HP Accelerator Collaboration Meeting

May 22, 2008 - Singapore

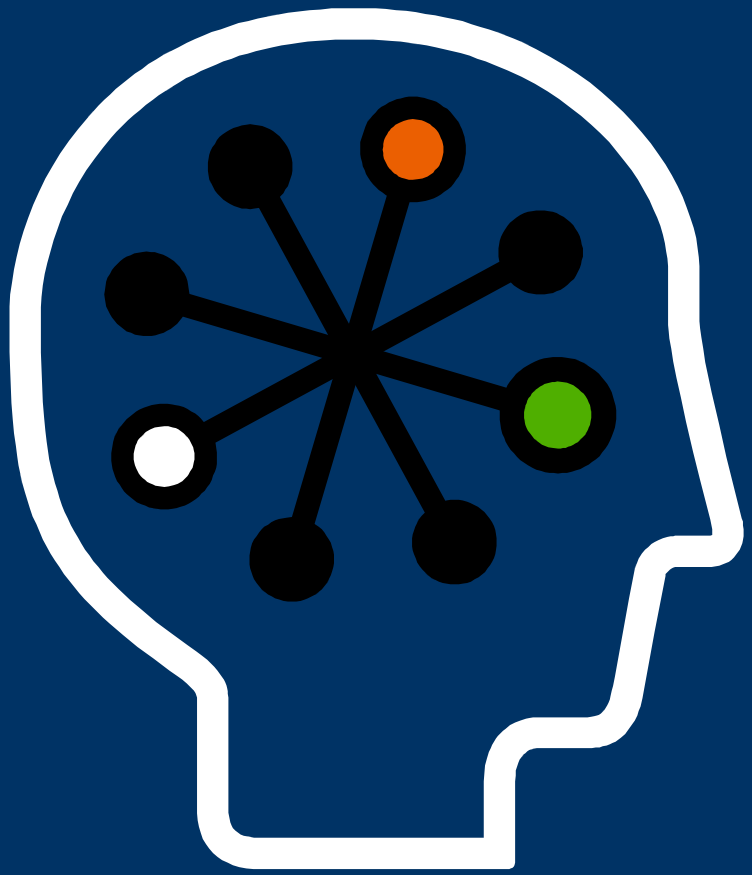

Technology for better business outcomes

Benchmarking

HP Accelerator Catalyst team

- Formed in early 2006.
- Extensive information gathering and hands on testing / benchmarking
- Internal and external White papers published in Jan. 07
- Working directly with accelerator vendors
- Qualifying and characterizing accelerators, both hardware and development environments
  - Benchmarking
    - SGEMM, DGEMM, FFT, Monte Carlo, Streams, others to come
  - Mechanical
  - Power and cooling

Accelerator hardware / software vendors

- Program announcement – Nov.1, 2007
- HPC Accelerator Program
  - Clearspeed e620 qualified in blades and ProLiants
  - Celoxica RCHTX qualified in DL145 G3
  - NVIDIA Tesla Server qualified with DL160 G5
  - ATI Firestream – pending
  - FPGA PCIe cards and socket replacements - pending
  - Rapidmind, Mitrionics contributing innovative programming technologies
- As we go forward, more accelerators and platforms will be qualified.

Benchmarking efforts

- Using common industry benchmarks
  - Matrix-Matrix Multiply
  - FFT 1D & 2D
  - Monte Carlo Black-Scholes
  - Streams
- Bandwidth testing
  - HP internal or vendor supplied test tools
- Others – as needed and as available

Benchmarking Status

- Matrix-Matrix Mutiply
  - NVIDIA GPUs
  - Clearspeed E620
  - Various blades and Proliant servers
- FFT
  - NVIDIA GPUs
  - Various blades and Proliant servers
- Black-Scholes
  - Clearspeed E620 and CATS
  - Various blades and Proliant servers
  - NVIDIA GPUs – results pending
- Bandwidth testing
  - NVIDIA GPUs

Benchmark Futures

- Matrix-Matrix Mutiply
  - ATI GPGPUs
  - FPGAs
  - Various blades and Proliant servers
- FFT
  - ATI GPGPUs
  - FPGAs
  - Various blades and Proliant servers
- Black-Scholes
  - ATI GPGPUs
  - FPGAs
  - Various blades and Proliant servers
- Streams
  - All accelerators
  - Various blades and Proliant servers
- Bandwidth testing
  - ATI GPGPUs
  - FPGAs

Matrix-Matrix Multiply Benchmark

- Equation C = A * B + C
- Square matrixes only
- Dimensions – 100 through 10000
- Using vendor supplied libraries
- Single-precision -single-core, multi-core, GPU results
- Double-precision – single-core, Clearspeed E620 results

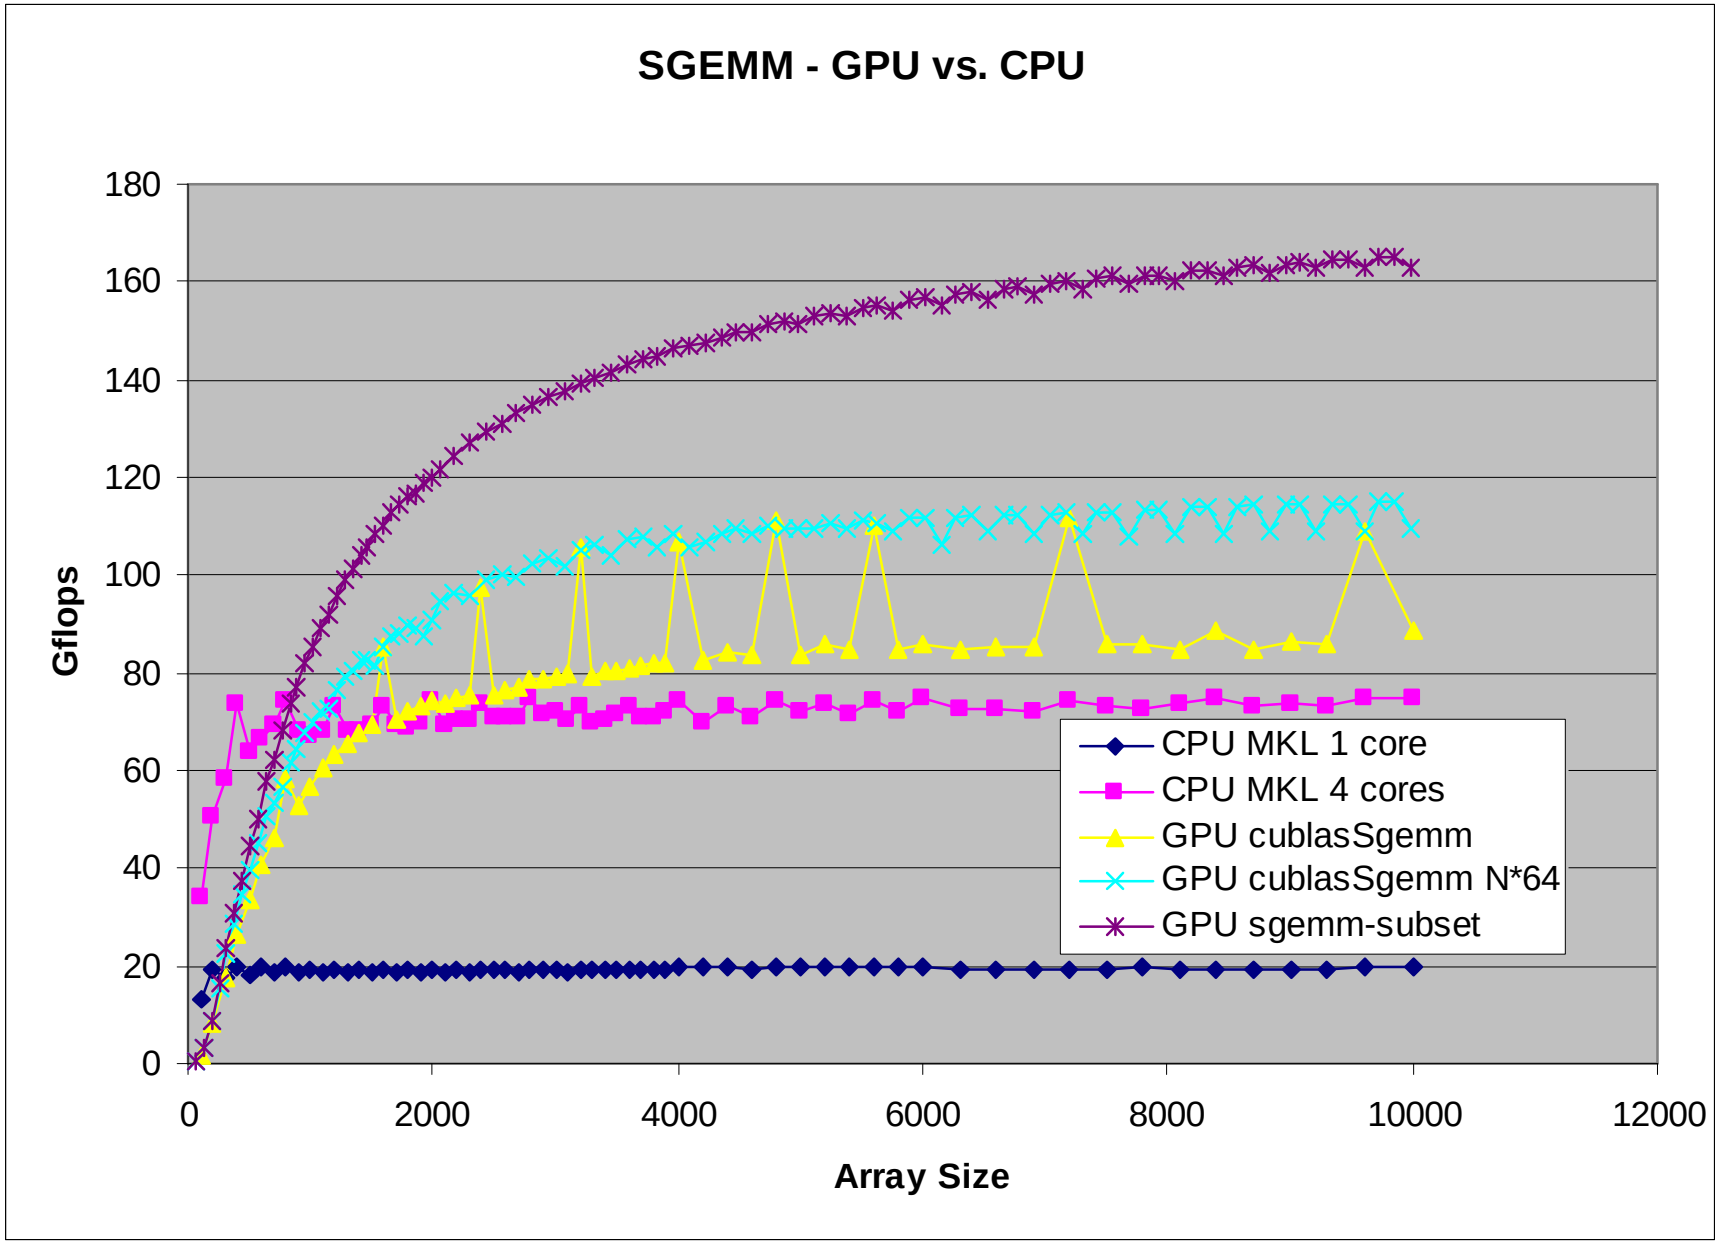
SGEMM - GPU vs. CPU
Gflops
Array Size
0
20
40
60
80
100
120
140
160
180
0
2000
4000
6000
8000
10000
12000
CPU MKL 1 core
CPU MKL 4 cores
GPU cublasSgemm
GPU cublasSgemm N*64
GPU sgemm-subset

DGEMM performance - CS e620 .vs. CPU

(CS card using SDK 2.51 .vs. CPU using MKL; ClearSpeed e620, HP DL380 G5 Xeon 5160 3.0GHz)

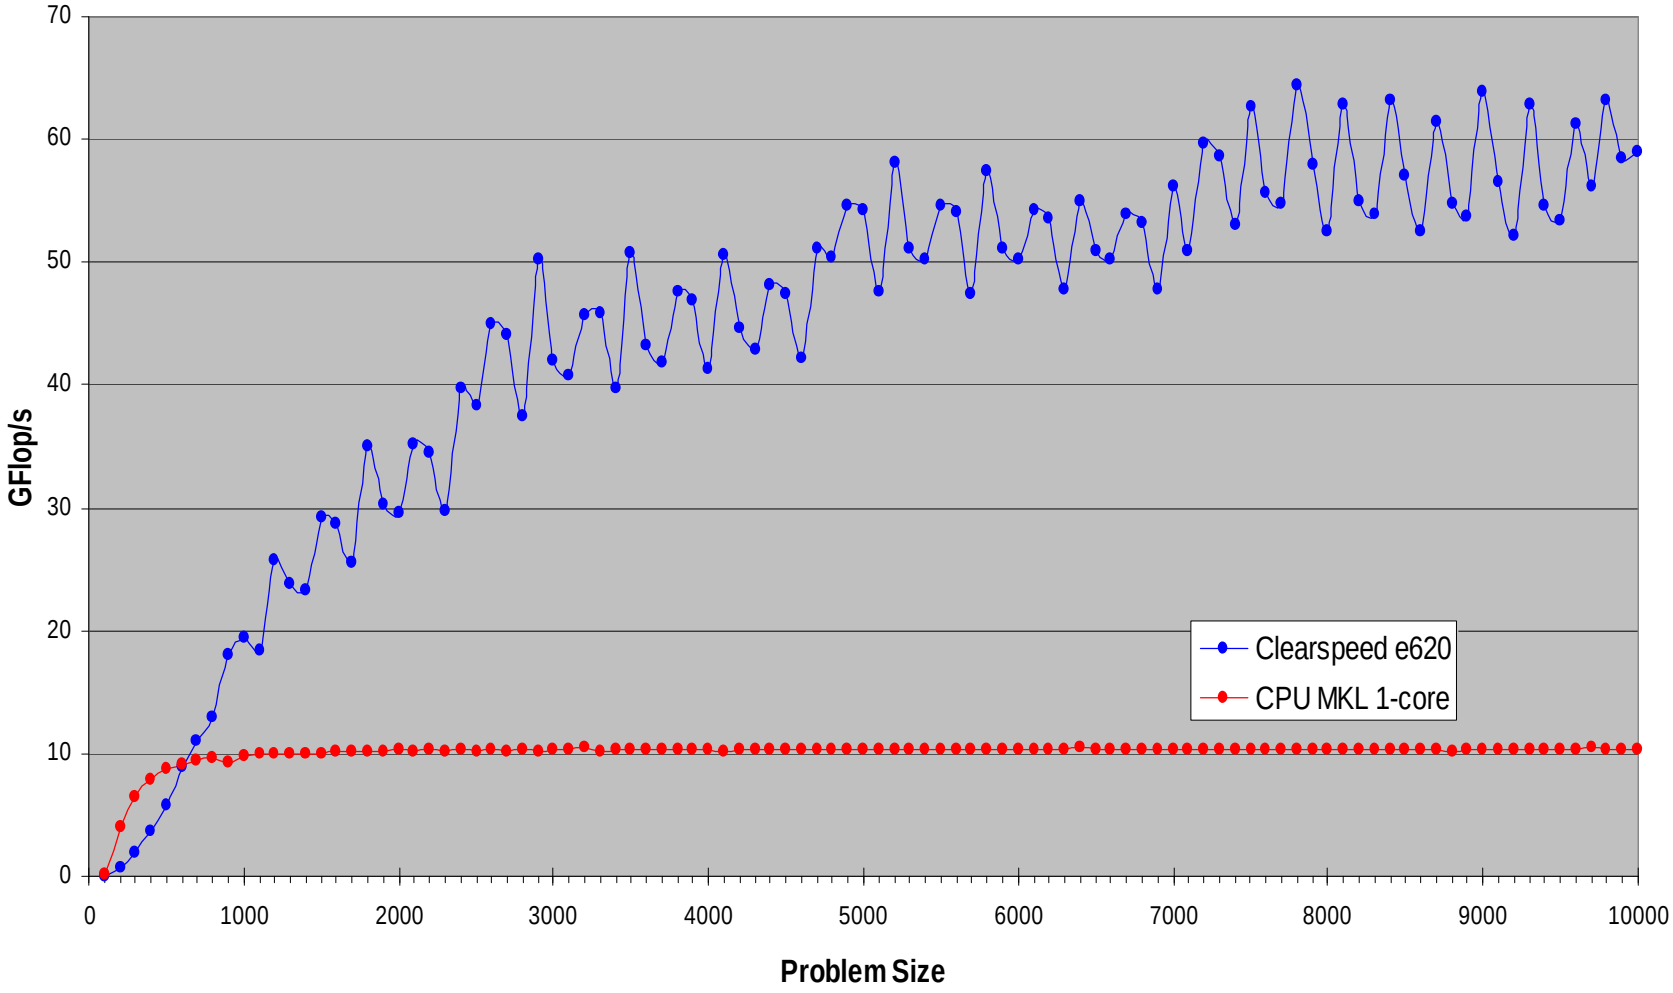

hp invent

Monte Carlo Black-Scholes Benchmark

- European Option pricing using Monte Carlo methods
- Pricing a single option, millions of experiments
- RNG is based on a Hammersley sequence, transformed to a normal distribution
- Computationally intensive, very little data movement
- NVIDIA GPGPU results are pending

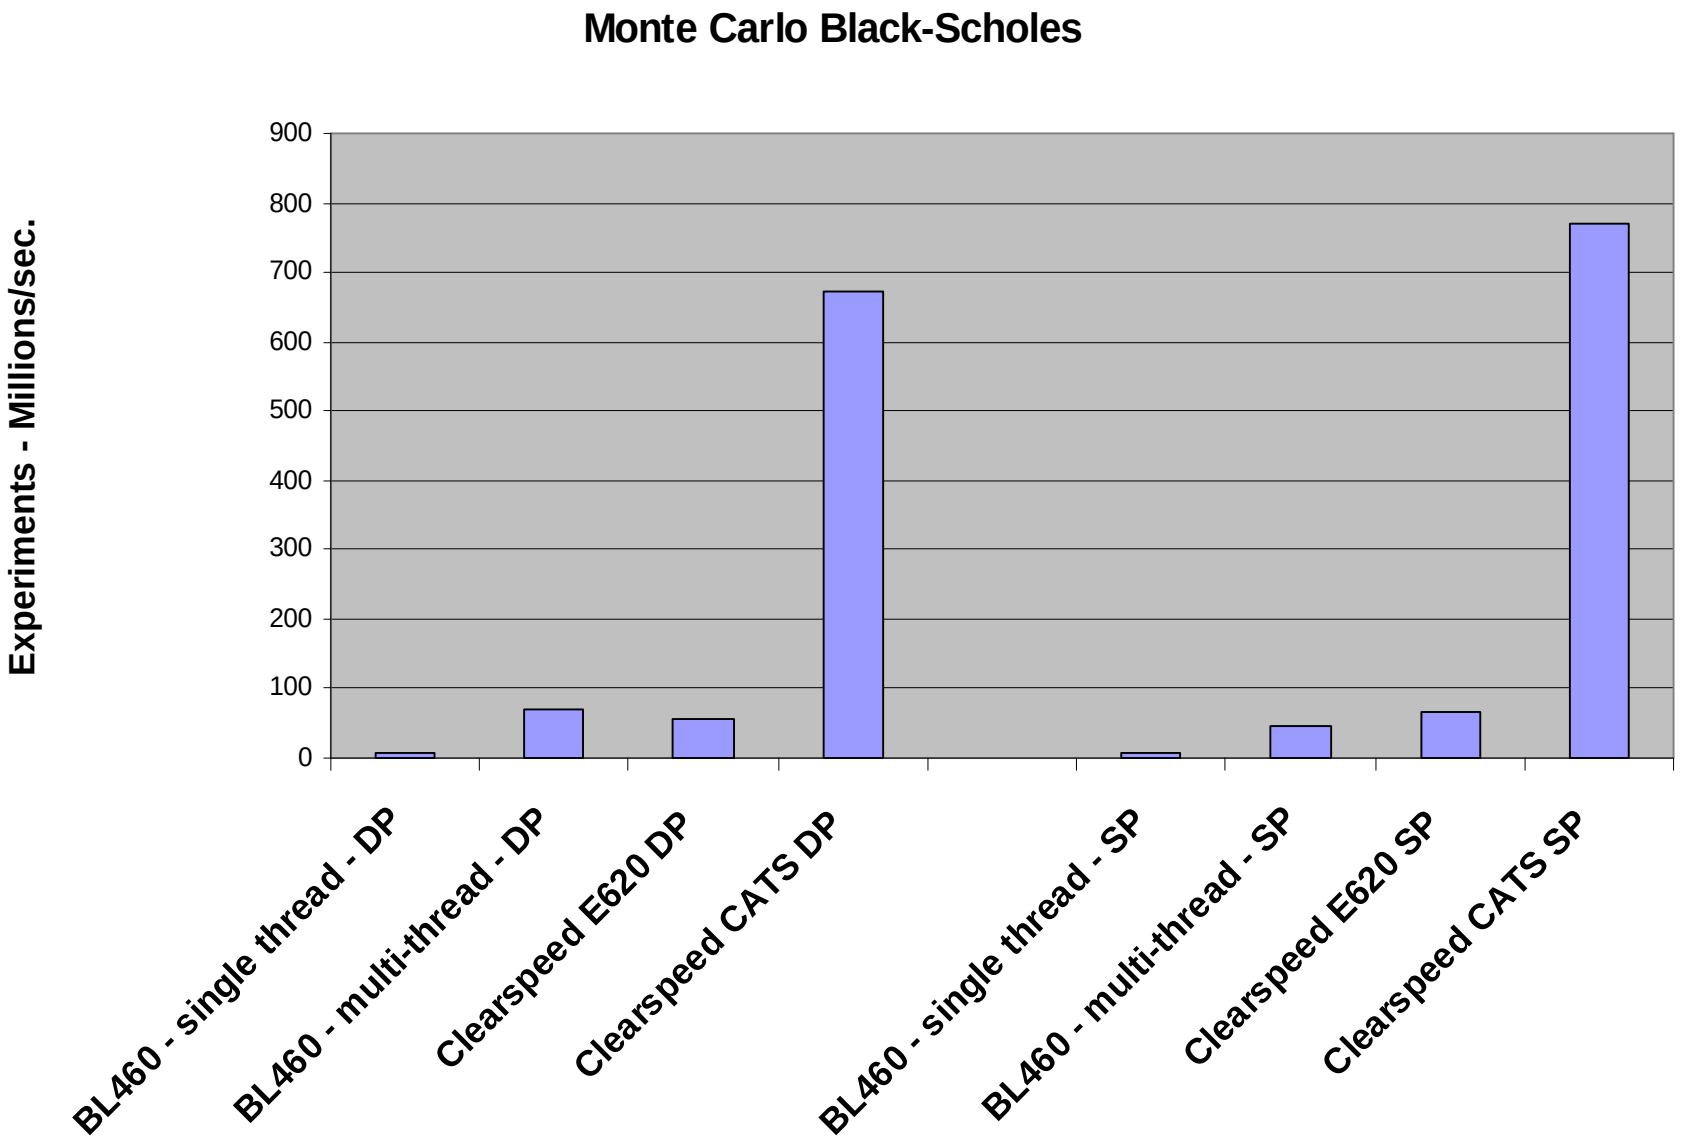
Monte Carlo Black-Scholes
Experiments - Millions/sec.
900
800
700
600
500
400
300
200
100
0
BL460 - single thread - DP
BL460 - multi-thread - DP
Clearspeed E620 DP
Clearspeed CATS DP
BL460 - single thread - SP
BL460 - multi-thread - SP
Clearspeed E620 SP
Clearspeed CATS SP

FFT Benchmarks

- 1D and 2D FFT benchmarks using vendor supplied libraries, MKL and cufft
- 1D FFT sizes are 2 to 1 million elements
- 2D FFT sizes are 2 to 4096 elements, each dimension
- Measured test consists of:
  - Forward FFT
  - Inverse FFT
  - Scale results
- Accelerator data transfers occur before forward fft and after scale of results.
- Accuracy check of results

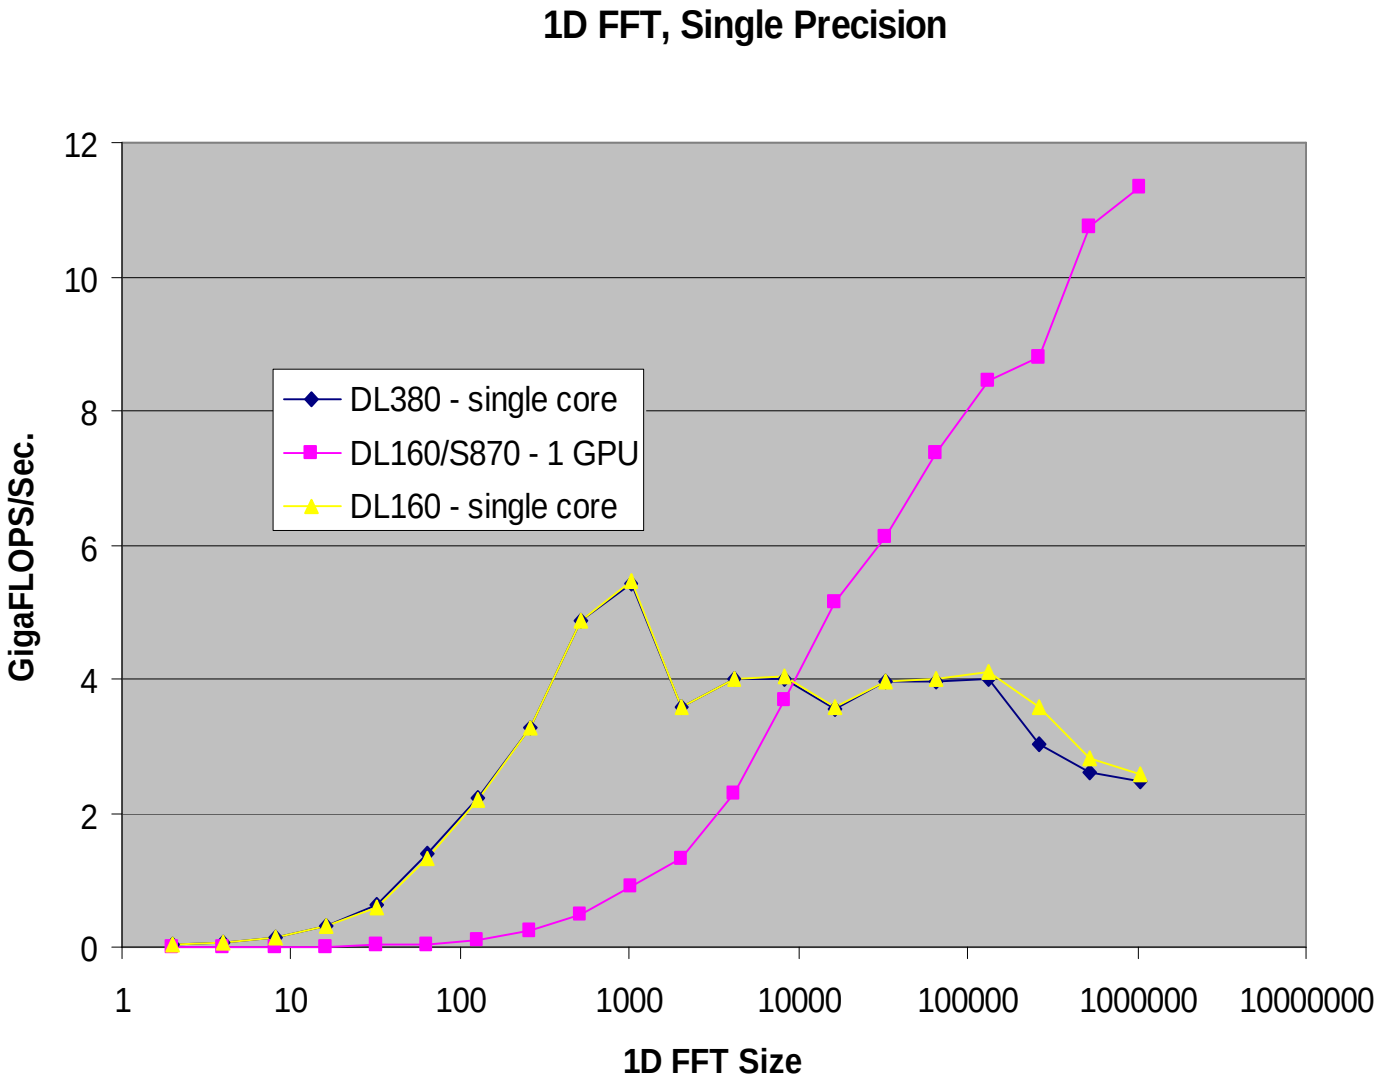
1D FFT, Single Precision
GigaFLOPS/Sec.
12
10
8
6
4
2
0
DL380 - single core
DL160/S870 - 1 GPU
DL160 - single core
1
10
100
1000
10000
100000
1000000
10000000
1D FFT Size

2D FFT, Single Precision

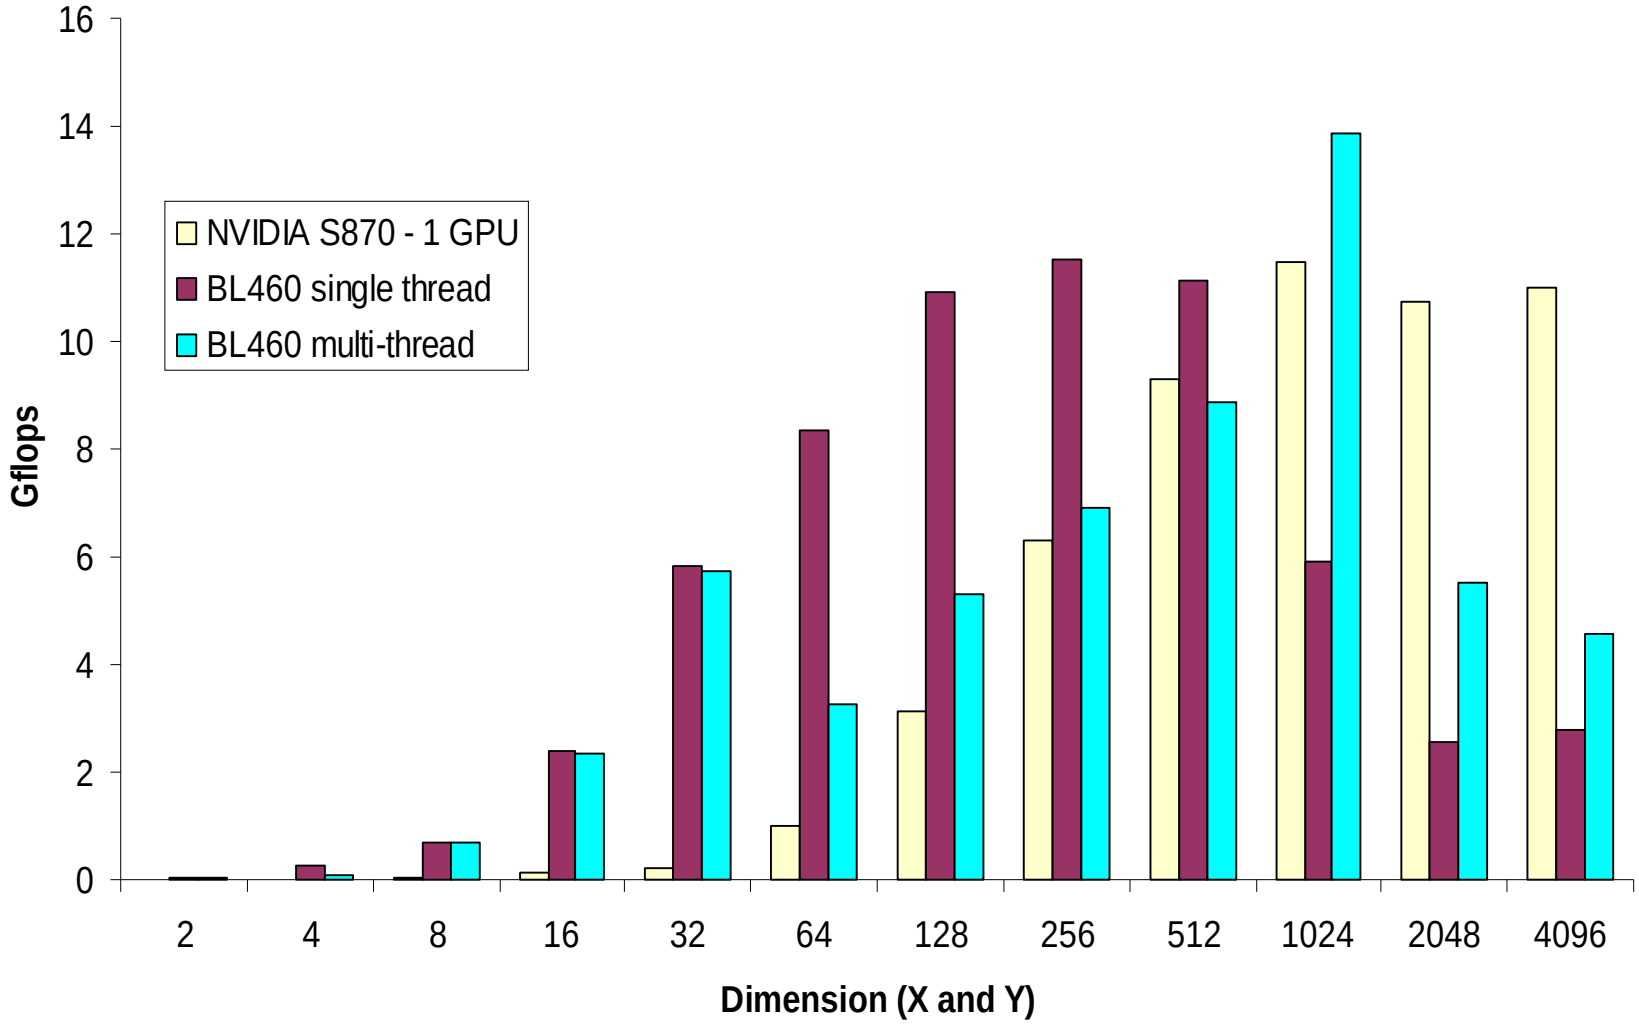

Bandwidth Testing

- Multi-GPU bandwidth test tool
  - Does not execute on the GPU
  - Host to Device, Device to Host, On Device, Read after write Modes
  - Transfer sizes up to GPU memory size
  - Thread per GPU
  - Thread may be synced at the start of a sub-test
  - Each sub-test executed multiple times, results averaged
- Data presented – sampled at the 50Million byte transfer size

DL160/S870, Host to Device, Pinned Memory

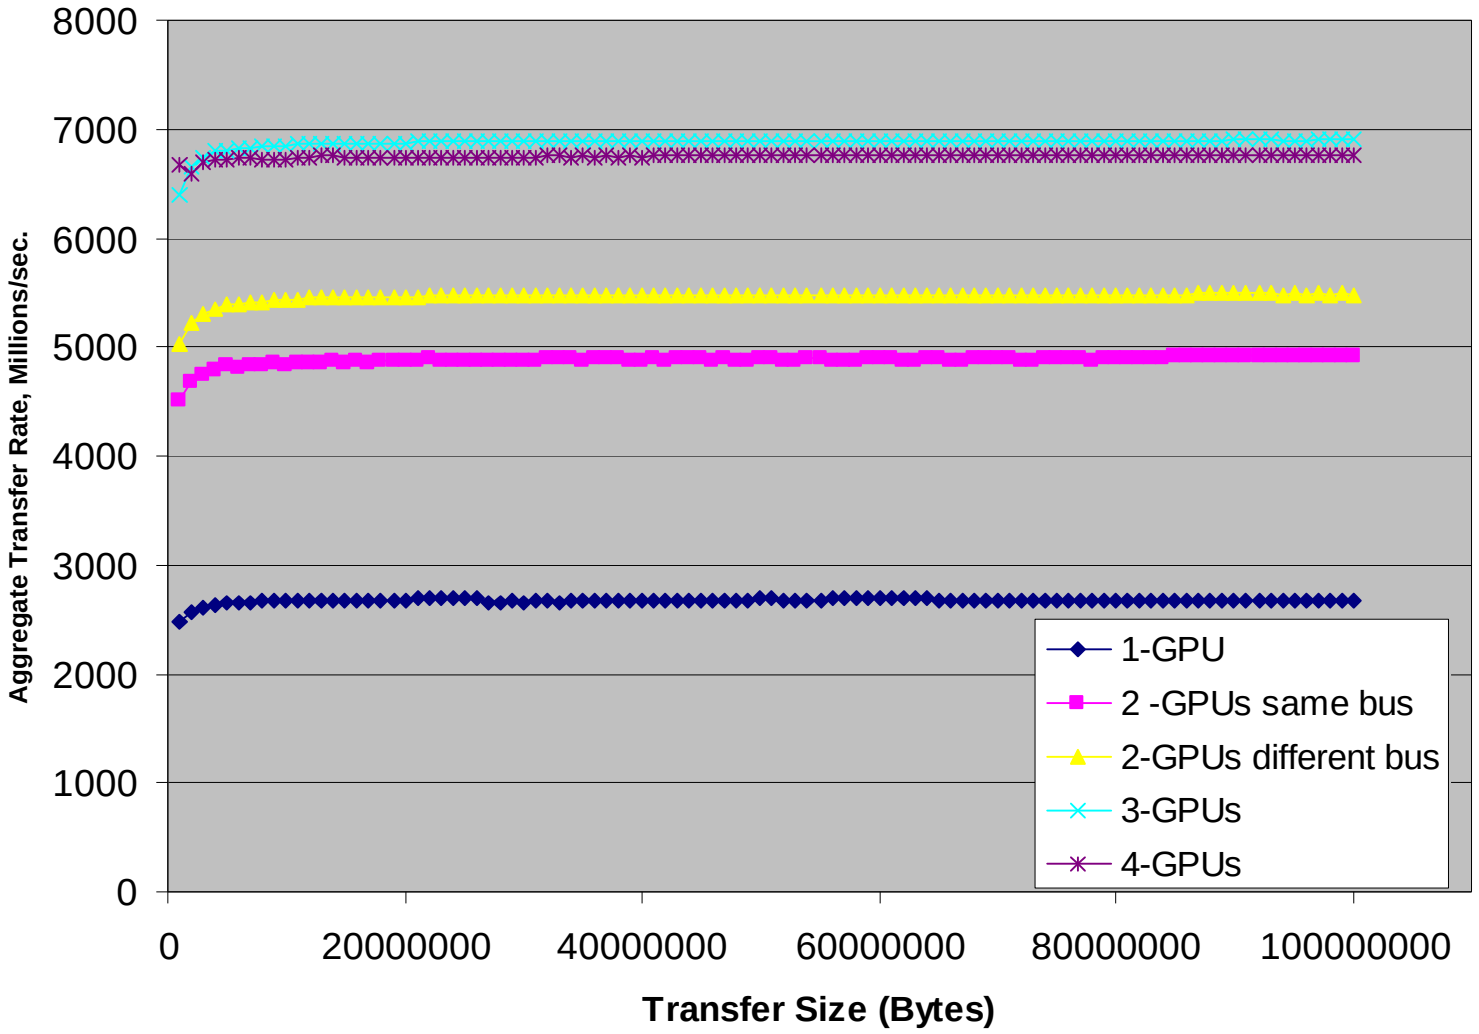

invent

DL160/S870, Host to Device, Pinned Memory

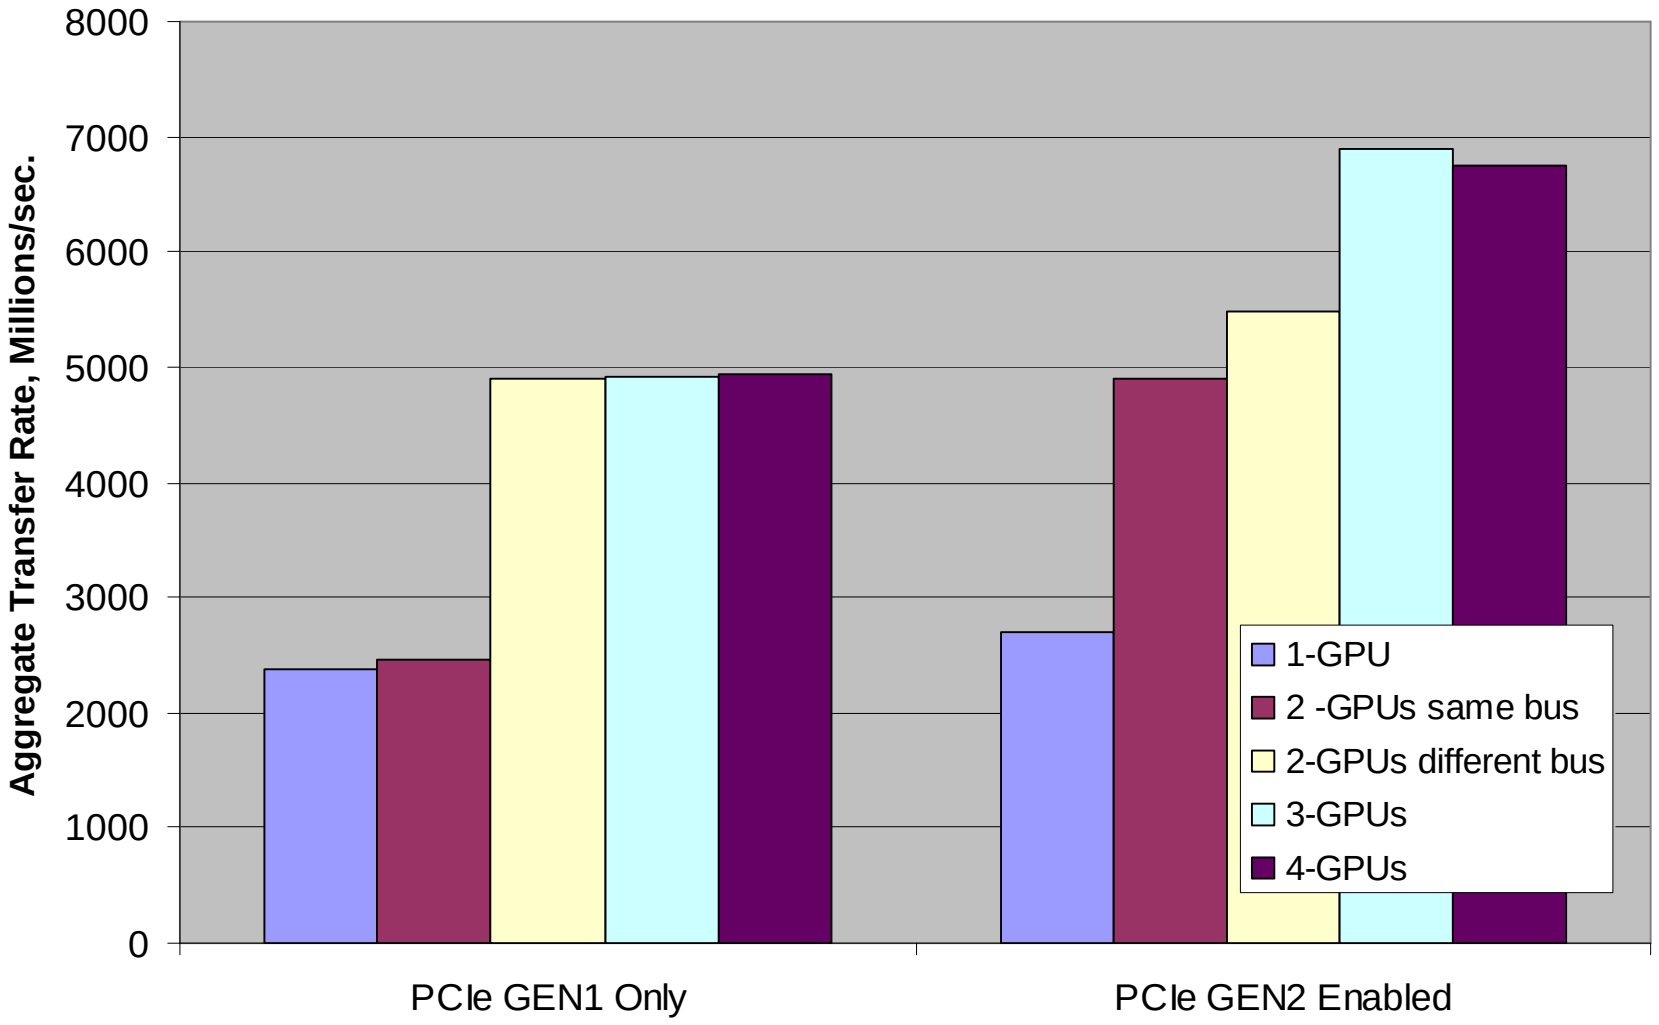

invent

DL160/S870, Device to Host, Pinned Memory

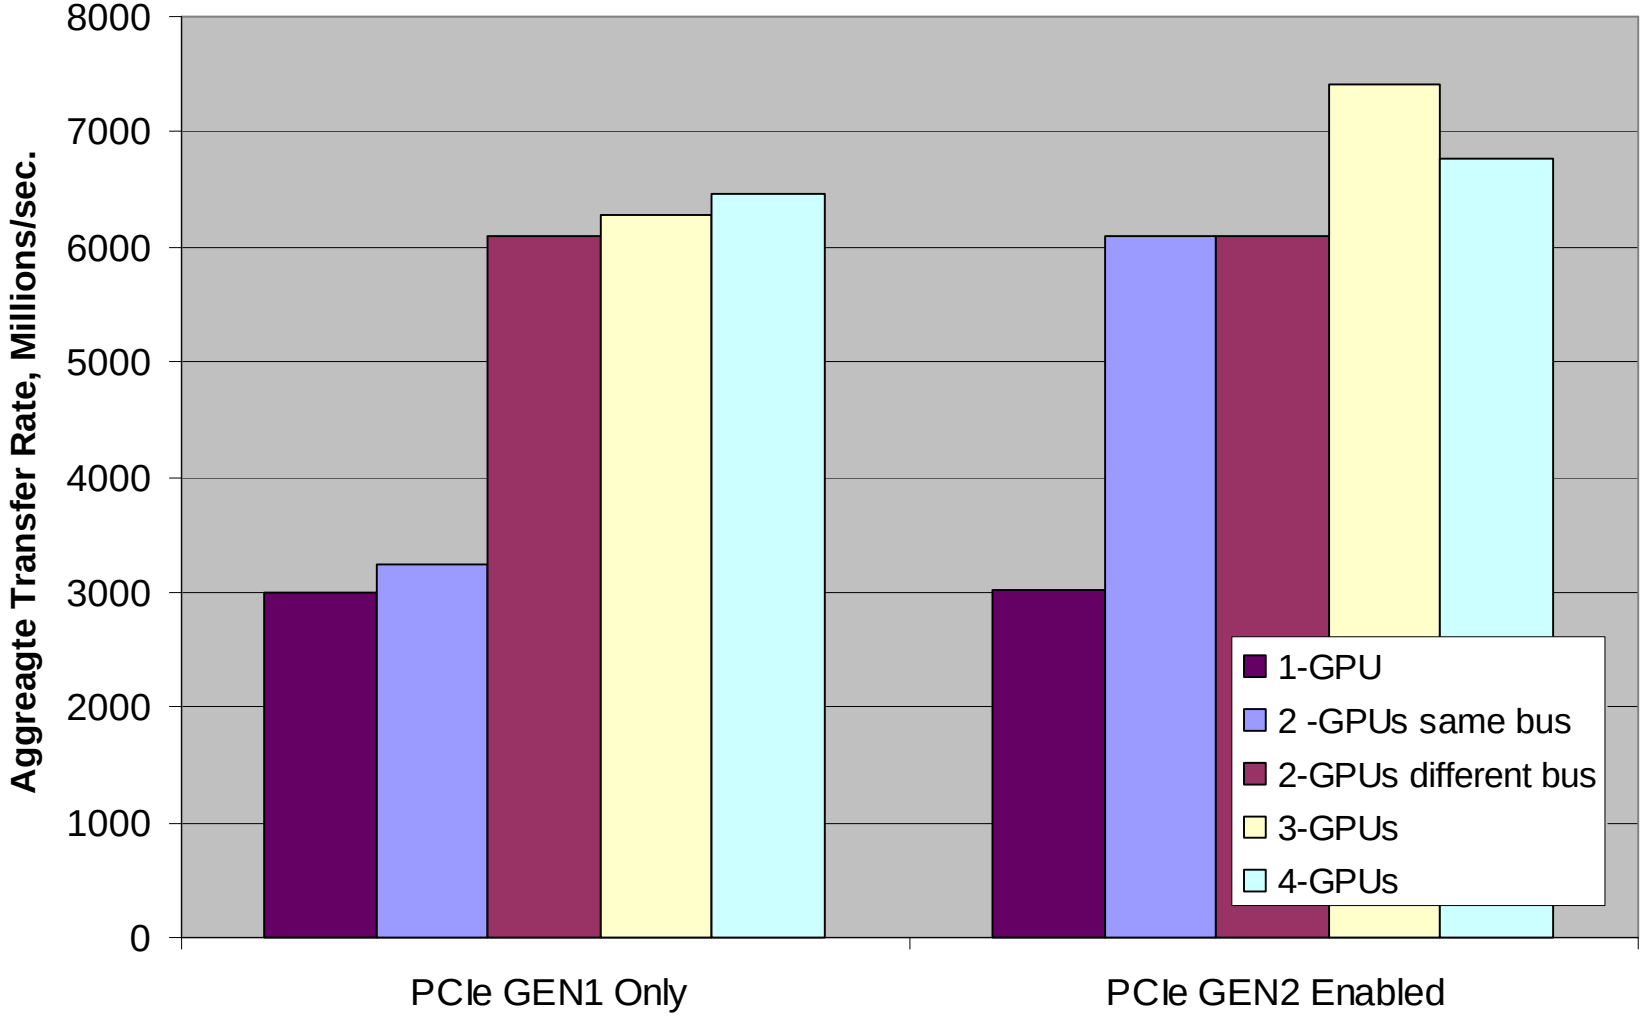

invent

DL160/S870, Host to Device, Paged Memory

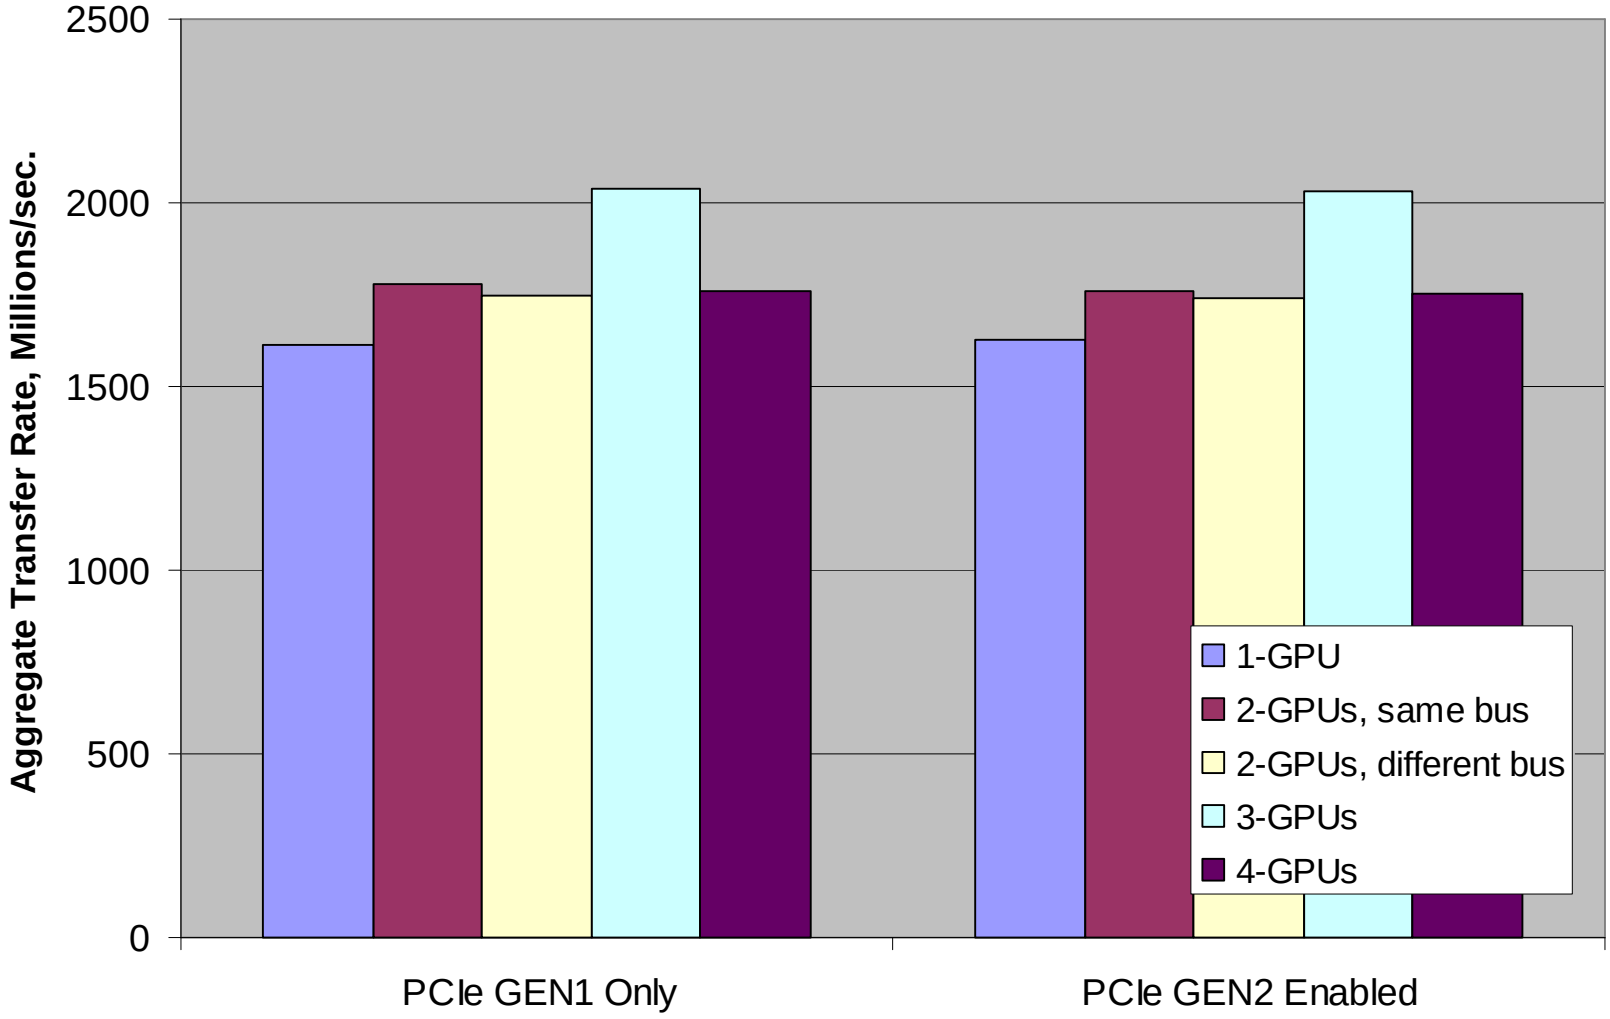

DL160/S870, Host to Device, Paged Memory

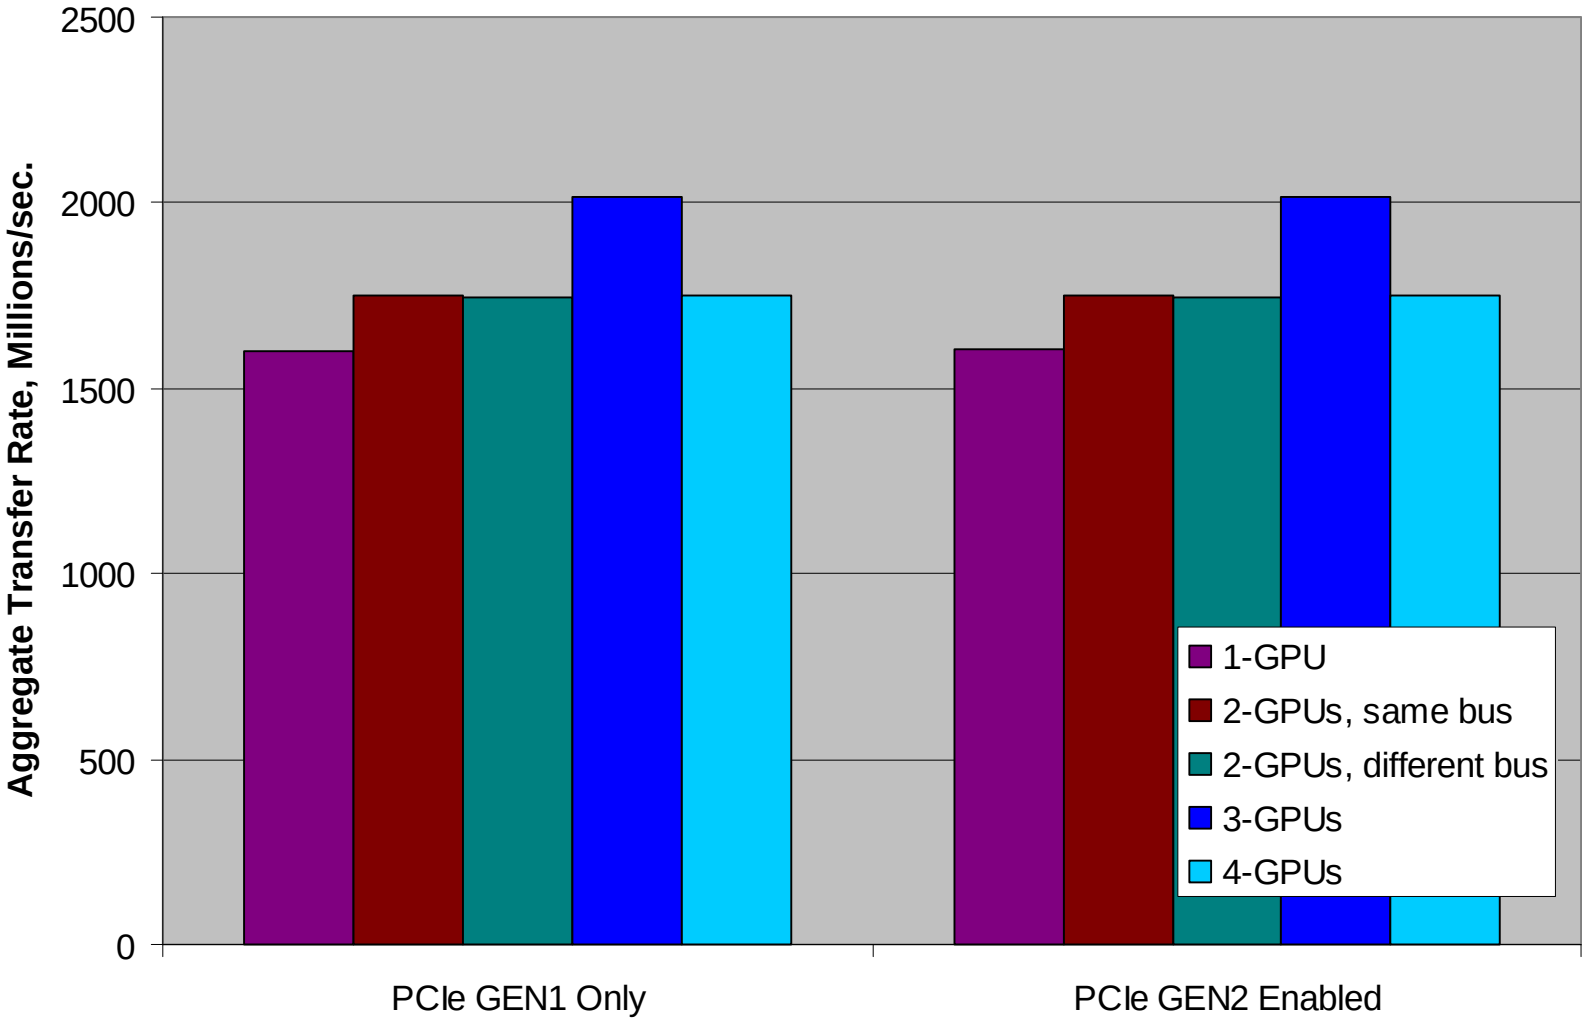

invent

HP Accelerator Collaboration

- Web pages - www.hp.com/go/accelerators
- You need to sign up to be a member
- HP Accelerator CCN Advocate
  - Don Thulin (donald.thulin@hp.com)