Dataflow & Tiled Architectures

WaveScalar and TRIPS

- Irene Lin & Kevin Rohan

WaveScalar

- Motivation
- Solution & Implementation
- Results
- Conclusion & Discussion

Motivation

- Scaling up superscalar is hard
 - Circuit complexity, fast transistors, slow wires, communication infrastructure
- Von Neumann means sequential
 - Sequential fetch (PC) and memory
- Untapped dataflow locality
 - predictability in the dynamic data dependencies

Solution & Implementation

- Chunk CFG into waves
- Every data value carries a tag, every wave has a wave number.
- Total ordering of memory operations
 - From wave number and memory instruction sequence number
- Increment wave number with WAVE-ADVANCE instruction
- Conditional split to steer data value to destination
 - Converting control dependencies into data dependencies

WaveCache

Grid of processing elements for:

- control instruction placement
- input output queues
- communication logic
- functional unit

Cluster of 4 PEs:

- L1 cache
- Store buffer

processor fetch is data-driven







Conclusion & Discussion

- "Scalable, low complexity, high performance"
 - Is it actually scalable?? Compiler Scalability??
- Dataflow driven rather than von neumann style linearity
 - Increase in parallelism
- Binaries are bigger
 - Maybe not relevant in the present scenario
- Miss is a heavyweight event
- Compiler vs programmer responsibility
 - Superscalar and WaveScalar both split up the program into blocks
 - Multi-core systems programmer's responsibility

TRIPS

- Motivation
- Solution & Implementation
- Results
- Conclusion
- Discussion

Motivation

- Pretty much the same as before
 - Scalability, reduce the circuit complexity etc.
 - Reduce burden on programmer
- Does not replace Von Neumann architecture

Solution & Implementation

EDGE ISA

- Block atomic execution
- Direct instruction communication within a block

Micro Arch

- Operand Network (bypass reg file and memory)
- Tiles (global control, execution, register, data, instruction) Compiler
- Conventional optimizations
- Translate code to TRIPS intermediate language and make TRIPS blocks
- Translate blocks into assembly



On-Chip Network

Processor 1

Figure 2. TRIPS die photo with tile overlays.

"Need as many as 2–4 times more instructions than the Alpha, due to aggressive predication."

Predication works by executing instructions from both paths of the branch and only permitting those instructions from the taken path to modify architectural state.



Figure 4. TRIPS instructions normalized to Alpha for compiled (C) and hand-optimized (H) benchmarks.

Memory Instruction \Rightarrow Register Instructions

Register Instructions ⇒ Direct Communication



Figure 5. Storage accesses normalized to Alpha for compiled (C) and hand-optimized (H) benchmarks.

Microarchitecture Evaluation : ILP Evaluation



Microarchitecture Evaluation : v/s Commercial



Figure 10. Speedup of TRIPS relative to the Core 2 on simple, EEMBC, and SPEC benchmarks.

Conclusion

"the performance and potential energy efficiency of EDGE designs may be sufficiently large to justify adoption in mobile systems or data centers, where <u>high</u> <u>performance at low power</u> is essential."

Polymorphic processor - every task can run on every unit

Discussion

WaveScalar vs TRIPS:

- Data-flow type of execution in both
- Inter-block communication:
 - TRIPS register file
 - WaveScalar WaveCache
- TRIPS is a Von Neumann architecture
- WaveScalar has 2.5 times more speedup as more parallelism