

GPU-based data processing for ultrafast X-ray imaging — KSETA Plenary Workshop 2014

Matthias Vogelgesang

Institute for Data Processing and Electronics

Synchrotron radiation from ANKA

- X-ray source producing beam with broad spectrum and high brilliance
- X-ray scattering and diffraction, high-resolution and high-speed radiography, **tomography** and laminography
- New IMAGE beamline used to investigate fast processes in life and material sciences

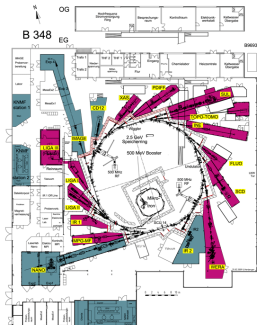




Figure: Floor plan of ANKA electron ring with 16 tangential *beamlines*

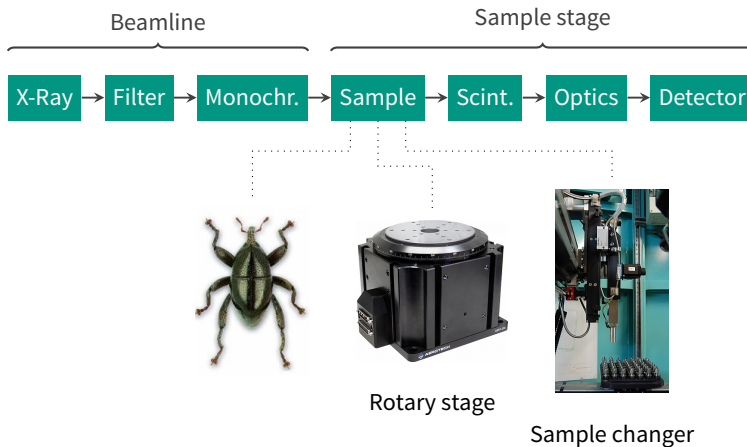
X-ray imaging experiments




X-ray imaging experiments



X-ray imaging experiments

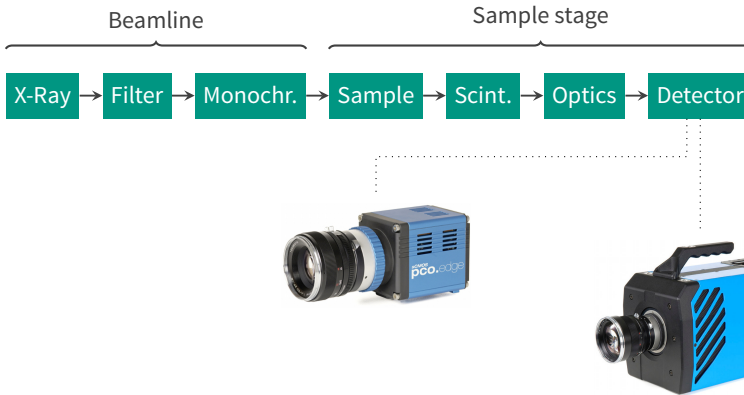


X-ray imaging experiments




Convert X-ray photons into visible light, magnify and detect them with CCD or CMOS.


X-ray imaging experiments



X-ray imaging experiments



Projections



Reconstructed slice



Visualization

Fast, high volume acquisition

- Up to hundreds of samples per experiment
- Many thousand frames per second per sample
- Up to eight mega pixels at 16 bit per frame

Fast, high volume acquisition


- Up to hundreds of samples per experiment
- Many thousand frames per second per sample
- Up to eight mega pixels at 16 bit per frame

Efficient, fast data processing

- On-site data processing for visual quality assurance
- Use reconstructed data for early feedback
- Use existing computing infrastructure


Problem

From a series of projections ...




Problem

From a series of projections ... reconstruct *unknown* slice information



Problem

From a series of projections ... reconstruct *unknown* slice information



Solutions

- Solve analytically using the Fourier-slice theorem (DFI)
- Filter projections and smear back into empty volume (FBP)
- Model detection as a linear system and solve algebraically (ART)

There is no “one-size-fits-all” solution

- Quality: DFI < FBP < ART
- Speed: DFI > FBP > ART¹

CPUs are still too slow

- Up to hours for very large data sets
- On-line and on-site inspection assessment inconvenient
- Fast feedback impossible

¹ $O(n \log n) \subset O(n^3) \subseteq O(cn^3)$

Why GPUs?

A superficial comparison

Processor	Purpose	Cores	Performance	Bandwidth	Price
CPU	General	12	0.5 TFLOPs	60 GB/s	> € 2200
GPU	Graphics	2880	5.0 TFLOPs	336 GB/s	€ 430


But ...

- Work differently than a CPU
- Cover restricted problem domains

More challenges

Heterogeneous computing


- Systems with different processors and topologies (CPUs, GPUs, FPGAs)



More challenges

Heterogeneous computing


- Systems with different processors and topologies (CPUs, GPUs, FPGAs)
- Needs precise scheduling and resource allocation as well as architecture-specific code
- Avoid unnecessary copies between devices



More challenges

Heterogeneous computing


- Systems with different processors and topologies (CPUs, GPUs, FPGAs)
- Needs precise scheduling and resource allocation as well as architecture-specific code
- Avoid unnecessary copies between devices



More challenges

Heterogeneous computing

- Systems with different processors and topologies (CPUs, GPUs, FPGAs)
- Needs precise scheduling and resource allocation as well as architecture-specific code
- Avoid unnecessary copies between devices




Data streaming

- Common in trigger systems
- Low latency and high throughput desirable
- Must consider clusters



Basic approach

- Define algorithmic computation *and* data flow as a directed graph





Basic approach

- Define algorithmic computation *and* data flow as a directed graph
- Determine local system of CPUs and GPUs




Basic approach

- Define algorithmic computation *and* data flow as a directed graph
- Determine local system of CPUs and GPUs
- Transform graph to accommodate for additional processors




Basic approach

- Define algorithmic computation *and* data flow as a directed graph
- Determine local system of CPUs and GPUs
- Transform graph to accommodate for additional processors
- Assign tasks to GPUs




Basic approach

- Define algorithmic computation *and* data flow as a directed graph
- Determine local system of CPUs and GPUs
- Transform graph to accommodate for additional processors
- Assign tasks to GPUs and CPUs




Basic approach

- Define algorithmic computation *and* data flow as a directed graph
- Determine local system of CPUs and GPUs
- Transform graph to accommodate for additional processors
- Assign tasks to GPUs and CPUs




- Use algorithmic description



Local master


- Use algorithmic description
- Replicate sub-graph



Local master

Scaling to clusters

- Use algorithmic description
- Replicate sub-graph
- Instantiate tasks on remote nodes




Local master



Remote slave

Scaling to clusters

- Use algorithmic description
- Replicate sub-graph
- Instantiate tasks on remote nodes
- Forward data and receive results

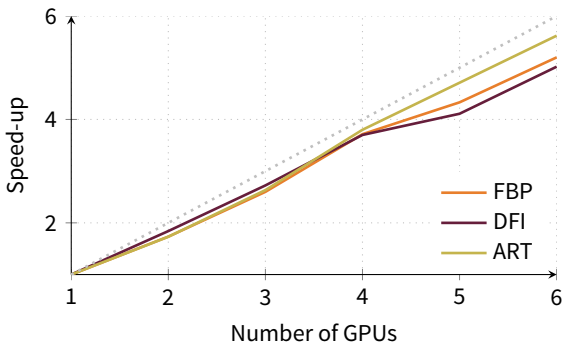


Local master

Remote slave

Improvements

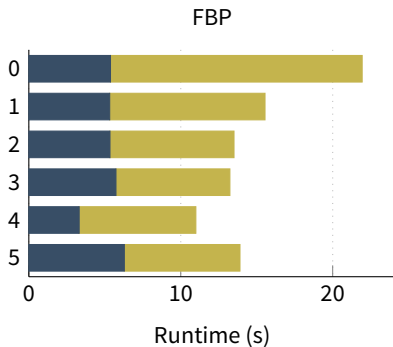
Single node multi-GPU



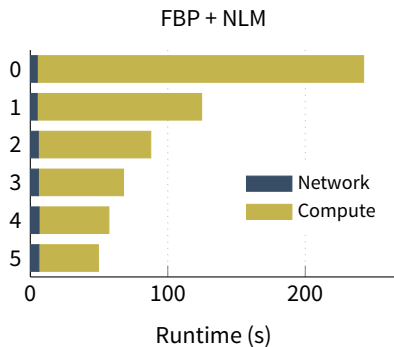
Good scalability with near linear speed-up for up to 6 NVIDIA GTX 580's.

Multi node cluster

Scalability strongly depends on “computation / data transfer” ratio ...



High network overhead



Almost linear speedup

Let's do better

Performance problems but ...

- Yet another compute language
- Hardware knowledge required for best performance
 - Work group layout,
 - Memory access patterns
 - Cache hierarchies etc. pp. ...

Performance problems but ...

- Yet another compute language
- Hardware knowledge required for best performance
 - Work group layout,
 - Memory access patterns
 - Cache hierarchies etc. pp. ...

GPU programming is **hard!**

Performance problems but ...

- Yet another compute language
- Hardware knowledge required for best performance
 - Work group layout,
 - Memory access patterns
 - Cache hierarchies etc. pp. ...

GPU programming is **hard!**

Use Python instead

- Straightforward: Use NumPy's vectorized expressions
- Clever: Run the same code on GPU/CPU/FPGA


Compiling Python into GPU code

- Python interprets statements on an abstract virtual CPU




Compiling Python into GPU code

- Python interprets statements on an abstract virtual CPU
- Parse Python code into syntax tree




Compiling Python into GPU code

- Python interprets statements on an abstract virtual CPU
- Parse Python code into syntax tree
- Optimize tree according to environment



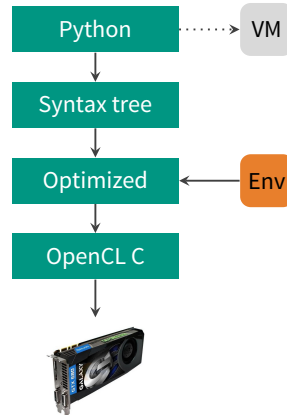
Compiling Python into GPU code

- Python interprets statements on an abstract virtual CPU
- Parse Python code into syntax tree
- Optimize tree according to environment
- Generate OpenCL C code



Compiling Python into GPU code

- Python interprets statements on an abstract virtual CPU
- Parse Python code into syntax tree
- Optimize tree according to environment
- Generate OpenCL C code
- Run generated code on GPU(s)



- Non-invasive changes

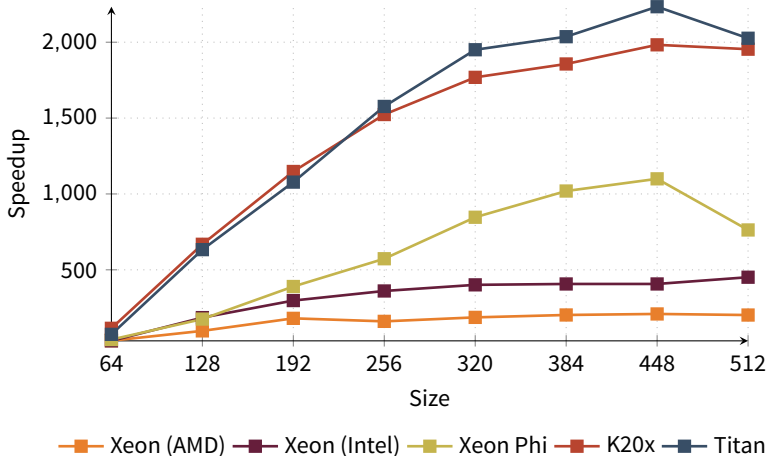
```
def calc(x, y):  
    return np.cos(x) + np.sin(y)
```

becomes

```
@jit  
def calc(x, y):  
    return np.cos(x) + np.sin(y)
```

- Additional optimization opportunities

Speedup compared to NumPy



- Heterogeneous systems can be utilized for streamed data

Conclusion

- Heterogeneous systems can be utilized for streamed data
- Scaling out to clusters works well

Conclusion

- Heterogeneous systems can be utilized for streamed data
- Scaling out to clusters works well
- No one needs to write GPU kernels

Thank you. Any questions?