NC STATE UNIVERSITY

Design of 3D-Specific Systems: Propsctives and Interface Requirements

Paul Franzon

North Carolina State University

Raleigh, NC

paulf@ncsu.edu

919.515.7351

Outline

- Overview of the Vectors in 3D Product Design

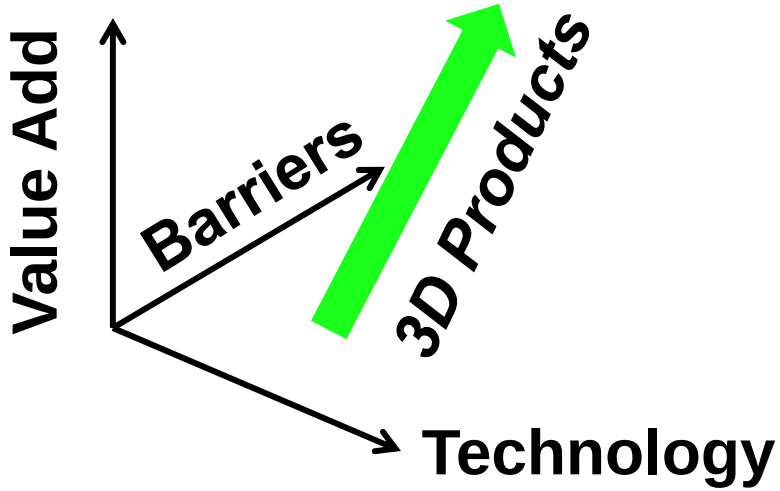

- Short term – find the low-hanging fruit
- Medium term – Logic on logic, memory on logic
- Long term – Extreme Scaling; Heterogeneous integration; Miniaturization
- Interface requirements

Future 3DIC Product Space

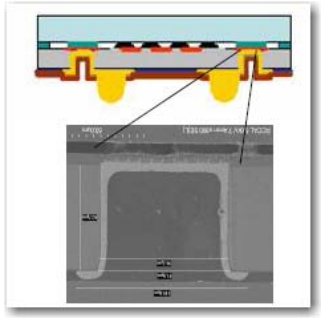

Image sensor

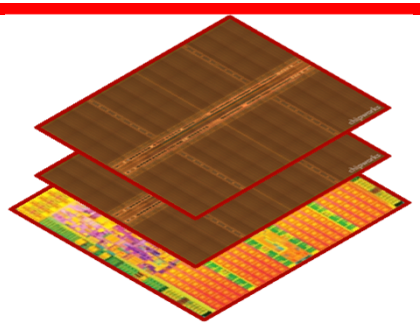

3D Mobile

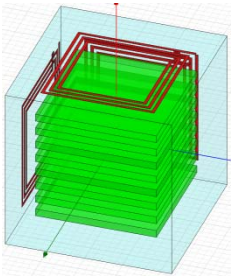

Sensor Node

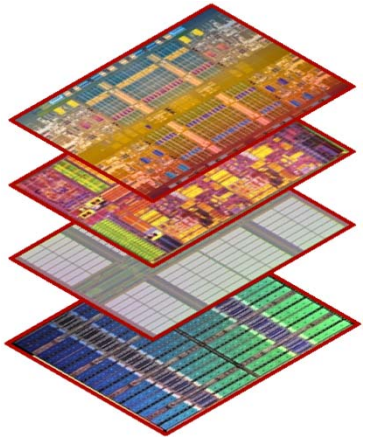

Heterogeneous

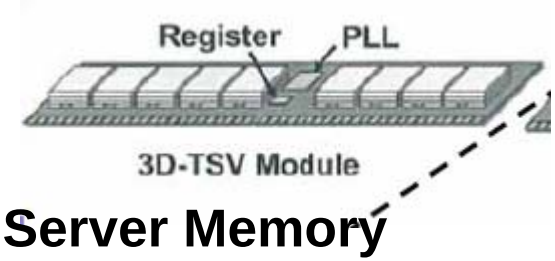

Server Memory

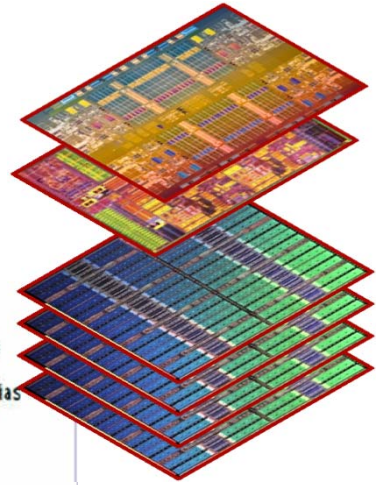

3D Processor

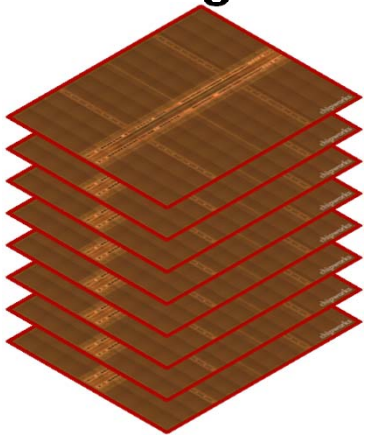

"Extreme" 3D Integration

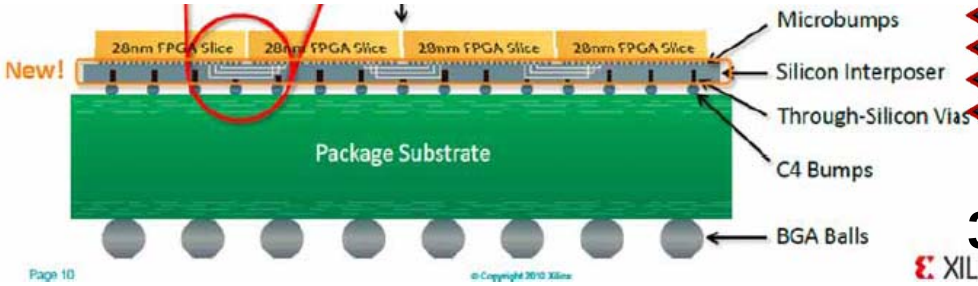

Interposer

Time

3DIC Technology Set

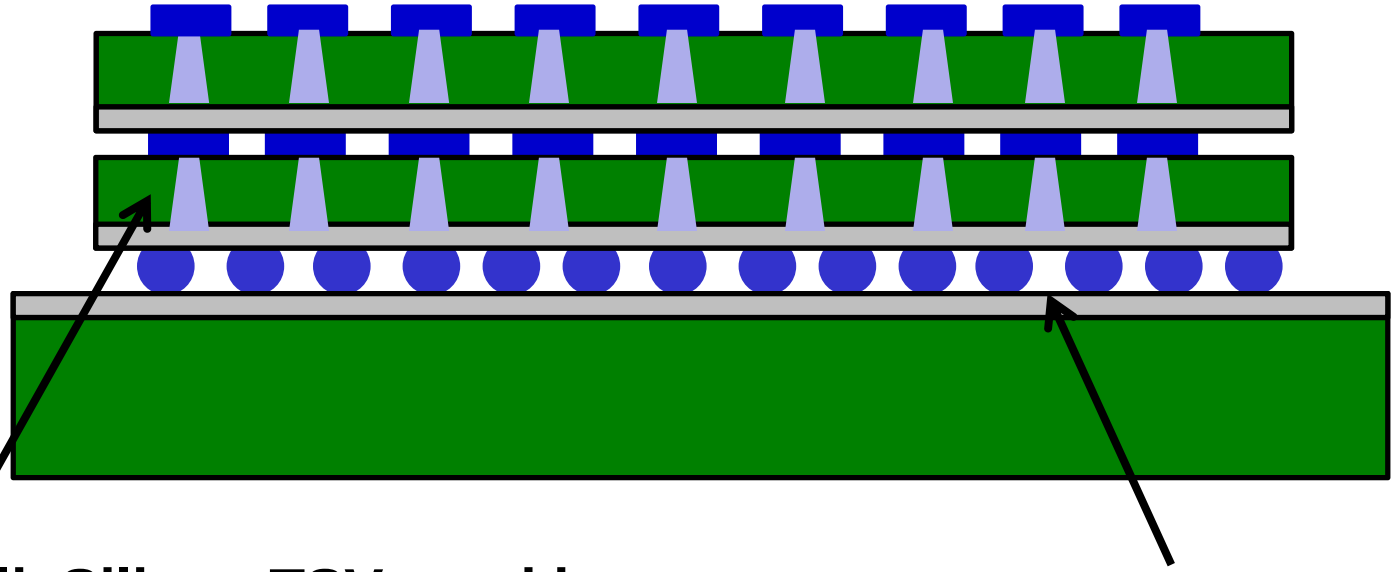

Bulk Silicon TSVs and bumps (25 - 40 μm pitch)

Face to face microbumps (1 - 30 μm pitch)

C_{TSV} ~ 30 fF

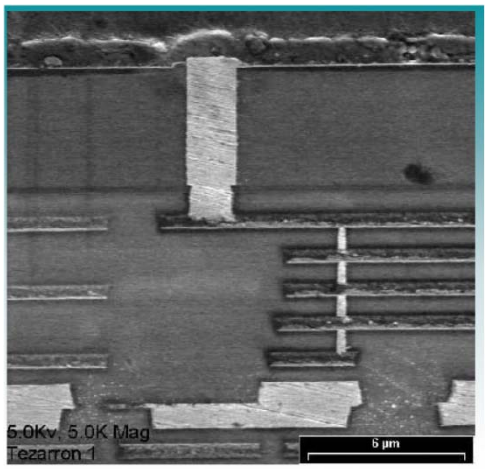

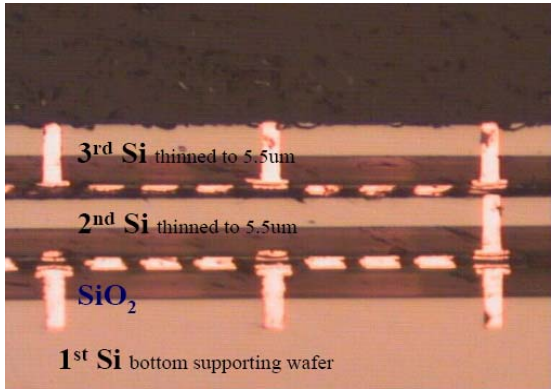

Tezzaron

Simplified Process Flow

1. Etch TSV holes in substrate
   - Max. aspect ratio 10:1 ➔ hole depth < 10x hole radius

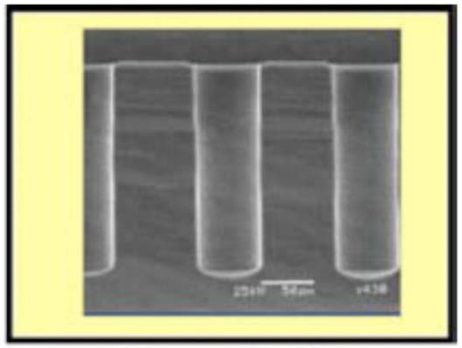

Deep RIE etching

2. Passivate side walls to isolate from bulk

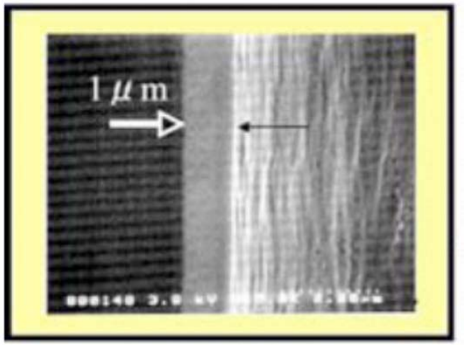

Insulator deposition

… Simplified Process Flow

3. Fill TSV with metal
   - Copper plating, or
   - Tungsten filling

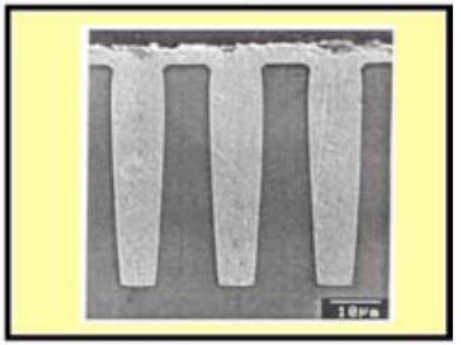

High aspect
Cu plating

4. Often the wafer is then attached to a carrier or another wafer before thinning

5. Back side grinding and etching to expose bottom of metal filled holes

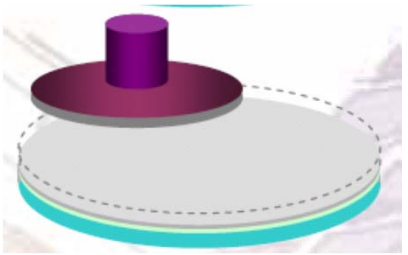

Wafer Thinning

6. Formation of backside microbumps

Top and back
bump formation

… Simplified Process Flow

7. Wafer bonding and (sometimes) underfill distribution

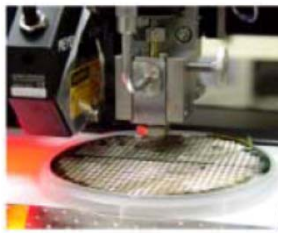

Underfill

➔ TSV enabled 3D stack

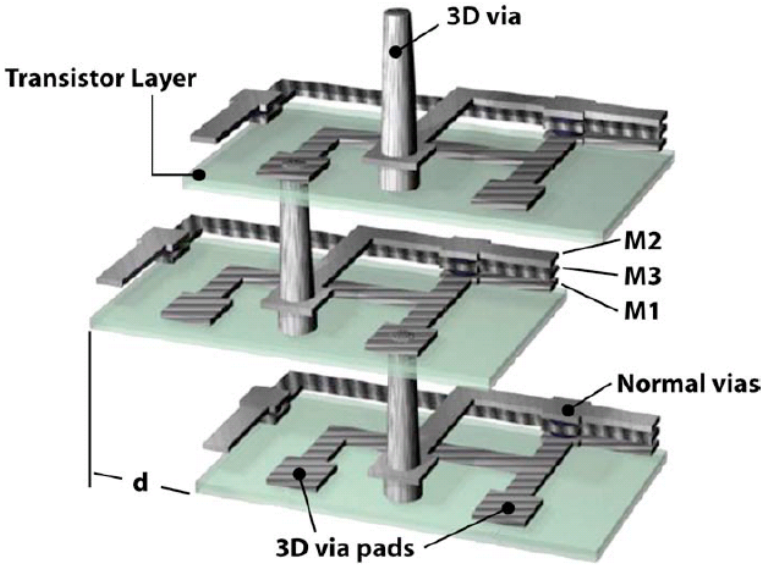

Source: University of Alberta.

3DIC Technology Set

- Interposers: Thin film or 65/90 nm BEOL

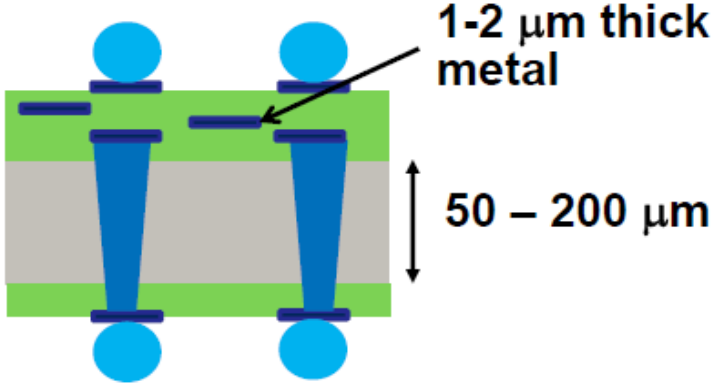

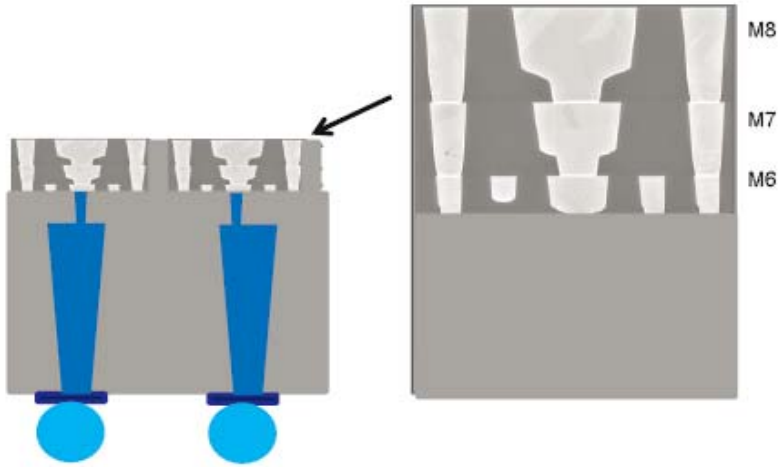

- Assembly : Chip to Wafer or Wafer to wafer

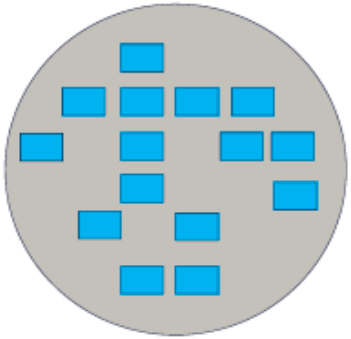

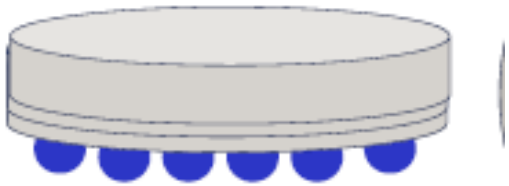

Value Propositions

Fundamentally, 3DIC permits:

- Shorter wires
  - consuming less power, and costing less
  - The memory interface is the biggest source of large wire bundles

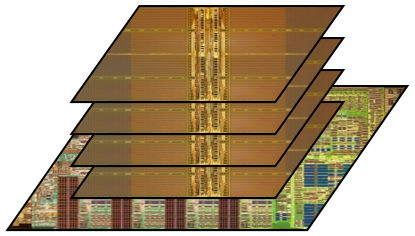

- Heterogeneous integration
  - Each layer is different!
  - Giving fundamental performance and cost advantages, particularly if high interconnectivity is advantageous

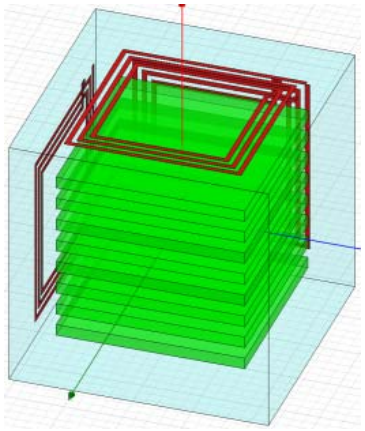

- Consolidated “super chips”
  - Reducing packaging overhead
  - Enabling integrated microsystems

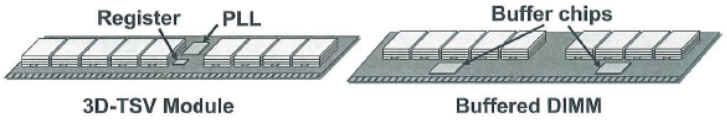

The Demand for Memory Bandwidth

Computing

MULTICORE AND REVERSE SCALING

	2004 Baseline	Multi-core Approach	Reverse scaling	Reverse scaling
Frequency	4 GHz	8 GHz	8 GHz	4GHz
No. of Cores	1 Core	4 Cores	16 Cores	16 Cores
Core rel. IPC	1	1	0.5	1
Total Flops	32 GFlops	256 GFlops	512 GFlops	512 GFlops
Supply	1.2V	1.0V	1.0V	1.0V
Power	84W	233W	233W	117-163W
Bandwidth requirement	32GB/s	256GB/s	512GB/s	512GB/s

Future microprocessors and off-chip SOP interconnect
Hofstee, H.P.;
Advanced Packaging, IEEE Transactions on [see also Components, Packaging and Manufacturing Technology, Part B: Advanced Packaging, IEEE Transactions on]
Volume 27, Issue 2, May 2004 Page(s):301 - 303

Similar demands in Networking and Graphics

Ideal: 1 TB / 1 TBps memory stack

Memory on Logic

Conventional

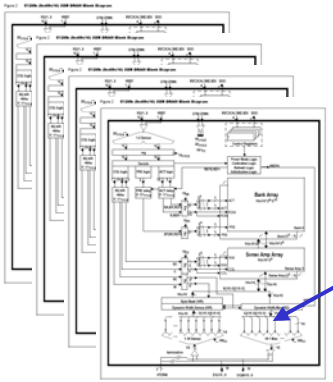

x32

to

x128

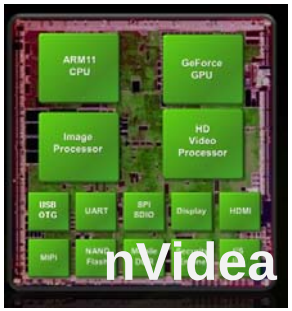

or

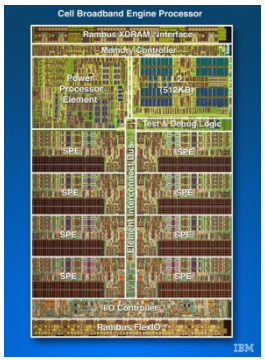

TSV Enabled

Less Overhead

Flexible bank access

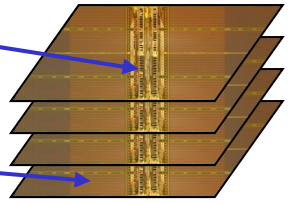

Less interface power

3.2 GHz @ >10 pJ/bit

→ 1 GHz @ 0.3 pJ/bit

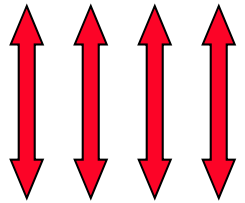

N x 128

“wide I/O”

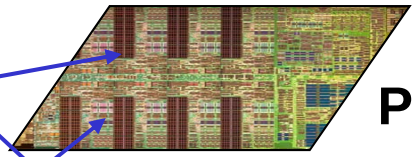

Processor

Flexible architecture

or

Mobile

Short on-chip wires

Mobile Graphics

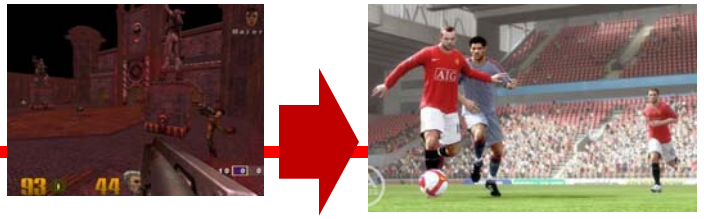

- **Problem**: Want more graphics capacity but total power is constrained
- **Solution:** Trade power in memory interface with power to spend on computation

POP with LPDDR2

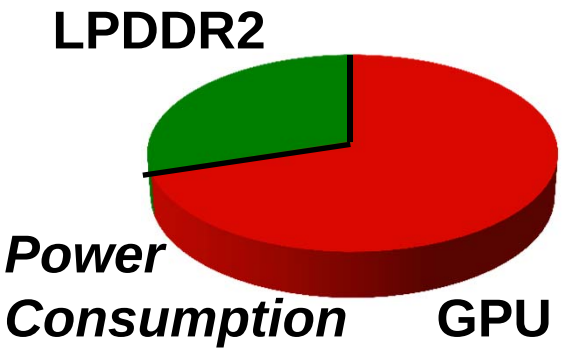

532 M triangles/s

TSV Enabled

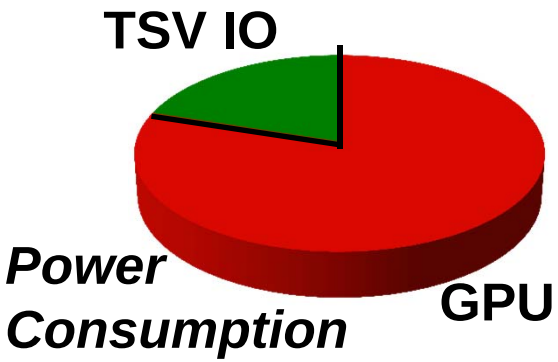

695 M triangles/s

Dark Silicon

- Performance per unit power
  - Systems increasingly limited by power consumption, not number of transistors
  - ➔ "**Dark Silicon**" : Most of the chip will be OFF to meet thermal limits

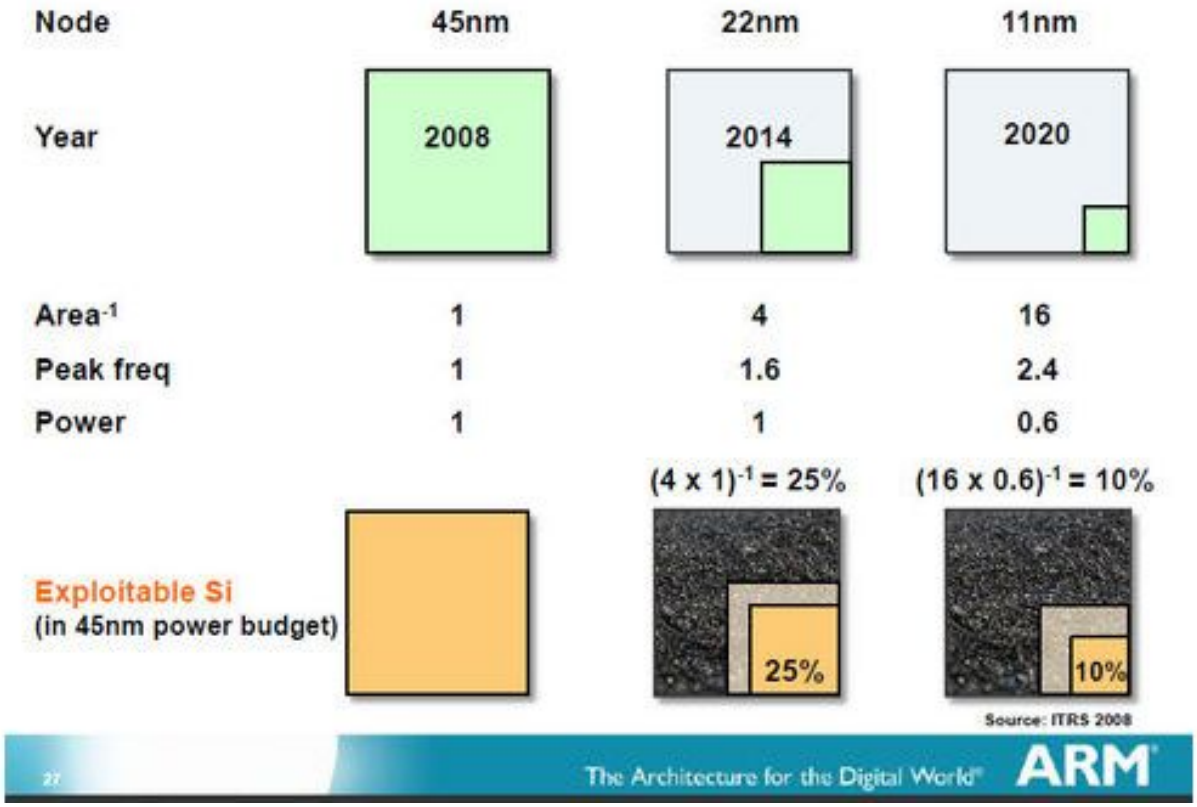

Energy per Operation

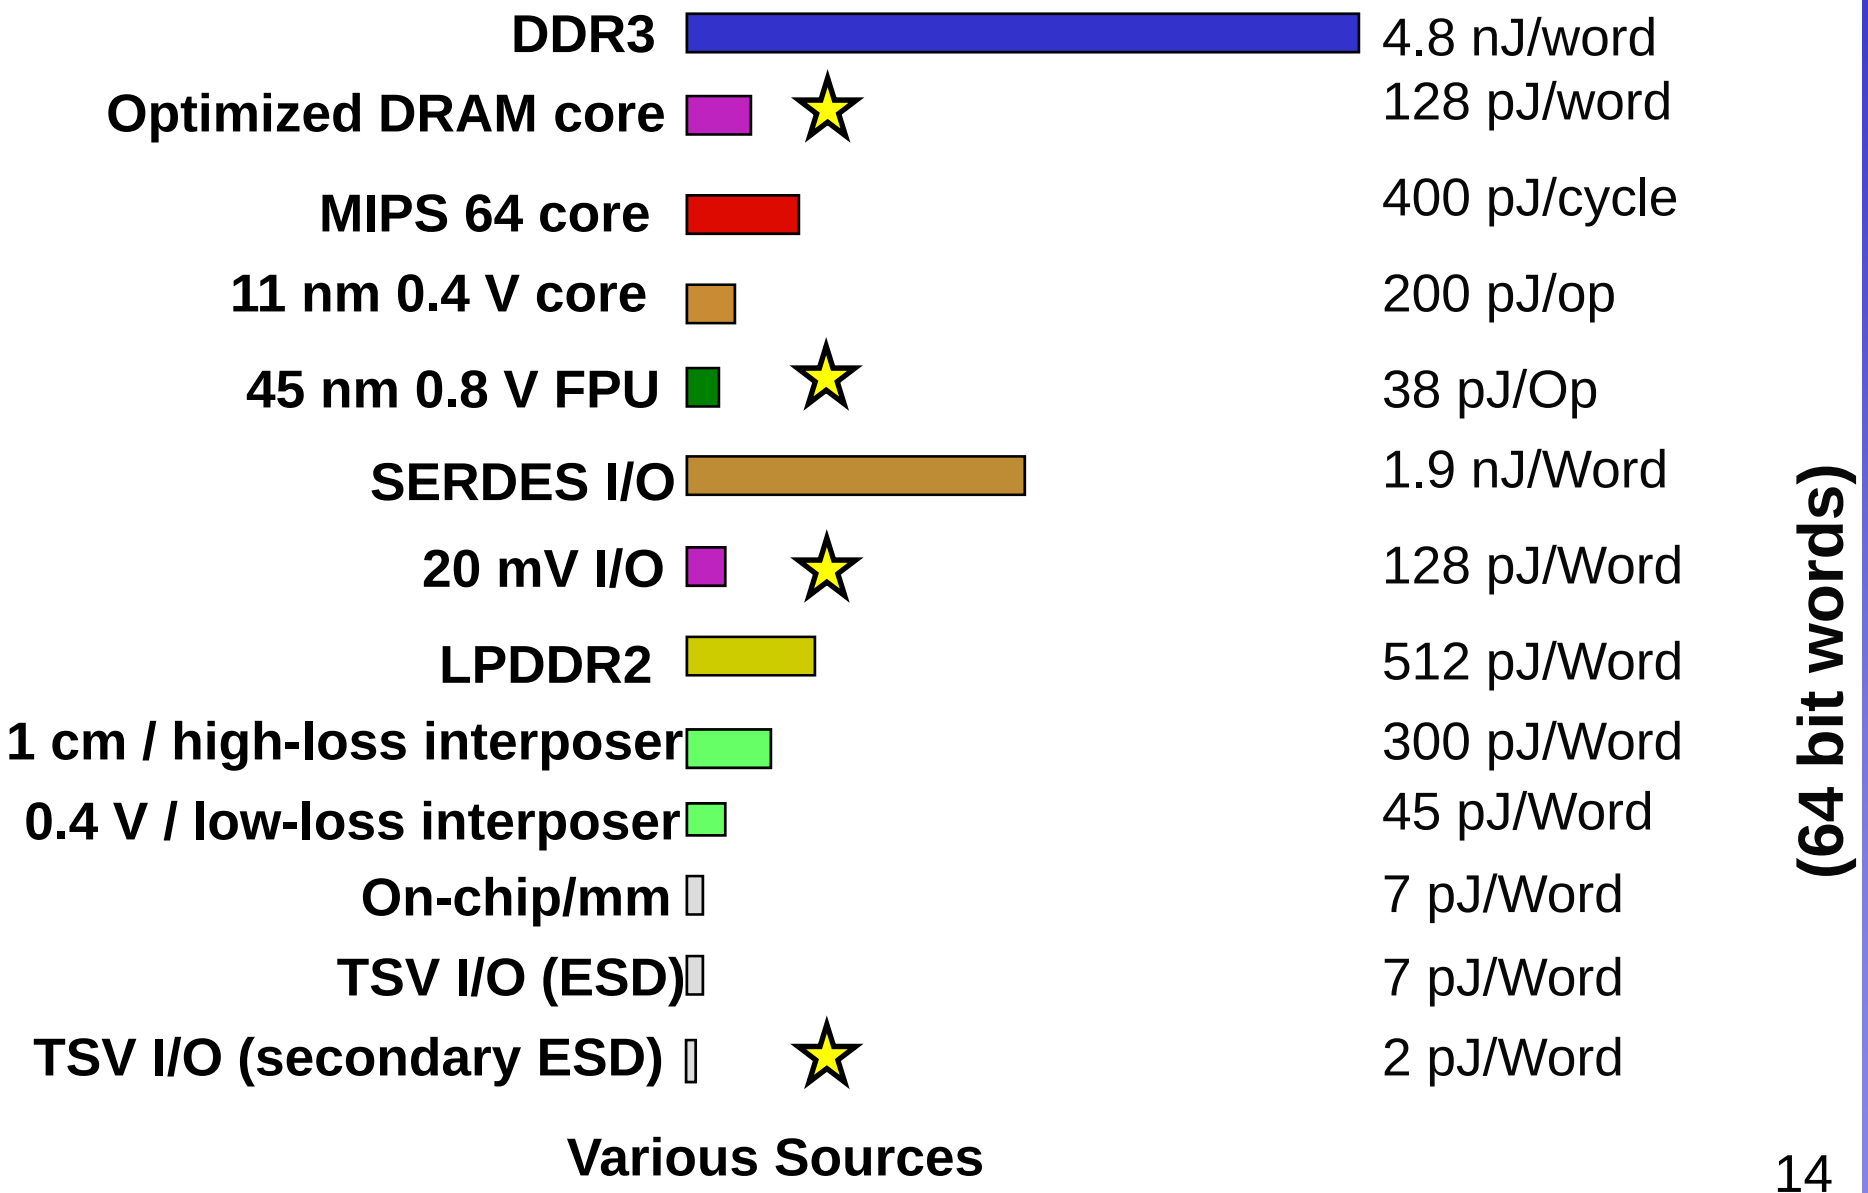

Synthetic Aperture Radar Processor

- Built FFT in Lincoln Labs 3D Process

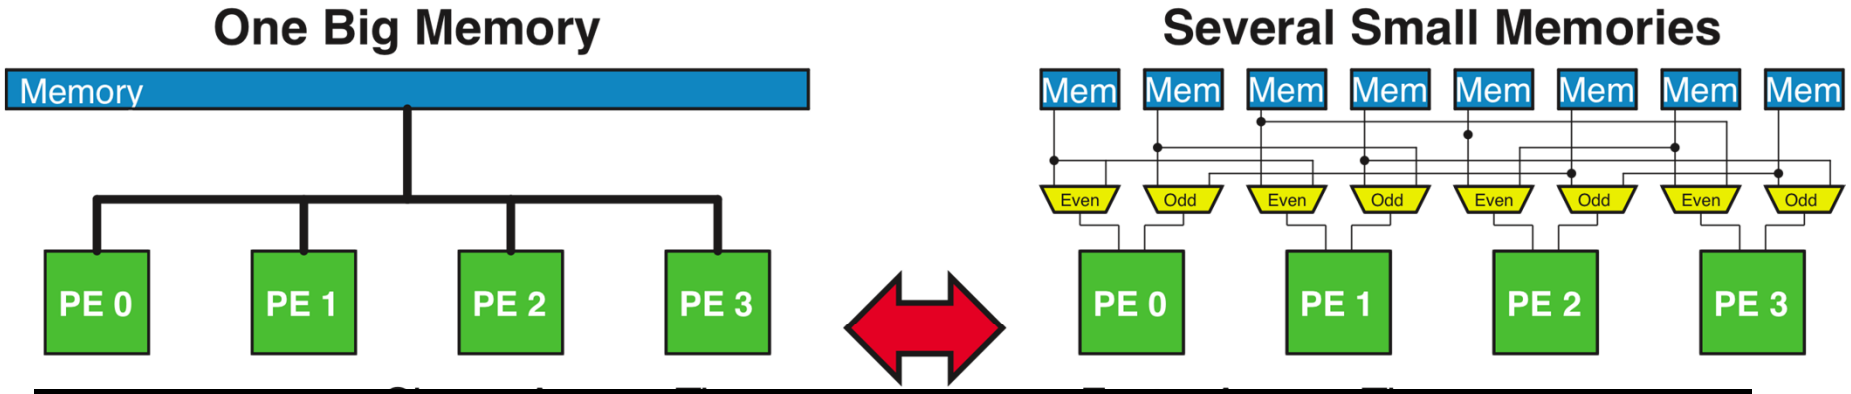

Metric	Undivided	Divided	%
Bandwidth (GBps)	13.4	128.4	+854.9
Energy Per Write(pJ)	14.48	6.142	-57.6
Energy Per Read (pJ)	68.205	26.718	-60.8
Memory Pins (#)	150	2272	+1414.7
Total Area (mm^2)	23.4	26.7	+16.8%

Thor Thorolfsson

3D FFT Floorplan

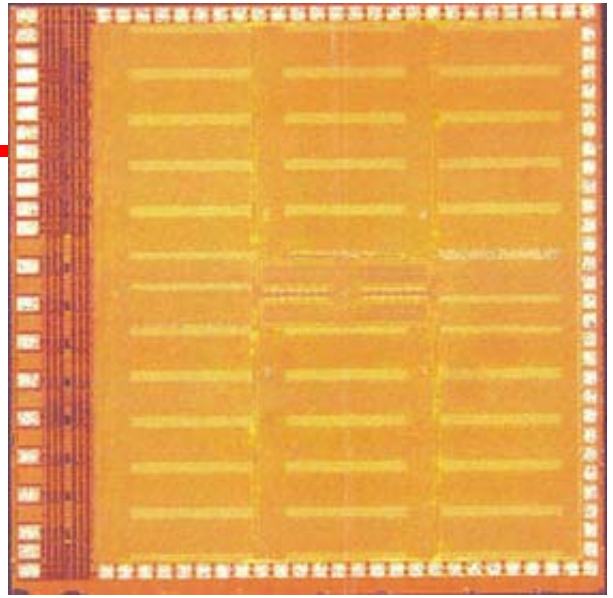

- All communications is vertical
- Support multiple small memories WITHOUT an interconnect penalty
  - AND Gives 60% memory power savings

RePartition FFT to Exploit Locality

- Every partition is a PE
- Every unique intersection is a memory

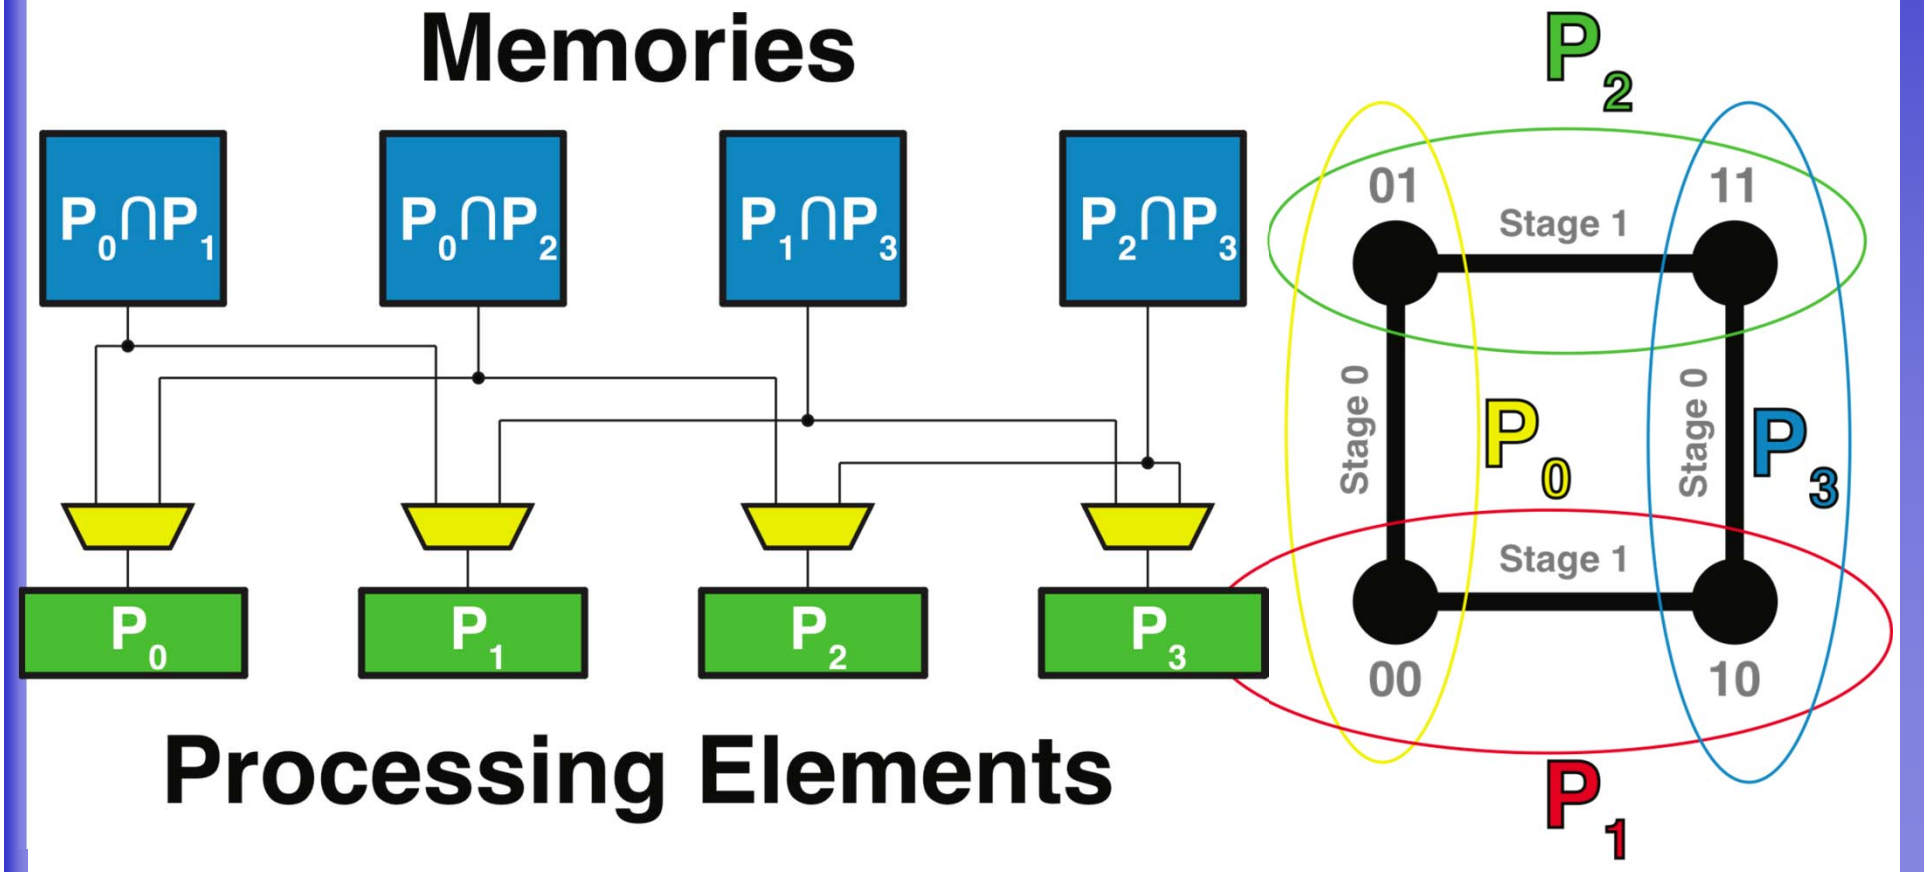

2DIC vs. 3DIC Implementation

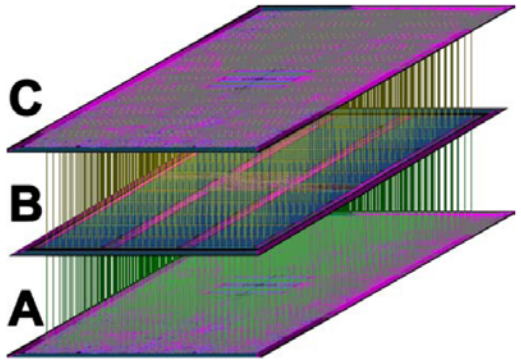

vs.

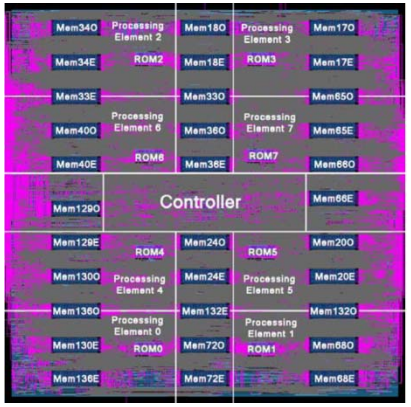

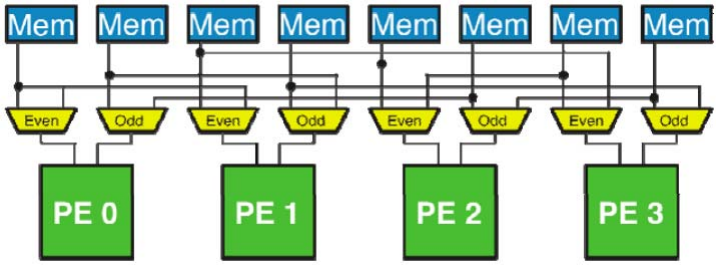

Metric	2D	3D	Change
Total Area (mm^2)	31.36	23.4	-25.3%
Total Wire Length (m)	19.107	8.238	-56.9%
Max Speed (Mhz)	63.7	79.4	+24.6%
Power @ 63.7MHz (mW)	340.0	324.9	-4.4%
FFT Logic Energy (µJ)	3.552	3.366	-5.2%

Thor Thorolfsson

Tezzaron 130 nm 3D SAR DSP

- Complete Synthetic Aperture Radar processor
  - 10.3 mW/GFLOPS
  - 2 layer 3D logic
- All Flip-flops on bottom partition
  - Removes need for 3D clock router
- HMETIS partitioning used to drive 3D placement

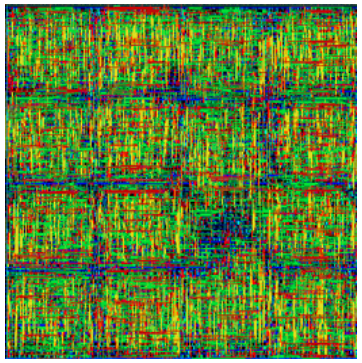

Logic only

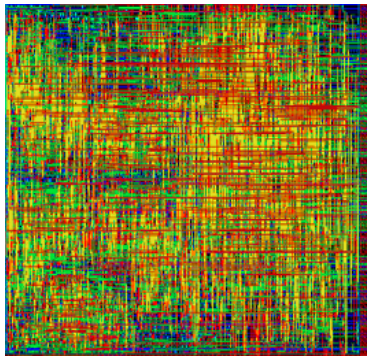

Logic, clocks, flip-flops

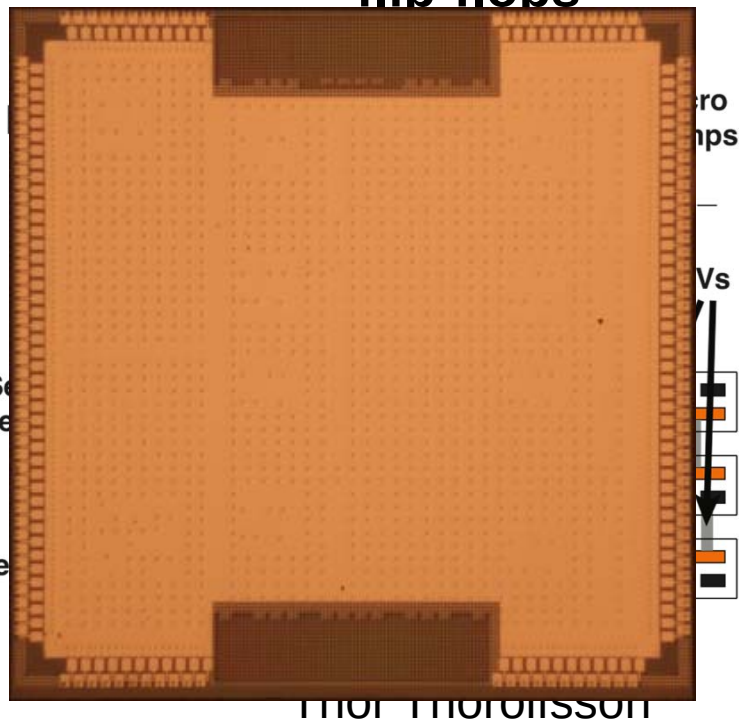

6.6 μm face to face

Sequential Commercial Tools

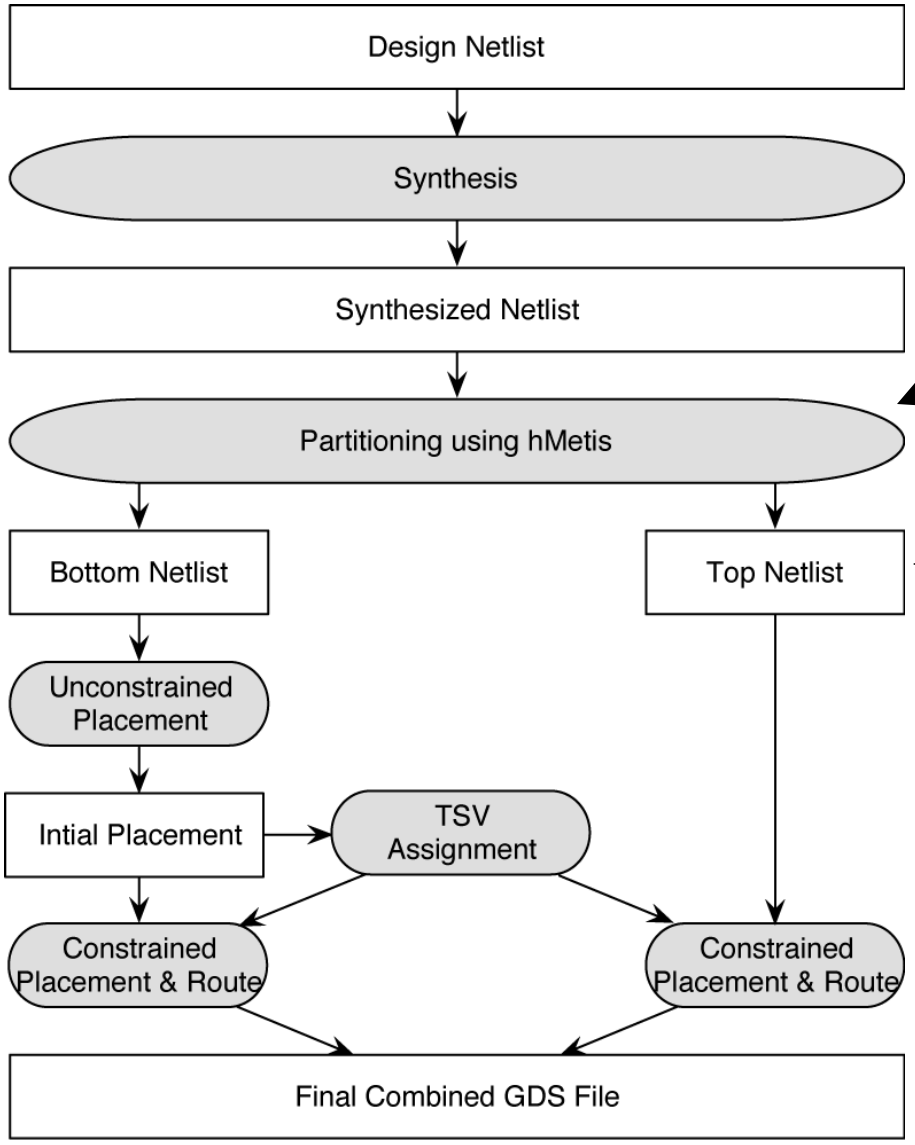

<u>hMetis</u>:

- **Balance top and bottom cell area**
- **Minimizing number of TSVs**

All clocked cells (Flip-flops) & IOs

(Clock distribution <u>entirely</u> in Top tier)

Shorter wires → Modest - Good Returns

- Relying on wire-length reduction alone is not enough

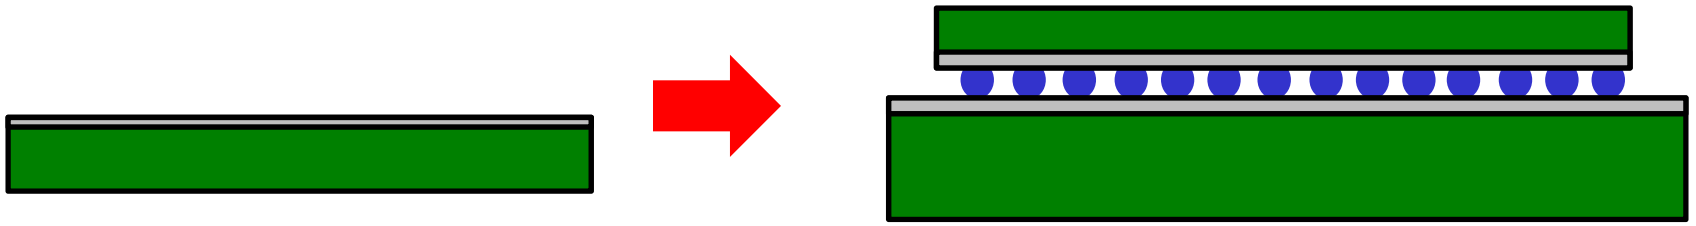

2D Design

0.13 μm Cell Placement split across 6.6 μm face-to-face bump structure

	Total Wire Length (% Change)	Max Frequency (% Change)	Parasitic Pow	
PE 3D Seq.	-17.1%			
PE 3D Sim.	-17.7%			-7.7%
PE 3D True	-21.0		45.2%	-12.9%
AES 3D Seq.	-8.0%		-19.6%	-2.6%
MIMO 3D Seq.	+216.1%	+17.1%	-34.9%	-5.1%

18% - 35% improvement in Power.Delay (21% for SAR)

Results get less compelling with technology scaling, as the microbumps don't scale as fast as the underlying process

Thorlindur Thorolfsson□ , Guojie Luo† , Jason Cong† and Paul D. Franzon□
□ Department of Electrical & Computer Engineering, North Carolina State University, Raleigh, NC 27695
† Computer Science Department, University of California, Los Angeles, CA 90095
Email: gluo@cs.ucla.edu and cong@cs.ucla.edu

Increasing the Return

1. 3D specific architectures

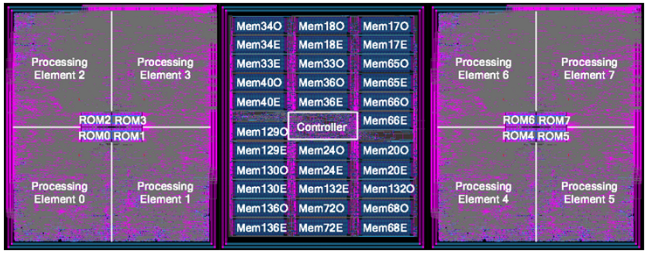

2. Exploiting Heterogeneity

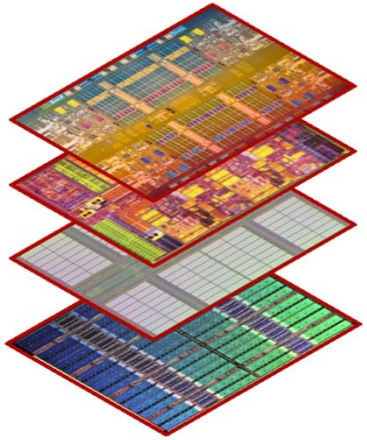

High Performance
Low Power/Accelerato
Specialized RAM
General RAM

3. Ultra 3D Scaling

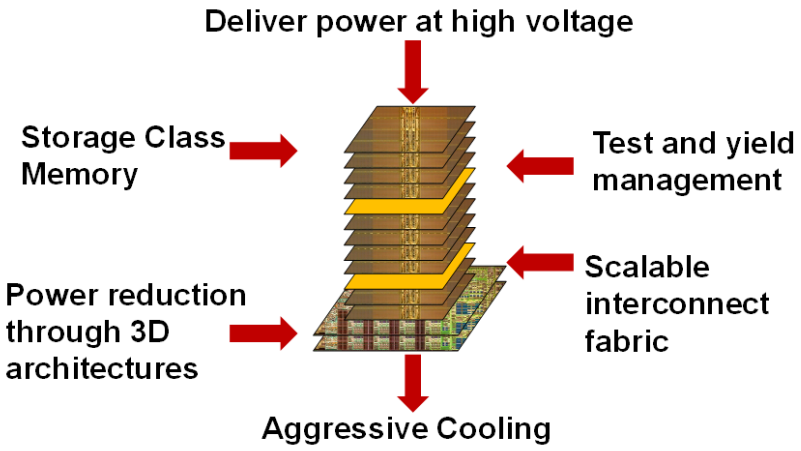

Extreme Integration

Motivation: Database Servers; High End DSP

Deliver power at high voltage

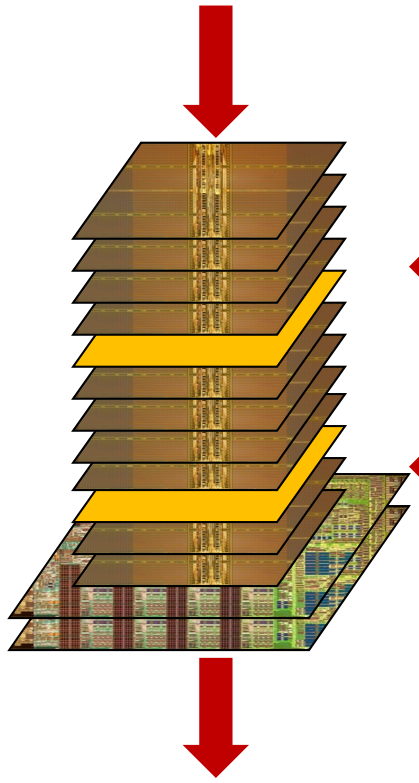

High Capacity Memory

Test and yield management

Scalable interconnect fabric

Power reduction through 3D architectures

Aggressive Cooling

3D Miniaturization

Miniature Sensors

- mm^3 scale - Human Implantable (with Jan Rabaey, UC(B))
- cm^3 scale - Food Safety & Agriculture (with KP Sandeep, NCSU)

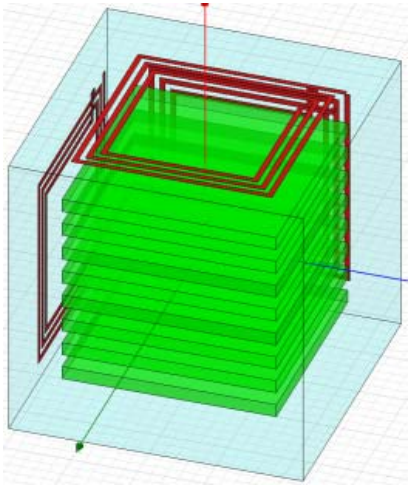

- Problems:
  - Power harvesting @ any angle (mm-scale)
  - Local power management (cm scale)

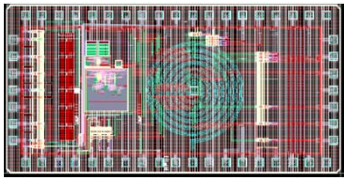

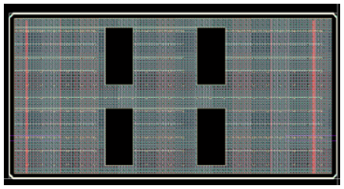

Peter Gadfort, Akalu Lentiro, Steve Lipa

Chip-Scale 3DIC Sensor

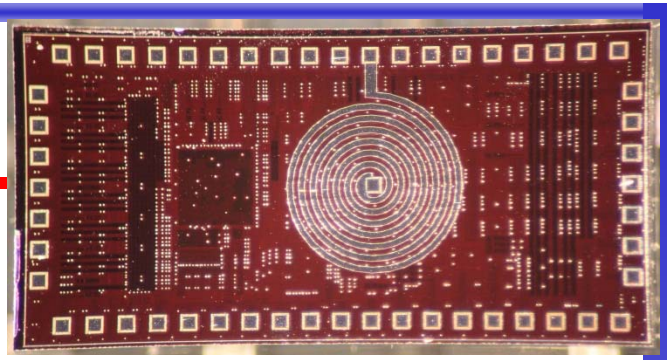

- Two tiers:
  - Processing tier and power-storage tier

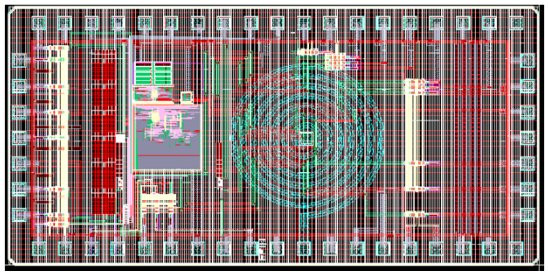

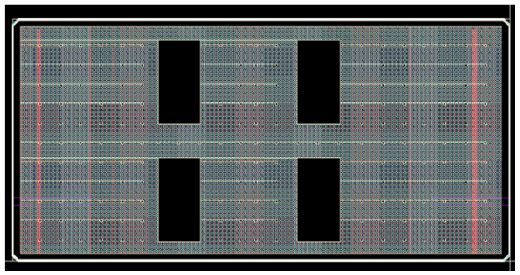

- **Sensor & ADC**
- **SRAM Memory**
- **RFID coil**
- **Power Interfaces to capacitor; battery and RFID power harvesting**
- **Back-scatter RFID communications interface**

- **Stacked Capacitor: 128 nF**
- **Absorbs short RFID power cycles and stabilizes Vdd**
- **Could be larger in a specialized technology**

Integrated with two 0.3 μAhr batteries

"True" 3D Integration

- **Orientation of mm-scale sensor will be random**
  - **Building antenna "through" 3DIC chip stack on edge will be very lossy**
  - **Need power harvesting on all 3 sides**
  - **Developed packaging integration flow to achieve this**

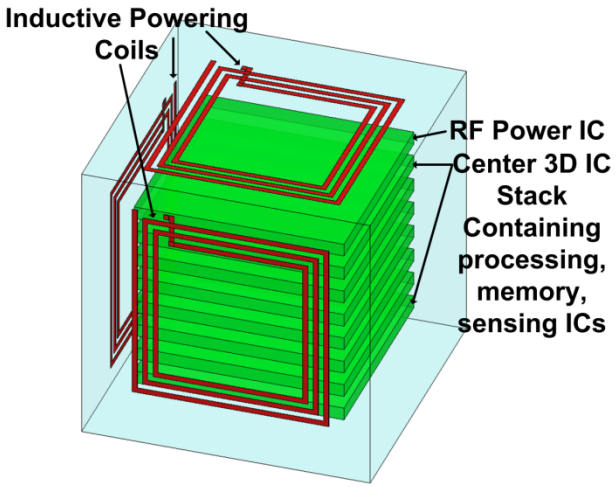

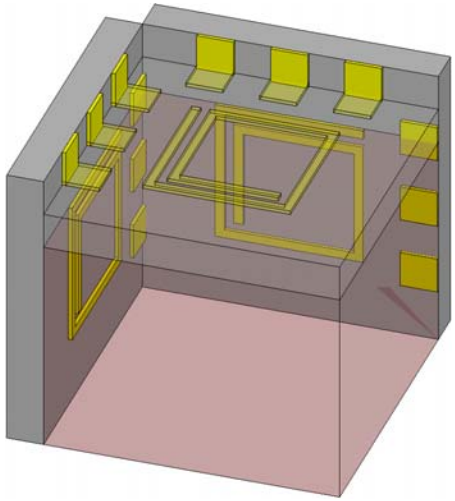

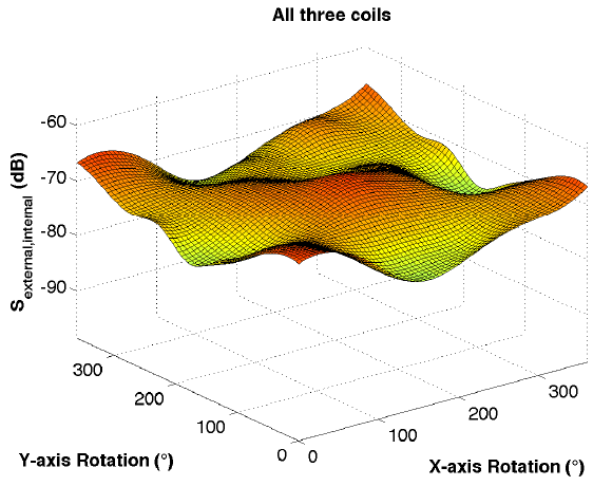

Peter Gadfort

“True” 3D Integration

- Cubes built at 3 mm and 5 mm scale
- 5 mm coils measured results:

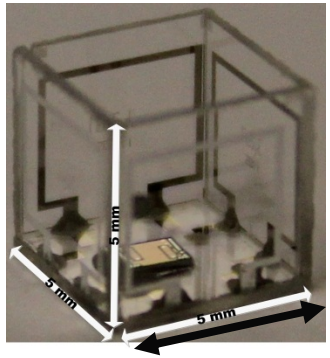

5 mm

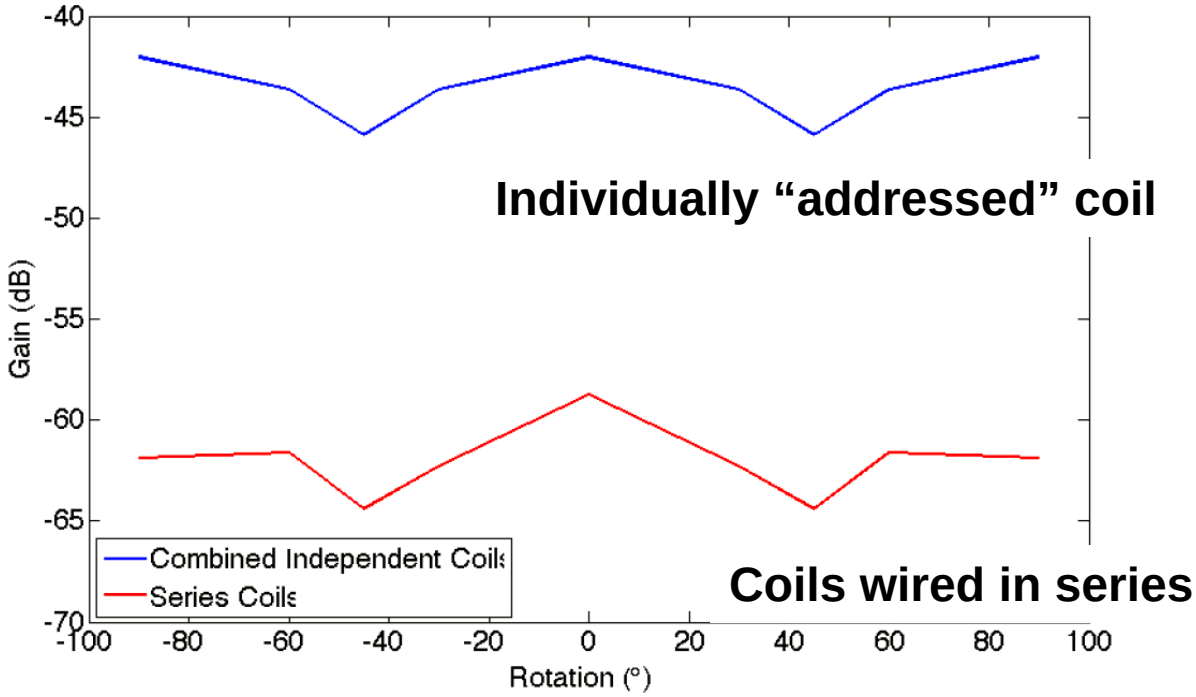

Individually “addressed” coil

Coils wired in series

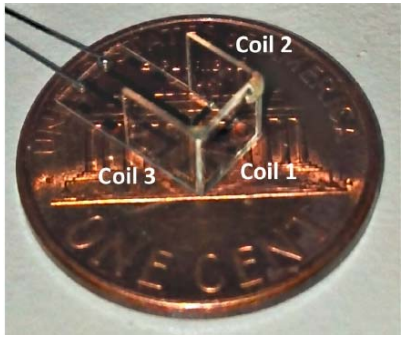

Peter Gadfort

Mid-term Barriers to Deployment

Barrier	Solutions
Thermal	Early System Codesign of floorplan and thermal evaluation
	DRAM thermal isolation
Test	Specialized test port & test flow
Codesign	“Pathfinding” in SystemC
	CAD Interchange Standards
Cost & Yield	Supporting low manufacturing cost through design

Long-term Barriers to Deployment

Barrier	Solutions
Thermal & Power Deliver	3D specific temperature management
	New structures and architectures for power delivery
Test & Yield management	Modular, scalable test and repair
Co-implementation	Support for Modularity and Scalability
Cost & Yield	Supporting low manufacturing cost through design

Early Codesign: Thermal Flow

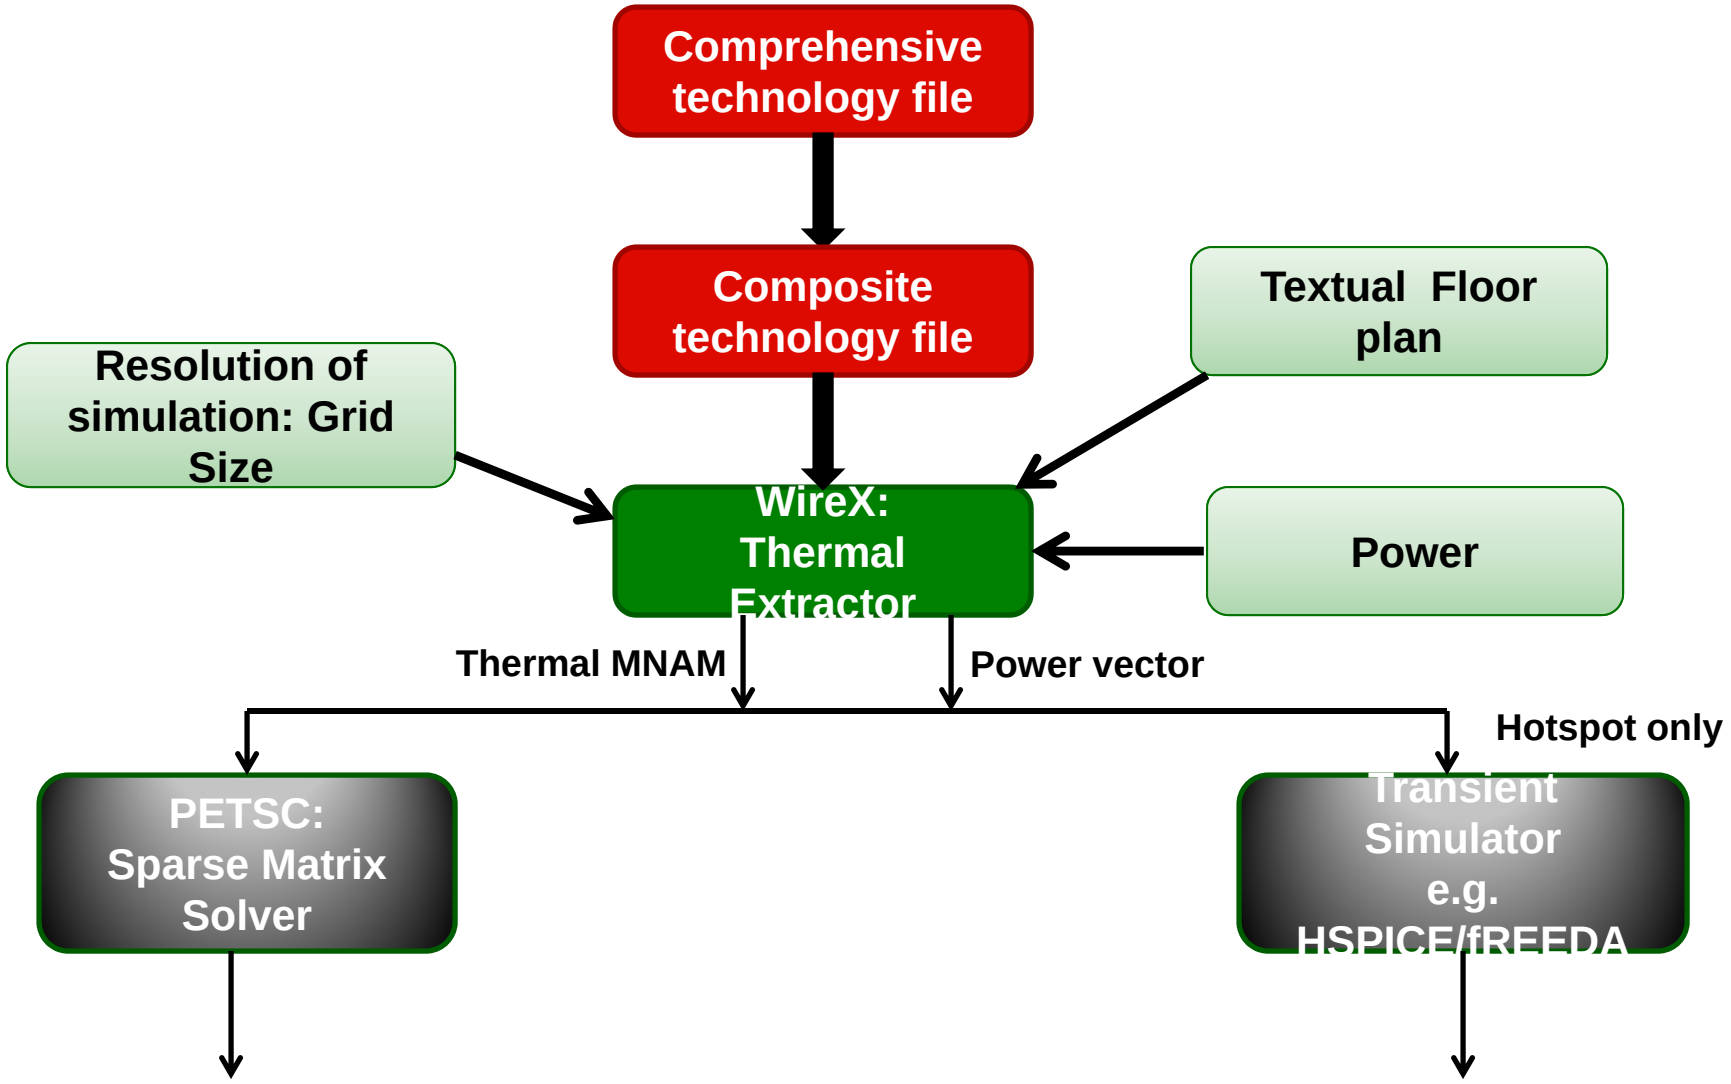

Pathfinder 3D:

- Goals:
  - Electronic System Level (ESL) codesign for fast investigation of performance, logic, power delivery, and thermal tradeoffs
  - Focus to date: Thermal/speed tradeoffs – static and transient
- Test case : Stacking of Heterogeneous Cores

Test case

- Stacking of a 4-wide average core and a 2-wide large window core on 45 nm
  - Rise in channel temperature on top tier.

Transient Thermal Profile

Transient Junction profile of two stacked processor one running “mcf” and other running “bzip” :

Modular Interfaces: Problem Statement

- Goal: “Plug and Play” 3D integration

- Despite:
  - Different technologies
  - Different process nodes
  - Different clock frequencies
  - Complex temporal requirements
  - Unknown power requirements
  - Unknown thermal constraints

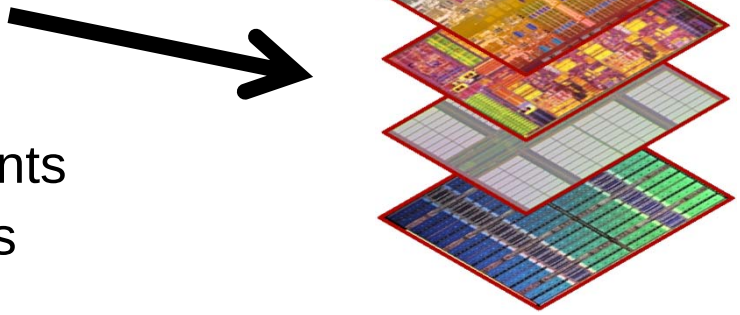

Intel Tick-Tock Development Model

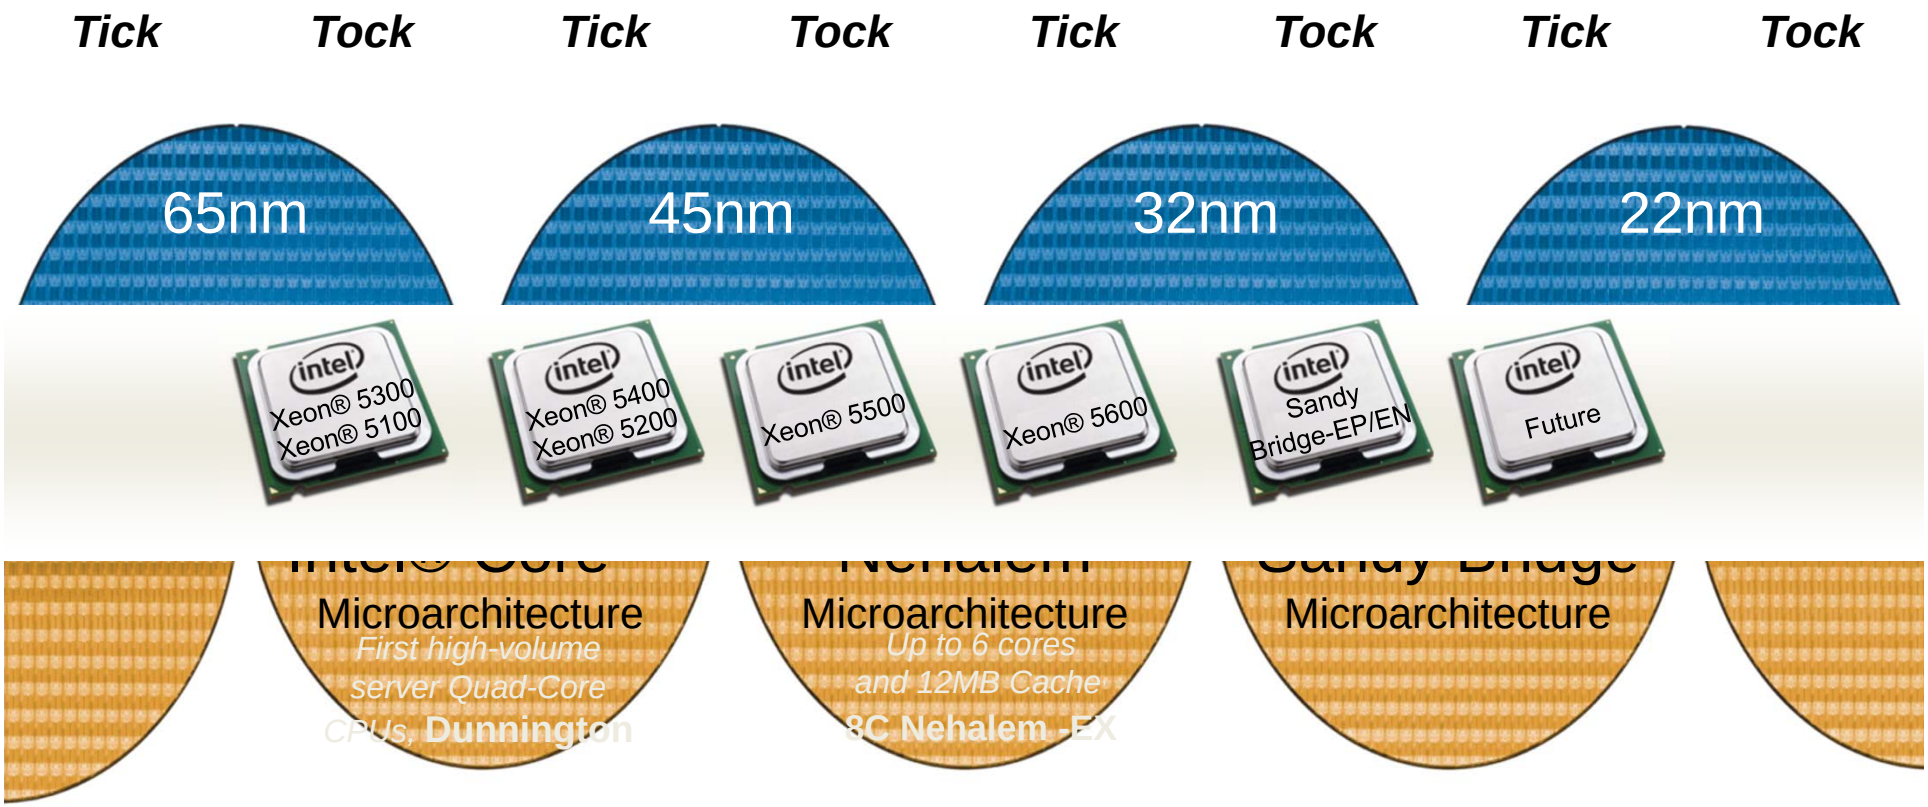

- Substantial design shrink on “tock”
- New architecure on “tick”

3D issues:
Heterogeneous Processors
➔ Hard to define 3D interconnect in advance

Opportunity

- "Plug and Play" 3D integration with very high bandwidth

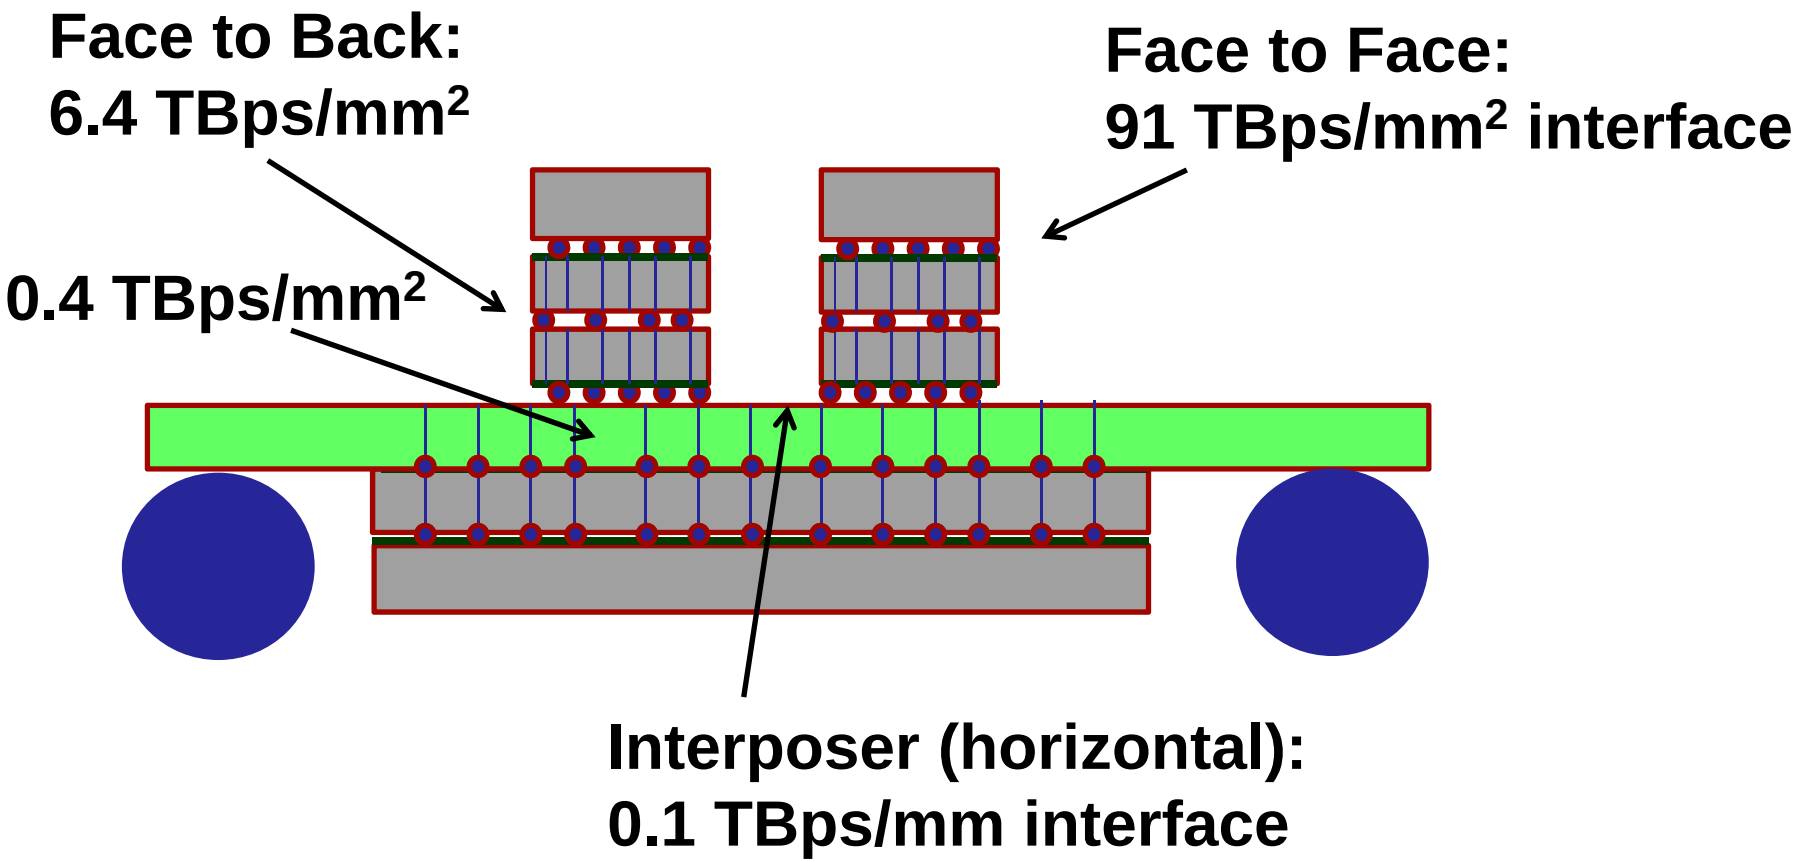

3D Specific Interface IP

Proposal:

Open Source IP for 3D and 2.5D interfaces

An interface specification that supports signaling, timing, power delivery, and thermal control within a 3D chip-stack, 2.5D (interposer) structure and SIP solutions

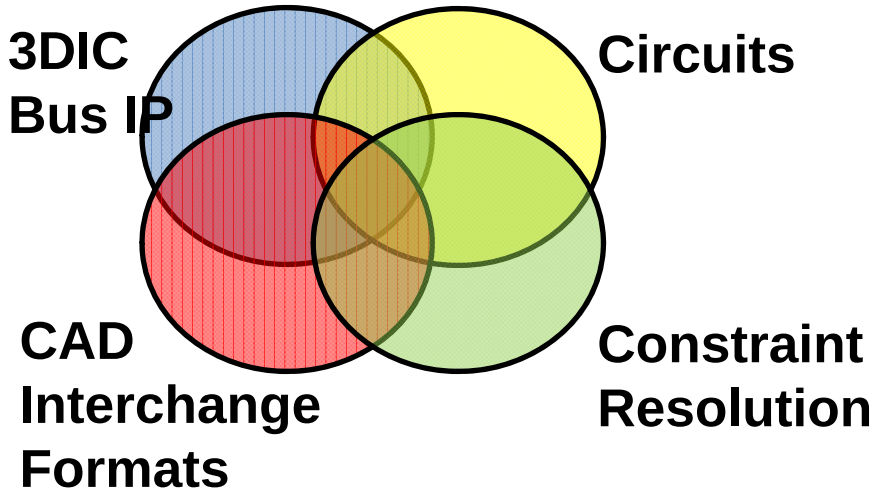

(Proc. 3DIC 2011)

Open Bus IP

- Amba style split cycle bus set

Circuits

- Tier-to-tier data forwarding without common clocks
- Requirements:
  - Fast – low latency, high bandwidth
  - Testable
  - Low-power
  - Reliable

CAD Interchange Formats

- Open standard formats to propagate EDA information from design to design, with a focus on
  - Floorplan constraints
  - Pin locations
  - Thermal constraints
  - Power delivery requirements
  - Design for Test

Constraint Resolution

- In-situ resolution of key constraints
  - E.g. Collaborative solutions for Thermal Mitigation

Upcoming Events

- Eworkshop on Open Standards for 3DIC IP
  - Please contact me if you are interested
  - paulf@ncsu.edu

- 2013 IEEE 3DIC Conference
  - Subscribe to email list
  - Email list
  - Email to mj2@lists.ncsu.edu
  - With “subscribe 3dic_conf” in main body

Conclusions

- Three dimensional integration offers potential to
  - Deliver memory bandwidth power-effectively;
  - Improve system power efficiency through 3D optimized codesign
  - Enable new products through aggressive Heterogeneous Integration

- Main challenges in 3D integration (from design perspective)
  - Effective early codesign to realize these advantages in workable solutions
  - Managing cost and yield, including test and test escape
  - Managing thermal, power and signal integrity while achieving performance goals
  - Scaling and interface scaling

NC STATE UNIVERSITY

Acknowledgements

Faculty: William Rhett Davis, Michael B. Steer, Eric Rotenberg,
Professionals: Steven Lipa, Neil DiSpigna,
Students: Hua Hao, Samson Melamed, Peter Gadfort, Akalu Lentiro,
Shivam Priyadarshi, Christopher Mineo, Julie Oh, Won Ha Choi,
Zhou Yang, Ambirish Sule, Gary Charles, Thor Thorolfsson,
Department of Electrical and Computer Engineering
NC State University