

# The Artificial Scientist - Leveraging In-transit Machine Learning for Plasma Simulations

**Sunita Chandrasekaran**

David L. & Beverly J.C. Mills Career Development Chair  
Associate Professor, Dept. of Computer & Information Sciences  
Co-Director, AI Center of Excellence  
State of Delaware, Vice Chair AI Commission

University of Delaware

[schandra@udel.edu](mailto:schandra@udel.edu)

Feb 19, 2025, MulticoreWorld XII



KIA ORA TĀTOU

GREETINGS ALL

KO YOSEMITE TE MAUNGA

YOSEMITE IS THE MOUNTAIN

KO GANGES TE AWA

GANGES IS THE RIVER

NŌ INDIA AHAU

I AM FROM INDIA


KO CHANDRASEKARAN

TŌKU WHĀNAU

CHANDRASEKARAN IS MY FAMILY

KO SUNITA TŌKU INGOA


MY NAME IS SUNITA



Laser particle accelerator, painted using gouache and acrylic.  
Displayed at SC24 conference "Art of HPC".  
Sending a copy to Schloss Dagstuhl, Germany, as well. ☺




Enabling scientific advancements by migrating real-world applications on novel and powerful hardware



Advancing foundational computer science by building effective compiler tools



Understanding physics simulations with ML-enhanced models





Exploring GenAI/Large Language Models to complement manual software testing



# A Grand Challenge: Extracting meaningful insights from complex data streams



- Simulation or scientific instruments produce complex data at massive scale
- Storage of such data for offline analysis is impractical
- Physical constraints of the file system pose a massive challenge
- Data reduction will simply NOT suffice
- A complete view of full data is NEVER available
- Need solutions for online analysis of data generated at high rate and volume - to extract meaningful information




# Climate & Weather Modeling



- Earth System Models (ESMs) dataset is multi-dimensional, diverse, high-resolution including structured and unstructured.
  - Models can generate TBs of data, with long-term simulations spanning centuries producing PBs of data.
- Destination Earth project – Digital Twin
- ECMWF, ICON, MPAS-A, MPAS-O

# CERN's Large Hadron Collider (LHC)




- CERN's Large Hadron Collider (LHC) detectors generate over 1 billion collisions per second, with only a fraction of them being recorded and analyzed due to the sheer volume of data. This still results in TBs of data per day.
- Discovery of the Higgs boson, the LHC produced around 30 PB of data per year from collisions

# Plasma Physics – governed by the physics of charged particles




Diverse data properties




Microscopic particle interactions  
to macroscopic magnetic field



Every simulation tracks millions  
to billions of particles




Plasma Instabilities



Credit: Felix Meyer (former HZDR, now NVIDIA)  
Real-Time Vector Field Visualization test using  
HZDR Hemera Cluster with 4 NVIDIA V100.

# Plasma Instabilities – Implications



Sudden violent reaction in fusion research



Solar flares and enormous quantities of radiation




Geomagnetic Storms and cosmic rays



Magnetic reconnection

# What do we want to learn?

- Complete reconstruction of phase space from observations to get a detailed view of the growth and dynamics of **instabilities**
- Automatic detection of correlations between plasma dynamics and emitted radiation spectrum
  - Otherwise this requires extensive post-processing and analysis spanning years
- Extracting meaningful info out of complex heterogeneous data being generated from source or simulation at an exponential rate




Credit: Rene Widera, HZDR

Real-Time Vector Field Visualization test using HZDR Hemera Cluster with 4 NVIDIA V100.

**PIC**on **GPU** 

# The Plasma-In-Cell on GPU (PIConGPU)team



*S. Chandrasekaran<sup>2,3</sup>, A. Debus<sup>1</sup>, T. Kluge<sup>1</sup>, R. Widera<sup>1</sup>, K. Steiniger<sup>1</sup>, M. Garten<sup>4</sup>, M. Werner, J. Kelling<sup>1</sup>, R. Pausch<sup>1</sup>, B. Hernandez<sup>7</sup>, F. Meyer<sup>7</sup>, V. Gutta<sup>3</sup>, F. Mora<sup>3</sup>, F. Pöschel<sup>1,2</sup>, J. Young<sup>2,5</sup>, B. Worpitz, A. Huebl<sup>4</sup>, D. Rogers<sup>6</sup>, G. Juckeland<sup>1</sup>, M. Bussmann<sup>1,2</sup>*



<sup>1</sup> Helmholtz-Zentrum Dresden-Rossendorf, Dresden, Germany

<sup>2</sup> CASUS, Center for Advanced Systems Understanding, Goerlitz, Germany

<sup>3</sup> University of Delaware, Newark, Delaware, USA

<sup>4</sup> Lawrence Berkeley National Laboratories, Berkeley, CA, USA

<sup>5</sup> Georgia Institute of Technology, Atlanta, GA, USA

<sup>6</sup> Oak Ridge National Laboratory, Knoxville, TN, USA

<sup>7</sup> NVIDIA

*Thank you very much HPE Cray/AMD Center of Excellence (COE) for your tremendous hardware/software support!*

*This research partially used resources of the Oak Ridge Leadership Computing Facility (OLCF) at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.*


PIConGPU is funded by the Plasma-PEPSC EuroHPC Center of Excellence by the European Union through Grant Agreement No. 101093261.



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

*We would like to acknowledge the Gauss Centre for Supercomputing e.V. ([www.gauss-centre.eu](http://www.gauss-centre.eu)) for funding this project by providing computing time through the John von Neumann Institute for Computing (NIC) on<sub>9</sub> the GCS Supercomputer JUWELS at Jülich Supercomputing Centre (JSC).*

# Particle In Cell on GPU (PIConGPU) Laser WakeField Acceleration (LWFA)



Acceleration of charged particles within plasmas

# PIConGPU applicability




Compact table top X-Ray sources of high brightness, e.g. Free-Electron Lasers to create snapshots of ultrafast processes in material science





Extend plasma-based electron accelerators from multi-GeV towards TeV electron energies



Applications in radiation therapy of cancer.




Fundamental studies of warm-dense matter and high energy density physics



© Huebl (HZDR), Matheson (ORNL)

# ORNL's Center for Accelerated Application Readiness (CAAR)

- *To stress test Frontier's hardware & software stack*




ACK: Felix Meyer (NVIDIA, former HZDR), Richard Pausch, HZDR  
Still image from an uncut LWFA simulation video using OLCF Summit and 48 NVIDIA V100s using ISAAC 1.5 (in-situ library)

$\sim >= 4 \times$



**Vs Summit  
at ORNL**



OLCF Frontier's AMD EPYC  
CPU + AMD Radeon Instinct  
MI250x GPU

12

$$\text{FOM} = \frac{(90\% \times \text{particle updates} + 10\% \times \text{cell updates})}{\text{second}}$$


# PIConGPU targets....

- Since ORNL's TITAN supercomputer NVIDIA's K20s GPUs
- ORNL's Summit supercomputer NVIDIA's V100 GPUs
- LBNL's Perlmutter NVIDIA A100 GPUs
- JSC's JEWELS Booster NVIDIA A100 GPUs
- NVIDIA's H100 GPUs
- Julich's H200 GPUs
- **Frontier's AMD MI250x GPUs**
- Ampere computing Altra Q80 6—bit CPUs (based on Arm Neoverse N1)


# ORNL's Frontier supercomputer



- **Compute Node 1** 64-core AMD “Optimized 3rd Gen EPYC”  
CPU 4 AMD Instinct MI250X GPUs = **606,208 cores**
- **GPU Architecture** AMD Instinct MI250X GPUs, each feature 2 Graphics Compute Dies (GCDs) for a total of 8 GCDs per node = **37,888 Instinct GPUs**
- **System Interconnect** 4-port HPE Slingshot 200 Gbps (25 GB/s) NICs providing a node-injection bandwidth of 800 Gbps (100 GB/s)
- **Storage** 700 PB HDD+11 PB Flash Performance Tier, 9.4 TB/s and 10 PB Metadata Flash Lustre
- **System Size** ~9400 nodes
- **Ranking** No. 1 in the Top500 as of June 2024




# PIConGPU Exascale challenges



ACK: Benjamin Hernandez, ORNL

LWFA Simulation. Using Summit's 8 nodes (48 V100 GPUs) with ~2 billion particles using ISAAC v1.5.1 running on OLCF's cloud environment (SLATE)




**Portability:** Run code on different compute architectures (single-source, run everywhere)



**Performance:** Cannot lose performance while maintaining portability



**Scalability:** Code profiling & scaling tests to ensure science cases scale to Frontier




**Visualizations:** Create and develop tools to visualize PIConGPU on the new system




**Exascale workflows:** Extend I/O capabilities, provide in-situ analysis, data reduction and visualization workflows

# alpaka software stack



# FOM baseline run on OLCF Summit (2019) TWEAC case study (Single Precision)




$$FOM = \frac{(90\% \times \text{particle updates} + 10\% \times \text{cell updates})}{\text{second}}$$

Nº timesteps	1000
Nº NVIDIA V100 GPUs	27600 (4600 nodes)
Nº cells total	404 billion
Nº cells per GPU	14.6 million
Nº particles total	10.1 trillion
Nº particles per GPU	365 million
Nº simulation data	324 TB

Particle Data	313.36 TB
Cell Data	14.52 TB
Particles Processed	16.2 Trillion particles/sec
Cells Processed	656 Billion cells/sec
GPU Kernel Calls	9 Million kernels/sec


# Major Improvements to PIConGPU since Summit run (2019)

- Algorithmic improvements
  - Optimized laser functor [TWTsfast]
  - New field background algorithm [SuperPusher]
  - New laser algorithm [IncidentField] – 2 years' worth work
- Performance optimizations
  - GPU-aware MPI
  - Optimized particle assignment
  - Enhance device utilization




No. of commits: Autumn 2019: cupla 136, Alpaka 1057, PIConGPU 1278, mallocMC 93  
Red Queen's race – staying in the same place is falling behind

Hardware is only as good as its software and tools and a close  
COMMUNICATION between developers & users!



# Exascale FOM runs for TWEAC case study




Note, these results use single precision.

8.5x for double precision

■ Summit'19 4600 nodes baseline  
■ Frontier'22 8704 nodes

■ Summit'22 projected  
■ Frontier'22 9216 nodes

# FOM run on Frontier (2023) TWEAC case study (Single Precision)



PIConGPU on 98% of Frontier  
Peak Power: **23 MW**  
Average Power: **21.5 MW**

Peak FOM: **78.3 TeraUpdates / s**  
Average FOM: **65.7 TeraUpdates / s**


$$\text{FOM} = \frac{(90\% \times \text{particle updates} + 10\% \times \text{cell updates})}{\text{second}}$$

<b>Nº timesteps</b>	1000
<b>Nº AMD GCDs</b>	73,728 (9216 nodes)
<b>Nº cells total</b>	1.1 Trillion
<b>Nº cells per GCD</b>	14.6 Million
<b>Nº particles total</b>	27 trillion
<b>Nº particles per GCD</b>	365 million

<b>Particle Data</b>	760.7 TB	41% more
<b>Cell Data</b>	35.3 TB	41% more
<b>Particles Processed</b>	72 Trillion/sec	22.5% more
<b>Cells Processed</b>	2.6 Trillion/sec	25% more
<b>GPU Kernel Calls</b>	24 Million kernels/sec	37% more


# Weak Scaling FOM case on Frontier

- Nº Iterations: 1000
- Runtime: ~10 mins  
~ 0.37 secs per iteration
- FOM Science case
- Scaling:
  - 6 nodes → 9,216 nodes
  - 48 GCDs → 73,728 GCDs
  - 24 GPUs → 36,864 GPUs
  - 96-98% GPU utilization



# There is more to do....


- With Frontier we can **ONLY** perform a few long-running simulations where we will hopefully observe acceleration of electrons to highest 100GeV scale energies, **BUT** the parameter space of laser electron acceleration is **HUGE**; we are still just able to catch a **TINY** bit of it.
- Need to explore **LARGER** parameter range, need a **sophisticated WORKFLOW** to further advance science
- Need to close gap between simulation scenarios on supercomputers and the experimental setup in the labs



ACK: Vincent Gerber, HZDR, Germany  
Using **In-Situ** viz library for **Animation of Accelerated Computations (ISAAC)**


# PIConGPU data volume

- Forward calculation incurs a heavy computational cost prohibitively expensive for a wide range of simulation parameters.
- Simulations of plasma behavior, involve solving complex nonlinear equations for trillions of particles creating **TBs** of data
- Due to the scale of the data, we cannot save all the raw data to disk
- We need a different solution!




# What do we want to learn?

- Complete reconstruction of phase space from observations to get a detailed view of the growth and dynamics of **instabilities**
- Automatic detection of correlations between plasma dynamics and emitted radiation spectrum
  - Otherwise this requires extensive post-processing and analysis spanning years




# Extracting knowledge from large scale simulations is a challenge!!



Particle data streaming



Learning



Learning Outcomes

Jeffrey Kelling.....Michael Bussmann, Sunita Chandrasekaran, "The Artificial Scientist - Leveraging In-transit Machine Learning for Plasma Simulations" Accepted to 39th IEEE International Parallel & Distributed Processing Symposium (IPDPS) 2025, Best Paper Nomination

# Case Study: Kelvin-Helmholtz instability (KHI) in PIConGPU

- Well known shear surface instability observed in fluids and plasma
- When 2 layers exhibit different velocity/density
- Viz using ISAAC



Richard Pausch, HZDR, Germany; Uses 4 V100 GPUs

*In plasmas, the KHI is driven by a self-amplifying cycle of small density or velocity fluctuations that lead to a growing magnetic field at the shear surface, which further amplifies the initial fluctuations as depicted*

# openPMD and ADIOS2

- Particle data is streamed from the simulation (PIConGPU) using a custom implementation of openPMD thus avoiding file system limitations
  - Data standard for **particle mesh data**
- ADIOS2 allows implementations of TCP (non-scalable fallback), libfabric, ucx and the `MPI_Open_port()` API of MPI
- Combination of openPMD-api and SST data engine of ADIOS2 allows direct in-memory transport



Poeschel, Franz, et al. "Transitioning from file-based HPC workflows to streaming data pipelines with openPMD and ADIOS2." *Smoky Mountains Computational Sciences and Engineering Conference*. Cham: Springer International Publishing, 2021.


Huebl, A., Lehe, R., Vay, J. L., Grote, D. P., Sbalzarini, I., Kuschel, S., ... & Bussmann, M. (2015). openPMD: A meta data standard for particle and mesh based data. URL <https://doi.org/10.5281/zenodo.591699>.

# Challenges - System Constraints



- **20 to 30 TB/s** parallel throughput of particle data @ 5.86GB/node \* no. of nodes
- BUT only **10 TB/s** bandwidth of Orion file system on Frontier
- PICoGPU scaling to just 25% of Frontier system would produce 1 PByte of data every time step
  - 1000 time steps with each lasting 0.1s to 1s
  - We often see 1 to 10 Petabytes/s of particle data requiring disk volumes on the order of 10 Exabyte
- **Implication:** Entire file system can be exhausted in 100 to 1000 seconds

We need to circumvent the file system!!!!!!

# Three in-transit workflow approaches on Frontier



Streaming without going through storage unlocks more bandwidth





Reducing simulation data close to the producer lowers bandwidth requirements


Distributed producer and consumer, system topology presents communication paths with vastly different bandwidths which must be reconciled with the loosely-coupled application's communication requirements.

# Data streaming using MPI data plane and libfabric using 9126 Frontier nodes

- The amount of particle data produced by PICoGPU KHI is 5.86 GiB per compute node and time step.
- Parallel throughput (w/o using file system) for 4096 nodes
  - Libfabric per-node achieved per-node throughput of 3.5 ~ 4.7 GB/s
  - MPI data plane yields a per-node throughput from 2.6 ~ 3.7 GB/s
- Maximum parallel throughput of 20 - 30 TB/s which compares outstandingly against the 10 TB/s bandwidth of the parallel Orion filesystem
- The libfabric's lower-level control can bring performance improvements over the more managed MPI implementation




# In-transit continual learning




- Model is continuously trained online from subsequent snapshots of the evolution of plasma and radiation data **without** storing every data point to disk
- Continuous learning circumvents lack of adequate disk capacity and bandwidth by enabling data to reside and distribute in-memory via network interconnects.

# ML model architecture




# Data parallelism and scaling on Frontier

- ML model fits into 1 GCD
- Each copy of the model receives different chunks of data to train on
- Asynchronously train the model with that data
- Scaling depends on the optimization of all-to-all communication in PyTorch DDP
  - using N/RCCL backend – hits a wall after 100 Frontier nodes = 400 AMD GPUs
  - Need to explore libfrabic backend for N/RCCL or PyTorch DDP's MPI backend



# Weak scaling and observations

- Increasing time spent in PyTorch DDP for larger runs
- **Low Efficiency** - Inevitable all-to-all communication between PyTorch ranks taking place to average gradients during each backward pass – deficit of 30%
- **Low Efficiency** - Lack of availability of PyTorch distributed primitive for matrix dot product to evaluate INN
- In-transit training at very large batch size
  - Hyperparameter needs to happen at scale; doesn't transfer well from small scale experiments
- Loss functions to compare point cloud – CD vs EMD
- Comprehensive studies between batch sizes, block learning rate and weights need to be studied at scale with streamed simulation




- In-transit training from 8 to 96 nodes (32 to 384 GCDs)
- Reaches around 35% at 96 nodes


# Take aways

- The model **clearly learned to partition the latent space** into regions for different flow directions and vortex regions, which both the encoder and the inversion network learned to map to.
- Achieved partial reconstruction of the plasma distribution, the ML prediction still clearly identifies the **instability regions**
- A promising **first** result towards large scale in-transit learning for a non-steady state processes
- Ran problem at scale and discovered challenges/opportunities along the way

Automated learning of physical relationships from in-transit data of large-scale simulations and the identification of regions of instability is already a major step forward




# Extracting knowledge from large scale plasma simulations




# So much more to do....

- Build surrogate models of simulations of different configurations – classically these experiments would be very expensive
- Collect many time steps
- Intelligently reduce data
- Learning in-transit valuable information from more than just one simulation
- More sophisticated decoder to generate higher-fidelity depictions of particle configurations
- Encoders incorporating point transformer blocks and a deeper network around the bottleneck to better extract latent information from point-features

# Need for WORKFLOWS to tackle complex, data-intensive real-world applications!








# AI Center of Excellence (2022) leading to AI Institute (2025)

## AI Center of Excellence Seed Funding Program

The center's seed funding program aimed to provide support to various AI projects at UD through funds, Research Software Engineer support, or undergraduate researcher support.



### AI Research Areas of Interest



Financial Equity



Complementing Human Impact



Medicine, Science and Engineering



Build Software Infrastructure

AICOE Faculty

60 Affiliated faculty members

7 Colleges affiliated

21 Different Departments

### 2023 - 80 Participants, 11 case studies

Udel, Kendal Corp, FMC, Chemours, Community Health Options

### 2024 - 74 Participants, 11 case studies from

ATOM Team and  
UDel, Delaware State, Bowie State, Penn State,  
Northeastern, Berkeley, ATOM, Frederick Lab



### Artificial Intelligence Graduate Certificate

9 Credits of grad-level courses

#### Why Learn AI?


This rapidly evolving field is integrated into so many aspect of daily life. Skills in building and training AI systems are highly sought-after across industries.

This certificate can open the door to a wide range of career opportunities and provide a competitive edge to your skillset.


#### Sectors looking for AI experts




Finance



Healthcare



Technology



Manufacturing

Aug 2023

Computer Vision

Data Mining

Machine Learning

Natural Language Processing

Multi-Agent Systems

## Our Mission

Democratize access to research software engineering (RSE) by creating an educational pipeline for RSE, building a team of RSE professional and trainees, and engaging them in the advancement of computational, data-intensive, and AI-enabled science projects in the Mid-Atlantic region and beyond.

### Project Goals




Establish a graduate level course and RSE pilot certificate



Identify and create a sustainable and scalable pipeline of RSEs to accelerate and enable domain sciences, especially in the SBE and CSI domain areas



Connect RSEs from other initiatives such as NSF ACESS, SCiPE and US-RSE to foster a broader network



Explore how data-driven ML/AI methods can be applied for problems in domain science

## What is an RSE?

### Research Software Engineer (RSE)

Someone who combines professional software engineering expertise with an intimate understanding of research.

### Who would be considered an RSE?

- Researchers who spend a lot of time programming
- Software engineers who write code to solve research problems
- Someone who feels they are in between a researcher and a software engineer!

## Current Projects

### Coastal Science

#### Urban and Coastal Flood Modeling

- ML to create classifications of land cover from aerial imagery and LiDAR to assist with flood modeling

#### Sub-Model for Settling Velocities of Cohesive Sediment

- ML for a surrogate model to predict time series of floc size distribution subject to environmental forces

### Social Science

#### Multi-Model Topic Models of Foreign Policy Speech Texts and Photographs

- Large scale collection and storage of text-as-data and image-as-data, and subsequent application and scaling of transformer-based models for text- and multi-model content.



PI



Sunita Chandrasekaran  
Associate Professor  
Computer & Info Sciences



Rudi Eigenmann  
Professor  
Electrical & Computer  
Engineering



Tian-Jian Hsu  
Associate Professor  
Computer & Info  
Science



Benjamin Bagozzi  
Professor  
Political Science &  
Internat. Relations



John Huffman  
Director  
Research Computing

## UD Leadership

