SurrogateTrain: Drastically Improving Performance of Data Loading for Training Scientific Surrogate Models

Baixi Sun[1], Dr. Xiaodong Yu[2], Dr. Kamil A. Iskra[2], Dr. Dingwen Tao[1]

[1]Indiana University Bloomington,
[2]Argonne National Laboratory

INDIANA UNIVERSITY BLOOMINGTON

Argonne NATIONAL LABORATORY

Background & Motivation

Ptychography Imaging

APS-U
High-Energy X-Ray Beams Experiments

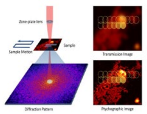

Data Scale
Hundreds of **Peta Bytes** brain initiatives

Surrogate Models

PtychoNN
Predicts the 2D amplitude and phase images.

AutoPhaseNN
Predicts the 3D amplitude and phase patterns.

BraggNN
Predicts the positions of diffraction peaks.

- Ptychography imaging techniques aim to **increase the resolution** of images beyond x-ray optics.
- Traditional iterative algorithms are **computationally expensive.**
- Advanced Photon Source Upgrade (APS-U) will provide **immense data**.
- **Surrogate models** are designed to achieve the task more efficiently.

From Single Device to Distributed Training

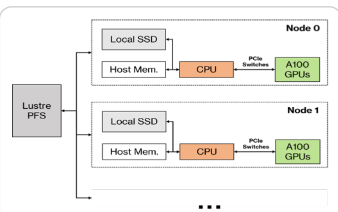

The Architecture of DGX-A100 Nodes

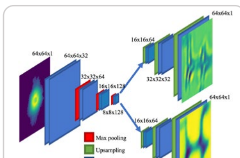

PtychoNN

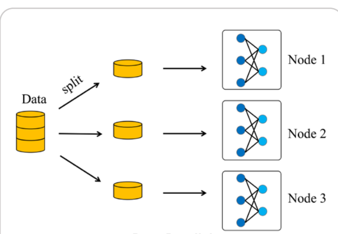

Data-Parallel

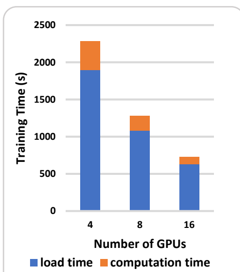

Training Time Breakdown

- To train surrogates on large datasets:
  - Utilize **supercomputers** like ThetaGPU.
  - Utilize **data parallelism** in distributed training.
- Data loading takes **>80%** of training time!

Proposed Design

Part I: Offline Scheduling

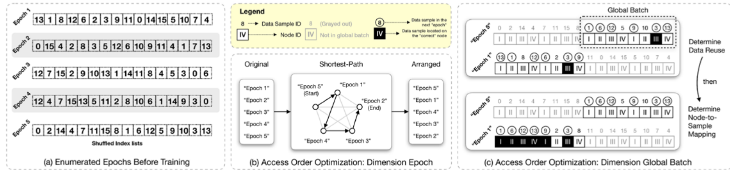

- (a & b) Epoch order optimization.
- (c) Data locality optimization + Load balancing scheduling

Part II: Runtime Buffering

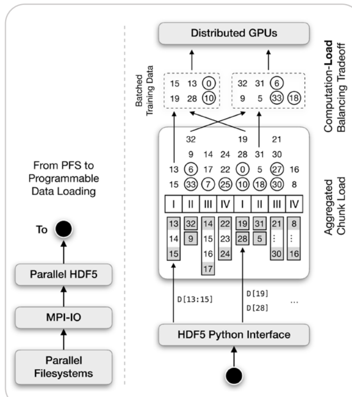

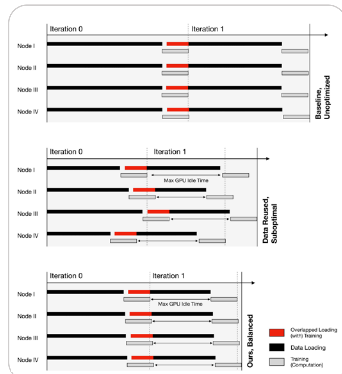

- Determine buffered data eviction.
- Perform **balanced loading**.
- Aggregated chunked loading

1. **Idle time** while waiting for data loading.
2. **Load imbalance** incurred.
3. **Overall time** reduced from balancing the load.

Reference

[1] Mark A. Iwen, Aditya Viswanathan, and Yang Wang. Fast Phase Retrieval from Local Correlation Measurements. SIAM Journal on Imaging Sciences, 9(4):1655–1688, 2016.
[2] Mathew J Cherukara, Tao Zhou, Youssef Nashed, Pablo Enfedaque, Alex Hexemer, Ross J Harder, and Martin V Holt. 2020. AI-enabled high-resolution scanning coherent diffraction imaging. Applied Physics Letters 117, 4 (2020), 044103.
[3] Yao, Y., Chan, H., Sankaranarayanan, S. et al. AutoPhaseNN: unsupervised physics-aware deep learning of 3D nanoscale Bragg coherent diffraction imaging. npj Comput Mater 8, 124 (2022). https://doi.org/10.1038/s41524-022-00803-w
[4] Zhengchun Liu, Hemant Sharma, Jun-Sang Park, Peter Kenesei, Antonino Miceli, Jonathan Almer, Rajkumar Kettimuthu, and Ian Foster. 2022. BraggNN: fast X-ray Bragg peak analysis using deep learning. IUCrJ 9, 1 (Jan 2022), 104–113. https://doi.org/10.1107/S2052252521011258
[5] Shuo Ouyang et al. "Communication optimization strategies for distributed deep neural network training: A survey". In: *Journal of Parallel and Distributed Computing* 149 (2021), pp. 52–65.
[6] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. Advances in neural information processing systems 32 (2019).
[7] Chih-Chieh Yang and Guojing Cong. 2019. Accelerating data loading in deep neural network training. In 2019 IEEE 26th International Conference on High Performance Computing, Data, and Analytics (HiPC). IEEE, 235–245.
[8] Yue Zhu, Fahim Chowdhury, Huansong Fu, Adam Moody, Kathryn Mohror, Kento Sato, and Weikuan Yu. 2018. Entropy-aware I/O pipelining for large-scale deep learning on HPC systems. In 2018 IEEE 26th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS). IEEE, 145–156.12
[9] Nikoli Dryden, Roman Böhringer, Tal Ben-Nun, and Torsten Hoefler. 2021. Clairvoyant prefetching for distributed machine learning I/O. In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis. 1–15.

Evaluation

Setup

- **Environment**: ThetaGPU supercomputer
- **Baseline**: PtychoNN using Pytorch DataLoader.
- **Dataset**: In-house dataset from ANL APS. 262,896 images

Number of Data Loaded From PFS

V1	rank 0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	avg.	stddev	median	max
iteration 1	67	107	69	54	57	58	41	69	75	48	68	74	54	79	58	55	64.6	15.4	62.5	107
2	56	75	82	69	96	63	75	62	69	88	74	57	60	87	59	81	72.1	12.3	71.5	96
3	67	56	71	71	85	105	65	72	79	71	60	81	84	53	80	87	74.2	13.1	71.5	105
4	110	70	91	92	68	77	61	59	64	79	79	62	79	85	89	68	77.1	13.9	78.0	110
5	84	89	94	80	73	83	75	70	86	100	110	81	74	73	79	77	83.0	10.8	80.5	110
																total:	74.2	14.2	74.0	110

V2	rank 0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	avg.	stddev	median	max
iteration 1	72	51	55	67	59	62	52	70	77	53	59	75	54	66	62	66	64.6	8.3	62.0	77
2	59	81	83	73	65	63	61	64	66	77	55	64	73	55	74	58	72.1	8.9	64.5	83
3	66	61	71	80	66	68	63	68	68	73	79	63	83	62	73	72	74.2	6.6	68.0	83
4	72	71	73	64	75	75	65	72	62	81	69	64	81	80	74	76	77.1	6.1	72.5	81
5	67	64	75	93	72	83	69	71	63	82	78	81	76	64	78	71	83.0	8.2	73.5	93
																total:	69.1	8.6	69.0	93

- The baseline method loads 512 images on each rank in each step.
- Compared to Pytorch data loading, we reduced data loading by **6.7x**.

	1	2	3	4	☺
Full Randomization	•	•		•	•
Epoch Order Optimization					•
Node-Sample Mapping (Locality)		○	○	○	•
Load Balancing		○	•	•	•
Aggregated Chunk Load					•

• technique in use
○ data loading achieved via inter-node communication

1 PyTorch DataLoader [6]
2 Locality Aware [7]
3 DeepIO [8]
4 NoPFS [9]
☺ SurrogateTrain (Ours)

Performance on Each Optimization Step

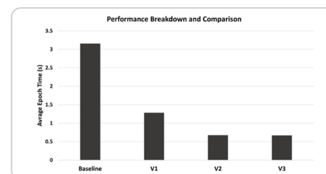

(V1) Data access order optimization,
(V2) V1 + load balancing optimization,
(V3) V2 + chunked loading optimization.

Optimization	Speedup
V1	2.45x
V2	4.60x
V3	4.70x

Acknowledgment

This research was supported by the U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research (ASCR), under contract No. DE-AC02-06CH11357 and by the National Science Foundation through Award OAC-2003709 and. We gratefully acknowledge the computing resources provided on Theta/ThetaGPU supercomputer operated by the Argonne Leadership Computing Facility.