Comparison of 10GigE and InfiniBand Interconnects

Colby Boyer (UC – Berkeley)

Joseph Lopez (NMHU)

Kimberly Jo McCormick (UNM)

Student Symposium

Los Alamos National Laboratory

August 2-3, 2007

Los Alamos NATIONAL LABORATORY EST. 1943

Goals

- 10GigE vs InfiniBand 4x
- Comparison of Results
- Cost Analysis
- Future Research

Key Issues:

- High Performance Demands
  - Low Latency & High Bandwidth
- Scalability
  - Integration into Existing Networks
  - Maintenance
- Cost Effectiveness

Equipment & Benchmarks

- Equipment
  - InfiniBand 4x
  - 10 GigE

    Fujitsu switch

    NetEffect NIC

- Benchmarks:
  - OSU
    - MPICH2
    - MVAPICH2

Latency Results

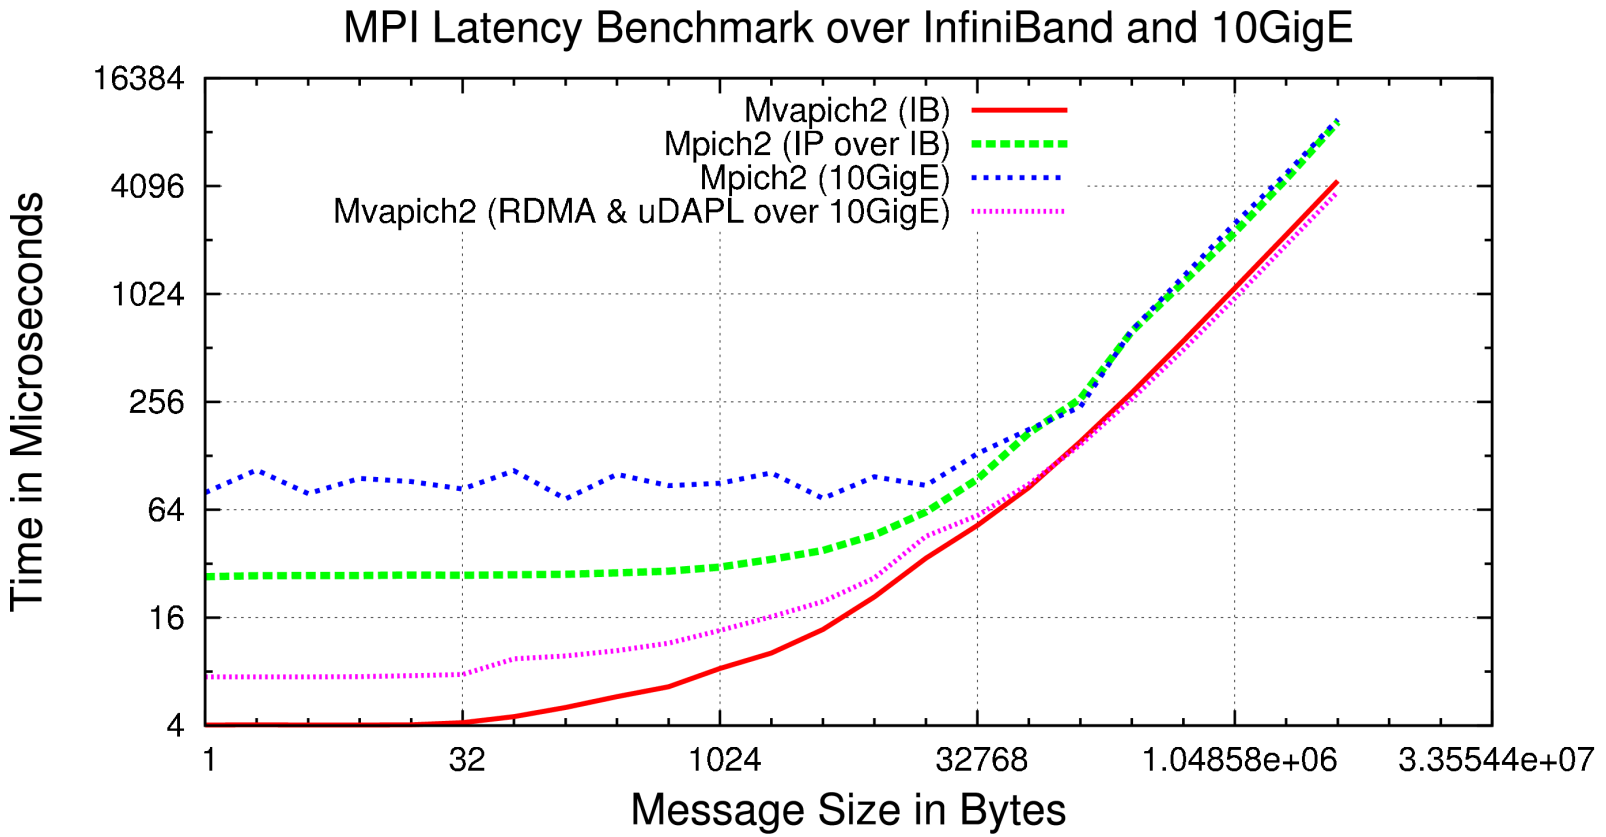

Bandwidth Results

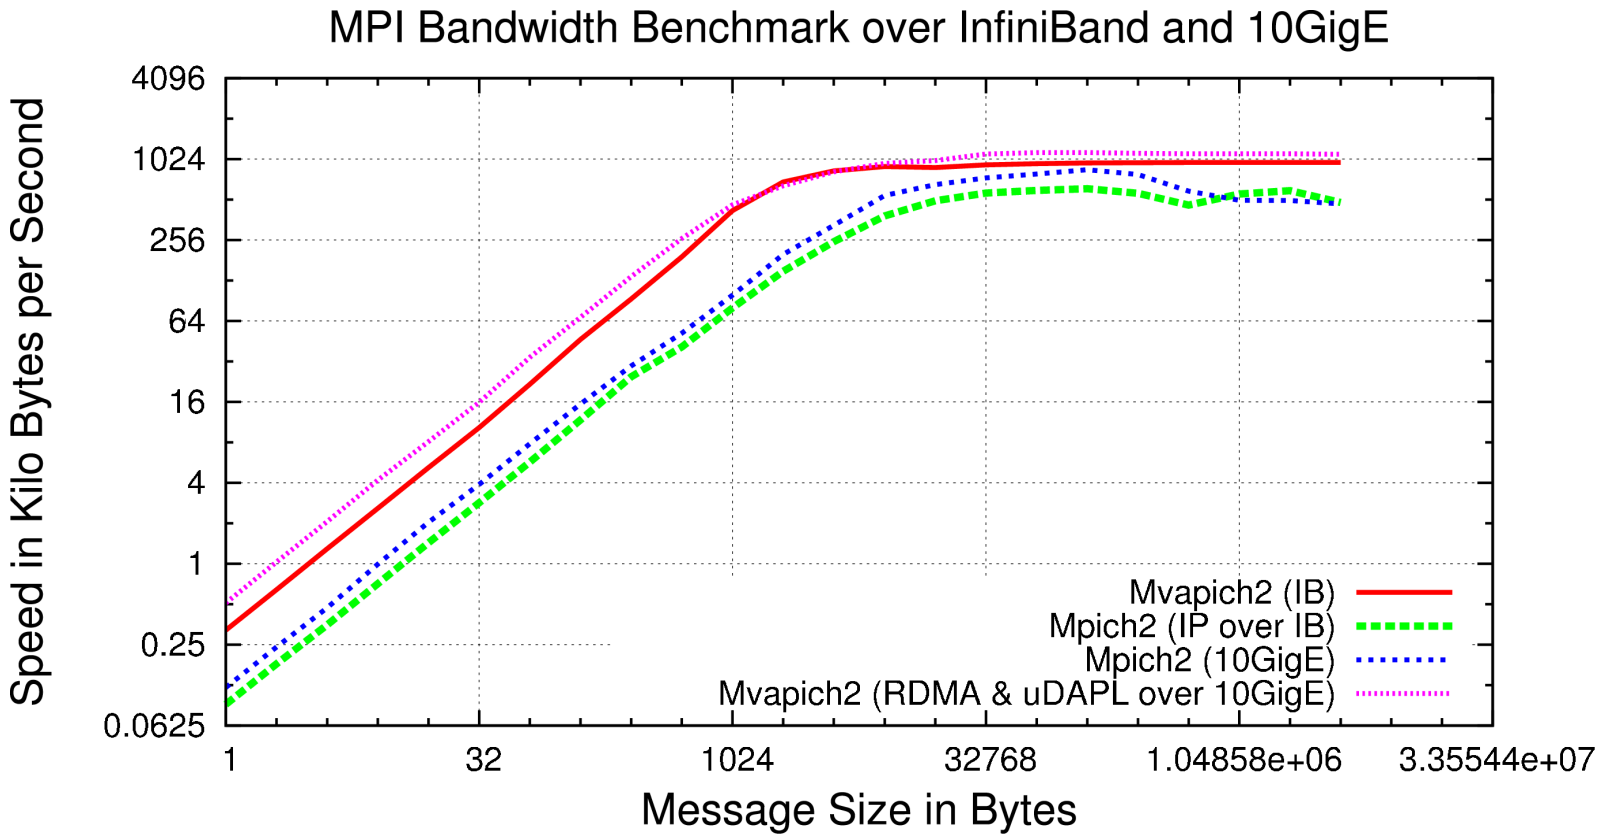

Bidirectional

Bandwidth Results

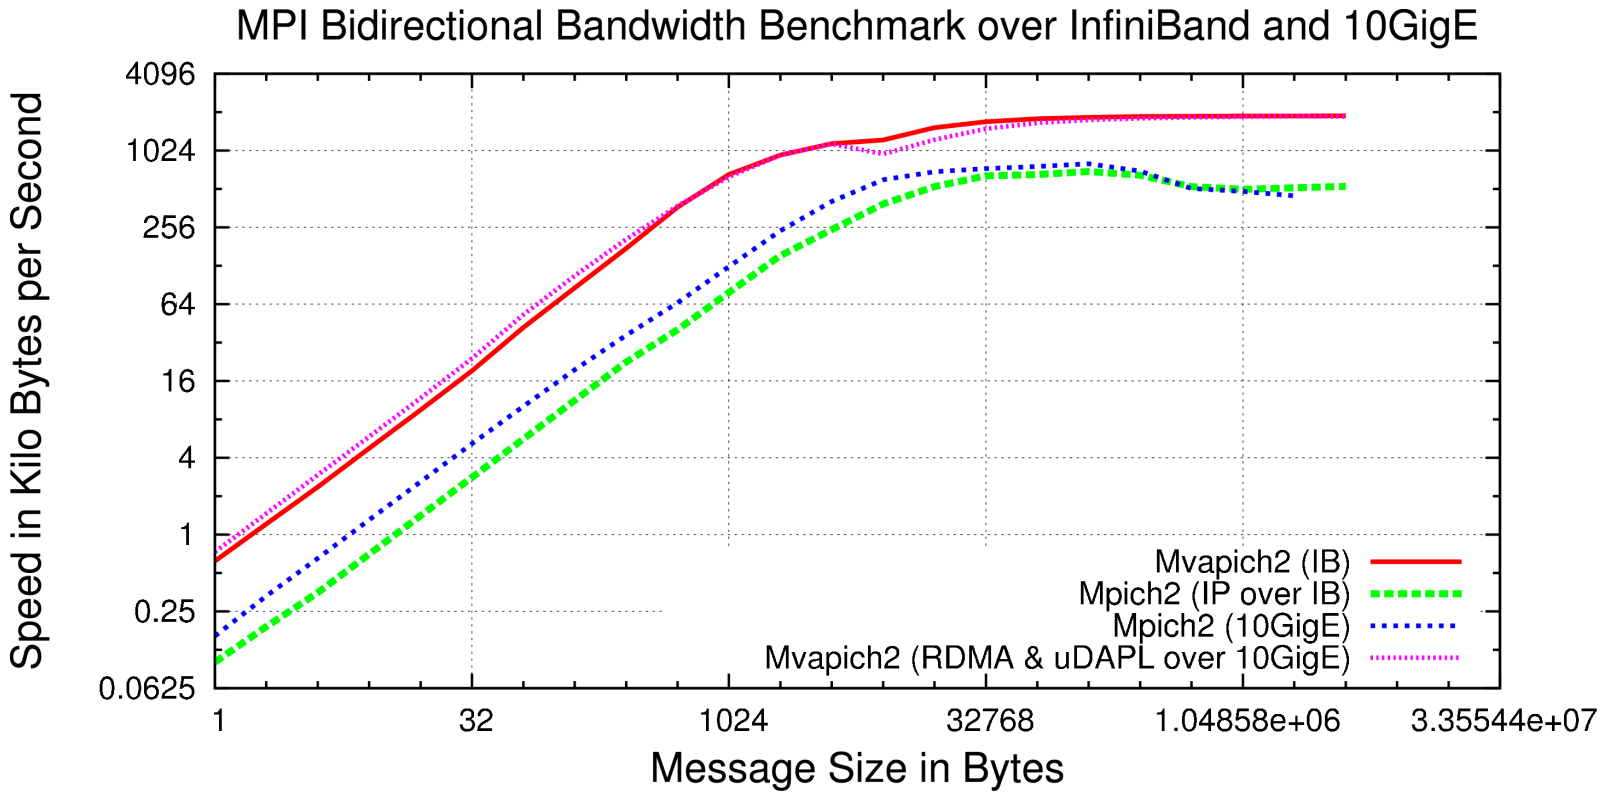

Cost Analysis

For a small cluster using copper connections

- 24 Port Infiniband switch ~$6100 plus InfiniBand cards at ~$500 - $700/ea
  - Per port cost ~$750-$950
- 24 Port 10GigE switch ~$13,000.00 plus cards at ~$700.00/ea (fiber optic cards available for ~ $1200-$1900/ea)
  - Per Port cost ~ $1250

Conclusions

Performance

- InfiniBand achieved lower latency
- 10GigE achieved higher peak bandwidth
  - RDMA(via uDAPL) is key

Cost

- InfiniBand has a clear advantage over 10GigE

**Good point for 10GigE is that it can be easily integrated into existing TCP/IP networks

Future Research

- Test the use of uDAPL with large codes
  - Test the effect that the latency difference has on production programs
- Increased Scale of Testing
  - Use of multiple switches (Larger Clusters)
  - Integrate with external network systems
- Continue to Analyze Cost
  - Prices Will Continue to Change

References

Mentors:

Todd Bowman (CTN-5)

HB Chen (HPC-5)

Andree Jacobson (UNM)

Robert Martinez (CTN-5)