Towards end-to-end data-management for large scale x-ray facilities

Professor Brian Vinter

Niels Bohr Institute
University of Copenhagen

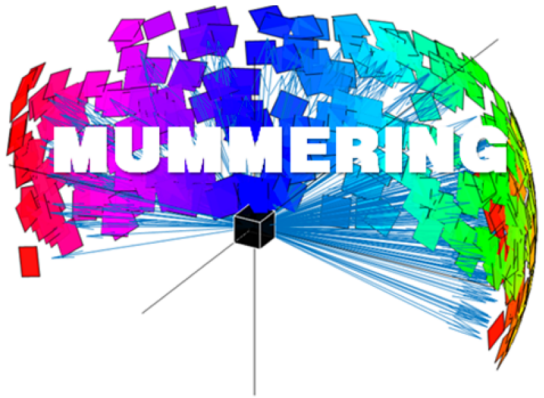

MSCA-ITN-2017 under the European Union's H2020 program

Grant No.765604

Towards end-to-end data-management for large scale ~~x-ray~~ facilities

Professor Brian Vinter

Niels Bohr Institute
University of Copenhagen

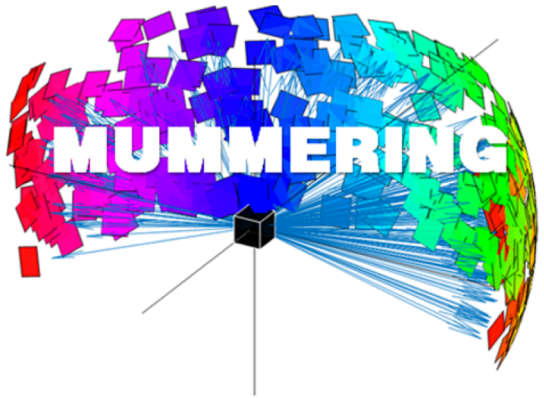

MSCA-ITN-2017 under the European Union's H2020 program
Grant No.765604

Tons of good news in x-ray technology

- Sources are getting brighter
- Sources are getting still more stable
- Spatial detector resolution grows exponentially
- Temporal detector resolution grows exponentially
- The x-ray user community is growing

The downside

- The size of each detection grows exponentially
- Frequency of detection grows exponentially
- Frequency of experiment grows rapidly

The challenges

- We need to store much more data
- Individual datasets are too large for a PC to store
- Individual datasets are too large for a PC to process
- Hand-me-down Matlab scripts are not usable for such large datasets
- Many of the new user communities are not computing natives

Moonshot proposal

Can we build a software framework that supports huge datasets, has a user friendly interface, offers an easy-to-use compute service, and facilitates cross-organizational collaboration?

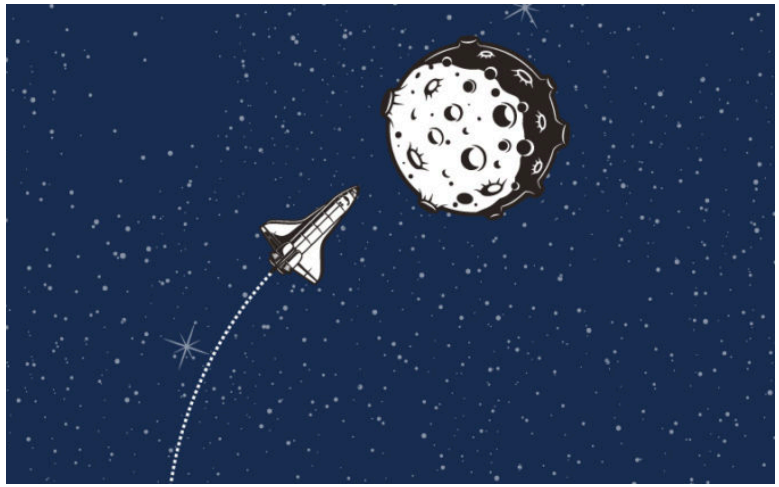

Initial Requirements

- Large data-storage
  - Fast storage
  - Data Management features
    - Online Inspection
    - Archiving
- Build-in processing support
  - Interactive
  - Batch Processing
- Cross Organization Support

Current Prototype

- Imaging Data Management System
  - Not really imaging specific so poor choice of name!
- Alternative entry to UCPH ERDA system
  - 10 PB storage
  - File system
  - Project sharing
  - Folder Synchronization
  - Jupyter Interactive Processing
  - MiG grid-backend

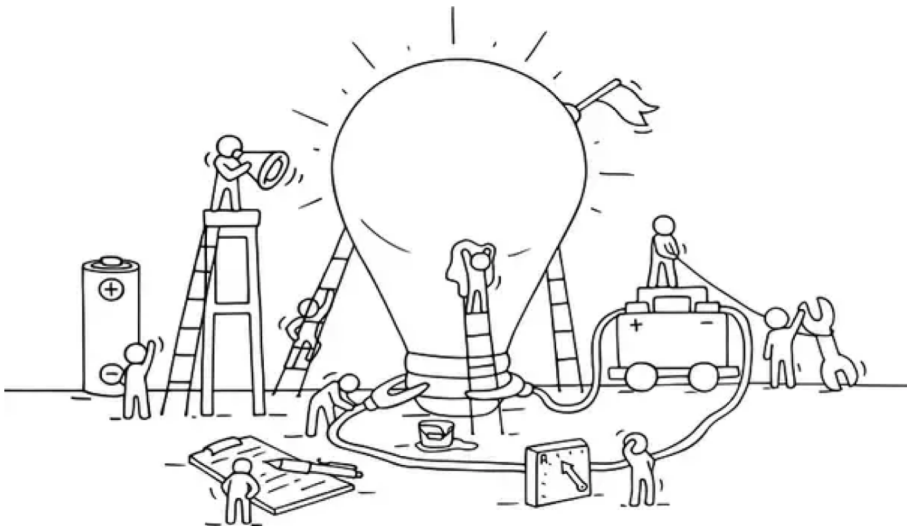

Large Data Storage

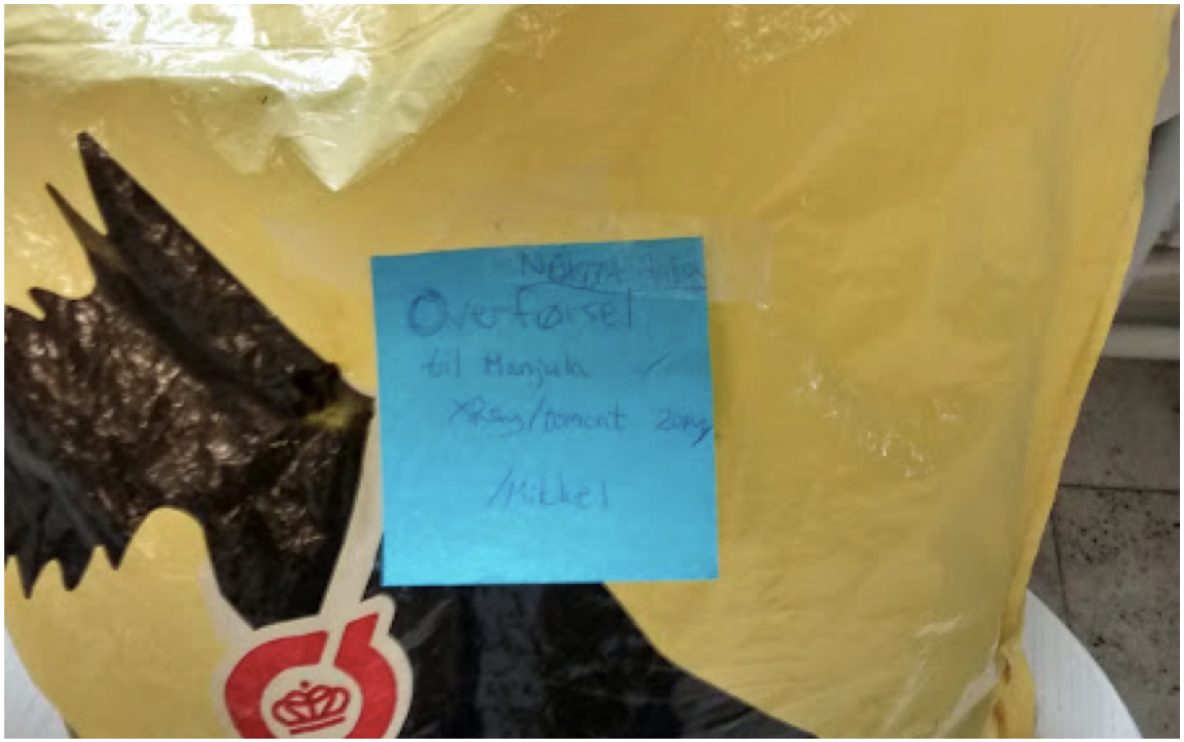

Large Storage

- Dirt cheap
- Much larger disk systems per node
- Disk redundancy
- Server redundancy
- 2 x 100 Gb input

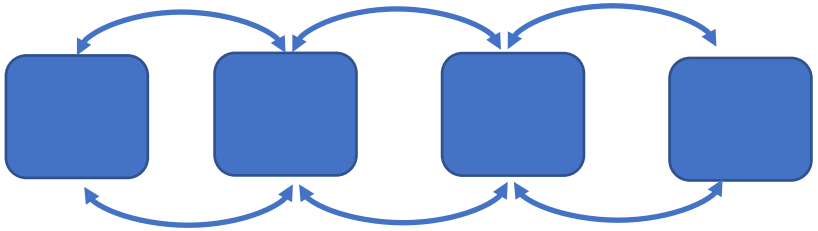

Data Management

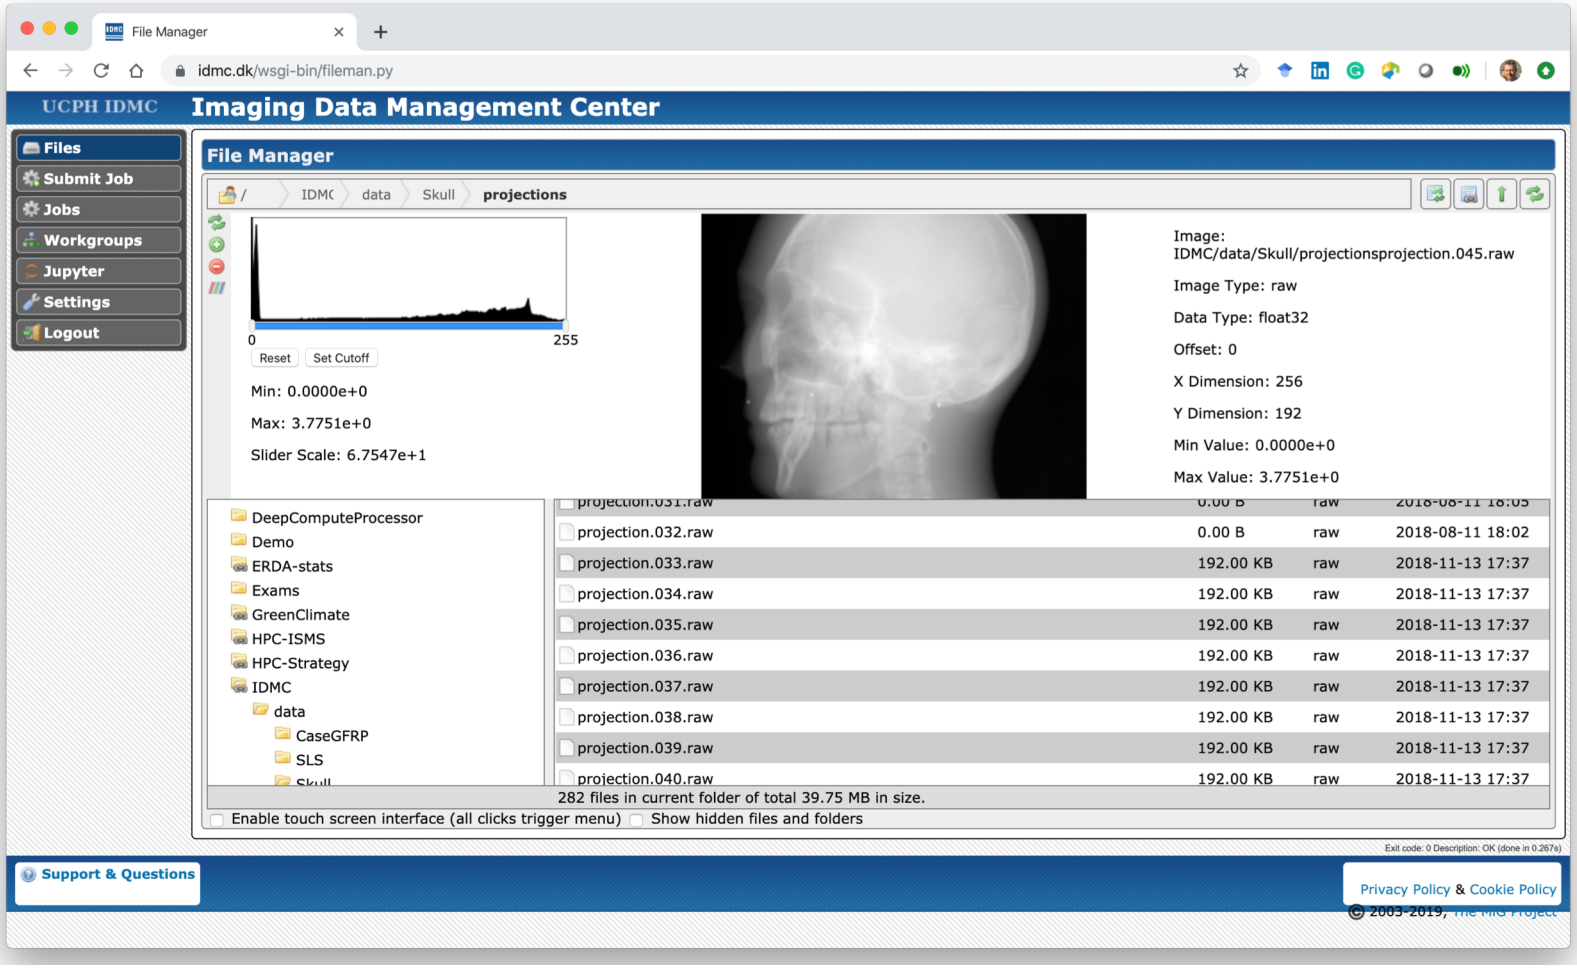

Interactive Processing

- Jupyter based interface
  - Python
  - R
  - C++
  - C#
  - (others are possible)
- Three kinds of resources
  - DAG (64 cores 256 GB memory)
  - HEL (DGX-1)
  - MODI (Cluster: 512 cores, 1TB memory)

Interactive Processing

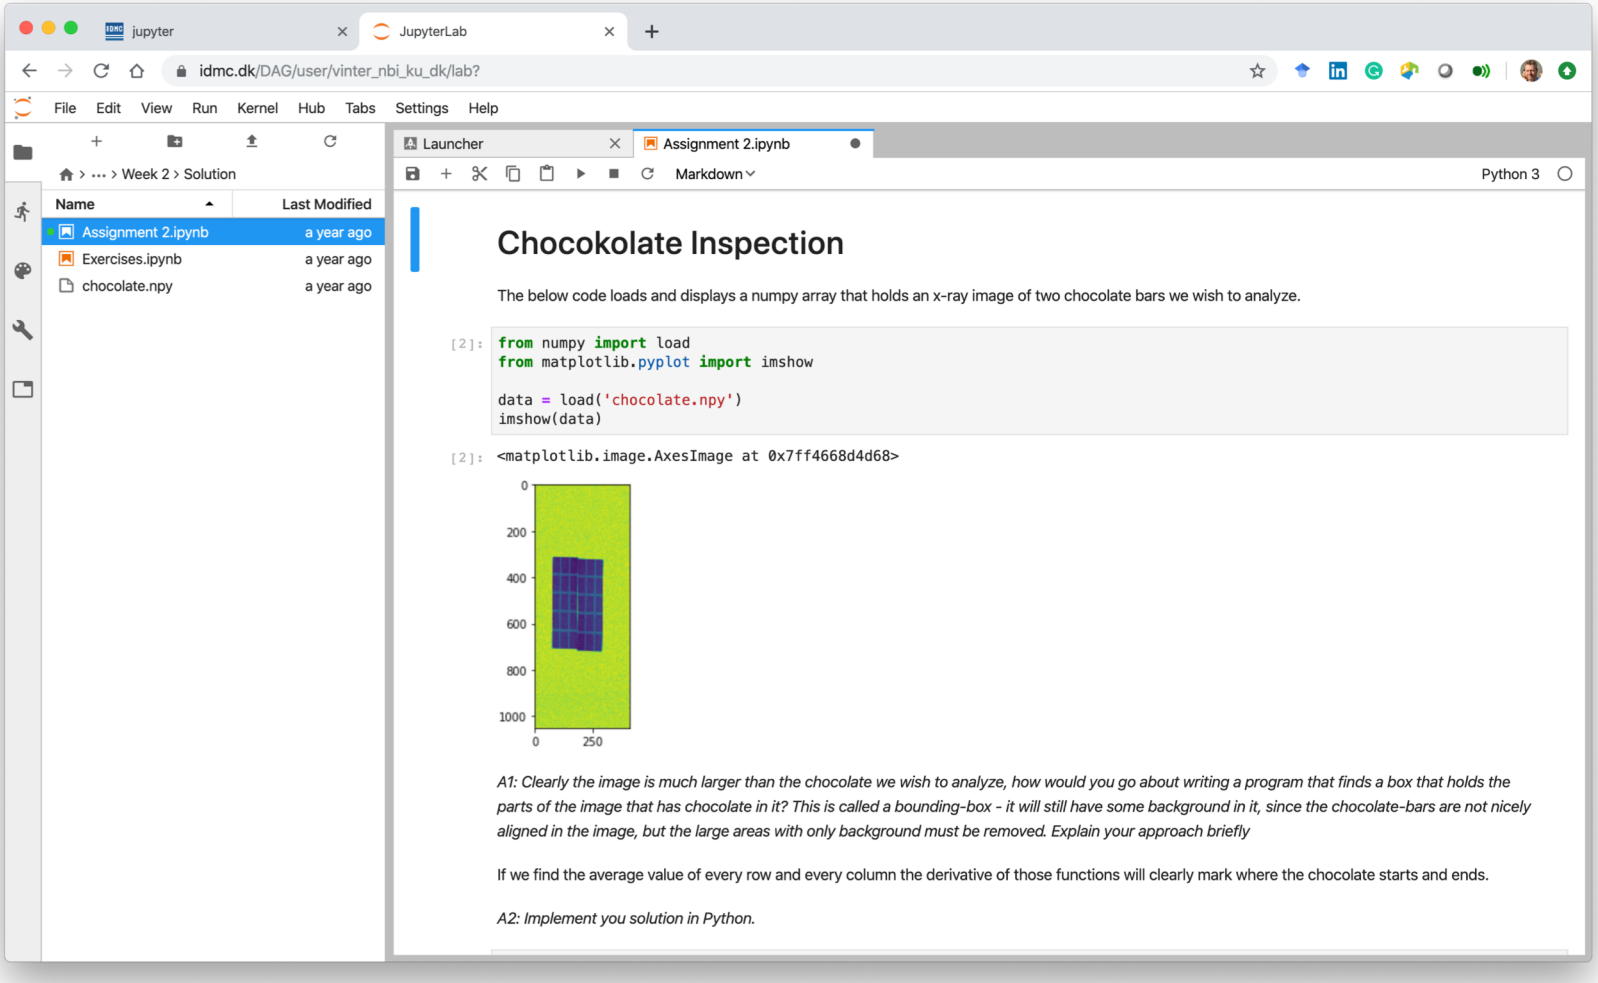

Batch Processing

- Run completely at user-level
- No custom grid software
  - On Unix based systems
- Resource Owners decide which projects can use their computers
  - And when
- Automatic error recovery

Batch Processing

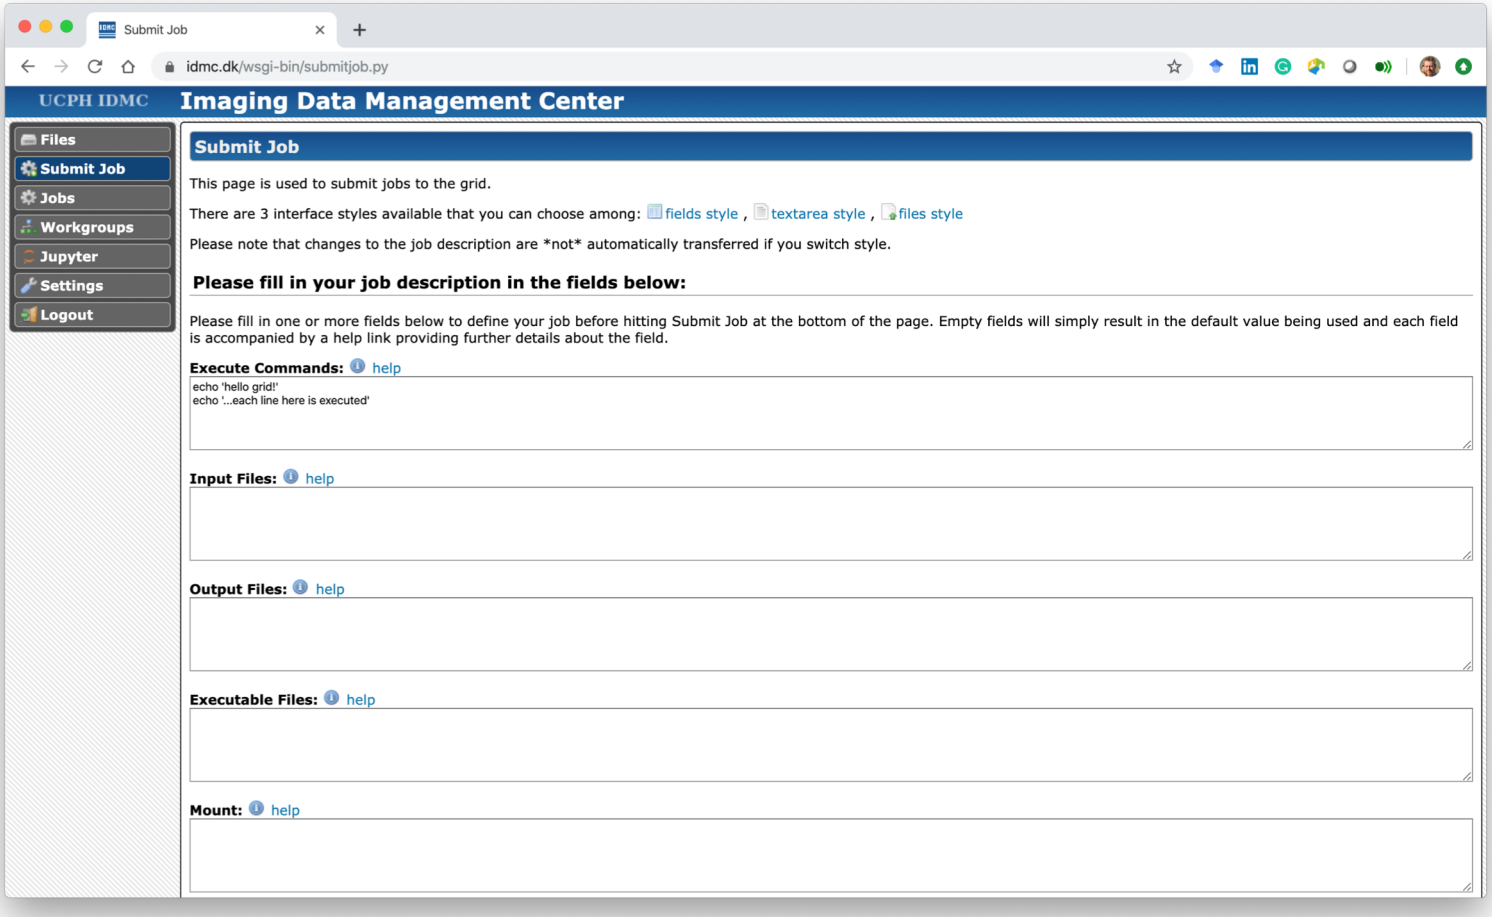

Cross Organization Support

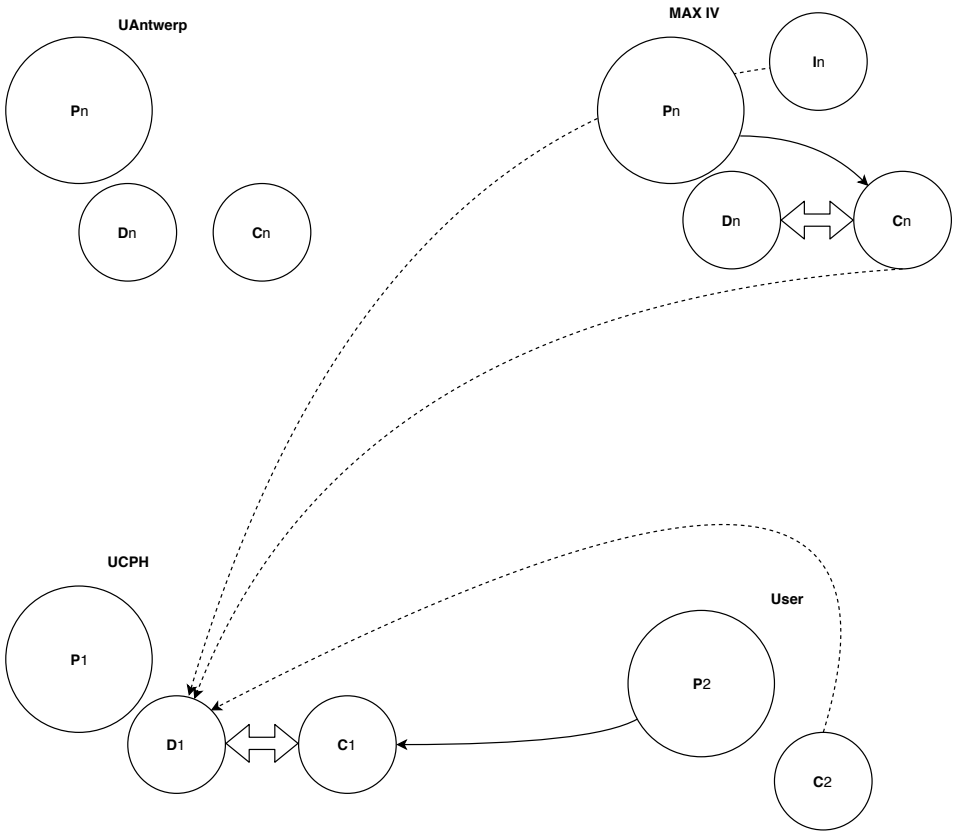

Cross Organization Support

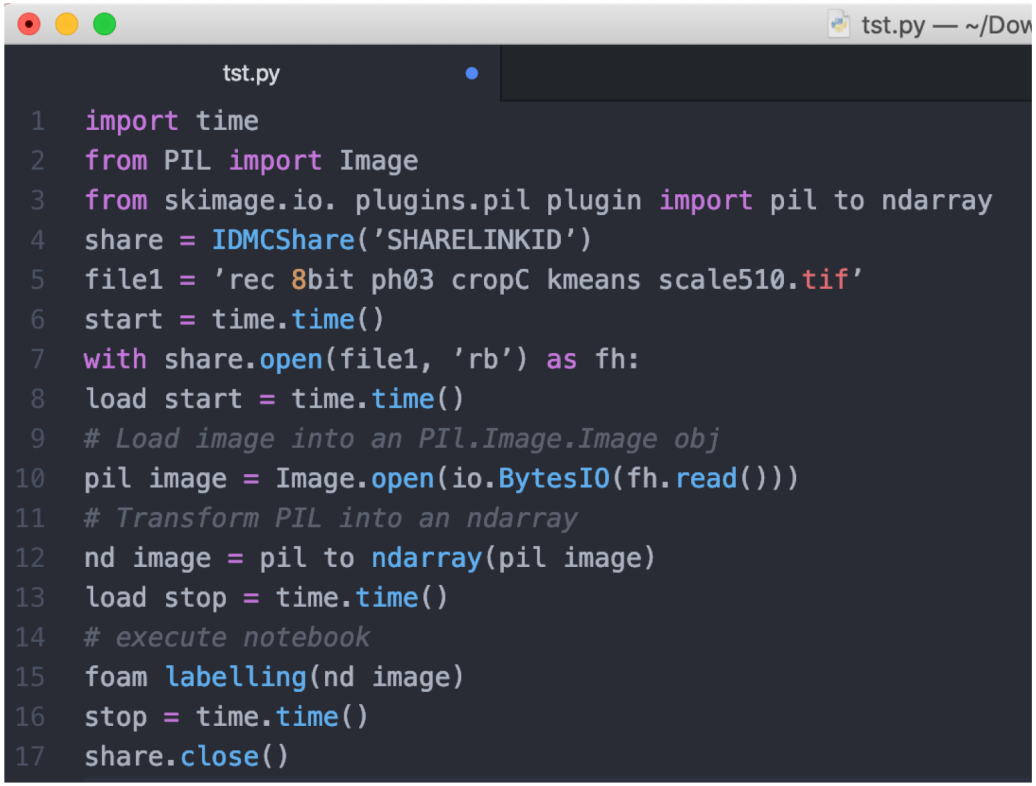

Initial Requirements

- ✓Large data-storage
- ✓Fast storage
- ✓Data Management features
  - ✓Online Inspection
  - ✓Archiving
- ✓Build-in processing support
  - ✓Interactive
  - ✓Batch Processing
- ✓Cross Organization Support

Future Developments

- High Speed Real Time Data Analysis
- Usage of low power storage
- Integrating batch-setups in Jupyter
- Securing data integrity with signing

High Speed Real Time Data Analysis

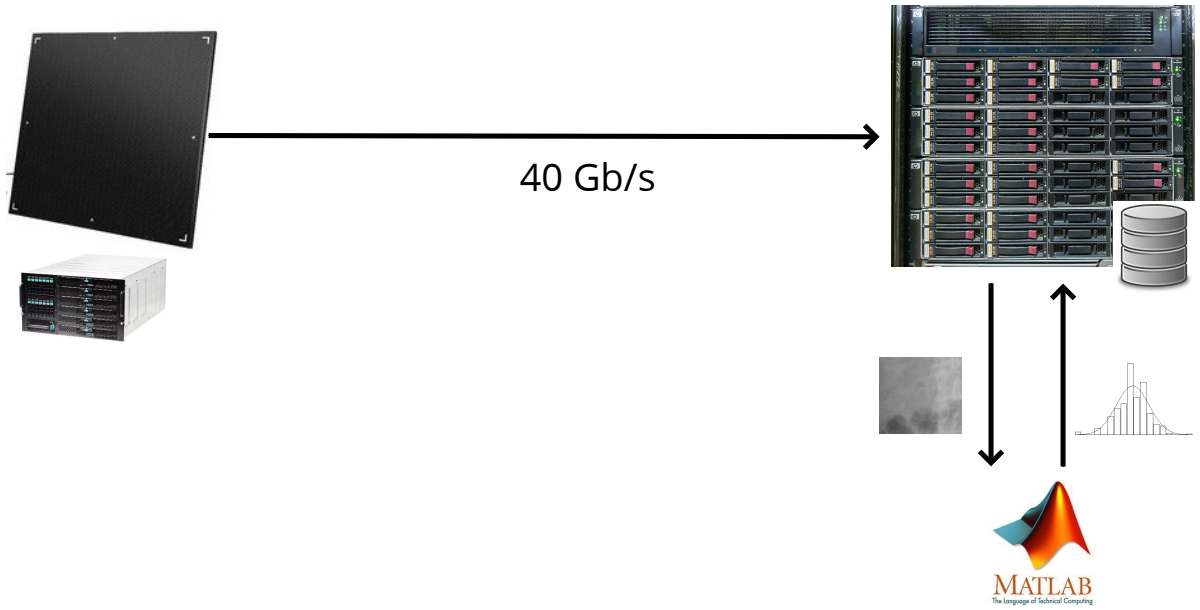

High Speed Real Time Data Analysis

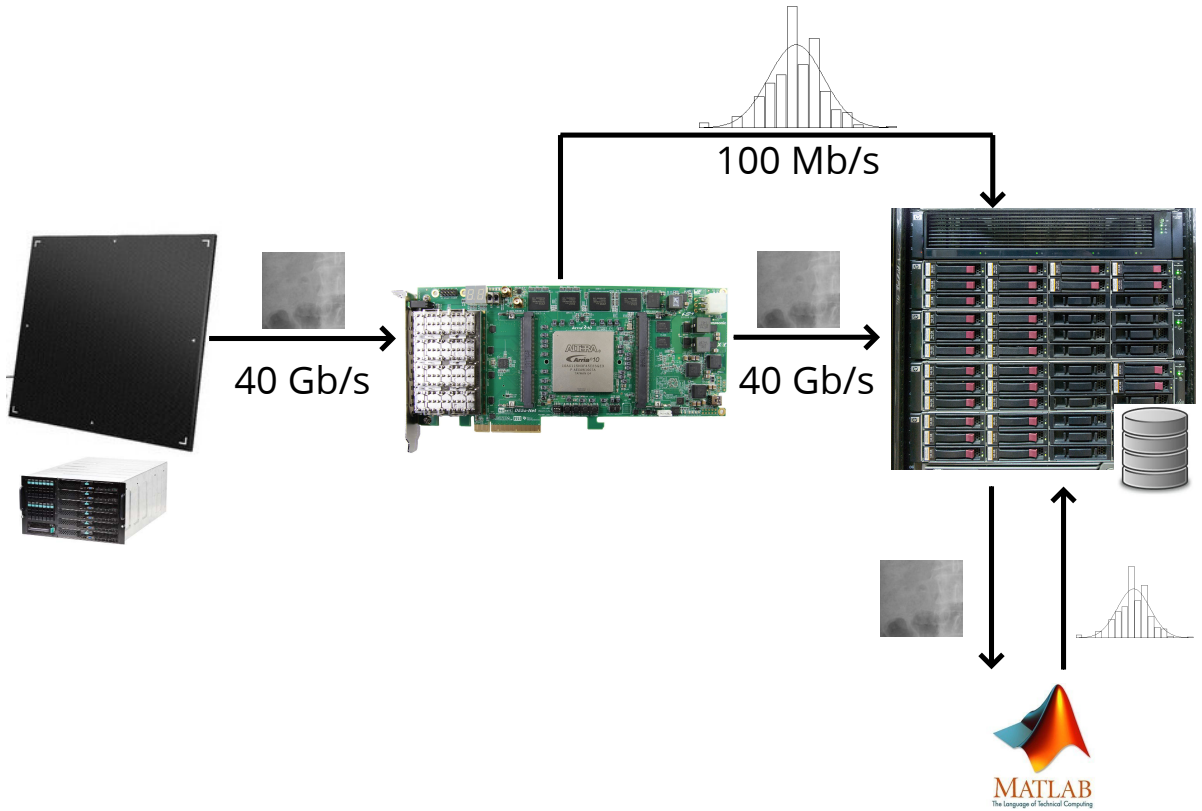

High Speed Real Time Data Analysis

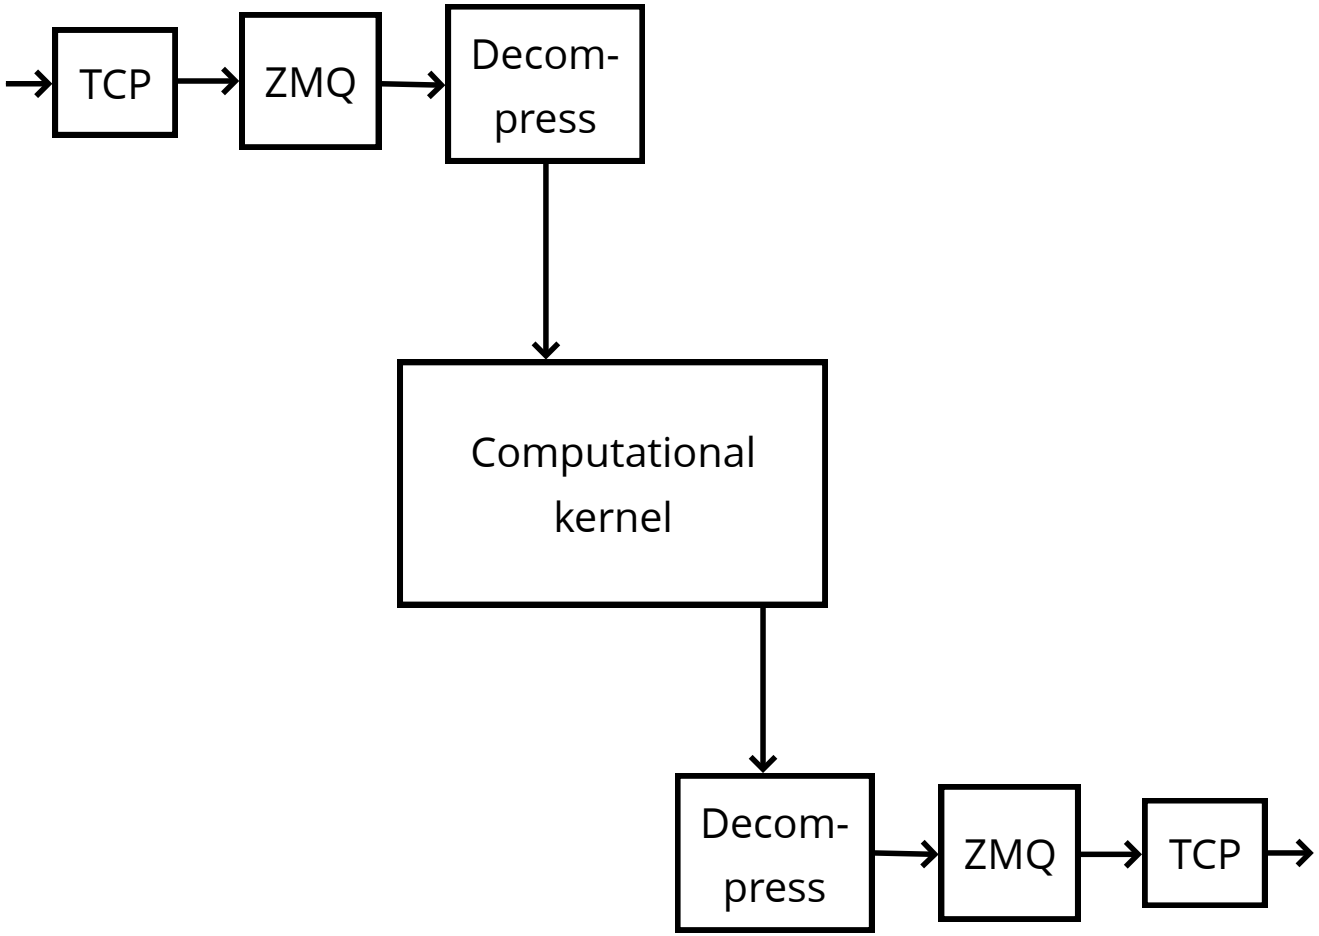

High Speed Real Time Data Analysis

```
def tth2Dsimple(delta,N,M,params):
    # get parameters
    n0 = params['n0'] # n0 - detector zero
    m0 = params['m0'] # m0 - detector zero
    wn = params['wn'] # wn/L
    wm = params['wm'] # wm/L
    phi = params['phi'] # rotation around detector axis
    # calculate pixel coordinates in the lab ref. system
    # apply detector phi-rotation
    c = np.cos(phi)
    s = np.sin(phi)
    tN = c*(N-n0)*wn - s*(M-m0)*wm
    tM = s*(N-n0)*wn + c*(M-m0)*wm
    # main axis rotation
    c = np.cos(delta)
    s = np.sin(delta)
    X = c - s*tN
    Y = s + c*tN
    Z = tM
    R = np.sqrt( X**2 + Y**2 + Z**2 )
    tth = np.arccos(X/R)
    return tth
```

Bohrium

- Bohrium provides automatic acceleration of array operations in Python/NumPy, C, and C++ targeting multi-core CPUs and GP-GPUs.

```
> python compute.py
```



```
> python -m bohrium compute.py
```

Bohrium

```
import numpy as np

a = np.arange(10)
b = (a + 2) * a
c = np.histogram(b)
```

```
a = [10]
t0 = [10]
arange a, 10
add, t0, a, 2
mul b, t0, a
hist c, b
```

OpenMP OpenCL CUDA

Synchronous Message Exchange

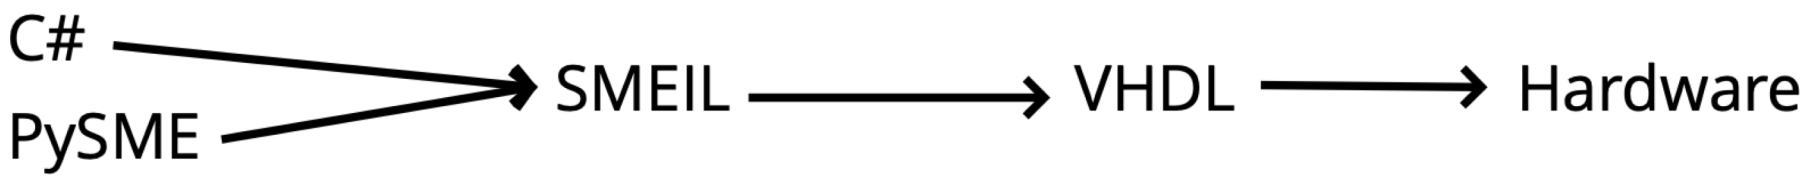

Synchronous Message Exchange

- Simple testing and debugging
- Human readable VHDL
- Automatic testbench

```
public interface ICounterControl : IBus {
  [Initial(false)] bool Valid { get; set; }
  [Initial(false)] bool Reset { get; set; }
}

public interface ICounterData : IBus {
  [Initial(0)] int Value { get; set; }
}

public class Counter {
  private readonly ICounterControl Control = CreateBus<ICounterControl>();
  private readonly ICounterData Data = CreateBus<ICounterData>();

  public void OnReady() {
    if (Control.Reset) {
      Data.Value = 0;
    } else if (Control.Valid) {
      Data.Value++;
    }
  }
}
```

Vision

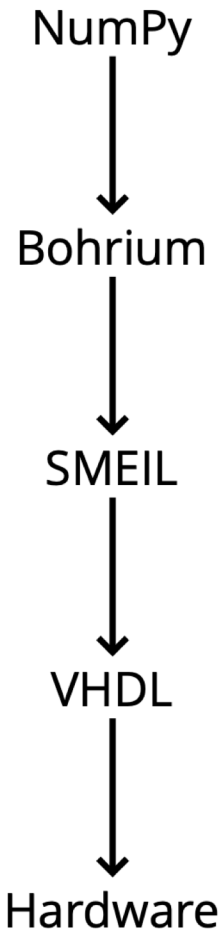


```
import numpy as np

def kernel(a, c):
  a[:] = np.arange(10)
  b = (a + 2) * a
  c[:] = np.histogram(b)
```

Grand Vision

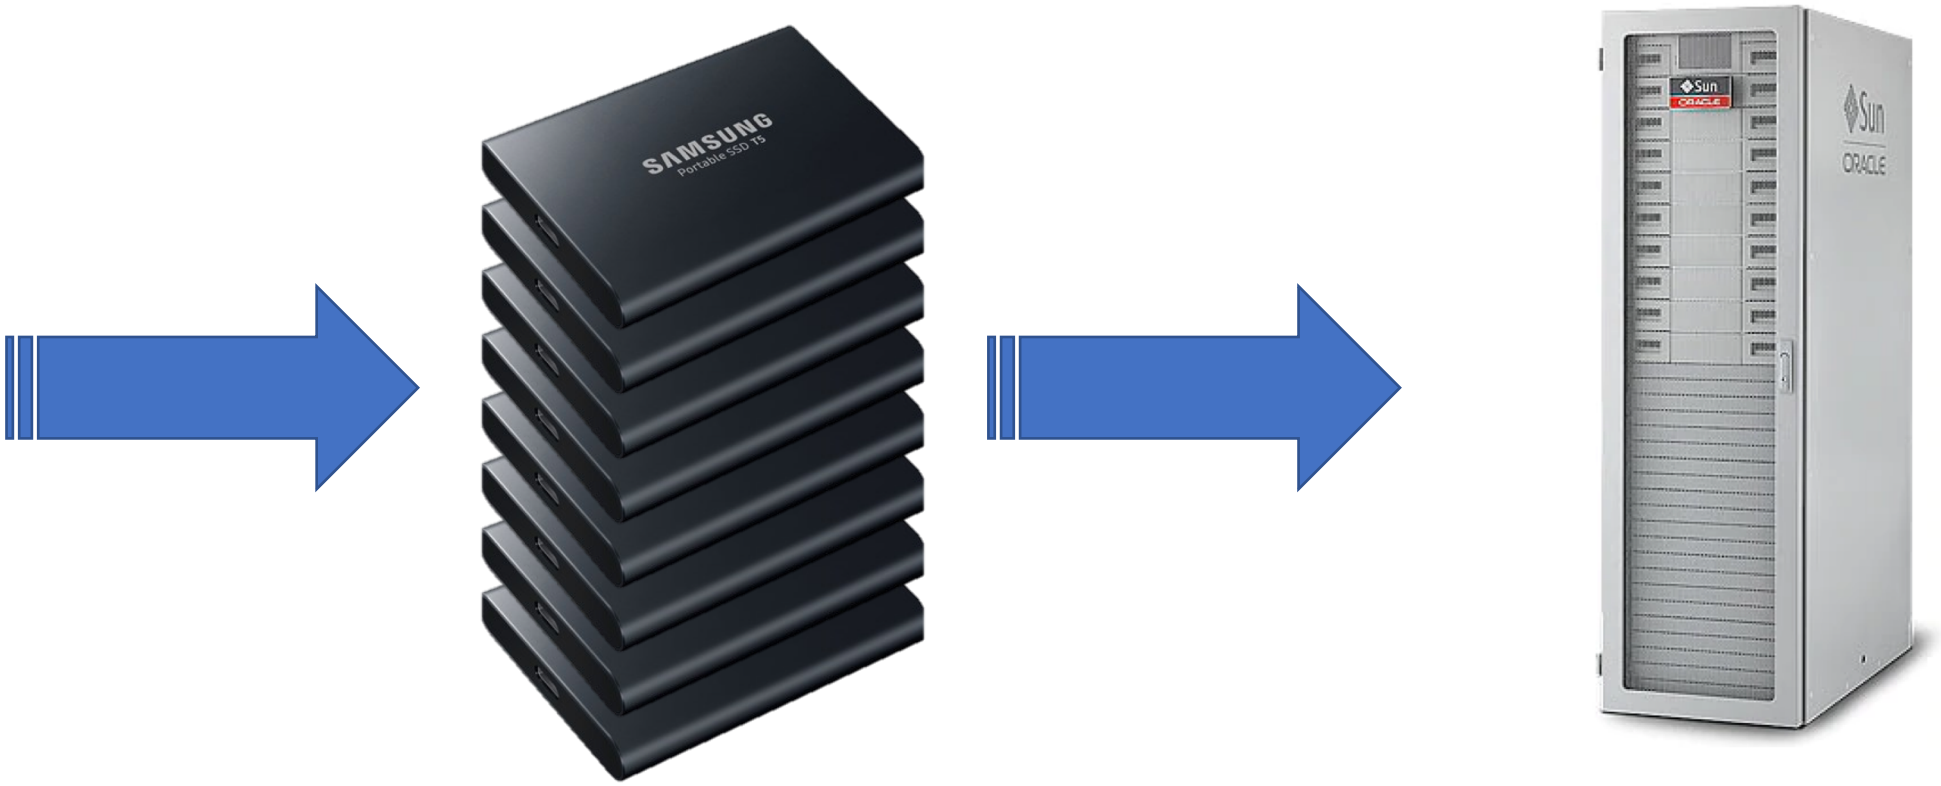

Grand Vision

Proposal: Fighting Scientific Misconduct

- Scientific misconduct is a problem
- With FAIR it may become worse (or not!)
- Proposal: We establish a scientific block-chain
  - Instruments signs the raw data
  - Software that is provided by the facility also sign the result
- Outcome: We can trace the validity of data until the researcher runs untrusted software on the data
  - Which makes it very clear where the problem arises


That's all Folks!