Basilisk: Remote Code Execution by Laser Excitation of P–N Junctions Without Insider Assistance

Joe Loughry
Netoir.com
joe@netoir.com

Kasper Rasmussen
University of Oxford
kasper.rasmussen@cs.ox.ac.uk

Abstract

Inadvertent photosensitivity of P–N junctions has been known for a long time, but most of the attacks that have been demonstrated are covert channels, requiring an adversarial presence on the device. We show not only how it is possible for an external attacker to bias a P–N junction with a low power laser, without any kind of insider assistance, but also how this kind of attack can be used to perform logic level attacks on the target device and thus interfere with the device's operation. The technique requires precision but is feasible in practice with off the shelf hardware, as long as the attacker has a line of sight to the target. It can result in attacks that include crashing a computer, change memory contents, alter the instruction stream of a running program, alter messages on a shared communication bus, insert new messages, or prevent communication. Most of these attacks have never been demonstrated before without insider assistance. We demonstrate that under the right circumstances the attack can lead to arbitrary code execution on the target device. We show a working proof of concept including remote code execution, and quantitative measurements leading to testable predictions. Mitigation of this vulnerability is challenging and countermeasures will in most cases require hardware changes.

1 Introduction

Semiconductors are the substrate on which most modern electronics are built. Semiconductors can be "doped", *i.e.*, contaminated, with elements that makes it either positively charged or negatively charged, with the interface between positively- and negatively charged sections called a P–N junction. These junctions are the basic building blocks of all transistors and diodes (and by extension logic gates, IC and LEDs) and are thus present in basically every device.

In some cases these components are exposed to the outside world, either by design in the case of LEDs, or accidentally in case any part of a circuit board can be seen through vent holes or other openings in a device. Exposed P–N junctions that are electrically connected to a shared communication bus are vulnerable to being optically "pumped" by a modulated laser beam. The effect is to reverse the LED, turning it into a current source, or to bridge the diode, turning it into a conductor, thereby affecting the circuit that the component is connected to.

The effect depends on the type of component and the surrounding electric field. In photovoltaic mode, an LED in forward bias will be reversed if sufficient optical power is pumped into it at the right wavelength; this causes the LED to generate a photocurrent that runs backwards through the connected circuits of the computer, driving the connected circuit *high*; in photoconductive mode, *e.g.*, an electrostatic discharge (ESD) protection diode in reverse bias, pumped by an infrared laser, conducts current in the opposite direction, grounding the bus and driving it *low*. The effect is transient, and leaves no evidence behind.

We explored the feasibility of conducting such attacks in practice and the extent to which this effect can be used by an attacker to affect a victim system. We found it is in fact possible to affect P–N junctions using an external light source, in practical conditions. Further more it is possible to achieve a range of different effects depending on what the exposed LED or diode is connected to.

We conducted a number of experiments to determine the parameters for a successful attack. These include the angle at which the laser hits the diode, the power level needed and the modulation of the laser. All experiments are done with low cost devices that are easily available to anyone, *e.g.*, a laser module from a Blu-ray player.

We present two working proofs of concept. First an attack on a live I^2C bus running between commercially available devices, and second, an attack on the CPU–memory bus on a pedagogically minimized CPU. The first shows that we can inject or alter messages on the I^2C bus that will be accepted as legitimate by other devices on the bus. The second gives the attacker arbitrary execution privilege (with some interesting constraints) on the CPU.

Mitigation of the vulnerability is not straightforward. In

most cases countermeasures will require hardware changes and we discuss when and how this can be done in practice.

2 Photosensitivity of P–N junctions

Semiconductor P–N junctions are photosensitive because photons generate electron–hole pairs when they hit a semiconductor under the right conditions, specifically the depletion layer that forms in the presence of an electric field between the p-type and n-type doped regions. Electron–hole pairs form when a photon with the right amount of energy is absorbed by a semiconductor atom [84, p. 80]. In zero-bias mode, the photocurrent generated is proportional to irradiance on the P–N junction [86, Chapter 6, p. 238].

The result of this conversion depends on the electric field, *i.e.*, how the junction is biased. If the P–N junction, say in an LED, is zero biased and illuminated by a suitable wavelength, it goes into what is called photovoltaic mode. Electrons are swept towards the anode, and holes are swept towards the cathode. This makes the cathode positive with respect to the anode, sending an electric current through any circuit attached to the LED but in the opposite direction from the way that usually makes the LED light up. The same thing happens if the LED is forward biased, *i.e.*, lit.

If the P–N junction is reverse biased, as in an ESD protection diode on a shared bus, it goes instead into photoconductive mode when illuminated by the laser. Here, electrons are swept by the electric field towards the cathode, immediately recombining with holes there, lowering the resistance to electric current and making the P–N junction into a conductor. Because the depletion layer is widened by the reverse bias voltage, making a larger volume where electron–hole pairs may be created by absorption of photons, photoconductive mode is more sensitive than photovoltaic mode.

The semiconductor used for the "P" and "N" parts of the junction influences the best laser frequency to use to excite it. The same 980 nm infrared photons that work well on silicon ESD protection diodes are too low in energy (too long a wavelength) to work effectively on, say, gallium arsenide (GaAs) doped LEDs. However a typical 520 nm green solid-state laser works well in that case.

3 Related Work

Unwanted photosensitivity of electronic components has been a known issue for a long time. In 1952, the first production IBM 701 mainframe computer failed at its unveiling when the flashbulbs of news photographers disrupted the Williams tube memory of the computer [10, 17, 35, 75]. Semiconductor memory chips, if not protected from visible or ultraviolet light, are similarly sensitive [51, 82, 89].

In 2015 a flip-chip voltage regulator on the Raspberry Pi 2 single-board computer (SBC) was found—once again when it was being photographed for a press release before the product introduction—to crash the computer whenever exposed to xenon strobe camera flashes [28, 87]. Raspberry Pi 3 was later found to suffer a similar problem with a much larger chip-scale package integrated circuit (IC) on the back side, a Broadcom BCM43438 Wi-Fi and Bluetooth chipset (U19) this time [70]. In both cases, the root cause was found to be failure to specify an opaque backside laminate (BSL) on the chips [61].

Glass-encapsulated small signal diodes have long been known to pick up noise from overhead fluorescent lighting fixtures [13]. Photosensitive LEDs have been used before for communication [23, 81] or light detection [14, 52–55] or power delivery [1, 26, 49, 62, 63].

Other researchers have proposed ways to make a covert channel out of this effect. All such efforts require an insider with the ability to run a listening process on the target device [41]. This is not an unreasonable assumption, as STUXNET surely proved [19, 40], but what separates those efforts from this paper is that Basilisk does not require participation or cooperation by an insider.[1]

Most of the remaining published research related to compromising optical emanations (optical TEMPEST) concerns information flow out of the computer system; only a few papers in the literature address information flow inwards [24, 39, 59, 74]. Perhaps closer are laser injection attacks on optical fiber components of a quantum key distribution system, but in that case the real target was an optical receiver actively listening for a signal [30, 73]. A good survey of signal injection attack vectors is [31].

Sugawara *et al.* demonstrated coupling between a relatively high power laser and microelectromechanical systems (MEMS), *e.g.*, microphones and accelerometers, a clear parallel to our work because it similarly depends on the physical principle of energy transfer from the laser to the target system in order to effect coupling to a subsystem that was not listening for optical signals directly [83]. Rampazzi *et al.* (2020) found evidence of both photomechanical and photovoltaic effects at work [71, §4.3] and this work was further extended [22] by Cyr *et al.* (2024) but MEMS attacks are always targeting sensors, with the intention of falsifying measurements; our work is targeting any device that has an exposed P–N junction and can give direct access to the internal state of the machine.

Basilisk is not precisely the same thing as optical fault injection; that work is done on decapsulated chips, not on conventionally exposed P–N junctions [3, 5, 25, 33, 36, 37, 64–67, 78–80]. Laser fault injection is a technique primarily used in chip design and manufacturing for reliability testing. Focused radio frequency, x-ray, subatomic particle, or

[1] Randal (2023) makes the useful distinction amongst (1) covert channels, where both sender and receiver are malicious, (2) side channels, where only the receiver of information is malicious, and (3) fault injection, where only the sender is malicious [72, §2.1].

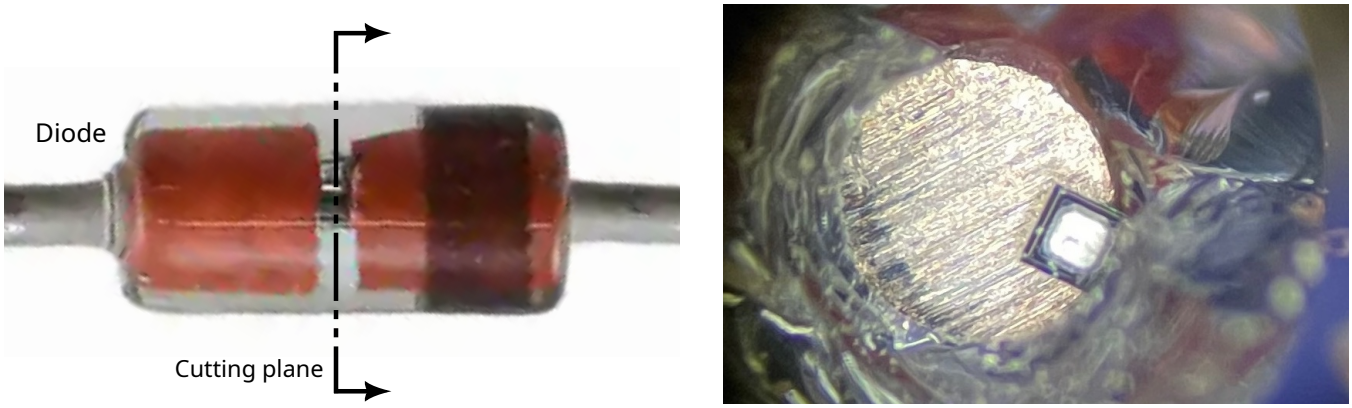

Figure 1: Example of an ESD protection diode. The actual silicon diode is the small square shape visible in the gap between the electrodes. It is often located off-center in the gap, decreasing the angle from which it can be hit.

laser radiation can be used to induce permanent or temporary changes in electronic circuit elements, leading to error states resulting in failures of the system [4, 90]. Light-induced voltage (or current) by means of a scanning laser is a test and characterization method used by semiconductor manufacturers to optimize process changes. It may be used in margin testing to assess reliability. Environmental conditions such as temperature, clock speed, or power supply voltage may be varied to induce faults; the latter is the basis for glitching attacks [29, 32].

Single event upset (SEU) testing (*e.g.*, cosmic rays) is important for space vehicles and devices designed to operate in the vicinity of a nuclear reactor. Electromagnetic interference (EMI) testing for electromagnetic compatibility (EMC) is a type of fault injection, and Basilisk may be considered intentional electromagnetic interference (IEMI).

Laser or photoflash fault injection at the basic component level is usually done on a decapped chip or bare die, under a microscope, allowing for precise placement and small spot size. As an attack vector, optical fault injection is normally used for key extraction—from TV set top boxes, electricity meters, smart cards, and payment terminals, in the course of hardware security module (HSM) testing, KG-type military link encryptors, or trusted platform module (TPM) chips [38, 48]. Bar-El *et al.* (2004) contains a comprehensive list of active protections in hardware—duplication of circuits, multiple redundancy with or without time shifting, and error-correction codes—and software such as are routinely used in spacecraft to guard against SEU events [5].

In contrast to all these localized events, we aim to show that ranged attacks are practicable on networking or programmable logic controller (PLC) equipment used in factory automation and control. Instead of extracting information, such as cryptographic keys, our capability is to take over control of the system.

4 Basilisk

To demonstrate the various effects of P–N junction excitation we design an injection attack framework called Basilisk. This attack framework allows an external attacker to compromise an air-gapped system, without the need for internal collaboration.

Fundamentally the low-level effect of a Basilisk attack is to pull an internal wire in the target device high or low, depending on the type of diode that it is connected to and the surrounding circuit.

This can be used to disrupt a number of higher level tasks, including chip-to-chip communication, assert error- or interrupt conditions or to send or alter commands on a communication bus. For example, as we demonstrate in detail in Section 7, it can be used to corrupt or modify instructions fetched from memory, change the data sent to a display or communication module, or simply crash a device and make it unresponsive.

For Basilisk to be effective an attacker must be able to target a P–N junction directly. In this explanation we will use ESD protection diodes as examples although the same applies to LEDs or any other P–N junction encased in a transparent material. Figure 1 shows a magnified view of a diode cut across the center to expose the silicon chip. Observe that the silicon chip is sandwiched in a gap between cylindrical metal electrodes, and the chip is often off-center in the gap, making it easier or harder to target with a laser beam.

4.1 System model

The system model for Basilisk is depicted in Figure 2. It is quite broadly applicable and only has a few requirements.

First of all, in order to conduct a Basilisk attack the target system must have an exposed P–N junction. This can come in many forms, *e.g.*, in the form of an indicator LED designed to be visible from the outside; or in the form of an ESD

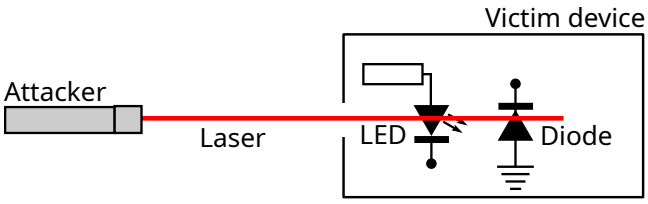

Figure 2: System and adversary model. The victim system has an exposed diode (or LED) that can be targeted by an adversary with a laser beam.

protection component (often a diode) mounted somewhere on a printed circuit board visible through an opening in the case.

The exposed diode must be connected to a useful target. The adversary is able to pull the wire connected to the diode either high or low depending on the type of diode and the surrounding circuit, but typically not both. The diode must be connected to a circuit that if pulled high (or low) will have an effect on the rest of the system. This is very often the case in practical systems and we will show a number of practical examples of such systems throughout the rest of the paper.

Finally, for some of the attacks to be effective the target device must be turned on and running code. This is not always strictly required, for example one can imagine an example where a soft power switch can be activated with a Basilisk attack or a reset line pulled low to reset a halted CPU, but in most practical cases we will assume that the target system is running.

The vulnerability exists whenever exposed PN junctions (like LEDs or ESD protection components) are connected directly or indirectly to electronic circuits carrying sensitive information.

4.2 Adversary model

The adversary must have line of sight to the target diode. This is not trivial in practice, but neither is it very difficult. In Section 5 we define a metric called active area that measures how precisely an attacker must aim to be effective. We show that it is indeed feasible to do even without specialist equipment.

For certain attacks it is necessary to be able to modulate the laser, *i.e.*, turn it off an on, say, in order to create valid packets. We assume the attacker is able to do this as fast as is needed for the attack. This is a fairly easy requirement in practice, as the modulation need not be any faster than the communication protocol under attack, *i.e.*, under 5 MHz for I^2C.

Although not always needed, we grant the adversary full knowledge of the timing of any messages that are transmitted across a targeted bus. This is somewhat of an over approximation, but it models the case where this information is available through other side channels. If this information is not available, attacks that require specific timing become probabilistic.

Basilisk works by pulling the wire connected to the diode

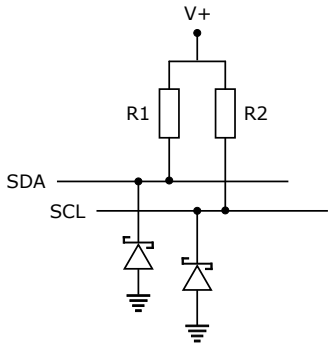

Figure 3: ESD protection diodes on the I^2C bus—when illuminated—are able to pull the bus down against the effect of pull-up resistors R1 and R2.

either low or high, but not both. For this reason the attacker can only change a binary 1 to 0, (or 0 to 1) but not both ways. This is not as much of a limitation as it first appears, because shared communication buses have pull-up (or pull down) resistors making the default bus state high (or low) as in Figure 3. The attacker can send arbitrary messages in that case by, say, pulling the bus low when needed and letting it go back up to high by just turning off the laser. Nevertheless it is a limitation that can come into play in some circumstances. We describe one example of this in Section 7.1.

Note that this technique does not permit the attacker to receive information from the circuit under attack, only send messages. If bidirectional communication is needed, another side channel must be used to read information. Such side channels are in fact often available, *e.g.*, when attacking a display as we demonstrate in Section 7.2, but the attacks we describe do not need to read from the device.

5 Diode Attack Surface

In this section we demonstrate the specific conditions under which it is possible for an external attacker to use a laser to execute a Basilisk attack. We introduce the measurement setup and then use it to drive ESD protection diodes or LEDs into photoconductive- or photovoltaic mode respectively, thus controlling the signal level of the connected circuits.

5.1 CMOS Logic Circuits

Before we continue we need to clarify the target voltage for our experiments. In CMOS, 3.3 V circuits signals are supposed to be either above $V_{IH} = 2.0\,V$ to indicate a logic high condition or below $V_{IL} = 0.8\,V$ to indicate logic low. The gap between V_{IL} and V_{IH} is not to be used to avoid ambiguity and to provide a buffer—called the noise margin—against small fluctuations in the electrical signal changing the logic state.

However CMOS logic is a binary system and the "undefined" state cannot actually be represented in hardware, so

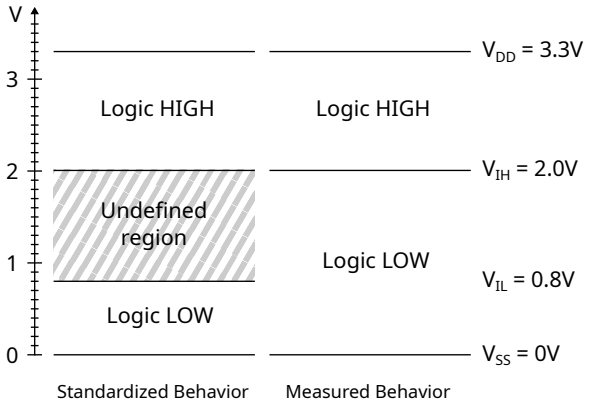

Figure 4: CMOS logic levels for 3.3 V circuits. Signals above 2 V are logic high, and below 0.8 V are logic low. In all our experiments we found that devices will default to a logic low condition in the undefined region, so although it ought never to be used, a signal below 2 V is sufficient for an attack.

in practice; in our experiments, anything below 2 V is interpreted as a logic low condition.[2] This difference is illustrated in Figure 4. This slightly higher value for a logic low condition is helpful for our attack since the lower the attacker wants to drive the signal, the more power is needed. Given the observed behavior we can use 2 V as the threshold for a successful attack.

5.2 Attack Measurements

To make sure our measurements and results are applicable to a real world system, we make all our measurements on an experimental setup consisting of two devices (bus controller and target) communicating over an I^{2}C-bus. The bus has ESD protection diodes and external pull-up resistors to allow us to experiment with different values. We can change the resistor values and bus voltage independently; this mimics the I^{2}C specification, which allows a wide variety of values to be used [60, 68]. A schematic of the test setup can be seen in Figure 5 and a photo in Figure 6.

The measurement setup consists of a pair of linear actuators at 90° to each other to allow for a systematic raster-scan of the diode under attack. The actuators are driven by stepper motors which enable precise repeatable measurements to be taken. Each raster scan of the diode moves the laser beam across a 1 mm^2 area of one of the ESD protection diodes on the I^{2}C bus, stopping every 20 µm to measure the voltage on the bus.

Basilisk attacks are ideally suited for a shared communication bus because such buses use open collector (or open drain) drivers. A device wishing to transmit will drive the bus low to send a binary 0 or simply release the bus and allow the

[2] Our determination of this value is supported by Lancaster (1974, 1977) [42–44].

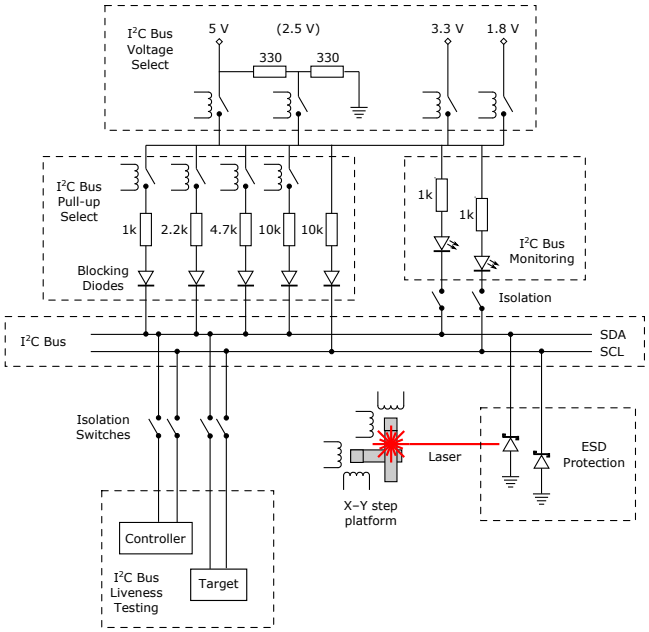

Figure 5: Schematic of the experimental apparatus.

pull-up resistor to return the bus to its logic high condition (binary 1) between zero bits or between transmissions.

Typical values for I^{2}C bus pull-up resistors are 2.2 kΩ or 4.7 kΩ; lower values pull the bus up more strongly, allowing faster communication; conversely, a higher value resistor like 10 kΩ is a weaker pull-up.

The ESD protection diodes tested are a common type of glass-encapsulated DO-35 size small signal diode, a type 1N34A equivalent silicon Schottky diode chosen for its fast recovery speed. The diode is connected in reverse bias with its anode at ground potential. Reverse bias makes the diode non-conductive under normal circumstances, so it doesn't affect the operation of the bus. If the voltage on the bus ever exceeds the reverse breakdown voltage of the diode, *e.g.*, during a power surge, the diode becomes conductive and shunts the power surge to ground, protecting the bus and the connected devices [47, 85].

We test several different lasers from near infrared (IR) devices at 780 nm, to a longer wavelength of 808 nm and 980 nm in order to identify the type best suited for a particular diode type. The lasers all have a fixed power (intensity) rating between 3–5 mW.

We use our measurement setup to trace a raster pattern with the laser, back and forth over the diode, to identify the best place to direct the laser during an attack. This is illustrated in Figure 7. The lasers are focused to the smallest achievable spot size at a working distance of 32–35 mm. This is not critical for the attack to work but it gives our measurements a higher resolution, allowing us to map out diode vulnerabilities in more detail. Note that even when focused, the beam from a semiconductor laser is slightly elliptical, so we repeated

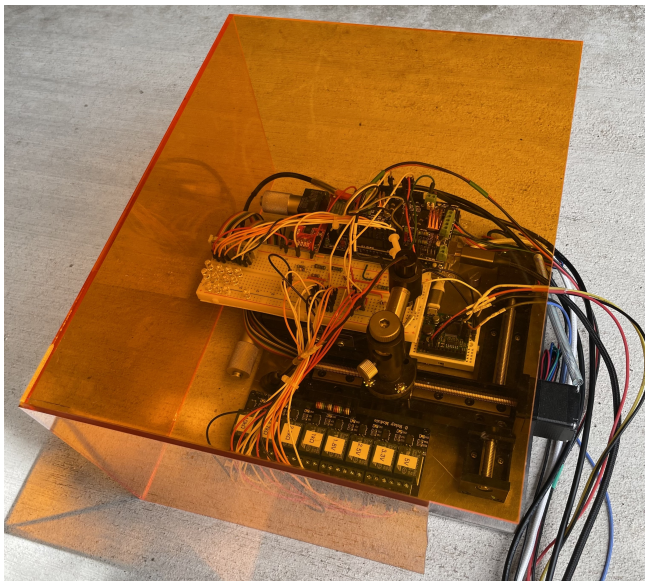

Figure 6: Experimental apparatus. Stepper motor linear actuators raster scan the laser, its elliptical beam axis marked, across the glass envelope of an ESD protection diode, obscured here by the laser mount. The color of the radiation shielding is wavelength-dependent, here for wavelengths shorter than 550 nm.

all the experiments with the long axis of the elliptical beam oriented parallel, perpendicular, and diagonally to the long axis of the gap between the electrodes in the ESD protection diode under test.

One run of the experiment consists of setting the bus voltage and pull-up resistor values, then scanning the laser spot over the glass body of the diode as shown in Figure 7. Voltage measurements are taken at fifty evenly spaced points along each scan line, for a total of 2500 measurements per run. Each run is repeated sixteen times for different combinations of bus voltage and pull-up resistor values, then the laser is rotated 45 degrees to shift the elliptical beam axis. We have experimented extensively with different wavelengths as well, however for most diode types we have one laser that is clearly the most efficient, so we only present the data for that set of experiments. See Section 8.2 for more details on laser wavelengths.

Voltage measurements are made with a 10-bit analogue-to-digital converter (ADC) channel 0 of an Arduino Uno single-board computer, against a 5 V reference. The ADC was allowed to settle for 500 ms every time the bus voltage was changed, and several voltage measurements were taken at every point, then averaged. The resolution of the ADC is 4.88 mV. All power supply voltages were verified with a Fluke 107 multimeter to be within 0.1 V of spec before the start of each run.

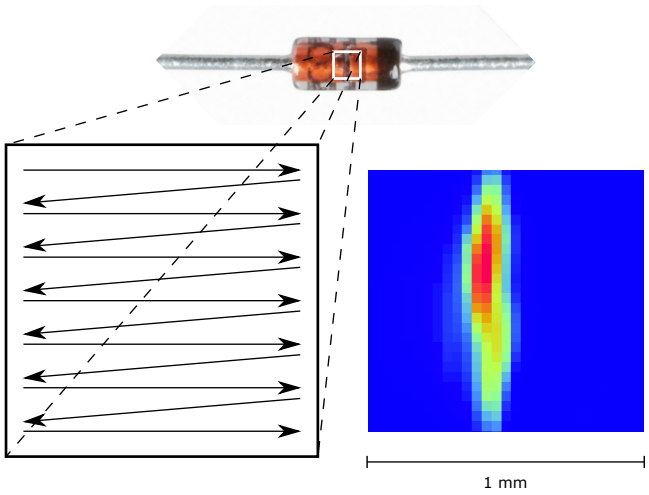

Figure 7: The raster pattern used to scan the diodes and LEDs to obtain repeatable results. Each run consists of 2,500 individual measurements.

6 Experimental Results

After a comprehensive series of tests on both ESD diodes and LEDs we have detailed results that show both options as a viable entry point for a Basilisk attack. In the following we present our findings for the two diode types.

In both cases the results of the experiments are a series of voltage measurements across the measurement area. This area is 1 mm^2 for the ESD diodes and 25 mm^2 for the larger LEDs.

6.1 ESD Diodes

A representative result is shown in Figure 8. The size and shape of the gap between the electrodes can be clearly seen. There is some indication of the size and location of the silicon chip. The black isovolt curve indicates the $V_{IH} = 2.0\,V$ level where the laser has forced the bus voltage below the logic threshold. Any hit from the laser inside the 2 V contour line will be seen as a logic low signal on the I^2C bus. We call this region the *active area* of the diode and it serves as a good metric for how easy it is to execute the attack for a specific set of parameters. If the active area is small for a chosen set of parameters, it means that those parameters make it more difficult to drive the signal low. Conversely if the area is large it means that it is comparatively easier.

To demonstrate how an attacker can best influence an exposed diode, and what circuit types are more vulnerable, we look at each of the parameters individually. These include beam axis rotation, bus voltage, and pull-up strength.

Beam rotation. Figure 9 depicts example results for three different beam rotations. The active areas (region inside the isovolt curve) look slightly different for the three rotations but

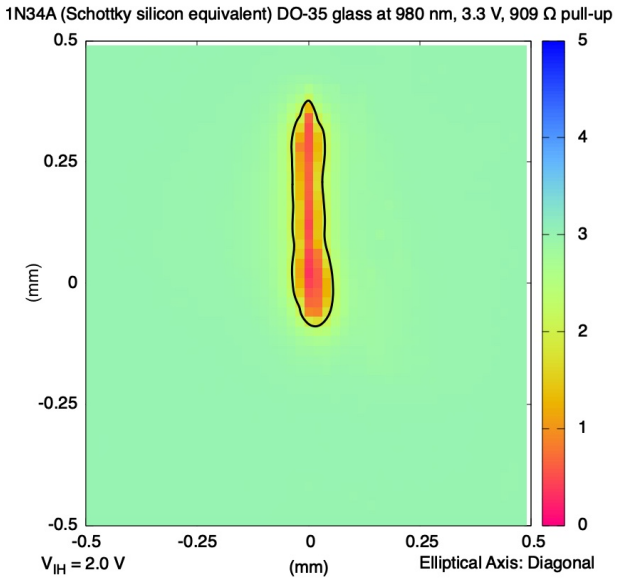

Figure 8: Voltage measurements from a 980 nm laser. The size and shape of the gap between the electrodes can clearly be seen. The black isovolt curve delineates the *active region* where the laser was able to force the bus voltage below the logic threshold V_{IH}, thereby imposing a binary zero on the bus. Simply releasing the bus, by turning off the laser, is sufficient to send a binary one.

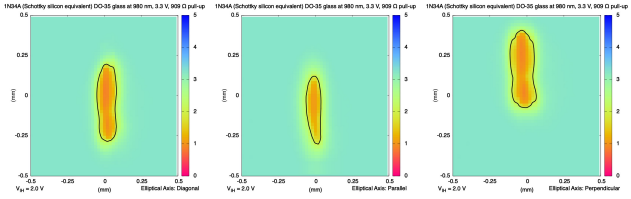

Figure 9: Effect of beam elliptical axis rotation: diagonal 45° (left), parallel 0° (middle), and perpendicular 90° (right).

it seems hard to draw a firm conclusion. To help with that we plot the average size of each of the active areas over 10 runs in Figure 10. Here it can be seen that we have a larger active area if the beam is perpendicular, *i.e.*, rotated 90°. While it is tempting to conclude that perpendicular beams are better, it is likely an artifact of the geometry of a specific diode, as the internal placement of the silicon die is believed to be random. However what we can learn from this is that rotation matters, and in a practical scenario rotating the beam might yield a bit of extra efficiency if the diode can only be targeted with an off center beam, or if the power of the laser is limited.

Bus voltage and pull-up resistor strength. Figure 11 shows the effect of bus voltage and pull-up resistor strength. Voltage varies from top to bottom and pull-up resistor strength varies from left to right. Observe how the active area becomes larger

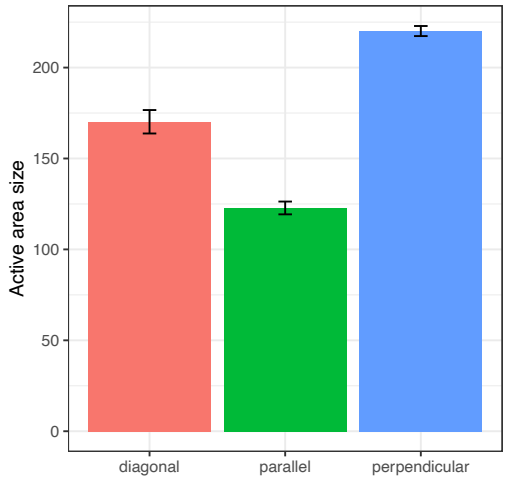

Figure 10: Average size of the active area (over 10 runs) for each of the three elliptical axis beam rotations. The error bars indicate standard deviation.

at lower bus voltages and with weaker pull-up resistors. This makes intuitive sense since with a lower bus voltage there is less of a voltage difference between a logic high and logic low signal, and thus it takes less to pull the signal down to a low state. Similarly for weaker pull-up resistors there is less current flow available to counteract the attempt to pull the signal level down to low.

For example, while at 5 V with a strong pull-up, attacking the I^2C bus is difficult, it is easy to attack a bus at 3.3 V with any reasonable pull-up resistor value.

The isovolt contour lines support the prediction that the attack will get easier in future as system voltages fall from 5 V TTL through 3.3 V CMOS to 2.5 V and 1.8 V LVCMOS [2, introduction]. In general, we can now predict whether a given combination of component type, bus voltage, pull-up resistor value, and laser wavelength is reversible.

6.2 Light Emitting Diodes (LEDs)

Doing the same set of experiments again for LEDs serve two purposes. First we want to show that Basilisk attacks are possible on LEDs as well, which is important since LEDs are often more available as targets. Second we want to investigate how different colored LEDs behave. The behavior will be different since the semiconductor is doped with different elements in order to produce different colored light. Furthermore, some LEDs are encased in a colored resin, which could influence the effectiveness of our attack laser.

An LED is a significantly bigger target than an ESD diode. This makes aiming the beam and beam rotation less important so we will not reproduce that part of the experiment here. We instead focus on the most important aspects that we found

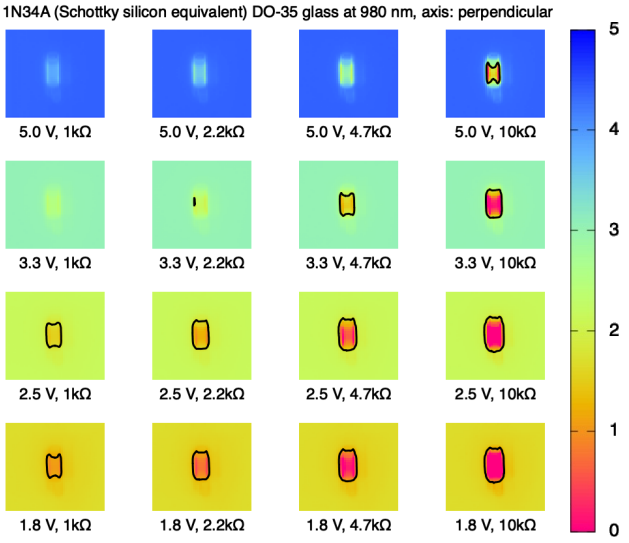

Figure 11: Voltage measurements from varying I^{2}C bus voltage and pull-up resistor strength as independent variables. All attacks are done with an 980 nm laser. The black isovolt lines are 2.0 V for 5 and 3.3 V logic, 1.7 V for 2.5 V logic, and 1.17 V for 1.8 V logic.

will determine the resulting bus voltage, namely the color of the LED that is being attacked, the bus voltage of the victim system, and the pull-up resistor strength.

LED color. We tested the Basilisk attacks on six LEDs (four different colors) with the same 405 nm (violet/blue) laser to get a measure for how color impacts the efficiency of the attack. The results can be seen in Figure 12.

We see that pink LEDs exhibit the weakest response which is likely due to a phosphor layer absorbing a lot of the shorter wavelength before it gets to the chip. The blue and green LEDs both respond fine to the attack, and it is clear that those LEDs are vulnerable enough to be used in practice. The most vulnerable LED color is white. We have not attempted to uncover the specific physical reason for these differences, just note that with the exception of the pink LED, all had fairly large active regions, making them easy to hit.

Bus voltage and pull-up resistor strength. Another thing to note from Figure 12 is that the voltage in the active region, *i.e.*, inside the isovolt curve, is not as low as it is for ESD protection diodes. While LEDs are large and easy to hit, they do not drive the voltage as far down with the same laser power, *i.e.*, they are less efficient. While an ESD protection diode easily drives the voltage down to zero if hit correctly, LEDs usually bottom out at 1–1.5 V. As mentioned in Section 4, this signal level is considered undefined for CMOS circuits but in practice it is enough to trigger a logic low condition, which means it is good enough for our purpose.

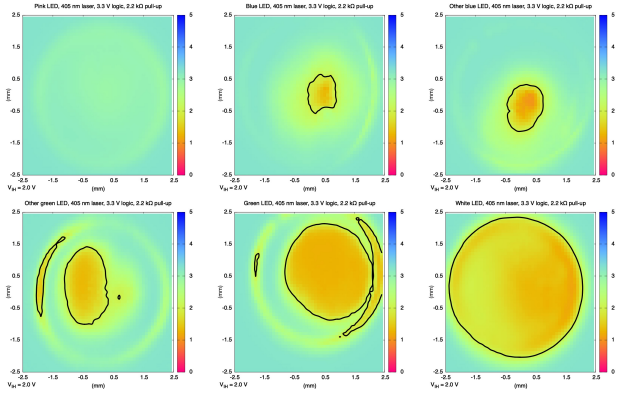

Figure 12: Voltage measurements of different colored LEDs. pink (top left), blue (top middle, top right), green (bottom left, bottom middle), and white (bottom right). All attacks are done with a 405 nm laser.

7 Case Studies

In the previous section we demonstrated that an adversary can change the logic level of an I^{2}C bus with a laser. But that falls short of demonstrating remote code execution on a live CPU, which requires the attacker to have precise control of timing.

We next show that the attack vector can be used to perform a meaningful attack. We do this through two case studies. In the first one we attack a small computer with a simple instruction set to demonstrate how we can effect arbitrary code execution. The second case study demonstrates that Basilisk attacks can be performed against off the shelf hardware.

The two case studies also allow us to demonstrate both types of Basilisk attack: photovoltaic and photoconductive. In the first case study the exposed diode is an LED connected to the memory bus so we can use photovoltaic attacks to drive the bus wires to their logic high state. In the second case we attack the ESD protection diodes on an I^{2}C bus so we use a photoconductive attack to drive the bus low.

7.1 Changing the running instruction stream

To demonstrate the practicality of remote code execution, at the same time keeping the complexity low so all details are visible, we constructed a minimal 4-bit computer called M5. It has status lights on the memory bus—a feature sadly lacking in most computers designed since the 1970s—and an instruction set architecture (ISA) small enough to make Figure 17 feasible. We wanted to be able to show the reachability analysis between instructions without the overwhelming complexity of a modern ISA like ARMv7 or RISC-V. The point is to show that because of an interesting constraint on the attacker—who may only be able to change a binary 0 to 1—but not the other way around—certain opcodes are reachable from certain other opcodes, but not always the one you want.

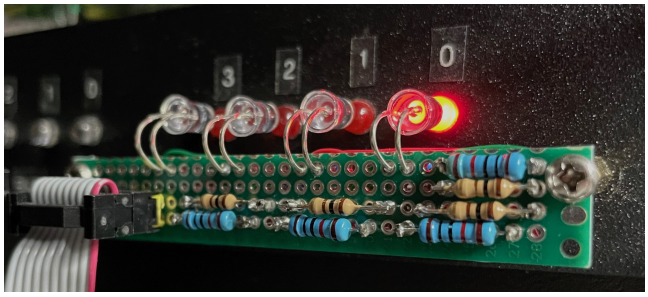

Figure 13: The attacker watching the accumulator to get a phase lock on the CPU at a cycle time of 250 µs.

M5, shown in Figures 15 (a) and 16, is a minimalist CPU intended not so much to show the practicality of the attack against real hardware but—quite the opposite—to highlight certain unique difficulties of the attack, beyond obvious ones like aiming and focusing.

The accumulator register is visible on the front panel (Figure 13)—visibility is key to establishing a phase lock on the CPU. It has a simple instruction set (Table 1) to make feasible the reachability analysis in Figure 17, and will halt if it decodes an illegal opcode. Cycle time is 250 µs but this was limited by the speed of the laser drivers, not the FPGA.

Both CPU and memory are implemented in the FPGA but the bus between them was routed externally to be accessible for probing outside the FPGA's internal interconnection fabric, as shown in the schematic of Figure 15 (a). The FPGA used is a Lattice Semiconductor iCE40-HX8K breakout board, configured in Verilog with the open source Project iCEstorm. All experiments in this section are done with a 405 nm, 5 mW laser.[3]

It is important to note this is not an FPGA vulnerability; the attack happens outside the FPGA's internal interconnection fabric, on external I/O pins, which are all standard CMOS.

M5 runs the microcode for each instruction on a sixteen-step cycle. Figure 14 shows a typical instruction, opcode mnemonic STA, which stores the value currently in the accumulator register to a specified memory location.

The attacker needs to know a lot about the CPU and the program that is currently running to perform the attack successfully. Firstly, the attacker needs to establish a phase lock on the internal state of the CPU, and from it to accurately measure the cycle time, because the entire attack is predicated on cycle counting.

Here, phase locking is accomplished by watching the accumulator display for changes, because the display always changes at a known microcode cycle. From the direction and magnitude of the change, the attacker can deduce what

[3]With the laser focused at infinity, any portion of the beam lying within a 5 mm diameter acceptance cone at the target will automatically be captured and index matched to the P–N junction, because optical systems—in this case, the lens of the LED—are time-reversible.

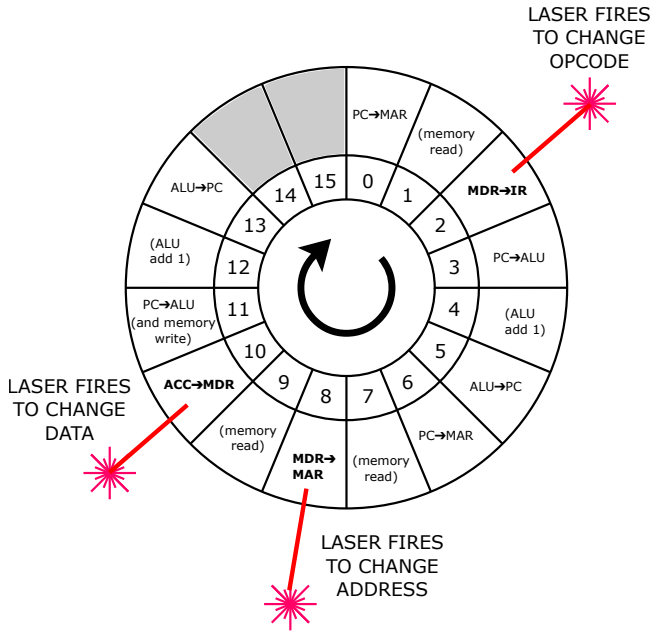

Figure 14: Timing diagram of the microcode showing where the laser fires during fetch and execution of a single instruction (here, STA for "store accumulator").

instruction was running (for example, INC or DEC if the accumulator value changed by one, or STA if the value changed by more than one). From the time between changes, the attacker can calculate the cycle time by dividing by the number of instructions executed between changes and looking up the cycle time of each instruction, which may be different.

All this could be accomplished a different way simply by watching the bus LEDs. We do it by means of the accumulator simply to illustrate the general principle that the attacker is not necessarily attacking the same LED as the one being watched.

After the attacker has established a phase lock and measured the cycle time, the attack proceeds by counting cycles into a predicted part of the fetch–execute cycle and firing the lasers at the instant when the desired value is known to be on the bus (Figure 14). Typically, the laser fires more than once during a particular instruction; for example, once to change the opcode, again four and a half cycles later to change the memory address, and five cycles after that to change the data being written to memory.

The result of the attack can be seen on the right side of Figure 15. Memory map (b) shows what the memory contents of M5 looked like before the attack; (c) shows what it looks like after. The attacker fired the lasers a total of twenty times in six seconds to force the normal program in locations 0–7 to write a new program in high memory at locations 9–14, and then forced a branch to it.

We did not collect any data on the probability of the attack being successful, because we found it to be reliable under the conditions we set up. The lasers are bolted in position, aimed

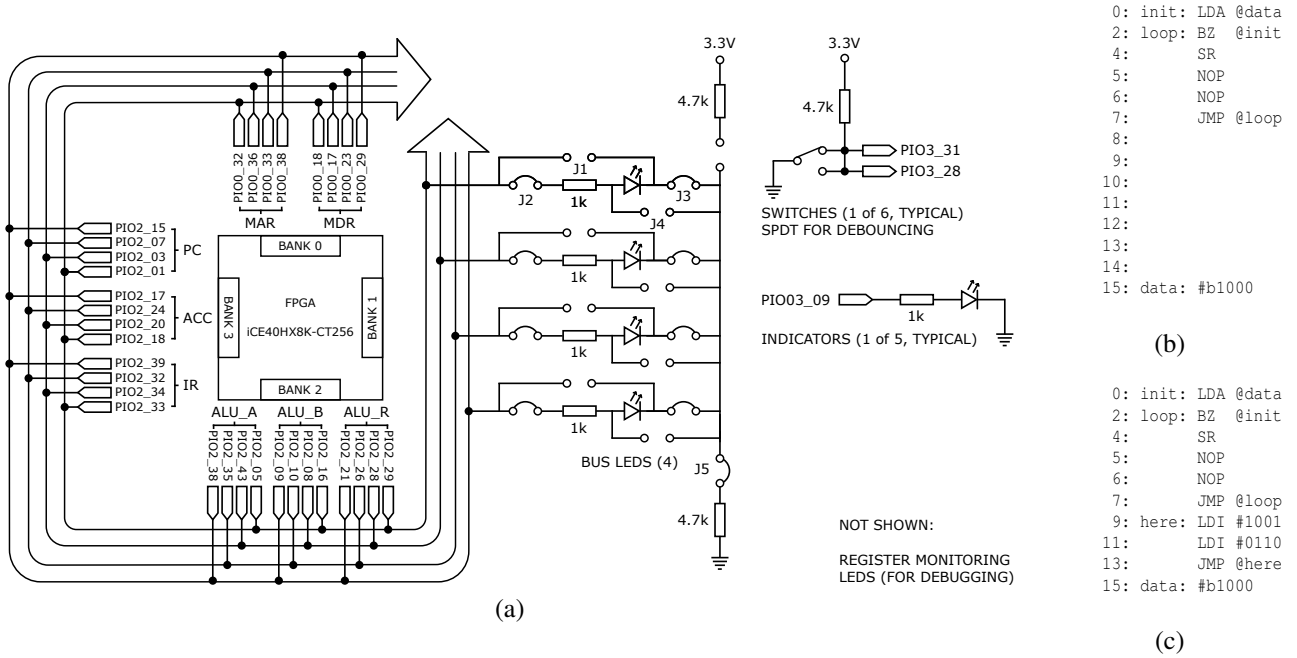

Figure 15: (a) Schematic of the M5 computer, built around a Lattice iCE40HX8K-CT256 FPGA. Jumpers J1–J4 form an H-bridge to allow flexibly reorienting the polarity of bus LEDs for experimentation. The bus is pulled down by jumper J5. (b) Memory map before the attack. (c) Memory map after the successful attack.

Figure 16: The M5 computer. Four lasers controlled by an Arduino with MOSFET laser drivers, are mounted aimed at the bus LEDs. The four LEDs on the right are the accumulator (A) register and the LEDs and switches on the left are used to interact with the computer.

at the bus LEDs from a range of 2 cm, because it removes an independent variable (aiming error) from the experiment.

The lasers used in this experiment were 405 nm near-UV diode laser modules of unknown power rating. The lasers were extracted from "cat toys" sold on Amazon.com at https://www.amazon.com/gp/product/B09Y4D7NFB/. In operation, they draw approximately 150 mA each from a 3.3 V supply, so their optical power must be < 500 mW and is probably considerably less, as they get warm in continuous operation. *These are absolutely not eye-safe and should never have been sold as cat toys.*

For safety when using 405 nm lasers, we recommend an enclosure made from #2422 transparent orange polycarbonate sheet 3 mm thick, as shown in Figure 6 when that apparatus was running at 405 nm.

The lasers are modulated by switching their power supply on and off with a MOSFET. Two important considerations apply to these lasers; firstly, they need 3.3 V and will burn out quickly at 5 V, but the MOSFETs won't switch a load less than their gate (control) voltage. So to make the MOSFETs work and avoid burning out the lasers, always switch 5 V through the MOSFET, and drop it down to 3.3 V with an LM317 voltage regulator between the MOSFET and the laser.

These lasers were chosen for use because they exhibit quick response when modulated in this way, typically $< 100\,\mu$s turn-on and turn-off latency. Many other laser modules from other sources, when measured, had a turn-on latency of more than 4000 µs, limiting modulation to < 0.25 kHz.

Note that the bus is pulled down by the 10 kΩ resistors. There is no particular reason why it needs to be that way; it's only to make clear that LEDs under laser illumination drive their signals in reverse.

Using the M5 computer we can now demonstrate a few different adversarial capabilities, the first and easiest being to crash the computer using a Basilisk attack, and the second being arbitrary code execution (with some minor constraints).

To crash the computer, the attacker directs an unmodulated

Mnemonic	Opcode	Instruction
NOP	0000	no operation
LDA	0001	load accumulator (addr)
INC	0010	increment accumulator
DEC	0011	decrement accumulator
STA	0100	store accumulator (addr)
BZ	0101	branch if $A \neq 0$ (addr)
JMP	0110	unconditional branch (addr)
SR	0111	shift right accumulator
LDI	1001	load immediate

Table 1: M5 instruction set. It consists of 9 instructions including NOP. This is intentionally simple but Turing complete.

laser to any of the bus LEDs for a few seconds. This has any of the following effects: (1) changing a valid opcode to an invalid one; (2) changing the value of a memory address; (3) changing the contents of a memory read or write operation; or (4) changing the control flow of the program.

To achieve arbitrary code execution, the attacker needs to know something about the internal state of the CPU in order to synchronize precisely. Watching the accumulator display (Figure 13) is sufficient to obtain a phase lock on the CPU's internal state because the display changes at a known cycle offset within the STA instruction.

Even with knowledge of the running program and the CPU timing, the attack is not trivial. By targeting some or all of the bus LEDs, the attacker can change instructions, or data, or addresses after they are fetched from memory but before they are executed by the CPU. However the attacker can only set bits, *i.e.*, change a 0 to a 1, not the other way around.[4] This means that for every instruction there is a certain set of different instructions the attacker can reach. This is illustrated in Figure 17. The level of effort is not dissimilar to finding "gadgets" in return-oriented programming (ROP) or selecting opcodes from the printable character subset of the Intel X86 instruction set architecture [12, 15, 16, 57, 58, 76]. Another analogy would be to looking a few moves ahead in a chess game.

The address of a load or store operation can be easily redirected to high memory addresses simply by setting the most significant bit. The attacker can use many loops through the program, setting a few bits here and there, gradually building up the desired program in high memory. Once done, the attacker can simply redirect a branch to the new code.

In case the code in high memory cannot be created exactly how the attacker wants, maybe because some bit combinations were inaccessible, the constructed code can run fixups on itself the first time it runs. This enables the attacker to use the available instructions to write the desired program. After that,

[4]This is in photovoltaic mode; in photoconductive mode, we might only be able to *reset* a bit.

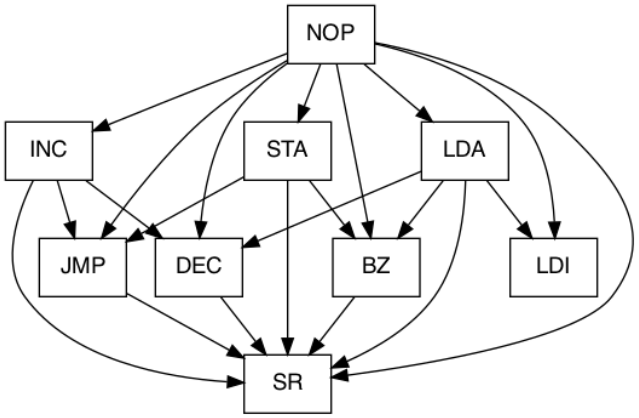

Figure 17: Allowable transitions in the instructions set of the computer defined in Table 1 if the attacker can only set bits, but not reset them.

the fixed-up code runs, and the attacker has full control.

The attacker must be very careful not to crash the running program. If the CPU ever halts, the attack is blocked.

7.2 Attacking an I^{2}C bus

We have demonstrated how a Basilisk attack can lead to arbitrary code execution on a toy computer. In this section we prove the viability by attacking an I^{2}C bus between commercially available devices at 100 kbit s^{-1}.

The I^{2}C bus is a serial, synchronously clocked communication bus that is widely used and has a tolerant specification we can abuse. It underlies the System Management Bus (SMBus) and PMBus, as well as being incorporated in many other real-world interfaces including HDMI, PCIe, DVI, and VESA [18]. A timing diagram of the I^{2}C bus can be seen in Figure 19. It consist of two wires: synchronous data (SDA) and synchronous clock (SCL). It is less expensive to attack than a parallel bus (*e.g.*, as used on M5), because the attacker needs to aim and synchronise at most two lasers. Timing is dictated by the sending device and can be arbitrarily slow. This further reduces the difficulty for an attacker since there is no need for synchronization with an existing clock.

I^{2}C is a shared bus with pull-up resistors making the default state of the bus logic high. It uses open collector drivers to pull the bus low when needed so that two devices trying to send at the same time will not cause hardware damage to each other. This makes it ideal as a target for Basilisk attacks.

The I^{2}C bus in our setup (see Figure 18) is provided with ESD protection devices (1N34A Schottky equivalent) on each of the two lines providing an entry point for a Basilisk attack. Figure 20 shows this attack working in practice—on LEDs, in this case. The figure shows an oscilloscope trace of the two I^{2}C wires with the attacker driving both wires to execute the attack. Note how, in the upper trace, the receiver acknowl-

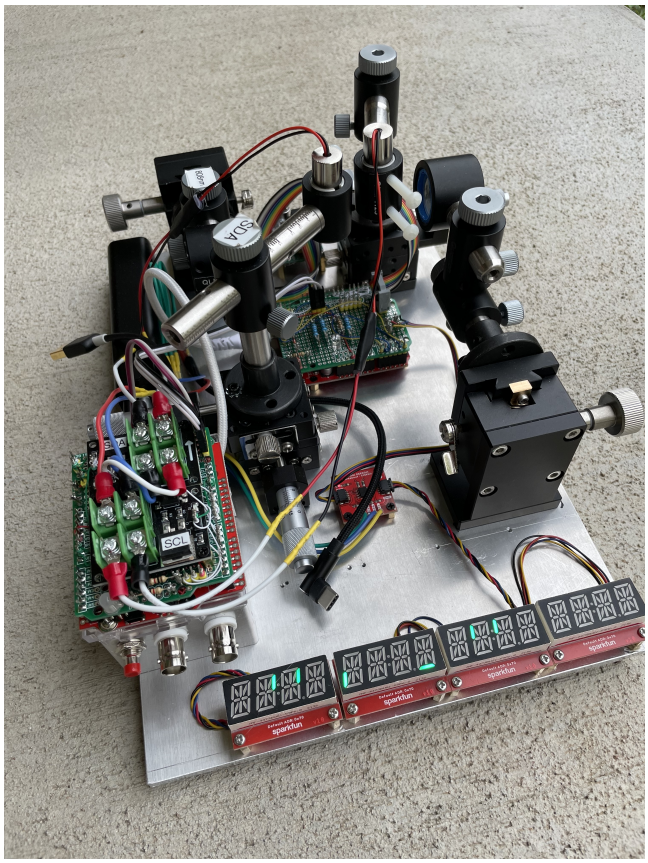

Figure 18: Proof of concept Basilisk attack on I^{2}C bus, 780 nm lasers irradiating 1N34A glass-enclosed ESD protection diodes. Representative devices on the bus include quad alphanumeric displays and nonvolatile memory.

edges each byte sent via the lasers by pulling the bus low. We can be sure that this is indeed the receiver pulling the bus low because it is pulled down all the way to 0 V, which the attacker cannot achieve. The LED can only pull the bus down to about 1.2 V (from V_{DD} = 3.3 V) but that is enough.

We identify three attacks: denial of service, message manipulation, and message insertion. The first and easiest attack is denial of service where the attacker simply sends a constant beam of light. Even if the attacker only controls one of the two bus lines, this attack effectively prevents any other communication from taking place.

The second attack is message manipulation where an attacker alters messages sent by other nodes, subject to analogous restrictions as in the M5 example—the attacker can only reset bits, not set them. This attack requires knowledge of the messages going across the bus and precise control over the timing of the attack. The sender of the message is able to detect the change, and will interpret it as bus contention, backing off automatically, but the intended receiver does receive the altered message.

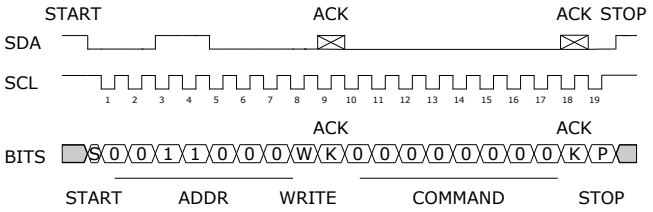

Figure 19: I^{2}C bus timing diagram. The I^{2}C protocol uses two wires, one for synchronous data (SDA) and one for a synchronous clock (SCL). Both of these are driven by the sender, except for the acknowledgement which is driven by the receiver.

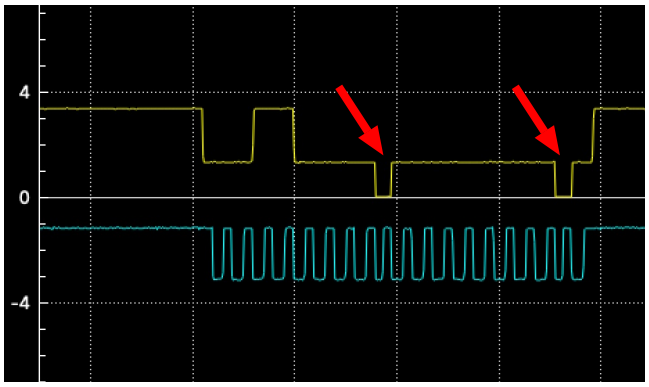

Figure 20: Oscilloscope trace showing successful communication with an I^{2}C device via laser illumination of status LEDs. Upper trace (yellow) is SDA; lower trace (blue) is SCL. The attacker is only able to pull the bus down to about 1.2 V, not all the way to ground. We know communication was successful because of the two negative-going acknowledgement pulses in the SDA trace, generated by the receiving device—one of the alphanumeric display modules, in this case—which is able to pull the I^{2}C bus all the way to ground.

The third and final attack is to transmit messages on an empty bus. It assumes that the attacker knows when the bus is empty, but beyond that there is no further need for precise synchronization with the existing devices.

8 Discussion

Introducing energy to the system—along with information—violates assumptions made by the hardware designer.

8.1 Laser power

In most cases accuracy and precision of aiming can be at least partially substituted with more power.

Our experiments used reclaimed laser modules from old CD, DVD, and Blu-ray players. This is in part to show that Basilisk attacks can be done with very inexpensive hardware,

but also that is what we could easily source. The downside of this is that we do not have specifications for all the lasers, specifically we do not have calibrated intensity measurements.

The lasers we have are sufficient to execute the attacks, but being able to precisely control the optical power and to focus the beam to a smaller spot size (*e.g.*, 10 µm) would allow for a more effective delivery of power to the P–N junction.

We observed standard laser safety protocols including use of warning signs, shielding, beam blocks, and protective glasses [6–8, 46, 88]. Hazards related to frequency-doubled 532 nm green laser pointers are well-known [11, 20].

There is no particular reason to think the light need be coherent, or even monochromatic. This is supported by the fact that xenon flash tubes are as effective as lasers.

8.2 Laser Wavelength

The lasers we have allowed us to test a variety of shorter wavelengths, 650 nm (red), 532 nm (green), and 405 nm (blue/violet) in addition to the infrared 780 nm, 808 nm, and 980 nm lasers we used. Shorter wavelengths tend to be more effective for LEDs and longer wavelengths are better for silicon.

8.3 Countermeasures

There are two main classes of countermeasures, active and passive. Active countermeasures rely on first detecting that an attack is taking place and then taking corrective action; attacks might be detected optically [9] or electrically [56]. Passive countermeasures minimize the attack surface, either by eliminating vulnerable components or shielding them from influence. Opaque chip packages can mitigate the attack, but this is no option for indicators that must remain visible.

Minimize the number of exposed LEDs or other photosensitive components on a device. Avoid connecting exposed P–N junctions directly to circuits carrying sensitive information.[5] Existing electromagnetic interference (EMI) reduction techniques may be effective against photovoltaic mode, but are unlikely to be effective in photoconductive mode.

It is interesting to note that the most effective wavelength for silicon ESD protection diodes (not LEDs) is in the infrared part of the spectrum, which makes the attacks stealthy.

8.4 Other Targets

There are a few areas where it is surprisingly common to have LEDs directly connected to sensitive circuits. We found that CAN bus devices commonly have them [34, 50]. Attacking a differential signaling bus like CAN is more challenging, as it will require the attacker to exploit both the photovoltaic and photoconductive methods on electrically adjacent components at the same time for it to work.

In certain circumstances it is possible that a Basilisk attack can damage the victim device. LEDs driven into photovoltaic mode tend to pull the circuit on their cathode sides more negative. This can, if the circuit on the cathode side of the LED is a low positive voltage, result in the circuit going below ground which could be damaging for sensitive electronics.

8.5 Commercially Available Hardware

We have been able to demonstrate the effect only on one piece of commercially available hardware: a 5 mm RGB color-changing LED often found in light-up toys [77].

The device was found to be disrupted by 405 nm, 520 nm, 650 nm, 780 nm, 808 nm, and 980 nm lasers, and the color changing sequence can be reliably reset to red at a distance of 25 cm by a xenon camera flash; this is consistent with the emission spectrum of xenon, which is rich in near-IR.

The effect was first reported in a comment on the article about the Raspberry Pi 2 glitch mentioned earlier [21, 28]. The chip inside is believed to be a CDT3447 or similar [69].

9 Conclusion

We present an attack framework we call Basilisk, after the mythical animal that could kill with a single glance [27, 45].

While the photosensitivity of semiconductor diodes is a known phenomenon, we demonstrate the practical requirements for an external attacker to use the effect as an attack vector. We show that Basilisk attacks are feasible in practice, both against ESD protection devices and LEDs, and can be performed as long as the attacker has line of sight access.

Our results go beyond a feasibility study. We show two concrete attacks that have serious consequences. Depending on an attacker's knowledge of the victim and equipment complexity, it is possible to achieve a number of effects, from denial of service to arbitrary code execution on an air-gapped computer system. Our results lead to testable predictions about the vulnerability of any shared bus that uses open collector (or open drain) tristate drivers.

Minimization is the most effective countermeasure, followed by buffering LEDs—which might be enough to block photovoltaic mode attacks only—but mitigation remains a challenge, especially for indicators that need to be exposed.

Availability

Verilog source code for FPGA implementation of the M5 CPU, Arduino source code for the attacker's equipment and the experimental apparatus, raw data, and scripts for data reduction and plotting are available on GitHub at https://github.com/jloughry/basilisk_artifacts.

[5] Indicators are not often found connected directly to a shared bus, because they load the circuit, slowing communication. LEDs are sometimes buffered by a transistor or op-amp driver, which has the advantage of a brighter indicator, but with the drawback of a slight overhead in cost and complexity.

References

[1] Rhett Allain. You can power a calculator with some LEDs, 11 March 2019. https://www.wired.com/story/you-can-power-a-calculator-with-some-leds/.

[2] Analog Devices. Low voltage logic interfacing. Tutorial MT-098, Analog Devices, Inc., 2009. https://www.analog.com/media/en/training-seminars/tutorials/mt-098.pdf.

[3] Ross Anderson and Markus Kuhn. Tamper resistance—a cautionary note. In *Second USENIX Workshop on Electronic Commerce*, pages 1–11, Oakland, California, 18–21 November 1996. https://www.usenix.org/conference/2nd-usenix-workshop-electronic-commerce/tamper-resistance-cautionary-note.

[4] Algirdas Avizienis, Jean-Claude Laprie, Brian Randell, and Carl Landwehr. Basic concepts and taxonomy of dependable and secure computing. *IEEE Transactions on Dependable and Secure Computing*, 1(1):11–33, January–March 2004. Available from https://www.nasa.gov/pdf/636745main_day_3-algirdas_avizienis.pdf.

[5] Hagai Bar-El, Hamid Choukri, David Naccache, Michael Tunstall, and Claire Whelan. The Sorcerer's Apprentice guide to fault attacks. IACR e-print 2004/100, International Association for Cryptologic Research, 2004. https://eprint.iacr.org/2004/100.

[6] Ken Barat. Eye safety in the laser lab: Using the humble beam block shows infinite wisdom. *Photonics Spectra*, August 2007.

[7] Ken Barat. Laser safety and the optical table. *Photonics Spectra*, October 2007.

[8] Kenneth L. Barat. In laser safety, little mistakes can have big consequences. *Photonics Spectra*, March 2005.

[9] Don Barber, Vikram Kanth, Zachary White, and John McEachen. Spatial frequency detection of optical signals embedded in the environment. In *2022 32nd International Telecommunication Networks and Applications Conference (ITNAC)*, 2022.

[10] Charles J. Bashe, Lyle R. Johnson, Emerson W. Pugh, and John H. Palmer. *IBM's Early Computers*. MIT Press, 1985.

[11] 'Brainiac75'. The issue with green laser pointers. https://www.youtube.com/watch?v=iR1Ku5dnbH8, 5 August 2018.

[12] Jerome Bruandet. Anatomy of the EICAR antivirus test file. Ninja Technologies Network, 26 August 2021. https://blog.nintechnet.com/anatomy-of-the-eicar-antivirus-test-file/.

[13] James Bryant. Glass diodes may see the light—and hum. *Analog Dialog*, 45, May 2009.

[14] James Bryant. LEDs are photodiodes too. Planet Analog, 5 August 2014.

[15] Erik Buchanan, Ryan Roemer, Stefan Savage, and Hovav Shacham. Return-oriented programming: Exploitation without code injection. In *Black Hat US 2008*, 2008.

[16] Erik Buchanan, Ryan Roemer, Hovav Shacham, and Stefan Savage. When good instructions go bad: Generalizing return-oriented programming to RISC. In *Proceedings of CCS 2008*, pages 27–28, Alexandria, Virginia, October 27–31, 2008.

[17] Christopher P. Burton. Replicating the Manchester Baby: motives, methods, and messages from the past. *IEEE Annals of the History of Computing*, 27(3):44–60, July–September 2005.

[18] Zitai Chen and David Oswald. PMFault: Faulting and bricking server CPUs through management interfaces. arXiv preprint arXiv:2301.05538 [cs.CR], Cornell University, 13 January 2023.

[19] Sean Collins and Stephen McCombie. Stuxnet: the emergence of a new cyber weapon and its implications. *Journal of Policing, Intelligence and Counter Terrorism*, 7(1):80–91, 2012.

[20] Annette Colton. Study of laser pointer safety. Study, Lucid Optical Services, 2010.

[21] 'Conundrum1885'. RE. Re. Xenon DEATH FLASH!!!!! *The Register*, 14 March 2015. https://forums.theregister.com/forum/all/2015/02/08/raspberry_pi_2_camera_flash_glitch/#c_2465254.

[22] Benjamin Cyr, Vedant Sumaria, Yan Long, Srinivas Tadigadapa, Takeshi Sugawara, and Kevin Fu. How lasers exploit photoacoustic and photoelectric phenomena to inject signals into MEMS microphones. Preprint, Research Square, 11 April 2024.

[23] Paul Dietz, William Yerazunis, and Darren Leigh. Very low-cost sensing and communication using bidirectional LEDs. Technical Report TR2003-35, Mitsubishi Electronics Research Laboratories (MERL), 201 Broadway, Cambridge, Massachusetts 02139, USA, July 2003.

[24] Jeroen Domburg. Optical mouse cam. http://spritesmods.com/?art=mouseeye, 2006.

[25] Mathieu Dumont, Pierre-Alain Moellic, Raphael Viera, Jean-Max Dutertre, and Rémi Bernhard. An overview of laser injection against embedded neural network models. arXiv preprint arXiv:2105.01403 [cs.CR], Cornell University, 4 May 2021.

[26] Arno Erzberger. Der LED fehlt der doppelpfeil. *Elektronik*, 21 June 2016. https://www.elektroniknet.de/power/energy-harvesting/der-led-fehlt-der-doppelpfeil.131470.html.

[27] Oliver Evans. Selections from the bestiary of Leonardo Da Vinci. *Journal of American Folklore*, 64(254):393–396, Ocotber–December 1951.

[28] Kelly Fiveash. 'Camera-shy' Raspberry Pi 2 suffers strange 'XENON DEATH FLASH' glitch. *The Register*, 8 February 2015.

[29] Curtis Franklin. Glitching: The hardware attack that can disrupt secure software. *Dark Reading*, 18 October 2019. https://www.darkreading.com/edge-articles/glitching-the-hardware-attack-that-can-disrupt-secure-software.

[30] Juan Carlos Garcia-Escartin, Shihan Sajeed, and Vadim Makarov. Attacking quantum key distribution by light injection via ventilation openings. *PLoS One*, 15(8):e0236630, August 2020.

[31] Ilias Giechaskiel and Kasper Rasmussen. Taxonomy and challenges of out-of-band signal injection attacks and defenses. *IEEE Communications Surveys & Tutorials*, 22(1), First Quarter 2020.

[32] Bret Giller. Implementing practical electrical glitching attacks. In *Black Hat Europe 2015*, Amsterdam, 10–13 November 2015. https://www.blackhat.com/docs/eu-15/materials/eu-15-Giller-Implementing-Electrical-Glitching-Attacks.pdf.

[33] D. H. Habing. The use of lasers to simulate radiation-induced transients in semiconductor devices and circuits. *IEEE Transactions on Nuclear Science*, 12(5):91–100, 1965.

[34] Microchip Technology Inc. *CAN Bus Analyzer User's Guide*, 2020.

[35] Karla Jennings. *The Devouring Fungus: Tales of the Computer Age*. W. W. Norton and Company, Inc., 1990.

[36] Stefanos Koffas and Praveen Kumar Vadnala. On the effect of clock frequency on voltage and electromagnetic fault injection. arXiv preprint arXiv:2310.13389 [cs.CR], Cornell University, 20 October 2023. Published in AIHWS workshop held for Applied Cryptography and Network Security Conference (ACNS 2022); DOI: https://doi.org/10.1007/978-3-031-16815-4_8.

[37] Thilo Krachenfels, Heiko Lohrke, Jean-Pierre Seifert, Enrico Dietz, Sven Frohmann, and Heinz-Wilhelm Hübers. Evaluation of low-cost thermal laser stimulation for data extraction and key readout. *Journal of Hardware and Systems Security*, 4(1):24–33, 2020.

[38] Thilo Krachenfels, Heiko Lohrke, Jean-Pierre Seifert, Enrico Dietz, Sven Frohmann, and Heinz-Wilhelm Hübers. Evaluation of low-cost thermal laser stimulation for data extraction and key readout. arXiv preprint arXiv:2006.06290 [cs.CR], Cornell University, 11 June 2020.

[39] Ireneusz Kubiak. *Specjalne fonty komputerowe w bezpieczeństwie elektromagnetycznym cyfrowych standardów graficznych: TEMPEST optyczny*. Wojskowa Akademia Techniczna, Warsaw, 2020.

[40] David Kushner. The real story of Stuxnet. *IEEE Spectrum*, 50(3):48–53, March 2013.

[41] Butler W. Lampson. A note on the confinement problem. *Comm. ACM*, 16(10):613–615, October 1973.

[42] Don Lancaster. Build Digiviewer II. *Popular Electronics*, 6(3):63–69, September 1974.

[43] Don Lancaster. *TTL Cookbook*. SAMS, Indianapolis, Indiana, 1974.

[44] Don Lancaster. *CMOS Cookbook*. SAMS, Indianapolis, Indiana, 1977.

[45] David Langford. comp.basilisk FAQ. *Nature*, 402:465, 2 December 1999.

[46] Laser Institute of America. Seven-year update drives national laser safety standard forward. *Photonics Spectra*, June 2007.

[47] Littelfuse.com. Application hints for transient voltage suppression diode circuits. Application Note AND8230/D, Littelfuse, Inc., 16 September 2016.

[48] Heiko Lohrke, Shahin Tajik, Thilo Krachenfels, Christian Boit, and Jean-Pierre Seifert. Key extraction using thermal laser stimulation: A case study on Xilinx Ultrascale FPGAs. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2018(3):573–595, 2018.

[49] Marcel Meli and Niklas Roth. LEDs und photodioden als energy harvester. *Elektronik*, 3 June 2016. https://www.elektroniknet.de/power/energy-harvesting/leds-und-photodioden-als-energy-harvester.130892.html.

[50] Microchip Technology Inc. *MCP2515 CAN Bus Monitor Demo Board User's Guide*, 2014.

[51] Micron Technology. *IS32 Optic Ram data sheet*. Boise, Idaho, USA, May 1984.

[52] Forrest M Mimms. Using LED's as light detectors. *Popular Electronics*, 11(5):86–88, May 1977.

[53] Forrest M Mims. *Light Emitting Diodes*. Howard W. Sams & Co., Inc., Indianapolis, Indiana, USA, 1973.

[54] Forrest M Mims. Bidirectional optoisolator puts two LEDs nose to nose. *Electronics*, 52(10):127, 10th May 1979.

[55] Forrest M Mims. How to use LEDs to detect light. *Make*, 36:136–138, 2014.

[56] Saleh Khalaj Monfared, Kyle Mitard, Andrew Cannon, Domenic Forte, and Shahin Tajik. LaserEscape: Detecting and mitigating optical probing attacks. arXiv preprint arXiv:2405.03632 [cs.CR], Cornell University, 6 May 2024.

[57] Tom Murphy. C with ABC! In *11th Special Interest Group on Harry Q. Bovik (SIGBOVIK)*, Pittsburgh, Pennsylvania, 31 March 2017. Association for Compting Heresy.

[58] Tom Murphy. Compiling C to printable x86, to make an executable research paper, 31 March 2017. https://www.youtube.com/watch?v=LA_DrBwkiJA.

[59] Ben Nassi, Adi Shamir, and Yuval Elovici. Oops!...I think I scanned a malware. *arXiv preprint*, 03 2017.

[60] NXP Semiconductors. *UM10204: I2C-bus specification and user manual, Rev. 7.0*, 1 October 2021.

[61] ON Semiconductor. Wafer-level chip-scale package (WLCSP) at ON Semiconductor. Application Note 5075/D, Fairchild Semiconductor, October 2018.

[62] Masashi Ono, Parthiban Santhanam, Wei Li Li, Bo Zhao, and Shanhui Fan. Experimental demonstration of energy harvesting from the sky using the negative illumination effect of a semiconductor photodiode. *Applied Physics Letters*, 114(16):161102, 2019. https://aip.scitation.org/doi/10.1063/1.5089783.

[63] Sung-Yun Park, Kyuseok Lee, and Hyunsoo Song. Simultaneous imaging and energy harvesting in CMOS image sensor pixels. *IEEE Electron Device Letters*, 39(4):532–535, April 2018.

[64] Dmytro Petryk, Zoya Dyka, Jens Katzer, and Peter Langendoerfer. Metal fillers as potential low cost countermeasure against optical fault injection attacks. arXiv preprint arXiv:2103.12436 [cs.CR], Cornell University, 17 January 2022.

[65] Dmytro Petryk, Zoya Dyka, and Peter Langendörfer. Sensitivity of standard library cells to optical fault injection attacks in IHP 250 nm technology. In *9th Mediterranean Conference on Embedded Computing (MECO)*, Budva, Montenegro, 8–11 June 2020.

[66] Dmytro Petryk, Zoya Dyka, Eduardo Perez, Mamathamba Kalishettyhalli Mahadevaiaha, Ievgen Kabin, Christian Wenger, and Peter Langendörfer. Evaluation of the sensitivity of RRAM cells to optical fault injection attacks. In *23rd Euromicro Conference on Digital System Design (DSD)*, Kranj, Slovenia, 26–28 August 2020.

[67] Dmytro Petryk, Zoya Dyka, Roland Sorge, Jan Schaeffner, and Peter Langendoerfer. Optical fault injection attacks against radiation-hard registers. arXiv preprint arXiv:2106.07271 [cs.CR], Cornell University, 18 January 2022.

[68] Phillips Semiconductor. *The I2C-Bus and How to Use It (including specifications)*, April 1995.

[69] Qipeng Semiconductor Co., Ltd. *CDT3447 3 LED Fade-in and fade-out output control IC*, 2008. http://www.bowin-ic.com.hk/IC/LED%20flasher%20IC/CDT3447.pdf.

[70] Paul Rako. EEVblog #901 - Raspberry Pi 3 photoflash problem. YouTube, 16 July 2016.

[71] Sara Rampazzi, Benjamin Cyr, and Daniel Genkin. Light commands: Hacking voice assistants with lasers. In *Black Hat Europe 2020*, Virtual, 7–10 December 2020. https://www.blackhat.com/eu-20/briefings/schedule/index.html#light-commands-hacking-voice-assistants-with-lasers-21731.

[72] Allison Randal. This is how you lose the transient execution war. arXiv preprint arXiv:2309.03376 [cs.CR], Cornell University, 6 September 2023.

[73] Khushboo Rani, Hansika Weerasena, Stephen A. Butler, Subodha Charles, and Prabhat Mishra. Modeling and exploration of gain competition attacks in optical network-on-chip architectures. arXiv preprint arXiv:2303.01550 [cs.CR], Cornell University, 3 March 2023.

[74] Melanie R. Rieback, Bruno Crispo, and Andrew S. Tanenbaum. Is your cat infected with a computer virus? In *Proceedings of the Fourth Annual IEEE International Conference on Pervasive Computing and Communications (PERCOM '06)*, pages 169–179, Pisa, Italy, 13–17 March 2006. IEEE Computer Society.

[75] D. E. Rosenheim. Installation of the first production 701. *Annals of the History of Computing*, 5(2):146–147, April–June 1983.

[76] Hovav Shacham. The geometry of innocent flesh on the bone: return-into-libc without function calls (on the x86). In *Proceedings of the 14th ACM conference on Computer and communications security*, Alexandria, Virginia, October 28–31, 2007.

[77] Shenzhen Xuancai Electronc Co., Ltd. *Color Changing LED*. Shen Zhen, China, 2011. http://cdn.sparkfun.com/datasheets/Components/LED/changingLED.pdf.

[78] Sergei Skorobogatov. Fault attacks on secure chips: from glitch to flash. In *Design and Security of Cryptographic Algorithms and Devices (ECRYPT II)*, Albena, Bulgaria, 29 May–3rd June 2011.

[79] Sergei P. Skorobogatov and Ross J. Anderson. Optical fault induction attacks. In *4th International Workshop on Cryptographic Hardware and Embedded Systems (CHES)*, Redwood Shores, California, USA, 13–15 August 2002.

[80] Sergei P. Skorobogatov and Ross J. Anderson. Optical fault induction attacks. In *4th International Workshop on Cryptographic Hardware and Embedded Systems (CHES 2002)*, pages 2–12, Redwood Shores, California, 13–15 August 2002.

[81] R. Stojanović and Dejan Karadaglić. Single LED takes on both light-emitting and detecting duties. *Electronic Design*, 55(16):53–54, 18th July 2007.

[82] Charles Stross. "Nothing like this will be built again". Charlie's Diary, 2010. https://www.antipope.org/charlie/blog-static/rants/nothing-like-this-will-be-buil.html.

[83] Takeshi Sugawara, Benjamin Cyr, Sara Rampazzi, Daniel Genkin, and Kevin Fu. Light commands: Laser-Based audio injection attacks on Voice-Controllable systems. In *29th USENIX Security Symposium (USENIX Security 20)*, pages 2631–2648. USENIX Association, August 2020.

[84] S. M. Sze and Kwok K. Ng. *Physics of Semiconductor Devices*. John Wiley & Sons, Inc., Hoboken, New Jersey, 2007.

[85] Toshiba. Basics of ESD protection (TVS) diodes. Applicatio note, Toshiba Electronic Devices & Storage Corporation, 27 May 2022.

[86] Endel Uiga. *Optoelectronics*. Prentice-Hall, Englewood Cliffs, New Jersey, 1985.

[87] Liz Upton. Xenon death flash: a free physics lesson. Raspberry Pi blog, 9 February 2015.

[88] John Wallace. Many laser-safety eyewear products do not meet specs for shielding light from ultrafast lasers. *Laser Focus World*, 29th November 2017. http://www.laserfocusworld.com/articles/2017/11/many-laser-safety-eyewear-products-do-not-meet-specs-for-shielding-light-from-ultrafast-lasers.html.

[89] D. G. Whitehead, I. Mitchell, and P. V. Mellor. A low-resolution vision sensor. *J. Phys. E: Sci. Instrum.*, 17:653–656, 1984.

[90] J. F. Ziegler, H. W. Curtis, H. P. Muhlfeld, C. J. Montrose, B. Chin, M. Nicewicz, C. A. Russell, W. Y. Wang, L. B. Freeman, P. Hosier, L. E. LaFave, J. L. Walsh, J. M. Orro, G. J. Unger, J. M. Ross, T. J. O'Gorman, B. Messina, T. D. Sullivan, A. J. Sykes, H. Yourke, T. A. Enger, V. Tolat, T. S. Scott, A. H. Taber, R. J. Sussman, W. A. Klein, T. D. Sullivan, and C. W. Wahaus. IBM experiments in soft fails in computer electronics (1978–1994). *IBM Journal of Research and Development*, 40(1):3–18, January 1996.