

Versatile Tiled-Processor Architectures: The Raw Approach

Rodric M. Rabbah, Ian Bratt, and Anant Agarwal
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, MA 02139

Recent advances in VLSI technology have created an increasing interest within the computer architecture community to build a new kind of “general purpose” processor that is able to run a broad class of applications including primarily those from the domain of embedded systems—graphics, wireless processing, networking, and various forms of signal processing. The interest in new architectures is compounded by a growing wire delay concern which limits the distance that information can travel in a single clock cycle. The realities of interconnect delay—and power consumption—seriously challenge the ability of microprocessor designers to fulfill the promise of Moore’s Law. As a result, new architecture designs are largely centered around scalable and distributed alternatives to current centralized microprocessor designs.

Several projects such as VIRAM [2] at Berkeley, Smart Memories [4] at Stanford, TRIPS [5] at UT-Austin, Raw [8] and SCALE [3] at MIT, and industrial efforts such as the Tarantula [1] extension to Alpha, have proposed architectures that organize silicon resources more effectively and as *tiled-processors* that are easily scalable. The DARPA program in Polymorphic Computing Architectures is also a research thrust in this new area, and emerging “polymorphic” architectures will eventually compete with traditional desktop processors (e.g., Pentium IV) not so much in better performance on desktop workloads, but in *versatility*, or the ability to run a broader class of applications more effectively. We also expect that architectures that are more versatile are also likely to run complex real-world applications more effectively, since complex applications are often comprised of diverse components. One such versatile, tiled-processor architecture (TPA), is the Raw microprocessor which was designed and implemented at MIT.

Raw divides the chip into a two-dimensional mesh of sixteen programmable tiles, and interconnects them through on-chip, point-to-point scalar operand networks (SON) [7]. The Raw processor can issue sixteen different floating-point, integer, load, store, or branch instructions each cycle. It also has a large set of registers and a distributed memory hierarchy. The SON is exposed to the Raw compilation infrastructure which orchestrates the flow of data within the network for streaming computation and fine-grained instruction-level parallel-processing.

The focus on TPAs and architectural versatility necessitates *new benchmark suites and metrics* to accurately reflect the goals of the architecture community. Toward that end, we propose both a new benchmark suite—*VersaBench*—and a new metric called *Versatility*. VersaBench is a collection of applications from three central tiers—desktops, servers, and embedded systems—encompassing traditional integer workloads, floating-point and scientific applications, server computing, stream processing, and bit-level computation. VersaBench thereby attempts to better characterize the broad set of workloads that the new tiled-processor architectures are required to run.

The Versatility metric is inspired by SPEC rates [6]. For example, the SPEC CINT89 rate for an architecture is the geometric mean of the speedups of that architecture relative to a reference machine (specifically, the VAX 11/780)¹ for each of the applications in the SPEC CINT89 suite. Computing the Versatility of an architecture is purposefully designed to mirror that of SPEC rates. Accordingly, like SPEC, Versatility takes the geometric mean of the speedups of an architecture for each of the applications in the VersaBench suite. Unlike SPEC rates however, the speedup of each application is not computed relative to a single reference machine, but rather relative to the architecture which provides the *best* performance for that application (in the 2004 time frame from known results at the time of this writing).

¹The reference machines have changed over time. While the VAX 11/780 was the reference machine for SPEC CINT89 and SPEC CINT92, the SPARCstation 10/40 was the reference machine for SPEC CINT95, and the Sun Ultra5-10 workstation with a 300MHz SPARC processor is the reference machine for SPEC CINT2000.

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE 01 FEB 2005	2. REPORT TYPE N/A	3. DATES COVERED -	
4. TITLE AND SUBTITLE Versatile Tiled-Processor Architectures: The Raw Approach		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) Computer Science and Artificial Intelligence Laboratory Massachusetts Institute of Technology Cambridge, MA 02139		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 See also ADM001742, HPEC-7 Volume 1, Proceedings of the Eighth Annual High Performance Embedded Computing (HPEC) Workshops, 28-30 September 2004. , The original document contains color images.			
14. ABSTRACT			
15. SUBJECT TERMS			
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified	UU
			18. NUMBER OF PAGES 32
			19a. NAME OF RESPONSIBLE PERSON

Table 1. Characteristics of the VersaBench workloads.

benchmark category	data type	parallelism	control complexity	temporal locality	spatial locality
Desktop Integer	integer	low	high	high	low
Desktop Float	float	medium	medium	medium	medium
Server	integer/float	high	medium to high	medium to high	medium to low
Embedded Stream	integer/float/bit	very high	low	low to high	very high
Embedded Bit	bit	very high	very low	very low	very high

Presentation Outline

The presentation will describe the Raw architecture, its implementation, and performance. We will focus on Raw’s ability to support a diverse set of applications (ranging from desktop to embedded workloads) and multiple forms of parallelism (including instruction-level-parallelism (ILP) for desktop applications, and stream parallelism for embedded computing) as represented by the VersaBench suite. We will also report detailed performance measurements that quantify the versatility of Raw compared to some widely deployed architectures. As a prelude, the measured versatility of the Raw processor is 0.7, while that of the Pentium III is 0.1. The Pentium’s relatively poor performance on stream benchmarks hurts its versatility. Although Raw’s versatility is better in comparison, the VersaBench suite highlights two clear areas that merit additional research. The first is in improving the architecture to better support embedded bit-level workloads: ASICs perform 2x-3x better than Raw. Another area of research focuses on desktop integer applications: Raw’s performance is 2x lower than a Pentium III for applications with low degrees of ILP.

VersaBench

The VersaBench suite consists of fifteen benchmarks that are grouped into five categories that span desktop, server, and embedded application workloads. The VersaBench applications are themselves drawn from various suites, and were selected because of the salient behavior and the properties they exhibit along various dimensions. For example, it is widely accepted that desktop applications have complex control structures, whereas streaming applications consist of relatively small computational kernels with simple control mechanisms. In all, we consider five property-dimensions when characterizing a benchmark; they are

- *predominant data type*: summarizes the predominant type-domain over which computation is performed,
- *parallelism*: quantifies maximum IPC (instructions per cycle) in a benchmark,
- *control complexity*: measures instruction temporal locality,
- *data temporal locality*
- *data spatial locality*

Intuitively, we believe the properties of each of the five benchmark-categories are as shown in Table 1. Accordingly, the VersaBench suite was created systematically by measuring the properties of numerous applications and selecting those that matched intuition. The presentation will include detailed results that map applications into the five-dimensional space.

Versatility Metric

We define the Versatility of an architecture as the geometric mean of the speedup of every application in the VersaBench suite relative to the architecture which provides the best performance for that application. Thus architectural versatility becomes quantitative, with ASICs (application specific integrated circuits) occupying the lowest end of the spectrum (i.e., $\text{Versatility}(\text{ASIC}) = 0$); as future process-technologies deliver higher clock frequencies, architectural versatility will increase beyond unity. Versatility measure servers to identify the areas where opportunities for architectural improvements are greatest, and those where further efforts will lead to marginal returns. For example, a new architecture that performs as well as a Pentium IV for desktop workloads has little to gain from further improvements targeted toward that application domain. In contrast, an architecture that fares poorly compared to the best streaming processor warrants attention that is focused on improving the performance of that architecture within the streaming context.

REFERENCES

- [1] R. Espasa, F. Ardanaz, J. Gago, R. Gramunt, I. Hernandez, T. Juan, J. Emer, S. Felix, G. Lowney, M. Mattina, and A. Seznev. Tarantula: A Vector Extension to the Alpha Architecture. In *Proceedings of the 29th International Symposium on Computer Architecture (ISCA)*, pages 281–292, 2002.
- [2] C. E. Kozyrakis and D. Patterson. A New Direction for Computer Architecture Research. *IEEE Computer*, 30(9), Sept. 1997.
- [3] R. Krashinsky, C. Batten, M. Hampton, S. Gerding, B. Pharris, J. Casper, , and K. Asanovic. The vector-thread architecture. In *Proceedings of the 31st International Symposium on Computer Architecture (ISCA)*, 2004.
- [4] K. Mai, T. Paaske, N. Jayasena, R. Ho, W. Dally, and M. Horowitz. Smart memories: A modular reconfigurable architecture. In *Proceedings of the 27th International Symposium on Computer Architecture*, pages 161–170, 2000.
- [5] R. Nagarajan, K. Sankaralingam, D. Burger, and S. W. Keckler. A Design Space Evaluation of Grid Processor Architectures. In *International Symposium on Microarchitecture (MICRO)*, December 2001.
- [6] STANDARD PERFORMANCE EVALUATION CORPORATION. <http://www.spec.org>.
- [7] M. Taylor, W. Lee, S. Amarasinghe, and A. Agarwal. Scalar Operand Networks: On-chip Interconnect for ILP in Partitioned Architectures. In *International Symposium on High Performance Computer Architecture (HPCA)*, Feb 2003.
- [8] E. Waingold, M. Taylor, D. Srikrishna, V. Sarkar, W. Lee, V. Lee, J. Kim, M. Frank, P. Finch, R. Barua, J. Babb, S. Amarasinghe, and A. Agarwal. Baring It All to Software: Raw Machines. *IEEE Computer*, 30(9):86–93, Sept. 1997. Also available as MIT-LCS-TR-709.

Versatile Tiled-Processor Architectures The Raw Approach

Rodric M. Rabbah

with Ian Bratt, Krste Asanovic,
Anant Agarwal

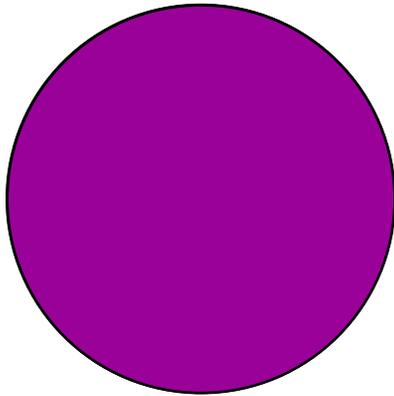


Processor Model

- Stable model for last few decades
 - Von Neumann architecture
 - Sequentially execute instructions
 - Simple abstraction
 - Easy to program

Change Is Around the Corner

- Processor performance not scaling as before
 - Wire delay and power



old view: chip looks small to a wire



chip size



distance signal can travel
in 1 cycle

new view: chip looks much bigger to a wire,
communication is expensive even on chip!

- How to effectively use transistors?

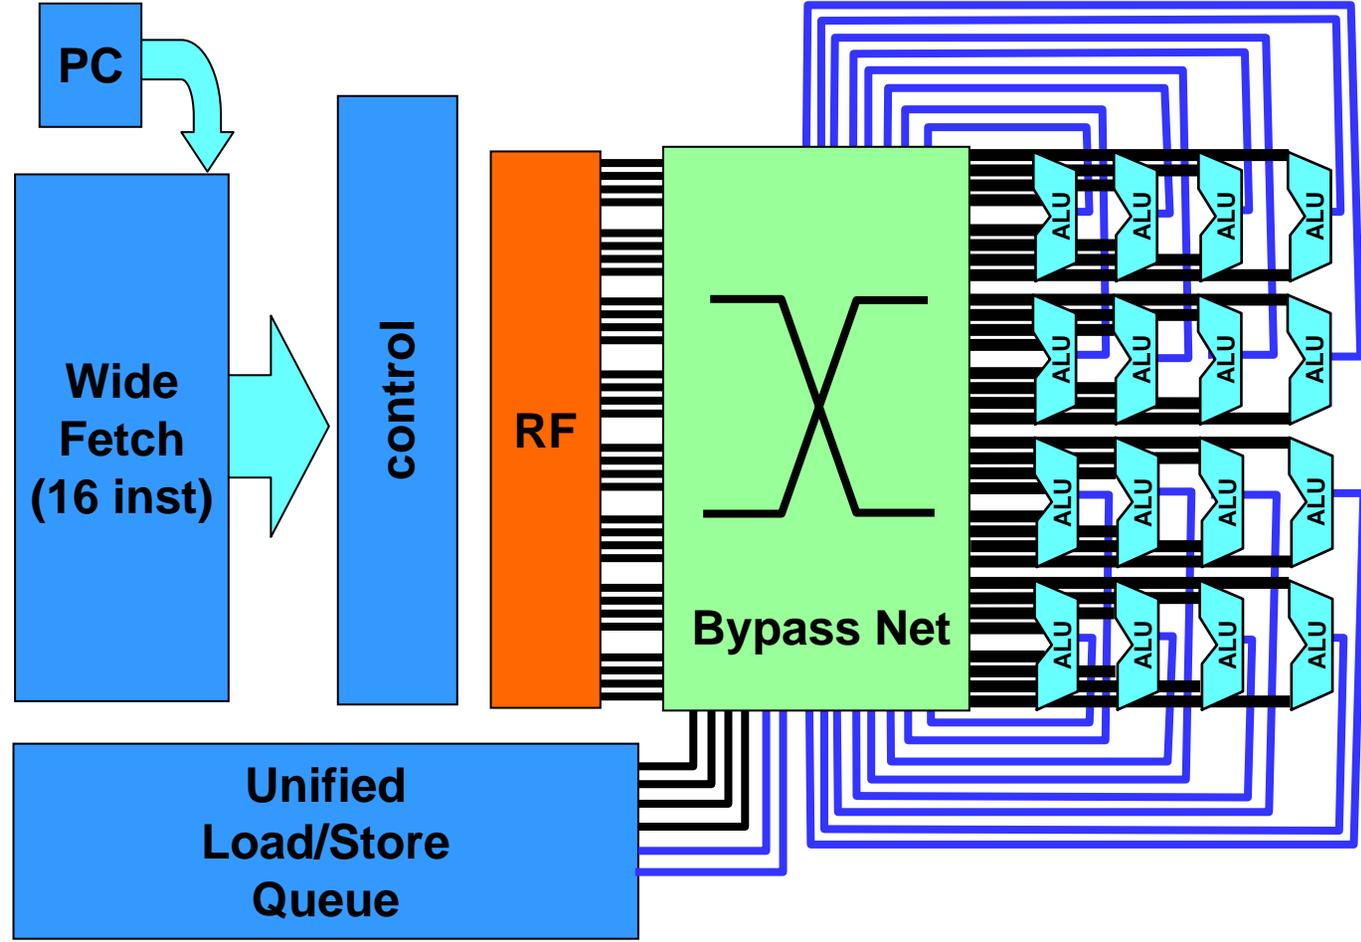
Spatially-Aware Architectures

- Many forward looking architectures are addressing the physical challenges
 - MIT Raw processor
 - MIT Scale processor
 - Stanford Imagine processor
 - Stanford Smart Memories processor
 - UC David Synchronscalar
 - UT Austin TRIPS processor
 - Wisconsin ILDP architecture
 - The original IBM BlueGene processor

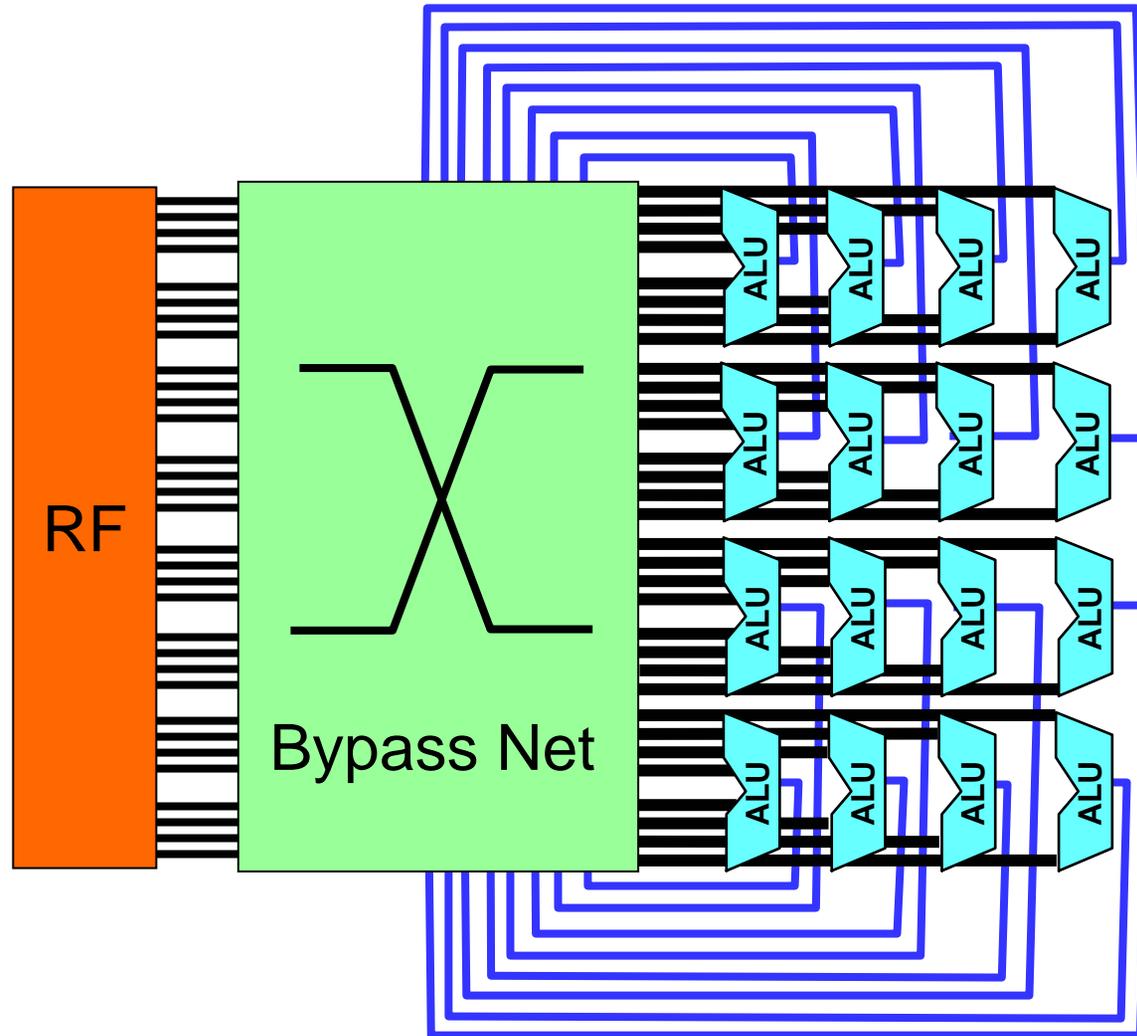
Problems with Monolithic Designs

- Super-wide general purpose processors are no longer practical

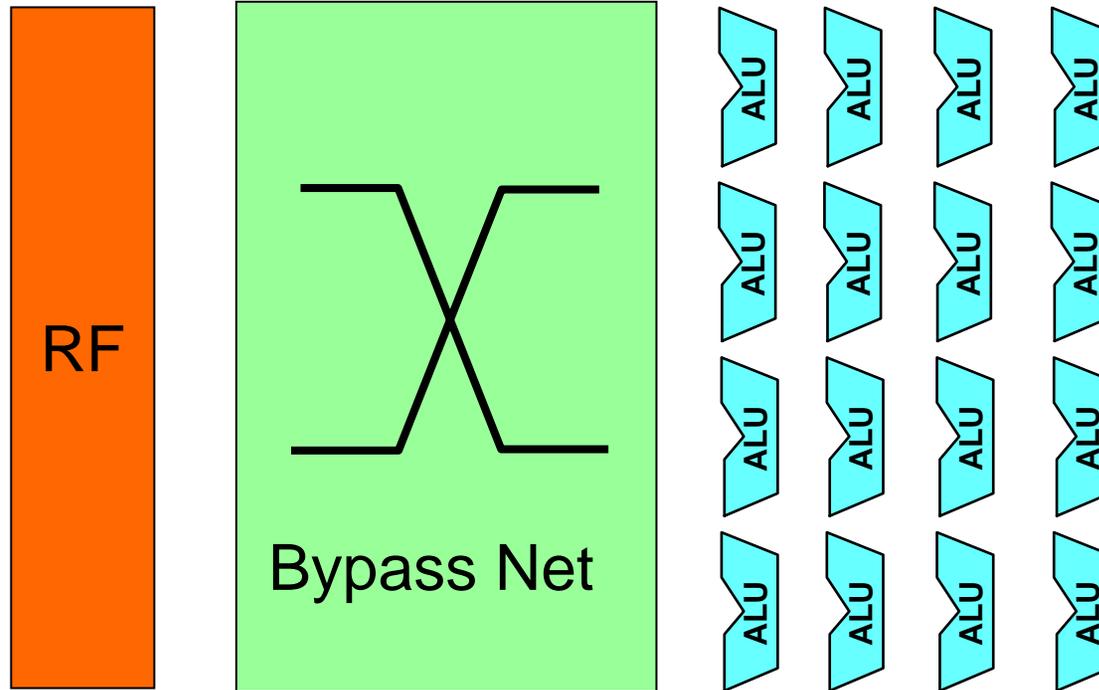
- Centralized control with global operand routing
- Area, power, and frequency concerns



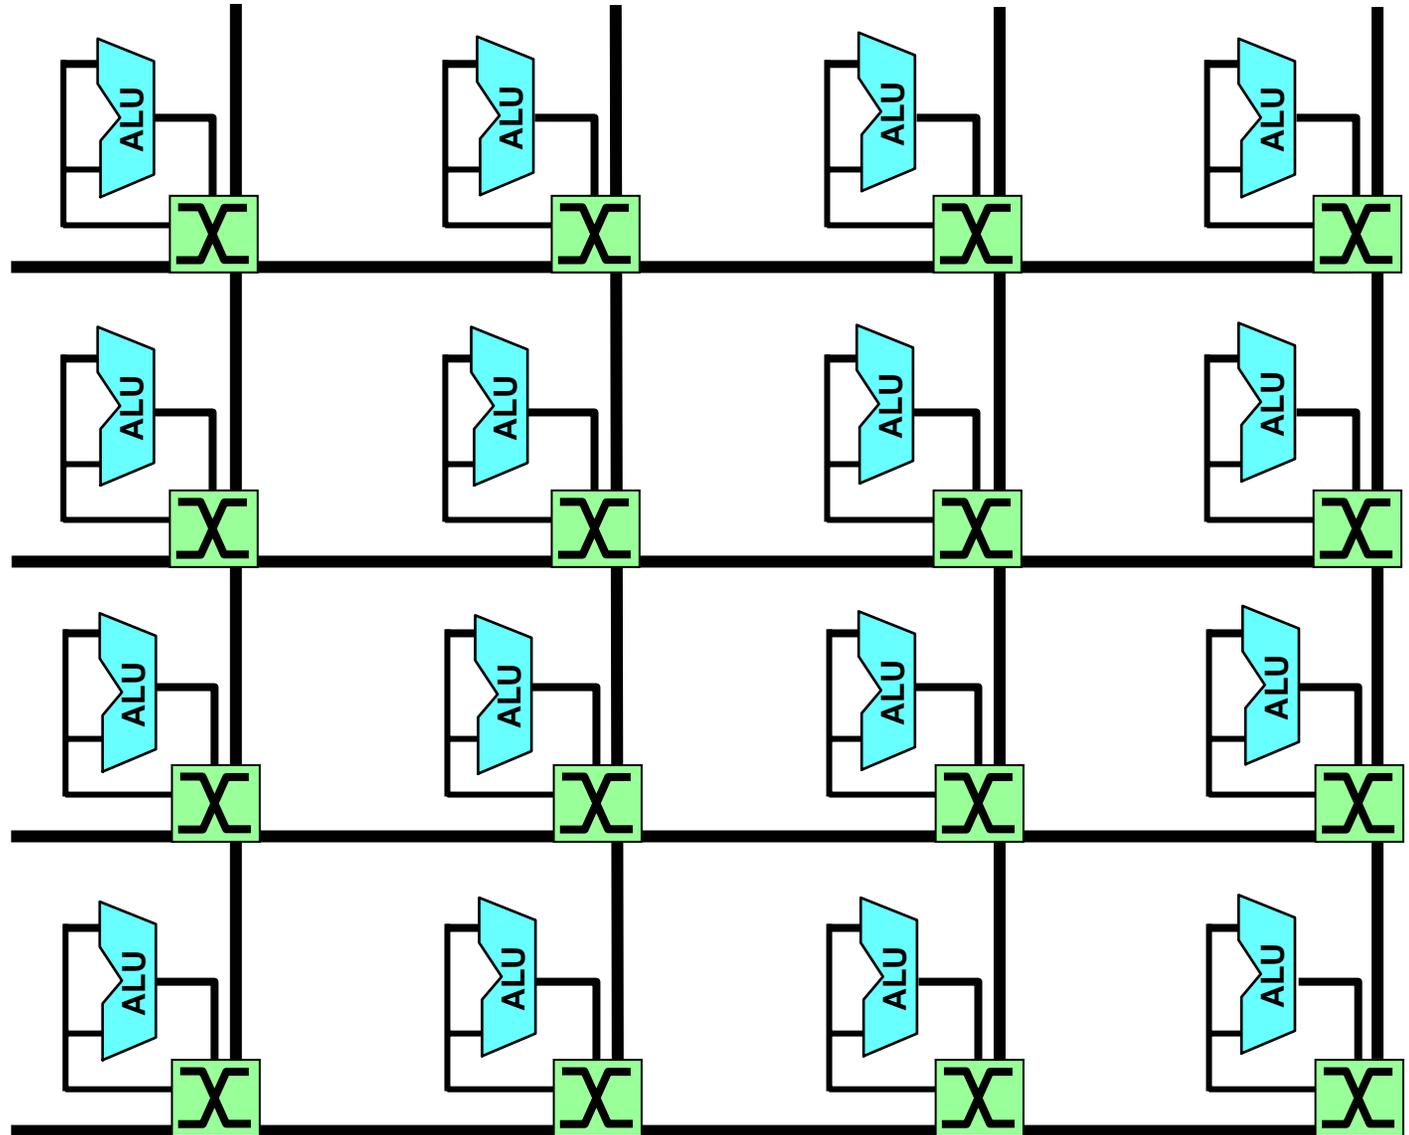
Spatial Architectures



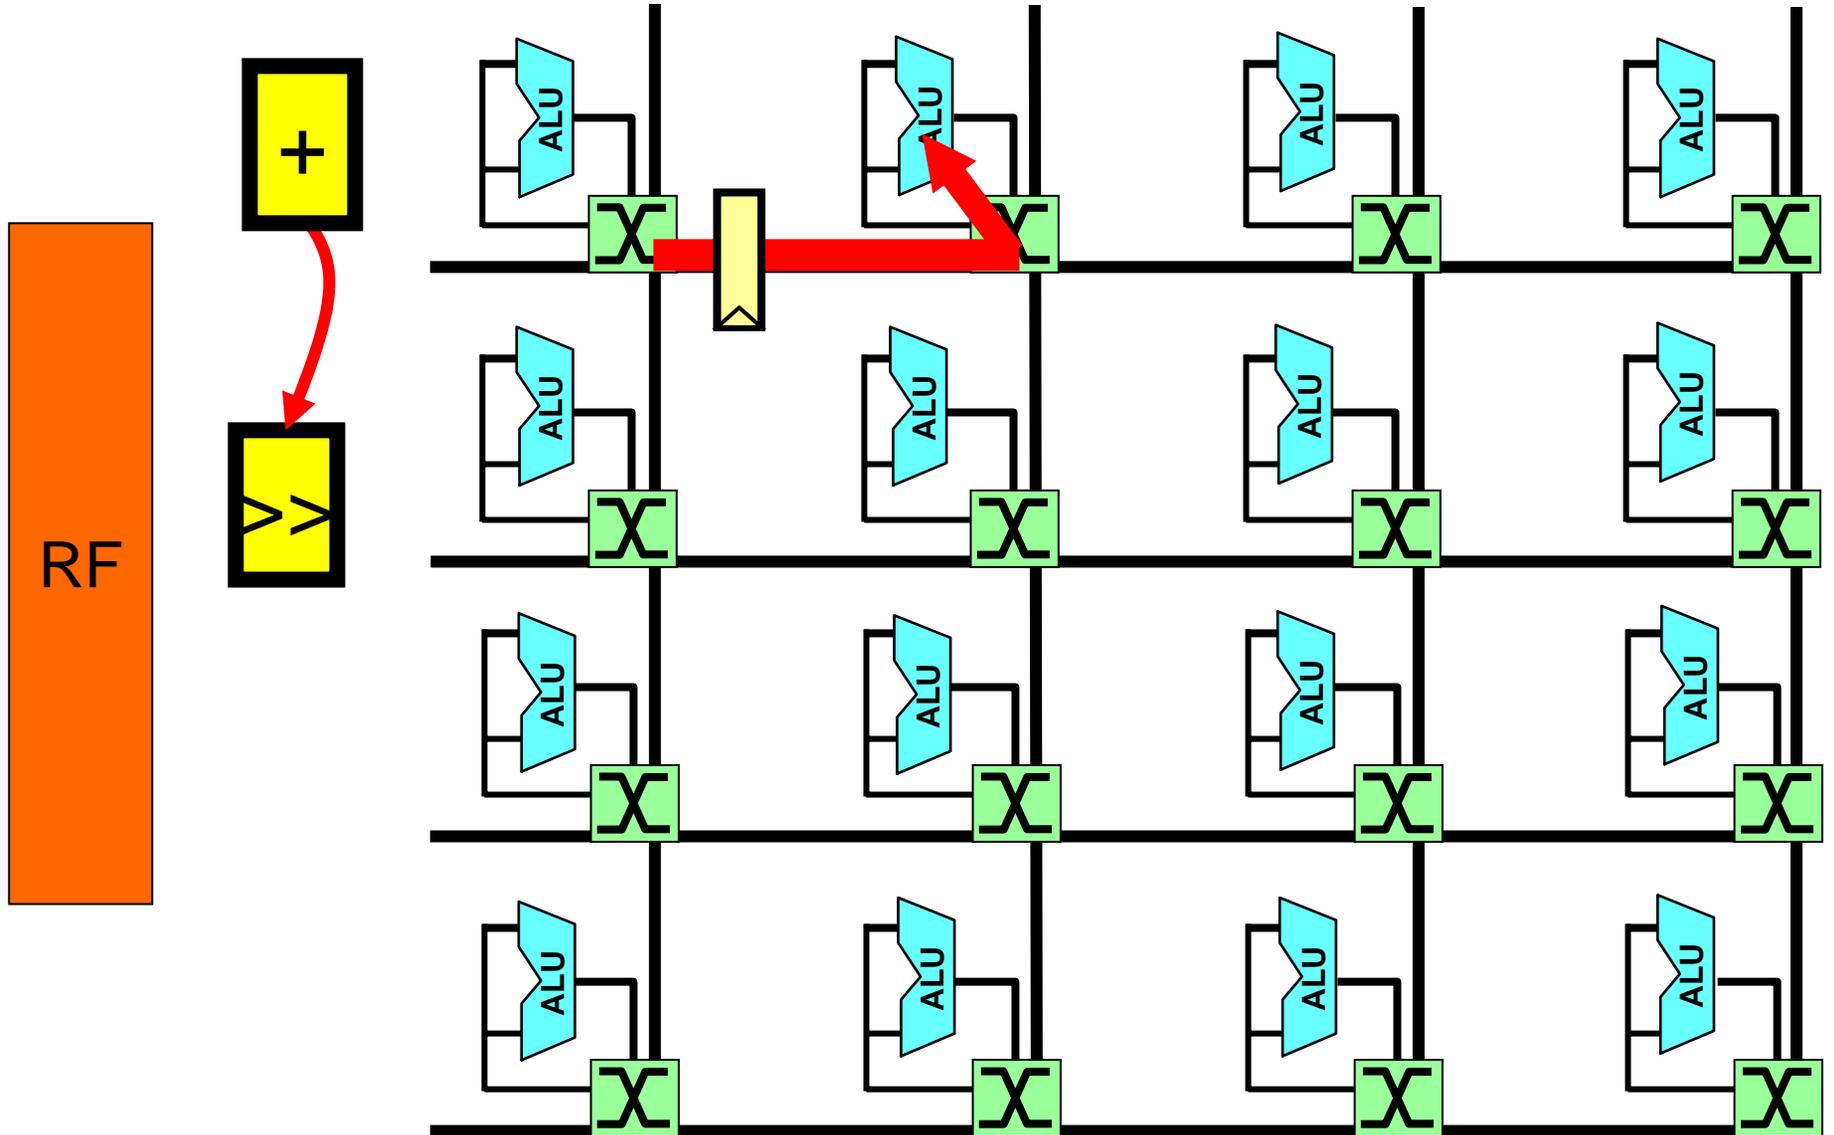
Spatial Architectures



Spatial Architectures



Exploiting Locality



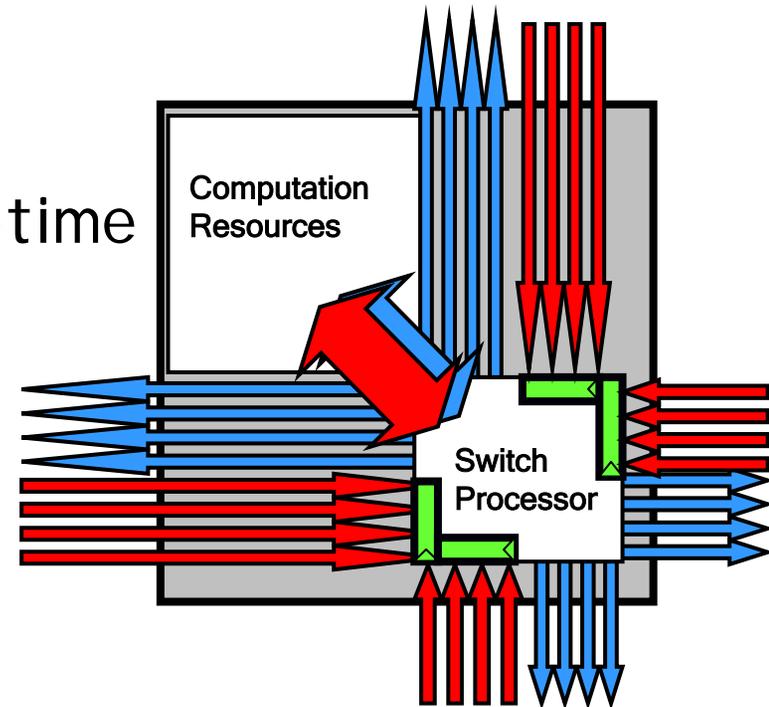
Raw On-Chip Networks

- 2 Static Networks

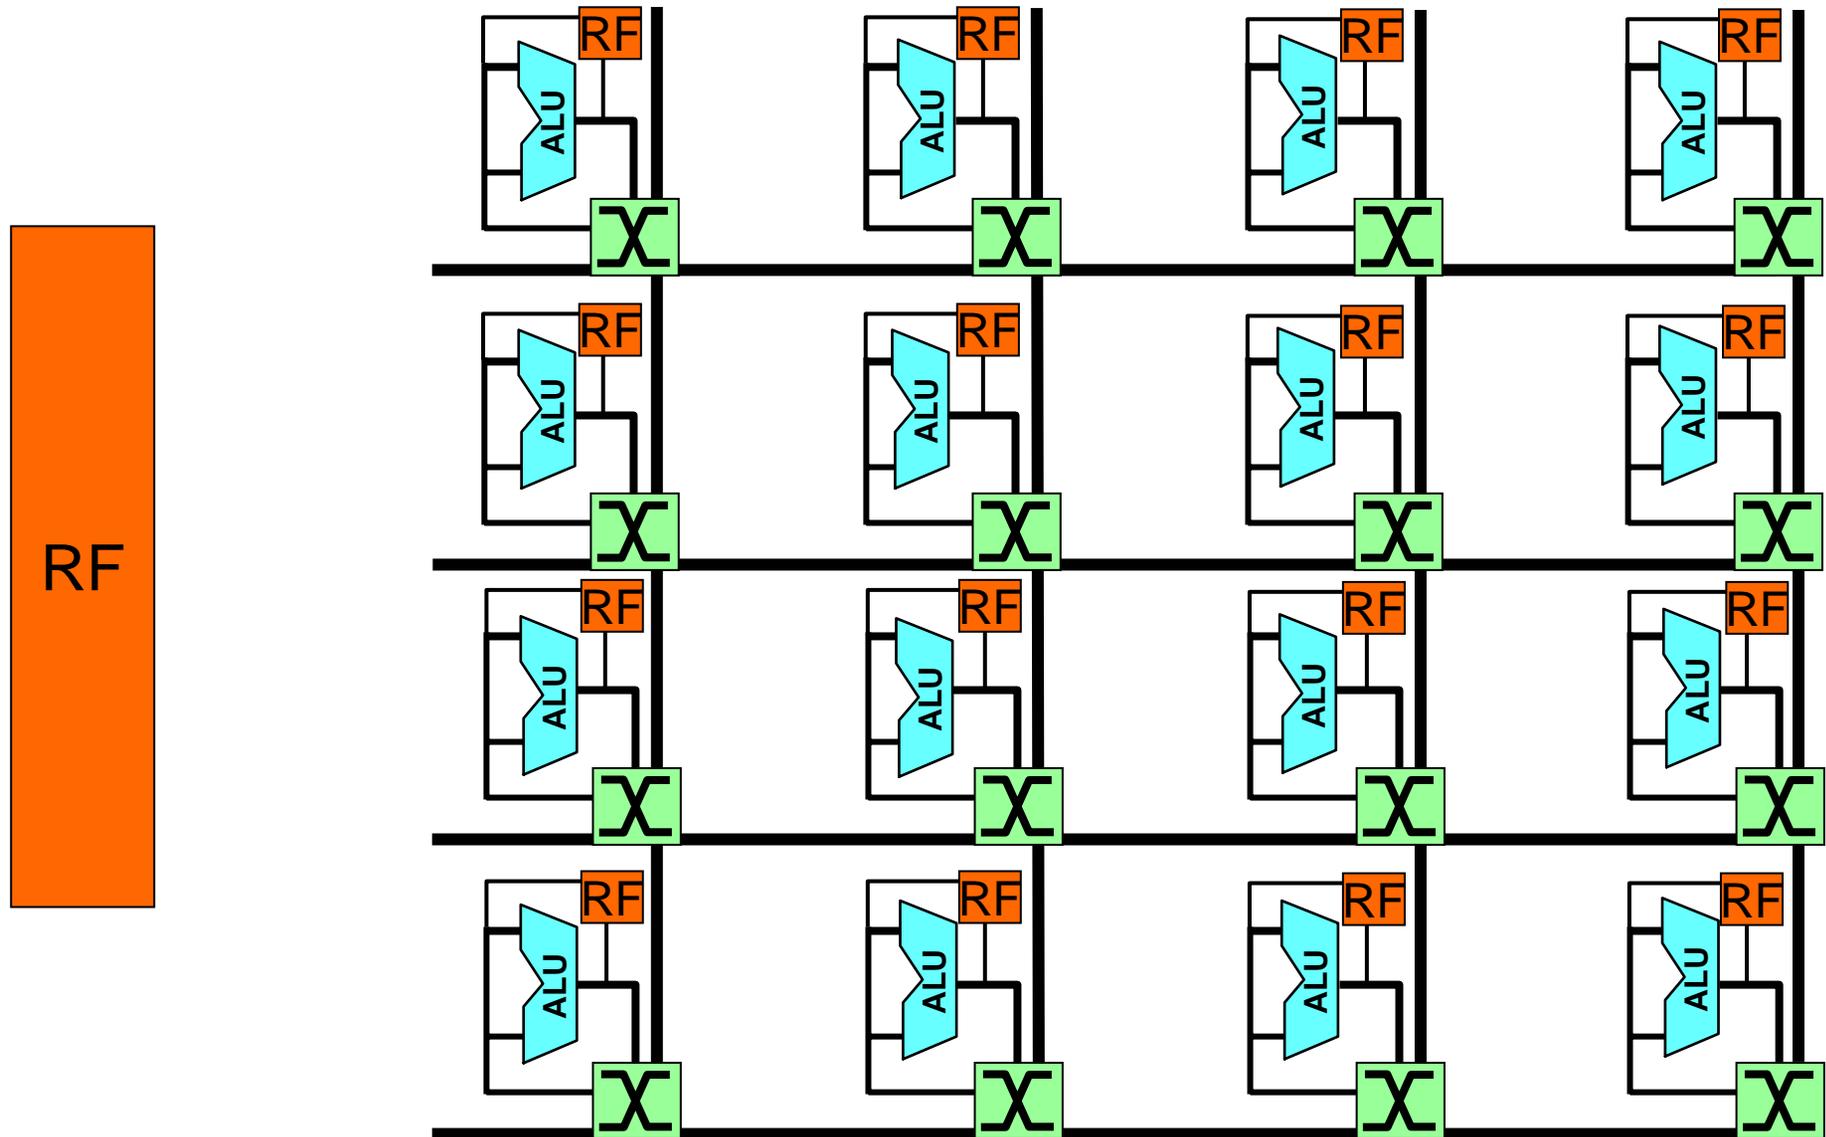
- Software configurable crossbar
- 3 cycle latency for nearest-neighbor ALU to ALU
- Must know pattern at compile-time
- Flow controlled

- 2 Dynamic Networks

- Header encodes destination
- Fire and Forget
- 15 cycle latency for nearest-neighbor



Distribute the Register File



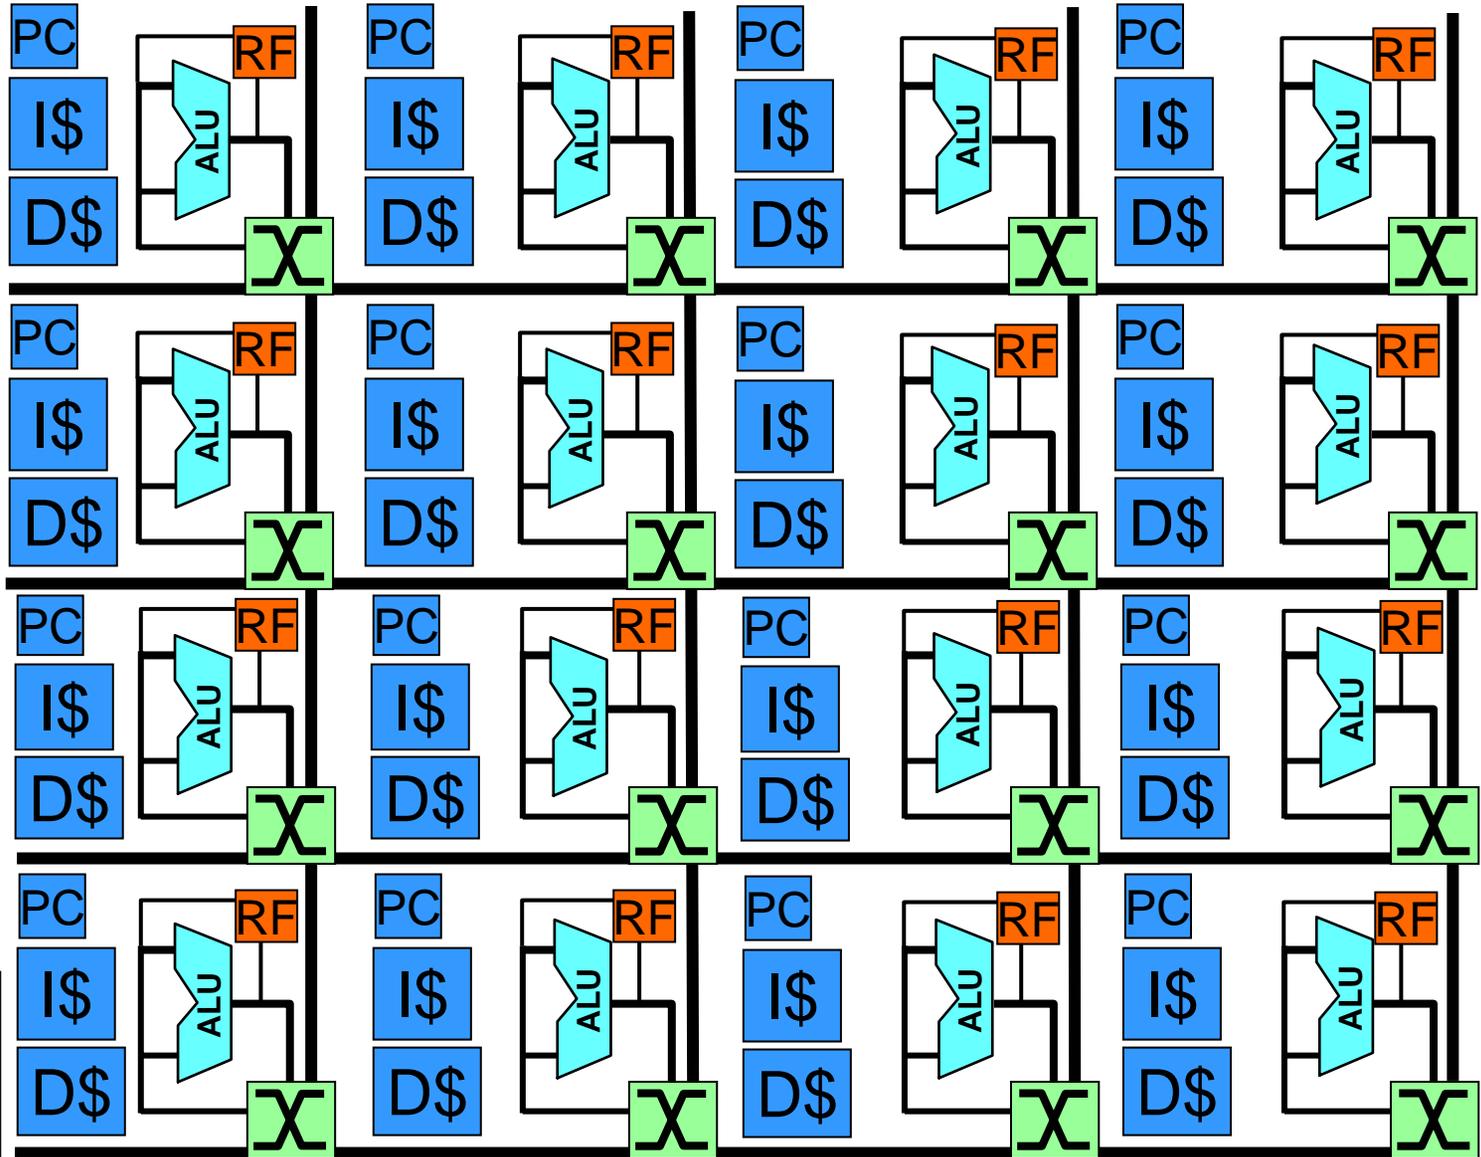
Distribute the Rest

PC

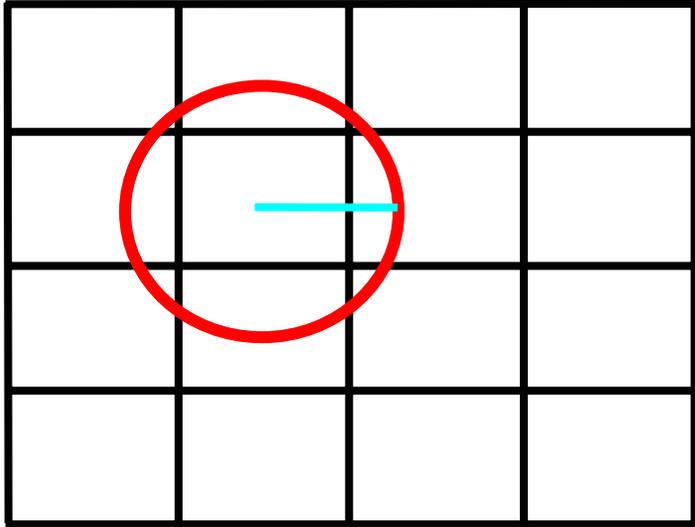
Wide Fetch (16 inst)

Control

Unified Load/Store Queue



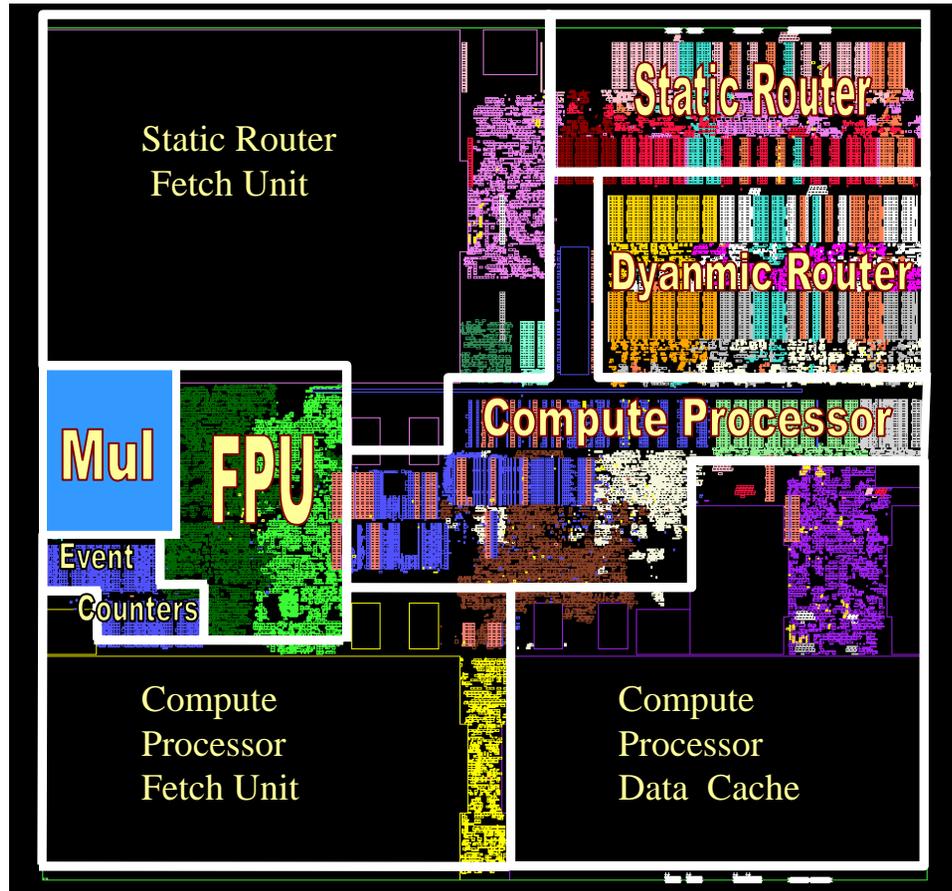
Tiled-Processor Architecture



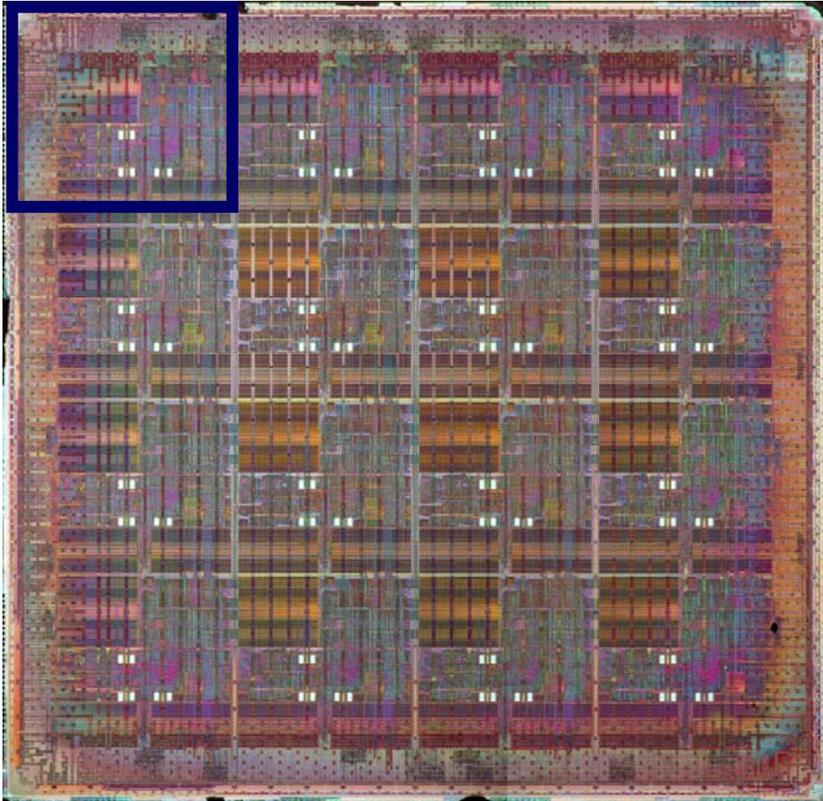
Make a tile as big as you can go in one clock cycle, and expose longer communication to the programmer

- Tile abstraction is quite powerful
 - e.g., power → resources used as necessary
- Easily scalable
 - All signals registered at tile boundaries, no global signals
 - Easier to Tune the Frequency
 - Easier to do the Physical Design
 - Easier to Verify

Close-up of a Single Raw Tile



The MIT Raw Processor



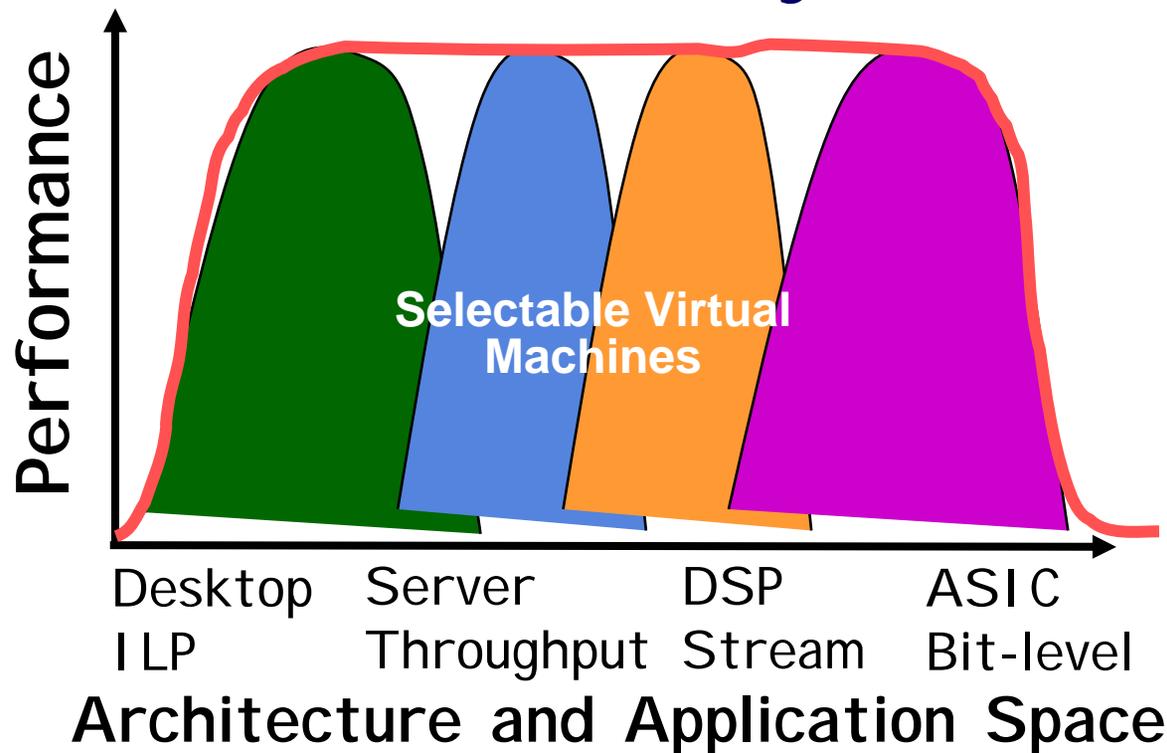
- 180 nm ASIC (IBM SA-27E)
- 16 tiles → 16 issue
- Core Frequency:
 - 425 MHz @ 1.8 V
 - 500 MHz @ 2.2 V
- Frequency competitive with IBM-implemented PowerPCs in same process
- 18 W (vpenta)

The Raw Goal

- Create an architecture that
 - Scales to 100's-1000's of functional units, memory ports
 - By exploiting custom-chip like features
 - Application-specific routing of operands
 - Is "general purpose" (*Versatile*)
 - Run ILP sequential programs, scientific computations, server-style processing, streaming systems, and bit-level applications
 - Support standard General Purpose Abstractions
 - Context switching, caching and instruction virtualization

The New Performance Goal

Versatility

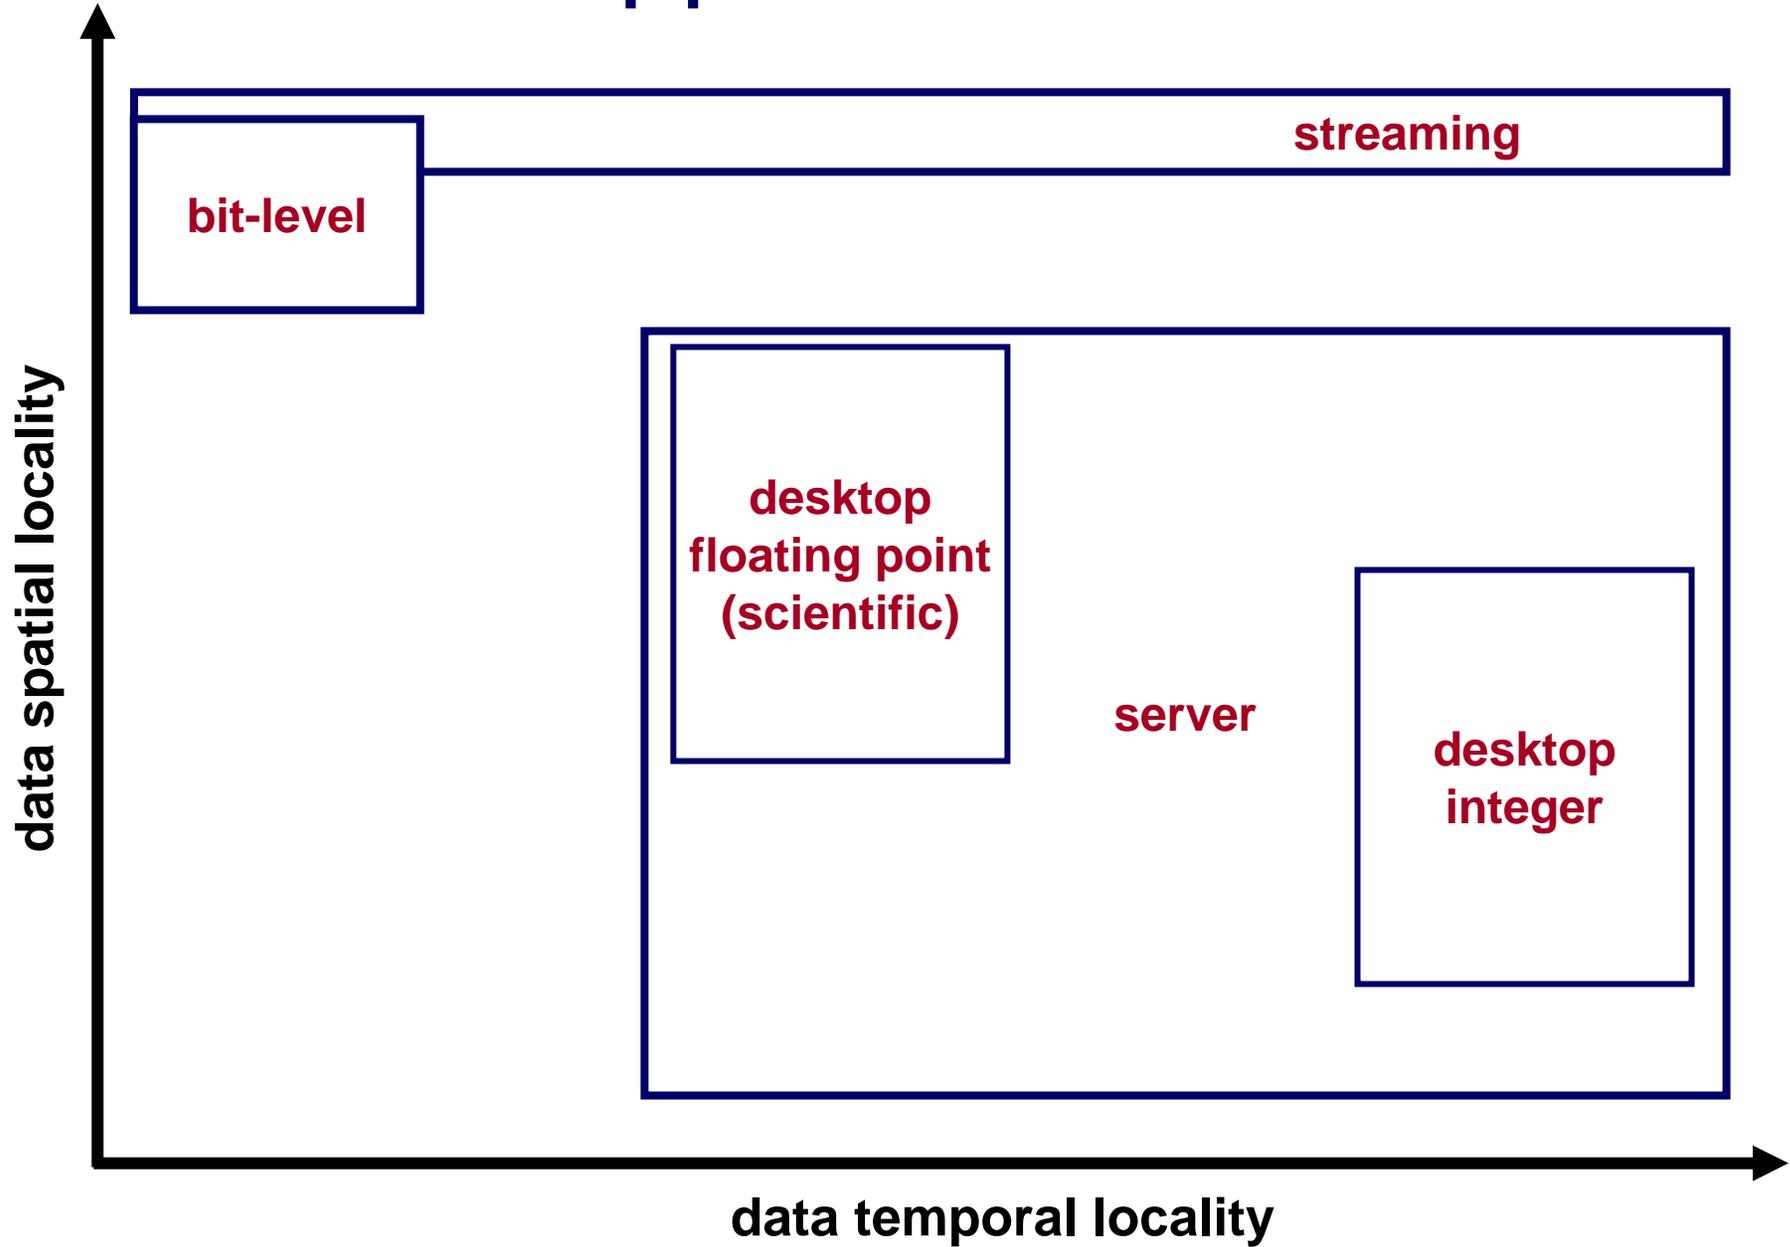


- Raw architecture as an “all-purpose” processor
 - Better SPECmark/Watt across the board
 - Higher SPECmark → think more MIPS compared to some reference machine (e.g., VAX 11/780)

Application Domains

- 5 market-dominant application domains
 - Desktop Integer
 - Desktop Floating (Scientific codes)
 - Server (Throughput Based)
 - Ergonomic simulations, Grid computation, Transaction processing
 - Embedded Streaming
 - Embedded Bit-Level

How Applications Differ



Distinguishing Application Domains

- Five **basis** properties
 - Data temporal locality
 - Quantify address reuse
 - Spatial temporal locality
 - Quantify address adjacency
 - Predominant data type
 - Parallelism
 - ILP, DLP, TLP, etc
 - Instruction temporal locality
 - Inverse of control complexity

Classifying Applications

- Quantitative metrics for the basis properties
 - Measure properties of different applications
- Cluster applications into domains
 - VersaBench

	Data Type	Parallelism	Instruction Temporal Locality	Data Temporal Locality	Data Spatial Locality
Desktop INT	integer	low	low	high	low
Desktop FLT	float	medium	medium	medium	medium
Server	integer/float	high	low to medium	medium to high	low to medium
Streaming	integer/float	very high	high	low to high	very high
Bit-Level	bit	medium-high	very high	low	very high

VersaBench Status

- 15 total benchmarks
 - 3 per category
 - Drawn from SPEC INT/FP, Raw, StreamIt, DIS (AAEC), USC ISI
 - Manageable size, encourages evaluation using the entire suite
 - Available online at <http://cag.csail.mit.edu/versabench>
- Benchmarks selected systematically
 - MIT Technical Memo 646, June 2004
Rabbah, Bratt, Asanovic, Agarwal

Proposed Metric: Versatility

Versatility (VersaBench)

Geometric Mean of Speedup relative to **best** performing machines

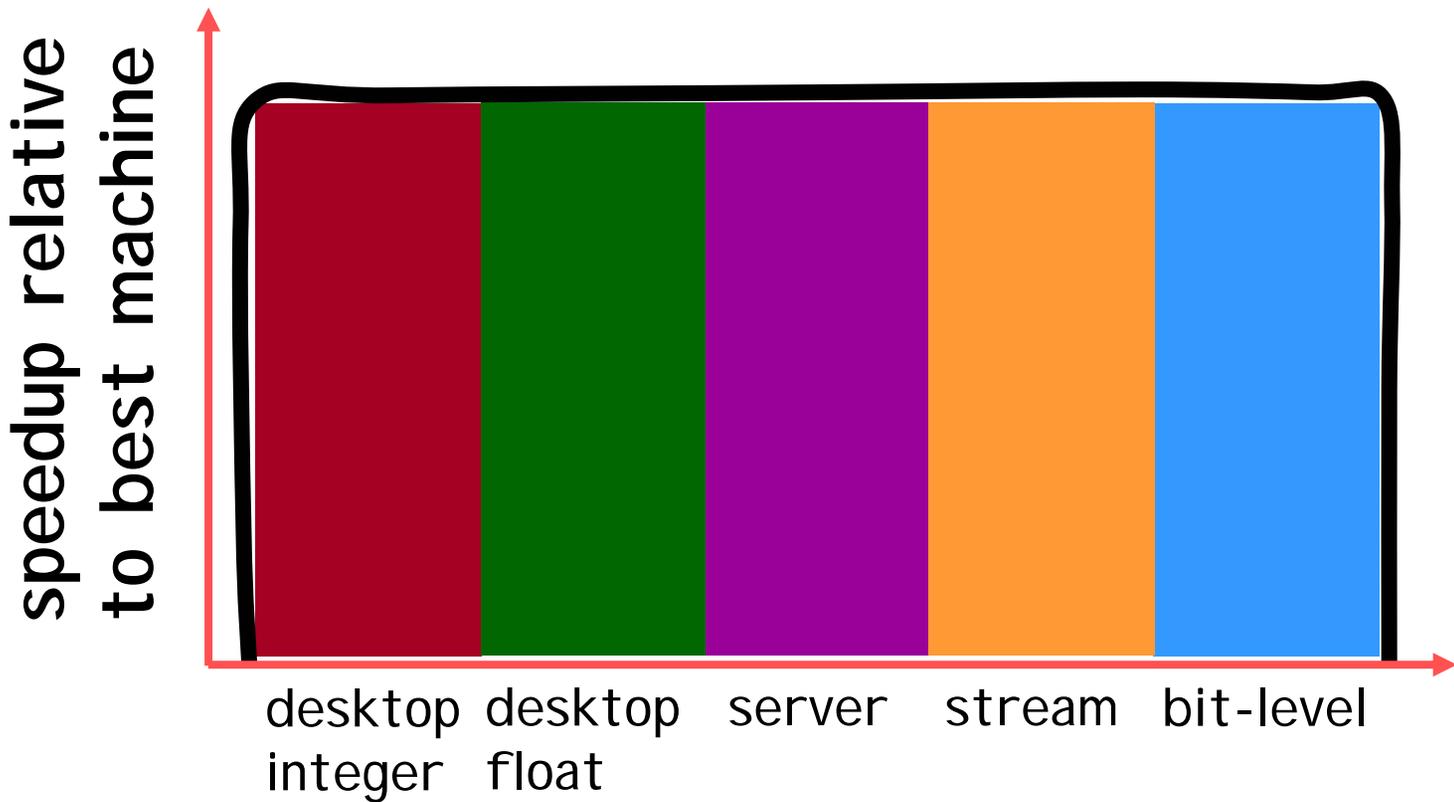
SPECmark (SPEC)

Geometric Mean of Speedup relative to a **single** reference machine

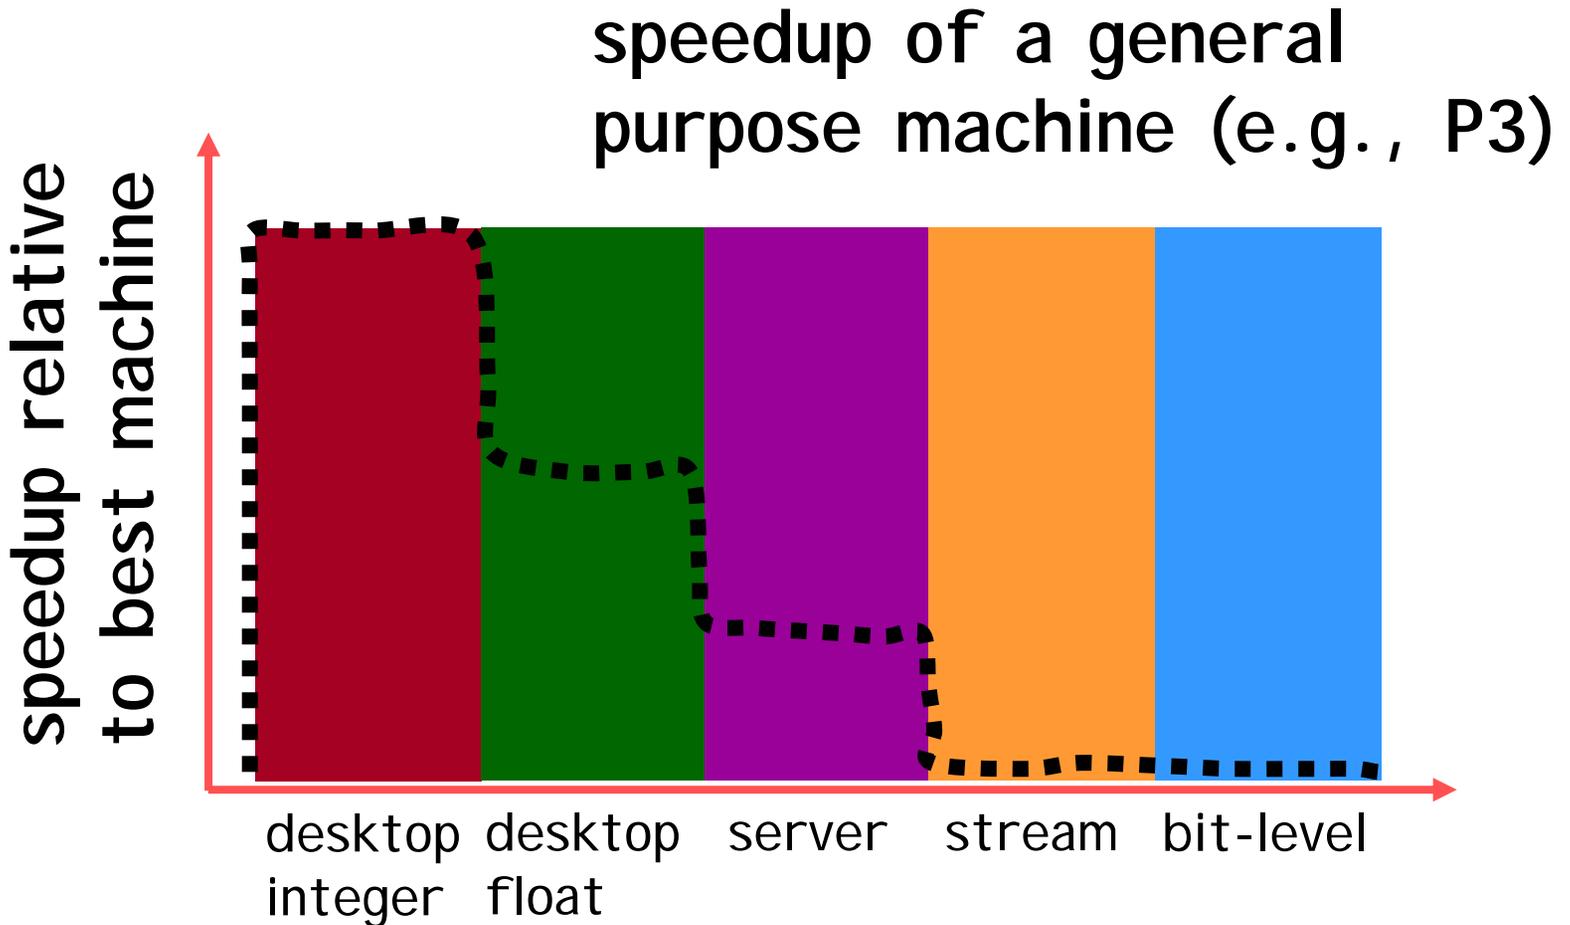
- **Normalization to the best performing machines identifies areas for improvements**
 - This is especially important → **VersaGraphs**
 - Not another mean over N benchmarks
- High Versatility mark implies architecture is good across the board

VersaGraph Example

speedup of an ideal machine

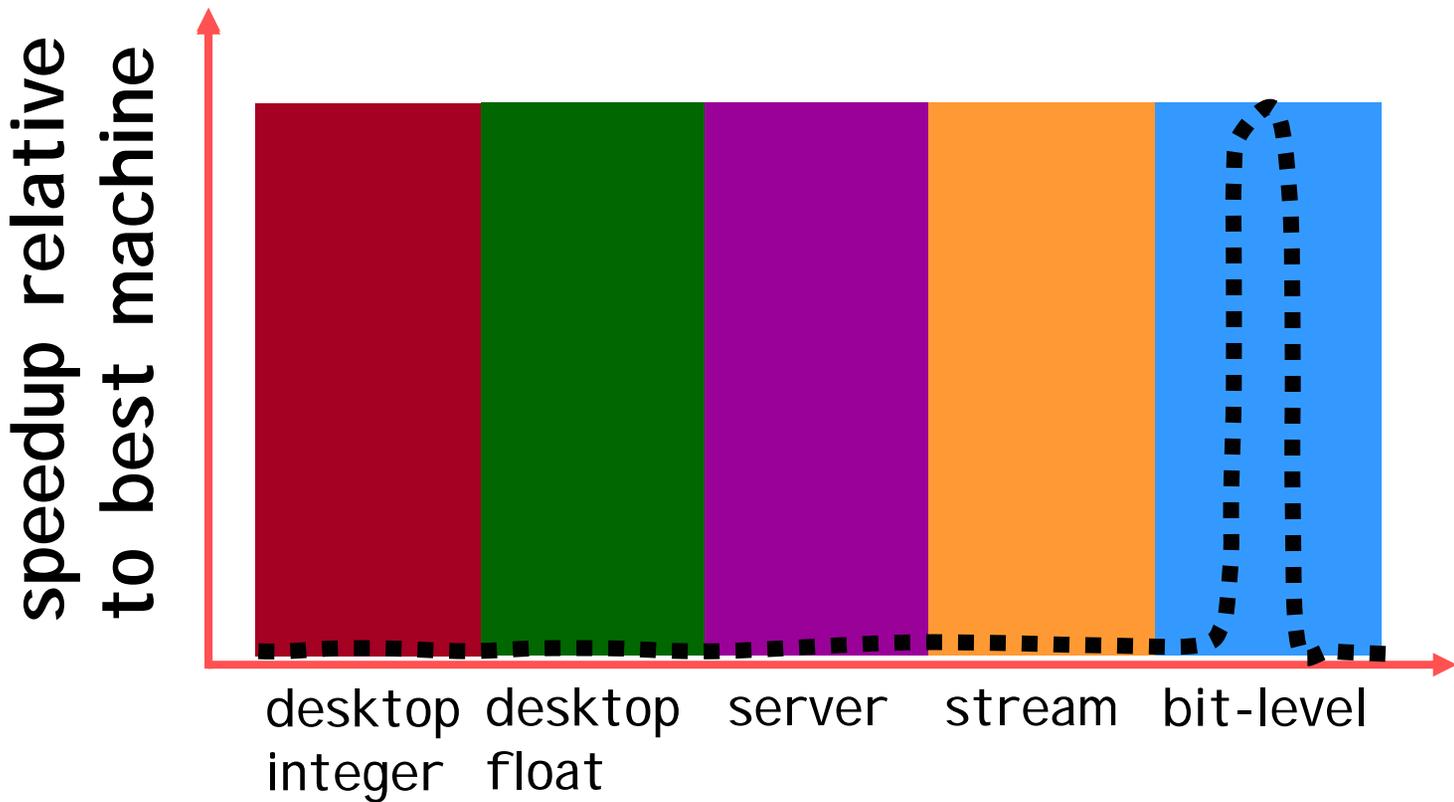


VersaGraph Example



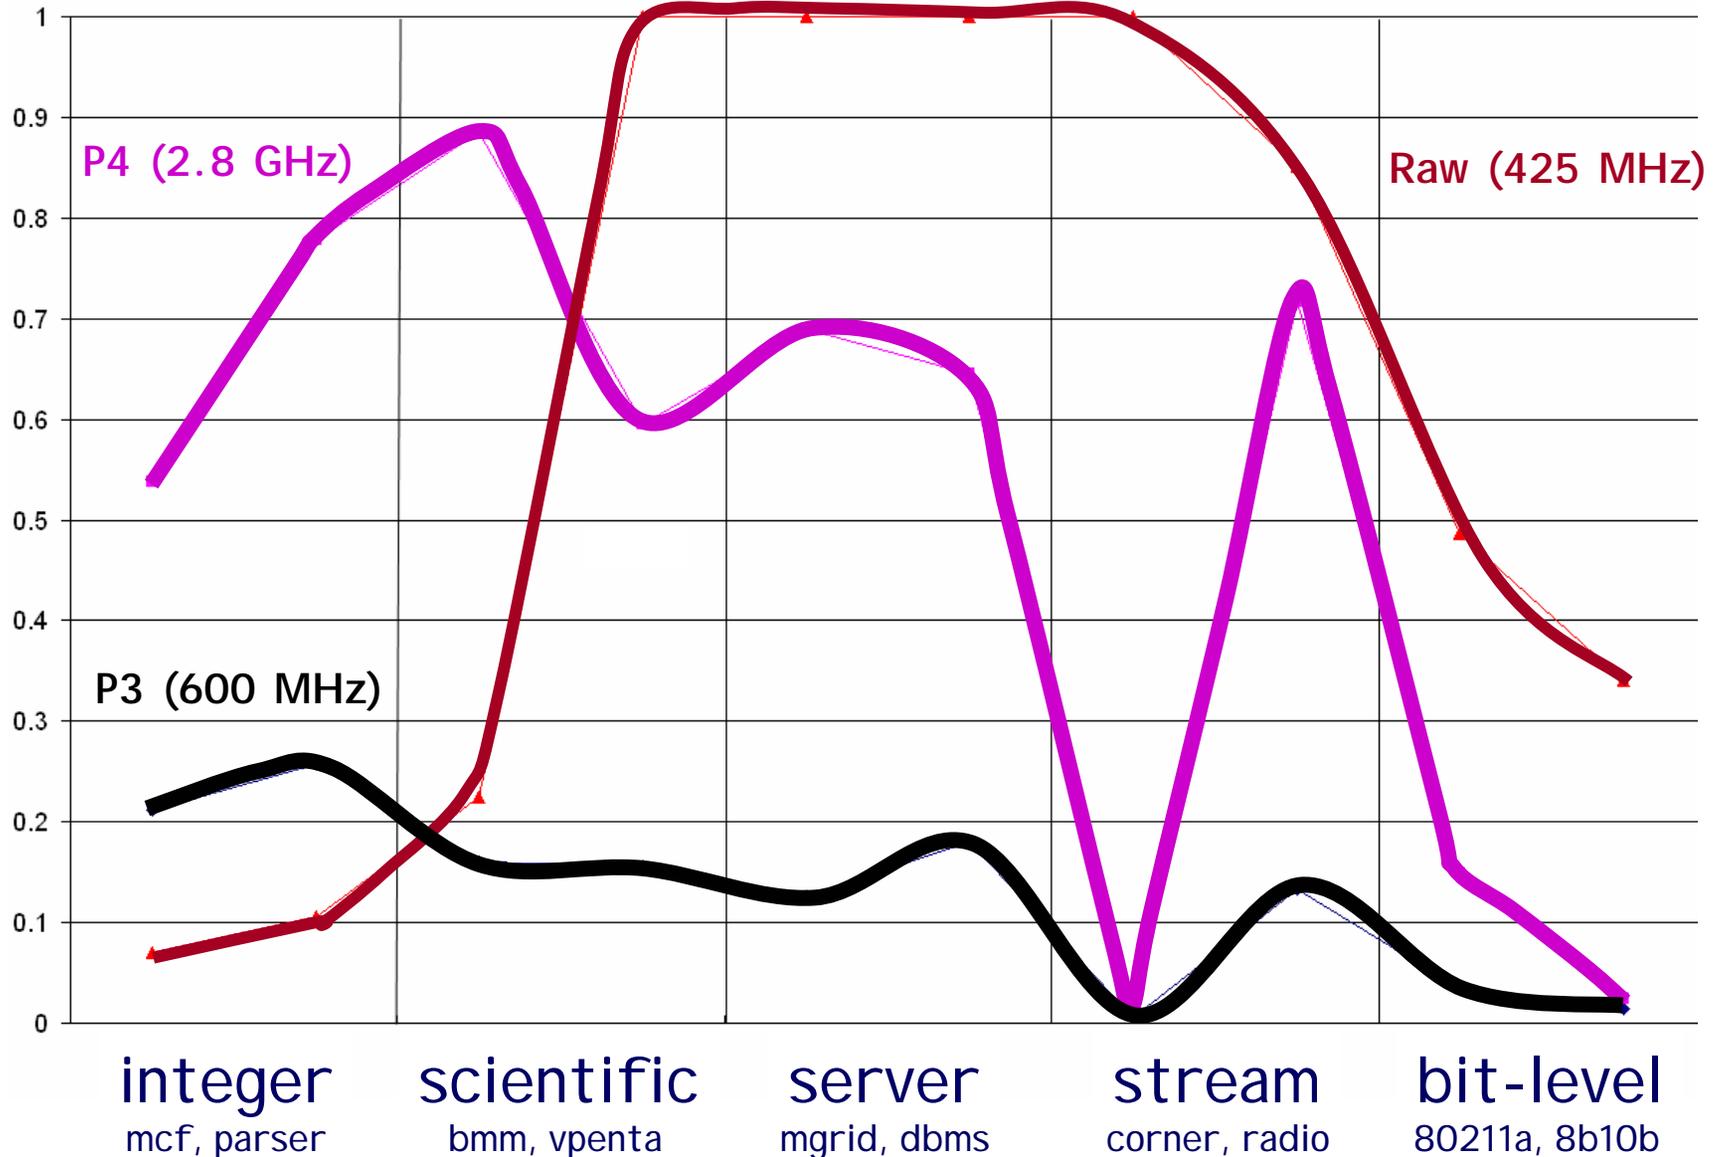
VersaGraph Example

speedup of an ASIC



VersaGraphs For Real Architectures

Also compared against Athlon 64 and Itanium 2



Raw Homepage

<http://cag.csail.mit.edu/raw>

download papers, benchmarks, ...