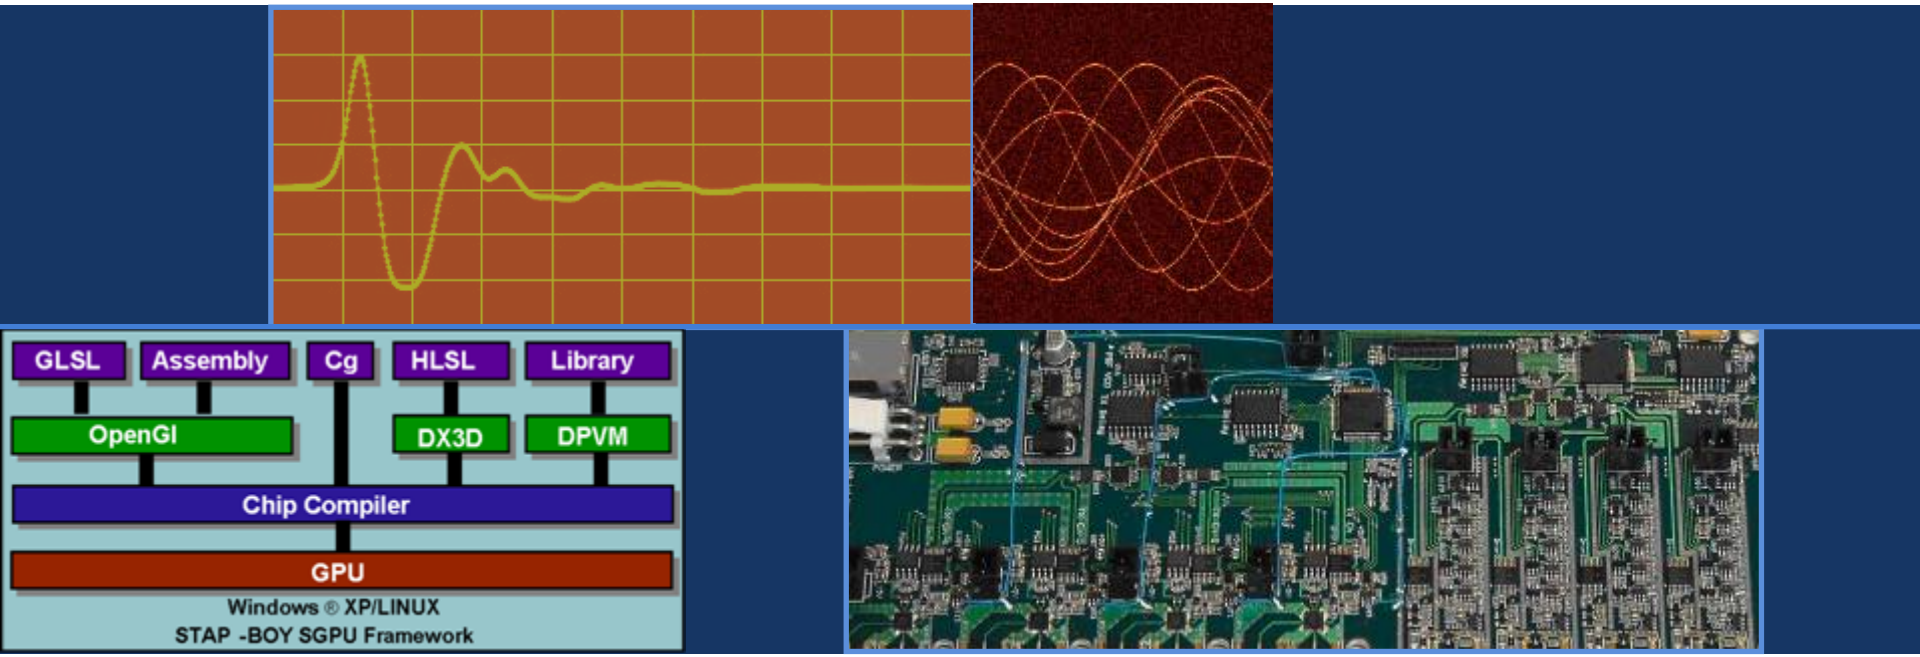


# Power Consumption of Desktop and Mobile GPUs for IRSTAP Applications

Michael Roeder, Jeremy Furtek, Nolan Davis,
Cezario Tebcherani, Masatoshi Tanida and Dennis Braunreiter

High Performance Embedded Computing (HPEC) Workshop, 23-25 September 2008

# STAP-BOY: Concept

## Problem:

- Complex sensor modalities and algorithms needed for smaller platforms (SAR, 3D-motion video, STAP, SIGINT*…)
- Low-cost platform constraints limit real-time on-board/off-board and distributed sensing algorithms and performance
- Timely distribution, visualization, and processing of mission-critical data not available to tactical decision makers

## STAP-BOY goal:

- Develop low-cost, scalable, teraflop, embedded multi-modal sensor processing capability based on commercial off-the-shelf (COTS) graphics chips

## STAP-BOY approach:

- Map complex algorithms to COTS graphics chips with open source graphics languages
- Prototype scalable parallel embedded computing architecture for handhelds to teraflop single card
- Demonstrate on available tactically representative sensor systems

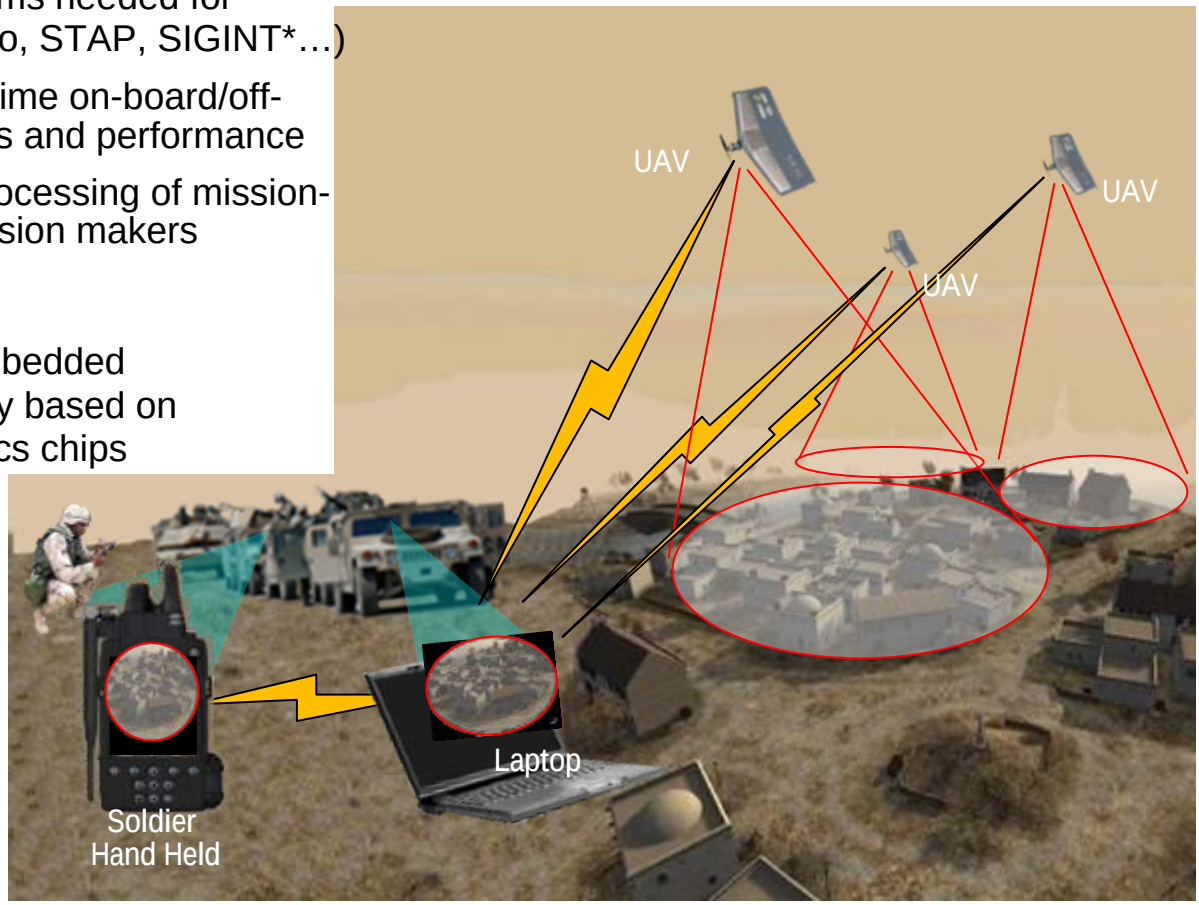


**Goal:**

**½ Teraflop**
**10 Mobile GPUs**
**100W Total Power**
**$<15K**



SAIC From Science to Solutions

# Applications Pull

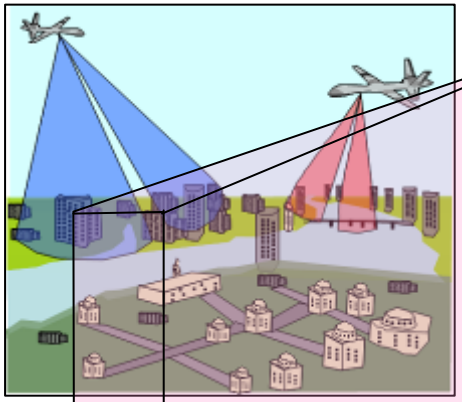

1000+
ASIC
20
10
1.5
1.0
0.5
0.1
Cost ($K)

10km, 32beam
1000Mpixel 2 Hz
267Mpixel 2 Hz
64km/1ft
67Mpixel 2 Hz
16km/1ft
4km/1ft
64km, 64beams
1km, 16beams
1km/1ft
16Mpixel 2Hz

CPU/DSP Systems

Phase III Handheld

50 75 100 200 350 500 1000 2000
GFLOPs

Power (Watts)
400 200 100 40 30 20 10

GPGPU

Phase II 500 GFLOP/100-200W

Image size
Frame rate

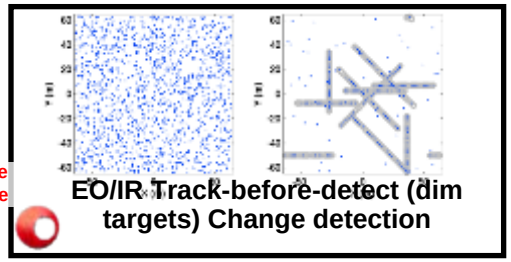

EO/IR Track-before-detect (dim targets) Change detection

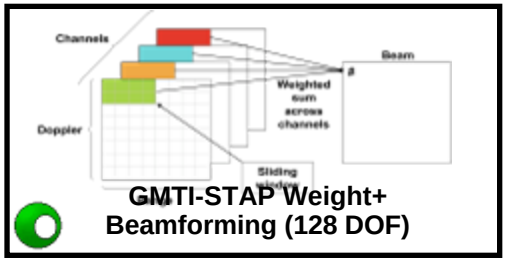


GMTI-STAP Weight+ Beamforming (128 DOF)

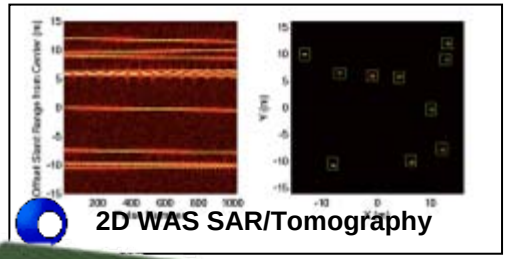

2D WAS SAR/Tomography

**Increasing Advanced Sensor Resolution Driving Processing Throughput Beyond CPU/DSP Capability In Communications and Cost-constrained Platforms**



SAIC From Science to Solutions

# NVIDIA® GPU Performance/Power Study

| | GeForce™ 280 GTX | GeForce™ 8800 GTX | GeForce™ 8800M GTX |
|---|---|---|---|
| Category | Desktop | Desktop | Mobile |
| Process (nm) | 65 | 90 | 65 |
| Transistor Count (Millions) | 1400 | 681 | 754 |
| Stream Processors | 240 | 128 | 96 |
| Shader Clock (GHz) | 1.30 | 1.35 | 1.25 |
| Memory Clock (MHz) | 1107 | 900 | 800 |
| Memory Amount (MB) | 1024 | 768 | 512 |
| Thermal Design Power (W) | 236 | 177 | 65 |
| Release Date | June 2008 | November 2006 | November 2007 |

# CPUs vs. GPUs: Head to Head (Floating Point)

| Intel® quad-core Q9650 | | NVIDIA® 280 GTX |
|---|---|---|
| **820 million** | Transistors | **1400 million** |
| **3.0 GHz** | Clock Speed | **1.3 GHz** |
| **4** | Number of Cores | **240** |
| **Serial** | Programming Model | **Highly parallel** |
| **Minimize latency** | Design Goal | **Maximize throughput** |
| **Complex cores:**<br>**• Branch prediction**<br>**• Out-of-order execution** | Design Approach | **Simple cores:**<br>**• Smaller caches**<br>**• In-order execution** |
| **130 W** | Thermal Design Power (TDP) | **236 W** |
| **96 GFLOPS** | Theoretical Max. Computation Rate (single precision) | **933 GFLOPS** |

# IRSTAP Application

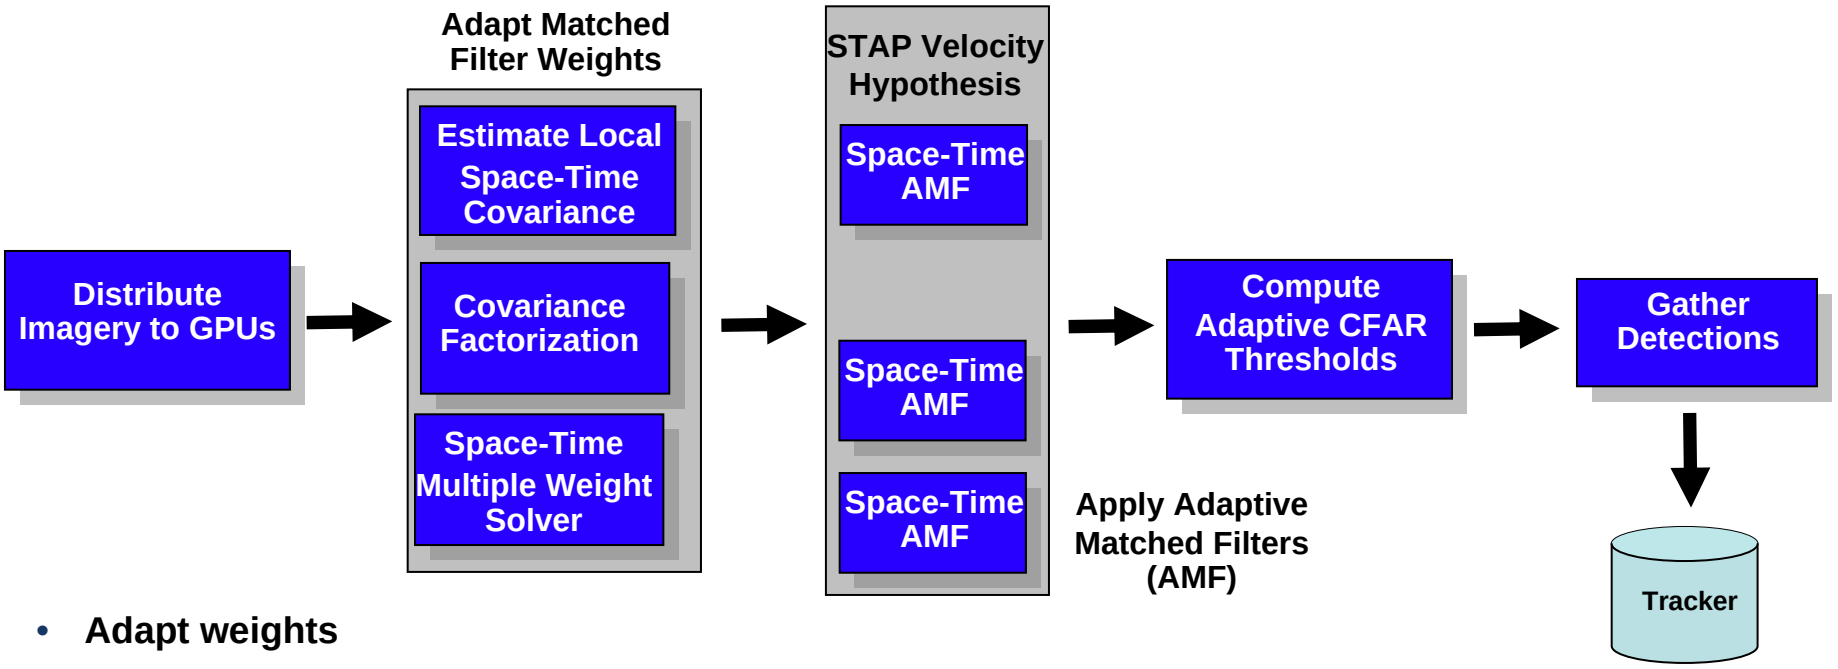


- **Adapt weights**
  - Covariance estimation
  - Covariance factorization
  - Find LMS solution
- **Apply AMF - convolution**
- **Compute CFAR thresholds**
  - Estimate standard deviation for each block
  - Adapt local standard deviations, excluding CUT
- **Gather detections into dense detection list**

# IRSTAP Kernels

- **Image preprocessing**
  - Registration
  - Clutter subtraction
  - Spatial demeaning
- **Covariance calculation**
  - Ensemble generation
  - Covariance calculation
- **Weight adaptation**
  - Cholesky Factorization
  - Back and forward substitution
- **Weight application**
  - Convolution
- **Adaptive CFAR threshold adaptation**
  - Localized standard deviation
  - Complex mean kernel
- **Detection reduction**

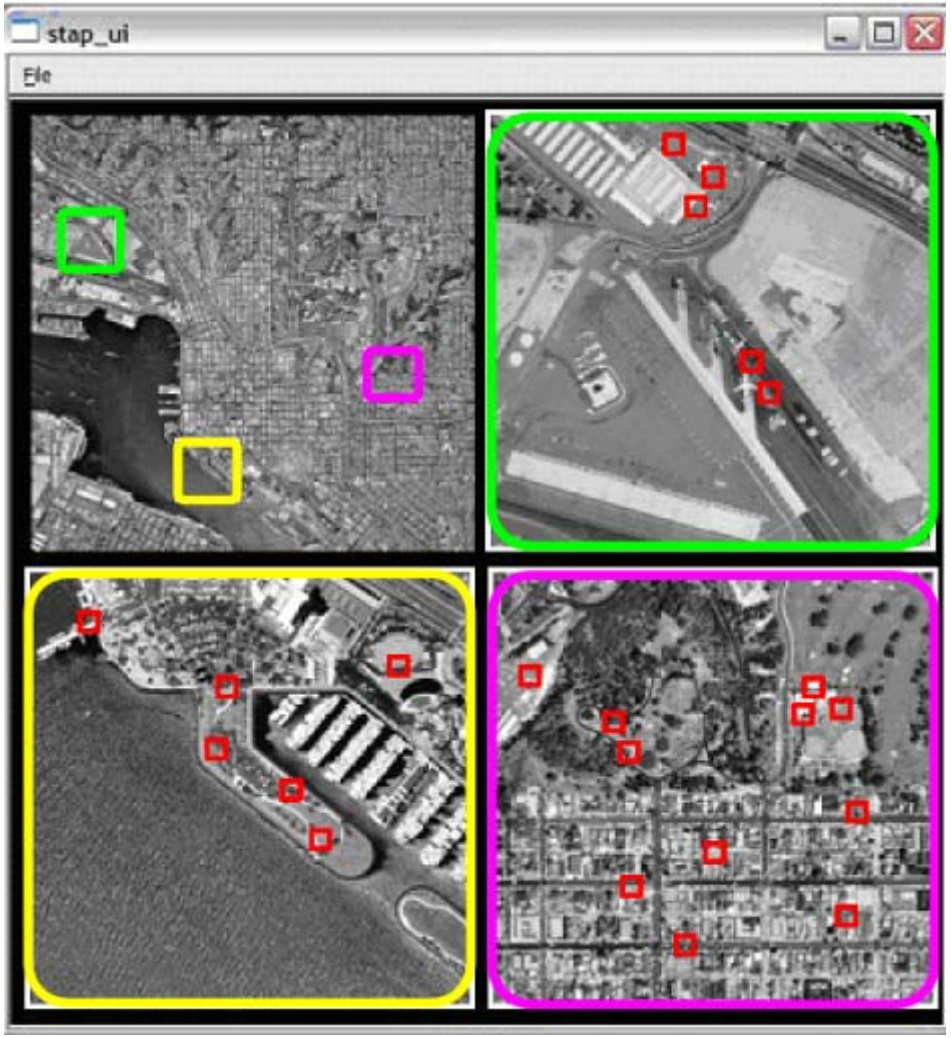

# IRSTAP Benchmarks GFLOP/s

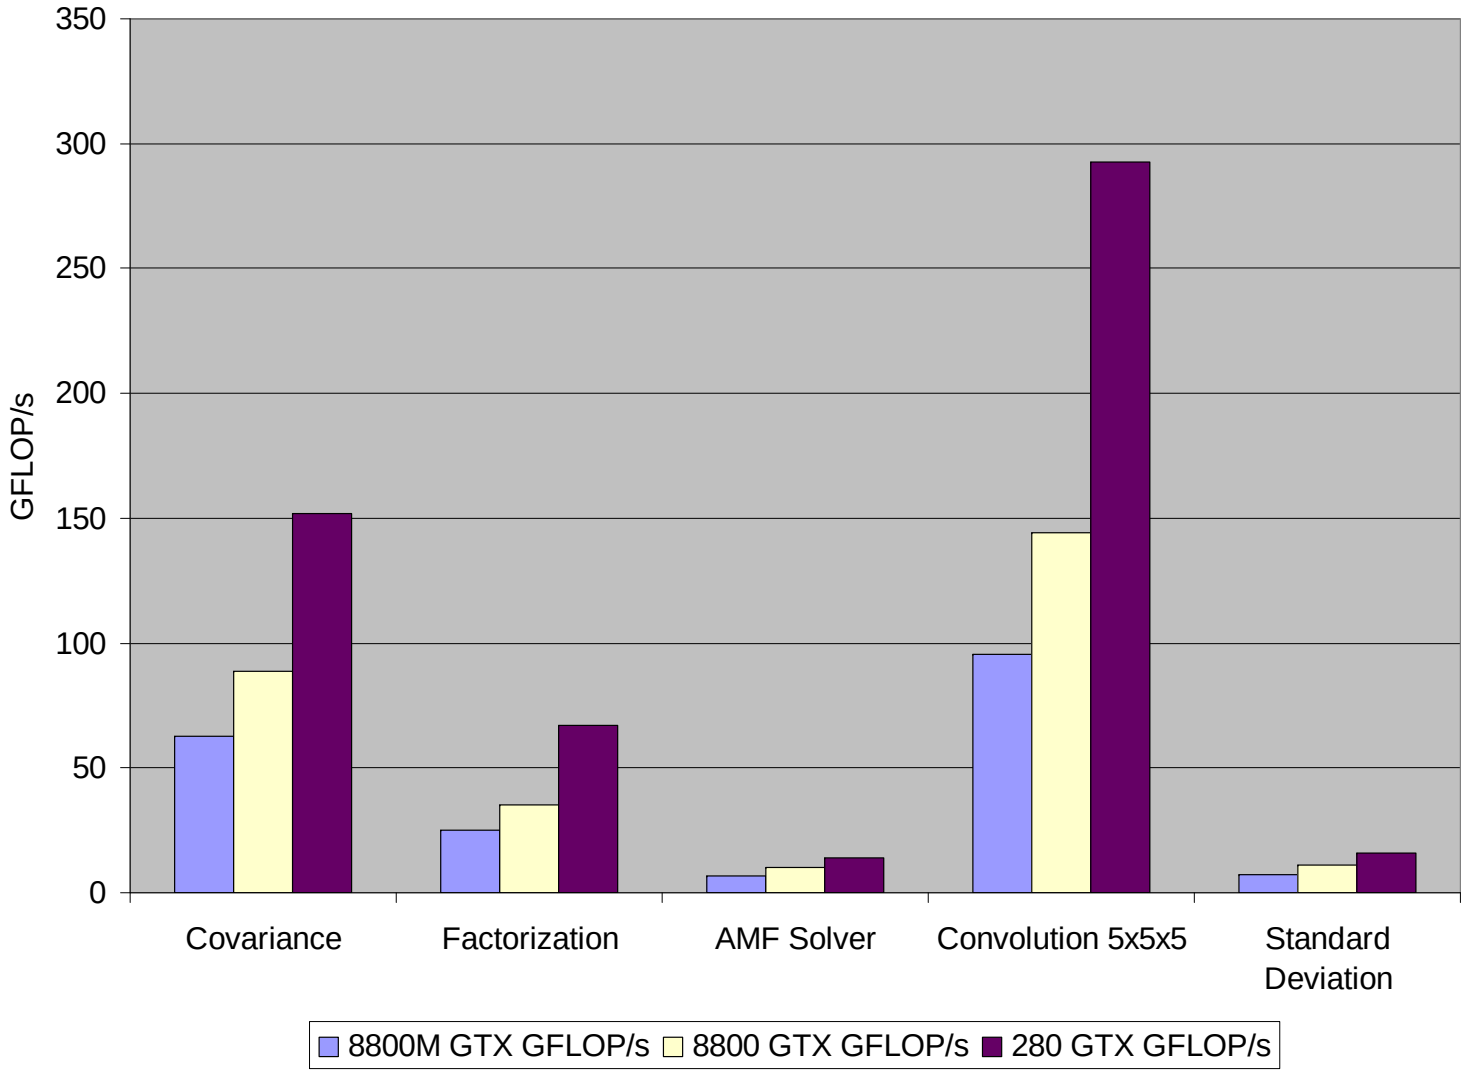

# IRSTAP Benchmarks GFLOP/W

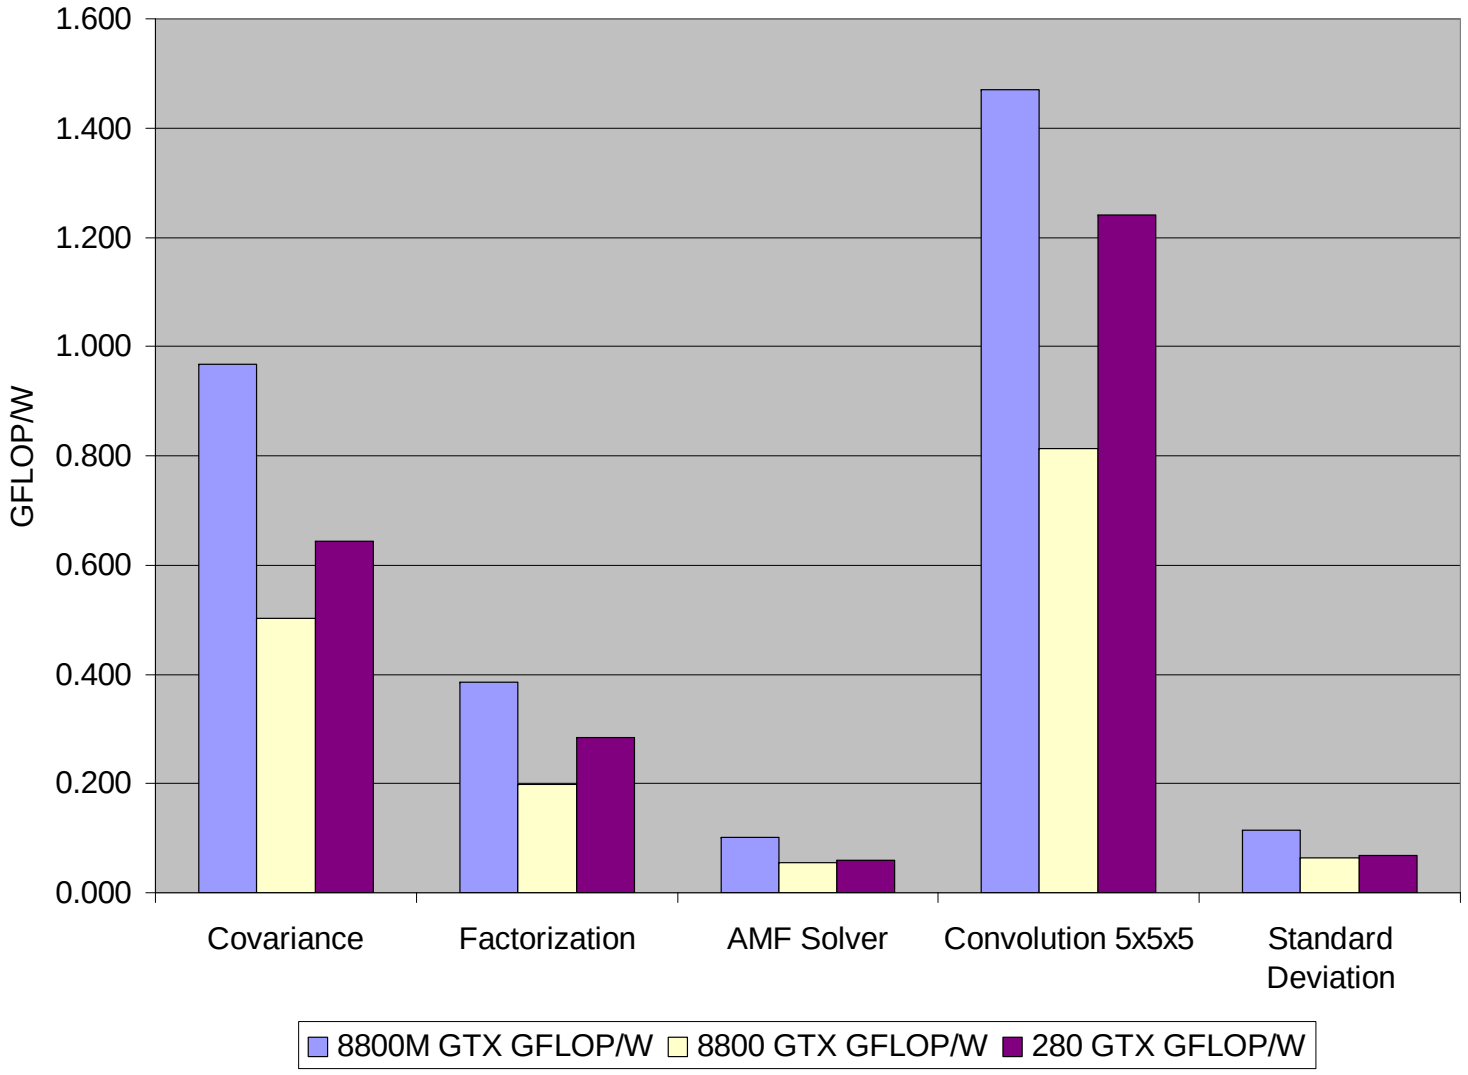

# Matrix Multiplication GFLOP/s

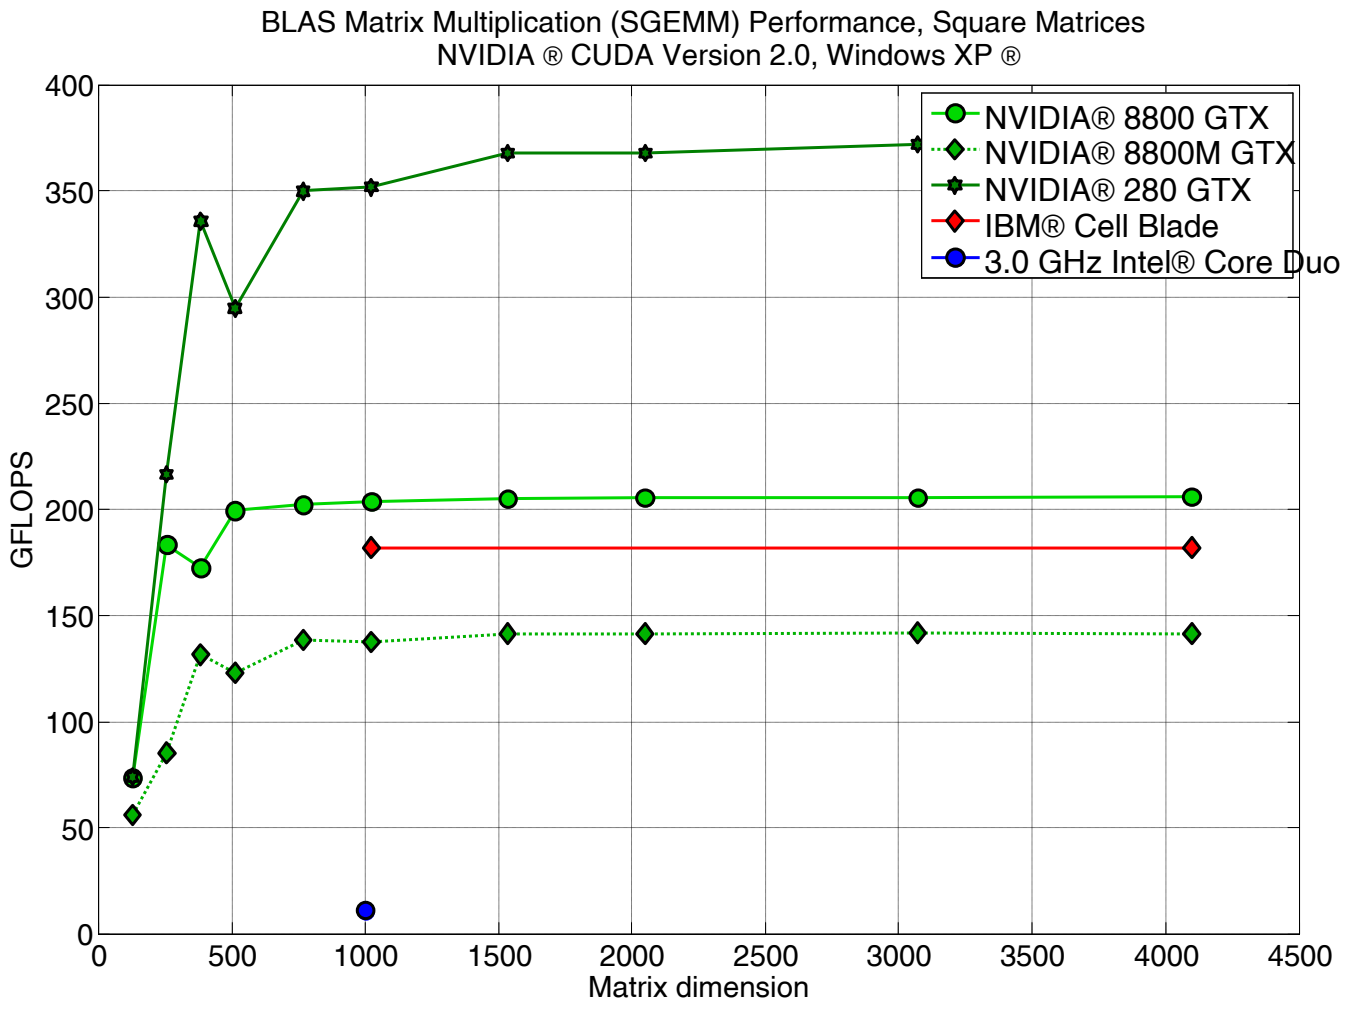


**GPU Multiplication is SGEMM is from the NVIDIA® library**

# Matrix Multiplication GFLOP/W

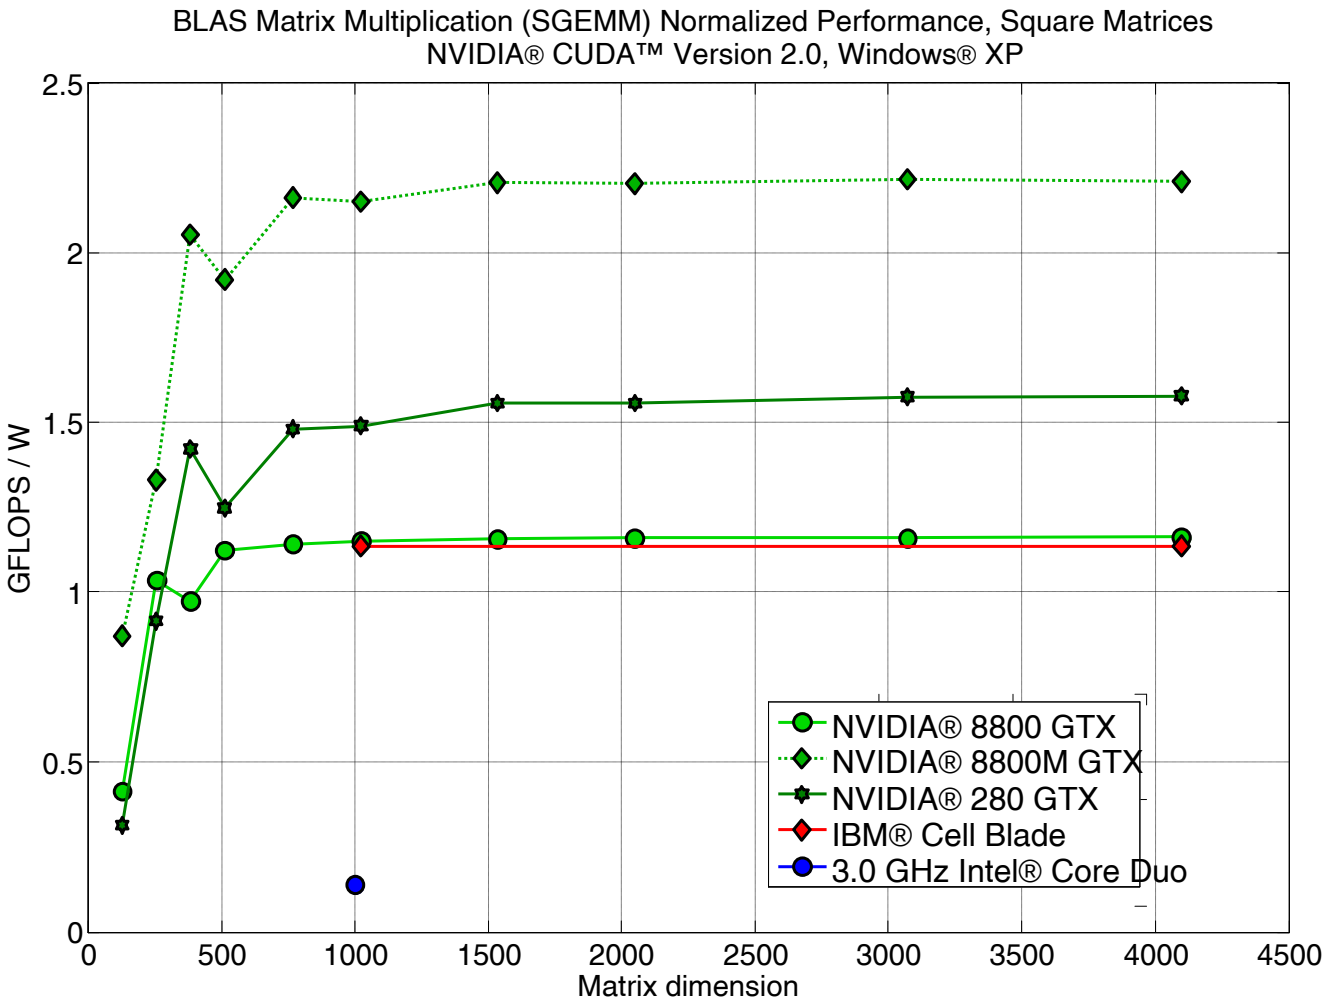

# 2D Complex FFT GFLOP/s

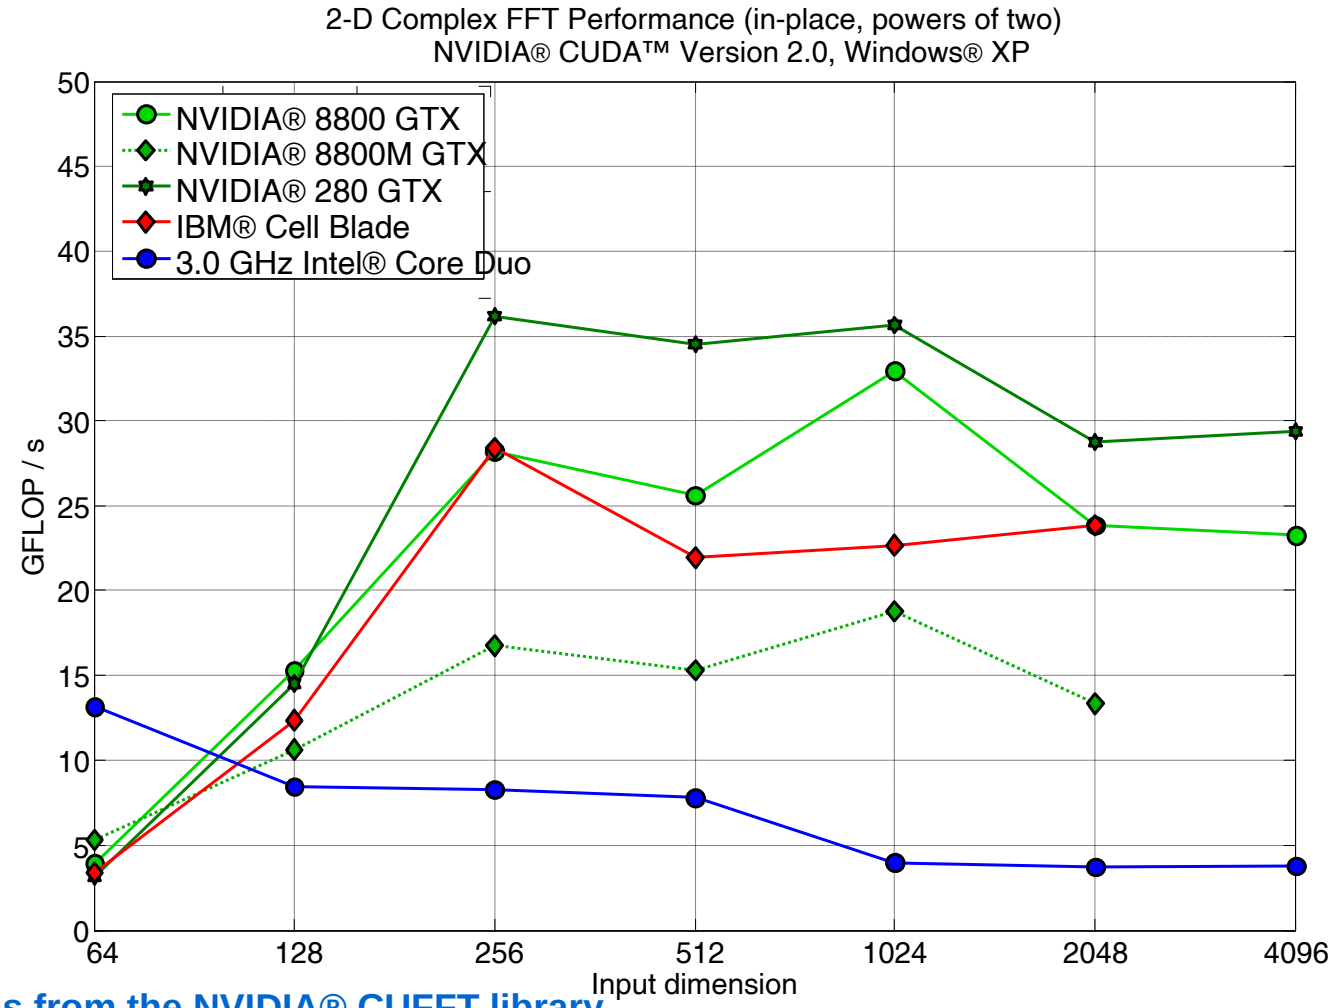


**GPU FFT is from the NVIDIA® CUFFT library**

# 2D Complex FFT GFLOP/W

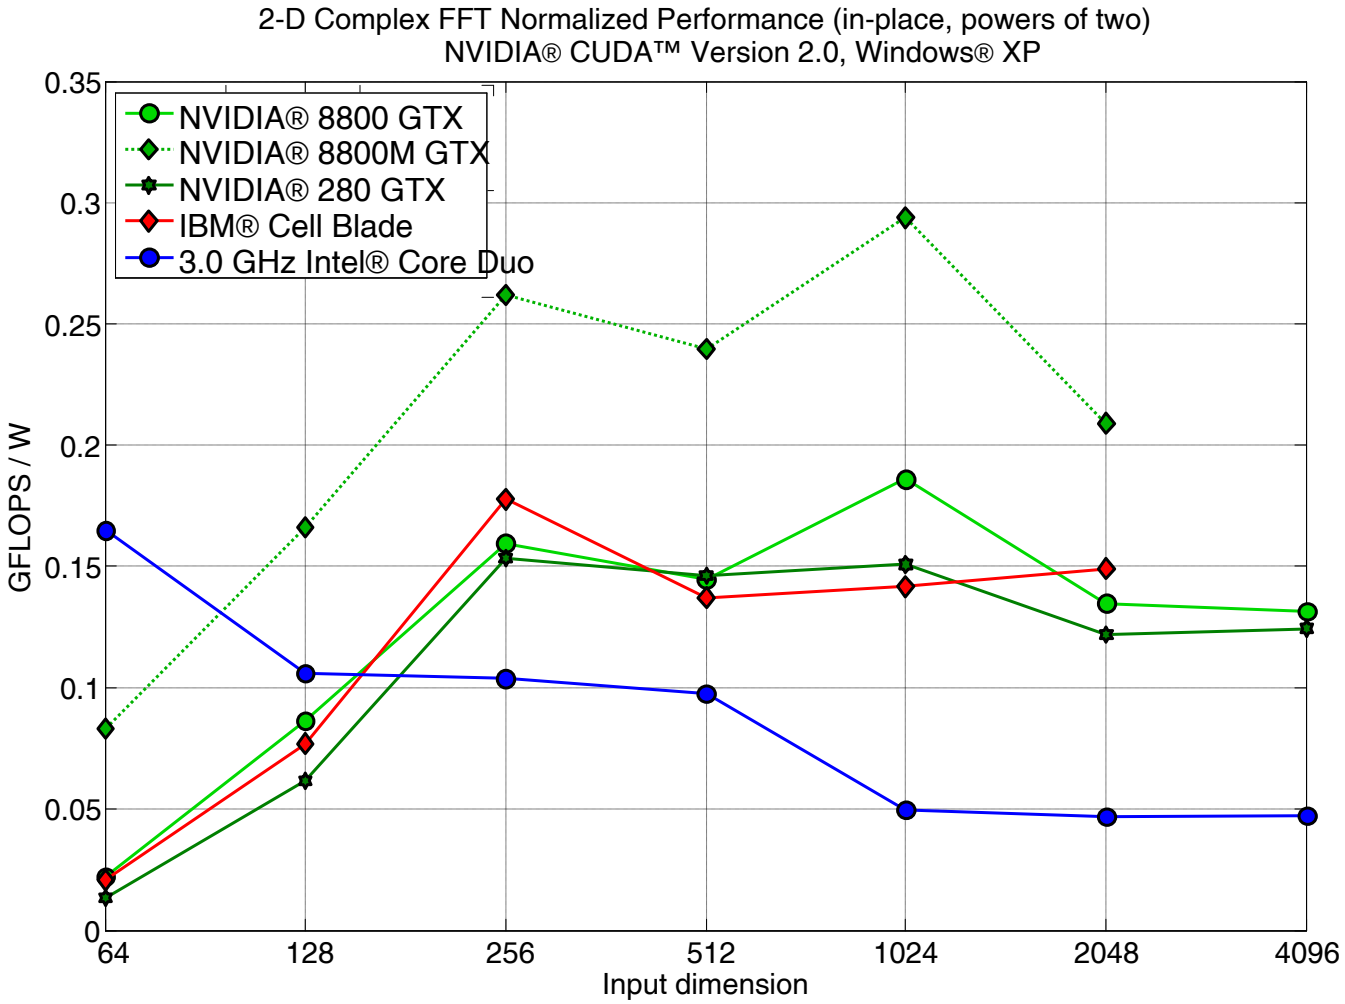

# Cholesky GFLOP/s

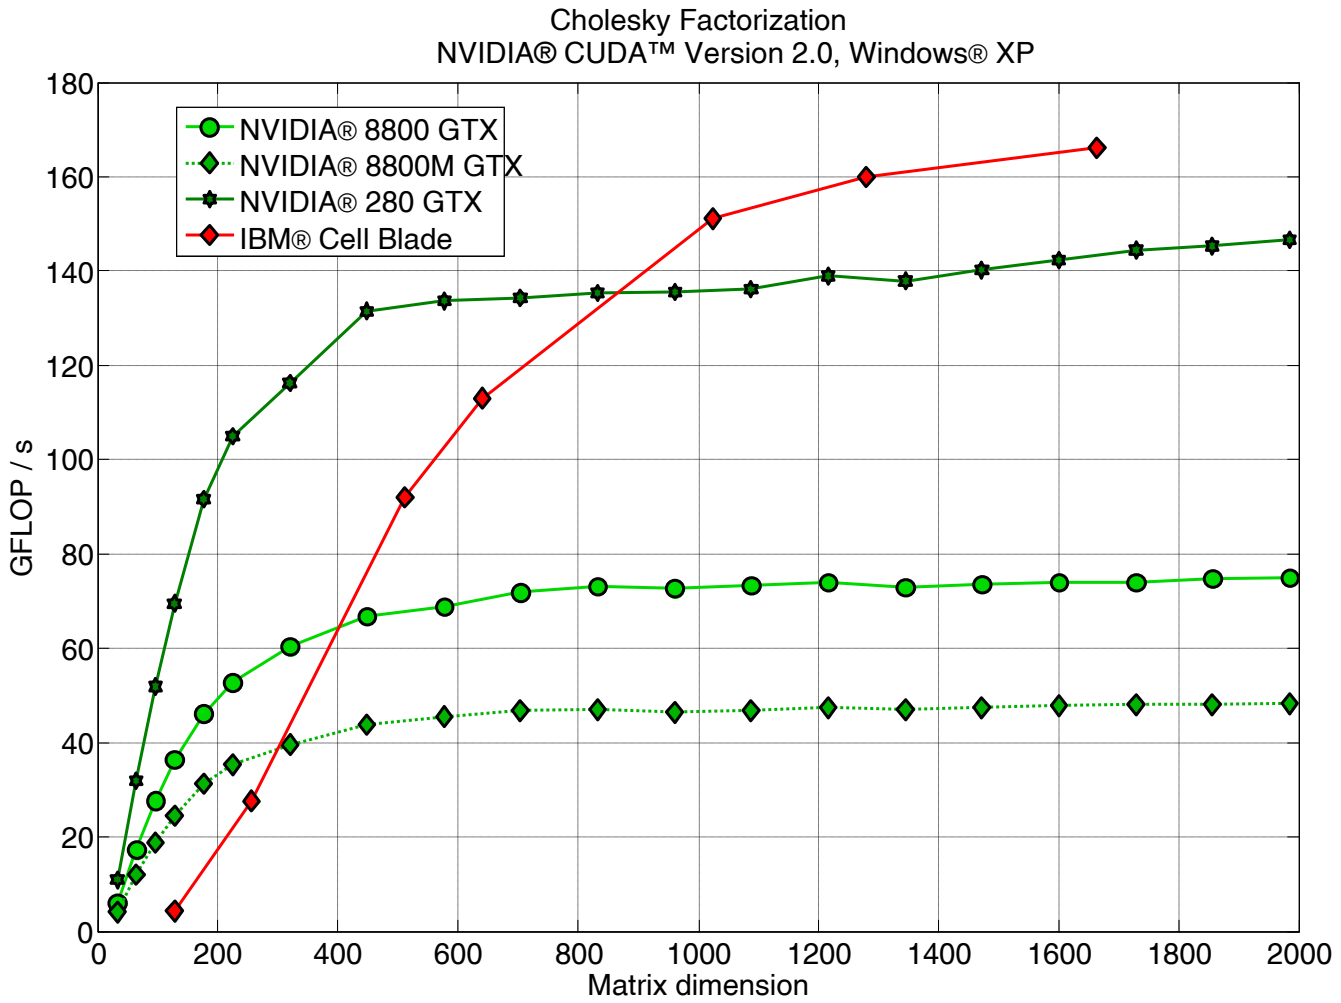

# Cholesky GFLOP/W

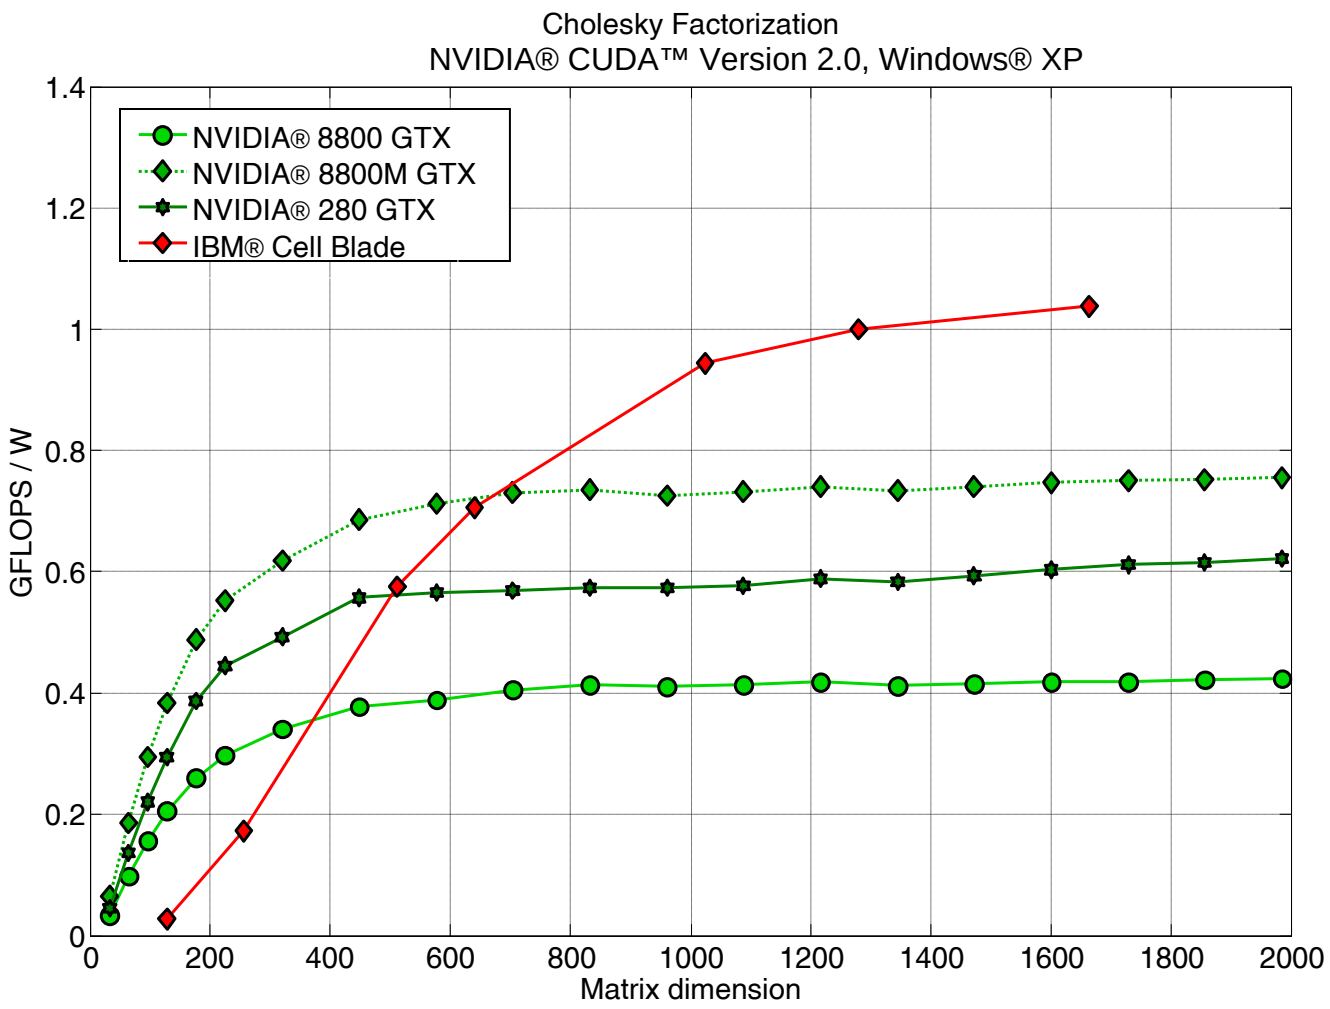

# Cholesky Efficiency

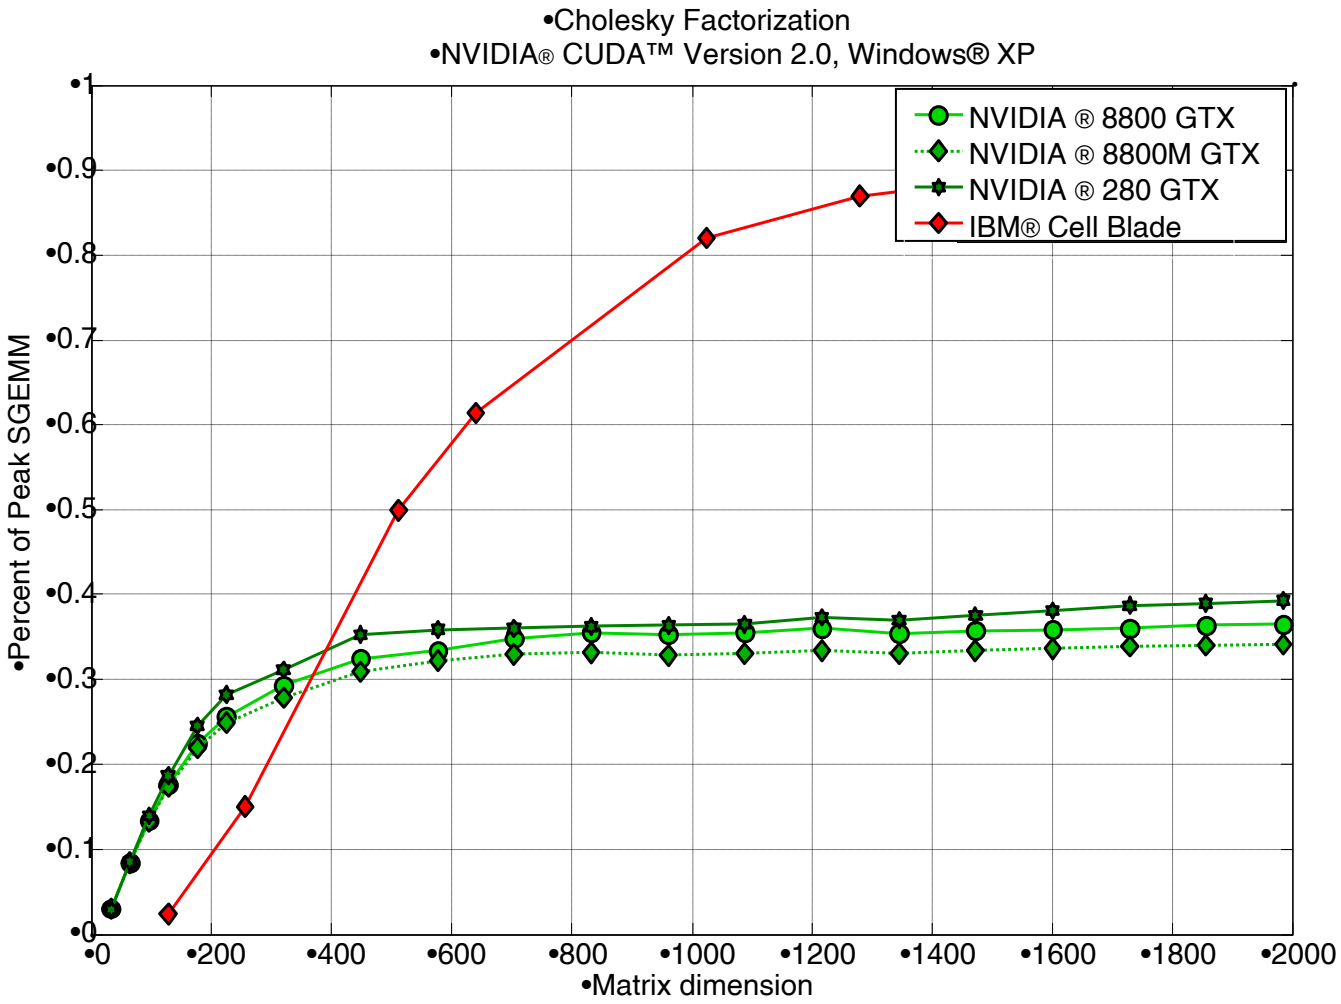

# 2D Convolution GFLOP/s

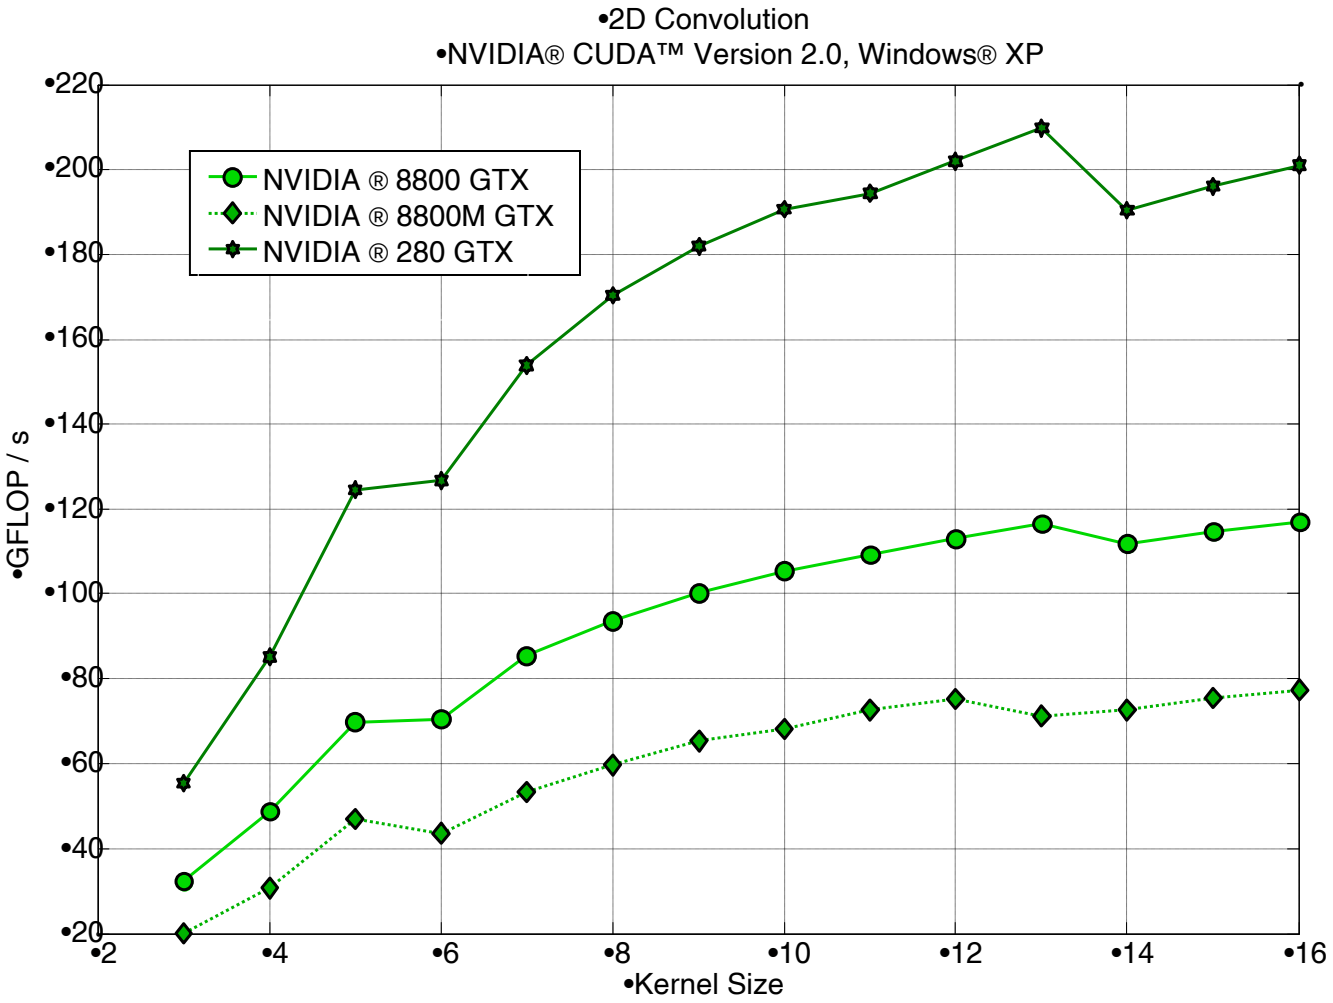

# 2D Convolution GFLOP/W

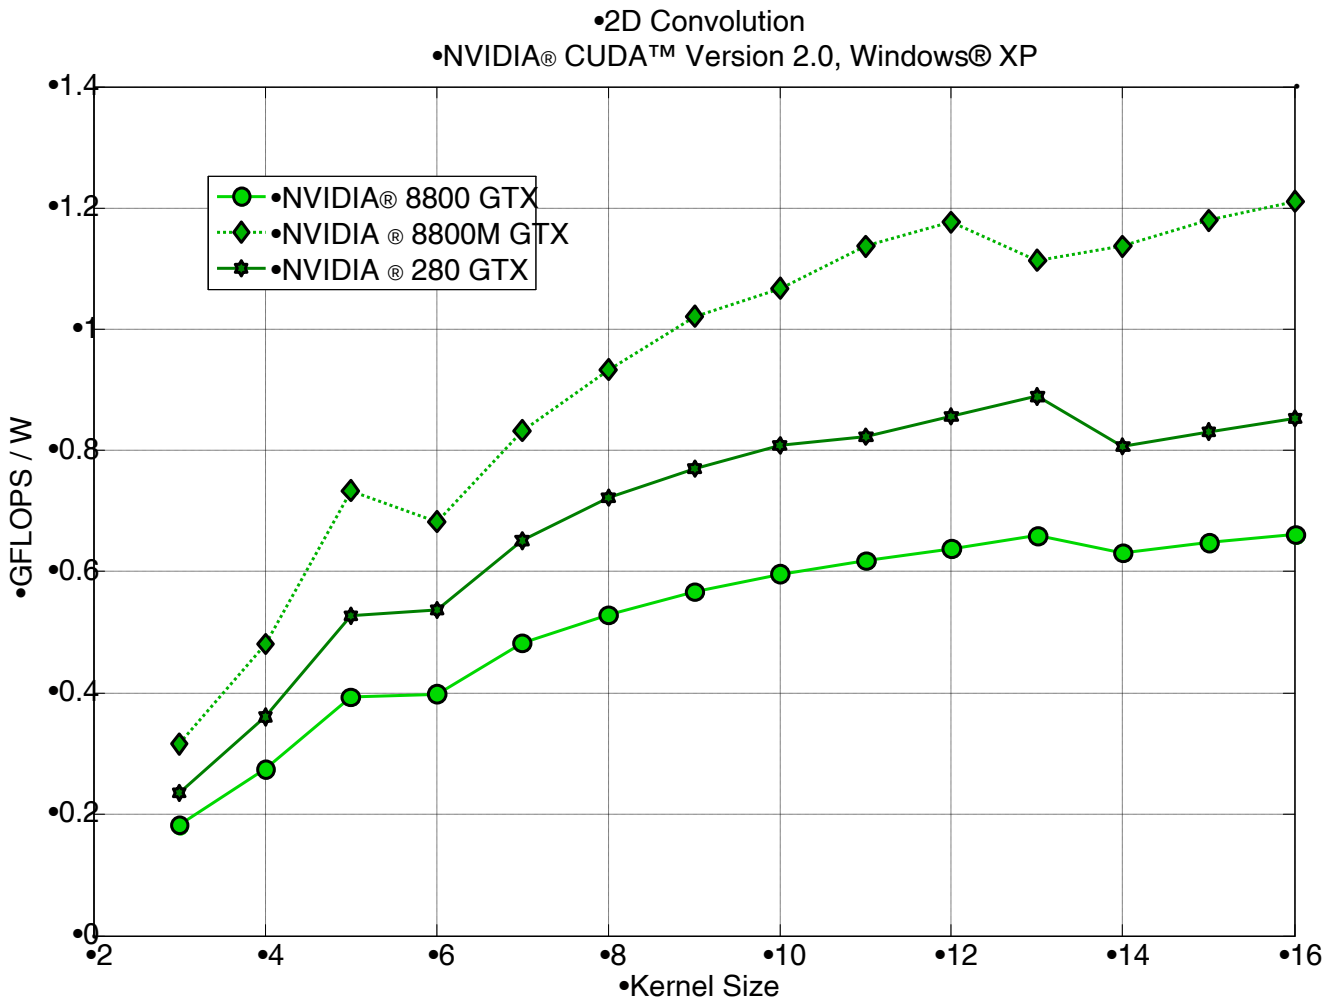

# Power Estimation

- **Using thermal design power (TDP) for each device**
  - NVDIA® 8800 GTX 128SP 1.35Ghz – TDP: 177W
  - NVDIA® 8800M GTX 96SP 1.25Ghz – TDP: 65W
  - NVDIA® 280 GTX 240SP 1.30Ghz – TDP: 236W
- **Cell – estimating TDP: 160W with 90nm process**

# Questions ?

# Just for comparison…

**Assume Cell Processor power consumption is approximately 160 W**

**2 Cell Processors per blade**

7000 ft at 32°C

**Power**

Cell Accelerator Board 2 with 4-GB DDR2 162W

Power is provided through the use of a single cable connector in addition to the 75W power provided through the PCI Express edge connector.

- Power consumption:
  - QS20: 315 watts maximum
  - QS20: 330 watts with 1 IB (#2945)
  - QS20: 345 watts with 2 IB (#2945)