Confession Session: Learning from Others Mistakes

Contributors: P. Abshire, A. Bermak, R. Berner, G. Cauwenberghs, S. Chen, J. B. Christen, T. Constandinou, E. Culurciello, M. Dandin, T. Datta, T. Delbruck, P. Dudek, A. Eftekhar, R. Etienne-Cummings, G. Indiveri, M.K. Law, B. Linares-Barranco, J. Tapson, W. Tang, Y. Zhai.

Organizers and Editors:
Bernabe Linares-Barranco (Sevilla Center for Microelectronics) and
Tobi Delbruck (University of Zurich and ETH Zurich).

With thanks to: JED Hurwitz and Sayed Danesh, Bassen Lande, Yulia Massoud & Sergio Netta

JED Hurwitz' call for "misteakes" at the 2009 ISSCC that inspired our confession session

Forewarned is Four-Armed:
Classic Analog Misteakes to Avoid

Welcome to the Analog Misteakes to Avoid Page. At this year's IEEE ISSCC conference in San Francisco, there will be a special evening session titled Forewarned is Four-Armed: Classic Analog Misteakes to Avoid. It is a rare occasion for people to talk about the things that go wrong within analog chip design at a conference and we hope it will be a memorable evening. To help it we have created this webpage with the aim of gathering statistics and stories of classic mistakes made by the greater analog design community.

Submit Story

Do you have a great story on how things went wrong for you or maybe a friend of yours when working on a chip? Maybe a simple analog mistake, tool problems, misunderstanding between designers or interface errors. We would love to hear the story. Please use the form below to submit a story with some statistics.

Unfortunately the mistakes were never published and are lost except in memories of attendees

Confession session was advertised on front page of ISCAS

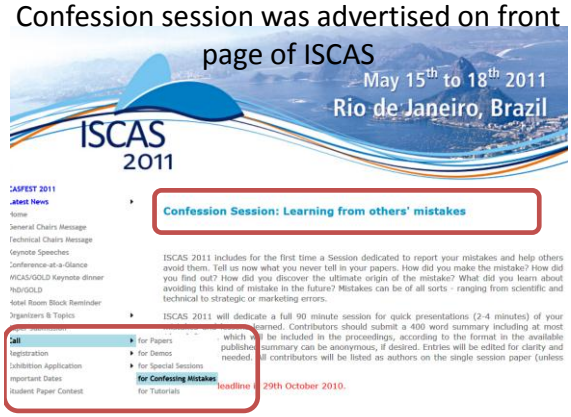

We collected 26 confessions.
(All were from our friends in the Sensory Systems and Neural Systems and Applications TCs)

- 3 mistakes in planning
- **6 parasitics not considered or improperly modeled**
- 2 process design errors
- 3 (not) looking at the bigger picture
- 2 strategic errors
- 4 logic goofs
- 3 LVS confessions
- 3 analog goofs

MISTAKES IN PLANNING

Timir Datta, Marc Dandin, and Pamela Abshire, University of Maryland, College Park,

1. TINY DIES ARE HARD TO HANDLE! (POST-CMOS PROCESSING FOR MEMS INTEGRATION)

Tiny dies are hard to handle!

Must leave "buffer" zone for handling
Costly, and scales with die area

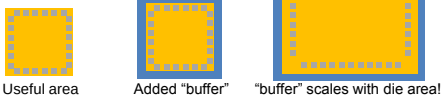

Edge bead becomes a significant problem with tiny dies

Marc Dandin and Pamela Abshire, University of Maryland, College Park

2. COORDINATING DIFFERENT INSTRUMENTS CAN BE NASTY WORK

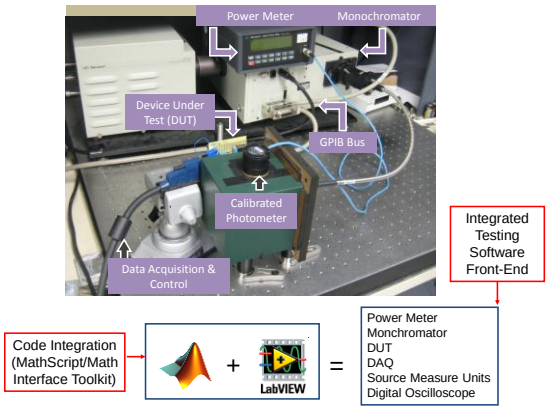

Yiming Zhai and Pamela Abshire, University of Maryland, College Park

3. DIP40 PACKAGES ARE GOOD THERMAL MASSES

DIP40 packages are good thermal masses!

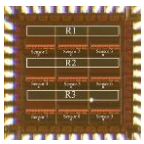

Chip: array of heaters and temp sensors
Bright idea: heat the chip and observe ΔT

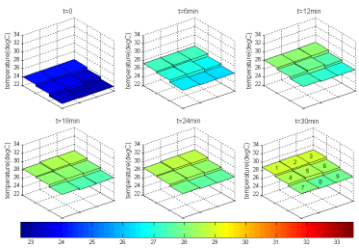

PARASITICS NOT CONSIDERED OR IMPROPERLY MODELED

Timothy Constandinou, Imperial College London

4. BUS CAPACITANCE PARASITICS INCORRECTLY CALCULATED FOR PRE-CHARGED OUTPUT

The Design (4096 element address bus output)

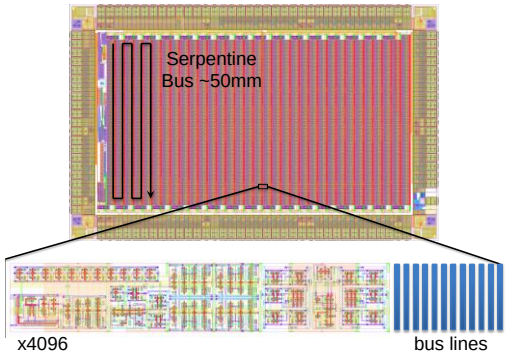

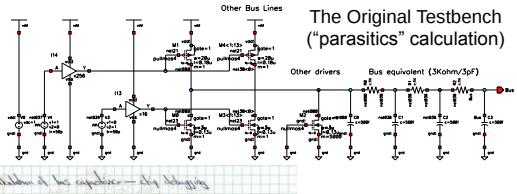

The Original Testbench ("parasitics" calculation)

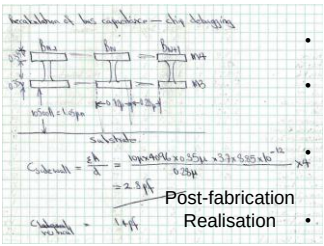

Post-fabrication Realisation

- 48mm 12-bit bus operating on pre-charge scheme
- Asynchronous timing with 50% margin (based on simulations from above)
- Underestimated parasitics resulted in incomplete precharging
- FIBed to adjust timings

Lateral capacitance can be significant, especially in deep submicron processes. The hand-calculated estimated values were only 1/3 of the actual values.

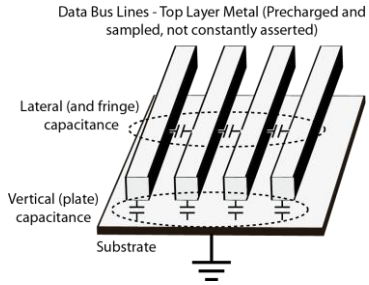

Piotr Dudek, The University of Manchester

5. FLOATING PIN IMPROVES PERFORMANCE

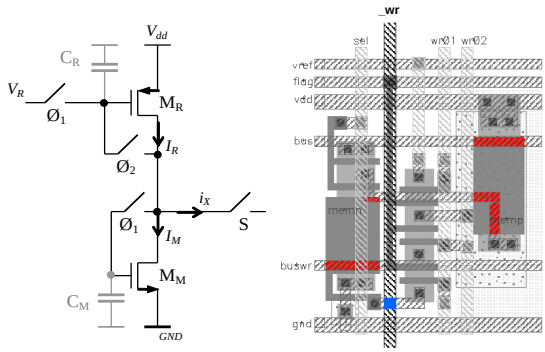

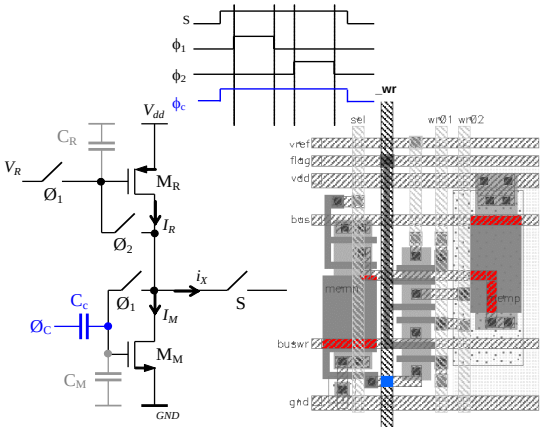

Piotr Dudek, The University of Manchester

6. LAYOUT EXTRACTION TOOL DOES NOT UNDERSTAND MAXWELL'S EQUATIONS

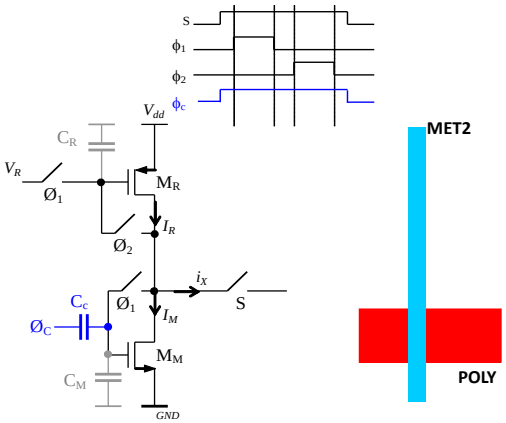

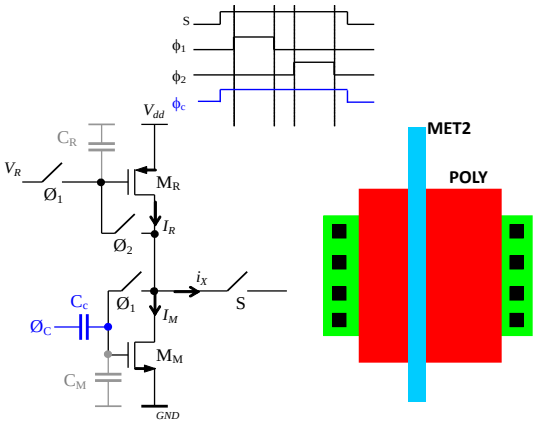

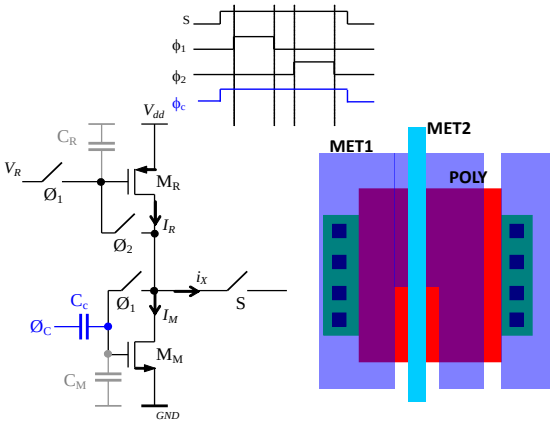

3D structure of the poly to metal2 capacitor.
There is no metal 1 directly in the designed capacitor area, but it is all around, and it makes all the difference!

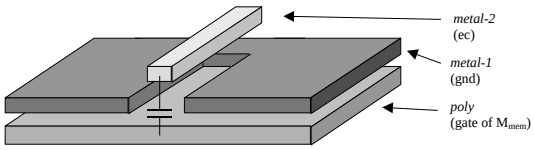

$C_{extracted}$ = 0.723fF (M2-POLY area, peri)

C_{actual} = 0.035fF (Maxwell, 3D)

Gert Cauwenberghs, University of California San Diego

7. ONE OF THE MOST COMPLEX 3-BIT DAC AND ADCS EVER REALIZED

My first VLSI chip design was a memorable and rewarding experience as part of taking Carver Mead's analog VLSI and neural systems course at Caltech in 1989. Inspired by Hopfield and Tank's analog neural model of an optimization network for analog-to-digital conversion, I decided to give it a try and implement this neural model in silicon. Having learned about the problems of mismatch with transconductance-based circuits from the labs, I further decided to do the whole thing with switched capacitors, and figured out an architecture and clocking scheme to get rid of capacitance mismatch altogether. Tobi Delbruck, my TA, thought I was nuts, but was supportive and encouraging nevertheless. The Mathematica symbolic calculations and the Analog transistor-level simulations proved that I was right: the circuit converged to arbitrary precision limited only by the number of binary stages, and by switch injection noise. Not deterred by any potential source of imprecision, I further implemented one of the fancy schemes of switch injection noise cancellation that I found in the literature. With some polygon "pushing" effort I was able to cram the entire circuit of a 16-bit DAC and ADC onto a single Tiny (really tiny!) chip in 3um 2P2M CMOS technology. The layout looked beautiful, certainly when staring at the Wolcomp screen after a 48-hour shift on the Chipmunks! When the chip came back from MOSIS after summer, I was elated to give it the carefully designed test sequence to measure its performance. My enthusiasm was contained by the three effective bits of INL and DNL observed on the oscilloscope. Even though the switch injection cancellation circuits worked as advertised, the mismatch was severe because of one critical oversight: the capacitors were not floating (as modeled in Analog), but had significant backplate capacitance to the substrate. ***Moral*** – As Carver told us all along, listen to the technology and find out what it is telling you. No device in silicon, no matter how elegantly modeled in Mathematica and Analog, is far away from the electrons and holes in the substrate.

Afterthought: The circuit may work a lot better in Silicon-on-Sapphire! If anyone is interested in giving it a try, contact me and you're in for another great experience.

Bernabe Linares-Barranco, Sevilla Microelectronics Institute

8. GROUND AND POWER CONNECTIONS TO CORE ARE TOO SLIM

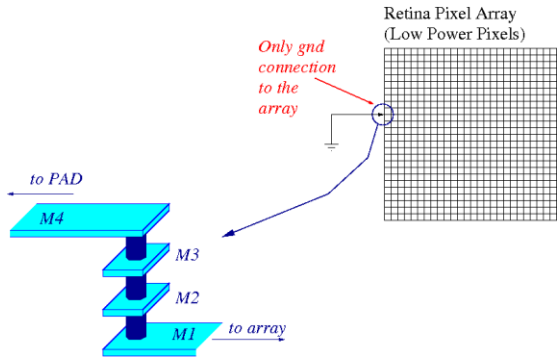

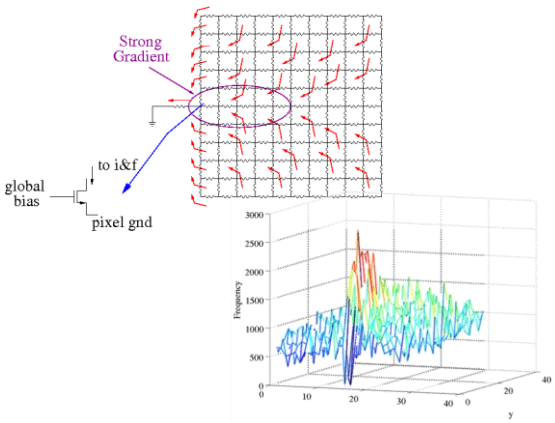

M. K. Law and A. Bermak, Hong Kong University of Science and Technology

9. BEWARE OF PARASITIC PHOTODIODES IN CMOS IMAGE SENSOR DESIGN

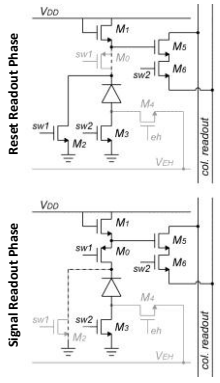

Measurement shows that output signal power is reduced because the reset voltage is light sensitive!!

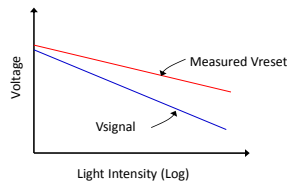

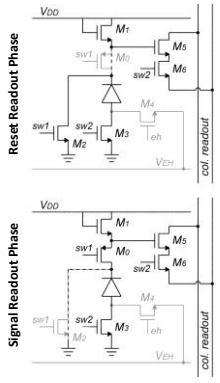

Readout voltages

$$V_{reset} = V_{DD} - V_{th,M1} - nV_T \ln\left(\frac{I_{leak,M0}}{I_0}\right)$$

$$V_{signal} = V_{DD} - V_{th,M1} - nV_T \ln\left(\frac{I_{ph} + I_{leak,M2}}{I_0}\right)$$

$$\Delta V = nV_T \ln\left(\frac{I_{ph} + I_{leak,M2}}{I_{leak,M0}}\right)$$

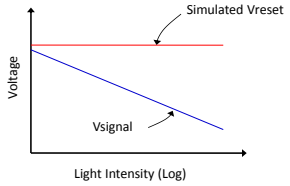

Because the transistors were not shielded from light, their parasitic photodiodes reduced performance of the FPN correction mechanism.

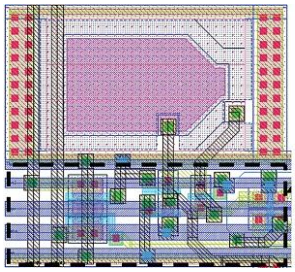

Solution:

Cover transistors with metal to block light from parasitic source / drain photodiodes. Nfets need additional lateral shielding to block diffusing minority carriers.

Tobi Delbruck, University of Zurich and ETH Zurich

10. A BIPOLAR IMAGER WITH ONE GIANT PIXEL

Our pulsed bipolar pixels

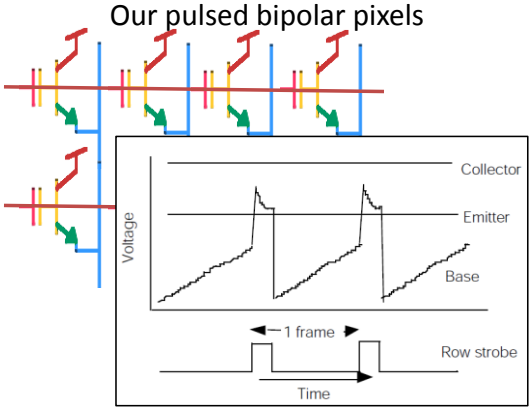

Because the bipolar p-base implant was set to **400keV** instead of **40keV**, it penetrated through the field oxide, creating **a single giant pixel**.

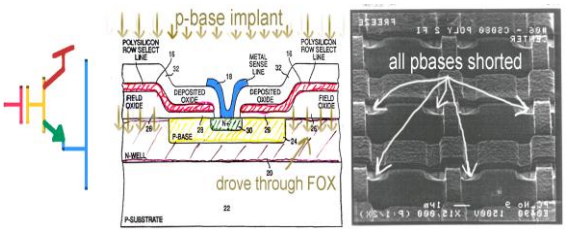

Tobi Delbruck, University of Zurich and ETH Zurich

11. METAL DENSITY RULES ARE THERE FOR A REASON

Our Foveon imager:
9 Mpixel pulsed bipolar array.
3um pixels in 0.25um technology

Close-up of pixel layout

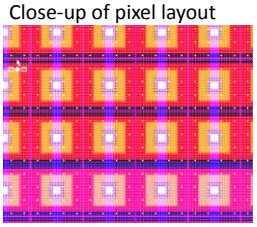

We only have thin M1 wires

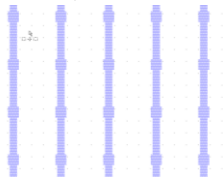

Result: CMP polishes array to a different height than rest of wafer and many wires have opens!

LOOKING AT THE BIGGER PICTURE

Jennifer Blain Christen, Arizona State University

12. I WAS A PID CONTROLLER

Marc Dandin and Pamela Abshire, University of Maryland, College Park

13. OPTICAL ORDER SORTING FILTERS MATTER!

Monochromator System

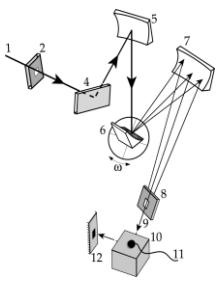

1– White Light
12– Device Under Test (DUT)

Light Incident on the DUT

Monochromatic signal of wavelength λ

But harmonics are present:
λ/2, λ/4, λ/8, etc…

Solution

Place long-pass filter with cut-on wavelength greater than λ/2 between 2, and 4.

Such a filter is called an order sorting filter.

Example:

For $\lambda = 1000\ \mu m$
λ/2, λ/4, etc.., would be present
Order sorting filter: 535 nm cut-on

Jon Tapson, Department of Electrical Engineering, University of Cape Town

14. PAY ATTENTION TO THE STATISTICS OF EXTREME EVENTS

In his books *The Black Swan* and *Fooled By Randomness*, Nassim Nicholas Taleb has highlighted that people are bad at making common sense judgments about extremely unlikely events, and particularly bad at allowing for the consequences of these events; we always underestimate both the likelihood and the consequences.

As a grad student or lab researcher, we tend to work with small numbers of components, and so we seldom see the full range of weirdness which can be exhibited by otherwise sensible designs, when the extremes of the range of component parameters is explored. One gets in the habit of reading only the "Typ." column of the datasheet, without considering the "Min." and "Max." columns. After all, the consequences of getting an extreme component aren't high. On the other hand, when you produce a system for industrial use and it gets deployed in large numbers, you *will* see every possible value in the entire parameter space, as well as some values which are well outside the stated limits. You had better be prepared for the consequences of industrial-scale malfunction.

This happened to me in 1999 when I was using the first generation of MEMS accelerometers to measure vibration on factory machinery. We built thousands of MEMS devices into robust little packages and stuck them on machinery from Seoul to Sydney. One of the issues we had to contend with, was that these MEMS devices had an undesirable mechanical resonance frequency at about 10x the measurement bandwidth (resonating at about 40kHz with a measurement bandwidth DC-4kHz). We figured we didn't need to worry about it because in the first place, we had a low-pass anti-aliasing filter which cut off more than a decade lower than the resonance, and secondly the Q of the resonance was so high it was unlikely there would be an vibration source at exactly the right frequency, and thirdly big industrial machines don't vibrate with any real energy at 40 kHz, do they?

Well, if you put enough devices on enough machines, you find that some of them do indeed vibrate at exactly the wrong frequency around 40kHz, and if the resonance Q is really high, a two-pole low-pass filter at 1kHz doesn't help much at all. I had had a big argument with a colleague as to whether we should go for a higher order cut-off filter, and every time a report came in from the field of another rogue accelerometer, he would shout across the lab "Yeah, we don't need no stinkin' filter!" We wound up having to replace a great many of the units. It would be nice to say that since then I have never again been caught by this kind of mistake, but it would be a lie...

STRATEGIC ERRORS

Tobi Delbruck, University of Zurich and ETH Zurich

15. DON'T COUGH UP YOUR CORE TECHNOLOGY

Giacomo Indiveri, University of Zurich and ETH Zurich

16. YOU SHOULD START WITH THE BIG PICTURE

LOGIC GOOFS

Raphael Berner, University of Zurich and ETH Zurich

17. CONFUSING ACTIVE HIGH AND ACTIVE LOW SIGNALS

What went wrong in in this asynchronous logic

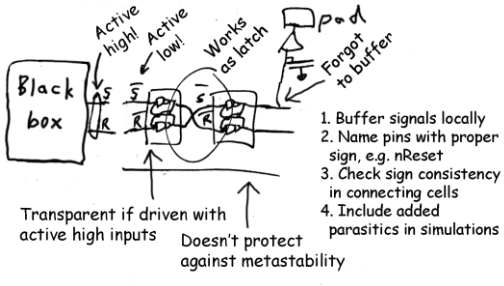

T. Serrano-Gotarredona and B. Linares-Barranco, Sevilla Microelectronics Institute

18. DON'T JUST IMPROVE ASYNCHRONOUS CIRCUITS BY ELECTRICAL SIMULATION

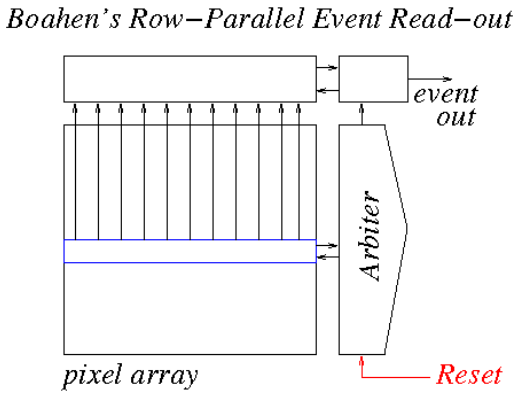

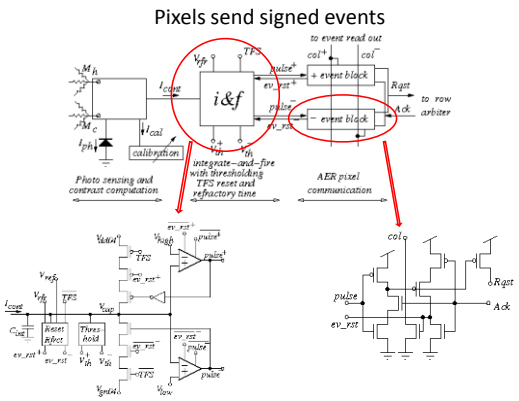

Ralph Etienne-Cummings, Johns Hopkins University

19. A 7 AMP SRAM MEMORY - SMALLER IS NOT ALWAYS BETTER

Incorrectly sizing the FETs and using pass gates in the left SRAM cell led to a huge static power consumption at the system level. A better way is shown on the right.

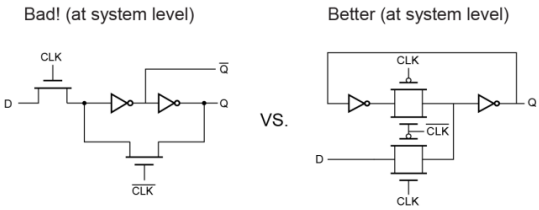

Shoushun Chen, Nanyang Technological University

20. ADDRESS DECODING GLITCHES RESET PIXEL

Glitching from decoder logic ruined this design for a motion detection image sensor

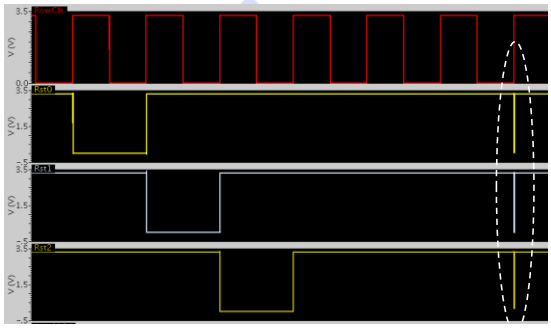

LAYOUT VS. SCHEMATIC (LVS) CONFESSIONS

Amir Eftekhar, Imperial College London

21. WRONG PADS USED ON THE PAD-RING

WITH NO SLEEP AND A LOOMING DEADLINE THE DIGITAL INPUT BONDPADS WERE MADE OUTPUTS…

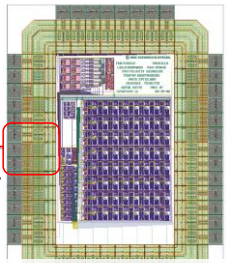

…chip worked, just couldn't control it

Lesson: Plan, plan, plan! When you have several people helping out (not all involved with design), not much sleep and no common document to work from mistakes will happened.

A simple checklist is all we needed

Tobi Delbruck, University of Zurich and ETH Zurich

22. CHIP LACKS A GLOBAL RESET

The core reset was never tied to a pad because it wasn't put in the top schematic or labeled clearly on the layout

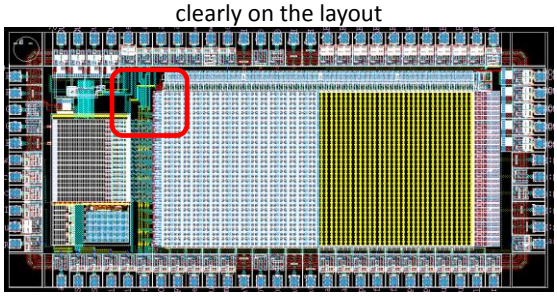

Lessons: Put your IO in top level schematic early.
Label your ports clearly and bring them to periphery of cells.

Rafael Serrano-Gotarredona, T. Serrano-Gotarredona and B. Linares-Barranco, Sevilla Microelectronics Institute

23. PULL-UP TRANSISTORS ARE MISSING

Boahen's Row–Parallel Event Read–out

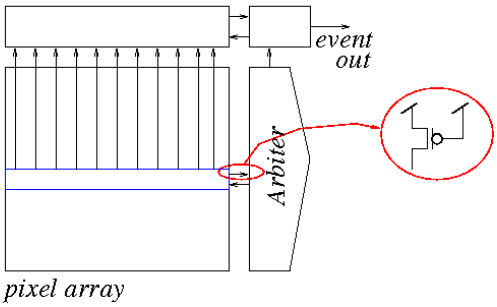

- Behavioral description of Analog Blocks (AHDL)
- Digital blocks described with Verilog (VHDL)
- Pull-ups were not included in Schematics, but behavioral simulations worked fine
- Pull-ups forgotten in layout
- LVS checked fine
- Chip did NOT produce outputs
- Fortunately, there were 2 spare pads -> FIB

ANALOG GOOFS

Wei Tang and Eugenio Culurciello, Yale University

24. PROBLEMS OF ASYNCHRONOUS DELTA MODULATOR USED IN BIO-POTENTIAL SIGNAL RECORDING

Subtraction-based level-crossing sampling requires fast reset. Because the main op-amp is slow, the original circuit suffers from distortion during resetting. A buffer fixes the problem.

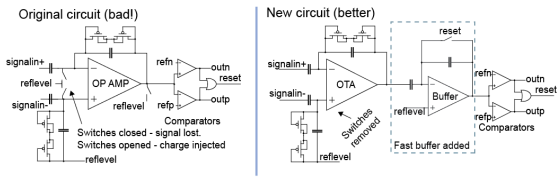

Raphael Berner and Tobi Delbruck, University of Zurich and ETH Zurich

25. DON'T GIVE UP YOUR FREEDOMS BEFORE YOU KNOW YOU DON'T NEED THEM

We were building a pixel with a 2-stage switched-cap amplifier.

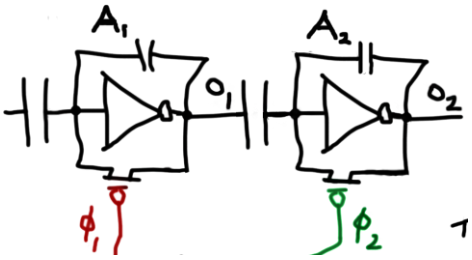

Because of our fancy scheme for generating φ_1 from φ_2, o_1 is still changing when o_2 is leaving reset, causing a huge offset in o_2.

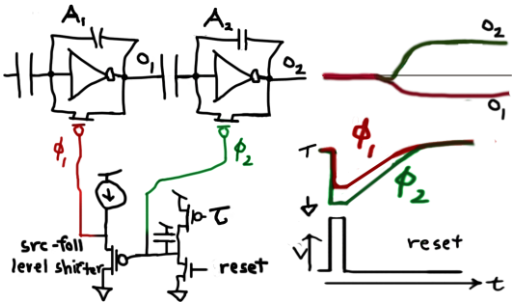

The anti-bump circuit generates a nicely-shaped "squaring" response. Ideally this shape doesn't depend on common mode voltage.

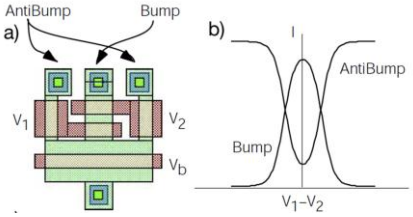

Tobi Delbruck, University of Zurich and ETH Zurich

26. BLINDED BY THEORY TO REALITY FOR 9 LONG YEARS

Reality: The shape of the "anti-bump" bowl changes dramatically with common mode DC level.
Cause: Short and narrow channel effects.
Solution: Use unit transistors with local bulk voltage

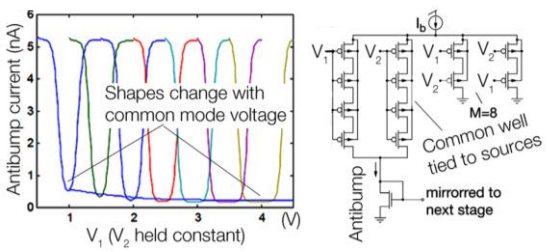

Discussion and conclusions

- Every confession comes just from members of the Sensory Systems Technical Committee. Other TCs **must** have their own confessions.
- Maybe this session can be a rotating session between active TCs?

Award: The winner for **"favorite confession"** goes to . . .

3-way tie between

Confession 5: **"Floating pin improves performance"** from Piotr Dudek

Confession 8: **"Ground and power connections are too slim"** from Bernabe Linares-Barranco

Confession 12: **"I was a PID controller"** from Jennifer Blain Christen