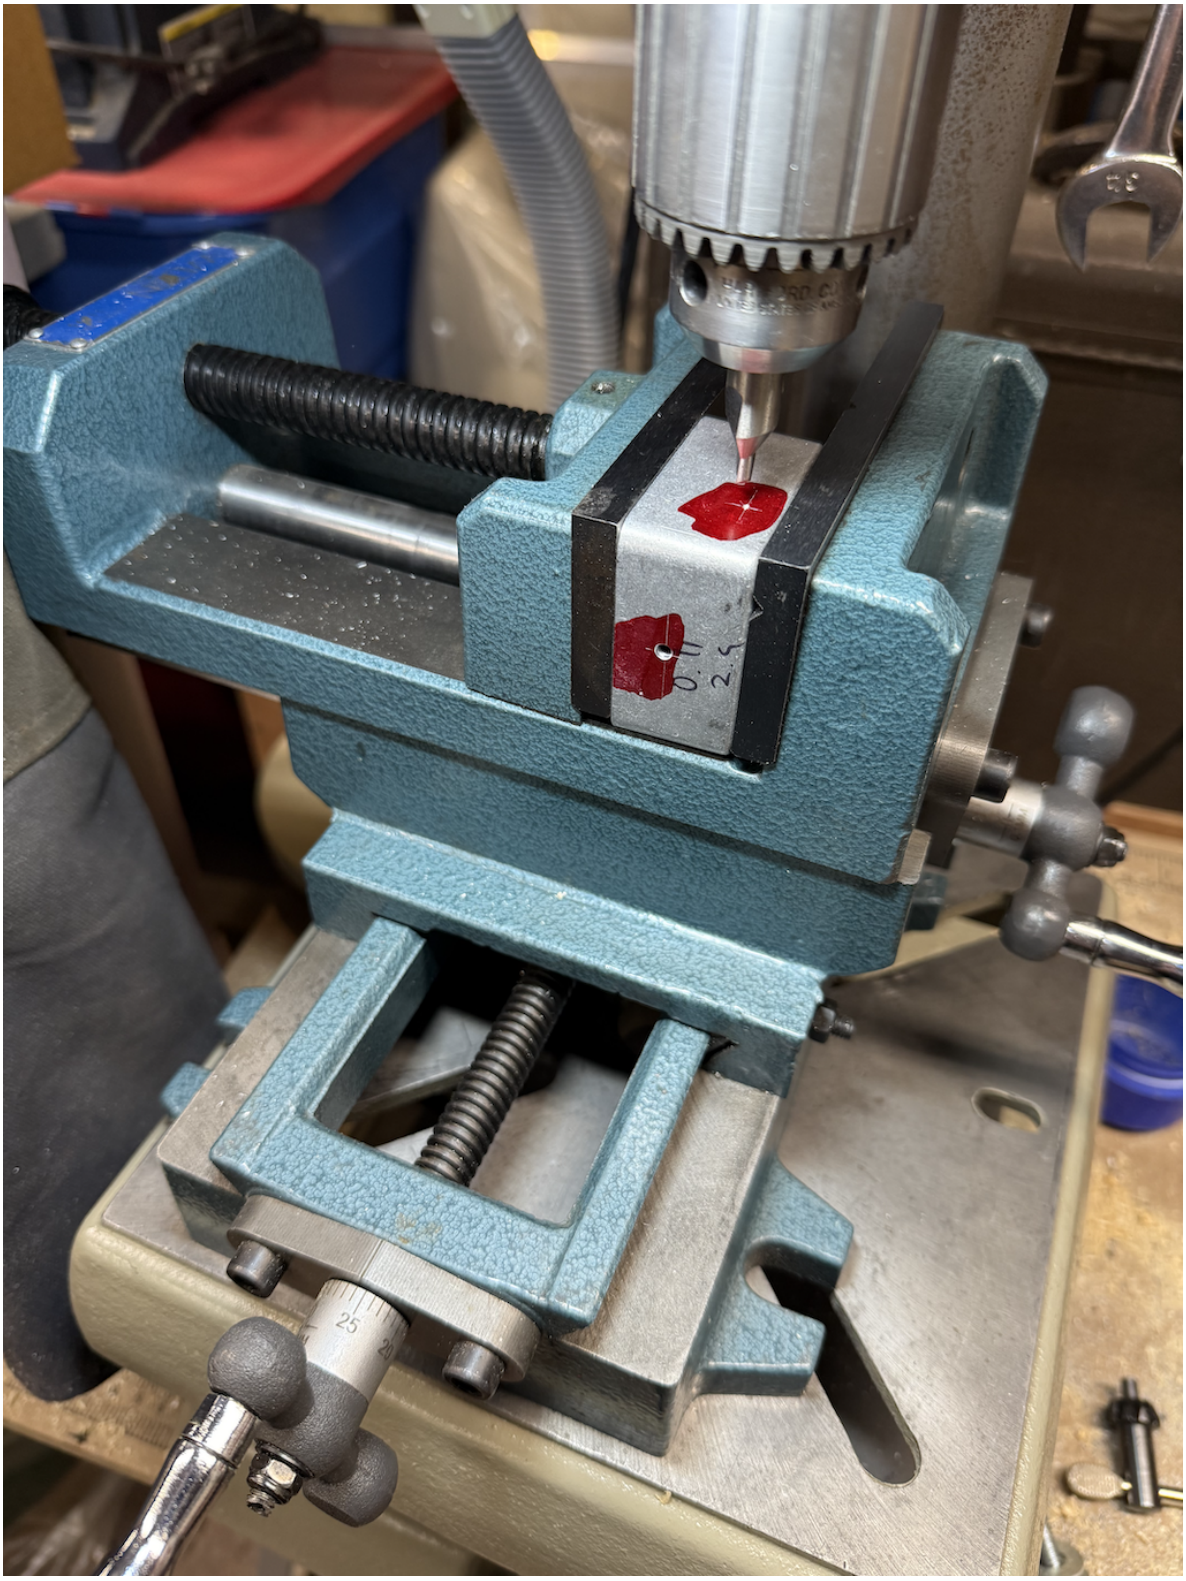

Figure 5: Chassis fabrication

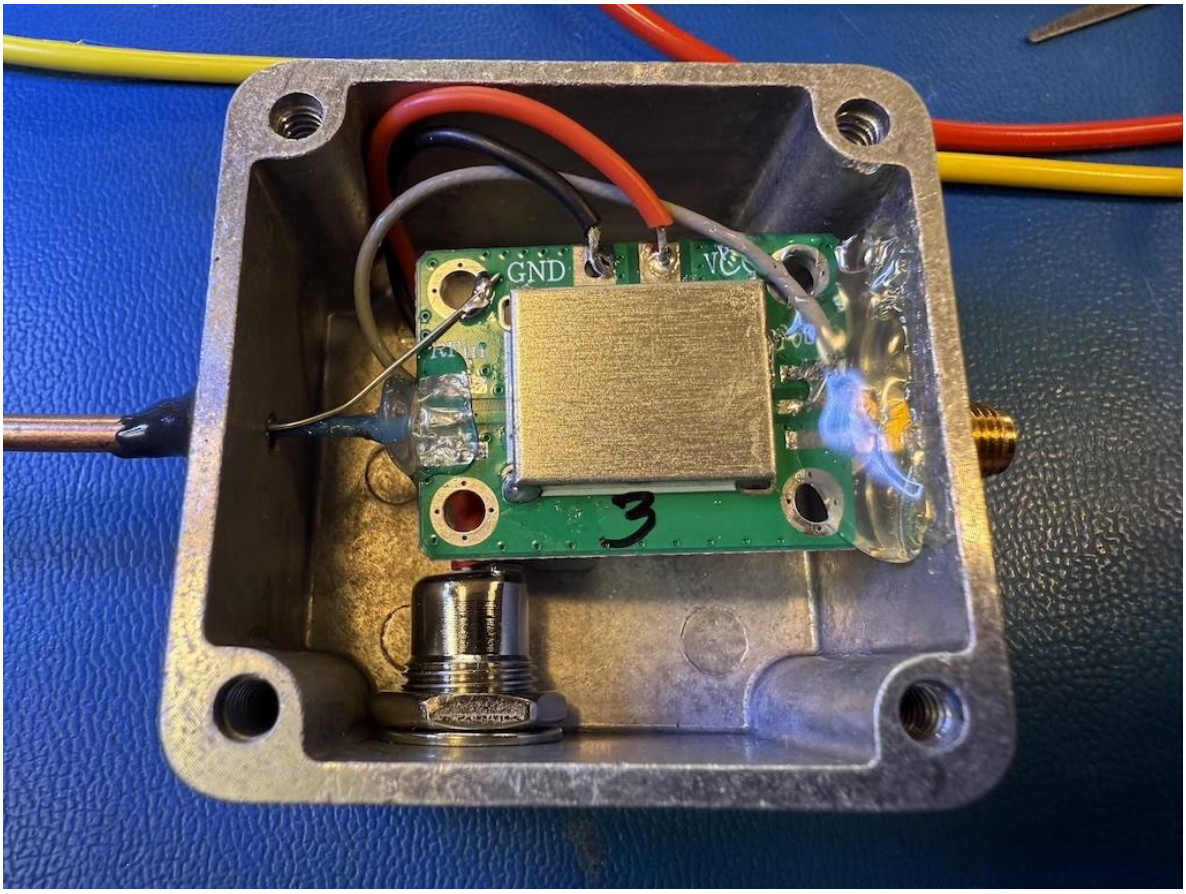

Figure 6: Internal components

Test Fixture Setup

Sniffing the electromagnetic signals emanating from a chip requires extremely precise positioning of any probes, as the structures on the IC are vanishingly small. While we are not trying to intercept a specific signal from a specific wire inside the chip (instead we are sensing the aggregate EM field of many conductors), we still need careful and repeatable control over where the probe is positioned.

For this, we used an XYZR micropositioning stage. Fine control over 3 axes of translation and one axis of rotation are made via controls on the stage, similar to a machinist's micrometer. This stage was affixed to a metal platform which had a vertical post to which equipment could be secured.

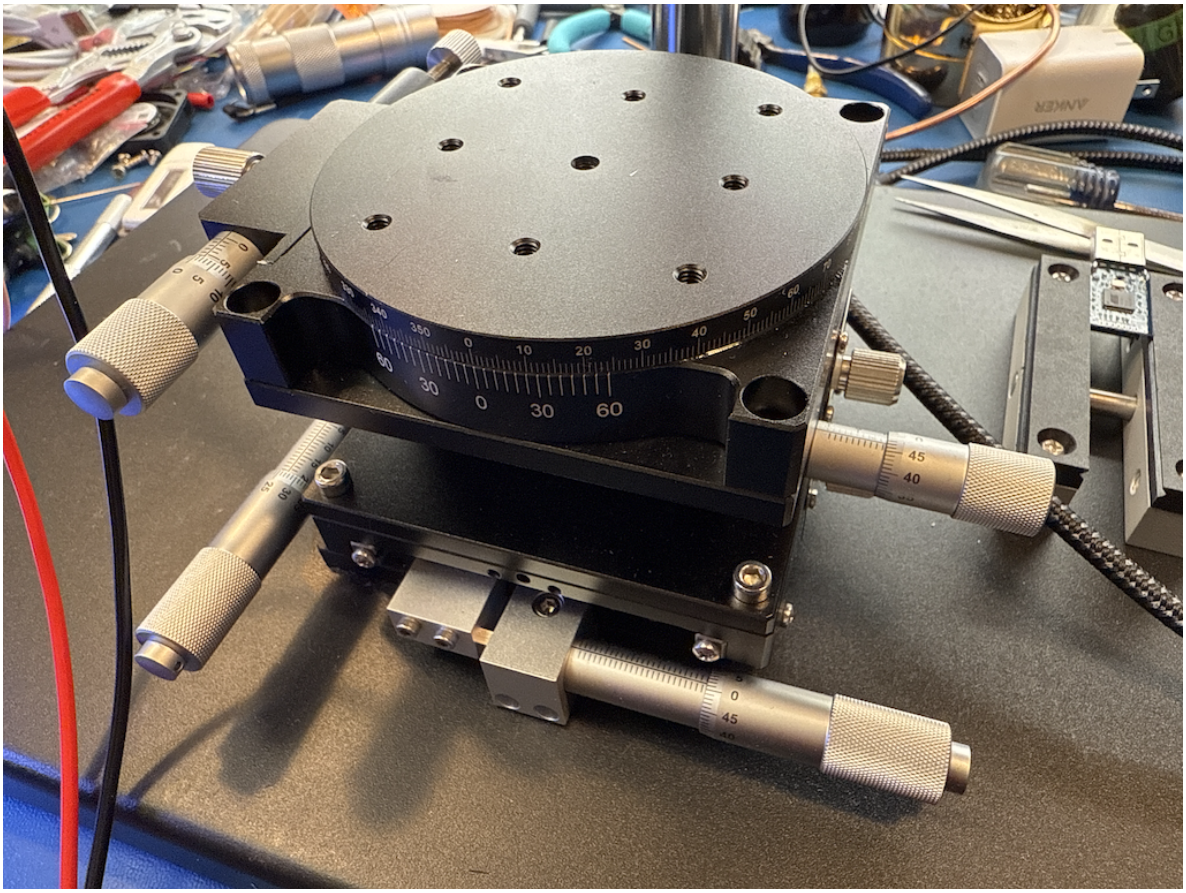

Figure 7: XYZR Micropositioning stage

An articulating arm with clamps allowed for free positioning of the probe assembly in space, and a stereo microscope allowed for close observation. The target device is a tiny security key that fits in a USB-C socket, akin to the YubiKey Nano. We used a right-angle USB-C adapter and a custom 3D-printed clamp to affix the socket and cable firmly to the micropositioning stage. The device could be inserted and removed as needed while maintaining positioning accuracy to a good degree. A stereo microscope was angled in closely to get a good look at probe alignment.

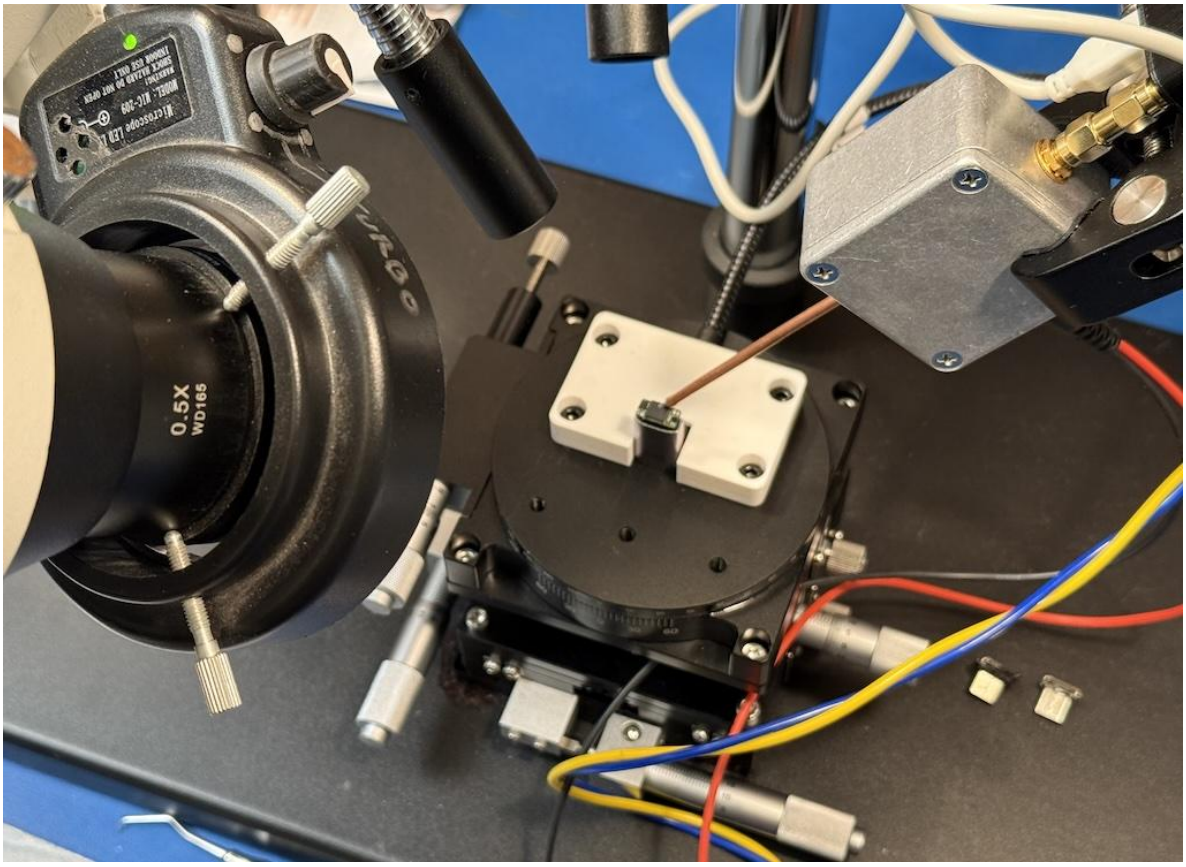

Figure 8: Stage setup

Signal Capture

The process for capturing signals of interest involves several steps in sequence:

- Prerequisite steps
- ◦ Register a new ECDSA FIDO2 credential to the security key
- ◦ Save a copy of the public key for use in the next steps
- ◦ Configure oscilloscope amplitude and timing settings as appropriate
- Signal capture
- ◦ Set oscilloscope to single-trace capture, and stopped
- ◦ Send a FIDO verification request to the security key
- ◦ Wait for the key to ask for the user to touch-activate it
- ◦ Activate the simulated user touch relay
- ◦ Delay for a set time
- ◦ Activate the scope trigger input
- ◦ Wait for scope to capture a trace
- ◦ Download the trace data
- ◦ Repeat as needed

Support Electronics

Automating trace captures required some external electronics controlled by software on the host machine. This was needed to simulate the capacitive touch user gesture required to complete a FIDO verification.

We first tried a simple Arduino with a relay shield, hoping to use the relay contacts to attach a length of wire to the cap-touch input to simulate the capacitive load of a user touch gesture. This didn't work, as even the extra parasitic load of the wire from the input pin to the relay module confused the sensor. So, we had to minimize the extra load until the relay was triggered.

We did this using a small reed relay module mounted directly on the security key, with one contact soldered to the cap-touch input, and the other contact having roughly 30cm of wire attached, leading to nothing. This worked reliably to simulate a user touch gesture when the relay was closed.

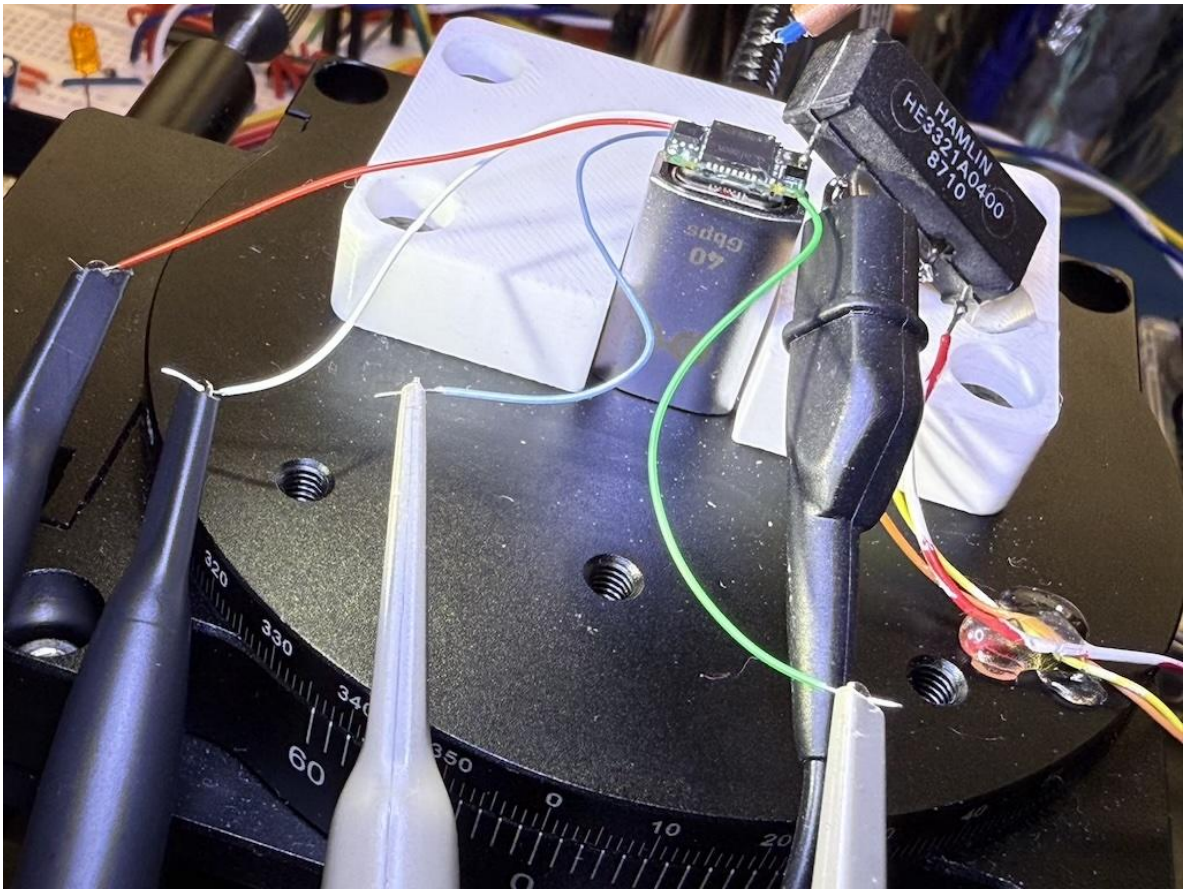

Figure 9: Reed relay setup

Unfortunately, we were extremely crunched for time at this point, so we wired the reed relay's coil into one of the relays on the relay shield. The mechanical nature of relays created a small unpredictable jitter in the timing between the activation signal and the relay closing. This complicated the signal capture process quite a bit, and it took several activations to get usable traces. Additionally, we discovered an input on the device used for factory testing which directly activated the touch circuit without having to simulate the human body. This would have been a far better choice both for ease of implementation and timing consistency. Hindsight, as always, is 20-20.

The next module added was a second Arduino with a knob and an LED display, to serve as an adjustable signal delay. This was a workaround for a firmware bug in the oscilloscope. Ideally, the scope would have been set with a trigger delay - once the trigger signal is seen, it waits a specified amount of time, then captures a sweep. But the feature simply didn't work correctly, and we chose to implement the delay externally. While this could have been done on the Arduino running the relays, again, for expediency, we simply deployed a second

module. It took the signal to the relay shield as input, delayed a set amount of time, and activated the output.

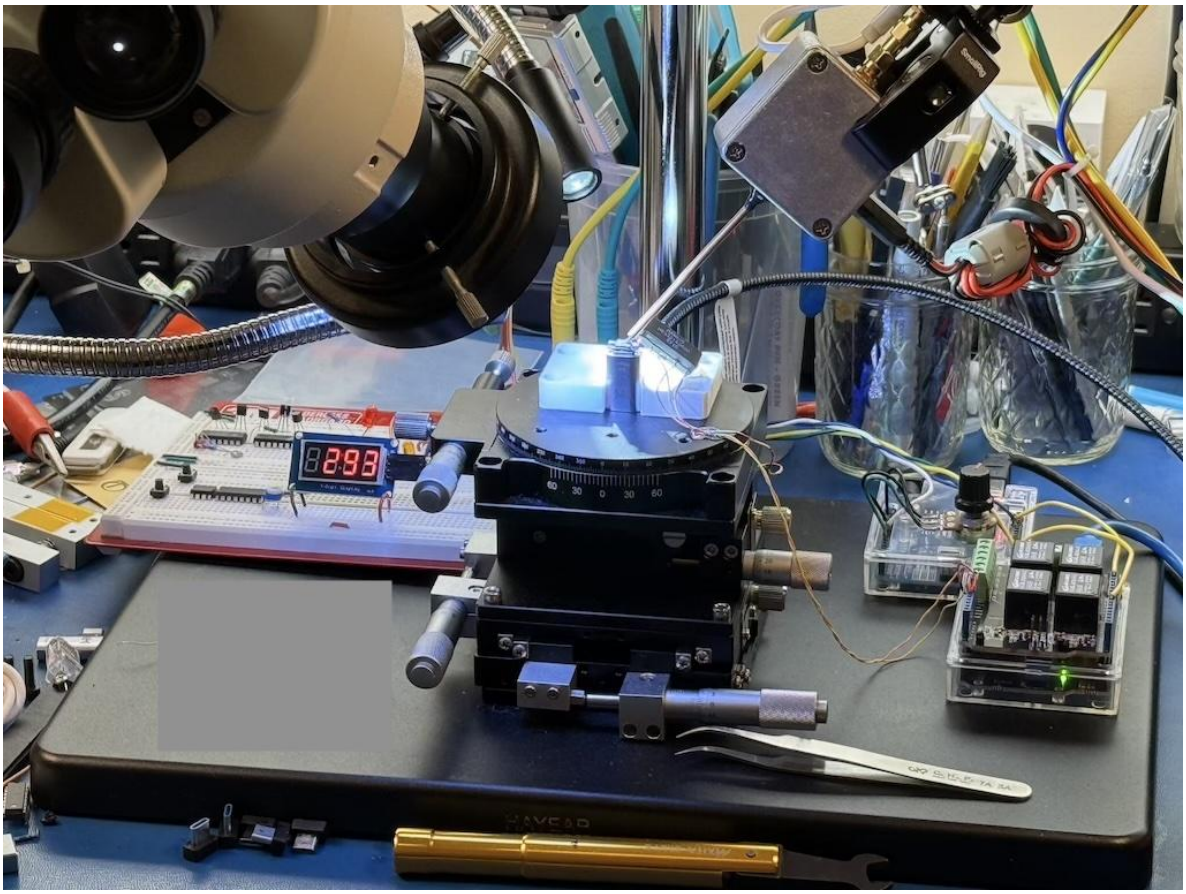

Figure 10: Delay module and improvised differential line driver

The output went through an improvised differential line driver to make use of the logic analyzer input on the scope (a mixed-domain scope), which proved the most reliable way to trigger. The knob and LED display was used to quickly set the time delay and read out the value.

Probe Configuration

The signals coming from the H-field probe are extremely small, and the front-end of this particular oscilloscope does not deal well with these minuscule signals. In order to get a strong enough signal, we needed to use three LNAs for a total boost of about 60dB. There was a major problem, however. Connecting two of our LNA modules together, the system would break into oscillation, and we didn't have time to diagnose why this was. We had two different types of modules, and no combination of them seemed to work, despite them all having good specifications verified with a vector network analyzer sweep prior to use. Thinking we were sunk, we tried a commercially-packaged LNA as the second stage in-between the other LNA modules, and it worked great! We were still curious why they were misbehaving, but again, time was a very pressing factor here. The mystery will be lost to time, unless this engineer gets nerd-baited into figuring it out some day!

The probe was then affixed to the platform via an articulating arm and positioned over the target device. Fine adjustments were made using the XYZR platform and a stereo microscope. We were ready to go.

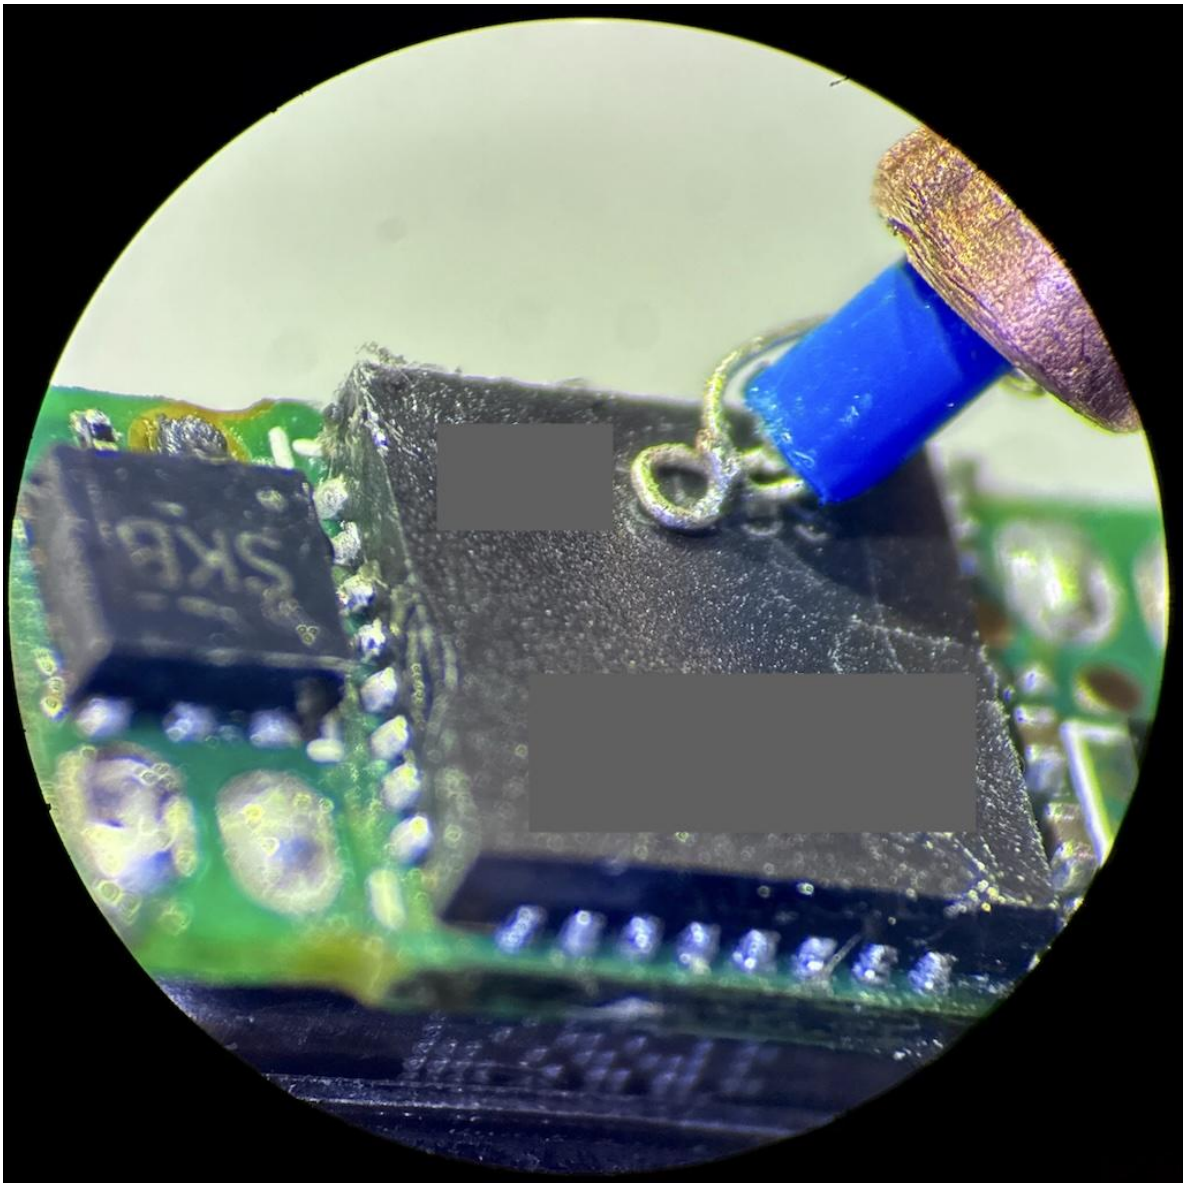

Figure 11: Probe tip over target

Signal Capture

Using our custom automation scripts, we began capturing traces, and very carefully adjusting the probe setup to find the signals of interest. Here we see the full trace of a validation, showing the same areas of interest discovered in the original Eucleak research.

The operations at the tail end employ modular inversion operations, where the code error caused the data leak.

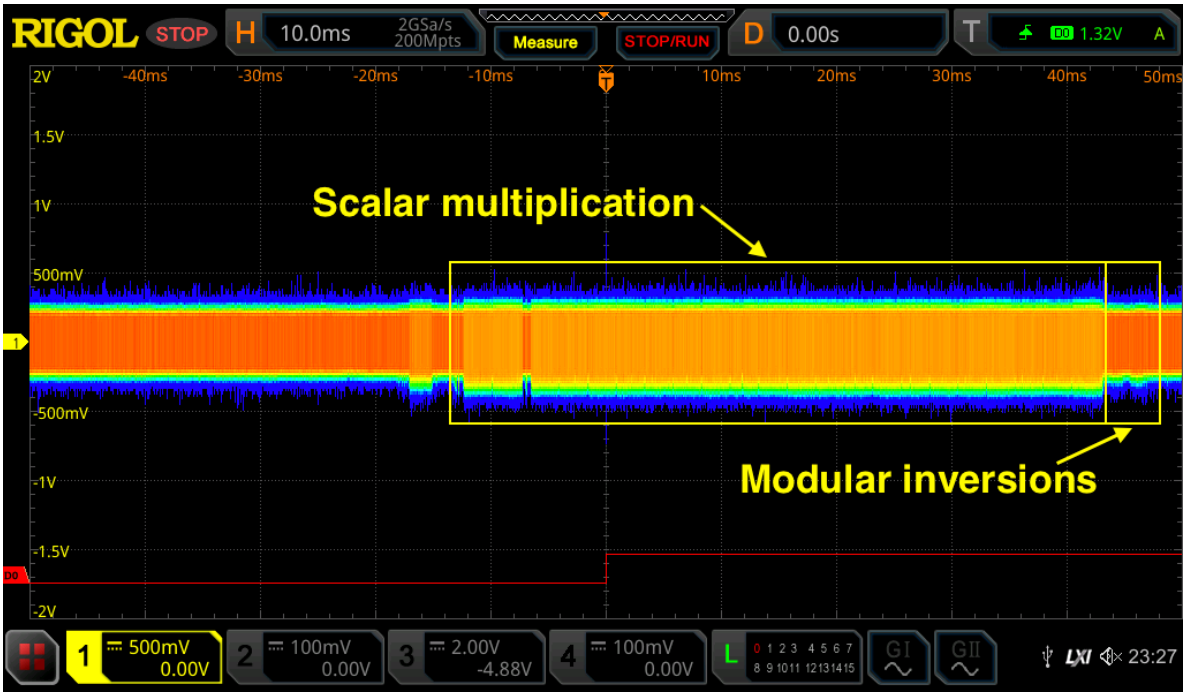

Figure 12: A successful capture!

Project Conclusion

The details which leak the key information are quite subtle, and can be observed when zoomed quite far in. While we were able to make out these features, further work was needed to dial in the equipment to capture them with a smidgen of greater fidelity, and then extensively post-process the data to recover the private keys. At this point we had consumed all of the time available to us and had to end the project.

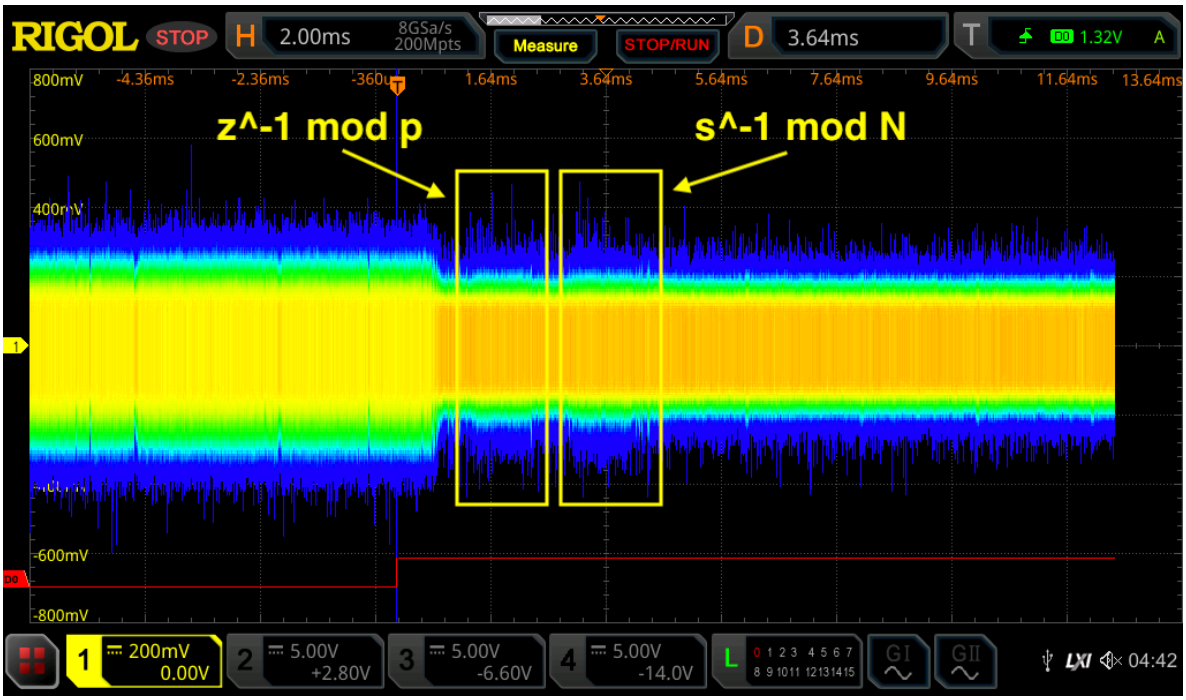

Figure 13: Zoomed trace showing modular inversions

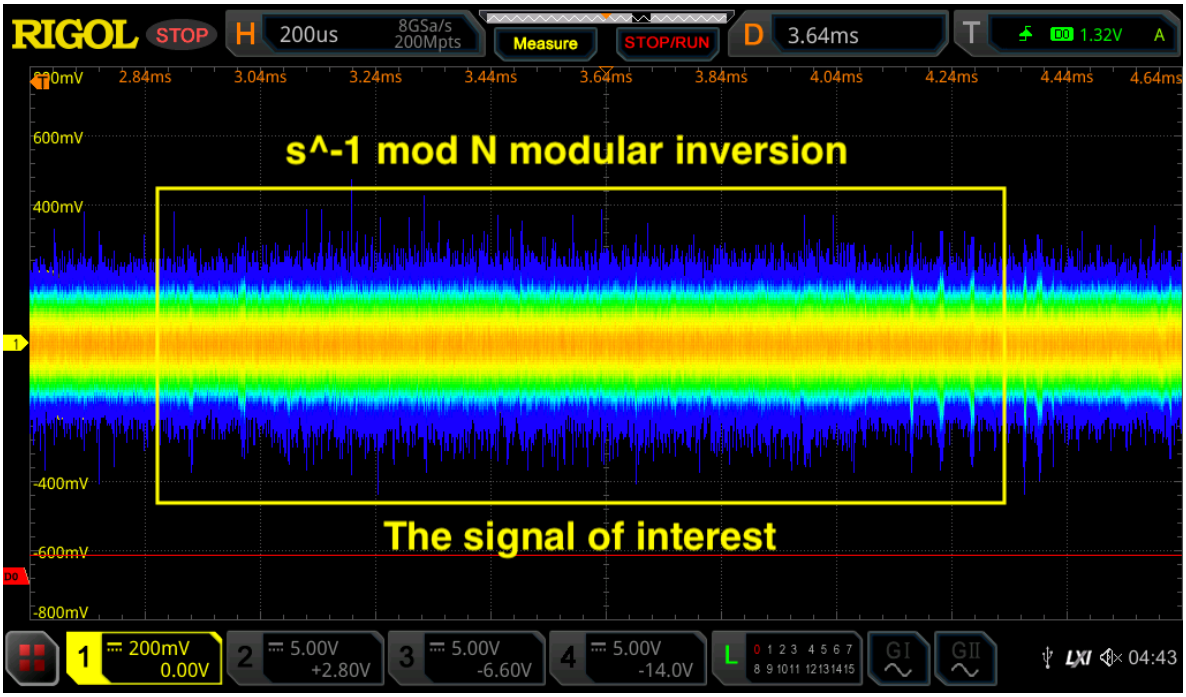

Figure 14: The signal of interest

We consider this a success, as we were able to replicate the signal capture from the original research in only two weeks, using lab-fabricated equipment and a number of expediencies. Employing the same equipment on other targets should be straightforward, and likely easier, given that the Infineon security element is one of the hardest targets out there.

Big thanks to Thomas Roche and others at NinjaLab who performed the original research and made it possible for us to replicate it in such a short period of time.

References

Below are links to some of the devices and equipment used in this project.

- XYZR micropositioning stage: https://www.amazon.com/dp/B07H3NP1L8
- Articulating arm with clamps: https://www.amazon.com/SmallRig-Adjustable-Friction-Articulating-Monitor/dp/B087T4T8D5/
- Nooelec LaNA: https://www.nooelec.com/store/lana.html
- Generic LNA: https://www.aliexpress.us/item/3256809974107292.html (or similar, we had a few models at hand and tested them on a VNA for best performance)
- Hammond 1590MM Enclosure: https://www.hammfg.com/part/1590MM

2 Contact Info

The team from NCC Group has the following primary member:

- Aaron Kondziela – Researcher
  aaron.kondziela@nccgroup.com