ESAPCA: Enabling the Analysis of Extremely Large Data Sets by Scalable and Hardware-Accelerated PCA and DMD

Visit us on GitHub

Fabian Hoppe, Philipp Knechtges, and Alexander Rüttgers

German Aerospace Center (DLR) – Institute of Software Technology – department High-Performance Computing – Cologne (Germany)

+ various colleagues from DLR, Research Center Jülich, and Karlsruhe Institute of Technology, and open-source contributors (for Heat-development in general)

The project ESAPCA

- Funded by the European Space Agency (ESA) as so-called Early Technology Development Project over 18 months (2024-2025)
- **Goal:** develop a parallel, GPU-accelerated implementation of singular value decomposition (SVD) and related data science techniques, …
  - … capable of running on high-performance computing (HPC) systems, and …
  - … with a focus on easy usability and interoperability within the Python (NumPy/SciPy/scikit-learn) data science ecosystem
- Possible **applications** at ESA include, e.g.:
  - Data-driven modelling and prediction of thermospheric density
  - analysis of in situ measurements of powder bed solidification
- Implementation will build on the existing infrastructure for multi-node array computing within the open-source **Heat** research software library and will become part of this library

ESAPCA

Truncated SVD
Hierarchical SVD,
Randomized SVD,
Incremental SVD

„Full" SVD
SVD via Zolotarev-Polar Decomposition [1],
Tall-skinny SVD

High-level interfaces for PCA, DMD, and DMDc

Integration into Heat

Parallel I/O
Linear Algebra
Array manipulation
Statistics
…

PCA: Principal Component Analysis – **DMD:** Dynamic Mode Decomposition – **DMDc:** DMD with control

Background: The library Heat [2]

- Open-source Python library for massively-parallel array computing and machine learning on CPU/GPU-clusters
- **Vision:** *make array computing and machine learning/data analytics as simple on a supercomputer as it is on a workstation*
- Developed by DLR, Research Center Jülich, and Karlsruhe Institute of Technology (KIT) since 2018

Multi-node- and GPU-capabilities

Operations can be performed in a multi-node-/multi-GPU-setting (e.g., on several nodes of a GPU-cluster)

Simple API and usage

The simple API mimics NumPy/SciPy/scikit-learn and allows for rapid prototyping or adaptation of existing workflows also by HPC-non-experts

Platform independence / Interoperability

As Heat is mainly based on PyTorch and MPI (via mpi4py) under the hood, it is interoperable, portable, and supports hardware of different vendors (e.g., GPUs by Nvidia and AMD).

Scientific background

Heat (primarily) targets usage by scientists and offers the opportunity of collaboration/joint publications with users.

- Basic data type: **`DNDarray`**, a distributed-memory- and GPU-capable counterpart of NumPy's **`ndarray`**
- Array creation, manipulation, and analysis routines, linear algebra operations, and classical machine learning algorithms, adapted to the parallel setting, of course.

Behind the scenes:

- SPMD programming model, hybrid-parallel and bulk-synchronous
- automatic distribution of data across available processing resources by slab decomposition (split along one axis of the array)
- eager execution

➔ Results in [2,3]: possible advantages w.r.t. RAM and time consumption compared to the task-based parallelization of Dask.

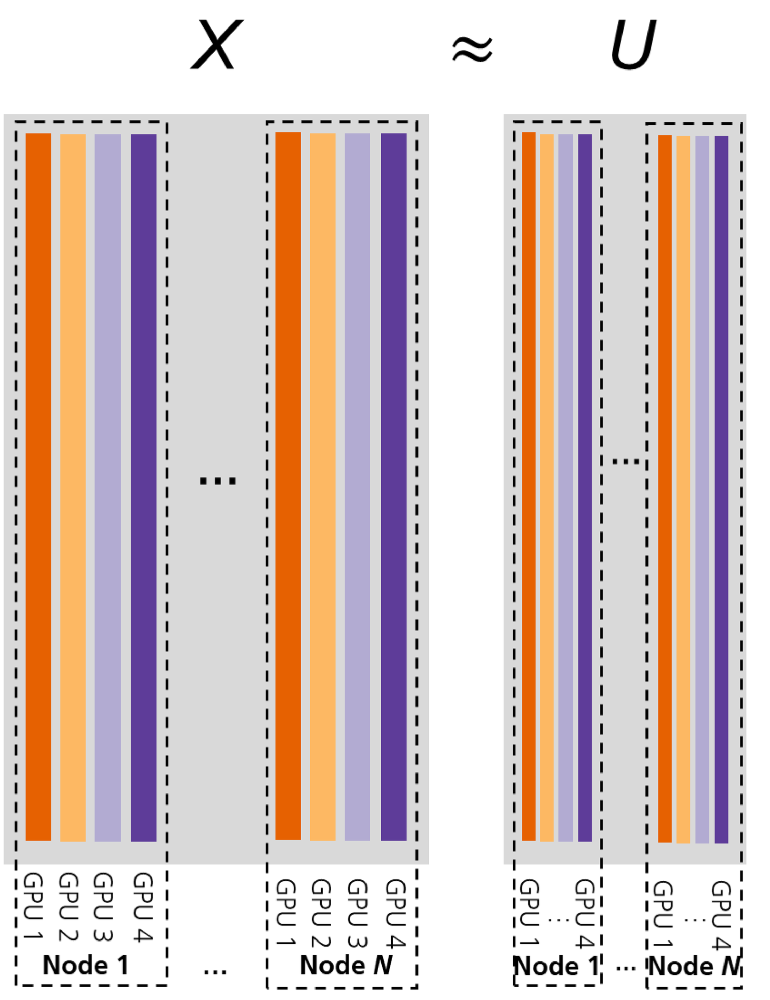

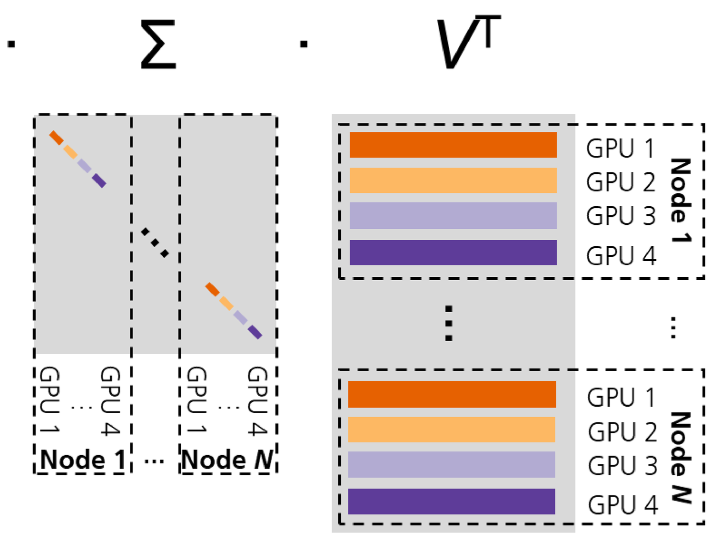


```python
import heat as ht

# load data from file
data = ht.load_hdf5('my_data.h5','data_name',
                    split=0, device='gpu')

# compute truncated PCA
PCA = ht.decomposition.TruncatedSVD(n_components=10,
                                    algorithm='hsvd')
data_reduced = PCA.fit_transform(data)
```

Illustration of the planned functionality: a simple scikit-learn-like API allows to compute the SVD of a huge matrix using several nodes of a GPU-cluster.

Preliminary results

Preliminary experiments with tall-skinny matrices show promising results, e.g., for matrix-matrix multiplication (as, e.g., required for the computation of the Gramian matrix) and SVD both on Nvidia A100 and AMD MI250 GPUs:

Weak scaling of runtime, memory, and energy consumption for matrix-matrix multiplication and SVD of tall-skinny matrices in Heat on up to 16 GPUs

Experiments

- matrix-matrix product A^TB of two tall-skinny matrices A and B, SVD of a tall-skinny matrix
- fraction of data¶ = 1 corresponds to matrices of shape 1959552 x 40824 (matmul, 320GB per matrix) and 1239312 x 25819 (SVD, 128GB), respectively.
- "A100 (1)": weak scaling study with 40GB/8GB (matmul/SVD) of data per A100-GPU
- "A100 (2)"/"MI250": weak scaling study with 20GB/4GB (matmul/SVD) of data per A100- or MI250-GPU
- 1 MPI-process per GPU

Hardware and Measurements

- 4 nodes of DLRs cluster terrabyte with 4 NVIDIA HGX A100 80GB 500W and 2 Intel Xeon Gold 6336Y 24C 185W 2.4GHz per node
- 2 nodes of KITs Future Technology Partition with 4 AMD MI250 and 2 AMD EPYC 7713 2.0GHz per node; since the MI250s are build as Multi Chip Modules (MCM) each MI250 is handled as two GPUs with 64GB RAM each
- Measurements using the hardware counters and the Python package `perun` [4]

Software

- OpenMPI 4.1, PyTorch 2.5.1 with ROCm 6.2.2 (MI250) or CUDA 11.8 (A100), Heat 1.6-dev

¶The reason for scaling by "fraction of data" instead of by number of nodes or number of GPUs is due to the fact that we compare systems with a different number of GPUs per node and GPUs with a different amount of RAM per GPU.

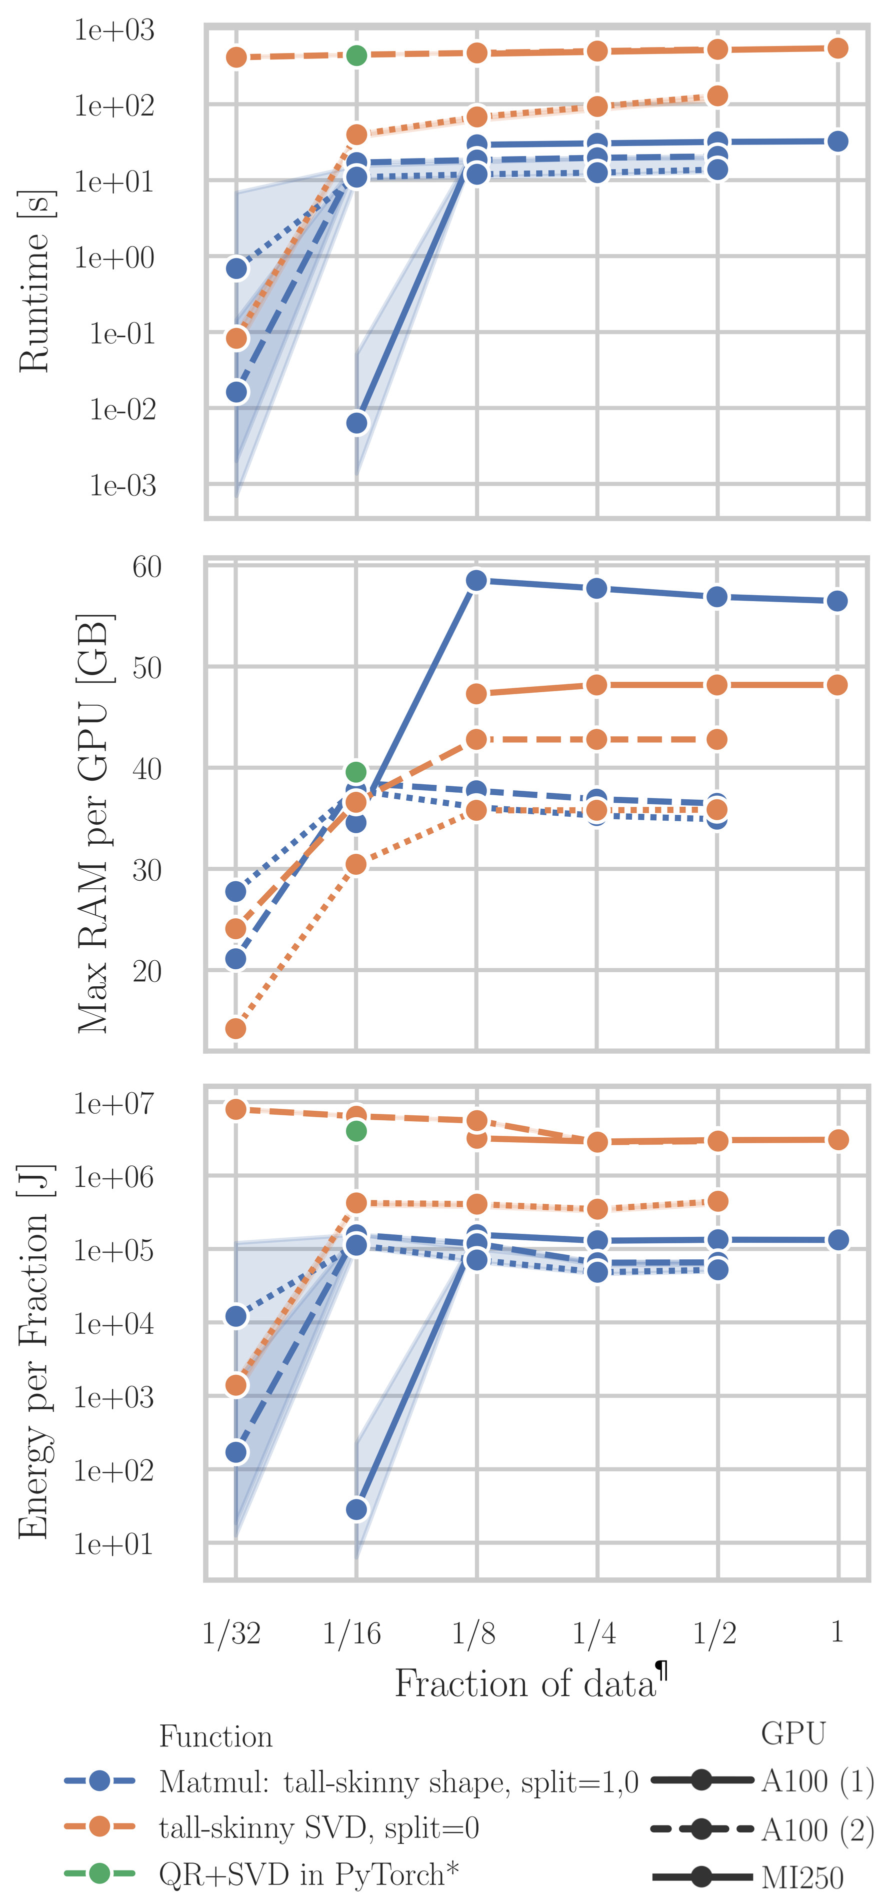

*as `torch.linalg.svd` fails for a 77457x25819-matrix, we compute the QR-decomposition first and then the SVD by computing the SVD of the R-factor.

References

1. Y. Nakatsukasa, R. W. Freund. *Computing fundamental matrix decompositions accurately via the matrix sign function in two iterations: The power of Zolotarev's functions.* SIAM Review 58, 3, 2016.
2. M. Götz et al. *HeAT – a Distributed and GPU-accelerated Tensor Framework for Data Analytics*. In: 2020 IEEE International Conference on Big Data, 2020.
3. F. Hoppe, J.P. Gutiérrez Hermosillo Muriedas, M. Tarnawa, P. Knechtges, B. Hagemeier, K. Krajsek, A. Rüttgers, M. Götz, and C. Comito. *Engineering a large-scale data analytics and array computing library for research: Heat*. Accepted at ECEASST, 2024.
4. J. P. Gutiérrez Hermosillo Muriedas et al. *perun: Benchmarking Energy Consumption of High-Performance Computing Applications.* In: Euro-Par 2023: Parallel Processing, 2023.

This research is supported by the European Space Agency through the Open Space Innovation Platform (https://ideas.esa.int) as a Early Technology Development Agreement and carried out under the Discovery Program ESA Early Technology Development (Research Agreement No. 4000144045/24/NL/GLC/ov). Disclaimer: The view expressed in this publication can in no way be taken to reflect the official opinion of the European Space Agency.

https://activities.esa.int/4000144045

This work was performed on the NHR@KIT Future Technologies Partition testbed funded by the Ministry of Science, Research and the Arts Baden-Württemberg and by the Federal Ministry of Education and Research. Moreover, the authors gratefully acknowledge the computational and data resources provided through the joint high-performance data analytics (HPDA) project "terrabyte" of the German Aerospace Center (DLR) and the Leibniz Supercomputing Center (LRZ).