

Confession Session: Lessons Learned the Hard Way

Contributors: Tobi Delbruck, Ibrahim (Abe) M. Elfadel, Shahzad Muzaffar, Germain Haessig, Bo Wang, Amine Bermak, Rui Graca, Luis Camuñas-Mesa, Bathiya Senevirathna, Pamela Abshire, Saeed Afshar, Bernabe Linares-Barranco, Shih-Chii Liu, Runchun (Mark) Wang, Piotr Dudek, Stephen Carey, Jose de la Rosa, Marc Dandin, Sheung Lu, Vincent Frick, Teresa Serrano-Gotarredona, Paula Lopez, Melika Payvand, Advait Madhavan, Eric Fossum, J. Camilo Vasquez Tieck, Ian Williams, Yan Liu, Timothy Constandinou, Alexander Serb, Ricardo Carmona-Galan, Robert Nawrocki and Walter D. Leon-Salas.

Organizers and Editors: Bernabe Linares-Barranco (Seville Microelectronics Institute), Teresa Serrano-Gotarredona (Seville Microelectronics Institute), Tobi Delbruck (University of Zurich and ETH Zurich), and Walter D. Leon-Salas (Purdue University)

Abstract—“Fail often to succeed sooner” is a common mantra that we are told is the secret to success. When reporting research results, however, scholars rarely write about their failed attempts and only focus on the successful ones. Perhaps the source of this disconnect between what we preach and what we do can be found in the underlying assumption that published work is meant to move the field forward and failed attempts supposedly do not. The goal of the confessions presented in this paper is to show that even failed attempts are genuine and valuable contributions to our field provided that we learn from our mistakes and correct them. The 27 confessions span from planning oversights, digital and analog design errors, misunderstanding of devices, overlooked parasitics, LVS errors, and troubles in testing.

I. RATIONALE

Most technical papers present only successful results, while little is said about the unsuccessful attempts and troubleshooting pains that the authors went through until they succeeded. The implicit assumption is that unsuccessful attempts do not move the field forward and therefore are not worth publishing or sharing with the broad engineering community. In practice, however, mistakes and failures are an integral part of doing research and of the learning process. It is our hope that the publication of the mistakes, failures and the solutions presented in this paper will help other researchers avoid costly mistakes. Mistakes at the microchip level, for instance, are difficult, if not impossible to fix, resulting in unforeseen delays of months, additional costs of tens of thousands of dollars, and delays in student graduation and publications. We expect students to be the greatest beneficiaries of these confessions because they will hear from others painful experiences and understand that failure is part of the learning process. They will see that even the most seasoned and experienced designers make mistakes, and hopefully will learn techniques to help avoid making their own.

Confession Sessions have been organized before in IEEE events. For the 2009 ISSCC, Jed Hurtwitz organized a live forum called *Forewarned is Forearmed: Classic Analog Mistakes to Avoid*. It included a website to which designers

uploaded their confessions. Although this forum was not published, during the conference each confessor was allotted few minutes and one slide to explain what went wrong in their designs. Inspired by Jed’s forum, two of the coordinators of this year’s session organized the first ISCAS confession session at the 2011 ISCAS in Rio de Janeiro. 25 confessions were contributed by members of the Sensory Systems Technical Committee. This session proved popular with a full room of at least 80 attendees. The resulting confessions paper [1] by all 21 authors has only been cited 10 times but hopefully had a larger impact. The slides from this session are still available¹.

The host city Seville, with its many magnificent Catholic churches and its rich history, provides a suitable backdrop to a confession session like this one. Similar to the format we had in Rio in 2011, this year each confessor will have three minutes and three slides to present his or her confession. These presentations will highlight the troubleshooting approach, namely, how they found the problem and how they fixed it. The audience will be polled for their favorite confessions and awards will be presented to winners. The slides will be posted on the web.

History repeats itself. Compared with [1], we have the same categories and the same types of errors. This time, we gathered a total of 29 confessions. We organized them into the following categories: three troubleshooting pains, two analog goofs, five logic design errors, four layout-vs-schematic goofs, two process design oversights, five not accounting or understanding parasitics, and eight planning errors. They follow in this order. The paper concludes with a checklist that might help to catch some of these types of mistakes.

II. TROUBLESHOOTING PAINS

Confession 1: Two problems masked each other
(Tobi Delbruck, UZH-ETH Zurich)

¹http://sensors.ini.uzh.ch/news_page/confession-session-2011.html

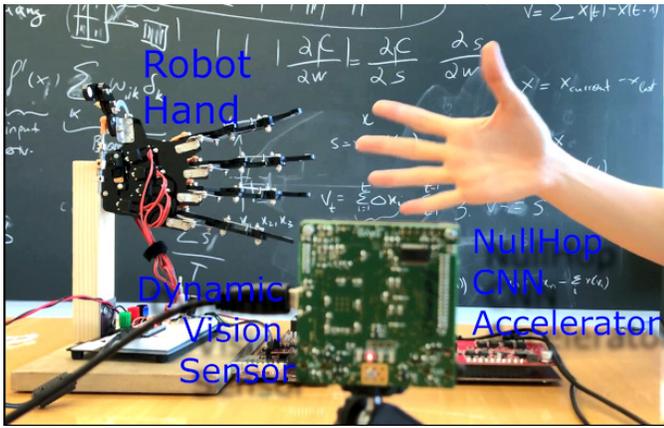


Fig. 1. The rock-scissors-paper system which had two bugs masking each other.

When trying to make the live rock-scissors-paper robot demo shown in Fig 1 out of our DVS event camera driving our CNN accelerator *NullHop* [2], we faced continual problems with FPGA lockup. The system would run for a while, then mysteriously stop. The original *NullHop* logic designer had long before left us, so we faced the challenge of trying to debug the logic ourselves. Since the design was done in a rush with a lot of spaghetti logic, even after 6 months we could not figure out the root cause of the problem. Finally we managed to reproduce it: If the input image from accumulated DVS events contained any row with all zero values, the logic would hang. Of course this bug was not exposed in simulation, since we never changed the input image presented to the CNN.

Although we could never figure out how to fix the logic, the workaround appeared to be simple: We made sure in software running on the SoC ARM core that no row was all zeros. But it would *still* hang, after apparently random times. NIPS was getting close and the teams in Zurich and Seville were yelling at each other. The Seville team didn't get the hangs but we in Zurich did. Finally we had a visit from the AVENT Xilinx representative. He said, "Oh yes, we know this power module has problems with the Zynq MMP. Use this other one instead." That fixed it and finally the demo would run stably.

Moral: I supposed it is to consider that there might be two bugs that result in the same symptom. The best clue was that the teams had slightly different hardware setups. Also, it never hurts to talk to the suppliers for clues to problems that might not be written down. Finally, your test bench needs to test real world (variable) inputs.

Confession 2: When workbenches look similar but are not the same

(Abe Elfadel and Shahzad Muzaffar, Khalifa University)

Workbenches in our lab are usually placed next to each other. One workbench is typically assigned to each student/staff for experimental work. Each of these workbenches is equipped with six power outlets. In one of our projects, we had a rather large experimental testbed that needed to be spread up over two workbenches. In particular, powering the instruments, the desktops, the boards, and displays of the

testbed required more than the six power outlets available on a single workbench. The testbed included two computers (Laptop+CPU+LCD), two FPGA boards, oscilloscope, function generator, multi-meter, power supply, and an LCD for the FPGA display. One of the FPGA's was an expensive one of the high-end type – a recent addition to our digital lab that we were quite proud of. Like two tables in a restaurant brought together for a big party, it was no brainer to bring two of our workbenches together to spread out our experimental setup and use the power outlets of both workbenches to power up all our nicely arranged boards and instruments. Little did we know about the dangers that were lurking within such trivial combination of similar workbenches.

Incidents: Three incidents with increasing intensity occurred during experimental work on one of our IoT projects [3]. In the first incident, a spark occurred when we connected the USB port of the FPGA board (powered from Workbench 2) to the laptop (powered from Workbench 1). As a result, the computer motherboard was burned out, and we of course blamed ourselves for some sloppy wiring of ours that caused a short circuit in our setup. We got the motherboard replaced and went back to the workbench for debugging but could not find anything wrong in our wiring. So we proceeded to connect an LCD (powered from workbench 2) to the FPGA board (powered from Workbench 1) and again a spark occurred, the second incident, with immediate tripping of the circuit breaker in our building. Again, we could not find anyone to blame but ourselves, and we thought that perhaps a short circuit was crated while we were trying to connect the display cable to the FPGA DVI port. The FPGA board and the LCD were completely burnt out. We got them replaced and went back to the workbenches for further debugging and experimental work. The third incident was the most intense and it occurred when we tried to connect the USB of the FPGA board (powered from Workbench 2) to the laptop (powered from Workbench 1). A huge spark with a cracking sound suddenly flashed, followed by an immediate tripping of the circuit breaker in our building. This time, luckily, the laptop was fine, but the FPGA, the expensive one we were proud of, was completely burnt out. This spark was so intense that the USB port on the FPGA board melted from one side, as shown in Figs 2. Luckily the brave student was not in direct contact with any metal of the FPGA board. Otherwise, he might have been electrocuted.

As it turned out, the fault lay in the workbench power connections. But to figure this out, we had to go through lots of soul searching, head scratching, and friendly hand holding by collegial colleagues of ours. In one instance, we asked the IT folks if they knew of any similar incidents involving the brand of laptops we were using and invited them to come over and check our laptop connections. Interestingly and in a typical IT style, they told us they did not need to come over because they could help us "remotely." The aha moment happened when one of our down-to-earth electronics colleagues brought over his multi-meter and measured the ground-to-ground voltage on the power outlets of one workbench and then the ground-to-ground voltage on two outlets, each belonging to a different workbench. The former was 0 V but the latter was 240 V! In

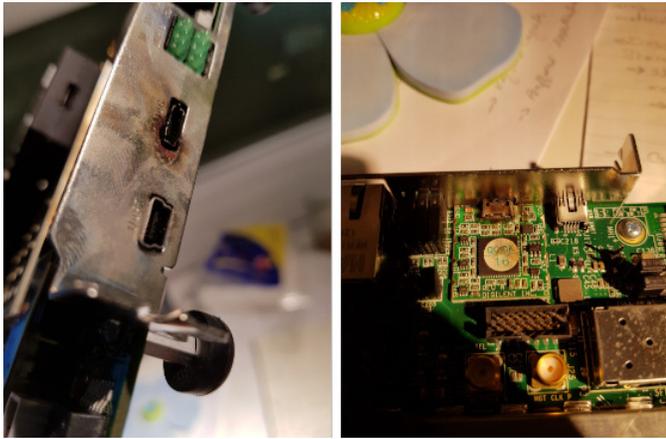


Fig. 2. USB Port Meltdown on the FPGA Board

other words, while ground on Workbench 2 was actual ground, it was a live wire on Workbench 2. Our connectors from a device powered from Workbench 2 to another powered from Workbench 1 with their ground-pin-to-ground-pin wires were causing the shorts to happen and our laptops, LCD, and boards to fry. At this point we informed facilities of what happened. They were extremely pleased to know that the damage was restricted to lab equipment and not people.

Moral:

- 1) Inform facilities on the first spark in your lab. The third spark might be too late.
- 2) Be humble and ask for help.
- 3) Don't take your lab infrastructure for granted.
- 4) Lab workbenches are not cafeteria tables.
- 5) Be down-to-earth and stand on solid ground.

Confession 3: Linear regulators waste a lot of power, especially if you do this!

(Germain Haessig, UZH-ETH Zurich)

We were building up more samples of the SEEBETTER project DAVIS event camera lab testing boards for use in a couple of our new projects. Tobi Delbruck had previously realized that it was a dumb choice to use a linear regulator to power 1.8V logic from 5V USB power since it made the cameras hotter which killed low light performance. Federico Corradi had designed a tiny interposer that replaced this 1.8V regulator with a more efficient switching regulator. But there were still some anecdotal reports of poor power supply regulator reliability on these boards. We had not replaced the regulators for the sensor analog supply voltages because we knew that switching regulators introduce a lot of noise. Even though the analog regulators still seemed to be very hot considering the small output current they needed to supply to the sensor, the cameras worked, and so no one considered it further.

I was looking through the board schematic one day, when I noticed something strange (Fig. 3). It seemed like the regulator output V_{out} was connected to ground. How was this possible? It turns out that for this regulator, the large bottom pad, that someone many years ago had assumed was ground,

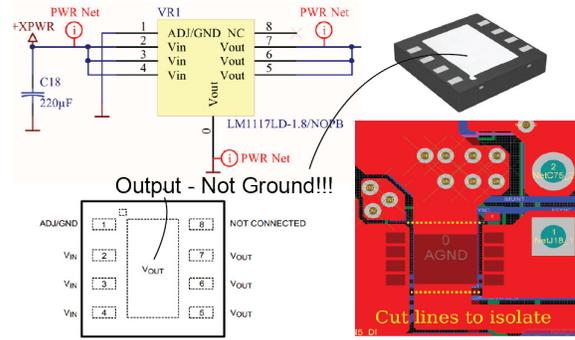


Fig. 3. This regulator output is shorted to ground, because the bottom solder pad is **not** ground as someone had assumed. The workaround was to resolder the regulator after cutting through the pad ground connections with razor.

was actually the regulated voltage output. As a result, these regulators were shorted, supplying a huge current to drive this resistive connection to the regulated output voltage. This problem was not easy to fix; the chip scale package needs to lay flat on the board, so to insulate the bottom pad, I had to desolder the regulator with a heat gun, and very carefully cut across the ground layer connections to the pad. After resoldering the regulator, we don't have any more problem and the voltage regulators do not get burning hot!

Moral: Not all bottom pads are ground. Don't take things for granted and always carefully check the component datasheets!

III. ANALOG GOOFS

These sometimes subtle errors were the subject of JED Hurwitz's original ISSCC forum. They can occur because of 2nd-order behavior of transistors or their bulk PN junctions, for example, that are not considered (or modeled properly in simulation). Or they can result from simple inexperience.

Confession 4: A simple switch can kill the whole system

(Bo Wang and Amine Bermak, CSE, Hamad Bin Khalifa University)

We designed an ultra-low-power sensor interface SoC with multiple low-power circuits on the same substrate. Because we do not want to include the leakage current of the ESD devices into our measurement, only a bare pad was used for the power supply. We did this many times and it turned out to be safe during lab characterization. Unfortunately, this time, we added a few power switches in order to measure the power consumption of the building block separately [Fig. 4(a)]. Conceptually, this is very simple by adding a PMOS transistor and an enable signal, as shown in Fig. 4(a). However the measured system power consumption was on the order of tens of mA, which is way larger than the designed $<10 \mu\text{A}$ budget. We tested several samples and observed that there is a linear current-voltage relationship in the supply even when all the sub-circuits were disabled.

After careful circuit review, we identified a risk in the power switch design. As shown in Fig. 4(b), for the PMOS switch, a +P-sub (V_{ss}) and $-V_{dd}$ ESD event can induce a large surge current flowing via its substrate resistor and parasitic diode D_1 . In our design, the switch is small and so is the parasitic diode,

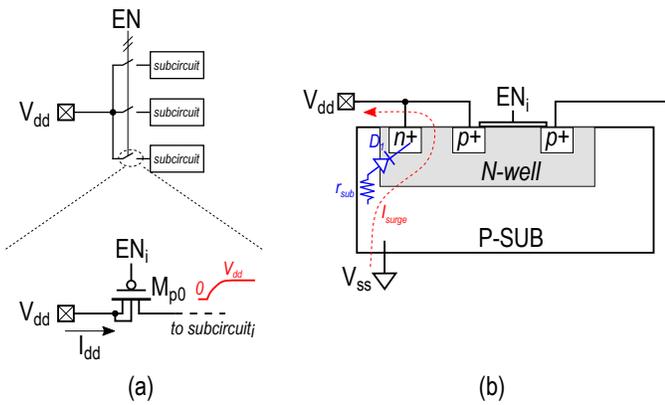


Fig. 4. (a) Switches added in the system for individual blocks' power measurement; (b) Illustration of switch diode breakdown.

so it cannot survive large current and burns out. Even worse, since we designed to have low resistance substrate tapping at the top level, the path resistance between the substrate and the V_{dd} pad is very low. As a result, the surge current can be huge, causing thermal-runaway and parasitic diode breakdown. A low impedance path then developed between V_{dd} and the substrate, resulting in the observed large current draw.

After understanding this, we revisited our previous designs that also did not include supply ESDs. In those designs, the V_{dd} pads were connected to the source of hundreds of PMOS transistors, so more parasitic diodes like D_1 exist and the path resistances between the P-sub to V_{dd} path are larger as a result of less substrate contacts. This helped to limit the amplitude of the surge current, helping those chips survive ESD events.

Because of this, we could test neither power consumption nor function of this chip since the circuits don't get the correct supply. A new tape-out solved this issue by deleting the power gates. Separate pads were used to perform current measurements for different building blocks.

Moral: Never under-estimate the importance of understanding the detailed process and parasitic influence on the designed circuit, even for a simple switch.

Confession 5: What the hell is this GMIN anyhow?

(Rui Graça, UZH/ETH Zurich)

At the beginning of my PhD, I was simulating the leakage current of cut-off PMOS transistors, which I wanted to use as pseudo-resistors for very low frequency high-pass filter [4], [5]. The results I was getting were strange: For the small voltages I cared about, the current increased linearly with drain to source voltage rather than exponentially.

I knew that these results could not be correct, but with my inexperience with the tools and the models, I could not figure out what was wrong. My advisor Tobi Delbruck pointed out that I needed to define a very low value for GMIN, which is the conductance that SPICE adds across all PN junctions and across FET source-drains [6]. I had been using the simulator default value of 1pS, which perfectly explained the results I was getting: When the leakage current was very low, all the measured current was due to the linear GMIN current.

After reducing GMIN to 1e-30, the simulation results started

to make sense.

Moral: Understand the simulator configuration options before running simulations.

IV. DIGITAL LOGIC MISTAKES

Confession 6: A missing bit ruins operation of AER Convolution Chip

(Luis Camuñas-Mesa, IMSE-CNM, Sevilla, Spain)

We had designed our nice first fully digital 32x32 pixel Convolution chip [7] and it was working great. Fig. 5(a) shows the floor-plan of such a chip. It included a digital controller (described and synthesized using vhdl), an array of computing pixels, an AER read out circuit, a kernel block RAM, and a left/right lines shifter. For each incoming Address Event, the Convolution kernel is copied from the RAM to the pixel array line by line. Depending on the event x-coordinate, the lines connecting the RAM to the pixel array go through a switch-matrix based left/right shifter to align the kernel weight lines properly with the incoming x-coordinate. We then designed an improved version with a much more compact pixel and improved kernel capabilities, resulting in a 64x64 kernel Convolution Chip with multi-kernel capability. However, when we received our test chip back and started to operate it, it showed a pair of nasty vertical lines, as in Fig. 5(b). After several days of test, changing samples, looking at the schematics, re-simulating parts (it was not viable to do a full chip transistor-level simulation), we found out that one of the control lines going from the controller to the left/right shifter was not connected. It was also not connected in our high level schematics. Since this mistake would only produce errors for some events (those which needed a shift of 32), we never detected the fault during our behavioral simulations. We corrected the error quickly, resubmitted the chip for fab and at the end everything worked fine [8]. But this little mistake cost us at least 6 months and thousands of Euros of fab money to fix.

Moral: if you have a little bug in your highest level schematics, it may never show up during design verification. Watch out carefully your high level descriptions and try to verify them exhaustively.

Confession 7: Wrong initialization condition causes chip failure

(Bathiya Senevirathna and Pamela Abshire, University of Maryland, College Park)

We recently fabricated an application specific integrated circuit (ASIC) chip to interface with a set of gas sensors, which were essentially passive resistors arranged in a standard bridge network. The chip has a programmable gain amplifier (PGA) which amplified the bridge network's output voltage. This was followed by an analog-to-digital conversion stage, and a digital logic controller that controlled sampling rates, applied gain settings, coordinated data communication off-chip, and performed other housekeeping tasks. One of the major design goals was to overcome technical challenges associated with the

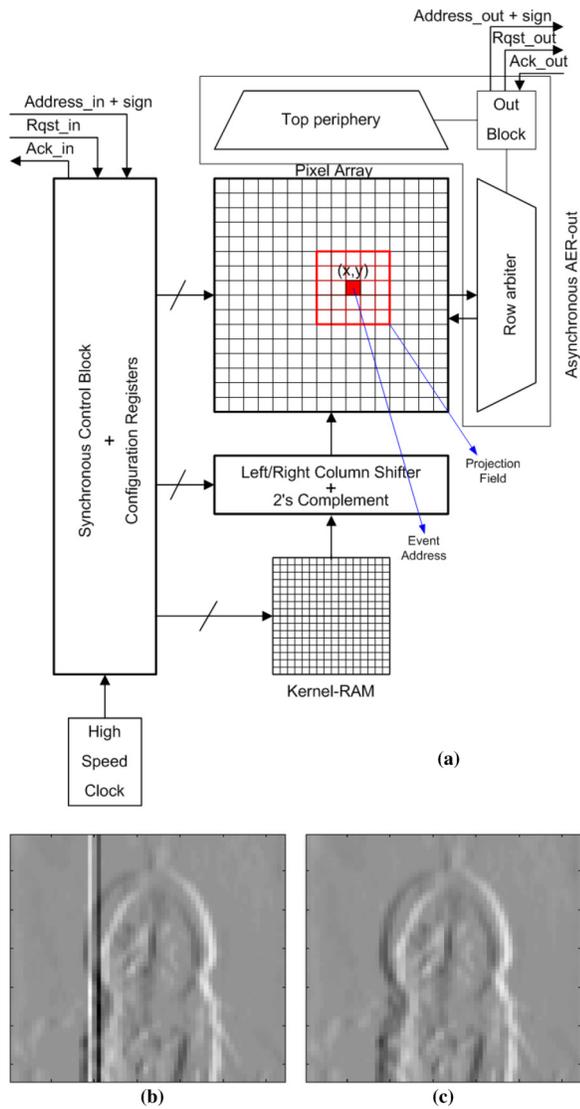


Fig. 5. (a) Convolution chip floor-plan. Output of event-driven convolution processing before (b) and after bug correction (c).

unknown baseline resistances of the sensing elements, which can diminish the PGA's output dynamic range, even causing the signal to saturate when no gas is present. Therefore, we implemented a calibration procedure to automatically drive the PGA's output to mid-rail during power-on, regardless of the intrinsic baseline resistances of the sensing elements. The PGA's uncalibrated output is first compared to the target level (mid-rail) in order to set the direction of correction. Then a series of on-chip current sources are used to automatically trim PGA's output to the desired level using a binary search algorithm. The trim current settings are then stored in a calibration setting register. Validating the calibration procedure and the PGA using circuit simulations was a little tricky since the logic was created as part of the much larger controller module (synthesizing and simulating everything was a non-starter). So we ended up modifying the calibration module to be self-contained, synthesized it, and then simulated it with the PGA. Everything seemed to look good.

Fast-forward a couple months and we received our fabri-

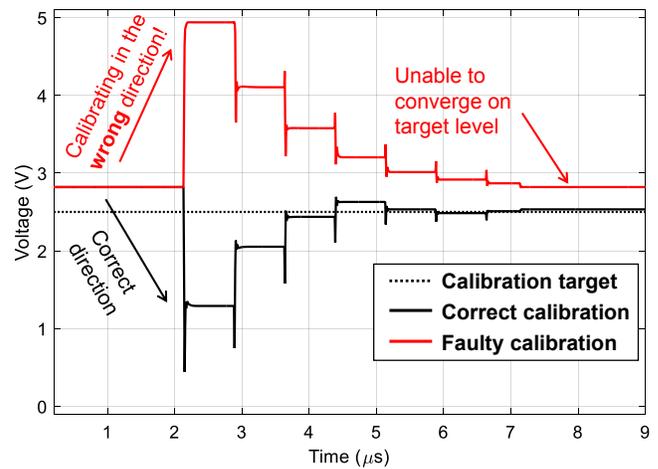


Fig. 6. Transient plot of the chip's output during the offset calibration procedure, and the effect of the wrong initialization condition. The output is initially at 2.8 V but needs to be trimmed down to 2.5 V. With the correct calibration procedure, the trimming aims to push the output voltage downwards, allowing it to eventually reach the 2.5 V target. On the other hand, the incorrect initialization causes the trimming to only push the output upwards, making it impossible to converge on the target.

cated chips. As an initial test the chip was tested with ideal voltage sources instead of the bridge network. However to our dismay, the ASIC was consistently pushing out saturated values. Not only that, but it looked like calibration procedure was making things worse, i.e. if there was a minor negative input differential, the chip saturates to GND and if there was a positive input differential, the chip hits VDD/Vref. After some debugging it appeared that a single bit in the calibration module was initialized to a 0 instead of a 1. This had the effect of flipping the direction of correction that was set during the very first step of calibration. So for instance, if the nominal PGA output was 100 mV *higher* than mid-rail, the initial trimming direction was set to go *even higher* (see Fig. 6). This behavior essentially made the chip a space heater. How did it work in simulation then? Well it turns out this problem was indeed seen in circuit simulations, and the initialization issue was fixed immediately. However the fix didn't propagate up to the final top-level HDL which had been undergoing separate modifications in parallel. Fortunately the calibration setting register were designed to be read/write and so after the buggy automatic sequence is finished, we could performed a second corrected calibration sequence by implementing the same binary search algorithm in software.

Moral: Keep good track of changes to your source files and use version control, especially in mixed-signal design where design/simulations have to be done using multiple disparate tools.

Confession 8: Using accumulators when counters work better.

(Saeed Afshar, ICNS, Western Sydney University, Australia)

We designed a digital hardware model of combined STDP and Spike Time Dependent Delay Plasticity (STDDP) which modelled the entire synapto-dendritic coupling between neu-

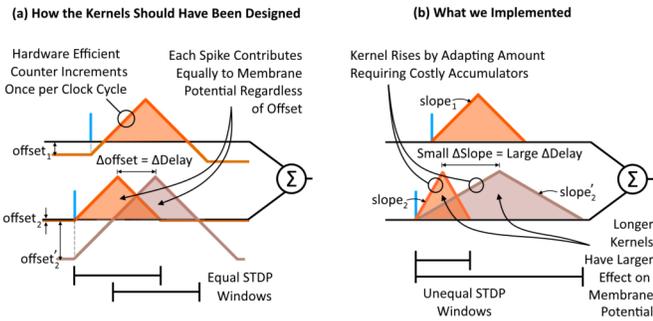


Fig. 7. Using accumulators instead of counters to implement STDDP and STDP. (a) Offset adaptation, good. (b) Slope adaptation, bad.

rons using triangular kernels whose slopes adapted via a simple input-output spike time plasticity rule [9][10]. Our aim was to make “super synapses” which modelled the entire coupling of multiple synapses and dendrite delay lines via a single hardware efficient mechanism. Trying to avoid digital multipliers, we used triangular kernels implemented using accumulators such that only the kernel slope needed to be adapted. (An accumulator is a counter that increments by a configurable integer quantity, in this case specifying the slope.) After coming up with the model we quickly published multiple papers, used the model for hand gesture recognition in hardware and even had others work on and extend the model [11][12][13][14], all before we realised we could have significantly improved the performance and hardware resource costs by just using counters which increment by one at each time step instead of by a variable amount.

As shown in Fig. 7 replacing the adaptive kernels with counters of slope one and an adaptive offset was a better design choice. Using adapting slopes for delay adaptation not only reduced the dynamic range of delays that could be encoded by the kernel, it also increased circuit complexity relative to kernel offset adaptation. Finally the slope adapting method meant that later spikes contributed less to the neuron membrane potential in general and in the presence of timing error in particular making them far less effective than early spikes regardless of their relative information content. This meant that the system was biased toward early spikes. Now, we could spin this as a feature [15][16]. But really, it was a bug. Our aim with the STDDP rule was to provide timing invariant inter-neuronal kernels for the STDP to weigh, not to tune neurons to early spikes. This oversight in design (adaptive slopes instead of offsets) was too small to warrant a publication to rectify, but too big to allow a clear conscience, until now.

Moral: Take more time to work through and optimise the fundamental blocks of a system before rushing off to build on it.

Confession 9: Late CDR (Clock Data Recovery) reset due to excessive delay

(Bernabe Linares-Barranco, IMSE-CNM, Sevilla, Spain)

We designed a bit serial high speed AER (Address-Event-Representation) interface optimized to transmit burst-mode digital traffic with low latency and low power. The request

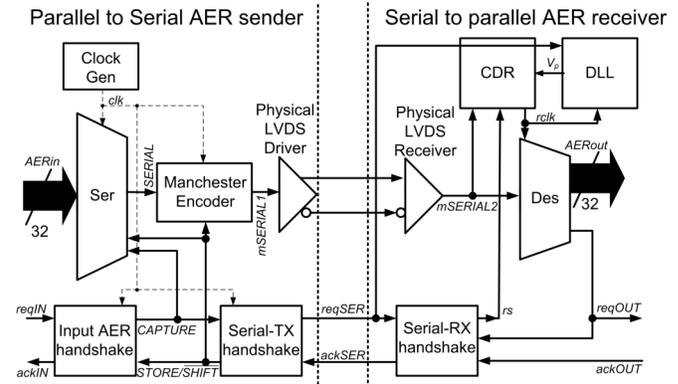


Fig. 8. Clock-Data Recovery (CDR) Circuit Block Diagram. Skew between the handshake and LVDS signals corrupted early parts of each transmitted word.

and acknowledge handshaking signals were used to reset the CDR status before any incoming event (see Fig. 8). Basically, a double edge triggered flip-flop is reset when the request signal presents a certain digital level. System level simulation were carried out to check the link performance at transistor level. When the chip came back and we started to test it, we observed that the first bits of the serial transmissions were lost. After studying the phenomenon, we noticed that the data stream was arriving at the receiver nano seconds earlier than the Request signal, which was triggering the whole process. The reason was that the Request signal was sent through normal pads, while the data was transmitted through high speed LVDS pads. In our simulations we did not include the delays associated to digital pads and PCB traces nor parasitic elements of the higher level circuit connections. Hence, the request signal had rising/falling times in the order of ns. Because a 500 Mbps Manchester-encoded serial bit-stream was transmitted through the channel, the request signal was arriving at the receiver when the bit serial circuits had already transmitted some bits and the digital words were not decoded properly at the Receiver. We discovered the error when we simulated the whole experimental set-up including all off-chip parasitics. We corrected the design to compensate for this, resubmitted the design, and in the end everything worked fine [17].

Moral: Do not forget any parasitic in your simulations, even your off-chip ones.

Confession 10: Can 2 wrongs make a right?

(Shih-Chii Liu, UZH-ETH Zurich)

In one of our first spiking cochlea designs, we made a couple of key mistakes. The first one was in the AER Chip REQ/ACK signals that were brought off chip. Following the AER protocol, the Chip REQ should come from the request in the 2nd dimension (Col REQ) for a 2D array. Here we made the mistake of bringing the Chip REQ from the Row REQ of the 1st dimension. It meant that the chip ACK (from the receiver) would come back too quickly most likely before or during the arbitration for the 2nd dimension. We mitigated some of the effect by including a delay in the chip ACK but this alone was not sufficient to guarantee proper transmission

of the address of the pixel that would have been selected.

However, there was a second mistake on the chip which helped us to overcome this problem. On this chip, we had included local latches for a KillBit in the cochlea channels, so that the AER circuits of individual channels can be shutdown if certain channels are not needed for the processing.

The second mistake was in the polarity of the latch signal which meant that when the Latch signal was activated, the Killbit was loaded into all channels *except* for the chosen channel. We used this second mistake to offset the mistake with the Chip Req.

The solution was to pulse the Killbit signal to switch off all channel requests for a short time after the chip ACK is received. That makes it likely for the most active channel to transmit its address before any of the other channels. These corrections allowed us to report on the functionality of the first spiking cochlea design [18].

Moral: Do not give up hope when you encounter the first mistake. By imaginative use of the existing logic, once in a while, two wrongs can make an approximate right.

V. LVS CONFESSIONS

Layout-versus-schematic chip design errors are well-known to all experienced chip designers. The chip perfectly passes LVS, but doesn't work.

Confession 11: Flipping the bond pads

(*Runchun (Mark) Wang, Western Sydney Univ.*)

A few years ago, we were working on several neuromorphic chips, e.g., SNN processors and cochlear. They are all fully digital chips, so we followed the standard digital design flow: RTL design, verification, FPGA prototyping, synthesis, Place and Route, conformal verification and DRC. To improve the performance and ease the future design migration (for an advanced process), we used ARM standard cells and pads instead of the ones provided by the foundry. Everything went well after several months of hard work. When it came to the final stage, we need to place the pads, which consist of IO modules (drivers and ESD devices) and bond pads (just the top 2 metals). By its default setup, the bond pads are in the inner ring instead of the outer ring. Based on our previous experience, the bond pads are all in the outer ring, which is easier for bonding. So we manually flipped all the bonding pads from inner ring to the outer ring. Since we used the digital flow, all these cells are in the LEF format and we can't see the actual layout. We generated the GDS at the tapeout.

When we tested these chips, we found that they all have power shorts. Due to the fact these chips were designed by several designers separately, we suspected that it might be process problems. The engineer of the foundry tested the back up dice and told us that there are shorts in Metal 5 in the pad rings. Then we checked the GDS and confirmed this. Finally we realised that these bonding pads must be in the inner ring.

Moral: LVS with the pads is a MUST. Never change the default settings of the pads, unless you must, and you are quite sure what you are doing.

Confession 12: Getting away with wrong pads on the pad ring.

(*Piotr Dudek, Stephen Carey, The University of Manchester*)

When we were designing the next generation of the SCAMP mixed-signal sensor-processor chip [19], we decided to put quite a bit of control logic on chip. As you do with a prototype, we routed out some key signals from internal logic circuits to the output pads, so we can monitor these to make any debugging a little easier. As is unfortunately often the case, the standard-cell libraries available to us were not quite complete. We had I/O pad layouts, but not schematics. We decided we can live with this, since we correctly assumed the performance of the pads was not going to be critical for this chip. Clearly, without schematics the full LVS including the pad ring was not readily available, so we did the usual – LVS the core, and then just really, really, really, I mean, really carefully check all round that all core signals connect to the pads as they should. Check again. And again. It's not that hard. And no, we did not mess up our wiring. The chips came back, all pads were nicely connected to the core as they should. We started testing but the chip was behaving erratically. It took a while to figure this one out. Turned out, in one instance, instead of placing a digital output pad, we used an input pad! This was supposed to be used to monitor one key control signal (nicely buffered of course). This clearly is a bit of a problem as our buffer and the input pad now both drive the same net inside, shorting the core's Vdd and Vss if they don't agree, which breaks the chip. The workaround? We replicated all control logic that is there inside the chip on an FPGA, in order to externally drive what was supposed to be a test pad in exactly the same way that the signal we were supposed to monitor is changing. We had to carefully adjust the timings, to make sure all runs exactly in sync. With this, everything else on the chip functions okay. We need to be running the same logic in parallel on-chip and of-chip, and in fact we are using it like that in the camera systems [20] we build based on the chip.

Moral: Figure out how to do a full LVS. Including the pad ring. Or just really, really, really, I mean, really carefully check things all round!

Confession 13: Bias Generator Circuit not Working as Expected

(*Walter D. Leon-Salas, School of Engineering Technology, Purdue University, USA*)

Our confession comes from a bias voltage generator circuit that we implemented on a 0.5 μm CMOS process. This circuit was part of a larger chip that had an experimental energy-harvesting 64 \times 64 CMOS imager [21]. The chip also included analog read-out circuits and an ADC. The bias generator circuit, shown in Fig. 9, was meant to provide the appropriate biasing voltages for all the circuits on the chip. The core of the bias generator circuit is a standard beta multiplier (with a kick starter circuit) that generates output voltage V_{biasP} . From this voltage we needed to generate the other bias voltages needed in the chip. To accomplish this we decided to use a MOSFET-based voltage divider that used strings of diode-connected transistors of the appropriate sizes. The basic idea

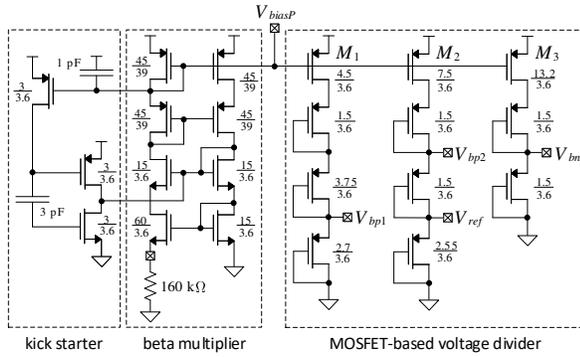


Fig. 9. Designed bias voltage generator circuit. Incomplete LVS and a bad connection threw the generated biases way off.

was to use the voltage V_{biasP} and transistors M_1 , M_2 and M_3 to generate currents of known values. As these currents flow through each diode connected transistor, a gate-to-source voltage proportional to the W/L ratio of each transistor would be developed. By adjusting each transistor size, the needed voltages could be generated. Following this approach and using computer circuit simulations, we found the transistor sizes shown in Fig. 9.

We expected to have a difference of about 5% to 10% between the designed and the measured voltages from the fabricated chip due to variations in the fabrication process. Table I shows the designed bias voltages (from pre-layout circuit simulations) and measured bias voltages from three fabricated chips (units are volts). The measured values for V_{biasP} and V_{bn} are within 10% of their designed values. However, for V_{bp1} , V_{bp2} and V_{ref} , the differences were much larger. After inspecting the layout we found a missing connection in the middle branch of the voltage divider circuit and a transistor with the wrong size in the first branch. Had we used an LVS that checked layout connectivity and W/L , or performed a thorough post-layout simulation we would have caught this mistake. On a positive note, we were sensible enough not to hardwire the voltage divider circuit to the other analog circuits on the chip but instead the bias voltages we brought to output pins on the chip. Thanks to this arrangement we were able to bypass the faulty bias generator circuit and provide bias voltages to the chip using an external DAC.

	V_{biasP}	V_{bp1}	V_{bp2}	V_{ref}	V_{bn}
Designed	2.65	0.97	2.58	1.6	2.92
Chip 1	2.891	0.377	0.053	0.052	2.905
Chip 2	2.884	0.387	0.052	0.056	2.90
Chip 3	2.915	0.422	0.054	0.059	2.932

TABLE I
DESIGNED AND MEASURED BIAS VOLTAGES

Moral: No matter how confident a layout designer might feel or how simple a circuit might look appear, it is always a good idea to run post-layout simulations and perform a thorough

LVS check. Also it is a good idea to design a chip in a modular manner such that individual sub-circuits can be bypassed or isolated in case they do not work as expected.

Confession 14: Pay Attention to Warning Messages given by Layout Verification Tools

(Jose M. de la Rosa, Instituto de Microelectronica de Sevilla, IMSE-CNM (CSIC/Universidad de Sevilla), Spain):

In 2003 we were designing a $\Sigma\Delta$ modulator for an automotive sensor A/D interface as part of a research project. It was the first time I was in charge of a design team in my institute, so I was very proud and honoured for my new role. We were using a $0.35\mu\text{m}$ CMOS technology which was under development those days, and our industrial partner was interested in testing the technology with high-performance analog cells. One of the consequences of using such a new technology was that the design kit was also under development, which added extra effort for designers. For instance, parasitic extraction was not available, and hence we had to be very careful and conservative during the design process, especially at the layout phase. Indeed, this was not the first time we were in that situation, because we usually work with cutting-edge technologies.

During the design phase, all verifications were carried out and the circuit was (apparently) working according to the required specifications. So, the chip was sent for fabrication according to the project planning, and it came back from the foundry a few months later. Everything was going well and I was very happy about that. However, as illustrated in Fig. 10(a), experimental measurements revealed a severe performance degradation as compared to the electrical (HSPICE) simulations. After a long and tedious debugging process, that involved manual checking of the entire layout (containing over 1000 transistors), we noticed there was a floating NWELL in one of the PMOS transistors of a CMOS switch – highlighted in Fig. 10(b). Obviously this error affected the signal transmitted from the first to the second integrator as well as the signal fed from the DAC back to the second integrator. We modelled this error at system level as an attenuation in the feedback DAC, and it gave very similar results to the experimental measurements.

After we found the problem, we were wondering how this error had been masked by DRC/LVS tools. This was explained during one of our multiple internal meetings that took place in those days. One of our PhD students admitted that she got a “warning message” related to the mentioned floating NWELL. However, she thought that this was not an error, but just a “warning message” and hence, since the circuit simulations were going well, she decided to ignore that message. In the end, we fixed the layout error and sent the chip for fabrication again and the new chip featured a state-of-the-art performance that was published [22].

Moral 1: Do not underestimate any warning message given by circuit verification CAD tools. It might happen that what you consider a minor issue may be the cause of a catastrophic failure. If you are an inexperienced analog designer, ask your supervisor.

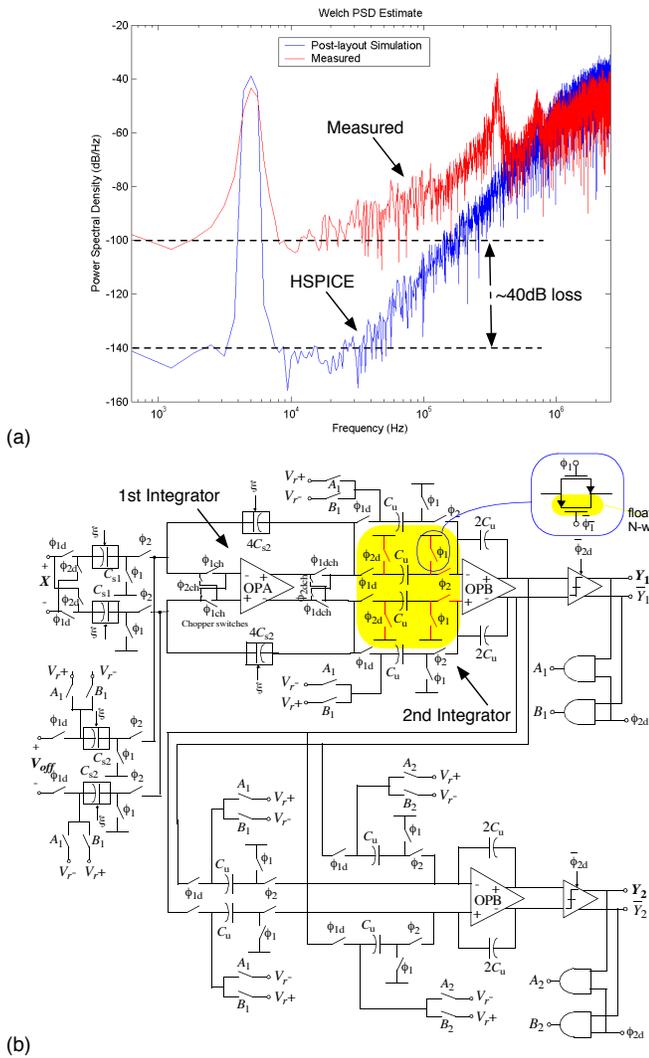


Fig. 10. Illustrating the performance degradation in a cascade 2-1 SC $\Sigma\Delta M$ caused by a floating NWELL. (a) Comparison between experimental and simulated output spectra, showing how the noise-shaping is severely degraded. (b) Circuit schematic highlighting the CMOS switch with the floating NWELL.

Moral 2: If you are a senior designer in charge of a design team where there are junior designers, do not assume that they know what you are supposed to know, and supervise their work in detail.

VI. PROCESS DESIGN OVERSIGHTS

These types of confessions are usually the result of not reading the process specifications carefully while doing the design.

Confession 15: Packaging Small ICs for Biosensing Applications: Beware of the Outer Metal Frame around Your Design!

(Marc Dandin, Carnegie Mellon University; Bathiya Senevirathna, Sheung Lu, and Pamela Abshire, University of Maryland, College Park)

Packaging single integrated circuit chips to enable biosensing applications is quite challenging, especially at the research

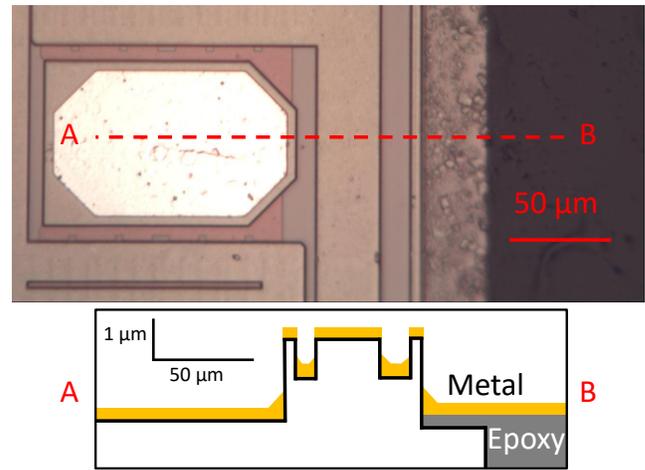


Fig. 11. Photomicrograph showing one side of an IC chip that is encapsulated in an epoxy substrate for biosensing applications. The metal ring surrounding the chip design introduces unwanted topography that causes a lot of difficulties during metallization; specifically, it will be difficult to deposit a metal conformally given this topography. See [23].

stage when prototypes are needed to validate novel biosensor technologies. Packaging challenges are exacerbated by the size of the integrated circuit chips that are typically available to researchers. For example, typical die sizes available to researchers can range from $1 \times 1 \text{ mm}^2$ to $5 \times 5 \text{ mm}^2$; these die sizes limit the number of I/Os that may be available for biosensing applications because wide pitches (100 to 200 μm) must be used in order to accommodate subsequent metallization and encapsulation processes that are needed to integrate a chip into a microfluidic platform. We have developed several processes for packaging such small ICs for biosensing applications [23], [24], [25]. Other groups have also addressed the issue of packaging small ICs [26].

We have demonstrated an epoxy process to encapsulate a small IC in a planar substrate for subsequent microfluidics integration and biosensing applications [23], [24]. While this process worked well for some of our initial ICs, we discovered problems for ICs where a metal frame is used to surround the pad frame. These metal frames are typically used by a foundry to demarcate dice from one another on a wafer. For biosensing applications, this outframe can be a nightmare, as it will cause issues in metallization when integrating the IC in a planar substrate. A typical scenario where this outer frame may cause problem is illustrated in Fig. 11, where the topographical features associated with the outer metal frame make it hard to deposition conformal metal layers for I/O traces.

Moral: Metal traces at the outer edge of an IC design layout can complicate integration. It is advisable to work with the foundry to avoid including those frames in layout; alternatively, plan to leave a significant buffer space between the metal frame and the pads in order to be able to remove the metal frame during wafer dicing. This reduces usable die size, but saves time and money in packaging.

Confession 16: Choosing the wrong capacitor can make an amplifier useless

(Vincent Frick, ICube-University of Strasbourg, Strasbourg, France)

A few years ago, one of our projects addressed an integrated front-end dedicated to the signal conditioning of an OFET-based gas sensor. The OFET device requires specific source to gate voltage V_{GS} control in the range of $\pm 15V$ to operate properly [27]. Therefore, we designed a 5-bit low-voltage digital to analogue converter (DAC) combined to a high-voltage amplifier. The latter would convert the DAC's output into the required voltage range. As shown in Fig. 12(a), we used a standard Miller OTA topology featuring the typical compensation capacitor. This was our first design in AMS 180 nm high-voltage CMOS technology. This technology features various types of capacitors including metal-insulator-metal (MIM). After a quick skim through the process parameters documents we chose the MIM for its compactness and ability to withstand voltages up to 50V. Yet, 50V is only for the bulk-to-terminal voltage whereas the maximum terminal-to-terminal voltage is only 5.5V. Unfortunately, since the deadline for fabrication was drawing close, we overlooked the DC voltage across the terminals of this capacitor, focusing exclusively on the AC characteristics of the amplifier, obsessed by optimal stability as we were. As a result, we ended up with an amplifier with very good AC characteristics but that was completely useless for our application because the compensation capacitor would breakdown as soon as the output voltage would be higher than 7V. The solution was pretty straight forward: as shown in Fig. 12, we replaced the MIM capacitor (Fig. 12(b)) by a high-voltage vertical capacitor (Fig. 12(c)), which is made of interdigitated vias and metal wires, and withstands terminal-to-terminal voltages up to 50V. The ironic point is that there was plenty of free silicon space on the initial chip. There was no point in trying to save room by using a compact capacitor in the first place... Anyway, the circuit works great now [27]!

Moral: Sit down, have a coffee, a tea, a whisky or whatever, and take a moment to carefully read the process parameters, no matter how short the deadline is. A stitch in time saves nine! Check whether space is really an issue of your design. Also, always double-check the transient and DC operating points, especially when you're happy with the AC performance. And last but not least, don't overlook the warnings in the simulation log file.

VII. PARASITICS

Many bugs are the result of parasitics from capacitive coupling, resistance, or unintended photocurrents. Unless the designer is aware of the possibility of these mistakes, modeling will not reveal them.

Confession 17: Parasitic Couplings of Test Lines ruins the Correct Operation of Working Circuit under Test

(Teresa Serrano-Gotarredona, IMSE-CNM, Sevilla, Spain)

Back in the early 2000s we were starting to design, test, and use integrate-and-fire analog neurons to build large scale neuromorphic arrays of different retinas and convolution chips. We designed a compact double sign integrate-and-fire

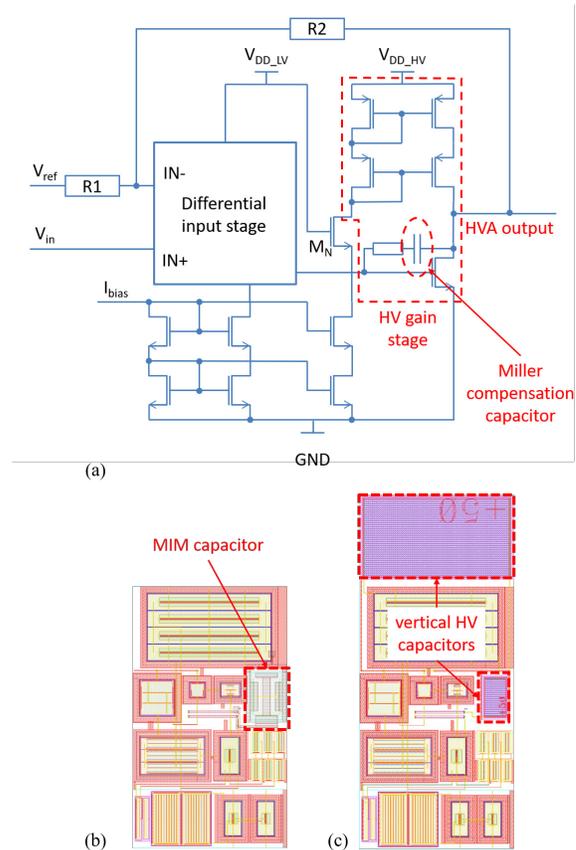


Fig. 12. (a) High-voltage Miller OTA. (b) Initial layout with MIM capacitor that breaks down at only 7V. (c) Corrected layout with vertical high-voltage capacitor.

neuron capable of operating with down to sub-pico-ampere currents [28], [29]. Therefore, we were very careful in drawing its layout, avoiding couplings, and simulating very carefully with all possible parasitics provided by the layout extractor. Everything was working great in simulation, but we wanted to fabricate and test carefully the circuit before building arrays of it. So we included it quickly in a little corner of a larger chip that was going out for fab shortly. We added simple test pads for all bias, all inputs and the two outputs (one for each sign). Outputs were buffered properly with analog buffers. When it came back, it would only produce outputs of one sign only. This was very strange, because the signed circuit was quite symmetric and operated very similarly for both signs. The simulations did not reveal any discrepancies between the positive and negative sign outputs. At some point, we decided to re-extract and re-simulate the full layout (including the test pads) with all possible parasitics, and observed that only outputs of one sign were being produced, as in the experimental chip. At this point we quickly spotted that it was due to a parasitic capacitive coupling between two metal lines routing two nodes from the circuit to the pads: one of the outputs to one of the inputs. By separating both lines a bit more, the coupling was gone and the circuit worked perfectly.

Moral: Check your full layout (up to the pads), including parasitic extraction and simulation verification.

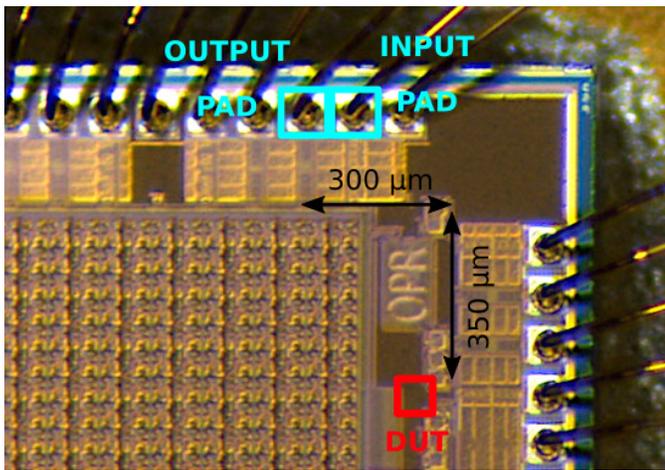


Fig. 13. A long thin metal line between the pads and device under test (DUT) caused problems with the LDO regulator.

Confession 18: Why is not a good idea to use long thin metal lines to power low dropout regulators

(Paula Lopez, Centro Singular de Investigacion en Tecnoloxias Intelixentes (Universidade de Santiago de Compostela), Spain)

Given the scarce funding typically received by universities, it makes a lot of sense to use the so called Multi Project Wafer services provided by several foundries. Even more, the allocated area has to be often shared between several projects from the same institution. Consequently, the pad ring has to be shared as well. Under these constraints it is often the case that long metal lines have to be used to connect the circuits to the available IO bondpads, as in Fig.13. This is OK as long as you have enough room to make these metal connections wide enough to not introduce a significant series resistance, especially when currents of several milliamps are expected to flow through them. In our case, we designed a low dropout (LDO) regulator and when performing the experimental characterization we realized that despite the good agreement between the simulated and experimental behaviors for the other figures of merit, the measured load regulation was systematically higher than the expected value. Using the sheet resistance values provided in the foundry documentation we estimated that the metal tracks that connect the input and output pins of the circuit with the corresponding IO pads had resistances of about $70\ \Omega$ and $40\ \Omega$ respectively. This was enough to cause an increase of the load regulation from $14\ \mu\text{V}/\text{mA}$ to $40\text{mV}/\text{mA}$.

Moral: Make sure that your long metal lines are wide enough.

Confession 19: Power grids are important for a reason!

(Melika Payvand, ETH/UZH Switzerland, Advait Madhavan, NIST, Maryland, USA):

During our PhD years, we were working on a project whose purpose was to demonstrate a monolithically integrated hybrid neuromorphic chip with CMOS neurons and memristive synapses. A specific feature we wanted to implement was an STDP based online learning scheme, inspired by the pulse shaping ideas presented in [30]. Because the device state was supposed to change depending on the relative spike timing,

the neurons had to be able to drive memristors into the write state, and provide the necessary current for the switching.

The constraints for the output drive of the neuron was determined by the size of the array and the lowest expected value of R_{on} . Note that the memristors were being designed in house and had a very conductive ON state. We were working with selector-less passive arrays and as a result, we had to take into consideration the scenario where the output stage would drive a whole column of the memristive array. This resulted in a design with a very high current output driver.

During the last few weeks of the project, there were personnel and deadline constraints that caused us to rush through certain aspects of the design, such as the power grid. We didn't thoroughly consider the total resistance of the traces dedicated for power distribution. As a result, the parasitic resistance of the power lines were not negligible. Hence a spiking event in a neuron would cause a high current to flow through 10s of ohms of parasitic power supply wire resistance and hence, the VDD would transiently drop by almost 500mVs. Such large transient fluctuations in power supply, rendered the circuits un-testable.

Moral: Pay attention to the power grid while doing the layout. Especially, if you have large amounts of current running through the circuits.

Confession 20: Look at the bigger picture when modeling parasitic capacitance

(Bathiya Senevirathna and Pamela Abshire, University of Maryland, College Park):

A few years ago we fabricated a capacitance sensor chip for detecting biological cell attachment on the sensor surface [31], and from past experience knew that limiting the number of I/O pads on the chip would make life much easier when it came to encapsulating the CMOS chip in a biocompatible package. The fewer the electrical connections to the outside world, the fewer possible points of failure. So we designed the chip to be as "self-contained" as possible, with the capacitance signal transduction and conversion, a system-level logic controller, and an I²C serial bus on chip. This approach reduced the number of I/O pads from ~ 40 down to 8. We ran simulations with post-layout extracted parasitics on the serial communication bus lines and ended up tripling up the data line drive capability based on what we expected the maximum bus capacitance to be. We got the chips back from the foundry and everything worked as expected on the bench. When it came time for *in vitro* experiments however, the chip started crashing sporadically when placed inside the incubator. And unfortunately, by trying to limit the number of I/O pads as much as possible there was limited ability to debug; everything relied on the I²C bus working. After some investigation, it turned out that the extra length of cabling needed to connect the chip to the readout equipment added just enough capacitive load on the serial communication lines to take the timing out of spec (see Fig. 14). We mitigated the issue by using shorter lengths of higher-quality shielded cables, but not before losing several days' worth of experimental time and countless cells.

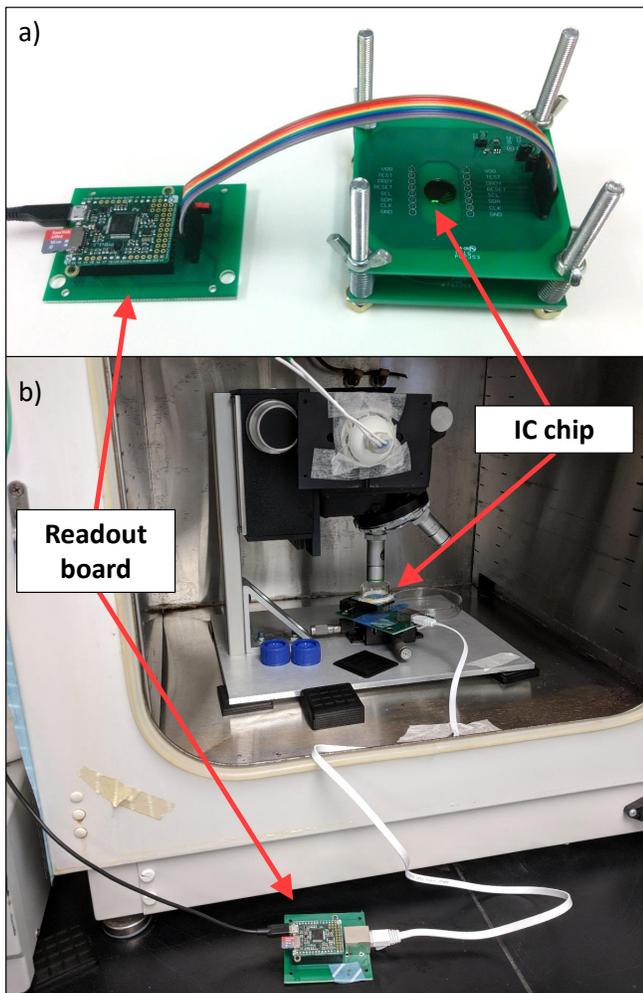


Fig. 14. Photos of the experimental setup to run the chip a) on the bench top and b) inside the incubator. The biosensor chip worked as expected on the bench top, but failed during *in vitro* experiments due to the extra length of wire needed to connect the chip, now inside the incubator, to the readout board.

Moral: Consider your entire testing setup, not just the chip itself, when estimating parasitic capacitance and signal driving requirements, especially if you need to do any non-standard packaging to perform experiments.

Confession 21: Light sneaks through the tiniest cracks (or, The Sun is Bright)

(Tobi Delbruck, University of Zurich and ETH Zurich)

Our prototype DAVIS event cameras designed and paid for by the EU SEEBETTER and VISUALISE projects were working fantastic and were selling like hotcakes from iniVation. All appeared great until we took the DAVIS out for tests on a military airfield in Florida as part of the FLA project. Suddenly there were mysterious bursts of events from the entire sensor array (Fig. 15(a-b)). What was the cause of these bursts? They seem to occur mainly when the camera was aimed at parts of the sky with bright white clouds.

Finally back in Zurich I traced down the source. The bursts were caused by bright light shining onto one edge of the chip. If I aimed the DAVIS at the sky and focused the sun on a spot just outside the pixel array, the whole sensor went wild.

In the design of our DAVIS event cameras we used our on-chip bias generator. Of course I knew all about parasitic photocurrents [1], and so we covered up the bias generator with overlapping upper metal, and were careful to surround it with an N WELL guard ring that would soak up the stray minority electrons generated outside the light shield. But the sun was so bright that either it penetrated the top metal, or its oblique path at the edge of the array allowed it to sneak through the dielectric between the top metal and the one underneath. The resulting photocurrents were large enough to disturb the small bias currents of less than 1 nA that are used for biasing the DVS pixels. I believe the effect was to disturb the ON and OFF thresholds enough to cause these bursts. Fig. 15(c-e) shows the setup we used with a microscope equipped with a bright light source that we could focus down to a 10 μm spot. We located the hot spot and the it corresponded to exactly these ON/OFF biases, so we are pretty sure of the cause.

Fig. 15(f) shows my quick solution to cover the edge of the array containing the bias generator with an umbrella formed from with a tiny piece of aluminum foil glued on top of the optical cover glass. (Black electrical tape was not sufficiently opaque). A better but much more expensive solution was for iniVation to arrange with our packaging partner to put down a black glob covering of the edge of the array. It was difficult to form this glob precisely. A production solution would be to use an special optical shielding metal provided by CMOS image sensor foundries.

Moral: The sun is bright! It has a luminance of more than a million Cd/m². Compare that with a screen luminance of a few hundred Cd/m². If you ever use your sensor outdoors, think about what focused sunlight shining onto the silicon might do.

VIII. PLANNING MISTAKES

Errors of this type occur in planning of a project.

Confession 22: Wafers tossed with trash before testing

(Eric Fossum, Dartmouth College)

The best expensive mistake story I know is when one of our Photobit engineers tossed out 12 new wafers from a prototype run in their shipping box, by mistake, by resting the box on top of his trash can. The cleaners threw it out with the rest of the trash.

Moral: A box of wafers costs \$100k. Would you put a \$100k painting on top of a trash can? Maybe it is good to have a planned policy to store wafers.

Confession 23: Understand the hardware before integrating it

(J. Camilo Vasquez Tieck, FZI Forschungszentrum Informatik Karlsruhe)

We work on different visuo-motor experiments as part of the neurorobotics sub-project of the HBP. A while ago, we were developing a new experiment using a simulated event-based camera [32] to move a simulated robotic arm [33] using spiking neurons. We started with simple building blocks and

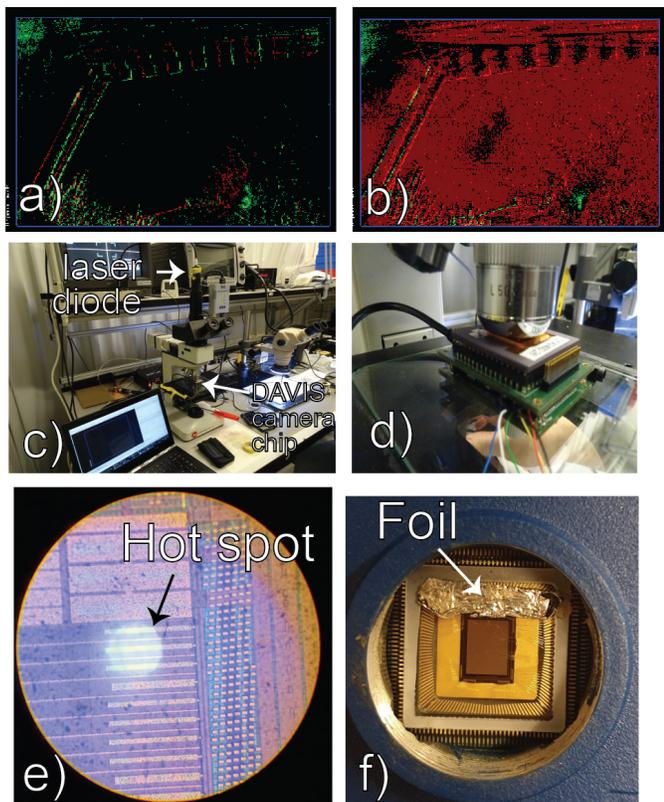


Fig. 15. Sun shining onto the bias generator makes DAVIS event camera go wild. (a) A normal output of DAVIS camera looking out the window. Green and red are ON and OFF events. (b) Focusing the sun just outside the pixel area to left side creates storms of OFF events. (c-e) Debugging location of sun hotspot using (c) microscope equipped with laser diode, by (d) focusing to find the (e) hotspot inside the bias generator on two particular biases for the DVS pixels. (f) Taping foil just over the bias generator helps fix the problem.

after some promising initial results we wanted to scale up the network. For this purpose SpiNNaker was perfect, and even more, if all worked out well, we could also use it with the real robot. Not that easy! By design, the Neurorobotics Platform (NRP) couples and synchronizes a neuro-simulator (NEST) with a physics simulation (Gazebo) by ticking them for 20 ms (typically) and waiting for both to complete. Our error was to think we could "just" replace the neuro-simulator with SpiNNaker and keep everything the same. Indeed, SpiNNaker is designed for real time control, but it was not meant to be used for stepped simulations with multiple runs. Still, we invested a lot of time and effort doing workarounds to connect the hardware to the current setup. We hacked both the NRP and the SpiNNaker software stack, and finally figured out a way to integrate them. The problem was that due to the stepped simulation and the communication overhead our solution had a bad performance. For academic purposes, we compared in Fig 16 running the same network in both scenarios. The bad performance is not caused by the hardware, it is a result of our design choices. We learned our lesson, and now we are working on a different solution for the synchronization mechanism using a continuous execution of the network in SpiNNaker, while streaming spikes in and out using UDP. That way, any overhead related to starting and stopping the

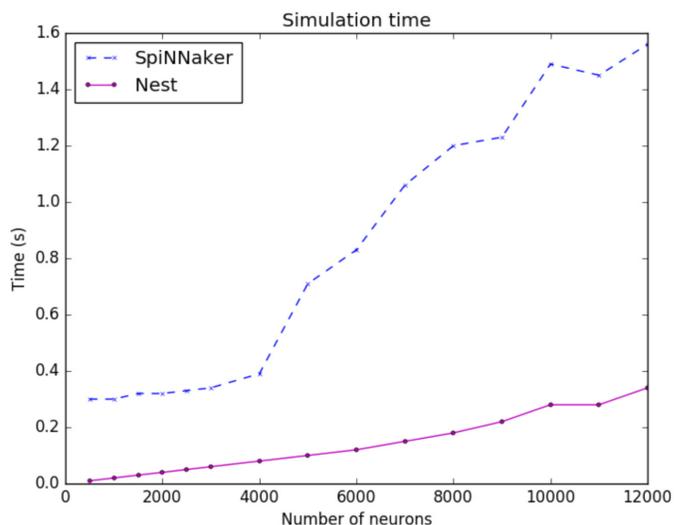


Fig. 16. Performance of the network using stepped execution setup with both SpiNNaker and NEST. Simulation time plotted against number of neurons. Times are averaged over 20 runs of 100ms each, the times do not include setup time.

simulation is no longer relevant, because it is always running.

Moral: Understand the capabilities and limitations of new components in your system. Analyse what is it good for, and consider redesigning the system as early as possible to take advantage of it.

Confession 24: Sure we've got space for some test circuits! (Ian Williams, Imperial College London):

A few years ago we set out to make our first full wafer run – pulling together designs from multiple projects in house and taking the step of designing the reticle set. The SenseBack neural interface chip [34] was one of these designs and due to some large dies the dicing strategy allocated SenseBack a $5 \times 4 \text{ mm}^2$ space to fit in. Partway through the chip design it became apparent that we could fit the design into a $4 \times 4 \text{ mm}$ square leaving us 4 mm^2 of area to play with. Naturally this was pounced on by our fellow designers to include some weird and wacky test circuits. Fast-forward a year and we were trying to get these devices packaged in a miniature, moulded package. Precision moulding rather than the cheaper glob-top alternative was needed because we were looking to embed the components inside the PCB (Embedded Component Technology). Unable to find a company willing to do small quantities of moulded packaging in the UK we sent these dies off to the US to be packaged in a 0.4 mm pitch QFN64. Having sent the company 50+ dies, we were disappointed months later to receive back 15 or so glob-top packaged devices and the rest of our dies moulded, rattling around in an envelope. The problem (shown in Fig. 17) was bondwires sticking out of the moulding! Although we had told the company about the test circuits and shown them on the bonding diagram, it seems no one had worked out how much the bondwires would need to be raised in order to clear them and it had taken more than 30 devices to find that out. Some of the partially moulded devices in the envelope are expected to work, but there are enough

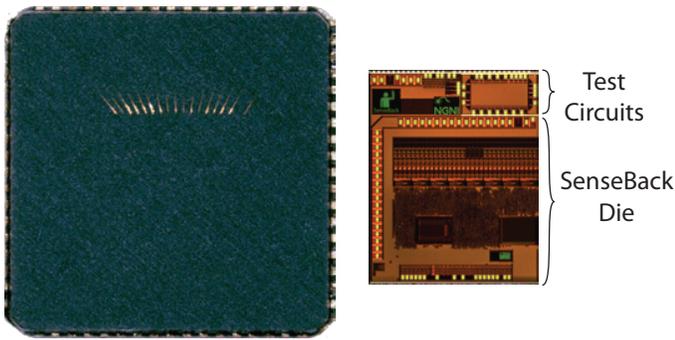


Fig. 17. Left - bondwires protruding through packaging of SenseBack chip. Right - the offending test circuits that raised the bond wires.

visible bondwire breakages to give us low expectations of the yield.

Moral: That prototyping jobs can trip up even expert companies; and that sometimes if your design easily fits into your allocated space, you should just not mention it to your colleagues, and instead make your padframe bigger and fill the space with extra supply bypass.

Confession 25: Finding a micro-controller to test chips is easy! right?

(Yan Liu, Timothy Constandinou, Imperial College London):

It is exciting to integrate the fabricated silicon into a miniaturized device and perform in-vivo experiments. A few years ago we designed and implemented a 64-channel neural recording chip with a standard SPI link and simple handshake [35]. During the design process, we decided to use some low-end small micro-controllers (e.g. NXP K64F) rather than FPGAs for system integration and in-vivo testing. We were very confident with our target testing system and our choice on the selection of MCU, since the overall communication speed is only 50 MHz, which is acceptable for the MCU according to the data sheet and maximum raw data throughput is only 16 Mbps, not to mention that the K64F also comes with a handy and cheap development platform and an SD card socket. Therefore, all the design constraints were defined accordingly and the test platform was finalized after sign-off.

After chips came back, we powered them up successfully and showed correct bit flipping patterns on the communication link. However when we were trying to run the chip on full power, i.e. with 64 channels fully on, the first thing we realized is that the SPI link can not reach 50 MHz with the test platform we designed. The cable and connectors added extra parasitics capacitance to the SPI wires such that the data can not be latched properly. Moreover the MCU could not accommodate the fast handshake together with SPI under maximum data throughput even without SD card interface. Even when the chip was under-clocked to 10s of MHz, 70% of the channels have to be shut down to avoid data collision. Finally, we used an FPGA board to mitigate all these problems with jumper wires, but the noise caused by these connections and the power supply lines corrupted the sensitive neural signal and apparently the bulky FPGA board with loose wiring is not that animal-friendly at all. It took extra time to re-design and

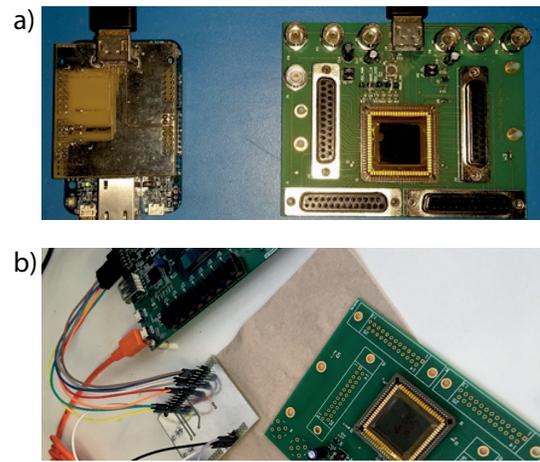


Fig. 18. a: the test platform with MCU and tidy connections. b. Plan b test platform with a FPGA board and dangling wires.

characterize the entire chip under full power. The only lucky thing is the implanted neural electrode only had 4 sites fully functionally during the initial in-vivo experiments, which our under-clocked MCU platform can easily handle.

Moral: Plan your testing strategy well before the design and carefully, even just for a slow and simple biomedical prototype device!

Confession 26: All eggs in one basket.

(Piotr Dudek, Stephen Carey, The University of Manchester)

On a recent project, we were trying to make use of the vision sensor chips that we made previously [19] on an MPW run. The chips were working very nicely (with the workaround described in our other confession), and seemed pretty useful. Initially we had ten packaged dies: a couple were faulty, we broke a couple during the tests, kept a couple in our older test PCBs, and used the rest to build a few functional camera systems, with an integrated controller, USB interface, etc. These systems were then used to explore potential applications in our lab, and by a few collaborators on various projects [36], [37], [38]. These collaborations were very successful, and we started to get more people interested in using the system, so we decided to package all the remaining twenty bare dies from the MPW run, that we still had in the cupboard, to build more systems for people to use. We sent the dies and the order to a carefully selected packaging house, specifying a glass-top package for this image sensor chip. Two months later we got an email saying "We are ready to ship, BUT we have used ceramic lids instead of glass lids. Sorry for the error". "Sorry" as an apology didn't quite deal with the issue in hand. The lids were stuck solid - they did a great job there! The contract, as is often the case, was limited to the value of the order, i.e. the packaging job, not the value of the silicon (even that would not quite cover what this silicon was worth to us at that point in the project!) After a few, not exactly very polite emails, the packaging house tried to fix what they broke. They damaged quite a few samples trying to figure out how to open the packages up, but finally came up with some laser-cutting scheme that managed to expose the dies again, alas with a

layer of sooty deposit right on the die! They put glass over it, and this is the best we got from it. We use some of these, as a “Grade B product”. They work okay for development, but the sensor is dirty.

Moral: If you have just a few precious bare dies left, don’t put them all into one envelope. No matter how reputable your subcontractors are.

Confession 27: Symmetry, incomplete seal rings and other tales from the crypt

(Alexander Serb, Univ. of Southampton and previously Imperial College UK)

For my PhD at Imperial I had to design some photodiode chips for testing how different layouts would shape the depletion layer of the junction (thus affecting photocurrent generation under a constant light source). To that end I designed a couple of chips: ‘Mr. Bean’ would have a bunch of large-area photodiodes and a few test circuits (charge pumps, ring oscillators and the like in order to run a power management system). On the other hand ‘Teddy’, was the much smaller companion to Mr. Bean and was exclusively made of photodiodes arranged in a beautiful perfectly symmetric pattern, even though each photodiode boasted its own special design. Now, people don’t often think of microchips as being designed perfectly symmetrically, but when the chip returns from the foundry and the person in charge of wirebonding asks the question: ‘Right, so which side is up?’ the less obvious disadvantage of perfect symmetry rears its sarcastic head. This was back in 2012-2013 and the solution we came up with is beginning to fade from my memory, but I believe it went along the lines of trying to figure out how different types of photodiode exhibit different colours under an optical microscope. Suffice it to say I’ve since made sure that some indication of what side is ‘up’ has been included on my chips.

Naturally, being some of my very first chips runs, Mr Bean and Teddy provided a wealth of other learning outcomes, most notably the error that cost me weeks of puzzlement and one chip burnt to a crisp: Somehow, and I to this day don’t know how, the seal ring ended up being conductive, i.e. having a conductive surface exposed and running around the entire chip. Naturally, this meant that when wirebonding the pads they would all short together. At the time, my idea of an experiment to see what was going on, was to run 0.5 A through it and see what would happen. I ran the experiment in a very dark anechoic chamber and predictably I would see the power supply bond wire glow red before all went dark again. I learnt nothing from that experiment but I had jolly good fun. I could go on, but I’m afraid this margin is too small to contain all the goofs.

Moral 1: Always include some mark in your chip that acts as a navigation aid. The mistake above was elementary, but large chips with many pads, chips made out of a spaghetti of synthesised logic etc. all may level challenges of checking which pad is which when you least expect them.

Moral 2: Never forget to use Occam’s razor. When all the

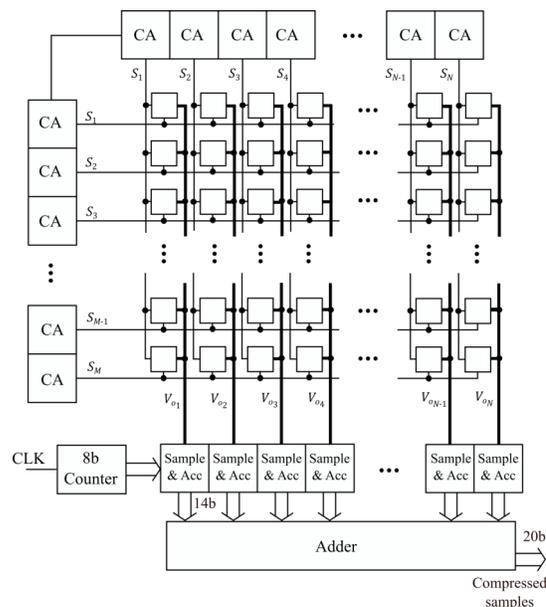


Fig. 19. Conceptual floorplan of the intensity-to-time vision sensor chip. Lack of row/column selection and column output makes testing the device much too hard.

likely scenarios have been exhausted whatever is left must be the answer, no matter how unlikely. Trust nothing and question everything. Who would have thought that a seal ring would ever expose a conductive path at its surface?

Confession 28: Test points not pains

(Ricardo Carmona-Galan, Institute of Microelectronics of Seville)

One of our last image sensor chips is an array of intensity-to-time pixels for concurrent focal-plane generation of compressed samples [39]. In order to do this, we have employed a 1D cellular automaton (CA) for the pseudo-random selection of pixels. This solution avoids both transmitting and storing the compressive strategy as it can be error-free reconstructed from the initial seed with an analogous CA implementation. In addition to this, pixel values need to be represented in a form that does not exhaust the available dynamic range when being added up. In order to do this, we are encoding pixel values in time. Addition is realized in the digital domain by asynchronously tagging events in a per-column basis. The architecture of the chip is depicted in Fig. 19.

The central element of the architecture is an array of 64×64 pixels. The peripheral circuit needs to implement the following functionalities: pseudo-random column and row selection, time-to-digital conversion of the pixel values, addition of the pixel values of the selected pixels. Events are generated asynchronously as function of light intensity. Then, they are transmitted through the column bus and arrive to block ‘Sample & Acc’ in Fig. 19. These pulses encode the pixel value in the period of time that has passed between the pixels reset and the arrival to the ‘Sample & Acc’. A straightforward method to translate all this pulses into digital codes is to use the pulses to activate the sampling of a global time counter (Fig. 19) activated by a clock signal and started with the global pixel reset —allocating some initial delay to allow the pulses

to reach the bottom of the array. Each time a pixel activation pulse arrives, the 8b of the counter are sampled and added to the already stored sum. After 256 clock periods, the pixel values have been accumulated at the ‘Sample & Acc’, which delivers a 14b word containing this sum, which is the result of adding up to 64 pixel values. After that, the 64 column sums are added up into a compressed sample of 20b.

This scheme employs rather simple blocks, but all of them need to be tested before the pixel array can be characterized and its compressive sampling behavior verified. We have reproduced an isolated pixel and the different blocks outside of the main array for test and characterization. However, the lack of agile column and row multiplexers for selection, and per-column test points for conveying the timing pulses generated at the pixels, convert the task of characterizing pixel properties in the complete array in a real pain. For instance, in order to characterize pixel non-uniformity, the procedure goes like this: first, the CA seed must be loaded and frozen to select just one particular pixel; then all the sampling, accumulate and addition is required to obtain the output of one single pixel. This process needs to be repeated 64×64 pixels times for each illumination level. Painful, isn't it?

Moral: When following bottom-up approaches, for instance in the design of smart pixel arrays, it is easy to arrive to a point in which there is no place for test points and the system has become so intricate that simple characterization becomes really difficult. Plan your testing strategy in advance.

Confession 29: Hurry up slowly, and put your phone away
(Robert Nawrocki, Purdue University)

There is a known Polish proverb that roughly translates as “hurry up slowly”. It was designed to convey the idea that you are more likely to make a costly mistake when trying to accomplish too much in a short time.

The process of manually micro fabricating organic electronics can be quite lengthy, tedious and labor intensive. Fabrication of a single device can take up to a week from start to finish. Fabrication typically starts with cleaning a substrate, followed by sequential deposition of individual layers of material using processes such as spin and spray coating, physical and/or chemical vapor deposition, etching (wet or dry), surface activation, and of course (photo)lithography. It is a sequential process that requires complete attention at every step, especially towards the end when people are most prone to fatigue-induced errors; these mistakes are the most costly because they nullify all the previous work.

Things get more complicated when we fabricate sub-micron thin organic transistors. To accomplish this we modified [40] the standard process to include a Teflon sacrificial delamination layer, so that the film can be delaminated from a glass or a silicon wafer substrate and then laminated on a target substrate, such as human skin (Fig. 20). After the layer of Teflon is deposited on a substrate, subsequent layers are deposited on top of it. After the entire fabrication process is complete, the Teflon layer allows for simple and effortless manual delamination and transfer of the film. However, this easy delamination also makes the process more prone to

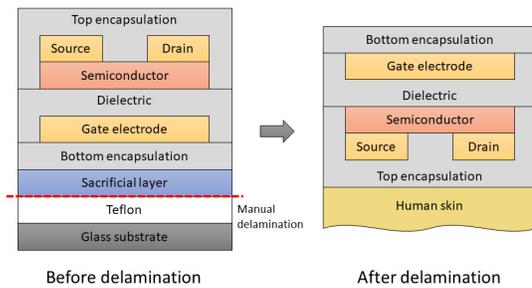


Fig. 20. The Teflon layer used in this organic semiconductor fabrication flow makes it very slippery!

failures because the Teflon layer makes things slippery.

Some time ago, towards the end of the fabrication of an ultra-thin organic transistor, I needed to clean the film surface before aligning the mask for the source/drain electrodes before thermal evaporation of gold film, by blowing a stream of N_2 gas from a gun. After a third consecutive 10-hour long day, hungry and tired and wanting to eat, holding the wafer, a tweezer, a mask, and the N_2 gun, I got momentarily distracted by a beep on my phone from an arriving message about lunch meeting, and I lost concentration for about a second. I squeezed the N_2 gun too hard just when its nozzle was right next to the edge of the chip. What followed was a movie-like, slow motion moment, with me watching the film delaminate from the chip, and landing crumpled on the floor of the clean room. In a momentary loss of attention I lost about four days' worth of hard work, all because I wanted to hurry up and meet friends for lunch.

Moral: Try not to lose focus, especially towards the end of a lengthy fabrication procedure. In fabrication there is no “backspace” or CTRL+Z. Turn your phone off and put away all other distractions.

IX. CONCLUSIONS

Based on the confessions here and in [1], we can see that many of the goofs were the result of not sufficiently looking outside the immediate low-level design aspects. It seems clear that students need to be better trained about the potential for parasitics that are not modeled, and need to clearly understand the process technology with its restrictions. Judging by the eight planning mistakes, it also seems clear that more time needs to be spent on planning a design including the target application.

By summarizing the confessions, we came up with following checklist that might be helpful for future projects:

- Does your test bench really cover reality? I.e. does it simulate random and noisy input?
- Have you talked to your supplier sales or technical representative to ask about undocumented problems?
- Are you powering everything from one wall plug? If not, are you sure there is a good common ground?
- Does LVS check W/L and bulk connections?
- Parasitic resistance in power lines OK?
- Parasitic capacitance between lines OK?

- Off-chip parasitics considered and modeled in simulation?
- Shining light onto the chip? Are circuits covered? Will the sun be focused? If so, have you covered the circuits with at least 3 layers of metal with huge overlaps?
- Does substrate leakage in transistor junctions matter? Especially if light is shining onto the chip?
- Has the entire chip including pads been LVS'ed and DRC'ed?
- Did you mess with the pads?
- Are all the warning messages understood?
- Have you double and triple checked the polarity of your control signals?
- Have the designers read the process specifications, particularly for special devices like capacitors and resistors?
- Will ESD protection in the pads affect the operation?
- Has packaging been considered?
- Does the IC have an ID mark that also shows orientation (i.e. does the mark indicate UL corner)?
- Are the power grids OK? Did you check for short circuit between the power supply rails?
- Will the die seal rings affect the project?
- Do you have a plan to inventory and store your precious dies and wafers?
- On your PCB, have you assumed that bottom pads are ground?
- If your chip needs a read-out controller, would a micro-controller provide enough read-out bandwidth or would you be safer going with an FPGA?

We hope that reading these confessions can have a significant preventative effect on future projects from ourselves and the wider circuits and systems community.

REFERENCES

- [1] P. Abshire, A. Bermak, R. Berner, G. Cauwenberghs, S. Chen, J. B. Christen, T. Constandinou, E. Culurciello, M. Dandin, T. Datta *et al.*, "Confession session: Learning from others mistakes," in *2011 IEEE International Symposium of Circuits and Systems (ISCAS)*. IEEE, 2011, pp. 1149–1162.
- [2] A. Aymar, H. Mostafa, E. Calabrese, A. Rios-Navarro, R. Tapiador-Morales, I.-A. Lungu, M. B. Milde, F. Corradi, A. Linares-Barranco, S.-C. Liu, and T. Delbruck, "NullHop: A flexible convolutional neural network accelerator based on sparse representations of feature maps," *IEEE Trans Neural Netw Learn Syst*, vol. 30, no. 3, pp. 644–656, Mar. 2019. [Online]. Available: <http://dx.doi.org/10.1109/TNNLS.2018.2852335>
- [3] S. Muzaffar and I. A. M. Elfadel, "Toward an integrated, low-power platform for continuous congestive heart failure monitoring," in *The IoT Physical Layer*, I. A. M. Elfadel and M. Ismail, Eds. Springer, 2018, ch. 19, pp. 327 – 352.
- [4] T. Delbruck and C. Mead, "Analog VLSI phototransduction by continuous-time, adaptive, logarithmic photoreceptor circuits," *CNS Memo No. 30, Calif. Inst. Technol.*, 1995.
- [5] R. F. Yazicioglu, P. Merken, R. Puers, and C. Van Hoof, "A 200 μ W eight-channel EEG acquisition ASIC for ambulatory EEG systems," *IEEE Journal of Solid-State Circuits*, vol. 43, no. 12, pp. 3025–3038, Dec 2008.
- [6] R. J. Baker, H. W. Li, and D. E. Boyce, *CMOS Circuit Design, Layout, and Simulation*. IEEE Press, 1998.
- [7] C. Z.-R. T. S.-G. L. Camuñas-Mesa, A. Acosta-Jiménez and B. Linares-Barranco, "A 32x32 pixel convolution processor chip for address event vision sensors with 155ns event latency and 20meps throughput," *IEEE Trans. Circ. and Syst. Part-I*, vol. 58, no. 4, pp. 777–790, April 2011.
- [8] L. Camuñas-Mesa, C. Zamarreño-Ramos, A. Linares-Barranco, A. Acosta-Jiménez, T. Serrano-Gotarredona, and B. Linares-Barranco, "An event-driven multi-kernel convolution processor module for event-driven vision sensors," *IEEE J. of Solid-State Circuits*, vol. 47, no. 2, pp. 504–517, February 2012.
- [9] S. Afshar, L. George, J. Tapson, A. van Schaik, and T. J. Hamilton, "Racing to learn: statistical inference and learning in a single spiking neuron with adaptive kernels," *Frontiers in Neuroscience*, vol. 8, no. November, pp. 1–18, nov 2014.
- [10] S. Afshar, L. George, C. S. Thakur, J. Tapson, A. van Schaik, P. De Chazal, and T. J. Hamilton, "Turn down that noise: synaptic encoding of afferent snr in a single spiking neuron," *IEEE transactions on biomedical circuits and systems*, vol. 9, no. 2, pp. 188–196, 2015.
- [11] R. J. Sofatzis, S. Afshar, and T. J. Hamilton, "The synaptic kernel adaptation network," in *2014 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE, 2014, pp. 2077–2080.
- [12] —, "Rotationally invariant vision recognition with neuromorphic transformation and learning networks," in *2014 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE, 2014, pp. 469–472.
- [13] E. Dumesnil, P.-O. Beaulieu, and M. Boukadoum, "FPGA minimal components SKAN model for classical and operant conditioning," in *2016 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE, 2016, pp. 2499–2502.
- [14] —, "Robotic implementation of classical and operant conditioning as a single STDP learning process," in *2016 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2016, pp. 5241–5247.
- [15] R. Guyonneau, R. VanRullen, and S. J. Thorpe, "Neurons tune to the earliest spikes through STDP," *Neural Computation*, vol. 17, no. 4, pp. 859–879, 2005.
- [16] T. Masquelier, R. Guyonneau, and S. J. Thorpe, "Spike timing dependent plasticity finds the start of repeating patterns in continuous spike trains," *PLoS one*, vol. 3, no. 1, p. e1377, 2008.
- [17] C. Zamarreño-Ramos, T. Serrano-Gotarredona, and B. Linares-Barranco, "An instant-startup jitter-tolerant manchester-encoding serializer/deserializer scheme for event-driven bit-serial lvds inter-chip aer links," *IEEE Trans. Circ. and Syst. Part-I*, vol. 58, no. 11, pp. 2647–2660, November 2011.
- [18] A. van Schaik and S.-C. Liu, "AER EAR: A matched silicon cochlea pair with address event representation interface," in *IEEE International Symposium of Circuits and Systems (ISCAS)*, 2005, pp. 4213–4216.
- [19] S. J. Carey, A. Lopich, D. R. Barr, B. Wang, and P. Dudek, "A 100,000 fps vision sensor with embedded 535GOPS/W 256x 256 SIMD processor array," in *2013 Symposium on VLSI Circuits*. IEEE, 2013, pp. C182–C183.
- [20] J. Chen, S. J. Carey, and P. Dudek, "SCAMP5D vision system and development framework," in *Proceedings of the 12th International Conference on Distributed Smart Cameras*. ACM, 2018, p. 23.
- [21] W. D. Leon-Salas, T. Fischer, X. Fan, G. Moayeri, and S. Luo, "A 64x 64 image energy harvesting configurable image sensor," in *2016 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE, 2016, pp. 1914–1917.
- [22] J. M. de la Rosa *et al.*, "A CMOS 110-dB@40-kS/s Programmable-Gain Chopper-Stabilized Third-Order 2-1 Cascade Sigma-Delta Modulator for Low-Power High-Linearity Automotive Sensor ASICs," *IEEE J. of Solid-State Circuits*, vol. 40, no. 11, pp. 2246–2264, November 2005.
- [23] S. Lu, B. Senevirathna, M. Dandin, E. Smela, and P. Abshire, "System integration of IC chips for Lab-on-CMOS applications," in *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*, May 2018, pp. 1–5.
- [24] T. Datta-Chaudhuri, E. Smela, and P. A. Abshire, "System-on-chip considerations for heterogeneous integration of CMOS and fluidic bio-interfaces," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 10, no. 6, pp. 1129–1142, Dec 2016.
- [25] M. Dandin, I. D. Jung, M. Piyasena, J. Gallagher, N. Nelson, M. Urdaneta, C. Artis, P. Abshire, and E. Smela, "Post-CMOS packaging methods for integrated biosensors," in *SENSORS, 2009 IEEE*, Oct 2009, pp. 795–798.
- [26] L. Li, H. Yin, and A. J. Mason, "Epoxy chip-in-carrier integration and screen-printed metalization for multichannel microfluidic Lab-on-CMOS microsystems," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 12, no. 2, pp. 416–425, April 2018.
- [27] V. Frick, P. Lévêque, U. Soysal, and T. Heiser, "Integrated high-voltage CMOS mixed-signal instrumentation system for OFET-based gas sensor," in *2016 IEEE International Conference on Electronics, Circuits and Systems (ICECS)*, Dec 2016, pp. 309–312.
- [28] B. Linares-Barranco, T. Serrano-Gotarredona, R. Serrano-Gotarredona, and J. Costas-Santos, "A new charge-packet driven mismatch-calibrated

-
- integrate-and-fire neuron for processing positive and negative signals in aer based systems," in *IEEE International Symposium on Circuits and Systems (ISCAS)*, vol. 5, May 2004.
- [29] R. Serrano-Gotarredona, T. Serrano-Gotarredona, A. Acosta-Jimenez, and B. Linares-Barranco, "A neuromorphic cortical-layer microchip for spike-based event processing vision systems," *IEEE Trans. Circuits and Systems, Part-I: Regular Papers*, vol. 53, no. 12, pp. 2548–2566, December 2006.
- [30] T. Serrano-Gotarredona, T. Masquelier, T. Prodromakis, G. Indiveri, and B. Linares-Barranco, "STDP and STDP variations with memristors for spiking neuromorphic learning systems," *Frontiers in neuroscience*, vol. 7, p. 2, 2013.
- [31] B. P. Senevirathna, S. Lu, M. P. Dandin, J. Basile, E. Smela, and P. A. Abshire, "Real-time measurements of cell proliferation using a Lab-on-CMOS capacitance sensor array," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 12, no. 3, pp. 510–520, June 2018.
- [32] J. Kaiser, J. C. V. Tieck, C. Hubschneider, P. Wolf, M. Weber, M. Hoff, A. Friedrich, K. Wojtasik, A. Roennau, R. Kohlhaas *et al.*, "Towards a framework for end-to-end control of a simulated vehicle with spiking neural networks," in *2016 IEEE International Conference on Simulation, Modeling, and Programming for Autonomous Robots (SIMPAP)*. IEEE, 2016, pp. 127–134.
- [33] J. C. V. Tieck, L. Steffen, J. Kaiser, D. Reichard, A. Roennau, and R. Dillmann, "Combining motor primitives for perception driven target reaching with spiking neurons," *International Journal of Cognitive Informatics and Natural Intelligence (IJCINI)*, vol. 13, no. 1, pp. 1–12, 2019.
- [34] I. Williams, A. Rapeaux, Y. Liu, S. Luan, and T. G. Constandinou, "A 32-ch. bidirectional neural/emg interface with on-chip spike detection for sensorimotor feedback," in *2016 IEEE Biomedical Circuits and Systems Conference (BioCAS)*. IEEE, 2016, pp. 528–531.
- [35] Y. Liu, S. Luan, I. Williams, A. Rapeaux, and T. G. Constandinou, "A 64-channel versatile neural recording soc with activity-dependent data throughput," *IEEE transactions on biomedical circuits and systems*, vol. 11, no. 6, pp. 1344–1355, 2017.
- [36] L. Bose, J. Chen, S. J. Carey, P. Dudek, and W. Mayol-Cuevas, "A camera that CNNs: Towards embedded neural networks on pixel processor arrays," in *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
- [37] C. Greatwood, L. Bose, T. Richardson, W. Mayol-Cuevas, J. Chen, S. J. Carey, and P. Dudek, "Perspective correcting visual odometry for agile MAVs using a pixel processor array," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 987–994.
- [38] J. N. Martel, L. K. Müller, S. J. Carey, J. Müller, Y. Sandamirskaya, and P. Dudek, "Real-time depth from focus on a programmable focal plane processor," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 65, no. 3, pp. 925–934, 2017.
- [39] M. Trevisi, H. C. Bandala, J. Fernández-Berni, R. Carmona-Galán, and A. Rodríguez-Vázquez, "Concurrent focal-plane generation of compressed samples from time-encoded pixel values," in *2018 Design, Automation Test in Europe Conference Exhibition (DATE)*, March 2018, pp. 1316–1320.
- [40] R. A. Nawrocki, N. Matsuhisa, T. Yokota, and T. Someya, "300-nm imperceptible, ultraflexible, and biocompatible e-skin fit with tactile sensors and organic transistors," *Advanced Electronic Materials*, vol. 2, no. 4, p. 1500452, 2016.