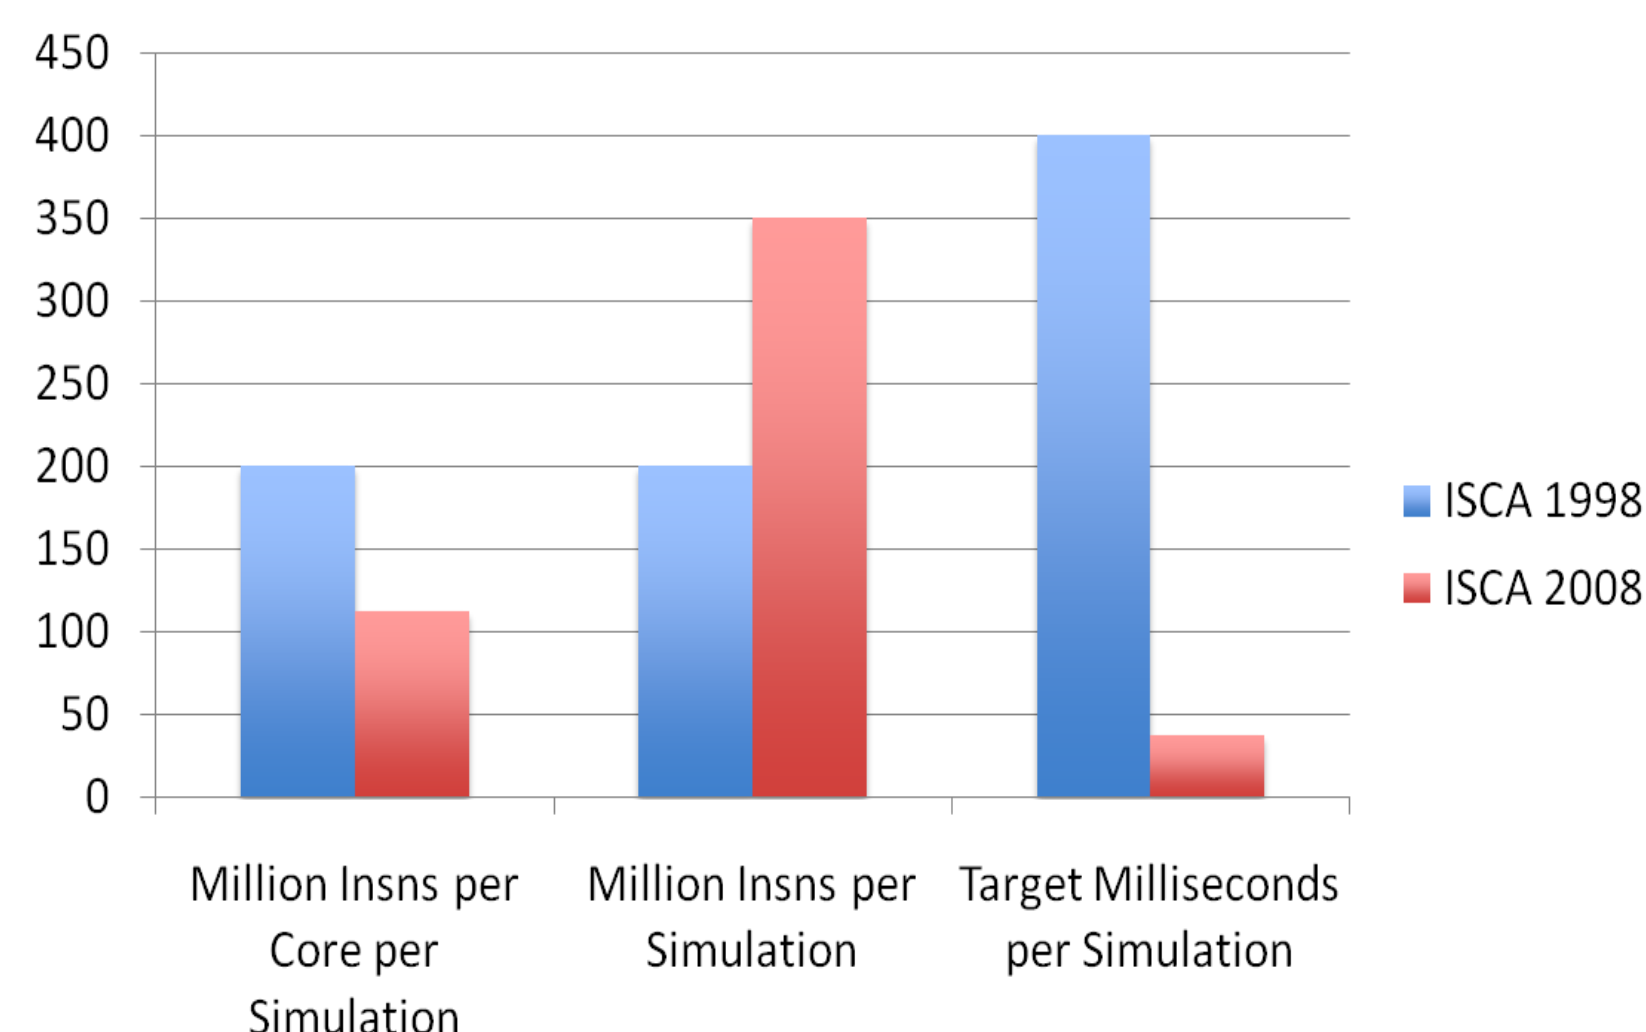


RAMP Gold Demo

Zhangxi Tan, Andrew Waterman, Rimas Avizienis, Yunsup Lee, David Patterson, Krste Asanovic

A Survey of μ Arch Simulation Trends



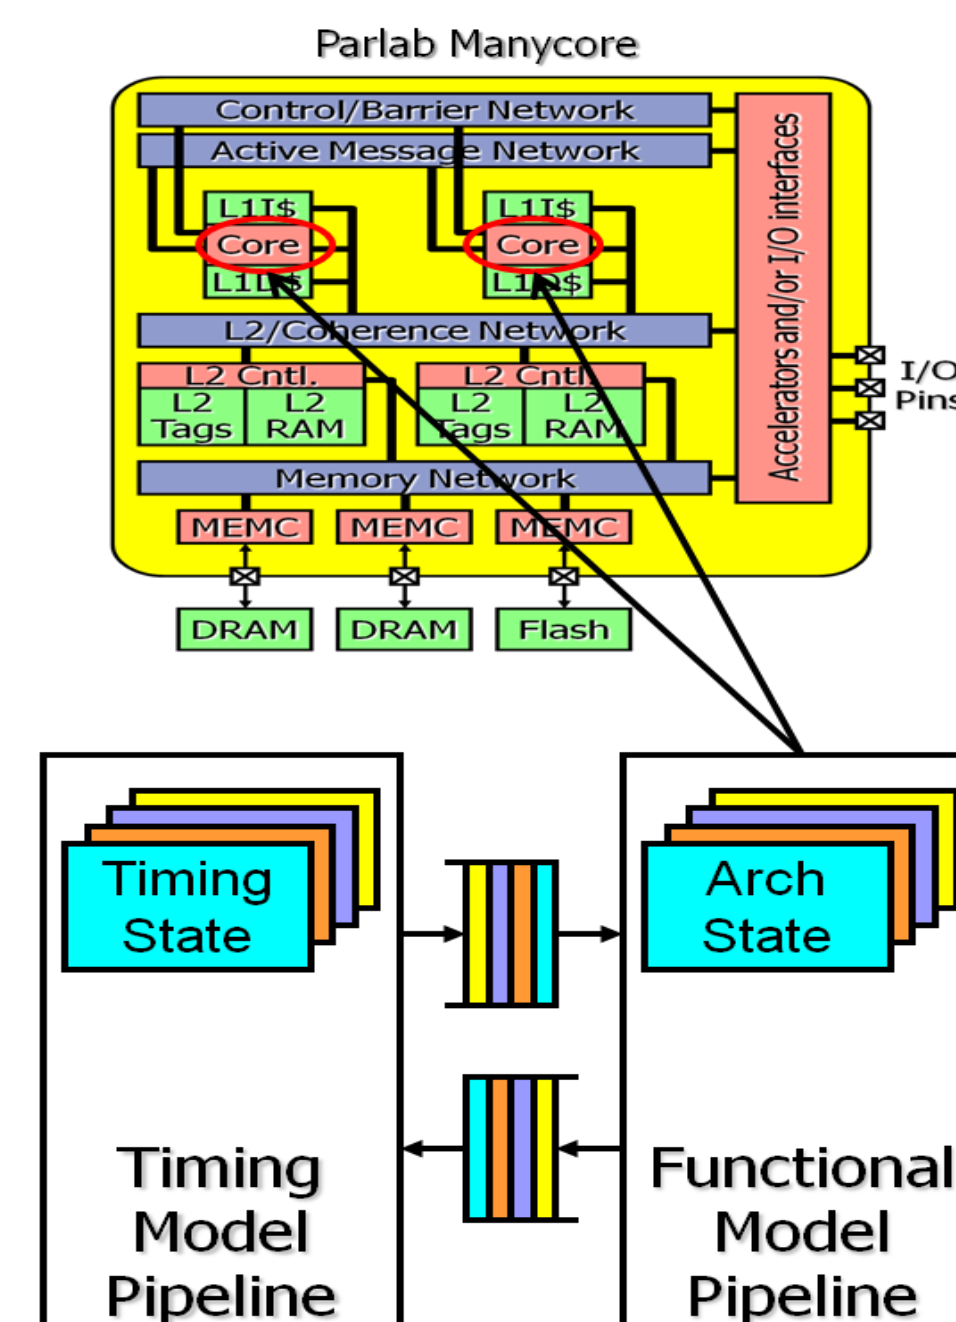
- A median ISCA '08 paper's simulations run for fewer than four OS scheduling quanta!
 - We run yesterday's apps at yesteryear's timescales
 - And attempt to model N communicating cores with $O(1/N)$ instructions per core?!
- The problem is that simulators are too slow
 - Irony: since performance scales as $\sqrt{\text{complexity}}$, simulated instructions per wall-clock second falls as processors get faster

Current Status

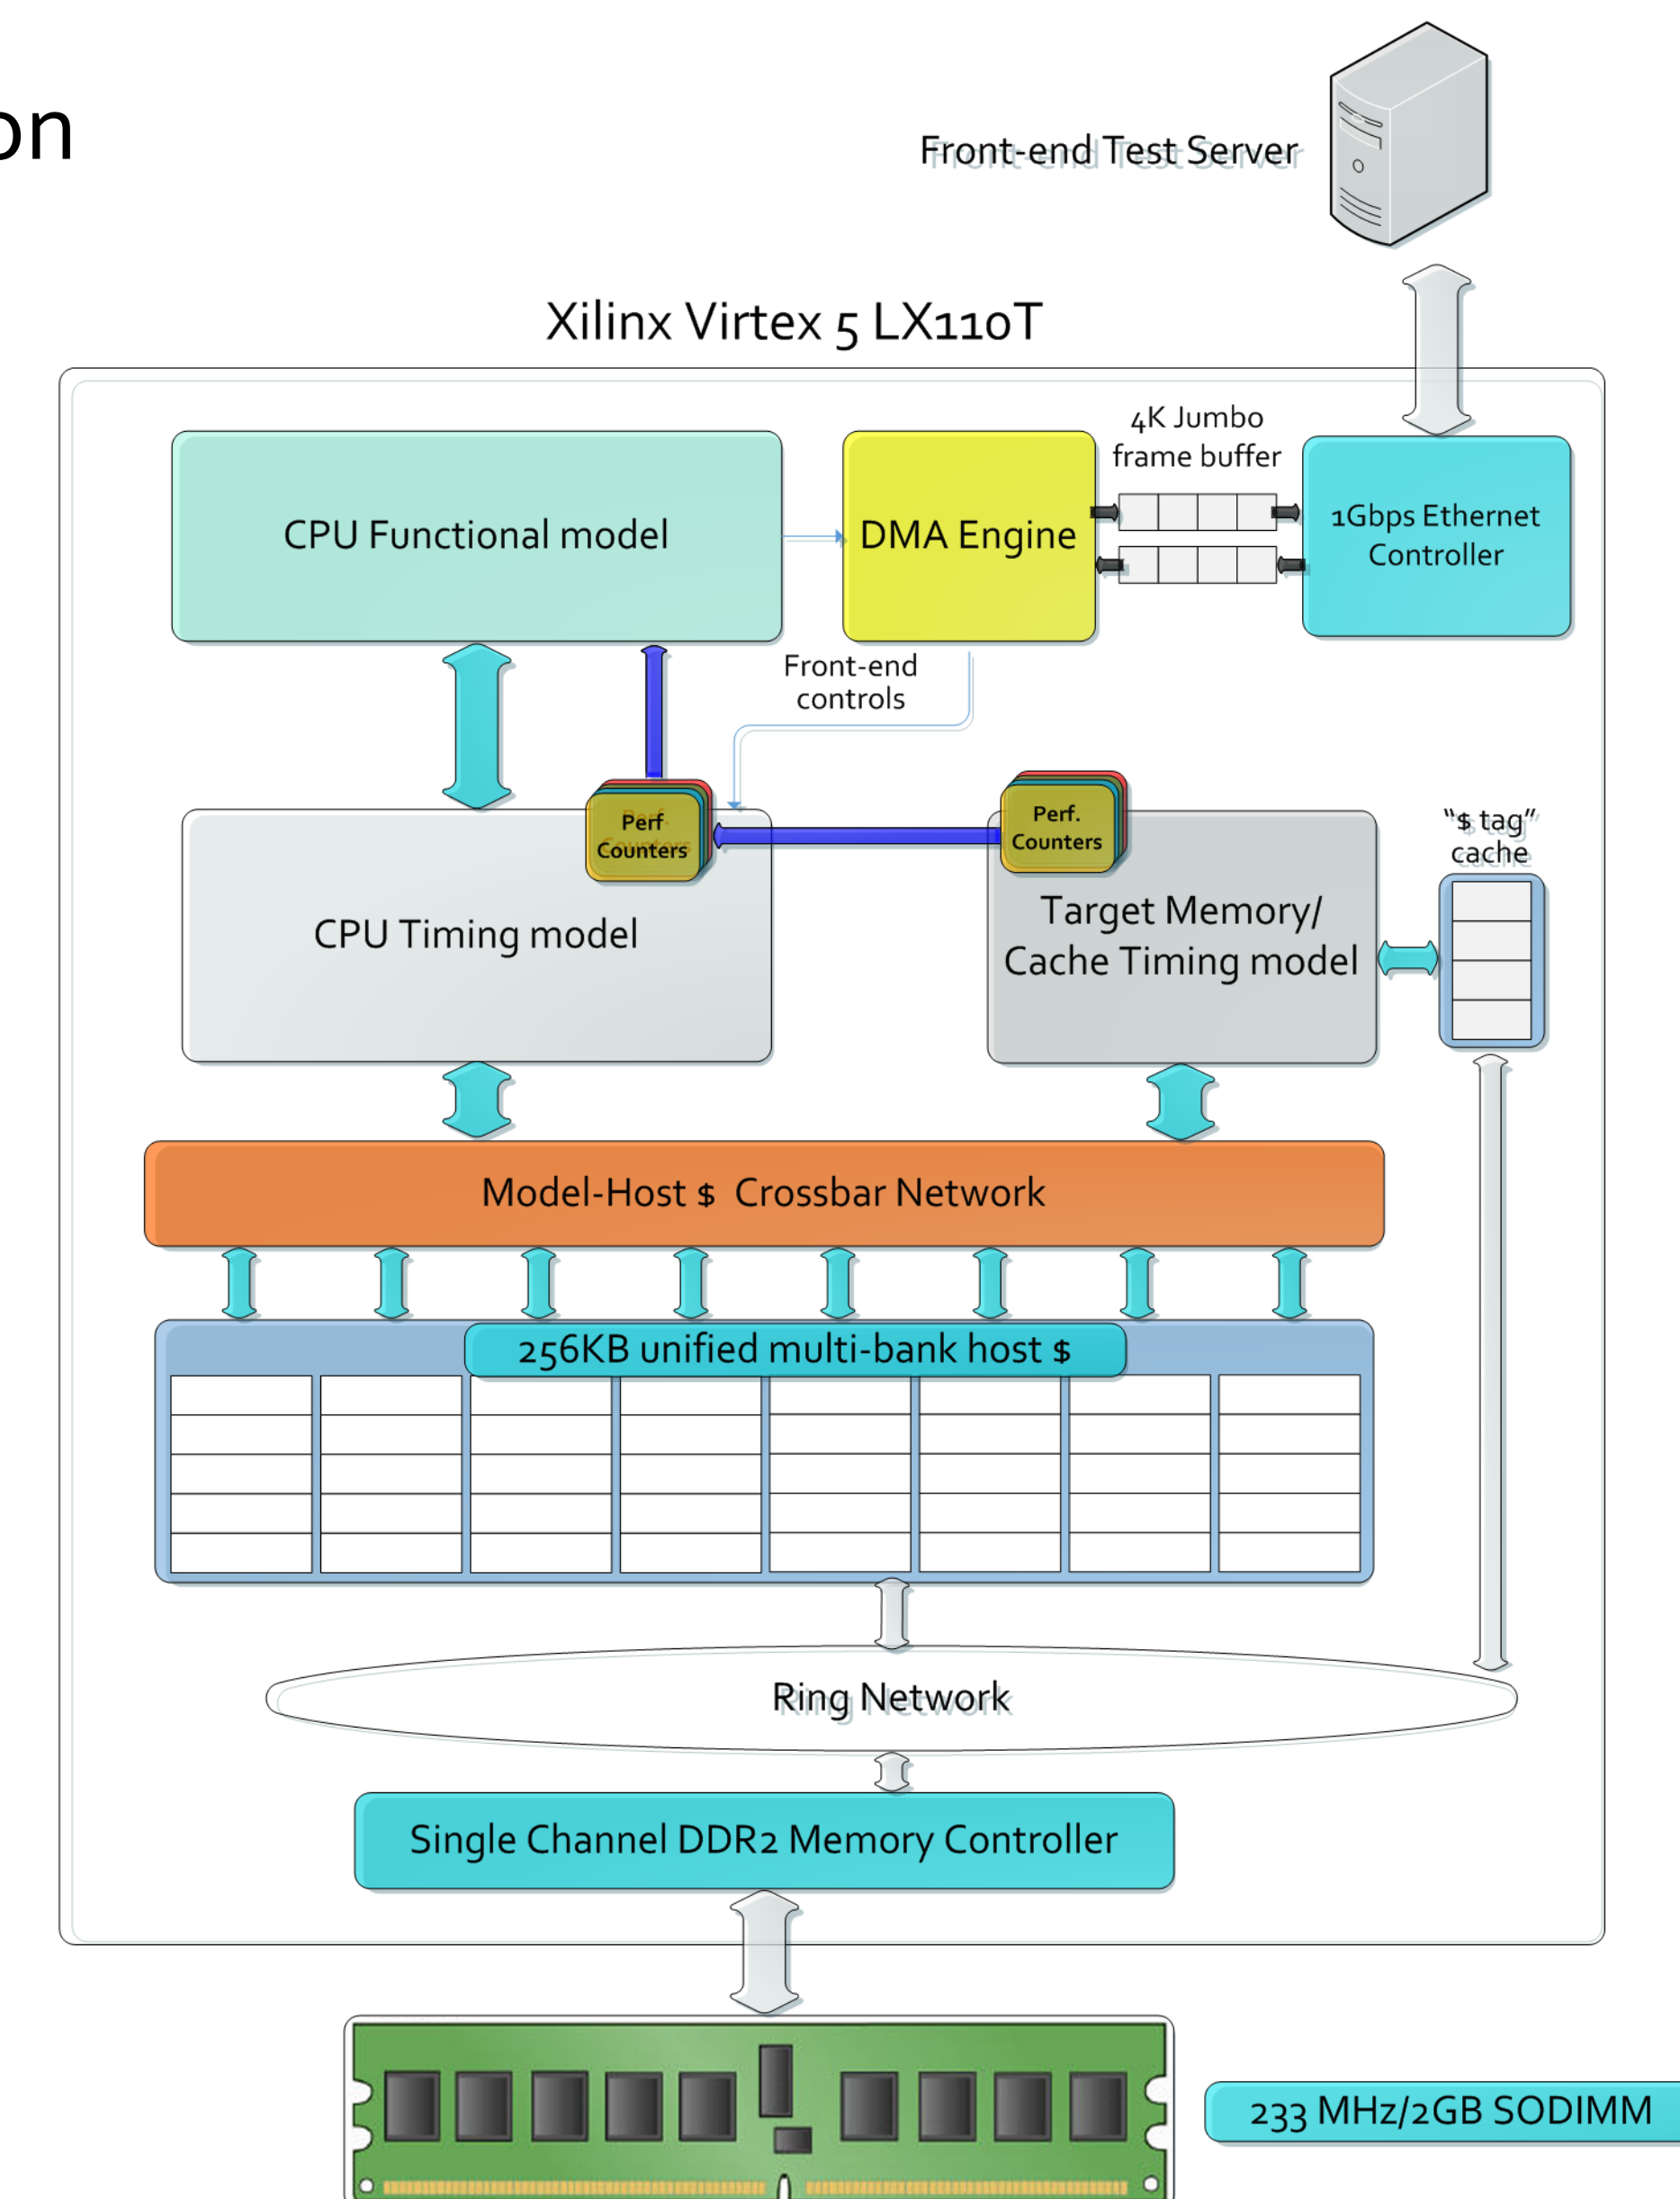
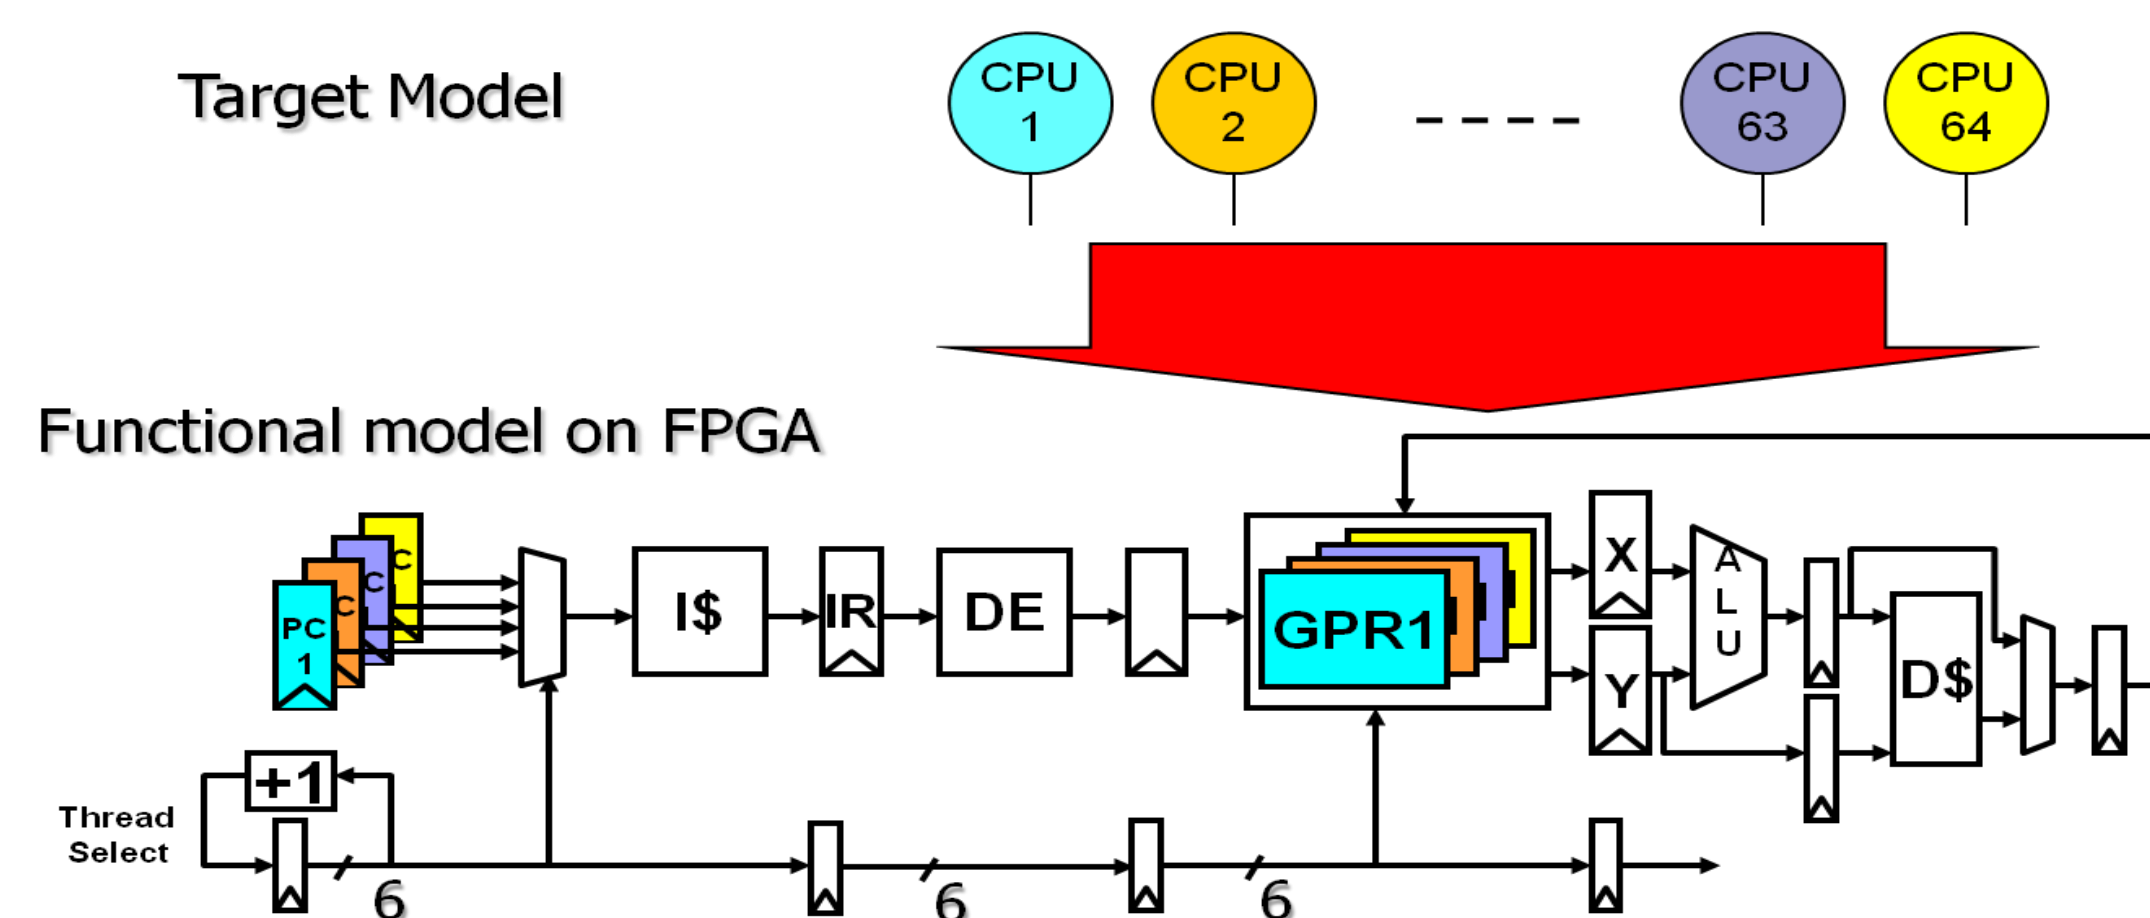
Jan 2009	This retreat
Single Threaded	64 Way Host Multithreaded
0.000032GB BRAM	2GB DDR2 SDRAM
"Hello World"	ParLab Damascene CBIR App, SPLASH2 + SPEC CPU2000
No Timing Model or Introspection	Runtime Configurable Cache Model, Performance Counters
No Floating Point	Hardware FPU Multiply/Add + Software Emulation

- Target Implementation
 - \$750 Xilinx Virtex 5 LX110T FPGA board
 - Conservatively running at 100 MHz
- Emulation Capacity
 - 64-core SPARC v8 SMP
 - Runtime configurable private L1 cache
 - Performance counters

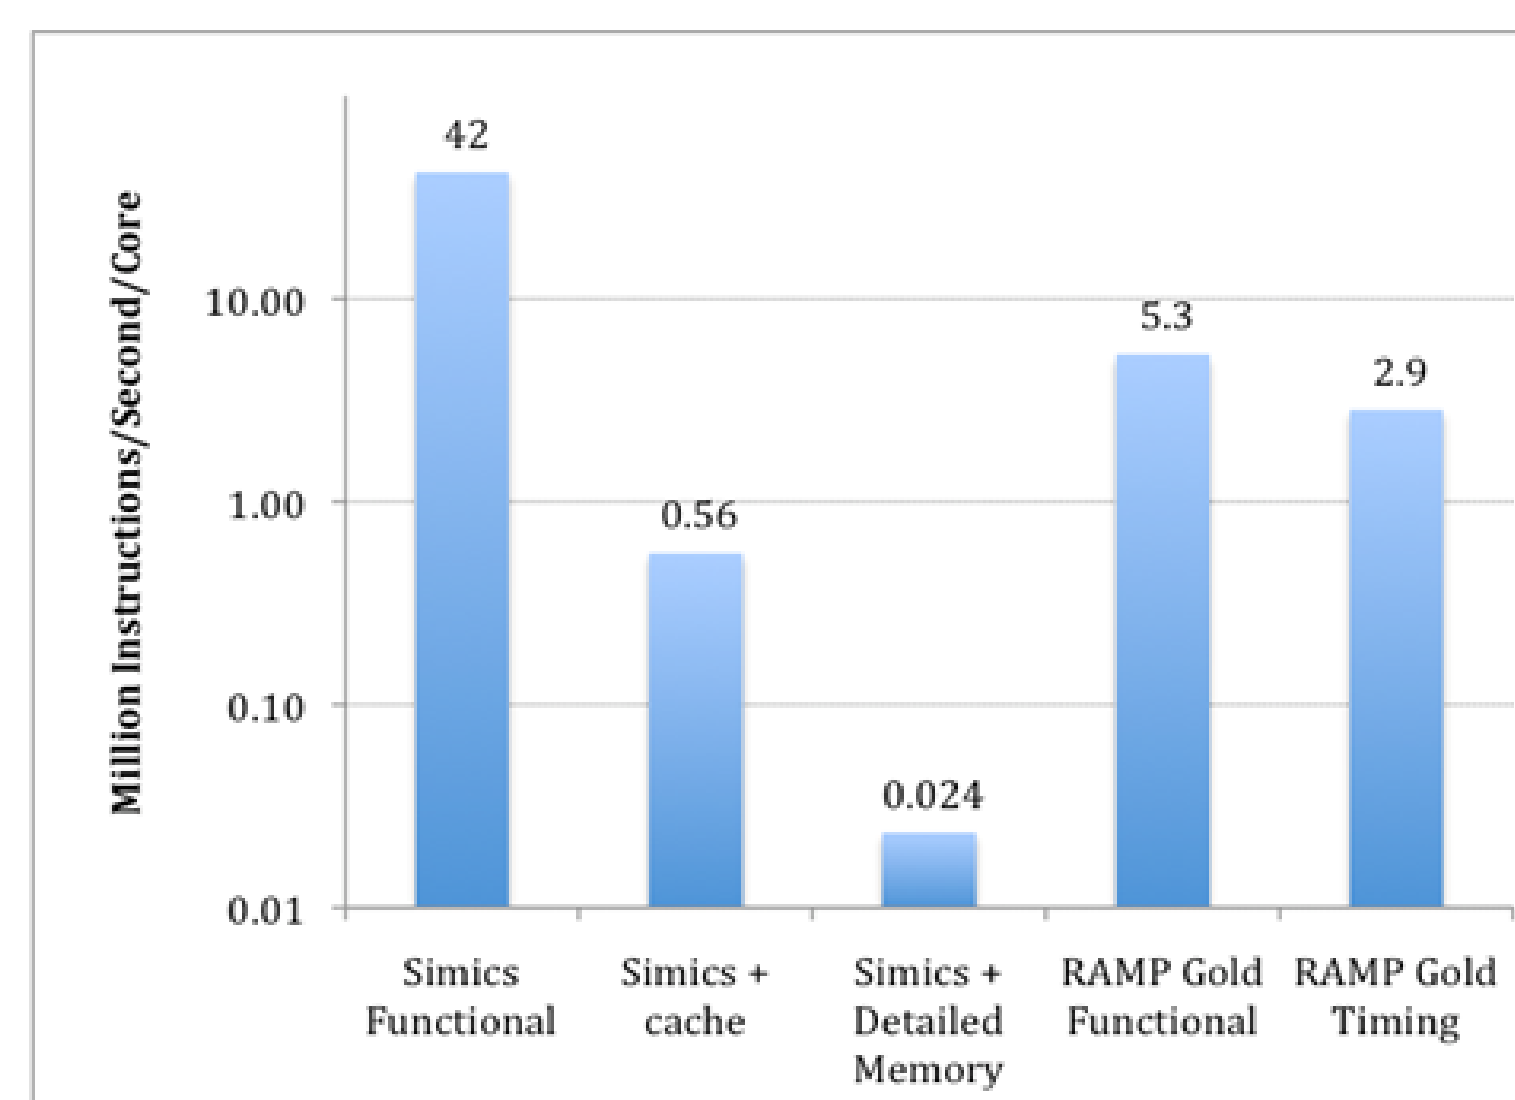
RAMP Gold: Our Solution



- RAMP Gold: A RAMP emulation model for ParLab manycore
 - Single-socket tiled shared memory manycore target
 - SPARC v8 (now)
 - ARM Thumb 2, SPARC v9 (future)
 - Split functional/timing model, both in hardware
 - Host multithreading of both functional and timing models



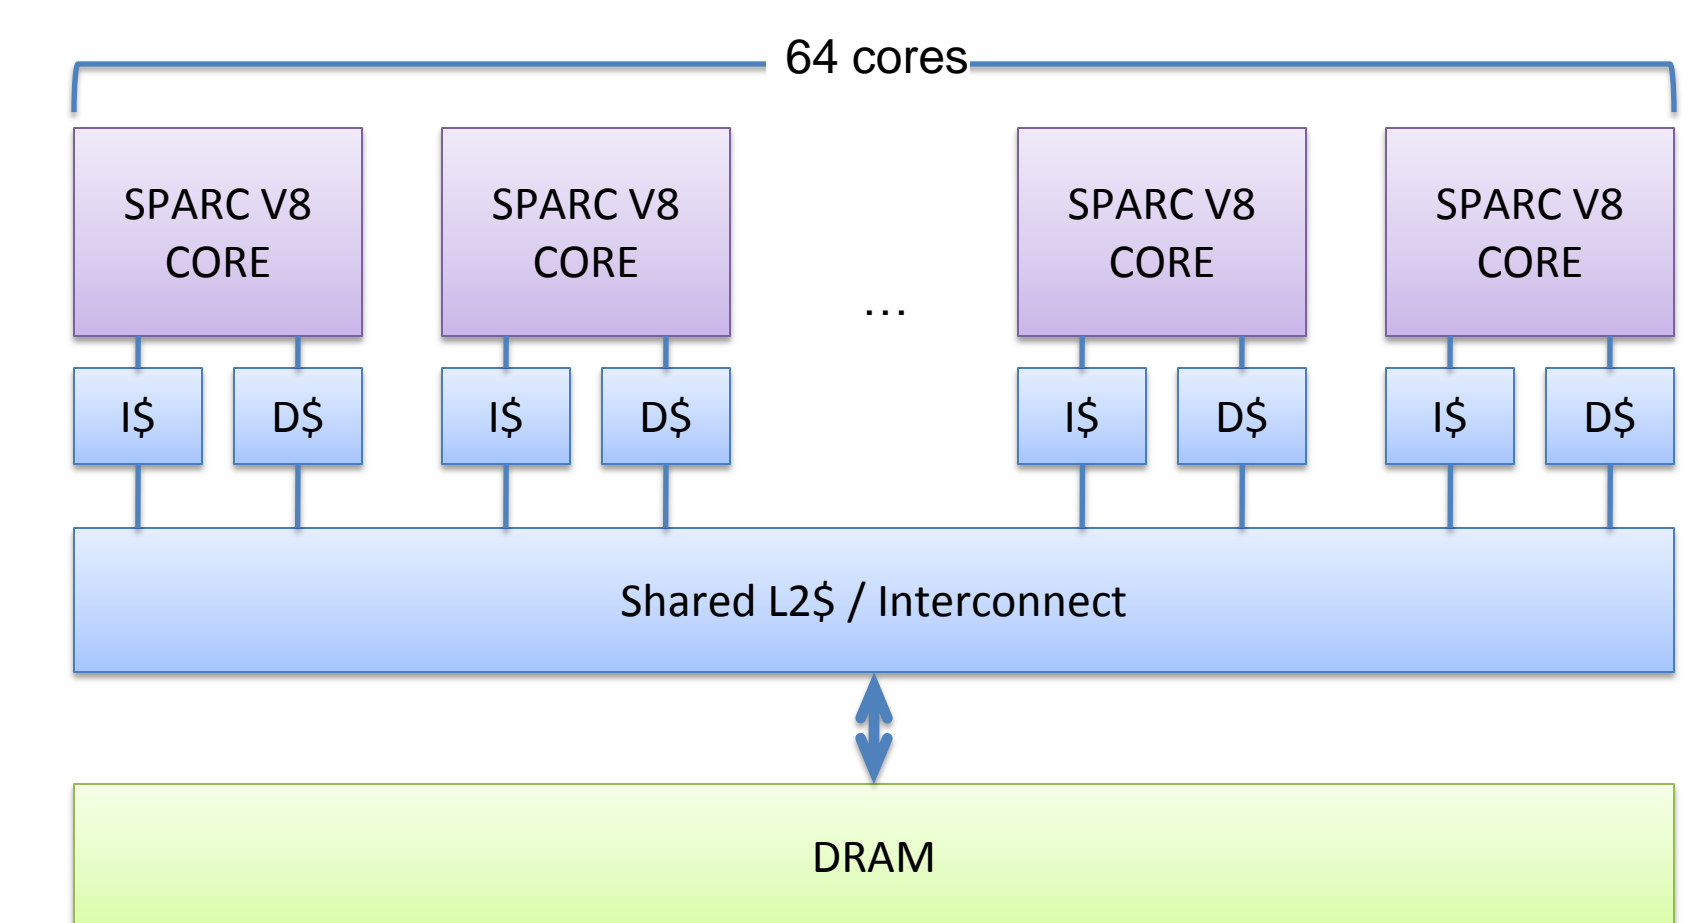
Simulator Performance



	Cost	Performance (MIPS)	Simulations per day
Software Simulator	\$2,000	0.1 - 1	1
RAMP Gold	\$2,000 + \$750	50 - 100	100

RAMP Gold Demo

- Simulated target machine



- SPEC/SPLASH-2 benchmarks
- Damascene (Image Contour Detector)
- Realtime Performance contours