Chip-Scale Energy and Power... and Heat

Dr. Wilfried Haensch

Thomas J. Watson Research Center, IBM

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE **MAR 2009**	2. REPORT TYPE	3. DATES COVERED **00-00-2009 to 00-00-2009**
4. TITLE AND SUBTITLE **Is Power the new Brick Wall?**		5a. CONTRACT NUMBER 5b. GRANT NUMBER 5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)		5d. PROJECT NUMBER 5e. TASK NUMBER 5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **IBM,Thomas J. Watson Research Center,Yorktown Heights,NY,10598**		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S) 11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**45**	

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

Is Power the new Brick Wall?

Wilfried Haensch,

Outline

- Introduction
- Classical CMOS
  - Devices
  - Power delivery
  - SRAM memory
  - I/Os
- Sub-threshold engineering
  - Does steep SS make sense?
- Adiabatic computing
- Summary
- Acknowledgements

Transistors, performance, power ….

- Density scales at traditional rate
  - Caches get larger and contribute significantly to transistor count
- Clock speed and power level off
  - Single thread performance is power limited
  - Need more efficient cooling solutions
- System performance can will pick up with parallelism
  - Many low power cores will deliver same throughput
  - Application dependent solution

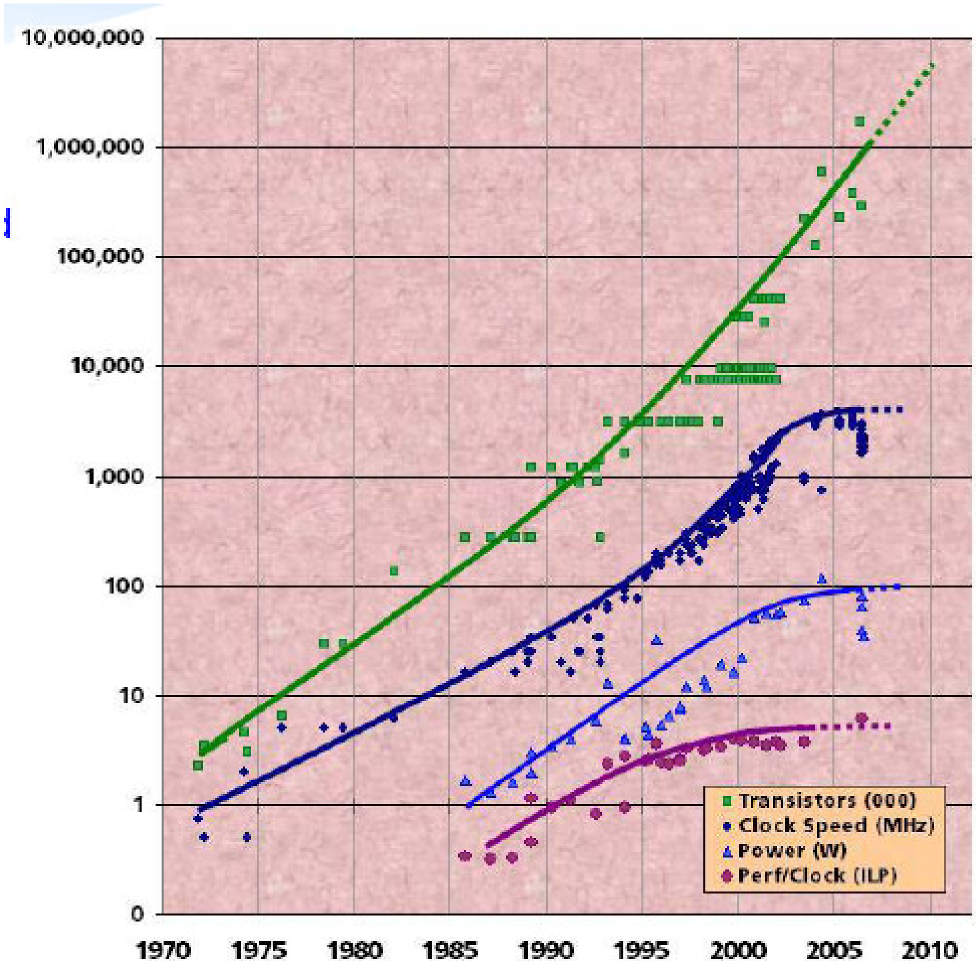

Source: **Kunle Olukotun, Lance Hammond, Herb Sutter, and Burton Smith**

Where does the power go?

Karthick Rajamani, Boulder TVC

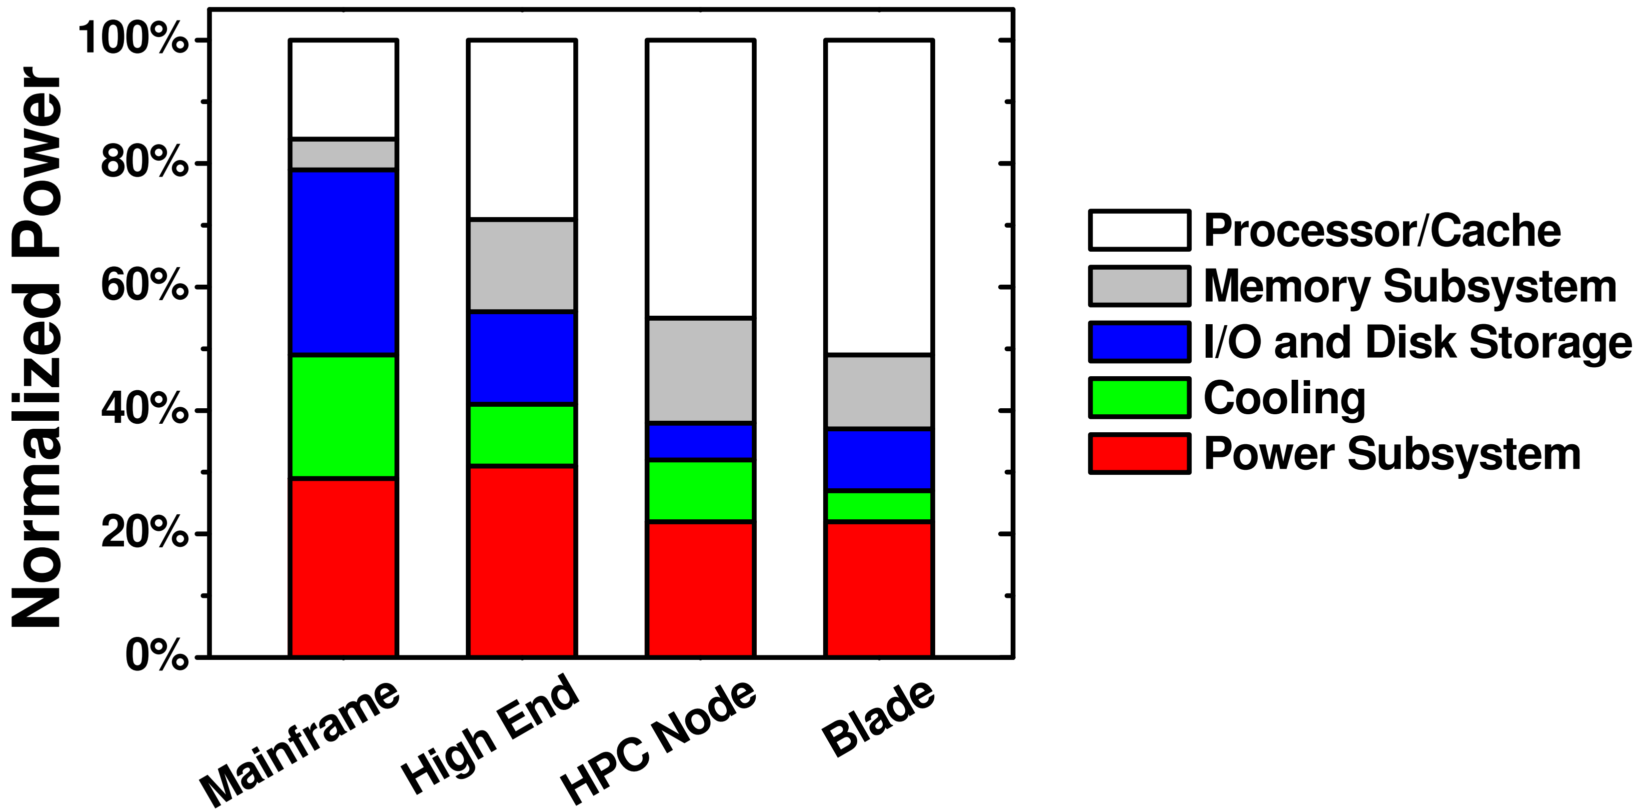

- Power is not only a processor problem. The power problem penetrates the whole system

Classical CMOS

General considerations

- Circuit

$$P \sim \frac{CV^2}{\tau} \quad \text{Power density}$$

$$\tau \sim \frac{CV}{I} \quad \text{Delay}$$

- Chip

$$P_{act} \sim fC_{eff}V^2 \quad \text{Power}$$

$$f \sim \alpha(V - V_0) \quad \text{Frequncy}$$

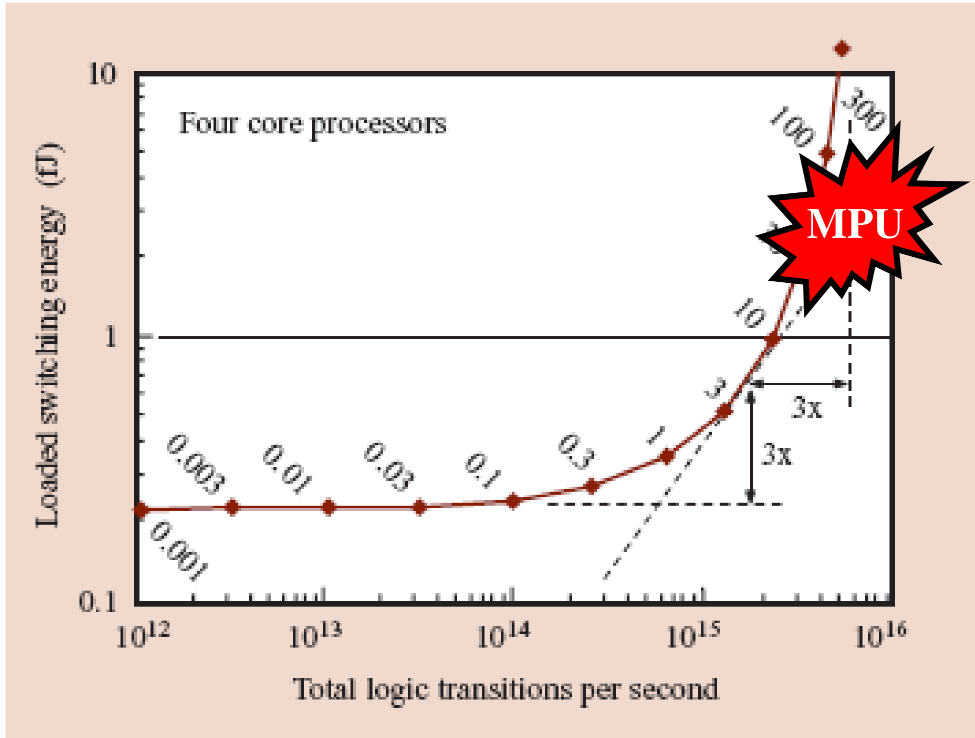

- Voltage scaling comes at the cost of performance

General considerations

- Circuit

$$P \sim \frac{CV^2}{\tau} \quad \text{Power density}$$

$$\tau \sim \frac{CV}{I} \quad \text{Delay}$$

- Chip

$$P_{act} \sim fC_{eff}V^2 \quad \text{Power}$$

$$f \sim \alpha(V - V_0) \quad \text{Frequncy}$$

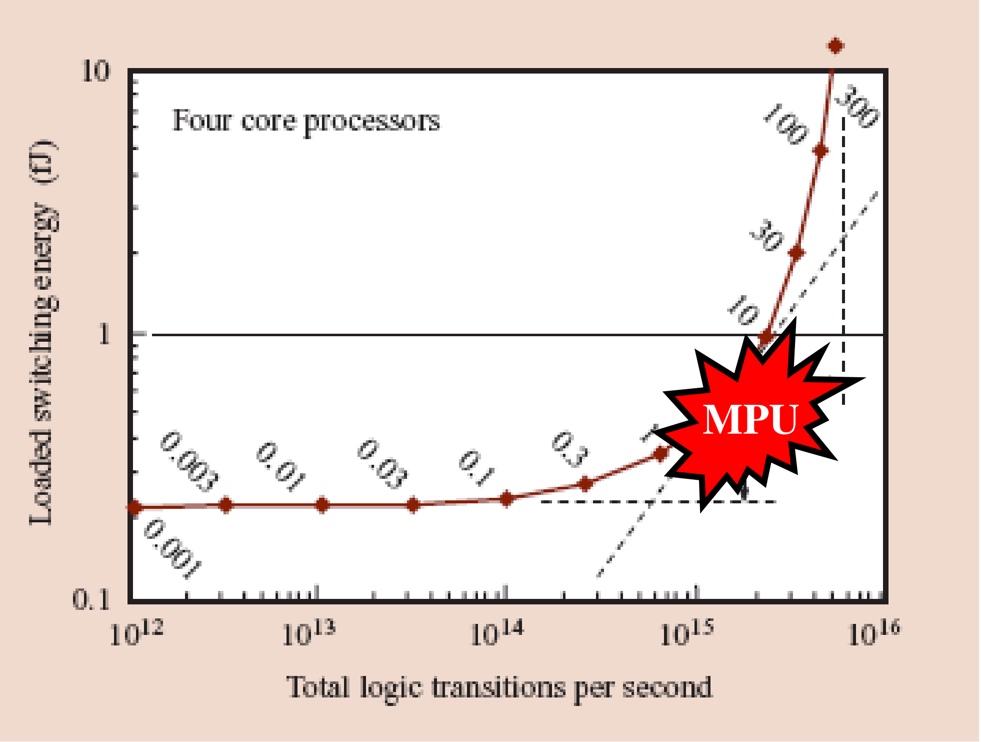

- Voltage scaling comes at the cost of performance
- Silver bullet: *Low V, low C, high I*

Scaling dilemma

$$I_{off} = I_{V_{th}} 10^{-\frac{V_{th}}{S}}$$

$$I_{on} = C_{inv} \upsilon_{sat} (V_{dd} - V_{th} - V_{sat})$$

I_{off} ~ 100 nA/µm
S > 60 mV/dec
I_{Vth} ~ 300 µm/L nA/µm
C_{inv} ~ 35 fF/µm^2 ~ EOT ~ 0.8nm
I_{on} ~ 1.5 mA/µm
V_{sat} ~ 0.25V

- Simple voltage scaling is challenging

$$V_{th} > 125\text{mV}$$

$$V_{dd} - V_{th} > 650\text{mV}$$

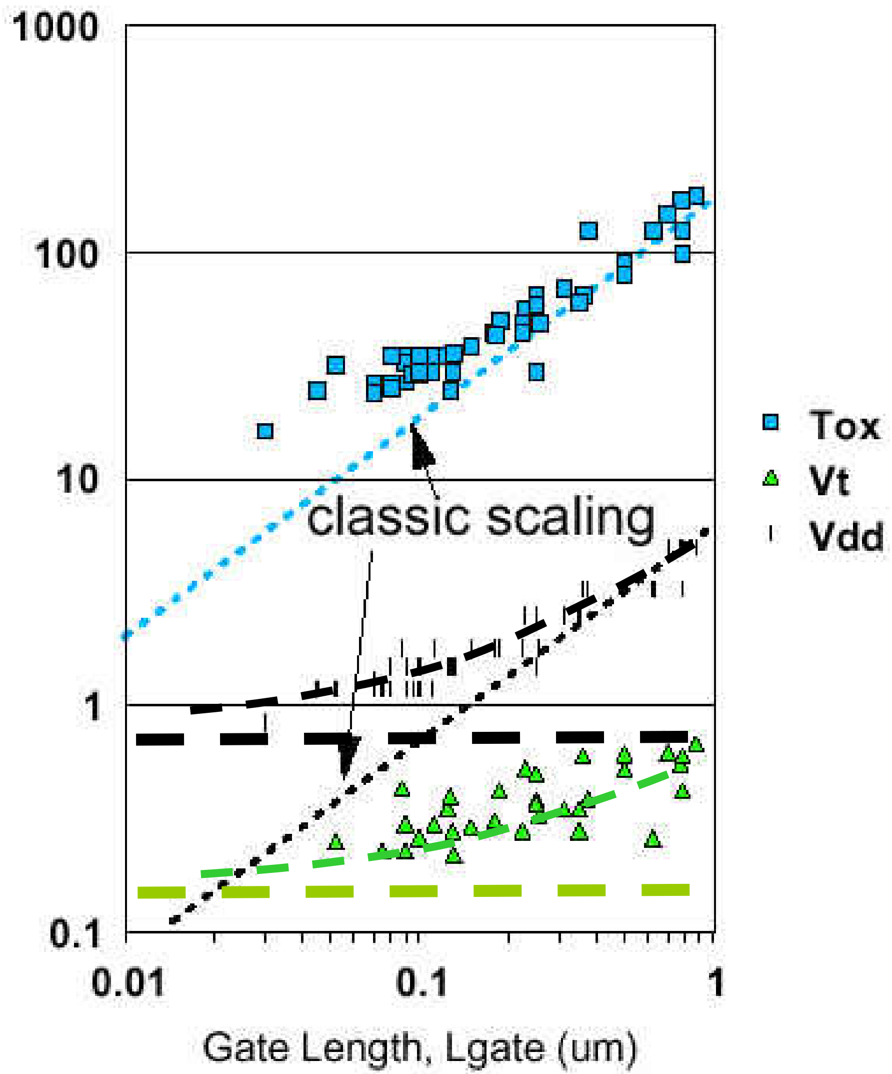

Device choices

Ion @ given Ioff

Ieff @ given Ioff

- Short channel effects will degrade I_{eff} faster than I_{on}
  - $I_{on} \sim (V\text{-}Vt_{sat})^{\alpha}$
  - $I_{on} - I_{high} \sim DIBL*V/2$
  - $I_{low} \sim (V/2\text{-}Vt_{sat})^{\alpha}$

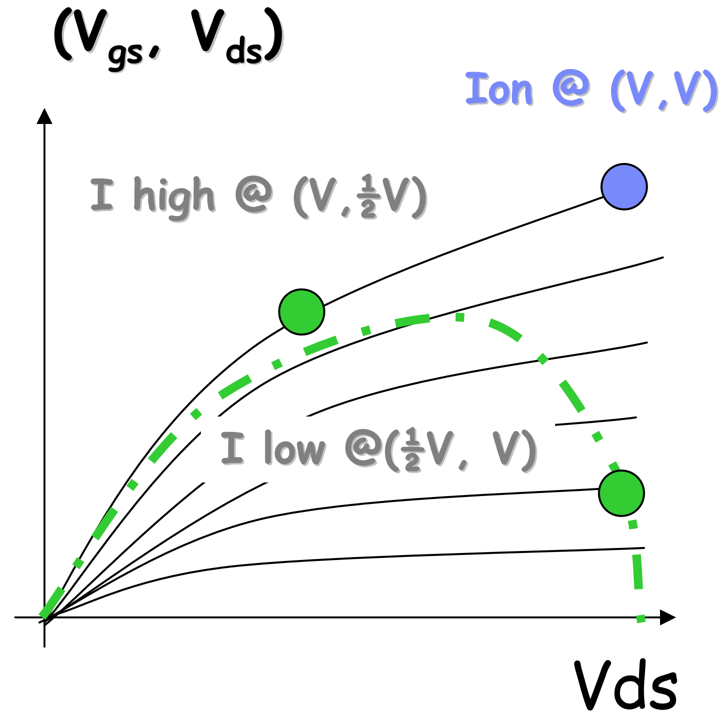

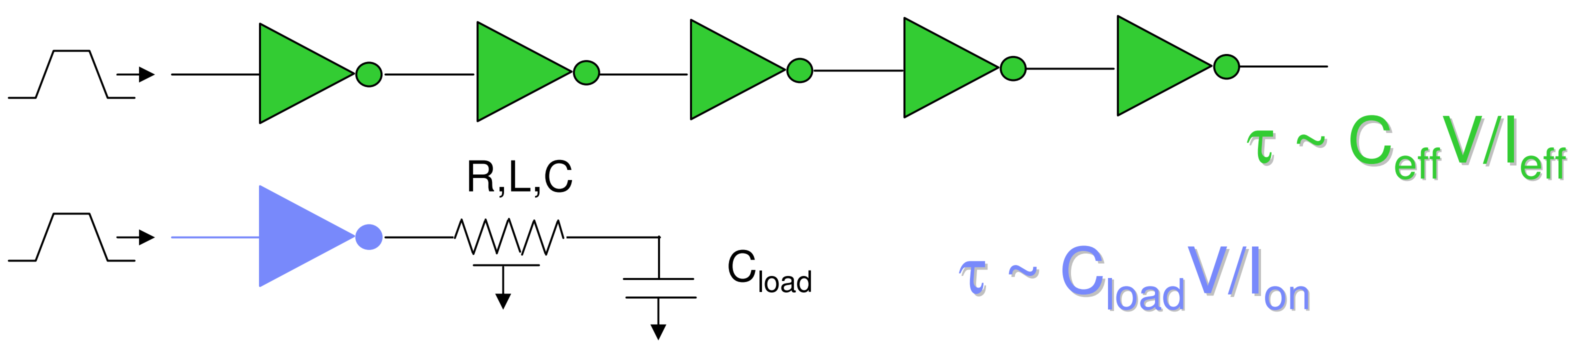

Si Device Evolution (I)

High μ options: liner stress, eSiGe, SMT, HOT, SSDOI, ..
Gate stack: high-k, Metal gate
Junction engineering

Body control SCE will enable length scaling without aggressive dielectric

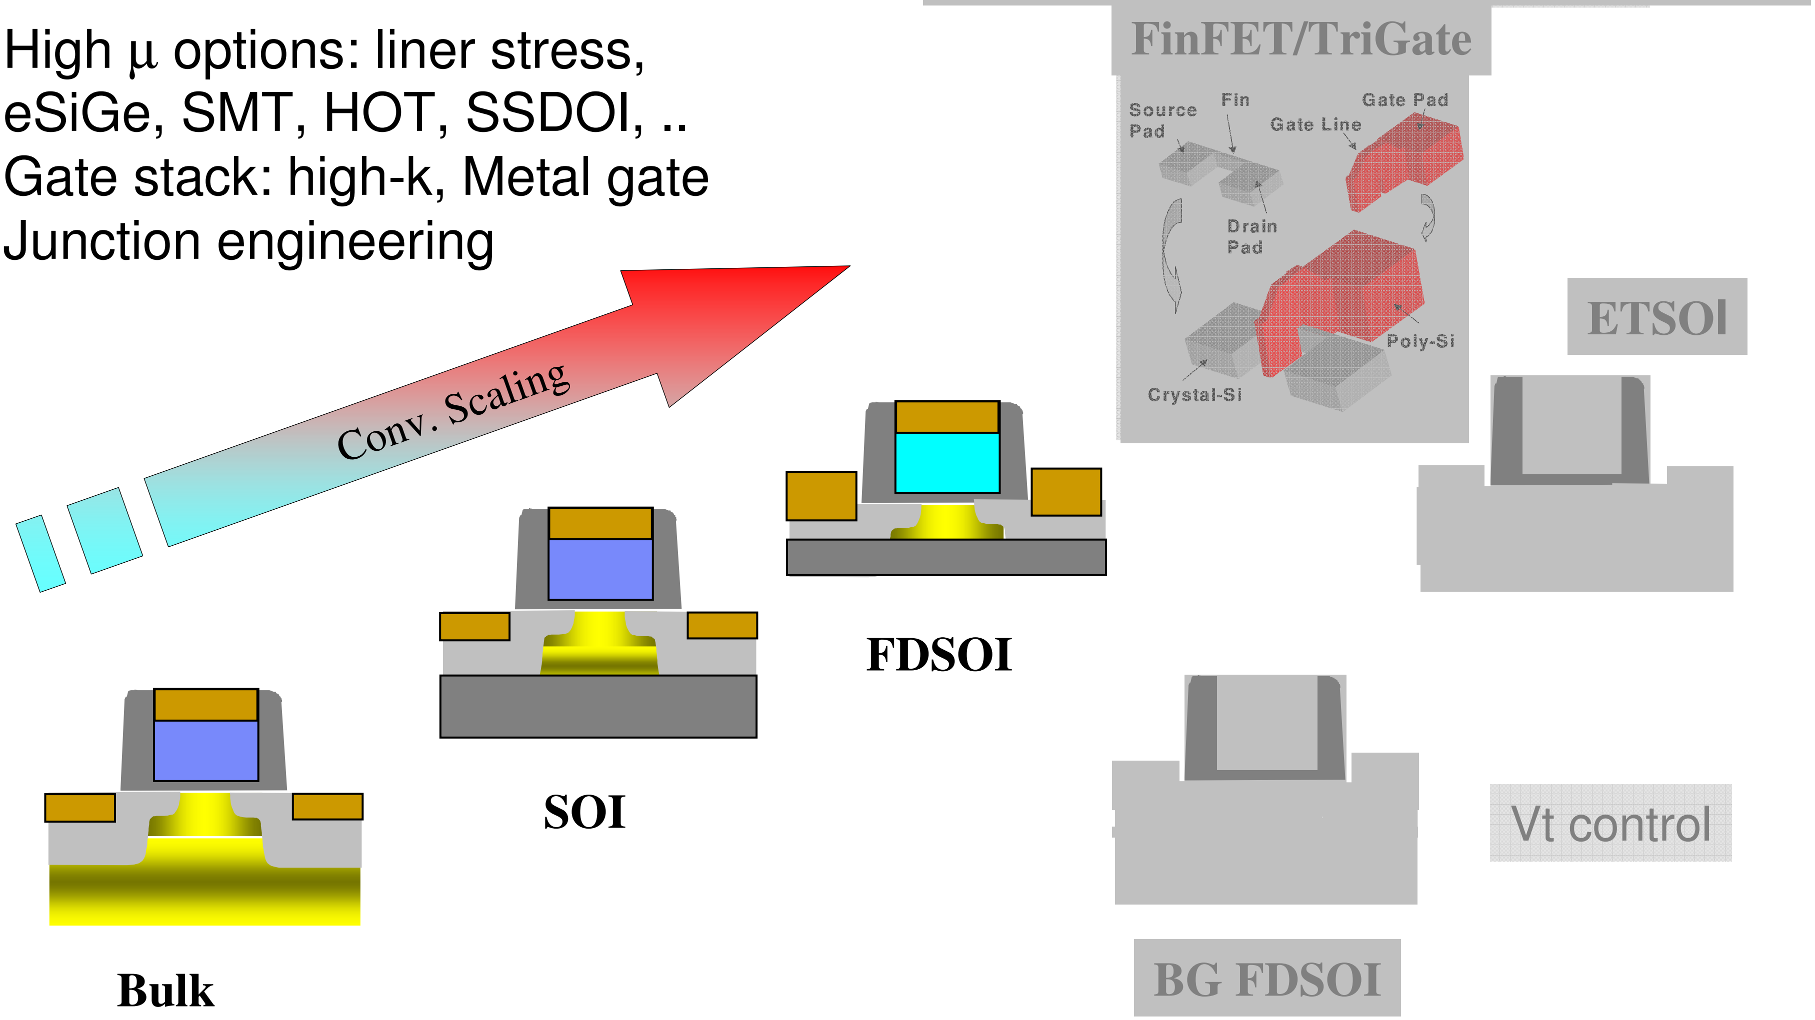

Si Device Evolution (II)

High µ options: liner stress, eSiGe, SMT, HOT, SSDOI, ..
Gate stack: high-k, Metal gate
Junction engineering

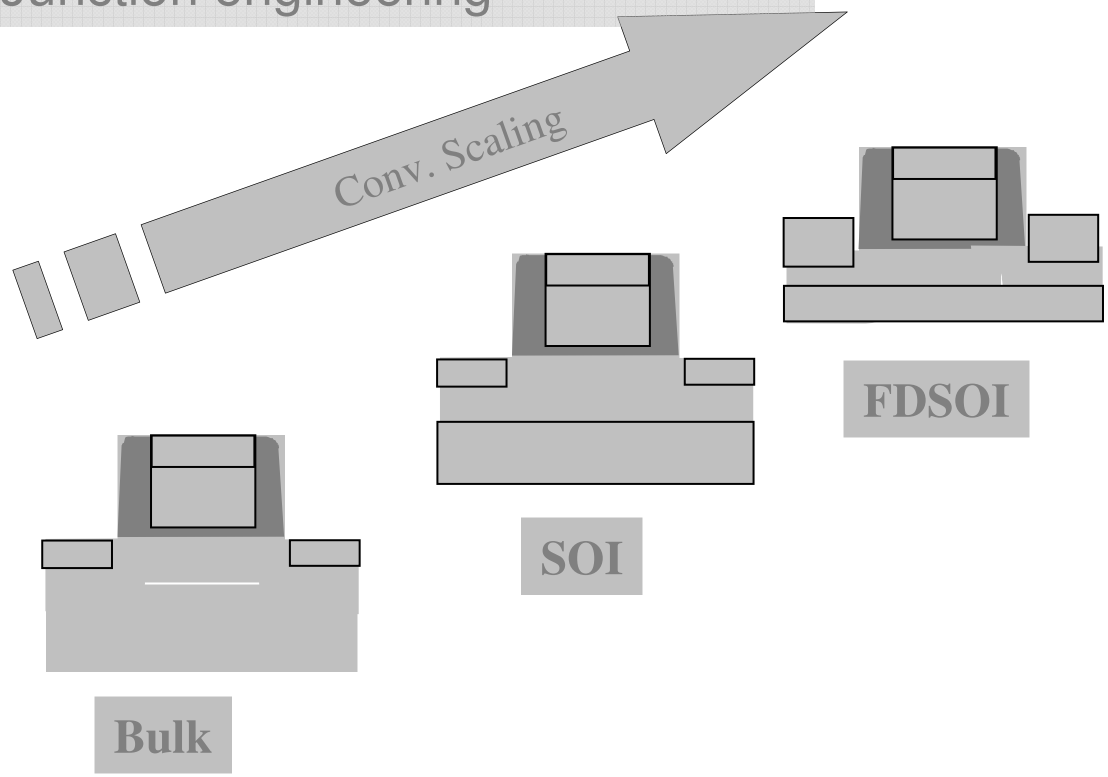

Body control SCE will enable length scaling without aggressive dielectric

FinFET/TriGate

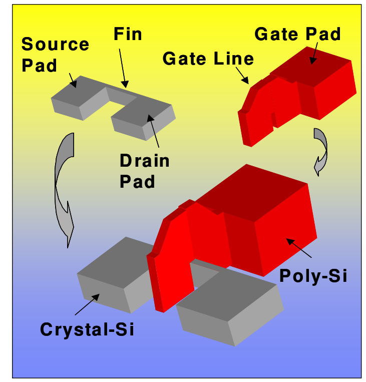

ETSOI

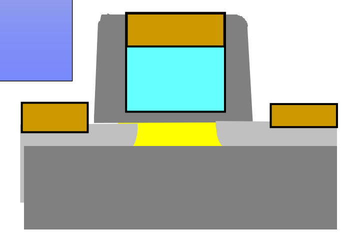

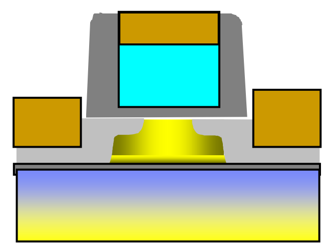

Vt control

BG FDSOI

Optimal technology - DIBL

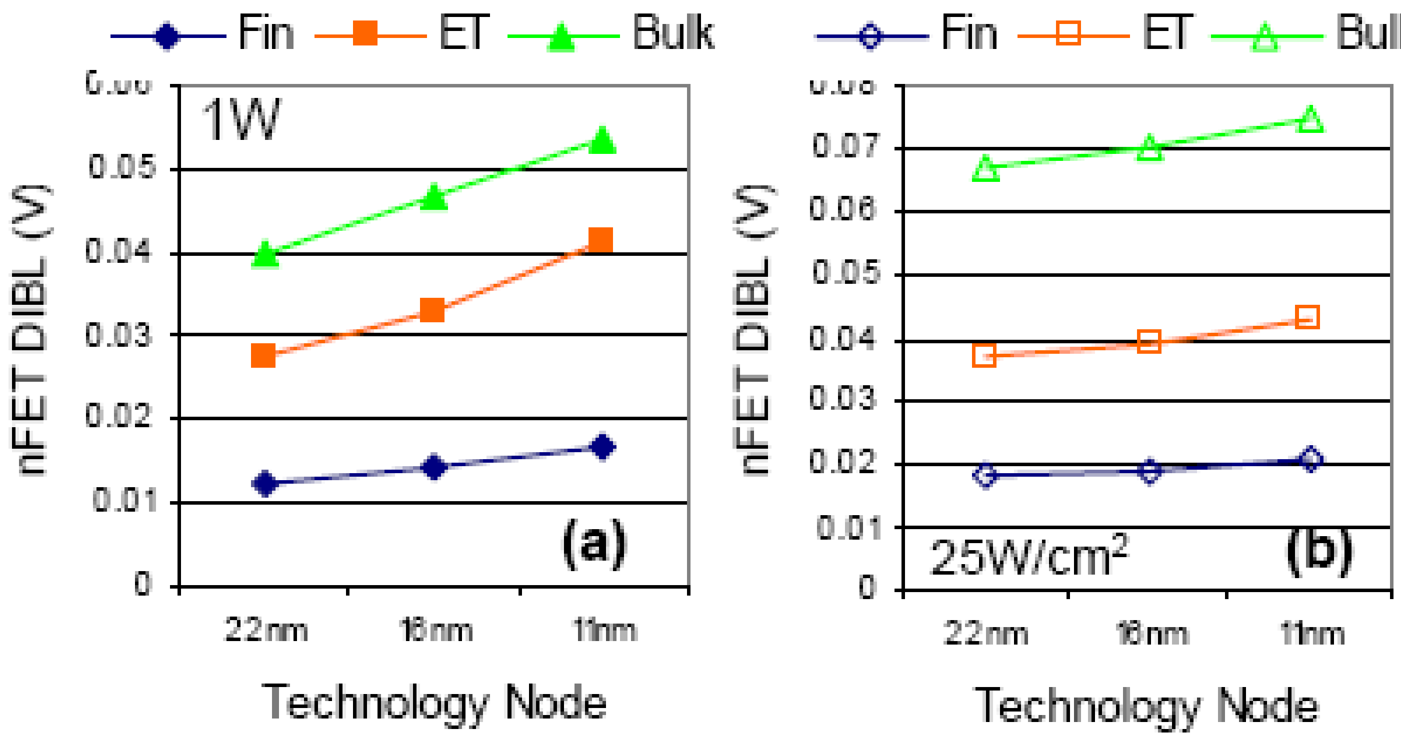

- Low DIBL is more advantageous for power optimized performance ……

Optimal technology – supply voltage

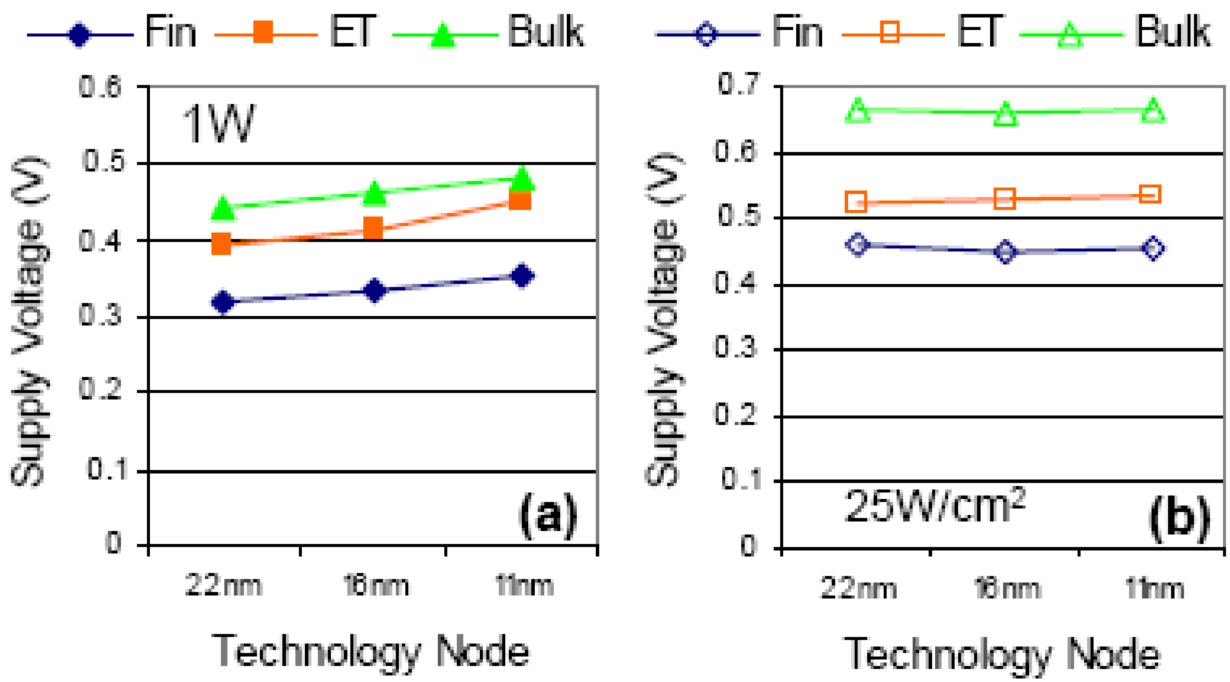

- …… because it allows reduction of supply voltage for maximum performance

Optimal technology – gate length

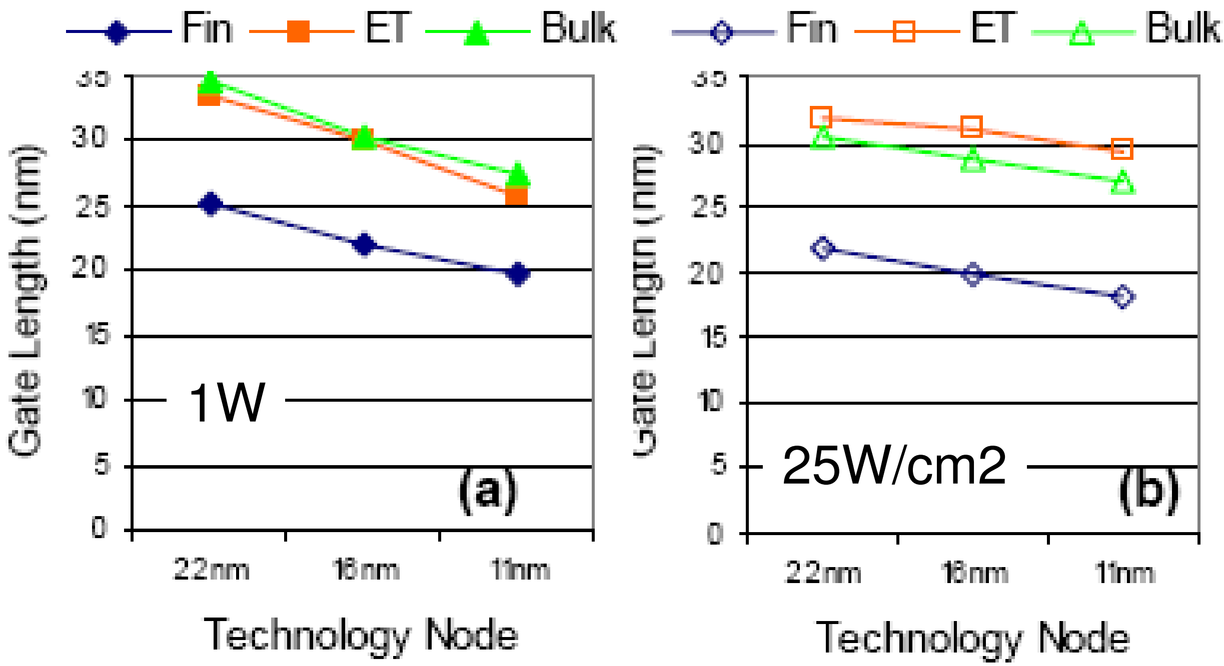

- Superior electrostatics allows shortest gate length for FinFET
- Gate length are longer than expected, no significant difference for low power or high performance

Optimal technology – density

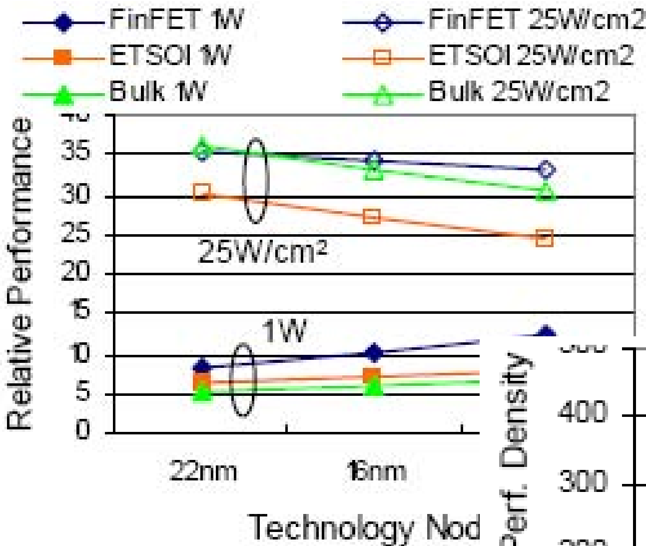

- Only moderate gains in single core performance for low power technology

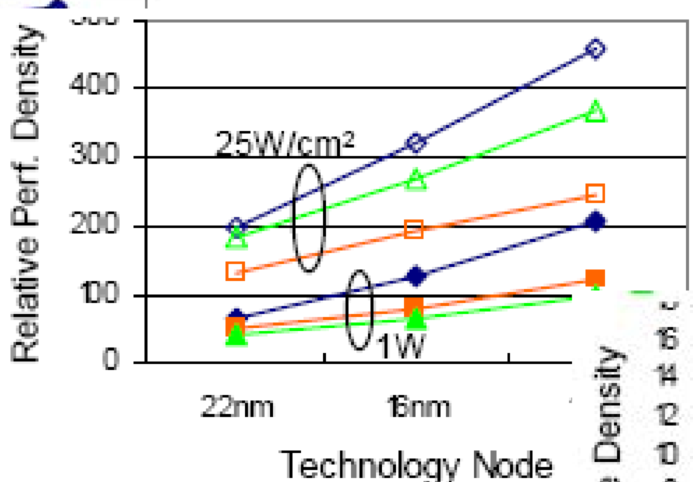

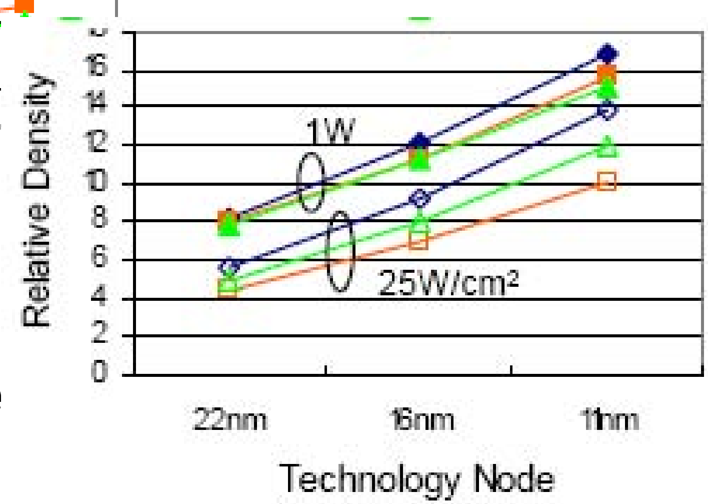

- Performance comes with parallelism and density advantage

Power delivery

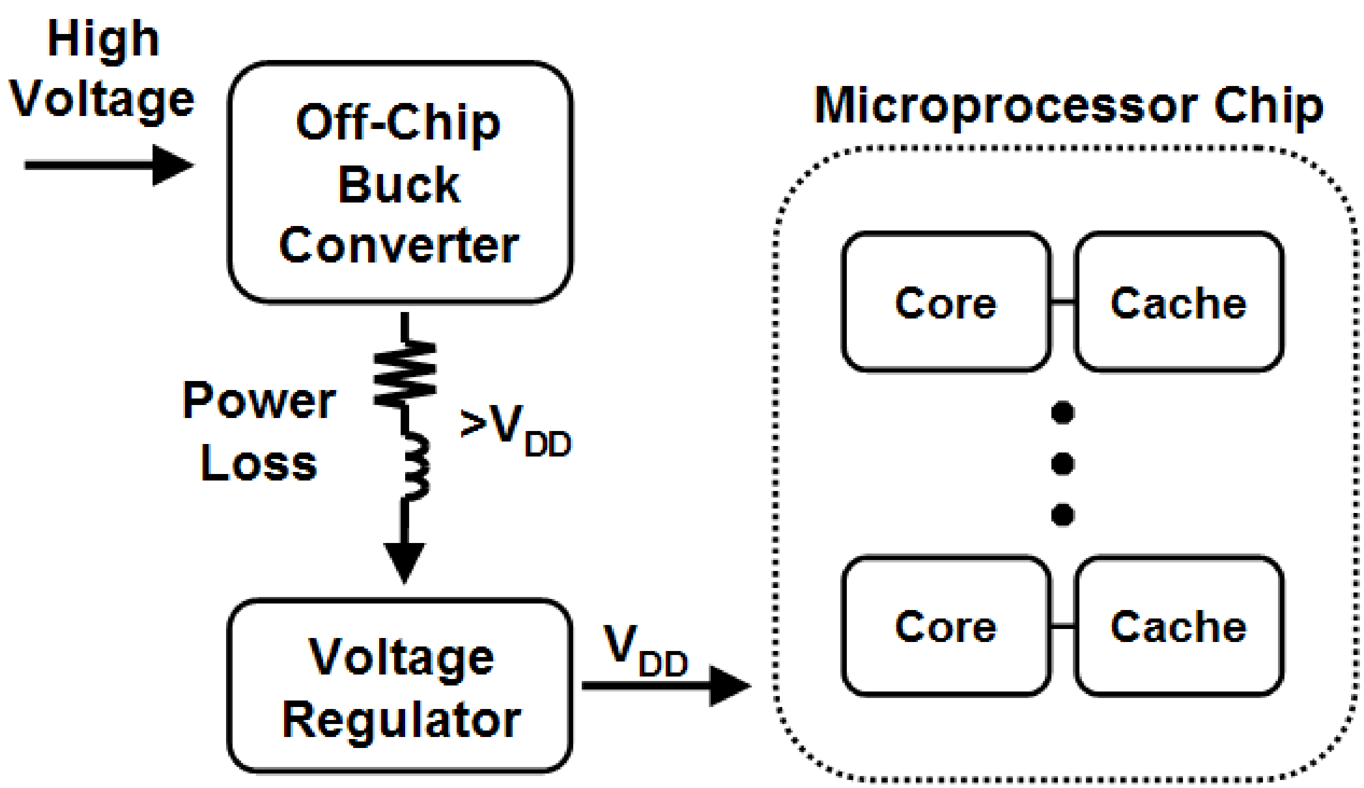

- Bringing high voltage as close to the chip would reduce power loss

The case for high voltage power delivery to ICs

- High voltage power delivery system ⇔ Put voltage down-converter close to the point of use
  - minimizes power loss
  - improves supply stability

- Power loss

$$\frac{P_{loss}}{P} = \frac{I^2 R}{P} = \frac{(P/V)^2 R}{P} = \frac{PR}{V^2}$$

- Supply stability

$$\frac{\Delta V_L}{V} = \frac{L\frac{dI}{dt}}{V} = \Omega \frac{LP}{V^2}$$

The case for high voltage power delivery to ICs

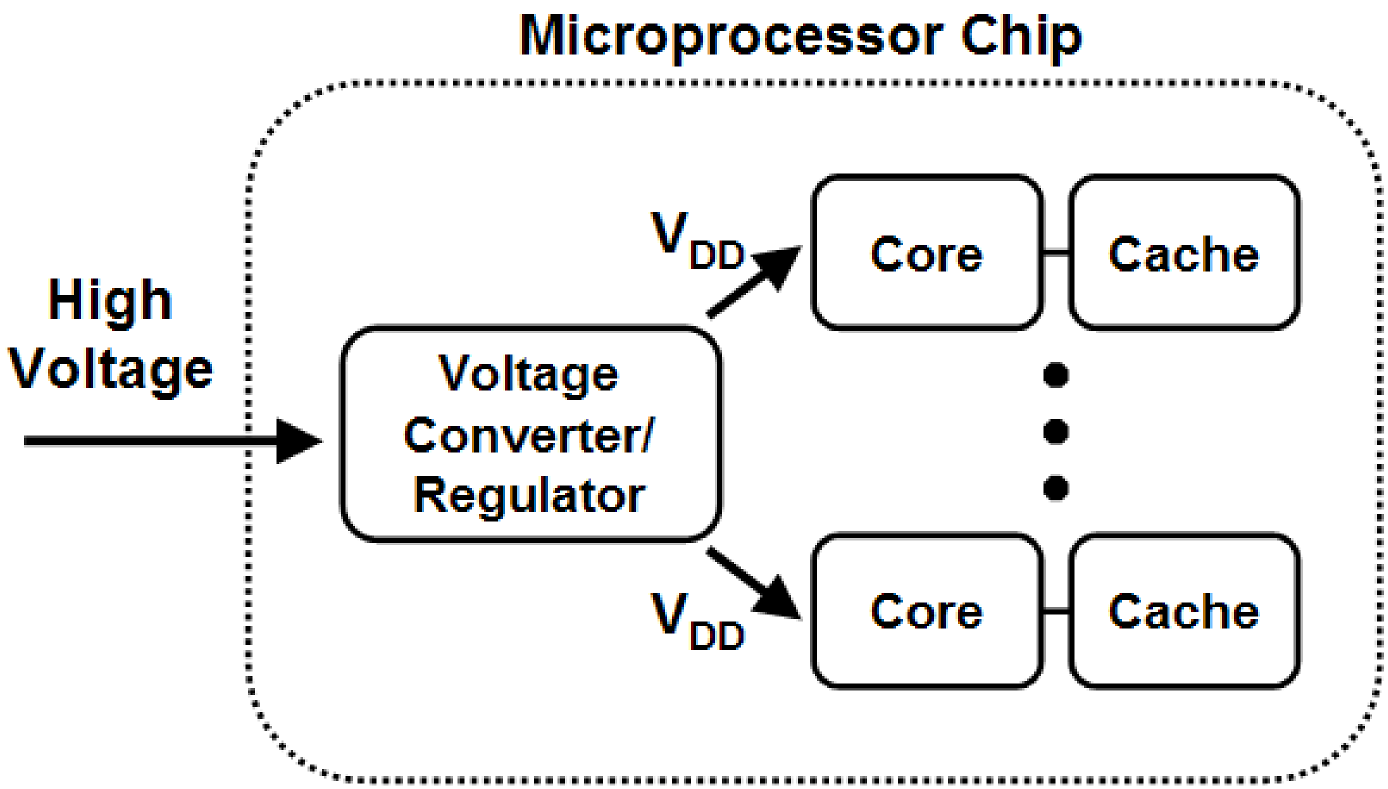

- Need efficient on-chip converters that are compatible with base CMOS technology

The ultimate vision

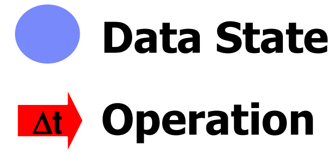

Homogeneous Multi-Processing

(e.g. 64 identical processors)

Critical Path = 5Δt

Total energy = 68

Heterogeneous Multi-Processing

(e.g. 1 HP + 63 LP processors)

Critical Path = 5Δt

Total energy = 15

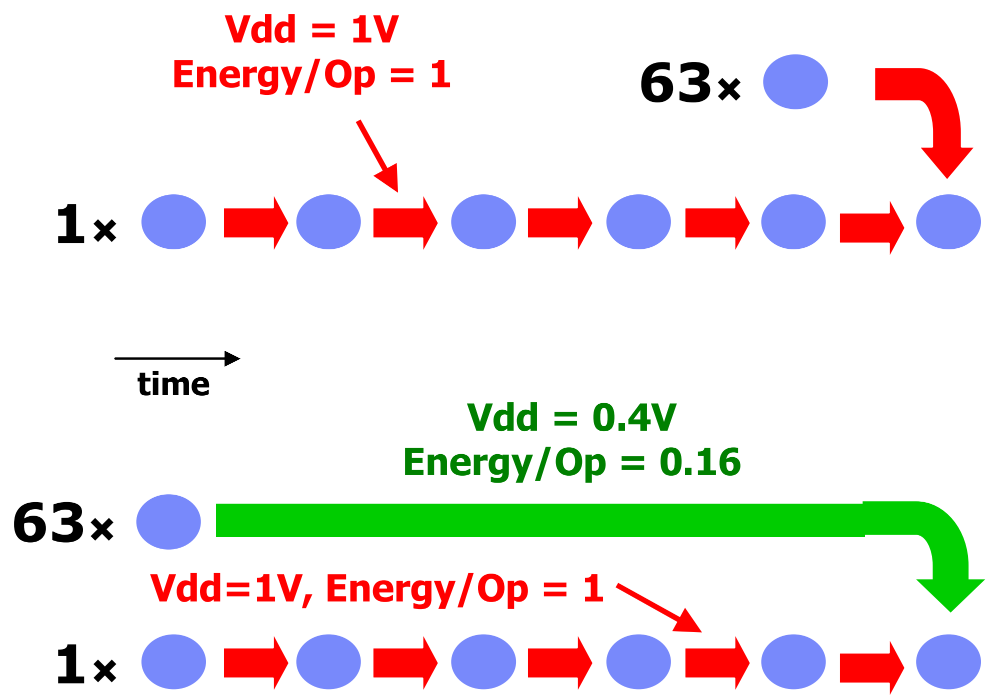

- Depending on load distribution, heterogeneous multi-processing can achieve dramatic power reduction

The case for high voltage power delivery to ICs

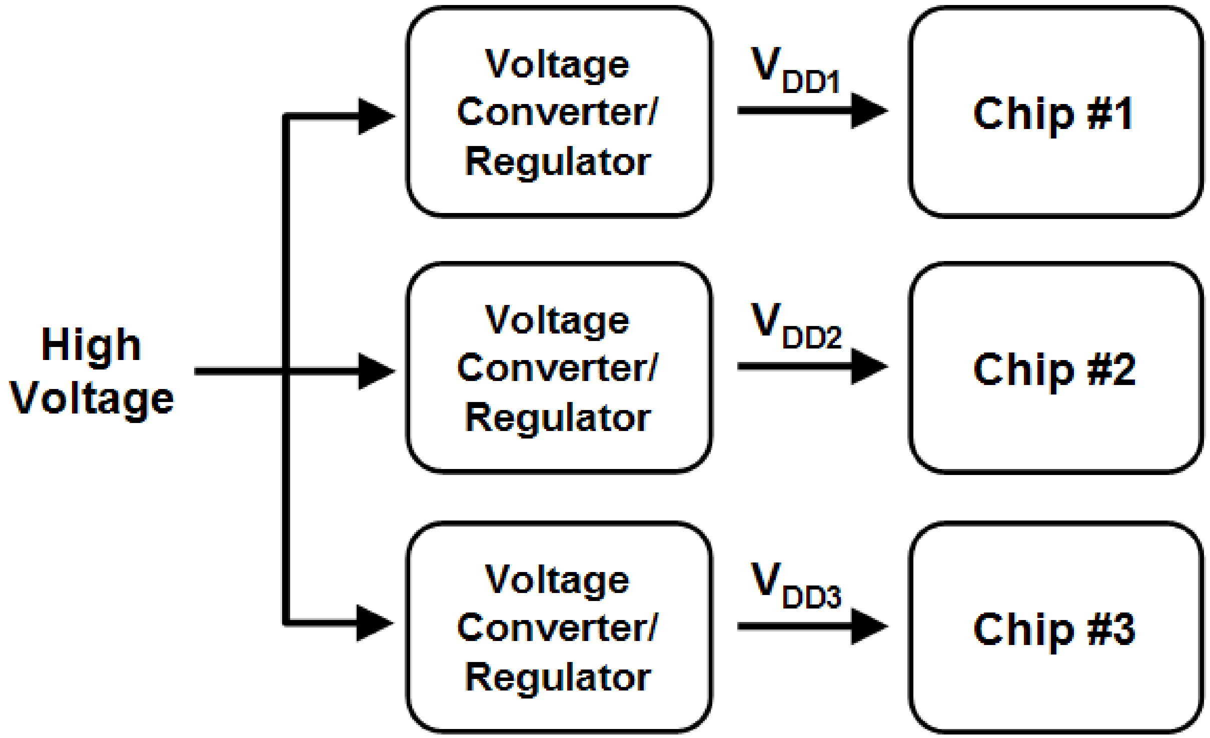

- Due to natural process variations (e. g. L_{gate}) need to regulate supply voltages for individual cores for performance matching

The case for high voltage power delivery to ICs

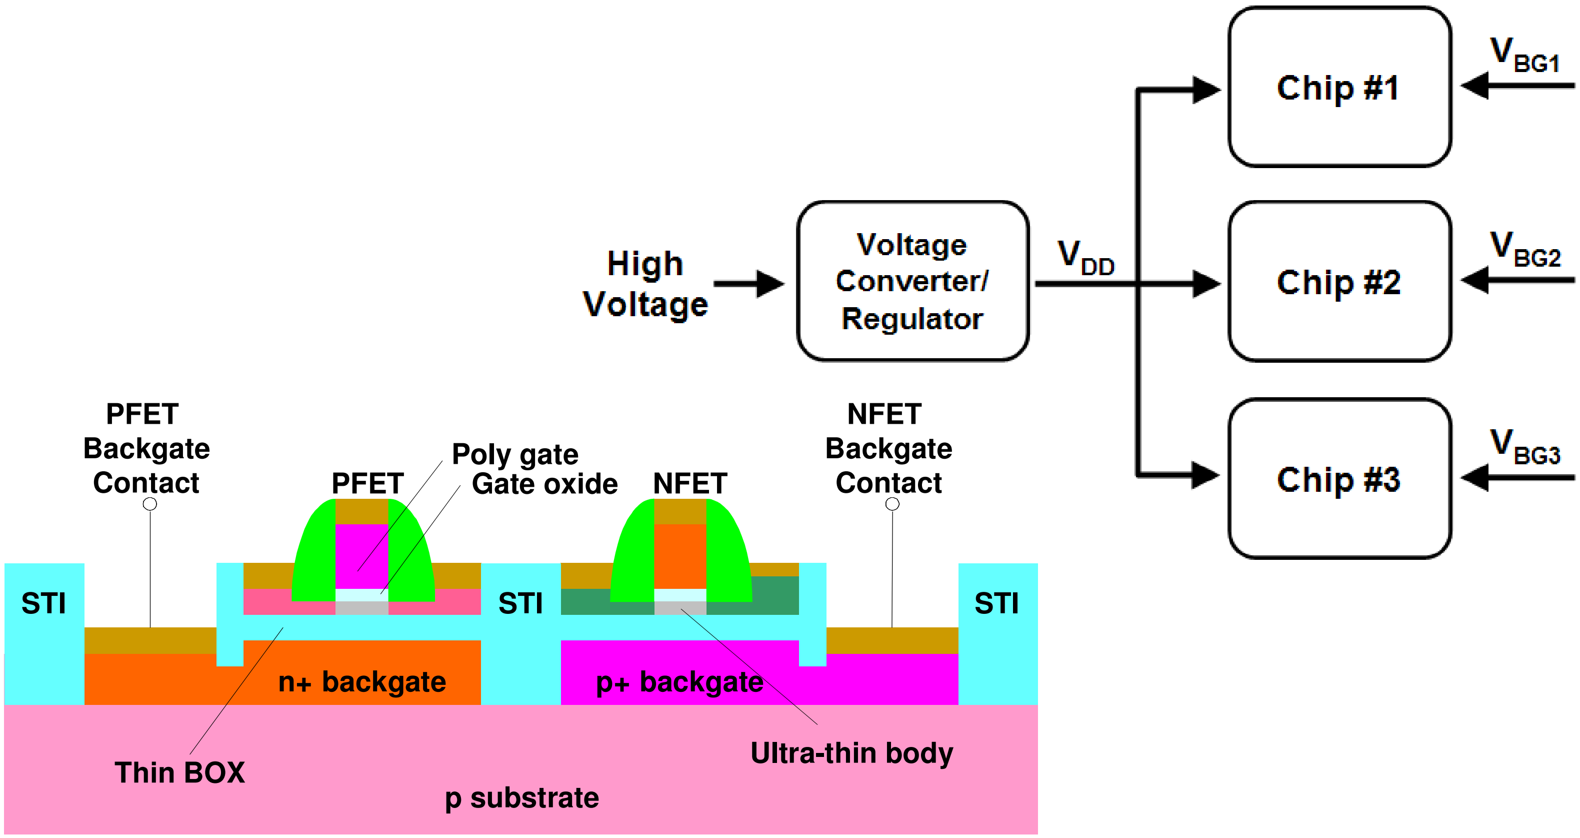

- Back-gate control would simplify voltage delivery

Obstacles to low-voltage operation

- Increased gate delays → *circuit design challenge*,
- Increased sensitivity to process variability → *device design challenge*.

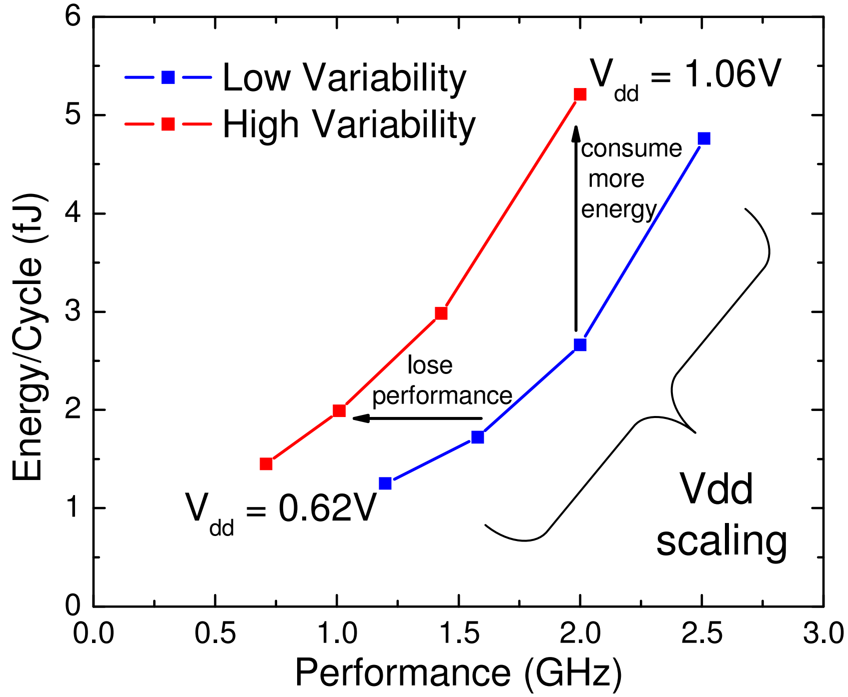

- Variability needs to be minimized for low voltage operation.

Noise

- Resistive noise
  - Current is a super-linear function of V_{DD}

$$\frac{I_{on}R}{V_{DD}} \propto \frac{V_{DD}^{1.5}}{V_{DD}} = V_{DD}^{0.5}$$

… scales faster than V_{DD}

- Capacitive noise
  - Charge is directly proportional to V_{DD}

$$\frac{CV_{DD}}{V_{DD}} \propto C$$

… scales with V_{DD}

- Inductive noise
  - Current is a super-linear function of V_{DD}

$$\frac{L\frac{\partial I}{\partial t}}{V_{DD}} \propto \frac{V_{DD}^{1.5}}{V_{DD}} = V_{DD}^{0.5}$$

… scales faster than V_{DD}

Low voltage circuit operation - SRAM

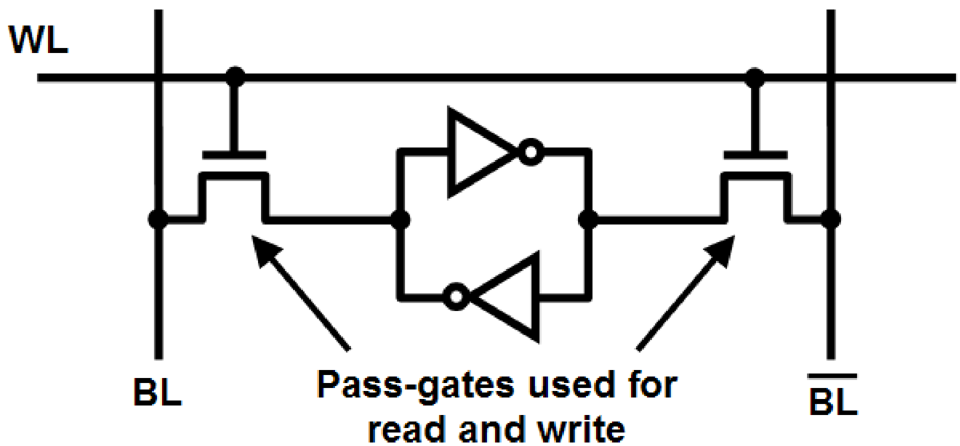

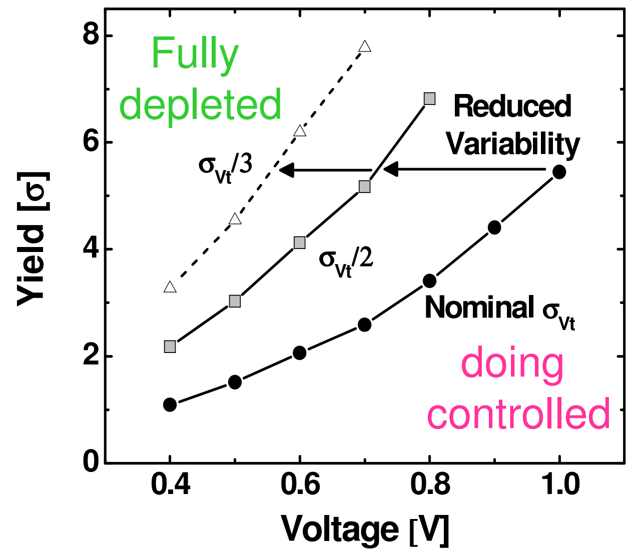

- Read and write operation required a delicate balance of cell devices and threshold mismatch will impact operation window, in particular at low voltages
- Device choices can significantly reduce operating voltage for 6T SRAM cell

Low voltage circuit operation - SRAM

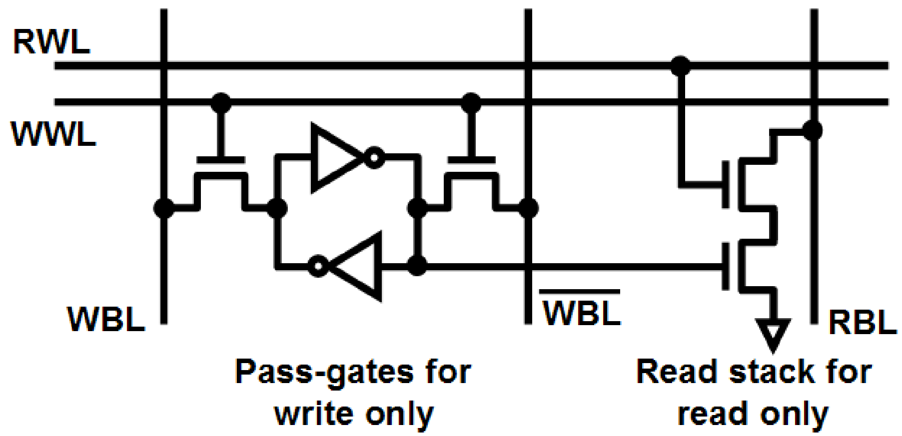

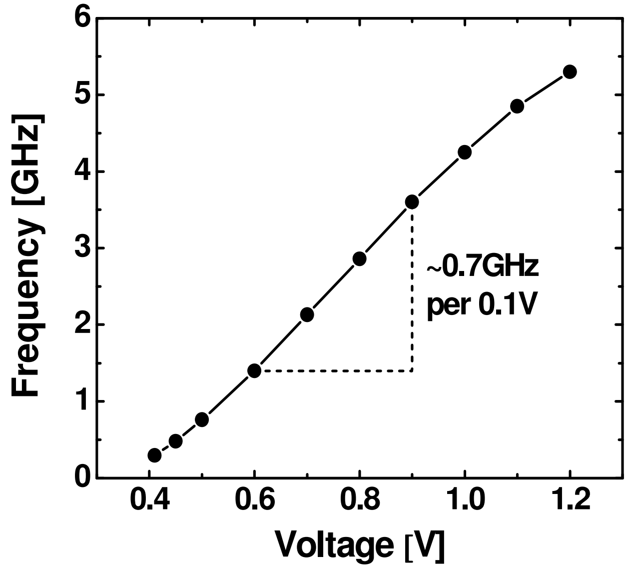

- 8T SRAM cell decouples read and write and opens up operating window for the cell
- Devices can be designed at minimum dimensions, which will benefit cell size

Off-chip bandwidth vs on-chip cache

- What is the real relationship?

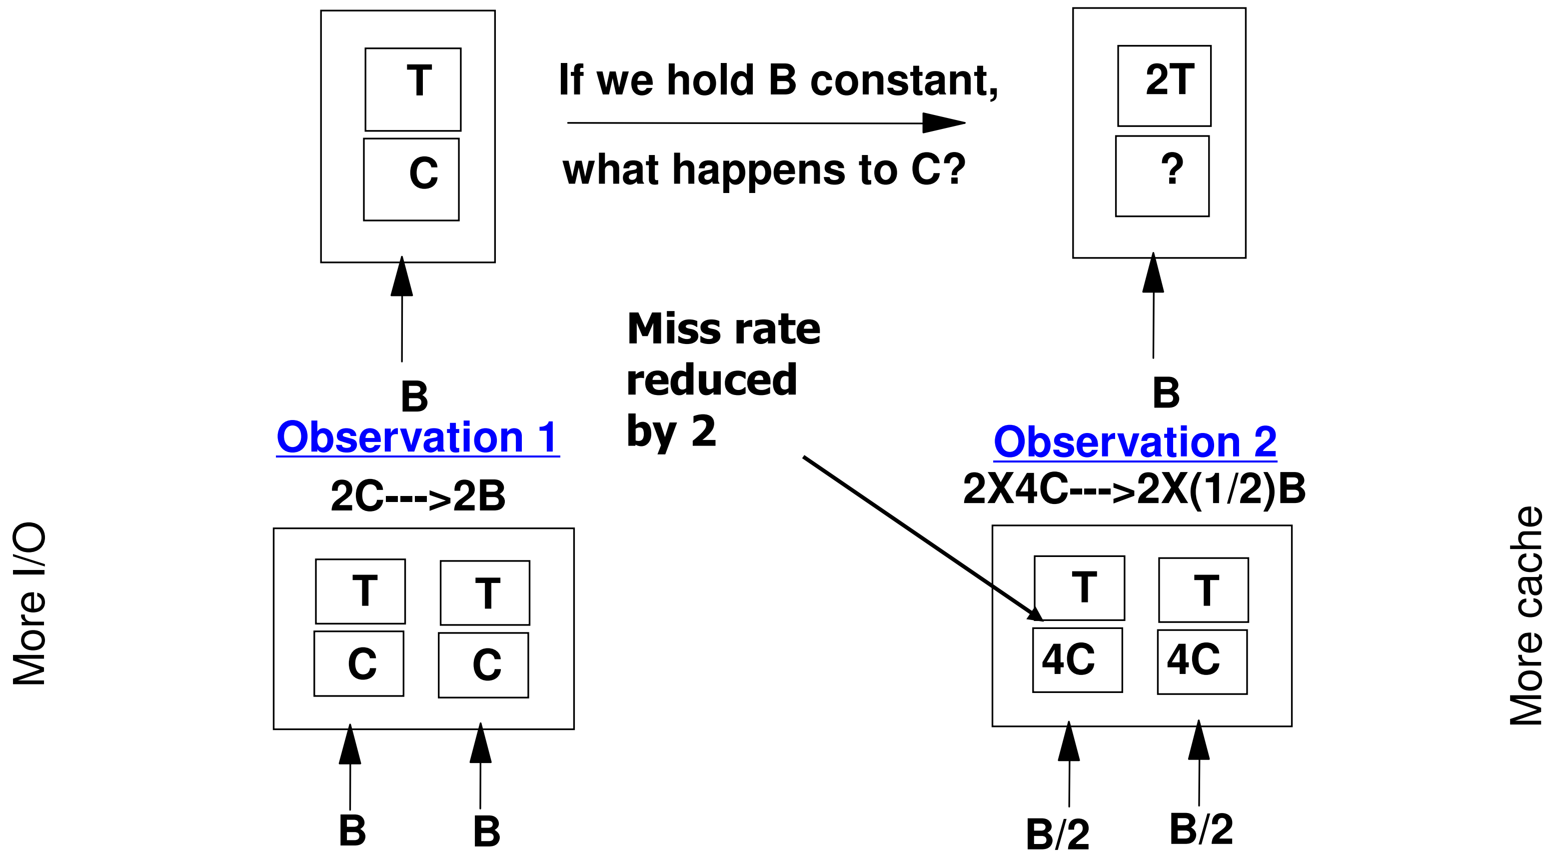

Off-chip bandwidth vs on-chip cache

- Architecture solution depends on application

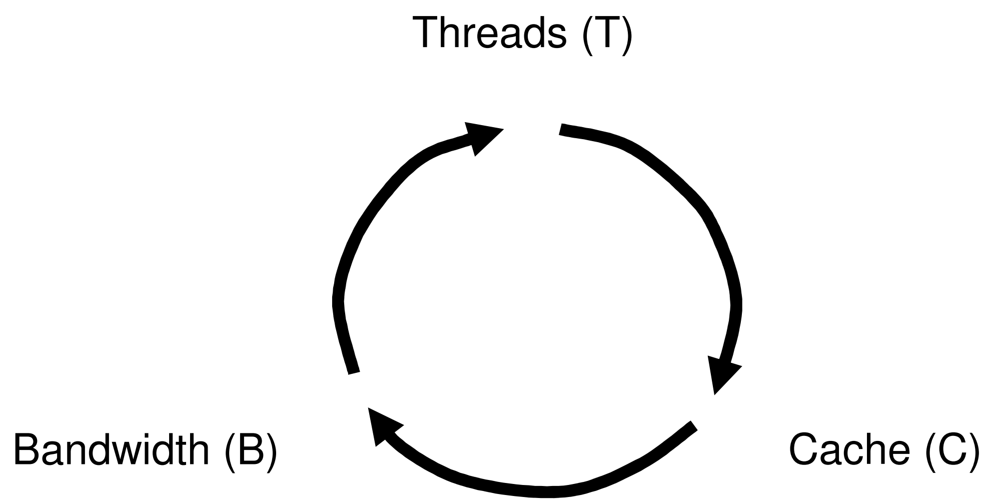

Emma's Law:

$$2^{\alpha+\beta} \mathbf{T} = (2^{\alpha} \mathbf{B}) \ \mathbf{X} \ (8^{\beta} \mathbf{C})$$

Low voltage signaling

- Low voltage (e.g. 0.25V) driver, significantly reduced power consumption on final driver: ~ $(0.25/0.925)^2 = 13.7$
  - Full Vdd on the gates of the nMOS drivers for high speed
- Gated diode receiver
  - Recover data back to full Vdd swing

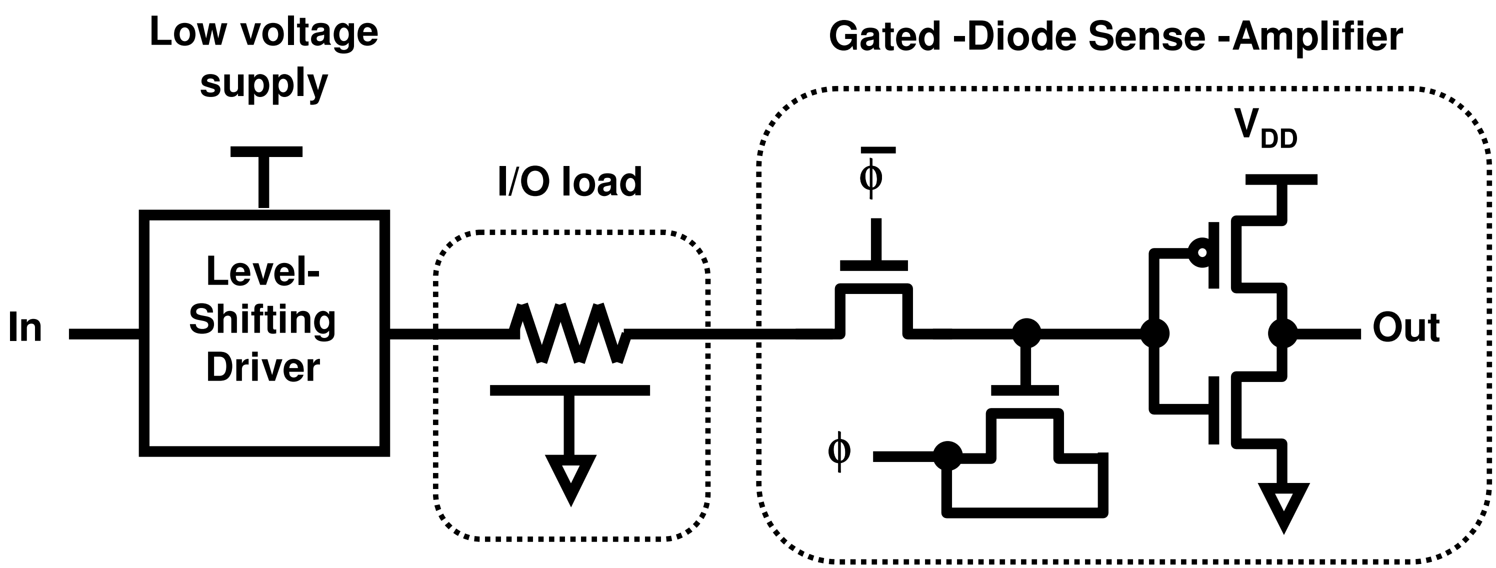

Three Dimensional Silicon Integration

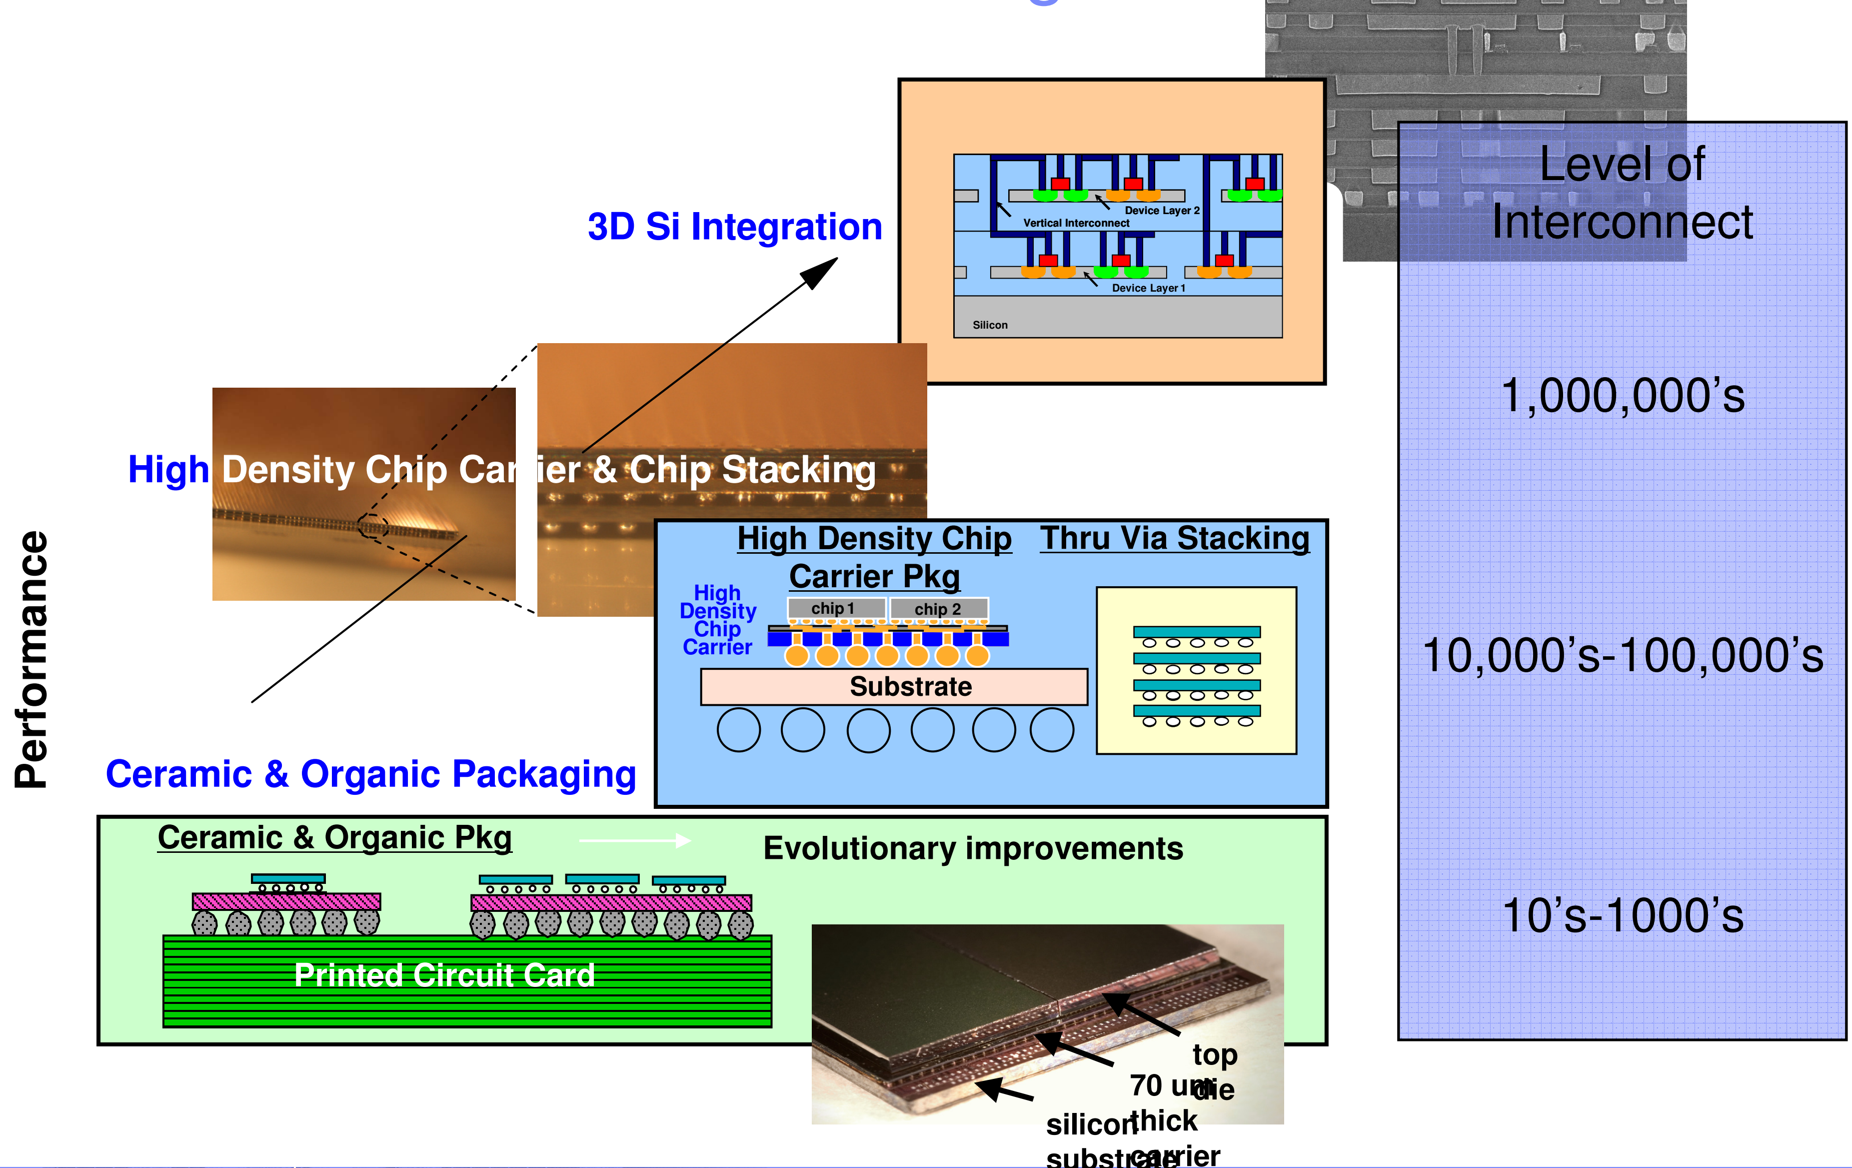

Low voltage scaling

Parameter	1→0.5V	Scaling Factor
Chip operating voltage	0.5	1
Chip operating frequency	0.33	κ
Device density	1	κ^2
Power density	0.056	κ^2
Power line energy loss	0.22–0.44	κ^2
Power line voltage instability	0.07–0.14	κ^2

Sub-Threshold Engineering

Beating the sub-threshold slope limit

$$T \sim \frac{E^2}{\sqrt{E_G}} e^{-B\frac{\sqrt{E_G^3}}{E}} \sim \frac{\sqrt{E_G^3}}{l^2} e^{-B\sqrt{E_G}\,l}$$

$$S \sim \frac{V_{gs}^2}{2V_{gs} + D(V_{gs}, V_{ds})}$$

Krishna K. Bhwalka et al. 2005

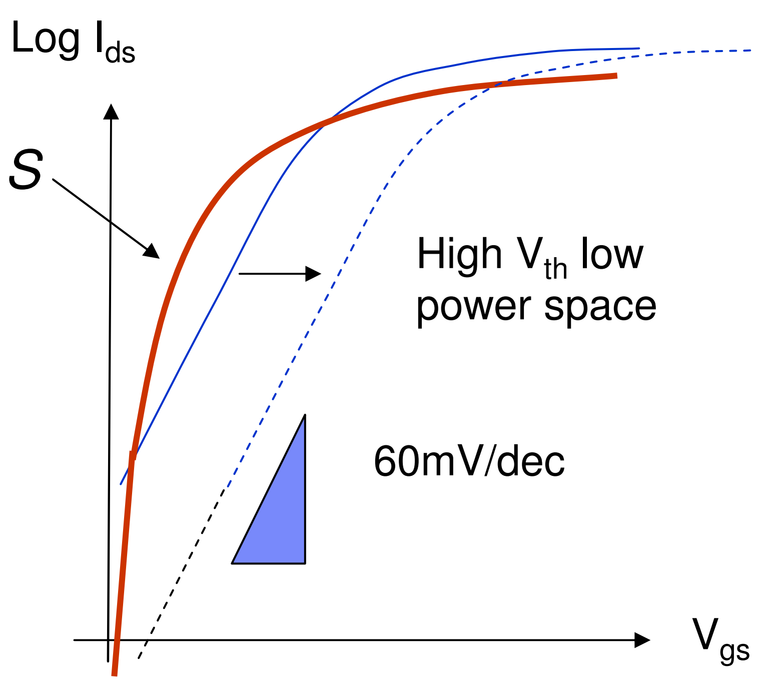

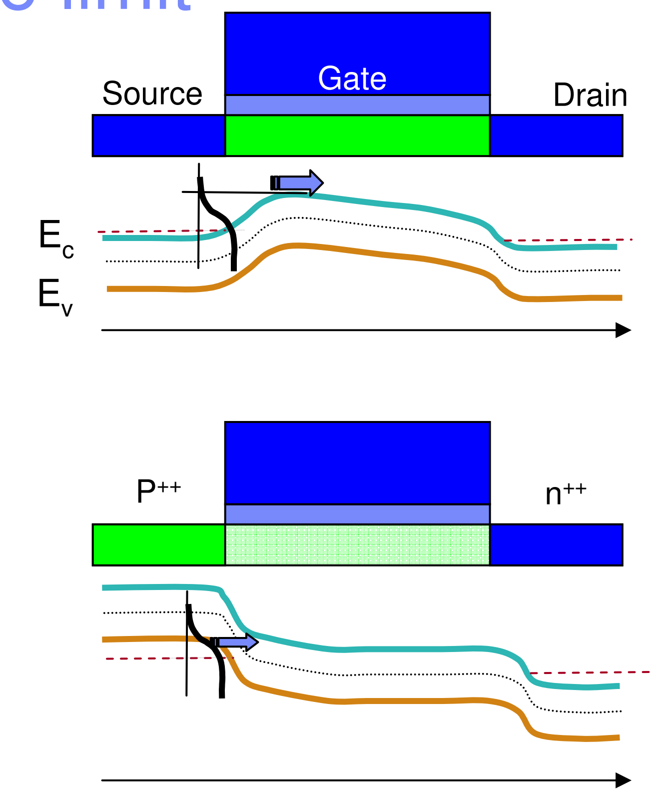

- Tunnel FETs show strong voltage dependence of sub-threshold slope
- On current not yet on par with conventional high performance FETs at comparable voltages

Beating the sub-threshold slope limit

$$T \sim \frac{E^2}{\sqrt{E_G}} e^{-B\frac{\sqrt{E_G^3}}{E}} \sim \frac{\sqrt{E_G^3}}{l^2} e^{-B\sqrt{E_G}l}$$

$$S \sim \frac{V_{gs}^2}{2V_{gs} + D(V_{gs}, V_{ds})}$$

Krishna K. Bhwalka et al. 2005

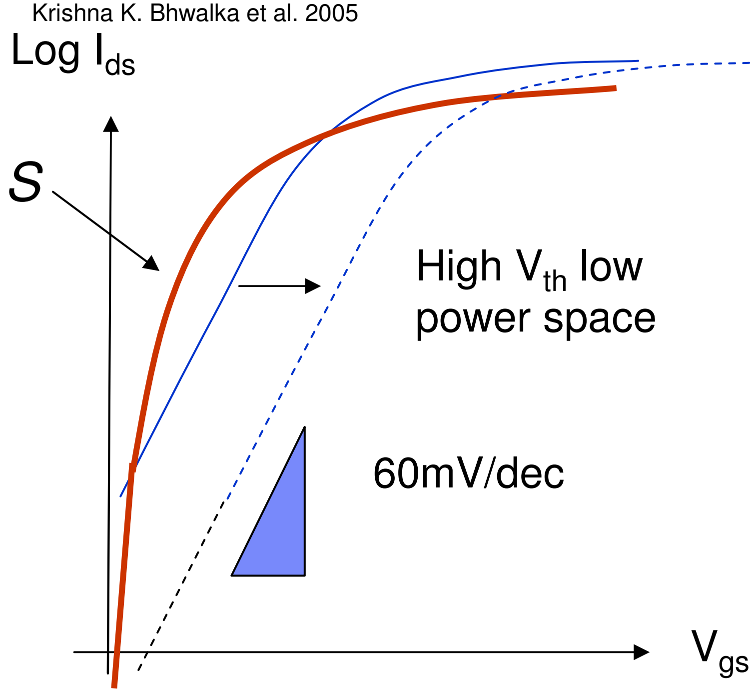

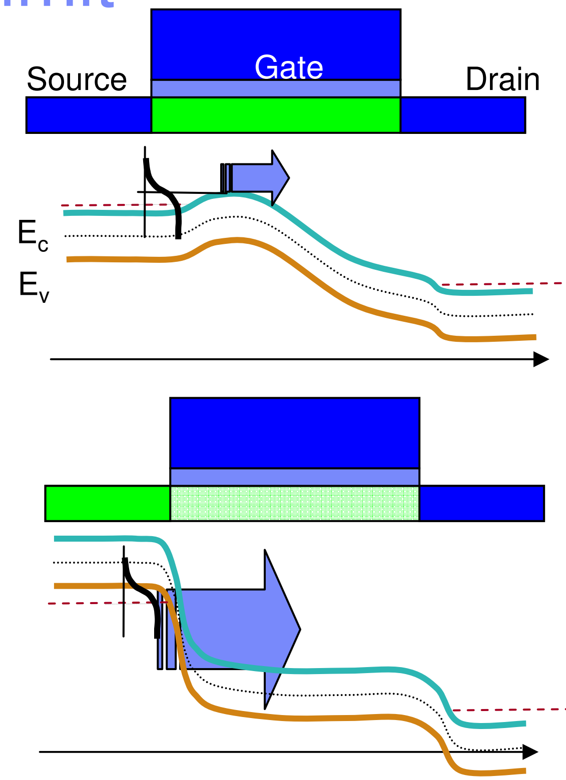

- Tunnel FETs show strong voltage dependence of sub-threshold slope
- On current not yet on par with conventional high performance FETs at comparable voltages

Supply voltage for max performance

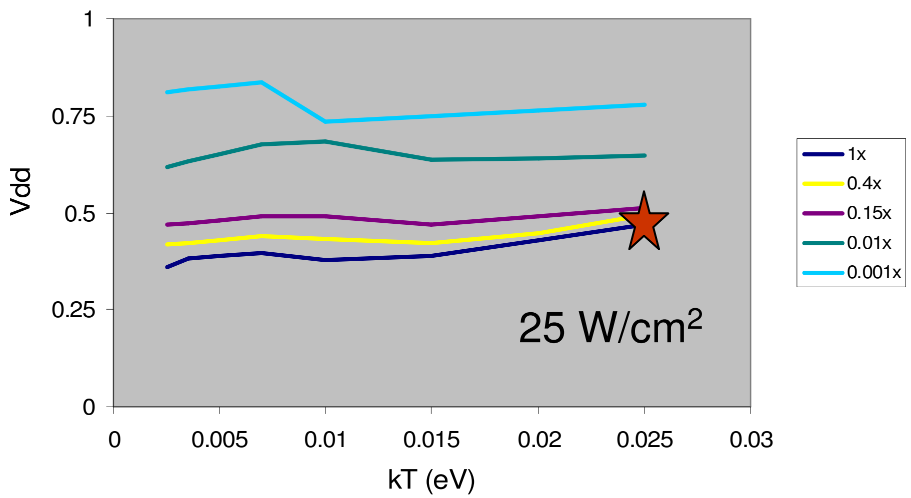

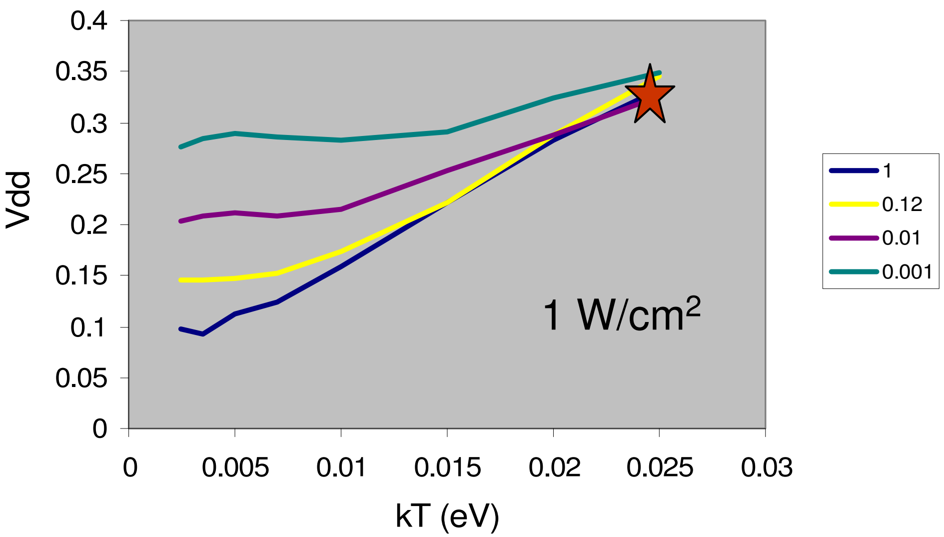

- For high power parts there is little advantage for supply voltage reduction due to steep sub threshold devices
- For low power parts there is a *significant power supply reduction* potential if *drive current equivalent* to conventional FET can be achieved

SS Engineering - 22nm node

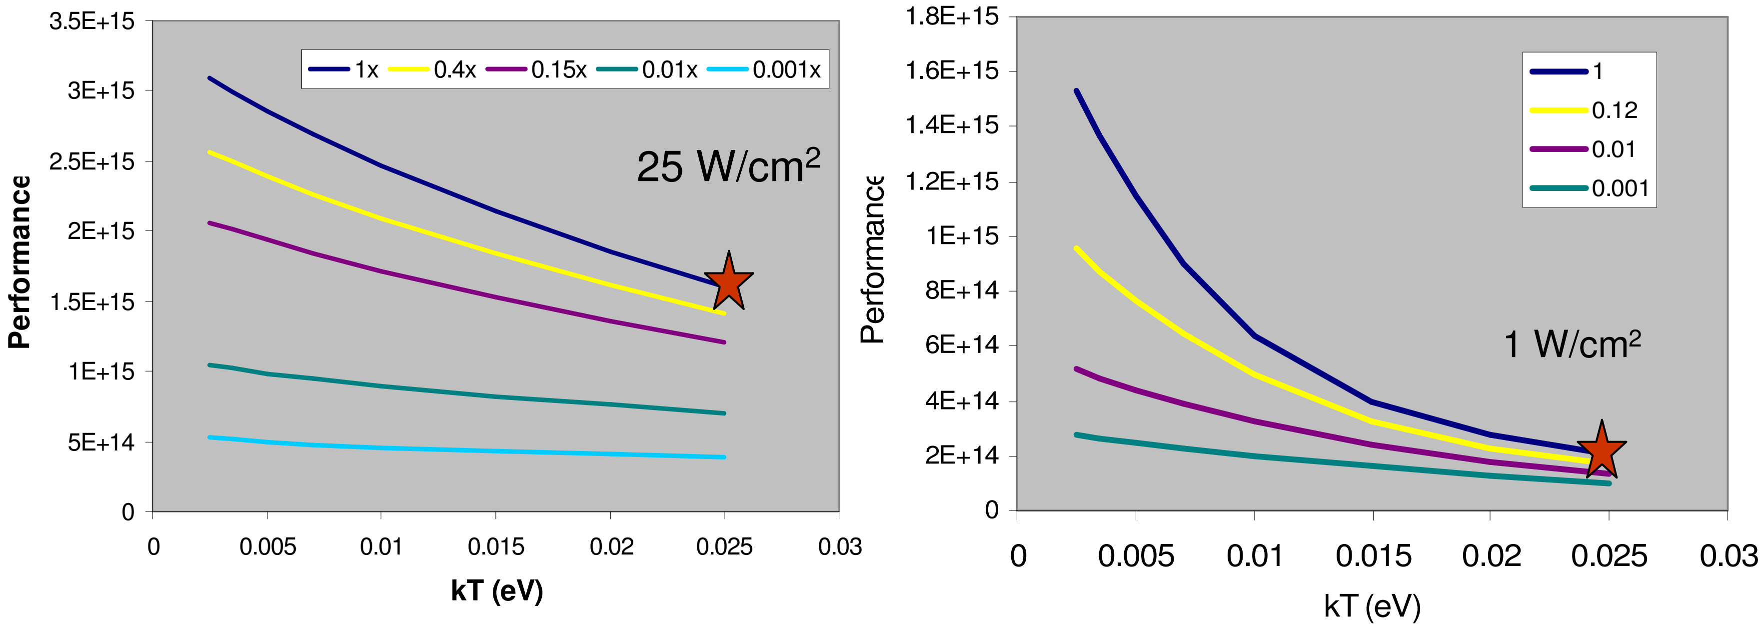

- Performance SS slope trade-off window small for high power case
- For low power space there is a large trade-off window performance versus SS-slope

On-current vs SS-slope, for constant performance

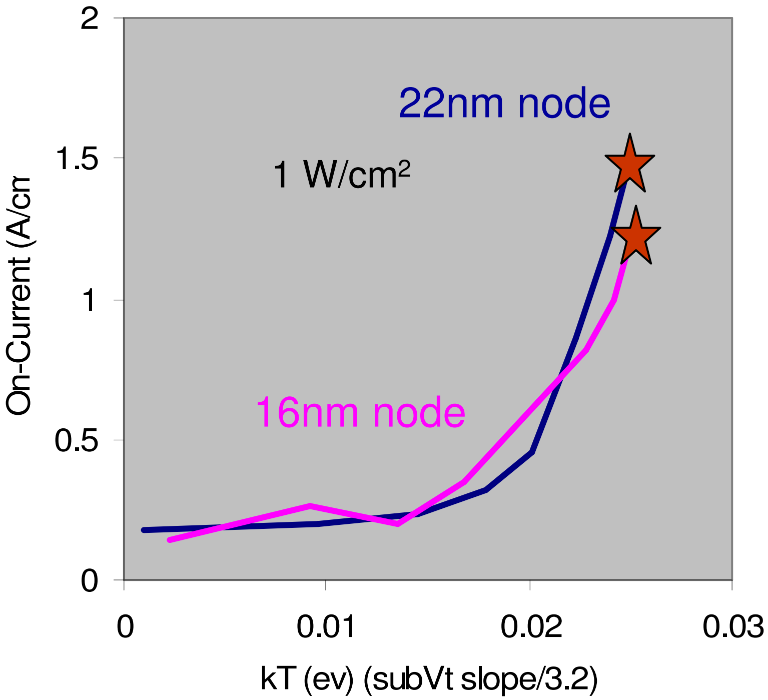

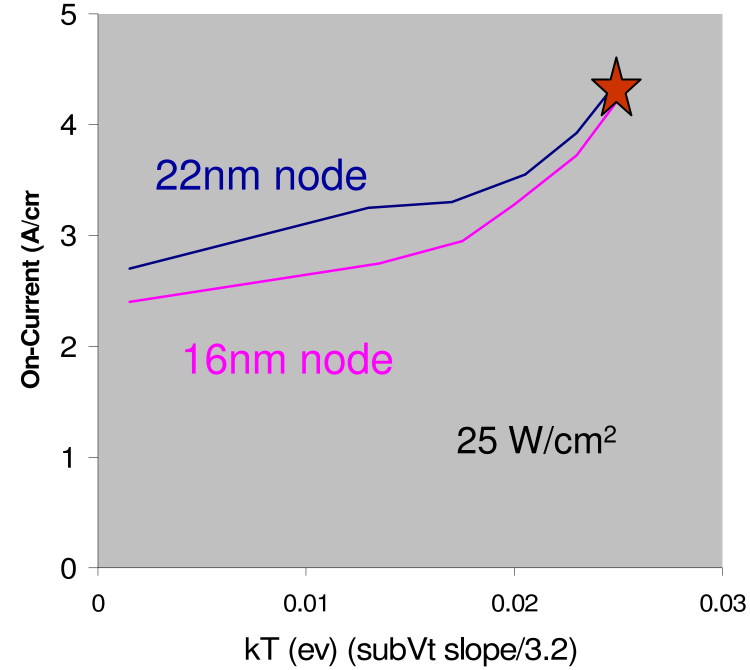

- High power/ performance: can only tolerate 30% drive current degradation
- Low power / performance: can tolerate x 10 degradation of drive current

Adiabatic Computing

Adiabatic charging

How much energy must be dissipated to charge a capacitor?

Abrupt method

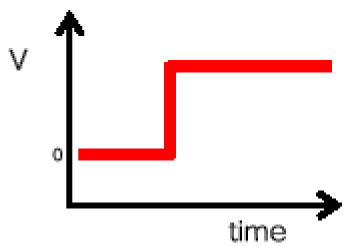

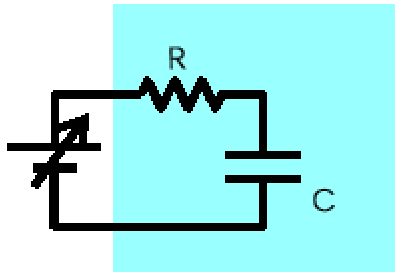

$$E = \frac{1}{2} CV^2$$

Quasi-static Charging (aka 'adiaba

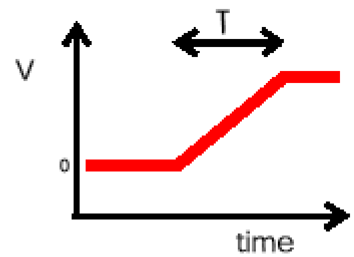

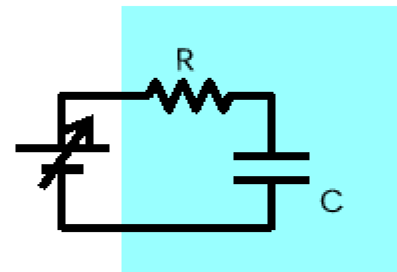

$$E = \frac{1}{2} CV^2 \left(\frac{2RC}{T}\right)$$

This assumes T >> RC.

(There's an extra factor of $\pi^2/8$ for a sinusoidal ramp.)

Quasi-static Charging + Superconductivity

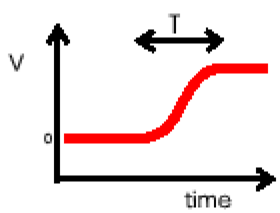

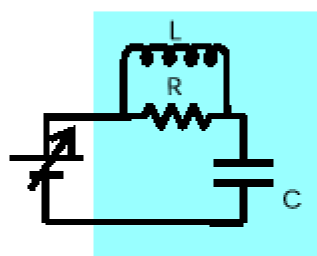

Charging through a superconductor, which behaves as an inductor and resistor in parallel.

$$E = \frac{\pi^4}{8} CV^2 \frac{RC(L/R)^2}{T^3}$$

This assumes T >> RC and T >> L/R.

(To implement logic in this system would require superconducting FETs. Such FETs are possible, but there have been very few experimental results.)

Adiabatic computing

- All logic transitions must be directly driven by a clock waveform passing through FETs, Output is cycled back to input
  - Transitions cannot ripple through statically powered gates as in conventional logic.
- The ramp rate of the clock waveforms must be low to save energy.
- Try never to turn on an FET while there is a voltage difference between source and drain, since this would result in dissipation.
  - Requires multiple clock waveforms

Clock Synchronization

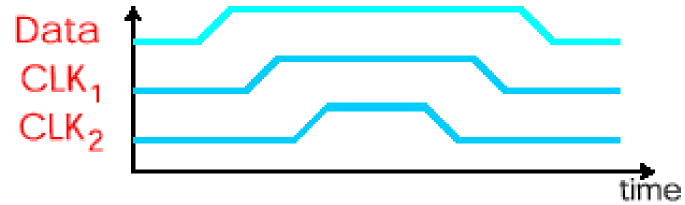

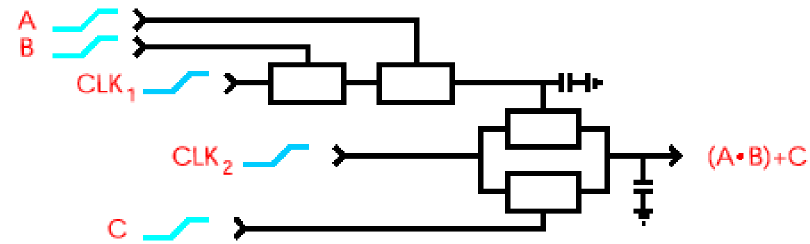

J. S. Hall 'Electroid' Switches)

Adiabatic computing

- If A = ‘1’ and B = ‘0’, or A = ‘0’ ar B = ‘1’: Charge oscillates betwee input and output.
  - In each cycle of the resonator (clock), the logic signal is created and then removed fro the output node capacitor.
- If A = B = ‘0’ or A = B = ‘1’: Output decoupled from input, and output kept to ground
  - Clock participates energy, per half cycle

$$E_{diss} = \frac{\pi}{8}\frac{C_L V^2}{Q} = \frac{\pi^2}{4} f R C_L^2 V^2$$

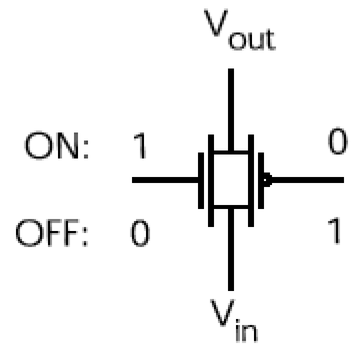

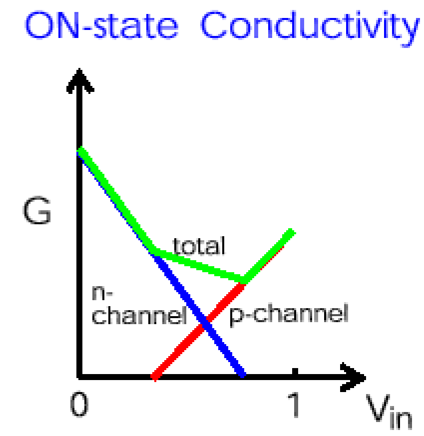

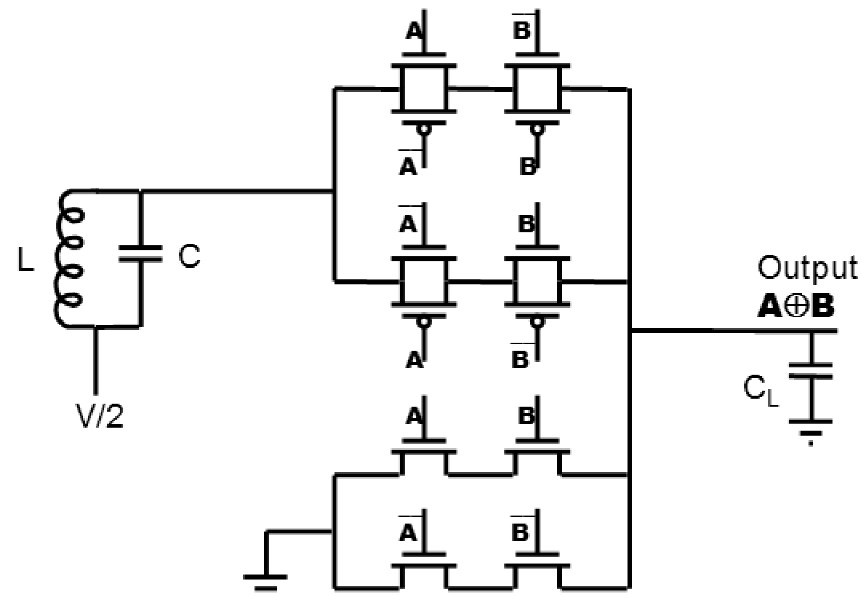

Adiabatic CMOS penalty factors relative to conventional CMOS

	#FETs	CPI	CV^2 Energy	comment
Dual Rail	2x	-	2x	Twice as many nodes
Logic density	¾x	-	¾x	Dual rail => no inverters
Uncompute	1.5x	1.5x	2x	Mixed retractile and pipeline
Generate control signals	4/3 x	-	1.5x	
Clock supply	1x	7x	1.5x	No ripple: drive every transition
$V_{DD} > 2V_T$	1x	-	4x	~double supply voltage
TOTALS	3x	10x	27x	

Incomplete reversibility also sets a lower limit on energy savings.

Compare to adiabatic to CMOS

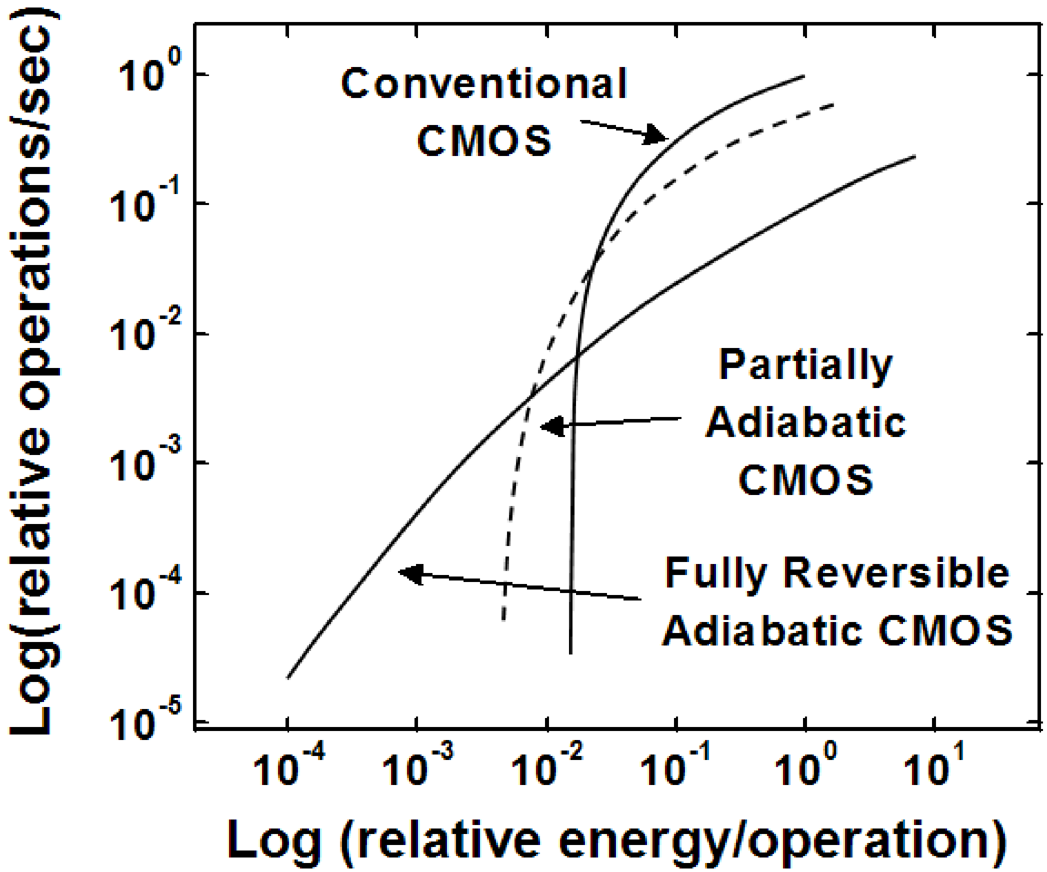

Summary

- Classical CMOS has still life left
  - Fully depleted devices will offer additional knobs for gate length scaling
  - Careful analysis of power delivery and I/Os show opportunities to reduce power losses in the system
  - Low voltage operation is biggest knob for power reduction, however need massive parallelism to get performance back and watch out for a low voltage cache solution
- Sub-threshold engineering
  - Steep sub-threshold devices have their place in the low power applications
  - Need to push devices to obtain at least 10% on current as conventional device
- Adiabatic computing
  - Concepts are there, not clear where application space is
  - This needs further research

Acknowledgement

Leland Chang

Steve Koester

Brian Ji

David Frank

Matt Wordemann