

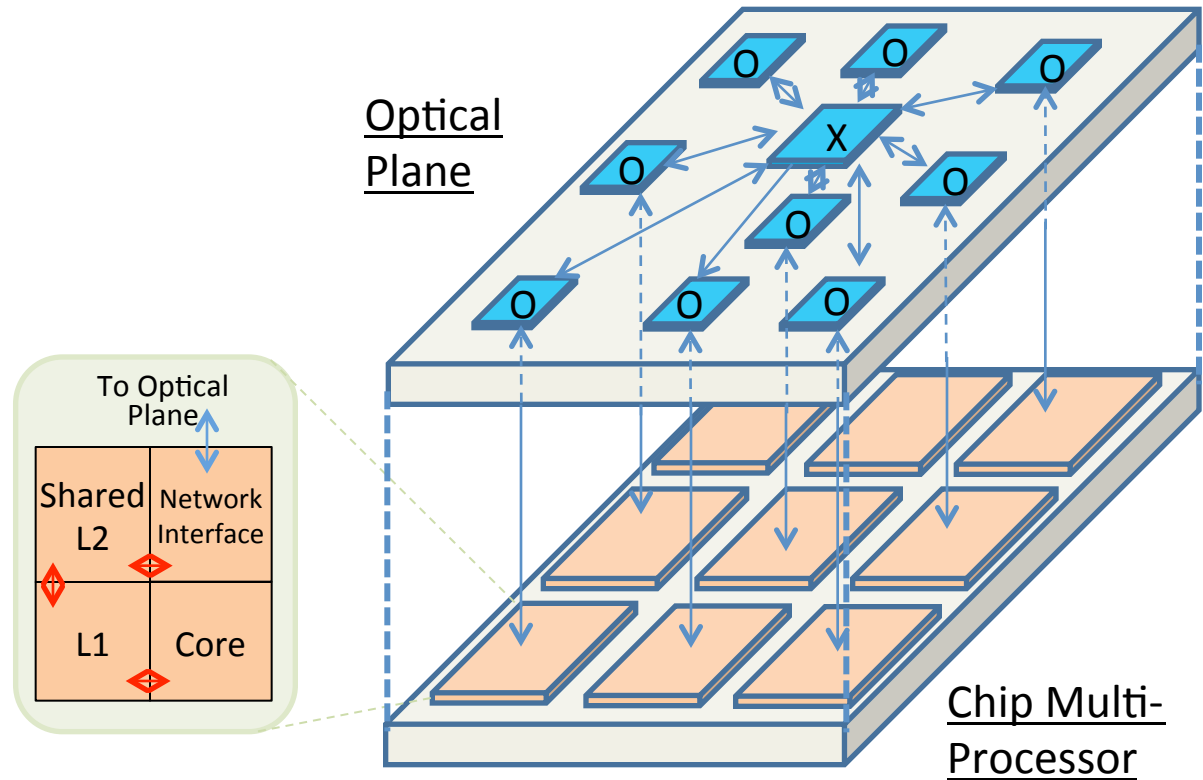
Low Latency Scheduling Algorithm for Shared Memory Communications over Optical Networks

Muhammad Ridwan Madarbux, Anouk Van Laer,
Philip M. Watts

Electronic and Electrical Engineering Department
University College London


MOTIVATION

- Scaling chip multi-processors (CMP) is increasing thermal issues
 - negative impact on performance
- Photonic NoCs have been shown to have lower power consumption



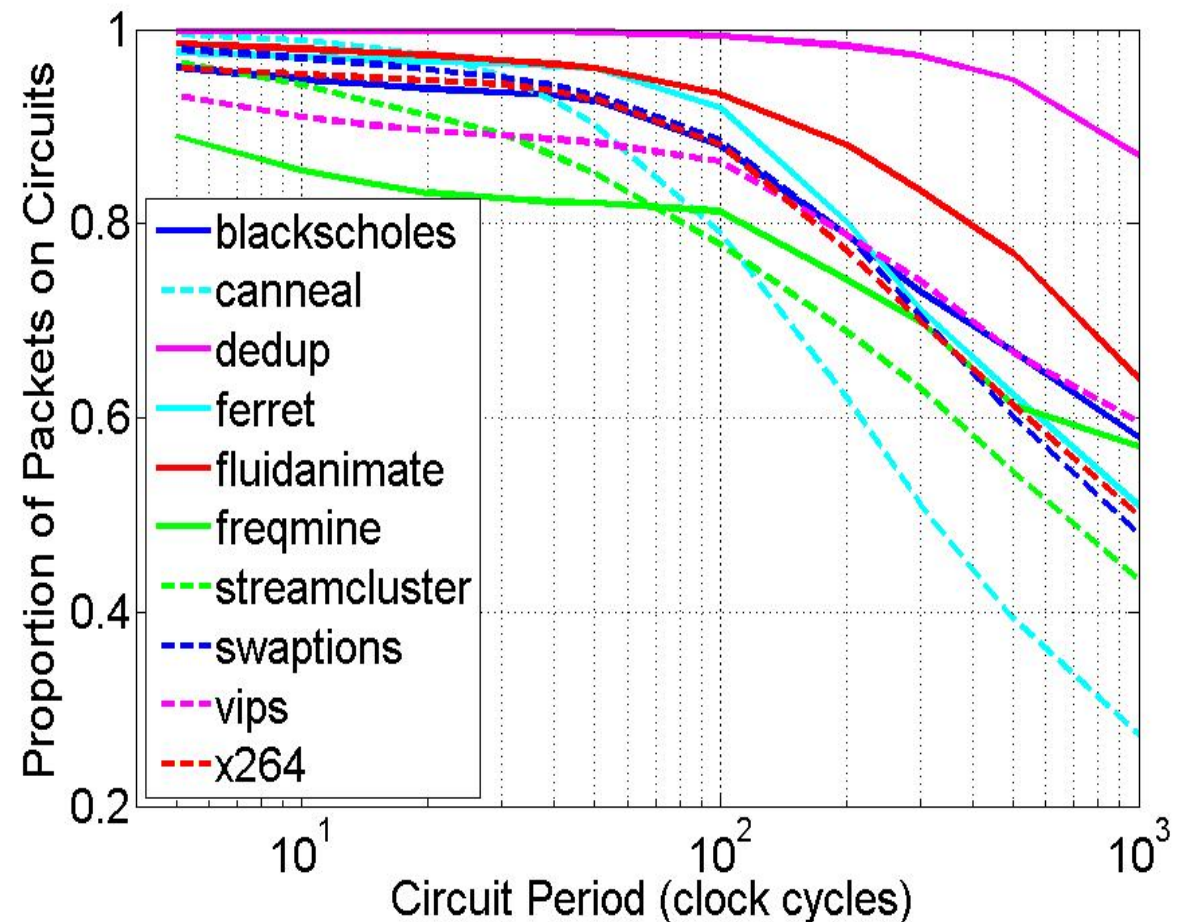
MOTIVATION

- Latency in switched photonic networks is dominated by scheduling
 - Request, arbitration and grant
- Scheduling generates a significant overhead in shared memory systems
 - 8B control messages
 - 16-256B data messages



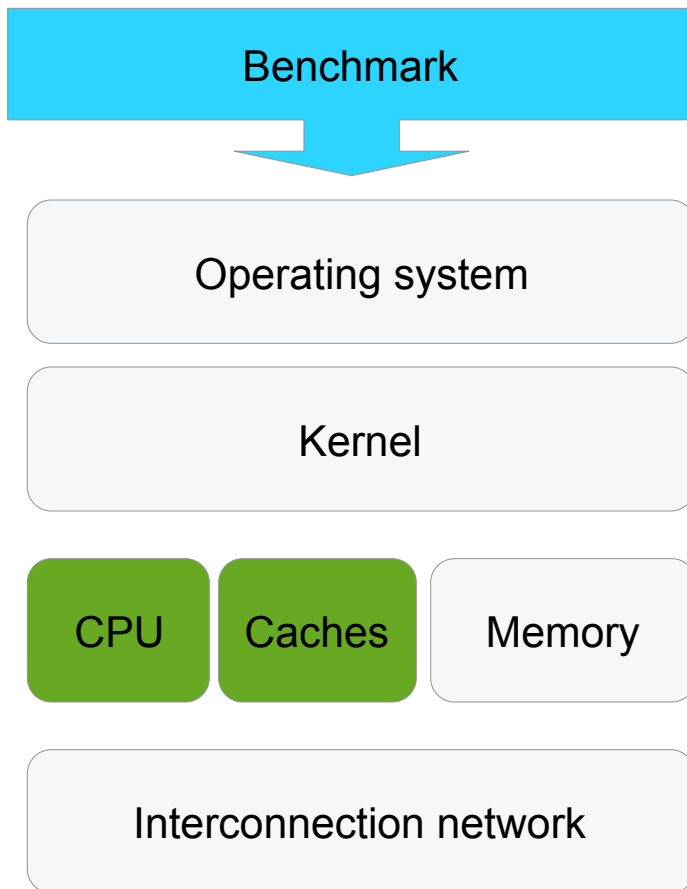
CIRCUIT SWITCHING FOR SHARED MEMORY

- Reduce scheduling latency by circuit switching for large flows of data between two cores
- Maximise proportion of messages on circuits to minimise contention and improve latency
 - Reduces for circuit periods >100 clock cycles
- Backup network required



This work proposes a new scheduling algorithm which intelligently uses information from the cache hierarchy to setup optical circuits

SIMULATION PARAMETERS

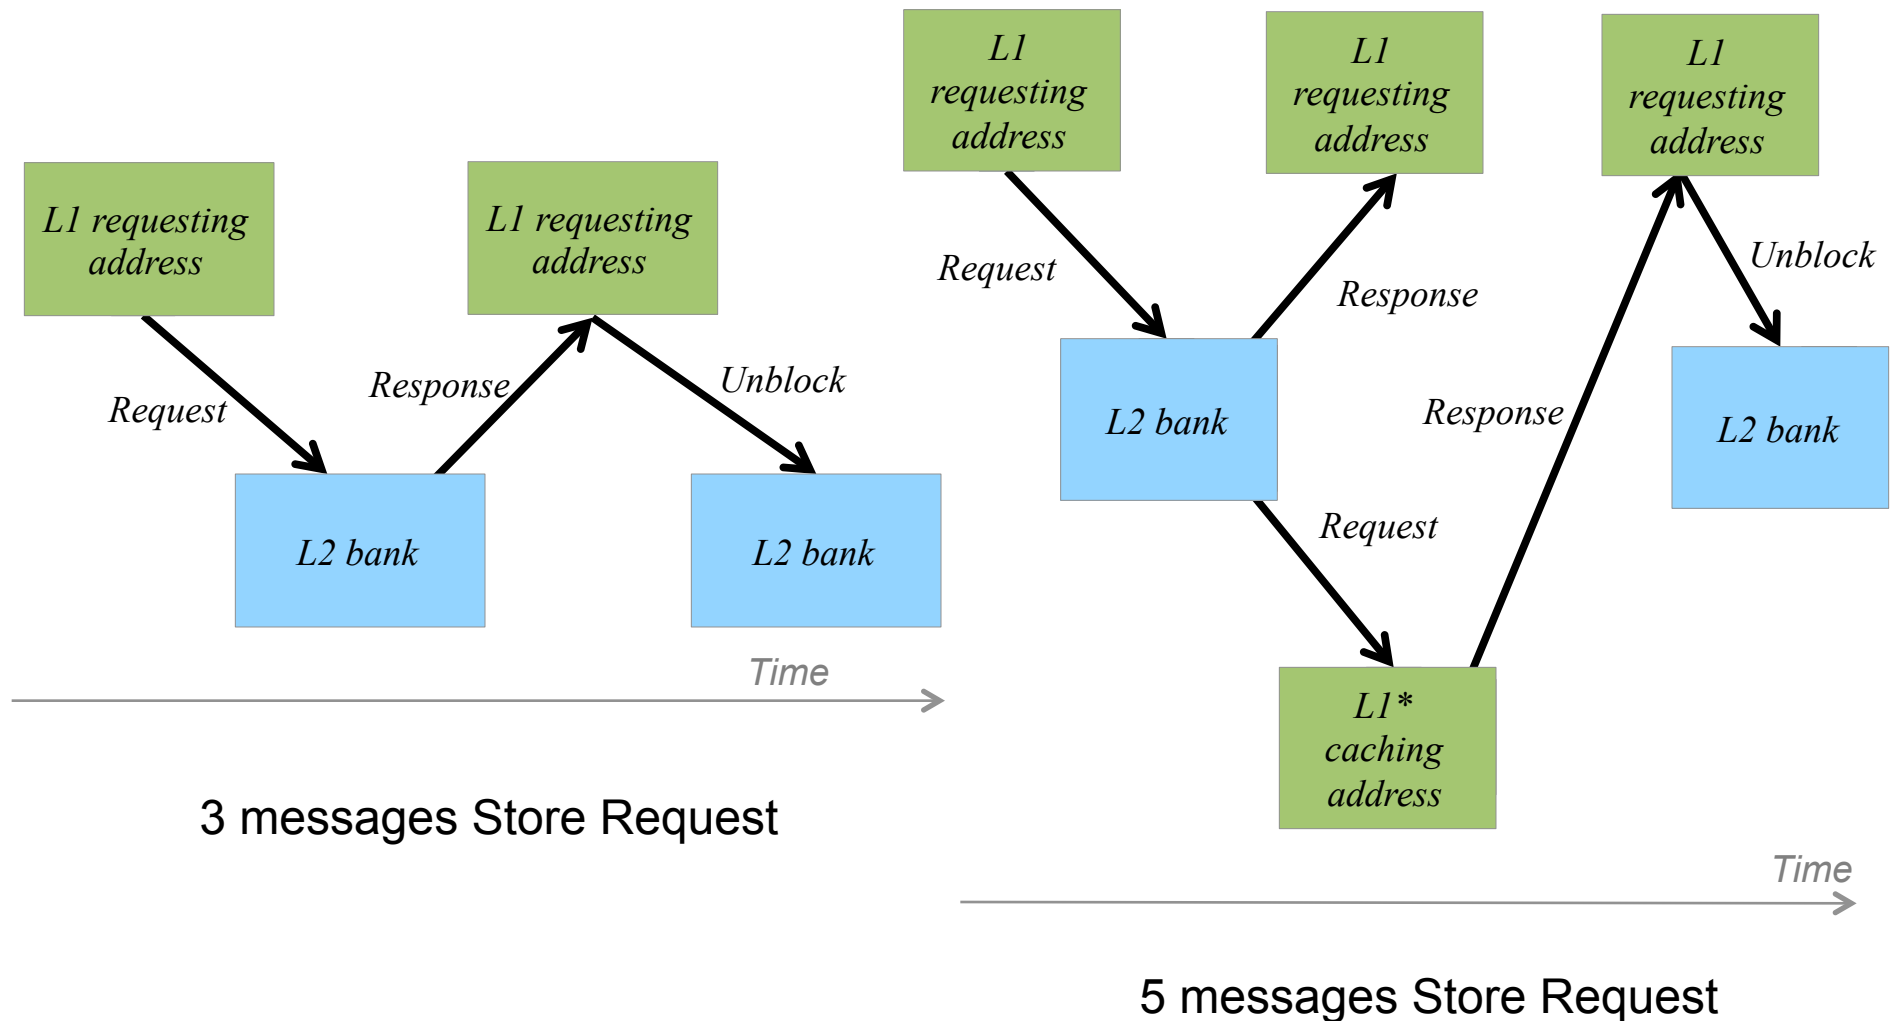


PARSEC benchmark suites

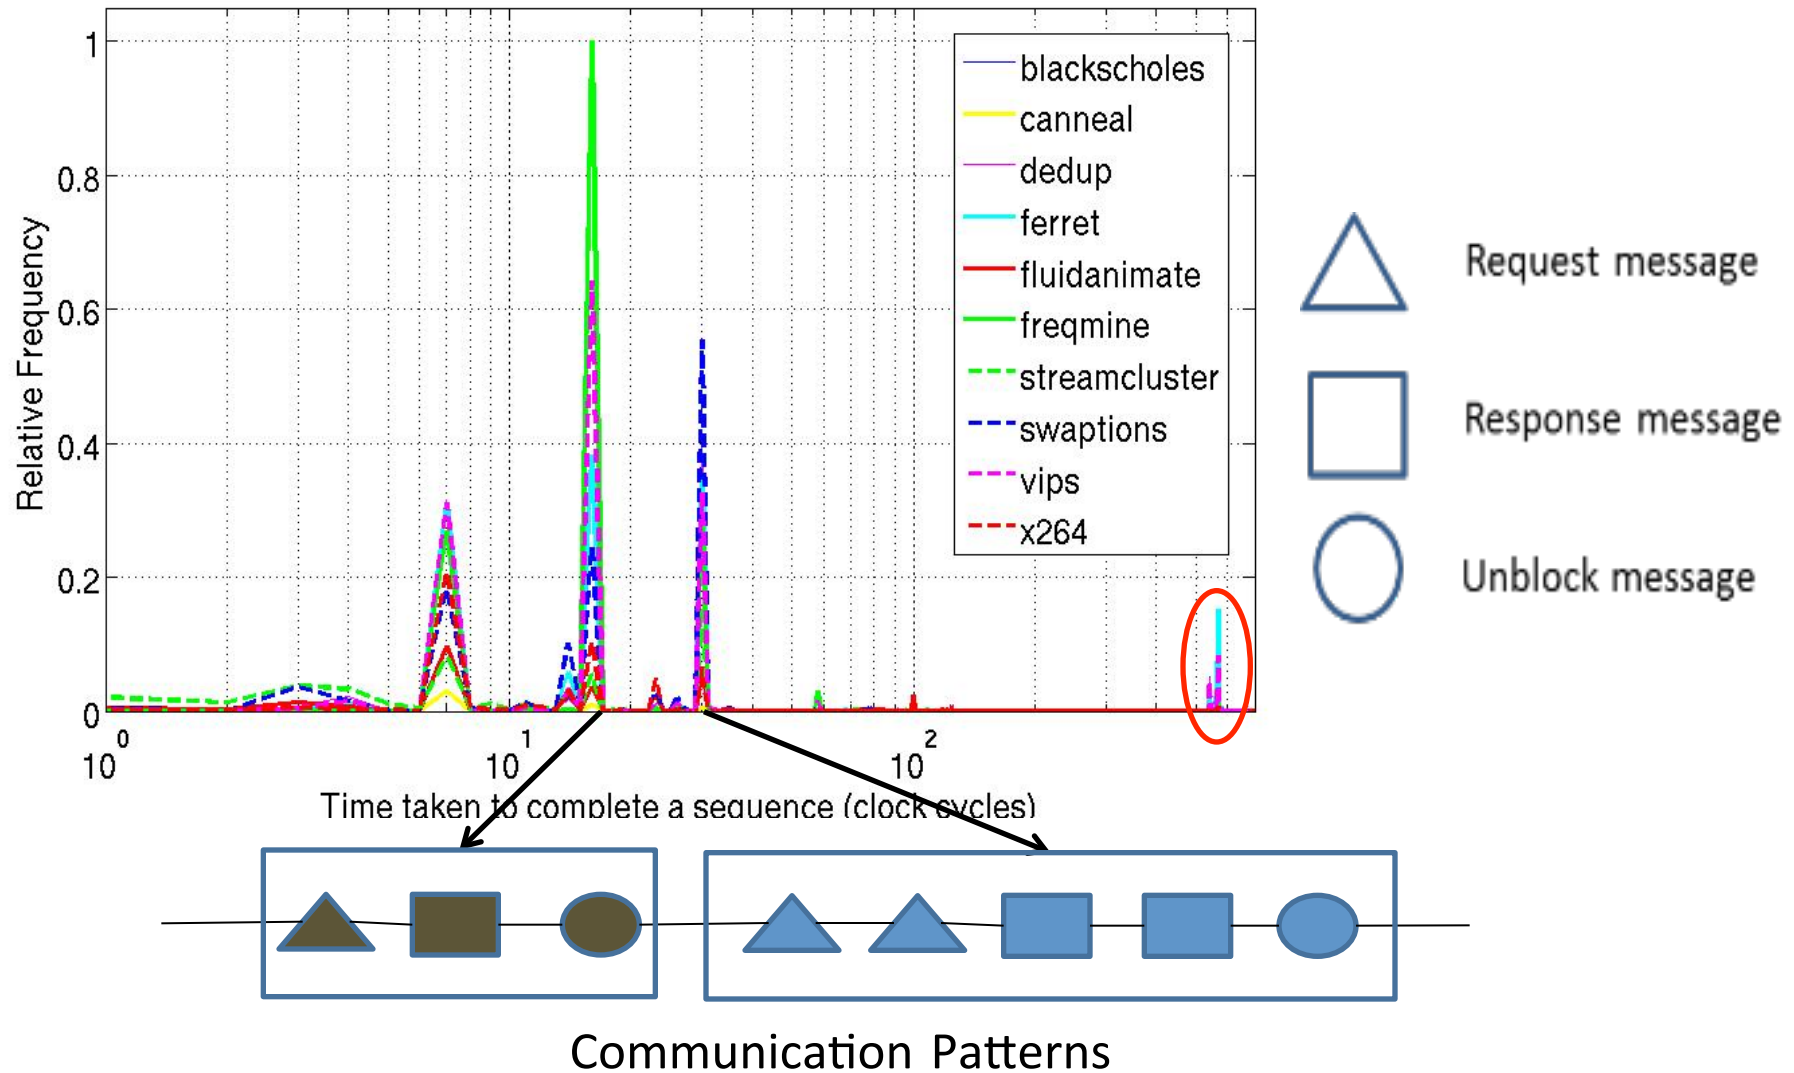
- Blackscholes – Calculate price of options
- Canneal – Optimization of routing in chip design
- Dedup – Compression using data deduplication
- Ferret – Content similarity search server
- Fluidanimate – Simulates fluid for animations
- Freqmine – Data mining
- Streamcluster – Online clustering of input streams
- Swaptions – Calculates price of swaptions (MC)
- Vips – Image processing
- X264 – Video encoder

- 32 cores @ 1.2 GHz
- Private L1 = 16 kB
- Shared L2 = 1 MB in total (distributed among the tiles)
- Cacheline size = 64B
- Cache coherence protocol = MESI

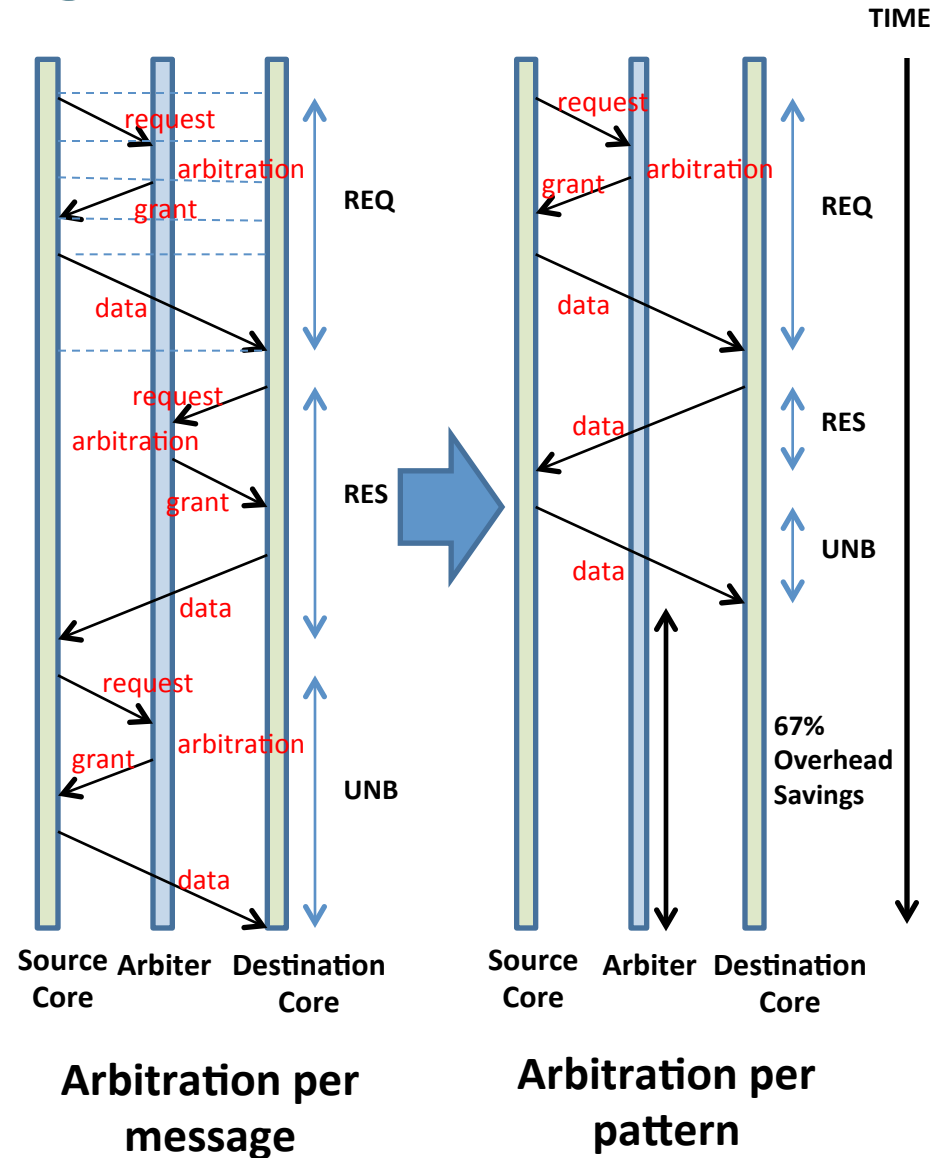
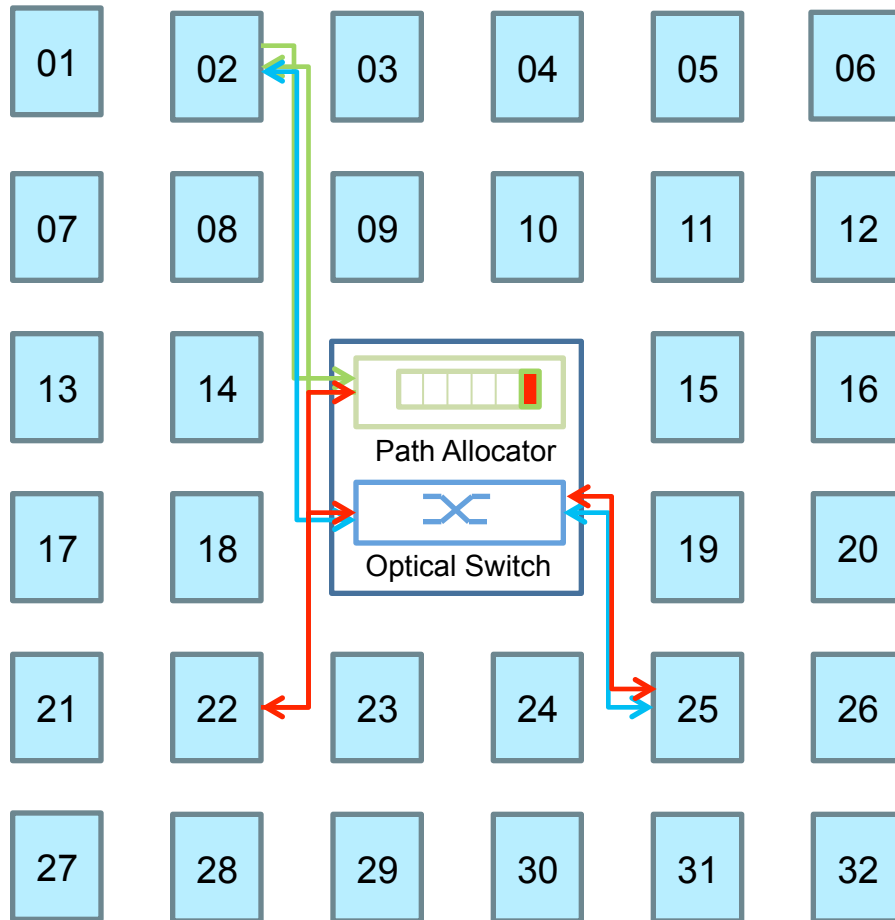
COMMON COMMUNICATION PATTERNS IN THE MESI COHERENCE PROTOCOL



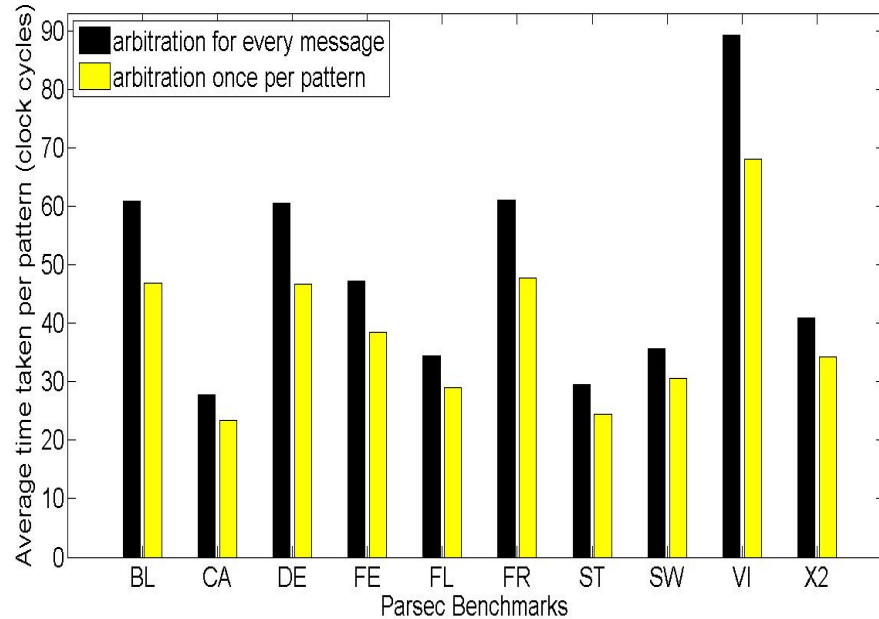
CACHE COHERENCE PROTOCOL BASED COMMUNICATION PATTERNS




LATENCY IMPROVEMENTS



LATENCY IMPROVEMENTS



Time taken per pattern



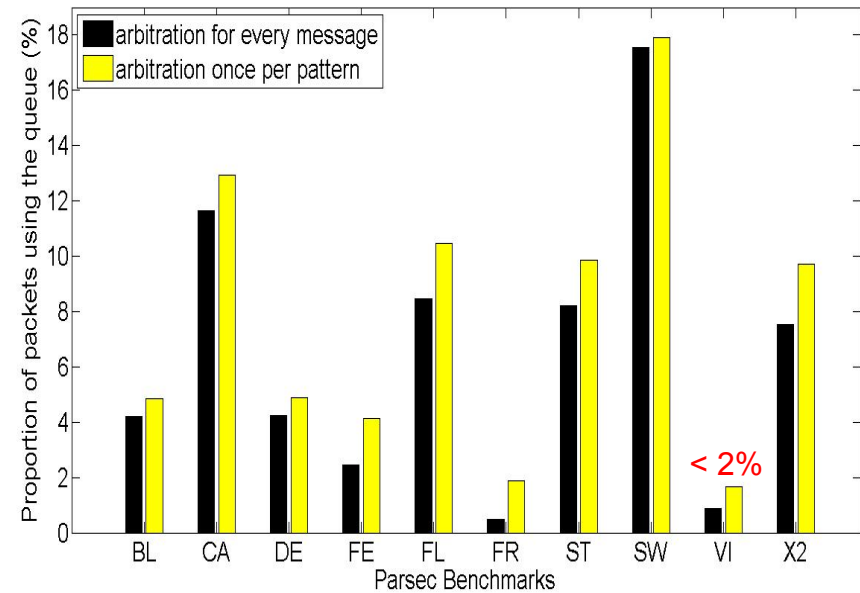
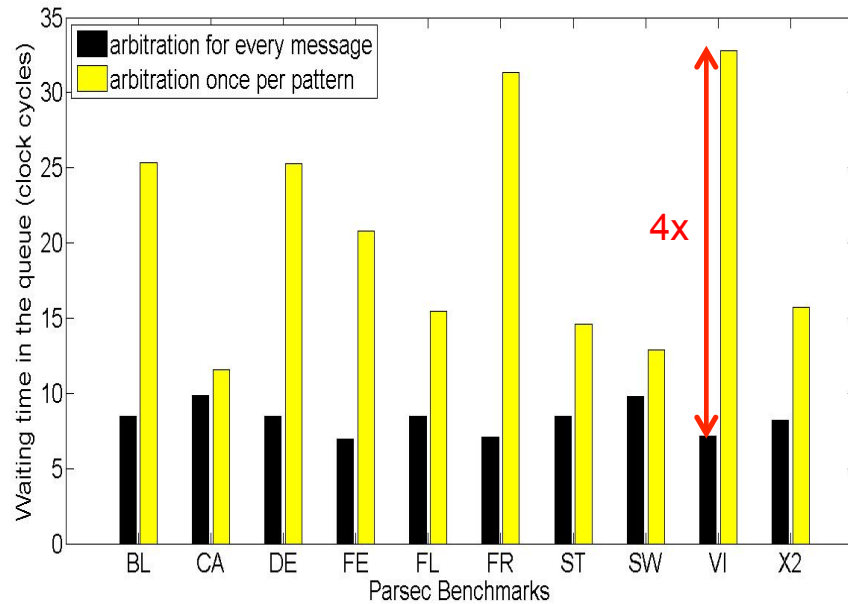
Head latency per message

BL – Blackscholes
 CA – Canneal
 DE – Dedup
 FE - Ferret

FL – Fluidanimate
 FR – Freqmine
 ST – Streamcluster

SW – Swaptions
 VI – Vips
 X2 – X264

EFFECT OF CONTENTION



BL – Blackscholes
 CA – Canneal
 DE – Dedup
 FE - Ferret

FL – Fluidanimate
 FR – Freqmine
 ST – Streamcluster

SW – Swaptions
 VI – Vips
 X2 – X264

CONCLUSION

- We have proposed an algorithm that intelligently uses information from the cache hierarchy to setup optical paths
- The algorithm provides significant latency reductions of up to 70.6% for Swaptions
- Results shown are for on-chip networks. Larger networks with longer time of flight will benefit more from this algorithm
 - Examples: Multiple socket servers or rack-scale networks

FUTURE WORK

- Considering the implications of this algorithm on the energy consumption of the allocator and control circuits
- Looking at the performance of the algorithm for heavier traffic loads considering that PARSEC benchmarks lightly load the network
- Measuring the performance improvements in full system gem5 simulation

THANK YOU FOR YOUR ATTENTION