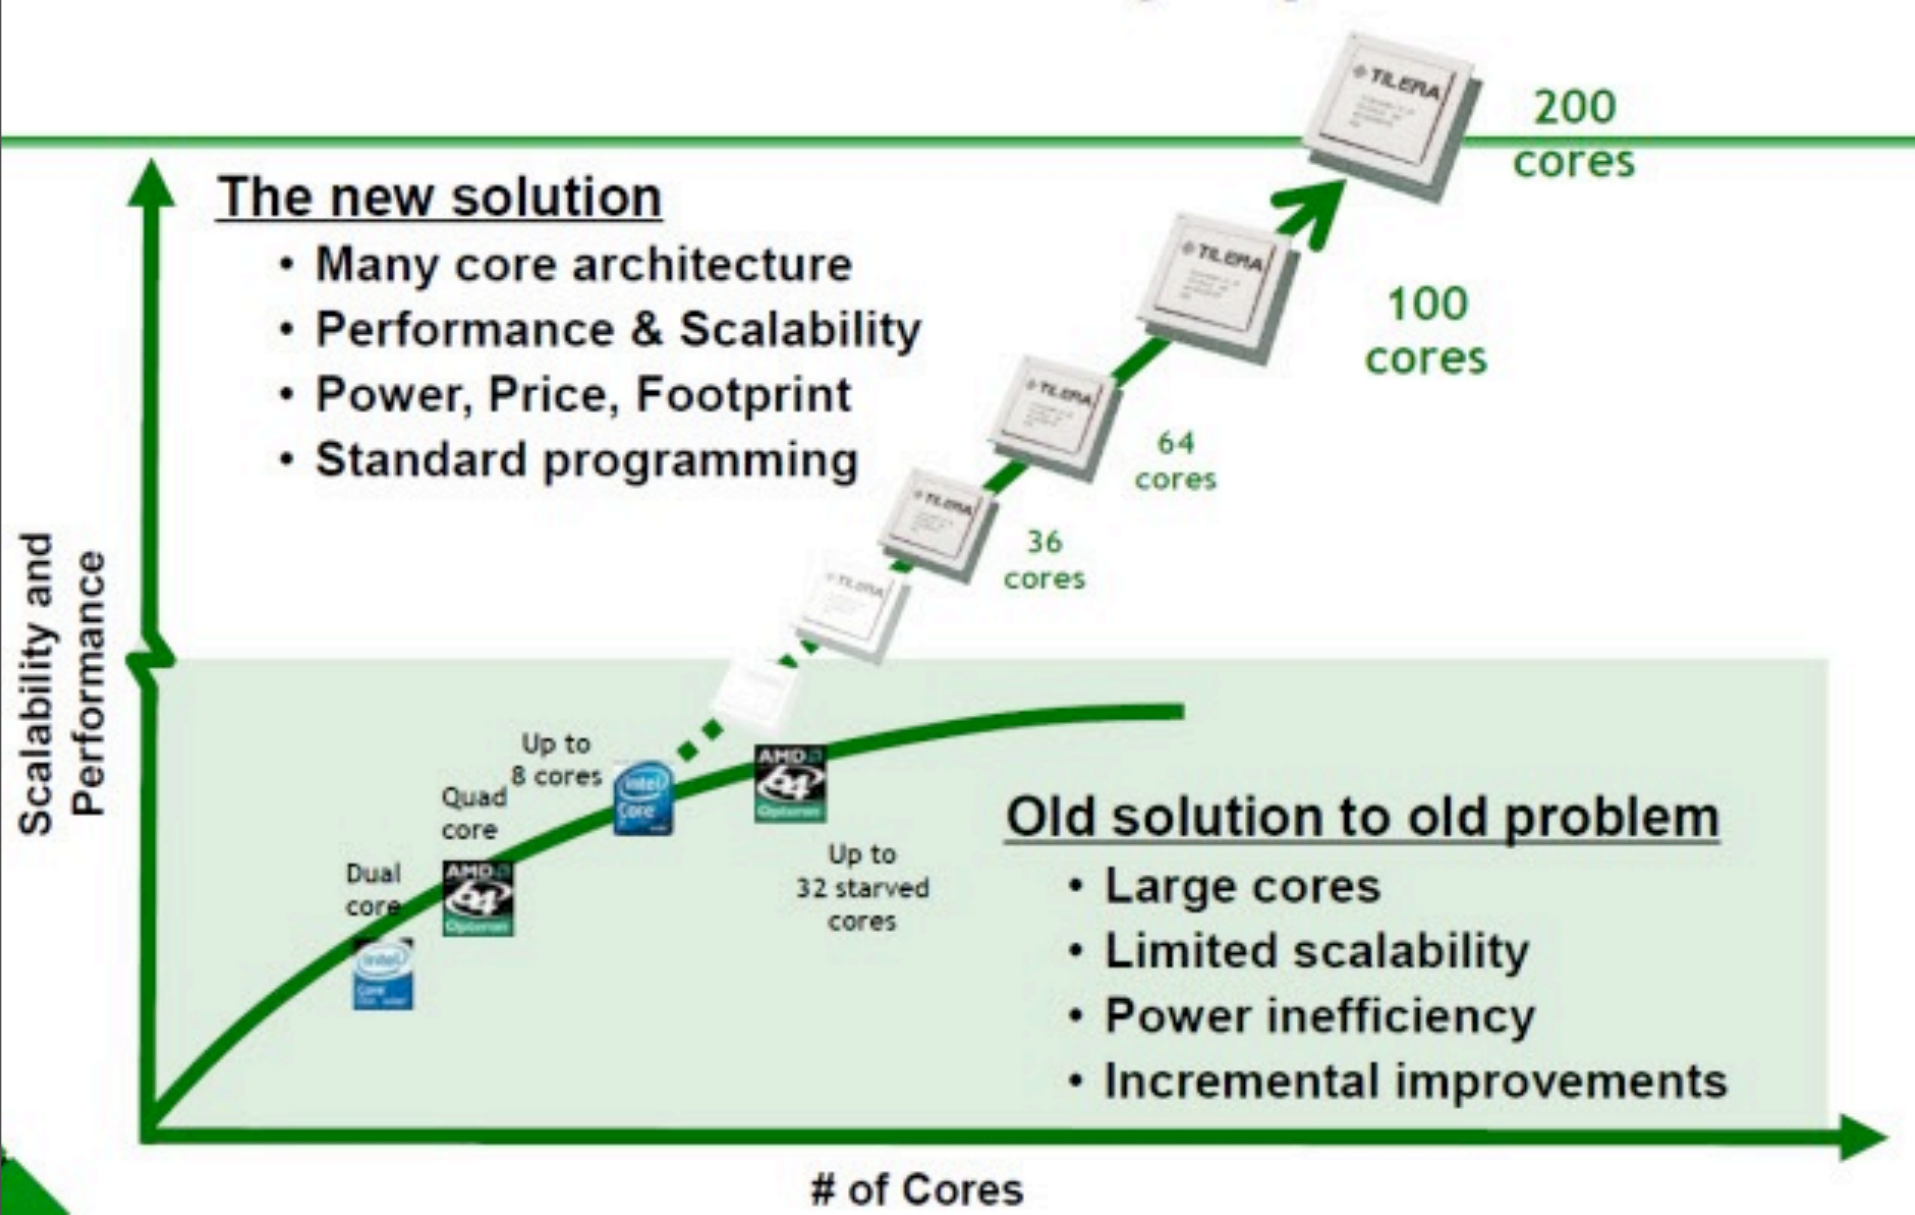


A microscopic view of a Tiler processor die, showing a grid of small, square blocks. The blocks are arranged in a regular pattern and are colored in various colors, including red, yellow, green, blue, and purple. The die is divided into several large, rectangular sections by thin lines.


Tiler (RAW) Processor

Jason Mars and Robbie Hott

Tilera: the solution to today's problems



RAW Beginnings

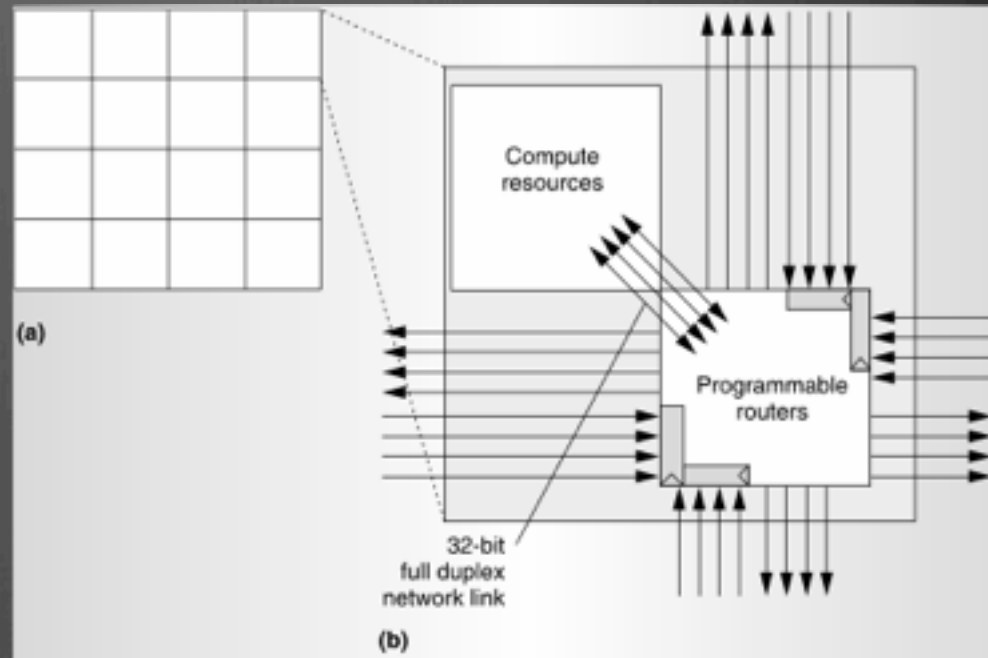


RAW Philosophy

Table 1. How Raw converts physical resources into architectural entities.*

Physical entity	Raw ISA analog	Conventional ISA analog
Gates	Tiles, new instructions	New instructions
Wire delay	Network hops	None
Pins	I/O ports	None


*Conventional ISAs attempt to utilize increasing gate quantities through the addition of new instructions (like parallel SIMD instructions) and through dynamic mapping of operations to a small number of architecturally invisible ALUs. Wire delay is typically hidden through pipelining and speculation, and is reflected to the user in the form of dynamic stalls for non-fast-path and mispredicted code. Pin bandwidth is hidden behind speculative cache-miss hardware prefetching and large line sizes.



Some networks are mem mapped (registers), others use explicit messages

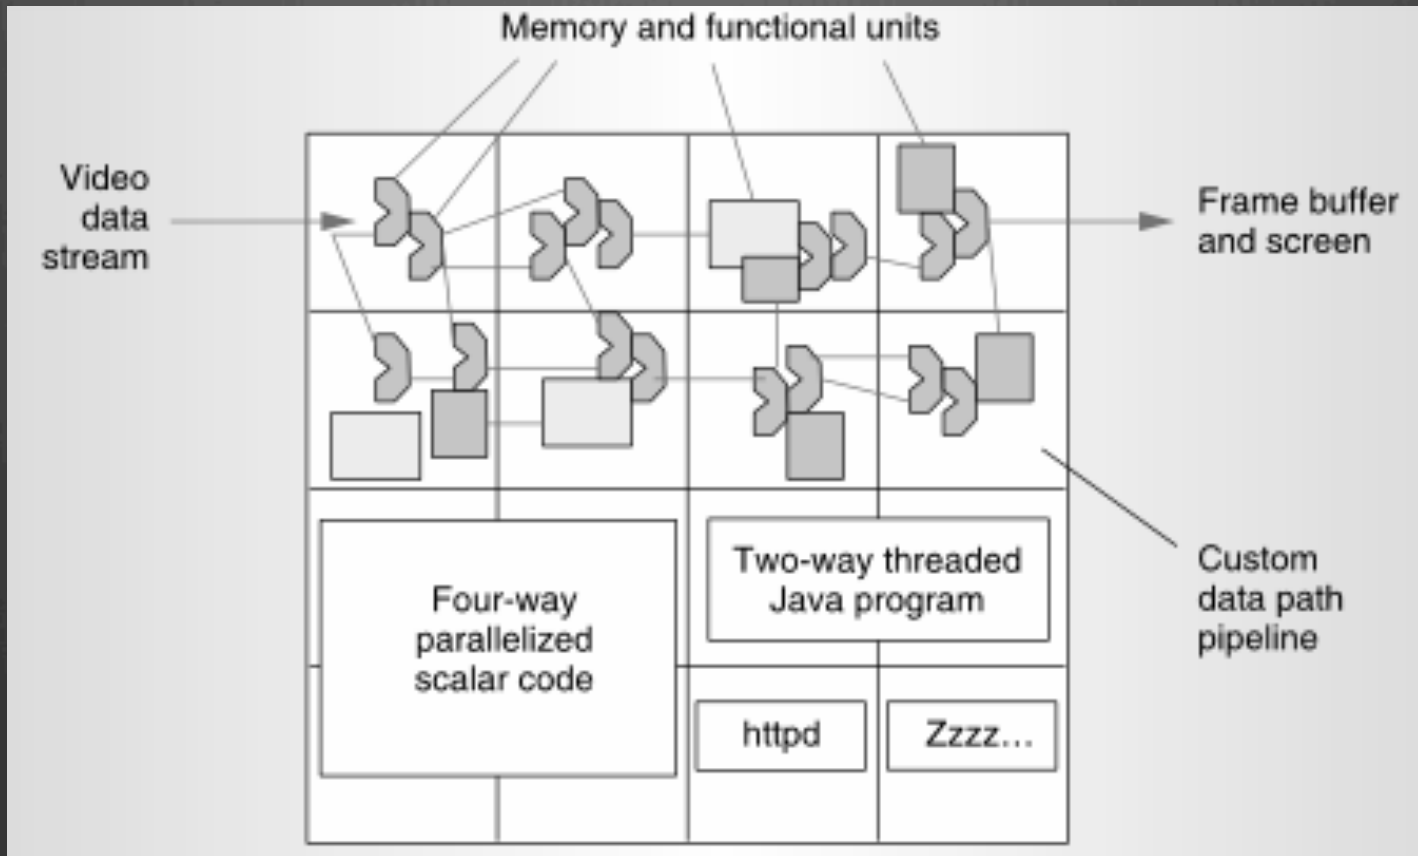
Routers are programmable

16 Tile RAW



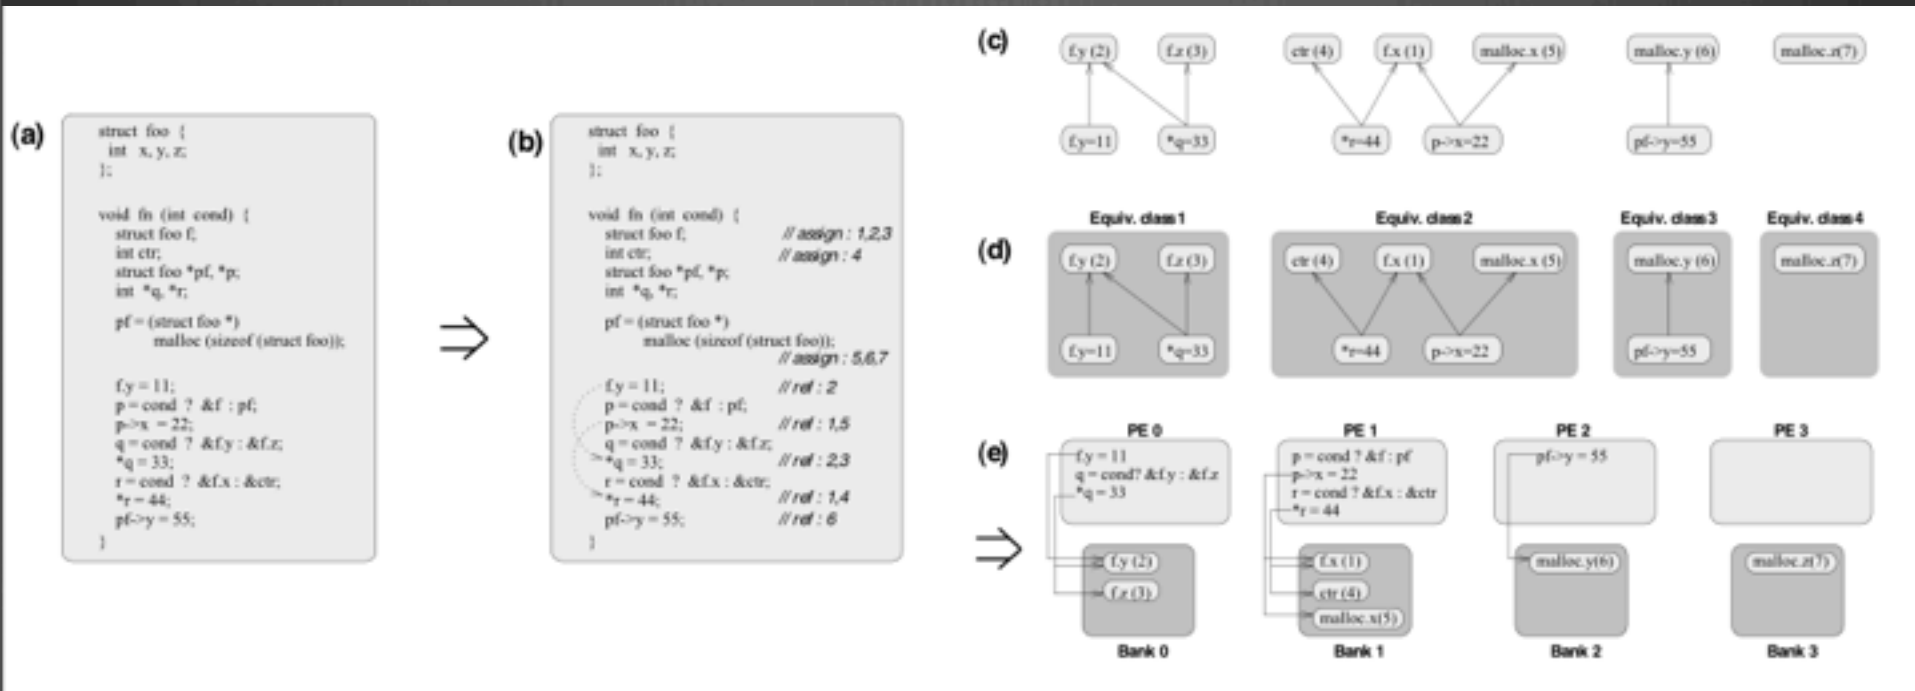
3 cycles between tiles

“Software” ASICs



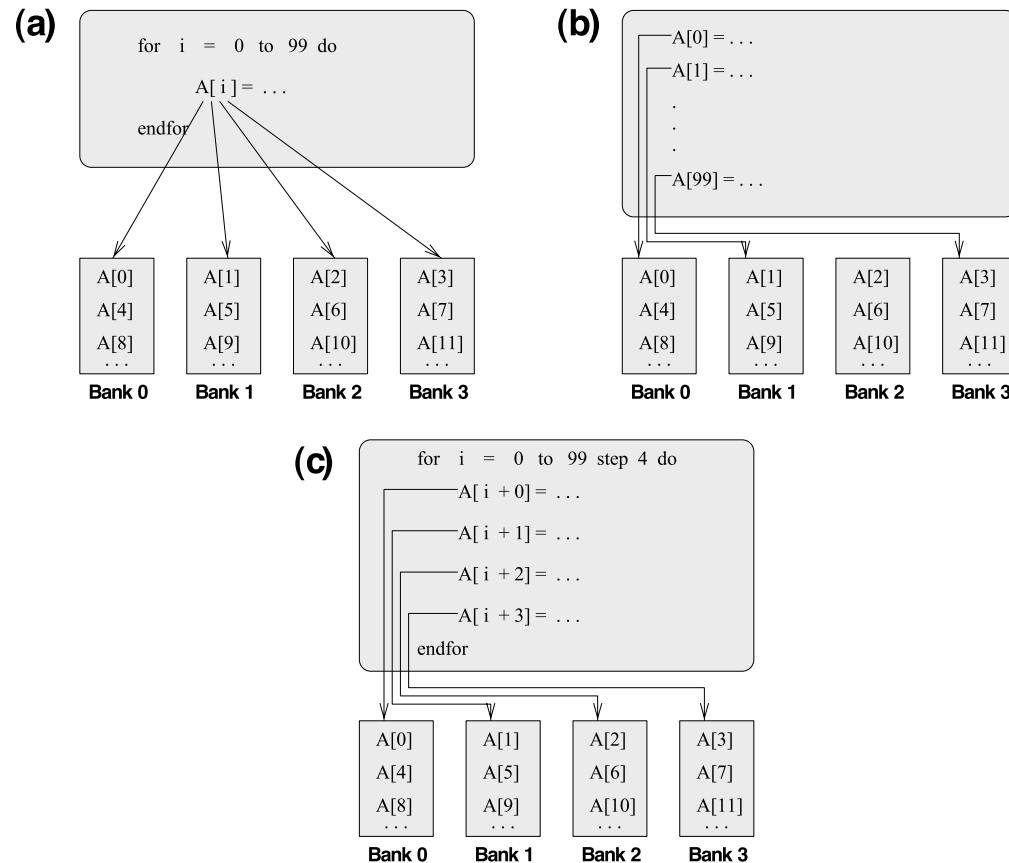
Place and route software circuits (video)
5x faster than 700mhz PIII

Compiling for RAW: Equivalence Class Unification



Based on pointer analysis

Compiling for RAW: Modulo Unrolling



Basic Parallelization

Comparing RAW to PIII

Parameter	Raw (IBM ASIC)	P3 (Intel)
Lithography Generation	180 nm	180 nm
Process Name	CMOS 7SF (SA-27E)	P858
Metal Layers	Cu 6	Al 6
Dielectric Material	SiO ₂	SiOF
Oxide Thickness (T _{ox})	3.5 nm	3.0 nm
SRAM Cell Size	4.8 μm ²	5.6 μm ²
Dielectric <i>k</i>	4.1	3.55
Ring Oscillator Stage (FO1)	23 ps	11 ps
Dynamic Logic, Custom Macros (SRAMs, RFs)	no	yes
Speedpath Tuning since First Silicon	no	yes
Initial Frequency	425 MHz	500-733 MHz
Die Area ²	331 mm ²	106 mm ²
Signal Pins	~ 1100	~ 190
Vdd used	1.8 V	1.65 V
Nominal Process Vdd	1.8 V	1.5 V

Win Some / Lose Some



Benchmark	Source	# Raw Tiles	Cycles on Raw	Speedup vs P3	
				Cycles	Time
<i>Dense-Matrix Scientific Applications</i>					
Swim	Spec95	16	14.5M	4.0	2.9
Tomcatv	Nasa7:Spec92	16	2.05M	1.9	1.3
Btrix	Nasa7:Spec92	16	516K	6.1	4.3
Cholesky	Nasa7:Spec92	16	3.09M	2.4	1.7
Mxm	Nasa7:Spec92	16	247K	2.0	1.4
Vpenta	Nasa7:Spec92	16	272K	9.1	6.4
Jacobi	Raw bench. suite	16	40.6K	6.9	4.9
Life	Raw bench. suite	16	332K	4.1	2.9
<i>Sparse-Matrix/Integer/Irregular Applications</i>					
SHA	Perl Oasis	16	768K	1.8	1.3
AES Decode	FIPS-197	16	292K	1.3	0.96
Fpppp-kernel	Nasa7:Spec92	16	169K	4.8	3.4
Unstructured	CHAOS	16	5.81M	1.4	1.0

Benchmark	Number of tiles				
	1	2	4	8	16
<i>Dense-Matrix Scientific Applications</i>					
Swim	1.0	1.1	2.4	4.7	9.0
Tomcatv	1.0	1.3	3.0	5.3	8.2
Btrix	1.0	1.7	5.5	15.1	33.4
Cholesky	1.0	1.8	4.8	9.0	10.3
Mxm	1.0	1.4	4.6	6.6	8.3
Vpenta	1.0	2.1	7.6	20.8	41.8
Jacobi	1.0	2.6	6.1	13.2	22.6
Life	1.0	1.0	2.4	5.9	12.6
<i>Sparse-Matrix/Integer/Irregular Applications</i>					
SHA	1.0	1.5	1.2	1.6	2.1
AES Decode	1.0	1.5	2.5	3.2	3.4
Fpppp-kernel	1.0	0.9	1.8	3.7	6.9
Unstructured	1.0	1.8	3.2	3.5	3.1

Benchmark	Source	# Raw Tiles	Cycles on Raw	Speedup vs P3	
				Cycles	Time
172.mgrid	SPECfp	1	.240B	0.97	0.69
173.applu	SPECfp	1	.324B	0.92	0.65
177.mesa	SPECfp	1	2.40B	0.74	0.53
183.equake	SPECfp	1	.866B	0.97	0.69
188.ammmp	SPECfp	1	7.16B	0.65	0.46
301.apsi	SPECfp	1	1.05B	0.55	0.39
175.vpr	SPECint	1	2.52B	0.69	0.49
181.mcf	SPECint	1	4.31B	0.46	0.33
197.parser	SPECint	1	6.23B	0.68	0.48
256.bzip2	SPECint	1	3.10B	0.66	0.47
300.twolf	SPECint	1	1.96B	0.57	0.41

Benchmark	Cycles on Raw	Speedup vs P3		Efficiency
		Cycles	Time	
172.mgrid	.240B	15.0	10.6	96%
173.applu	.324B	14.0	9.9	96%
177.mesa	2.40B	11.8	8.4	99%
183.equake	.866B	15.1	10.7	97%
188.ammmp	7.16B	9.1	6.5	87%
301.apsi	1.05B	8.5	6.0	96%
175.vpr	2.52B	10.9	7.7	98%
181.mcf	4.31B	5.5	3.9	74%
197.parser	6.23B	10.1	7.2	92%
256.bzip2	3.10B	10.0	7.1	94%
300.twolf	1.96B	8.5	6.1	94%

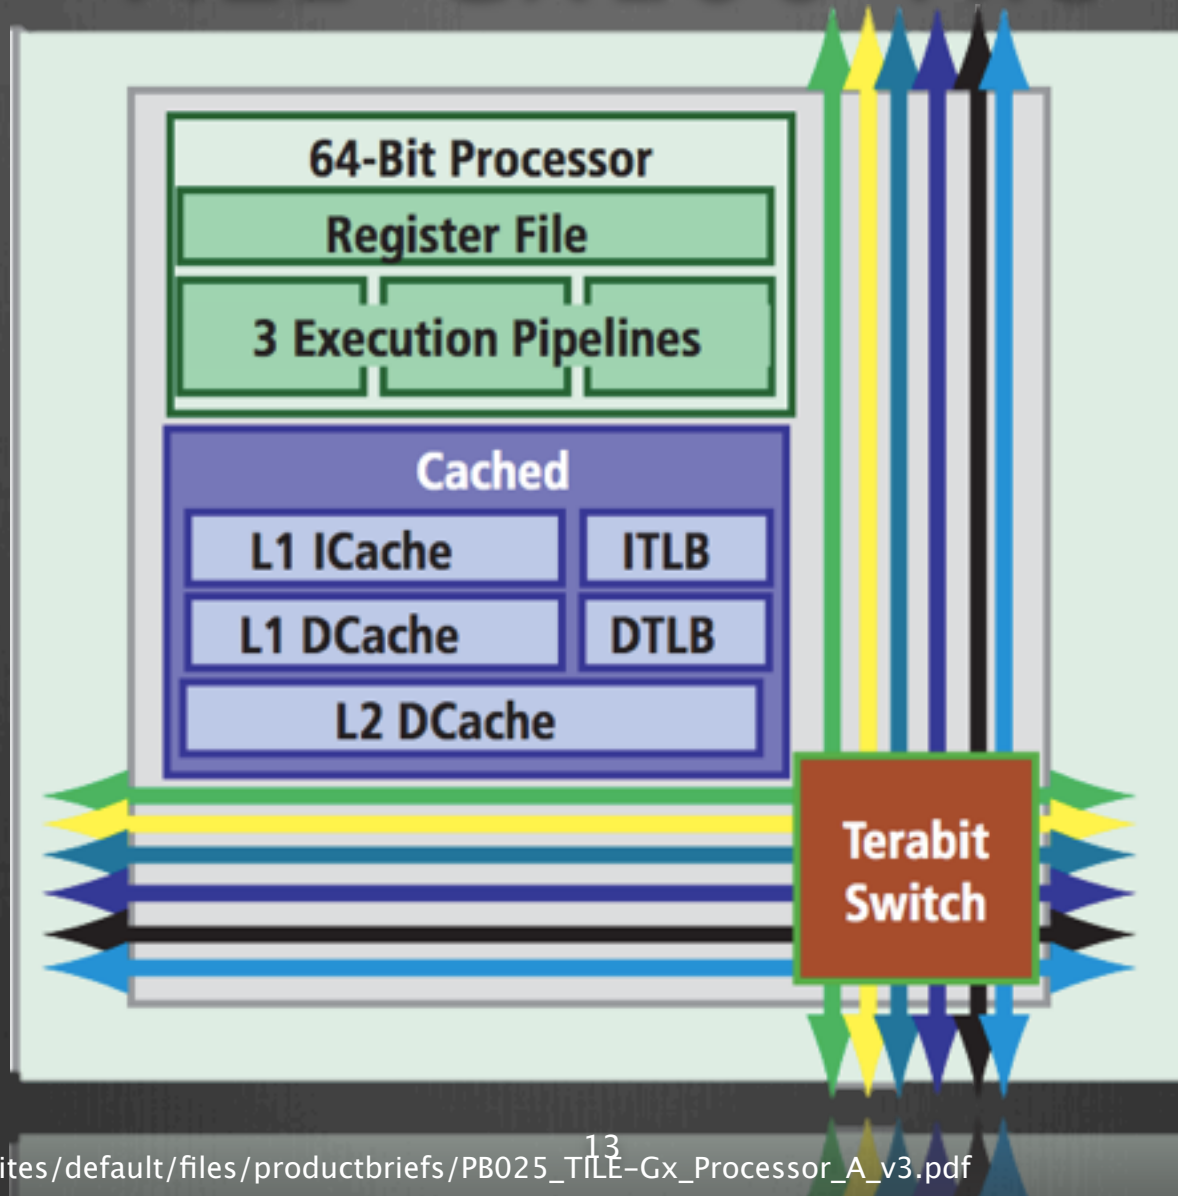
RAW to Tiler




Tilera TILE-Gx Series

- Low-power multi-core RISC architecture
 - 16, 36, 64, and 100 core models
 - Up to 1.5 GHz clock frequency
 - On-chip tiled mesh network
 - Each tile could operate as an individual processor
 - Multiple tiles can be used to run SMP Linux
 - TILE-Gx Tiles include:
 - 64-bit VLIW cores, 3-wide pipeline
 - 64-entry register file
 - DSP and SIMD extensions

TILE-Gx100 Tile



TILE-Gx100 Chip Overview



TILE-Gx100 Memory Hierarchy

- Scalable caching system: 32MB total on-chip cache
 - Each tile has L1 and L2 cache
 - 32 KB L1 Instruction Cache
 - 32 KB L1 Data Cache
 - 256 KB L2 Cache
 - Access through on-chip network
 - “patent pending DDC™ (Dynamic Distributed Cache) technology provides a fully coherent shared cache system across an arbitrarily-sized array of tiles”
 - No large centralized cache
 - TileDirect: coherent I/O directly into the tile caches

Tilera Extra Features


On-chip Connections


-  4 DDR3 Memory controllers

-  PCIe, USB, and Network interfaces

-  MiCA (Multistream iMesh Crypto Accelerator) for encryption, hashing, public key ops, at 40Gbps / 50,000 RSA ops/sec

On-chip network eliminates on-chip bus interconnect

-  Information must flow between processor cores or between cores and the memory / I/O controllers

-  iMesh network provides each tile with more than 1Tbps of interconnect bandwidth

Power Consumption

Processor	Number of Cores	Frequency	Avg Power Consumption
TILEPro36	36	500MHz	9–13W
TILE64/Pro64	64	700MHz/866MHz	15–23W (all cores)
TILE-Gx36	36	1.25GHz/1.5GHz	10–55W
TILE-Gx64	64	1.25GHz/1.5GHz	10–55W
TILE-Gx100	100	1.25GHz/1.5GHz	10–55W
Intel Core i7-920	4	2.66GHz	130W (max TDP)

Package Size

Processor	Number of Cores	Package Size
TILEPro36	36	40mm x 40mm
TILE64/Pro64	64	40mm x 40mm
TILE-Gx36	36	35mm x 35mm
TILE-Gx64	64	45mm x 45mm
TILE-Gx100	100	45mm x 45mm
Intel Core i7-920	4	42.5mm x 45mm

Dark Silicon?

- With so many cores, something has to be off?
 - TILEPro64 draws up to 23W with ALL CORES running
- Individual idle cores can be turned off
 - Q: How best to configure 100 cores?

Intended Market

⦿ General purpose processor market

- ⦿ TILE-Gx can run multiple OSes and applications simultaneously

⦿ Four main categories

- ⦿ Networking Machines (monitoring, firewall, vpn)
- ⦿ Wireless
- ⦿ Multimedia Production (streaming, conferencing)
- ⦿ Cloud computing (servers)

⦿ Server Market

- ⦿ Around 10,000 cores in an 8kW rack
- ⦿ Quanta S2Q Server: 8 TILEPro64 chips (512 cores) at 400W


Further Questions

- Interesting research areas into how to best control 100 cores on chip
- How best to organize cache data in the distributed cache model
- TILE-Gx100 vs Intel Core/Xeon/Atom benchmarks
- How might the cores have changed since RAW?
- What would it take to displace Intel? Which markets?

Further Reading

- Tiler Homepage: www.tilera.com
- RAW Publications: groups.csail.mit.edu/cag/raw/documents
- M. Taylor et al. “The Raw Microprocessor: A Computational Fabric for Software Circuits and General-Purpose Programs,” IEEE Micro, March 2002.
- J. Kim, M. B. Taylor, J. Miller, and D. Wentzlaff, “Energy Characterization of a Tiled Architecture Processor with On-Chip Networks,” in 2003 ISLPED, 2003, pp. 424–427.

Questions?



25+ connector

25+ connector