The Sensitivity of Communication Mechanisms to Bandwidth and Latency

Frederic T. Chong, Rajeev Barua,
Fredrik Dahlgrenz, John D. Kubiatowicz, and
Anant Agarwal

presented by: Scott Beamer

Overview

- Comparison of message passing and shared memory with respect to the latency and bandwidth of the network
- Done on hardware (Alewife) with real apps
- Also considers optimizations such as prefetching, polling, and bulk transfers

Supported options

- Message Passing
  - with interrupts - active message style
  - with polling - app checks queue
  - with bulk transfer - DMA to network
- Shared Memory
  - without prefetching - normal
  - with prefetching - done in background

Performance Models

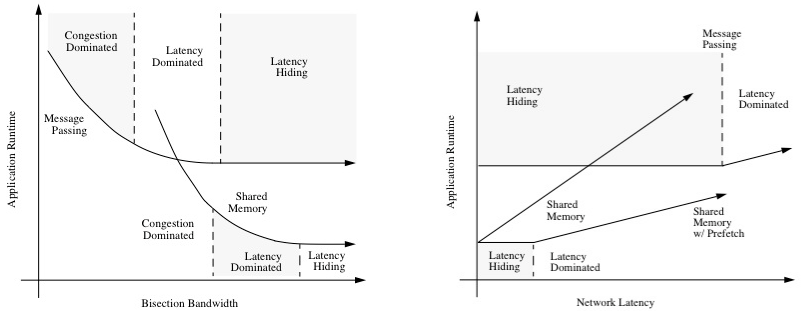

EM3D

- Models properties of electromagnetic waves through 3d objects

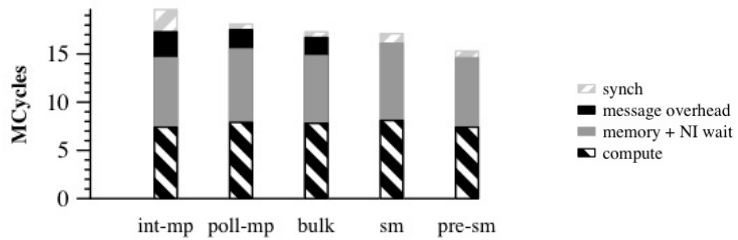

UNSTRUC

- Simulates fluid flows over unstructured meshes in 3D

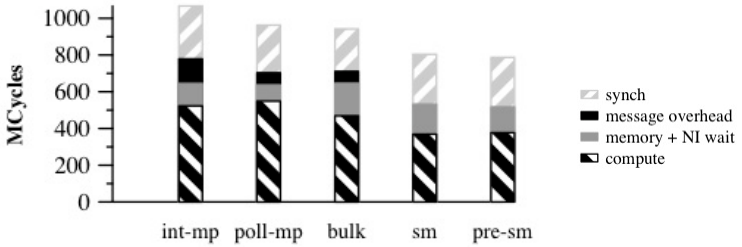

ICCG

- General iterative sparse matrix solver using conjugate gradients with preconditioning

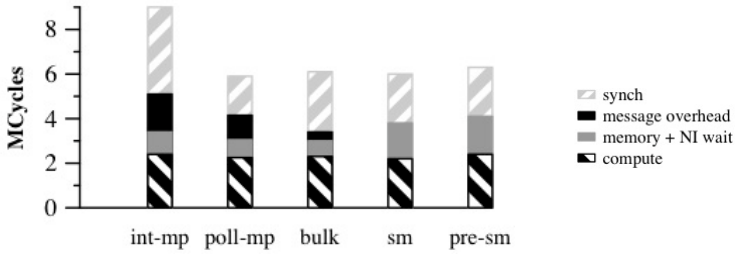

MOLDYN

- Computes interactions between molecules within a cut-off distance

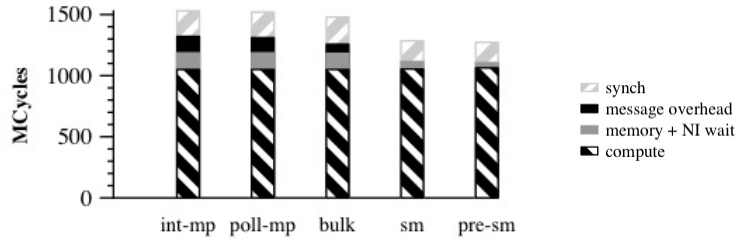

Communication Volume

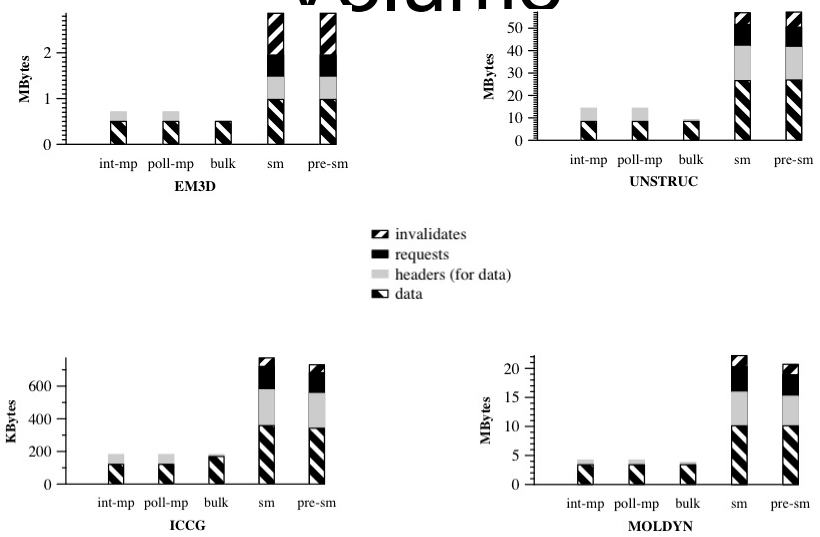

Bandwidth Experiment

- Controlled amount of I/O traffic across 2D mesh to reduce available bandwidth
- I/O message size didn't affect MP, bigger messages improved SM's performance

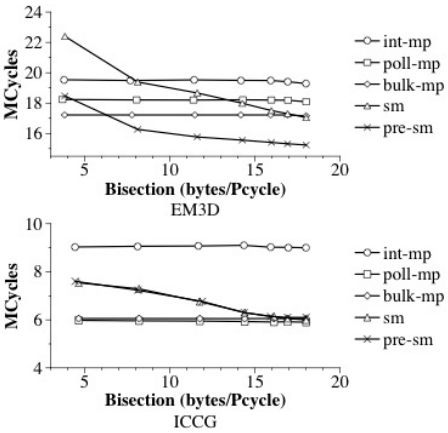

Latency Experiment

- To decrease latency, slowed down processor clock
- To increase latency, context switches to a delay loop

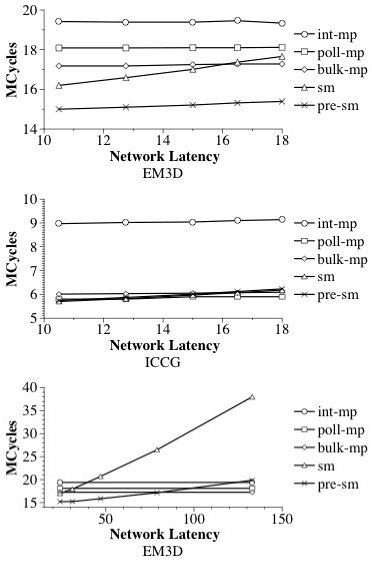

Conclusion

- Shared memory required less coding effort
- Message passing less sensitive to increased latency or reduced bandwidth
- Shared memory can provide good performance if there is enough bandwidth