Breaking the Memory Bottleneck with an Optical Data Path

Jason Fritts

Assistant Professor

Department of Computer Science

Co-Author: Roger Chamberlain

Overview

- **Memory Bottleneck**
- **Optical Technology**
  - VCSELs
  - Optical Links
- **Evaluation Environment**
- **Optical Bus Simulation**
  - Optical Bus with Current Technology
  - Optical Bus in Future Processors
  - Optical Bus with Additional Prefetching
- **Conclusions and Future Work**

2

Memory Bottleneck

- **Growing Processor-Memory Gap**
  - Longer Memory Latencies
  - Lower Memory Bandwidths
  - Memory Latency is Major Bottleneck
  - Memory Bandwidth is next Major Bottleneck
    - Doug Burger (1997)
      - examined processor and memory performance over a 2 decade period
      - found bandwidth is now becoming a critical bottleneck
- **Existing Solutions**
  - Primarily Focus on Memory Latency
  - Aggressive Latency Hiding Techniques
    - lockup-free caches
    - data speculation
    - cache-conscious load scheduling
    - hardware and software prefetching

Memory Bandwidth is next Major Bottleneck

3

Optical Technology

- **Traditionally used in long-haul communications**
  - very high bandwidth
  - low latency
  - moderate cost
- **Beginning to use in short-distance communications**
  - multiprocessor networks (R. Chamberlain and M. Franklin)
- **VCSELs – Now can take fiber directly to chip!**

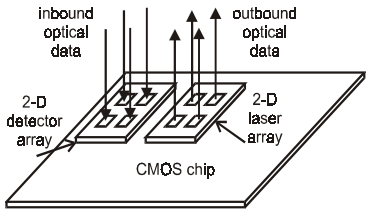

4

Optical Processor-Memory Bus

- **Feasibility**
  - VCSEL current < 1mA => Low Power!
    - Initial esimations indicate power likely much lower than existing electrical buses
  - Metal-Semiconductor-Metal (MSM) technology for photodectors fairly mature
  - Area for laser driver/receiver < ½ area for electrical pad and drive circuitry
  - Simple, effective fiber alignment methods have been demonstrated
- **Very High Bandwidth**
  - Orders of magnitude greater than traditional electrical buses
  - 16x16 arrays of VCSELs possible
    - 32 x 32 arrays available soon
  - Each VCSEL supports > 1 Gb/s rates
  - Can achieve at least 256 Gb/s
    - 1 Tb/s will be available soon
- **Comparable transmission latency**

5

Optical Bus Links

- **Rigid free-space optical links**
  - Intra-board communications

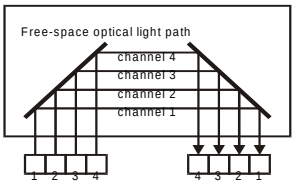

- **Fiber image guide optical links**
  - Intra-board communications, and
  - Board-to-board communications

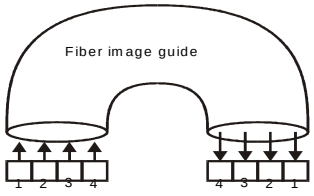

6

How Does Near-Infinite Bandwidth Affect Processor-Memory Performance?

8 - 10 Gb/s

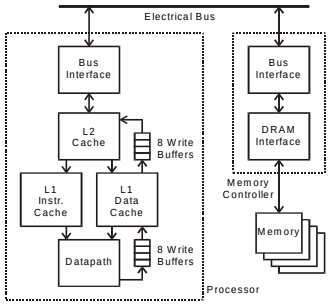

=>

256 Gb/s – 1 Tb/s

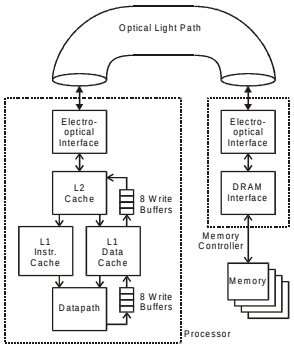

7

Evaluation Environment

8

IMPACT Compilation/Simulation Environment

- **Aggressive ILP research compiler**
- **Architecture-independent evaluation**
  - Large, generic instruction set
  - Retargetable back-end
- **Cycle-accurate trace-driven simulation**
  - Models variety of architectures:
    - in-order superscalar
    - out-of-order superscalar
    - VLIW
  - Highly parameterizable
  - Supports sample/skip simulation
    - 400,000 instructions per sample

$$max\_skip\_size = \max\left(\left[\frac{\min\left(1x10^{9}, trace\_size\right)}{50} - sample\_size\right], 0\right)$$

9

Benchmark

- **Benchmark of Image/Video Decompression Algorithms**
  - Large volumes of data
  - Few computations per memory element
  - More likely to be memory bound
  - MPEG-4 and image/video decoding apps from MediaBench

[CLee97] "MediaBench: A Tool for Evaluating and Synthesizing Multimedia Communication Systems," MICRO-30, 1997.

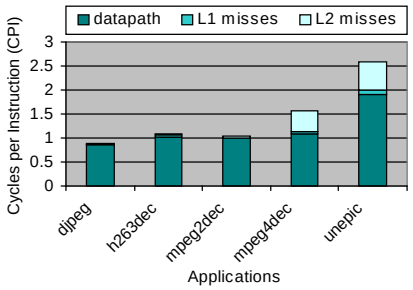

10

Benchmark Characteristics

Program	# Dynamic Instrs	Skip Size	% Program Simulated	Simulation Time (min)	Simulation Time (max)
djpeg	3M	0	100	0.43 min	0.48 min
h263dec	60M	1.5M	23.6	1.10 min	1.66 min
mpeg2dec	95M	2.5M	14.8	2.64 min	3.00 min
mpeg4dec	1400M	9.6M	4	6.20 min	9.14 min
unepic	5M	0	100	1.25 min	2.92 min

Program	# Static Instrs	Input-1		Input-2	
		File Size	*# Dynamic*	*File Size*	*# Dynamic*
djpeg	19,397	5756	3M	31,074	25M
h263dec	8721	20,364	60M	19,338	65M
mpeg2dec	9520	34,906	95M	1,593,409	720M
mpeg4dec	108,273	39,213	1400M	503,060	500M
unepic	3767	7432	5M	10,129	5M

11

Base Architecture Model

- **Architecture model**
  - 4-issue media processor
  - 1 GHz processor frequency
  - pipeline: 1 fetch, 2 decode, 1 write back, variable execute stages
  - 64-bit processor-memory bus
- **L1 Cache**
  - 16 KB direct-mapped L1 instruction cache w/ 256 byte lines
  - 32 KB direct-mapped L1 data cache w/ 64 byte lines
- **On-Chip L2 Cache**
  - 256 KB 4-way set associate w/ 64 byte lines
- **External Memory**
  - 80 ms – memory access and controller time
  - 64 ms – memory transfer time

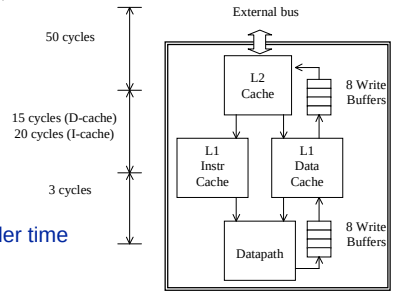

12

Optical Processor-Memory Bus Evaluation

13

Optical Bus with Current Technology

- **Memory Access Time:** $T_{L2\_miss} = T_A + \frac{T_T}{X_B}$
- **Optical bus virtually eliminates memory transfer time**

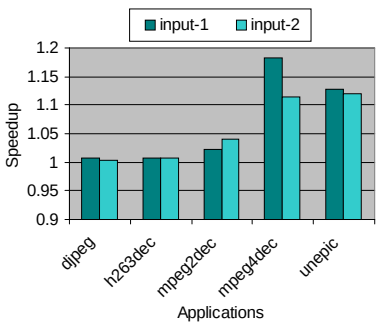

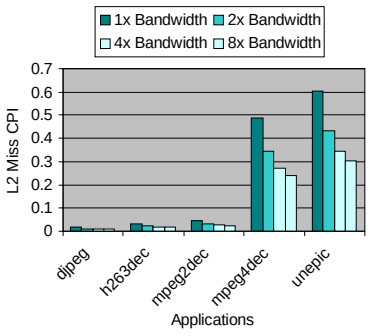

14

Optical Bus in Future Processors

- **Memory Access Time:**

$$T_{L2\_miss} = X_L * \left(T_A + \frac{T_T}{X_B} \right)$$

- **Improvement consistent across all benchmarks**
  - varied by degree of memory use
- **Up to 3.3x improvement**
  - mpeg4dec
  - unepic

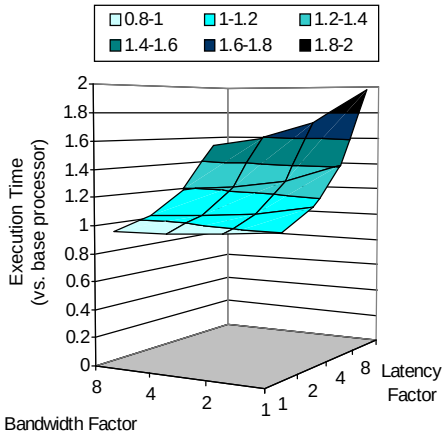

15

Optical Bus in Future Processors

- **Average reduction memory CPI is 50%**
- **However, close examination shows smaller reduction in L2 miss at higher latency factors (51%, 45%, 42%, 38%)**

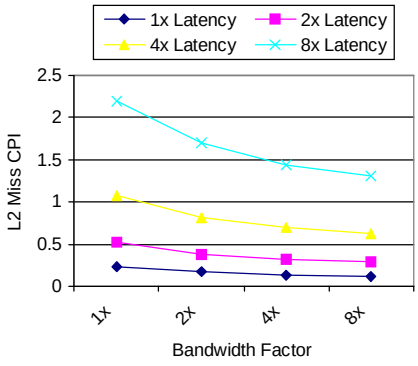

16

Optical Bus with Additional Prefetching

(see paper)

17

Conclusions

- **Optical processor-memory bus is feasible**
- **Optical bus virtually eliminates memory transfer time**
- **Optical bus is increasingly effective at higher memory latencies**
- **Benchmarks were not memory-bound, but on memory-bound applications, performance will be much more dramatic**
- **Significant opportunities for optical buses**
  - More aggressive latency hiding
  - Coalescing multiple buses
  - Multi-hop optical buses

18