GPU based software correlators - perspectives for VLBI2010

Hobiger T.[1], Kimura M.[1], Takefuji K.[1], Oyama T.[2],

Koyama Y.[1], Kondo T.[1], Gotoh T.[1], Amagai J.[1]

(1) Space-Time Standards Group, NICT, Japan

(2) NAOJ, Japan

Content

- Graphics processing units (GPUs)
- Programming GPUs
- Scientific applications based on GPUs
- GPUs for the VLBI software correlator?
- Benchmarking
- Results
- Conclusion and outlook

CPU vs. GPU performance

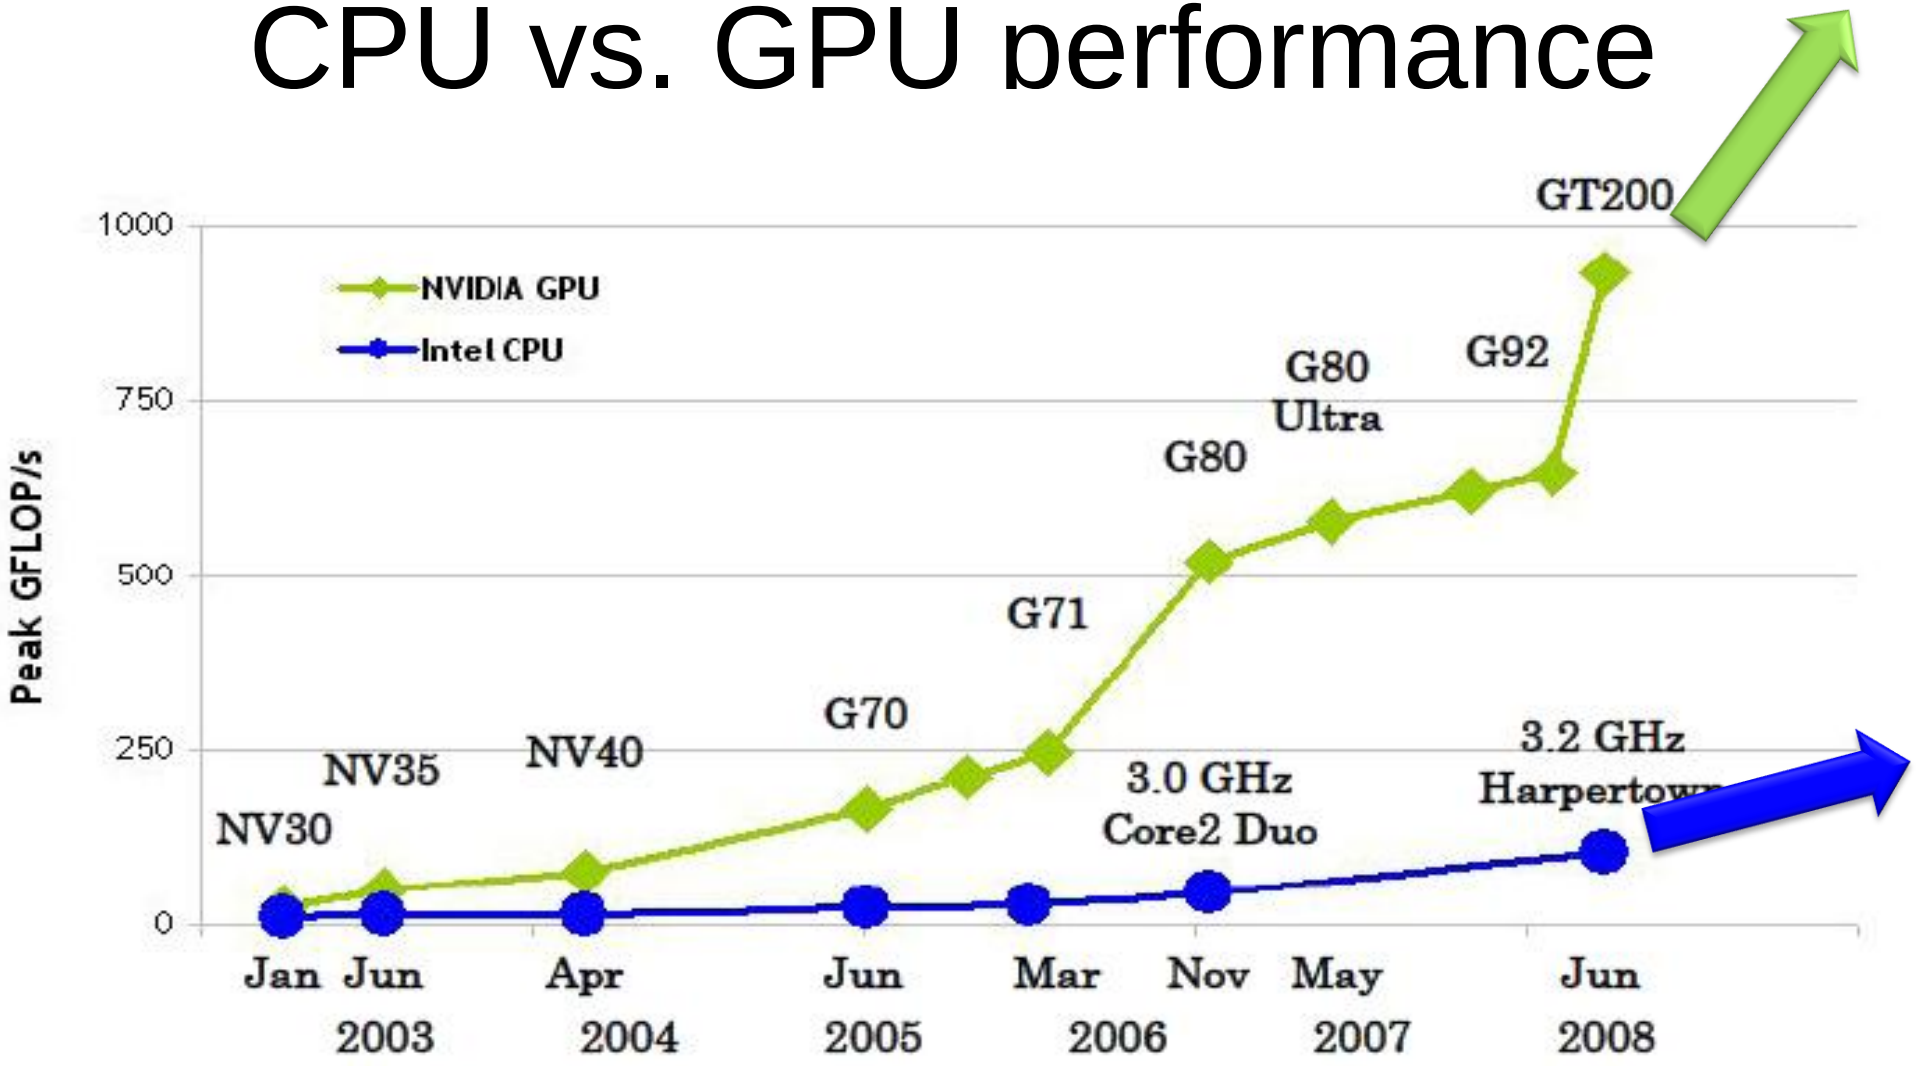

Originally no scientific purposes!
Development mainly driven by game industry

NiCT

GPU a “perfect” parallel processor

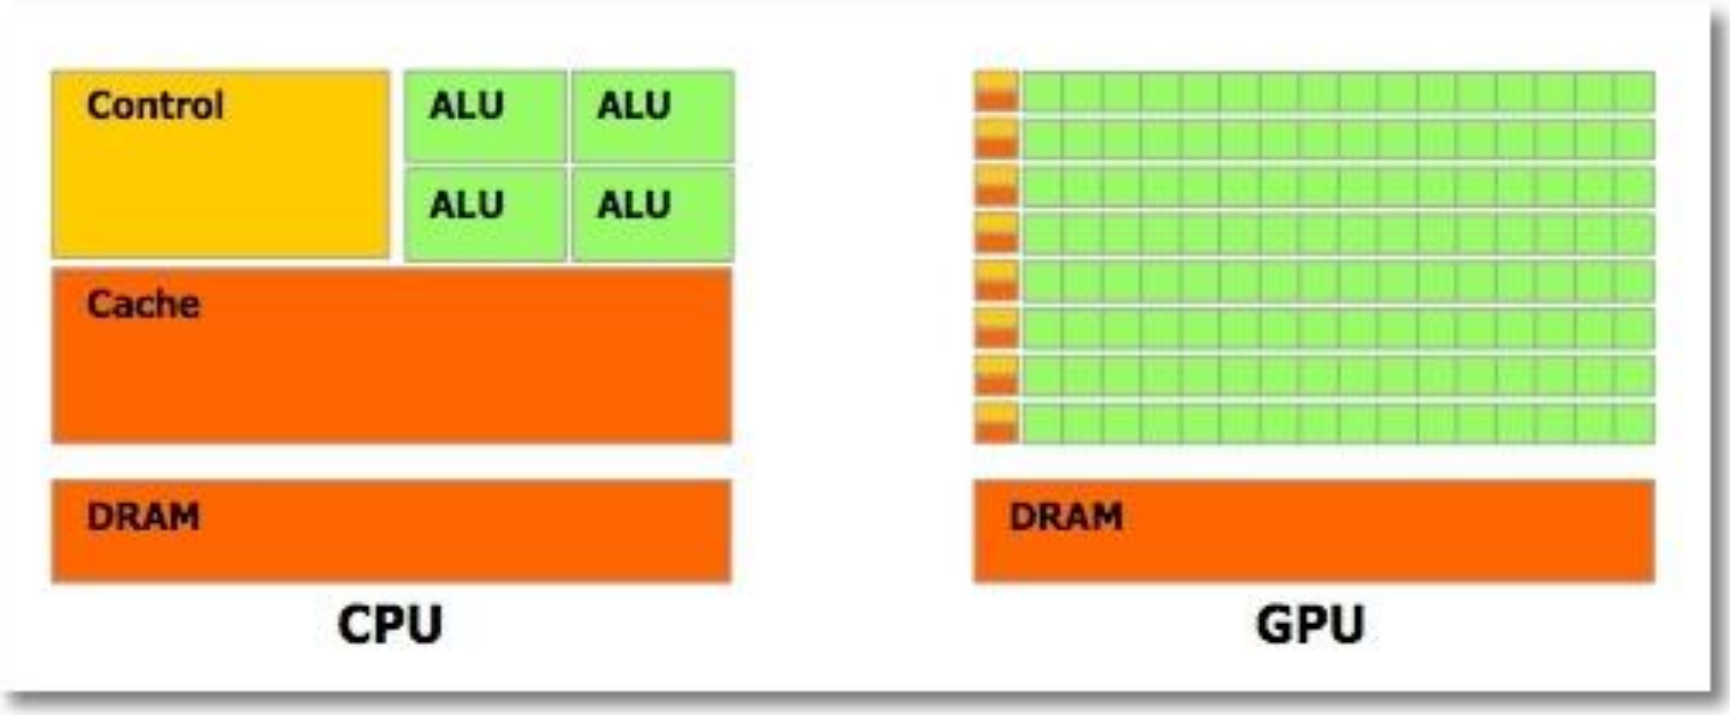

(Multi-core) CPU: Thread parallel
vs.
GPU: Data parallel architecture

GPU programming

- **OpenGL** – very graphics oriented
- **CUDA** – NVIDIA
  - Currently the most mature development platform
- **Brook+**, CAL – AMD/ATI
- **OpenCL**
  - platform independent initiative, now available for NVIDA and AMD GPUs plus AMD multicore CPUs
- Libraries: BLAS, FFT ...

GPUs for radio astronomy

- Murchison Widefield Array (MWA)
  - Real-time post-processing pipeline implemented on GPUs ("Diesel Powered GPU Computing “)
  - First tests with GPUs as correlators
- GPU accelerated radio astronomy signal convolution (Harris et al., ExA, 2008)
- Many other institutes have started to play with this tool
- NICT runs GPUs as GPS receivers and utilizes them for time and frequency transfer experiments
- NICT has started to develop VLBI correlators based on GPUs

Our test equipment

- Card used for this study: NVIDIA GTX 280
  - Processor Cores 240
  - Processor Clock 1296 MHz
  - Memory 1GB GDDR3
  - Price ~ 300 $
  - Power consumption ~ 200 Watt
  - Double and single precision support
- Deployed in off-the-shelf DELL PC
- Extendable up to 3 cards/PC

Benchmarking the GPU

- Simulate single-baseline FX correlation engine (w.o. delay tracking)

	FLOPS*
– Do FFT on station X data	$5 \bullet N \bullet \log_2 (N)$
– Do FFT on station Y data	$5 \bullet N \bullet \log_2 (N)$
– Compute cross-spectrum	$6 \bullet N$
Flops per FX engine:	$10 \bullet N \bullet \log_2 (N) + 6 \bullet N$

- Test serial and parallel execution for different FFT sizes
- Measure performance (Gflops) and theoretical throughput (Msps)
- Measure throughput for CPU→GPU and GPU→GPU data transfer

* Following the “FFT Benchmark Methodology” (http://www.fftw.org/speed/)

FX performance (Gflops)

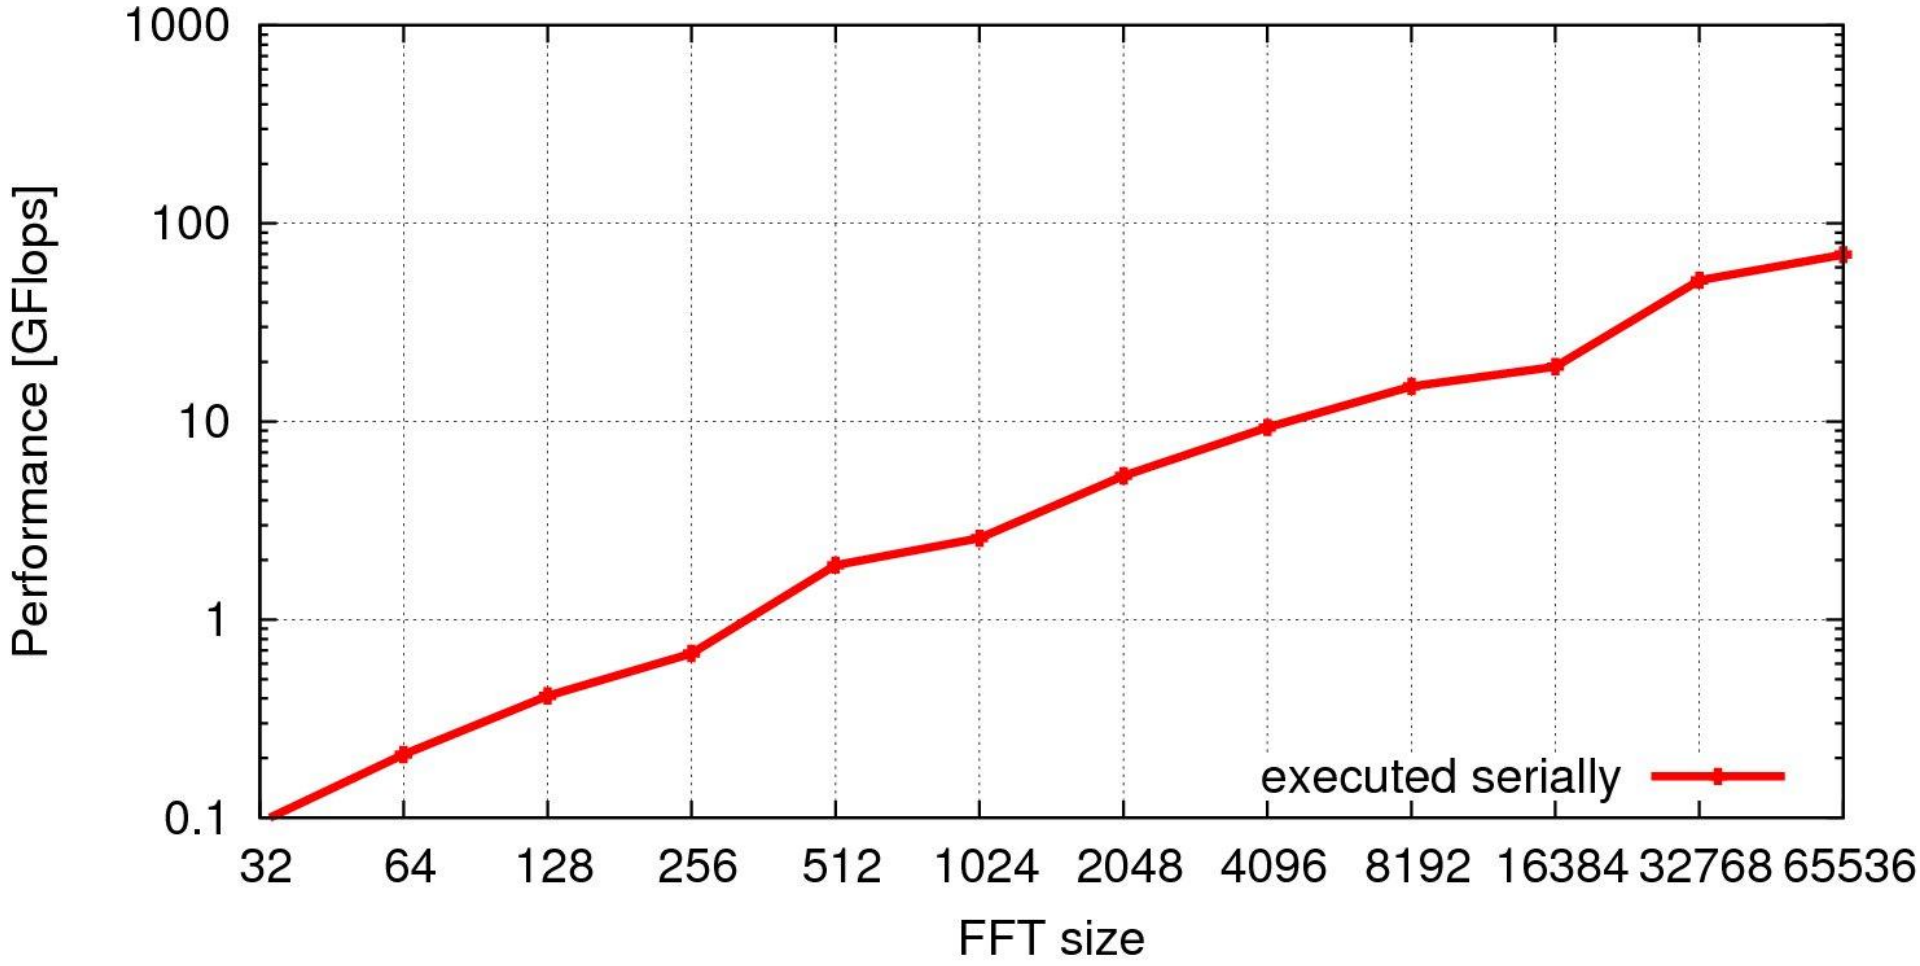

FX performance (Gflops)

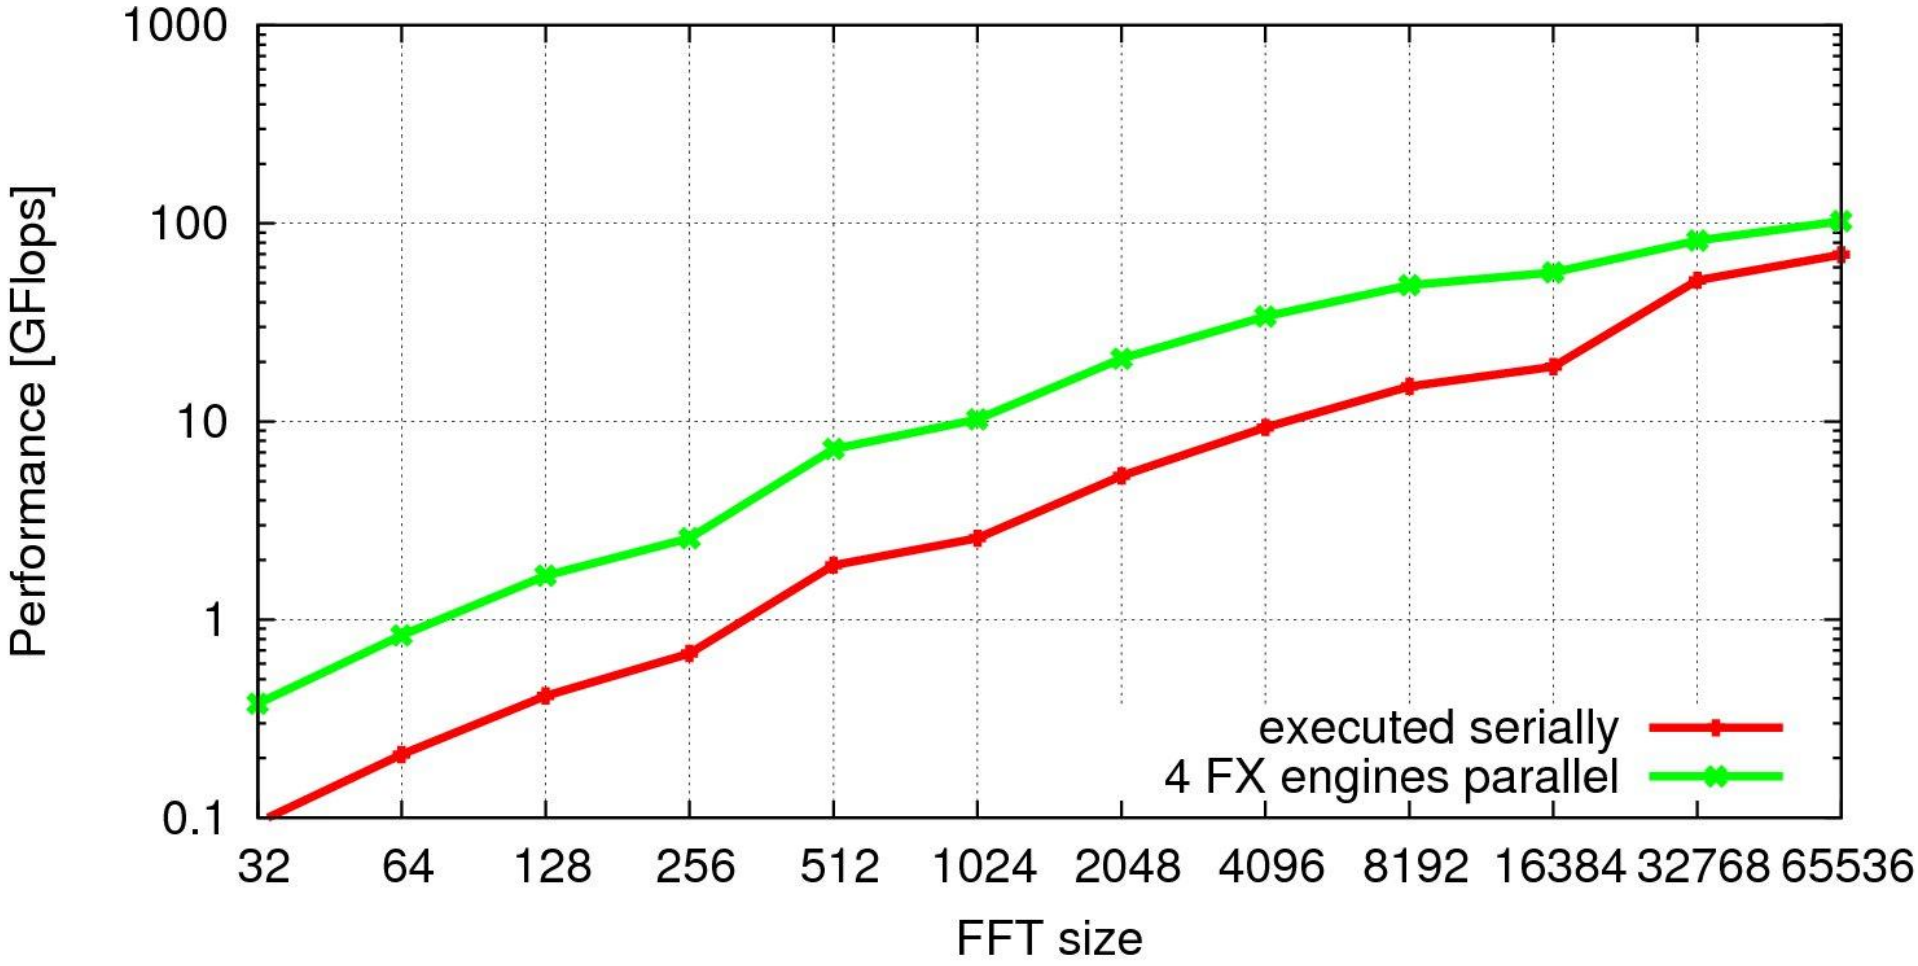

FX performance (Gflops)

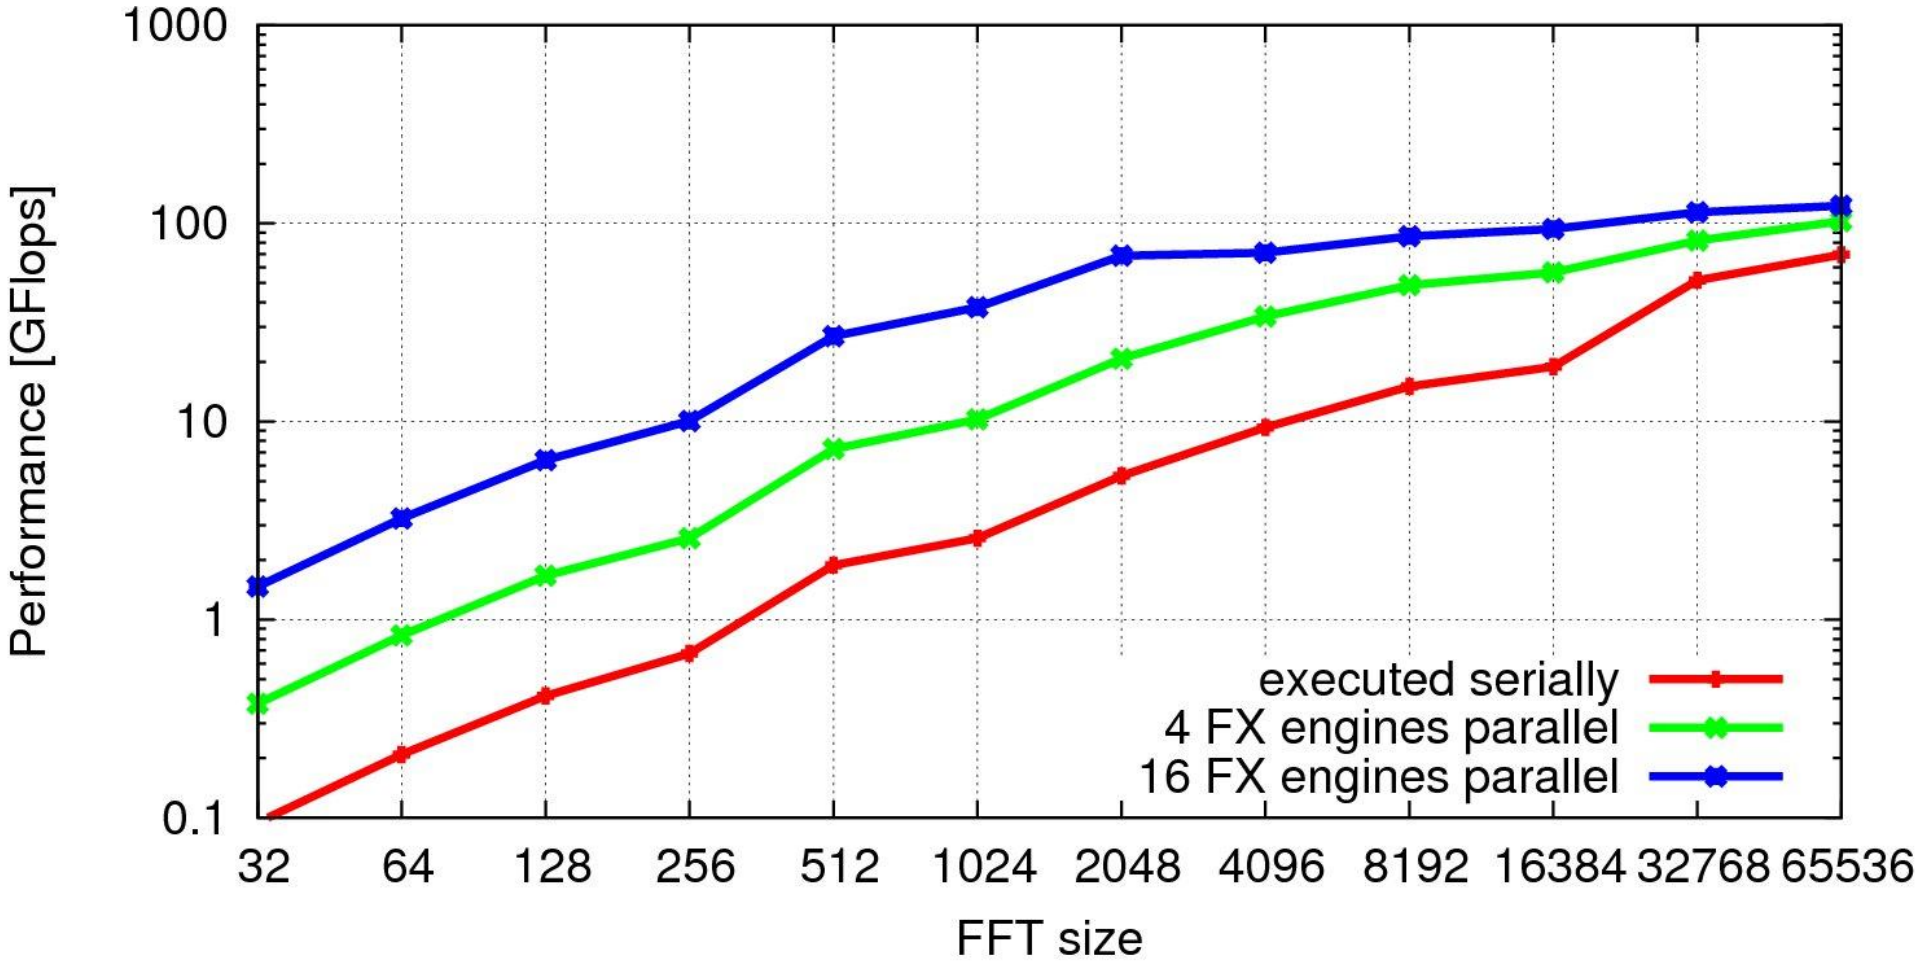

FX performance (Gflops)

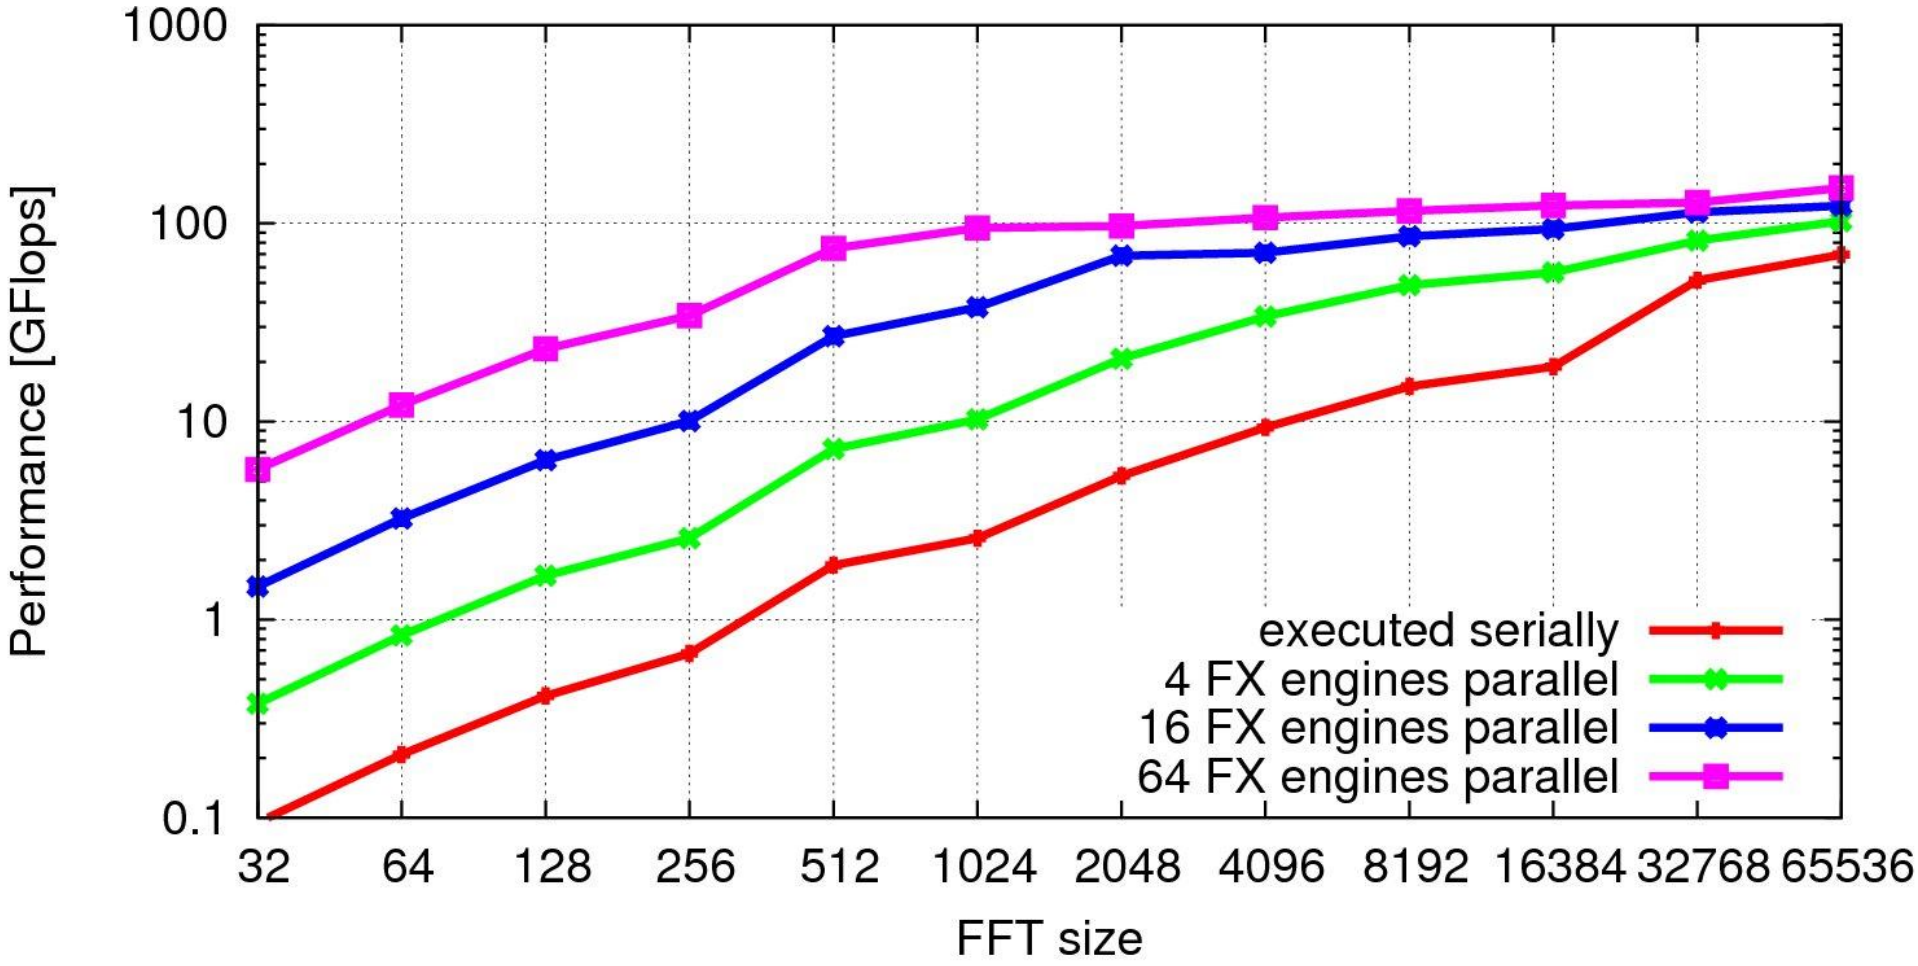

FX performance (Gflops)

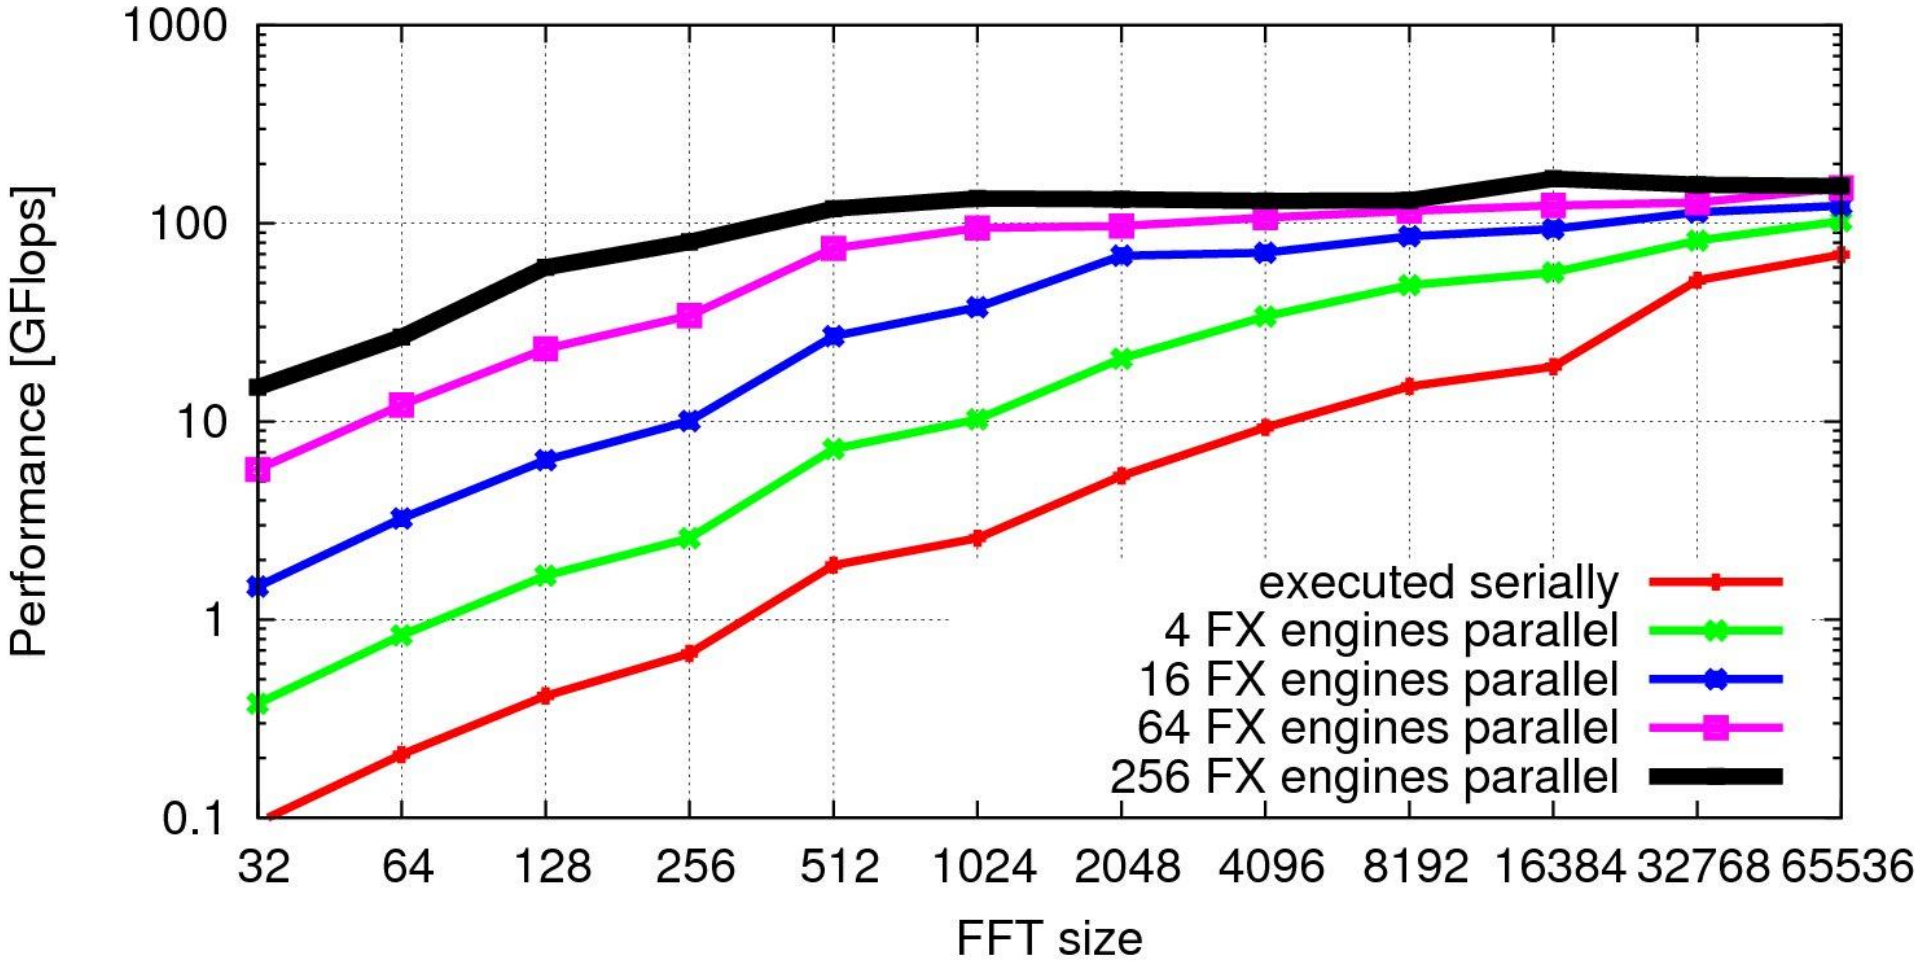

Performance (Msps)

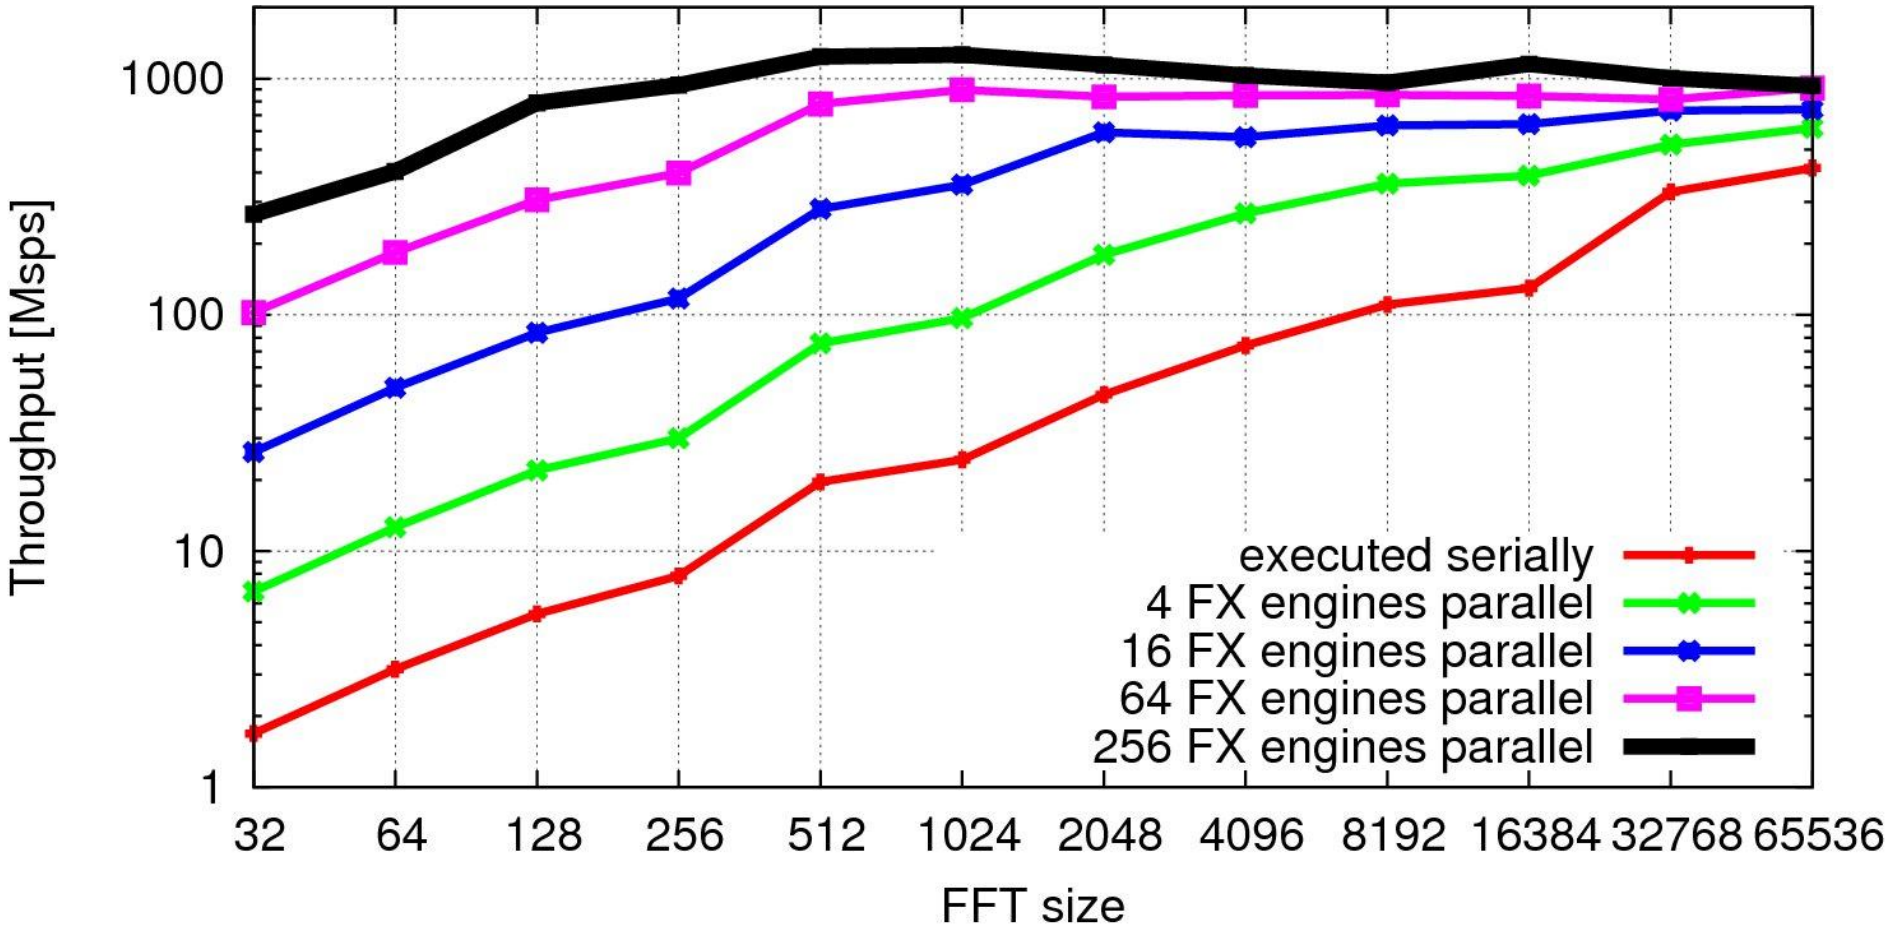

Data transfer

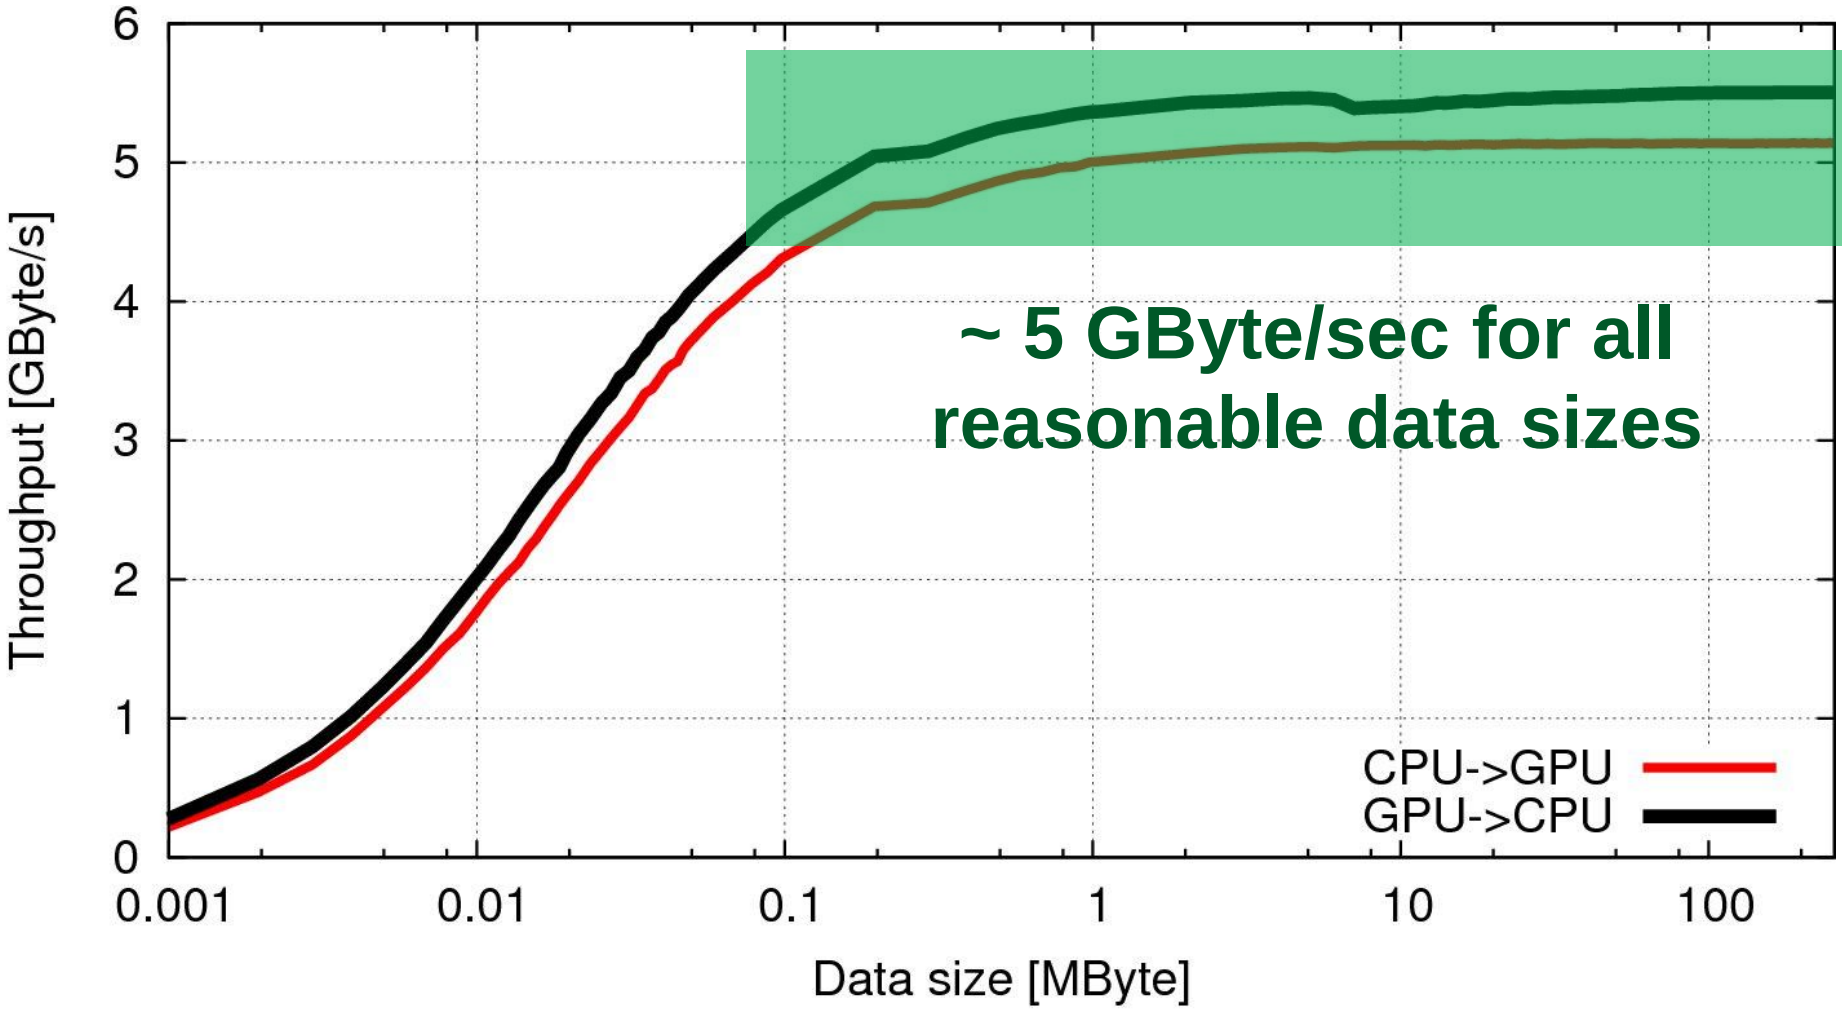

Implementing the correlator

- work done by **Kimura**-san in cooperation with NOAJ
- 2 GTX295 (NVIDIA) cards,
- total price ~ 800 $
- Implementing the “full correlator”
  - Copy data CPU → GPU memory
  - Unpacking + fringe stopping
  - FFT
  - Delay tracking
  - Correlation + Integration
- Testing with 1024 Msps / 2 bit / 1 channel
- Single-station (autocorrelation) mode
- Multi-station mode (correlating all baselines)

Results (Msps)

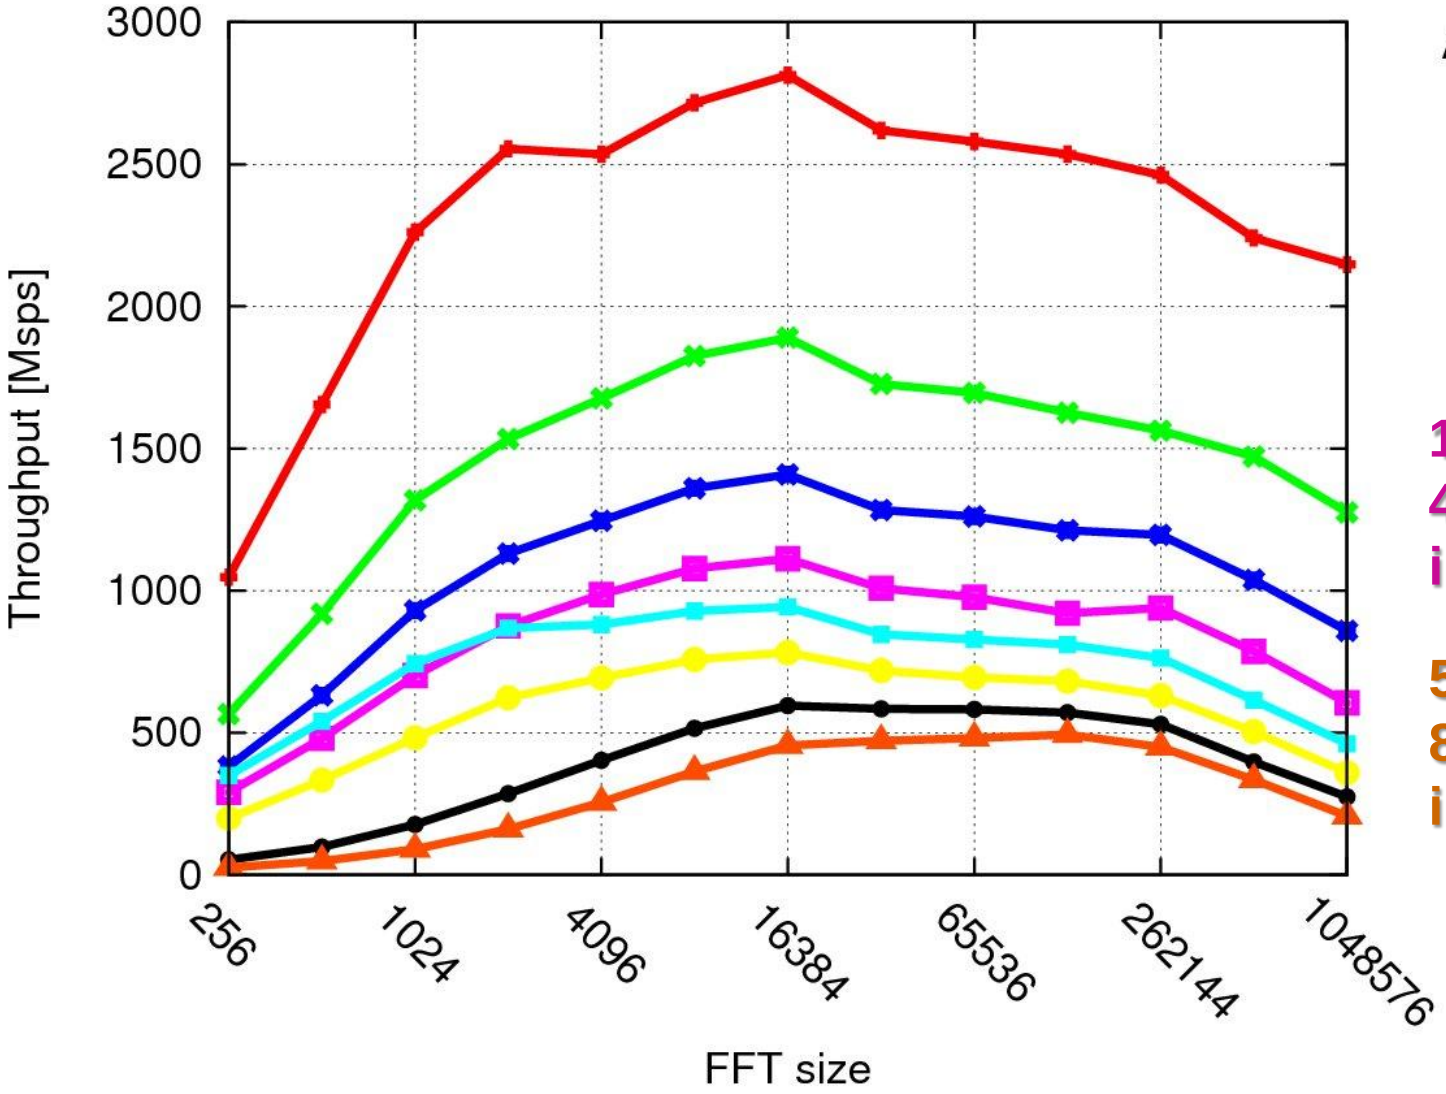

1 Gsps (2bit),
4 stations (6 basel.)
in real-time

500 Msps (2bit),
8 stations (28 basel.)
in real-time

Summary

GPU	CPU (3.60 GHz Pentium 4)
~ 200 Gflops (un-optimized)	~ 12 Gflops (based on best FFTW score)
~ 0.5 Gflops /$	~ 0.1 Gflops / $
~ 0.8 Gflops /Watt	~ 0.2 Gflops / Watt

- Data transfer CPU ↔ GPU is not a severe restriction (moreover, latest CUDA drivers allow processing while transferring data between GPU and CPU)
- Real bottleneck: VLBI2010 raw data, 8–16 Gbps/ station → does not go through 10G ethernet in real-time
- Programming as simple as on the CPU (more simple than writing SSE instructions)

Outlook

- Next generation GPUs announced for Q1 2010
  - two times more computing cores
  - Larger shared memory + L2 cache → speeds up the FFTs
- NICT has started the development of a software correlator based on GPUs
- VLBI2010@Home ?
  - SETI@Home and Folding@Home successfully use the computing power of thousands of idle GPUs worldwide

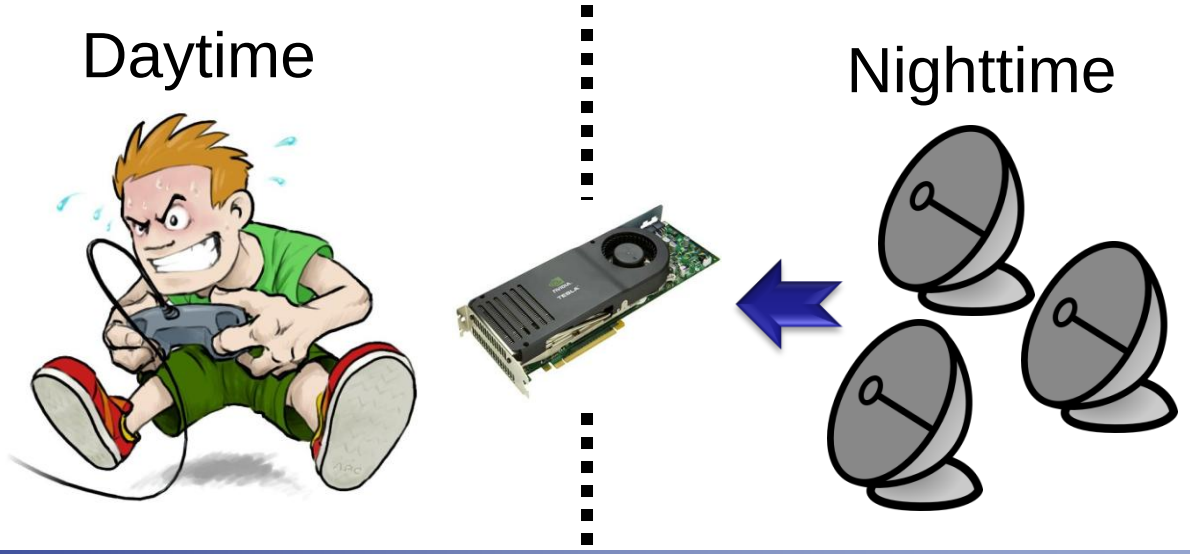

Thank you very much for your attention!

Contact:

hobiger@nict.go.jp

This work was supported by a Grant-in-Aid for scientific research (KAKENHI, No. 21241043)