Vector IRAM

A Media-oriented Vector Processor with Embedded DRAM

Christoforos Kozyrakis, Joseph Gebis, David Martin, Samuel Williams, Ioannis Mavroidis, Steven Pope, Darren Jones*, David Patterson, and Katherine Yelick

Computer Science Division
University of California at Berkeley

IBM Microelectronics
Hopewell Junction, NY

* MIPS Technologies Inc.
Mountain View, CA

Vector IRAM Overview

- Microprocessor prototype with
  - A 256-bit vector coprocessor
  - A 16-MByte embedded DRAM memory system
- 150 million transistors, 290 mm^2
- 3.2 Gops at 200 MHz
- 1.6 Gops/Watt at 2 W
- Industrial strength vectorizing compiler
- Implemented by 6 graduate students for architecture, design, simulation and testing

Motivation and Goals

- Processor features for PostPC systems:
  - High performance on demand for multimedia without continuous high power consumption
  - Tolerance to memory latency
  - Scalable
  - Mature, HLL-based software model
- Design a prototype processor chip
  - Complete proof of concept
  - Explore detailed architecture and design issues
  - Motivation for software development

Key Technologies

- Vector processing
  - High performance on demand for media processing
  - Low power for issue and control logic
  - Low design complexity
  - Well understood compiler technology
- Embedded DRAM
  - High bandwidth for vector processing
  - Low power/energy for memory accesses
  - “System on a chip”

Vector Instruction Set

- Complete load-store vector instruction set
  - Uses the MIPS64™ ISA coprocessor 2 opcode space
  - Data types supported: 64b, 32b, 16b (and 8b)
  - Architecture state
    - 32 general-purpose vector registers
    - 32 vector flag registers, 16 scalar registers
  - 91 instructions
    - Arithmetic, logical, vector processing
    - Sequential, strided and indexed loads and stores
- Not specified by the ISA
  - Maximum vector register length
  - Functional unit datapath width

Vector ISA Enhancements

- DSP support
  - Fixed-point shift, multiply and multiply-add, saturated arithmetic, rounding modes
  - Simple instructions for intra-register permutations for reductions and butterfly operations
    - High performance for dot-products and FFT without the complexity of a random permutation
- Compiler and OS support
  - Conditional execution of vector operations
  - Support for software speculation of load operations
  - MMU-based virtual memory
  - Restartable arithmetic exceptions
  - Valid and dirty bits for vector registers

VIRAM Prototype Architecture

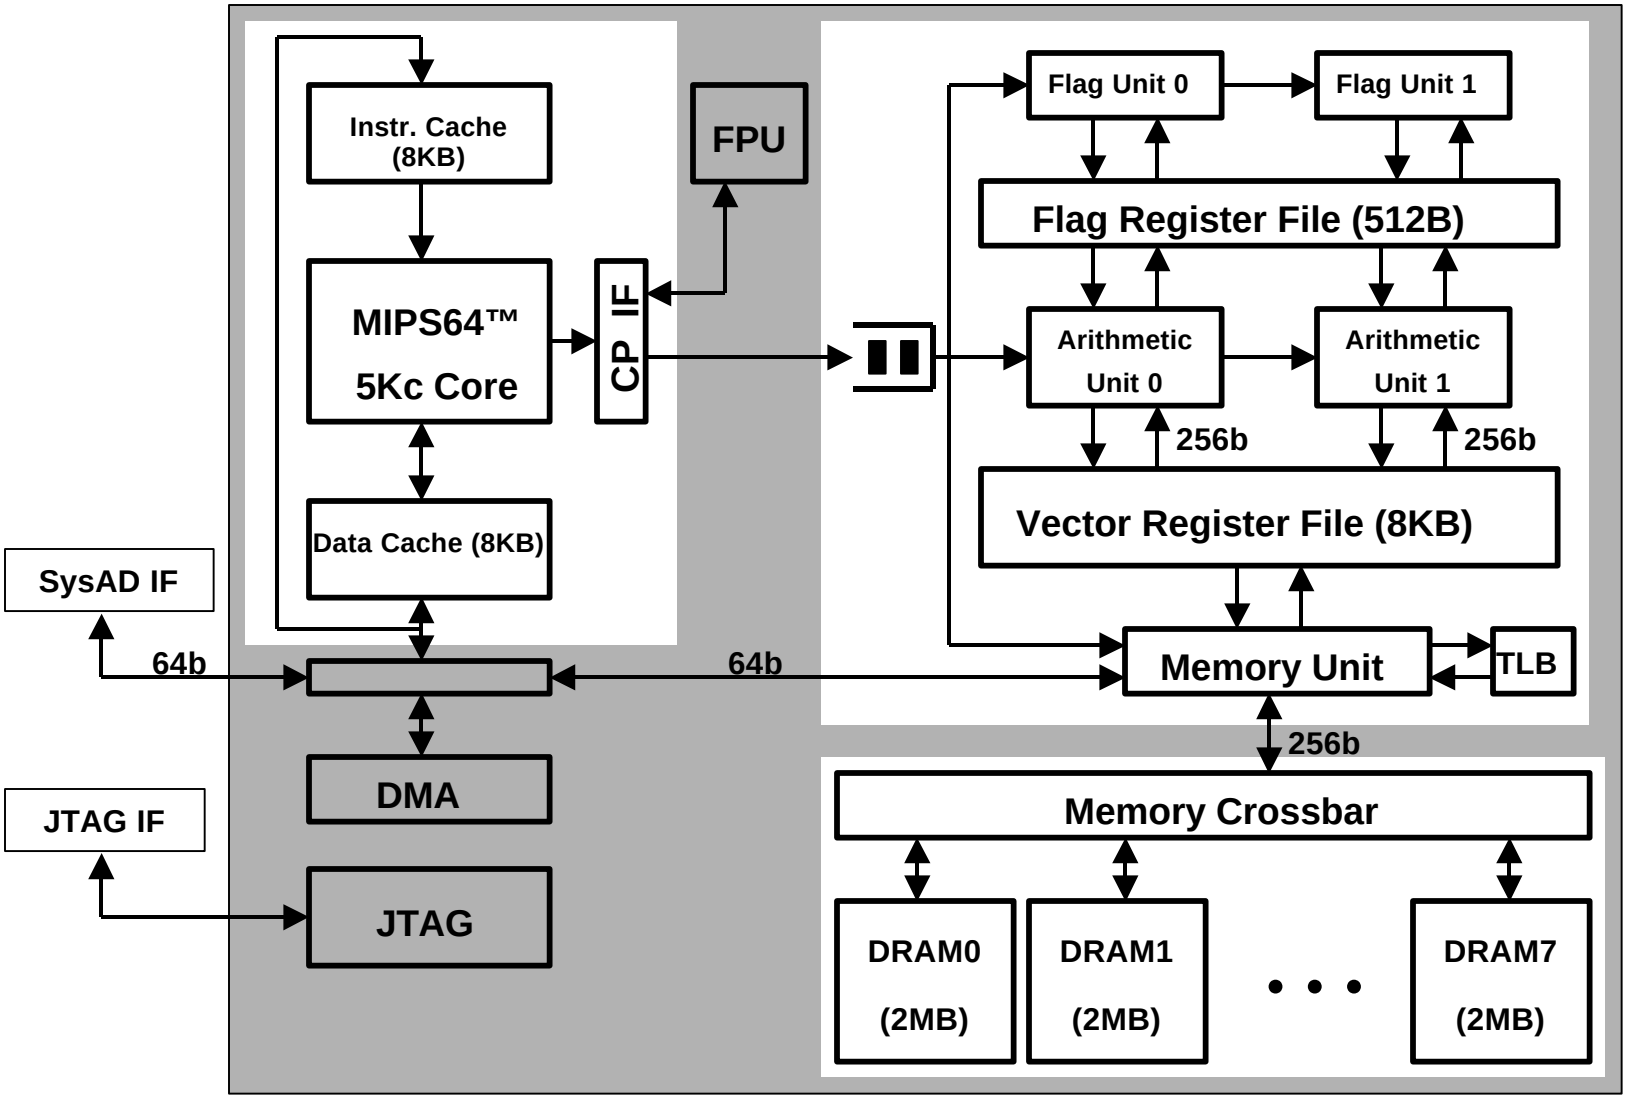

Vector Unit Pipeline

- Single-issue, in-order pipeline
- Efficient for short vectors
  - Pipelined instruction start-up
  - Full support for instruction chaining, the vector equivalent of result forwarding
- Hides long DRAM access latency
  - Random access latency could lead to stalls due to long load→use RAW hazards
  - Simple solution: “delayed” vector pipeline

Delayed Vector Pipeline

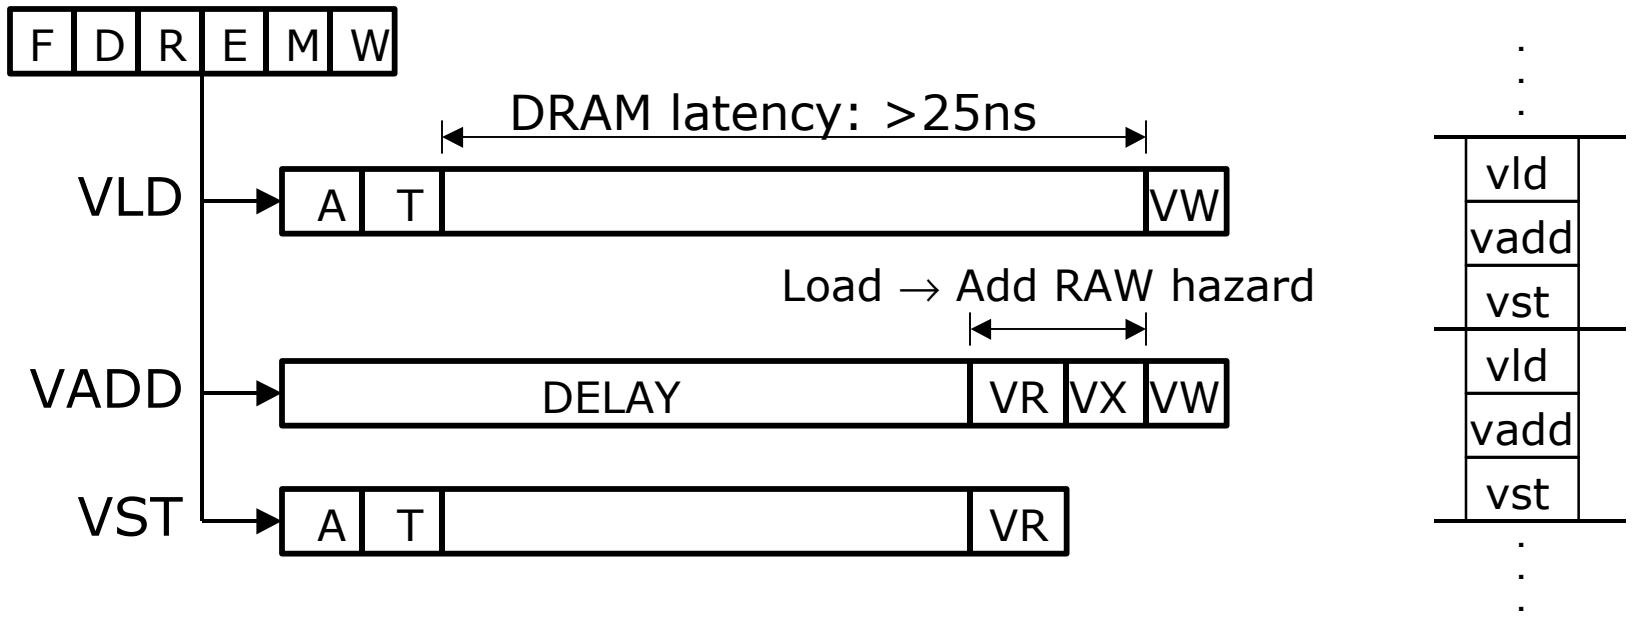

- Random access latency included in the vector unit pipeline
- Arithmetic operations and stores are delayed to shorten RAW hazards
- Long hazards eliminated for the common loop cases
- Vector pipeline length: 15 stages

Modular Vector Unit Design

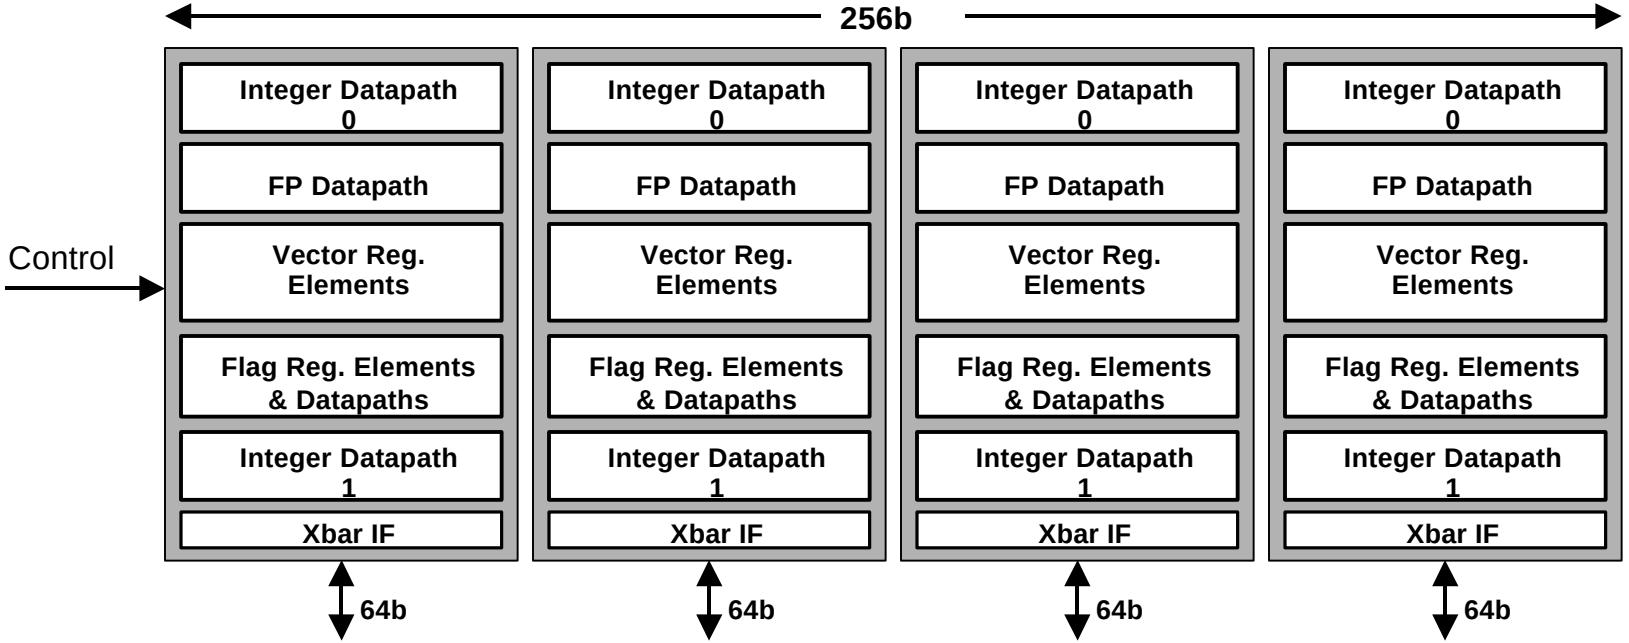

- Single 64b "lane" design replicated 4 times
  - Reduces design and testing time
  - Provides a simple scaling model (up or down) without major control or datapath redesign
- Most instructions require only intra-lane interconnect
  - Tolerance to interconnect delay scaling

Floorplan

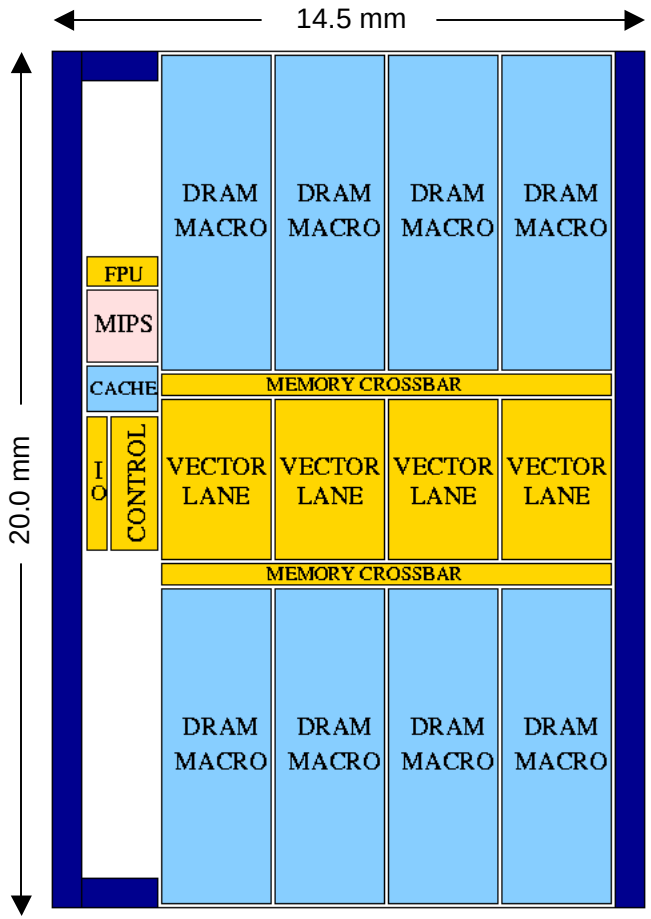

- Technology: IBM SA-27E
  - 0.18µm CMOS
  - 6 metal layers (copper)
- 290 mm^2 die area
  - 225 mm^2 for memory/logic
  - DRAM: 161 mm^2
  - Vector lanes: 51 mm^2
- Transistor count: ~150M
- Power supply
  - 1.2V for logic, 1.8V for DRAM
- Peak vector performance
  - 1.6/3.2/6.4 Gops wo. multiply-add (64b/32b/16b operations)
  - 3.2/6.4 /12.8 Gops w. multiply-add
  - 1.6 Gflops (single-precision)

Power Consumption

- Power saving techniques
  - Low power supply for logic (1.2 V)
    - Possible because of the low clock rate (200 MHz)
    - Wide vector datapaths provide high performance
  - Extensive clock gating and datapath disabling
    - Utilizing the explicit parallelism information of vector instructions and conditional execution
  - Simple, single-issue, in-order pipeline
- Typical power consumption: 2.0 W
  - MIPS core: 0.5 W
  - Vector unit: 1.0 W (min ~0 W)
  - DRAM: 0.2 W (min ~0 W)
  - Misc.: 0.3 W (min ~0 W)

Software Tools

- VIRAM compiler
  - Vectorizing compiler with C/C++/Fortran front-ends
  - Based on the Cray's PDGCS production environment for vector supercomputers
  - Extensive vectorization and optimization capabilities including outer loop vectorization
  - No need to use special libraries or variable types for vectorization
- Other software tools
  - ISA simulator and performance model
  - Assembler, disassembler, debugger

Performance: Efficiency

	Peak	Sustained	% of Peak
Image Composition	6.4 GOPS	**6.40 GOPS**	100%
iDCT	6.4 GOPS	**3.10 GOPS**	48.4%
Color Conversion	3.2 GOPS	**3.07 GOPS**	96.0%
Image Convolution	3.2 GOPS	**3.16 GOPS**	98.7%
Integer VM Multiply	3.2 GOPS	**3.00 GOPS**	93.7%
FP VM Multiply	1.6 GFLOPS	**1.59 GFLOPS**	99.6%
Average			89.4%

Performance: Comparison

	VIRAM	MMX	VIS
Image Composition	0.13	-	2.2 (17.0x)
iDCT	0.75	3.75 (5.0x)	-
Color Conversion	0.78	8.00 (10.2x)	-
Image Convolution	1.23	5.49 (4.5x)	6.19 (5.1x)
QCIF (176x144)	7.1M	33M (4.6x)	-
CIF (352x288)	28M	140M (5.0x)	-

- QCIF and CIF numbers are in clock cycles per frame
- All other numbers are in clock cycles per pixel
- MMX and VIS results assume no first level cache misses

Performance: FFT (1)

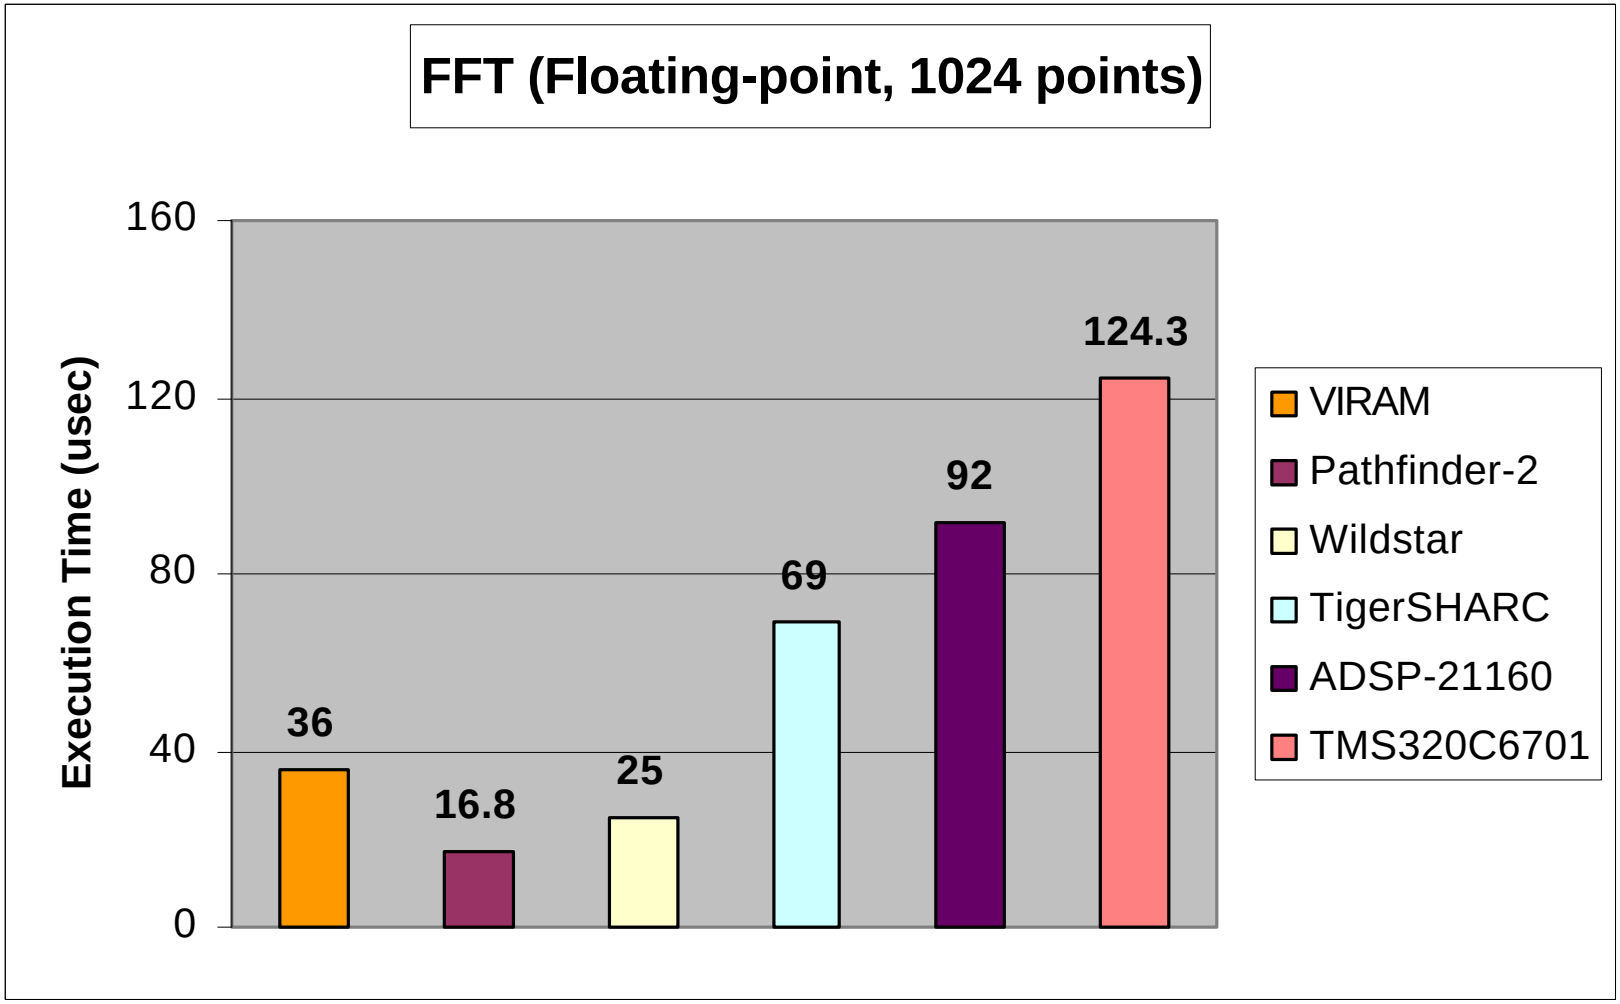

Performance: FFT (2)

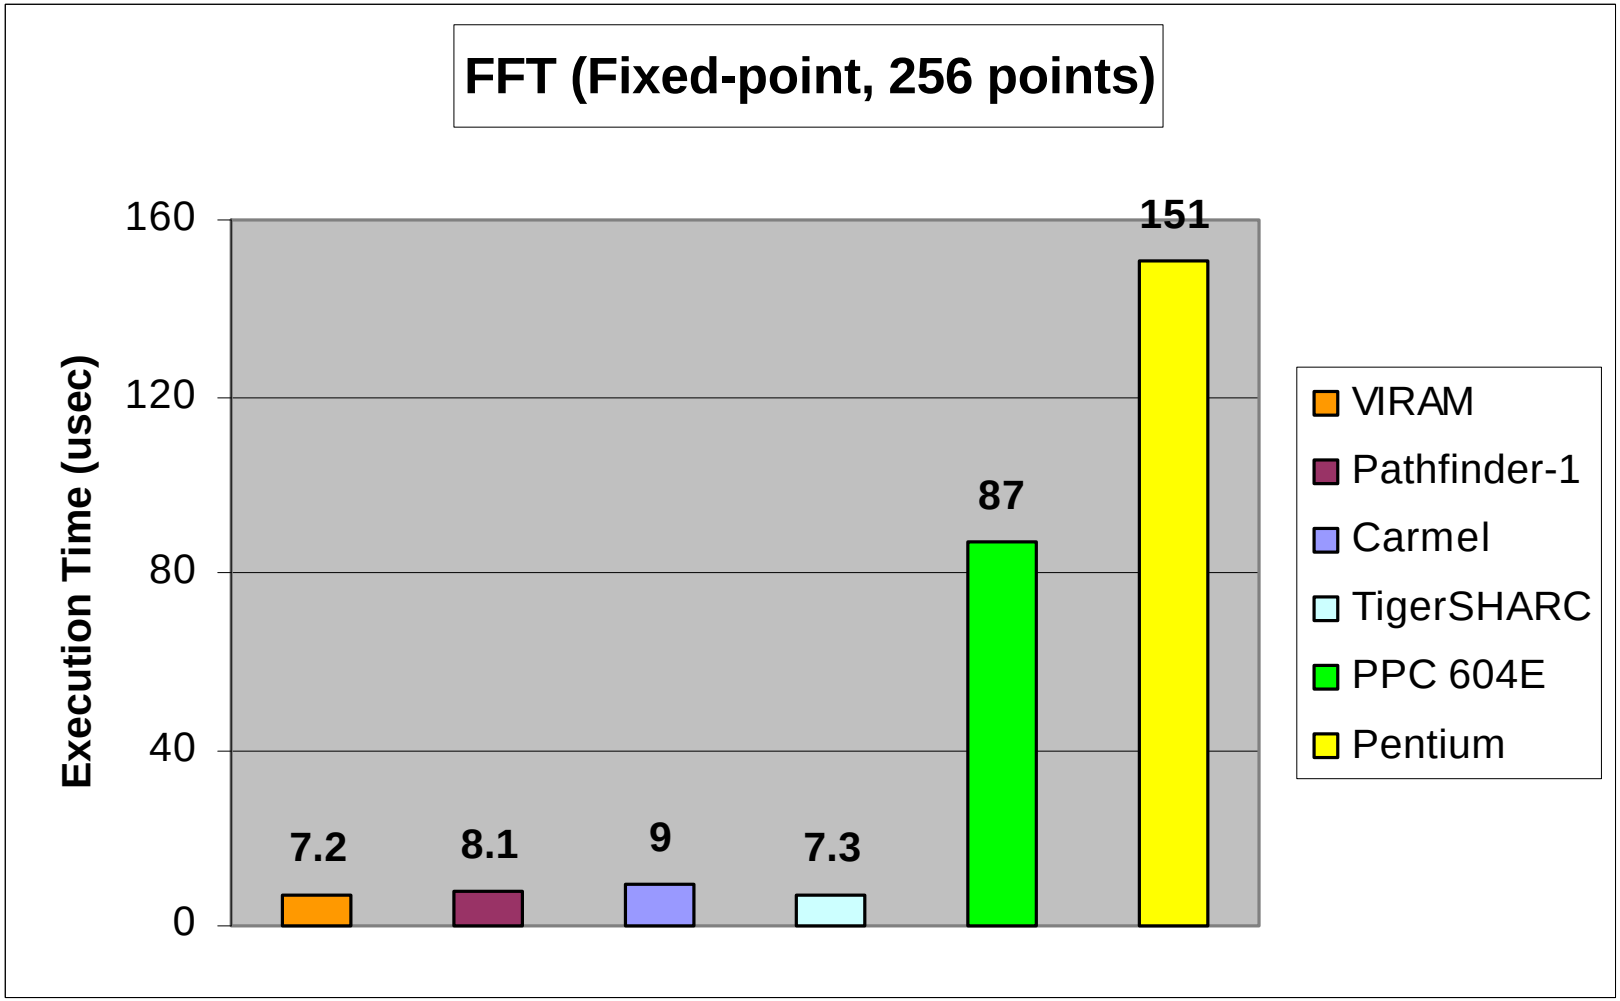

Conclusions

- Vector IRAM
  - An integrated architecture for media processing
  - Based on vector processing and embedded DRAM
  - Simple, scalable, and efficient
- Prototype
  - 16 MBytes DRAM, 256b vector unit
  - 150M transistors, 290 mm^2
  - 3.2 Gops, 2 W at 200 MHz
- Prototype status
  - RTL model completed
  - Back-end design and verification in progress
  - Design tape-out in late Fall 2000

Acknowledgments

- IBM Microelectronics
- MIPS Technologies Inc.
- Cray Inc.
- Raw research group (MIT)

- Randi Thomas, Thinh Nguyen, Dave Judd, James Beck, Krste Asanovic, and Richard Fromm

- This work has been sponsored by DARPA, NSF, and the California State MICRO program