# Wall's Paper

Limits of instruction-level parallelism.

Wall, David W.

WRL-93-6
November 1993
External

68 Pages

Seminal paper. Does it still hold ? What is the way forward then ?

# Going Beyond ILP

- Diminishing returns for further effort extracting ILP from a single thread?

- *System-level* parallelism
  - some workloads naturally parallel
    - multi-user machine
    - application plus XServer
    - application plus asynchronous I/O

- *Process/Thread-level* parallelism
  - Some applications already multithreaded

    - database, HTTP server, NFS server
    - fork, pthreads
  - may have smaller cache footprint
  - may be same Virtual address space

- *Loop-level* parallelism
  - generated by auto-parallelizing compilers
  - co-operative threads
  - need fast synchronization, communication, fork

# Exploiting Parallelism

- Multiple CPUs on a chip
  - – Exploit thread/process level parallelism
  - – Use traditional SMP mechanisms
  - ✘ Need correspondingly bigger caches and external memory bandwidth
  - – IBM Power4 2-way SMP on a chip
- **Multi-threading**
  - – Use one CPU to execute multiple threads
  - – Replicate PCs, architectural register file
  - – Different virtual address spaces?
- Static multi-threading
  - – Round-robin issue from a large # threads
  - ✔ No instruction dependencies
  - ✔ Hides memory latency
    - ∗ No expensive caches
  - ✔ Fast synchronization / fork possible
  - ✘ Requires many register files

- ✘ Progress of an individual thread is slow
  - * Poor SPEC marks (great SPEC Rate)
  - – Tera/Cray MTA, 128 threads

- Course-grained multi-threading
  - – Switch between threads on a major stall
  - – e.g. cache miss on Stanford SPARCLE

# Simultaneous Multi-Threading (SMT)

- Work on a small number of threads at once, aiming to keep all function units busy
- Duplicate architectural state
- Duplicate instruction fetch units
- Need to control allocation of resources
  - priority . fair share
  - (prioritising can be counter productive)

✔ Progress of individual threads is pretty good

✔ Cooperating threads may have smaller cache footprint than independent ones

✔ Potential for register-register synchronization and communication

✔ Potential for lightweight thread create

- Pentium IV Xeon uses 2-way "hyperthreading"
  - 2 virtual CPUs per chip
  - looks like SMP - separate VM contexts
  - Staticly partitions resources if both active
  - SMT `halt` and `pause` instructions
  - OS scheduler should understand SMT

# Sun Niagra T1

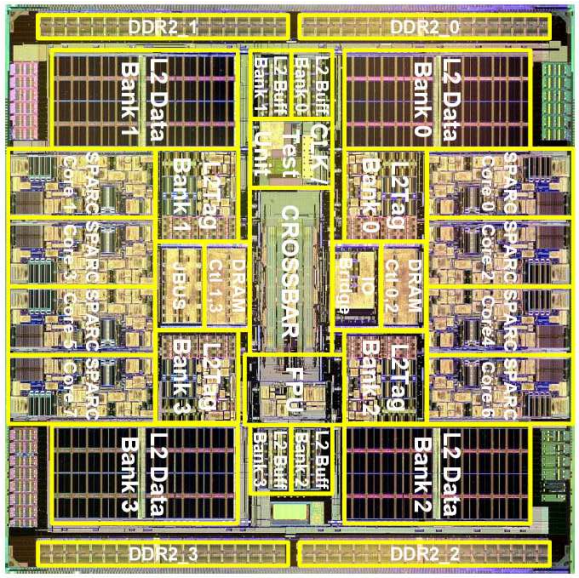


- Fine Grain Multiplexing (Barrel Processor)
- Five stage pipeline only
- A new, non-halted thread each issue.
- Four DRAM controllers on chip
- Shared FPU

# Texas TriMedia

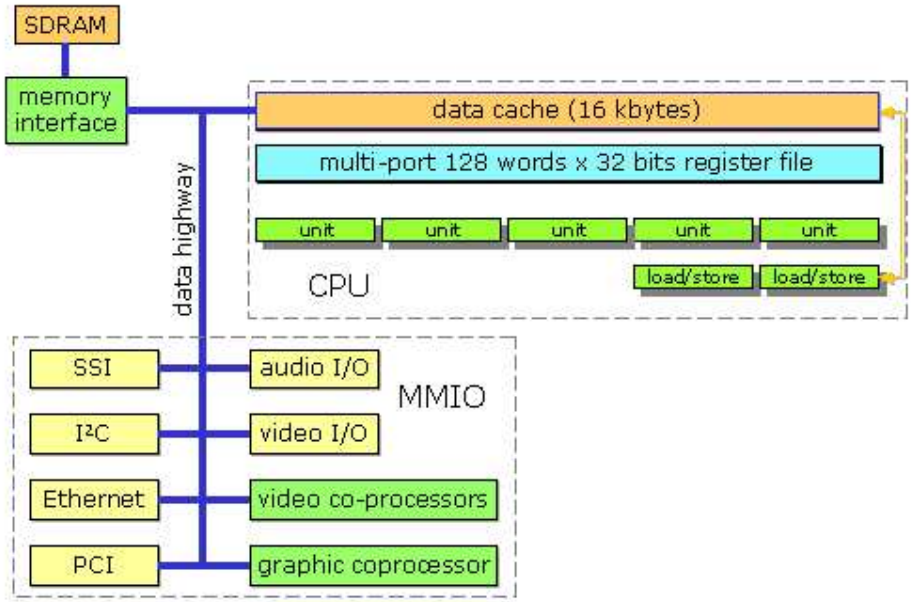


- 5 issue slots filled with up to 43 functional units
- 128 32-bit General Purpose Registers
- SIMD & DSP operations
- 32-bit IEEE 754 floating point operations
- 8/16/32/64 KB Instruction cache, 8/16/32/64/128 KB Data cache
- Separate memory and peripheral bus interfaces.

# Custom Silicon

- Compilation of SystemC to hardware.
- New HDLs, such as System Verilog.
- Can make special purpose co-processors (steadycam, MPEG).
- Compile legacy code to FPGA for acceleration.
- Greaves CTOV datapath inference algorithm.
- Ramp Blue, PicoChip and Clearspeed.

# State of the Art

| Processor | Alpha 21364 EV-78+ | AMD Opteron 254 | AMD Dual-core Opteron 280 | HP PA-8900 | IBM Power4+ | IBM Power5 |
|---|---|---|---|---|---|---|
| Processor Arch | 64-bit | 32/64-bit | Dual 32/64-bit | Dual 64-bit | Dual 64-bit | Dual, MT 64-bit |
| Clock Rate | 1.30GHz | 2.8GHz | 2.4GHz | 1.16GHz | 1.7GHz | 1.9GHz |
| Cache (I/D/L2/L3) | 64K/64K/ 1.75M | 64K/64K/ 1M | 2 x 64K/64K/1M | 1.5M/1.5M/ 64M | 64K/32K/ 1.5MB | 64K/32K/ 1.92MB/36MB |
| Issue Rate/Core | 4 issue | 3 x86 instr | 3 x86 instr | 4 issue | 8 issue | 8 issue |
| Pipeline Stages | 7/9 stages | 9/11 stages | 9/11 stages | 7/9 stages | 12/17 stages | 12/17 stages |
| Out of Order | 80 instr | 72ROPs | 72ROPs | 56 instr | 200 instr | 200 instr |
| Rename Regs | 48/41 | 36/36 | 36/36 | 56 total | 48/40 | 48/40 |
| BHT Entries | 4K x 9-bit | 4K x 2-bit | 4K x 2-bit | 8K x 2-bit | 3x16K x 1-bit | 3x16K x 1-bit |
| TLB Entries | 128/128 | 280/288 | 280/288 | 2 x 240 unified | 2x1,024 unified | 2x1,024 unified |
| Memory B/W | 12GB/s | 6.4GB/s | 6.4GB/s | 6.4GB/s | 12.8GB/s | 12.8GB/s |
| Package | FC-LGA-1443 | PGA-940 | PGA-940 | LGA-544 | MCM | MCM |
| IC Process | 0.18µm 7M | 0.13µm 6M | 0.09µm 7M | 0.13µm 7M | 0.13µm 7m | 0.13µm 7m |
| Die Size | 397mm² | 193mm² | 199mm² | 304mm² | 267mm²** | 389mm²** |
| Transistors | 135 million | 106 million | 233 million | 300 million | 184 million** | 276 million** |
| Est Die Cost | $180 | $79 | $85 | $96 | $144** | $200** |
| Power (Max) | 155W | 92W(MTP)* | 95W(MTP) | 103W | 100W** | 120W* |
| Availability | 3Q04 | 4Q05 | 4Q05 | 3Q03 | 2Q03 | 4Q05 |
| Configuration | 2–64 way | 1–2 way | 1–2 way | 1–128 way | 2–32 way | 2–32 way |
| SPEC_int2000(base) | 904 | 1,817 | 1,499 | N/A | 1,077 | 1,470 |
| SPEC_fp2000(base) | 1,279 | 2,132 | 1,752 | N/A | 1,598 | 2,839 |
| **Processor** | **Intel Itanium 2** | **Intel XeonMP** | **Intel Xeon** | **MIPS R16000** | **Fujitsu SPARC64 V** | **Sun UltraSPARC VI+** |
| Processor Arch | 64-bit | 32/64-bit | 32/64-bit | 64-bit | 64-bit | Dual 64-bit |
| Clock Rate | 1.66GHz | 3.66GHz | 3.8GHz | 700MHz | 2.16GHz | 1.5GHz |
| Cache (I/D/L2/L3) | 16K/16K/ 256K/9M | 12K/8K/ 1M/1M | 12K/512K/ 2M | 32K/32K | 128K/128K/4M | 64K/64K/ 2MB/32MB |
| Issue Rate/Core | 6 issue | 3 ROPs | 3 ROPs | 4 issue | 8 issue | 8 issue |
| Pipeline Stages | 8 stages | 22/24 stages | 22/24 stages | 6 stages | 9 stages (int) | 14 stages |
| Out of Order | None | 126 ROPs | 126 ROPs | 48 instr | 112 instr | None |
| Rename Regs | 328 total | 128 total | 128 total | 32/32 | 32/32 | None |
| BHT Entries | 512 x 2-bit | 4K x 2-bit | 4K x 2-bit | 2K x 2-bit | 16K x 2-bit | 2 x 16 x 2-bit |
| TLB Entries | 32L1I/32L1D/ 128L2I/128L2D | 128I/64D | 128I/64D | 64 unified | (2,048+32)I/ (2,048+32)D | 2 x (512+16)I/ 2 x (1,024+16)D |
| Memory B/W | 10.6GB/s | 5.3GB/s | 6.4GB/s | 1.6GB/s | 4.3GBs | 4.8GB/s |
| Package | mPGA-700 | mPGA-604 | mPGA-604 | FCBGA-1153 | LGA-908 | FC-LGA 1368 |
| IC Process | 0.13µm 6M | 0.09µm 6M | 0.09µm 6M | 0.11µm 7M | 0.09µm 10M | 0.09µm 9M |
| Die Size | 432mm² | 130mm²* | 145mm²* | 110mm² | 294mm² | 336mm²** |
| Transistors | 592 million | 125 million* | 175 million* | 7.2 million | 400 million | 295 million** |
| Est Die Cost | $165 | $22 | $24 | $60 | N/A | $125** |
| Power (Max) | 130W | 140W(TDP) | 130W(TDP) | 17W | 65W | 90W* |
| Availability | 3Q05 | 2Q05 | 4Q05 | 1Q03 | 2Q05 | 3Q05 |
| Configuration | 1–256 way | 1–8 way | 1–2 way | 1–512 way | 1–128 way | 4–72 way |
| SPEC_int2000(base) | 1,490 | 1,388 | 1,810 | N/A | 1,456 | N/A |
| SPEC_fp2000(base) | 2,801 | 1,314 | 1,909 | N/A | 1,808 | N/A |

Source: vendors, except *In-Stat estimates. Estimated manufacturing cost does not include external cache chips. ** Contains two processors on one die. n/a = not available.

# Other techniques

- Data-flow processors
  - Fine-grained control-flow, course-grained data-flow (opposite of standard super-scalar)
  - Begin execution of a block of sequential instructions when all inputs become available
  - ✘ Inputs are memory locations. The matching store required to figure out when all inputs are ready is large and potentially slow. (matching is easier with a small number of registers *a la* out-of-order execution)