inst.eecs.berkeley.edu/~eecs251b

EECS251B : Advanced Digital Circuits and Systems

Lecture 26 – Finale

Borivoje Nikolić

Startup Rivos raises $250 million to develop RISC-V AI chips

April 16 (Reuters) - Chip startup Rivos said on Tuesday it raised $250 million in a funding round that will enable it to manufacture its first server chip geared for artificial intelligence.
Rivos is tight-lipped about the specifics of the product, but has disclosed that its plans include designing chips based on the RISC-V architecture, which is an open source alternative to the architectures made by Arm, Intel, and Advanced Micro Devices. Instruction set architectures such as RISC-V are the building blocks of semiconductor designs, and using the open source alternative means Rivos does not have to pay a license fee to Arm.

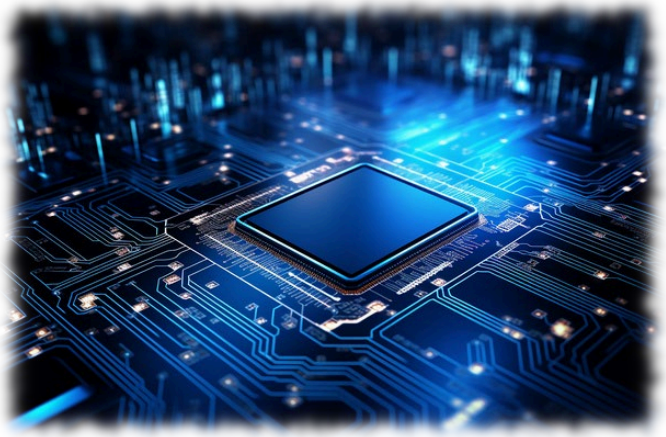

rivosinc.com

Announcements

- Homework 5 due today
- Project
  - Pay attention to integration with other teams!
  - Final presentations: May 2, 9am-12pm, BWRC
- Final exam: Thursday, April 26, in class
  - Two classrooms!

Delay-Locked Loops

Clock Generation

Delay-Locked Loop (Delay Line Based)

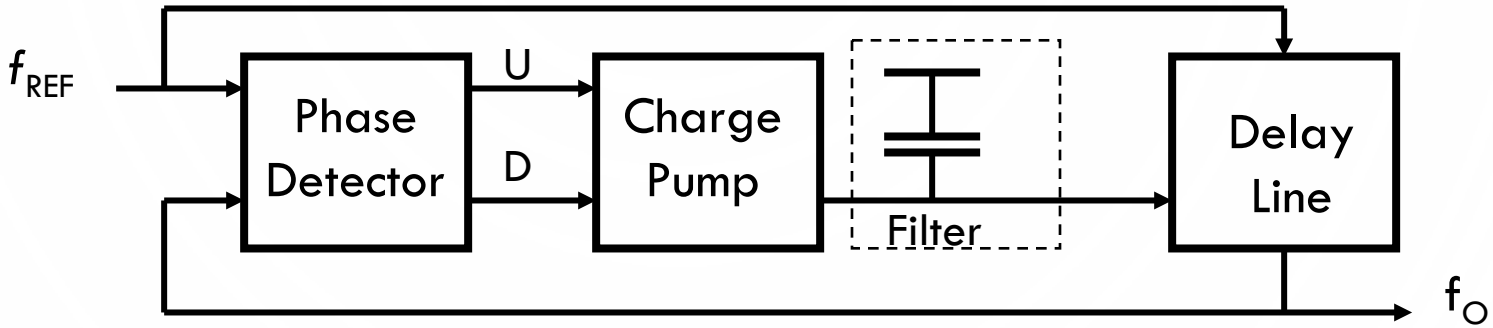

Phase-Locked Loop (VCO/DCO-Based)

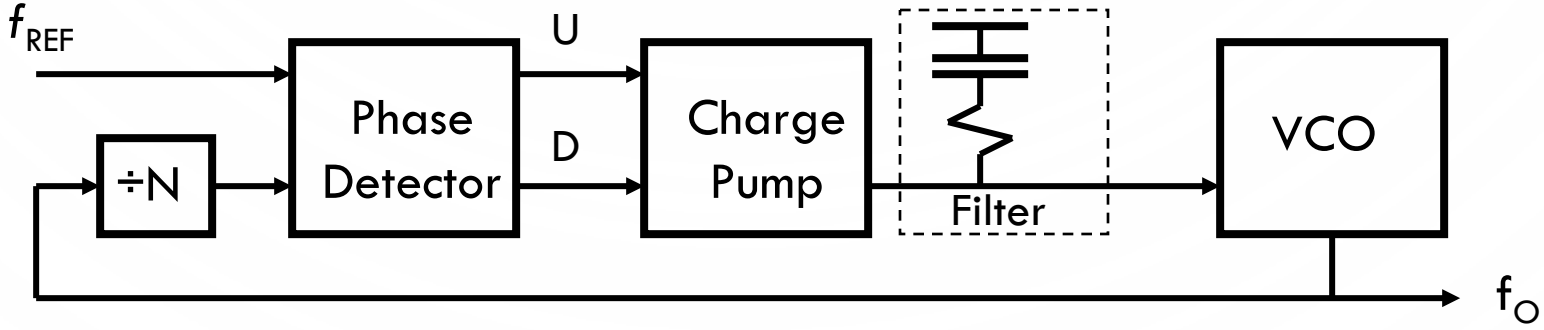

Delay-Locked Loop

- First order loop: inherently stable
- No filtering of input jitter
- Constant frequency (no synthesis)
- No phase error accumulation

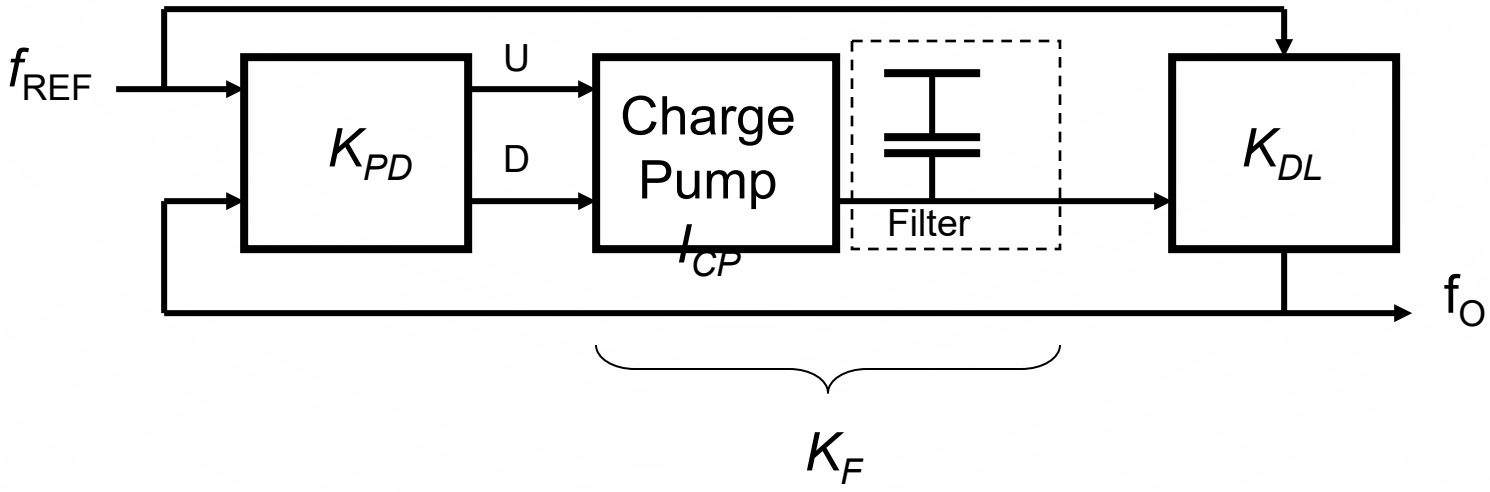

DLL Locking

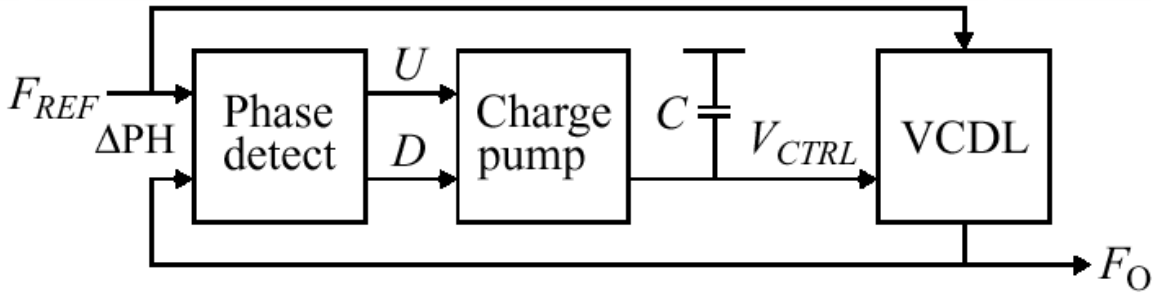

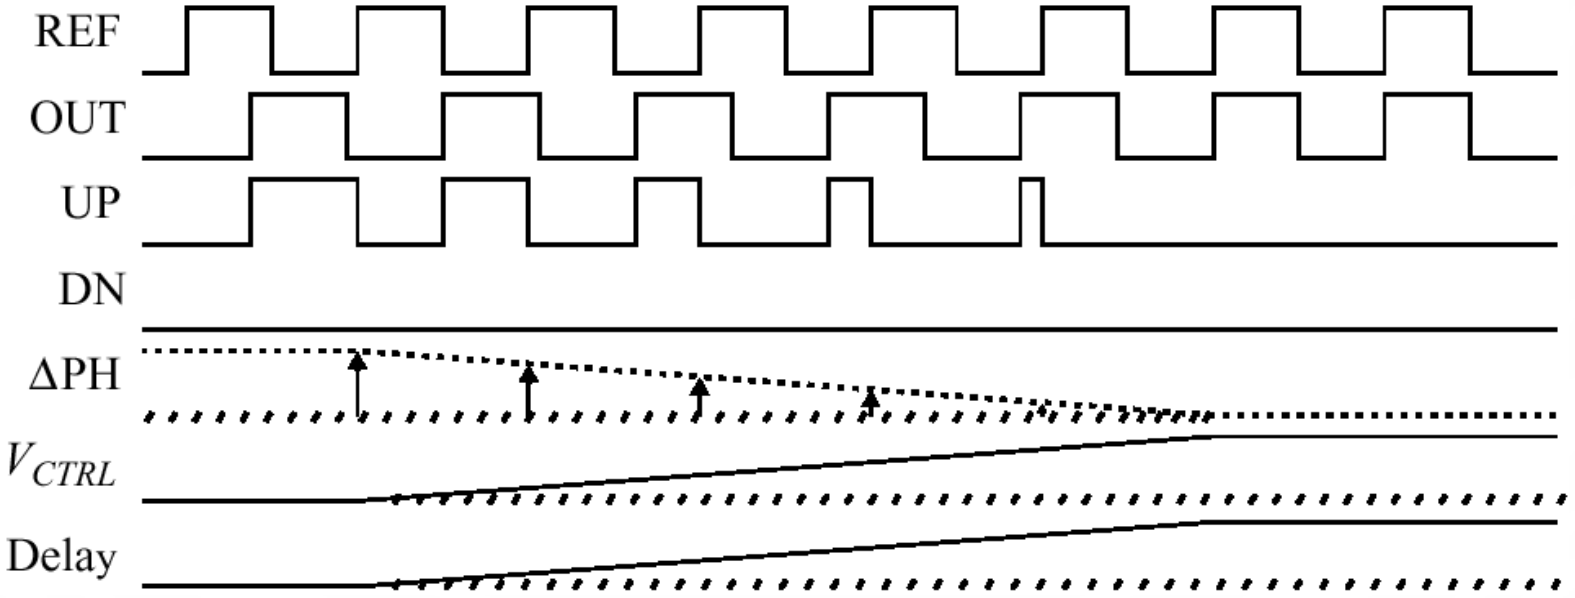

Courtesy of IEEE Press, New York. © 2000

Delay-Locked Loop

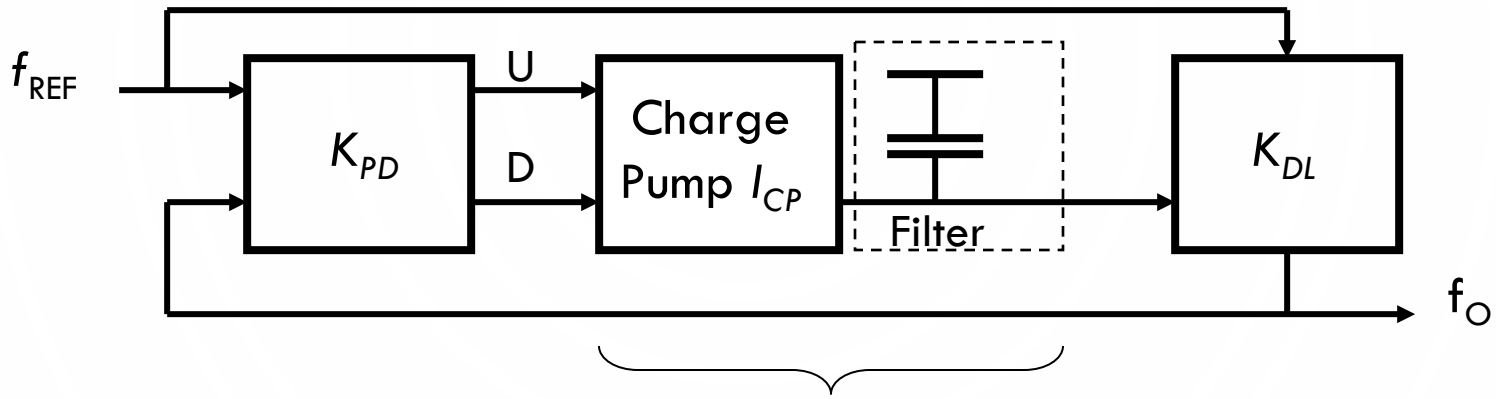

- Open loop transfer function

$$\frac{D_O(s)}{D_I(s)-D_O(s)} = K_{PD}\frac{1}{sC}I_{CP}K_{DL}F_{REF} = \frac{1}{s}K_{PD}K_FK_{DL}$$

- Closed loop transfer function

$$H(s) = \frac{D_O(s)}{D_I(s)} = \frac{K_{PD}K_FK_{DL}}{s + K_{PD}K_FK_{DL}}$$

Delay-Locked Loop

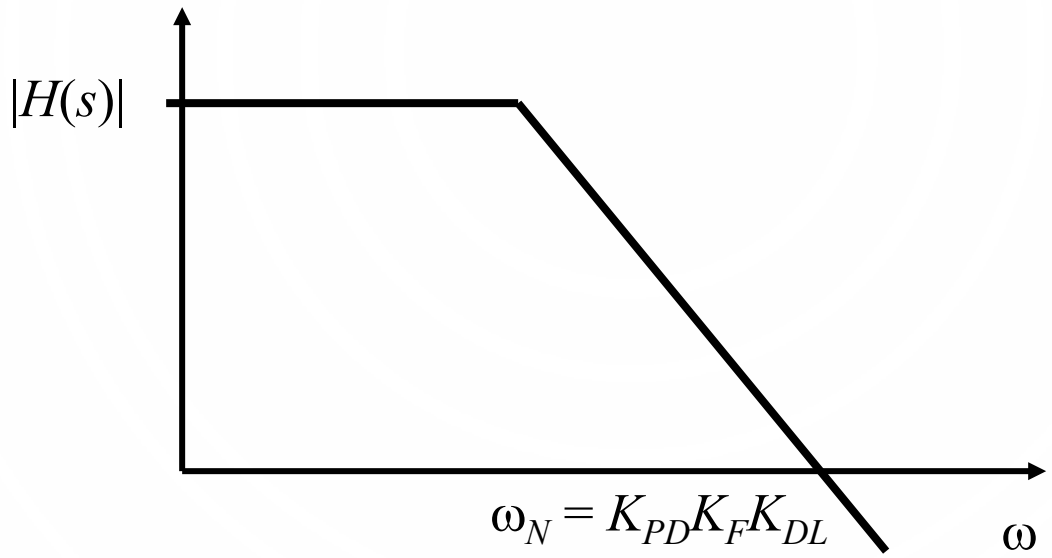

- ω_N > an order of magnitude below F_{REF}
- Use of DLLs requires low-jitter input
- VCDL must span adequate delay range + reset to min delay
- Noise sources:
  - Delay line (Supply sensitivity)
  - Clock buffers that follow
  - Device noise (small)

Clock Distribution

Clock Distribution

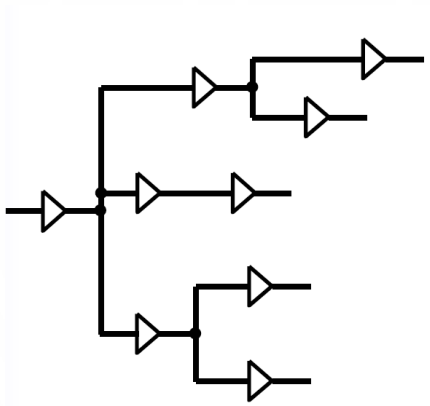

Tree

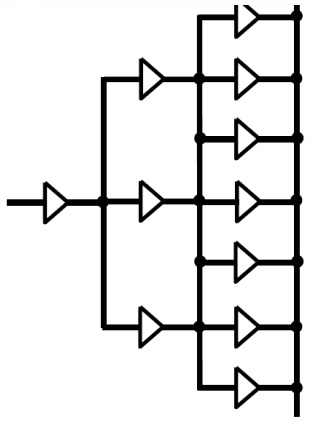

Mesh

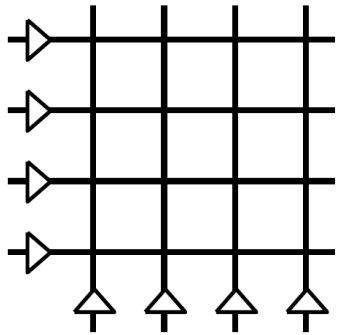

Grid

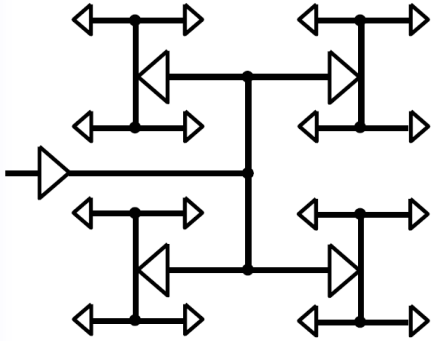

H-Tree

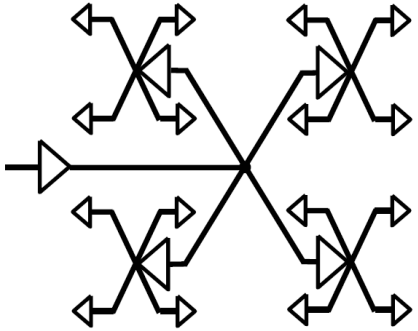

X-Tree

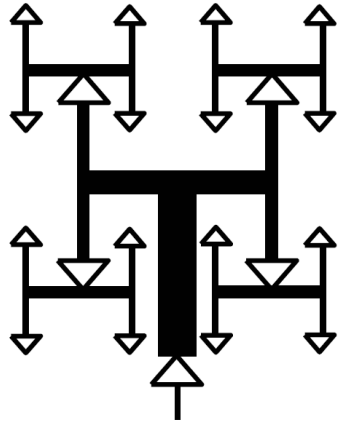

Tapered H-Tree

Example (Older) Clock System

- IBM Power 4

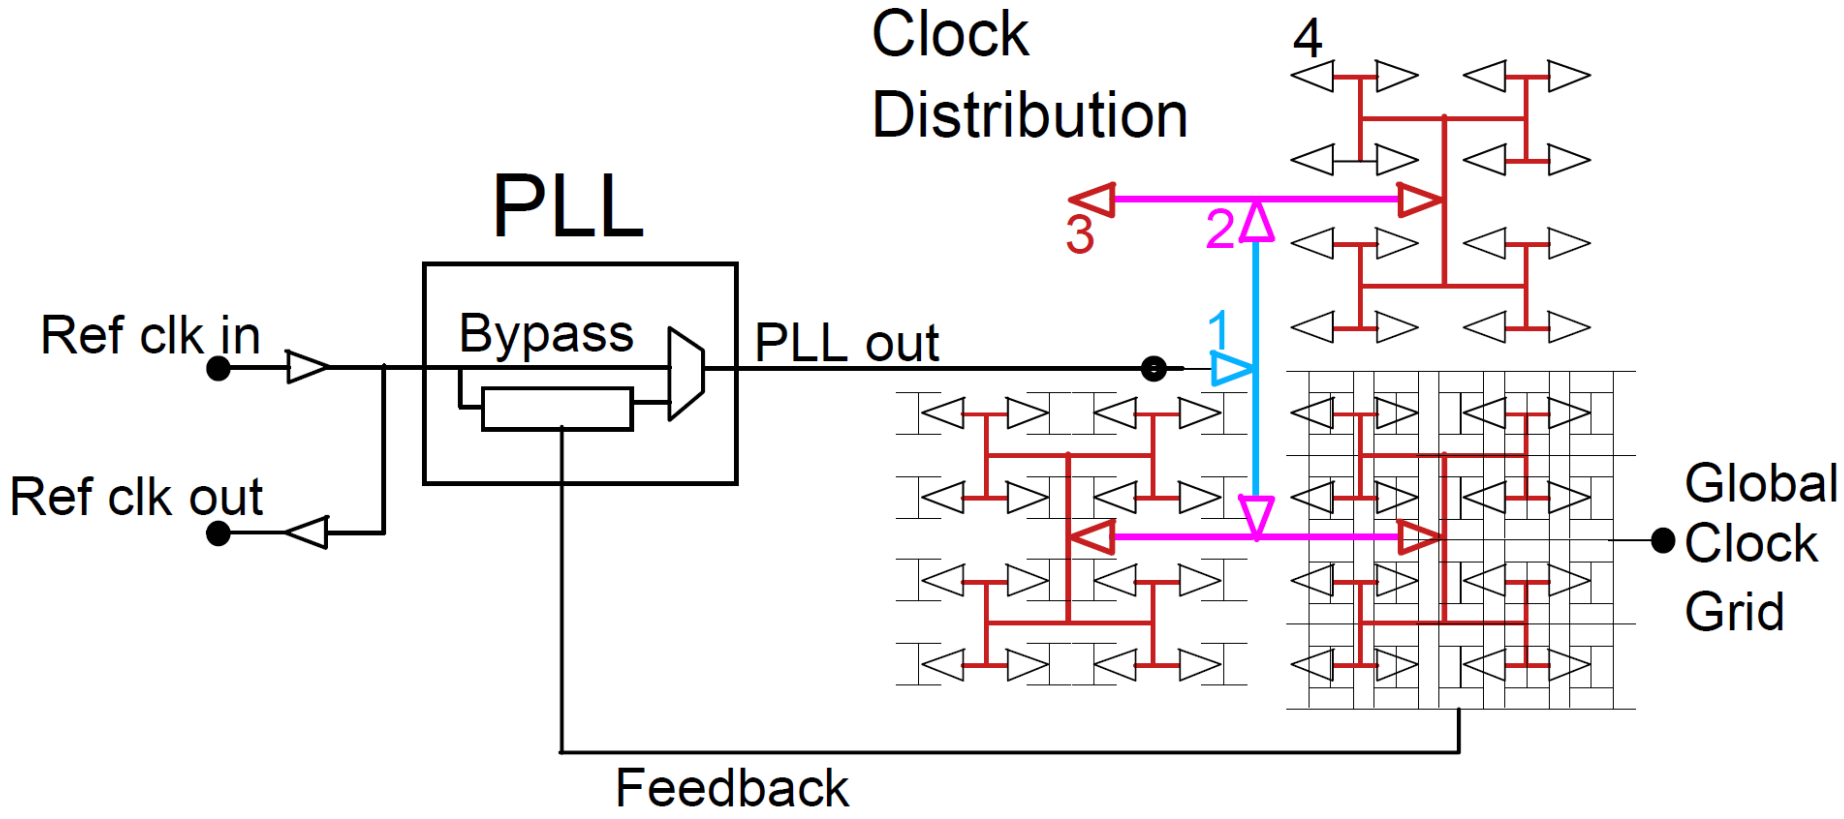

Restle, ISSCC'02

Clock Grid

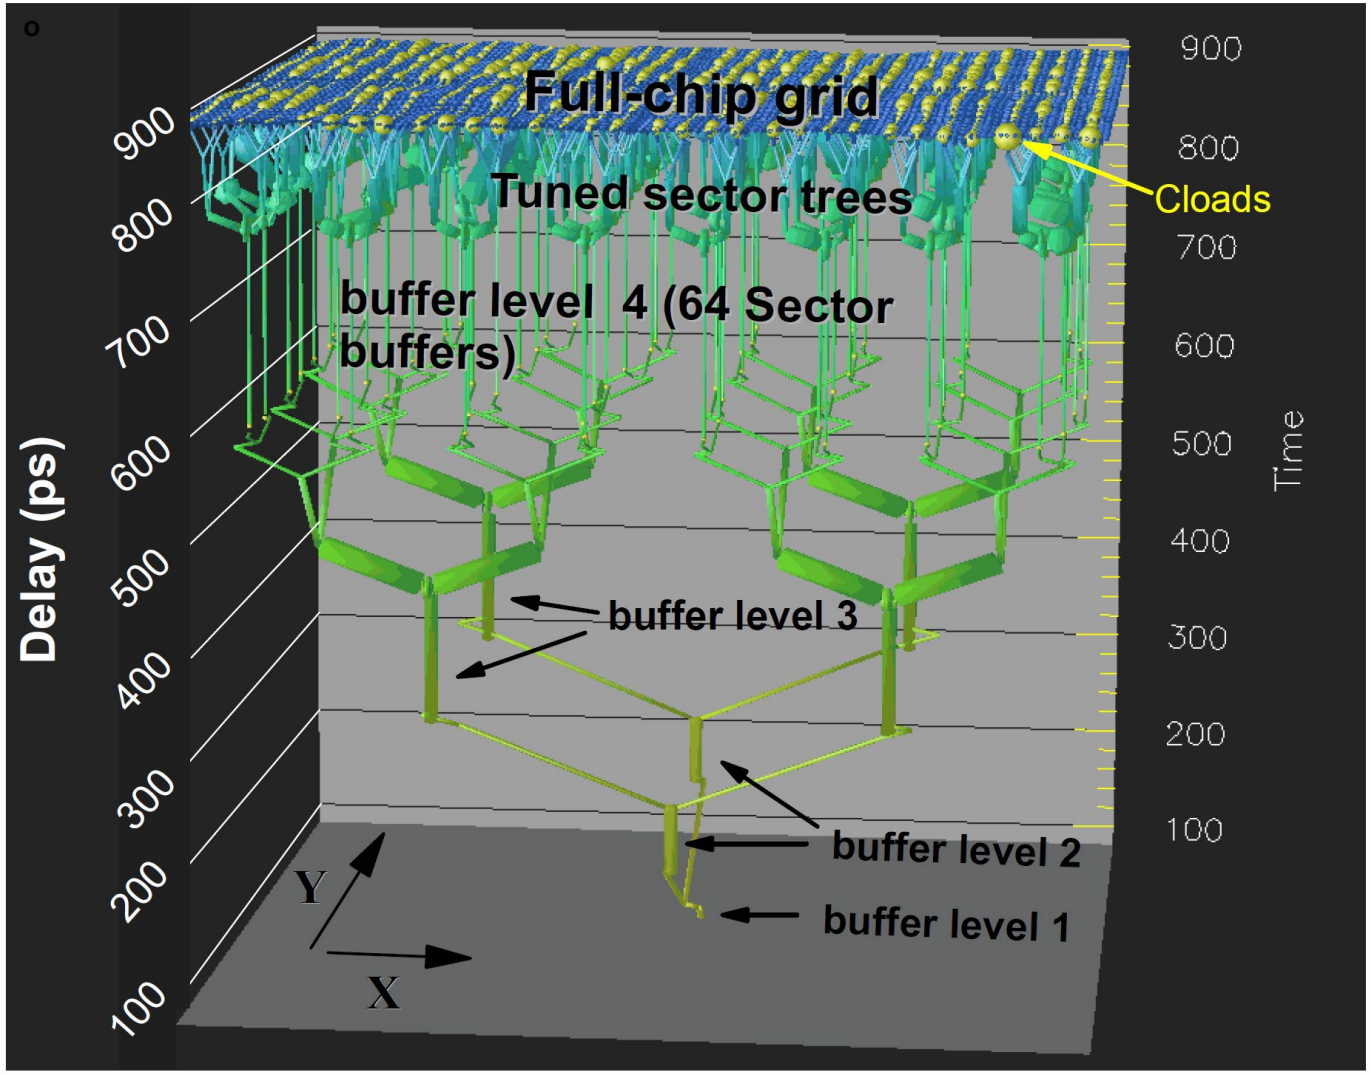
Full-chip grid
Tuned sector trees
Cloads
buffer level 4 (64 Sector buffers)
buffer level 3
buffer level 2
buffer level 1
Delay (ps)
Time
900
800
700
600
500
400
300
200
100
Y
X

One PLL with multiple DLLs

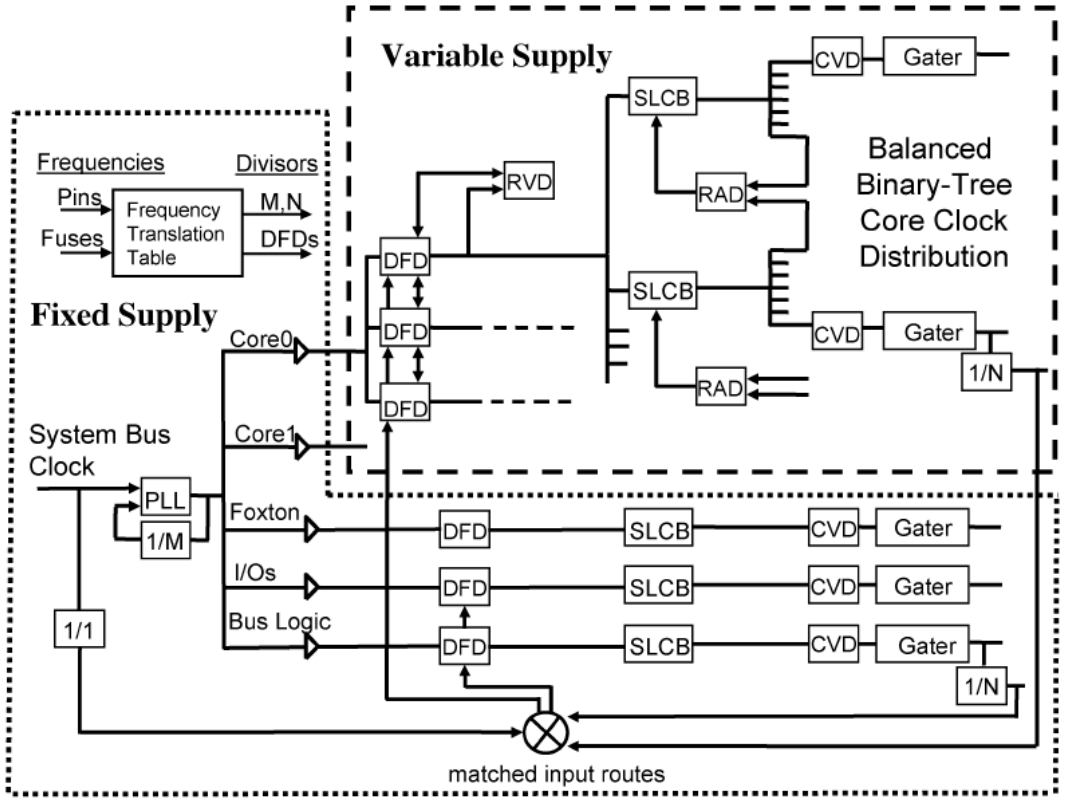

- Single PLL, and two cores vary frequency through digital frequency dividers (DFDs) and DLLs
  - SLCB: Second-Level Clock Buffer
  - CVD: Clock Vernier Device – fine (static) delay tuning

Fischer, JSSC 1/06

Deskewing and Synchronization

Clock Domain Synchronization

Type	Frequency	Phase
Synchronous	Same	Same
Mesochronous	Same	Constant offset
Plesiochronous	Small difference	Slowly varying
Asynchronous	Different	Arbitrary

Deskew System (Mesochronous)

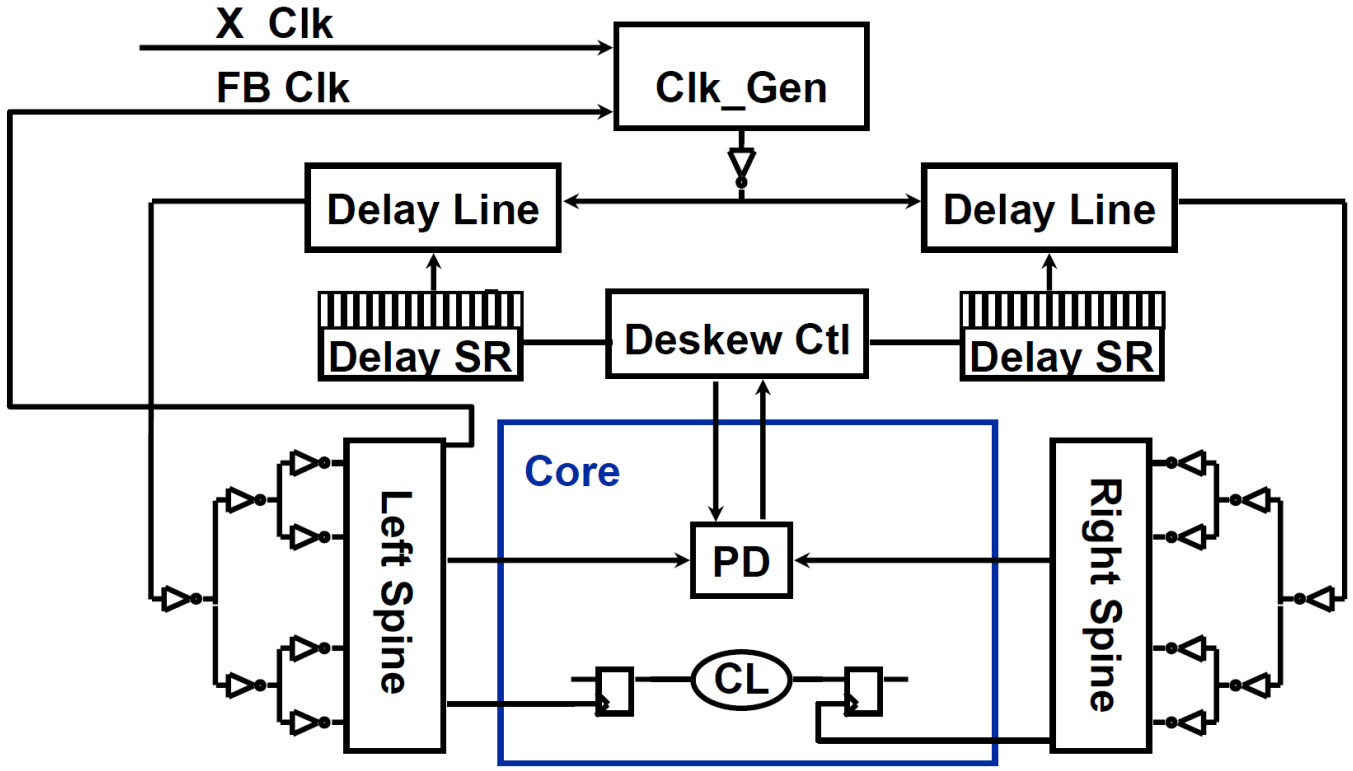

Geannopoulos, ISSCC'98

Deskew System

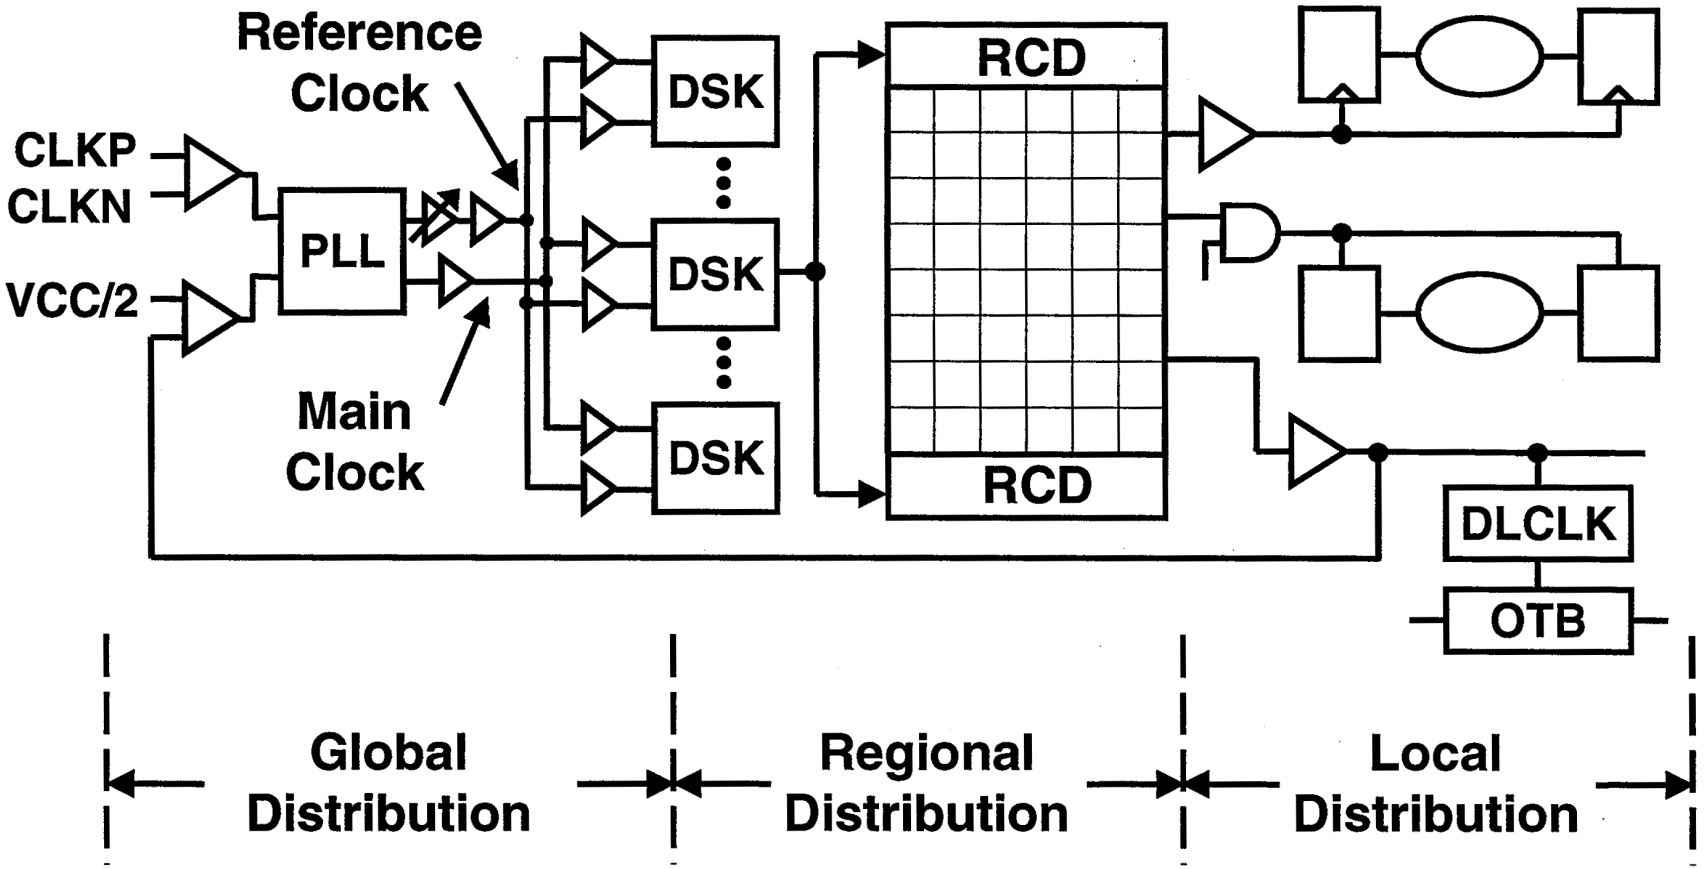

- Rusu, ISSCC'00

Deskew Buffer

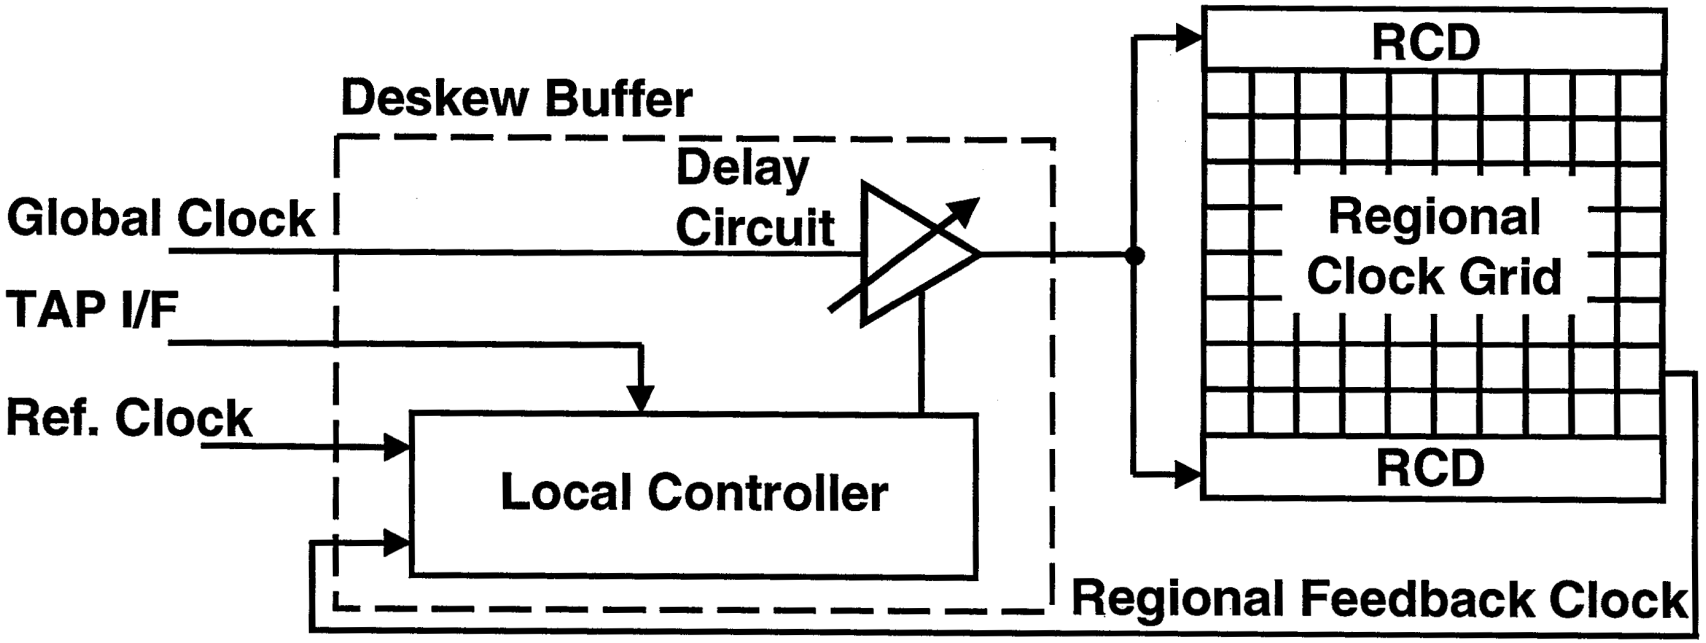

- Essentially a DLL to align regional clock with ref. clock

Clock Subsystem

- Intel Xeon – Bowhill, ISSCC'15
  - Independent clocks for 4-18 cores
- Self-biased ring-osc (SB) and LC PLLs

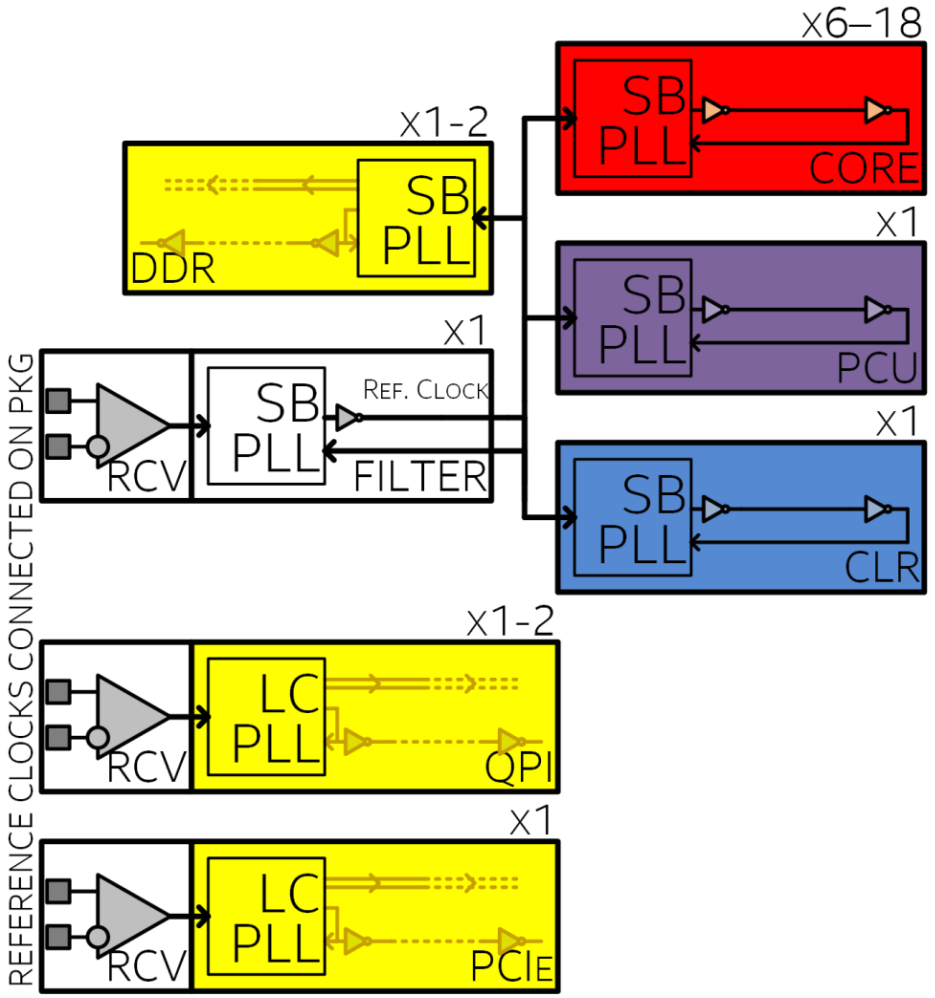

Clock Domain Crossings

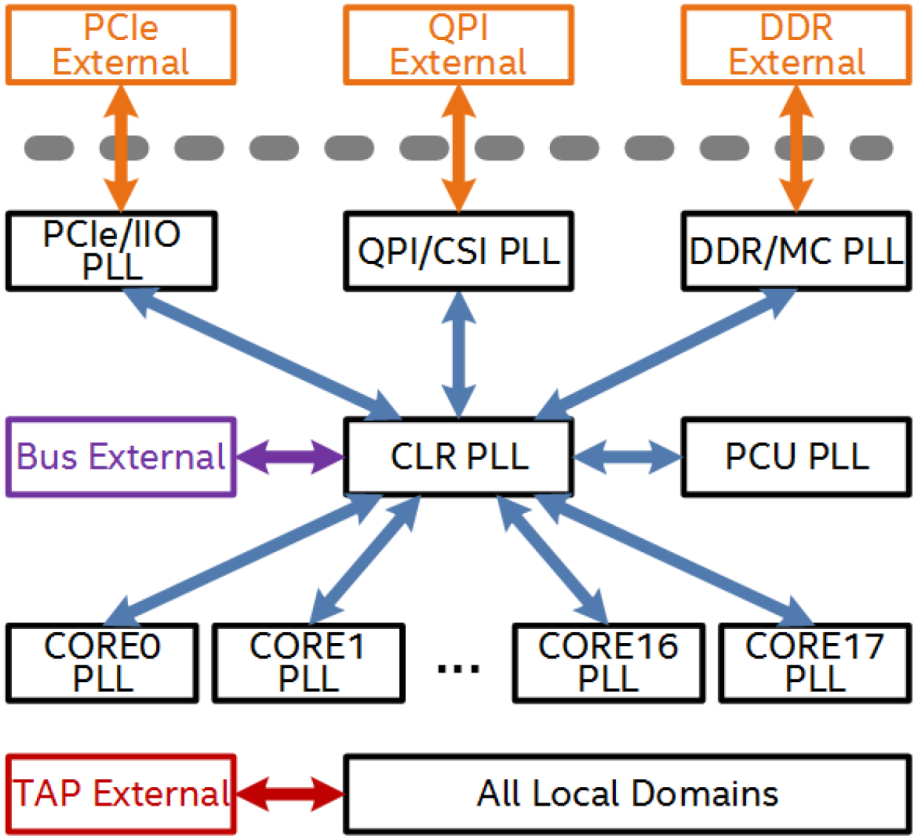

- Bowhill, ISSCC'15

Brute-Force Synchronizer

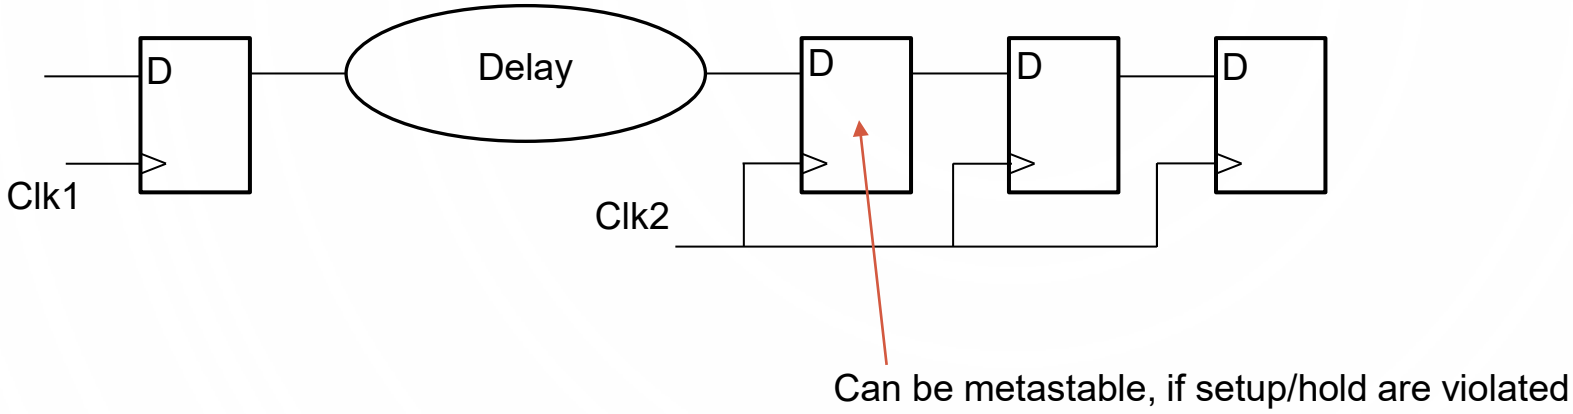

- Cascaded flip-flops reduce the probability of metastability

Clock Crossing FIFOs

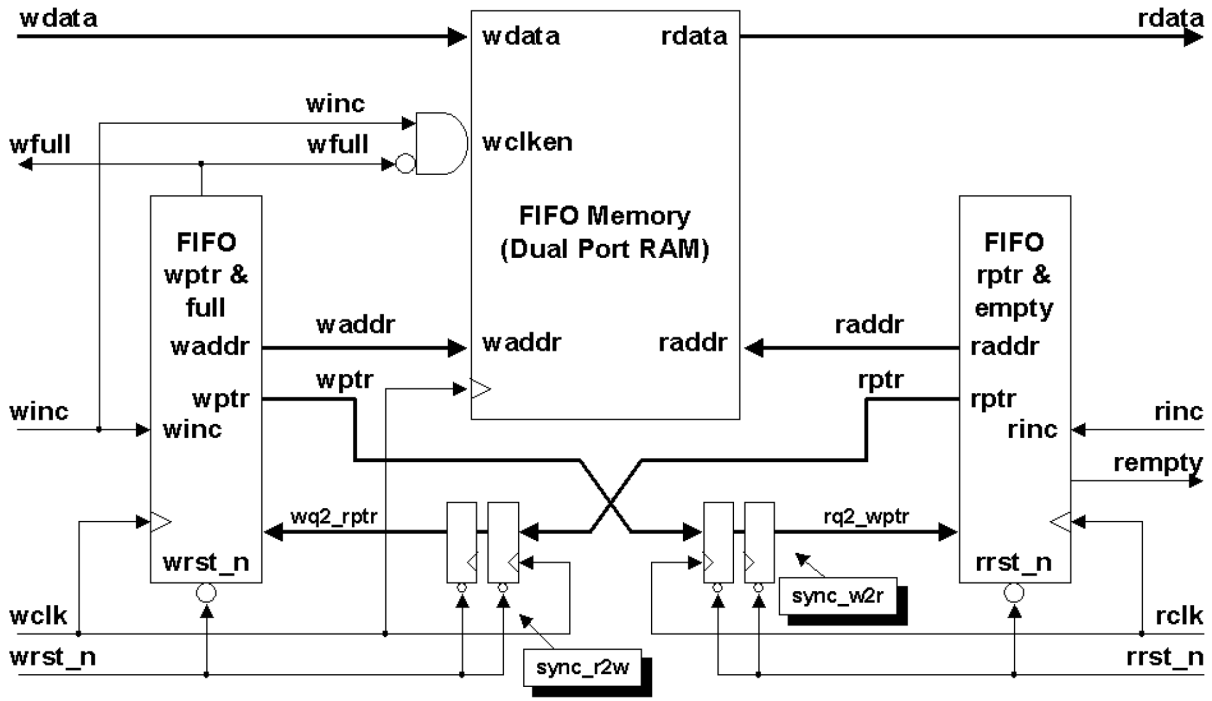

- FIFO for clock crossings

http://www.sunburst-design.com/papers/CummingsSNUG2002SJ_FIFO1.pdf

Supply Generation

Supply Generation

- Linear
  - Series or shunt
  - Linear regulation
  - Quiet
  - Inefficient (unless Vin-Vout is small)
- Switching (Capacitive)
  - Limited efficiency
  - Poor regulation
  - Voltage ripples
- Switching (Magnetic)
  - Efficient
  - Require external components
  - Noisy

Linear vs. Switching Regulators

Linear

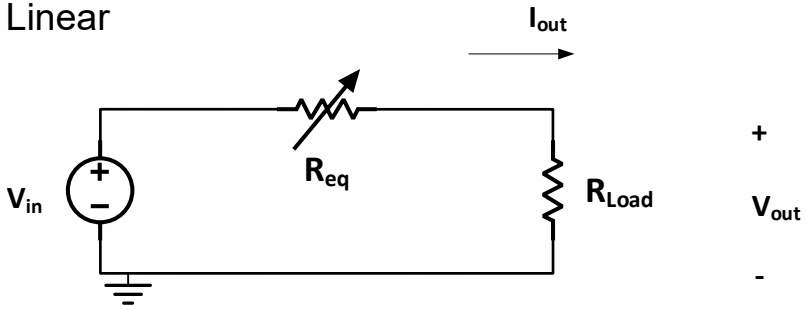

Efficiency η < Vout/Vin

Switching

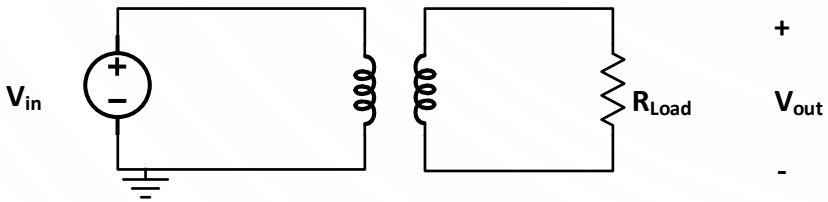

Linear Voltage Regulator

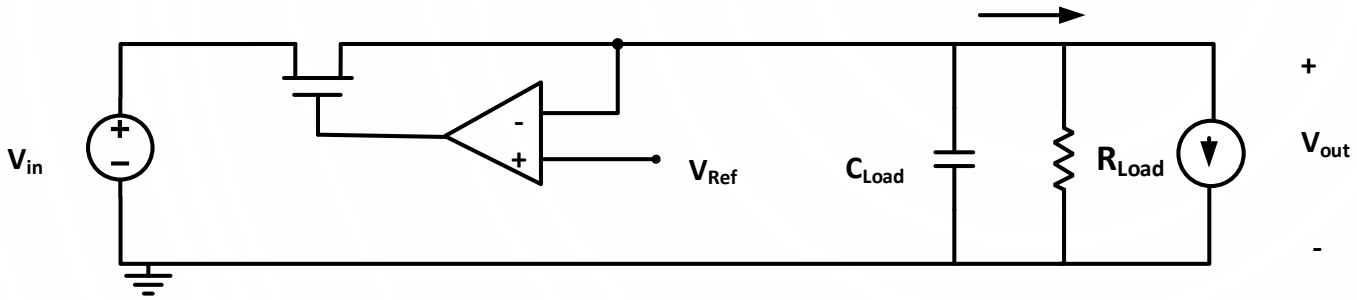

Negative feedback sets low supply resistance
Voltage regulated to desired level

E.g. IBM Power7 has 48 linear regulators

Switching Supply

- Buck Converter

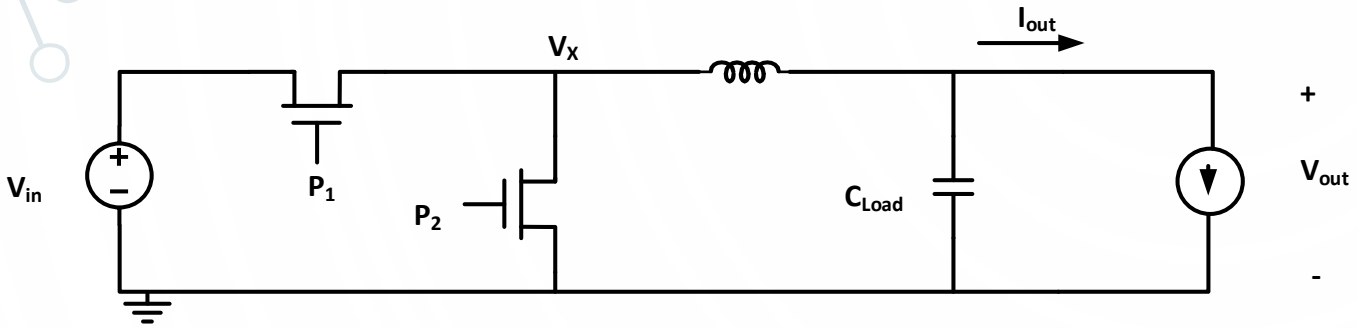

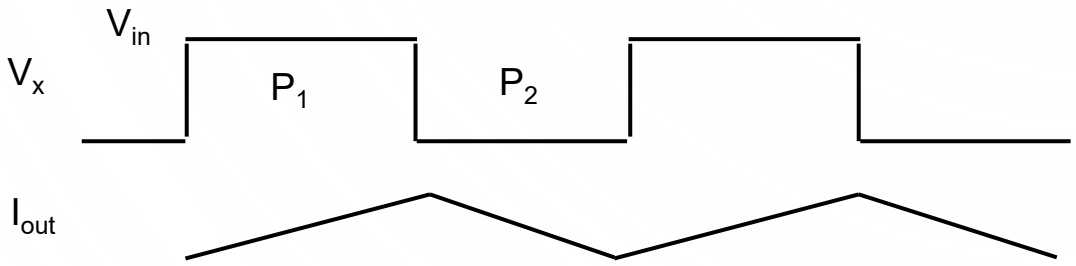

Pulse-Width Modulation (PWM) regulates V_{out}

High switching frequency, interleaving reduce ripple

Inside Haswell

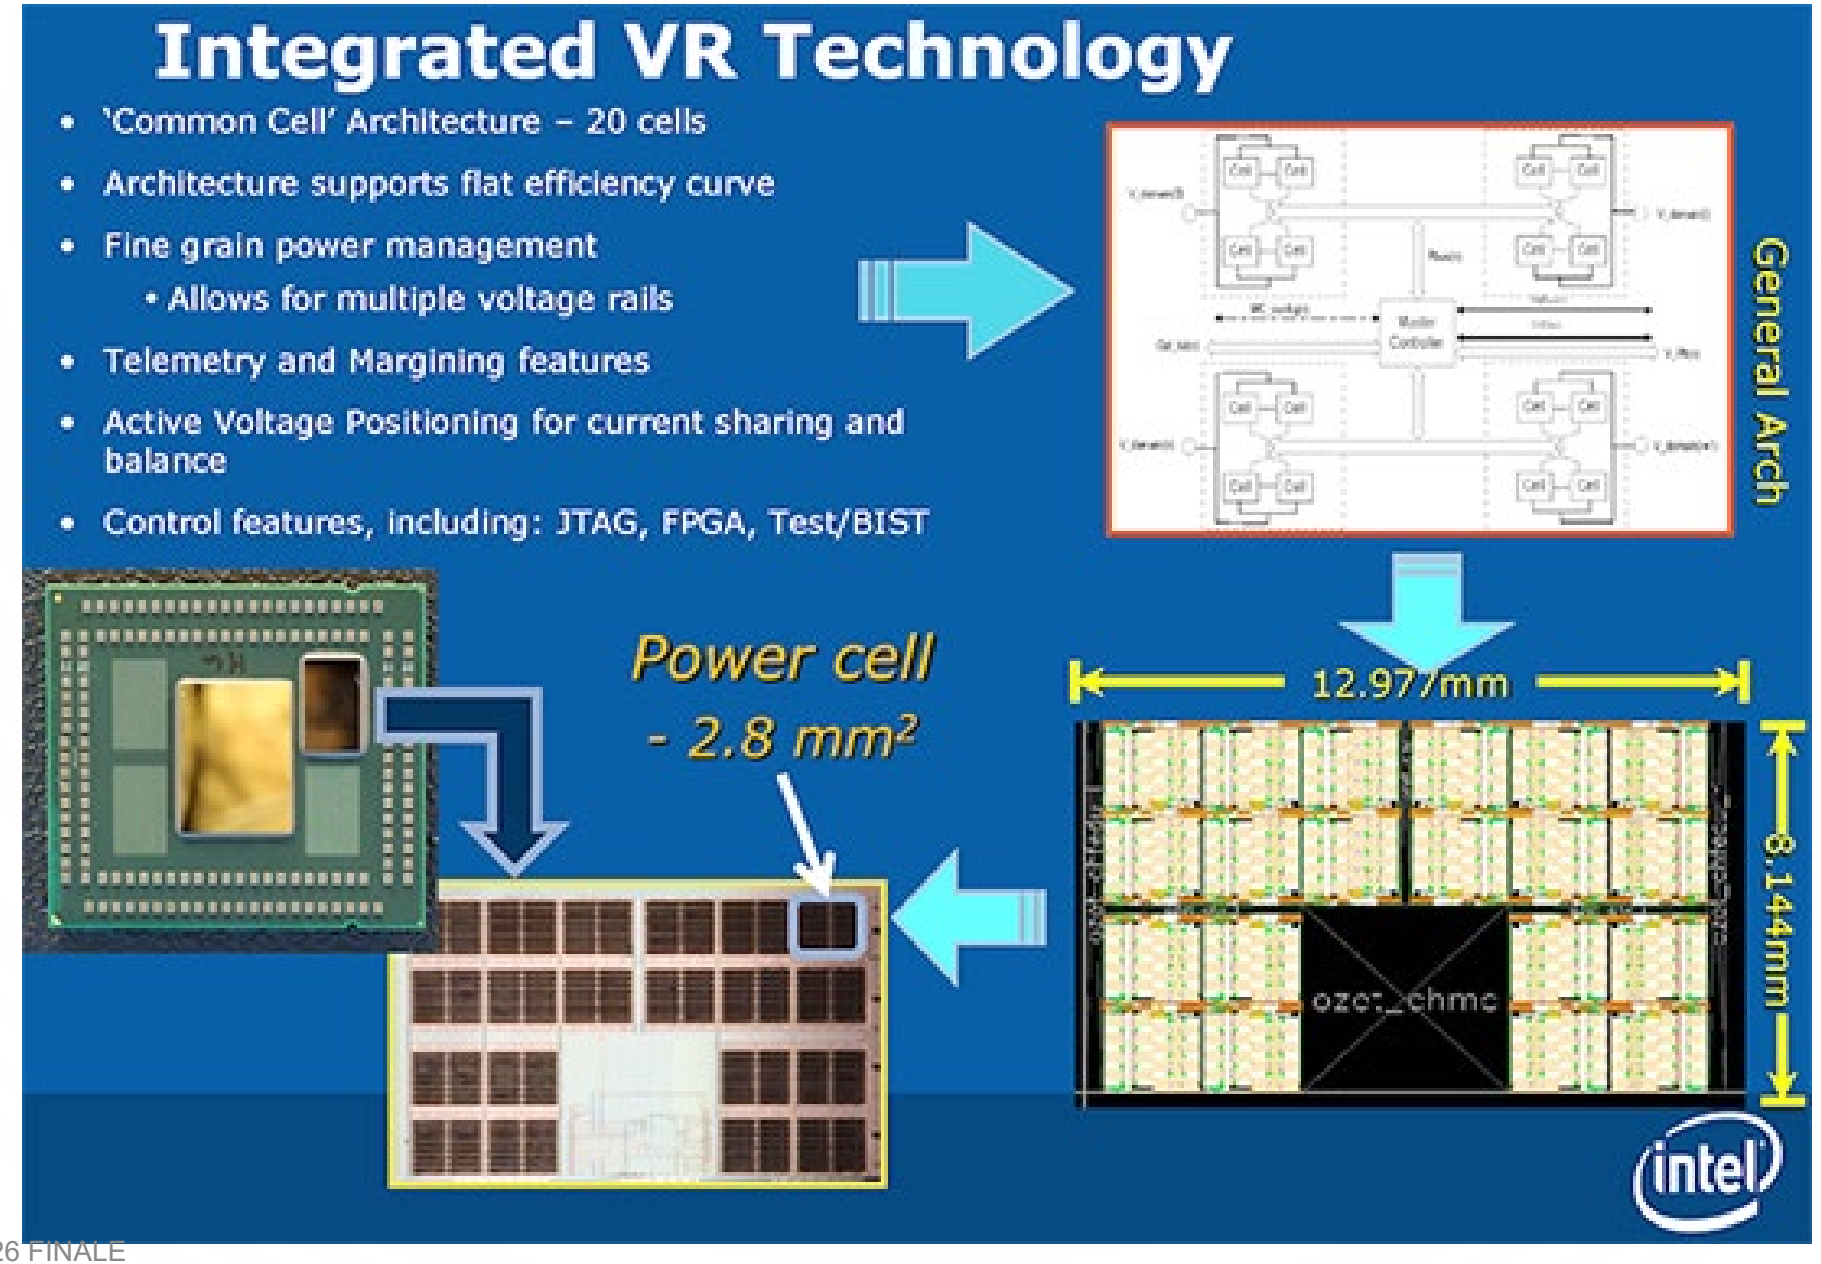

Inside Haswell

Review: Power Cell Architecture

- Each Power cell = Mini VR
  - Up to 25A rating* - tested
  - Programmable switching frequency 30MHz to 140MHz
  - Ring coupled inductor topology
- 16 phases per power cell, 320 phases per chip
  - High phase count reduces noise, ripple
  - High granularity
  - Cell shedding
  - Bridge shedding
- BIST
  - Self-load and characterization system.

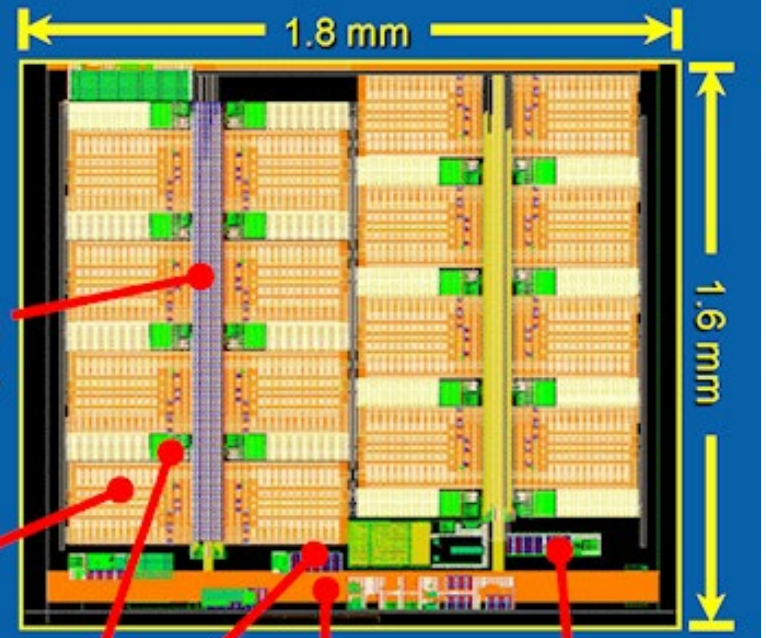

* Thermally constrained

Intel Broadwell

- Inductors moved to a small PCB

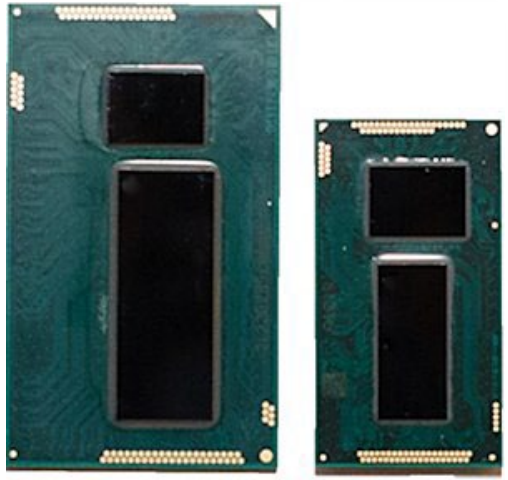

HSW U/Y
40x24x1.5mm

BDW-Y
30x16.5x1.04mm

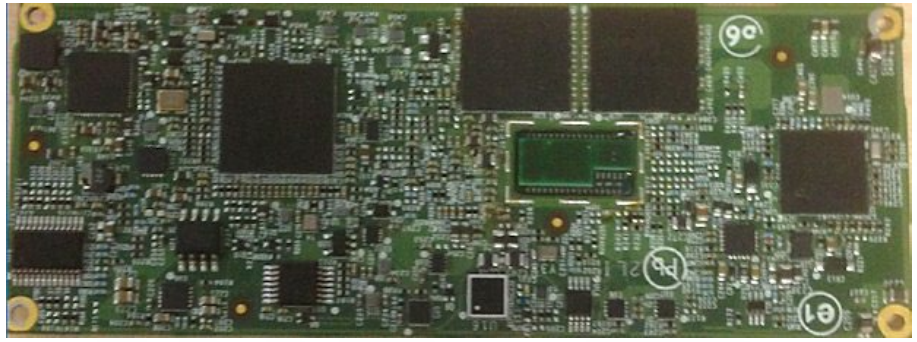

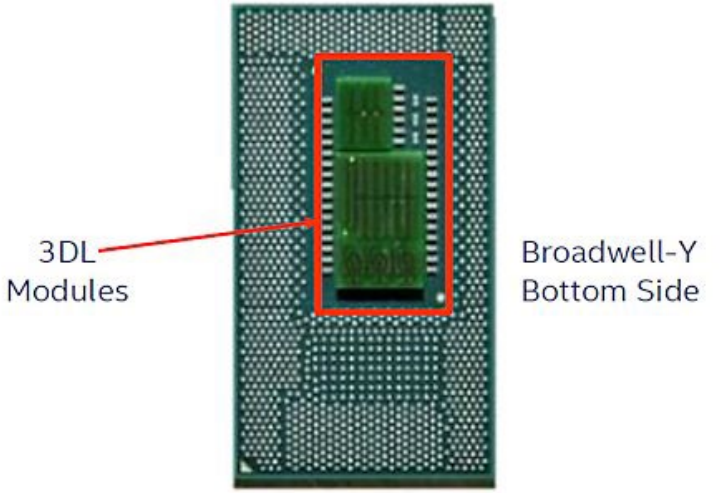

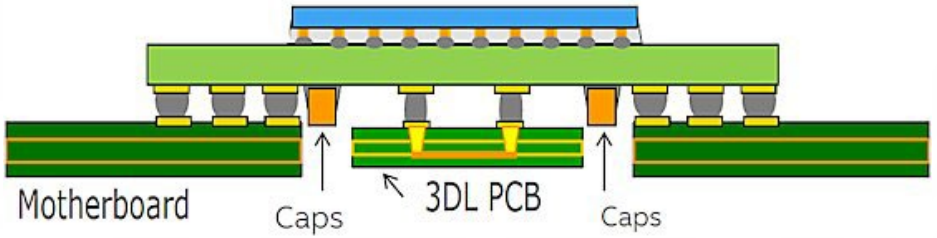

Switched-Capacitor Supply

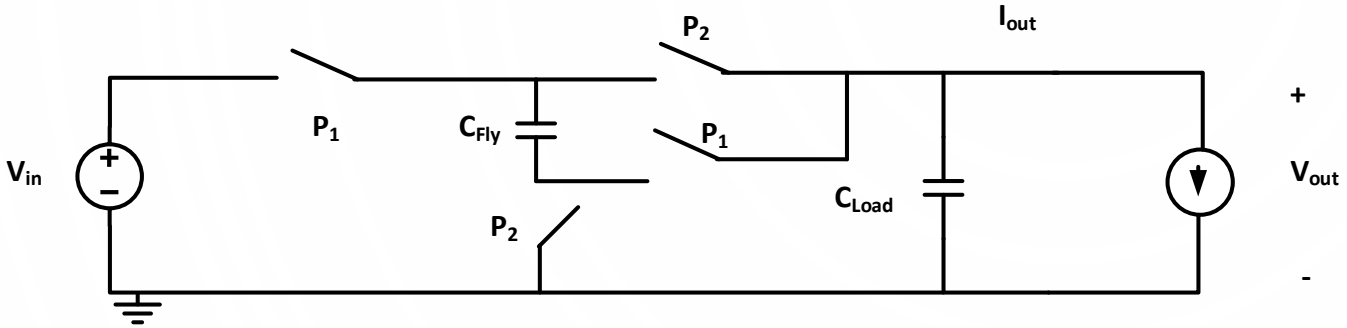

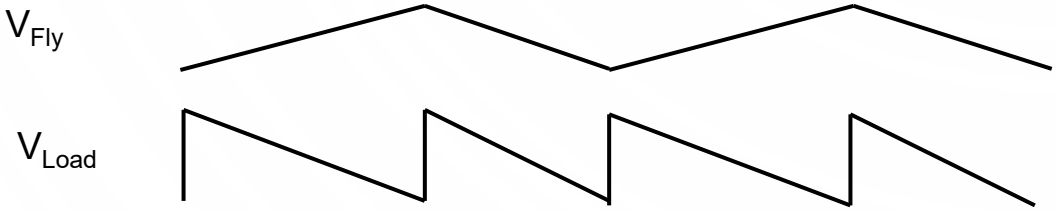

Interleaving reduces ripple, but lowers efficiency

What happens with supply when load changes?

http://www.research.ibm.com/people/r/restle/Animations/DAC01top.html

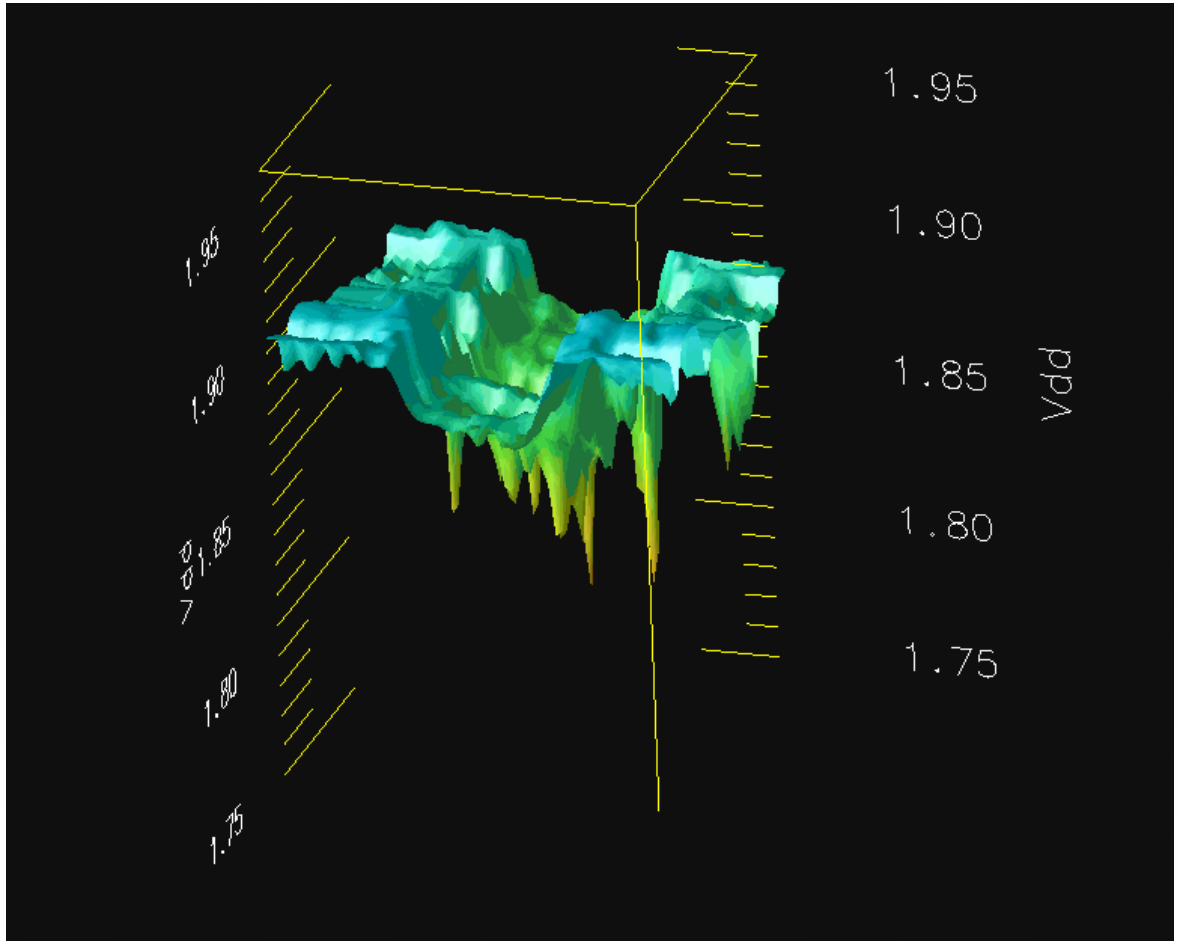

Phillip Restle, IBM

Power Delivery

- Typical model

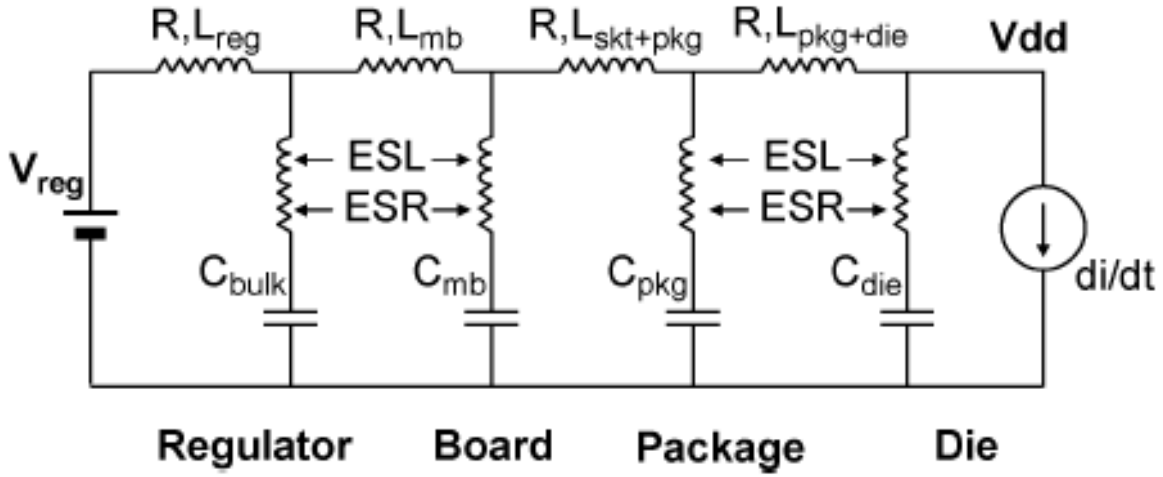

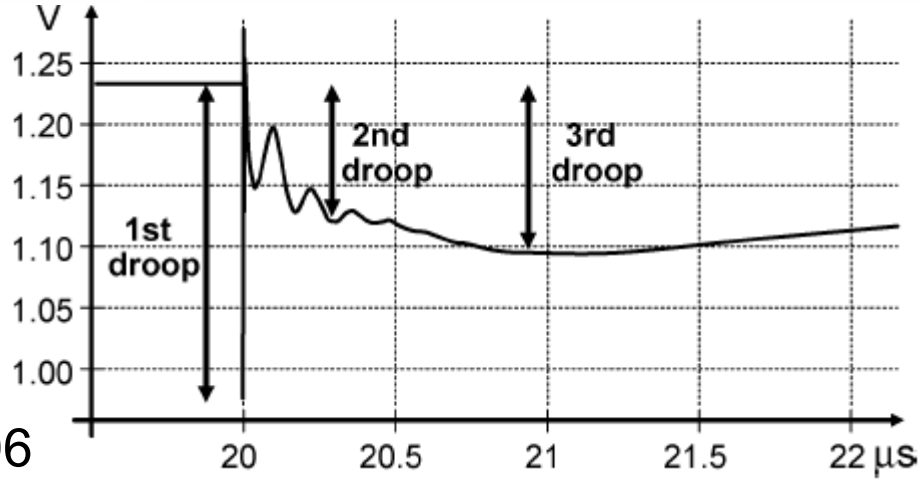

Wong, JSSC'06

Supply Resonances

- First droop
  - Package L + on-die C
- Second droop
  - Motherboard + package decoupling
- Third droop
  - Board capacitors

Clock and Supply

- Large digital systems can have large voltage transients
  - Can we filter impact of voltage on a clock generator?

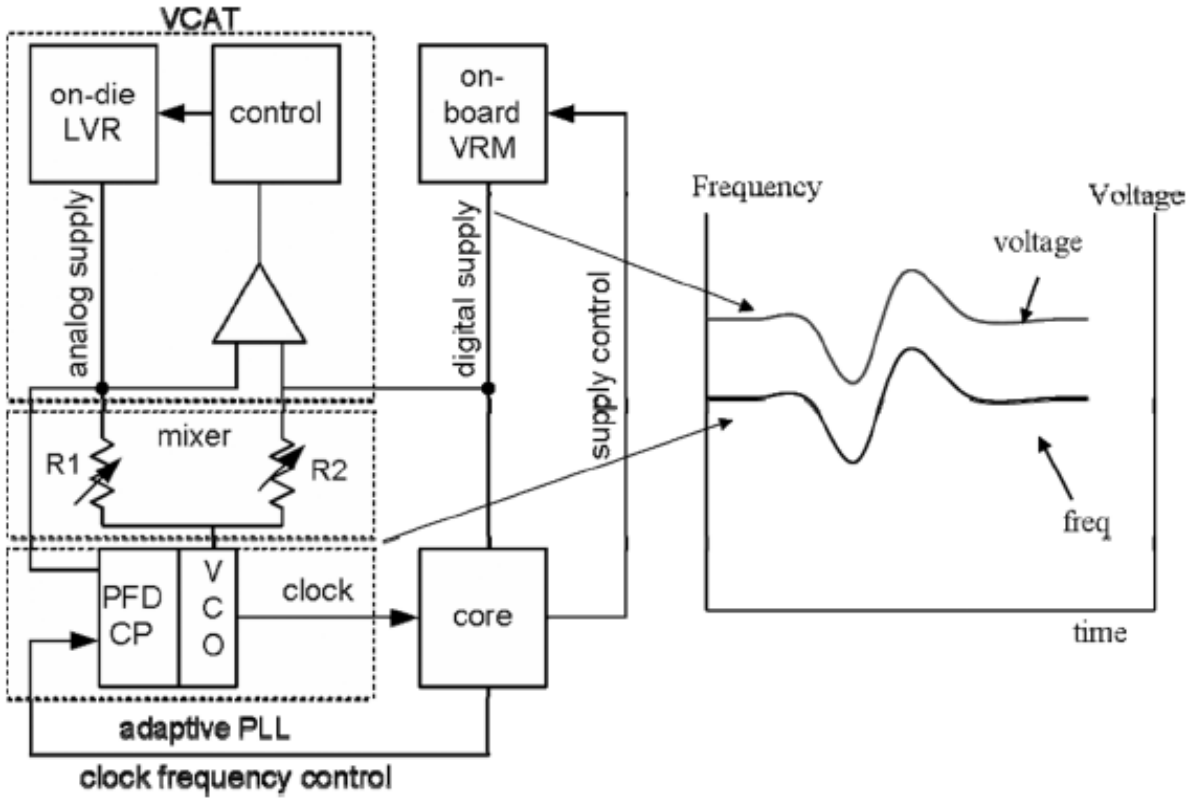

Kurd, JSSC'09

Clock and Supply

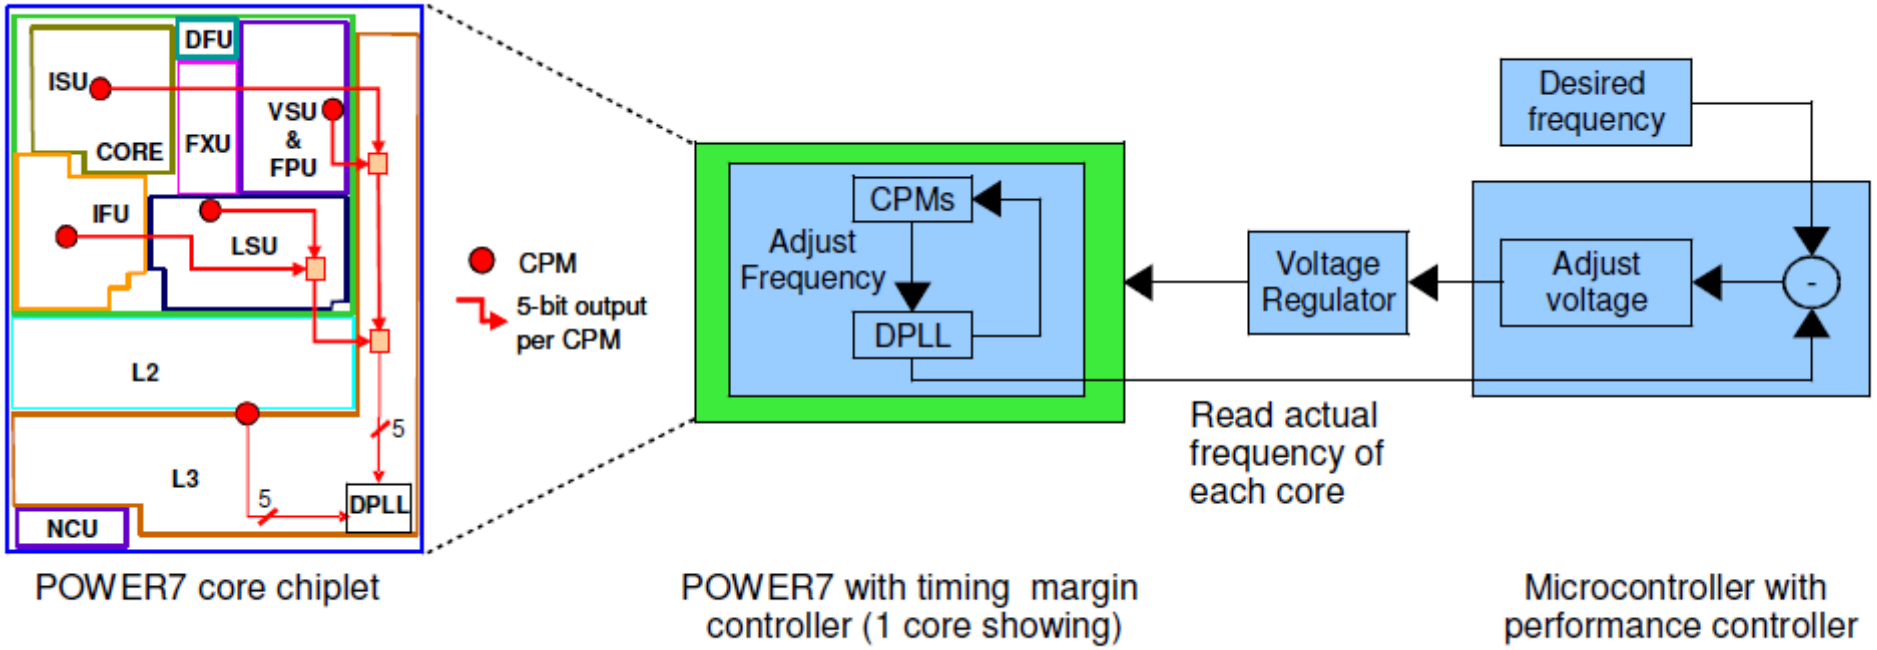

- IBM Power7, with one PLL per core

Lefurgy, MICRO'11

How to model

- Abstracted delay line

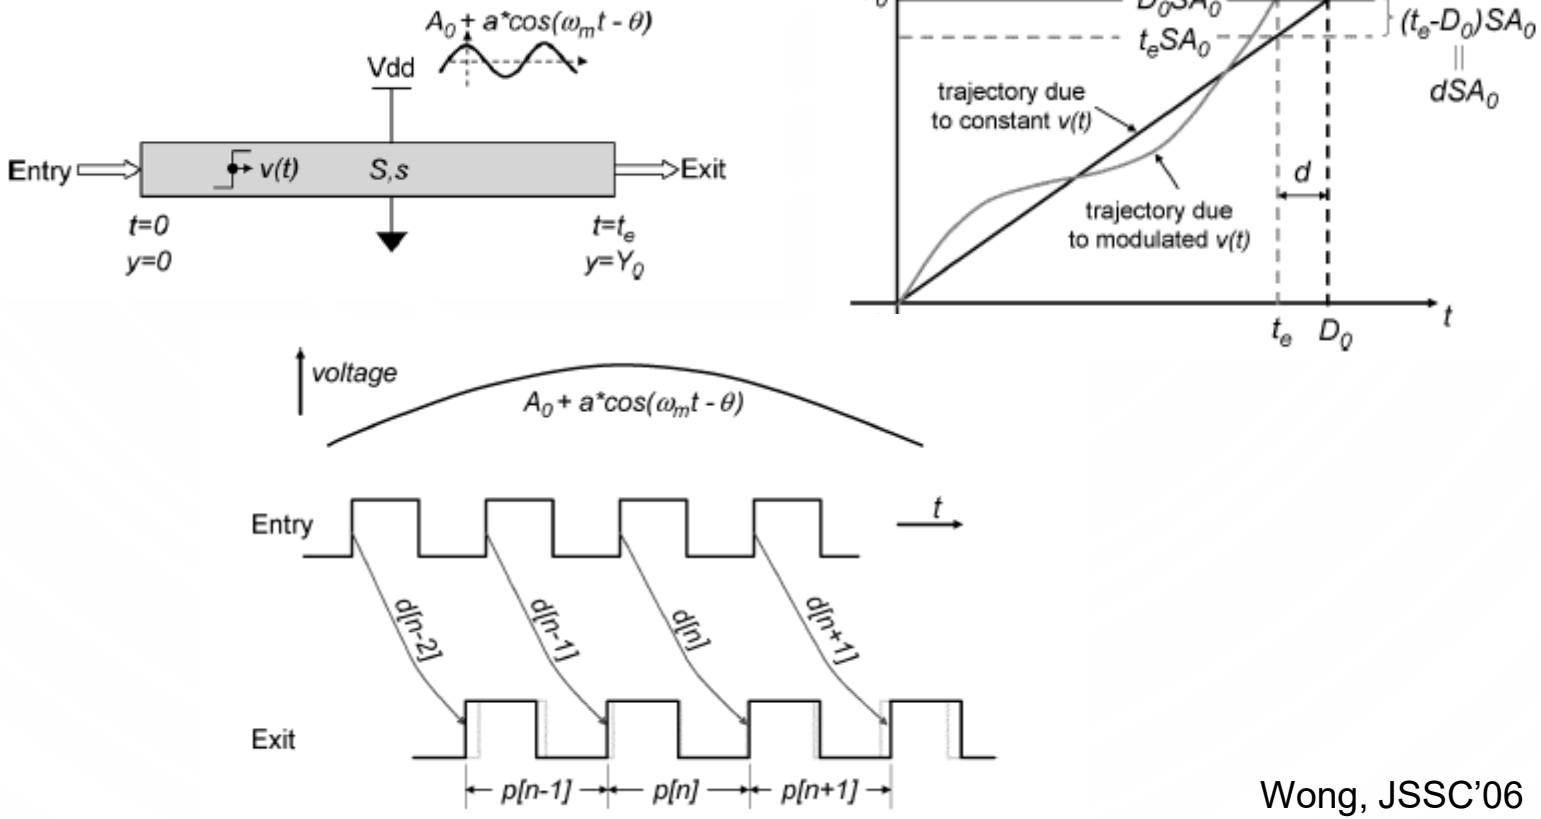

Wong, JSSC'06

Period modulation from successive modulated delays

Droop Detection

- Hashimoto, JSSC 4/18

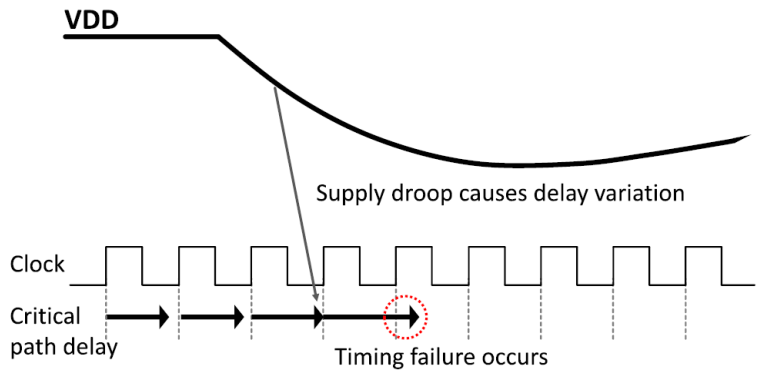

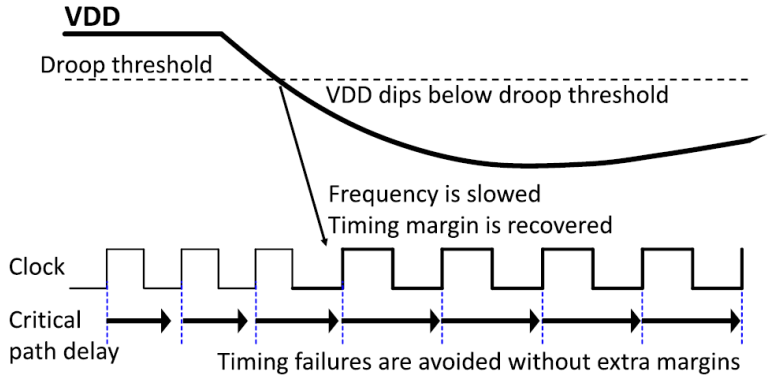

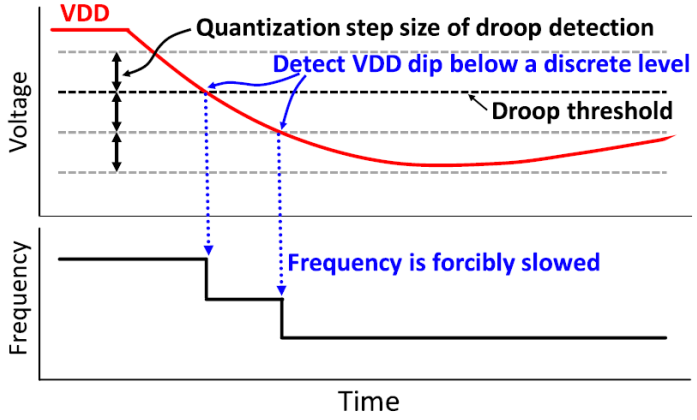

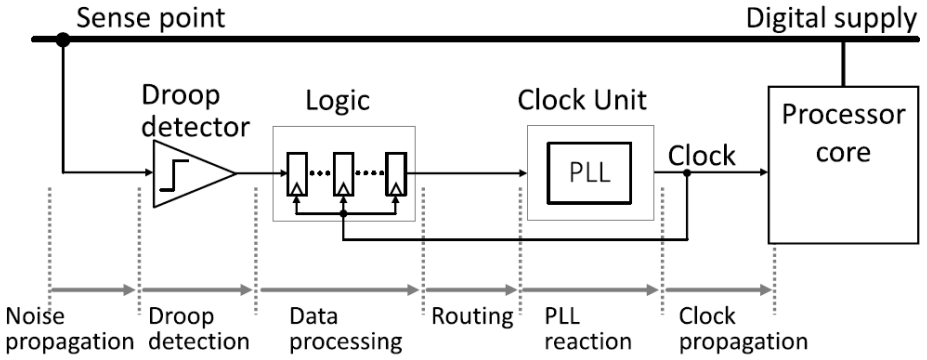

Wrap-Up

Many Topics Covered

- RISC-V SoC organization (Chipyard); interconnects; accelerators
- Productivity languages and flows for design and verification (SystemVerilog, Chisel, HLS)
- Technology features, variability
- Standard cells
- Design for performance: Timing, latch-based design
- Memory, SRAM
- Low-power design: Lowering supplies, DVFS, leakage control
- Clock generation
- Supply generation

This Class

- Put design choices in technology perspective
- The design constraints have changed and will be changing
  - Cost, energy, (power, leakage, …), performance
- Focused on SoC design, variability, power-performance tradeoffs, power management
- Did not cover arithmetic, domino, I/O, supply generation, packaging, …
  - Packaging will need to be added back in the course

Other Classes

- Integrated Circuits:
  - EE240B: Advanced Analog Circuits
  - EE240C: Data Conversion
  - EE242: Advanced RF
  - EE290C: Advanced Topics in Circuit Design
    - Tapeout Class
- Computer Architecture
  - CS252: Advanced Computer Architecture
  - EECS290: ML Hardware
- Systems
  - CS262: Advanced Computer Systems

Tapeout Class

2021:
18 students

TSMC 28nm
1mm x 1mm

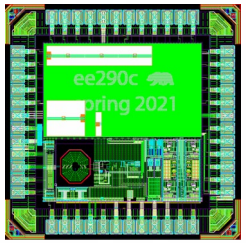

OSCIBear: 32b RISC-V + BLE + AES + Power

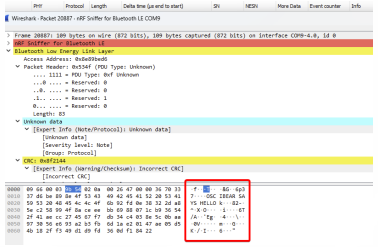

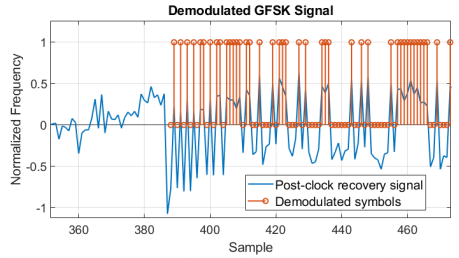

Processor, transmitter functional

HotChips'23

2022:
41 students

2x Intel 16
2mm x 2mm

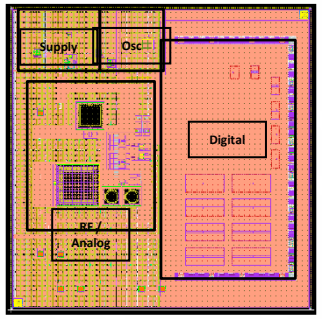

SCuM-V'22: 64b RISC-V core, BLE + 802.15.4, LDOs, references

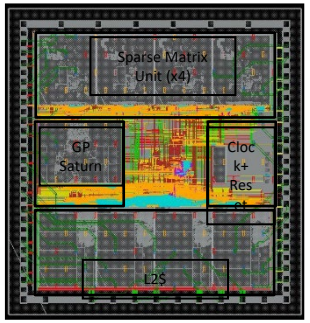

BearlyML'22: 5 RISC-V cores:
4 Rocket with custom sparse matrix acc,
Saturn-V, NoC, PLL, L2

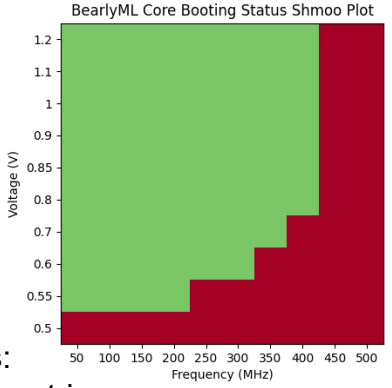

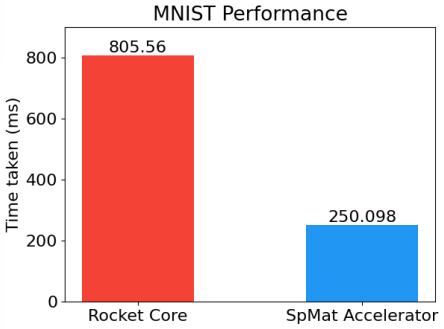

Shmoo plot, running MNIST

HotChips'23

2023:
54 students

3x Intel 16
2mm x 2mm

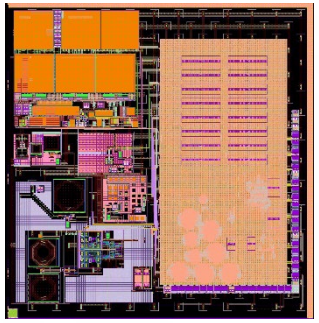

SCuM-V'23: 32b RISC-V core, BLE + 802.15.4, LDOs, references, radar

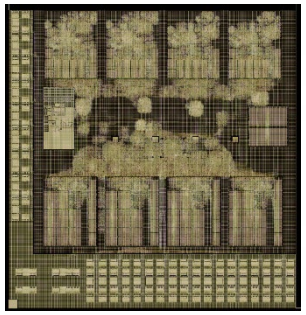

BearlyML'23: 4 RISC-V Rockets with custom sparse matrix acc, near-memory acc, NoC, L2$

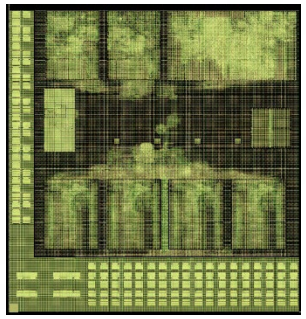

RoboChip'23: 2 RISC-V Rockets with Kalman, LQR acc, BooM + MTE, NoC, L2$

2024:
69 students

This Field

- Moore's law will end sometime during your (my?) career
  - 3nm in 2022 scales to 0.1nm by 2050 with 2-yr cycles (or to 0.5nm with 5-yr cycles)
- Physics will stop CMOS somewhere ~2nm (?)
  - Will we see a different (CMOS) device in the meantime
- Economics will likely stop it somewhere while still in single digits
  - And the nodes will be stretched out
- We will see multi-chip/packaging solutions
- Don't worry: Creativity is unlimited!
  - What can you build with 10B/100B/1 trillion transistors?
  - Even filling 10B-transistor chips with SRAM is not trivial!

Current Perspective for <5nm

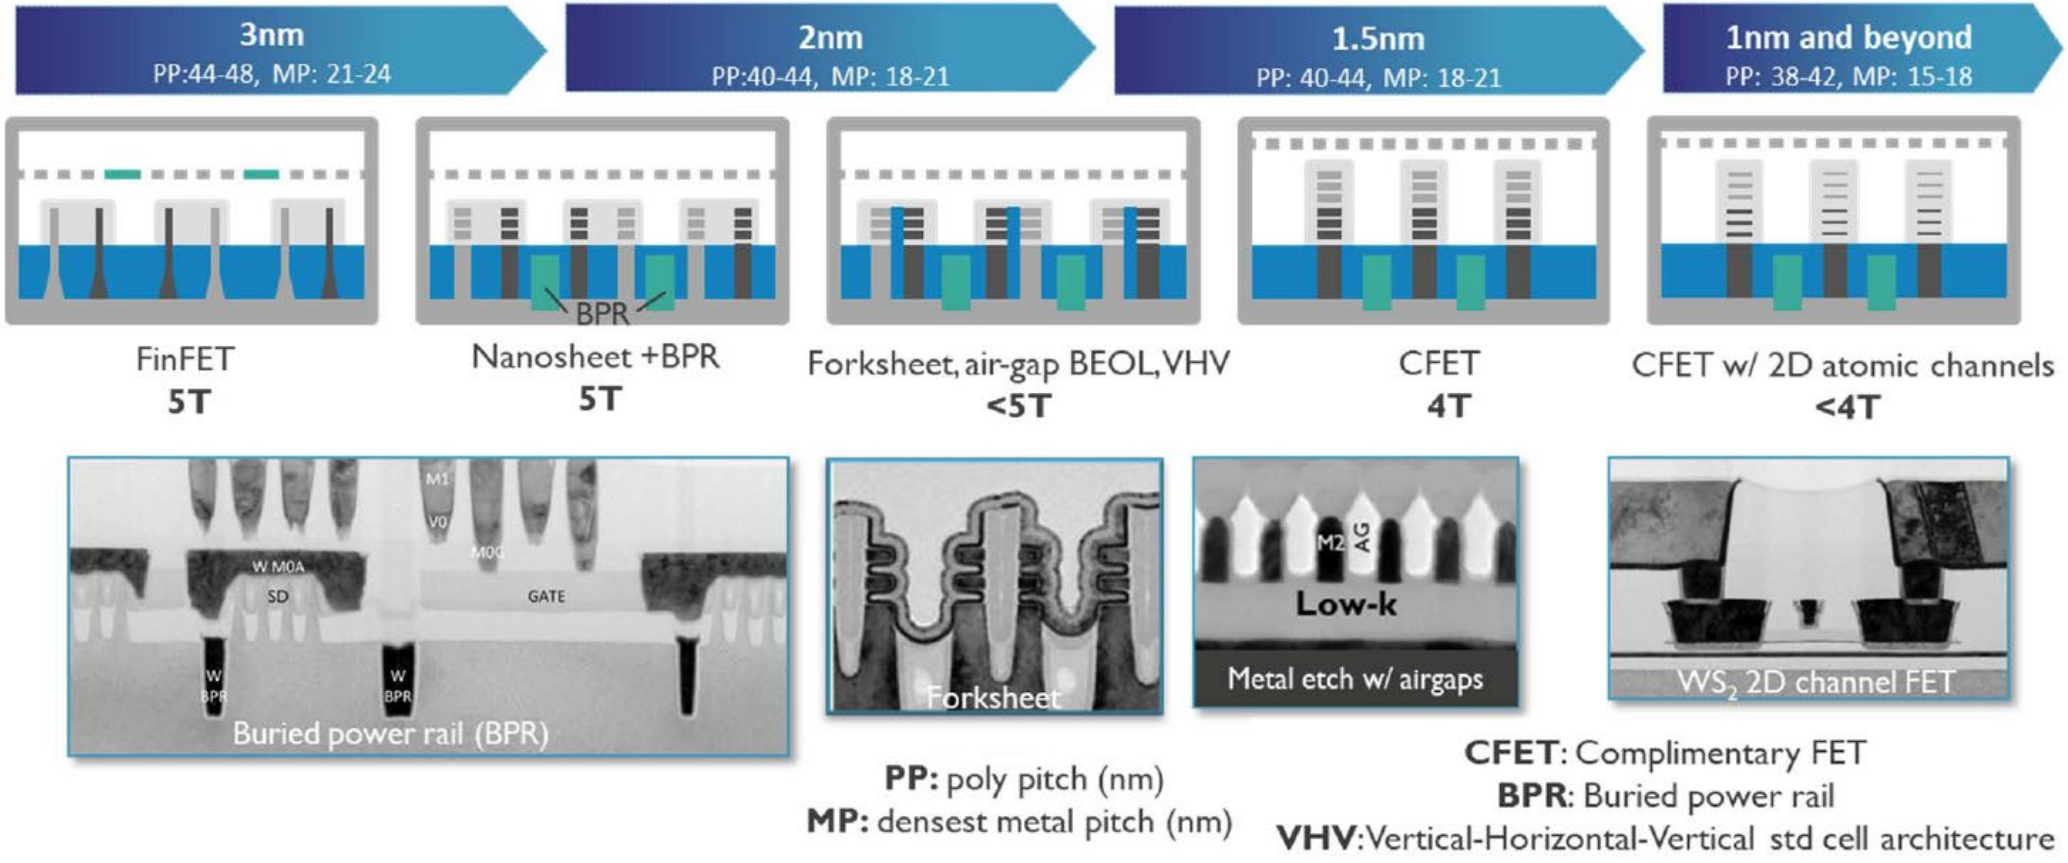

- Samavedam, et al, IEDM'20

The Era of Chiplets

UCIe
Universal Chiplet
Interconnect Express

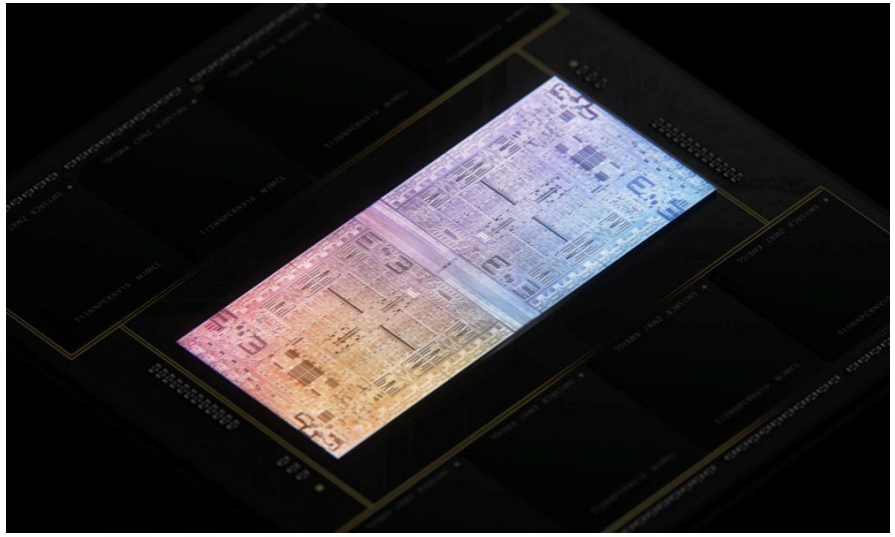

Apple M1 Ultra

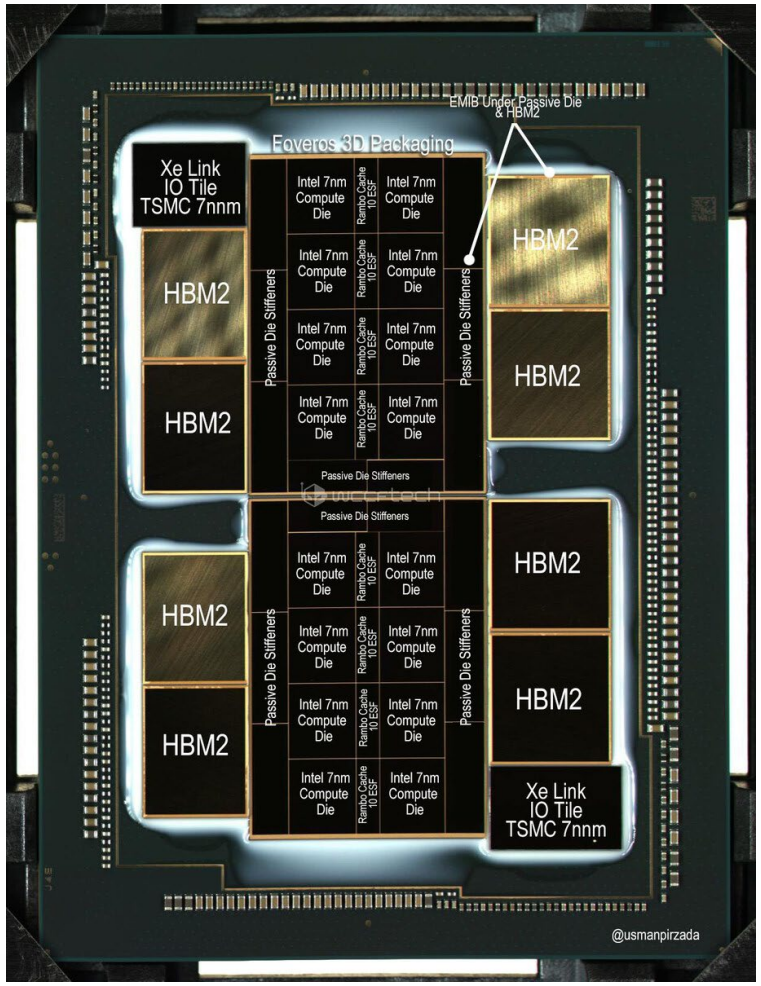

Intel Ponte Vecchio

The Era of Distributed, Domain-Specific Compute

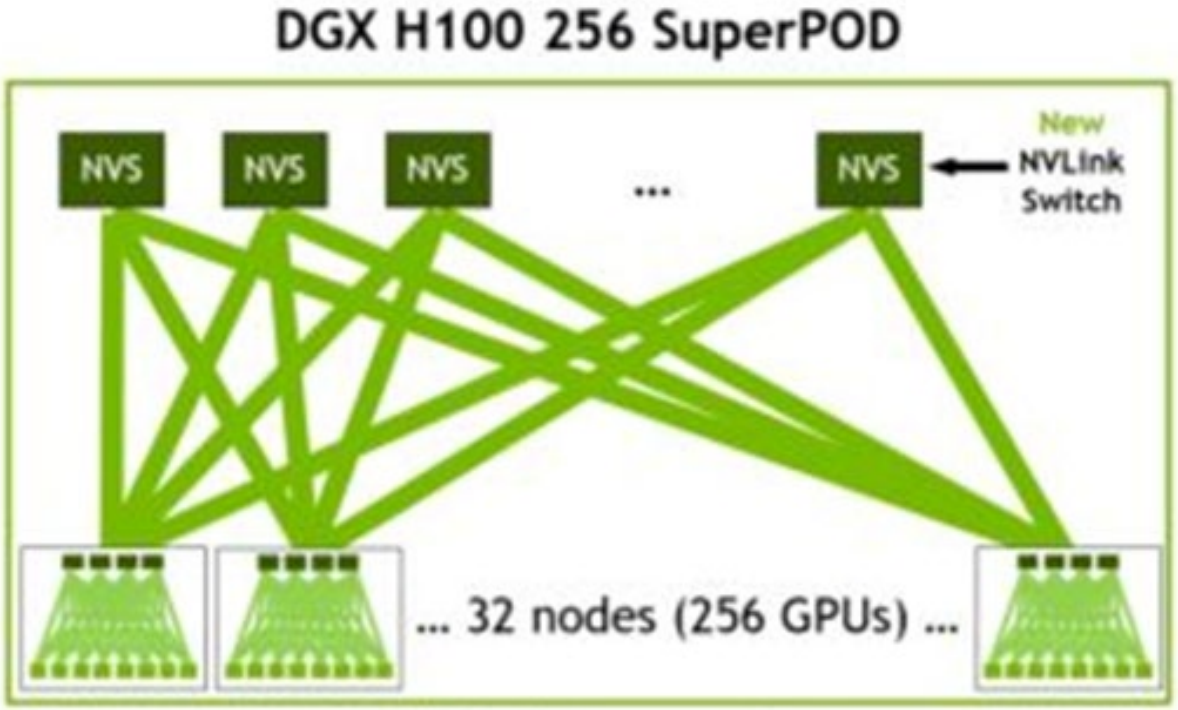

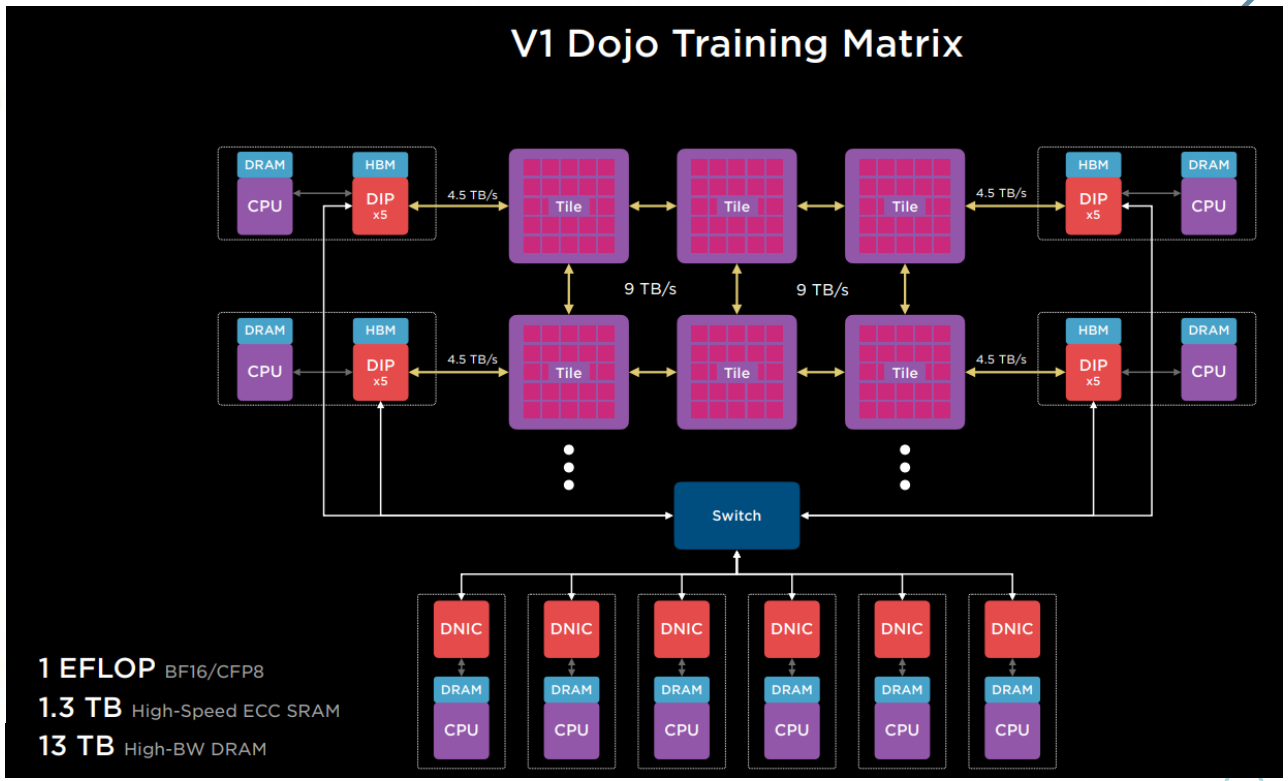

Technology Strategy / Roadmap

2000 2005 2010 2015 2020 2025 2030

Plan A: Extending Si CMOS

R D

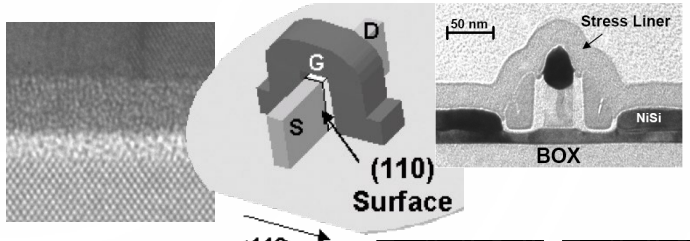

Plan B: Subsytem Integration

R D

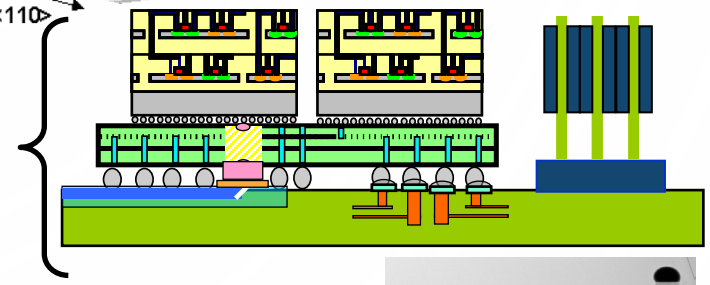

Plan C: Post Si CMOS Options

R R&D

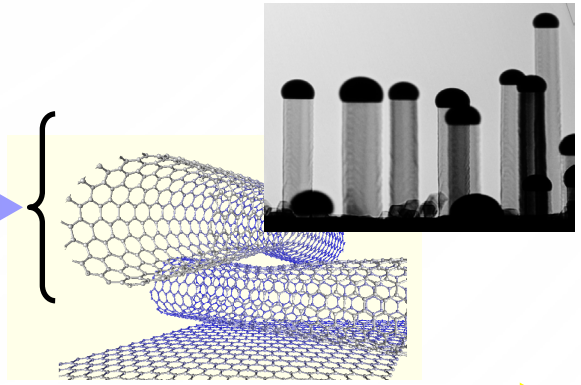

Plan Q: Quantum Computing

R D

T.C. Chen, Where Si-CMOS is going: Trendy Hype vs. Real Technology, ISSCC'06

This Field

- Focus on principles
- Watch out for opportunities
- Stay current!
  - Read
  - Keep learning new tools
  - Do