Commercial GPU-based KID readout

A. Kovács
L. J. Swenson, C. M. Bradford, C. D. Dowell, M. Hollister, C. McKenney,
R. Monroe, H. G. LeDuc, H. Yoshida, J. Zmuidzinas

1 Components

Signal generator and digitizer
Pentek 71650

2x 12-bit ADC and 2x 16-bit DAC
0 – 250 MHz operating range
(125 MHz max. DAC bandwidth + image)

Cost: **$13,000**
TPD: **18 W**

+

GPU with compute cabability

CUDA (*nVIDIA*)
or OpenCL (*ATI + Intel*)

Cost: **$300 – $2000**
TPD: **30W – 80W**

+

PC host with two PCIe v2.0 x16 slots

Cost: **$1,000 – $6,000**
TPD: **50 – 250 W**

=

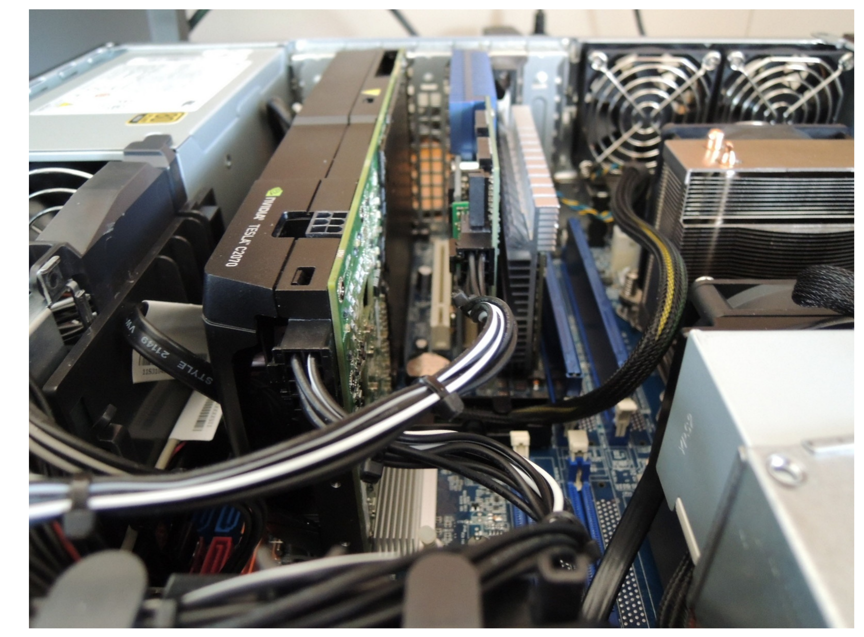

Integrated System

Cost: **$14,000 – $21,000**
TDP: **100 – 350 W**

=

No Compromise Spectal Processing

2000 – 4000 KIDs (2 readout lines)

0 – 250 MHz frequency coverage

Full FFTs (on 1 – 4 megasamples or more)

No spectral leakage

Unlimited tones and placement
(other than dynamic range considerations)

multiple waveforms

higher-level processing capability (e.g. chirps, resonance fitting or tracking)

Cost: $4 – $10 / channel
TPD: 25 mW – 200 mW / channel

2 Hardware Diagram

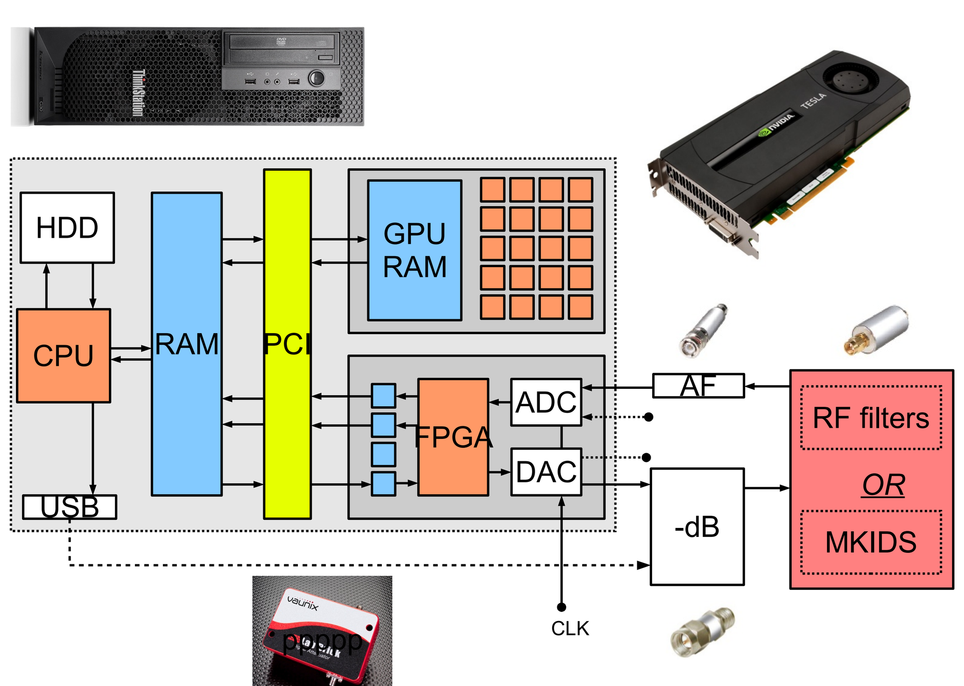

3 Data Bandwidth

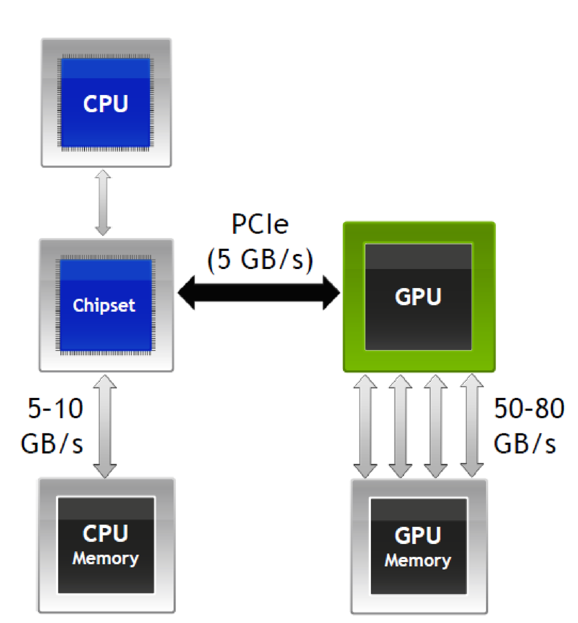

Current PCIe standard enables up to 2 GSPS sustained sampling (16-bit ADC). Enough for up to 4 KID lines with resonances up to 250 MHz.

PCIe throughput doubles every few years providing sufficient bandwidth for processing several GHz bandwidth directly, or processing more KID lines per readout system.

Within the GPU RAM bandwith is an order of magnitude higher.

4 PCIe Data Flow

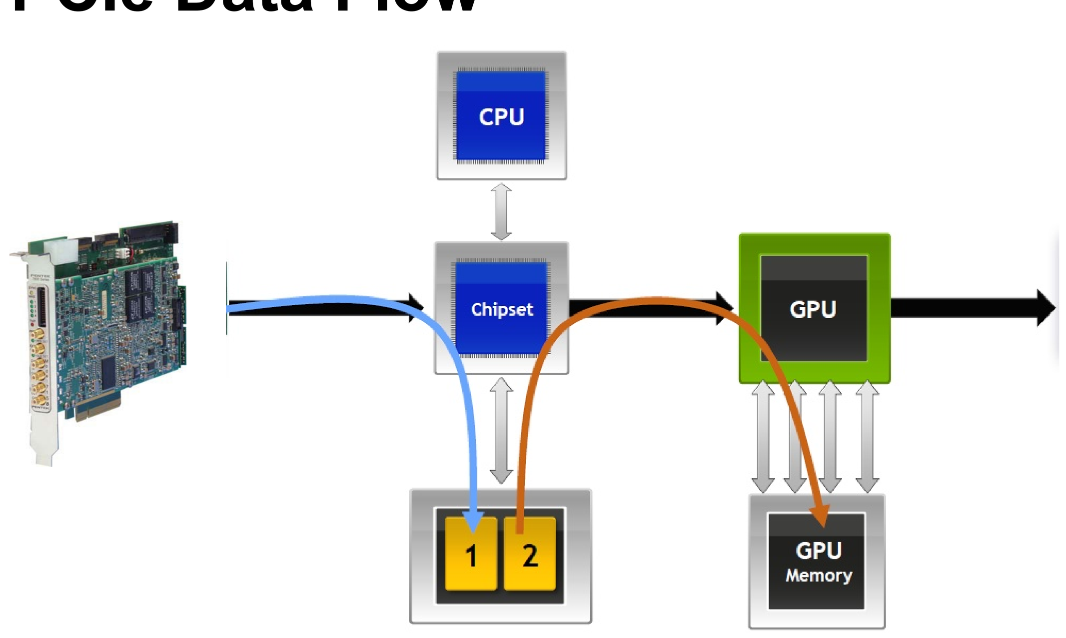

Currently, 2 PCIe transfers to get data to GPU (digitizer to RAM; RAM to GPU), but GPU allows for direct transfers from hardware that support it (currently only performance networking cards do, but perhaps digitizer cards will also in the future)

5 Spectral Processing Power

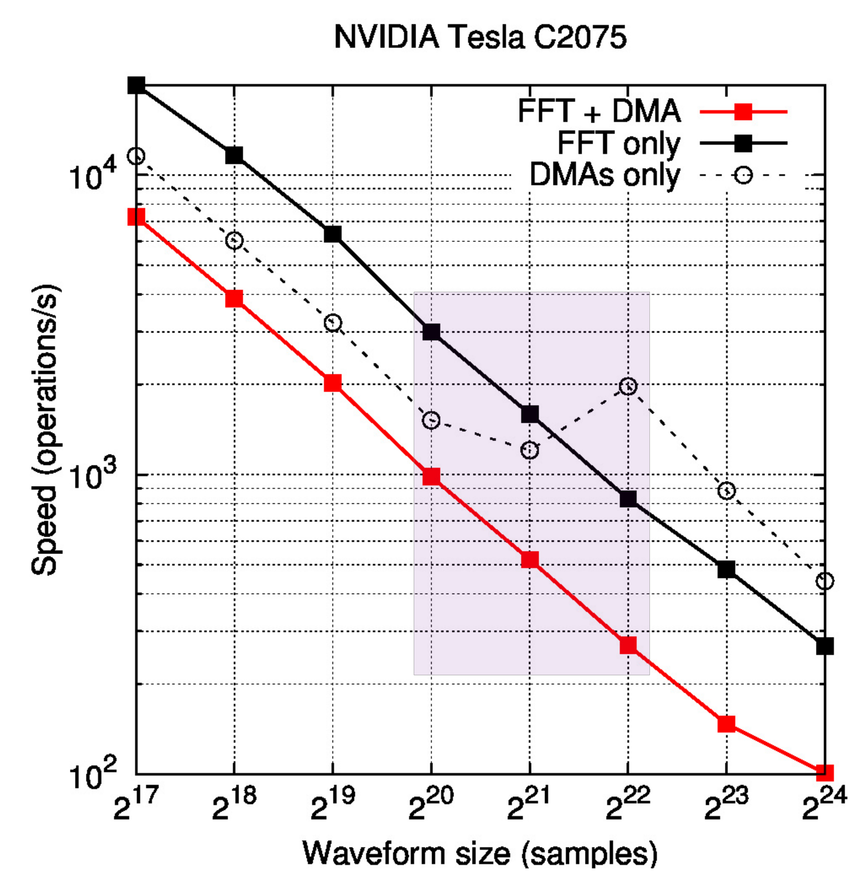

Q ~ 10^5 resonators desire waveforms with 1 – 4 M samples for tone placement. High-end GPU performance is limited by DMA transfers, not by FFTs. There is ample room for more complex processing (e.g. resonance fitting and/or tracking), and/or cheaper mid-range GPU with lower TPD. Latest 'Kepler' architecture is typically 3x faster than the Tesla tested above.

7 Noise Performance

100 MHz (200 MSPS)
2 tones ~ 95 MHz
Rubidium clock reference

NEC ~ 0.0025 counts / Hz$^{1/2}$
(on 16 bits!)
S/N ~ 100 dB/Hz (1000 tones)

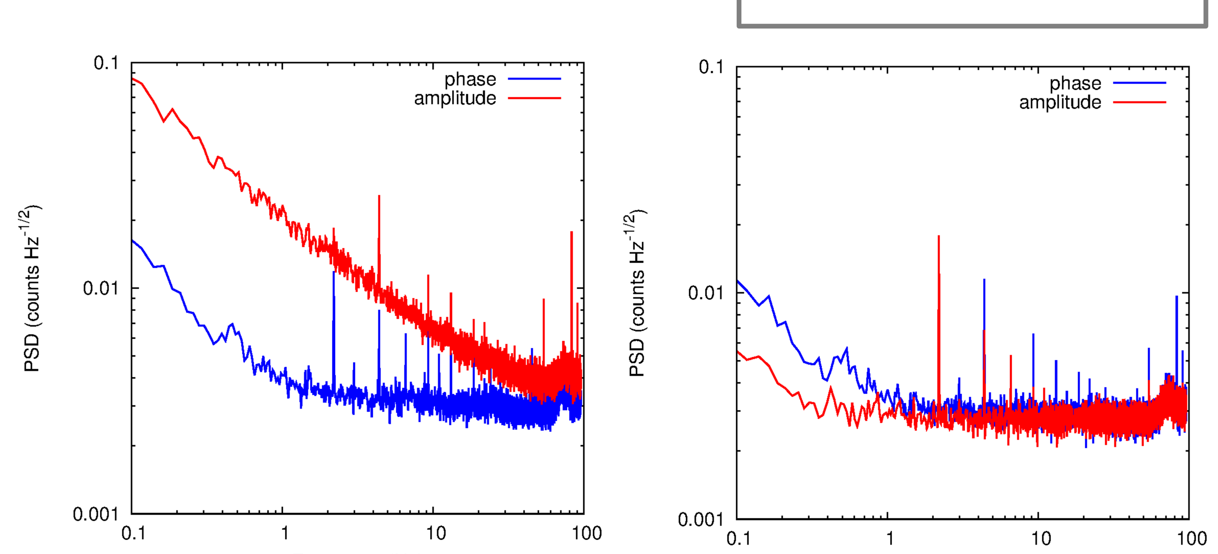

(*raw*) (*blind-tone calibrated*)

A practical limit on the number of tones used is set by the ADC dynamic range. To reach a signal-to-noise ratio of 100 dB/Hz (typical requirement for mm-wave KIDs) this means a around 1000-2000 tones (or resonators) per octave and readout line. Ultimately, better ADCs can support more tones and/or detectors.

6 Software Diagram

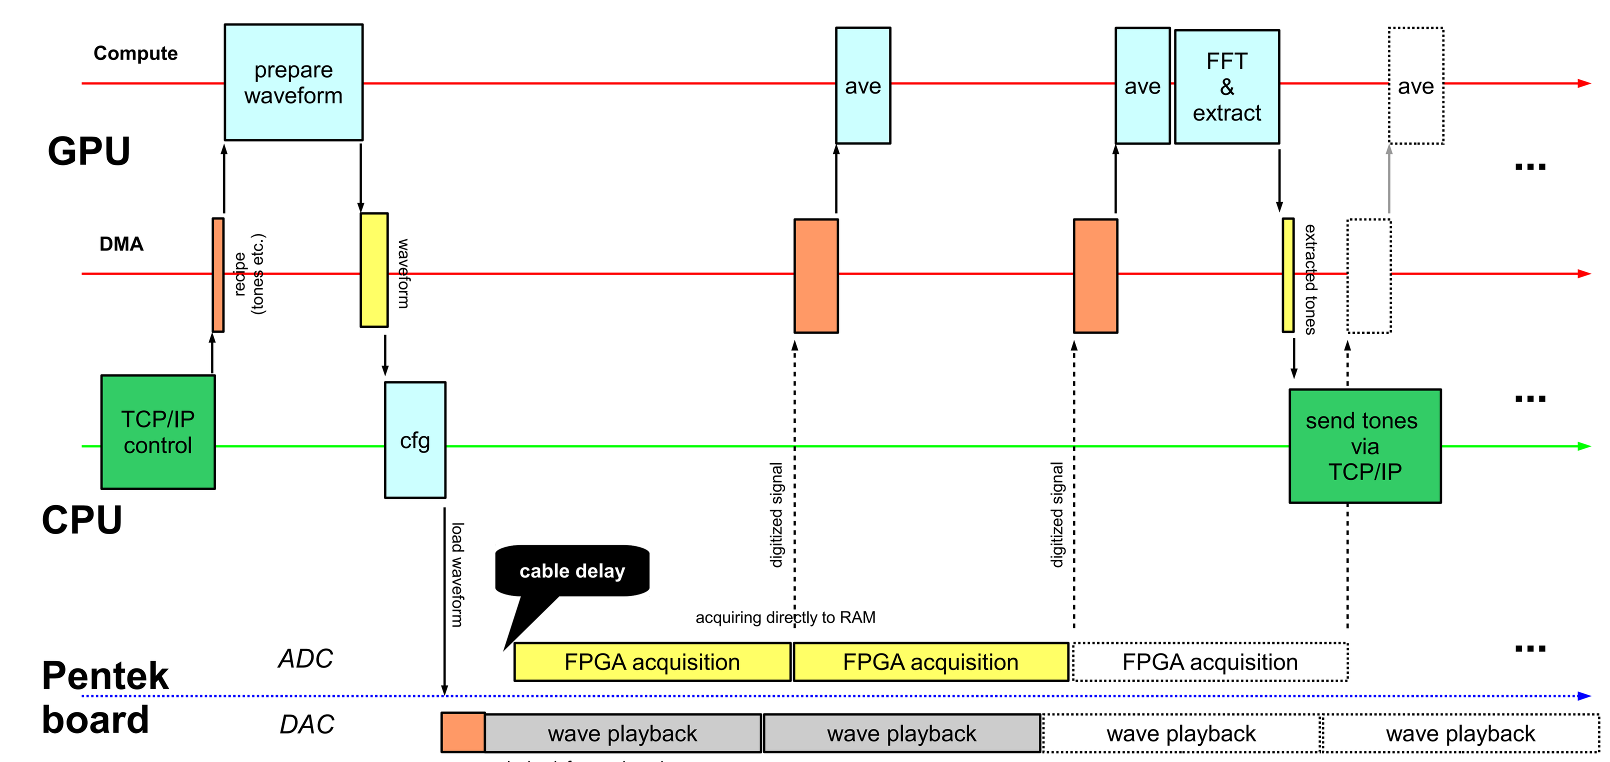

Digital acquisition via C/C++ libraries
GPU programmed in CUDA (nVIDIA) or OpenCL (all) – both C++ derivatives.
Control software in language of choice (not resource intensive)

8 Conclusions

Pros
Commercial GPU-based readout architecture offers a **powerful** and **flexible** spectral processing platform, especially for kinetic inductance detdectors (KIDs). C/C++ based development should mean fast turnover, great for **exploring** and **optimizing** readout strategies and algorithms.

Cons
Expensive and **power hungry** when compared to FPGAs (see poster by R. Monroe) but not prohibitively so.

8 Future Outlook

Technological progress means **cheaper and less power-hungry future implementations** of current capabilities and/or increasing processing power. We can anticipate a halving of both cost per detector and power dissipation per detector. Expect **$0.25 – $1 and 1.5 – 12 mW per channel by 2020**...

The expensive Pentek acquisition hardware may have cheaper future alternatives, or data can be digitized at the detector, and fed to the GPU via optical fibre, with optimized 1-step PCIe transfer. Costs could drop to **below $0.10 / channel**.