RAMP

Research Accelerator for Multiple Processors

David Patterson (Berkeley, CO-PI), Arvind (MIT), Krste Asanovíc (MIT), Derek Chiou (Texas), James Hoe(CMU), Christos Kozyrakis(Stanford), Shih-Lien Lu (Intel), Mark Oskin (Washington), Jan Rabaey (Berkeley), and John Wawrzynek (Berkeley-PI)

BWRC GSRC IBM Chip in! OpenSPARC™ Power.org™ XILINX®

Conventional Wisdom (CW) in Computer Architecture

- Old Conventional Wisdom:
  Demonstrate new ideas by building chips
- New Conventional Wisdom:
  Mask costs, ECAD costs, GHz clock rates mean
  $\approx$ researchers cannot build believable prototypes
  $\Rightarrow$ simulation only practical outlet

Conventional Wisdom (CW) in Computer Architecture

- Old CW: Power is free, Transistors expensive
- New CW: "Power wall" Power expensive, Xtors free (Can put more on chip than can afford to turn on)
- Old: Multiplies are slow, Memory access is fast
- New: "Memory wall" Memory slow, multiplies fast (200 clocks to DRAM memory, 4 clocks for FP multiply)
- Old : Increasing Instruction Level Parallelism via compilers, innovation (Out-of-order, speculation, VLIW, …)
- New: "ILP wall" diminishing returns on more ILP HW
- New: Power Wall + Memory Wall + ILP Wall = Brick Wall
  - Old CW: Uniprocessor performance 2X / 1.5 yrs
  - New CW: Uniprocessor performance only 2X / 5 yrs?

Uniprocessor Performance (SPECint)

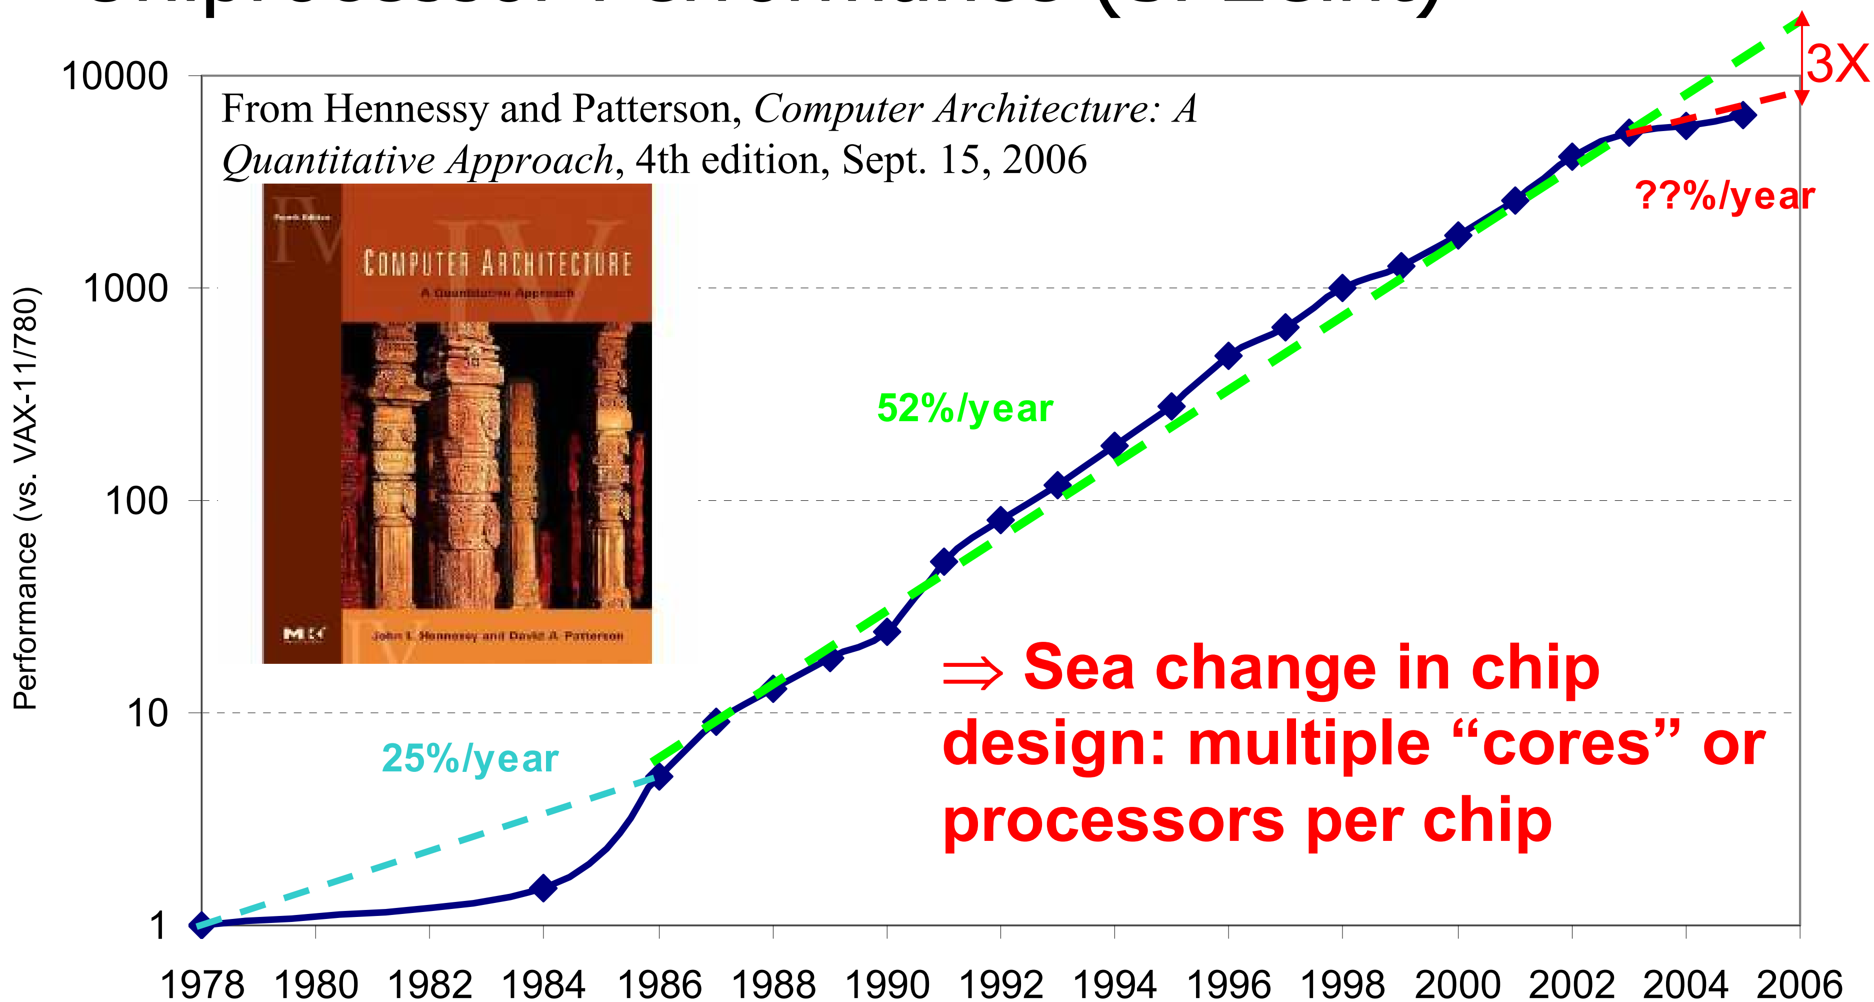

- **VAX : 25%/year 1978 to 1986**
- **RISC + x86: 52%/year 1986 to 2002**
- **RISC + x86: ??%/year 2002 to present**

Déjà vu all over again?

"… today's processors … are nearing an impasse as technologies approach the speed of light.."

David Mitchell, *The Transputer: The Time Is Now* (1989)

- Transputer had bad timing (Uniprocessor performance↑)
  ⇒ Procrastination rewarded: 2X seq. perf. / 1.5 years
- "We are dedicating all of our future product development to multicore designs. … This is a sea change in computing"

Paul Otellini, President, Intel (2005)

- All microprocessor companies switch to MP (2X CPUs / 2 yrs)
  ⇒ Procrastination penalized: 2X sequential perf. / 5 yrs

Manufacturer/Year	AMD/'05	Intel/'06	IBM/'04	Sun/'05
Processors/chip	2	2	2	8
Threads/Processor	1	2	2	4
Threads/chip	2	4	4	32

Outline

- The Parallel Revolution has started
- RAMP Vision
- RAMP Hardware
- Status and Development Plan
- Description Language
- Related Approaches
- Potential to Accelerate MP&NonMP Research
- Conclusions

Problems with "Manycore" Sea Change

1. Algorithms, Programming Languages, Compilers, Operating Systems, Architectures, Libraries, … not ready for 1000 CPUs / chip
2. ≈ Only companies can build HW, and it takes years
3. Software people don't start working hard until hardware arrives
   - 3 months after HW arrives, SW people list everything that must be fixed, then we all wait 4 years for next iteration of HW/SW
4. How get 1000 CPU systems in hands of researchers to innovate in timely fashion on in algorithms, compilers, languages, OS, architectures, … ?
5. Can avoid waiting years between HW/SW iterations?

Build Academic MPP from FPGAs

- As ≈ 20 CPUs will fit in Field Programmable Gate Array (FPGA), 1000-CPU system from ≈ 50 FPGAs?
  - 8 32-bit simple "soft core" RISC at 100MHz in 2004 (Virtex-II)
  - FPGA generations every 1.5 yrs; ≈ 2X CPUs, ≈ 1.2X clock rate
- HW research community does logic design ("gate shareware") to create out-of-the-box, MPP
  - E.g., 1000 processor, standard ISA binary-compatible, 64-bit, cache-coherent supercomputer @ ≈ 150 MHz/CPU in 2007
  - RAMPants: Arvind (MIT), Krste Asanovíc (MIT), Derek Chiou (Texas), James Hoe (CMU), Christos Kozyrakis (Stanford), Shih-Lien Lu (Intel), Mark Oskin (Washington), David Patterson (Berkeley, Co-PI), Jan Rabaey (Berkeley), and John Wawrzynek (Berkeley, PI)
- "Research Accelerator for Multiple Processors"

Characteristics of Ideal Academic CS Research Parallel Processor?

- Scales – Hard problems at 1000 CPUs
- Cheap to buy – Limited academic research $
- Cheap to operate, Small, Low Power – $ again
- Community – Share SW, training, ideas, …
- Simplifies debugging – High SW churn rate
- Reconfigurable – Test many parameters, imitate many ISAs, many organizations, …
- Credible – Results translate to real computers
- Performance – Fast enough to run real OS and full apps, get results overnight

Why RAMP Good for Research MPP?

	SMP	Cluster	Simulate	RAMP
Scalability (1k CPUs)	C	A	A	A
Cost (1k CPUs)	F ($40M)	C ($2-3M)	A+ ($0M)	A ($0.1-0.2M)
Cost of ownership	A	D	A	A
Power/Space (kilowatts, racks)	D (120 kw, 12 racks)	D (120 kw, 12 racks)	A+ (.1 kw, 0.1 racks)	A (1.5 kw, 0.3 racks)
Community	D	A	A	A
Observability	D	C	A+	A+
Reproducibility	B	D	A+	A+
Reconfigurability	D	C	A+	A+
Credibility	A+	A+	F	B+/A-
Perform. (clock)	A (2 GHz)	A (3 GHz)	F (0 GHz)	C (0.1 GHz)
GPA	C	B-	B	**A-**

Why RAMP More Credible?

- Starting point for processor is debugged design from Industry in HDL
- Fast enough that can run more software, do more experiments than simulators
- Design flow, CAD similar to real hardware
  - Logic synthesis, place and route, timing analysis
- HDL units implement operation vs. a high-level description of function
  - Model queuing delays at buffers by building real buffers
- Must work well enough to run OS
  - Can't go backwards in time, which simulators can
- Can measure anything as sanity checks

Can RAMP keep up?

- FGPA generations: 2X CPUs / 18 months
  - 2X CPUs / 24 months for desktop microprocessors
- 1.1X to 1.3X performance / 18 months
  - 1.2X? / year per CPU on desktop?
- However, goal for RAMP is accurate system emulation, not to be the real system
  - Goal is accurate target performance, parameterized reconfiguration, extensive monitoring, reproducibility, cheap (like a simulator) while being credible and fast enough to emulate 1000s of OS and apps in parallel (like a hardware prototype)
  - OK if ≈30X slower than real 1000 processor hardware, provided >1000X faster than simulator of 1000 CPUs

Example: Vary memory latency, BW

- Target system: TPC-C, Oracle, Linux on 1024 CPUs @ 2 GHz, 64 KB L1 I$ & D$/CPU, 16 CPUs share 0.5 MB L2$, shared 128 MB L3$
  - Latency: L1 1 - 2 cycles, L2 8 - 12 cycles, L3 20 - 30 cycles, DRAM 200 – 400 cycles
  - Bandwidth: L1 8 - 16 GB/s, L2 16 - 32 GB/s, L3 32 – 64 GB/s, DRAM 16 – 24 GB/s per port, 16 – 32 DDR3 128b memory ports
- Host system: TPC-C, Oracle, Linux on 1024 CPUs @ 0.1 GHz, 32 KB L1 I$, 16 KB D$
  - Latency: L1 1 cycle, DRAM 2 cycles
  - Bandwidth: L1 0.1 GB/s, DRAM 3 GB/s per port, 128 64b DDR2 ports
  - Use cache models and DRAM to emulate L1$, L2$, L3$ behavior

Accurate Clock Cycle Accounting

- Key to RAMP success is cycle-accurate emulation of parameterized target design
  - As vary number of CPUs, CPU clock rate, cache size and organization, memory latency & BW, interconnet latency & BW, disk latency & BW, Network Interface Card latency & BW, …
  - Least common divisor time unit to drive emulation?

1. For research results to be credible
2. To run standard, shrink-wrapped OS, DB,…
   - Otherwise fake interrupt times since devices relatively too fast

⇒ Good clock cycle accounting is high priority RAMP project

Why 1000 Processors?

- Eventually can build 1000 processors per chip
- Experience of high performance community on stress of level of parallelism on architectures and algorithms
  - 32-way: anything goes
  - 100-way: good architecture and bad algorithms
    or bad architecture and good algorithms
  - 1000-way: good architecture and good algorithms
- Must solve hard problems to scale to 1000
- Future is promising if can scale to 1000

RAMP 1 Hardware

- Completed Dec. 2004 (14x17 inch 22-layer PCB)

Board:

5 Virtex II FPGAs, 18 banks DDR2-400 memory, 20 10GigE conn.

1.5W / computer,
5 cu. in. /computer,
$100 / computer

Box:

10 compute modules in 8U rack mount chassis

1000 CPUs :
≈1.5 KW,
≈ ¼ rack,
≈ $100,000

BEE2: Berkeley Emulation Engine 2

By John Wawrzynek and Bob Brodersen with students Chen Chang and Pierre Droz

RAMP Storage

- RAMP can emulate disks as well as CPUs
  - Inspired by Xen, VMware Virtual Disk models
  - Have parameters to act like real disks
  - Can emulate performance, but need storage capacity
- Low cost Network Attached Storage to hold emulated disk content
  - Use file system on NAS box
  - E.g., Sun Fire X4500 Server ("Thumper")
    48 SATA disk drives,
    24TB of storage @ <$2k/TB

4 Rack Units High

the stone soup of architecture research platforms

Wawrzynek	Hardware
Chiou	Glue-support
Patterson	I/O
Hoe	Coherence
Kozyrakis	Monitoring
Asanovic	Cache
Oskin	Net Switch
Arvind	PPC
Lu	x86

Quick Sanity Check

- BEE2 4 banks DDR2-400 per FPGA
- Memory BW/FPGA = 4 * 400 * 8B = 12,800 MB/s
- 16 32-bit Microblazes per Virtex II FPGA (last generation)
  - Assume 150 MHz, CPI is 1.5 (4-stage pipeline), 33% Load/Stores
  - BW need/CPU = 150/1.5 * (1+ 0.33) * 4B ≈ 530 MB/sec
- BW need/FPGA ≈ 16 * 530 ≈ 8500 MB/s
  - 2/3 Peak Memory BW / FPGA
- Suppose add caches (.75MB ⇒ 32KI$, 16D$/CPU)
  - SPECint2000 I$ Miss 0.5%, D$ Miss 2.8%, 33% Load/stores, 64B blocks*
  - BW/CPU = 150/1.5*(0.5% + 33%*2.8%)*64 ≈ 100 MB/s
- BW/FPGA with caches ≈ 16 * 100 MB/s ≈ 1600 MB/s
  - 1/8 Peak Memory BW/FPGA; plenty BW available for tracing, …
- Example of optimization to improve emulation

* Cantin and Hill, "Cache Performance for SPEC CPU2000 Benchmarks"

Outline

- Parallel Revolution has started
- RAMP Vision
- RAMP Hardware
- Status and Development Plan
- Description Language
- Related Approaches
- Potential to Accelerate MP&NonMP Research
- Conclusions

RAMP Philosophy

- Build vanilla out-of-the-box examples to attract software community
  - Multiple industrial ISAs, real industrial operating systems, 1000 processors, accurate clock cycle accounting, reproducible, traceable, parameterizable, cheap to buy and operate, …
- But RAMPants have grander plans (will share)
  - Data flow computer (“Wavescalar”) – Oskin @ U. Washington
  - 1,000,000-way MP (“Transactors”) – Asanovic @ MIT
  - Distributed Data Centers (“RAD Lab”) – Patterson @ Berkeley
  - Transactional Memory (“TCC”) – Kozyrakis @ Stanford
  - Reliable Multiprocessors (“PROTOFLEX”) – Hoe @ CMU
  - X86 emulation (“UT FAST”) – Chiou @ Texas
  - Signal Processing in FPGAs (“BEE2”) – Wawrzynek @ Berkeley

RAMP multiple ISAs status:

- Got it: IBM Power 405 (32b), Sun SPARC v8 (32b), Xilinx Microblaze (32b)
- Sun announced 3/21/06 donating T1 ("Niagara") 64b SPARC (v9) to RAMP
- Likely: IBM Power 64b
- Likely: Tensilica
- Probably? (had a good meeting): ARM
- Probably? (haven't asked): MIPS32, MIPS64
- No: x86, x86-64
  - But Derek Chiou of UT looking at x86 binary translation

3 Examples of RAMP to Inspire Others

1. Transactional Memory RAMP
   - Based on Stanford TCC
   - Led by Kozyrakis at Stanford
2. Message Passing RAMP
   - First NAS benchmarks (MPI), then Internet Services (LAMP)
   - Led by Patterson and Wawrzynek at Berkeley
3. Cache Coherent RAMP
   - Shared memory/Cache coherent (ring-based)
   - Led by Chiou of Texas and Hoe of CMU

- Exercise common RAMP infrastructure
  - RDL, same processor, same OS, same benchmarks, …

RAMP Milestones

- September 2006 Decide on 1st ISA
  - Verification suite, Running full Linux, Size of design (LUTs/BRAMs)
  - Executes comm. app binaries, Configurability, Friendly licensing
- January 2007 milestones for all 3 RAMP examples
  - Run on Xilinx Virtex 2 XUP board
  - Run on 8 RAMP 1 (BEE2) boards
  - 64 to 128 processors
- June 2007 milestones for all 3 RAMPs
  - Accurate clock cycle accounting, I/O model
  - Run on 16 RAMP 1 (BEE2) boards and Virtex 5 XUP boards
  - 128 to 256 processors
- 2H07: RAMP 2.0 boards on Virtex 5
  - 3rd party sells board, download software and gateware from website on RAMP 2.0 or Xilinx V5 XUP boards

Transactional Memory status (8/06)

- 8 CPUs with 32KB L1 data-cache with Transactional Memory support
  - CPUs are hardcoded PowerPC405, Emulated FPU
  - UMA access to shared memory (no L2 yet)
  - Caches and memory operate at 100MHz
  - Links between FPGAs run at 200MHz
  - CPUs operate at 300MHz
- A separate, 9th, processor runs OS (PowerPC Linux)
- It works: runs SPLASH-2 benchmarks, AI apps, C-version of SpecJBB2000 (3-tier-like benchmark)
- **Transactional Memory RAMP runs 100x faster than simulator on a Apple 2GHz G5 (PowerPC)**

RAMP Blue Prototype (8/06)

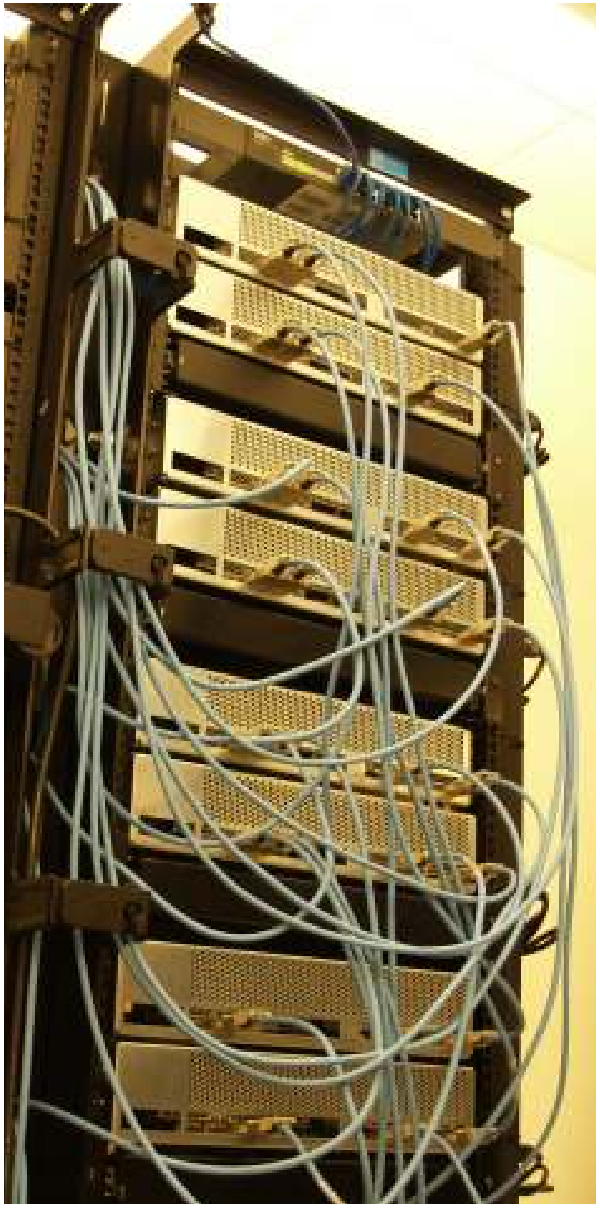

- 8 MicroBlaze cores / FPGA
- 8 BEE2 modules (32 “user” FPGAs) x 4 FPGAs/module = 256 cores @ 100MHz
- Full star-connection between modules
- Diagnostics running today, applications (UPC) this week
- CPUs are softcore MicroBlazes (32-bit Xilinx RISC architecture)
- Also 32-bit SPARC (LEON3)
  - Virtex 2 : 16 CPUs @ 50 MHz; Virtex 5: 60 CPUs @ 120 MHz
  - 30% reduction in number of LUTs from V2 to V5 (4- to 6-input)

RAMP Project Status

- NSF infrastructure grant awarded 3/06
  - 2 staff positions (NSF sponsored), no grad students
- IBM Faculty Awards to RAMPants 6/06
  - Krste Asanovic (MIT), Derek Chiou (Texas), James Hoe (CMU), Christos Kozyrakis (Stanford), John Wawrzynek (Berkeley)
- 3-day retreats with industry visitors
  - "Berkeley-style" retreats 1/06 (Berkeley), 6/06 (ISCA/Boston), 1/07 (Berkeley), 6/07 (ISCA/San Diego)
- RAMP 1/RDL short course
  - 40 people from 6 schools 1/06

RAMP Description Language (RDL)

- RDL describes plumbing connecting units together ≈ "HW Scripting Language/Linker"
- Design composed of **units** that send messages over **channels** via ports
- Units (10,000 + gates)
  - CPU + L1 cache, DRAM controller…
- Channels (≈ FIFO)
  - Lossless, point-to-point, unidirectional, in-order delivery…
- Generates HDL to connect units

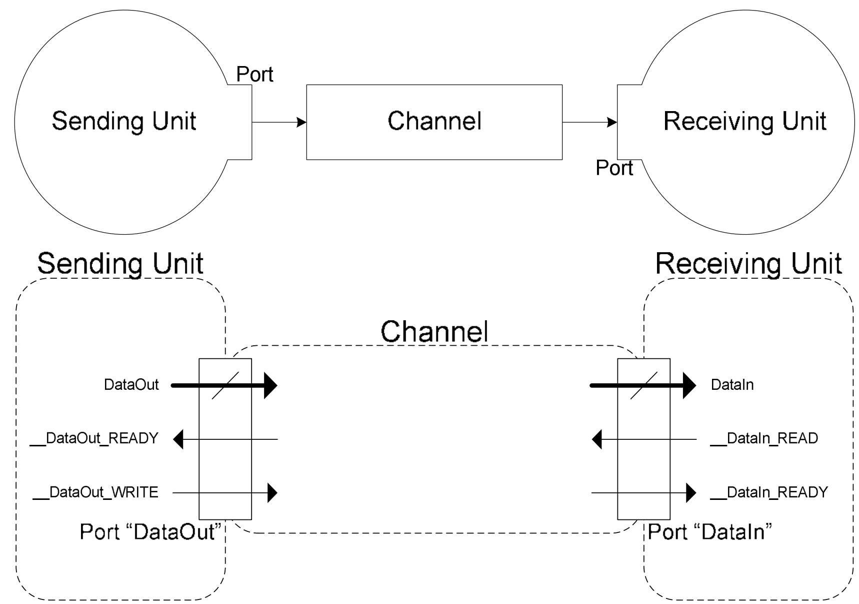

RDL at technological sweet spot

- Matches current chip design style
  - Locally synchronous, globally asynchronous
- To plug unit (in any HDL) into RAMP infrastructure, just add RDL "wrapper"
- Units can also be in C or Java or System C or …
  $\Rightarrow$ Allows debugging design at high level
- Compiles target interconnect onto RAMP paths
  - Handles housekeeping of data width, number of transfers
- FIFO communication model
  $\Rightarrow$ Computer can have deterministic behavior
  - Interrupts, memory accesses, … exactly same clock cycle each run

  $\Rightarrow$ Easier to debug parallel software on RAMP

RDL Developed by Krste Asanovíc and Greg Giebling

Related Approaches

- Quickturn, Axis, IKOS, Thara:
  - FPGA- or special-processor based gate-level hardware emulators
  - HDL mapped to array for cycle and bit-accurate netlist emulation
  - No DRAM memory since modeling CPU, not system
  - Doesn't worry about speed of logic synthesis: 1 MHz clock
  - Uses small FPGAs since takes many chips/CPU, and pin-limited
  - Expensive: $5M
- RAMP's emphasis is on emulating high-level system behaviors
  - More DRAMs than FPGAs: BEE2 has 5 FPGAs, 96 DRAM chips
  - Clock rate affects emulation time: >100 MHz clock
  - Uses biggest FGPAs, since many CPUs/chip
  - Affordable: $0.1 M

RAMP's Potential Beyond Manycore

- **Attractive Experimental Systems Platform: Standard ISA + standard OS + modifiable + fast enough + trace/measure anything**
  - Generate long traces of full stack: App, VM, OS, …
  - Test hardware security enhancements in the wild
  - Inserting faults to test availability schemes
  - Test design of switches and routers
  - SW Libraries for 128-bit floating point
  - App-specific instruction extensions (≈Tensilica)
  - Alternative Data Center designs
    - Akamai vs. Google: N centers of M computers

RAMP's Potential to Accelerate MPP

- **With RAMP: Fast, wide-ranging exploration of HW/SW options + head-to-head competitions to determine winners and losers**
  - Common artifact for HW and SW researchers ⇒ innovate across HW/SW boundaries
  - Minutes vs. years between "HW generations"
  - Cheap, small, low power ⇒ Every dept owns one
  - FTP supercomputer overnight, check claims locally
  - Emulate any MPP ⇒ aid to teaching parallelism
  - If HP, IBM, Intel, M/S, Sun, …had RAMP boxes
    ⇒ Easier to carefully evaluate research claims
    ⇒ Help technology transfer
- **Without RAMP: One Best Shot + Field of Dreams?**

Multiprocessing Watering Hole

Parallel file system Dataflow language/computer Data center in a box
Fault insertion to check dependability Router design Compile to FPGA
Flight Data Recorder Security enhancements Transactional Memory
Internet in a box 128-bit Floating Point Libraries Parallel languages

- Killer app: ≈ All CS Research, Advanced Development
- RAMP attracts many communities to shared artifact
  ⇒ Cross-disciplinary interactions
  ⇒ Ramp up innovation in multiprocessing
- RAMP as next Standard Research/AD Platform? (e.g., VAX/BSD Unix in 1980s)

RAMP Supporters:

- Gordon Bell (Microsoft)
- Ivo Bolsens (Xilinx CTO)
- Jan Gray (Microsoft)
- Norm Jouppi (HP Labs)
- Bill Kramer (NERSC/LBL)
- Konrad Lai (Intel)
- Craig Mundie (MS CTO)
- Jaime Moreno (IBM)
- G. Papadopoulos (Sun CTO)
- Jim Peek (Sun)
- Justin Rattner (Intel CTO)
- Michael Rosenfield (IBM)
- Tanaz Sowdagar (IBM)
- Ivan Sutherland (Sun Fellow)
- Chuck Thacker (Microsoft)
- Kees Vissers (Xilinx)
- Jeff Welser (IBM)
- David Yen (Sun EVP)
- *Doug Burger (Texas)*
- *Bill Dally (Stanford)*
- *Susan Eggers (Washington)*
- *Kathy Yelick (Berkeley)*

RAMP Participants: Arvind (MIT), Krste Asanovíc (MIT), Derek Chiou (Texas), James Hoe (CMU), Christos Kozyrakis (Stanford), Shih-Lien Lu (Intel), Mark Oskin (Washington), David Patterson (Berkeley, Co-PI), Jan Rabaey (Berkeley), and John Wawrzynek (Berkeley, PI)

Conclusions

- **Carpe Diem: need RAMP yesterday**
  - System emulation + good accounting (not FPGA computer)
  - FPGAs ready now, and getting better
  - Stand on shoulders vs. toes: standardize on BEE2
  - Architects aid colleagues via gateware
- **RAMP accelerates HW/SW generations**
  - Emulate, Trace, Reproduce anything; Tape out every day
  - RAMP⇒ search algorithm, language ***and*** architecture space
- "**Multiprocessor Research Watering Hole**"
  Ramp up research in multiprocessing via common research platform ⇒ innovate across fields ⇒ hasten sea change from sequential to parallel computing