

SIMD+ Overview


- Early machines
 - Illiac IV (first SIMD)
 - Cray-1 (vector processor, not a SIMD)
- SIMDs in the 1980s and 1990s
 - Thinking Machines CM-2 (1980s)
 - CPP's DAP & Gamma II (1990s)
- General characteristics
 - Host computer to interact with user and execute scalar instructions, control unit to send parallel instructions to PE array
 - 100s or 1000s of simple custom PEs, each with its own private memory
 - PEs connected by 2D torus, maybe also by row/column bus(es) or hypercube
 - Broadcast / reduction network

1

Fall 2007, SIMD+

Illiac IV History

- First massively parallel (SIMD) computer
- Sponsored by DARPA, built by various companies, assembled by Burroughs, under the direction of Daniel Slotnick at the University of Illinois
 - Plan was for 256 PEs, in 4 quadrants of 64 PEs, but only one quadrant was built
 - Used at NASA Ames Research Center in mid-1970s



2

Fall 2007, SIMD+

Illiac IV Architectural Overview

- CU (control unit) + 64 PUs (processing units)
 - PU = 64-bit PE (processing element) + PEM (PE memory)
- CU operates on scalars, PEs operate on vector-aligned arrays (A[1] on PE 1, A[2] on PE2, etc.)
 - All PEs execute the instruction broadcast by the CU, if they are in active mode
 - Each PE can perform various arithmetic and logical instructions on data in 64-bit, 32-bit, and 8-bit formats
 - Each PEM contains 2048 64-bit words
- Data routed between PEs various ways
- I/O is handled by a separate Burroughs B6500 computer (stack architecture)

3

Fall 2007, SIMD+

Illiac IV Routing and I/O


- Data routing
 - CU bus — instructions or data can be fetched from a PEM and sent to the CU
 - CDB (Common Data Bus) — broadcasts information from CU to all PEs
 - PE Routing network — 2D torus
- Laser memory
 - 1 Tb write-once read-only laser memory
 - Thin film of metal on a polyester sheet, on a rotating drum
- DFS (Disk File System)
 - 1 Gb, 128 heads (one per track)
- ARPA network link (50 Kbps)
 - Illiac IV was a network resource available to other members of the ARPA network

4

Fall 2007, SIMD+

Cray-1 History

- First famous vector (not SIMD) processor
- In January 1978 there were only 12 non-Cray-1 vector processors worldwide:
 - Illiac IV, TI ASC (7 installations), CDC STAR 100 (4 installations)



5

Fall 2007, SIMD+

Cray-1 Vector Operations

- Vector arithmetic
 - 8 vector registers, each holding a 64-element vector (64 64-bit words)
 - Arithmetic and logical instructions operate on 3 vector registers
 - Vector C = vector A + vector B
 - Decode the instruction once, then pipeline the load, add, store operations
- Vector chaining
 - Multiple functional units
 - 12 pipelined functional units in 4 groups: address, scalar, vector, and floating point
 - Scalar add = 3 cycles, vector add = 3 cycles, floating-point add = 6 cycles, floating-point multiply = 7 cycles, reciprocal approximation = 14 cycles
 - Use pipelining with data forwarding to bypass vector registers and send result of one functional unit to input of another

6

Fall 2007, SIMD+

Cray-1 Physical Architecture

- Custom implementation
 - Register chips, memory chips, low-speed and high-speed gates
- Physical architecture
 - “Cylindrical tower (6.5’ tall, 4.5’ diameter) with 8.5’ diameter seat”
 - Composed of 12 wedge-like columns in 270° arc, so a “reasonably trim individual” can get inside to work
 - World’s most expensive love-seat”
 - “Love seat” hides power supplies and plumbing for Freon cooling system
- Freon cooling system
 - Vertical cooling bars line each wall, modules have a copper heat transfer plate that attaches to the cooling bars
 - Freon is pumped through a stainless steel tube inside an aluminum casing

7

Fall 2007, SIMD+

Cray X-MP, Y-MP, and {CJT}90

- At Cray Research, Steve Chen continued to update the Cray-1, producing...
- X-MP
 - 8.5 ns clock (Cray-1 was 12.5 ns)
 - First multiprocessor supercomputer
 - 4 vector units with scatter / gather
- Y-MP
 - 32-bit addressing (X-MP is 24-bit)
 - 6 ns clock
 - 8 vector units
- C90, J90 (1994), T90
 - J90 built in CMOS, T90 from ECL (faster)
 - Up to 16 (C90) or 32 (J90/T90) processors, with one multiply and one add vector pipeline per CPU

8

Fall 2007, SIMD+

Cray-2 & Cray-3

- At Cray Research, Steve Chen continued to update the Cray-1 with improved technologies: X-MP, Y-MP, etc.
- Seymour Cray developed Cray-2 in 1985
 - 4-processor multiprocessor with vectors
 - DRAM memory (instead of SRAM), highly interleaved since DRAM is slower
 - Whole machine immersed in Fluorinert (artificial blood substitute)
 - 4.1 ns cycle time (3x faster than Cray-1)
 - Spun off to Cray Computer in 1989
- Seymour Cray developed Cray-3 in 1993
 - Replace the “C” shape with a cube so all signals take same time to travel
 - Supposed to have 16 processors, had 1 with a 2 ns cycle time

9

Fall 2007, SIMD+

Thinking Machines Corporation's Connection Machine CM-2

- Distributed-memory SIMD (bit-serial)
- Thinking Machines Corp. founded 1983
 - CM-1, 1986 (1000 MIPS, 4K processors)
 - CM-2, 1987 (2500 MFLOPS, 64K...)
- Programs run on one of 4 Front-End Processors, which issue instructions to the Parallel Processing Unit (PE array)
 - Control flow and scalar operations run on Front-End Processors, while parallel operations run on the PPU
 - A 4x4 crossbar switch (Nexus) connects the 4 Front-Ends to 4 sections of the PPU
 - Each PPU section is controlled by a Sequencer (control unit), which receives assembly language instructions and broadcasts micro-instructions to each processor in that PPU section

10

Fall 2007, SIMD+

CM-2 Nodes / Processors

- CM-2 constructed of “nodes”, each with:
 - 32 processors (implemented by 2 custom processor chips), 2 floating-point accelerator chips, and memory chips
- 2 processor chips (each 16 processors)
 - Contains ALU, flag registers, etc.
 - Contains NEWS interface, router interface, and I/O interface
 - 16 processors are connected in a 4x4 mesh to their N, E, W, and S neighbors
- 2 floating-point accelerator chips
 - First chip is interface, second is FP execution unit
- RAM memory
 - 64Kbits, bit addressable

11

Fall 2007, SIMD+

CM-2 Interconnect

- Broadcast and reduction network
 - Broadcast, Spread (scatter)
 - Reduction (e.g., bitwise OR, maximum, sum), Scan (e.g., collect cumulative results over sequence of processors such as parallel prefix)
 - Sort elements
- NEWS grid can be used for nearest-neighbor communication
 - Communication in multiple dimensions: 256x256, 1024x64, 8x8192, 64x32x32, 16x16x16x16, 8x8x4x8x8x4
- The 16-processor chips are also linked by a 12-dimensional hypercube
 - Good for long-distance point-to-point communication

12

Fall 2007, SIMD+

DAP Overview

- Distributed-memory SIMD (bit-serial)
- Cambridge Parallel Processing
 - International Computers Limited (ICL) built 1976 prototype, deliveries in 1980
 - ICL spun off Actime Memory Technology Ltd in 1986, became CPP Inc in 1992
- Matrix of PEs
 - One-bit PEs with 32Kb–1Mb of memory
 - 2D torus, plus column & row buses
 - 32x32 for DAP 500, 64x64 for DAP 600
- DAP system = host + MCU + PE array
 - Host (Sun or VAX) interacts with user
 - Master control unit (MCU) runs main program, PE array runs parallel code

13

Fall 2007, SIMD+

DAP MCU and PE Array

- MCU (Master Control Unit)
 - 32-bit 10 MHz CPU w/ registers, instruction counter, arithmetic unit, etc.
 - Executes scalar instructions and broadcasts instruction streams to PEs
- Processing Elements in PE array
 - 3 1-bit registers
 - Q = accumulator, C = carry, A = activity control (inhibit memory writes)
 - All bits of a register over all PEs is called a “register plane” (32x32 or 64x64 bits)
 - Adder
 - Two inputs connect to Q and C registers
 - Third input connects to multiplexor
 - Mux reads rom PE memory, output of Q or A registers, carry output from neighboring PEs, or data broadcast from MCU
 - PE outputs (adder and mux) can be stored in memory, under control of A reg

14

Fall 2007, SIMD+

Gamma II^{Plus}

- Fourth-generation DAP, produced by Cambridge Parallel Processing in 1995
- Gamma II^{Plus} 1000 = 32x32
Gamma II^{Plus} 4000 = 64x64
- PE memory: 128Kb–1Mb
- PE also contains an 8-bit processor
 - 32 bytes of internal memory
 - D register to transfer data to/from array memory (1-bit data path) and to/from internal memory (8-bit data path)
 - A register, similar to a 1-bit processor
 - Q register, like accumulator, 32 bits wide (any one of which can be selected as an operand), can also be shifted
 - ALU to provide addition, subtraction, and logical operations

15

Fall 2007, SIMD+