A PRACTICAL LOOK INTO

Ptychography with Extreme Data

13.05.25 – SOLEIL PtyPy Workshop

Yuri Rossi Tonin

What do we mean by extreme data?

$N > 10.000$ (Electrons)

\# of scan points

$$(N, Y, X)$$

\# of rows

\# of columns

$Y * X > 2000^2 \text{ pixels}^2$ (X-Rays)

Farfield

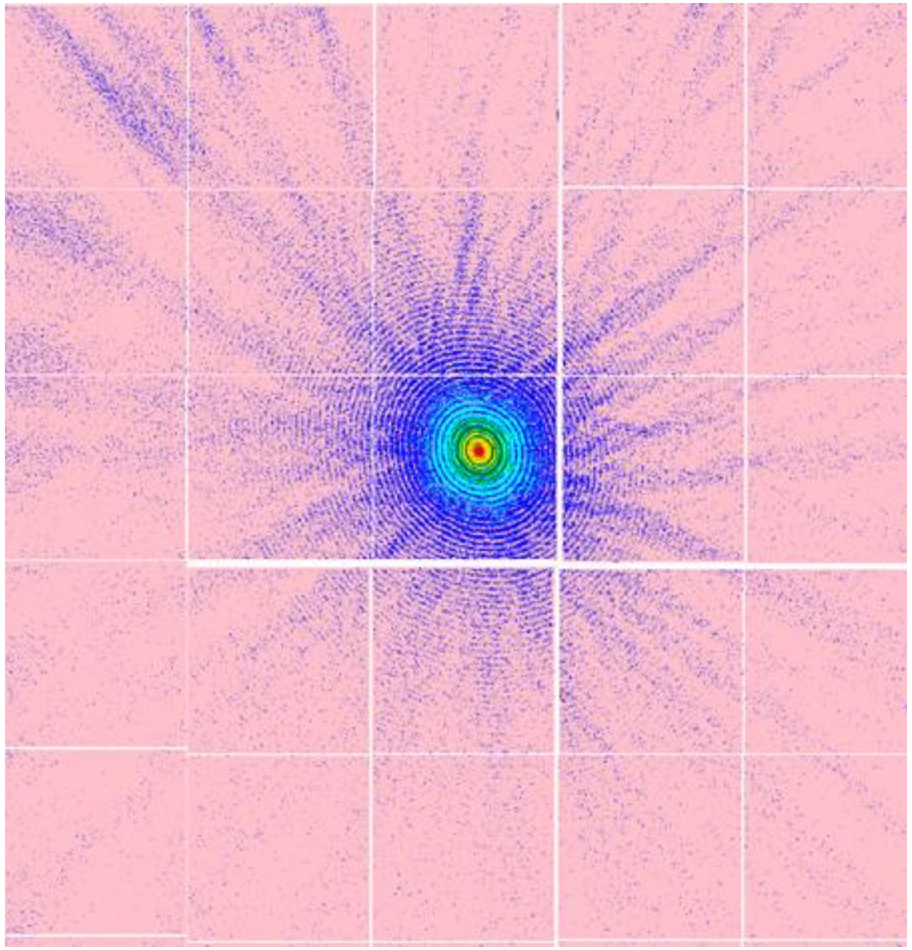

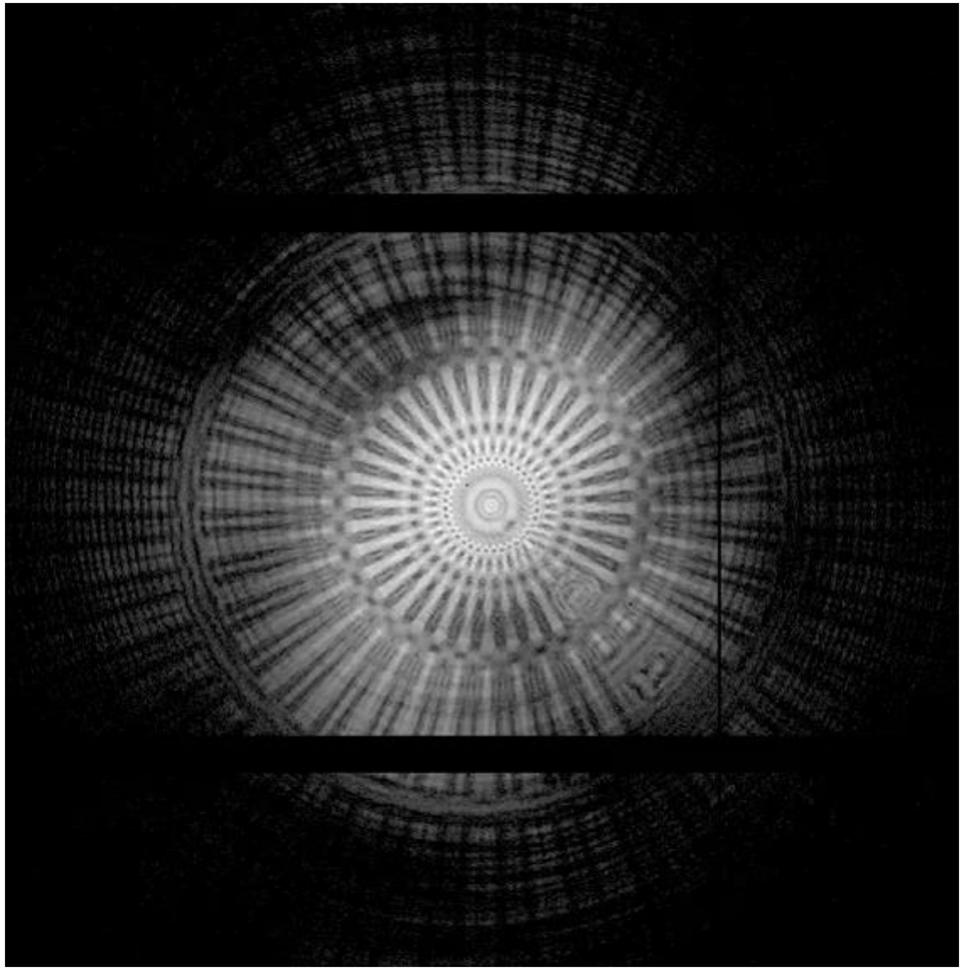

Nearfield

EIGER

4150x4371
75 µm
133 Hz

3110x3269
75 µm
238 Hz

2070x2167
75 µm
750 Hz

190x190
100 µm
30 – 120 KHz

SELUN

3072x3072
55 µm
up to 2kHz

PiMEGA

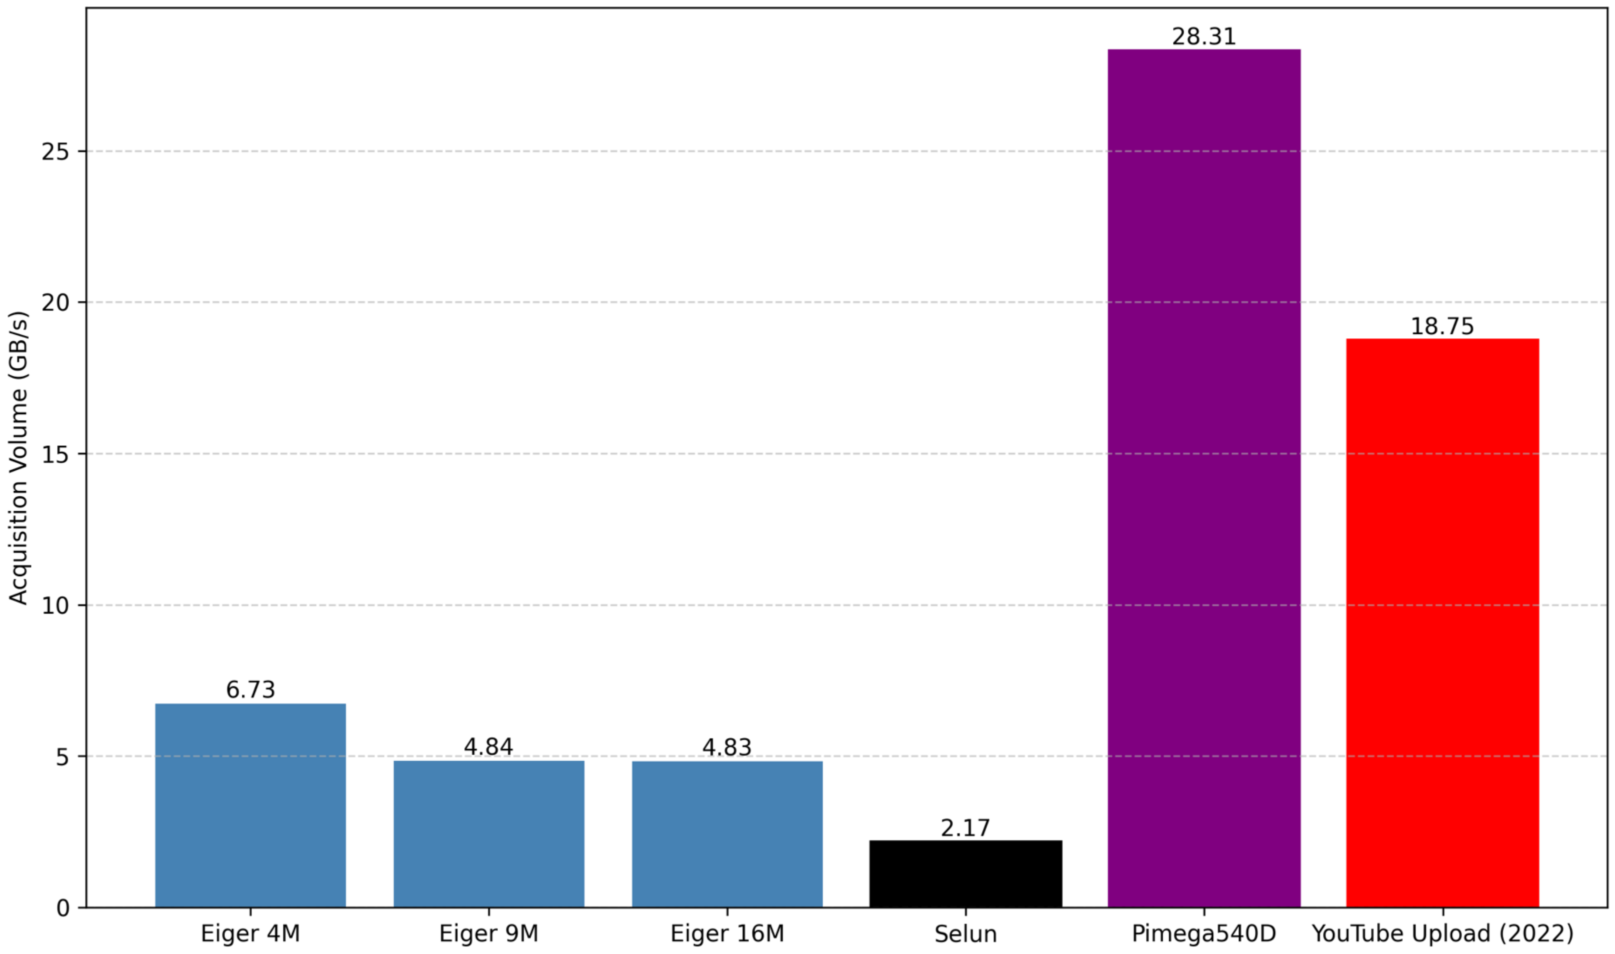
Acquisition Volume (GB/s)
0
5
10
15
20
25
6.73
4.84
4.83
2.17
28.31
18.75
Eiger 4M
Eiger 9M
Eiger 16M
Selun
Pimega540D
YouTube Upload (2022)

Ptychographic X-ray computed tomography of porous membranes with nanoscale resolution

Radosław Górecki [1,2], Carla Cristina Polo [3], Tiago Araujo Kalile[3], Eduardo X. S. Miqueles [3], Yuri R. Tonin[3], Lakshmeesha Upadhyaya[1,2], Florian Meneau [3,4] & Suzana P. Nunes [1,2,5]✉

(2552,768,768) @ 8 bytes = 12 GBs

498 projections in total
≅
6 TBs raw data

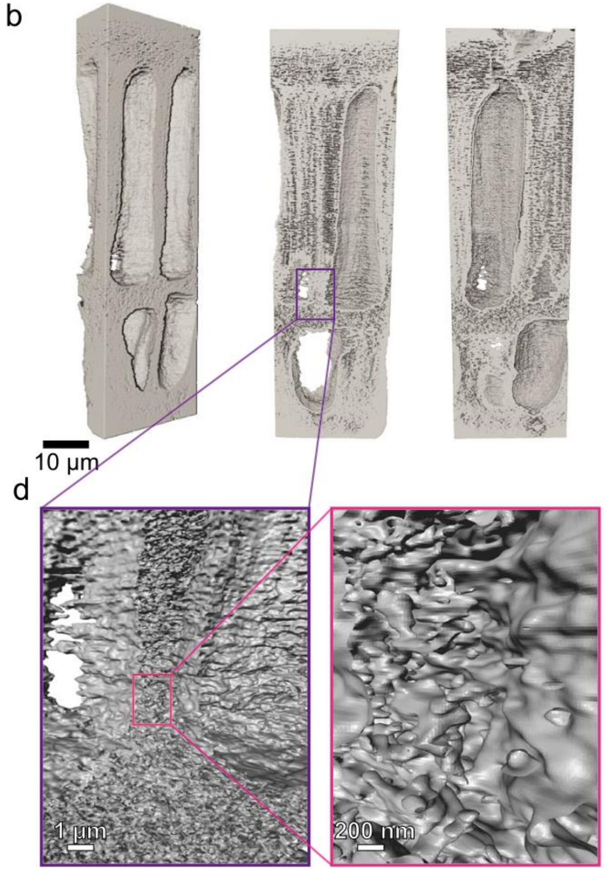

Górecki, R., *et al.* Ptychographic X-ray computed tomography of porous membranes with nanoscale resolution. *Commun Mater* **4**, 68 (2023).

Fast X-ray ptychography: towards nanoscale imaging of large volume of brain

Silvia Cipiccia[1,a], Michela Fratini[2,3,b], Ecem Erin[1,c], Marco Palombo[4,d], Silvia Vogel[5,e], Max Burian[5,f], Fenglei Zhou[1,6,g], Geoff J. M. Parker[1,7,h], Darren J. Batey[8,i]

facilities for high resolution imaging, with applications spanning from life science to nano-electronics. In the recent years there has been a great effort to make the technique faster to enable high throughput nanoscale imaging. Here we apply a fast ptychography scanning method to image in 3D $10^6 \mu m^3$ of brain-like phantom at 3 kHz, in a 7 h acquisition with a resolution of 270 nm. We then present the latest advances in fast ptychography by showing 2D images acquired at 110 kHz by combining the fast-acquisition scheme with a high-acquisition rate prototype detector from DECTRIS Ltd. We finally review the experimental outcome and discuss

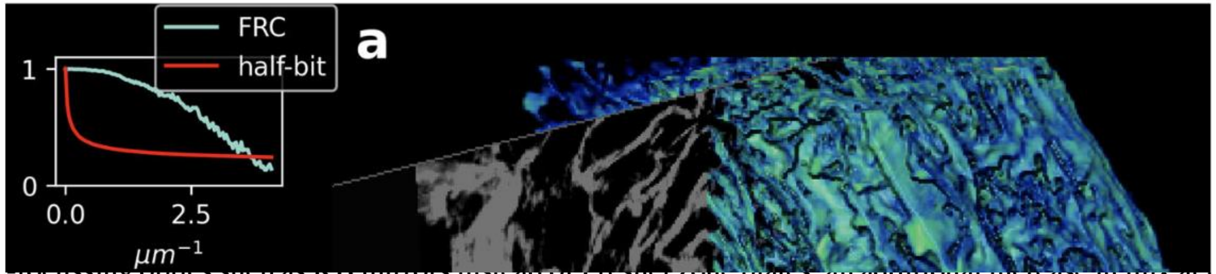

efficiency pre-focussing optics such as KB mirrors instead of Fresnel zone plates, an additional increase in flux at the sample of more than a factor 5 can be foreseen. Therefore, in total we expect a flux over an order of magnitude higher than required for matching Fig. 2. A study by Deng et al. [43] shows that an increase of one order of magnitude in average flux on the sample corresponds to an increase of a factor 2 in resolution. Based on this scaling, it is therefore reasonable to expect 100 nm resolution to be achievable for brain-like samples at 110 kHz ptychographic acquisition in the near future. To be noted that to remain within 200 Mphotons/pixel/s in the central disk, the higher future flux would need to be distributed across the whole detector active area, for example by using a

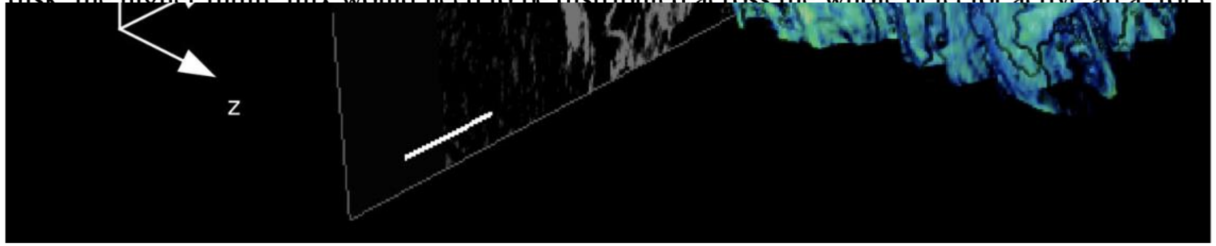

S. Cipiccia et. al., Eur. Phys. J. Plus (2024) 139:434

Article

High-performance 4-nm-resolution X-ray tomography using burst ptychography

https://doi.org/10.1038/s41586-024-07615-6

Received: 28 June 2023

Tomas Aidukas[1✉], Nicholas W. Phillips[1,6], Ana Diaz[1], Emiliya Poghosyan[1], Elisabeth Müller[1], A. F. J. Levi[2], Gabriel Aeppli[1,3,4,5], Manuel Guizar-Sicairos[1,4] & Mirko Holler[1✉]

c

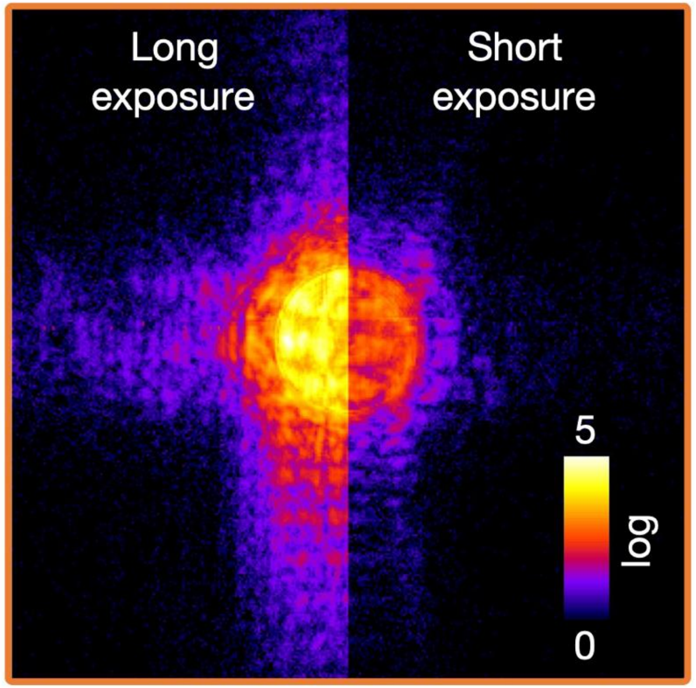

b Burst-frame position refinement

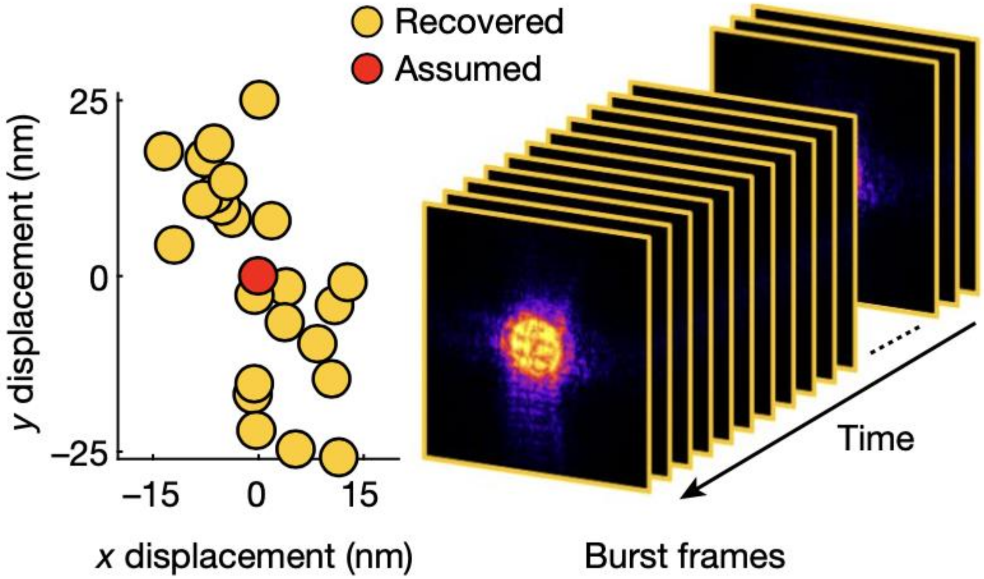

T. Aidukas et al. Nature Vol. 632. (2024).

```
O, P = ptychography(diffraction_data,
                    initial_object,
                    initial_probe,
                    scan_positions,
                    mask,
                    parameters)
```

Data types

storage processing

Diffraction data: (N,Y,X) uint32 -> float32

Object: (M_o, A, B) complex64

Probe: (M_p, Y, X) complex64

~~Positions: (N,2) float32~~

Mask: (Y, X) boolean

https://ptycho.github.io/tutorials/

Reconstruction Engines

Learning more about the different engines and their parameters.

PtyPy offers a range of different reconstruction engines that can be grouped into the 3 main categories of

- Projectional engines (DM, RAAR)
- Stochastic engines (ePIE, SDR)
- Gradient-based engines (ML)

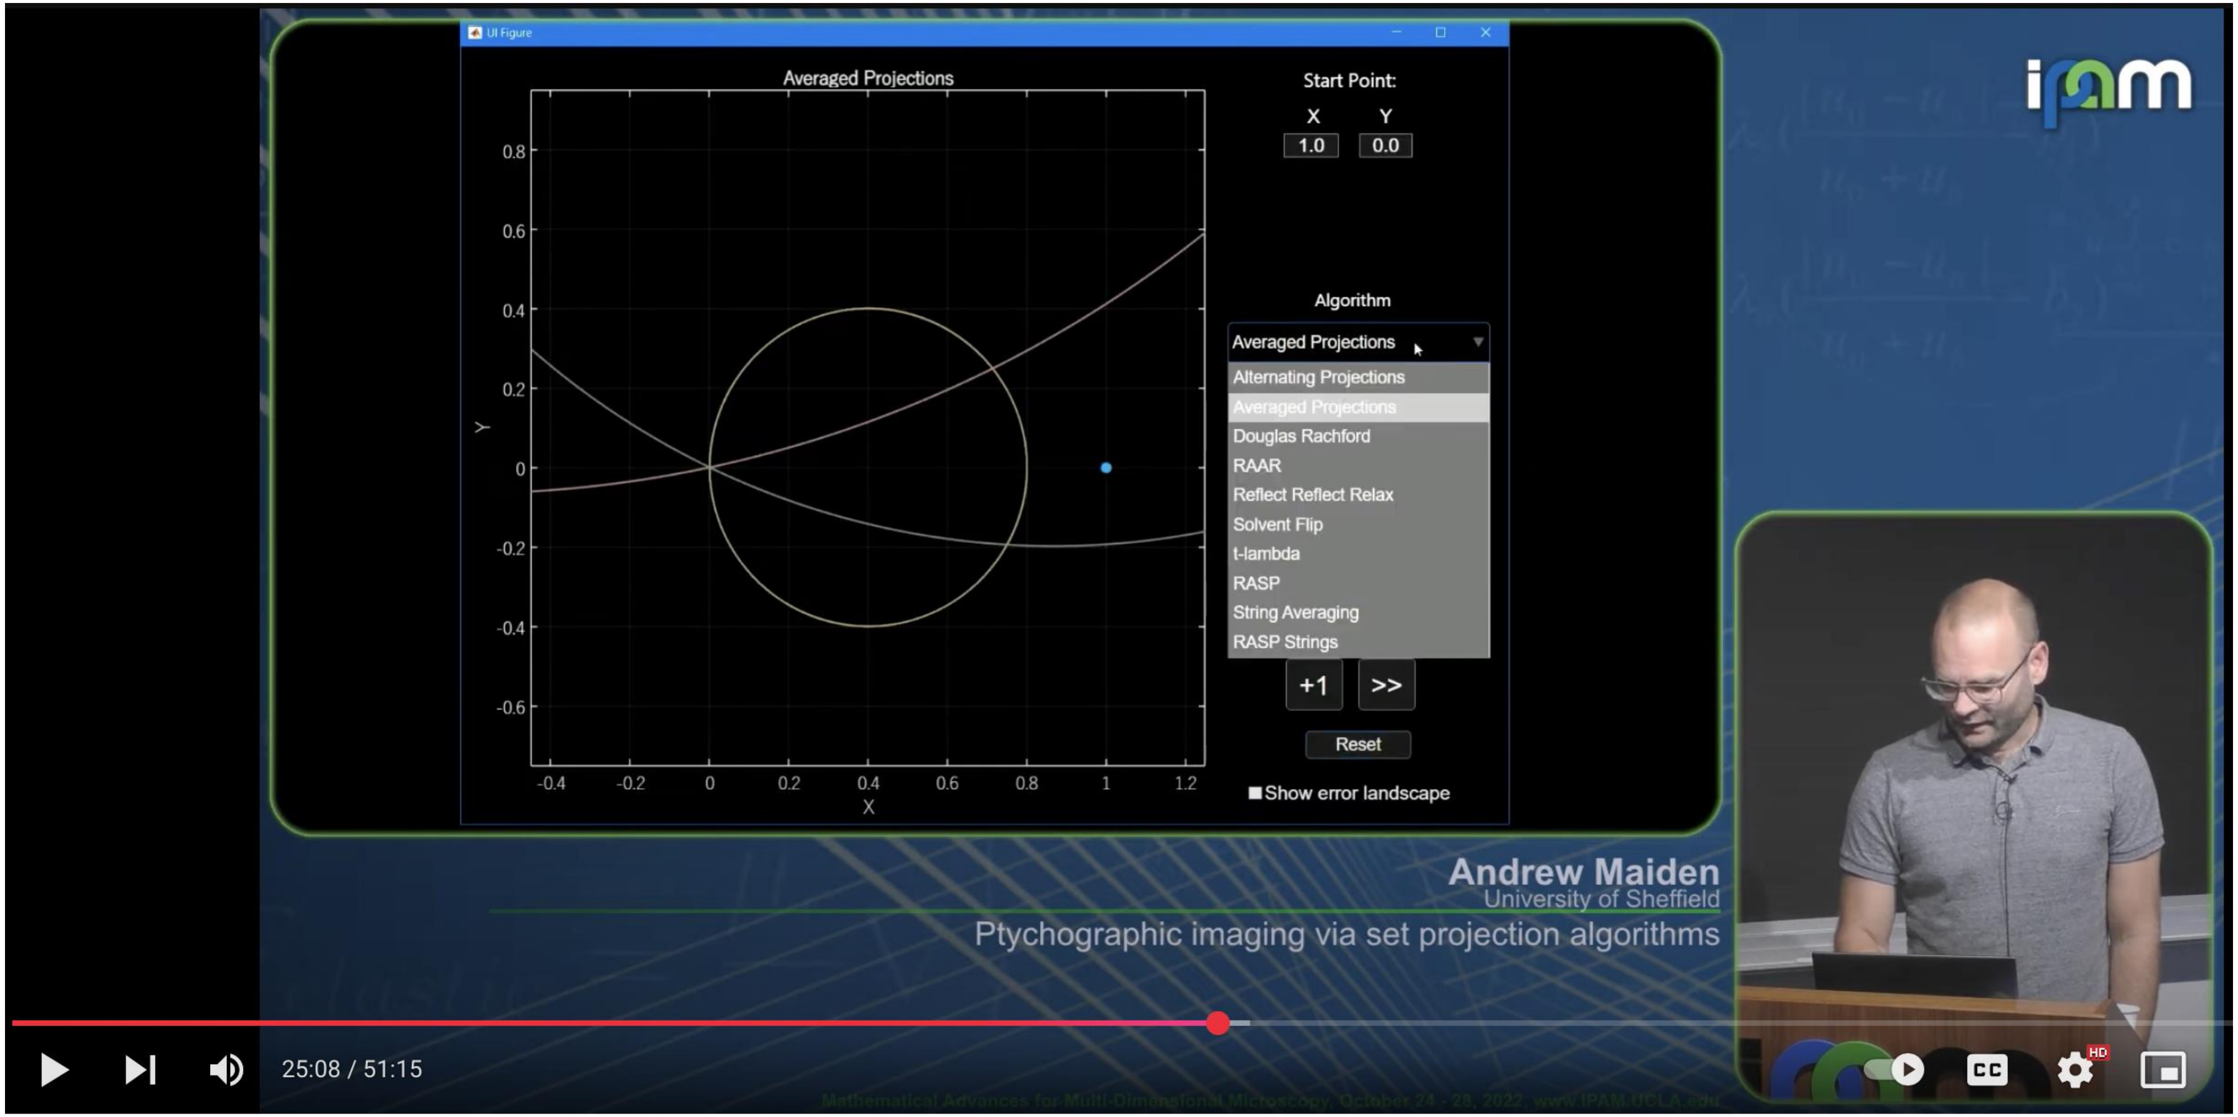

Andrew Maiden - Ptychographic imaging via set projection algorithms - IPAM at UCLA

Reconstruction Engines

Learning more about the different engines and their parameters.

PtyPy offers a range of different reconstruction engines that can be grouped into the 3 main categories of

- Projectional engines (DM, RAAR)
- Stochastic engines (ePIE, SDR)
- Gradient-based engines (ML)

`BlockFull` is another variation of `Full` designed for large data and accelerated engines. `BlockFull` requires `p.frames_per_block` to be set to a reasonable value.

DM, RAAR:

$$O'(\mathbf{r}) = \frac{\sum_i^N P^*(\mathbf{r}-\mathbf{r_i})\psi(\mathbf{r})}{\sum_i^N |P(\mathbf{r}-\mathbf{r_i})|^2},$$

$$P'(\mathbf{r}) = \frac{\sum_i^N O^*(\mathbf{r}+\mathbf{r_i})\psi(\mathbf{r}+\mathbf{r_i})}{\sum_i^N |O(\mathbf{r}+\mathbf{r_i})|^2}.$$

PIE:

$$O'(\mathbf{r}-\mathbf{r_i}) = O(\mathbf{r}-\mathbf{r_i}) + w_i^o(\mathbf{r})P^*(\mathbf{r})(\psi_i'(\mathbf{r}) - \psi_i(\mathbf{r})),$$

$$P'(\mathbf{r}) = P(\mathbf{r}) + w_i^p(\mathbf{r})O^*(\mathbf{r}-\mathbf{r_i})(\psi_i'(\mathbf{r}) - \psi_i(\mathbf{r})),$$

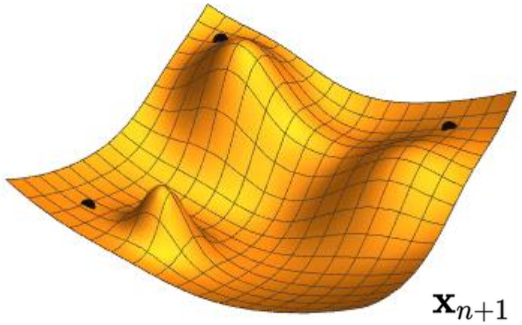

$$\mathbf{x}_{n+1} = \mathbf{x}_n - \gamma_n \nabla F(\mathbf{x}_n)$$

ML:

$$L_{\text{Poisson}} = \sum_i^N M_i(u,v)\left(\sqrt{|D_d\{\psi_i(\mathbf{r})\}|^2} - \sqrt{I_i}\right)^2$$

Can fit your data, but will probably be too slow.

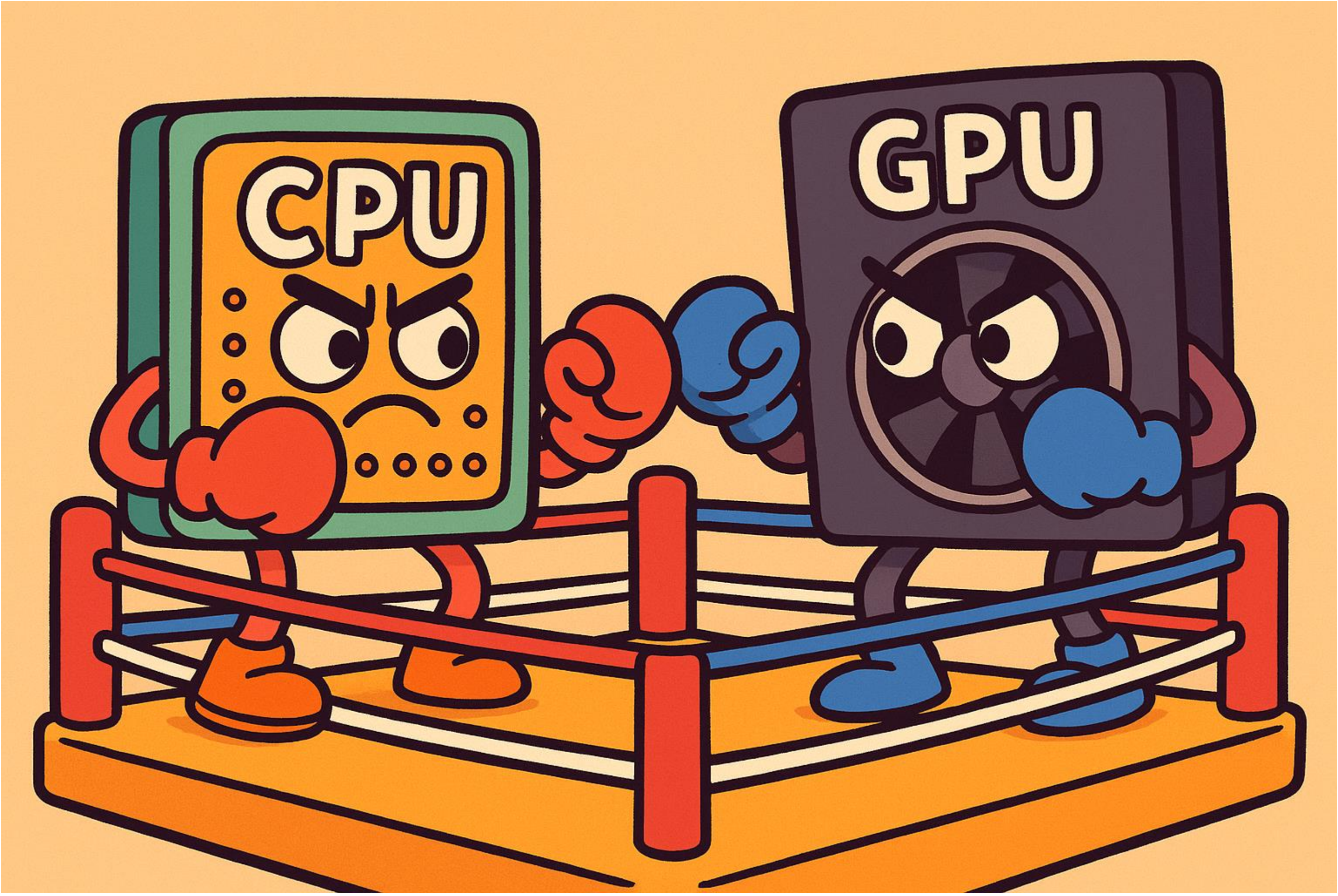

Is way faster, but may not be able to fit your whole data

NVIDIA DGX

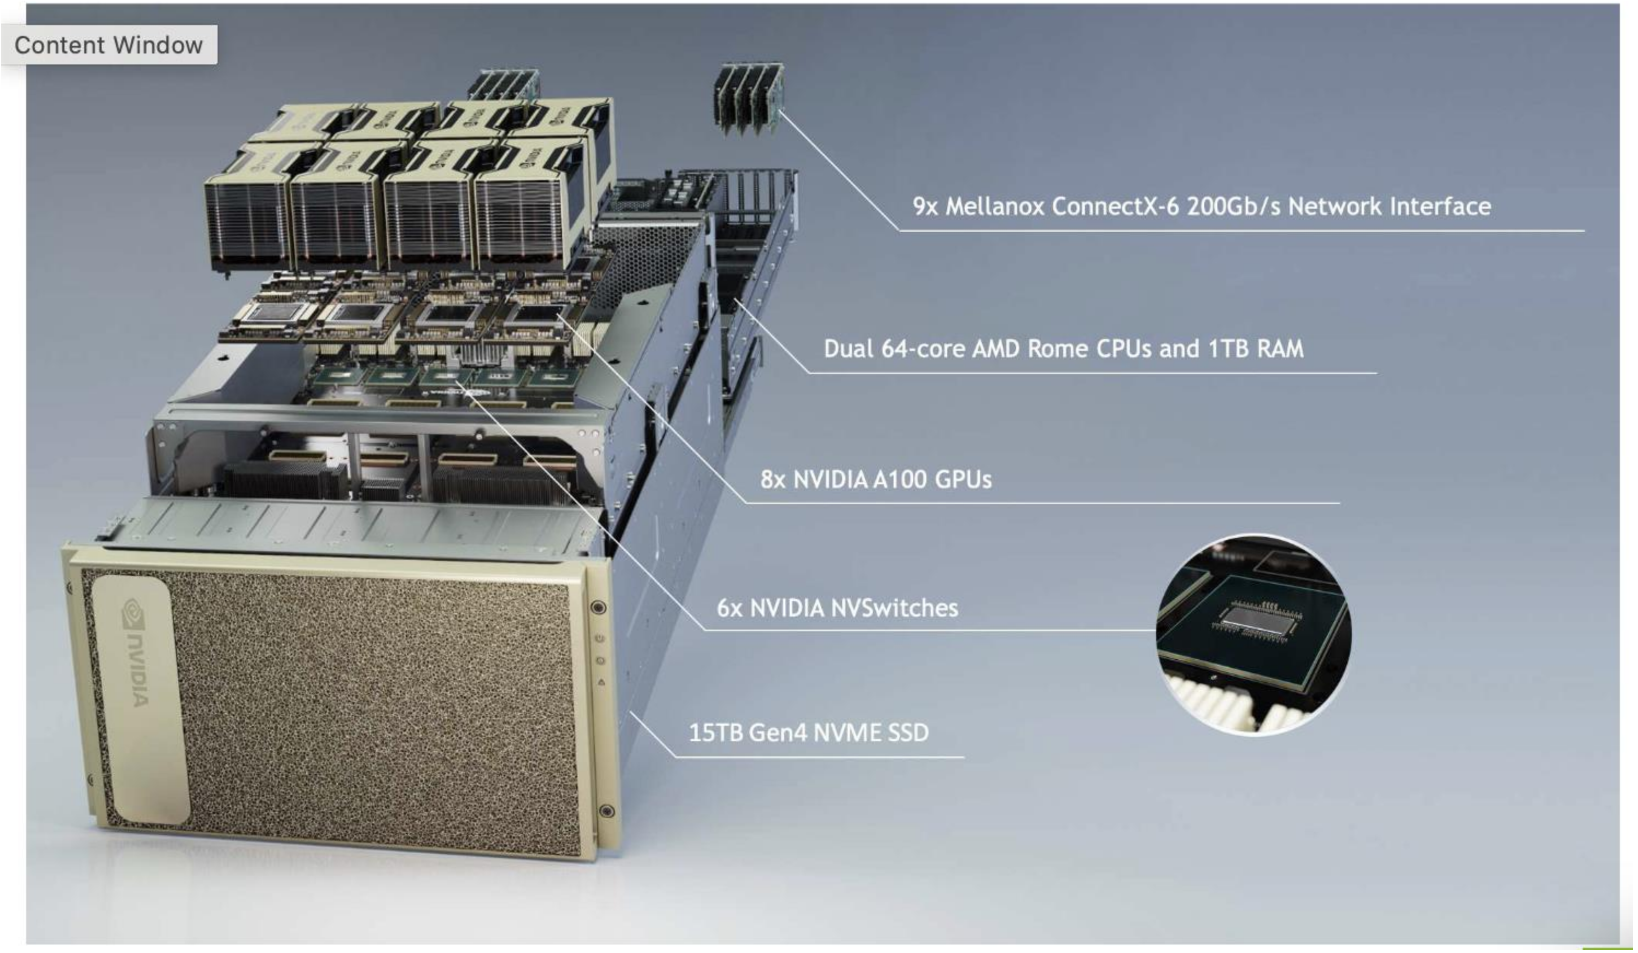
Content Window
9x Mellanox ConnectX-6 200Gb/s Network Interface
Dual 64-core AMD Rome CPUs and 1TB RAM
8x NVIDIA A100 GPUs
6x NVIDIA NVSwitches
15TB Gen4 NVME SSD

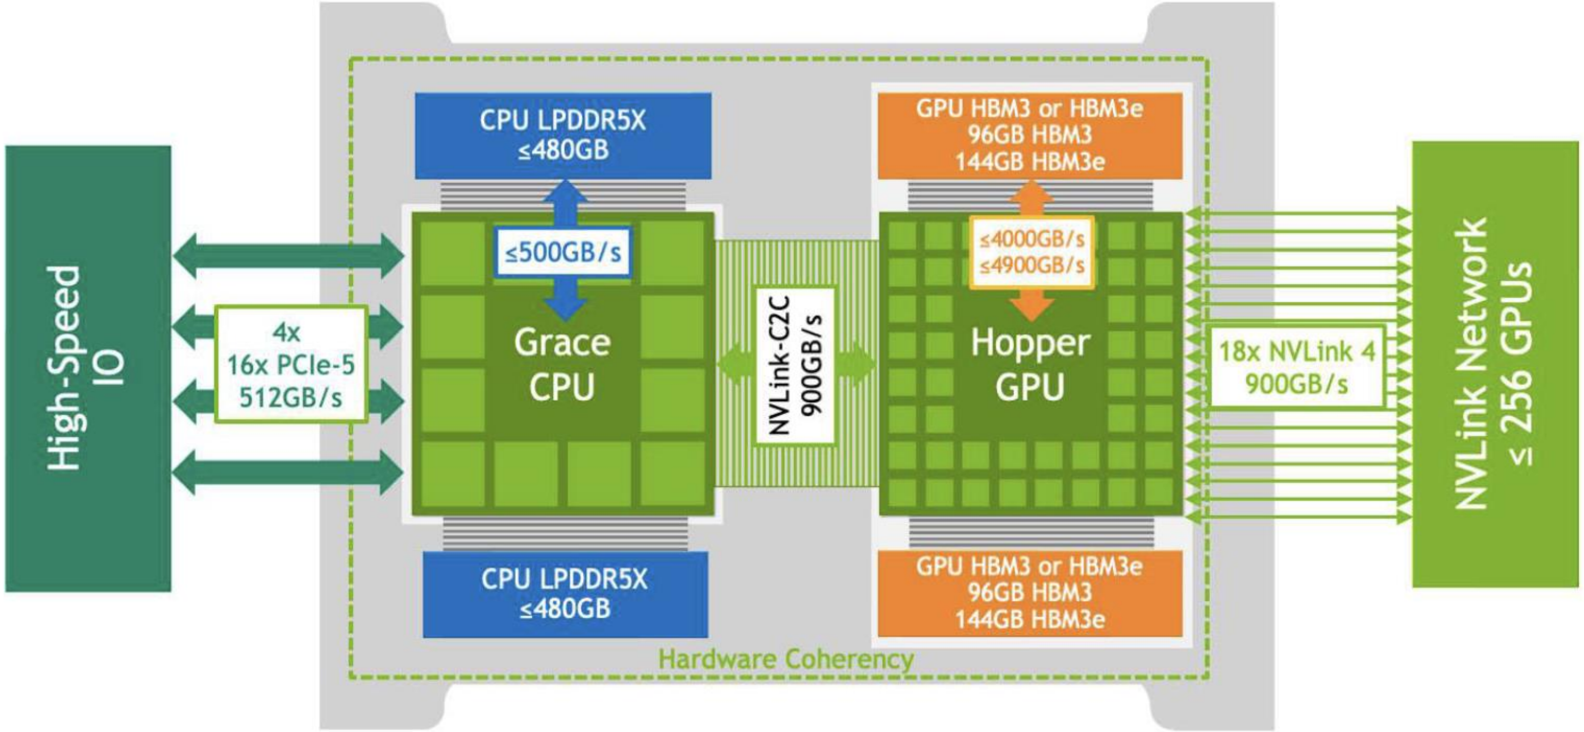
High-Speed IO
4x
16x PCIe-5
512GB/s
CPU LPDDR5X
≤480GB
≤500GB/s
Grace CPU
CPU LPDDR5X
≤480GB
NVLink-C2C
900GB/s
GPU HBM3 or HBM3e
96GB HBM3
144GB HBM3e
≤4000GB/s
≤4900GB/s
Hopper GPU
GPU HBM3 or HBM3e
96GB HBM3
144GB HBM3e
Hardware Coherency
18x NVLink 4
900GB/s
NVLink Network
≤ 256 GPUs

$$\psi_i'(\mathbf{r}) = D_d^{-1}\Pi_M D_d\{\psi_i(\mathbf{r})\}$$

$$O'(\mathbf{r}) = \frac{\sum_i^N P^*(\mathbf{r}-\mathbf{r_i})\psi(\mathbf{r})}{\sum_i^N |P(\mathbf{r}-\mathbf{r_i})|^2},$$

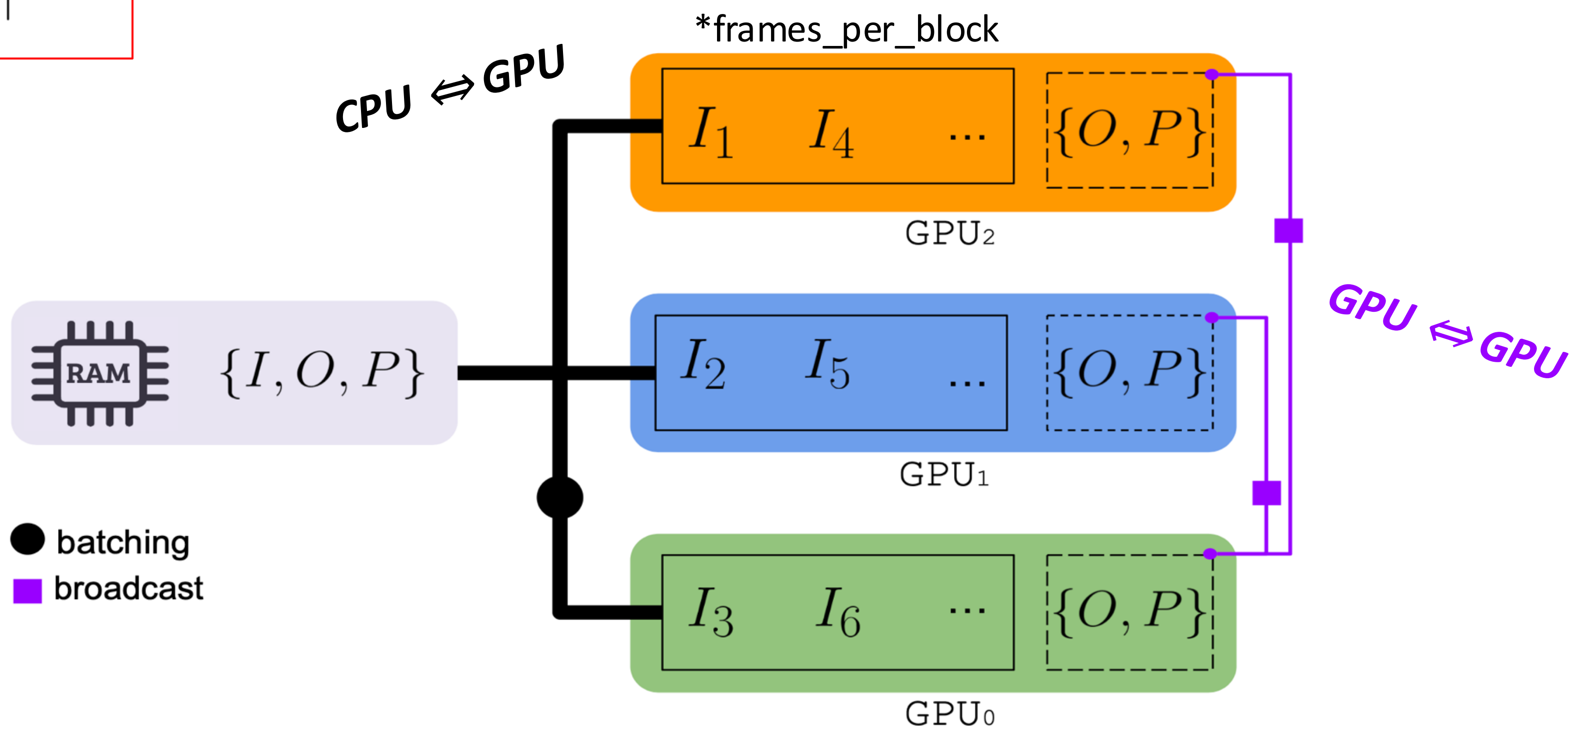

NVIDIA DGX

System Specifications

Partition	# of nodes	CPU	GPU	NIC
GPU	1536	1x AMD EPYC 7763	4x NVIDIA A100 (40GB)	4x HPE Slingshot 11
	256	1x AMD EPYC 7763	4x NVIDIA A100 (80GB)	4x HPE Slingshot 11
CPU	3072	2x AMD EPYC 7763	-	1x HPE Slingshot 11
Login	40	2x AMD EPYC 7713	1x NVIDIA A100 (40GB)	-

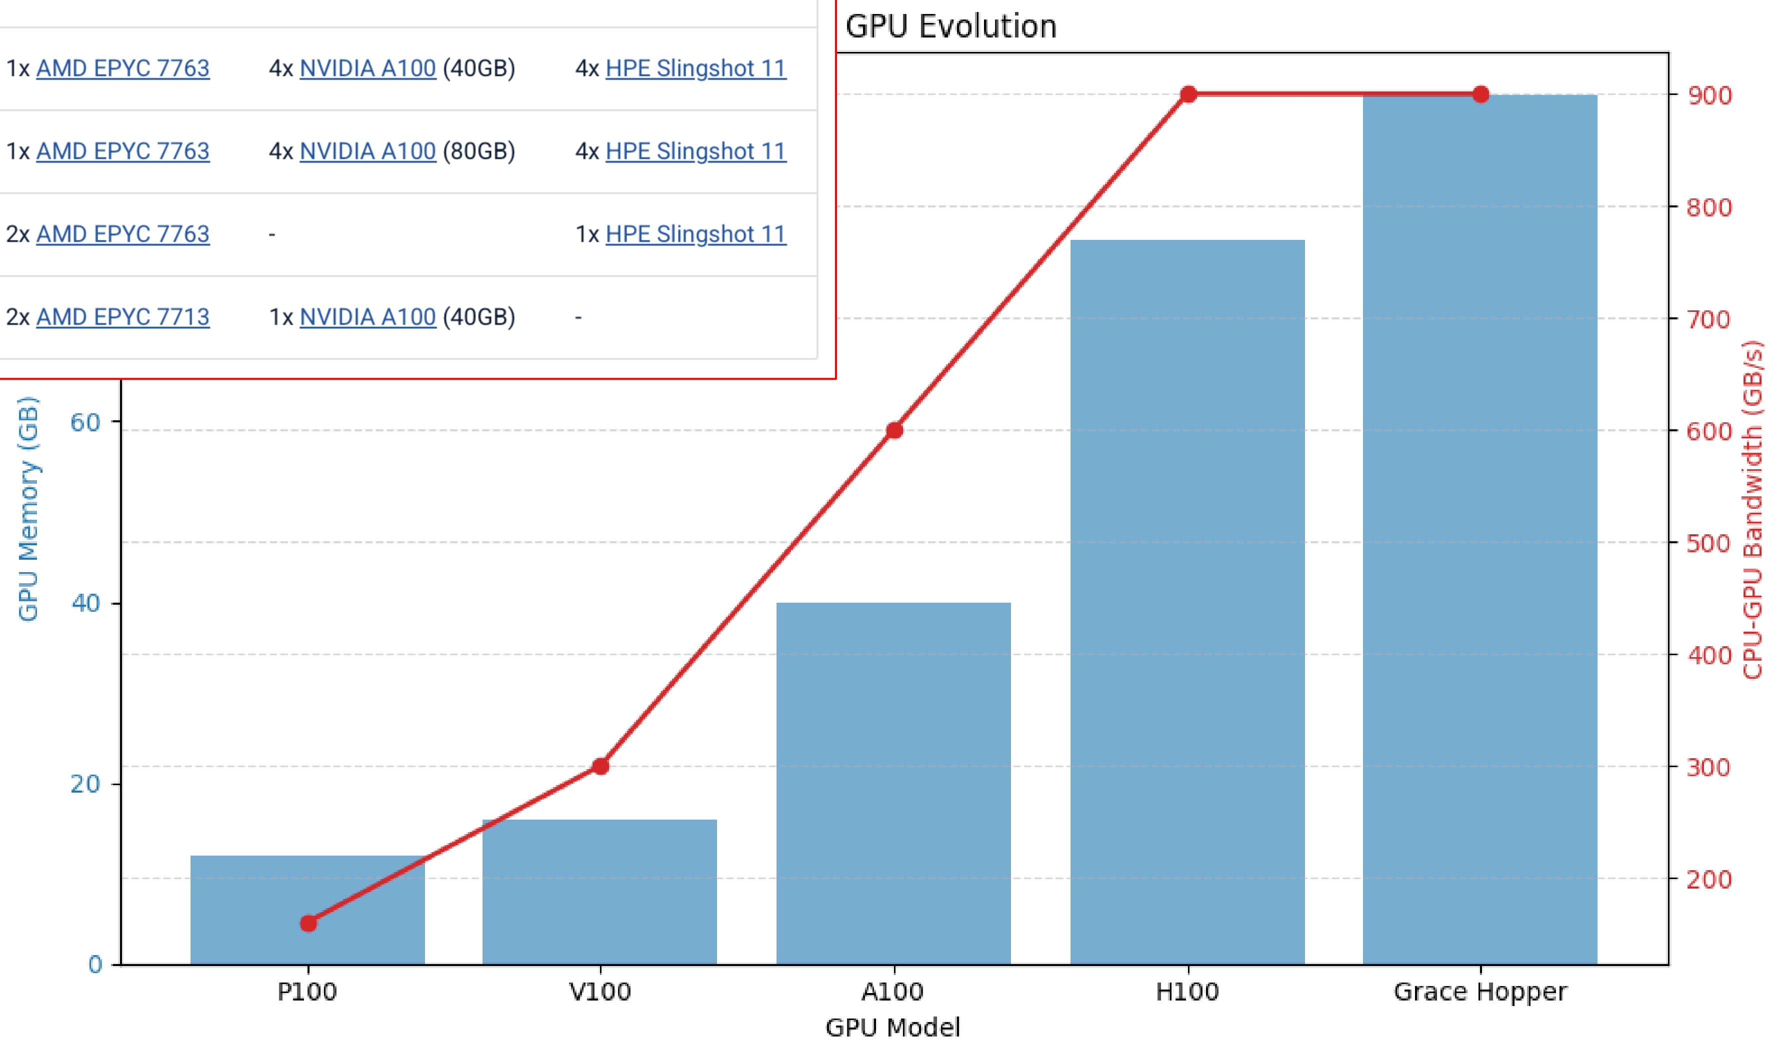

How much do we actually need?

PIE: $\psi_i'(\mathbf{r}) = D_d^{-1}\Pi_M D_d\{\psi_i(\mathbf{r})\}$

DM: $\psi_i' = \psi_i + \Pi_M\{2\Pi_O\{\psi_i\} - \psi_i\} - \Pi_O\{\psi_i\}$

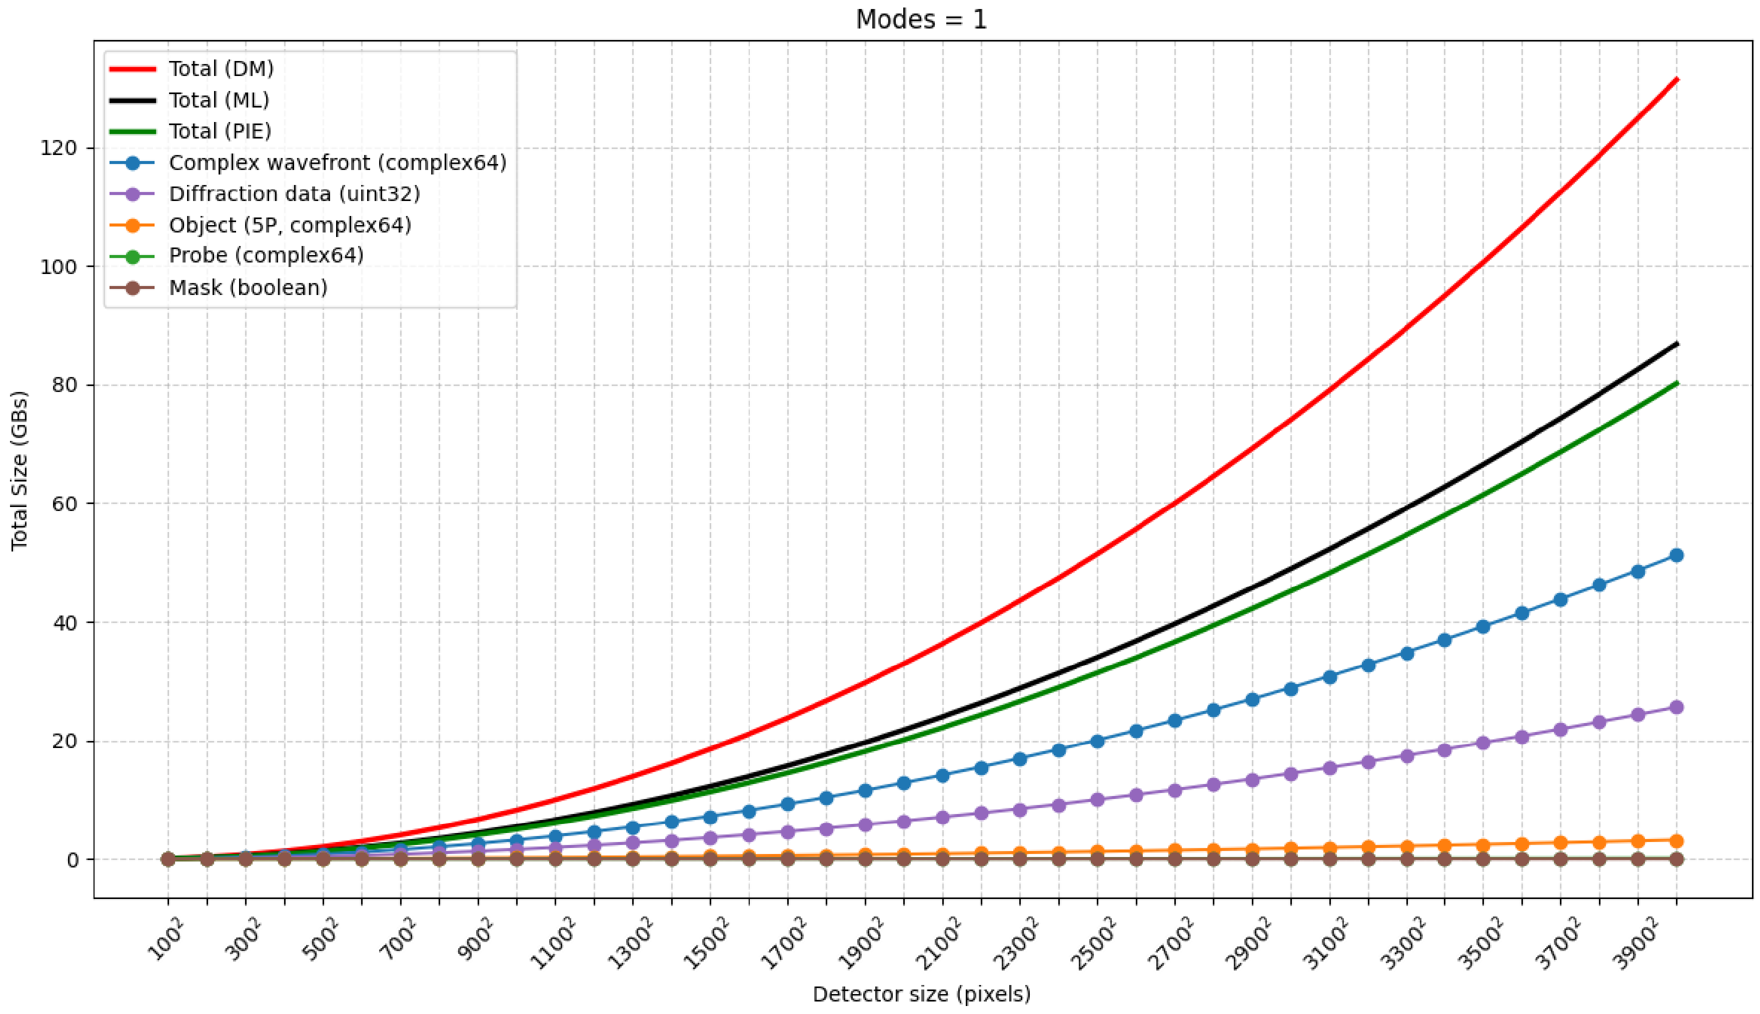

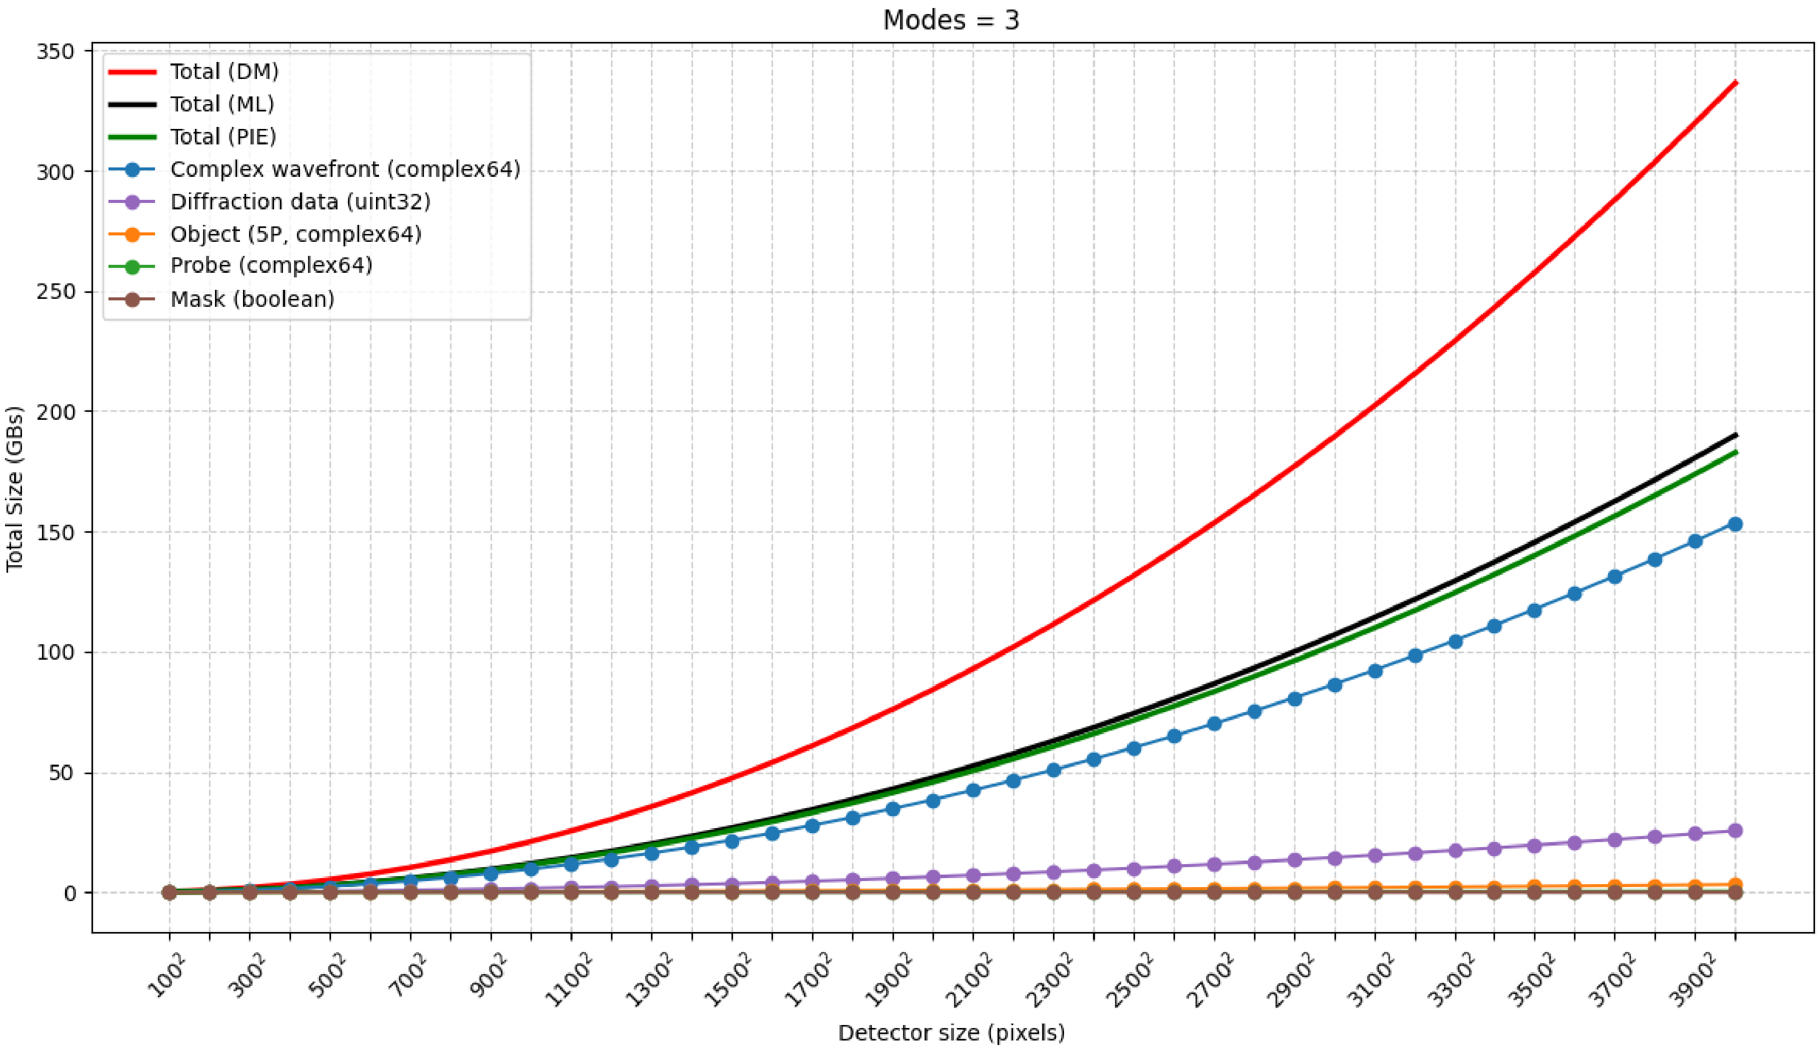

Modes = 3
Total (DM)
Total (ML)
Total (PIE)
Complex wavefront (complex64)
Diffraction data (uint32)
Object (5P, complex64)
Probe (complex64)
Mask (boolean)
Total Size (GBs)
Detector size (pixels)
0
50
100
150
200
250
300
350
100^2
300^2
500^2
700^2
900^2
1100^2
1300^2
1500^2
1700^2
1900^2
2100^2
2300^2
2500^2
2700^2
2900^2
3100^2
3300^2
3500^2
3700^2
3900^2

Strong and Weak Scaling

Strong Scaling

How much do I need to get it running as fast as I want?

$$S_{\max} = \frac{1}{(1-P) + \frac{P}{p}}$$

P: parallel fraction

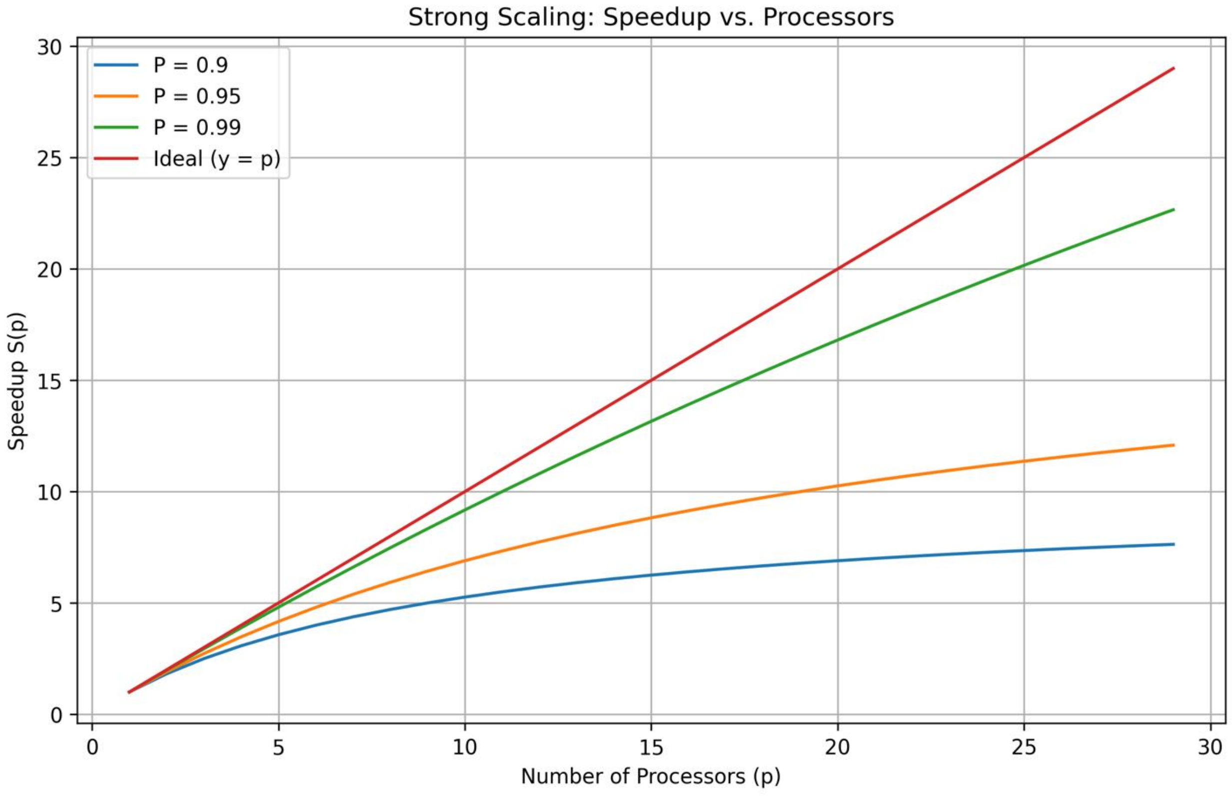

- Strong scaling: N MORE nodes does not necessarily means N x FASTER

Weak Scaling

Given I have time t, how large can I make my problem?

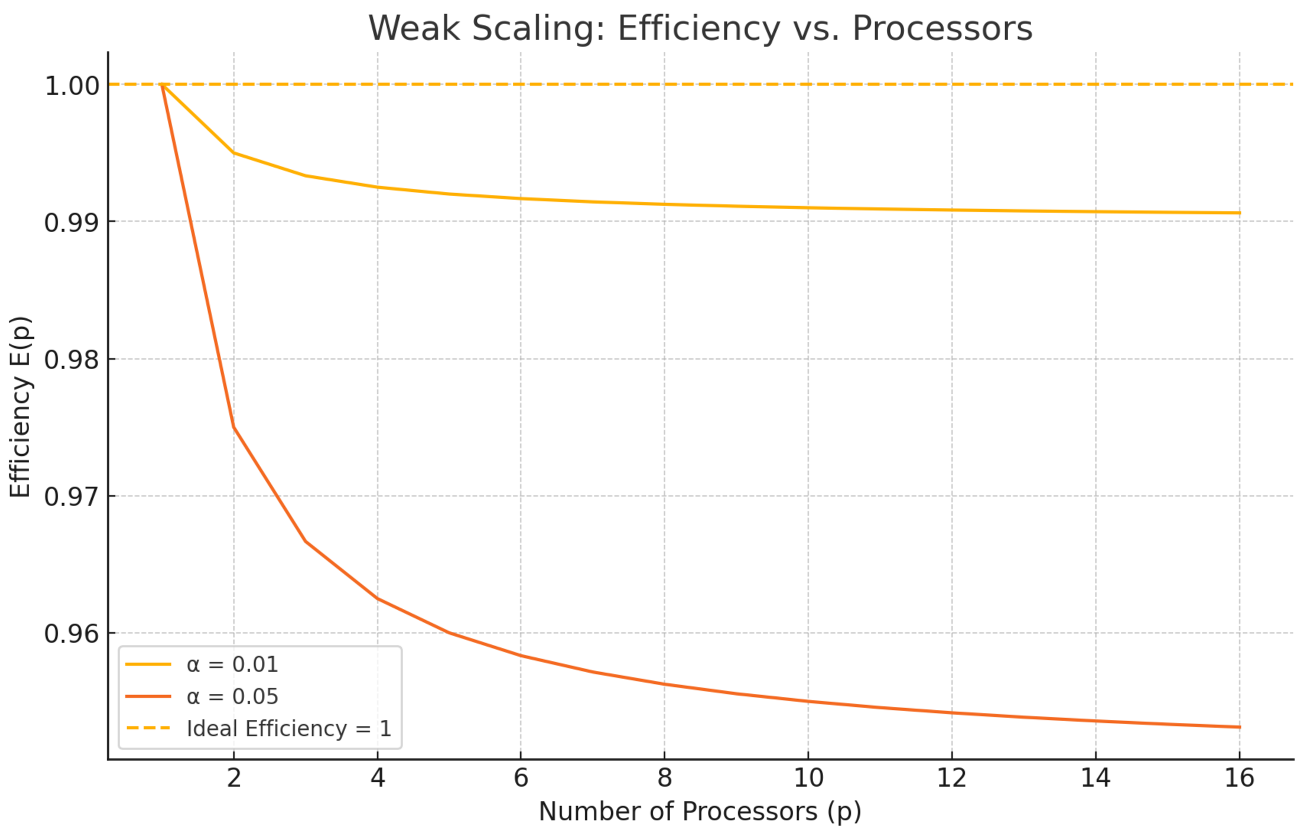

- ➢ Weak scaling: N MORE nodes does not necessarily means N x LARGER scan in the same time

Some benchmarks!

@hall_of_tech

single GPU

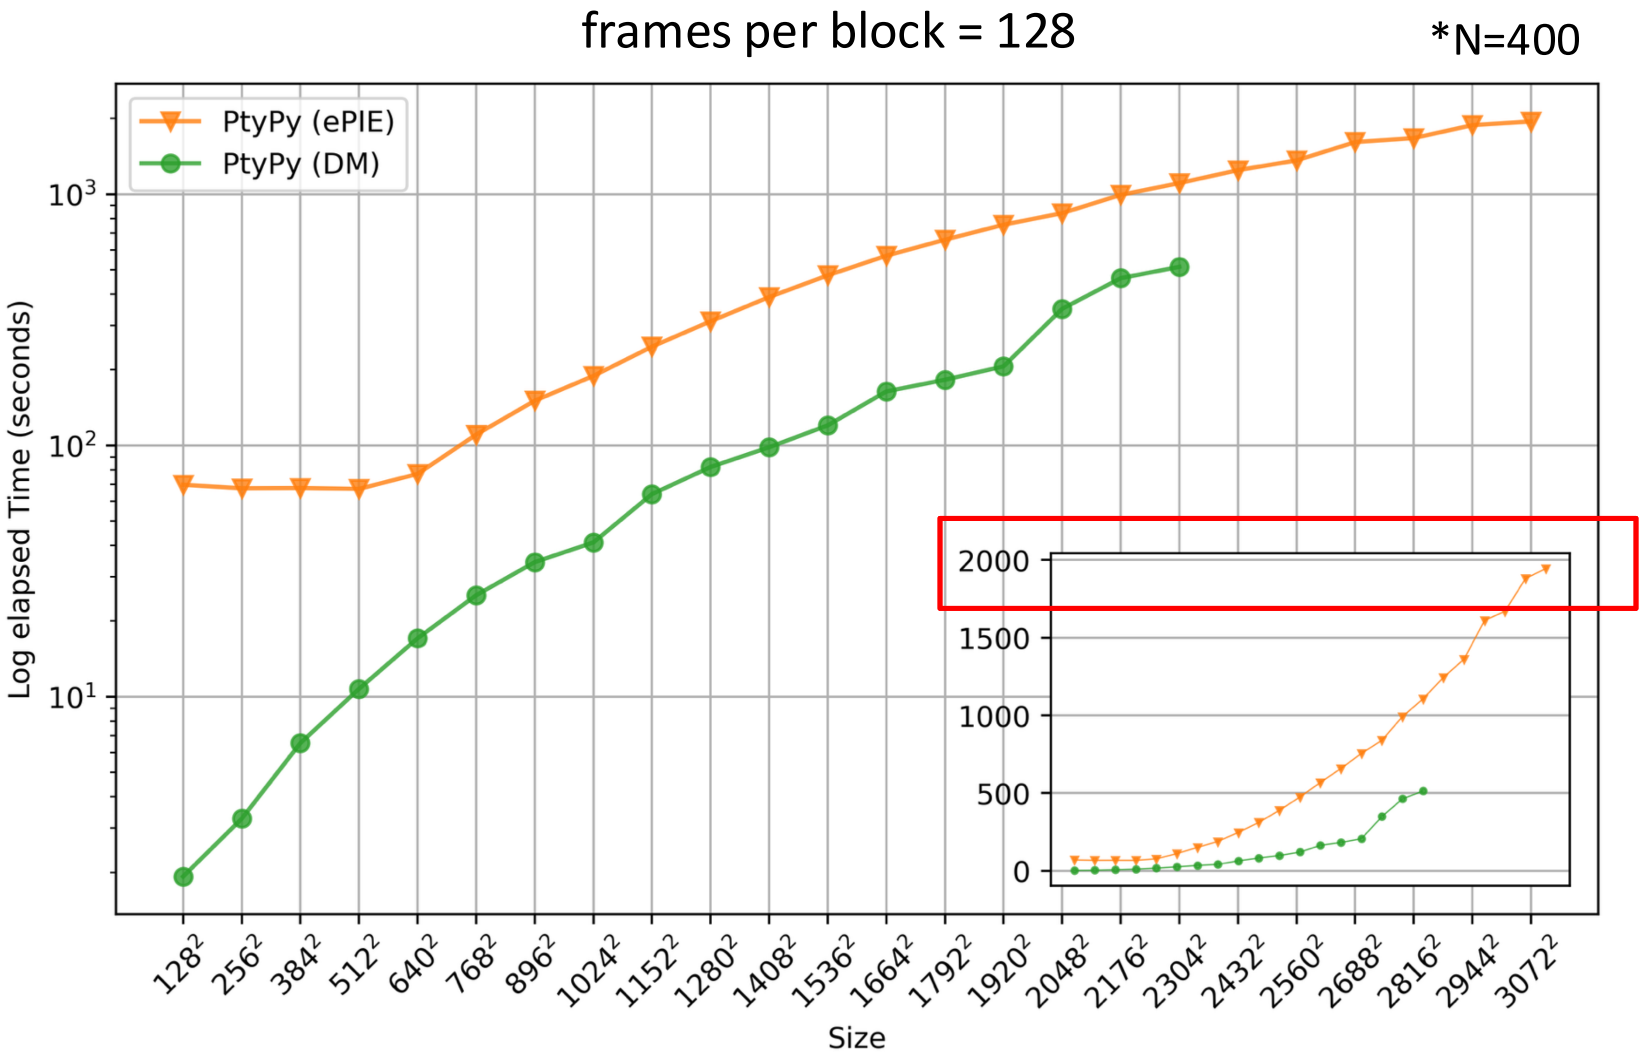

30 minutes/projection!

multi GPU

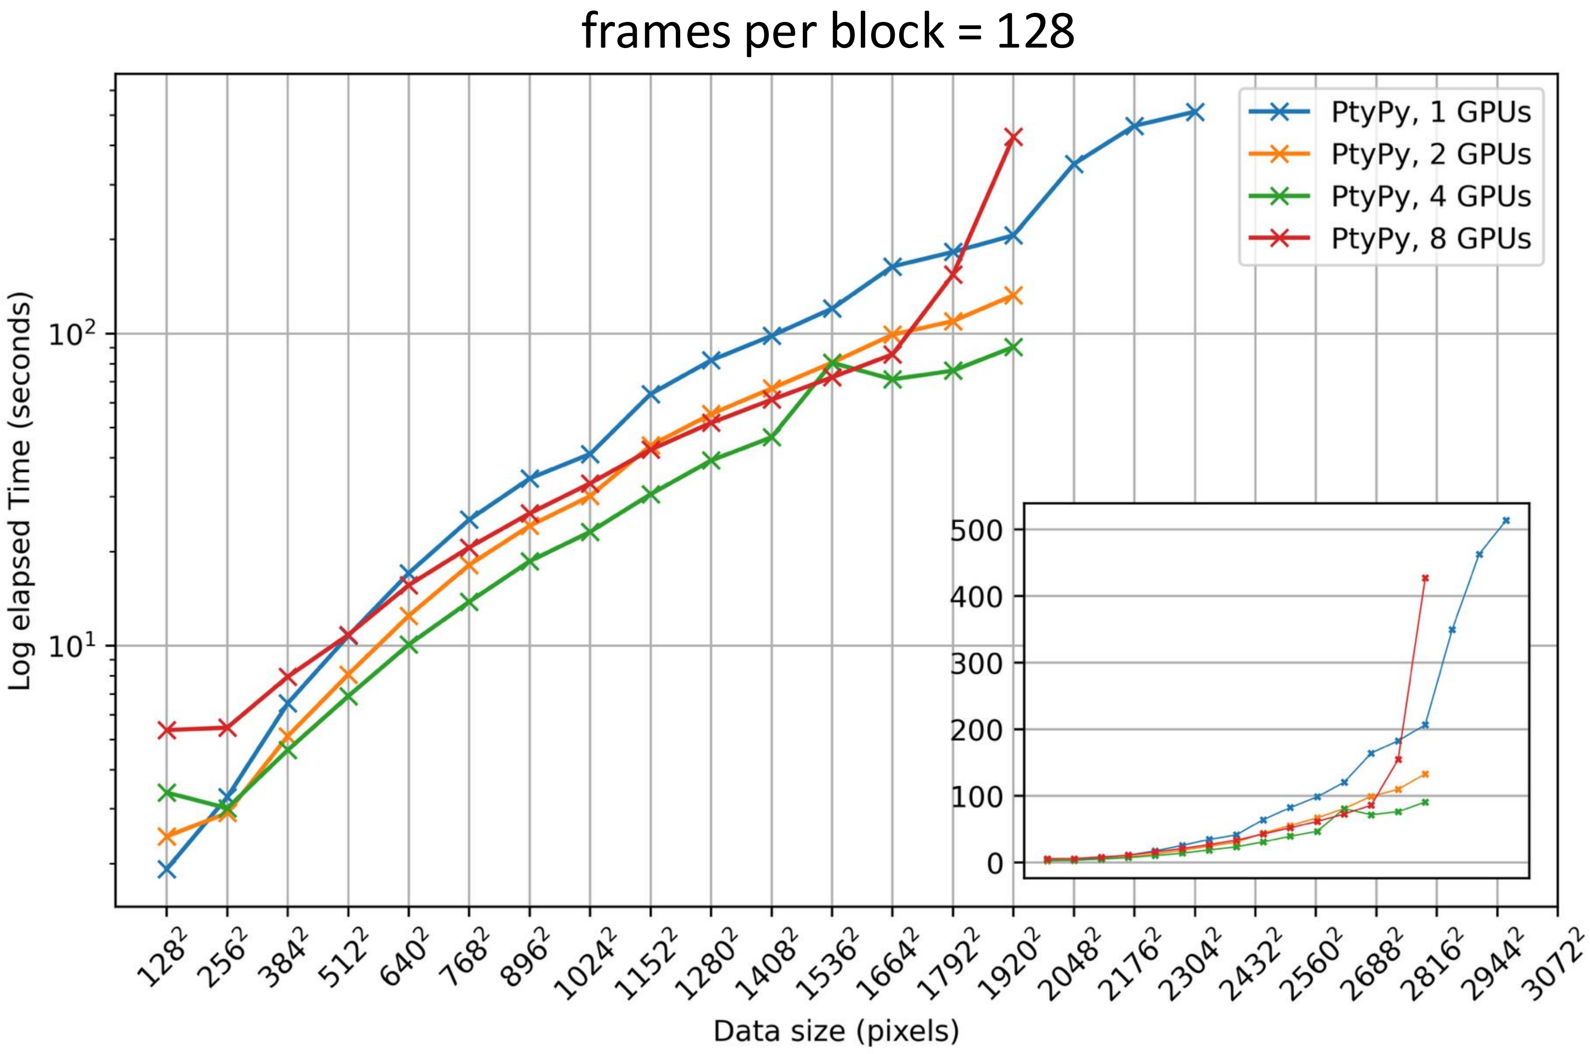

frames per block = 128
PtyPy, 1 GPUs
PtyPy, 2 GPUs
PtyPy, 4 GPUs
PtyPy, 8 GPUs
Log elapsed Time (seconds)
10^2
10^1
Data size (pixels)
128^2
256^2
384^2
512^2
640^2
768^2
896^2
1024^2
1152^2
1280^2
1408^2
1536^2
1664^2
1792^2
1920^2
2048^2
2176^2
2304^2
2432^2
2560^2
2688^2
2816^2
2944^2
3072^2
500
400
300
200
100
0

multi GPU

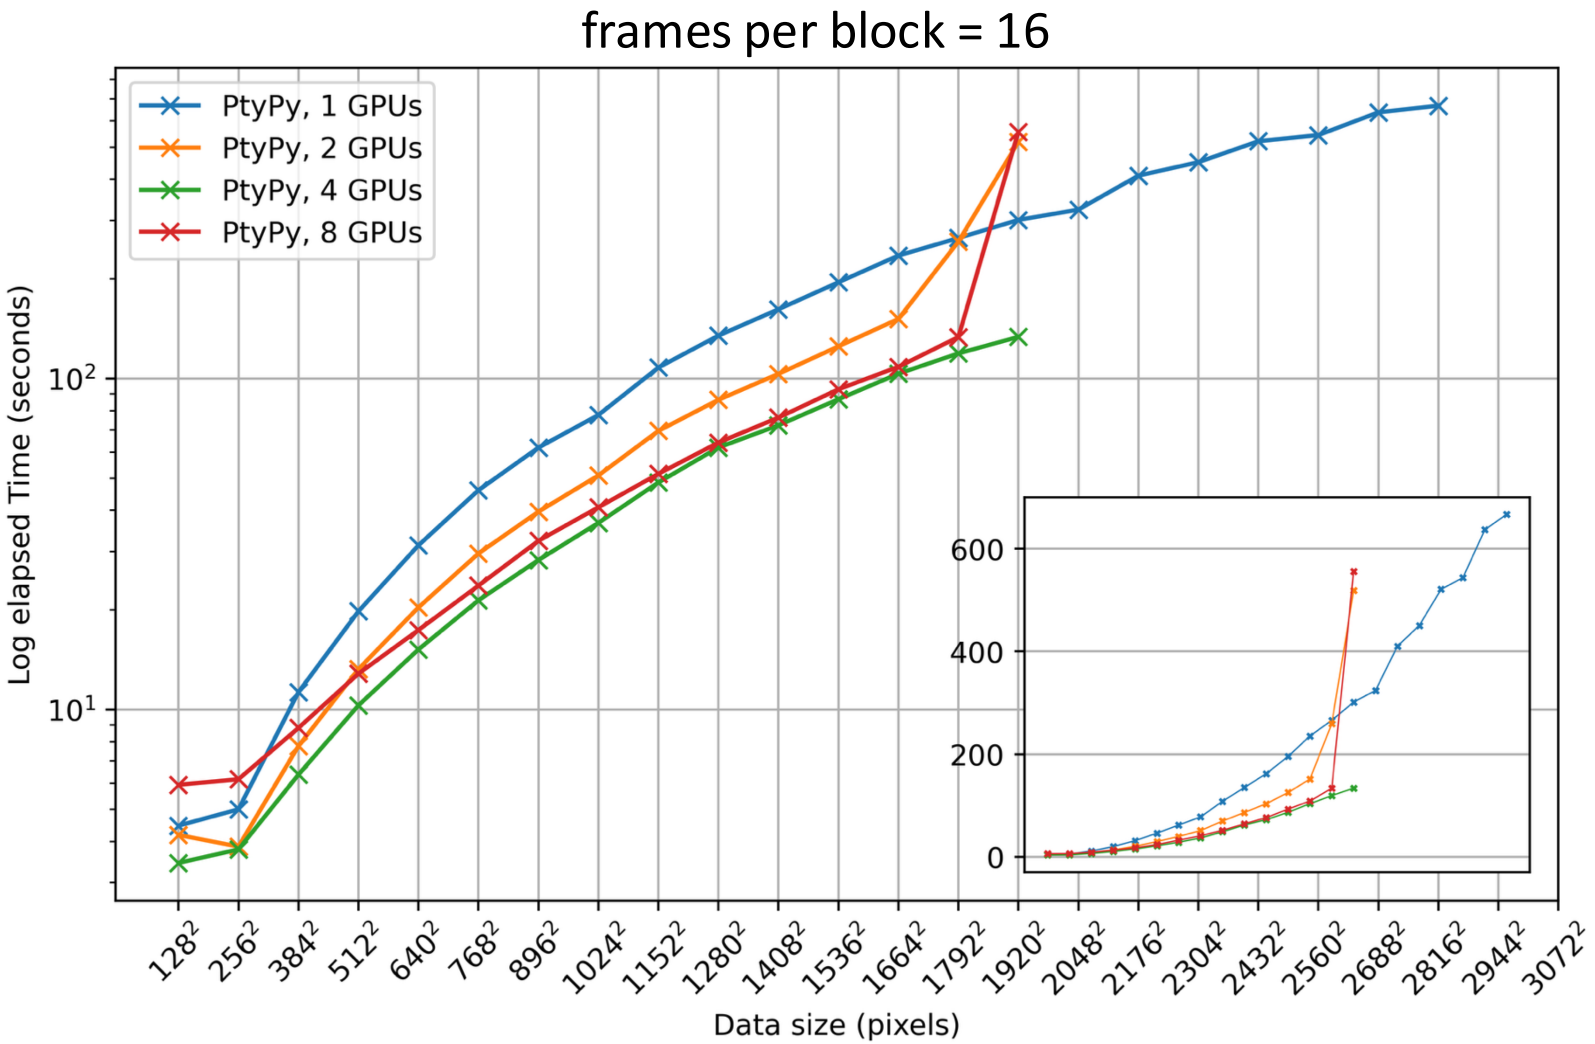

frames per block = 16
PtyPy, 1 GPUs
PtyPy, 2 GPUs
PtyPy, 4 GPUs
PtyPy, 8 GPUs
Log elapsed Time (seconds)
10^2
10^1
128^2 256^2 384^2 512^2 640^2 768^2 896^2 1024^2 1152^2 1280^2 1408^2 1536^2 1664^2 1792^2 1920^2 2048^2 2176^2 2304^2 2432^2 2560^2 2688^2 2816^2 2944^2 3072^2
Data size (pixels)
600
400
200
0

Live Iterative Ptychography

Dieter Weber[1,*], Simeon Ehrig[2,3], Andreas Schropp[4,5], Alexander Clausen[1], Silvio Achilles[4], Nico Hoffmann[2], Michael Bussmann[2,3], Rafal E. Dunin-Borkowski[1], and Christian G. Schroer[4,6]

LIVE ITERATIVE PTYCHOGRAPHY WITH PROJECTION-BASED ALGORITHMS

Simon Welker[1,2], Tal Peer[1], Henry N. Chapman[2,3,4], Timo Gerkmann[1]

[1] Si... Processing (SP), Universität Hamburg, Hamburg, Germany
...CFEL, Deutsches Elektronen-Synchrotron DESY, Notkestr. 85, 22607 Hamburg, Germany
...mburg Centre for Ultrafast Imaging, Hamburg, Germany
...ent of Physics, Universität Hamburg, Hamburg, Germany

Topology-aware Optimizations for Multi-GPU Ptychographic Image Reconstruction

Xiaodong Yu
Data Science and Learning Division
Argonne National Laboratory
Lemont, IL, USA
xyu@anl.gov

Tekin Biçer
Data Science and Learning Division
X-ray Science Division
Argonne National Laboratory
Lemont, IL, USA
tbicer@anl.gov

Rajkumar Kettimuthu
Data Science and Learning Division
Argonne National Laboratory
Lemont, IL, USA
kettimut@mcs.anl.gov

Ian T. Foster
Data Science and Learning Division
Argonne National Laboratory
Lemont, IL, USA
foster@anl.gov

`ssc-cdi`: A Memory-Efficient, Multi-GPU Package for Ptychography with Extreme Data

Yuri Rossi Tonin *, Alan Zanoni Peixinho, Mauro Luiz Brandao-Junior, Paola Ferraz and Eduardo Xavier Miqueles *

How should I approach m

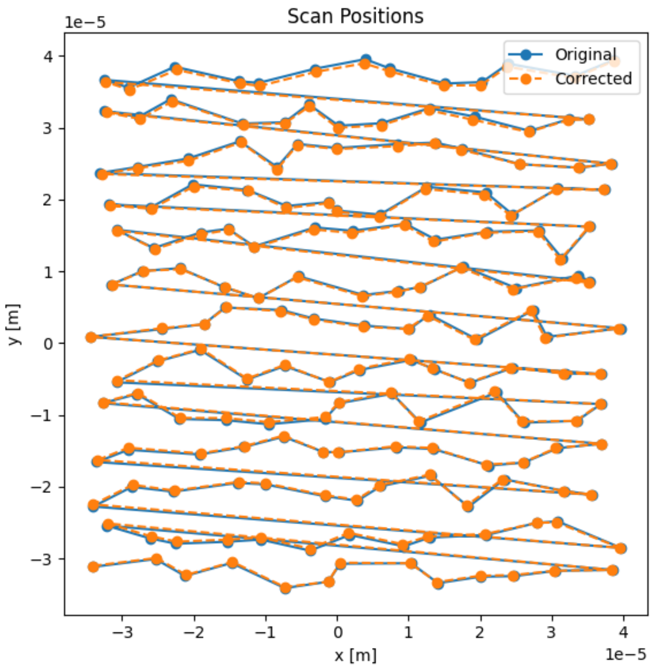

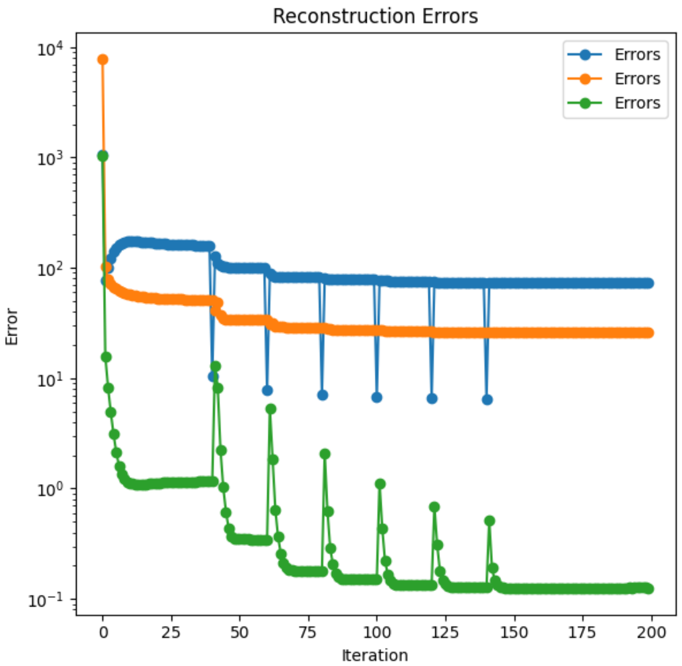

- If planning a large scan or a large volume, ma save you resources ☺
- Processing in GPUs ~~is necessary~~ helps a lot!
- Crop your diffraction data as much as possibl
- If farfield: how oversampled am I? Can I bin t
- Do you need position co ers, it might be helpful to save corrected positions
- Do you need modes? ~~If~~
- Combine algorithms. Tal es. Some iterations may not be helping much.
- Play a bit with the frame as possible before starting the full 3D reconstruction.

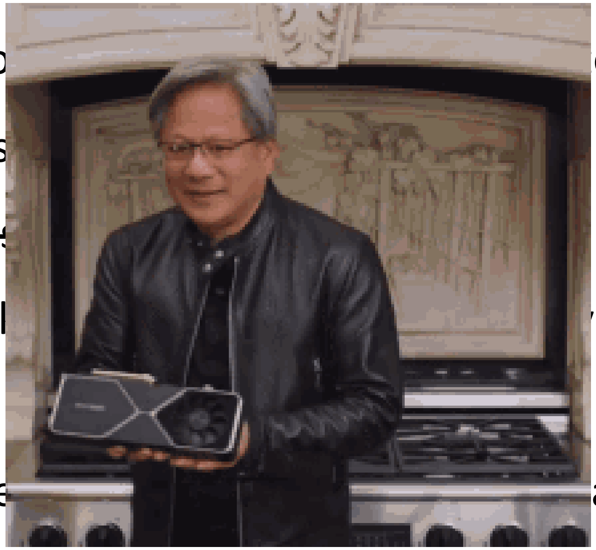

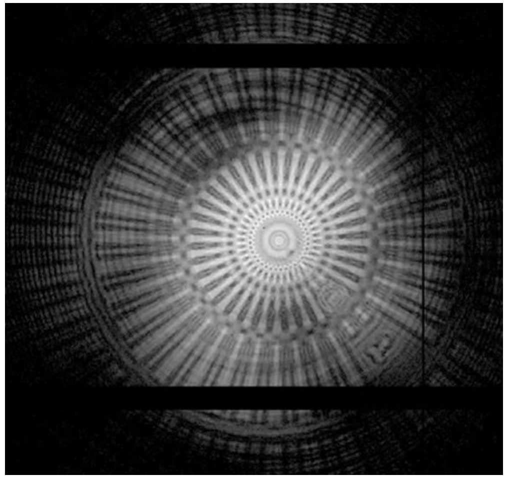

What comes next?

- Data will continue to come faster (better detectors, better controls, multi-beam,…)
- Will the hardware evolve as we'd like? Or do we need multi-node? (Which will be cheaper?)
- Is Machine-learning based reconstruction the future?

Ptychographic phase retrieval via a deep-learning-assisted iterative algorithm

Koki Yamada,[a*] Natsuki Akaishi,[a] Kohei Yatabe[a]

Article https://doi.org/10.1038/s41467-023-41496-z

Deep learning at the edge enables real-time streaming ptychographic imaging

Article | Open access | Published: 21 December 2023

Physics constrained unsupervised deep learning for rapid, high resolution scanning coherent diffraction reconstruction

023
r 2023
mber 2023

Anakha V. Babu[1,4,5], Tao Zhou[1,5], Saugat Kandel[1], Tekin Bicer[1], Zhengchun Liu[1], William Judge[2], Daniel J. Ching[1], Yi Jiang[1], Sinisa Veseli[1], Steven Henke[1], Ryan Chard[1], Yudong Yao[1], Ekaterina Sirazitdinova[3], Geetika Gupta[3], Martin V. Holt[1], Ian T. Foster[1], Antonino Miceli[1] & Mathew J. Cherukara[1]

Oliver Hoidn, Aashwin Ananda Mishra & Apurva Mehta

Questions?