HELMHOLTZ ASSOCIATION

Work package 2:

Real-time Data Processing

Andreas Kopmann (KIT)

HDRI / PanData Workshop 4.-5.3.2013, DESY

DESY

GSI

HZB Helmholtz Zentrum Berlin

Helmholtz-Zentrum Geesthacht
Zentrum für Material- und Küstenforschung

KIT Karlsruhe Institute of Technology

Real-time Data Processing: Overview

HELMHOLTZ ASSOCIATION

PCIe Acceptable data rates

1 DAQ hardware

2 PC-based processing

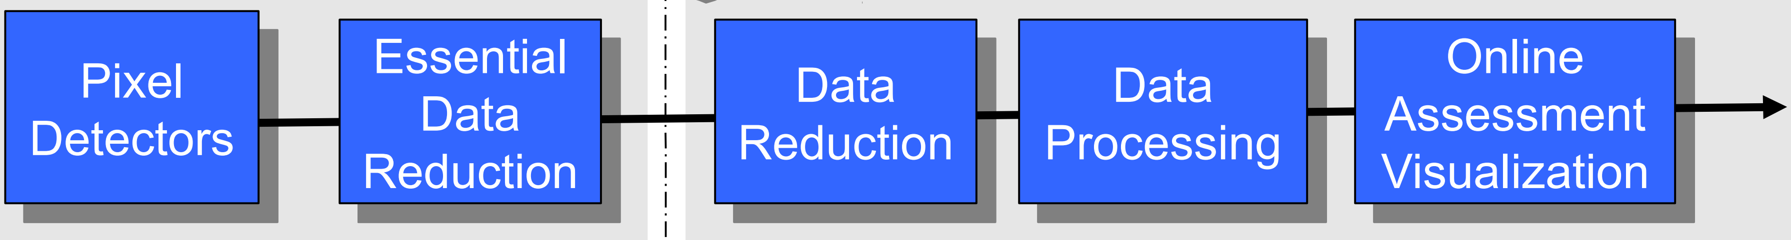

Technologies

Programmable hardware

FPGA, DSP, embedded GPUs

„Parallel Computing“:

PC + Coprocessing / GPU

Tasks

Selected commercial hardware
Standardized custom hardware

Hardware independant programming (e.g. OpenCL)
Library of common algorithms

Task 1: Dedicated Hardware

HELMHOLTZ ASSOCIATION

Actions:

- Survey of available solutions for programmable signal processing hardware
- Adopt existing solutions to the needs of PNI or join a running HGF or other development
- Introduce the solutions in PNI

Flexible high-throughput FPGA platform

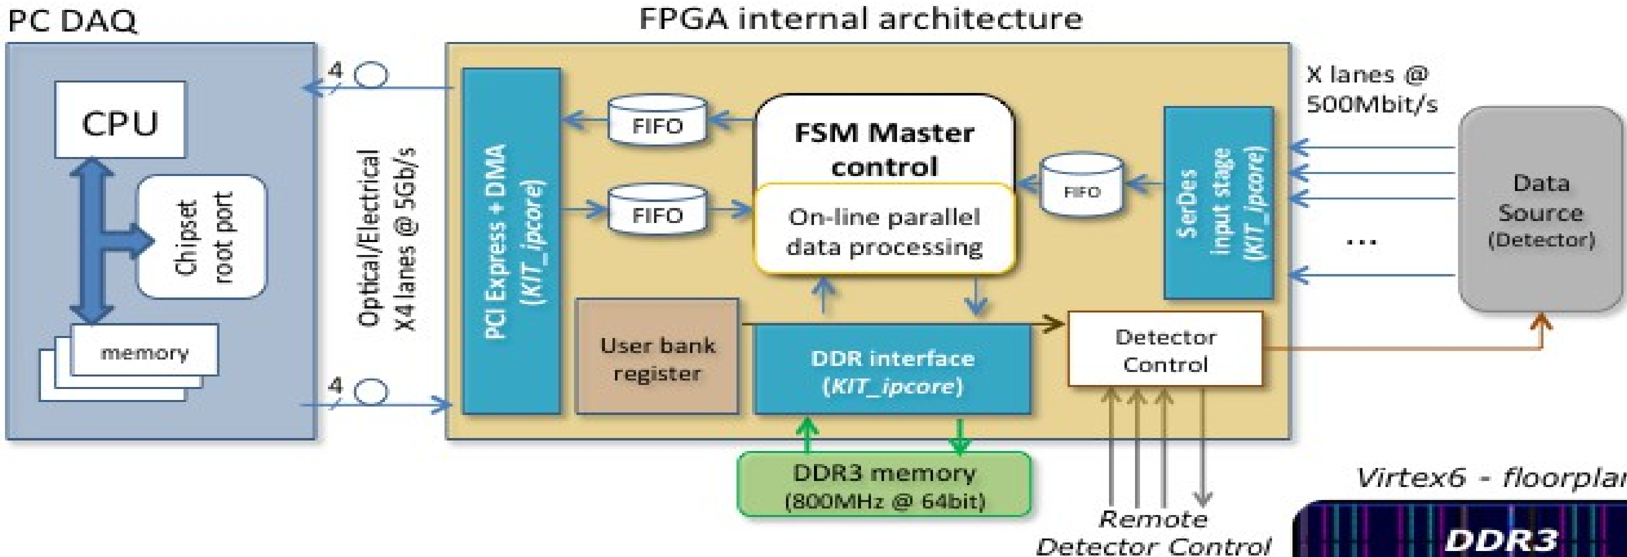

Virtex6 - floorplan

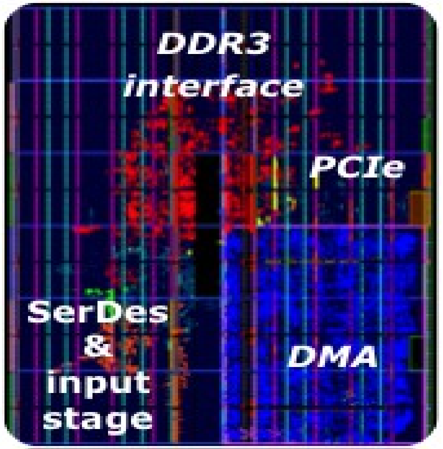

- ✓ Three logic cores have been developed for a flexible high-throughput platform
  - ✓ PCIe-Bus Master DMA readout architecture
  - ✓ Multi-port high speed DDR3 interface
  - ✓ Configurable 2..16 bits "SerDes" (Serializers /Deserializers) architecture
- ✓ PCI Express/DMA Linux 32-64 bits driver with ring buffer data management
- ✓ Integration in the parallel GPU/CPU computing framework
  - 64-bit Linux drivers

PCIe DMA IP-core

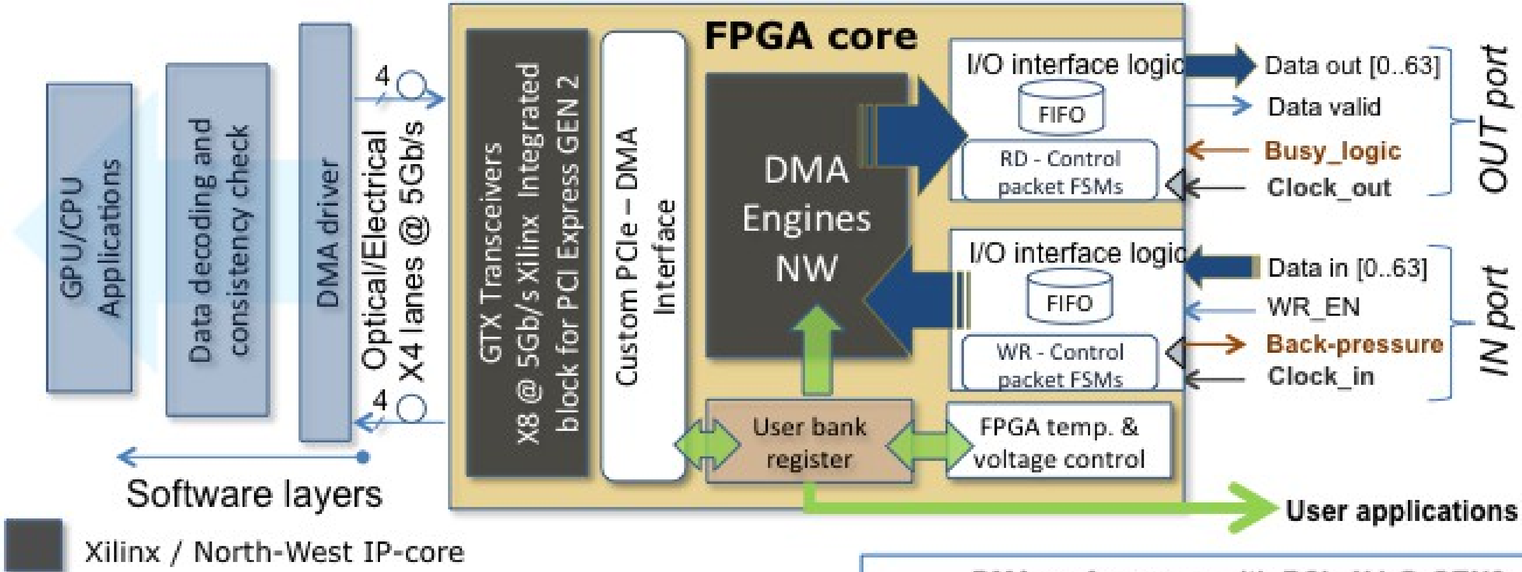

- ✓ Bus Master DMA operating with 4lanes PCIe @ Gen2 (250MHz)
- ✓ Two individual engines for write/read from FPGA (User logic) to PC centre memory
- ✓ IN and OUT FIFO-like interface (for User logic)
- ✓ FIFO used to decouple the time domain between DMA and User custom logic

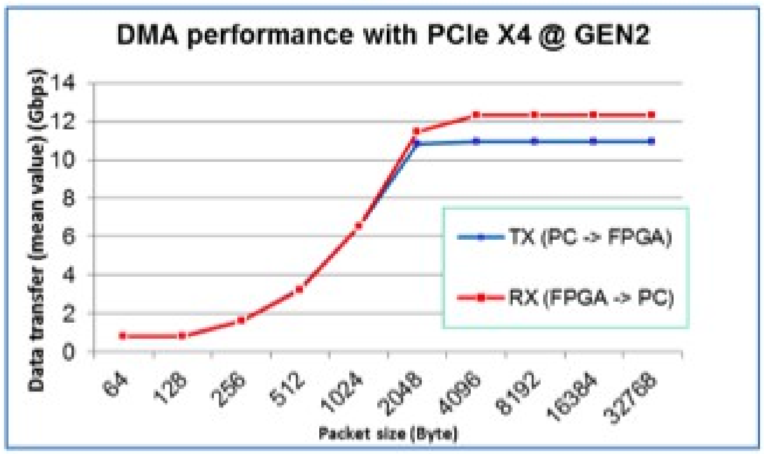

PCIe x8 up to 24Gb/sec + royalty-free IP-core under development

FPGA IP-cores

HELMHOLTZ ASSOCIATION

Disadvantages of commercial IP-cores :

- Expensive
  - North-West DMA → 20K€ (only netlist) and 35k€ (Source code)
  - PLDA → 10K€ (for netlist) and 60k€ for (source code) EZDMA/QuickPCIe
- Limited to unique FPGA family (e.g. Virtex 6, speed grade -2)

Portfolio detector technology starts project database

- 1 GbE UDP

detector technology

Collection of IP-Cores

The information of commonly available resources is collected on the wiki page

The sources will not be available on the wiki page, they can be get via the contact person

The structure of the entries in the databases is:

- Name of the project
- Type (IP core, FPGA project, FMC card, …)
- Description
- Platform (e.g. VERTEX 4)
- Status
- Contact person

High-throughput camera platform

Prototype

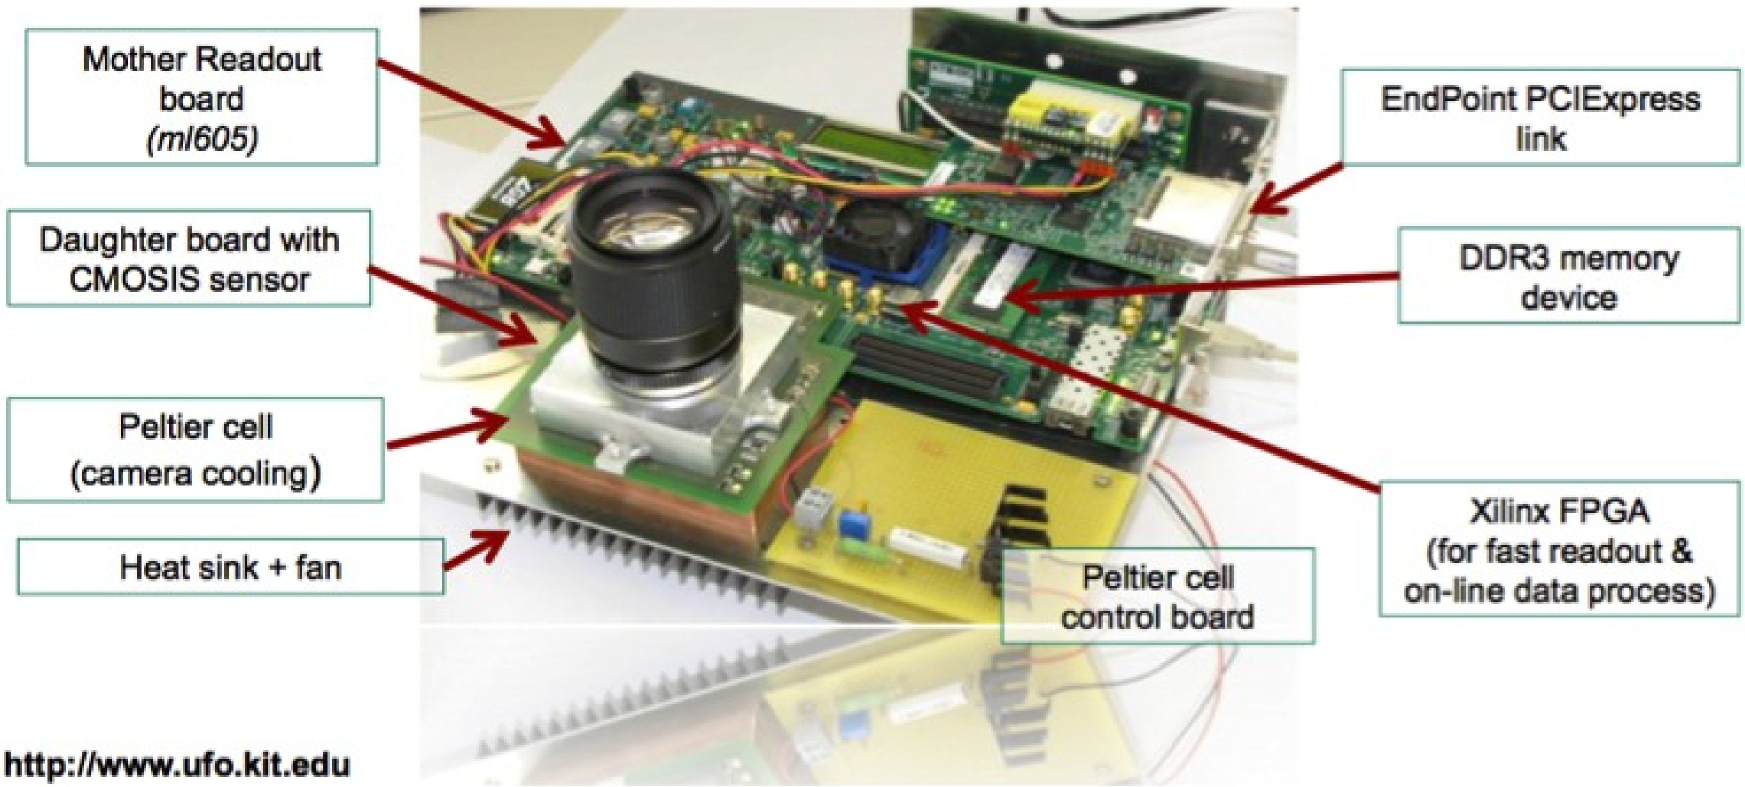

http://www.ufo.kit.edu

Application: Grating-based phase-contrast imaging

Improve the detector electronics: UFO - project

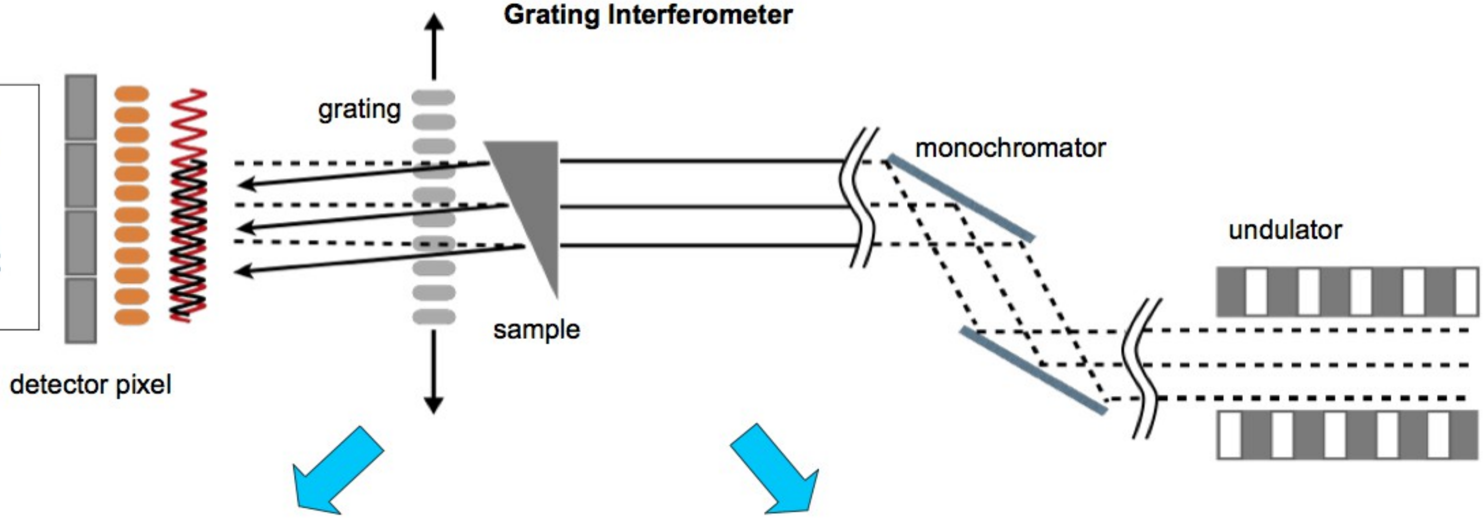

Current data acquisition:

- Set of interference patterns at fixed scanning positions needed
- Accurate piezo positions required (time consuming)
- Increase of number of scanning position (large raw data volume)

Optimized data acquisition:

- Scanning all the possible positions
- Data pre-processing by the intelligent detector

The new intelligent detector will allow us to vary:

- Statistics of a camera (CCD or CMOS)
- Phase-step parameters (equidistant and arbitrary and number of steps) according to the users needs

Piezo Drive:

- Directly connected with detector electronics.
- Allows control of the phase-step size during the experiment

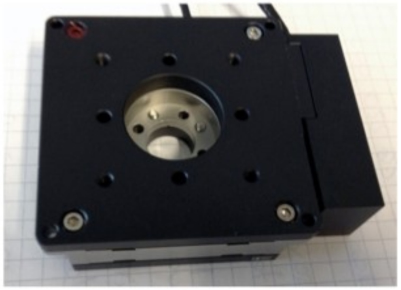

HZG + KIT

Application: Hot-electron bolometer (HEB)

Concept
- Signal splitter → 4 signal
- 4 fast ADC + Pico sec delay

Characterization of 1 branch with 500MHz Pulses

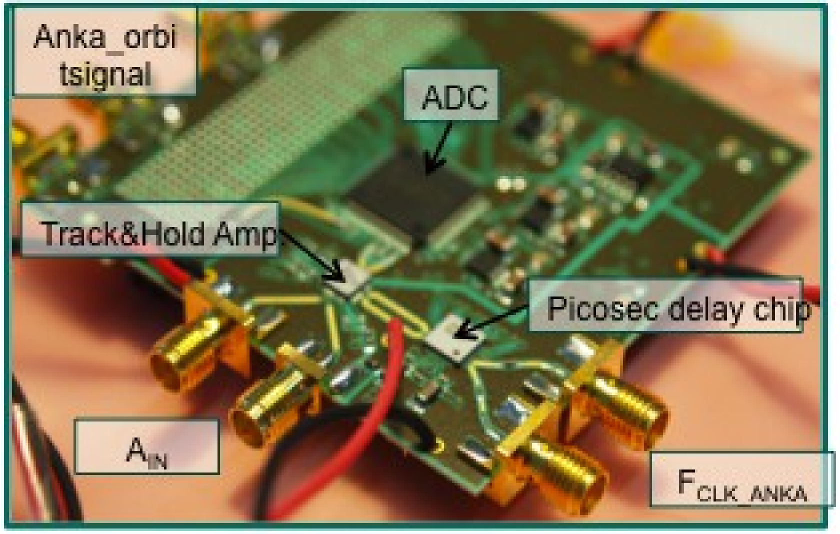

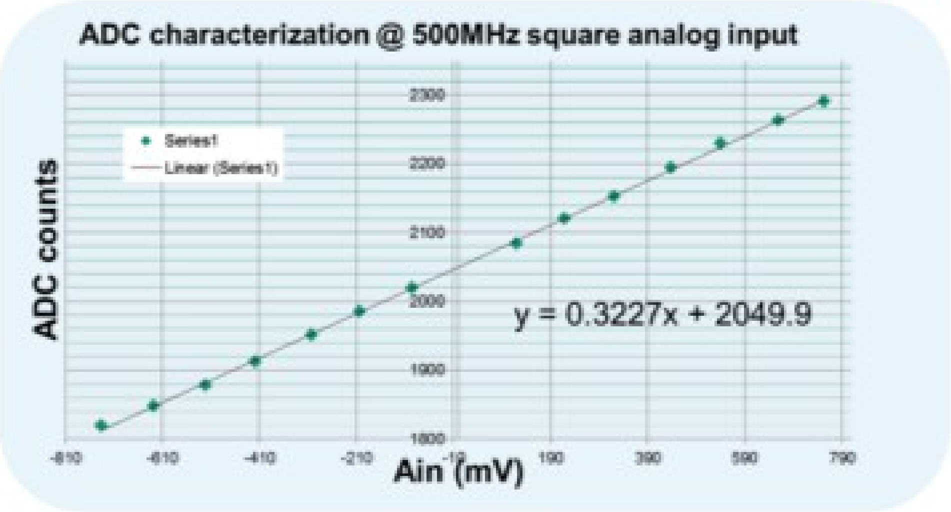

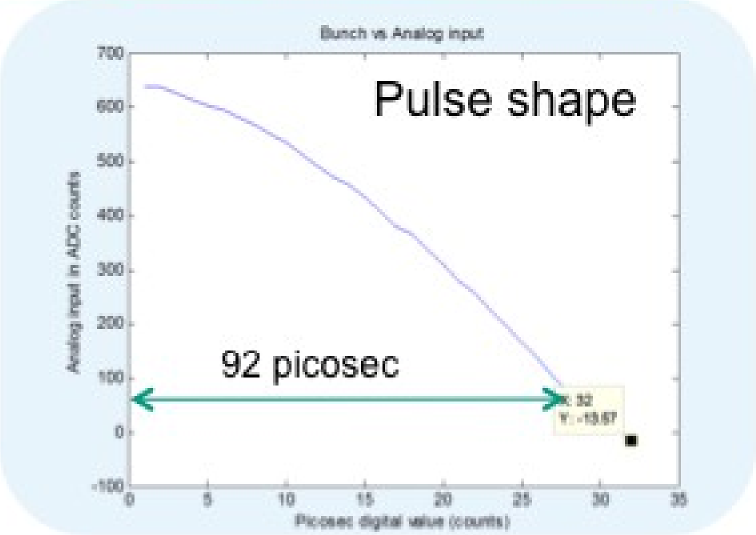

Application: HEB – ANKA long time bunch behaviour

Results with YBCO – terahertz detector:

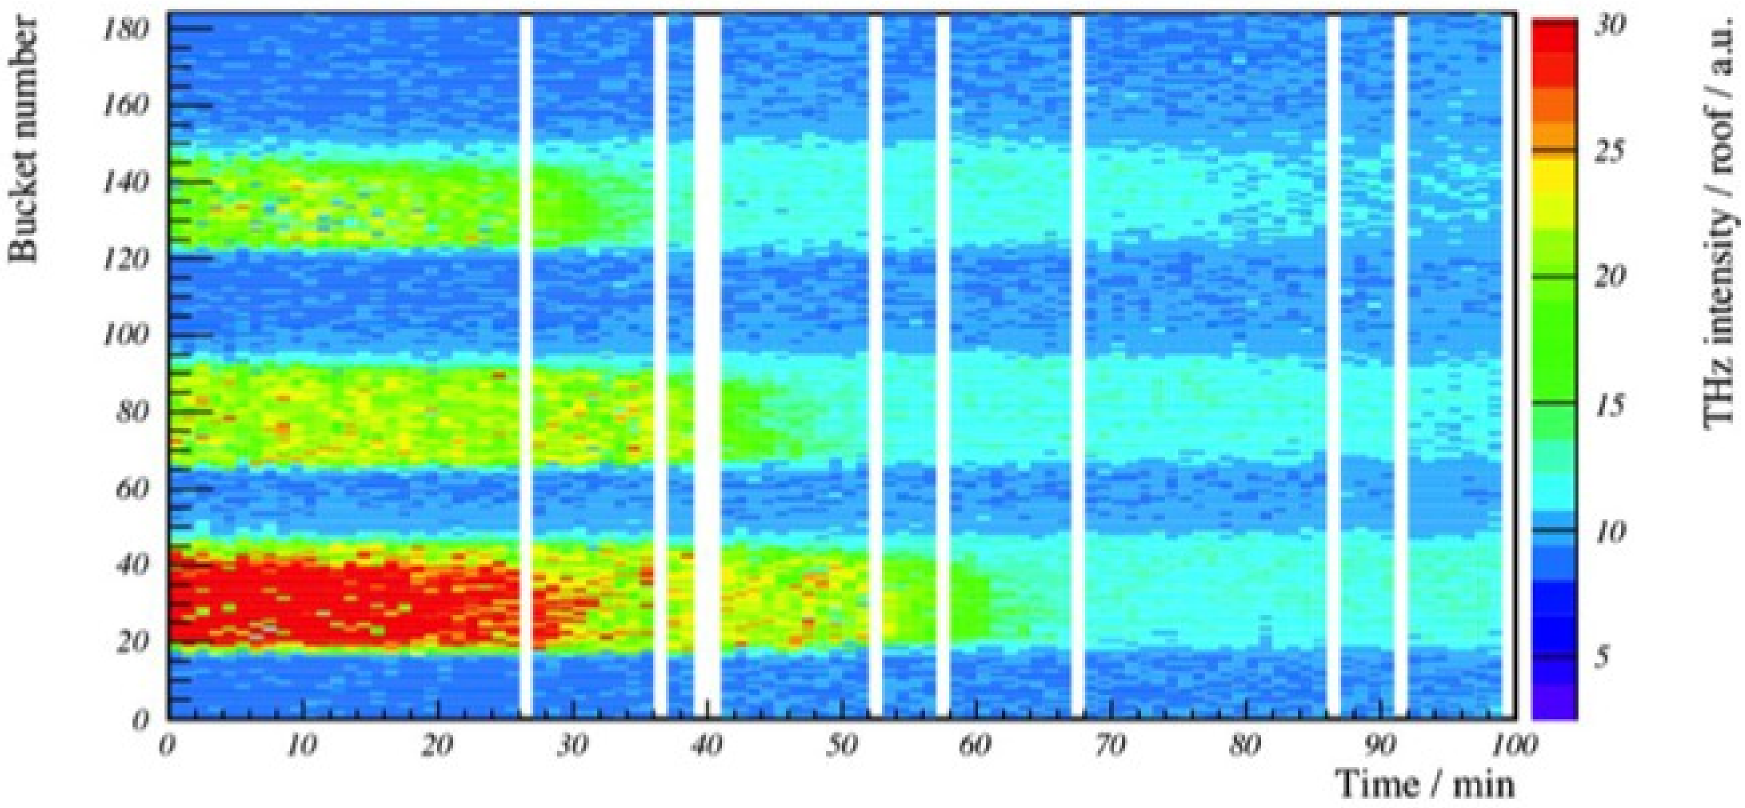

New tool for machine analysis
Contributes to portfolio ARD in program matter + technology

Flexible high-throughput FPGA platform

HELMHOLTZ ASSOCIATION

Summary:

- Data rates up to 2GB/sec
  Requirement for the next generation
  - Camera 5GB/sec
  - Hot electron bolometer 3GB/sec
- Involved IP-cores:
  - Serializer/Deserializers
  - DDR Memory
  - PCI Express w DMA
    - Preliminary version of next generation of DMA IP core ~3GB/sec
    - 64bit Linux drivers

Everybody is invited to join the
Development of the ultrafast DAQ readout system

Real-time Data Processing

1 **DAQ hardware**

PCIe Acceptable data rates

2 **PC-based processing**

Pixel Detectors → Essential Data Reduction → Data Reduction → Data Processing → Online Assessment Visualization →

Technologies

Programmable hardware	**„Parallel Computing“:**
FPGA, DSP, embedded GPUs	PC + Coprocessing / GPU

Tasks

Selected commercial hardware Standardized custom hardware	Hardware independant programming (e.g. OpenCL) Library of common algorithms

Task 2: Real-time Data Assessment

Actions:

- Computation system based on GPU co-processors
- Online tomographic reconstruction
- Prototype adoption of a complete crystallography data flow using the HDRI standard format
- General environment for parallel image processing
  - a) Independent from available hardware (e.g. OpenCL)
  - b) Library of standard algorithms
  - c) Easy adaption to new problems
- Further applications

GPUs for Ultrasound computer tomography

- Measurement Principle
  - Lying patient, breast is surrounded by many ultrasound transducers
  - Recording of interaction of ultrasound waves with tissue from many different angles
  - Reconstruction of volume image

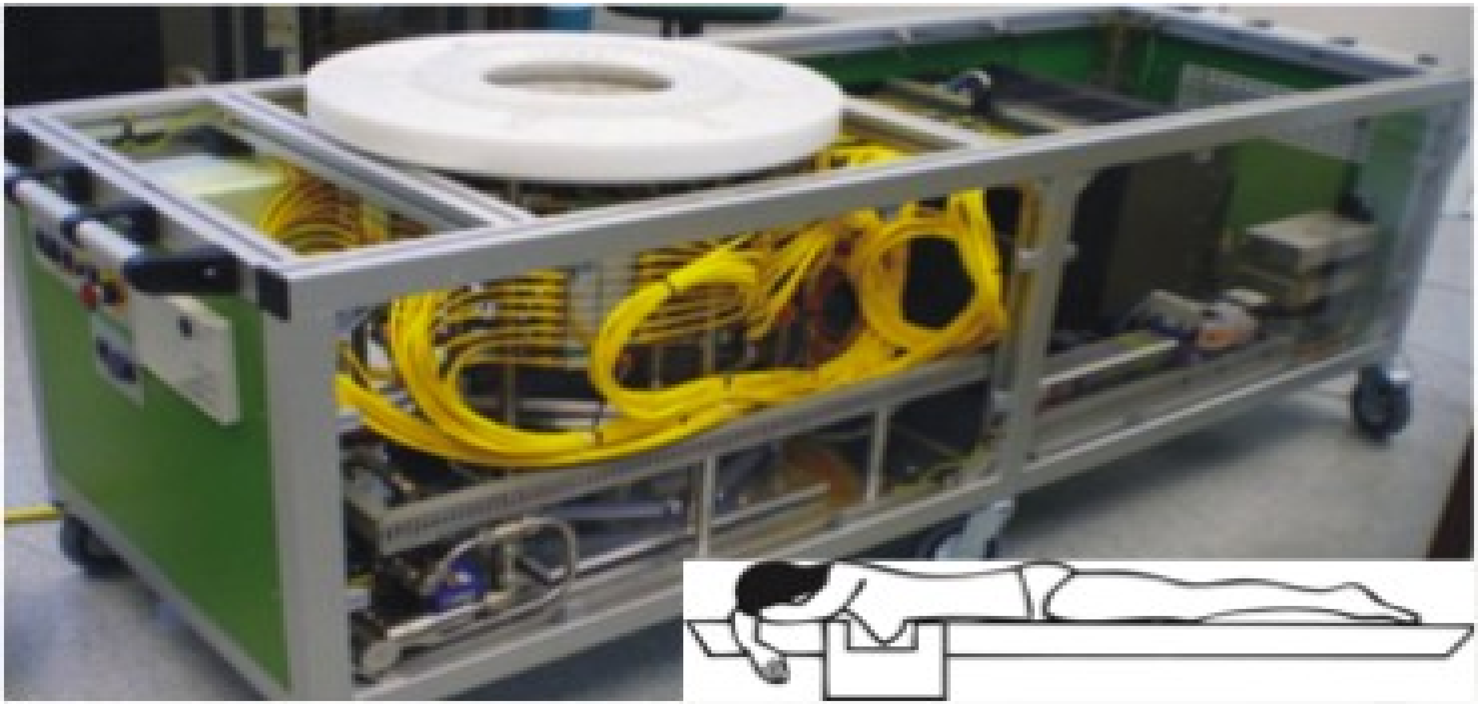

GPUs for Ultrasound computer tomography

- Measurement Principle
  - Lying patient, breast is surrounded by many ultrasound transducers
  - Recording of interaction of ultrasound waves with tissue from many different angles
  - Reconstruction of volume image

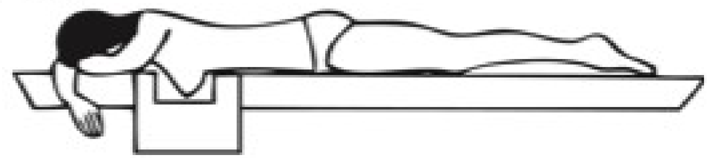

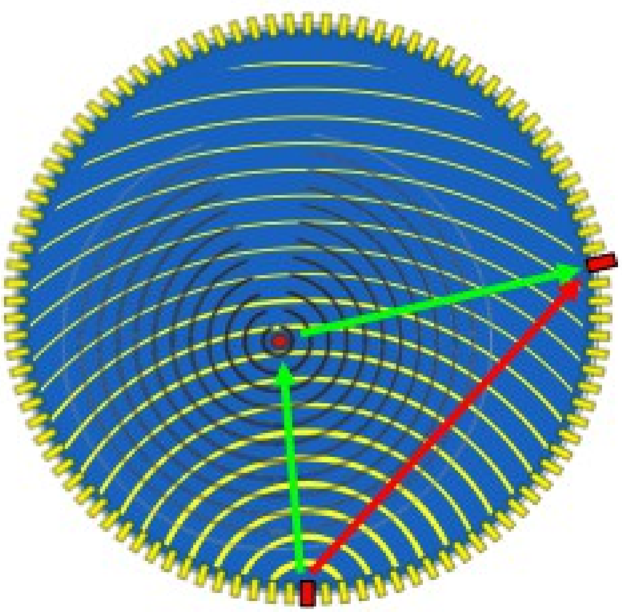

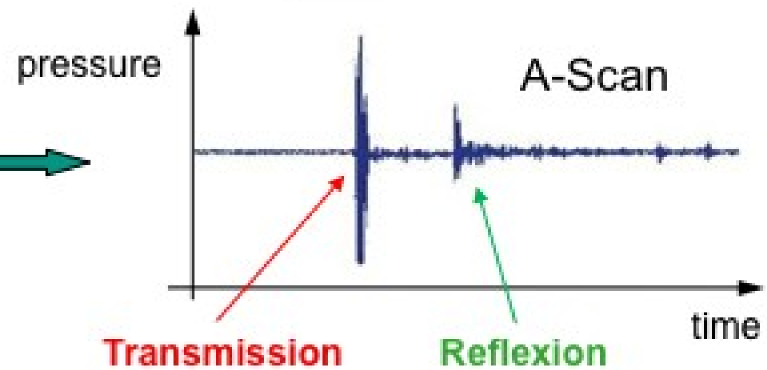

Reconstruction of
- Reflectivity
- Speed of sound
- Attenuation

"Image one, get two free"

GPUs for Ultrasound computer tomography

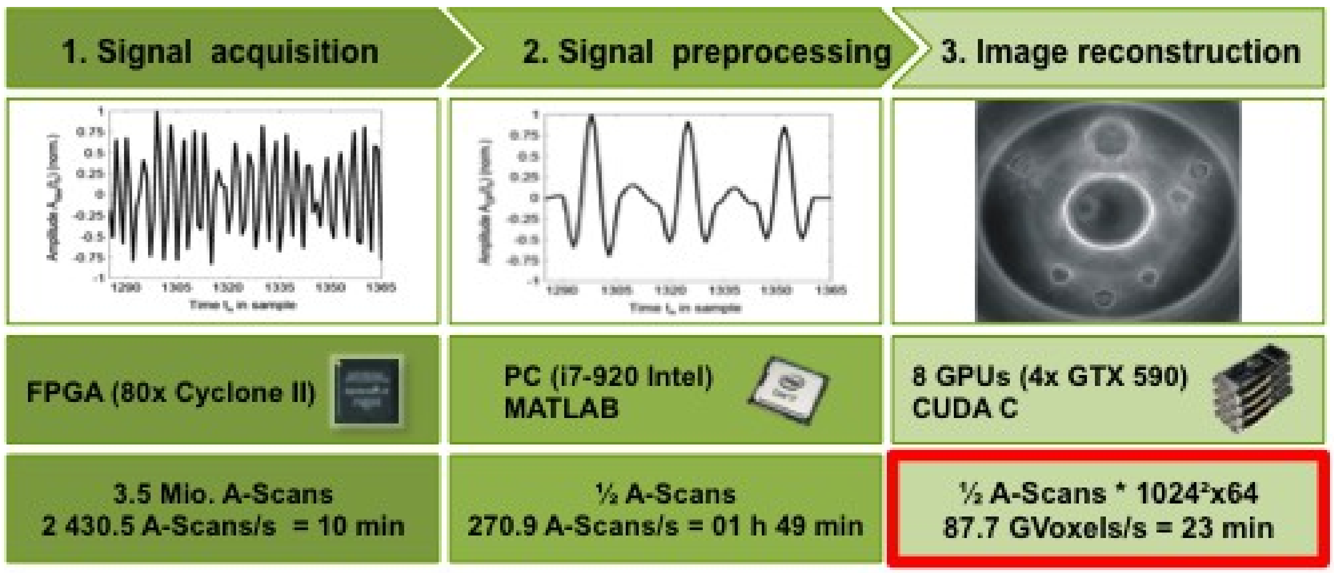

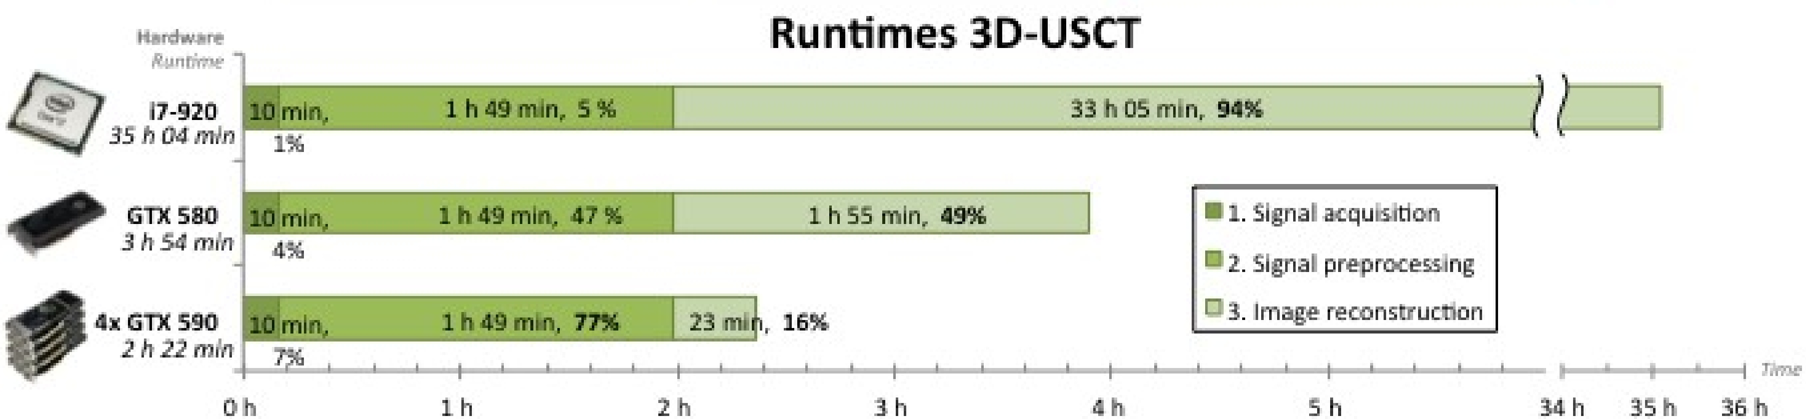

Scaling performance with multiple GPUs

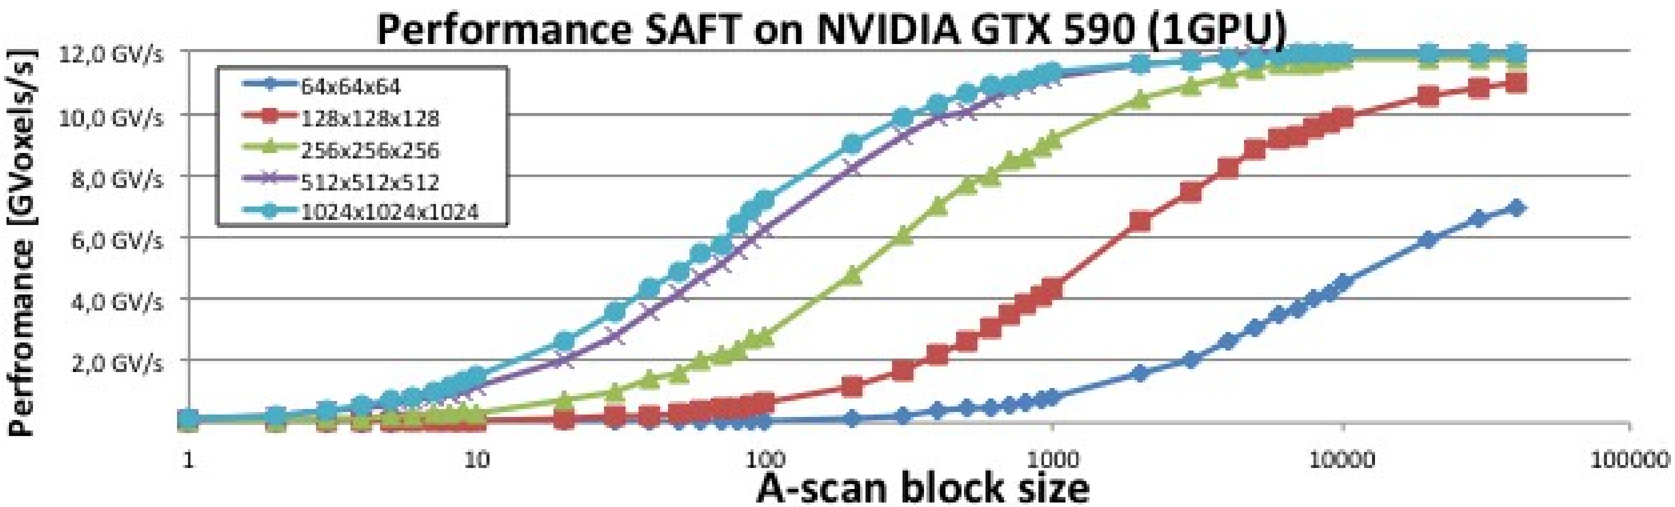

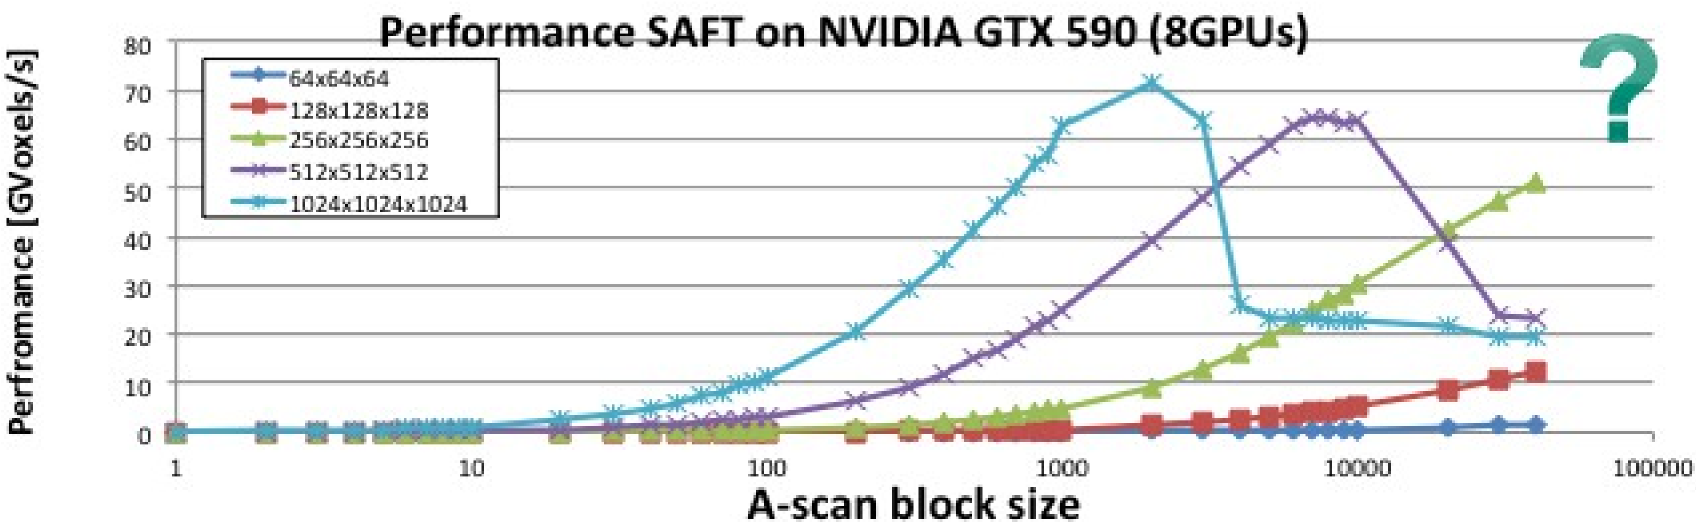

Dipl.-Ing. Ernst Kretzek

8-core GPU box: Optimization of air flow

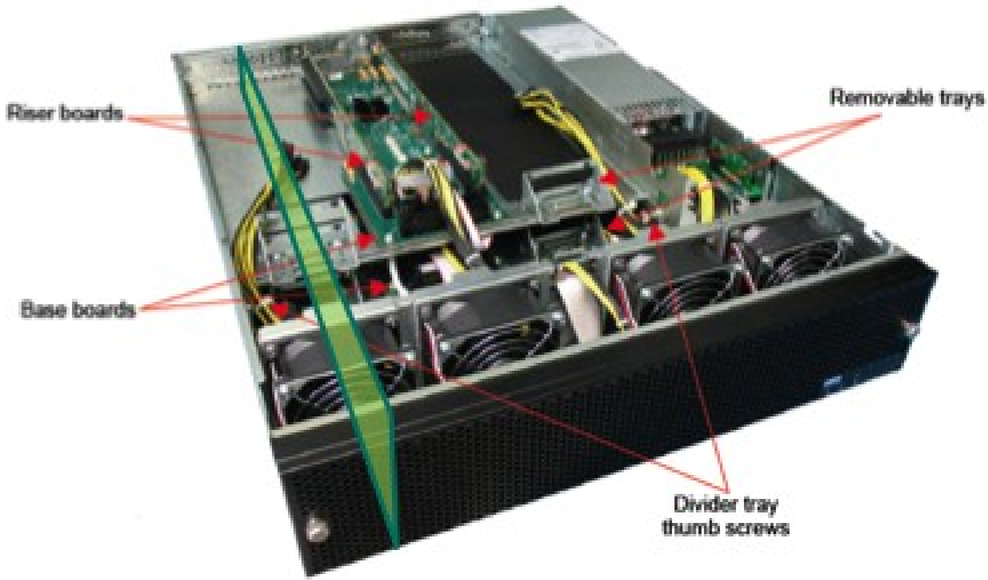

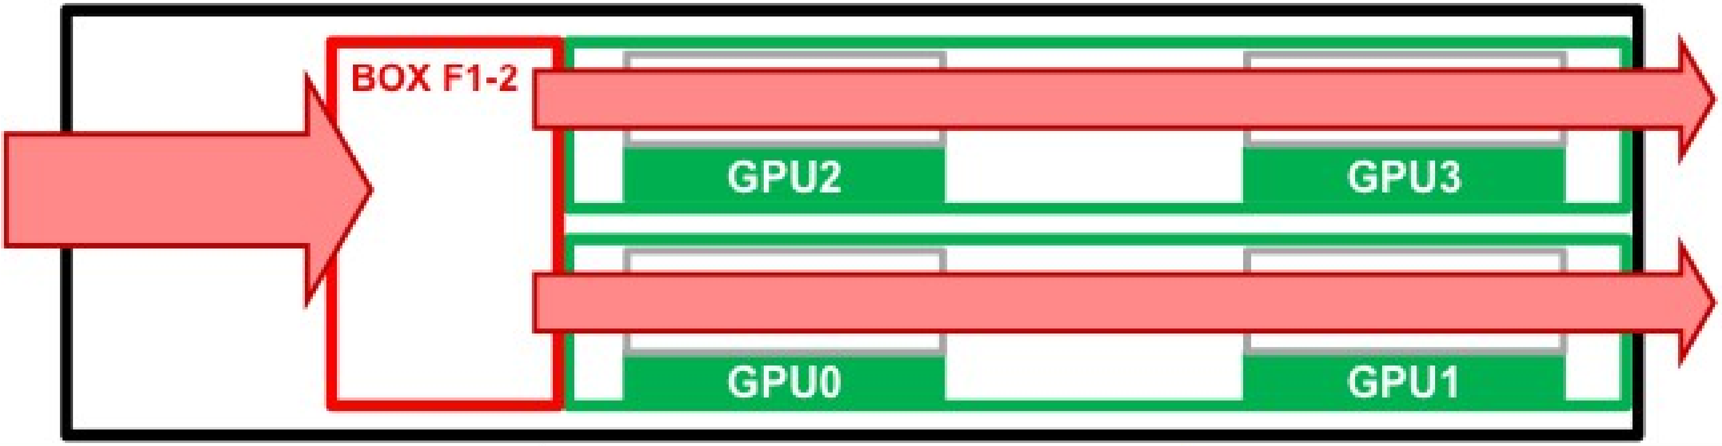

Scaling performance with multiple GPUs

HELMHOLTZ ASSOCIATION

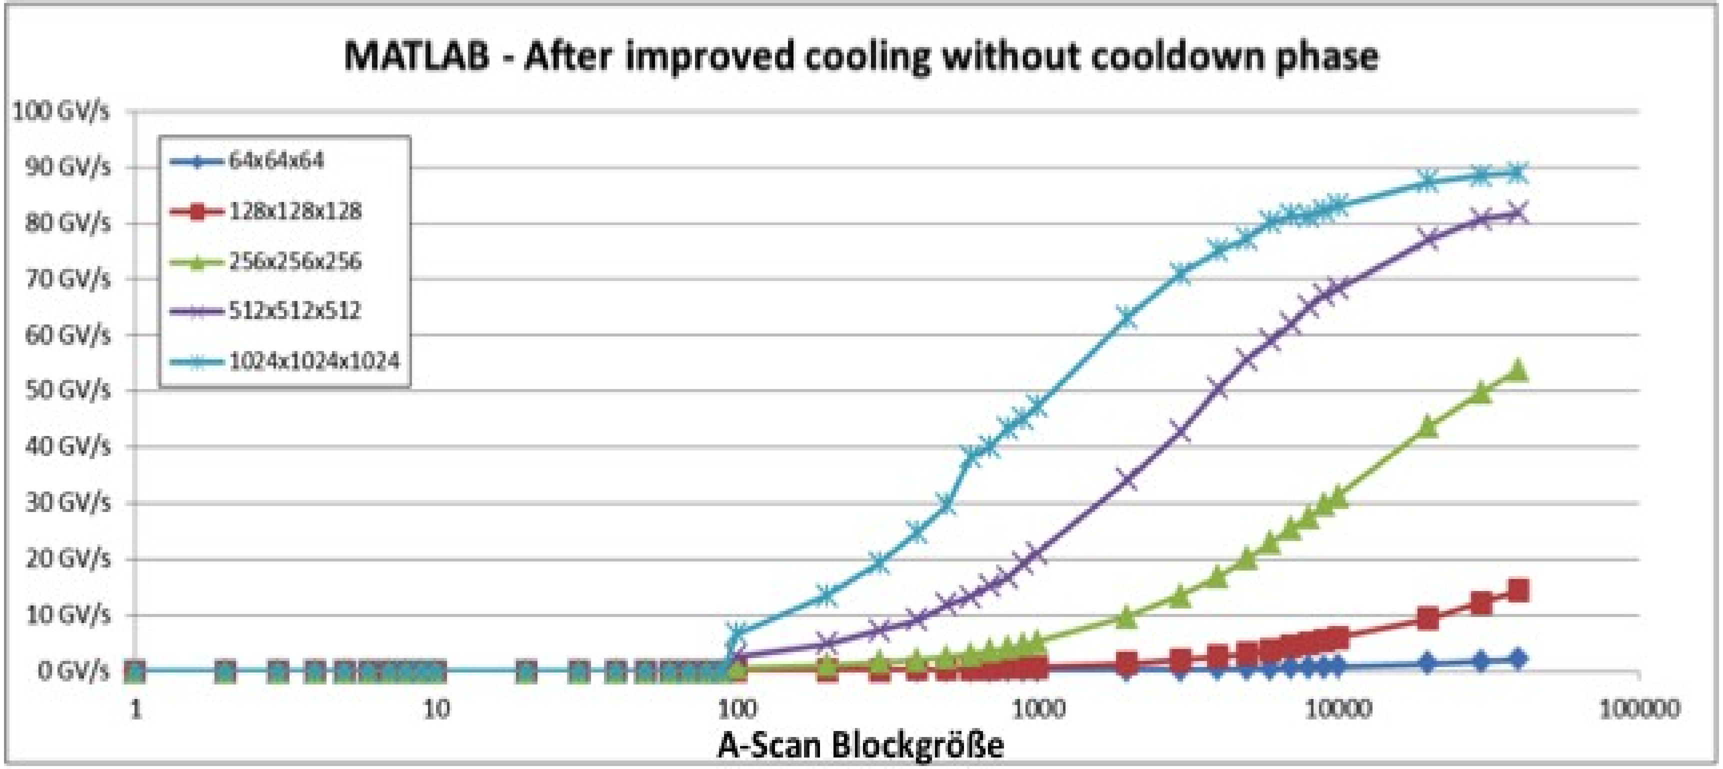

Hardware selection is crucial for high-performance GPUs
More on computing hardware → Talk: S Chilingaryan

Computing and power efficiency

- Supermicro X8DTG-QF GPU-workstation
- 2 × Intel XEON E5540 @2.53GHz (QPI-connected), 192 GB memory
- 4 × Fermi C2070 (PCI-Express 16 ×)

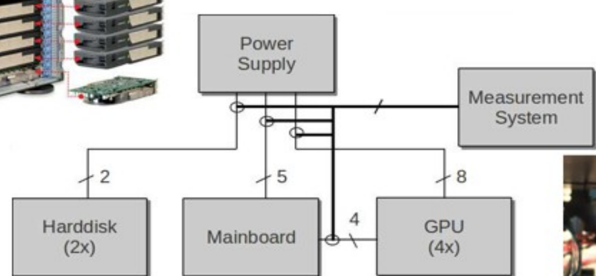

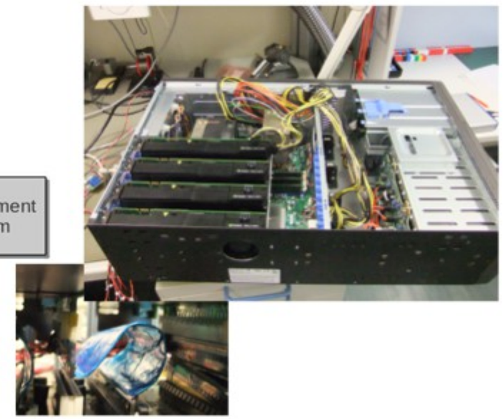

Independent, embedded measurement setup

- Powermeters monitor voltages and currents in lines powering **chipset**, the **hard- disks** and the **GPUs** (including PCI)
- 250 k Samples/second

Computing and power efficiency

Energy footprint of tomographic reconstruction, PyHST:

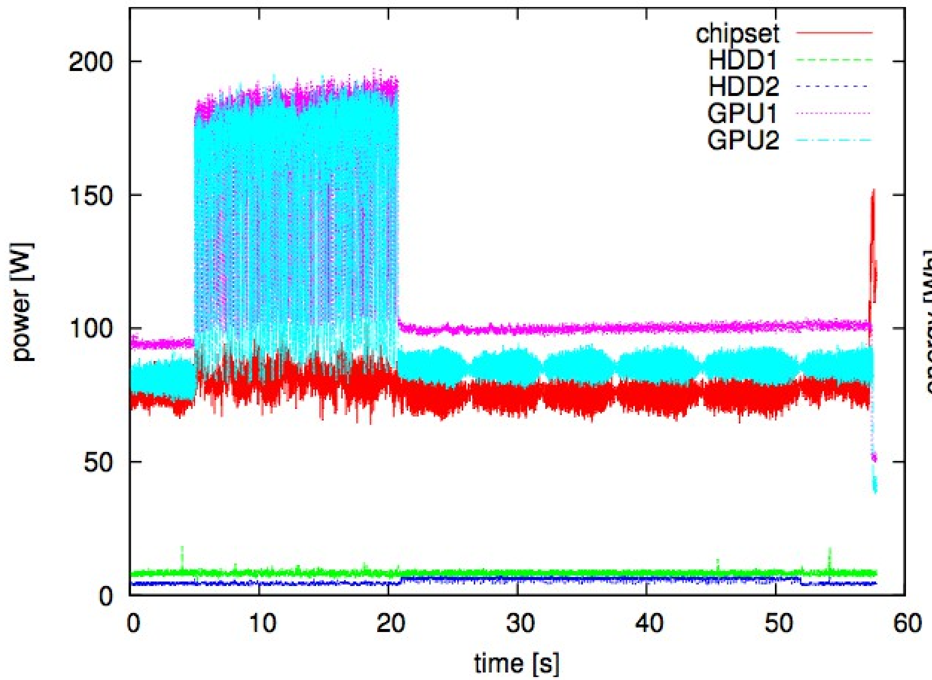

(i) Power profile

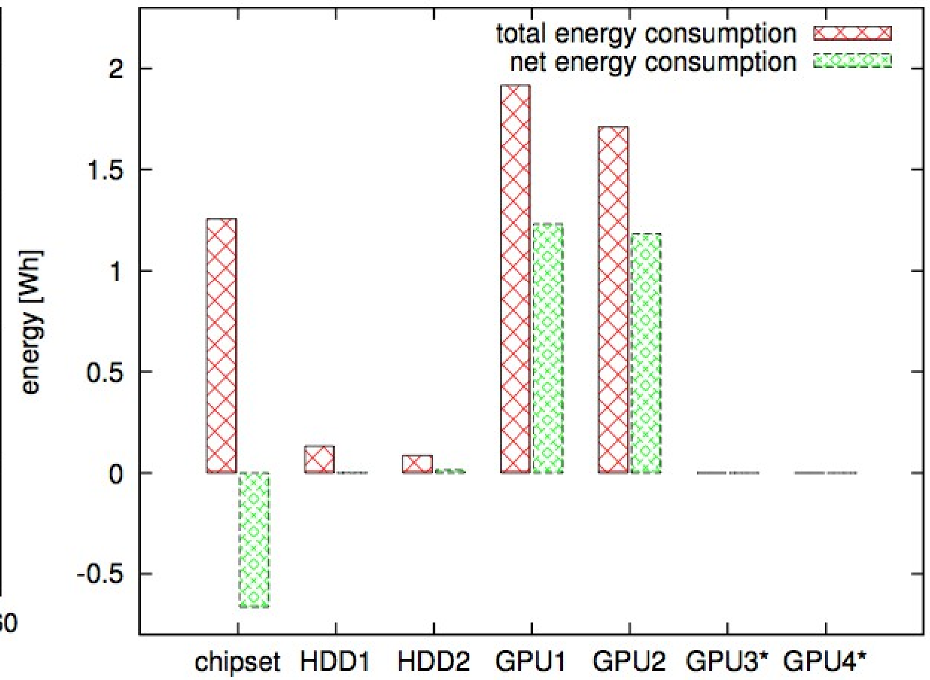

(j) Energy consumption

CPU: 22Wh, 455s
GPU: 5Wh, 58s

A. Anzt et at., ENA-HPC, 12-14.9.2012
Cooperation w Exa2Green project

Parallel Computing Framework

HELMHOLTZ ASSOCIATION

Applications
(E.g. tomographic reconstruction, ...)

Load balancing + management
(CPU ⇔ GPU ⇔ frame grabber / Single ⇔ double precision)

Primitives for image processing
(2d FFT, filters, Radon transform, image conversion, …)

Core functions, hardware access
(data transfer, file storage, camera buffering, ROI, …)

Framework: Basics

Image processing =
composition of filter nodes

- One thread per node
- Mapping to **CPU or GPU**
- Encourages **recycling** of tested components

Core system is built on top of **OpenCL** and **GLib/GObject**

- **Bindings to any language**

Documented using Sphinx and Gtk-Doc

→ **Talk: M. Vogelgesang**

Filter object = node

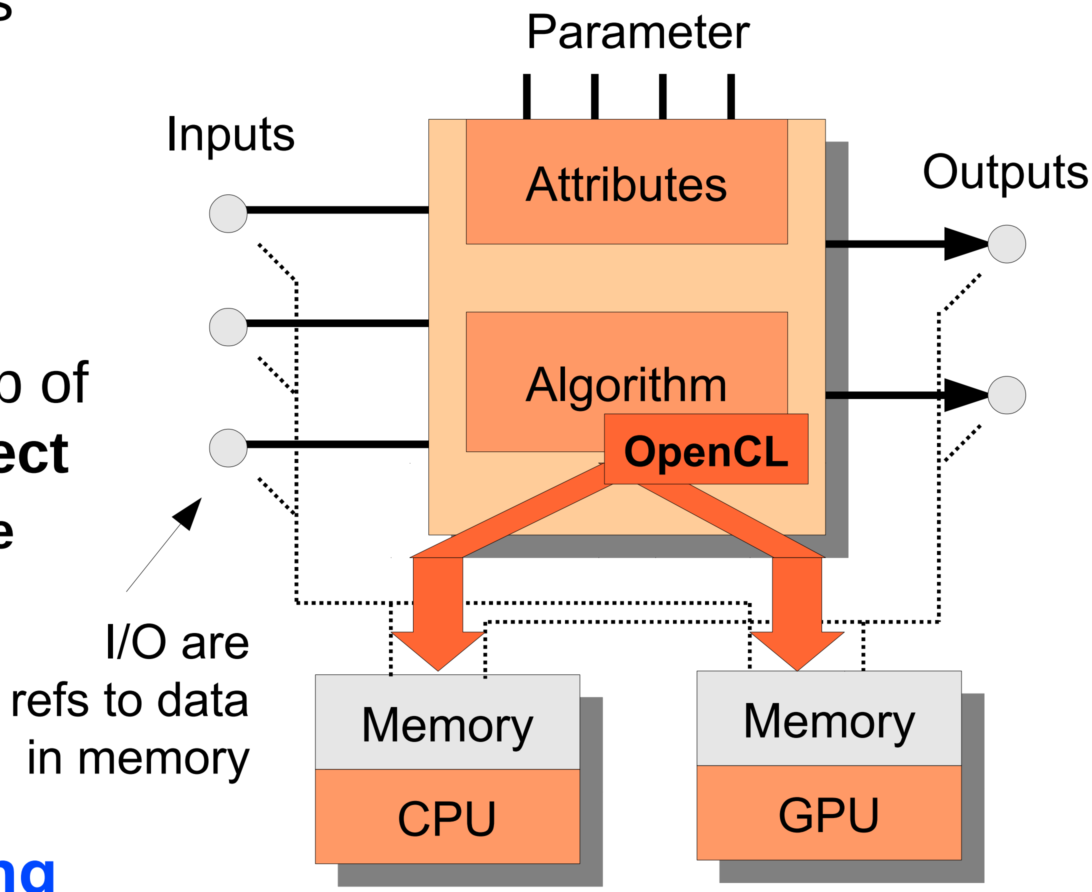

Framework: Algorithms

- Filtered back projection
- Precise calculation of the center of rotation
- Laminography
- De-noising filters
- Simulation of test data

Under development:

- Algebraic reconstruction (better results) → **Talk: Xiaoli**
  - Precise forward transform model and compressive sampling
- DFI-based reconstruction (faster processing)

Application: ROFEX – Rossendorf Fast Electron beam X-ray tomograph

HZDR

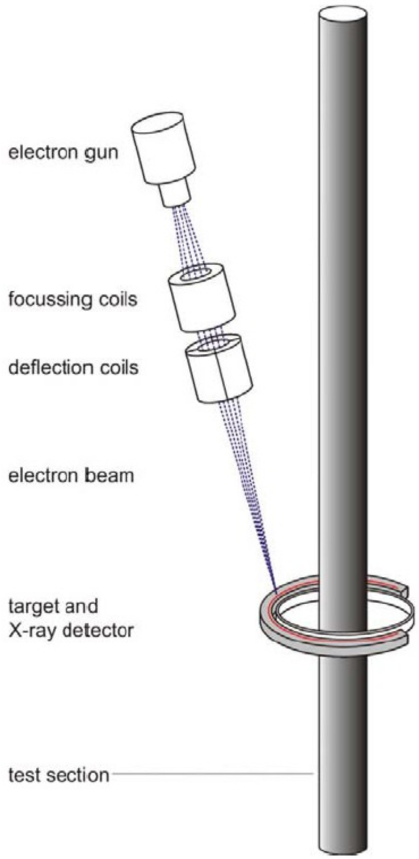

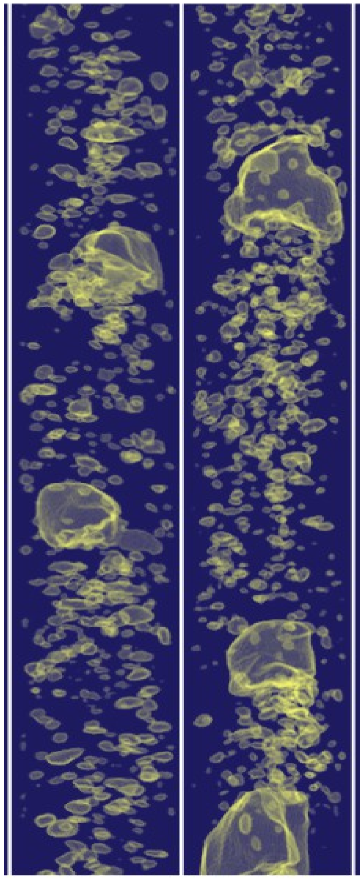

U. Hampel, et al., Flow Meas. Instrument. 16, pp. 85-90, 2005
Fischer et al., Meas. Sci. Technol. 19, pp. 094002, 2008.

DAQ frame rate: < 10.000 fps

Data rate ~ 1 GB/sec

Reconstruction with UFO-Framework

- Result: 256 x 256
- 600 frames → 0.35 sec
- Frame-rate: 1.7 kfps (1 GPU)

→ **RT-Operation possible**

Further task “Post processing“:

- Bubble detection
- Velocity measurement

Camera Abstraction

HELMHOLTZ ASSOCIATION

- Generalized access to streaming cameras (C-API)
- 64-bit linux support for PCO cameras
- Licensed under LGPL with permission from PCO
- TANGO driver

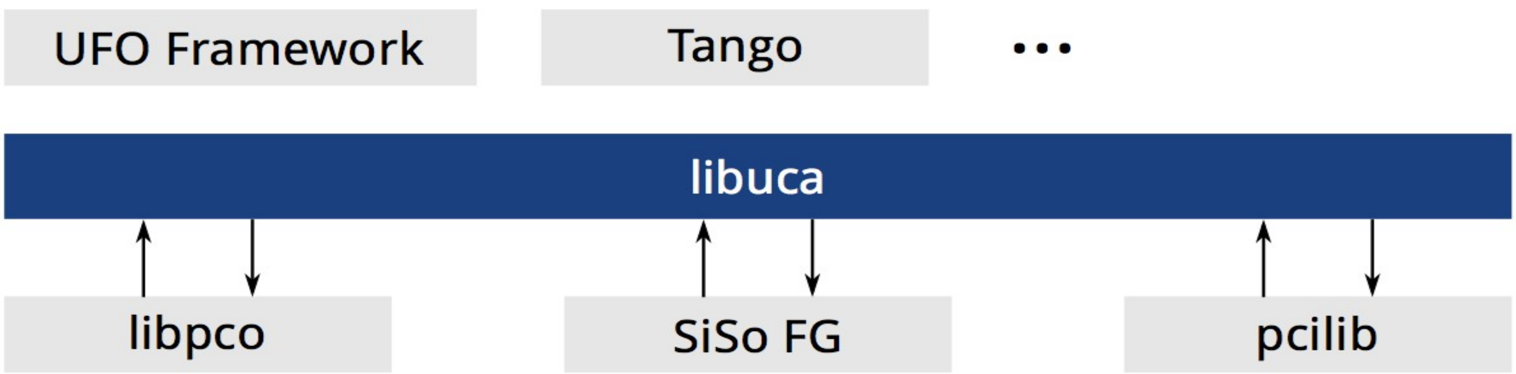

KIT high-throughput camera platform

→ Discussion on integration with LImA (ESRF, Soleil)

Summary

HELMHOLTZ ASSOCIATION

- Linux drivers for
  - High-throughput DAQ platform (KIT)
  - for PCO cameras (libuca by KIT)
    - Discussion with LImA developers (Soleil)
- Development + Optimization of GPU platforms
  - Cooling, power consumption
- Parallel processing framework for streamed data (KIT, HZG, Soleil)
  - Simplifies GPU programming
  - Possible standard interfaces for GPU algorithms ?
- Tomographic reconstruction (ESRF, KIT)
  - Advanced algorithms

Overlap with
- Supercomputing
- Matter and Technology
ARD, Detectors, LSDMA