Circuit- and System-Driven Requirements for Digital Logic Devices

Elad Alon

with Hei Kam, Yue Lu, Tsu-Jae King Liu

University of California, Berkeley

A Reminder

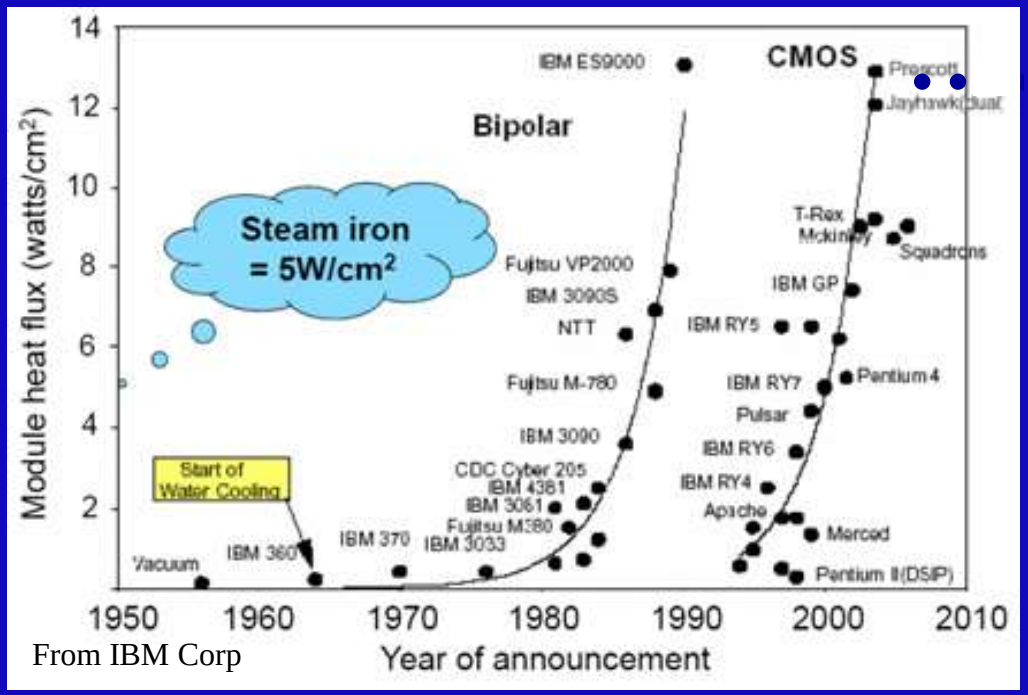

From IBM Corp

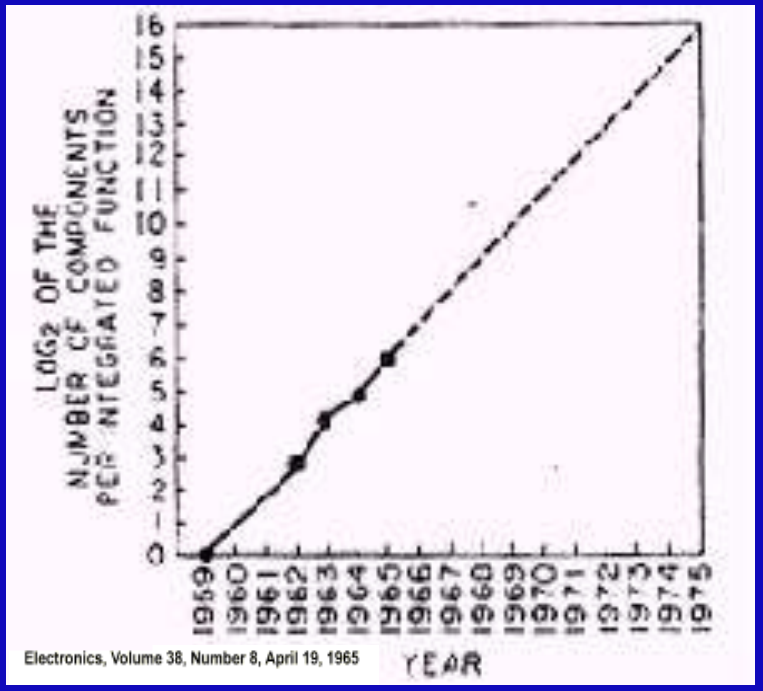

- **CMOS has hit the power wall**
  - **Non-scaling of kT/q and hence V_{dd}, V_{th}**
- **But, can't forget that reducing cost is still the underlying imperative**

My Assumptions (Implied by Cost)

- Stick to binary digital logic
  - Replacing full design/software stack generally too expensive
  - Devices need to have gain (noise margins)
- New device has yield and reliability comparable to CMOS
  - Today: chip with ~5 billion devices works for ~5 years
- New device has lower circuit/system-level energy over some range of performance and area…
  - Can translate all of these back to $

What Digital Chips Look Like

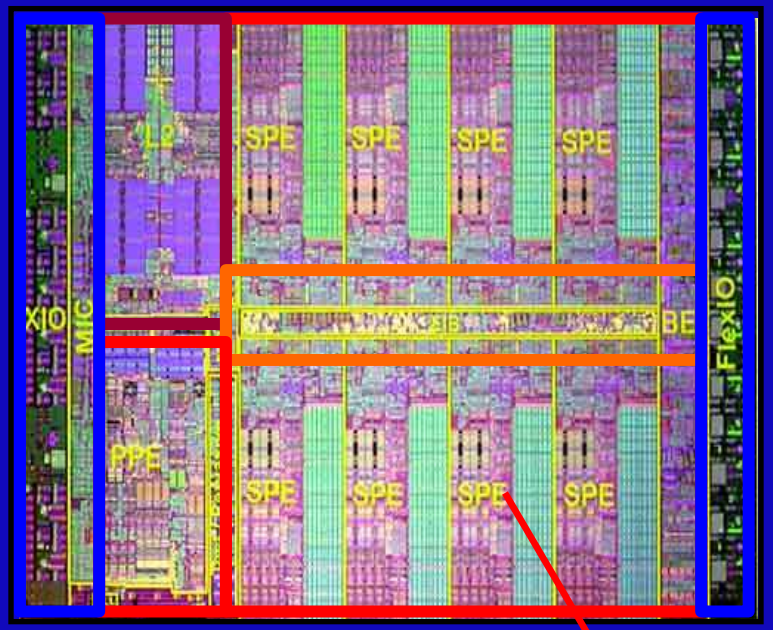

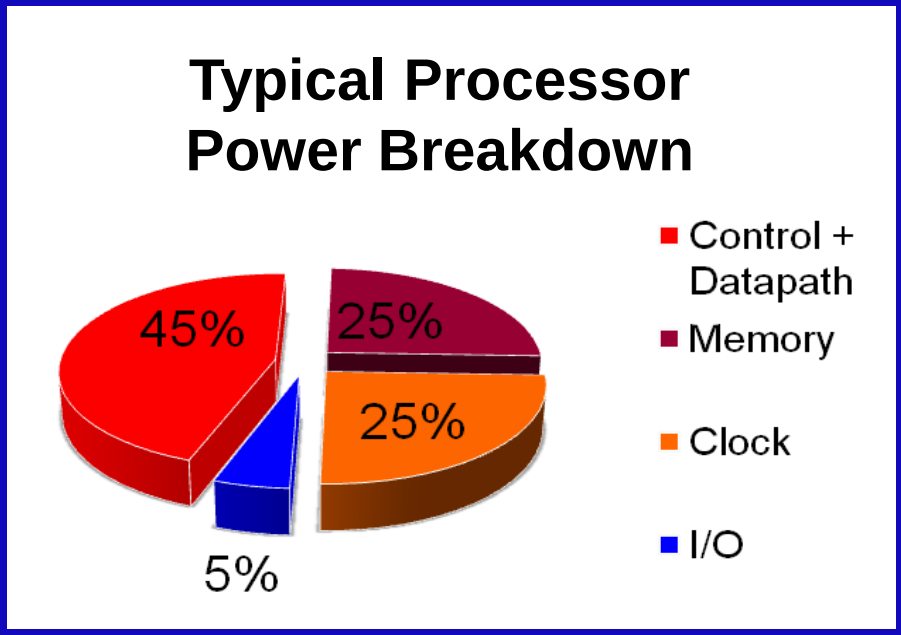

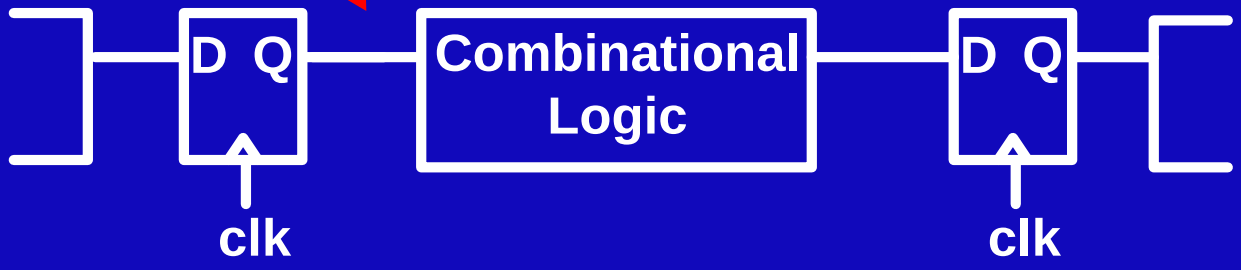

- Chip energy/perf. tracks datapath/control
- Clock frequency set by delay through CL

Logic Energy and Delay

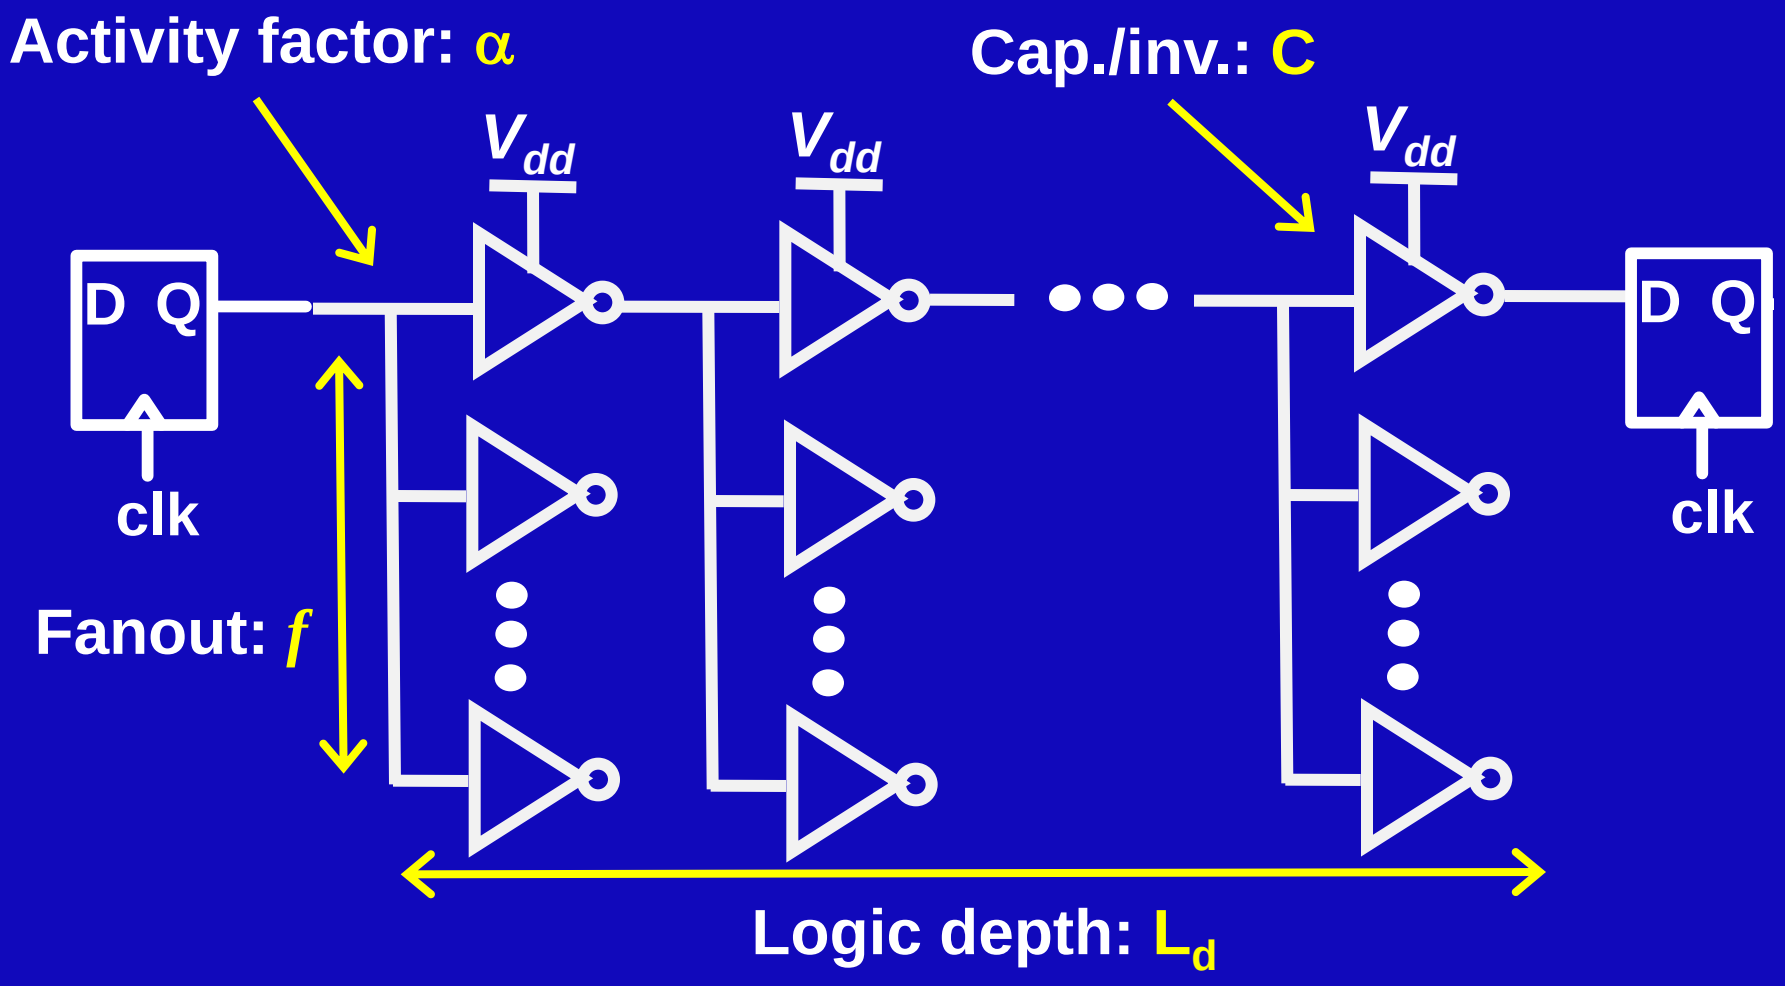

Logic Energy and Delay

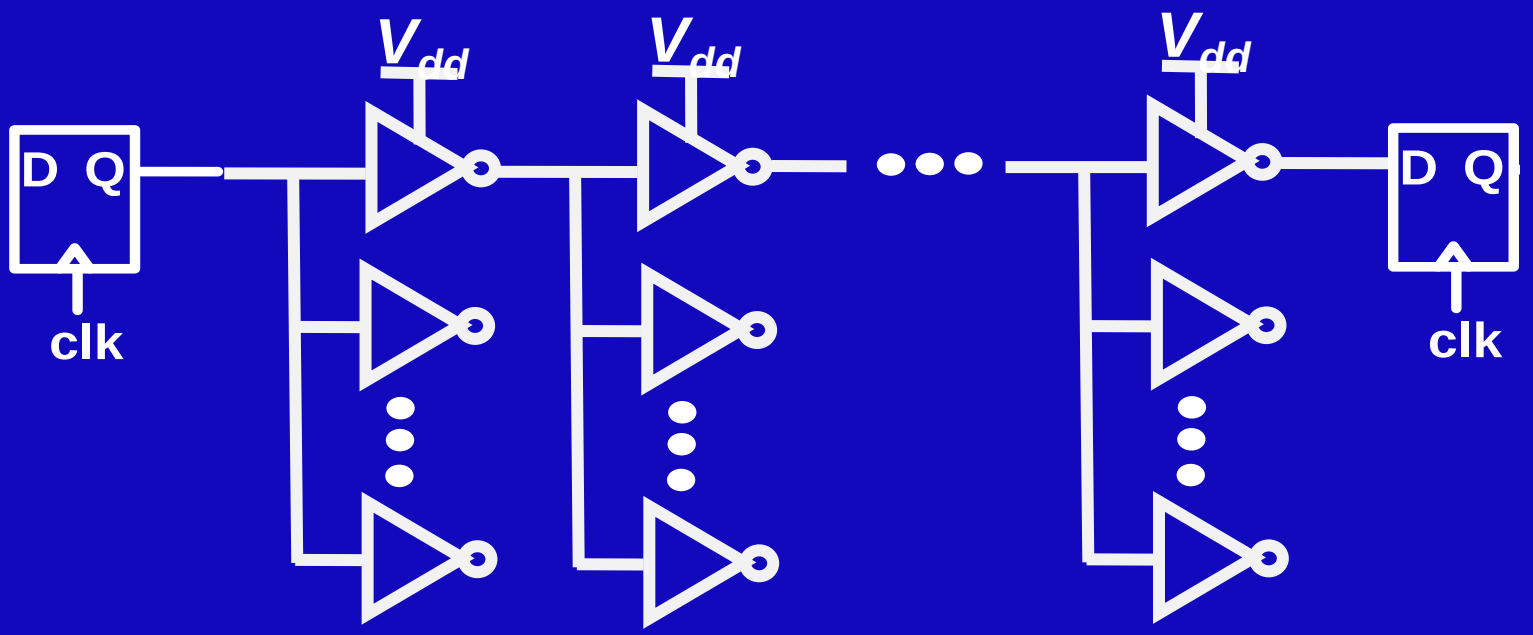

- $t_{delay} = L_d f CV_{dd}/(2I_{on})$
- $E_{dyn} + E_{leak} = \alpha L_d f CV_{dd}^2 + L_d f I_{off} V_{dd} t_{delay}$

- $E_{dyn} + E_{leak} = \alpha L_d f CV_{dd}^2(1 + (L_d f /2\alpha) / (I_{on}/I_{off}))$

Implications on Required I_{on}/I_{off}

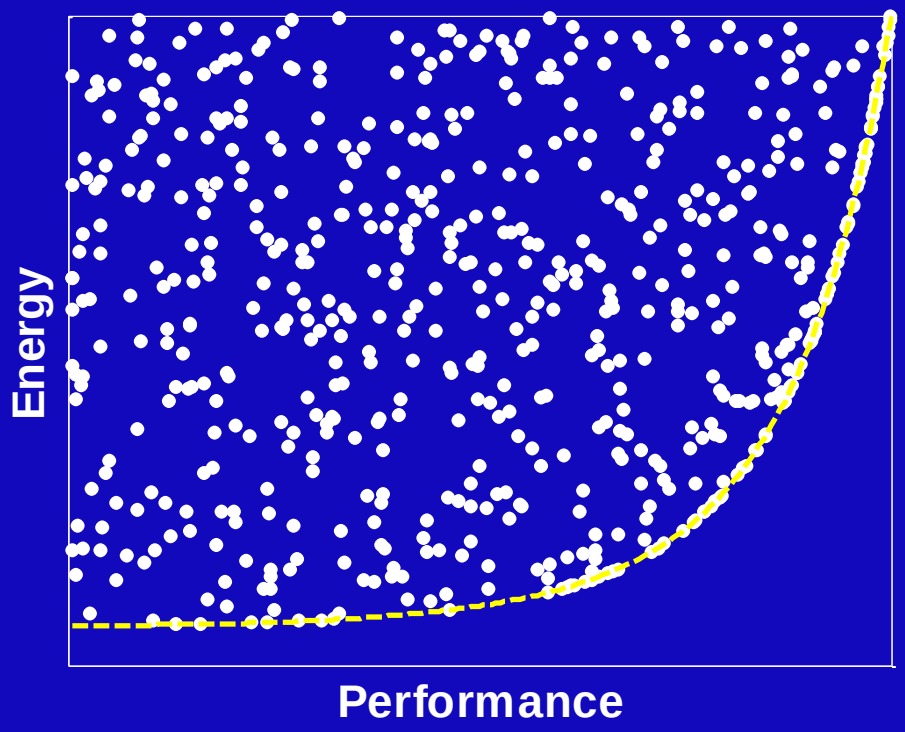

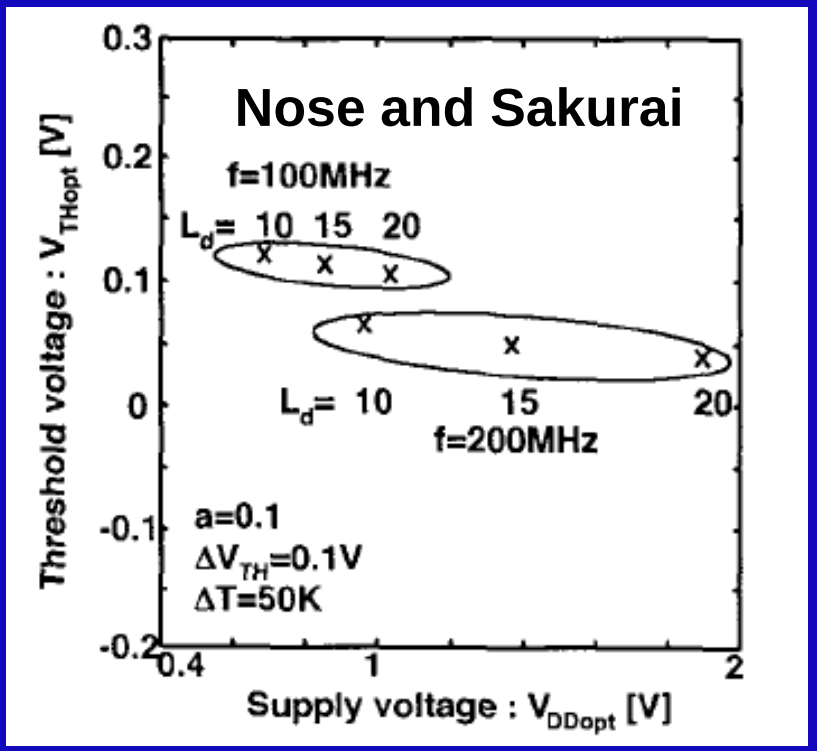

- Pick V_{dd}, V_{th} to minimize energy for given performance (1/delay)
  - Assuming work function (V_{th}) can be freely tuned
- Result: **optimal $I_{on}/I_{off} \propto L_d \cdot f / \alpha$**

Optimal I_{on}/I_{off} Insensitive to Device

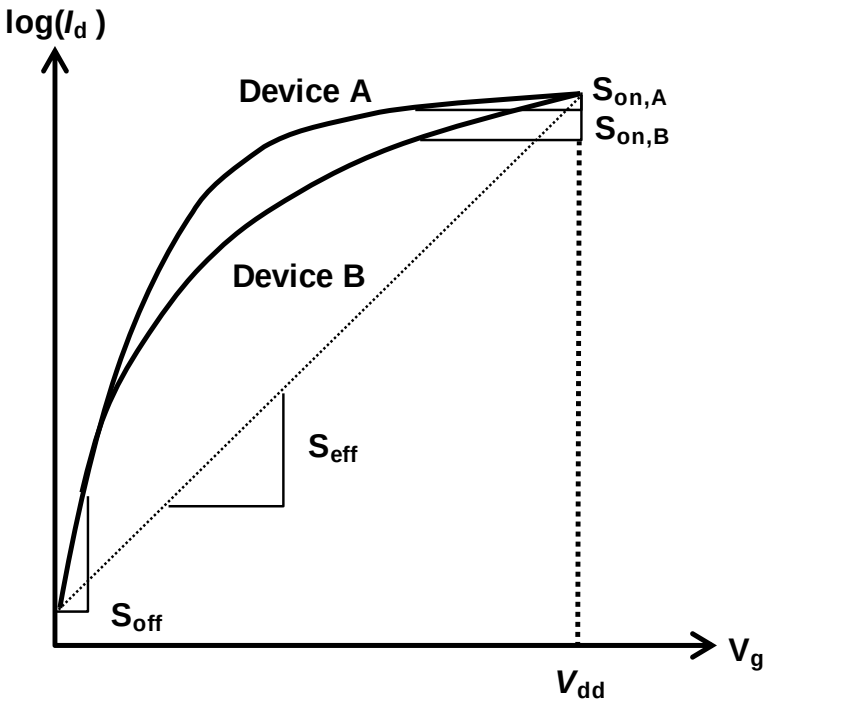

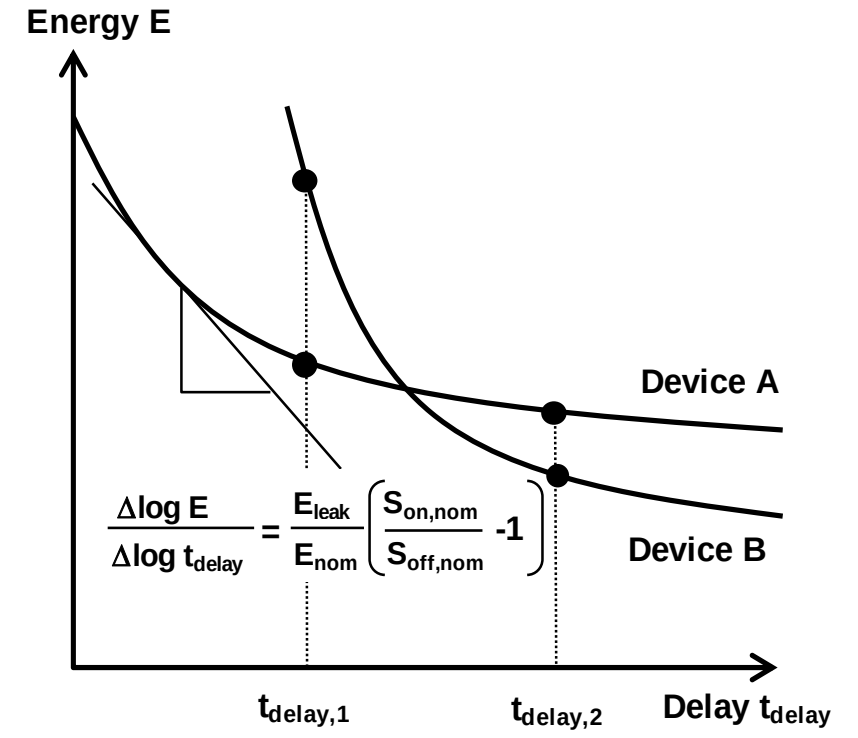

$$\frac{I_{on}}{I_{off}} = \beta \frac{L_d f}{4\alpha}, \quad \beta = -\frac{S_{eff}}{S_{off}} lambertW\left(-\frac{4\alpha}{L_d f} \frac{S_{off}}{S_{eff}} e^{\left(\frac{Soff + Son}{Seff} \right)} \right)$$

H. Kam, T.-J. King Liu, and E. Alon, "Design Requirements for Steeply Switching Logic Devices," to appear in *IEEE Trans. on Electron Devices*

Example Numbers

- Logic depth: L_d ~ 20 to 40
  - Can't be too small b/c of flip-flop and clocking overhead
- Activity factor: α ~ 1% to .01%
  - Most outputs unlikely to change in complex logic
- Fanout: f ~ 2 to 6
- So optimal I_{on}/I_{off} ~ $10^4 – 10^6$
  - This is really the logic switch requirement
  - I.e., power management doesn't change this…

Why Power Gating Doesn't Help

- Can indeed use another switch to turn off power
  - With higher I_{on}/I_{off} power switch, reduces E_{leak}
- But very hard to improve effective α
  - When to turn the power on?
- "Power managing" each gate = reproducing the logic…

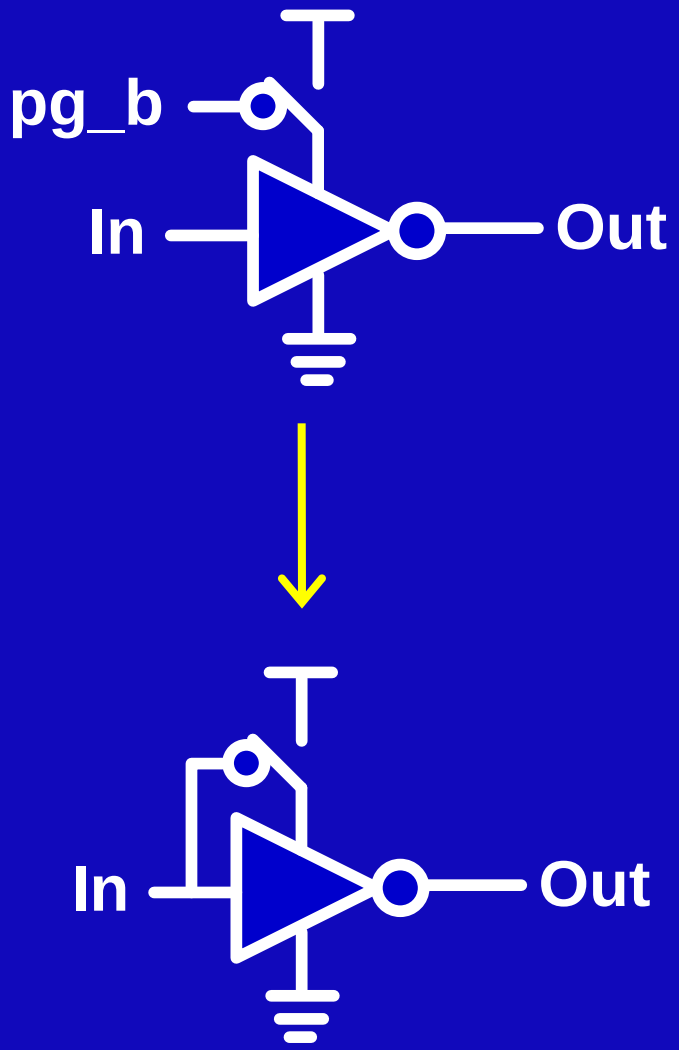

What Power Gating Is Good For

- Eliminate E_{leak} when system is off
  - I.e., when "obviously" not doing any work
  - So that knowing gating signal is nearly free
- Key point:
  - Power gating only reduces "system variability" penalty
  - Device variability?

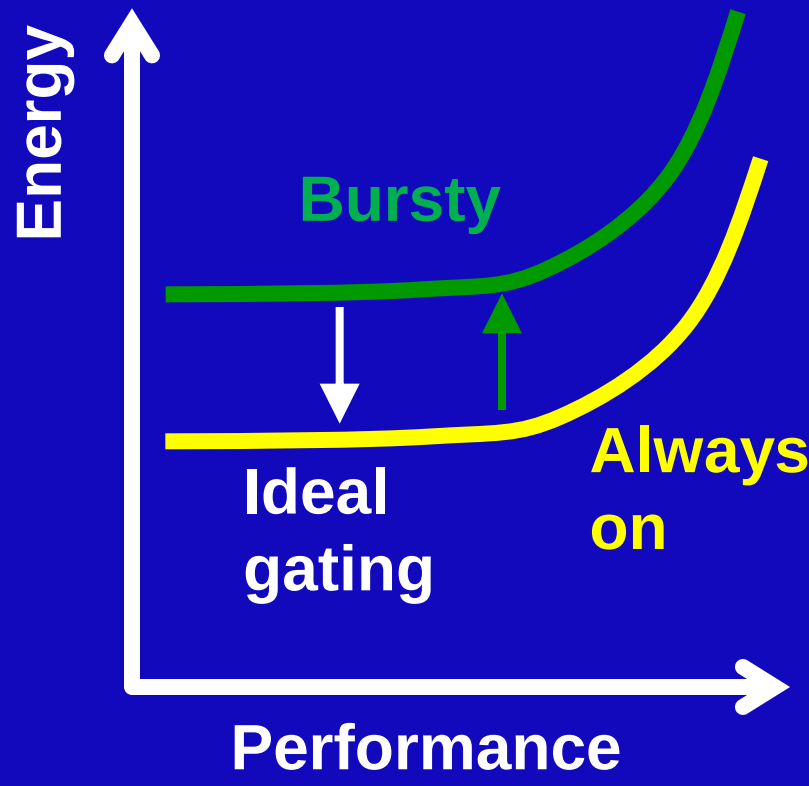

Implications of Device Variability

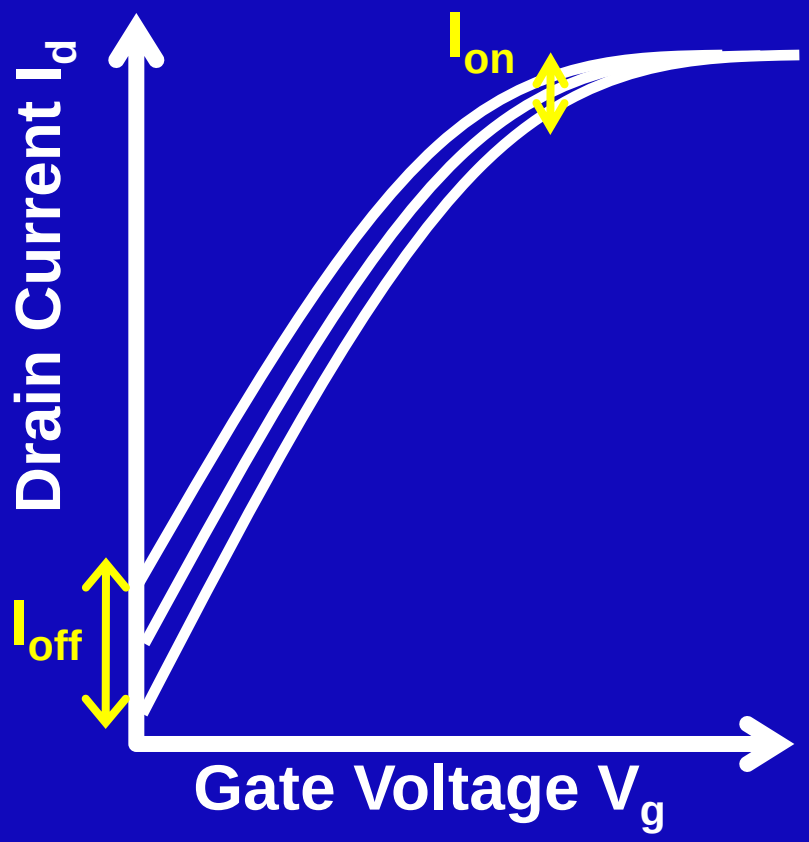

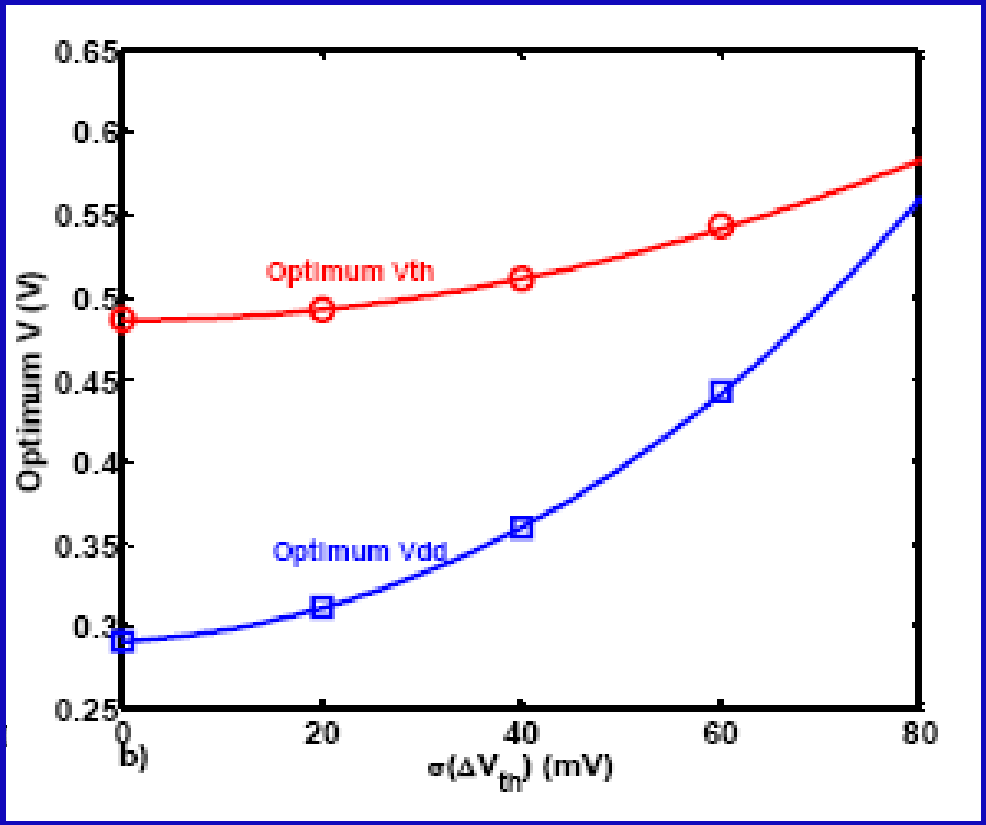

- Device variability hurts in two ways
  - Reduces effective I_{on} (delay set by worst-case)
  - Increase effective I_{off} (leaky devices dominate)
- Forces increase in nominal I_{on}/I_{off}…

Steep Switches Need Low Variation

- With steep device, leakage increases dramatically with $\sigma(\Delta V_{th})$
- For same variability:
  - Can even make "steep" switch worse than CMOS

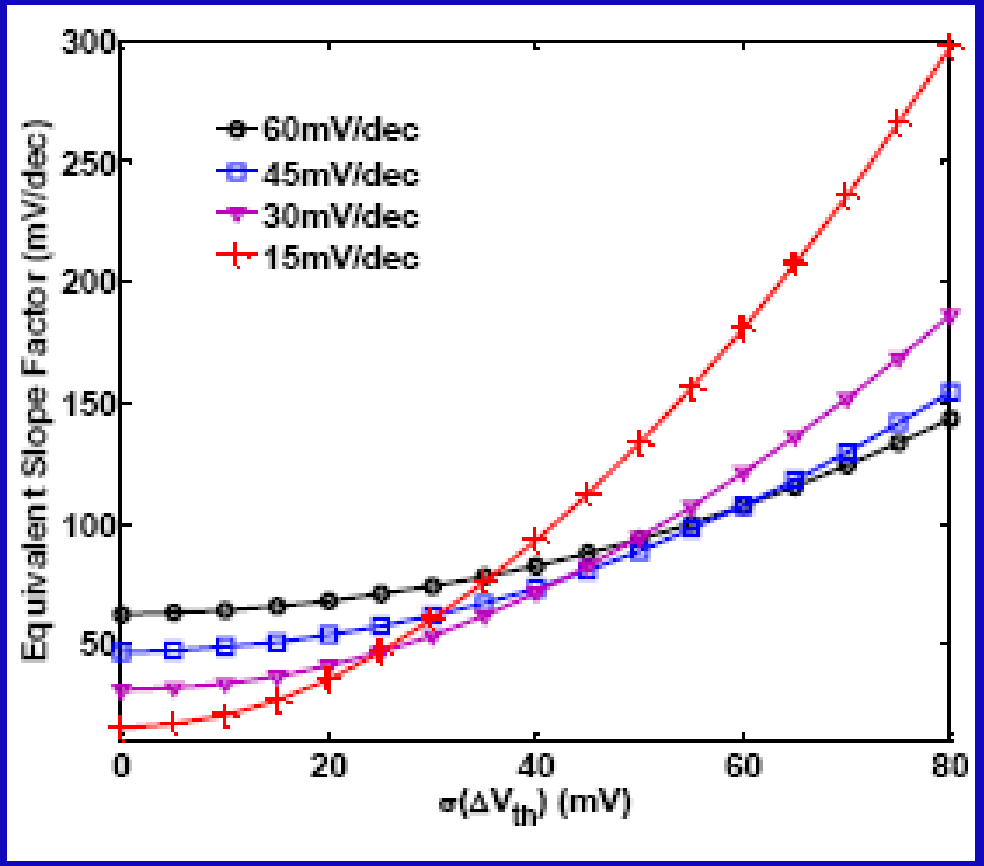

- Must consider and quantify device variability in advance

Another Issue: Wire Capacitance

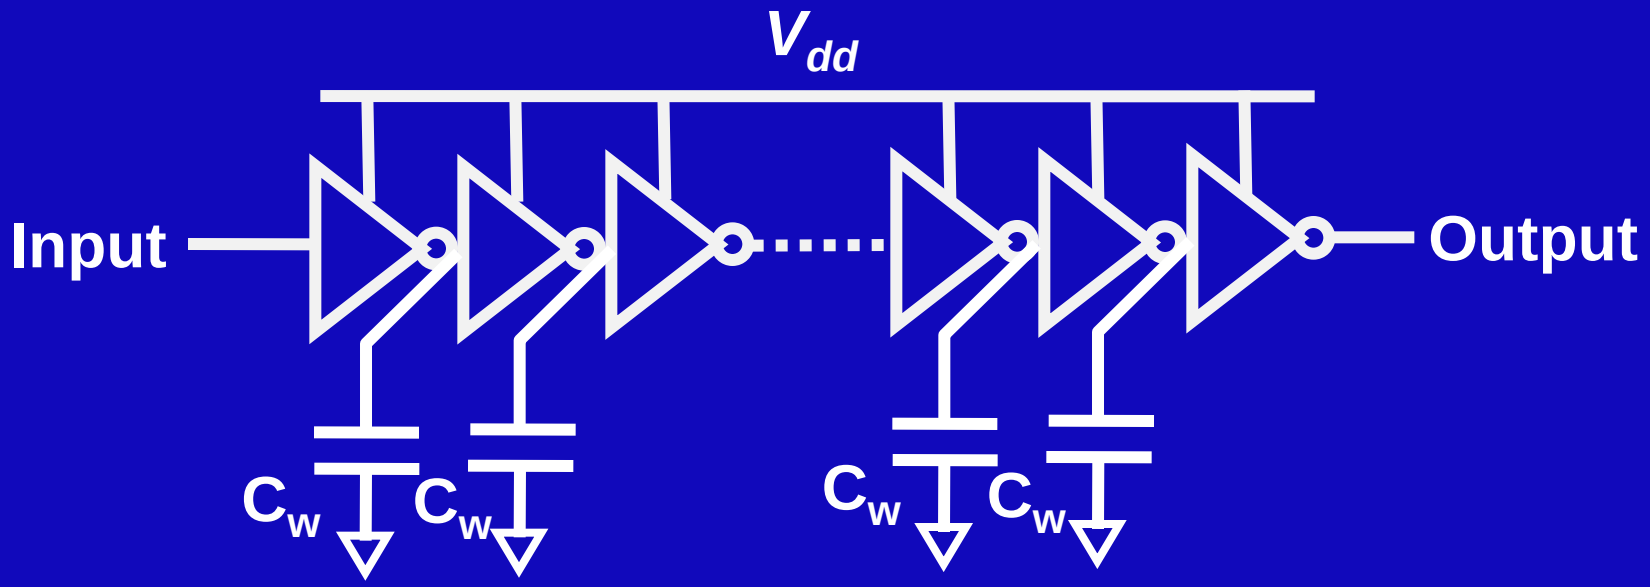

- **Wires critical to both delay and energy:**
  - **Minimum device C: ~0.1 fF**
  - **1μm wire C: ~0.2 fF**
- **Wires often set required device V_{dd}/ I_{on} (R_{on})**

Where New Devices Look Good

- Achieving sharp S^{-1} and low R_{on} looks really tough

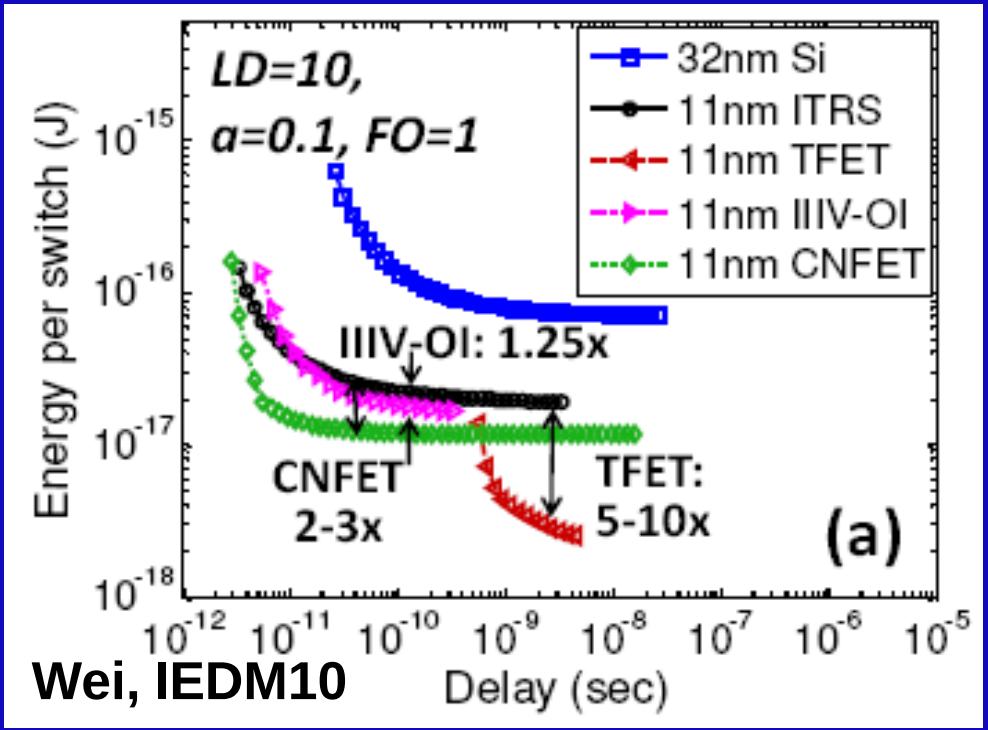

- But, even if new switch only improves energy at higher delay (higher R_{on})…

Parallelism

Perf. ∝ f_{clk}

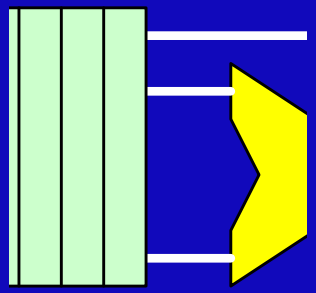

Perf. ∝ $2f_{clk}$, E/op ~const

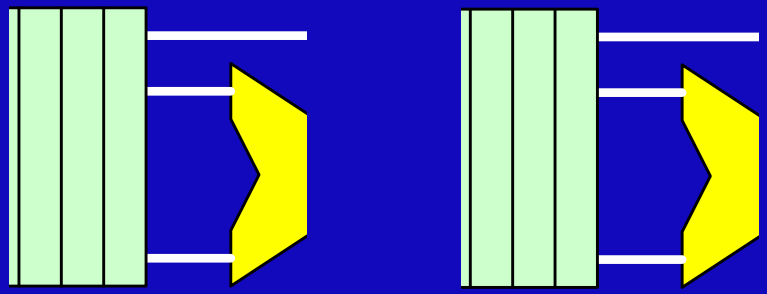

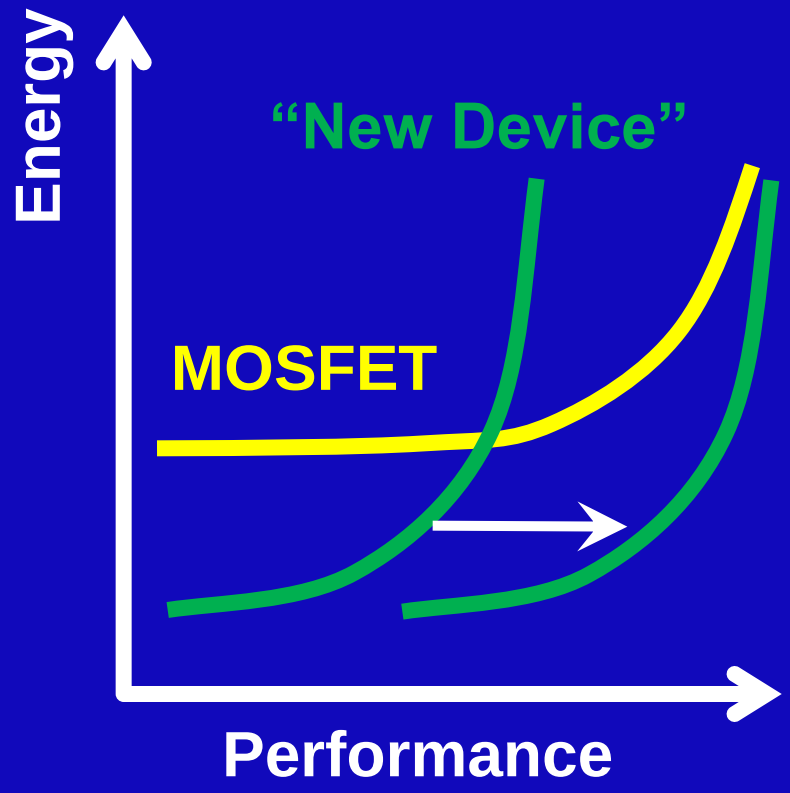

- Parallelism allows slower devices
  - Already applying parallelism to CMOS today

Parallelism (cont'd)

- Benefits of parallelism will eventually run out
  - CMOS has minimum energy/op
  - Set by min. V_{dd} to achieve optimal I_{on}/I_{off}

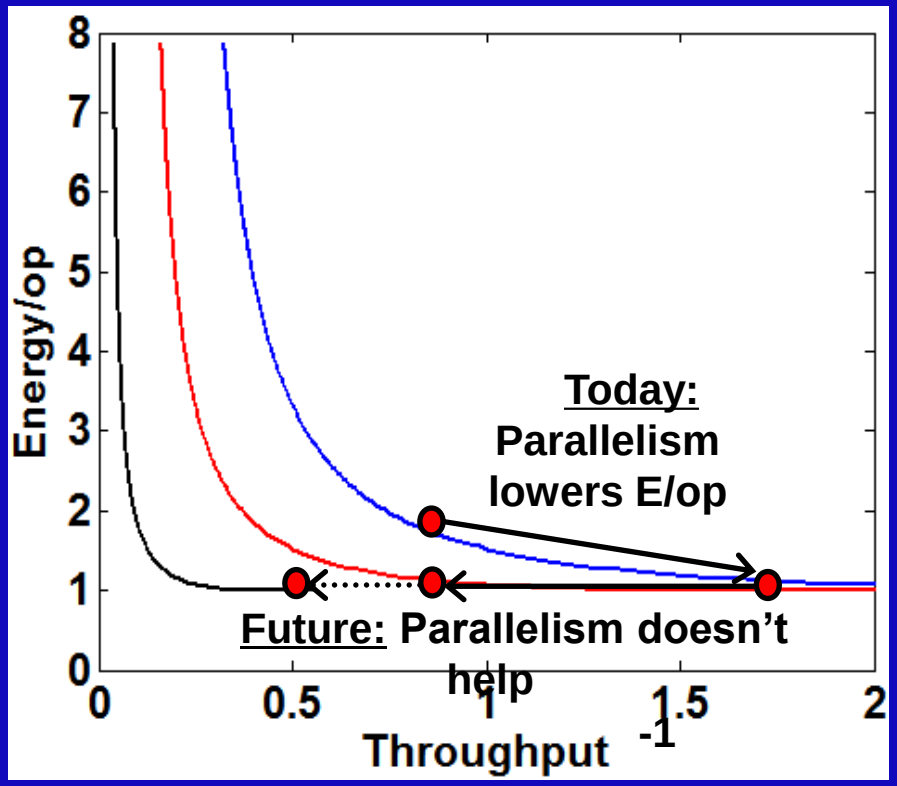

- Likely the main opportunity for new devices…
  - If achieve I_{on}/I_{off} of ~10^4 – 10^6 at (>10X) lower $C_{tot}V_{dd}^2$

Summary

- **Simple circuit/system models set device requirements**
  - **I_{on}/I_{off} set by logic depth, activity factor**
  - **Must consider variability and wires**
- **Parallelism limited by device E_{min}**
  - **Opportunity for new, low voltage devices**
- **Final plug: device/circuit co-design critical**
  - **Especially if alternate logic device is very different from CMOS**

Acknowledgements

- NSF Center for E3S
- Berkeley Wireless Research Center
- DARPA
- FCRP