An Update on CORBA Performance for HPEC Algorithms

Bill Beckwith

Objective Interface Systems, Inc.

bill.beckwith@ois.com

http://www.ois.com

September 2nd, 2003

Elements of Performance

- **Simplified (but accurate) execution model:**
  - Latency
    - End-to-end time to transfer one byte
  - Per Byte
    - Extra end-to-end time to transfer each additional byte
  - Total time
    - Latency + Per Byte * Bytes
- **Copies add to Per Byte time**
  - HPEC hardware transfer rates are competitive with local memcpy times (approx. one byte per clock cycle)
  - Result is that any copies kill throughput (but you knew that)

First Benchmarks: Zero Copy Affect on Windows

- **CPUs**
  - 2 GHz Pentium 4M laptop
  - 1 GHz Athlon desktop
  - (2 GHz P4M is 20% faster than 1 GHz Athlon)
- **Transports**
  - Shared memory on Windows
  - 100 Mb Ethernet

Memory Copy Performance

- **Number of Algorithms**
- **Performance varies depending on:**
  - Cache size
  - Cache line size
  - Bytes moved per operation
- **ORB uses most efficient copy algorithm we can discover**

Memory Copy Performance

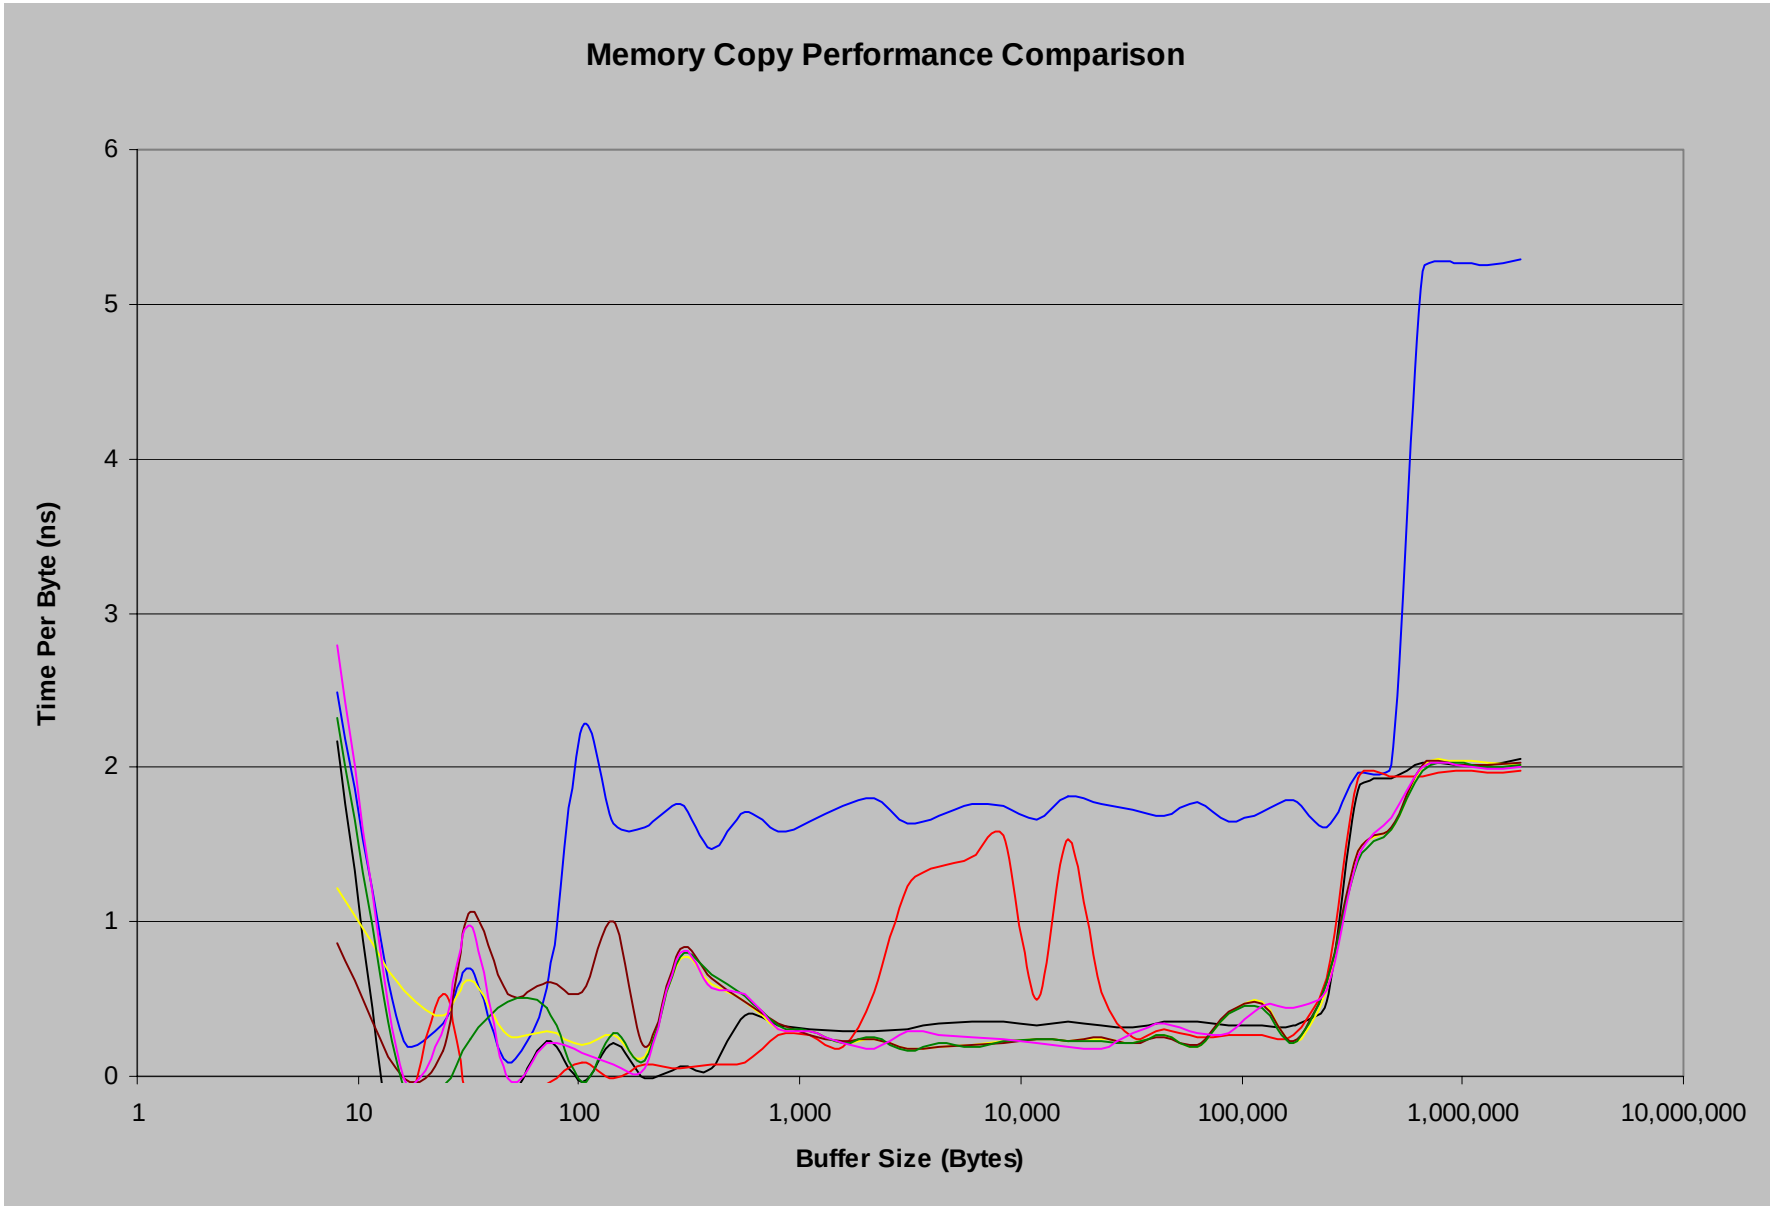

SHRMEM Latency

- **Reducing marshaling copies**
  - Decreases latency for large transfers
  - Increases latency for small transfers
- **Latency increase occurs because there are more system calls from the transport**
- **Scatter/Gather system calls would reduce the number of calls, and potentially the number of transport copies.**

SHRMEM Latency

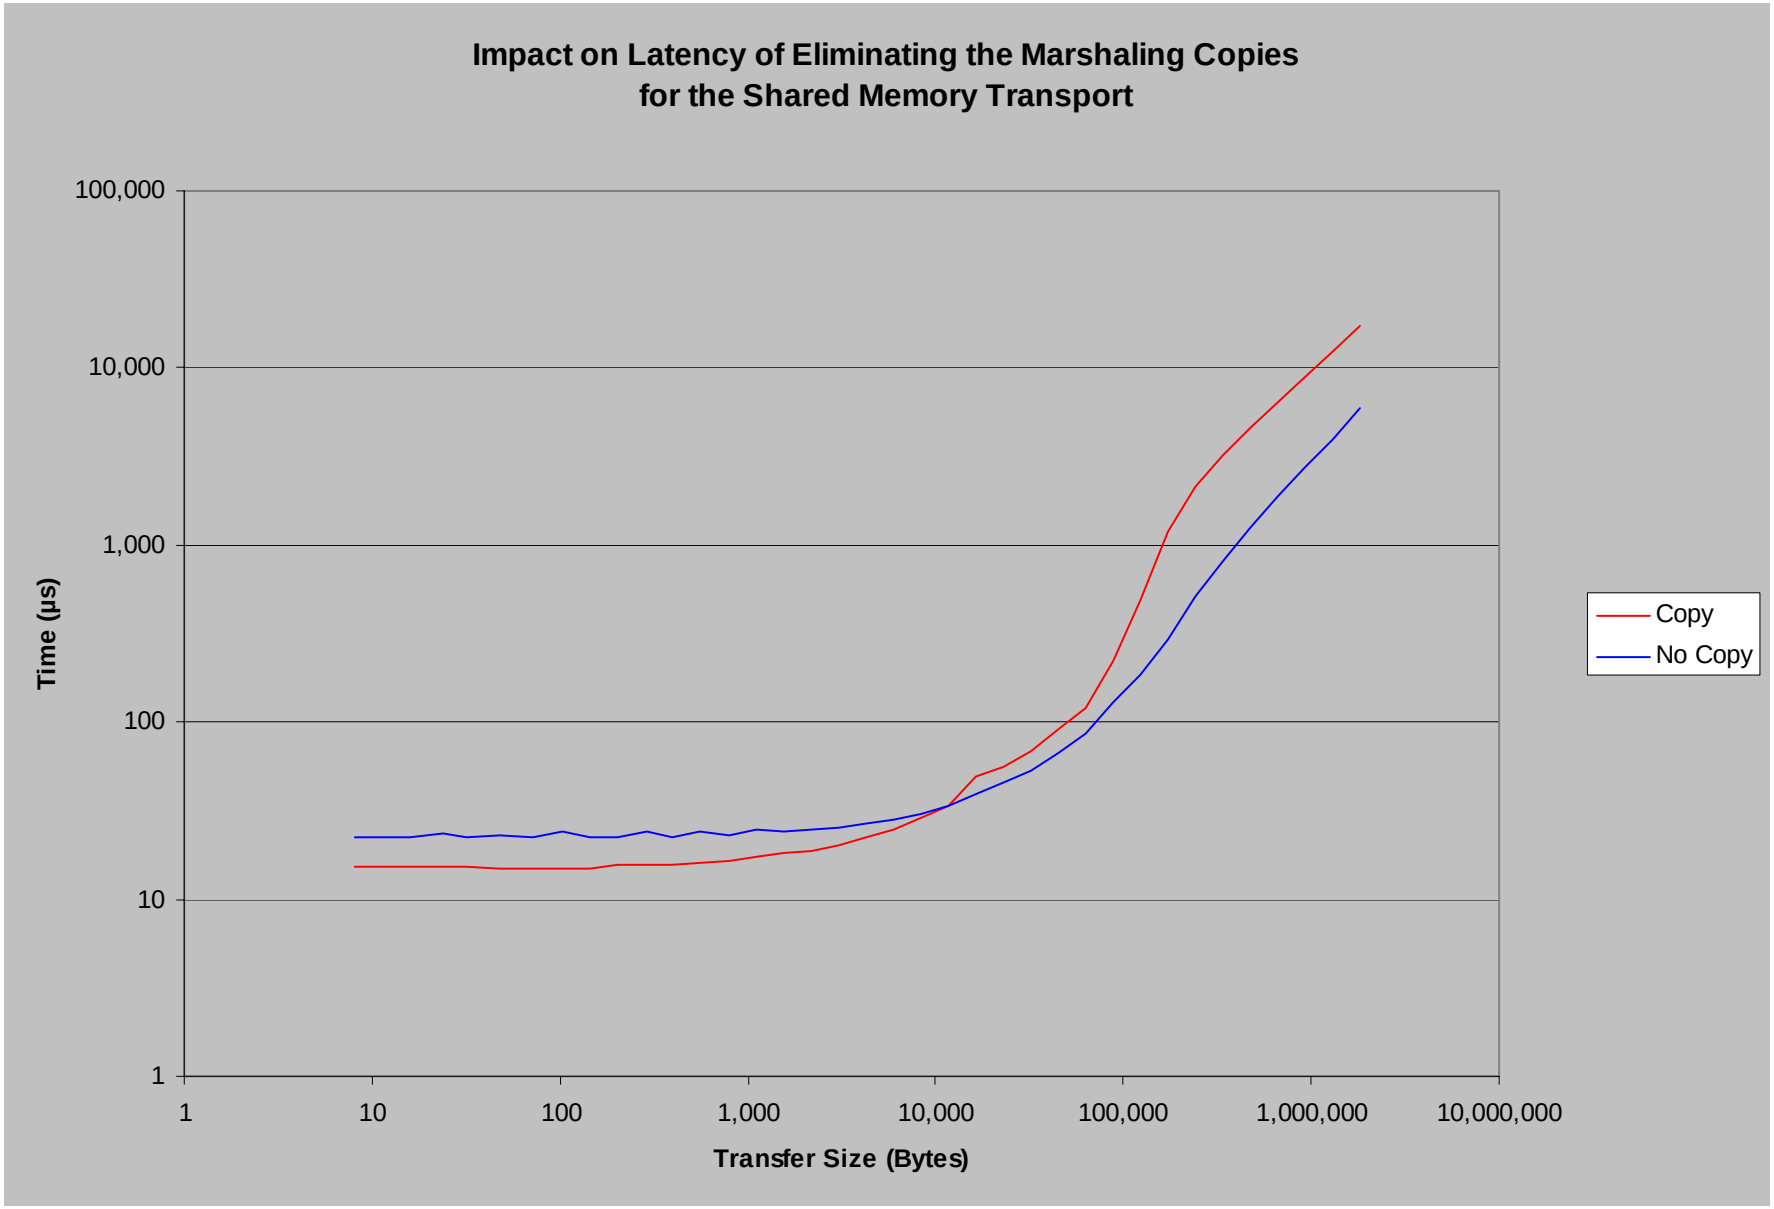

ORBexpress™

SHRMEM Bandwidth

- **Reducing marshaling copies**
  - Increases bandwidth for large transfers
  - Reduces bandwidth for small transfers
- **Bandwidth reduction occurs because there are more system calls from the transport**
- **Scatter/Gather system calls would reduce the number of calls, and potentially the number of transport copies.**

SHRMEM Bandwidth

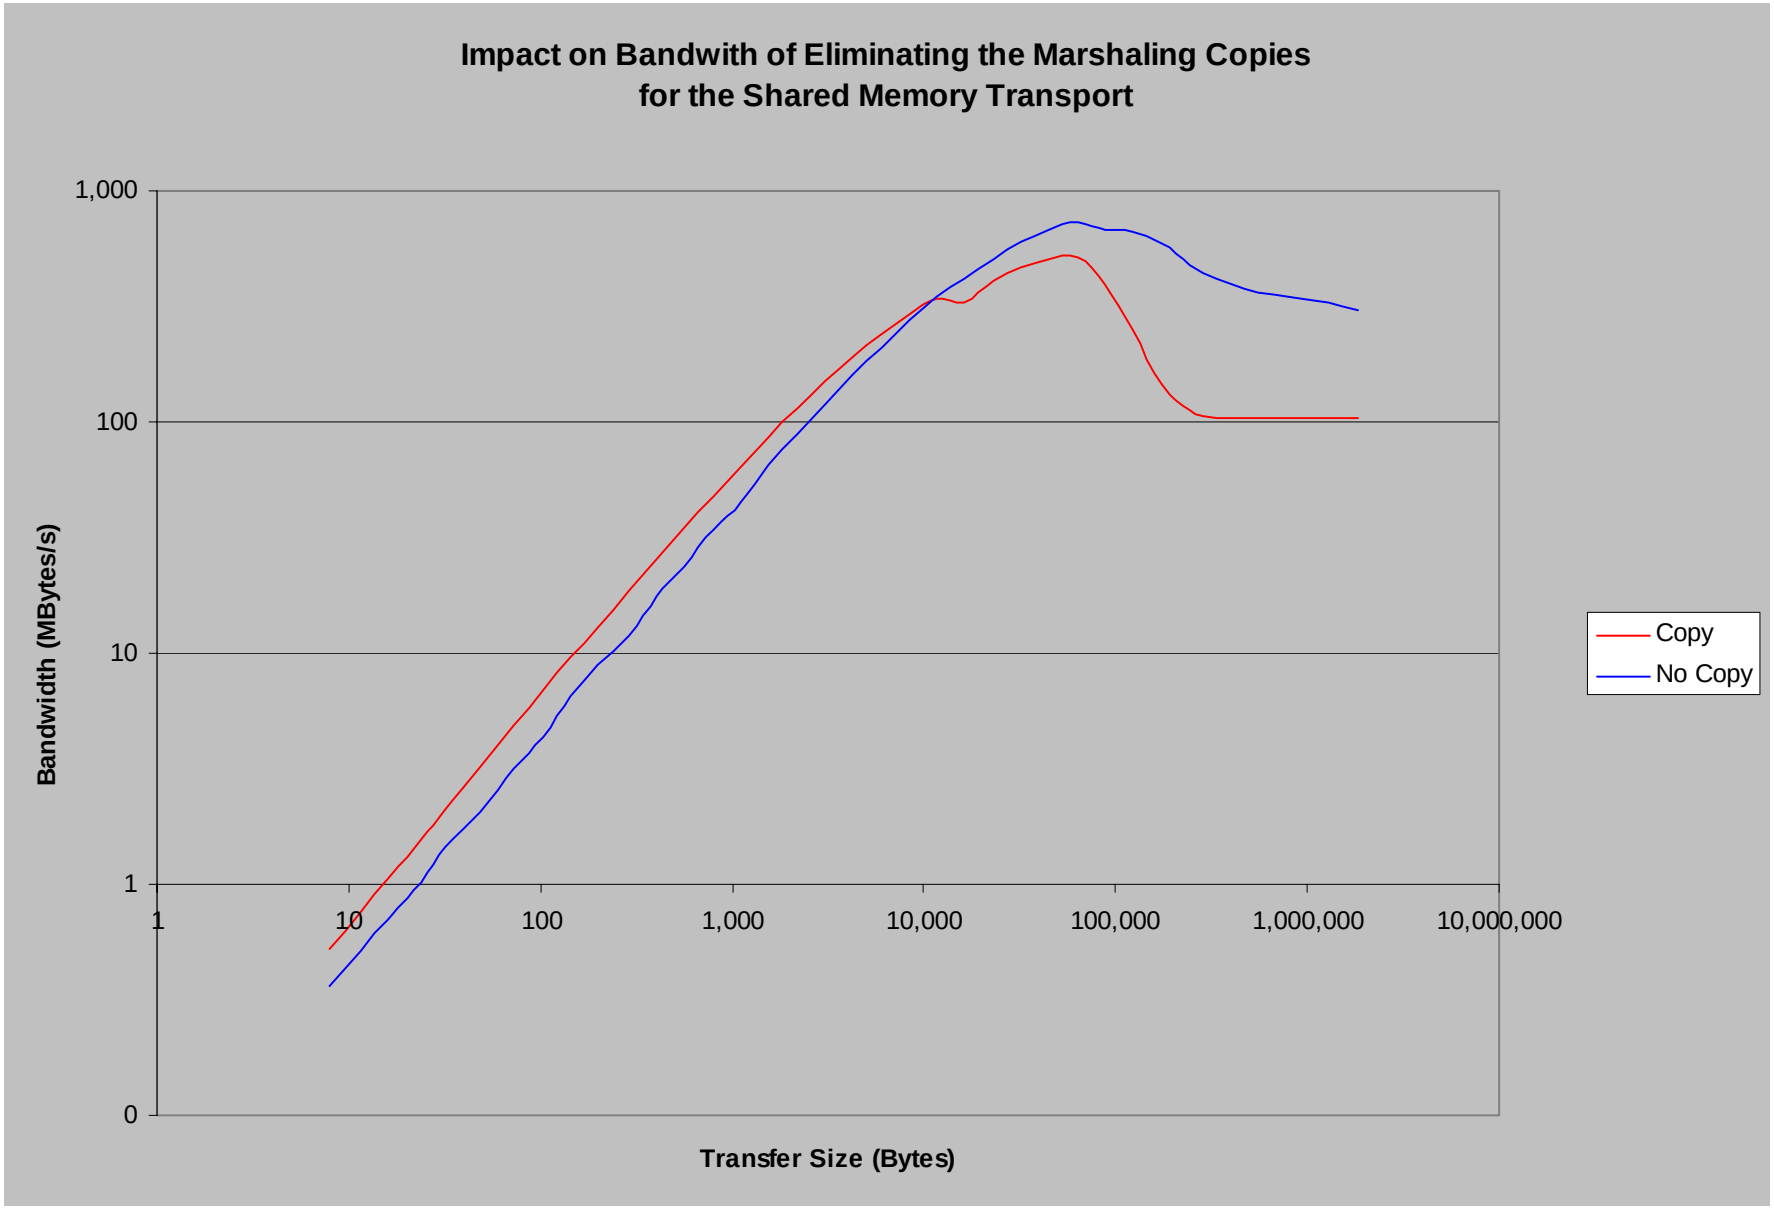

CPU Utilization

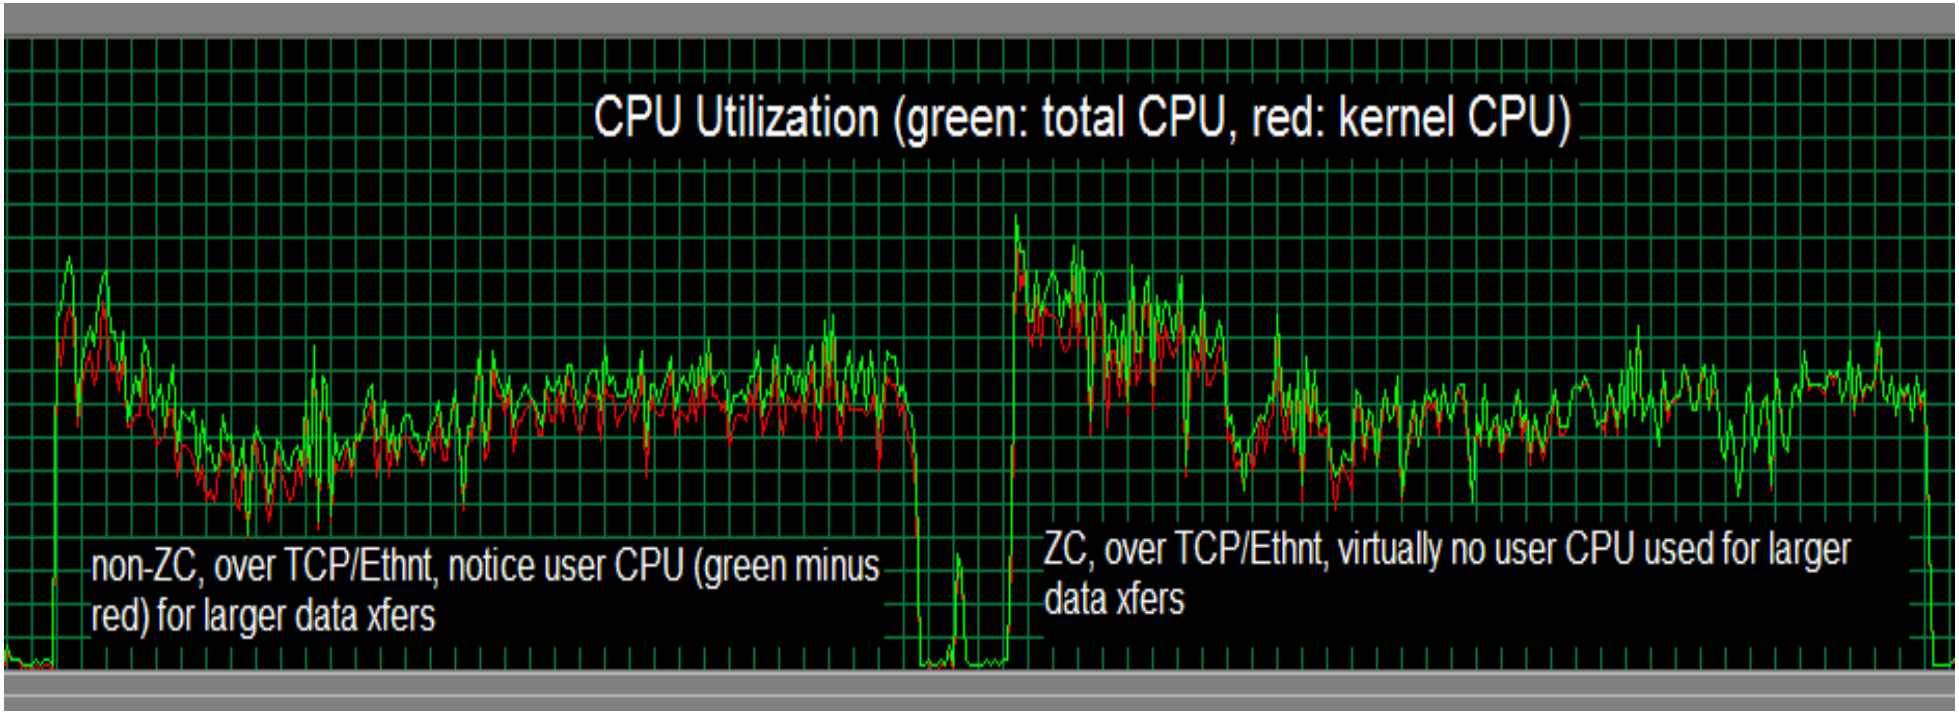

Network Utilization

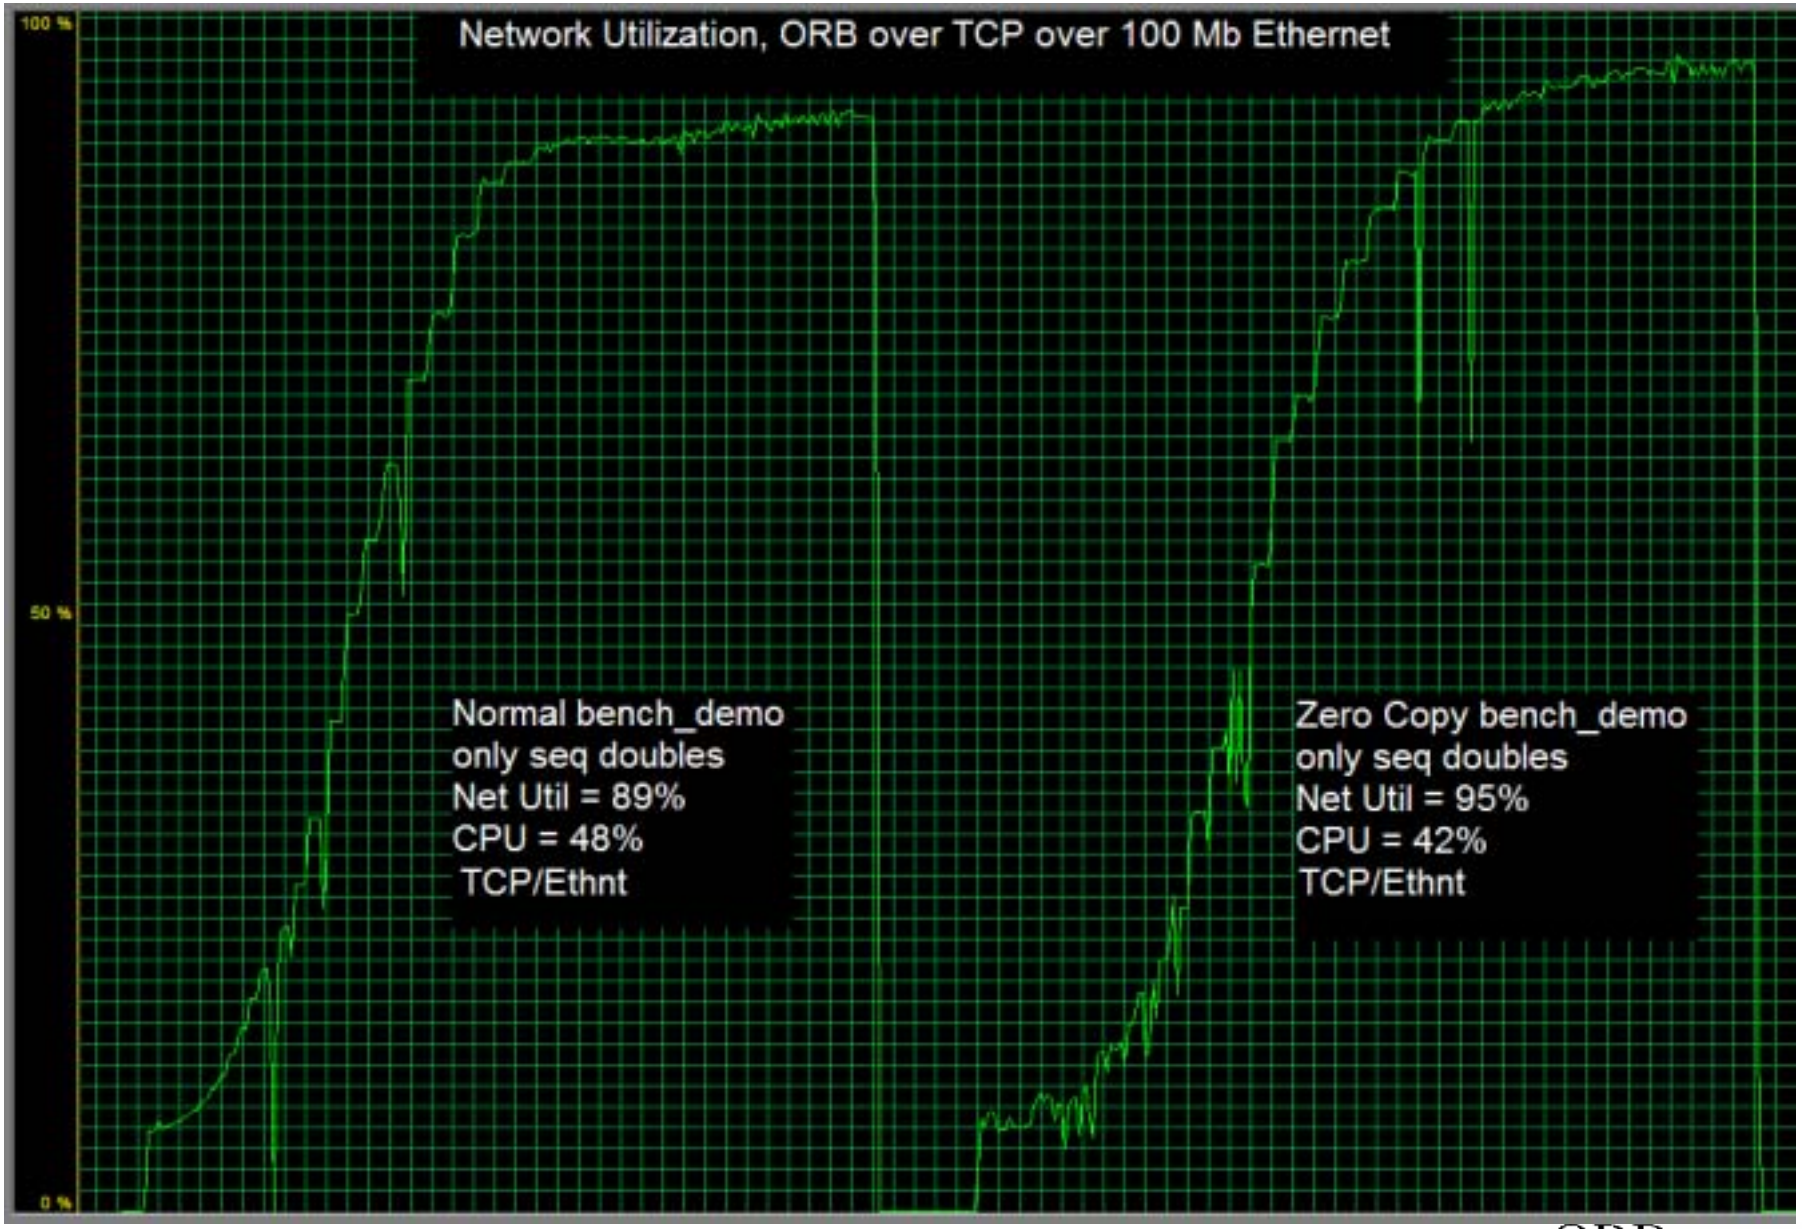

First Benchmarks: Zero Copy Affect on HPEC

- **Internal work-in-progress versions of ZC ORB**
  - Several suboptimal characteristics
    - Underlying transport
      - High latency
      - DMA transfers are 80K blocks
- **Mercury RACEway++**
  - VxWorks host
  - CE-to-CE communications

Comparing Copy Configurations

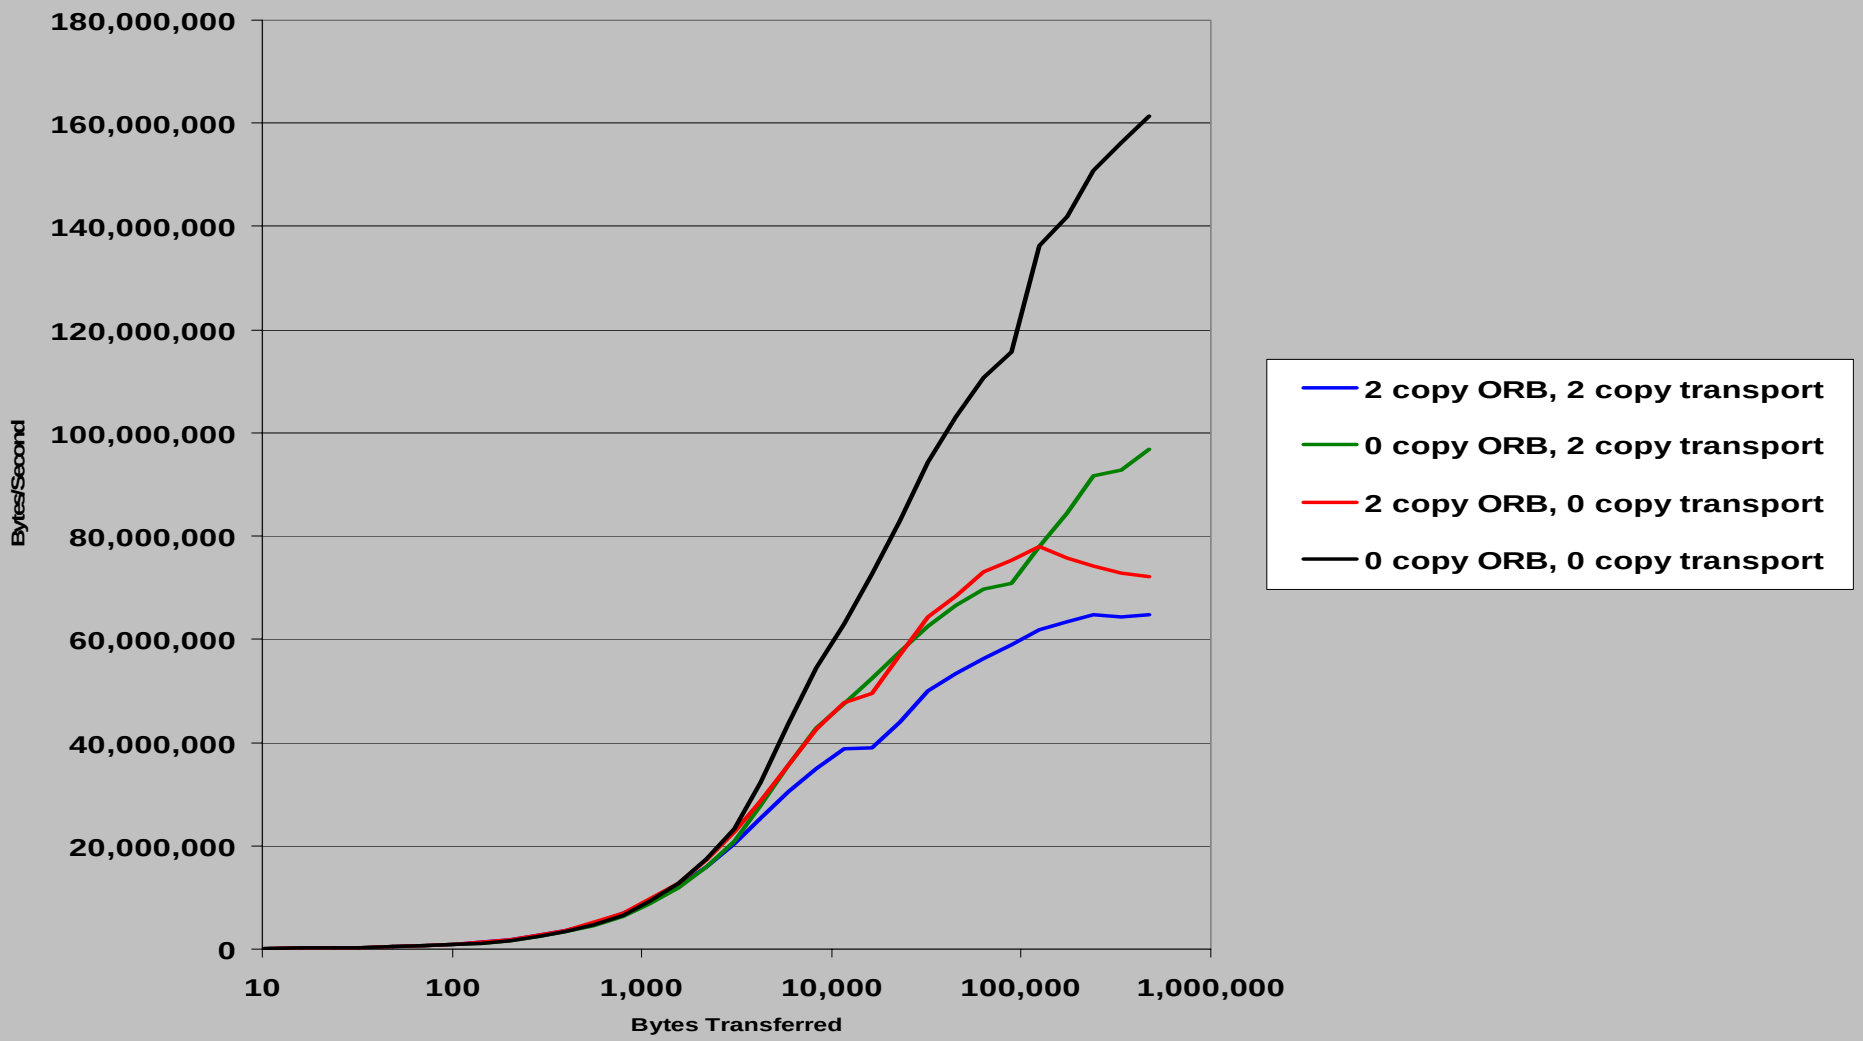

Work Left

- **Finalize Zero-Copy version of ORB*express***
- **Rewrite underlying transport, expectations:**
  - Better latency (> 10 usec)
  - More efficient use of DMA

Summary

- CORBA is progressing towards HPEC efficiency requirements
- Existing CORBA applications can take efficient advantage of HPEC hardware