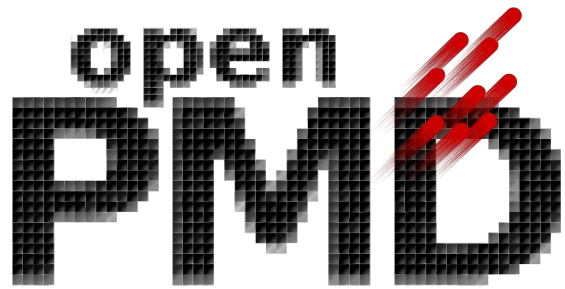


Open Standard for Particle-Mesh Data



HDF5 User Group Meeting

Columbus, OH
August 17th, 2023



CAMPA

Consortium for Advanced Modeling
of Particle Accelerators



U.S. DEPARTMENT OF
ENERGY

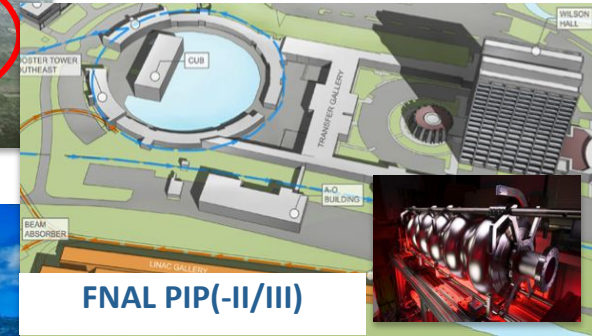
Office of
Science

Axel Huebl

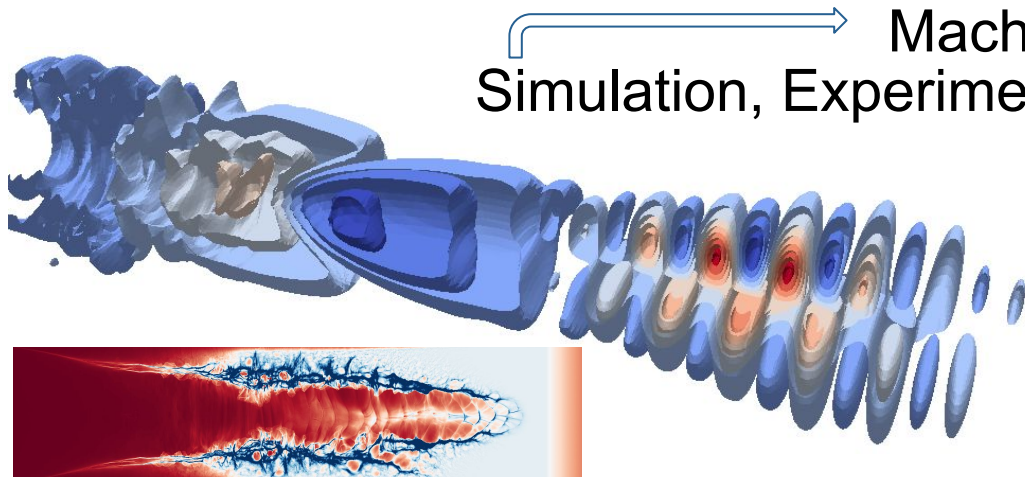
Lawrence Berkeley National Laboratory

On behalf of the openPMD Community incl. content from
Franz Poeschel (CASUS), Lipeng Wan (GSU),
Norbert Podhorszki (ORNL), Junmin Gu (LBNL),
Maxence Thévenet (DESY), Erik Schnetter (PITP), Remi Lehe (LBNL)

My Background: Particle Accelerator & Laser-Plasma Physics



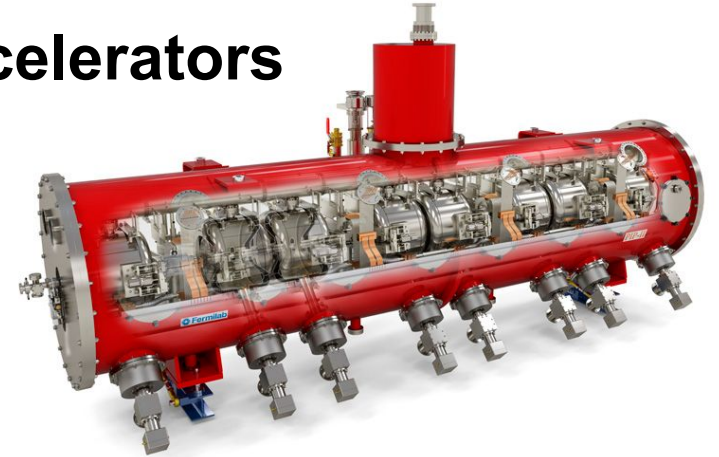
Laser-Plasma Physics



Machine Learning
Simulation, Experiment

Particle Accelerators

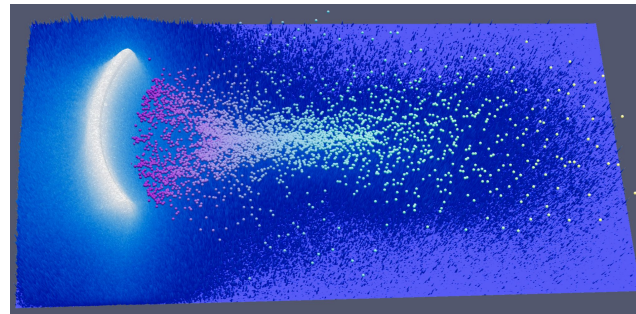
X-Ray Light Sources, FELs, Laser Pulses



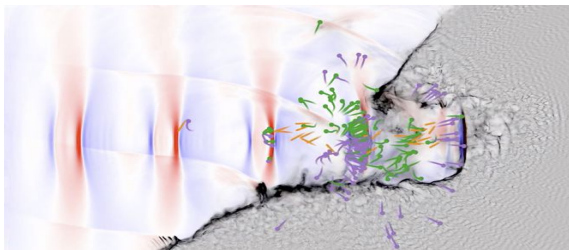
Fusion Energy Science

Astro-physical plasmas

Inertial Confinement Fusion
Laboratory Astrophysics
High-Energy Density Physics



High-Field Physics

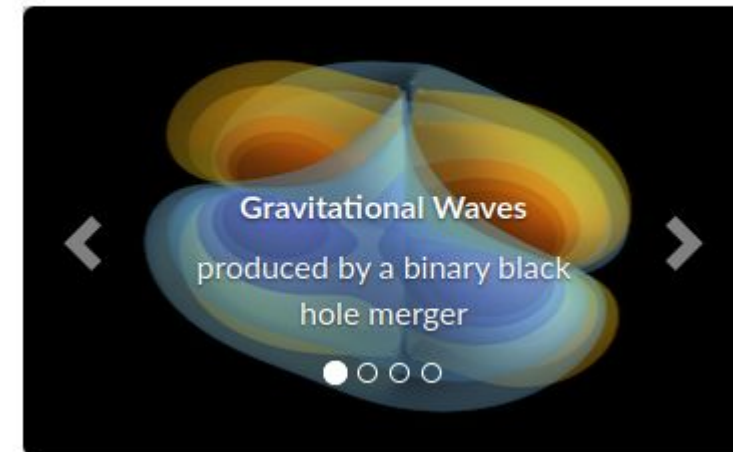


Life Sciences

Particle-Based Image Analysis

General Relativity

The Einstein Toolkit



Outline

- **openPMD Concepts**
 - design principles of our meta-data standard
 - components & modularity
 - example: sparse data for mesh-refinement
- **Open Ecosystem**
 - projects, libraries, tooling, integrations
 - open community
- **Selected R&D Highlights**
 - fast reads: data layouts & in situ statistics
 - from files to data streams

Design Principles

Open Standard for Particle-Mesh Data



- **high-level** description
- **minimal**: users can add more
- **human** readable & **machine** actionable

- file format **agnostic, portable**
- **scalable** from desktop to supercomputer
- versioned, **forward-updatable**:
start “strict“, relax later

h5py reader:
<400 lines
of code

📄	<code>__init__.py</code>
📄	<code>field_reader.py</code>
📄	<code>params_reader.py</code>
📄	<code>particle_reader.py</code>
📄	<code>utilities.py</code>



openPMD - A Self-Describing, FAIR Standard

Findable: Standardized metadata to identify the data producer

```
string    /author          attr    = "franz"  
string    /software       attr    = "PIconGPU"  
string    /softwareVersion attr    = "0.5.0-dev"
```

Accessible: Open standard, implementable in various formats



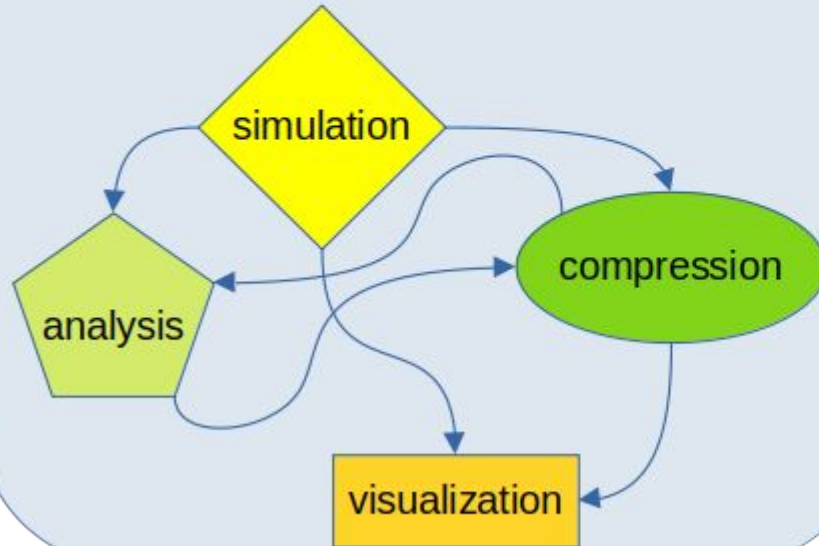
*currently implemented,
but not limited to

"The FAIR Guiding Principles for scientific data management and stewardship" (Mark D. Wilkinson et al.)

openPMD - A Self-Describing, FAIR Standard

Interoperable:

Data exchange spans applications, platforms and teams



Reusable:

Rich and standardized description for physical quantities

Name	Value
axisLabels	[b'z' b'y' b'x']
dataOrder	b'C'
fieldSmoothing	b'none'
geometry	b'cartesian'
gridGlobalOffset	[0. 0. 0.]
gridSpacing	[4.252342 1.0630856 4.252342]
gridUnitSI	4.1671151662e-08
position	[0. 0. 0.]
timeOffset	0.0
unitDimension	[-3. 0. 1. 1. 0. 0. 0.]
unitSI	15399437.98944343

"The FAIR Guiding Principles for scientific data management and stewardship" (Mark D. Wilkinson et al.)

openPMD Components & Modularity



- **markup** / schema for arbitrary hierarchical data formats
- truly, scientifically **self-describing**
- basis for **open data workflows**

openPMD standard (1.0.0, 1.0.1, 1.1.0)

the underlying file markup and definition

A Huebl et al., DOI:10.5281/zenodo.33624

base standard

general description

wavefronts, particle species, particle beams, weighted particles, PIC, MD, mesh-refinement, CCD images, ...

extensions

domain specific



openPMD-viewer

quick visualization

explore, e.g., in Jupyter

openPMD-api

reference library

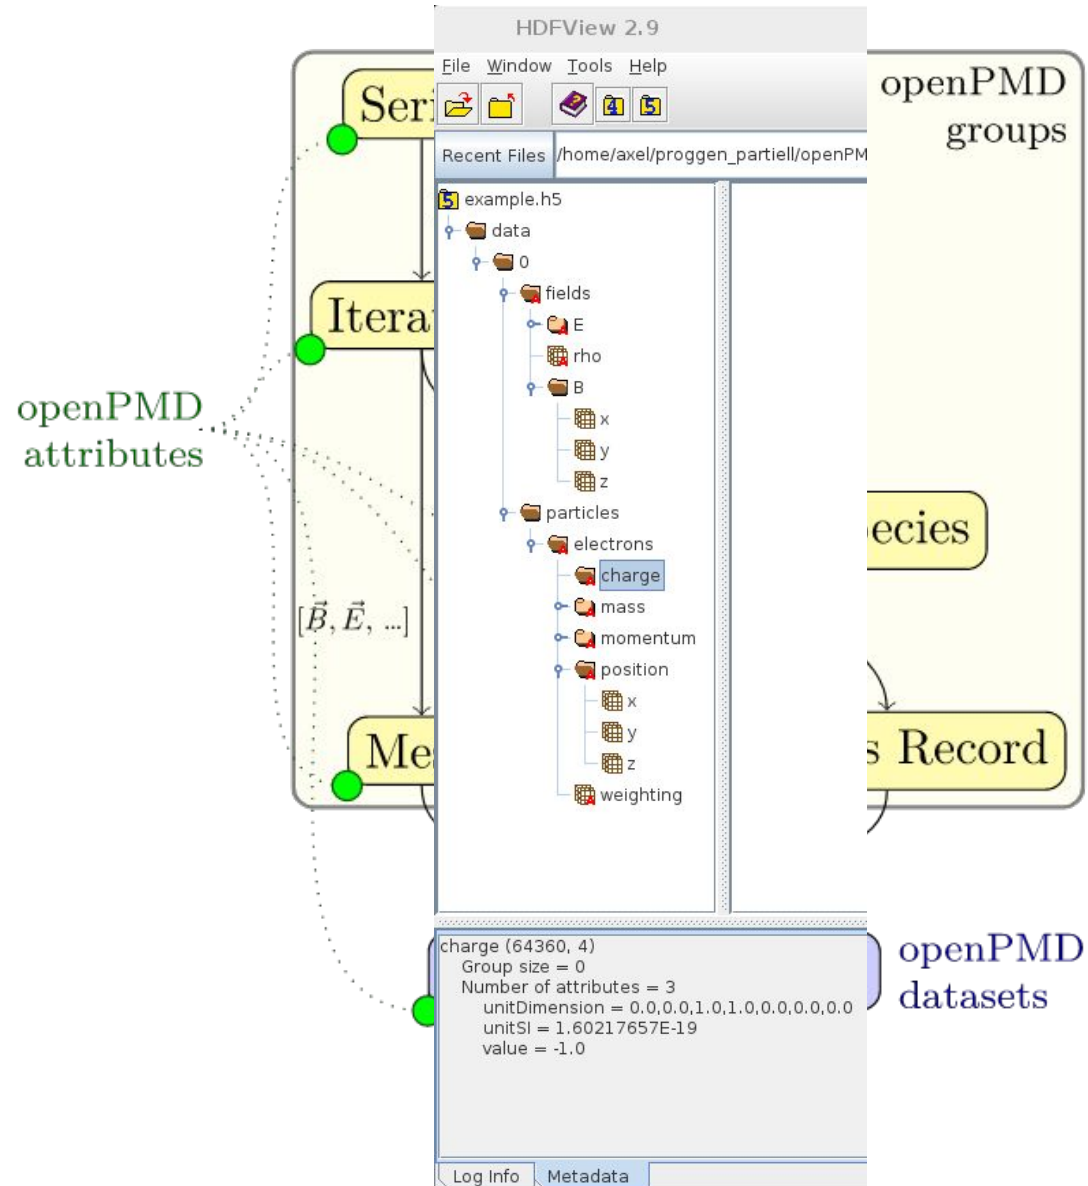
file-format agnostics API

openPMD-updater

auto-update to new standard, verify

openPMD-validator

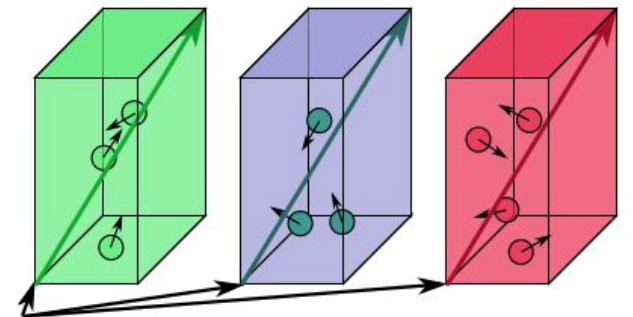
openPMD - Hierarchy for Data Series



- **Structure** for series & snapshots/iterations encoded in either:
 - files** new files per iteration
 - groups** reuse files
 - variables** reuse files & variables*
 - *ADIOS2-only via *steps*
- Records for **physical observables**
constants, mixed precision, complex numbers
- **Attributes:** unit conversion, description, relations, mesh geometry, authors, env. info, ...

- domain-decomposition

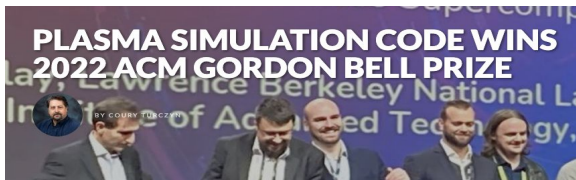
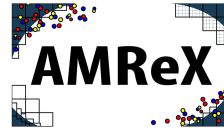
[0:3] particles [3:6] particles [6:10] particles



openPMD: Block-Structured Mesh-Refinement (WarpX)

WarpX Particle-in-Cell Code

implemented on AMReX
structured fields + many particles



WarpX was first PIC code to win prestigious ACM Gordon Bell award, the “Oscars” of Supercomputing.

open source 
ecp-warpX.github.io

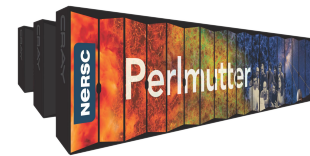
Multi-Node parallelization

- MPI: 3D domain decomposition
- dynamic load balancing



On-Node Parallelization

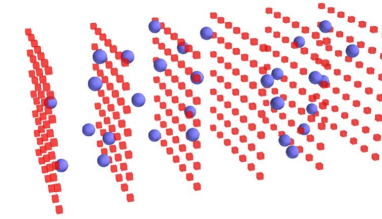
- GPU: CUDA, HIP and SYCL
- CPU: OpenMP



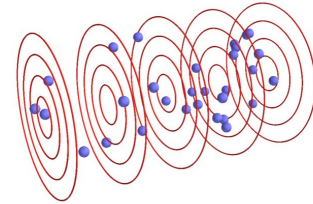
>100TB per output
Multi-PB per sim

Geometries

- 1D3V, 2D3V, 3D3V and RZ (quasi-cylindrical)



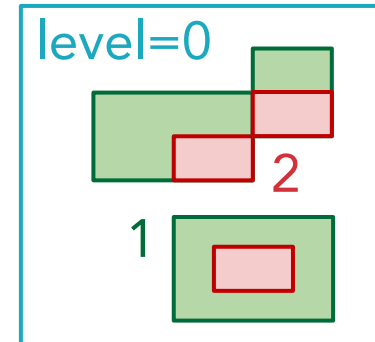
3D Cartesian grid



Cylindrical grid (schematic)

Block-Structured Mesh-Refinement in openPMD

- ADIOS2:
 - variable: 1 per level
 - fill with partial blocks
- HDF5:
 - one variable per patch, N per level
 - overhead...? see: Elena's *sparsity* proposal
- alternative: AMReX HDF5
1 variable per level, but loses HDF5 self-description (concatenated patches)



Outline

- **openPMD Concepts**
 - design principles of our meta-data standard
 - components & modularity
 - example: sparse data for mesh-refinement
- **Open Ecosystem**
 - projects, libraries, tooling, integrations
 - open community
- **Selected R&D Highlights**
 - fast reads: data layouts & in situ statistics
 - from files to data streams

First Steps

<https://github.com/openPMD>



openPMD

Open Standard for Particle-Mesh Data Files

<https://www.openPMD.org> axelhuebl@lbl.gov

[Overview](#) [Repositories 17](#) [Packages](#) [People 55](#) [Teams 5](#) [Projects](#) [Settings](#)

Pinned

Customize your pins

[openPMD-standard](#) Public ⋮
Open Standard for Particle-Mesh Data
☆ 52 20

[openPMD-projects](#) Public ⋮
Overview on Projects around openPMD
☆ 6 6

[openPMD-viewer](#) Public ⋮
Python visualization tools for openPMD files
Jupyter Notebook ☆ 45 36

[openPMD-api](#) Public ⋮
C++ & Python API for Scientific I/O
C++ ☆ 75 36

[openPMD-example-datasets](#) Public ⋮
HDF5 Example Files
Python ☆ 6 2

[openPMD-validator](#) Public ⋮
Validator and Example Scripts
Python ☆ 4 8

People



[View all](#)

[Invite someone](#)

Top languages

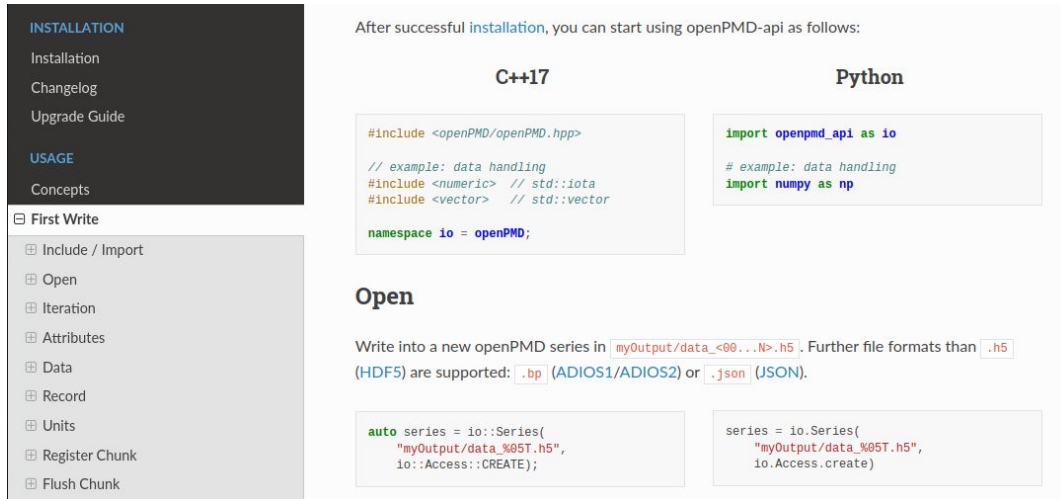
Python C++ Jupyter Notebook
 C CSS

Repositories

Reference Implementation in C++ & Bindings: Python, Julia

Online Documentation:
openpmd-api.readthedocs.io

Open-Source Development & Tests:
github.com/openPMD/openPMD-api



INSTALLATION

- Installation
- Changelog
- Upgrade Guide

USAGE

- Concepts

First Write

- Include / Import
- Open
- Iteration
- Attributes
- Data
- Record
- Units
- Register Chunk
- Flush Chunk

After successful installation, you can start using openPMD-api as follows:

C++17

```
#include <openPMD/openPMD.hpp>

// example: data handling
#include <numeric> // std::iota
#include <vector> // std::vector

namespace io = openPMD;
```

Python

```
import openpmd_api as io

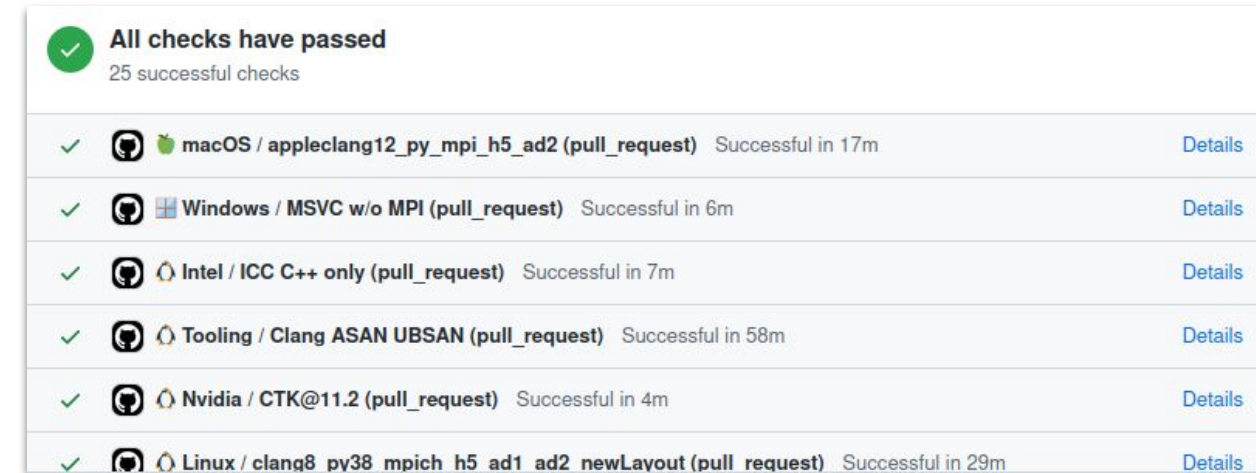
# example: data handling
import numpy as np
```

Open

Write into a new openPMD series in `myOutput/data_<00...N>.h5`. Further file formats than `.h5` (HDF5) are supported: `.bp` (ADIOS1/ADIOS2) or `.json` (JSON).

```
auto series = io::Series(
    "myOutput/data_%05T.h5",
    io::Access::CREATE);

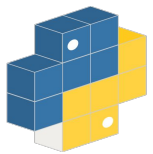
series = io.Series(
    "myOutput/data_%05T.h5",
    io::Access::create)
```



All checks have passed
25 successful checks

- macOS / appleclang12_py_mpi_h5_ad2 (pull_request) Successful in 17m [Details](#)
- Windows / MSVC w/o MPI (pull_request) Successful in 6m [Details](#)
- Intel / ICC C++ only (pull_request) Successful in 7m [Details](#)
- Tooling / Clang ASAN UBSAN (pull_request) Successful in 58m [Details](#)
- Nvidia / CTK@11.2 (pull_request) Successful in 4m [Details](#)
- Linux / clang8_py38_mpich_h5_ad1_ad2_newLayout (pull_request) Successful in 29m [Details](#)

Rapid and easy installation on any platform:



**python3 -m pip install
openpmd-api**



**brew tap openpmd/openpmd
brew install openpmd-api**



**cmake -S . -B build
cmake --build build --target install**



**conda install -c
conda-forge openpmd-api**



**spack install
openpmd-api**



module load openpmd-api

A Huebl, F Poeschel, F Koller, J Gu, et al.

"openPMD-api: C++ & Python API for Scientific I/O with openPMD" (2018) [DOI:10.14278/rodare.27](https://doi.org/10.14278/rodare.27)

openPMD is used in many Scientific Simulations

Simulation Codes

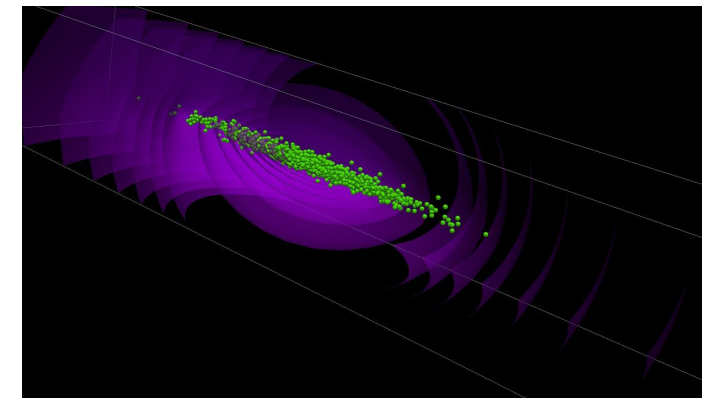
- WarpX
- PIconGPU
- HiPACE++
- ImpactX
- Synergia3
- Wake-T
- Warp
- FBPIC
- SimEx
- LUME
- CarpetX
- Osiris
- VPIC (prototype)
- ACE3P
- ParaTAXIS
- Bmad
- UPIC-Emma
- ...



openPMD/openPMD-projects

*scientific
simulations*

Many use openPMD via



Credit: M. Thévenet, A. Ferran Pousa

VisualPIC

Batteries Included: Bridges into Data Science & Visualization

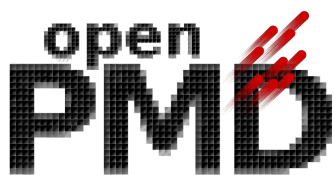
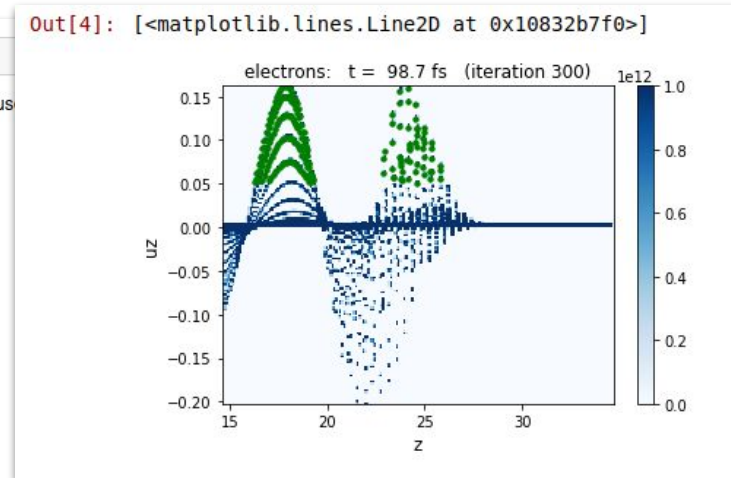
We integrate for scientific productivity

including data analytics frameworks & graphical user interfaces

In []: `ts_2d.slider()`

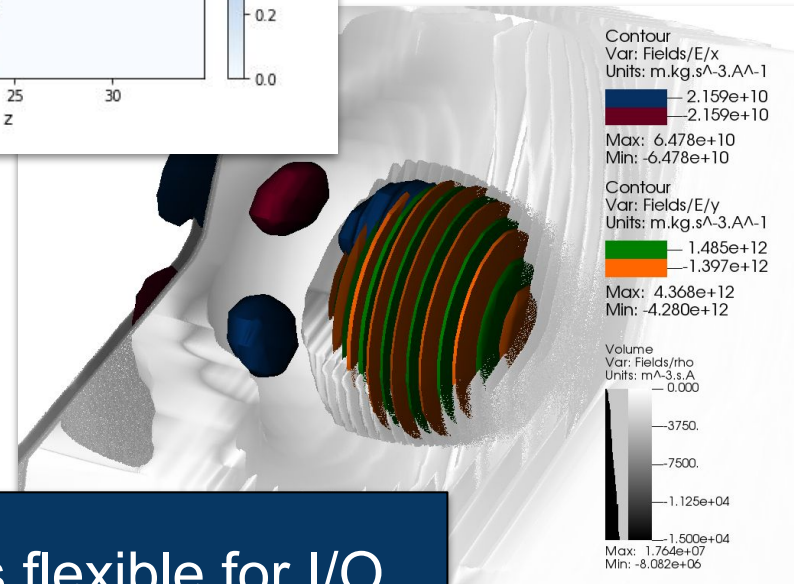
Calling this method will insert the following panel inside the notebook. (Note that the panel below is a **non-interactive image**, which is here for the use of this notebook online. Calling the `slider` method in a live notebook, will produce a truly **interactive** panel.)

The screenshot shows a Jupyter notebook cell with the command `ts_2d.slider()`. Below the code, there is a text explanation and a GUI panel. The GUI panel includes a slider for time t (fs) set to 33. It has sections for 'Field type' with options B, E, J, rho and coordinates x, y, z; 'Particle quantities' for 'Hydrogen1+' with options x, y, z, ux, uy, uz, w; and 'Plotting options' with 'Always refresh' and 'Refresh now!' buttons.



DASK

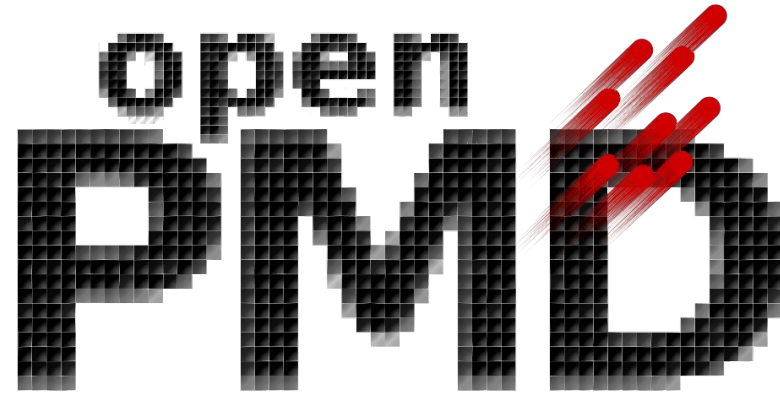
ParaView



RAPIDS



Open standardization, i.e. openPMD, makes us flexible for I/O libraries, tooling & domain-science needs.



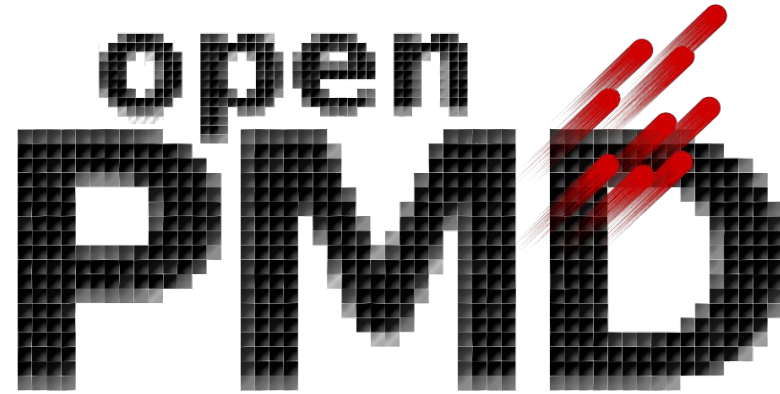
HZDR

HELMHOLTZ
ZENTRUM DRESDEN
ROSSENDORF



UCLA





The **openPMD standard** is co-authored by [Axel Huebl](#), [Rémi Lehe](#), Jean-Luc Vay, David P. Grote, Ivo F. Sbalzarini, Stephan Kuschel, David Sagan, Frédéric Pérez, Fabian Koller, [Franz Poeschel](#), Carsten Fortmann-Grote, Ángel Ferran Pousa, Juncheng E, [Maxence Thévenet](#), and Michael Bussmann.

The authors are thankful for the **community contributions** to libraries, software ecosystem, user support, review and integrations. Particularly, thank you to Yaser Afshar, Lígia Diana Amorim, James Amundson, Weiming An, Igor Andriyash, Ksenia Bastrakova, [Jean Luca Bez](#), Richard Briggs, Heiko Burau, Jong Choi, Ray Donnelly, Dmitry Ganyushin, Marco Garten, Lixin Ge, Berk Geveci, Daniel Grassinger, Alexander Grund, [Junmin Gu](#), Marc W. Guetg, Ulrik Günther, Sören Jalas, Manuel Kirchen, John Kirkham, Scott Klasky, Noah Klemm, Fabian Koller, Mathieu Lobet, Christopher Mayes, Ritiek Malhotra, Paweł Ordyna, Richard Pausch, [Norbert Podhorszki](#), David Pugmire, Felix Schmitt, [Erik Schnetter](#), Dominik Stańczak, Klaus Steiniger, Michael Sippel, Frank Tsung, Lipeng Wan, René Widera, and Erik Zenker!

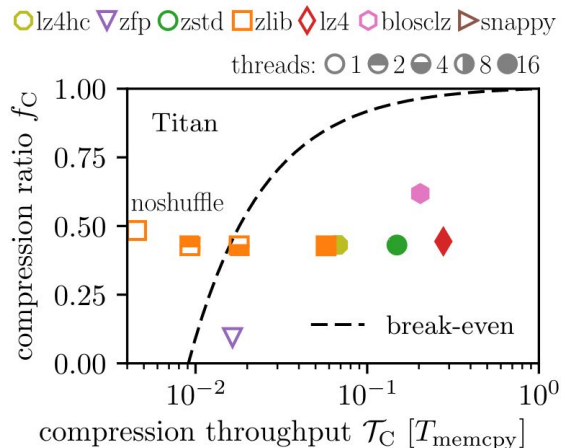
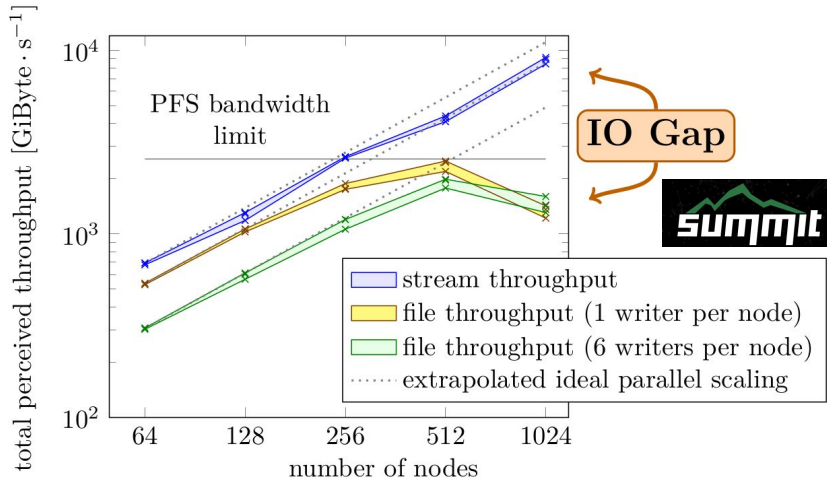
Outline

- **openPMD Concepts**
 - design principles of our meta-data standard
 - components & modularity
 - example: sparse data for mesh-refinement
- **Open Ecosystem**
 - projects, libraries, tooling, integrations
 - open community
- **Selected R&D Highlights**
 - fast reads: data layouts & in situ statistics
 - from files to data streams

Performance, Data Layouts and Fileless I/O

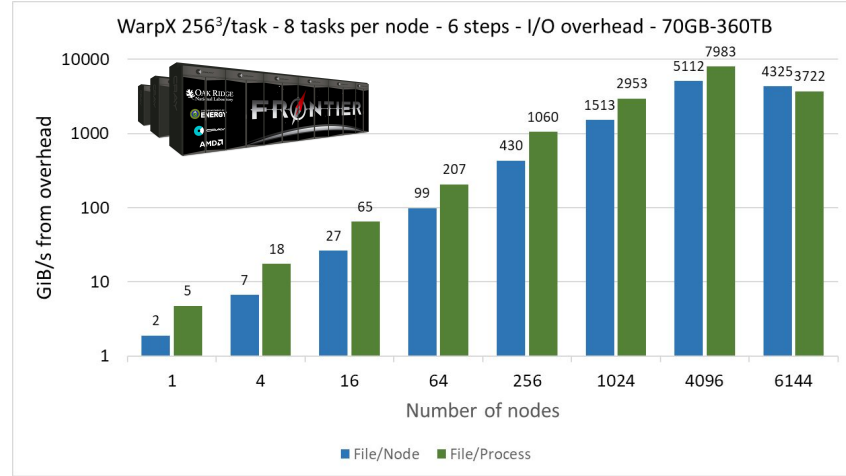
Application Challenges

- R&D in: scalable techniques, data layouts, libraries
- scientific data analysis & sharing



Fast Compressors Needed:
[DOI:10.1007/978-3-319-67630-2_2](https://doi.org/10.1007/978-3-319-67630-2_2)
 by A Huebl et al., ISC DRBSD-1 (2017)

ADIOS openPMD-api w/ WarpX

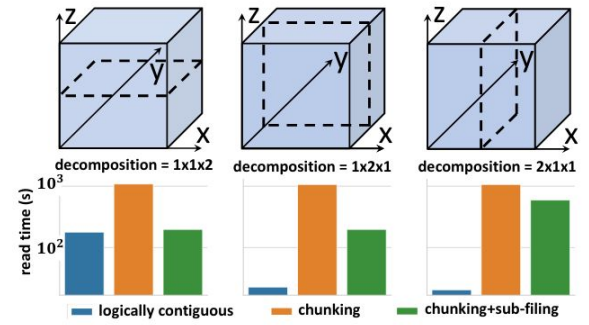
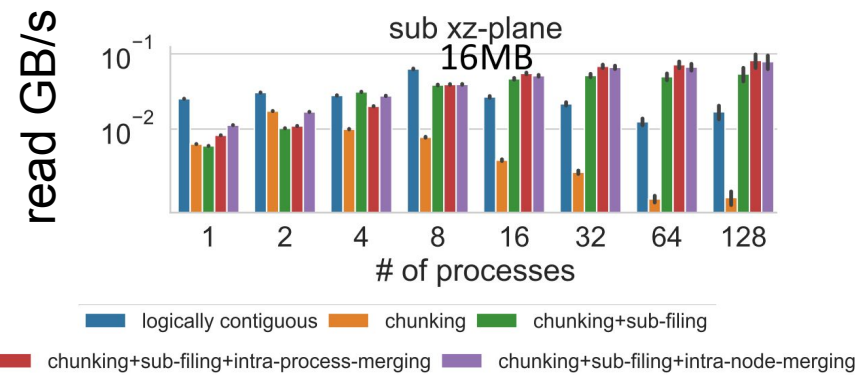


>5.5TB/s FS BW: two-tier lustre w/ high-performance storage & progressive files

Streaming Data Pipelines:

[DOI:10.1007/978-3-030-96498-6_6](https://doi.org/10.1007/978-3-030-96498-6_6)

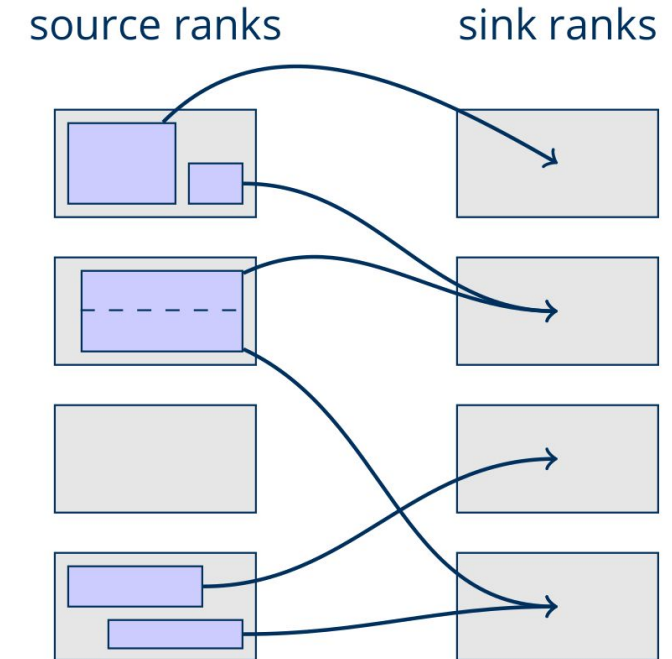
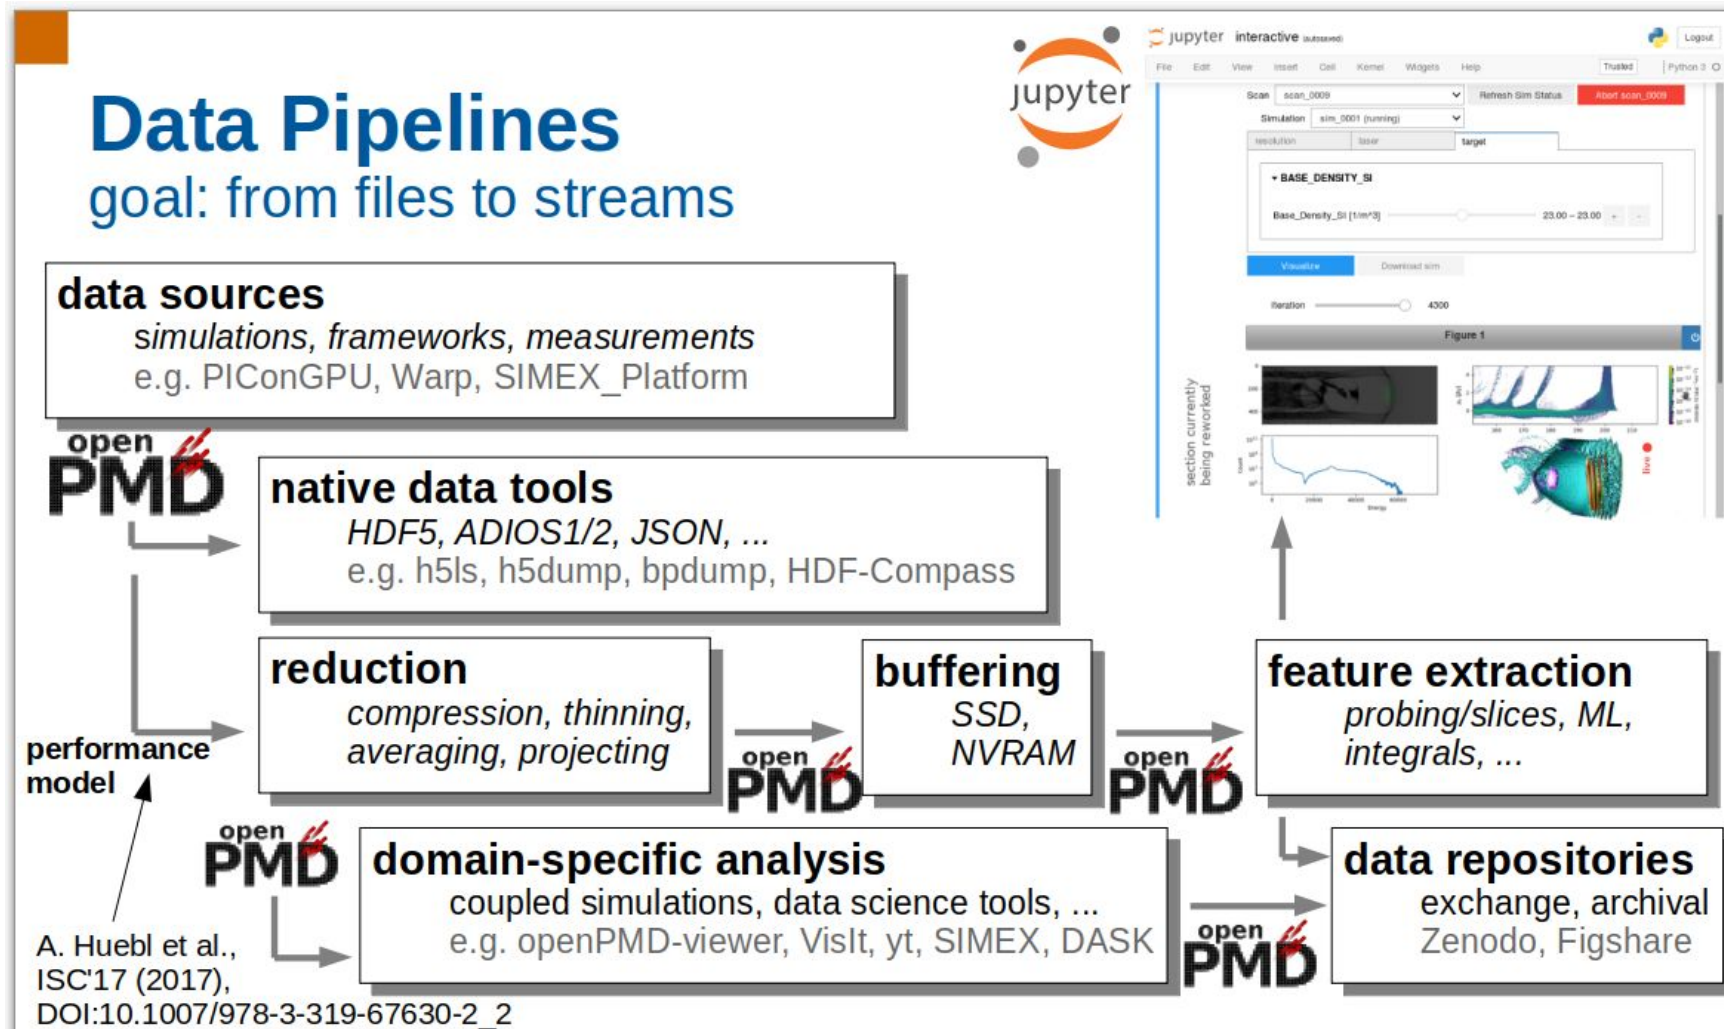
by F Poeschel, A Huebl et al., SMC21 (2022)



Impact of decomposition schemes when reading

Online Data Layout Reorganization:
[DOI:10.1109/TPDS.2021.3100784](https://doi.org/10.1109/TPDS.2021.3100784)
 by L Wan, A Huebl et al., TPDS (2021)

Streaming: Avoid Files Altogether



Right: For each arrow in the left figure, one could change the mapping between computing nodes (ranks).

Summary & Outlook

openPMD is

- a **standardized schema** and
- a large **open-source ecosystem**
 - documentation, example data, validation, scripts, integrations, reference libraries

Our community

- **evolves** and practices **open standardization**
- integrates multiple **state-of-the art** computer science **formats & tools**
 - we are not afraid of PByte-scale workflows
 - from parallel, in-transport data processing to file-less scripting workflows using RMDA

A growing number of international contributors

- try the tools, interact with us on what your needs are, contribute/share data & code



openPMD



openpmd.slack.com

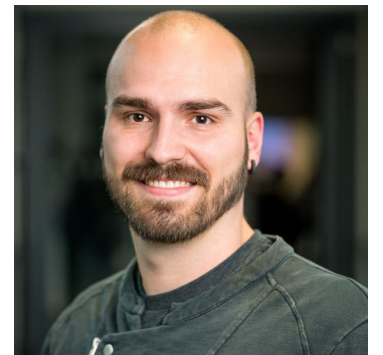


openpmd.org

Future Directions

- Integration with **experimental data acquisition** systems
- **Data Repositories / Portals**
 - AI/ML training, testing, validation, calibration
 - combine data sets
 - preservation, recasting and reinterpretation

presented by: **Axel Huebl (LBNL)**
✉ axelhuebl@lbl.gov



Acknowledgements

This project has received funding from the European Unions Horizon 2020 research and innovation programme under grant agreement No 654220. Supported by the CAMPA collaboration, a project of the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research and Office of High Energy Physics, Scientific Discovery through Advanced Computing (SciDAC) program. Previously supported by the Consortium for Advanced Modeling of Particle Accelerators (CAMPA), funded by the U.S. DOE Office of Science under Contract No. DE-AC02-05CH11231.



This work was partially funded by the Center of Advanced Systems Understanding (CASUS), which is financed by Germany's Federal Ministry of Education and Research (BMBF) and by the Saxon Ministry for Science, Culture and Tourism (SMWK) with tax funds on the basis of the budget approved by the Saxon State Parliament.



This research was supported by the Exascale Computing Project (17-SC-20-SC), a joint project of the U.S. Department of Energy's Office of Science and National Nuclear Security Administration, responsible for delivering a capable exascale ecosystem, including software, applications, and hardware technology, to support the nation's exascale computing imperative.

This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725 and of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231.

Backup Slides

The Open Standard for Particle-Mesh Data (openPMD) is a metadata standard for tabular (particle/dataframe) and structured mesh data in science and engineering. We show the basic components of openPMD, its extensions to specific domains and applications from laser-plasma physics, particle accelerators, light sources, astrophysics to imaging.

openPMD is implemented on top of portable, hierarchical data formats, especially HDF5 and ADIOS/ADIOS2. An extensive community ecosystem enabled productive workflows for developers and users alike, spanning Exascale simulations, in-transit data processing, post-processing, 3D visualization, GPU-accelerated data analytics and AI/ML. We will present the organization of this community, benefits and experience from supporting multiple data format backends, and future directions.

[1] Axel Huebl, Remi Lehe, Jean-Luc Vay, David P. Grote, Ivo F. Sbalzarini, Stephan Kuschel, David Sagan, Christopher Mayes, Frederic Perez, Fabian Koller, and Michael Bussmann. "openPMD: A meta data standard for particle and mesh based data," DOI:10.5281/zenodo.591699 (2015)

[2] Homepage: <https://www.openPMD.org>

[3] GitHub Organization: <https://github.com/openPMD>

[4] Projects using openPMD: <https://github.com/openPMD/openPMD-projects>

[4] Reference API implementation: Axel Huebl, Franz Poeschel, Fabian Koller, and Junmin Gu. "openPMD-api 0.14.3: C++ & Python API for Scientific I/O with openPMD," DOI:10.14278/rodare.1234 (2021) <https://openpmd-api.readthedocs.io>

[5] Selected earlier presentations on openPMD: <https://zenodo.org/search?page=1&size=20&q=openPMD&type=presentation>

[6] Axel Huebl, Rene Widera, Felix Schmitt, Alexander Matthes, Norbert Podhorszki, Jong Youl Choi, Scott Klasky, and Michael Bussmann. "On the Scalability of Data Reduction Techniques in Current and Upcoming HPC Systems from an Application Perspective," ISC High Performance 2017: High Performance Computing, pp. 15-29, 2017. arXiv:1706.00522, DOI:10.1007/978-3-319-67630-2_2

[7] Franz Poeschel, Juncheng E, William F. Godoy, Norbert Podhorszki, Scott Klasky, Greg Eisenhauer, Philip E. Davis, Lipeng Wan, Ana Gainaru, Junmin Gu, Fabian Koller, Rene Widera, Michael Bussmann, and Axel Huebl. Transitioning from file-based HPC workflows to streaming data pipelines with openPMD and ADIOS2, Part of Driving Scientific and Engineering Discoveries Through the Integration of Experiment, Big Data, and Modeling and Simulation, SMC 2021, Communications in Computer and Information Science (CCIS), vol 1512, 2022. arXiv:2107.06108, DOI:10.1007/978-3-030-96498-6_6

Organizing Scientific Records

- electric field $\vec{E}(\vec{r})$

/ ... / meshes / E /

x y z

- temperature $T(\vec{r})$

/ ... / meshes /

T

- electron position \vec{r}

/ ... / particles / electrons / position /

x y z



extensions to the base standard
define additional,
compatible keywords

```
charge (64360, 4)
Group size = 0
Number of attributes = 3
unitDimension = 0.0,0.0,1.0,1.0,0.0,0.0,0.0
unitSI = 1.60217657E-19
value = -1.0
```

Unit System

- **unitDimension**

automated description of physical dimension

only powers of base dimensions

length **L**, mass **M**, time **T**, electric current **I**, thermodynamic temperature **theta**,
amount of substance **N**, luminous intensity **J**

$$\begin{aligned} \text{magnetic field: } [B] &= M / (I * T^2) \\ &\rightarrow (0., 1., -2., -1., 0., 0., 0.) \end{aligned}$$

- **unitSI** (recommended)

relation to an *absolute* unit system



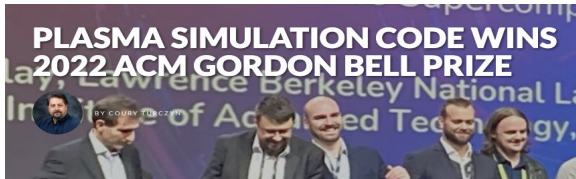
attributes

WarpX is a GPU-Accelerated PIC Code for Exascale

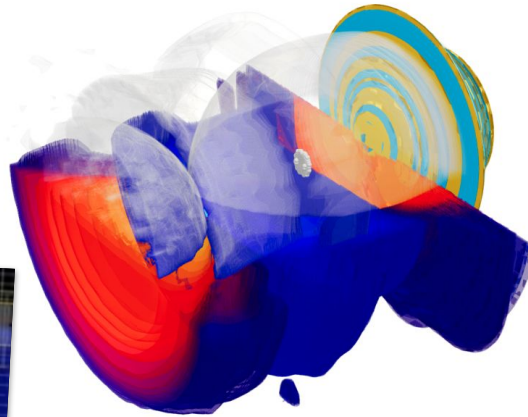
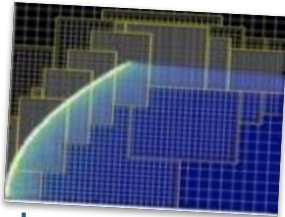


Particle-in-Cell Loops

electrostatic & electromagnetic (fully kinetic)



WarpX was first PIC code to win prestigious ACM Gordon Bell award, the “Oscars” of Supercomputing.



Advanced algorithms

boosted frame, spectral solvers, Galilean frame, embedded boundaries + CAD, MR, ...

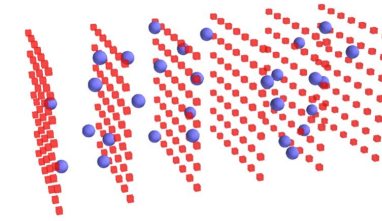
Multi-Physics Modules

field ionization of atomic levels, Coulomb collisions, QED processes (e.g. pair creation), macroscopic materials

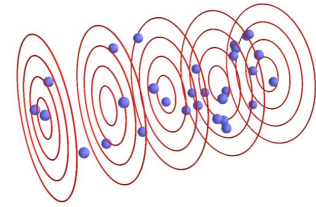


Geometries

- 1D3V, 2D3V, 3D3V and RZ (quasi-cylindrical)



3D Cartesian grid



Cylindrical grid (schematic)

Multi-Node parallelization

- MPI: 3D domain decomposition
- dynamic load balancing



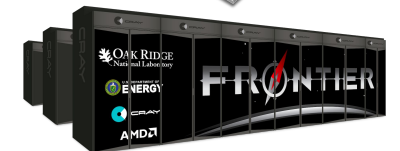
On-Node Parallelization

- GPU: CUDA, HIP and SYCL
- CPU: OpenMP



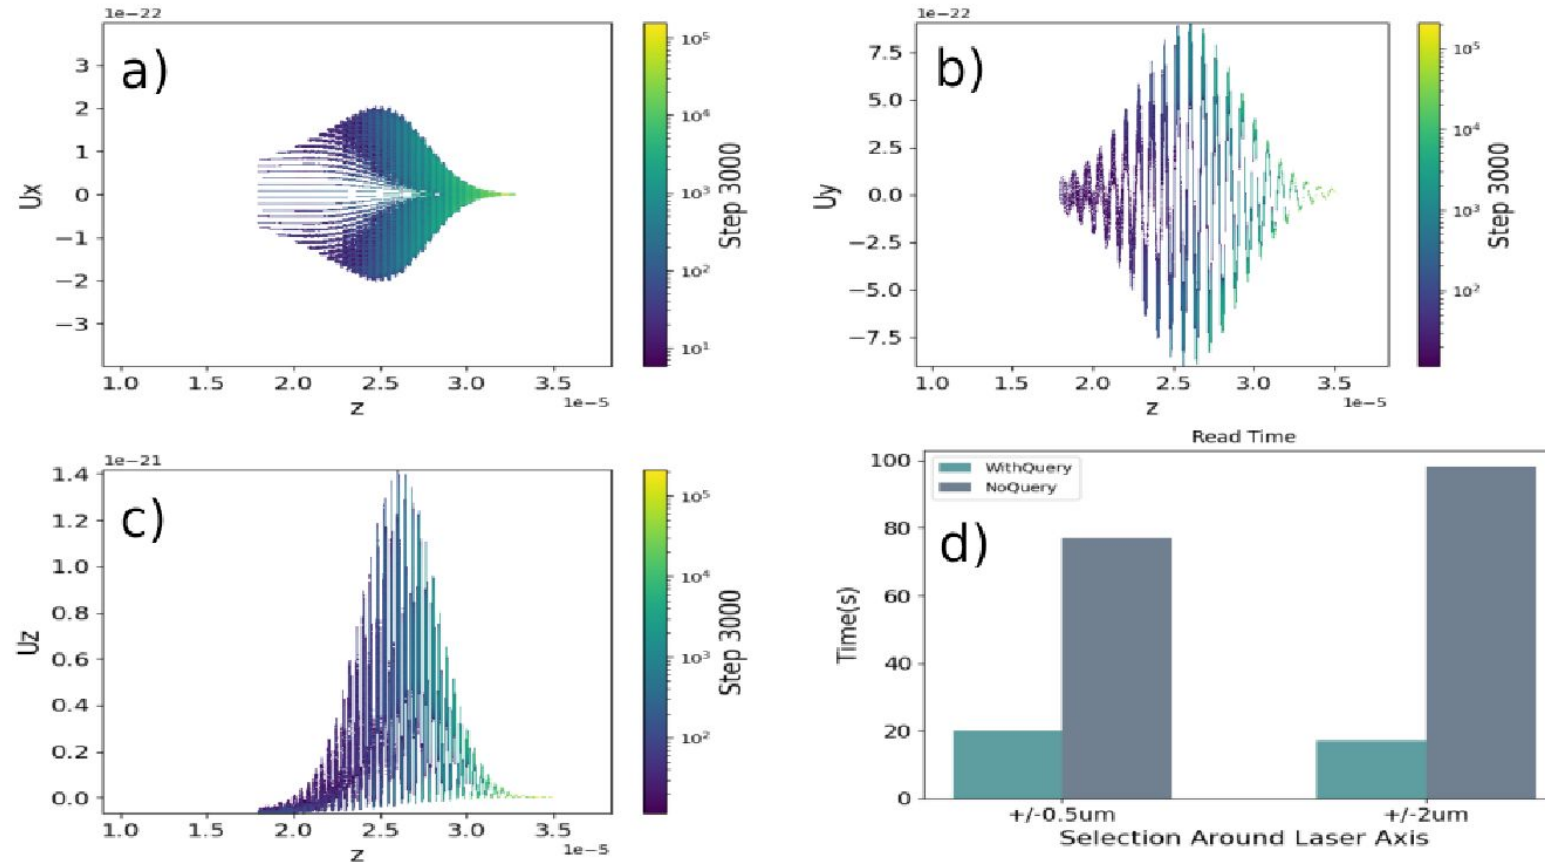
Scalable, Standardized I/O

- PICMI Python interface
- openPMD (HDF5 or ADIOS)
- in situ diagnostics



Open Source: 
ecp-warpx.github.io

In-Transport Data Processing: Data Statistics



Data analysis using region of interest filtering with ADIOS queries. a)-c) Phase space projections of plasma particles oscillating in a laser pulse, filtered close to the laser axis. d) Read time comparison between conventional reads and pre-filtered reads with queries.

Integration into the BELLA Control System

BELLA GEECS

- **G**eneralized **E**quipment and **E**xperiment **C**ontrol **S**ystem
- control and data acquisition system



openPMD/openPMD-CCD



GEECS-BELLA

openPMD-CCD

- **P**ython module for organizing CCD images with openPMD
- Optional integration with **LabView 2020+**

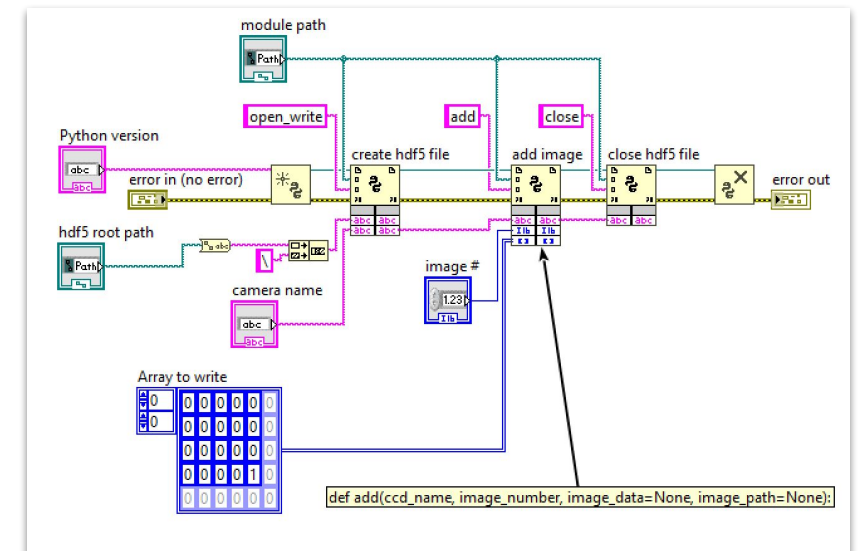
```
from openpmd_ccd import CCD

scan = CCD("defaultCam_scan001.h5", overwrite=True,
          name="Go Pro", model="HER08 Black", serial="12345678",
          operator="Axel Huebl <axelhuebl@lbl.gov>",
          # resolution=None, roi=None, exposure_time=None
)

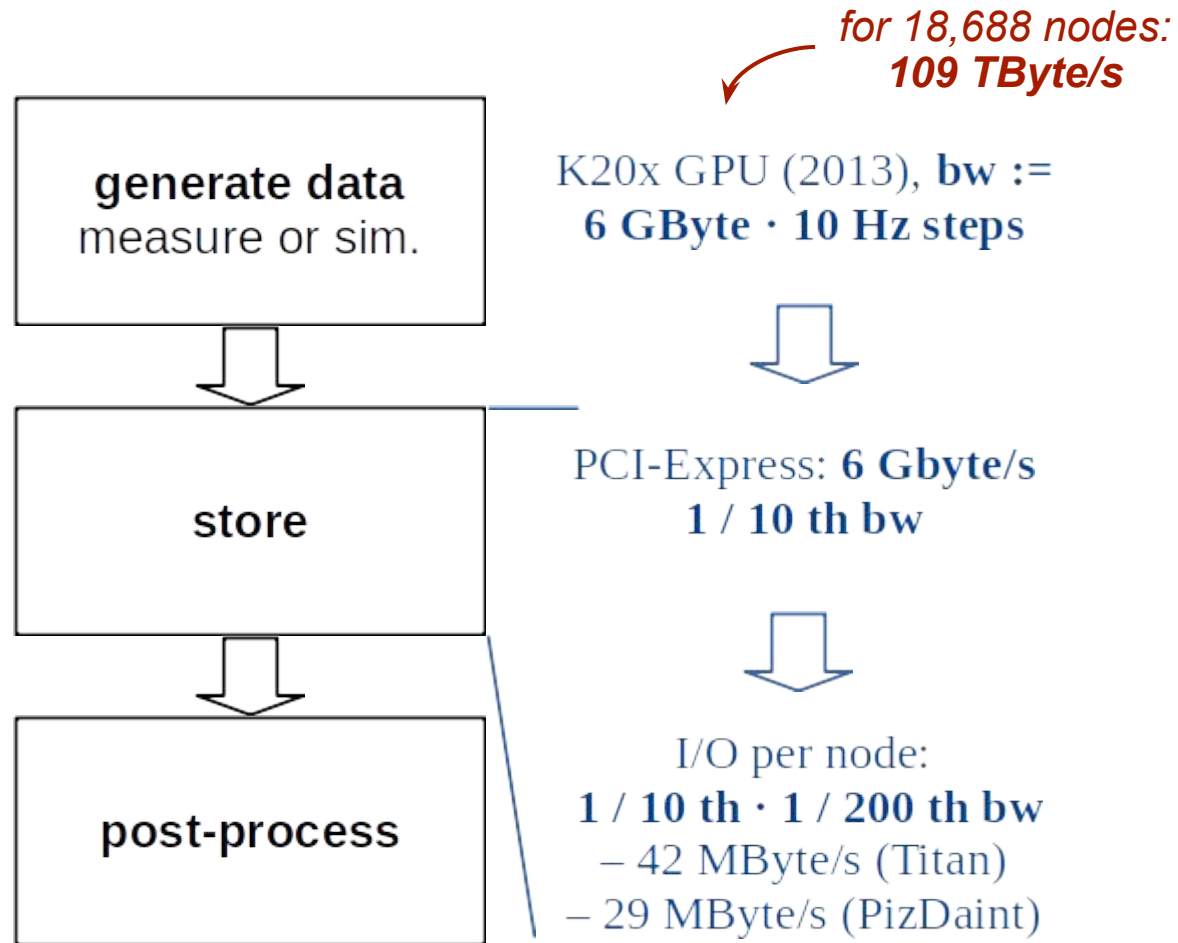
scan.add(0, "tests/data/Scan005_SimCam_001.png")
scan.add(1, np.array([[1., 2.], [3., 4.])))

# scan recalibrate(...)

scan.close()
```



HPC Background: Our Data Processing Bottlenecks Look Alike



Summit (ORNL, 2018): ratio 4x “worse” - gap of 10^4

Challenges

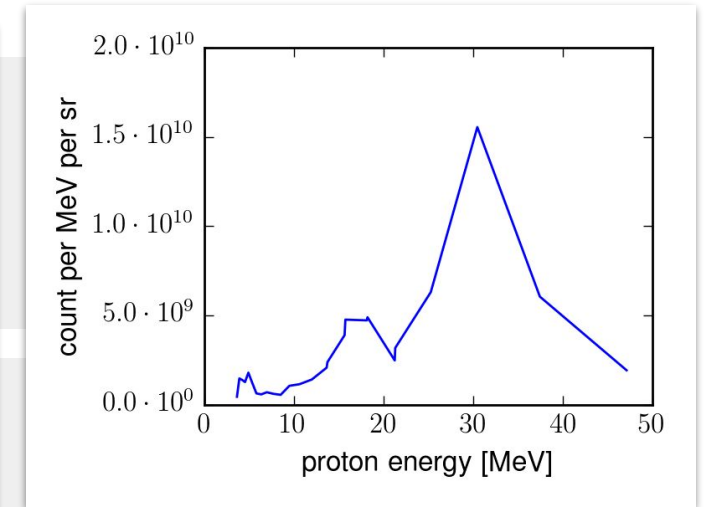
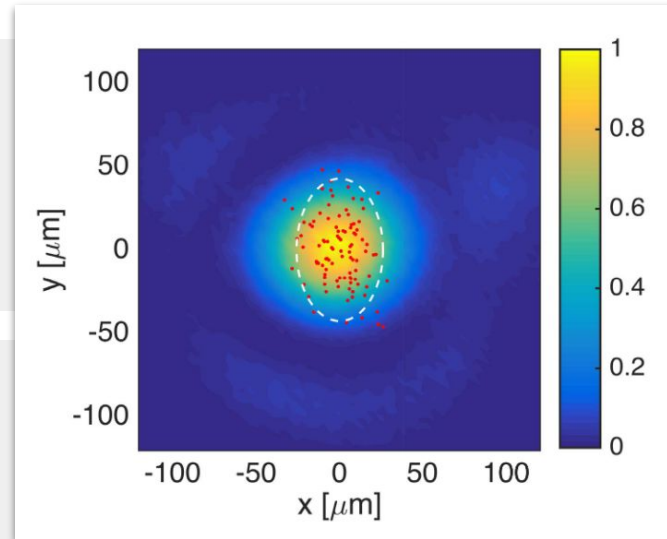
- **3 orders of magnitude gap** between producing devices and storage
- “store & analyze everything” is *unaffordable*

Opportunities

- analysis tasks have varying **fidelity** needs
- many common tasks can be done **in situ**
- manual steps: limit the sampling of raw data to **setup phase**

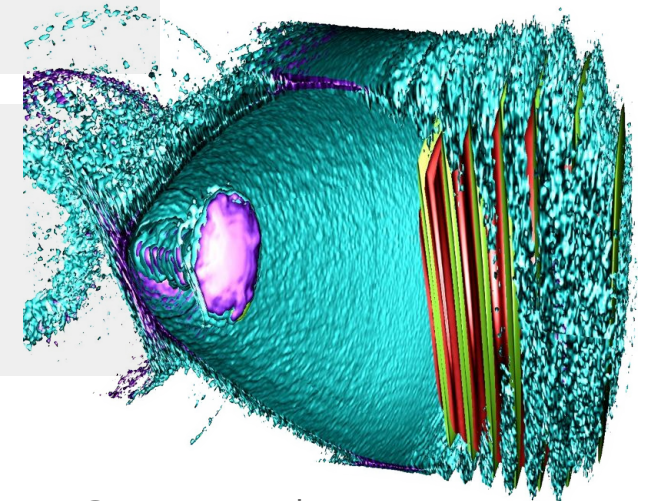
Data Processing & Reduction Examples

Binning of a spectrogram
Fitting of an ellipsoid



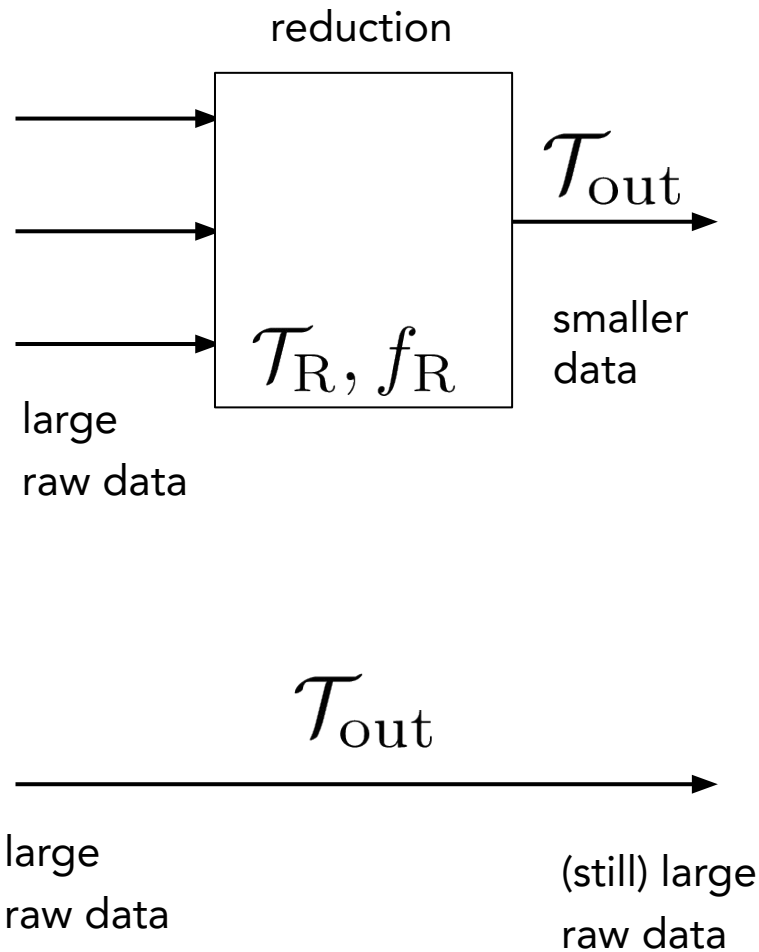
Compression (lossless/lossy)

Ray-casting 3D data,
training a neural network, etc.

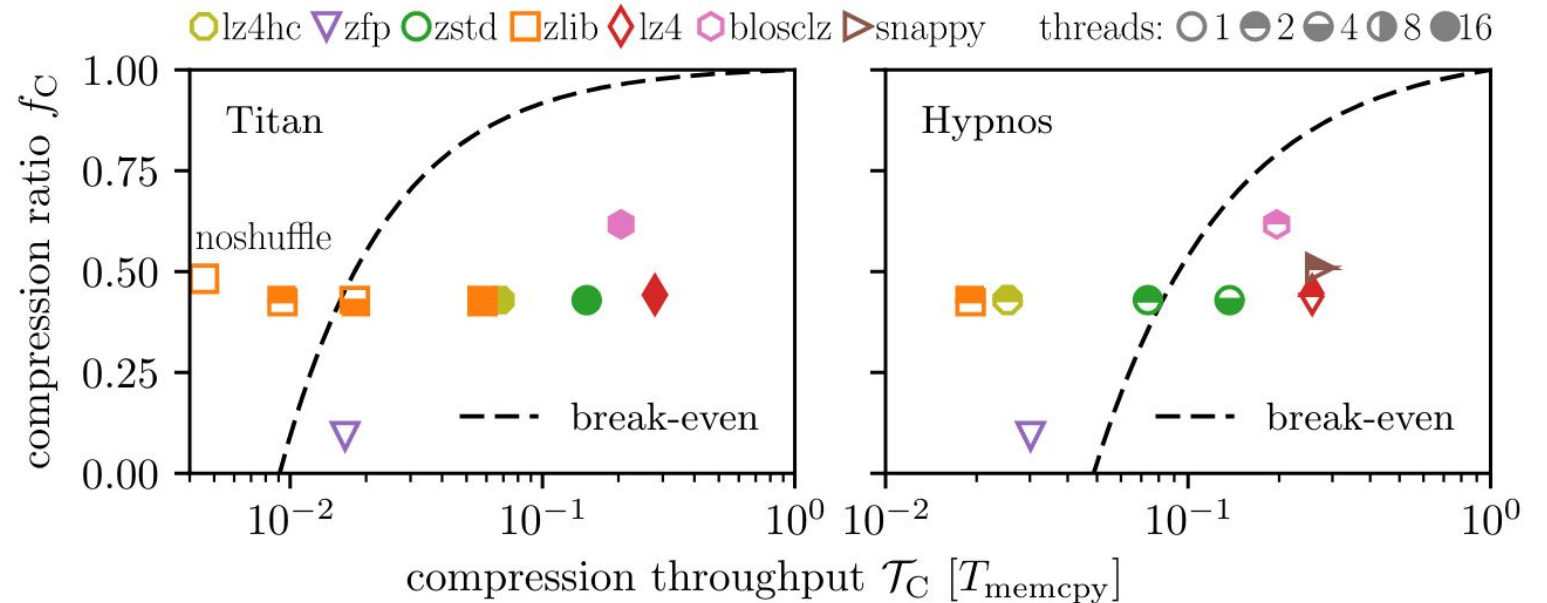


A Matthes, A Huebl et al., ISC 2017, DOI:10.14529/jsfi160403 (2017); K Nakamura et al., IEEE J. Quantum Electron, A Huebl et al., ISC 2017, DOI:10.1007/978-3-319-67630-2_2 (2017); DOI:10.1109/JQE.2017.2708601 (2019)

Avoid Backlog: Design Criteria for Data Reduction Pipelines

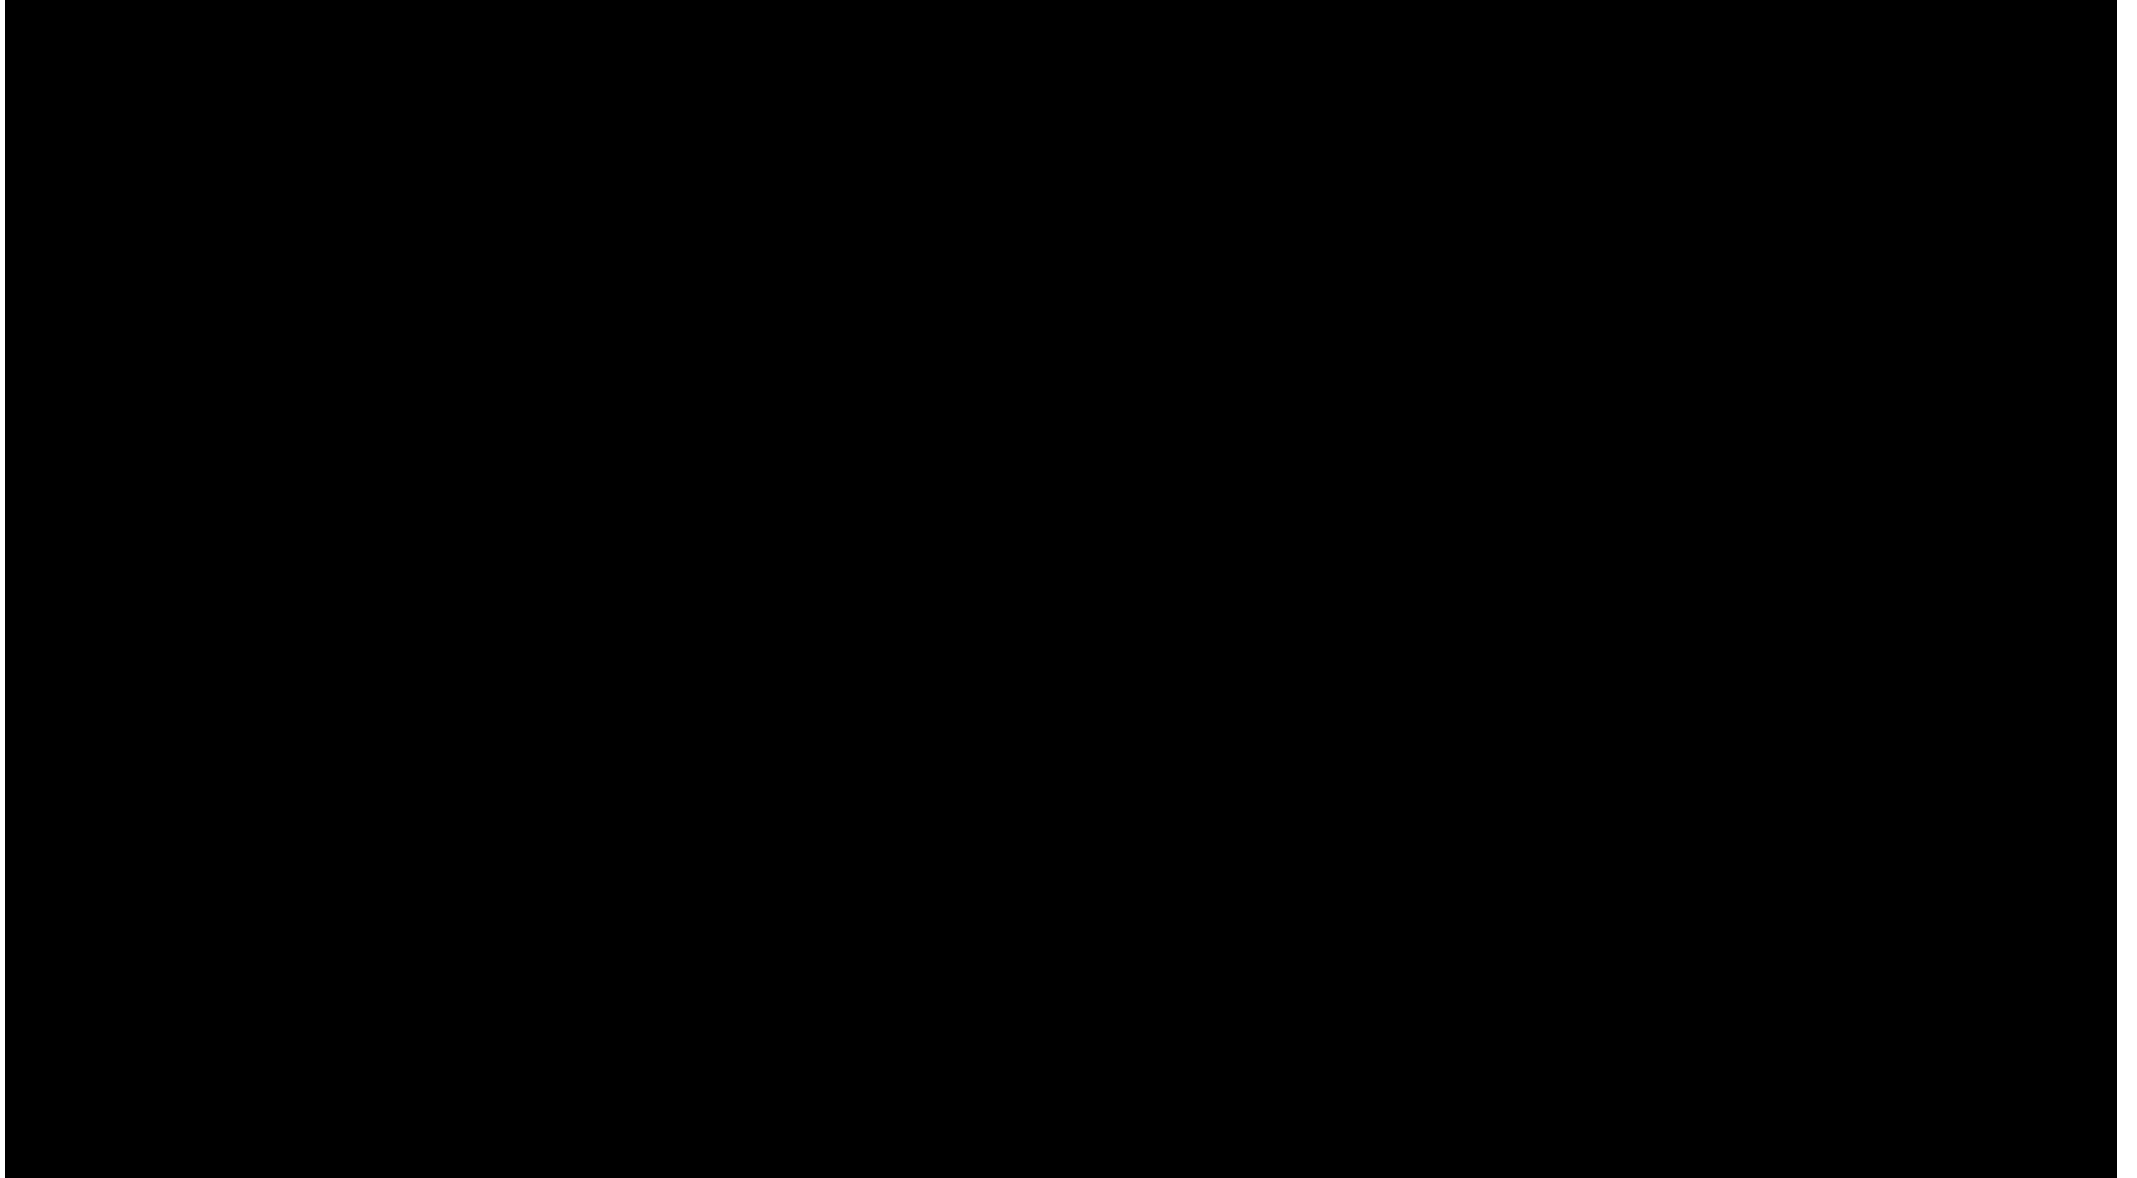


$$\frac{\mathcal{T}_R \times (1 - f_R)}{1 - \mathcal{T}_R} > \mathcal{T}_{out}$$



Result: **trade compute for throughput**, >100 GByte perceived application throughput and **280 GByte/s** peak parallel filesystem throughput

DASK Example: Modern Access to Parallel Processing



openPMD - Extension Example

Wavefronts, Particle Species, Particle Beams, Weighted Particles, PIC, MD, Mesh-Refinement, CCD images, ...



Convention for Specifying Particle Species

openPMD extension name: `SpeciesType`

Introduction

This convention is for standardizing the names of particle species, e.g. in particle physics.

Additional Record Attribute

The following additional attribute for openPMD `mesh records` and `particle groups` is defined in this extension:

- `speciesType`
 - type: *(string)*
 - scope: *optional*
 - description: particle species in this record. If there are multiple species to be specified, they can be specified using a semicolon separated list.
 - allowed values:
 - see the lists below and additionally
 - `other`: If none of the ones below applies, user are free to append a free text after a colon, e.g. `other:neutralino` or `other:cherry`
 - examples:
 - `electron` (e.g. on an electron `particle record` or an electron density `mesh record`)
 - `electron;proton;#12C` (e.g. on a `mesh record` for a plasma's local charge density)
 - `other:apple;other:orange` (for a `record` mixing apples & oranges)

This attribute can be used with any `record` (including `mesh records`).

Elementary Particles

Namings for fundamental fermions and their anti-matter particles.

Quarks:

- `up`, `anti-up`

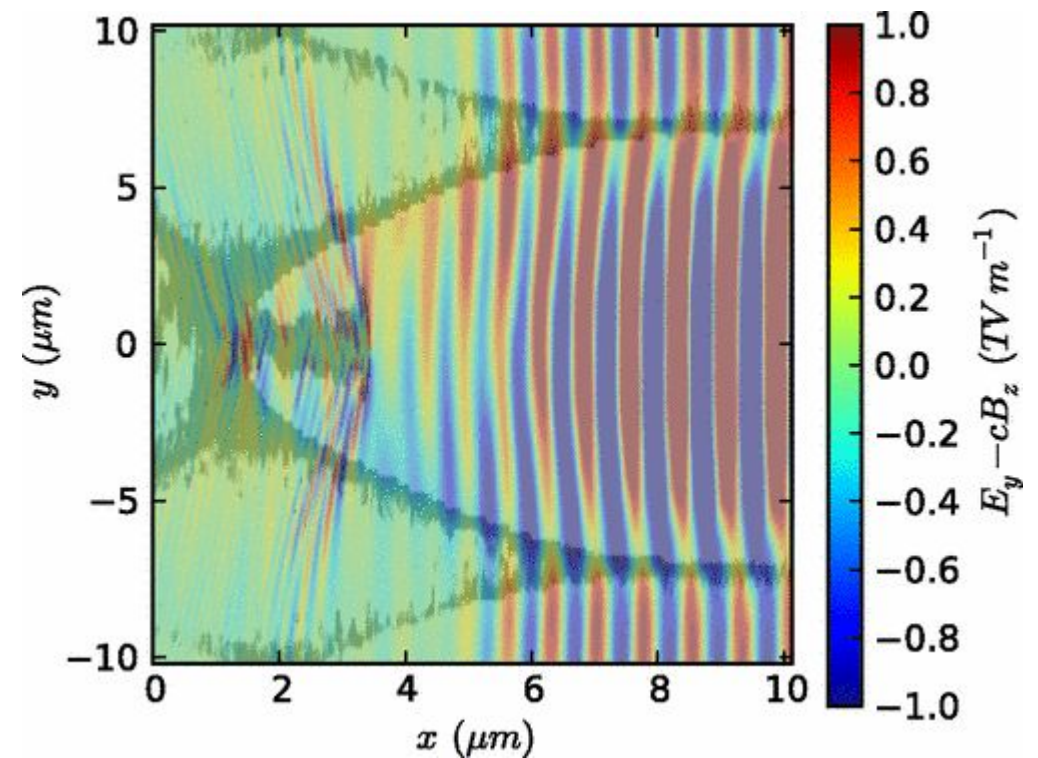
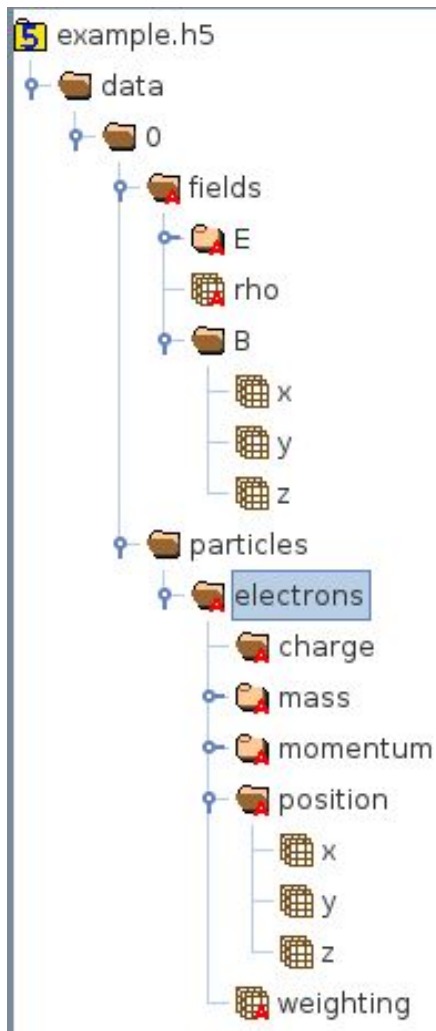


Image CC-BY 3.0: R. Lehe et al.,
PRSTAB 16, 021301 (2013), DOI:10.1103/PhysRevSTAB.16.021301

openPMD-api: Reference Implementation



Flagship implementation: **openPMD-api**

- **API** in C++ and Python (upcoming: Julia)
- Flexibly store to / read from interchangeable **backends**:
 - ADIOS1/2
 - HDF5
 - JSON (serial only)
- Applies **best practices & performance tuning**



openPMD/openPMD-api

Example users:

- Simulations: WarpX, PIConGPU, Wake-T, GPos
- Frameworks: SimEx, LUME, OASYS1-PaNOSC
- Visualization & analysis: ParaView, DASK, VisualPIC, APTools

A Huebl, F Poeschel, F Koller, J Gu and contributors (2018).

openPMD-api: C++ & Python API for Scientific I/O with openPMD. DOI:10.14278/rodare.27

Ecosystem



openPMD/openPMD-projects

openPMD standard (1.0.0, 1.0.1, 1.1.0)

the underlying file markup and definition

A Huebl et al., DOI:10.5281/zenodo.33624

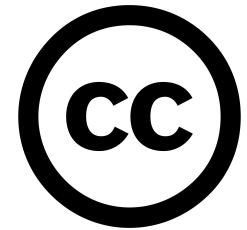
base standard

general description

wavefronts, particle species, particle beams, weighted particles, PIC, MD, mesh-refinement, CCD images, ...

extensions

domain specific



native data tools

HDF5, ADIOS1/2, NetCDF, ...

e.g. h5ls, h5repack, h5dump, bpdump

writers & converters

simulations, frameworks, measurements

e.g. PIConGPU, Warp, SIMEX_Platform

HDF Compass

HDF5 & ADIOS file explorer

open and explore file trees

readers

coupled simulations, post-processing frameworks, ...

e.g. SIMEX_Platform, VisIt, yt-project, openPMD-viewer

openPMD-updater

update to new standard

edit in- or new file

openPMD-api

I/O library abstraction

file format agnostic

data repositories

exchange and long-time archival

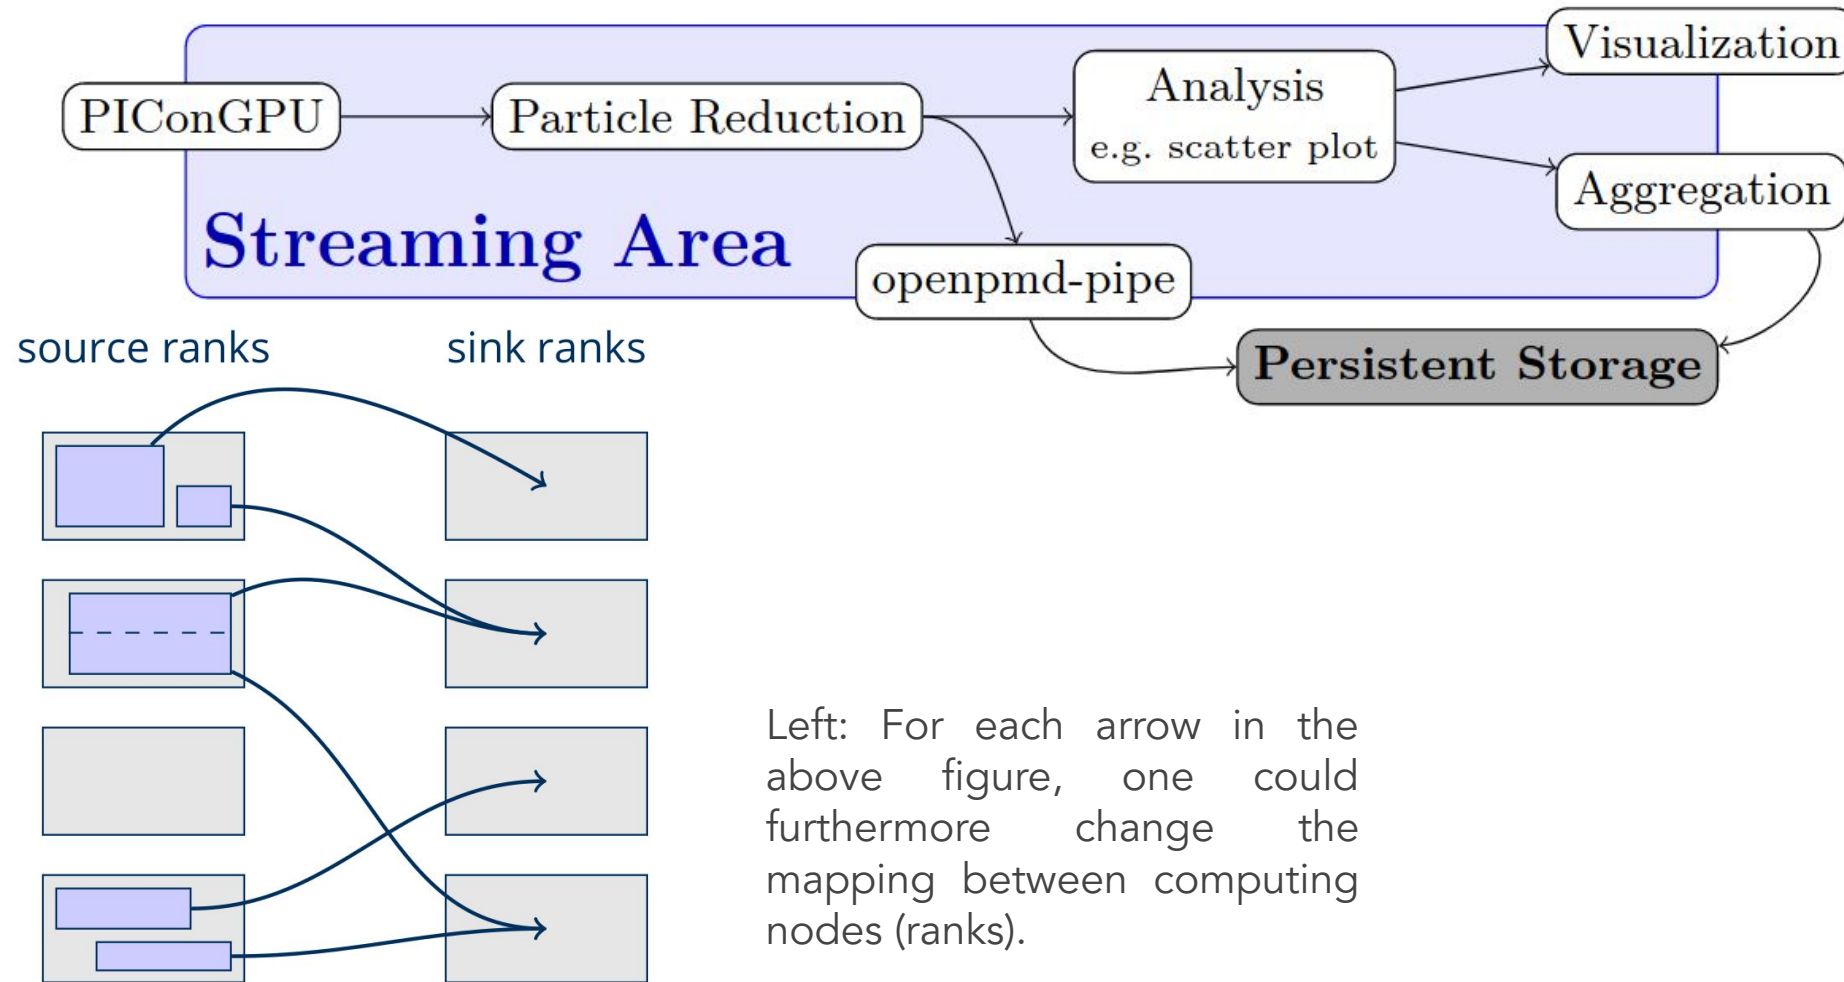
e.g. Zenodo, RODARE (HZDR)

A Rising Data Creation vs. Storage Gap

system	compute performance [PFlop · s ⁻¹]	parallel FS bandwidth [TiByte · s ⁻¹]	FS capacity [PiByte]	example storage requirements [PiByte]
Titan	27	1	27	5.3
Summit	200	2.5	250	21.1
Frontier	> 1500	5 - 10	500 - 1000	80 - 100

Table 1: System performance: OLCF Titan to Frontier. The last column shows the storage size needed by a full-scale simulation that dumps all GPU memory in the system 50 times.

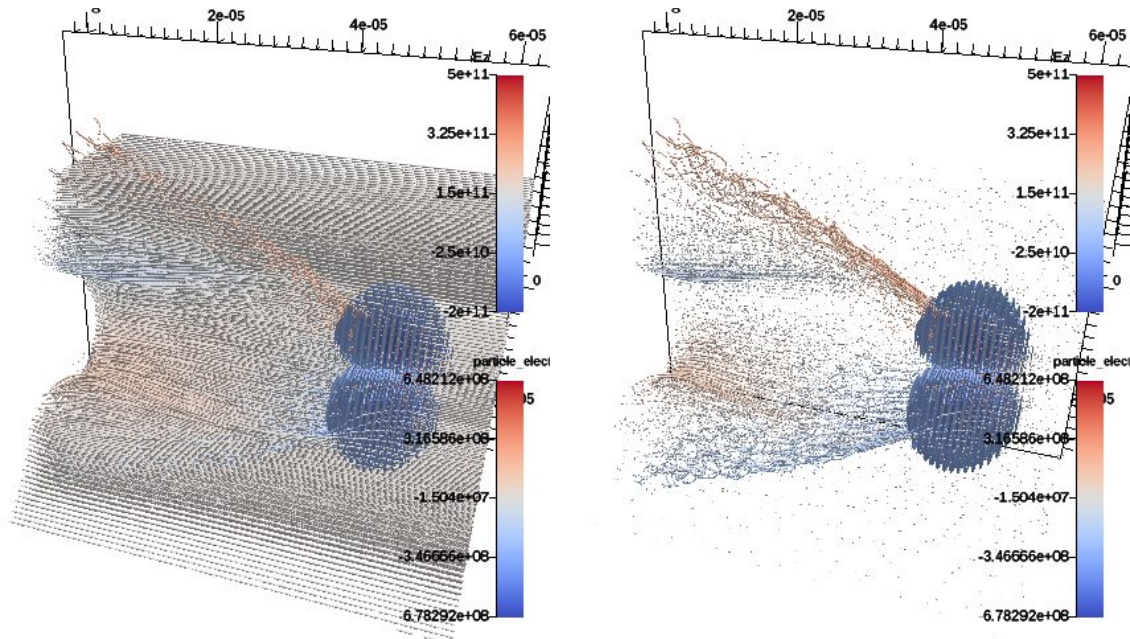
Streaming: Avoid Files Altogether



Novel Visualization Techniques

Particle Adaptive Sampling

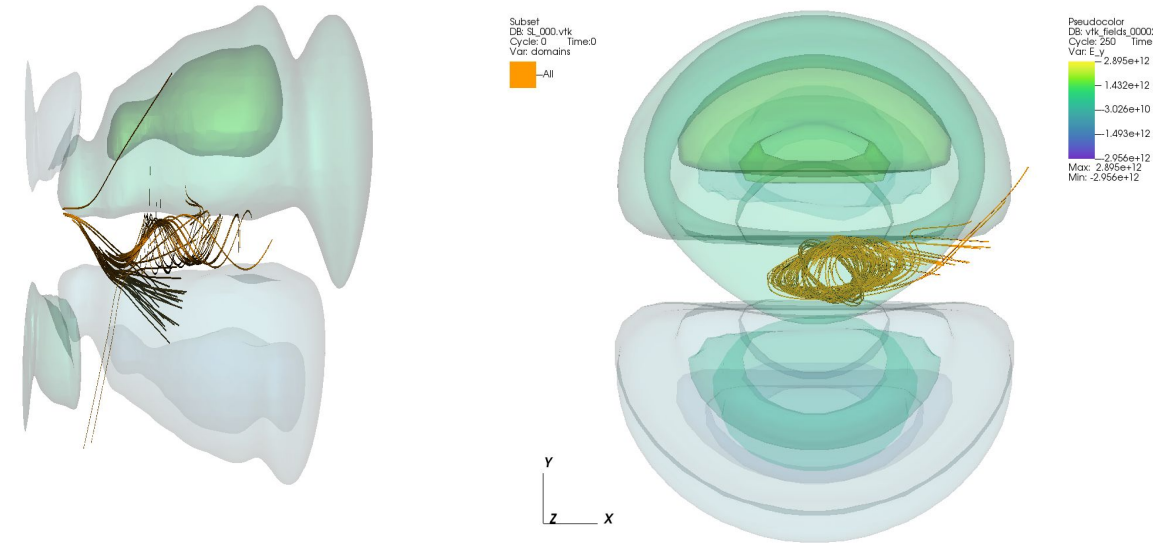
- **emphasis** on “uncommon” properties
- inverse sampling to incidence of a property



A. Biswas et al., “In Situ Data-Driven Adaptive Sampling for Large-scale Simulation Data Summarization,” ISAV18 @SC18 (2018)

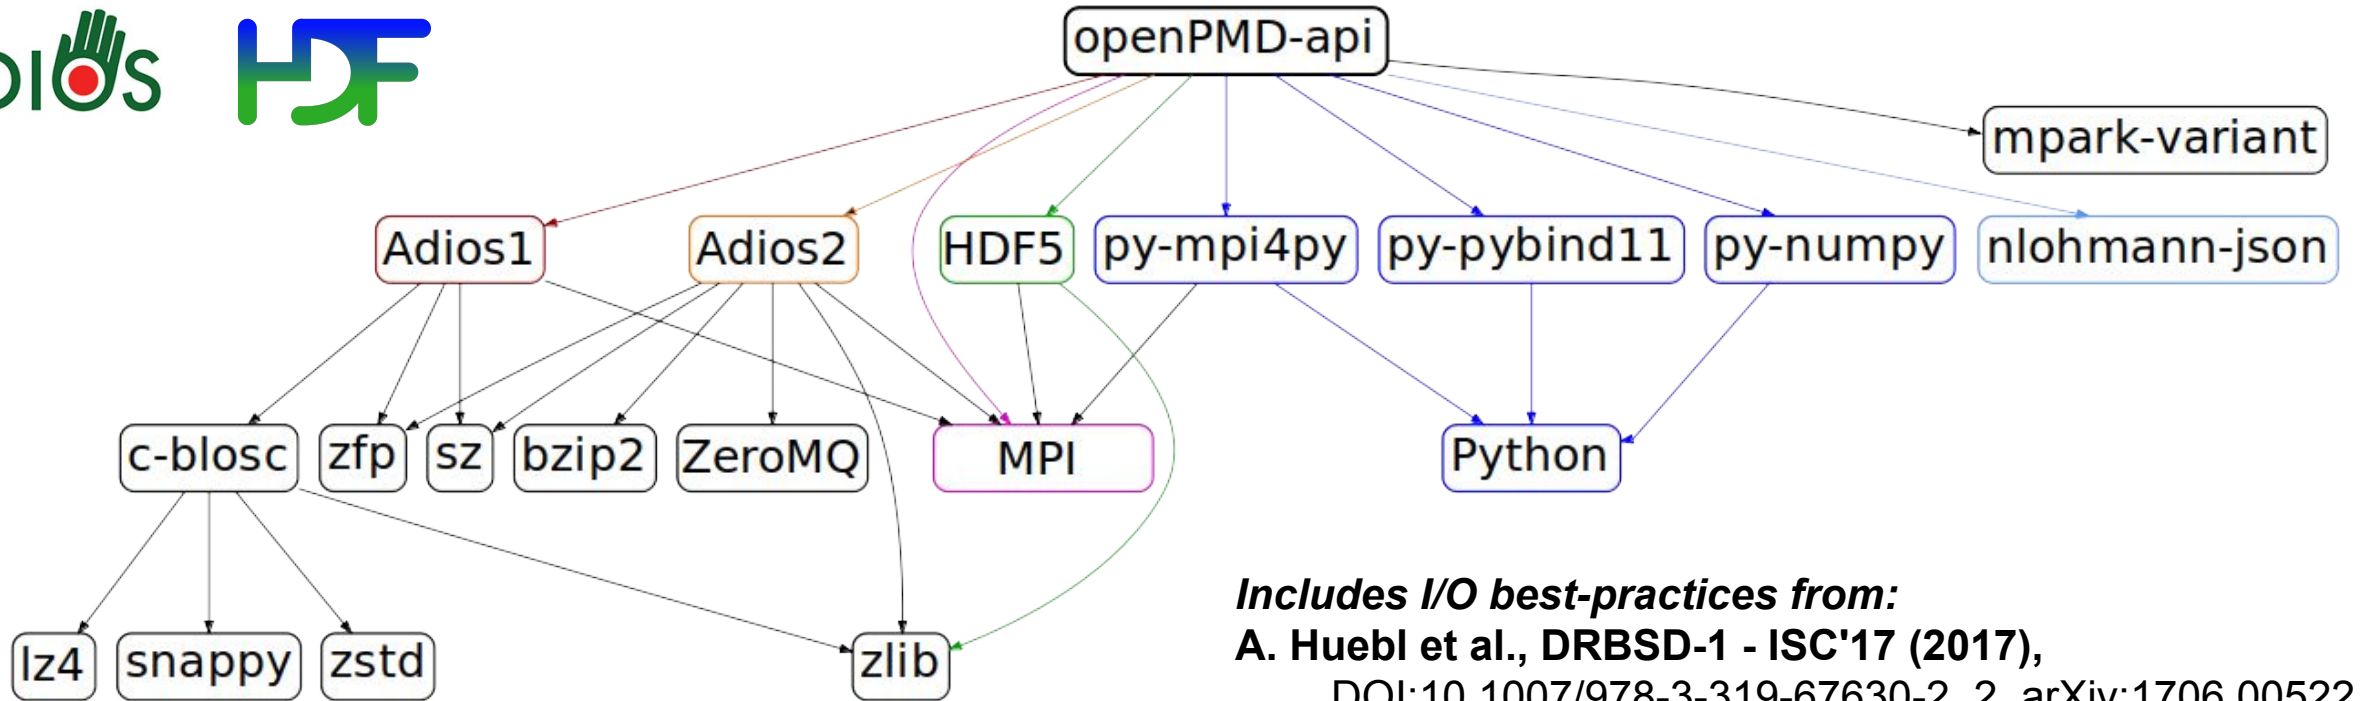
Physics-Informed Flow Tracelines

- traditional flow vis. depends only on *local field values*
- plasma particles:
 - **inert**: track *relativistic momentum* on a traceline
 - **Lorentz-Force**: 6 fields (electromag.), leap-frog
- chance to **significantly reduce particle I/O** in real-life workflows through savings on **temporal fidelity**



user: abhishek
Tue Dec 8 22:21:33 2020

openPMD-api: I/O Middleware Library



Includes I/O best-practices from:
A. Huebl et al., DRBSD-1 - ISC'17 (2017),
DOI:10.1007/978-3-319-67630-2_2, arXiv:1706.00522
F. Poeschel, A. Huebl et al., SMC21 (2021)
arXiv:2107.06108

Available via:



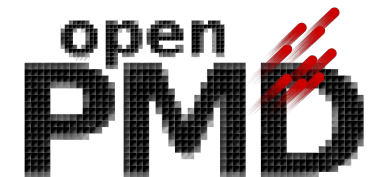
module load



Data Visualization & Analysis: Parallel, User-Facing Tooling

Continuously Improving User-Facing Data Workflows

- **Visit** integration: mainlining of an existing openPMD plugin (HDF5)
 - including quasi-cylindrical geometry (azimuthal decomposition)
Lobet, Huebl (LBNL), Pugmire (ORNL)
- **ParaView 5.9+** integration for openPMD (all formats, e.g. ADIOS2)
 - productivity: parallel plugin is fully written in Python (<500 LOC)
ParaView, openPMD-api, ADIOS2/HDF5 are all “lifting” in C/C++
 - next: azimuthal geometry, MR
Huebl (LBNL), Geveci (Kitware)
- **DASK/RAPIDS**
 - basis to parallelize Python analysis pipelines at NERSC & OLCF
 - potential candidate for further research with data staging (ADIOS)
Huebl (LBNL), Ganyushin (ORNL), Kirkham (Nvidia)



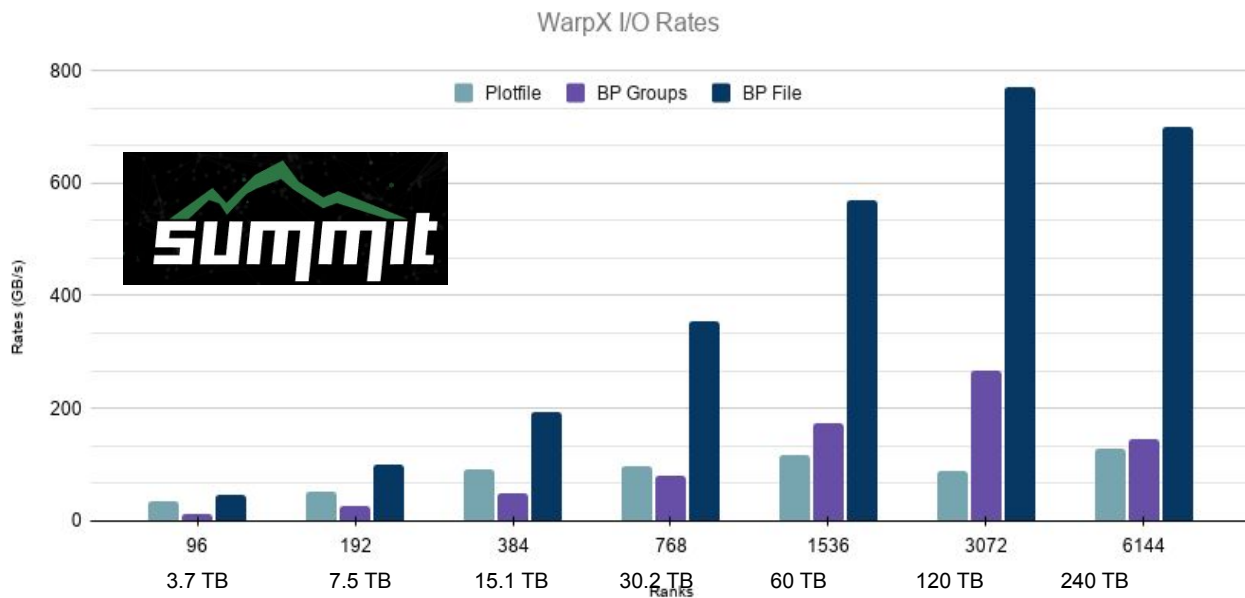
Open standardization, i.e. openPMD, makes us flexible for I/O libraries, tooling & domain-science needs.

Parallel I/O: integration of ADIOS2 through openPMD-api



Reducing I/O Risks in AMR Particle-in-Cell
baseline: subfiling & chunking for fast writes

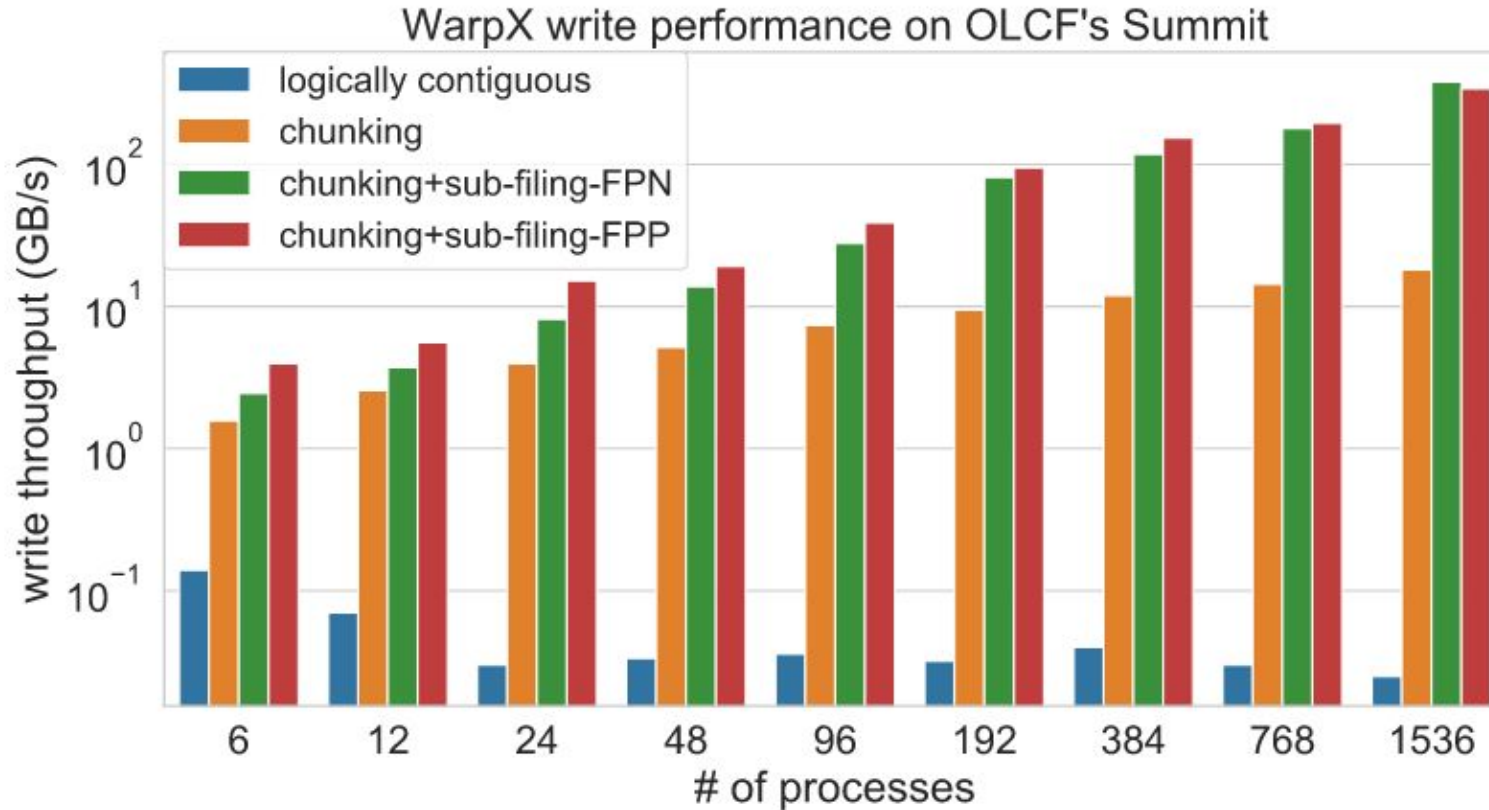
- aggregation of blocks, appending of steps
- representation of **sparsity in structured MR** meshes in **ADIOS blocks**



Write: plotfiles → ADIOS BP per rank & step → append to files

Aggregation: reduce number of files 6x

Parallel I/O: Logically Contiguous, Chunking, Sub-Filing



WarpX write performance under weak scaling tests

L. Wan, A. Huebl, J. Gu, F. Poeschel, A. Gainaru, R. Wang, J. Chen, X. Liang, D. Ganyushin, T. Munson, I. Foster, J.-L. Vay, N. Podhorszki, S. Klasky, "Data Layout Strategies for Parallel I/O: The Good, The Bad and The Ugly," submitted (2021)

Data Layout & Merging Improve Start-2-End Performance

Fast Writes

- Do-able



Fast Reads

- Very hard

Bridging both

- algorithm development
- library *implementations*
- integration in **scalable** analysis frameworks

Looking Ahead

- in-situ and in-transit pipelines
- smoothly transition domain scientists **from post-hoc to in-situ** scripts

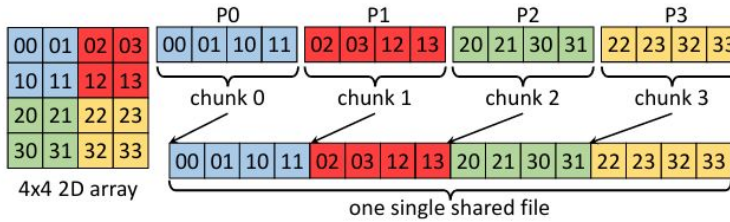
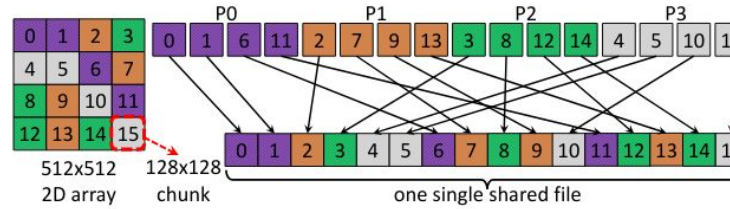
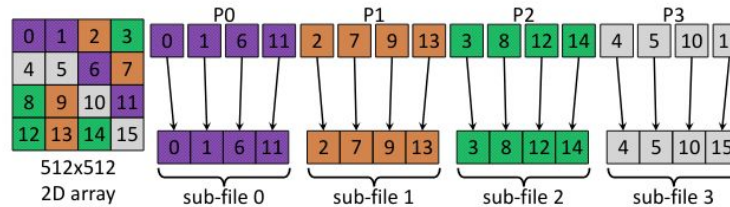


Fig. 2: Chunked Data layout: 1 chunk per process



(a) Without sub-filing



(b) With sub-filing

Fig. 3: Data layout: 4 chunks per process, 4 processors

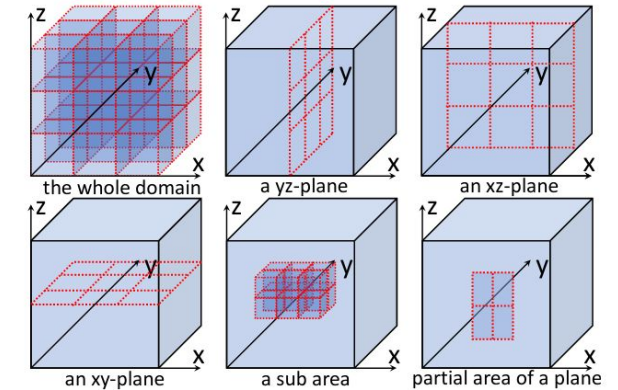


Fig. 5: Six common read patterns of a 3D mesh variable

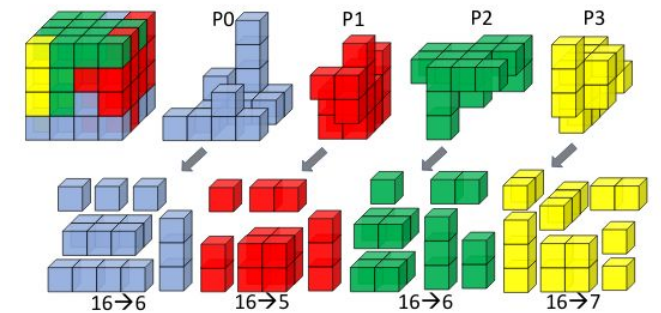
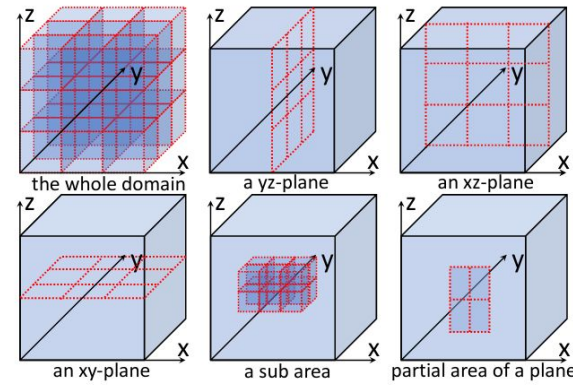
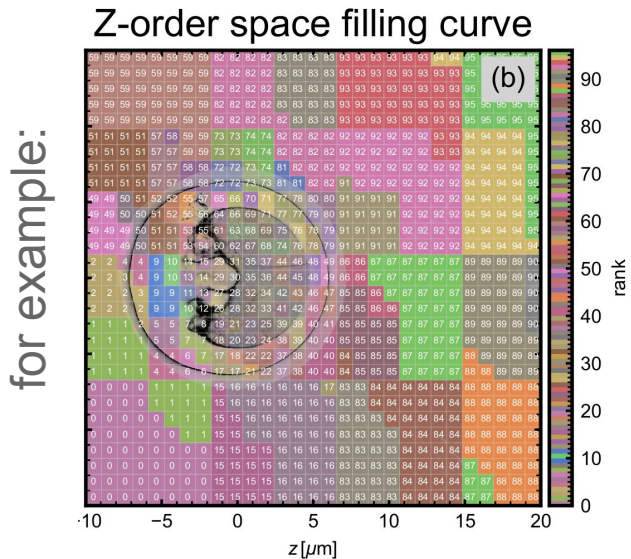


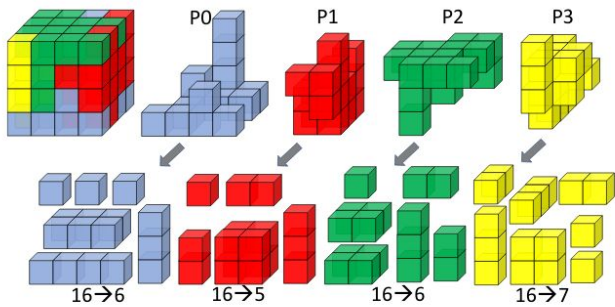
Fig. 8: Merging small blocks into bigger cuboid blocks

Block Meta-Data & Real-World Parallel Read Performance

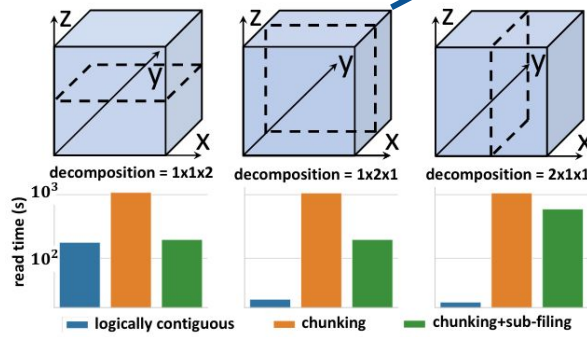
No one-fits-all data layout → bridging the gap & parallel reads



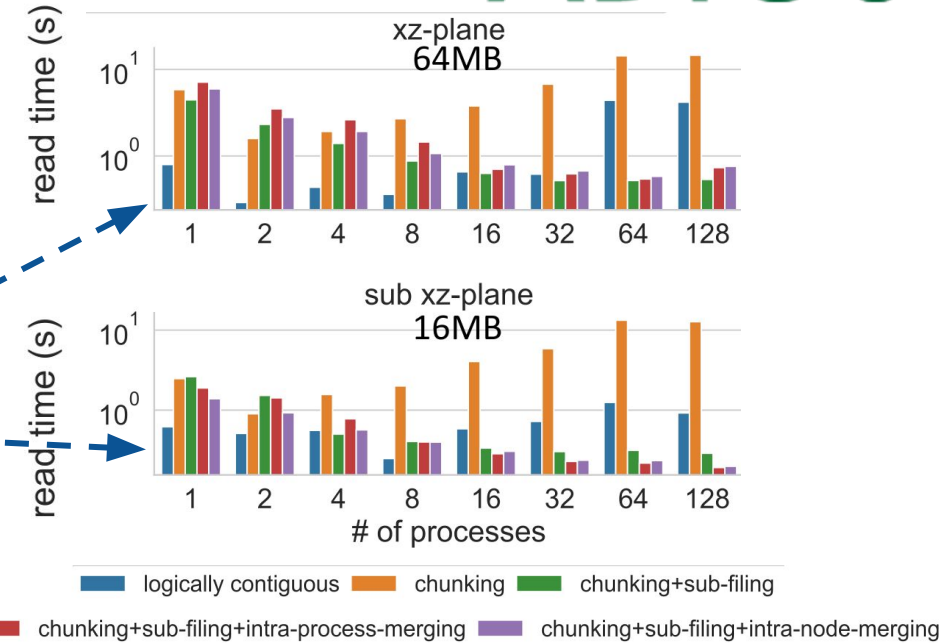
Six common read patterns of a 3D mesh variable



Merging small blocks into bigger cuboid blocks



Impact of decomposition schemes when reading



- ongoing: tuning MR fields +
 - lossy compressors: ZFP/MGARD
 - parallel post-processing; reordering
- orthogonal: async I/O (latency hiding)

Streaming: Avoid Files Altogether

Long-Term Strategy

- in-situ and in-transit pipelines
- smoothly transition domain scientists **from post-hoc to in-situ** scripts

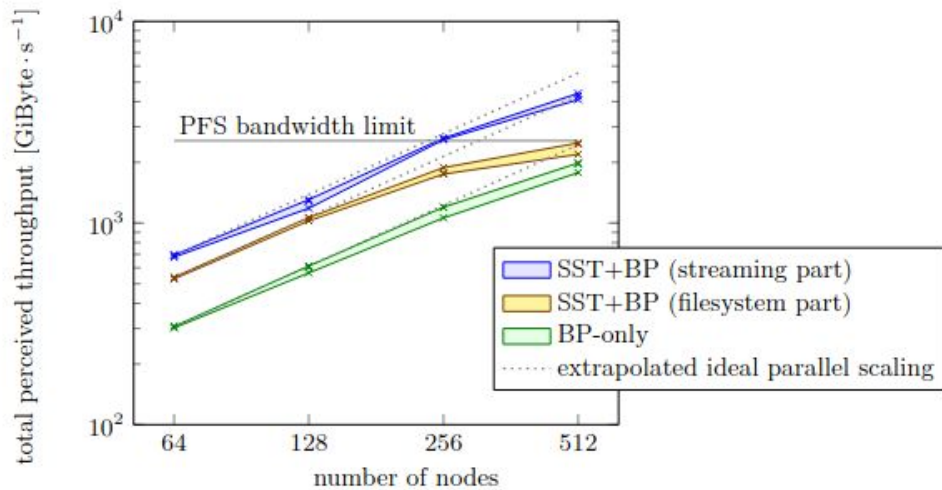


Fig. 5: Perceived total throughput. The file-based outputs (BP-only as well as SST+BP) are limited by the PFS bandwidth. At 512 nodes, the methods reach 4.15, 2.32, and 1.86 $\text{TiByte} \cdot \text{s}^{-1}$ on average, respectively.

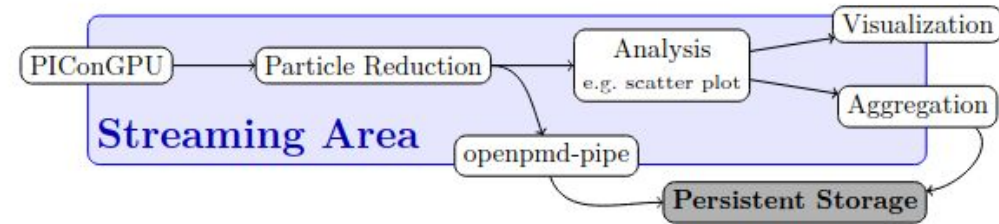


Fig. 2: A complex, loosely-coupled workflow: PICongPU is the data producer, a domain-specific particle reduction can conserve relevant ensemble properties, the analysis step might filter and bin, and aggregation might create a temporal integration from high-frequency data. At various sections of the workflow, visualization or data dumps might be generated from subscribers.

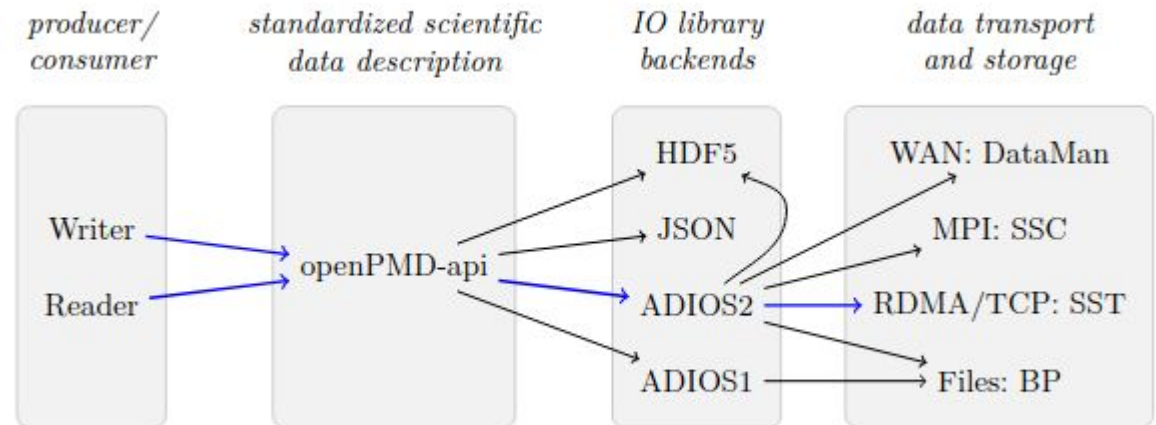


Fig. 3: IO software stack built by openPMD and ADIOS2

F. Poeschel, ... A. Huebl. "Transitioning from file-based HPC workflows to streaming data pipelines with openPMD and ADIOS2," accepted in SMC21, <https://arxiv.org/abs/2107.06108> (2021)

```
In [ ]: ts_2d.slider()
```

Calling this method will insert the following panel inside the notebook. (Note that the panel below is a **non-interactive image**, which is here for the users that are viewing this notebook online. Calling the `slider` method in a live notebook, will produce a truly **interactive** panel.)

The interface displays a time slider labeled "t (fs)" with a value of 33. Below the slider are two main control panels:

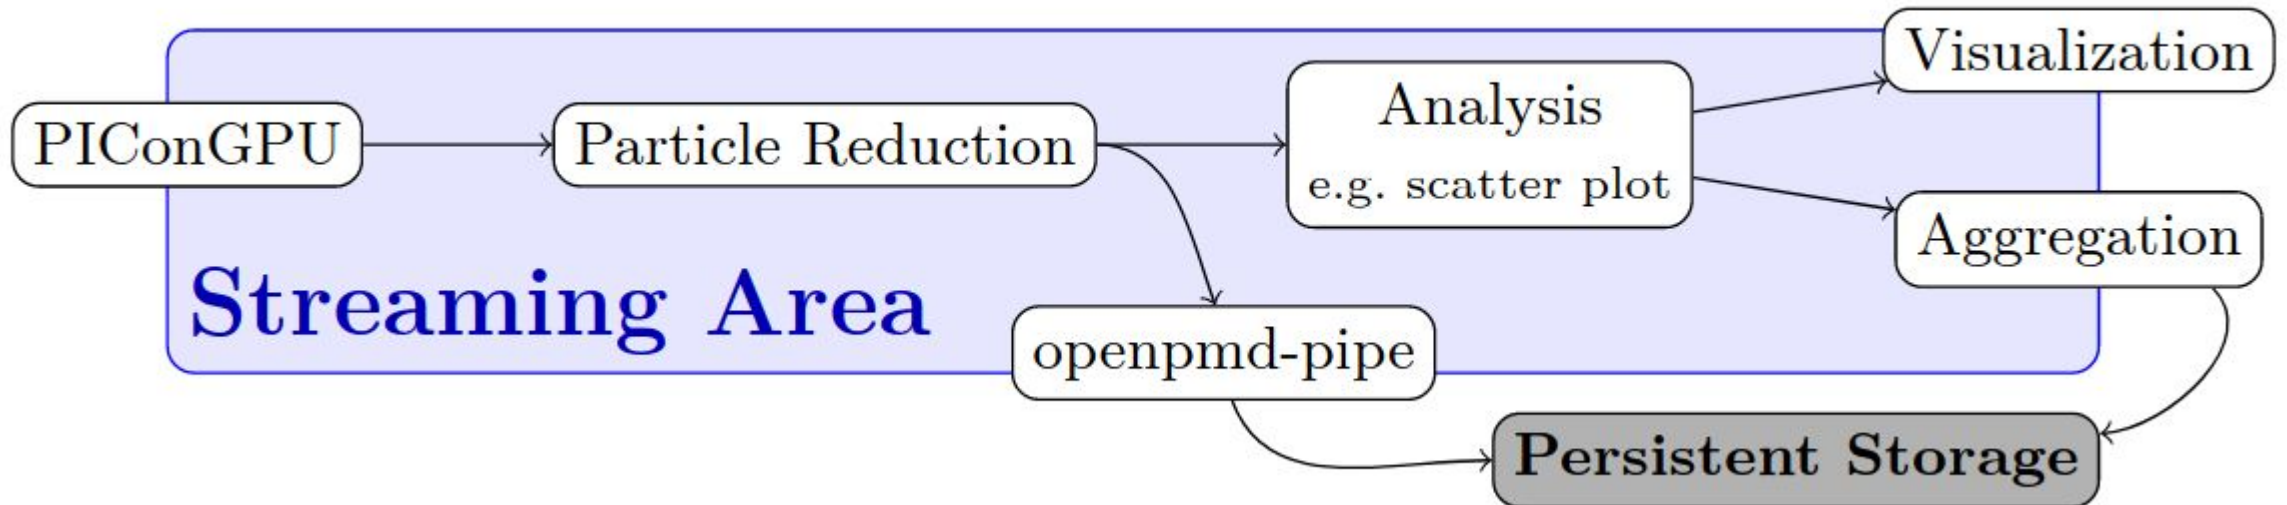
- Field type panel:** Contains a "Field:" section with buttons for B, E, J, and rho, and a "Coord:" section with buttons for x, y, and z.
- Particle quantities panel:** Features a dropdown menu set to "Hydrogen1+" and a grid of buttons for x, y, z, ux, uy, uz, w, and None.

Both panels include a "Plotting options" section and "Always refresh" / "Refresh now!" buttons at the bottom.

Check out the tutorial notebooks: github.com/openPMD/openPMD-viewer

Vision: Loosely coupled data processing pipeline

Loose coupling: Cooperate between independent applications, exchanging data
Streaming I/O between application bypasses PFS bottleneck:



provide a uniform, scientific I/O communication layer between coupled applications and data processing stages: *from creation to archival & reuse*