# Yet Another (Kind Of Bad) Hardware Random Number Generator

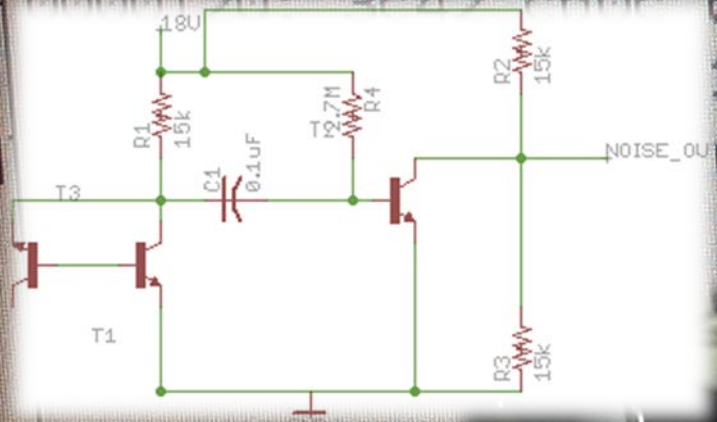


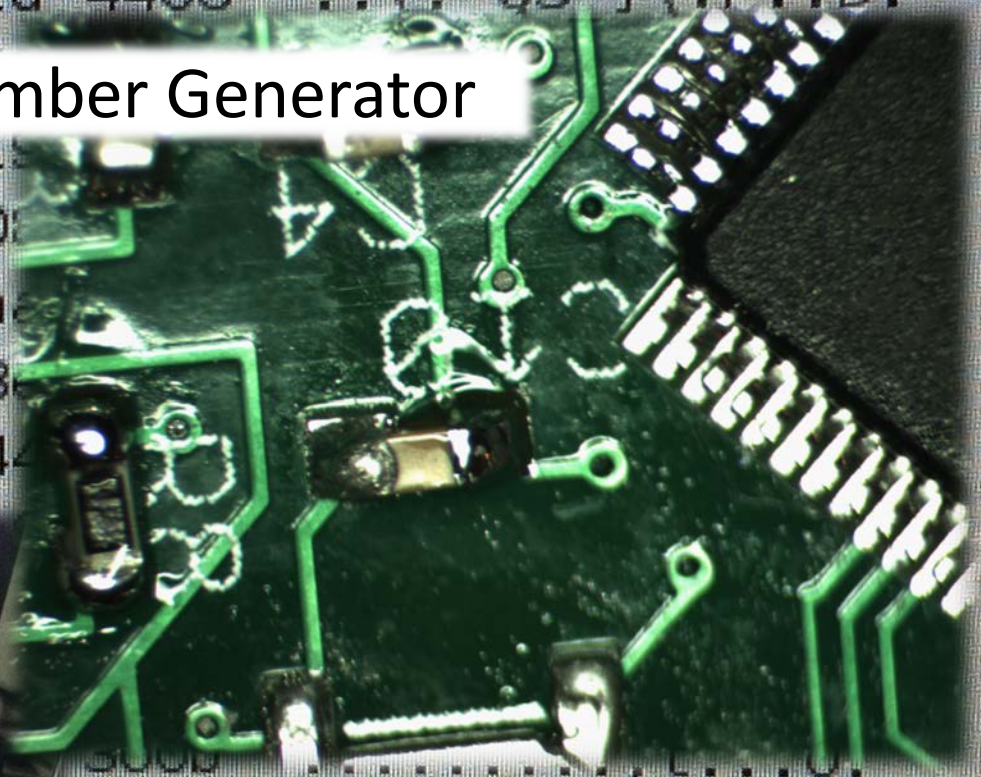

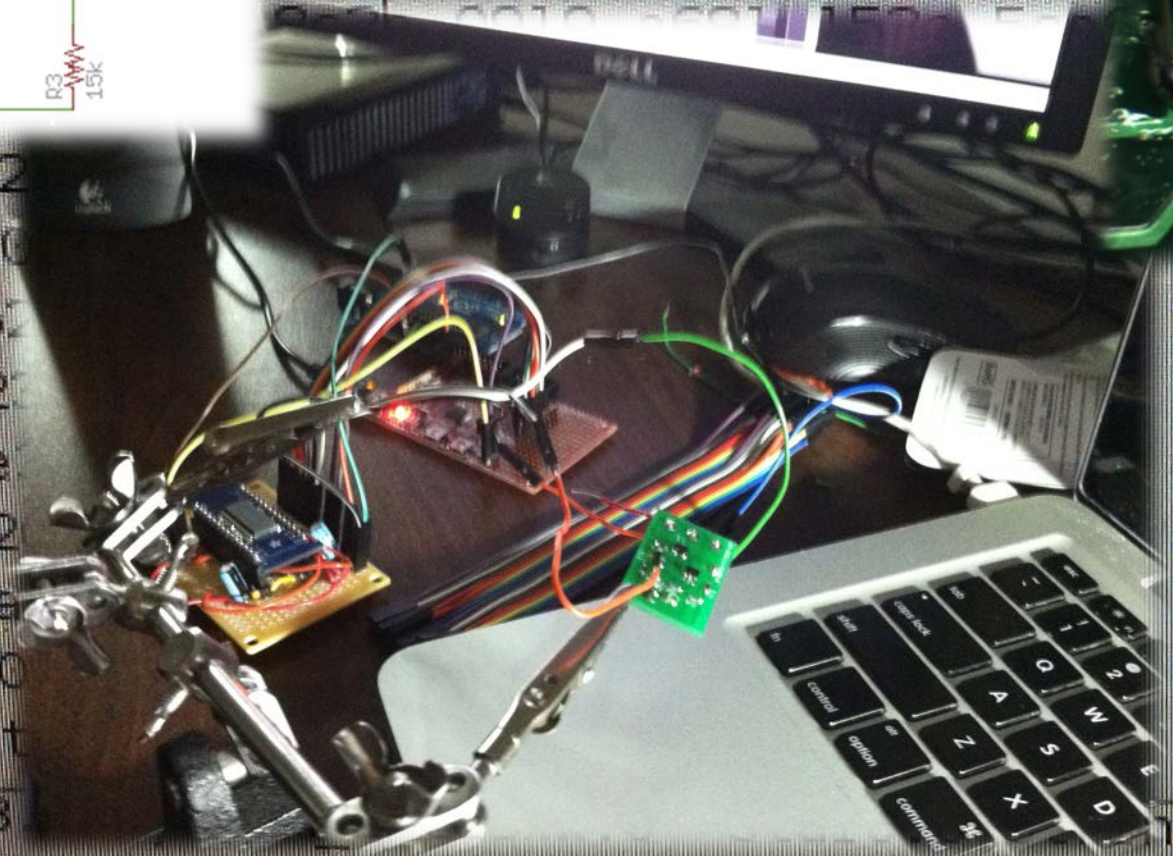

Ian Kilgore
iank@iank.org
2014-03-18

# I built a thing and I'm gonna talk about it

- Half theory
- Half case study
- ~~Half manbearpig~~

- Maybe I'll just yell for an hour about the so-called "moon"

**WHO KNOWS??**

# People are bad at random – 1/2

- We are pattern-matching cognitive-bias meat machines
- General handwaving from your presenter re: bias, influence, determinism
- Read a book

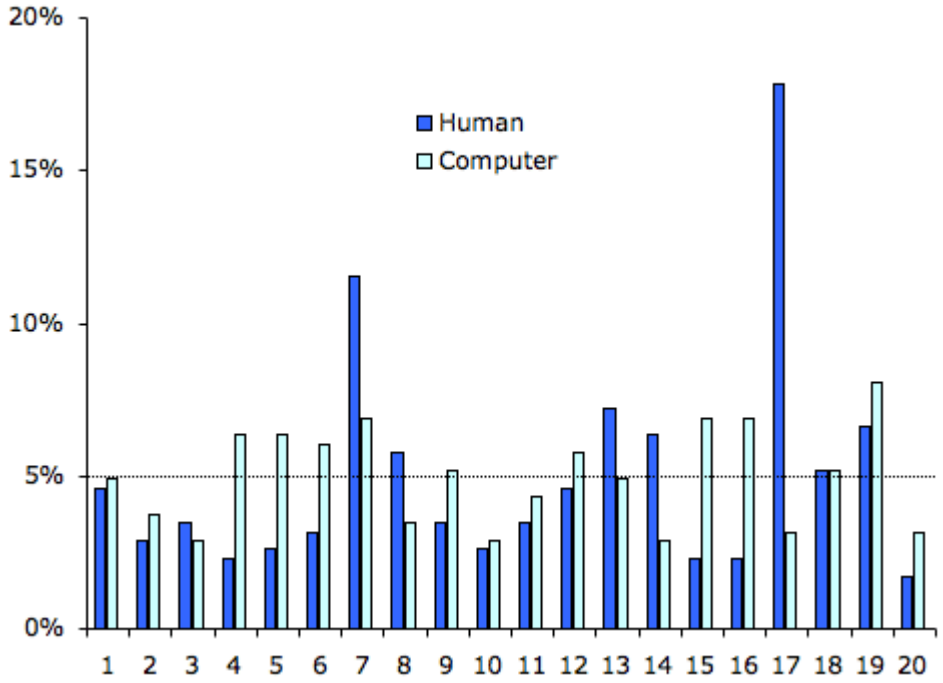


[1] Cognitive Daily, “Is 17 the ‘most random’ number?” by Dave Munger
http://scienceblogs.com/cognitivedaily/2007/02/05/is-17-the-most-random-number/

# People are bad at random – 2/2

- People are bad at FAKING random

- A: 01101001100110110101001 11
- B: 10001011011001111111 10001

# Computers are bad at random

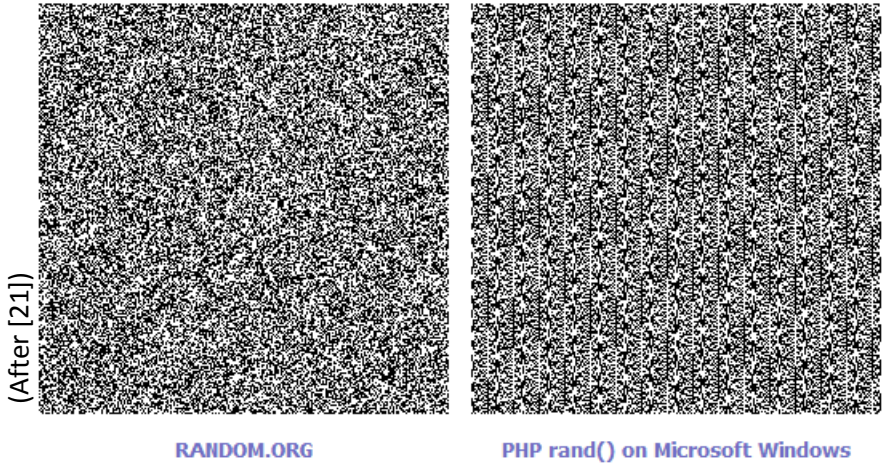


- Psuedo-Random Number Generators
  - Deterministic
    - Starting from an initial state (seed), algorithm will always produce the same sequence
    - Canonical seed is system time. (predictable). There are better ways
    - Number of embarrassing errors due to poorly-seeded PRNG
      - Netscape [16] (predictable seed: hashed time, PID, PPID)
      - Debian OpenSSL [17] (predictable seed: PID only, after entropy commented)
  - Periodic
    - Sequence does not grow beyond *n* bits
    - Finite internal state
- Entropy sources [15]
  - User input
  - Thermal
  - etc
- PRNG seeded with entropy source
- /dev/random, /dev/urandom

# Are there TRULY random sources?

- What does that even mean?
- No, seriously, what does that even mean?
- No, seriously
- What does that even mean

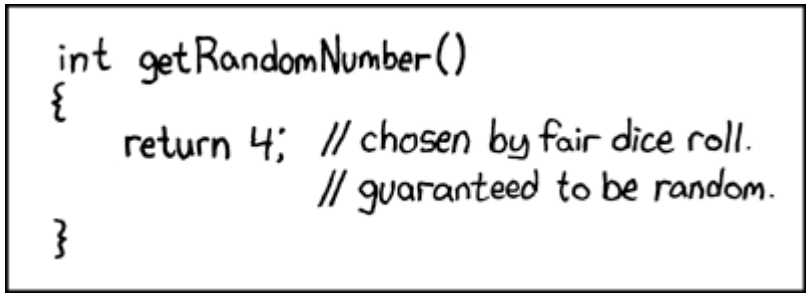



- Impossible to prove [20]
- Leave it to the philosophers

- "Good enough": cannot be predicted
  - **In the information theoretical sense**
  - Note *your inability* to predict does not mean unpredictable (stock market)
- Statistical tests: diehard [9], dieharder [10]
  - "Try it and see". Judge likelihood
  - Possible to flip a fair coin 1,000,000 heads in a row. Unlikely.

# Quantum sources

- Photon + beam splitter
- Radioactive decay
- Shot noise

# Chaotic sources - Theory

- Chaos Theory ≠ Heisenberg Uncertainty Principle
- mad about this:

- A chaotic system is:
  - theoretically deterministic
  - practically impossible to predict
    - extreme sensitivity to initial conditions
    - nonlinearity
    - topological mixing
- “Butterfly Effect”
- Weather is chaotic
  - Even ignoring quantum effects, weather models are extremely sensitive to initial conditions
  - Forecasters run 100s of models and consider most likely outcomes (e.g. “% chance of rain”)
- Nate Silver’s book [11]

# Chaotic sources - examples

- Dice-o-matic [19]
- Weather: Random.org lightning radios
- Lava lamp
- Thermal noise, Brownian noise

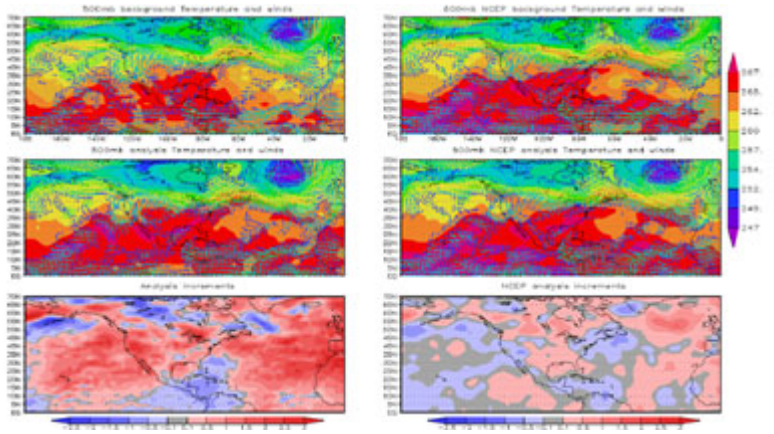

http://www.nasa.gov/centers/goddard/news/topstory/2005/esmf.html

***Credit: Carlos Cruz, Shujia Zhou, Northrop Grumman IT/GSFC; Arlindo da Silva, GSFC; Erik Kluzek, NCAR; Weiyu Yang, NCEP.***

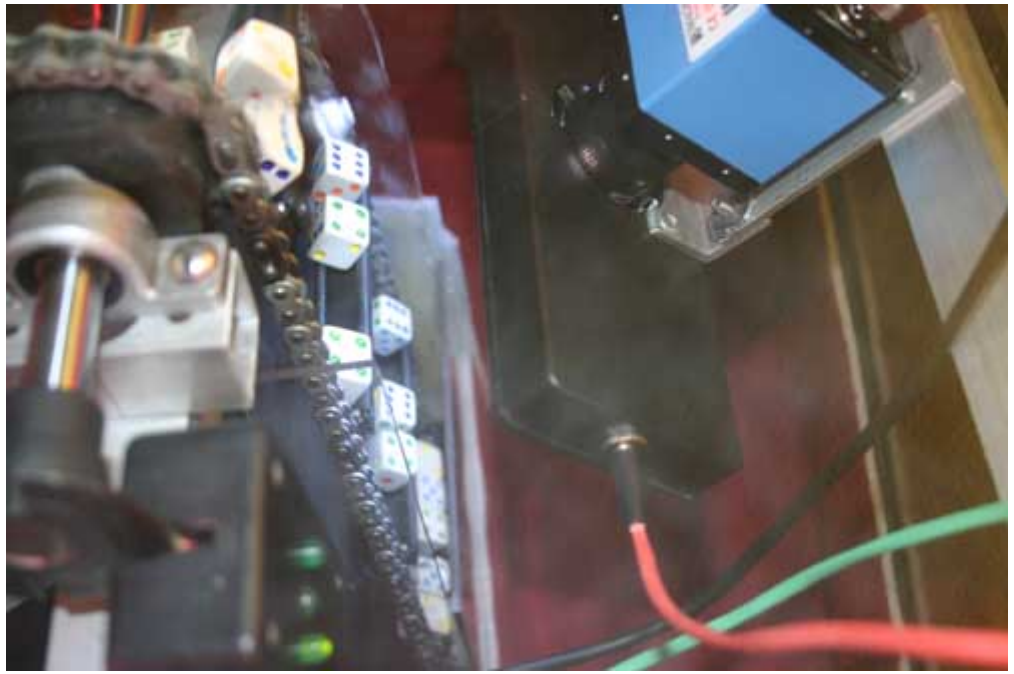

http://gamesbyemail.com/News/DiceOMatic

# Do we need good random numbers?

- Three use cases
  - High-volume
  - High-security
  - Make-believe
- Ex:
  - Scientific computation (sometimes)
  - Poker [8]
  - Lotteries [22, 23]
  - Debian package signing
  - Cryptography, esp. one-time pads
  - Sometimes the problem is just perceptual – Games by E-mail “Dice-o-Matic”

# An Aside About Threat models, or:

## Information security is fun to think about, but don't be insufferable about it

- A funny paper [5] that I'm otherwise mad at with describes two canonical threat models:
  - "Basically, you're either dealing with Mossad or not-Mossad."
  - "If your adversary is not-Mossad, then you'll probably be fine if you pick a good password […]."
  - "If your adversary *is* the Mossad, **YOU'RE GONNA DIE AND THERE'S NOTHING THAT YOU CAN DO ABOUT IT**"

# Previous slides aside, why *really* build this?

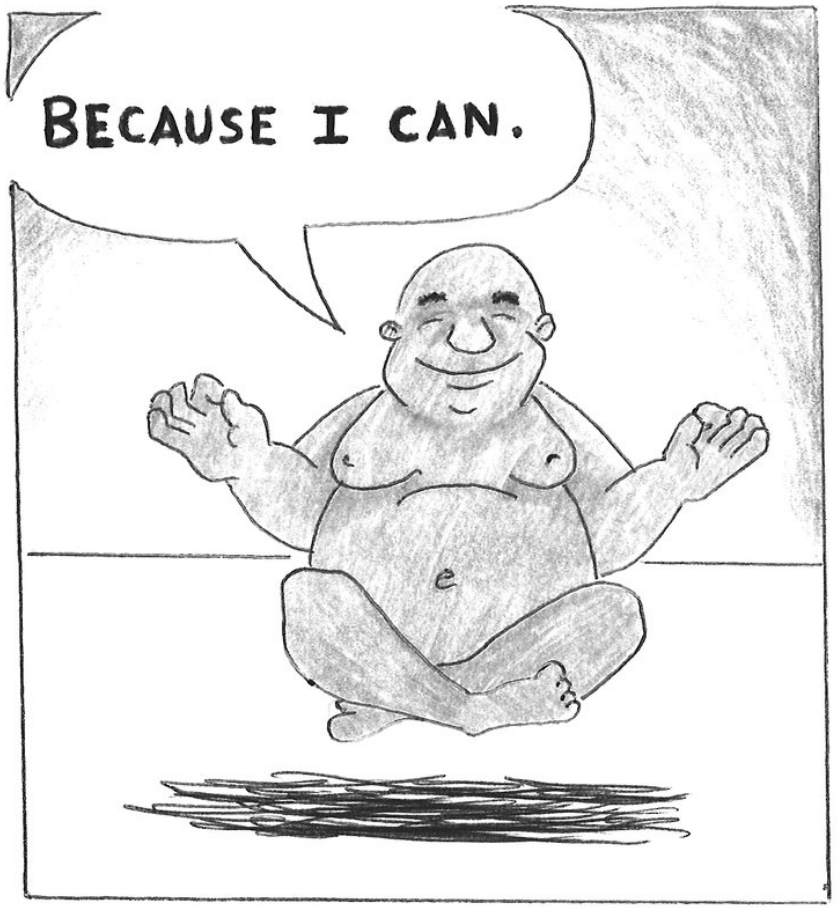


© Jacob Barrett http://www.realcostofprisons.org/comix/barrett/

# Avalanche noise - physics



- PN junction (diode)
- One-way current flow under normal conditions
  - Pause while ik draws furiously on the whiteboard about PN junctions (if time)
- Strong reverse-biased E field causes avalanche breakdown
- Impact ionization
  - energetic e- knocks another e- out of the valence band
  - creates another electron-hole pair
  - In presence of strong E field, this process can continue through width of depletion region- multiplication
  - In a diode, this is effectively multiplication of shot noise (& other phenomena)


- McIntyre, R. J. "Multiplication noise in uniform avalanche diodes." *Electron Devices, IEEE Transactions on* 13.1 (1966): 164-168.

# Avalanche noise – predictability (1/2)

- Wiki says avalanche noise isn't quantum


- Electron 'gas' is hardly classical

# Avalanche noise – predictability (2/2)

- Still, suppose that electrons in Si are classical gas
- Chaotic phenomenon which requires nearly-perfect knowledge of N initial states
  - (Heisenberg is getting antsy, but we're pretending it's classical; Uncertainty is a wave thing)
- Good discussion on wiki talk page [2]
- Let's do some Fermi/Napkin math:
  - Mean Free Path $\lambda$ on the order of:
    - 10 angstrom [3]
    - Largest measured ~135 angstrom [13]
  - Mean Free Time $\mathrm{T} = \lambda / V_d$
  - Drift velocity $V_d \sim 10^4\ m/s$
  - $\sim 10^{13}$ collisions per electron per second
  - Conservative order of magnitude estimate of N ~ $10^{10}$ electrons in our PN junction
    - (difference in between 1 $\mu m^3$ and molar volume of intrinsic Si)

So even if we consider the phenomenon to be classical (which it isn't):

**There's a !@#$' lot of Really Wild !#$@ going on**

# Avalanche noise – TRNG design

Will Ware, 1995 [4]

Also basically everyone alive has built one of these

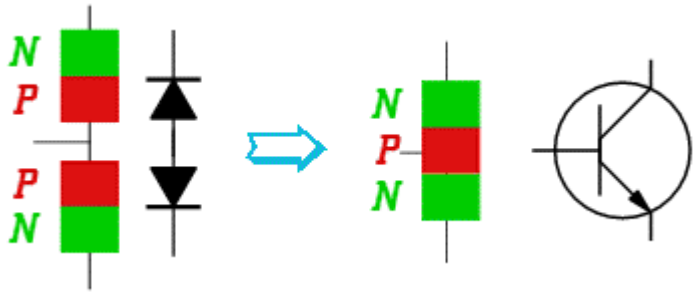


(After [14])

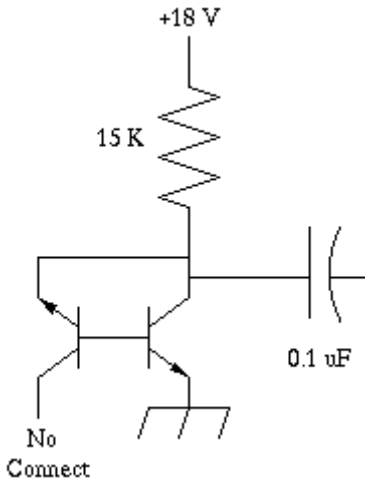


- First transistor
  - Collector N/C
  - One PN junction between base & emitter
  - Like a reverse-biased diode
- Why not diode?
  - Rectifying diodes designed for high breakdown voltage.
  - Zener diodes & others with low breakdown voltage designed to minimize avalanche noise
- Second transistor is C-E amplifier
- Capacitive coupling (high-pass filter) at output

# Prototype

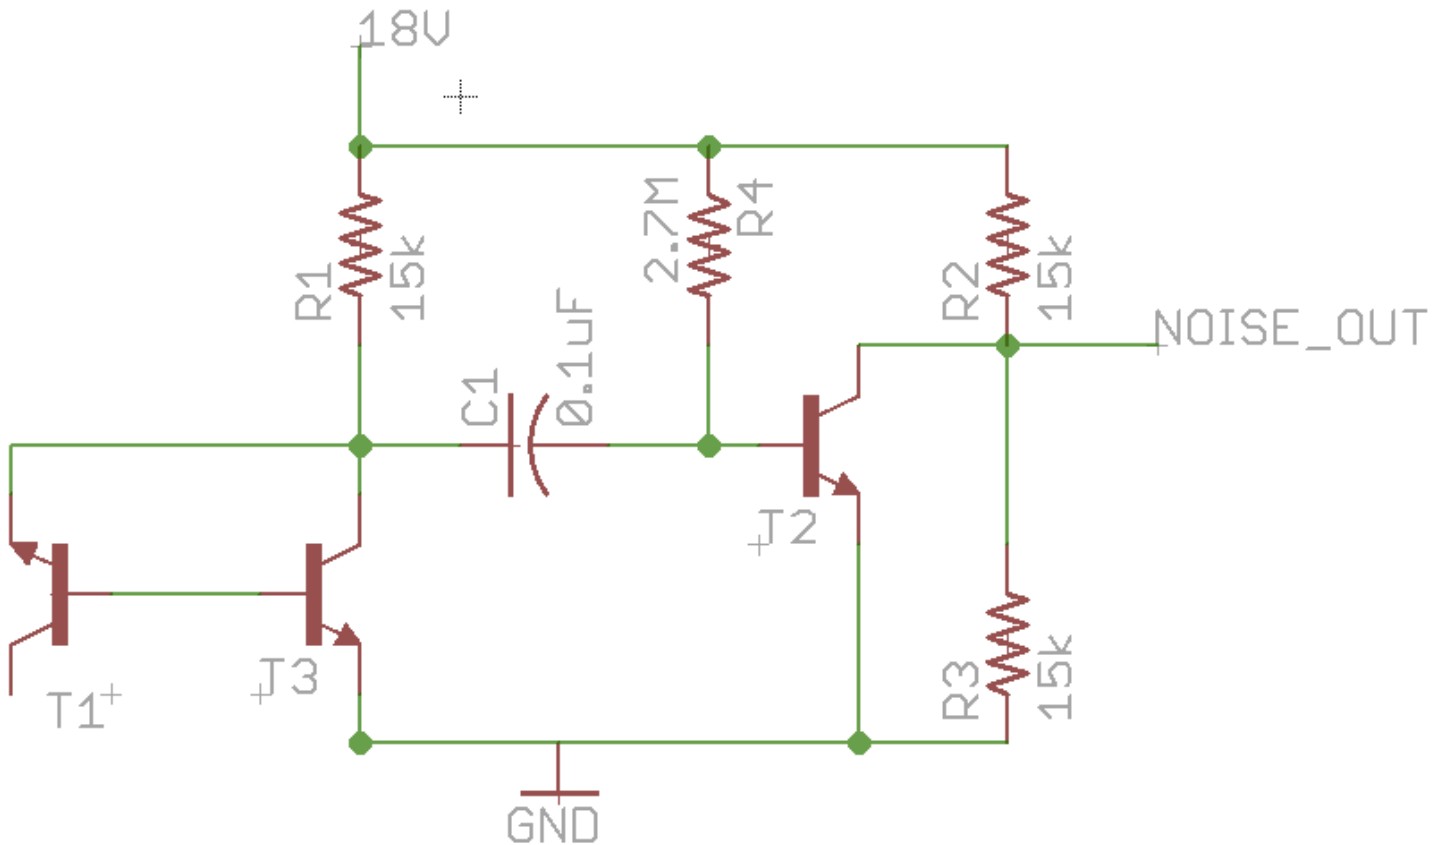

# Breadboard Prototype (failure)

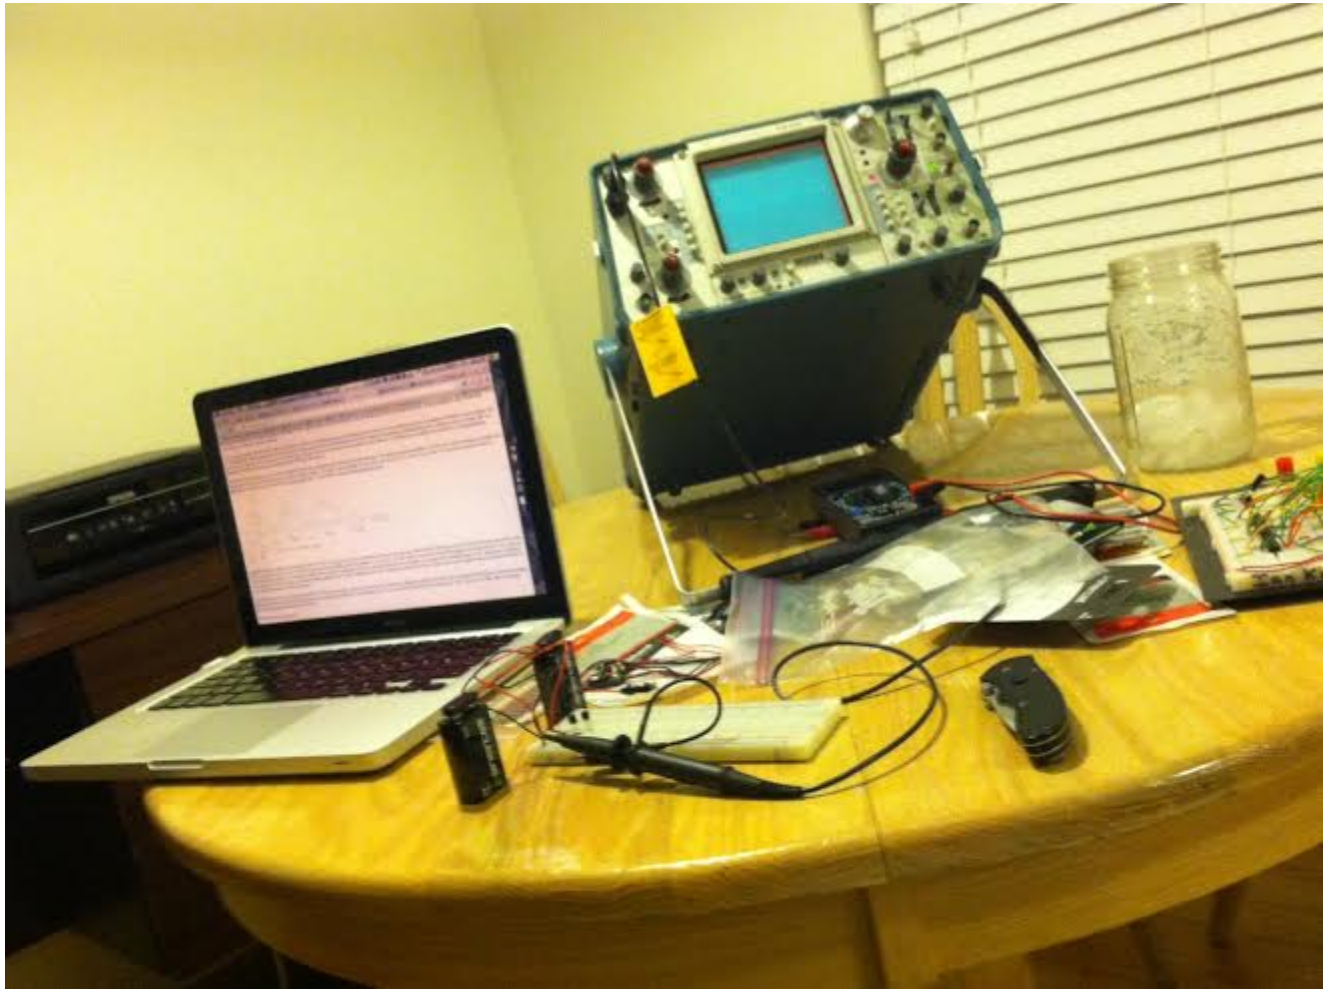

- Breadboard is a big box of parasitic reactance
- Also/or probably I hooked something up wrong

# Perfboard Prototype

- Perfboard makes me mad
- More on this later

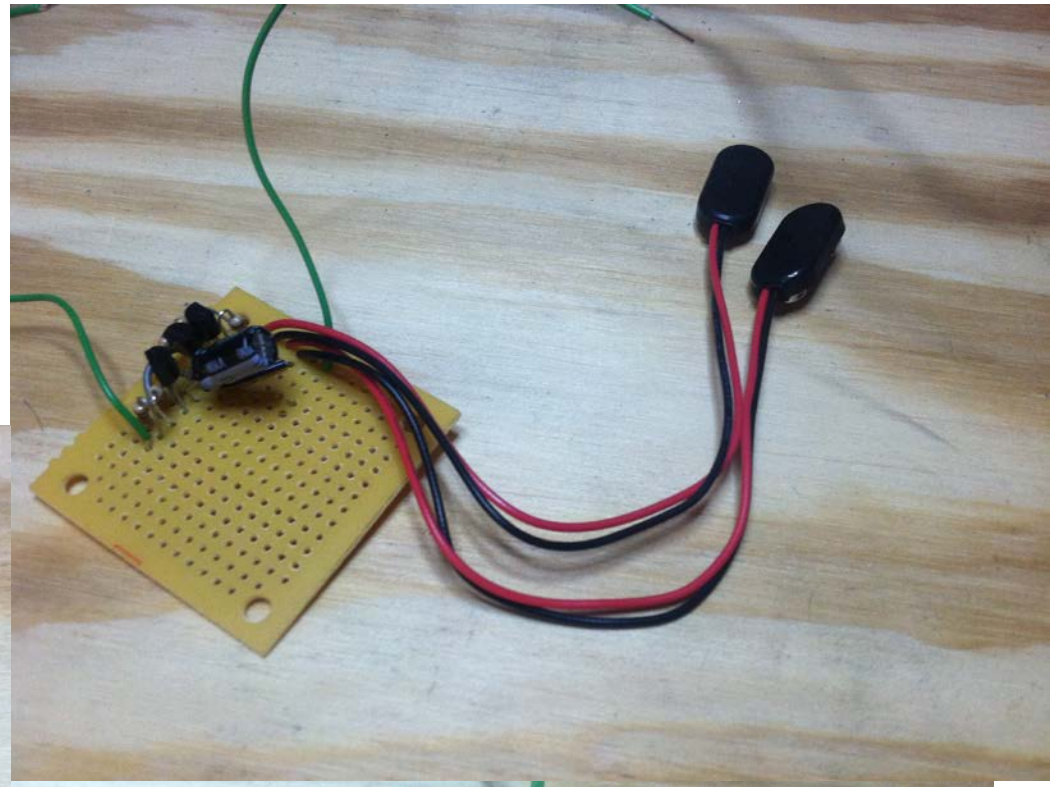

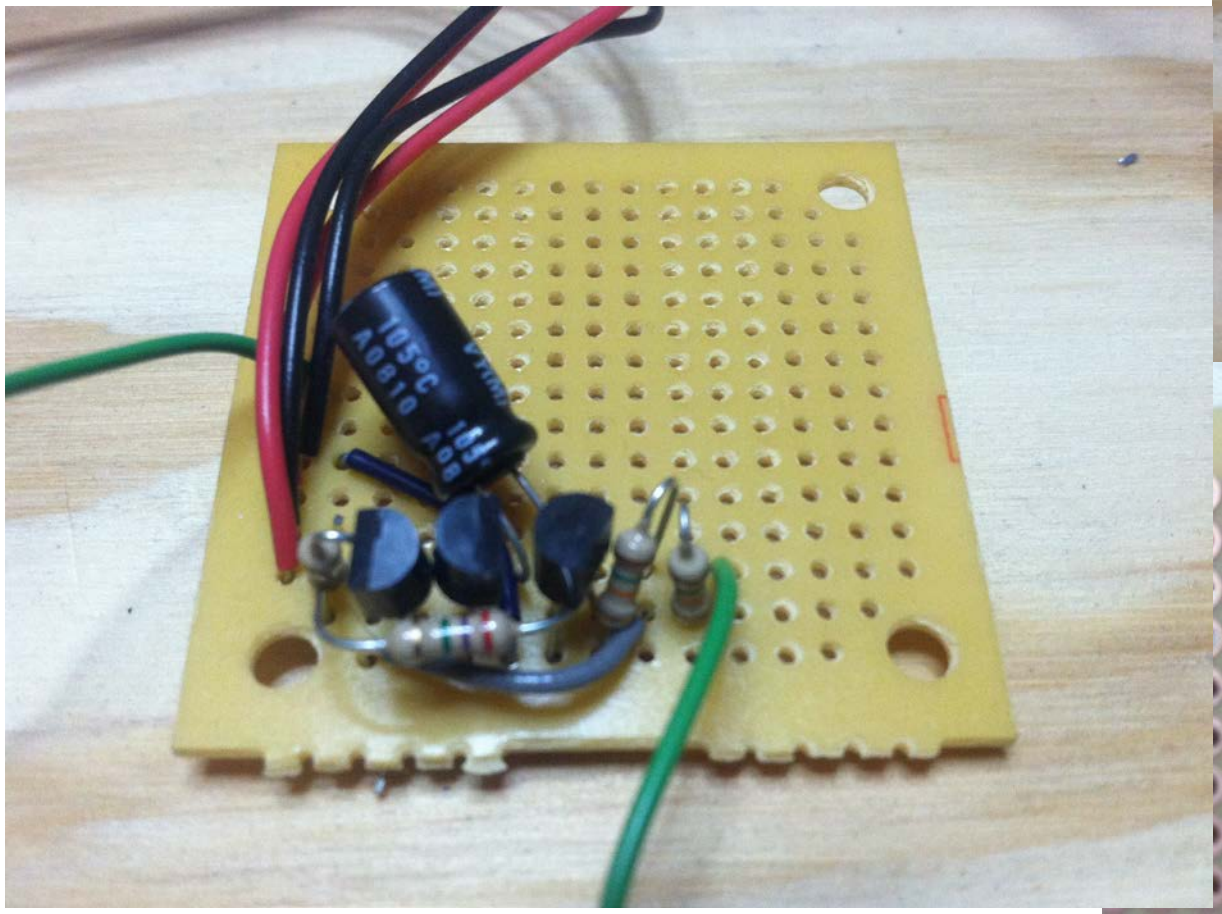


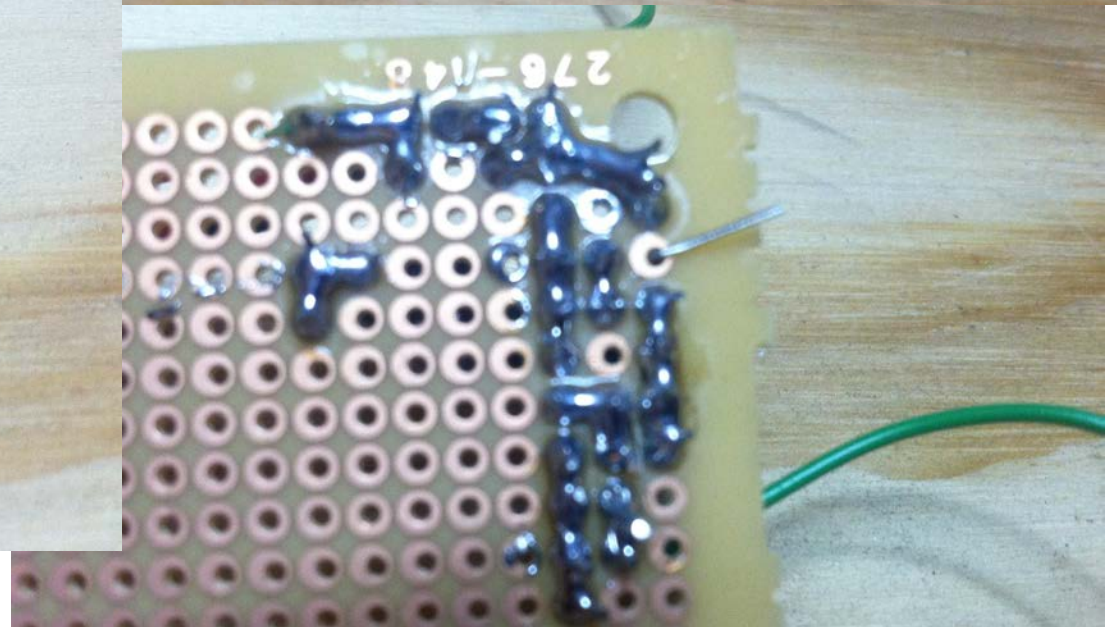

# Interface

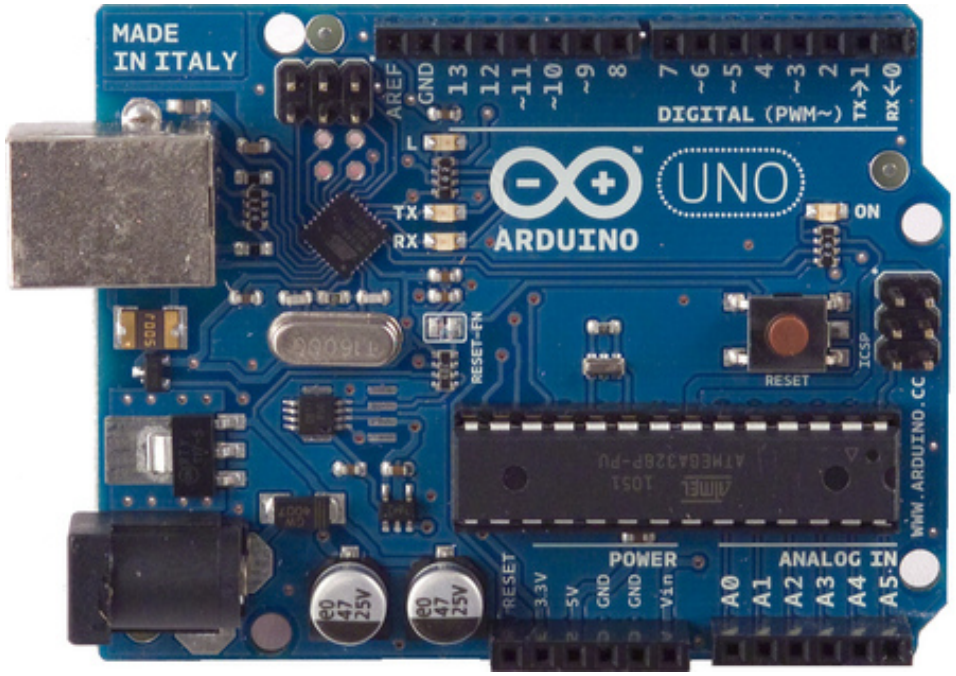


10-bit ADC, ~70k samples/sec
(~15 µs)

# Distribution/bias/correction

- Distribution not uniform
  - Still random (we hope- more later), just biased
  - Consider sum of a pair of dice
  - Can expect to get 7 more often than 12
- Software whitening
  - von Neumann's algorithm: take bits two at a time
  - 00 -> discard
  - 01 -> take 0 (or 1)
  - 10 -> take 1 (or 0)
  - 11 -> discard
  - (Think about it)

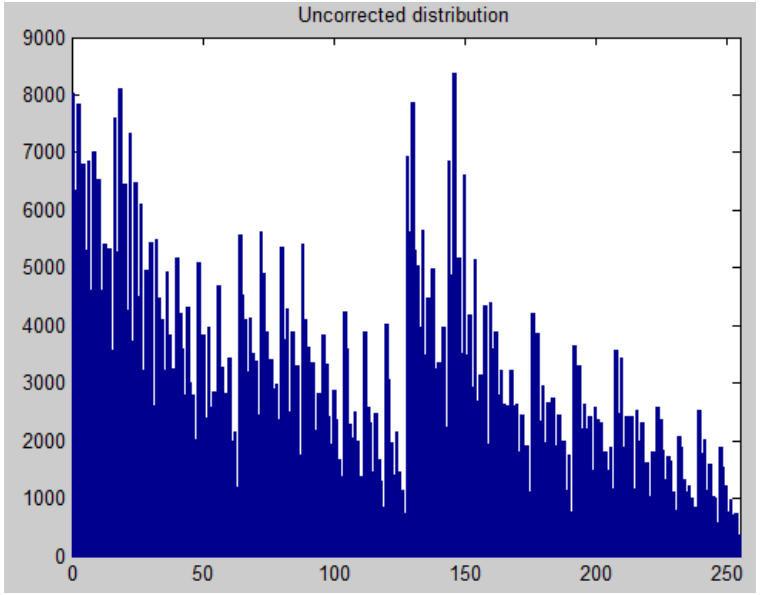


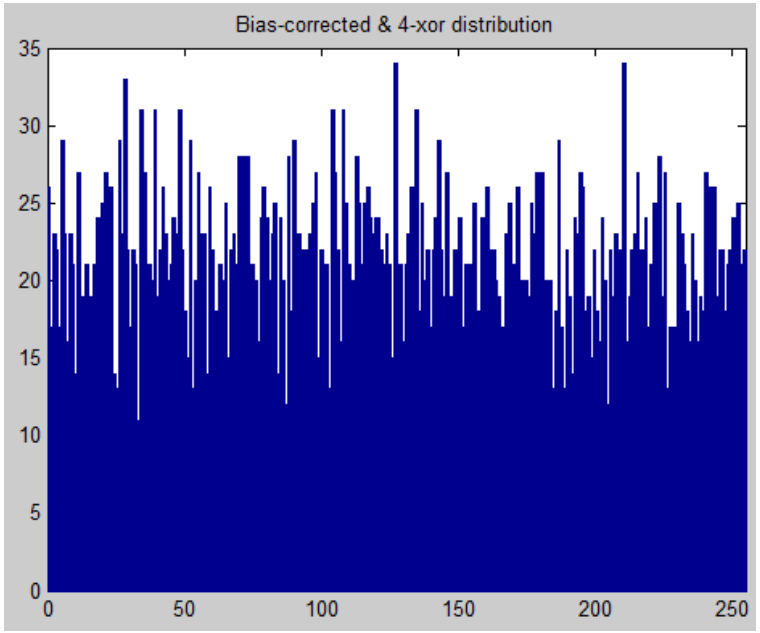


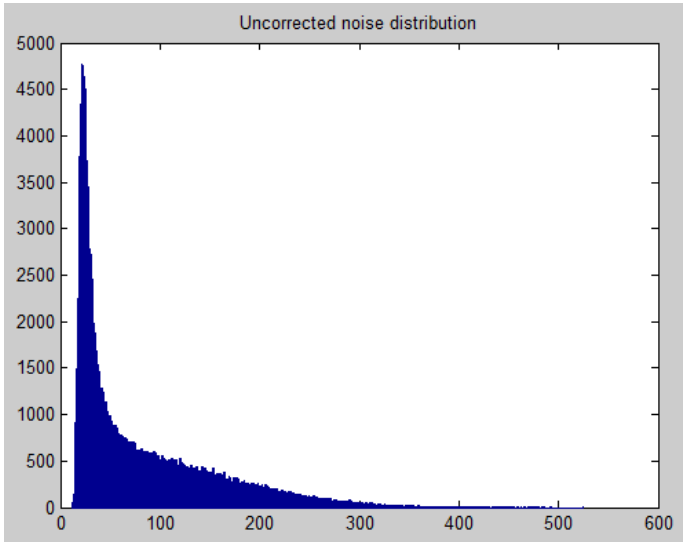

# Iteration

- Most hacks end at the PoC/prototype
- Good reasons for this
  - Do A Thing vs “plumbing”
  - Prototype took several hours and several dollars
  - Final revision: 100s of man-hours, 6 months wall-clock time, ??? dollars
- Goals
  - Single board, USB
  - Avoid “kits” / breakout boards in final revision
  - Speed
  - Also it should work

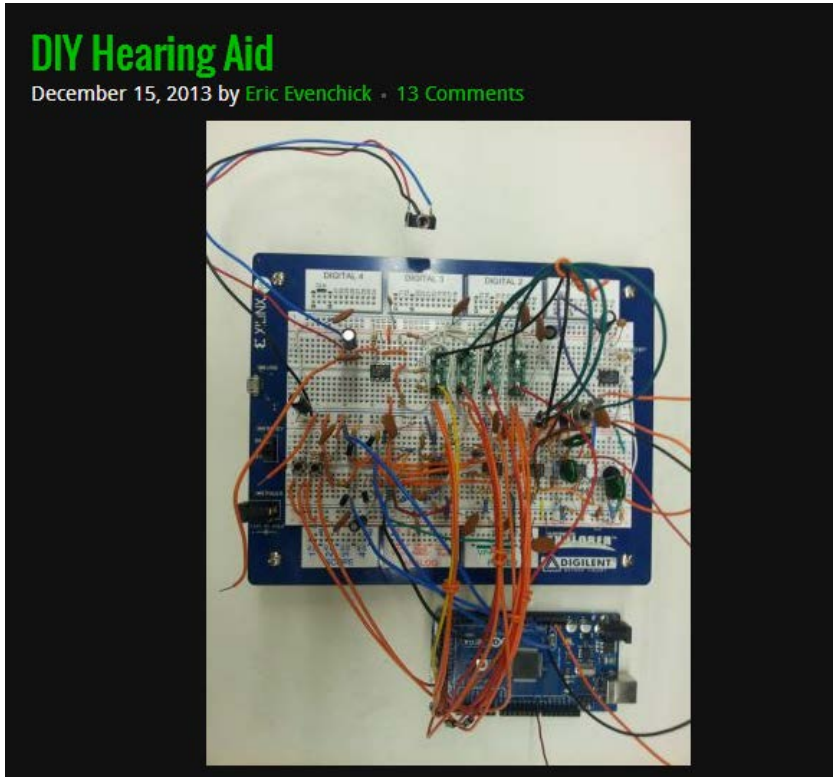


http://hackaday.com/

# Iteration – Power

- Eventually, run on USB (5V).
- Need 18V
- TPS61040 DC/DC boost converter
- Applications in datasheet

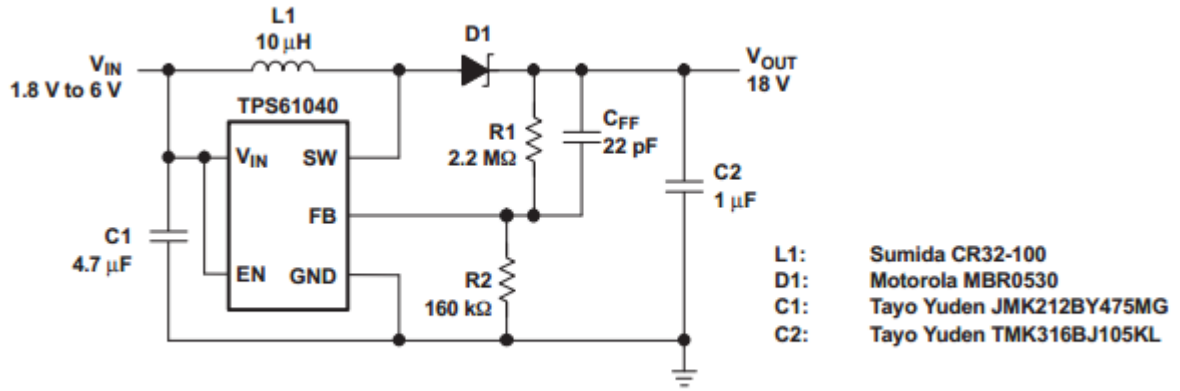


**Figure 16. LCD Bias Supply**

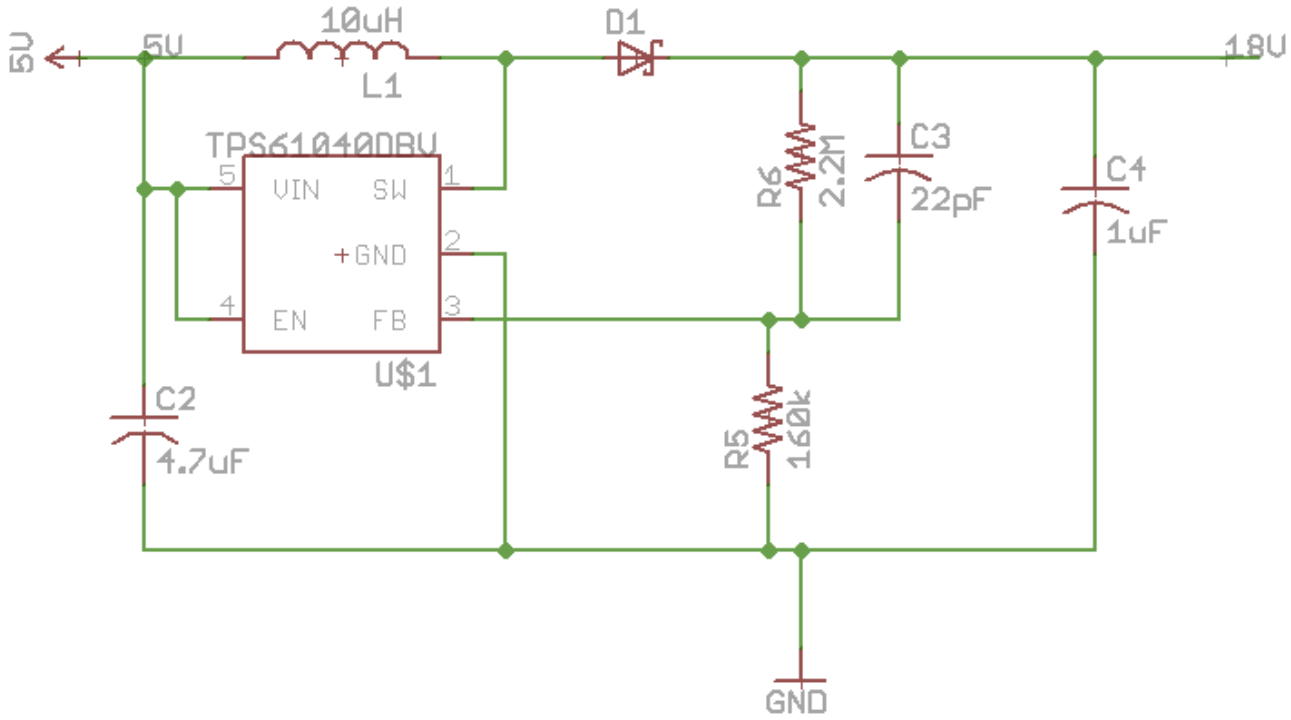

# Iteration - PCB

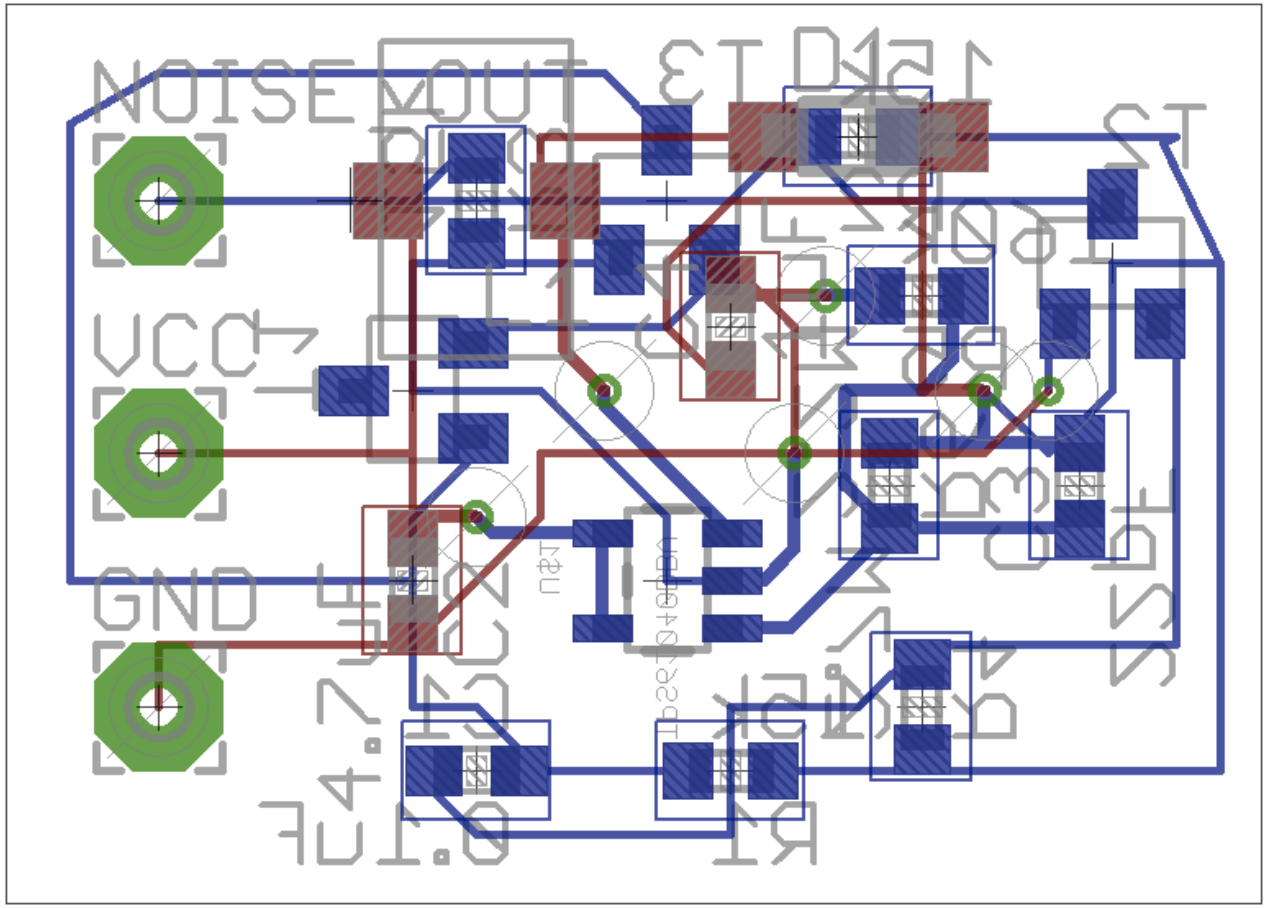


- Integrate physics mechanism & power supply
- PCBs manufactured by seeedstudio Fusion PCB service
- Real cost is time (~1 month)
- Errors are “expensive”

# Iteration - PCB

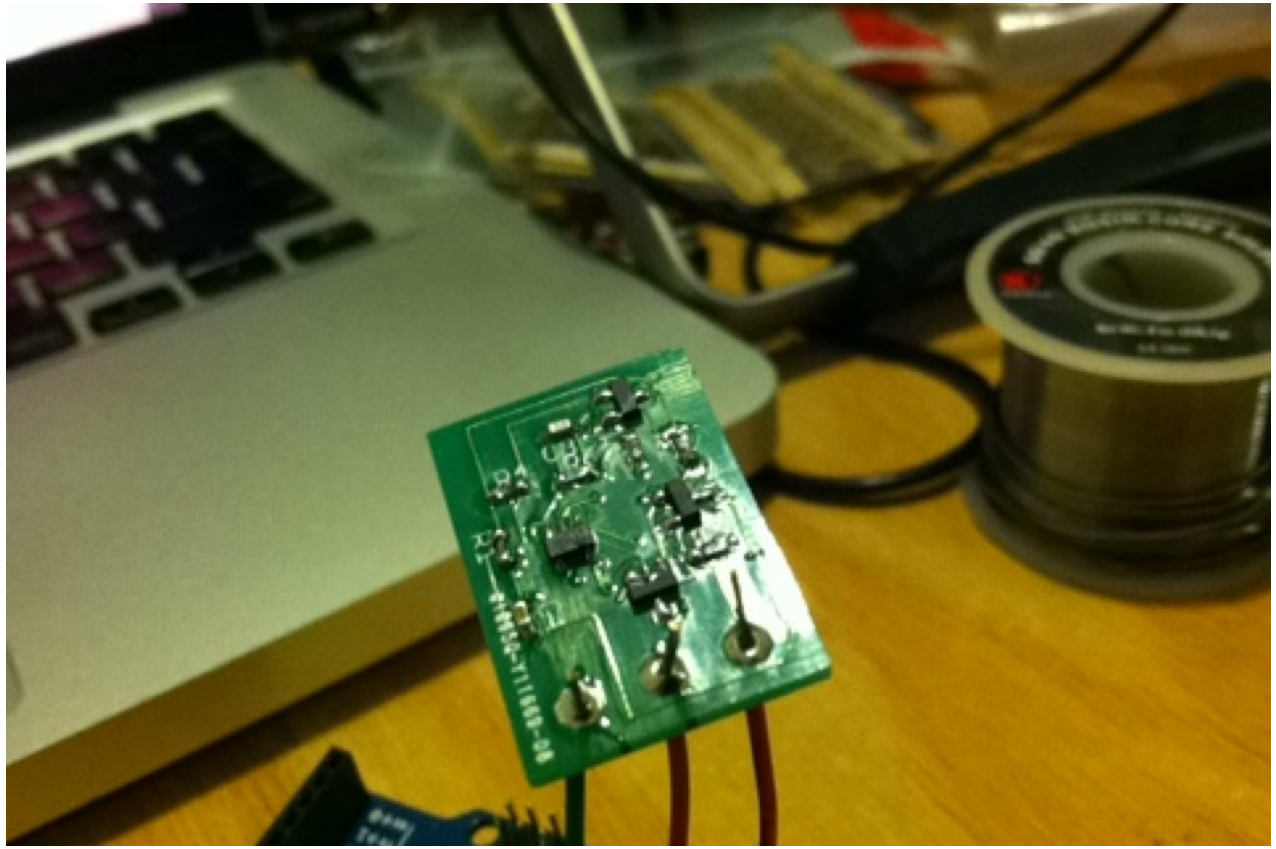

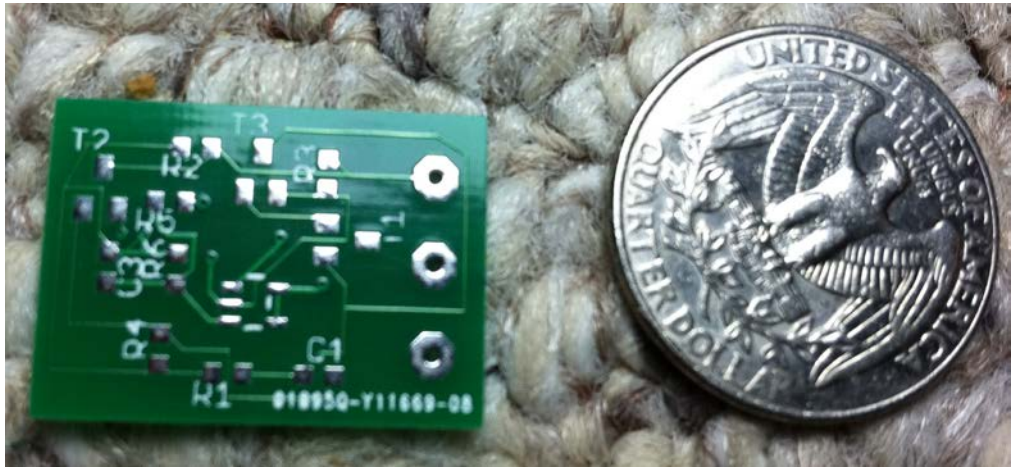


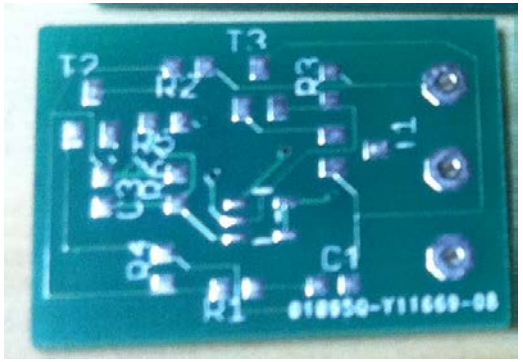


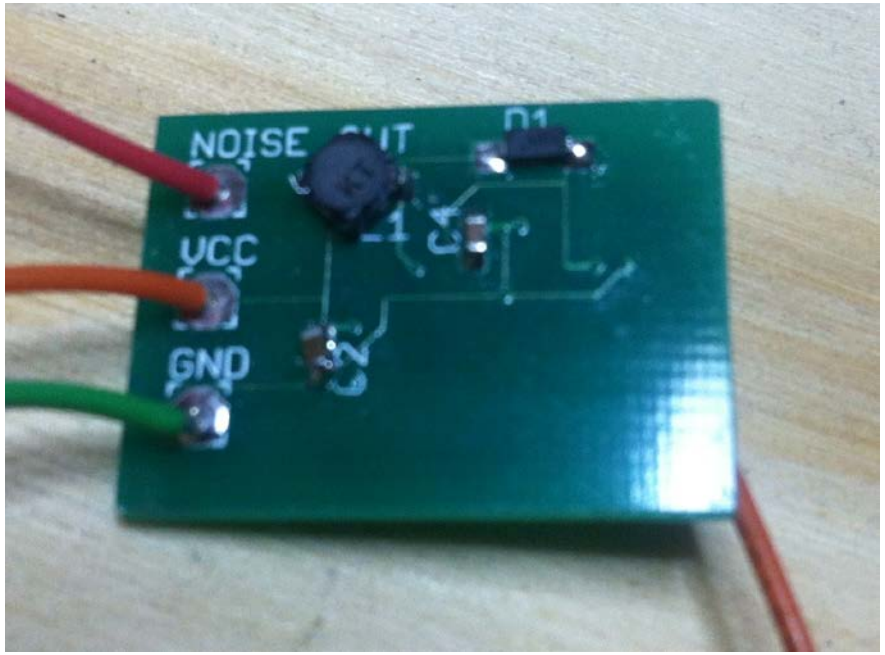

# Iteration – LPC1343

- Olimex dev board [7]
- ARM microprocessor, up to 72MHz
- Faster ADC
- Native USB 2.0
  - !!
  - vs. FTDI, V-USB, …
  - bootloader
- Wrote firmware: RNG enumerates as a USB mass-storage device. Writes fail, reads return random data.
- 40x faster (vs Arduino) (notes 2013-09-20 (1))

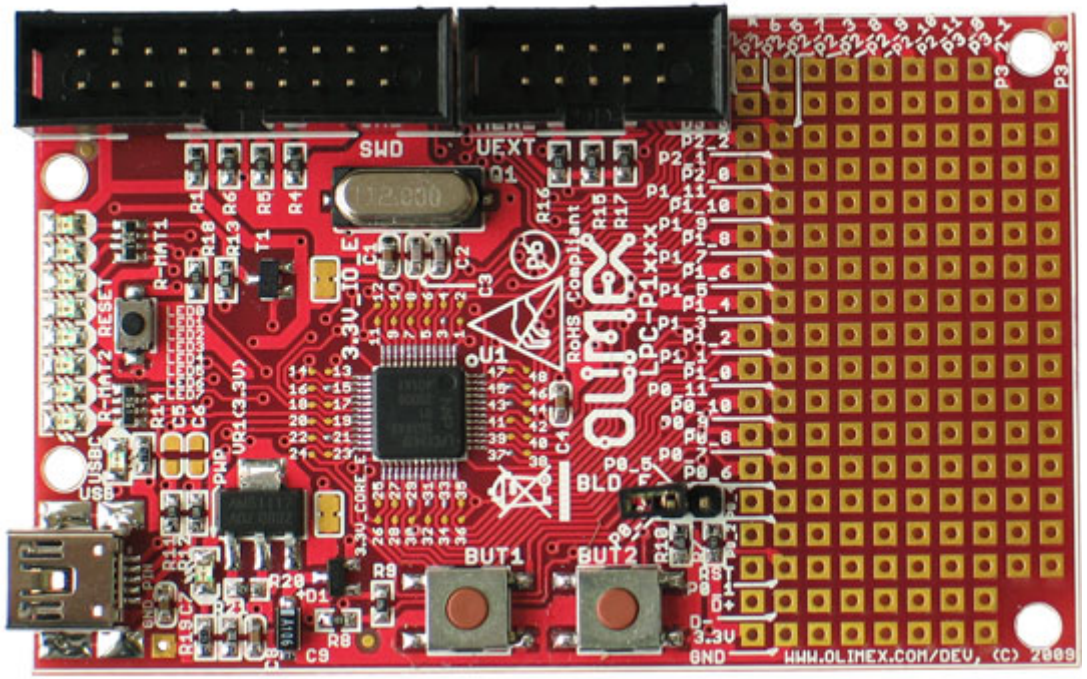

# Iteration – ADC – 1/3

- Intersil HI5767/2CBZ [18]
- Running at 12 MSPS (notes 2013-09-21 (1))
- Pin-compatible drop-in replacements up through 60MSPS available.
- Limited now by MCU's –digital– I/O sample rate, and the USB bus
- End-to-end: "whitened" random bits at 50 kilobytes / sec
- 400x speedup vs prototype
- (notes 2013-10-11 (1))

# Iteration – ADC – 2/3

- First PCB + Separate ADC on perfboard + LPC1343

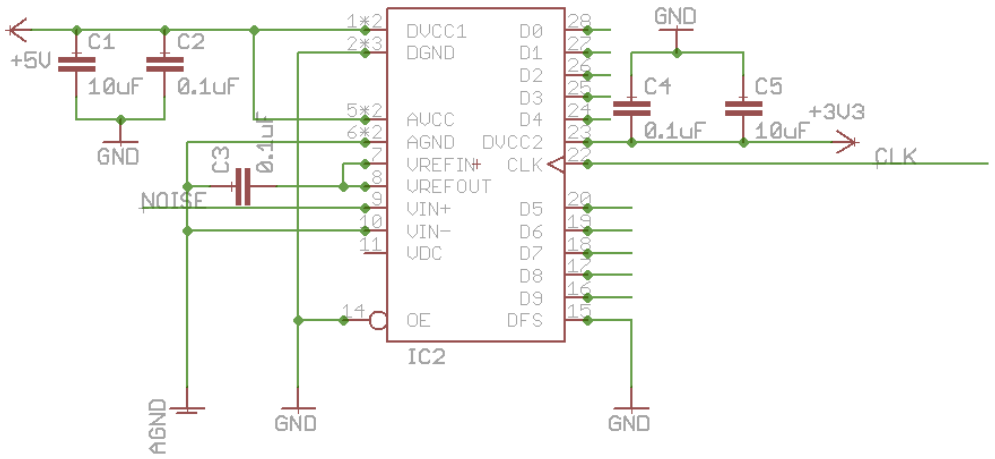


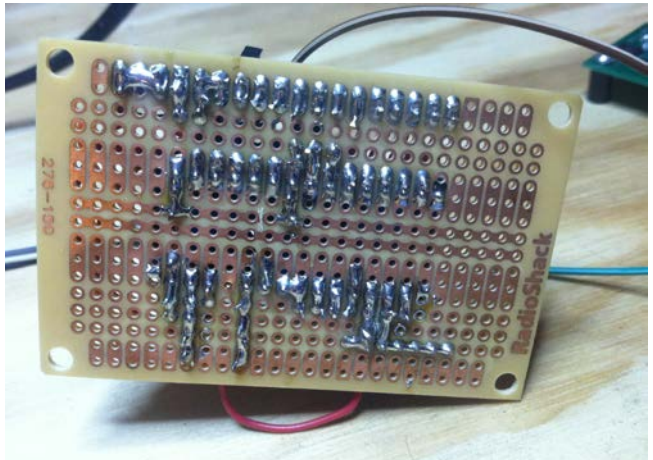


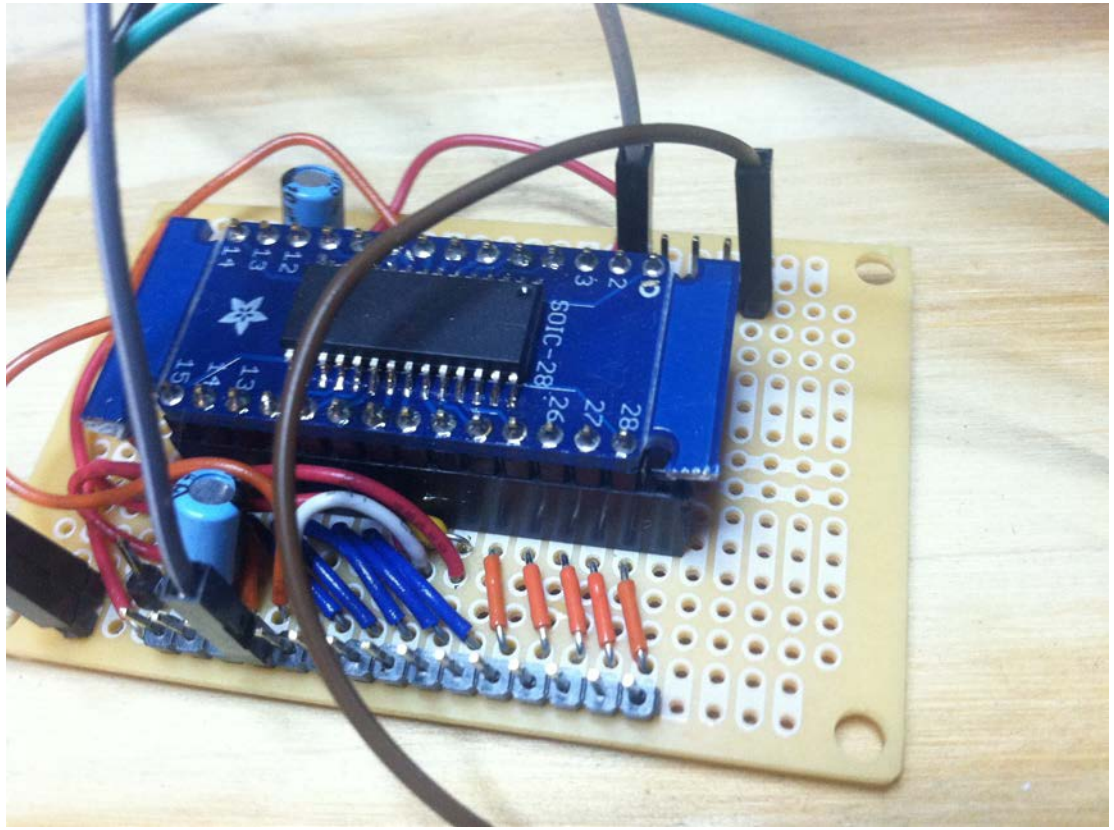

# Iteration – ADC – 3/3

- PROBLEM! Get rid of one amplifier b/c ADC expects 1V p-p
- Jerry-rigged initial PCB

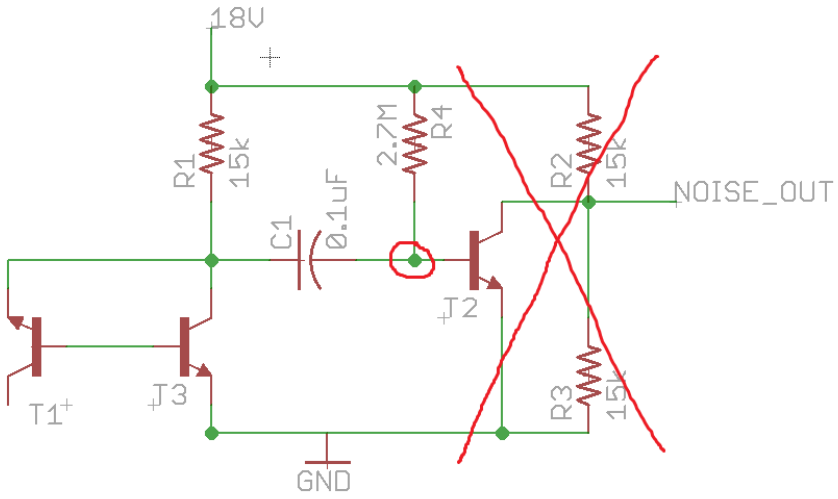


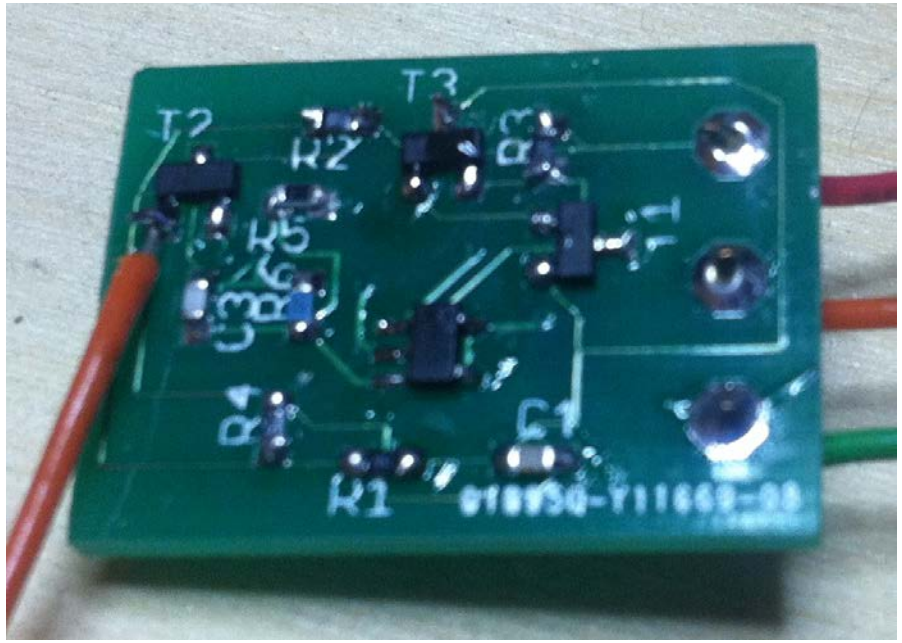

# Integration

- Hacked-up PCB (5V)
  - Physics mechanism
  - 5V -> 18V power supply
- ADC “breakout board” (5V)
- Olimex LPC1343 dev board (3.3V)
- Arduino uno (providing 5V)
- A whole mess of wires/problems
  - Running out of physical pins
  - Common ground
  - etc

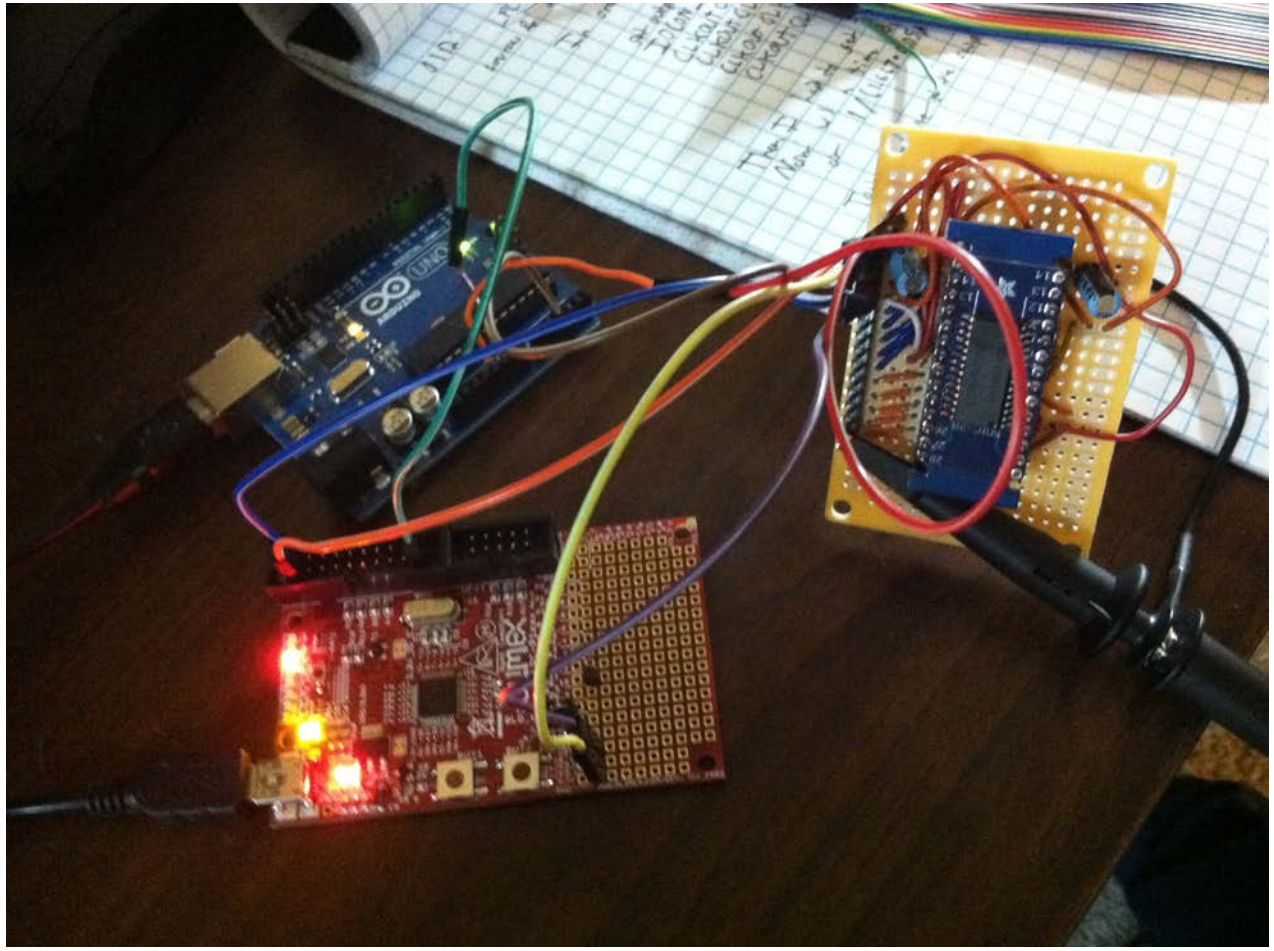

# Integration

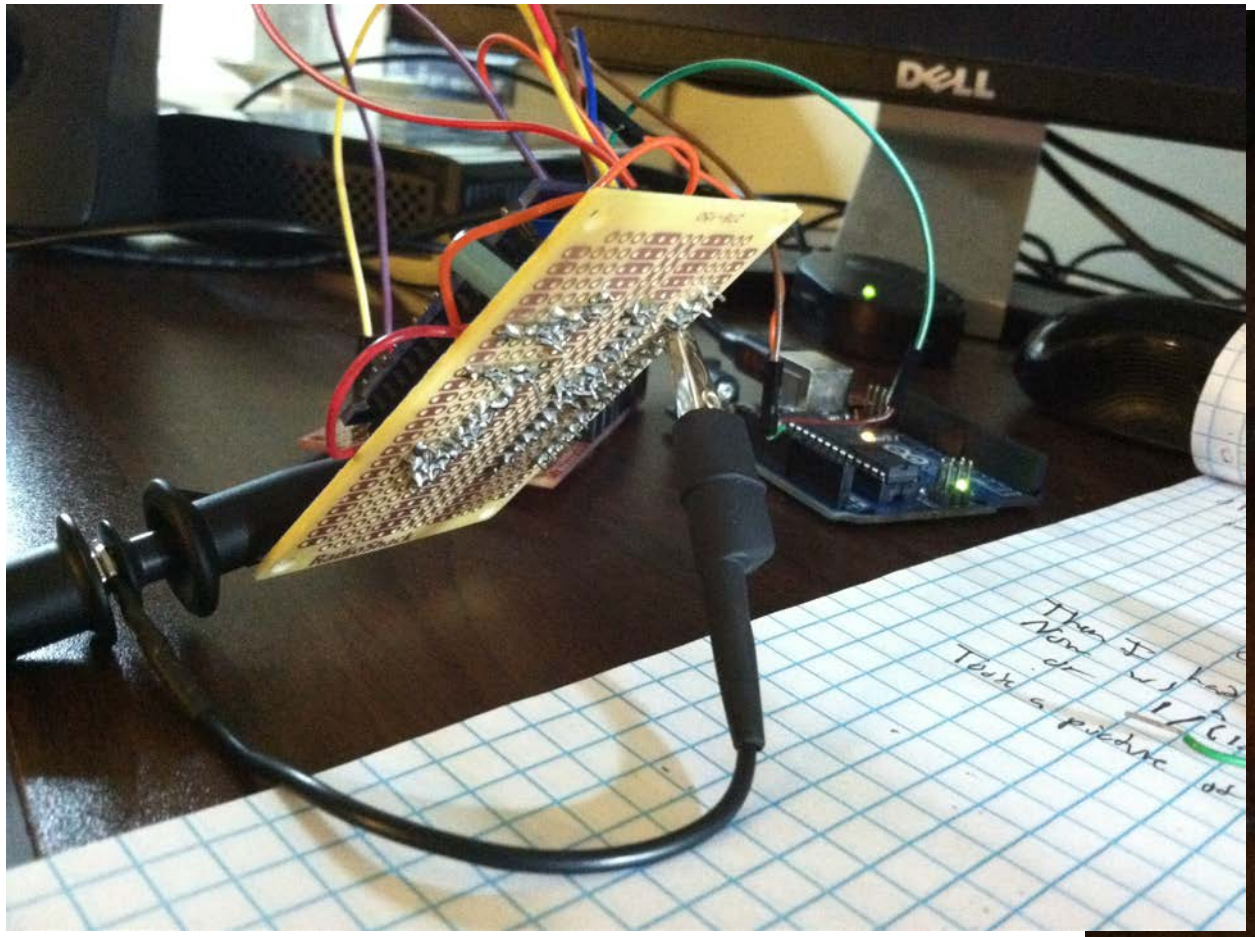


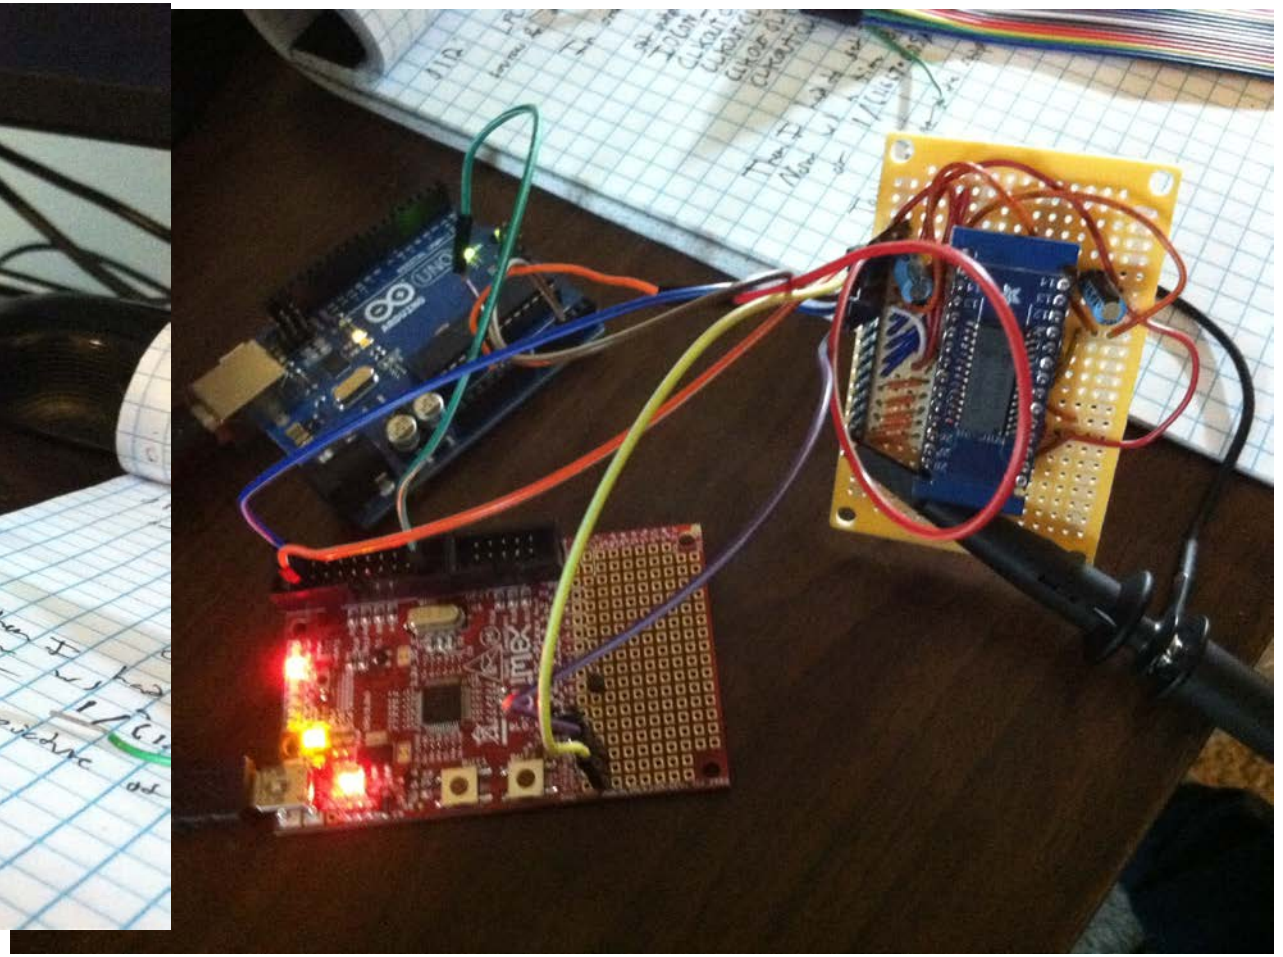

# Integration

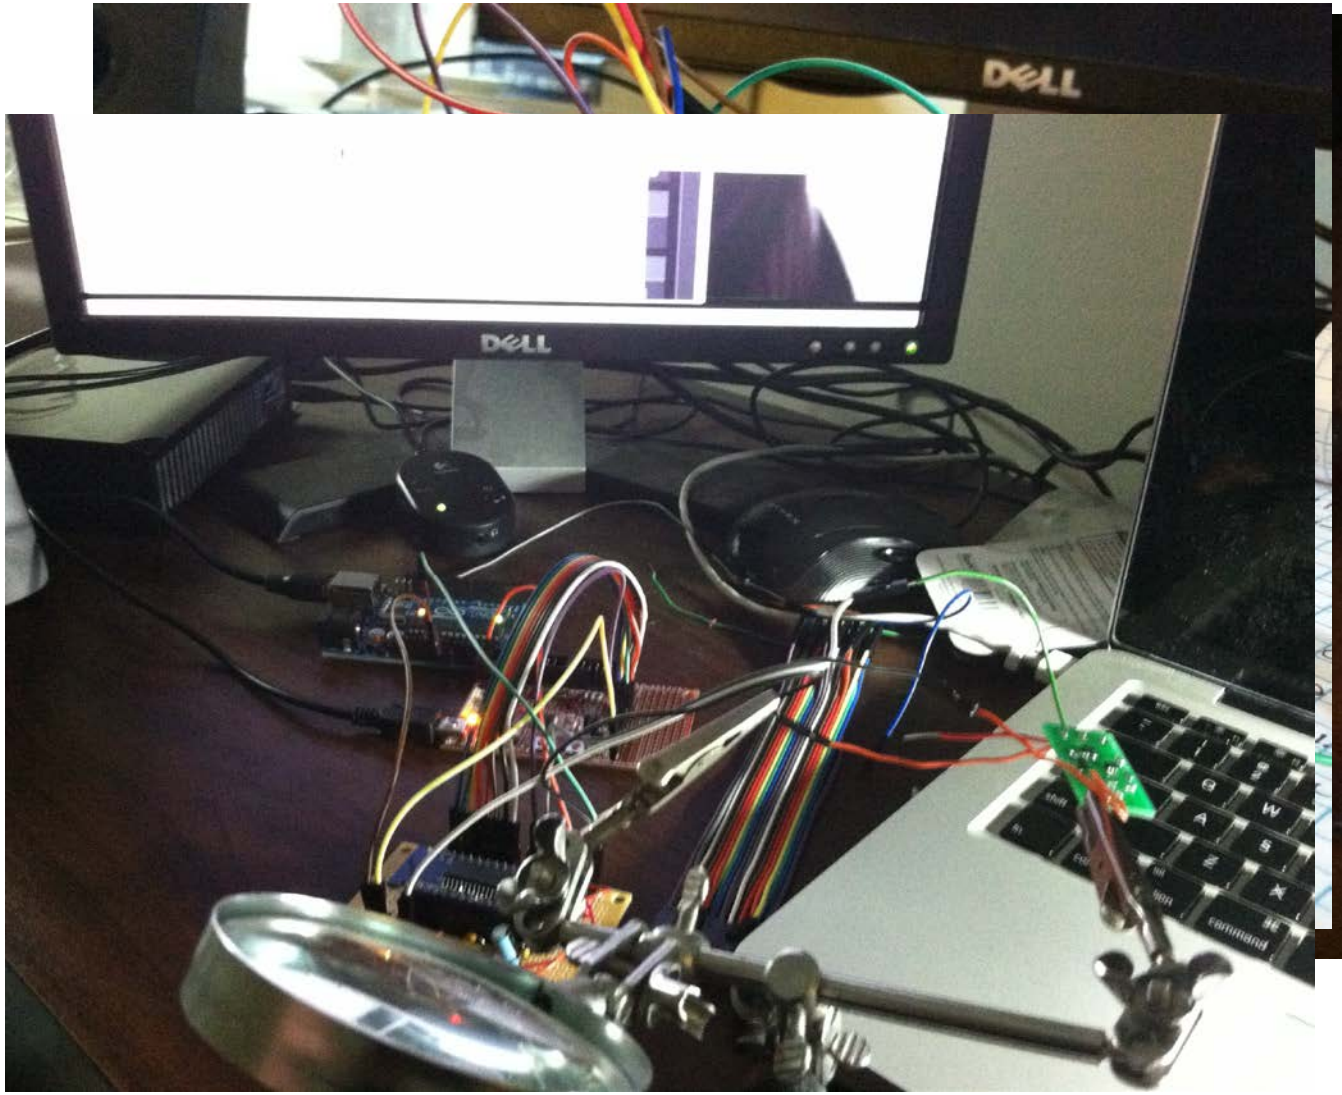


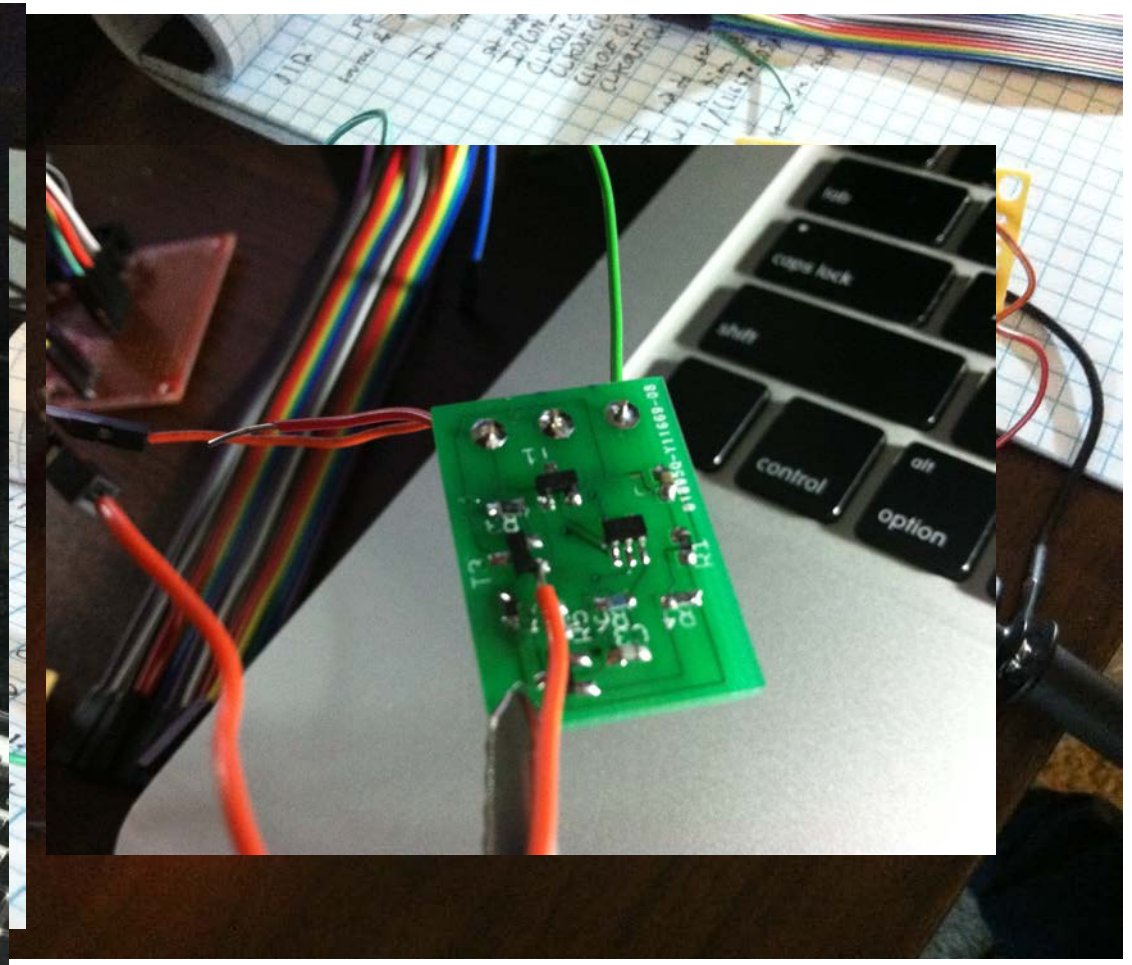

# Integration

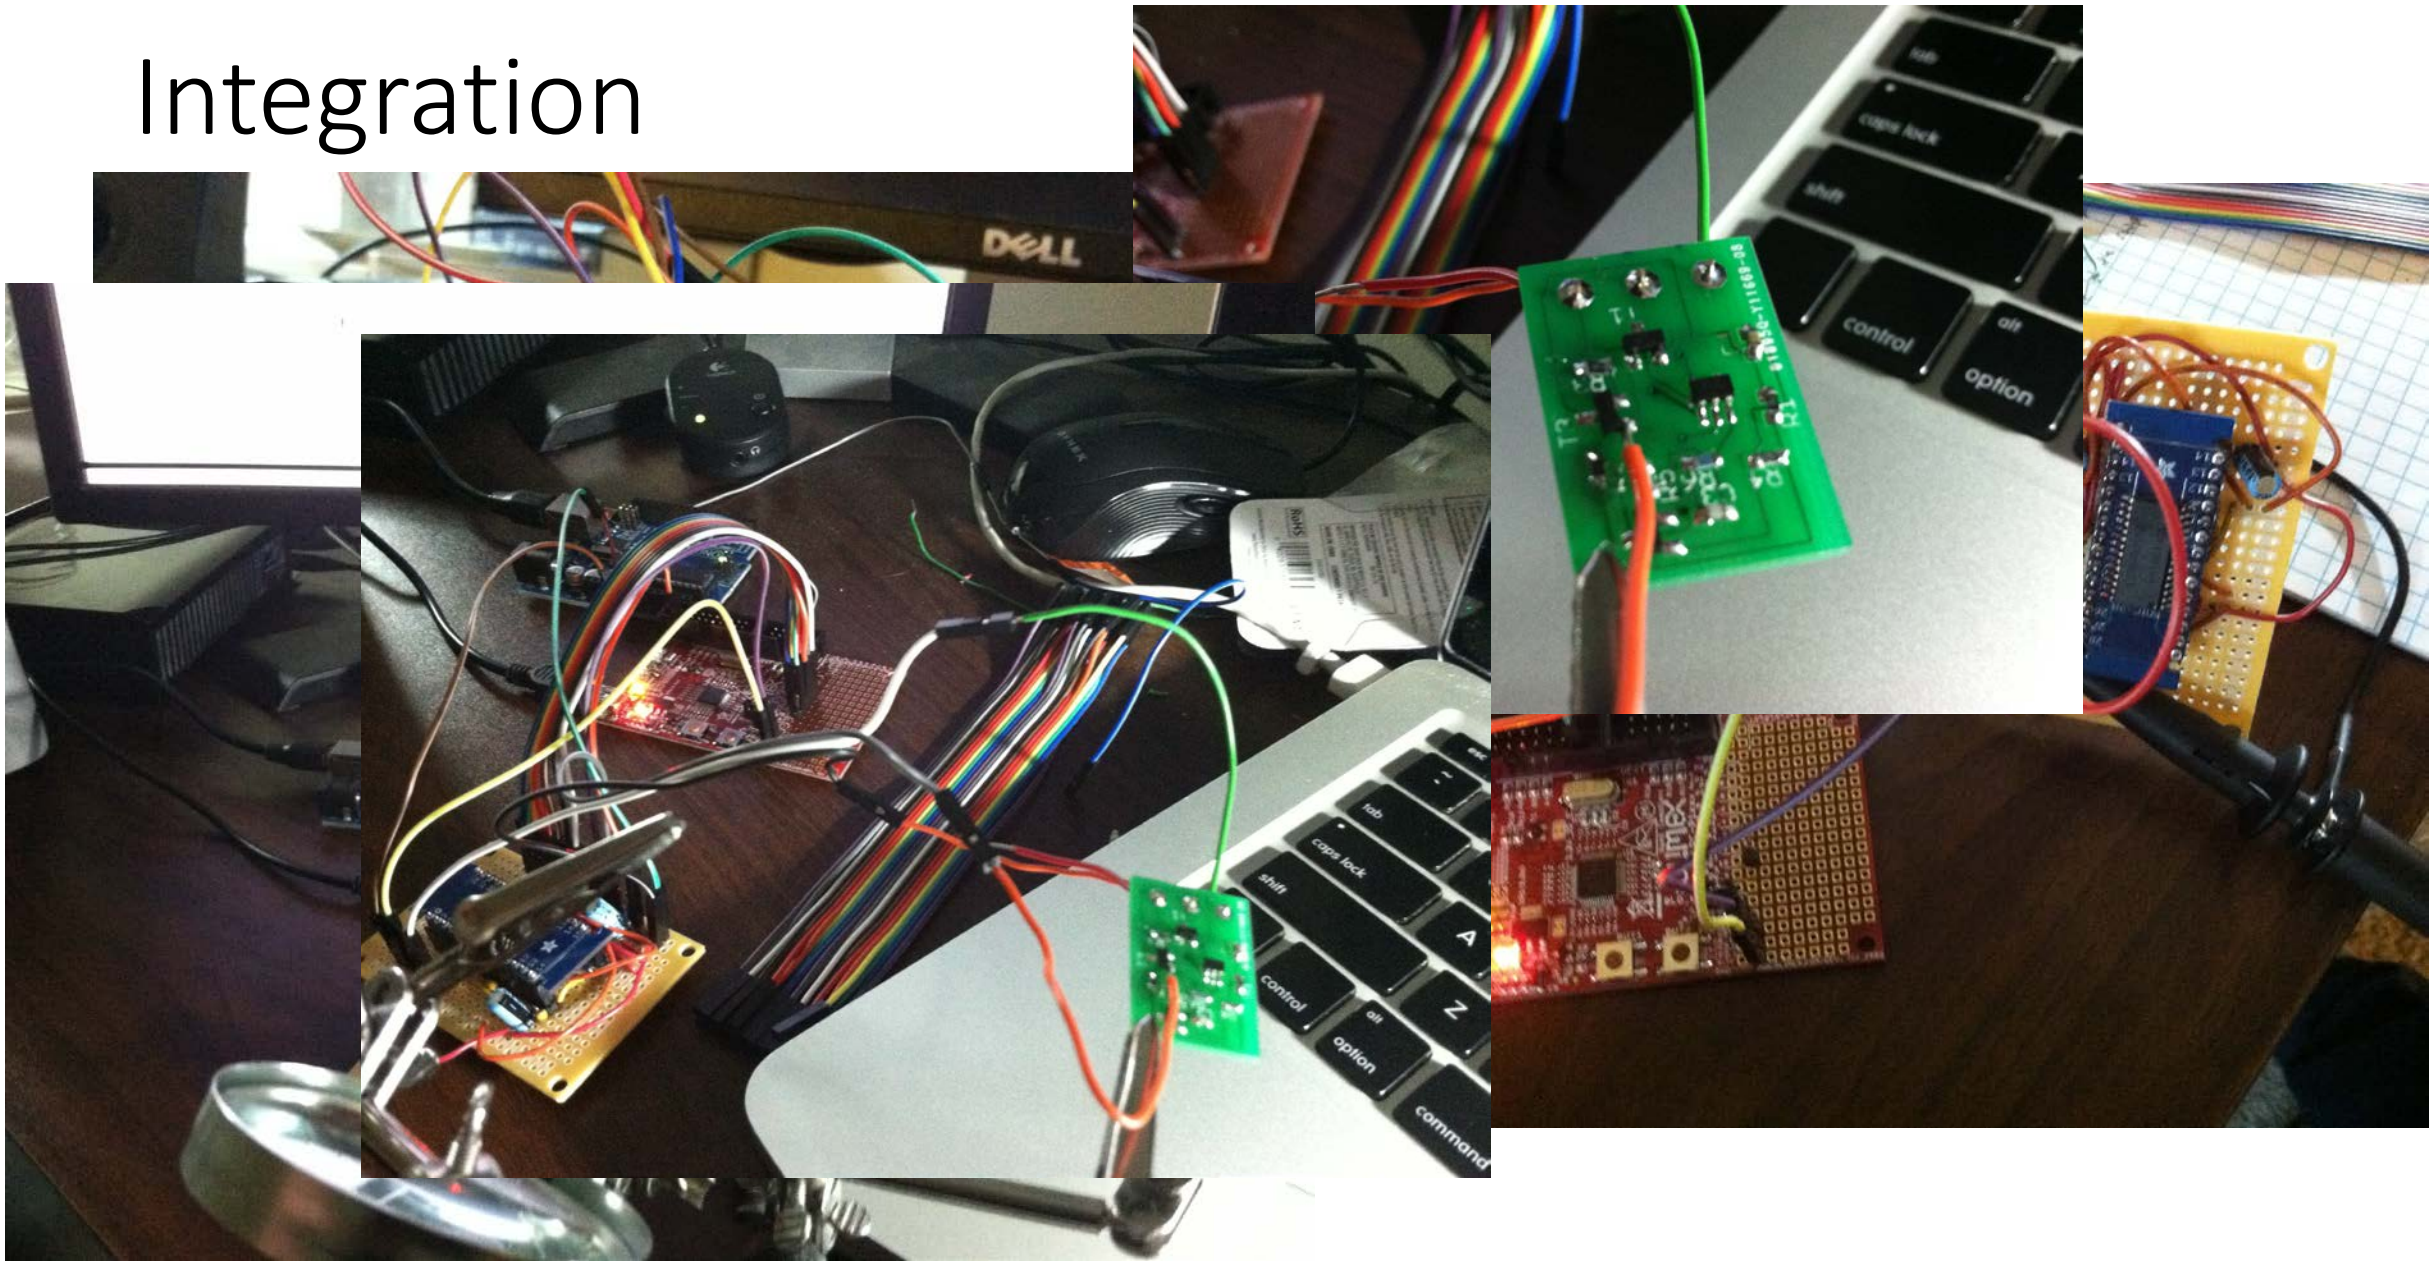

Inte

I AM BEAR OF VERY LITTLE BRAIN, AND MANY WIRES BOTHER ME

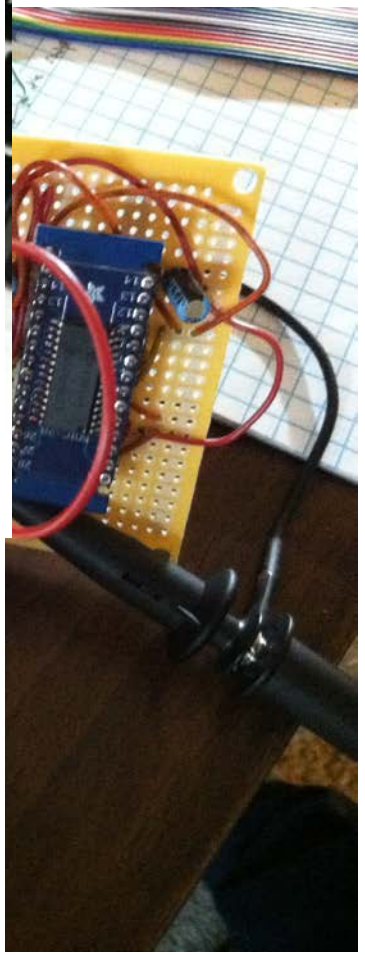

# Integration

- Second PCB: physics, PSU, ADC
- Still external LPC1343 dev board

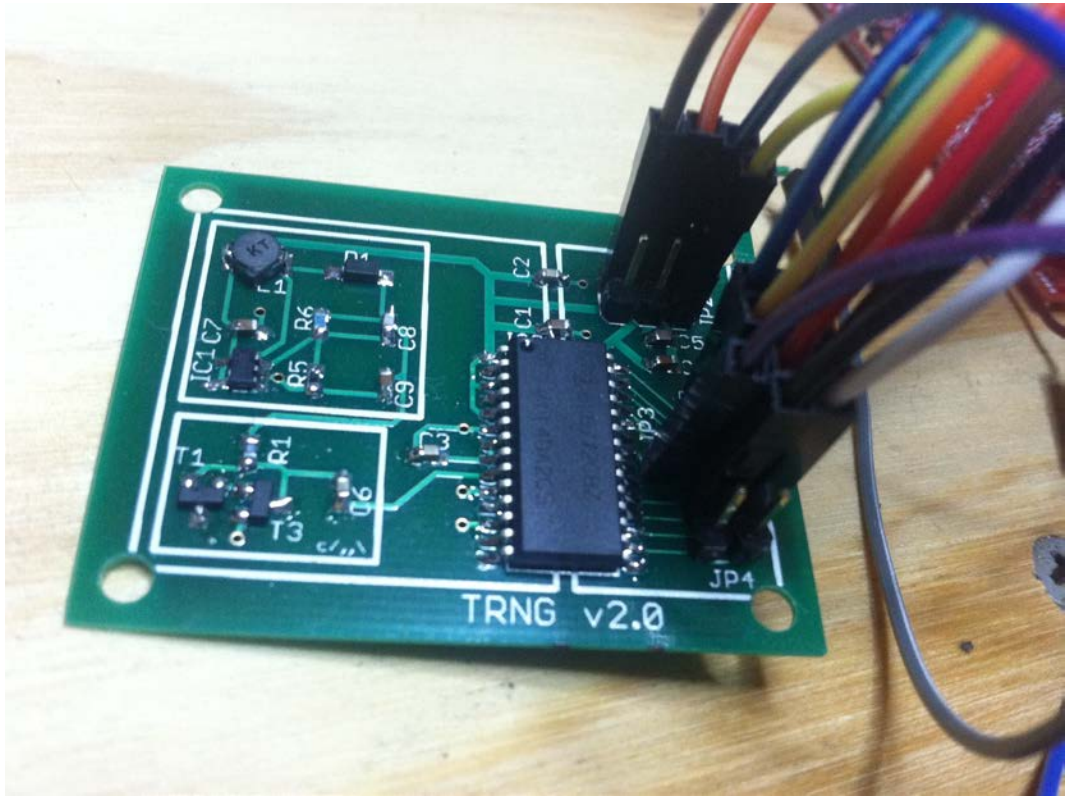


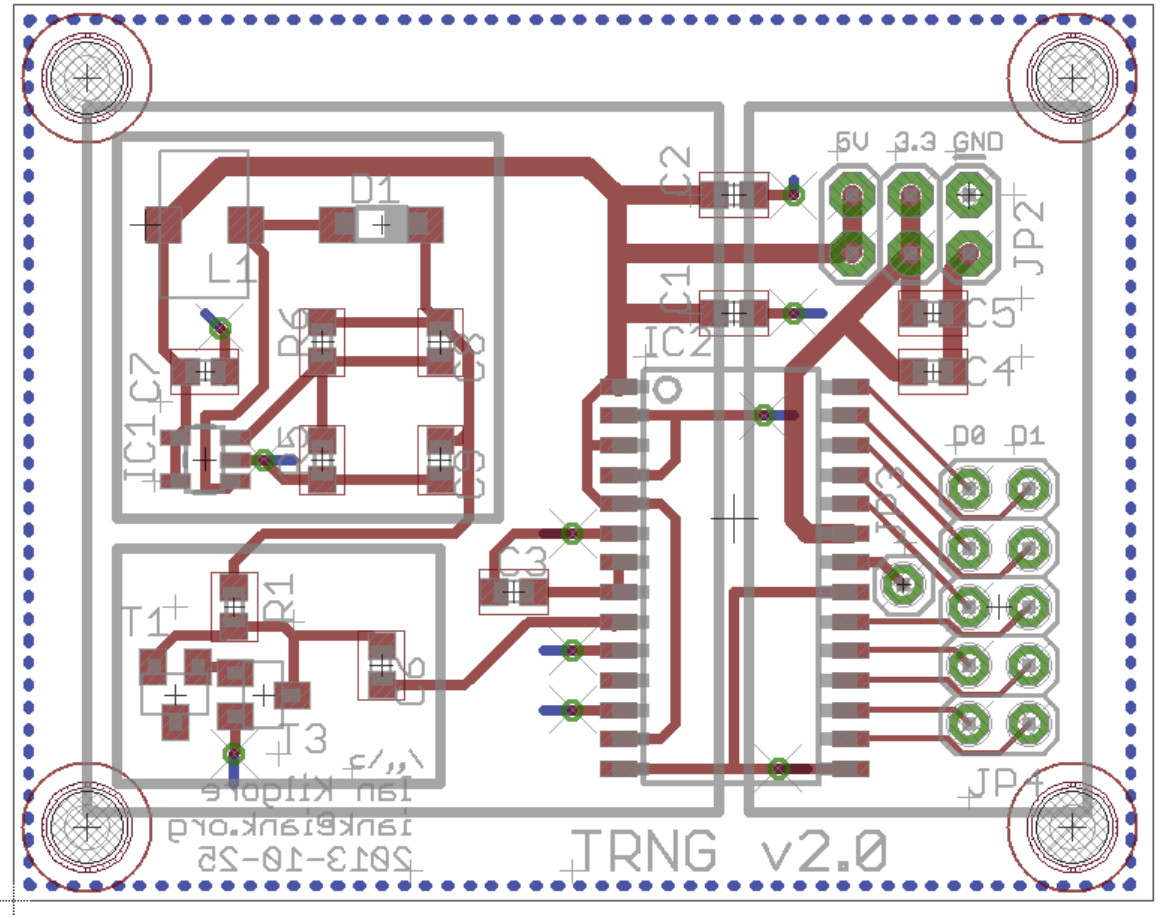

# Integration

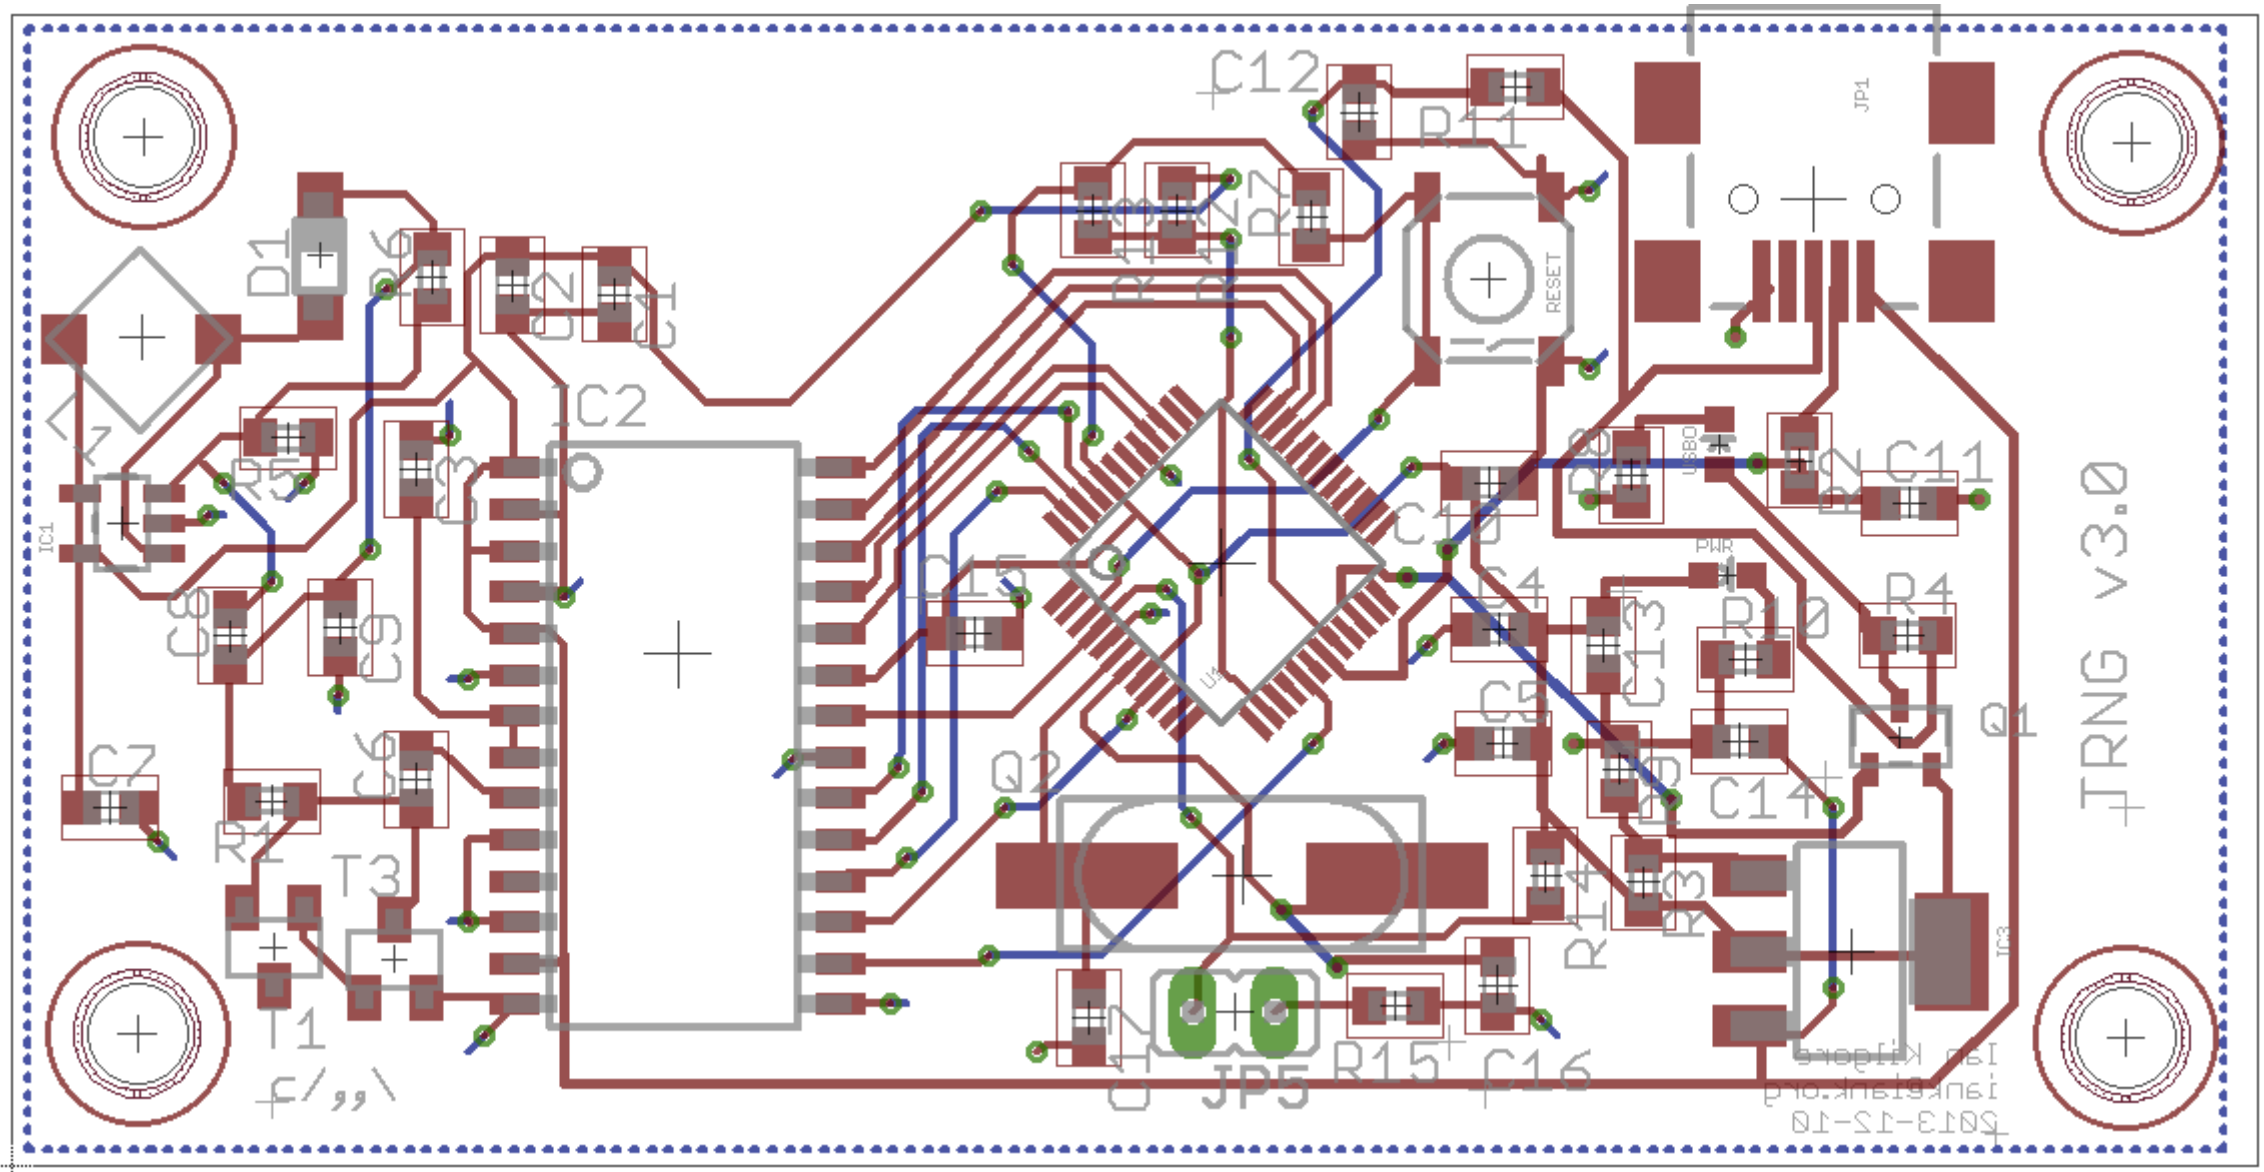

# Final PCB, Assembly

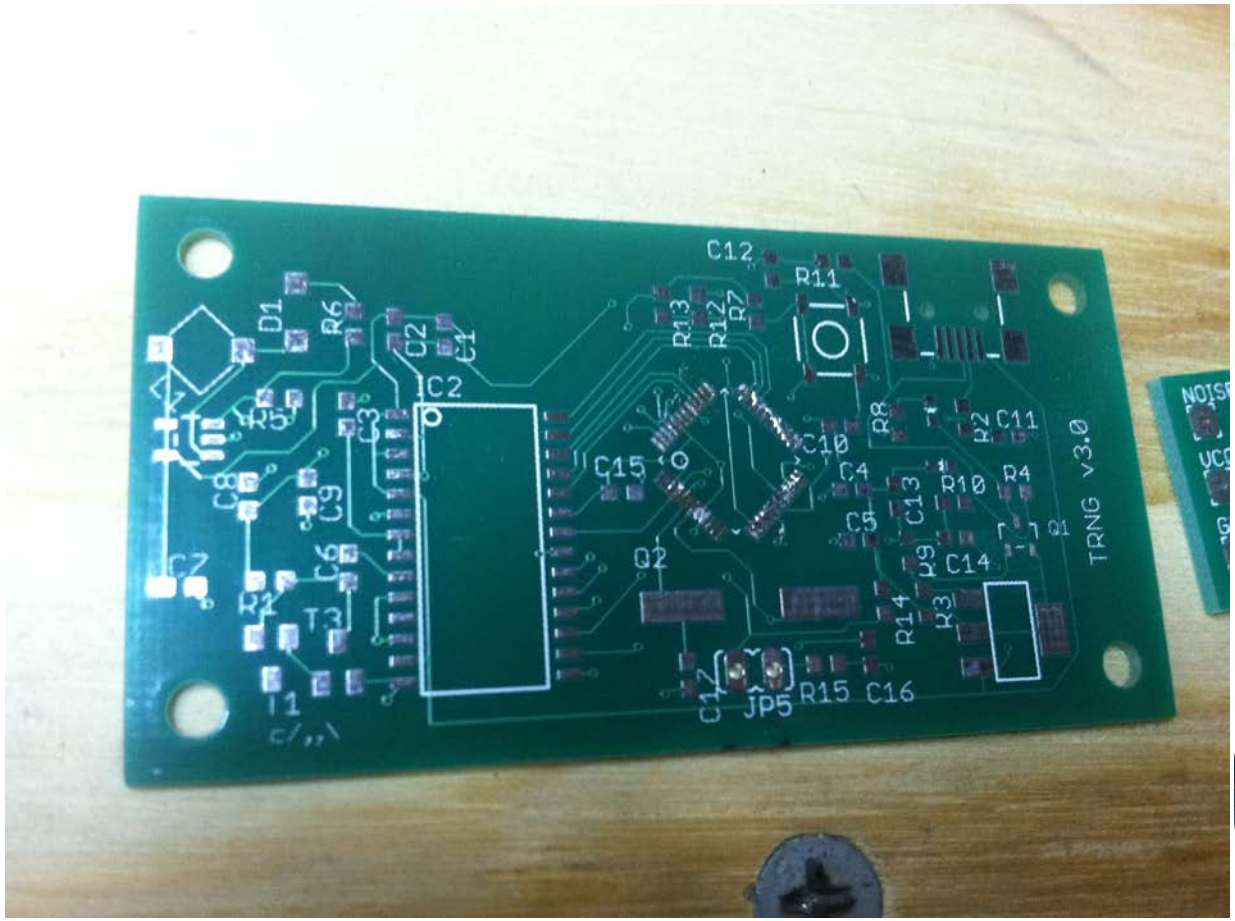


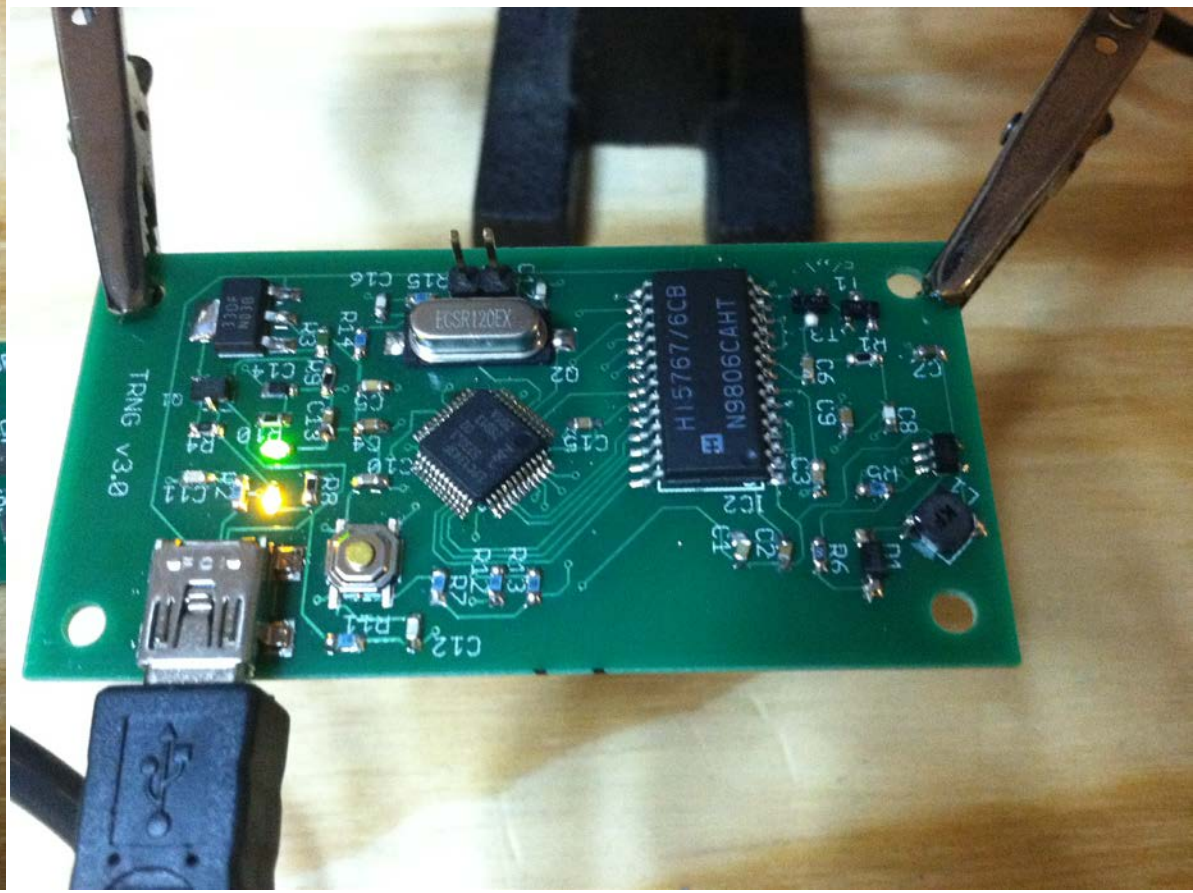

# Final PCB, A

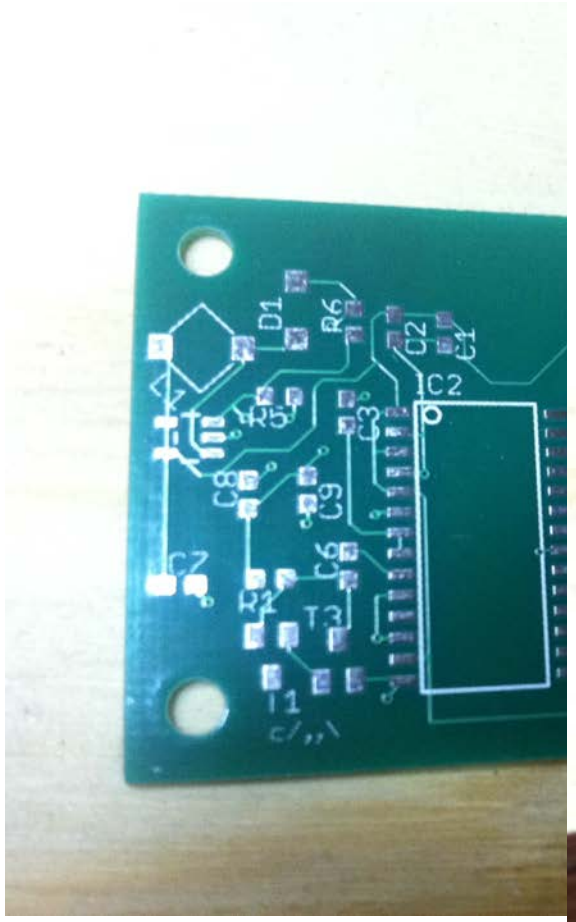

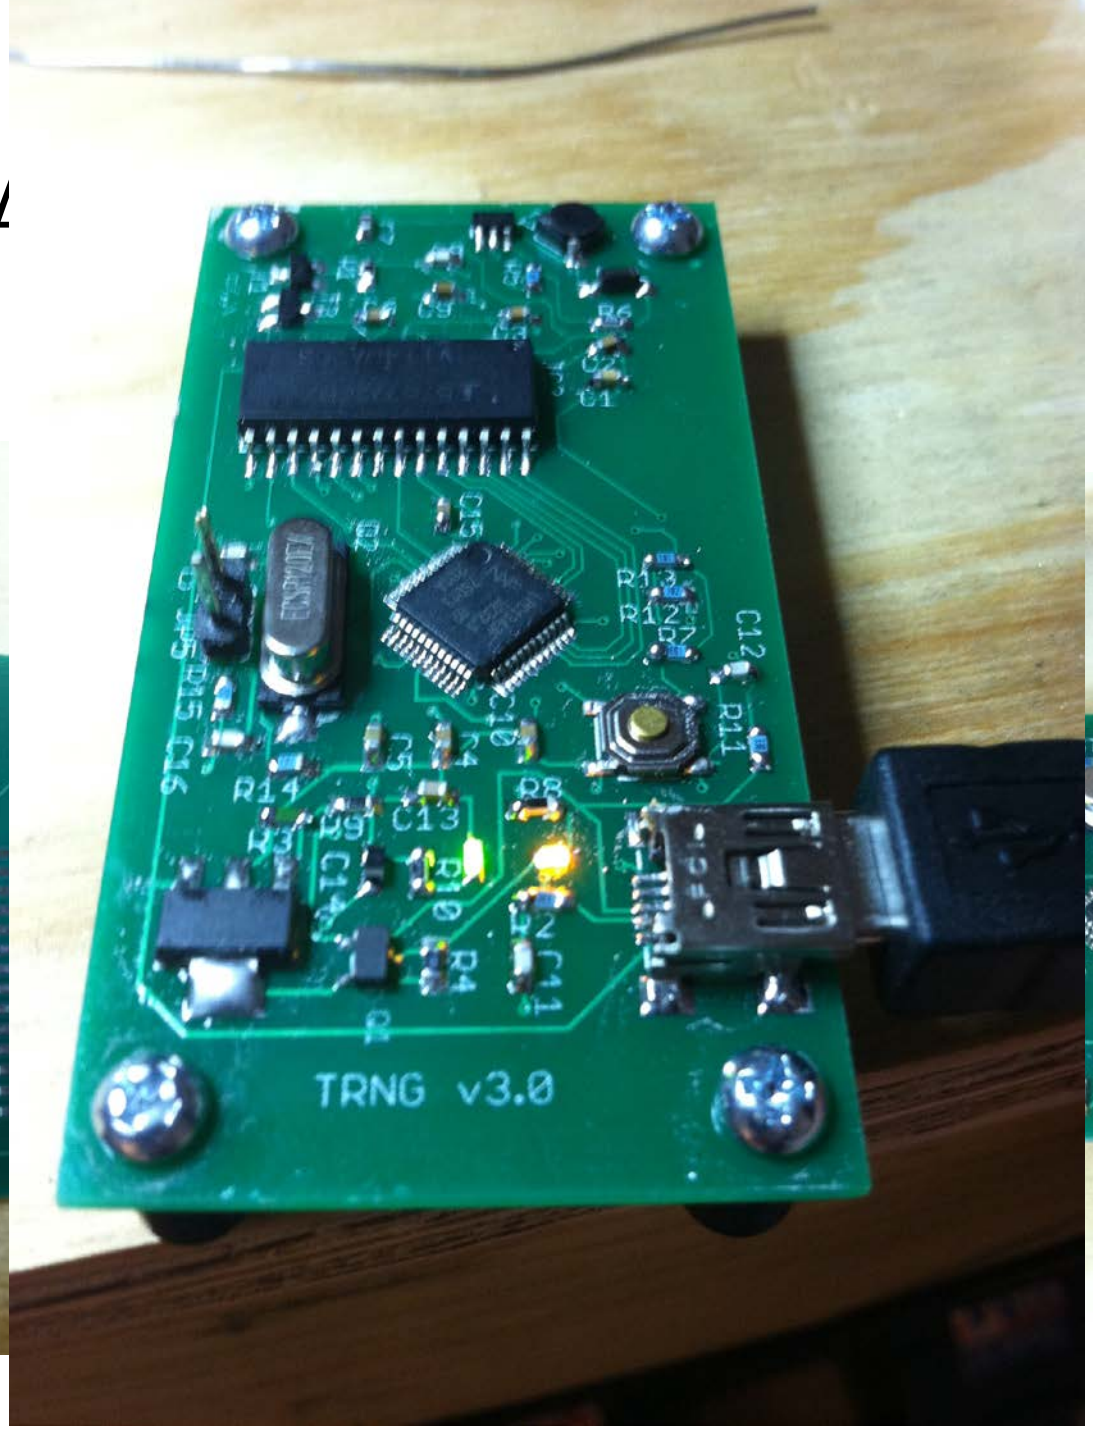


# Aside: Good tools are important

- Fine-pitch SMD components
- We've been well beyond $10 soldering "gun" territory for a while now
- Spend the $$ or get access to a lab/hackerspace/friend

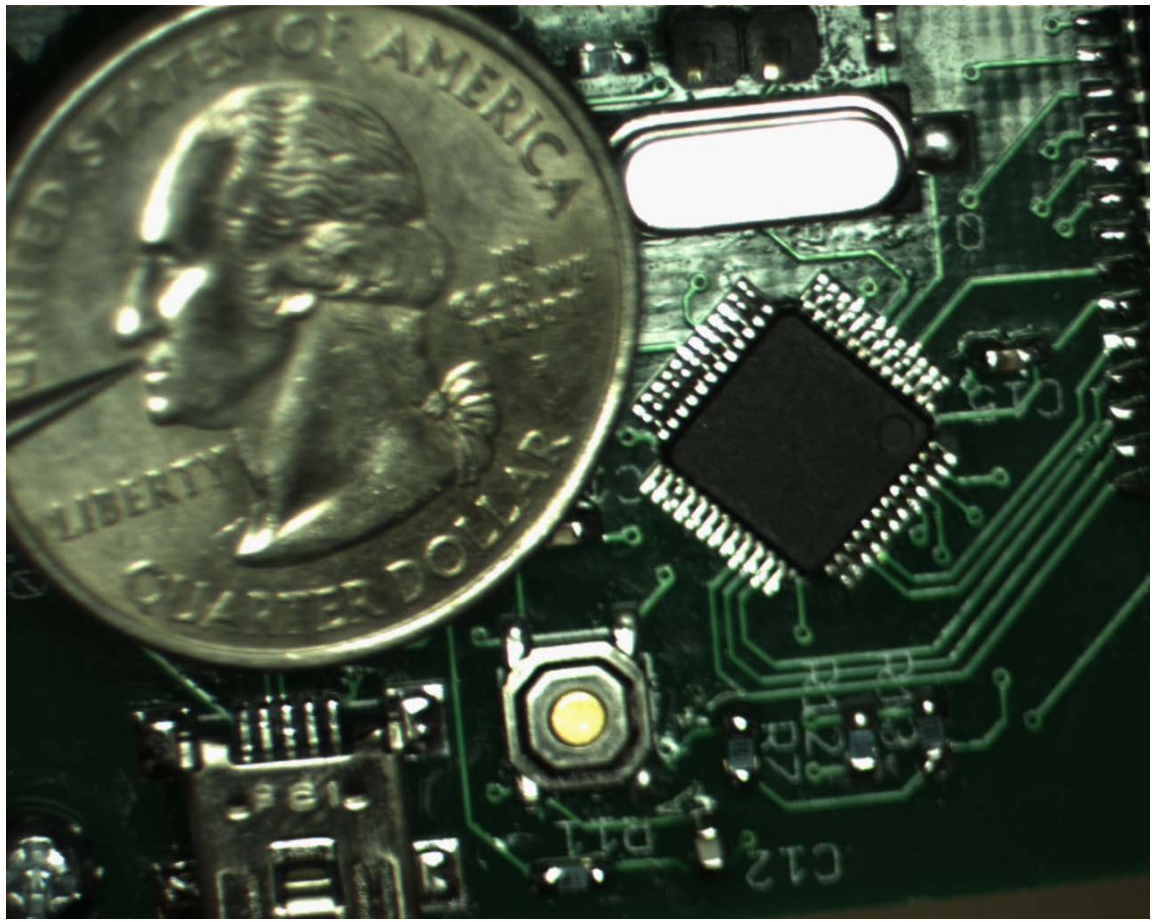
LIBERTY
QUARTER DOLLAR

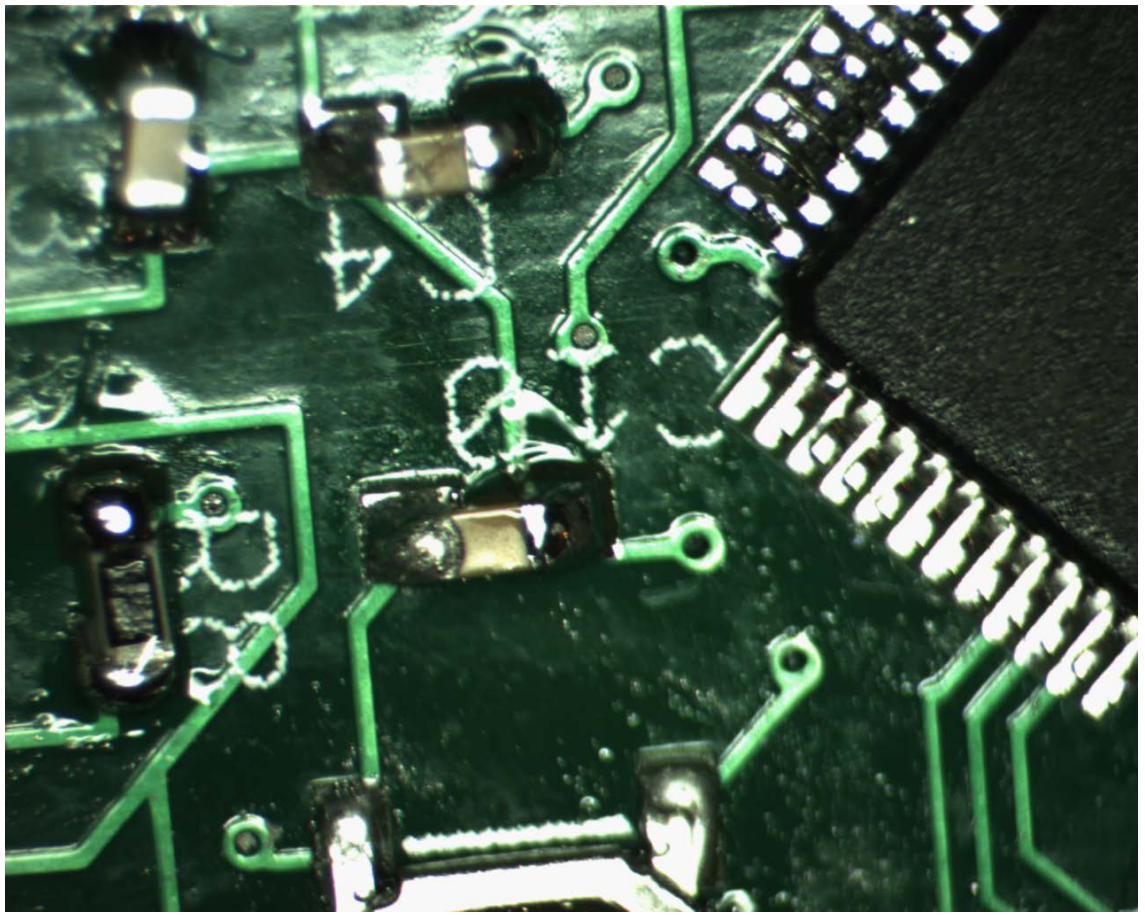

# Results

- Software whitening done in firmware
- Data rate reduced to 9 kilobytes / sec

http://iank.org/trng/2014-01-16_dieharder_test.txt

# Future Work

- Code for: I’m done with this now, but here’s some things I could have done incrementally better
- Reduce board area 30-50%
  - Reduce 0603 to 0402 or smaller
  - Reduce wasted space between components
- Increase ADC sample rate
- Implement whitening in CPLD

# Never use this


Photo © Tomás Fano CC BY-SA 2.0

- Not a differential circuit. Poor PSRR
- Easily-influenced externally
- In general don’t roll your own crypto, this applies to hardware too (especially?)
- There’s like a million of these, seriously. don’t use mine.

- Random is HARD and I got a lot of things wrong too
- Good news is it’s crypto-hard, which means nobody really cares

# Questions?

# Resources

- Detailed writeup (eventually) and these slides: http://iank.org/trng.html


- [1] Cognitive Daily, “Is 17 the ‘most random’ number?” by Dave Munger http://scienceblogs.com/cognitivedaily/2007/02/05/is-17-the-most-random-number/
- [2] http://en.wikipedia.org/wiki/Talk:Hardware_random_number_generator
- [3] http://www.nims.go.jp/research/organization/hdfqf1000000isjt-att/hdfqf1000000ispa.pdf
- [4] http://web.jfet.org/hw-rng.html
- [5] http://research.microsoft.com/en-us/people/mickens/thisworldofours.pdf
- [6] https://www.olimex.com/Products/ARM/NXP/LPC-P1343/resources/LPC-P1343-schematic.pdf
- [7] https://www.olimex.com/Products/ARM/NXP/LPC-P1343/
- [8] http://www.pokerstars.com/poker/rng/
- [9] http://www.stat.fsu.edu/pub/diehard/
- [10] http://www.phy.duke.edu/~rgb/General/dieharder.php
- [11] Silver, Nate. *The Signal and the Noise*. New York: Penguin Books Limited, 2012.
- [12] McIntyre, R. J. "Multiplication noise in uniform avalanche diodes." *Electron Devices, IEEE Transactions on* 13.1 (1966): 164-168.
- [13] Verwey, J.F.; Kramer, R.P.; De Maagt, B.J., "Mean free path of hot electrons at the surface of boron-doped silicon," *Journal of Applied Physics* , vol.46, no.6, pp.2612,2619, Jun 1975
doi: 10.1063/1.321938
- [14] http://fourier.eng.hmc.edu/e84/lectures/ch4/node3.html
- [15] http://eprint.iacr.org/2006/086.pdf
- [16] http://www.cs.berkeley.edu/~daw/papers/ddj-netscape.html
- [17] http://www.debian.org/security/2008/dsa-1571
- [18] http://www.intersil.com/content/dam/Intersil/documents/hi57/hi5767.pdf
- [19] http://gamesbyemail.com/News/DiceOMatic
- [20] http://www.random.org/analysis/
- [21] http://boallen.com/random-numbers.html
- [22] http://www.newscientist.com/article/mg21128264.900-lottery-wins-come-easy-if-you-can-spot-the-loopholes.html#.Uyi_HPldU3k
- [23] http://harpers.org/archive/2011/08/0083561