How to plan for continued device-level energy scaling

Erik P. DeBenedictis, Center for Computing Research, Sandia

SC 15, November 18, 2015

Approved for Unclassified Unlimited Release
SAND2016-3164C

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Overview

- Moore's Law will end when it becomes trendy for it to end
  - Because marketing forces can redefine Moore's Law
- The physics of where practical scaling ends for supercomputing is clear, but not the timetable
  - The development of the "millivolt switch" needed to continue scaling is considered inevitable
  - No time table is known, nor is the assurance we would be told if there were one
- Path forward
  - Maintain at least a "contingency plan" for both outcomes
  - Develop a endpoint design targets and develop software to those endpoints
    - We can project the design of computers at the end of scaling, but we don't know how many times we'll rewrite our software along the way

Current physics issue

- Chip power is CV^2f, where
  - C is capacitance between conductive sheets, simplifying to two sheets
  - V is supply voltage = signal voltage
  - $\frac{1}{2}CV^2$ is energy on a capacitor
  - f is frequency
- If the sheet has grooves, total capacitance hardly changes
  - Like wires
  - Line width doesn't matter
- If V doesn't change, we might expect f to be flat under scaling (which is what happens now)

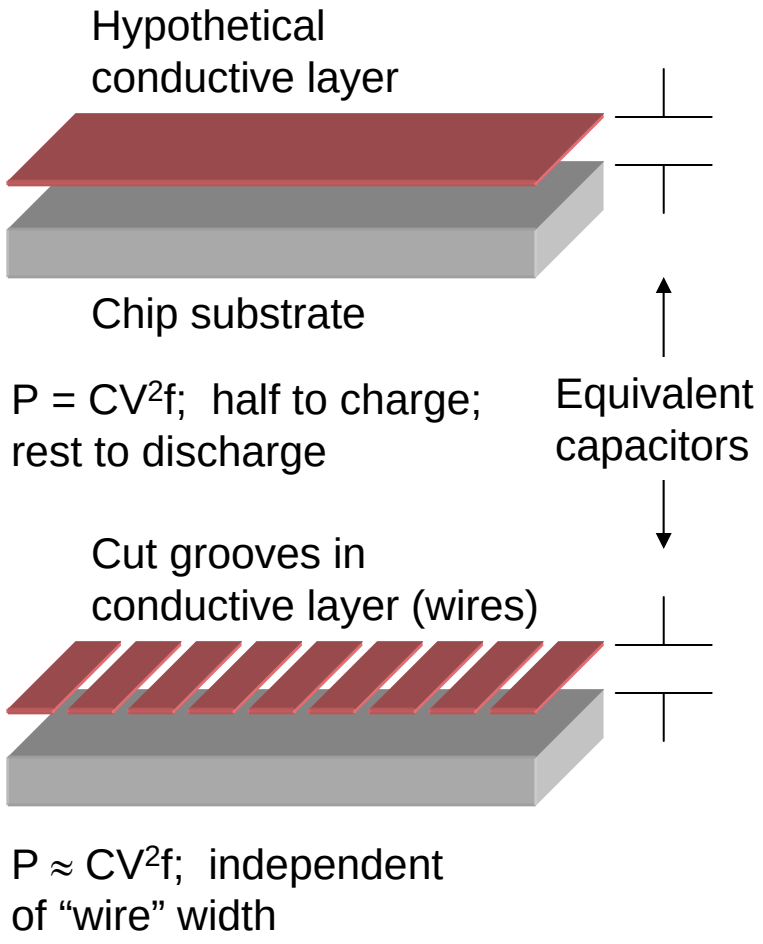

Reducing V

Not for MOSFET

- Barrier height = supply voltage
- Electrons energy distribution has an inverse-exponential tail
- Approximately proportion e^{-N} of electrons have energy >NkT
- As supply voltages lowers, leakage current grows exponentially IRRESPECTIVE OF LINEWIDTH

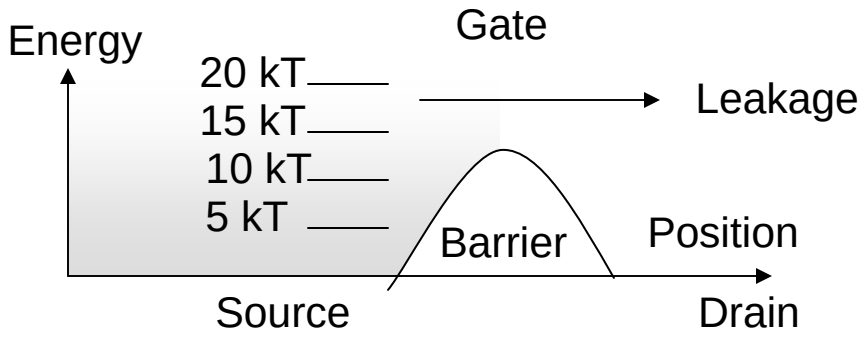

There are many other options

- Tunnel FET
- Negative capacitance on gate to boost gate voltage
- Use electron spin to shut off gate better
- Piezotronic transistor replacement (mechanical)
- Superconducting Josephson junctions

Roadmap for von Neumann architecture

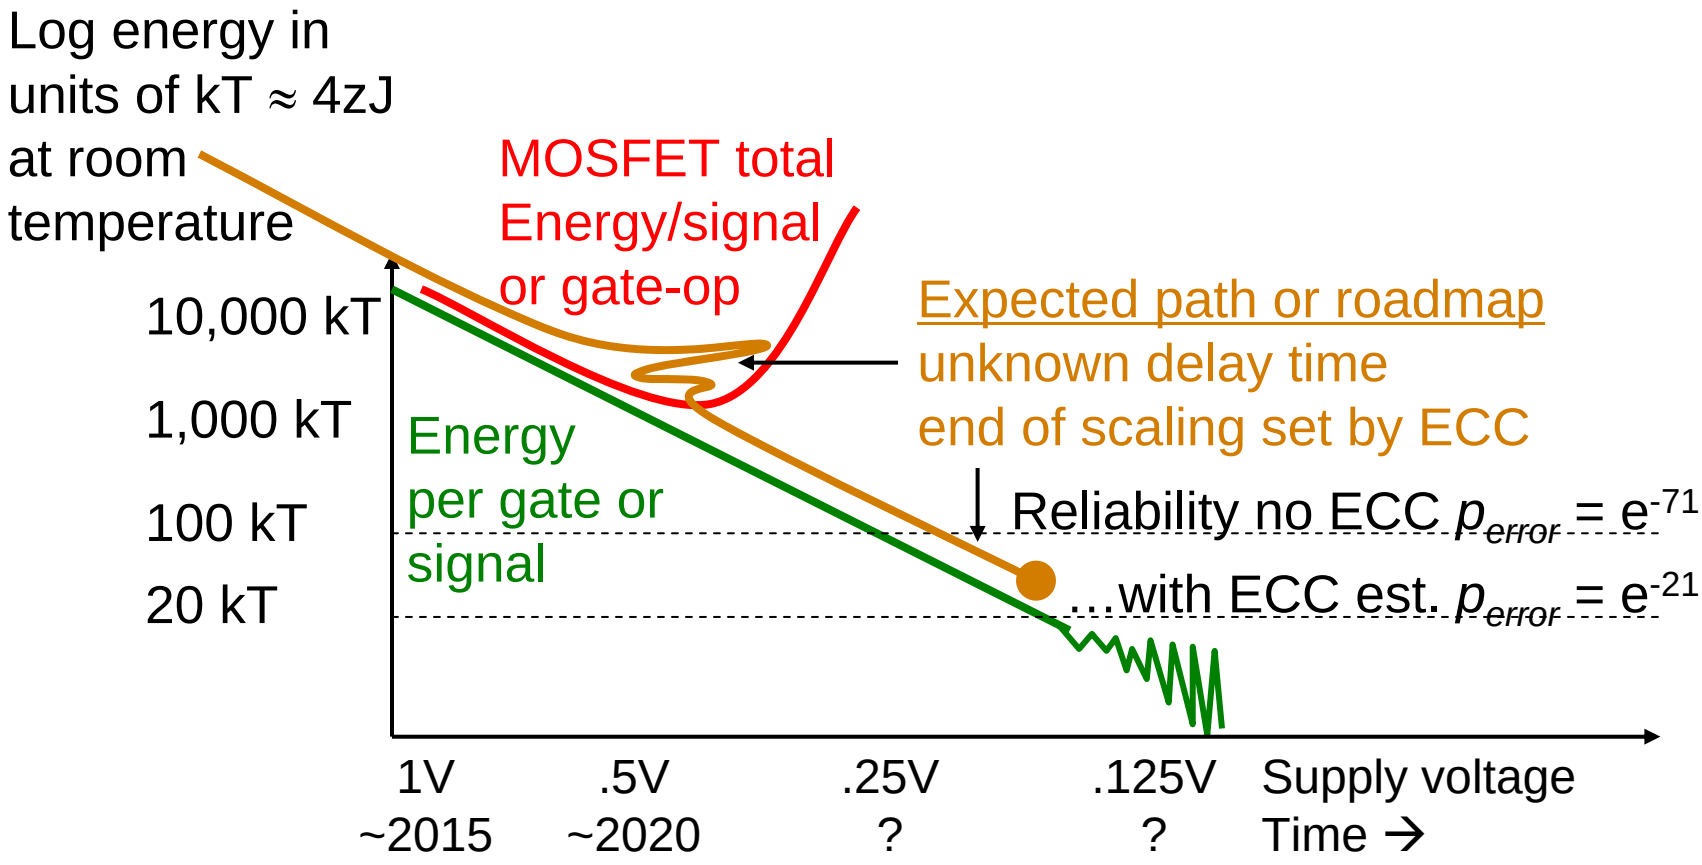

Need for error handling in semiconductor scaling

- Logic scaling has been connected quantitatively to redundancy and error correction
  - See →
  - See also Mike Frank
- We have queried the authors, but have not found
  - Examples of the needed error correction technique
  - A Turing-complete architecture

- Theis and Solomon*

1) *Conventional Logic:* Reduce the stored energy $(1/2)CV^2$. For conventional FETs, as V approaches a small multiple of kT/e, we must accept reduction in switching speed. New device concepts, discussed below, may allow more significant reduction in V and facilitate the reduction of stored energy towards kT. As thermal voltage fluctuations become significant, we must incorporate redundancy and error correction in the logic to keep the error rate in bounds. Refrigeration can reduce T but in a power-constrained environ-

ment with Meindl and Davis. Since Johnson–Nyquist voltage noise is Gaussian with a standard deviation of V_n, a stored logic voltage of m standard deviations, or a stored energy of m^2kT, would be needed to achieve a reliability of $(1/2)\mathrm{Erfc}[m/\sqrt{2}]$. (Eight standard deviations give an error probability of $\sim 10^{-15}$.)

Note that,

$p_{error} = \frac{1}{2}\mathrm{Erfc}[m/\sqrt{2}] \approx \exp(-E_{signal} / kT)$

*Theis, Thomas N., and Paul M. Solomon. "In Quest of the" Next Switch": Prospects for Greatly Reduced Power Dissipation in a Successor to the Silicon Field-Effect Transistor." *Proceedings of the IEEE* 98.12 (2010): 2005-2014.

End of Moore's Law in 1970s (for DRAM)

- Radiation induced soft errors grew more troublesome with increased DRAM size
  - This is essentially the Cosmic ray bit flip problem we have today
  - It was going to stop Moore's Law somewhere around 16K DRAMS
  - However, it wasn't actually Cosmic Rays, but radioactive impurities in the solder creating an alpha particle
  - Solved by a coating

- (Oops, false alarm)

High K dielectric early 2000s

- Moore's Law was going to end (for DRAM) due to inability to control leakage current
  - Solution was high-K dielectric
  - Solution was competitive between fabs, so nobody talked about their work in advance
  - Messed up planning

- (False denial there was a solution until last minute)

Memristors succeed Flash

- It’s been just a couple years away for a long time

- (Overoptimistic projections)
- (Entrenched competitor Flash improved)

Tunnel FET

- Expected to be a millivolt switch due to sharper turn on/off
- Unfortunately, manufacturing imperfections (surface states) cause lowering of gain

- (Looks tough)

Proposed plan I

- Inevitably, a new switch will be invented
  - Support and monitor R&D
  - There could be competitive pressure to keep it secret
- Should start a scaling process for energy efficiency improvement – possibly increased speed
- Good for 1-2 orders of magnitude – not enough to change the world, but enough to shift competitive landscape
  - Quite possible scaling will continue thermal limit, at which point error correction could be used for more scaling
- The new switch could be superconducting Josephson junction

Proposed plan II

- It ought to be possible to make technology projections based on two endpoints
- Design software to scale to the two endpoints
  - Avoid rewriting all the world's code more times than necessary

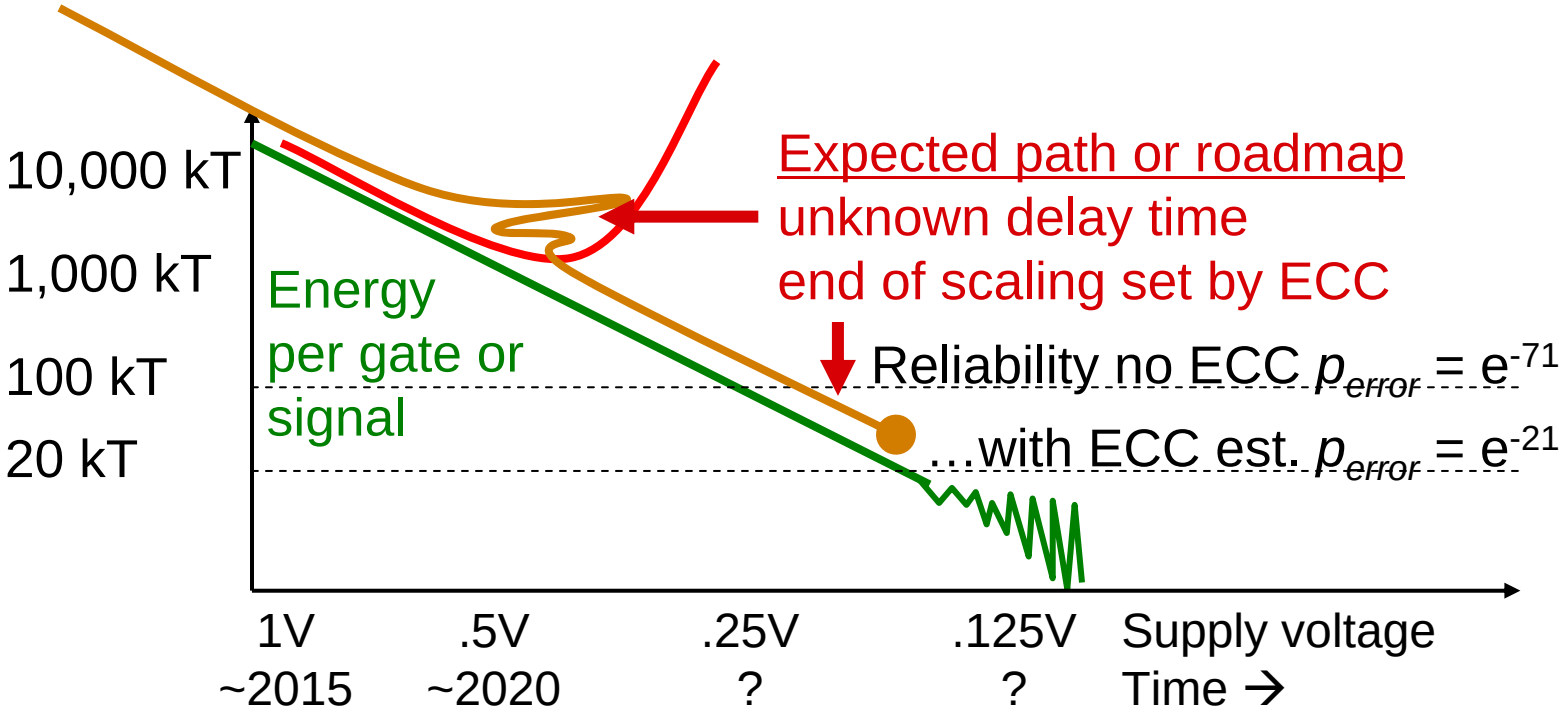

Backup: ECC boost about one generation

- Scaling will not stop abruptly, but it will be stopped by an exponential rise in error rate with declining energy
- But how much energy efficiency improvement is possible if we can tolerate errors? Spreadsheet →
  - No ECC 71 kT
  - ECC scenarios 24 kT – 28 kT
  - 2:1 after overhead, +/-
- A trillion dollar question

Exascale reliability requirement

100000	Gates-ops per floating point op where an error would cause a wrong answer
1.00E+18	ops/second (definition of Exascale)
60	seconds per minue
60	minutes per hour
24	hours per day
365	days per year
3	years for a computer's lifetime (before it becomes obsolete)
9.46E+30	number of gate operations per lifetime where an error would cause a wrong answer
71.33211	If we have Esignal equal this many kT's, error rate will be inverse of previous line

Say an operation is this many gate-ops	1000	20000	1.00E+05	1.00E+06	1.00E+18
Steps in lifetime (serial and parallel)	9.46E+27	4.73E+26	9.46E+25	9.46E+24	9.46E+12
RRNS using system in Watson and Hastings					
Gate ops per residue (four non-redundant residue	250	5000	25000	250000	2.5E+17
perror target for exaflops over lifetime	1	1	1	1	1
perror per step	1.06E-28	2.11E-27	1.06E-26	1.06E-25	1.06E-13
perror per residue; 3 errors in a step must go und	7.02E-09	1.91E-08	3.26E-08	7.02E-08	7.02E-04
Es = this many kTs will meet reliability in line abov	24.30	26.30	27.37	28.90	47.33
Energy savings	2.94	2.71	2.61	2.47	1.51
However, we need 6 total residues, not 4	1.96	1.81	1.74	1.65	1.00

Additional beneficial factors
Fixes Cosmic Ray hits
Fixes weak and aging components
Could support overclocking; i. e. catches an "excessive overclocking" error