



# Exhibitor Forum Take Over!

---

*May 22, 2023*

# Welcome!



**Frank Herold**  
CEO, BeeGFS



**Deb Goldfarb**  
Director of Product and  
Strategy for HPC & Batch,  
AWS



**Cristin Merritt**  
CMO, Alces Flight

# Agenda - Part One

- **Welcome**
- **SPECIAL ANNOUNCEMENT**
- **WHPC Chapters + Affiliates** - *Cerlane Leong and Sadaf Alam*
- **In Support of Diversity + Inclusion** - *Frank Herold, BeeGFS*
- **Early Career Speakers!**
  - *Carla Santana, Universidade Federal do Rio Grande do Norte*
  - *Stefanie Reuter, University of Cambridge*
  - *Elisabeth Ortega, HPCNow!*
  - *Jenny Wong, University of Birmingham*

# Agenda - Part Two


- **In Support of Diversity + Inclusion** - *Deb Goldfarb, AWS*
- **Early Career Speakers!**
  - *Paula Esposito, La Laguna University*
  - *Aleksandra Kowalczyk, University of Warsaw*
  - *Ekaterina Zossimova, University of Exeter*
  - *Helena Vela Beltran, Barcelona Supercomputing Center*
- **Close!**



# Supporting Institutions - Anchor



# Supporting Institutions - Career




**Hewlett Packard  
Enterprise**



**WHPC**  
WOMEN IN HPC

# Supporting Institutions - Volunteer



# Interested in Sponsoring Us?

---

Talk to us today about how you can fund our global initiatives!



# How to get involved!

## **Join Women in HPC for FREE!**

- Select a Chapter/Affiliate for local events
- Let us know if you would like to get involved:
  - SC23 - Denver, USA
  - ISC 2024

**womeninhpc.org**





# Poster Reception

---

May 23, 2023

Foyer D-G

3:30 - 5:30 PM





Grab your swag!



# Special Announcement

# Chapters and Affiliates

## ISC2023

May 22nd, Hamburg



WHPC works with Chapters and Affiliates to support and promote the work of women in their organizations, develop crucial role models, and assist employers in the recruitment and retention of a diverse and inclusive HPC workforce.

Around 20 Chapters are active today worldwide



7 Chapters are present in Europe, out of which 3 started in 2023:

-JuWinHPC at JSc in Juelich  
contact : Ruth Schoebel and Claire Wyatt

-BSC in Barcelona  
contact: Marta Garcia

-IDEAS4HPC in Switzerland  
not-for-profit association  
contact : see next slide



## Contact

Maria GIRONI, Director CERN OpenLab,  
Vice President

Cerlane LEONG, Senior Computer  
Scientist at ETHZ CSCS, Treasurer

Prof. Florina CIORBA, Head of HPC group  
of University of Basel, Secretary

Marie-Christine SAWLEY, HPC Advisor  
ICES Foundation Geneva, President



<https://elle-it.ch>




womeninhpc.org

**Women in High Performance Computing (WHPC)**

Stewarded by EPCC at the University of Edinburgh, UK

The University of Edinburgh is a charitable body, registered in Scotland  
with registration number SC005336





<https://www.jlab.org/conference/CHEP2023>



**exhibitor forum take over**  
**22 May 2023**

**PA  
SC 23**  
**Davos** | **26-28 June 2023**  
**Switzerland**

**Mini-Symposium:**  
**Breaking the Silos to Enhance HPC Impact**  
**27 June 2023, 16:00—18:00**



[womeninhpc.org](http://womeninhpc.org)

**Women in High Performance Computing (WHPC)**  
Stewarded by EPCC at the University of Edinburgh, UK  
The University of Edinburgh is a charitable body, registered in Scotland  
with registration number SC005336




**Frank Herold**  
CEO, BeeGFS

# Early Career Speakers



**Carla Santana**

Universidade Federal do  
Rio Grande do Norte



**Stefanie Reuter**

University of Cambridge



**Elisabeth Ortega**

HPCNow!



**Jenny Wong**

University of Birmingham

# DeLIA: Dependability Library for Iterative Applications

Carla dos Santos Santana, Idalmis Milian Sardina,  
Hervé Chauris, Claude Tadonki, Samuel Xavier de Souza





# Summary

- Motivation
- DeLIA Features
- Case Study: 3D Full Waveform Inversion
- DeLIA Usability
- Next steps





# Motivation

- One of the primary purposes of an HPC implementation is **scalability**.
- However, with more nodes, the probability of failure is higher.
- An execution of HPC application can be interrupted
  - Some failure in the HPC environment (e.g. network failures).
  - Preemptive circumstances (e.g. cloud instances preemptives).
- Therefore, dealing with faults is crucial for a scalable HPC application [4].





# Motivation

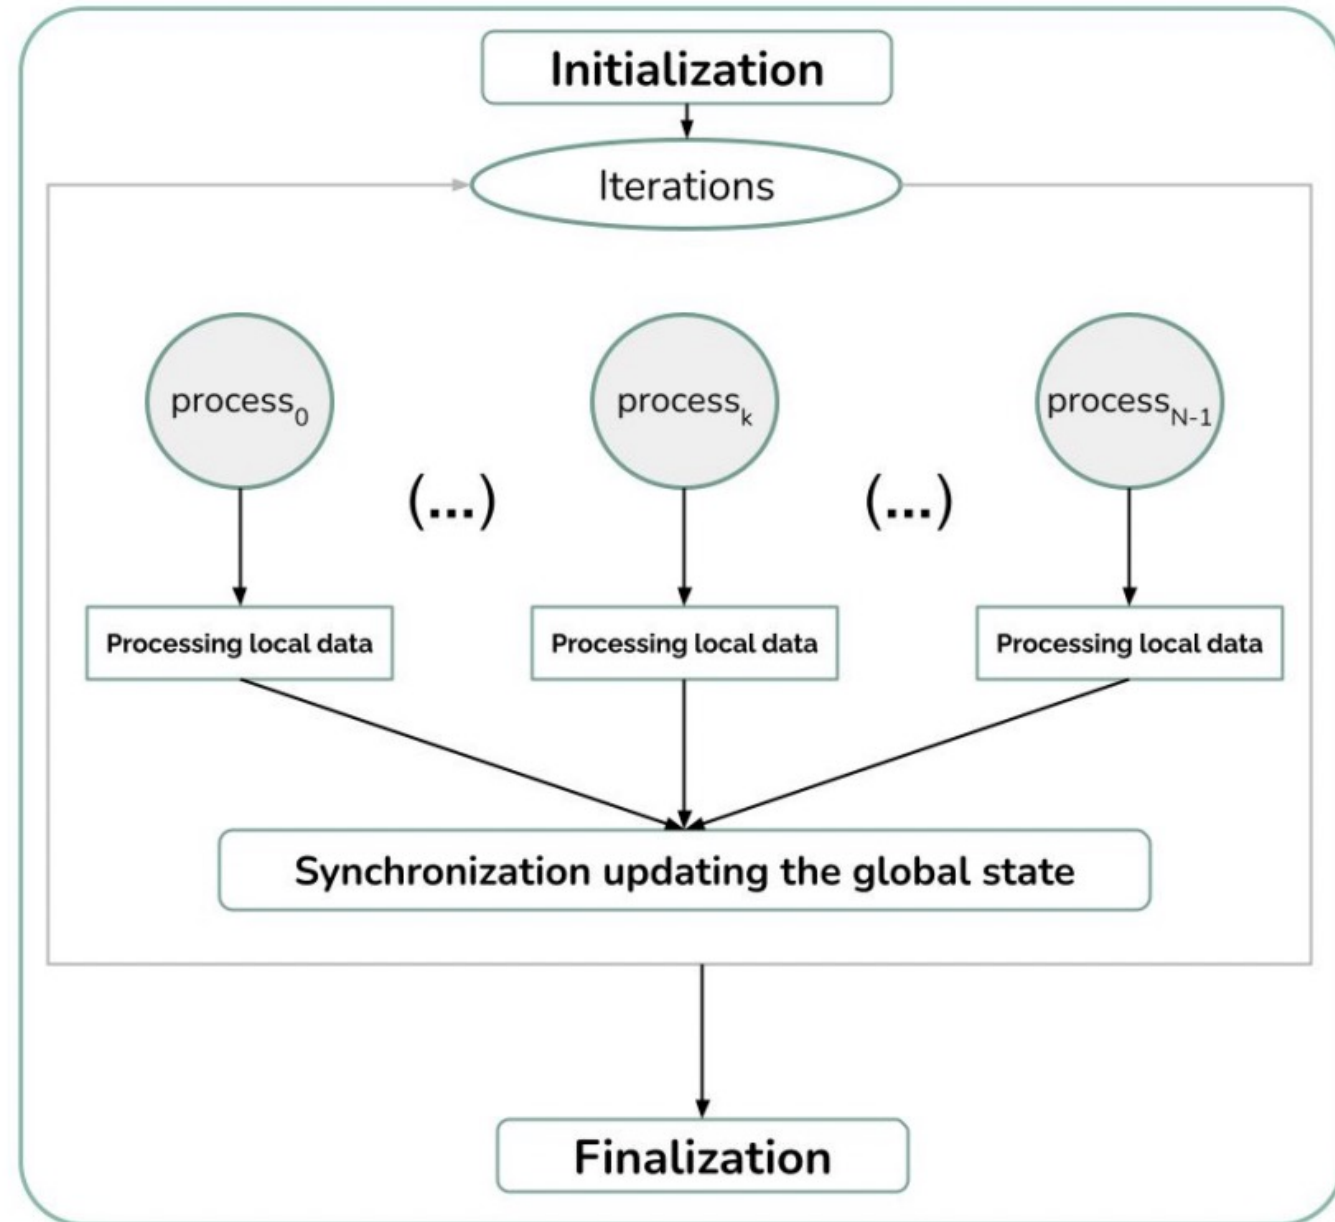
Developing fault tolerance techniques that  
**do not cause significant overheads.**

- HPC Applications
  - Many data
  - Many calculations
  - Spends a lot time to be complete
  - Fast



# Application

## Application Behavior





# DeLIA Features

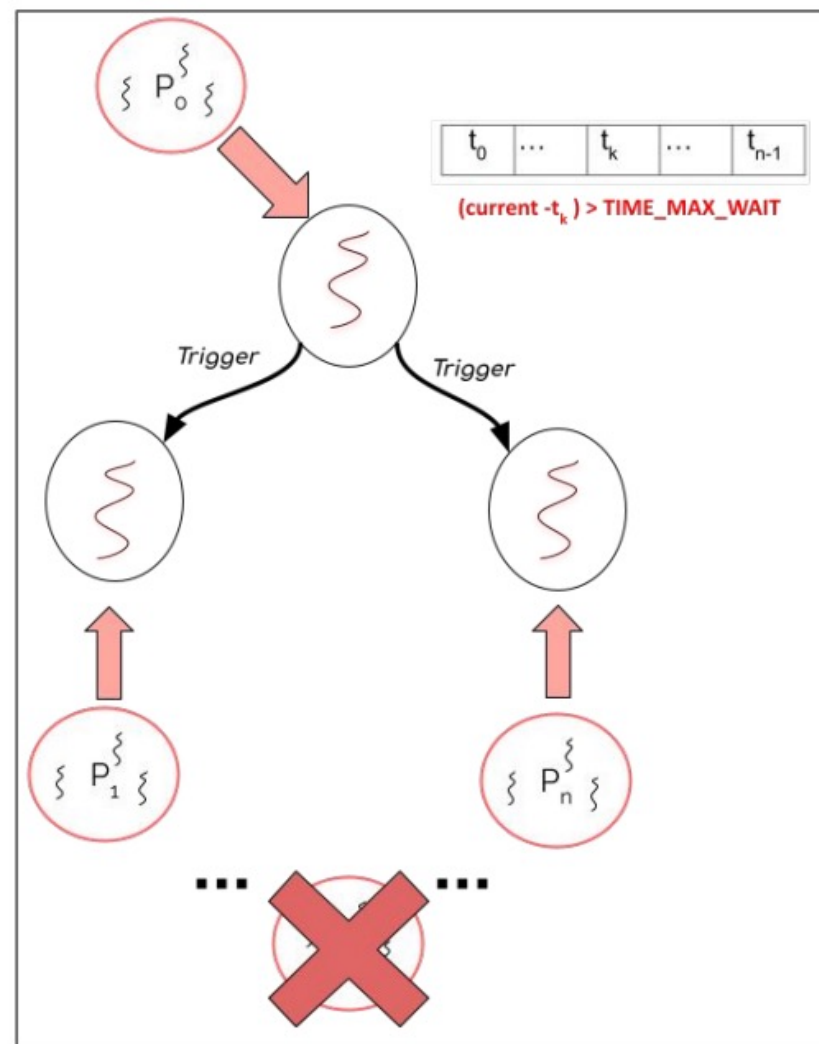

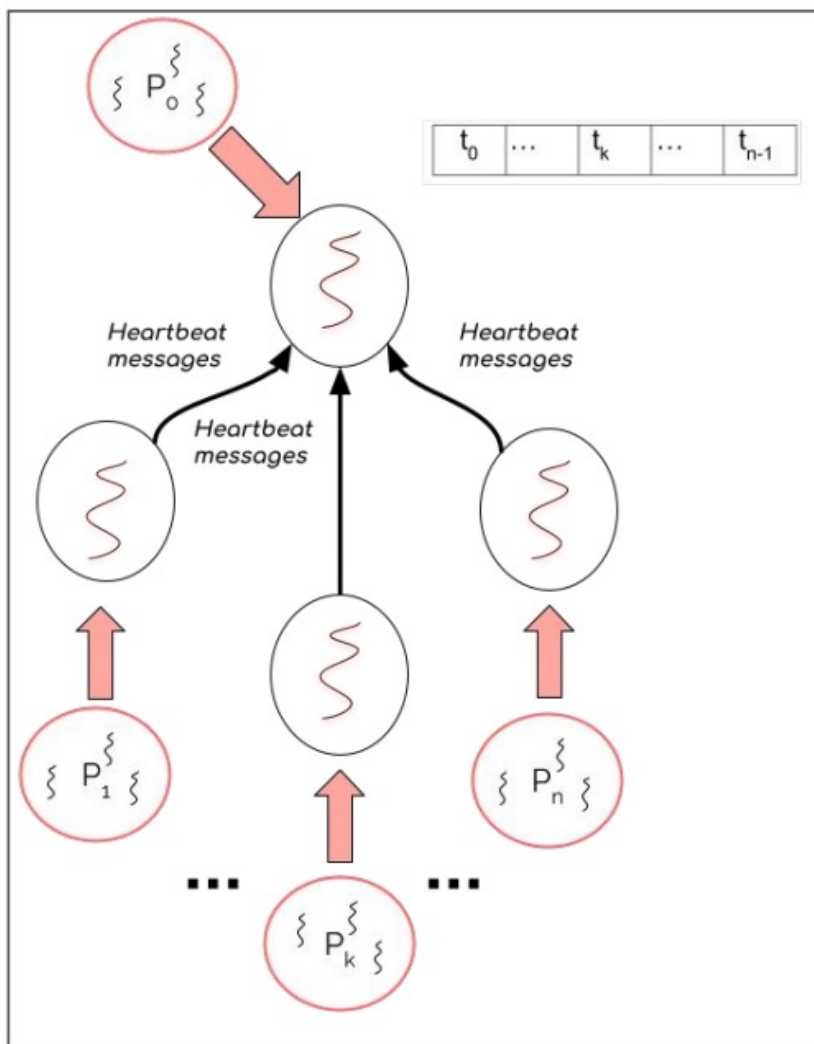
## Interruption Detection

- Fault detection with the method of Heartbeat Monitoring
- Detection of termination signals.

If **there is a possibility of interruption**,  
a trigger is sent to the nodes, and each one saves its local data

# DeLIA Features

## Heartbeat Monitoring





# DeLIA Features

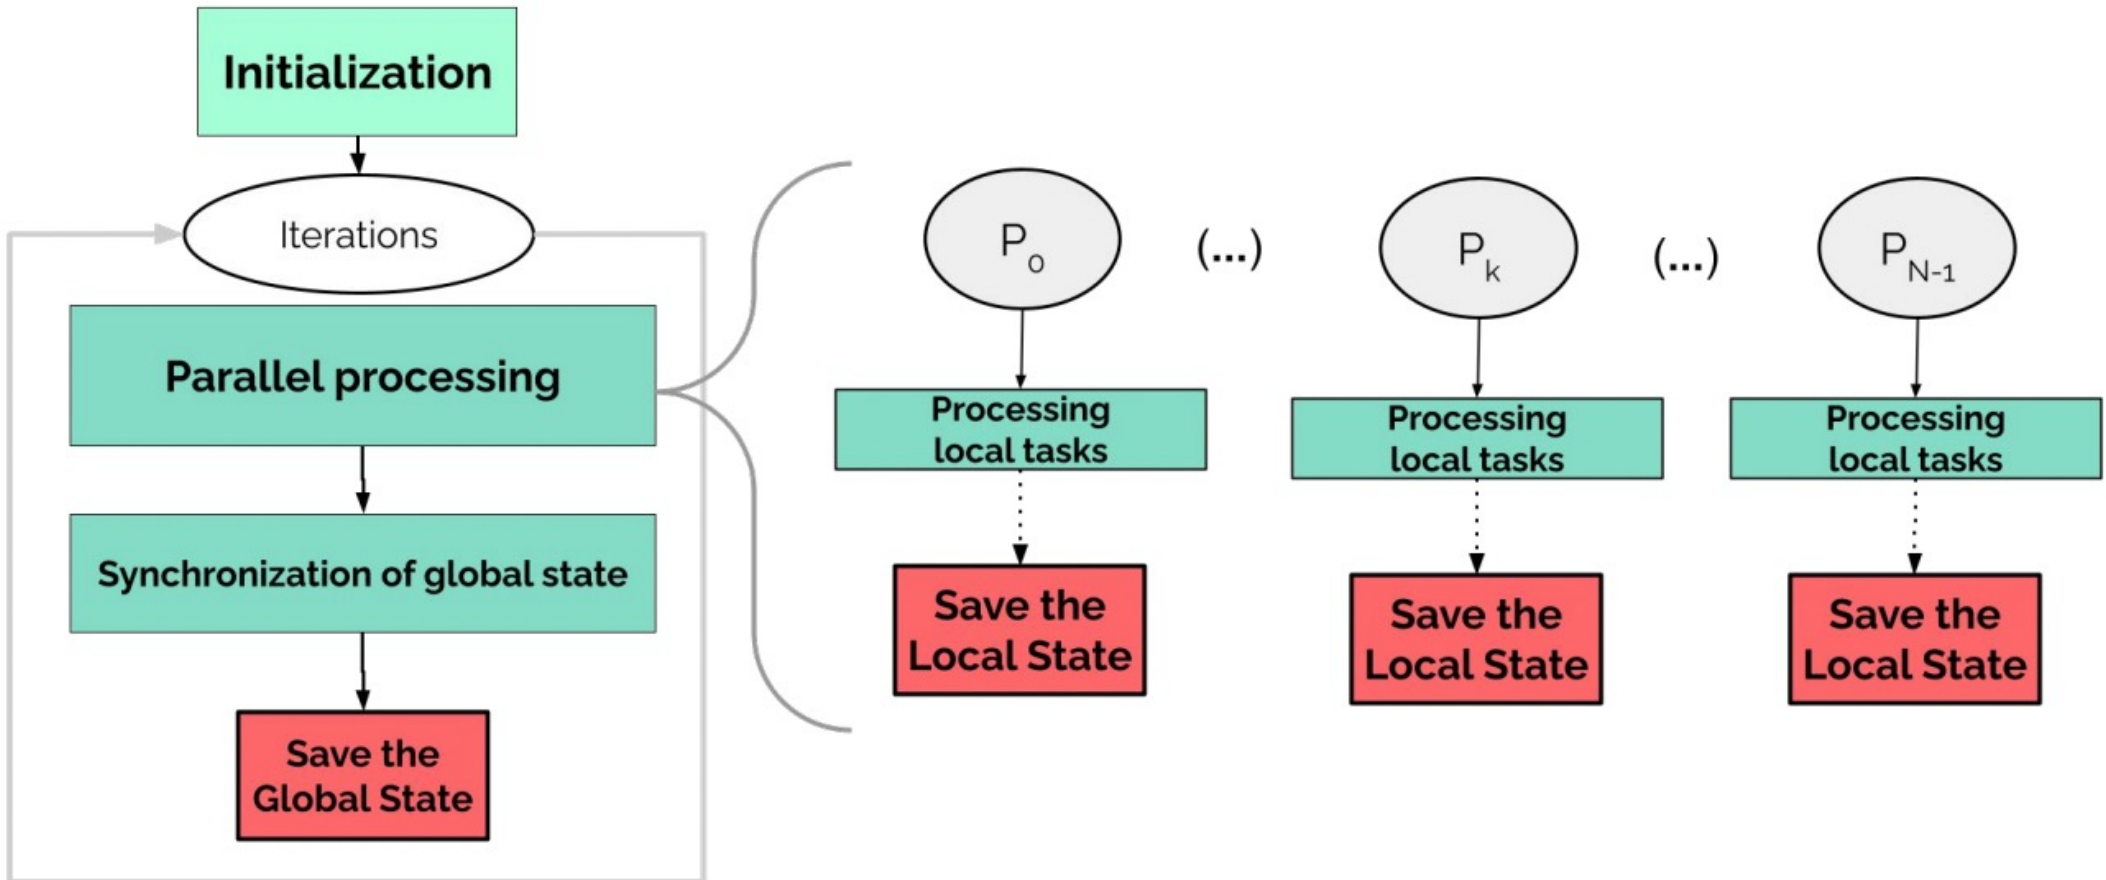
## Detection of Termination Signals

- Interruption Signal are applied to advise a process that will probably terminate at some moment.
- Some supercomputers and cloud systems use them to notify a job that it will finish for some reason.
- In DeLIA:
  - When an termination signal is detected a trigger is send to the process.

# DeLIA Features

## Checkpointing and Rollback

Checkpointing is the technique to save the data in a determined state. This can be used to resume the process later [5].

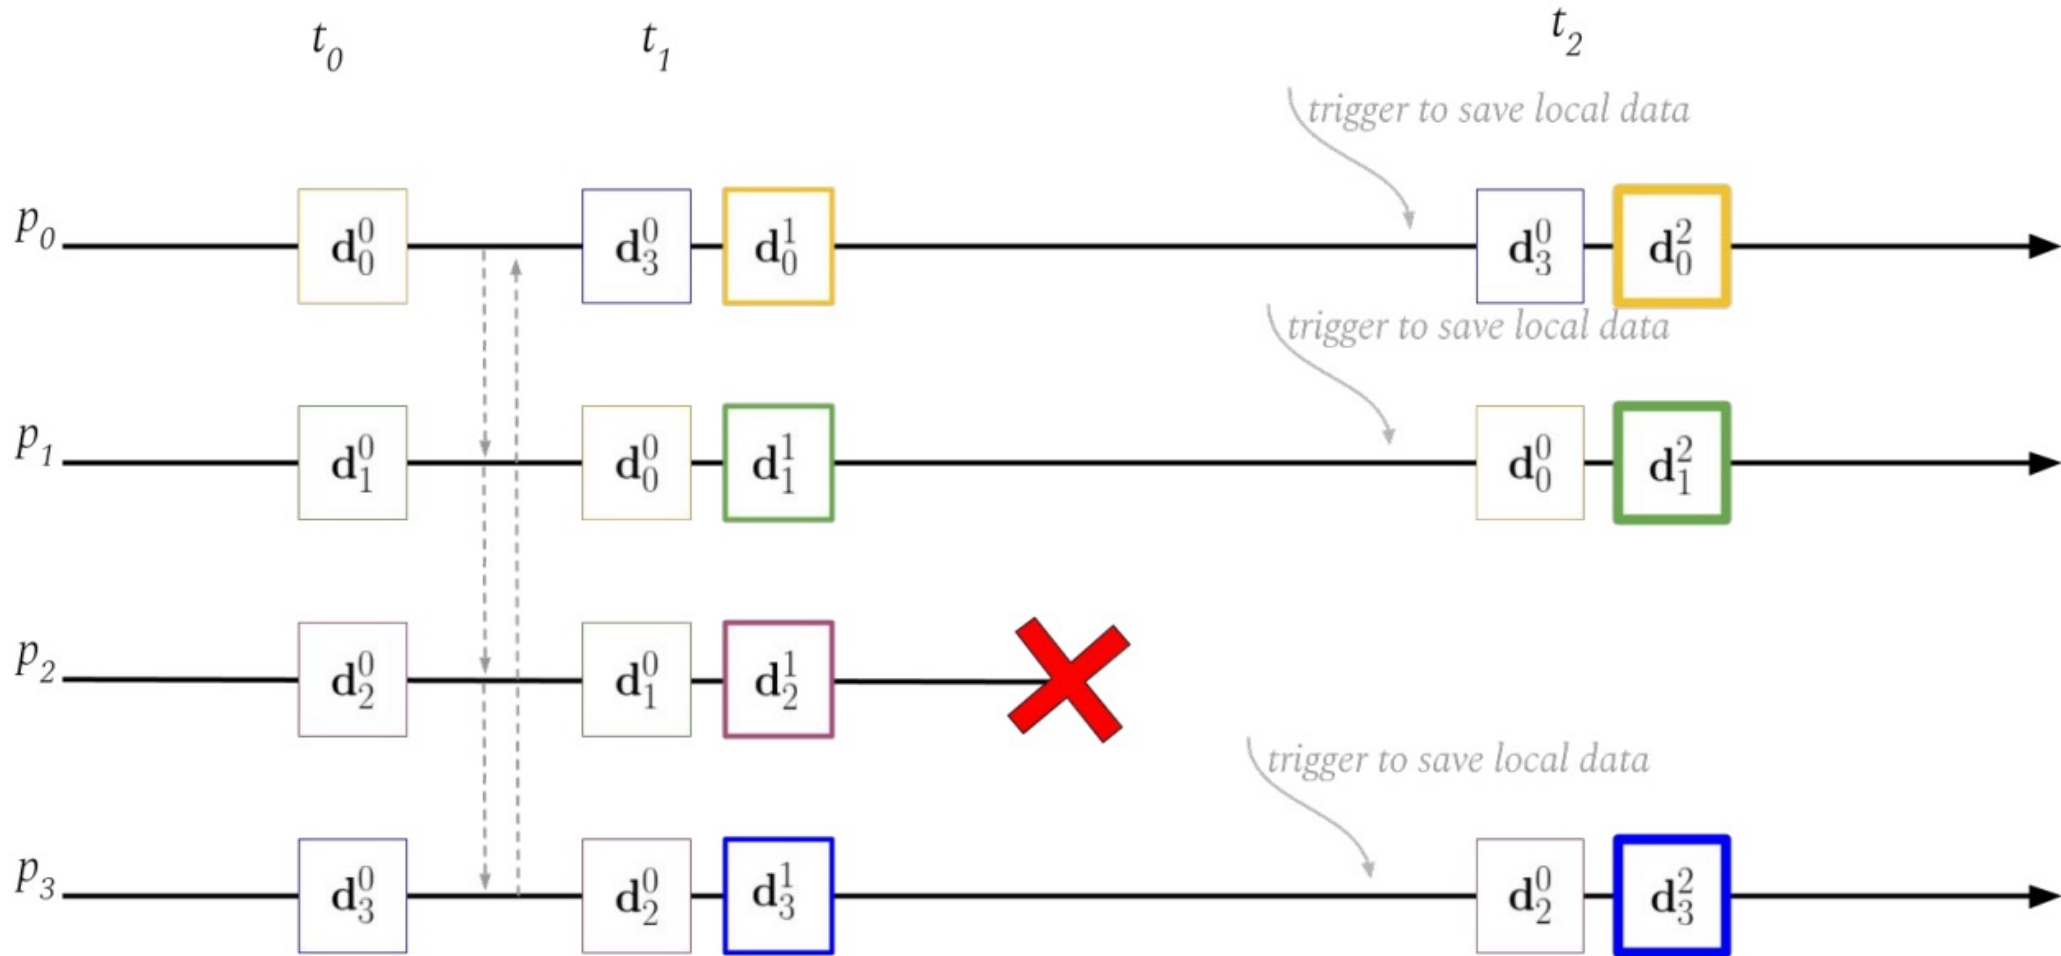




# DeLIA Features

## Replication

Replication of data is the most typical means to provide high availability and fault tolerance in distributed systems [2]



# Case Study: 3D Full Waveform Inversion

## Overview




Figure from [3]

# Case Study: 3D Full Waveform Inversion Overview

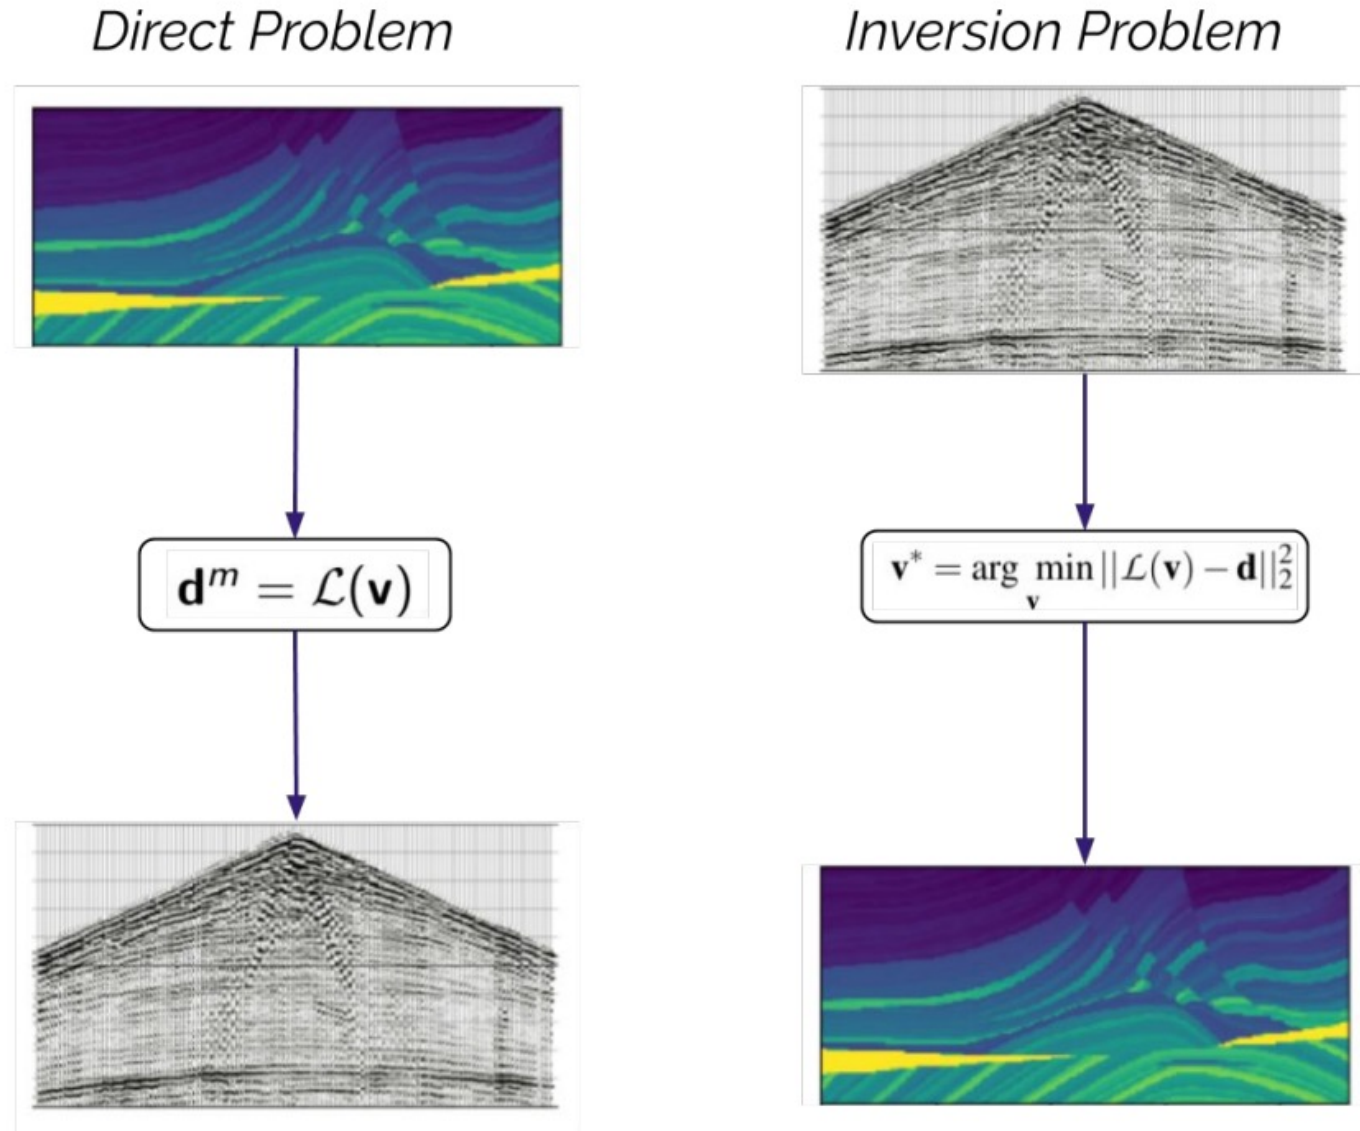
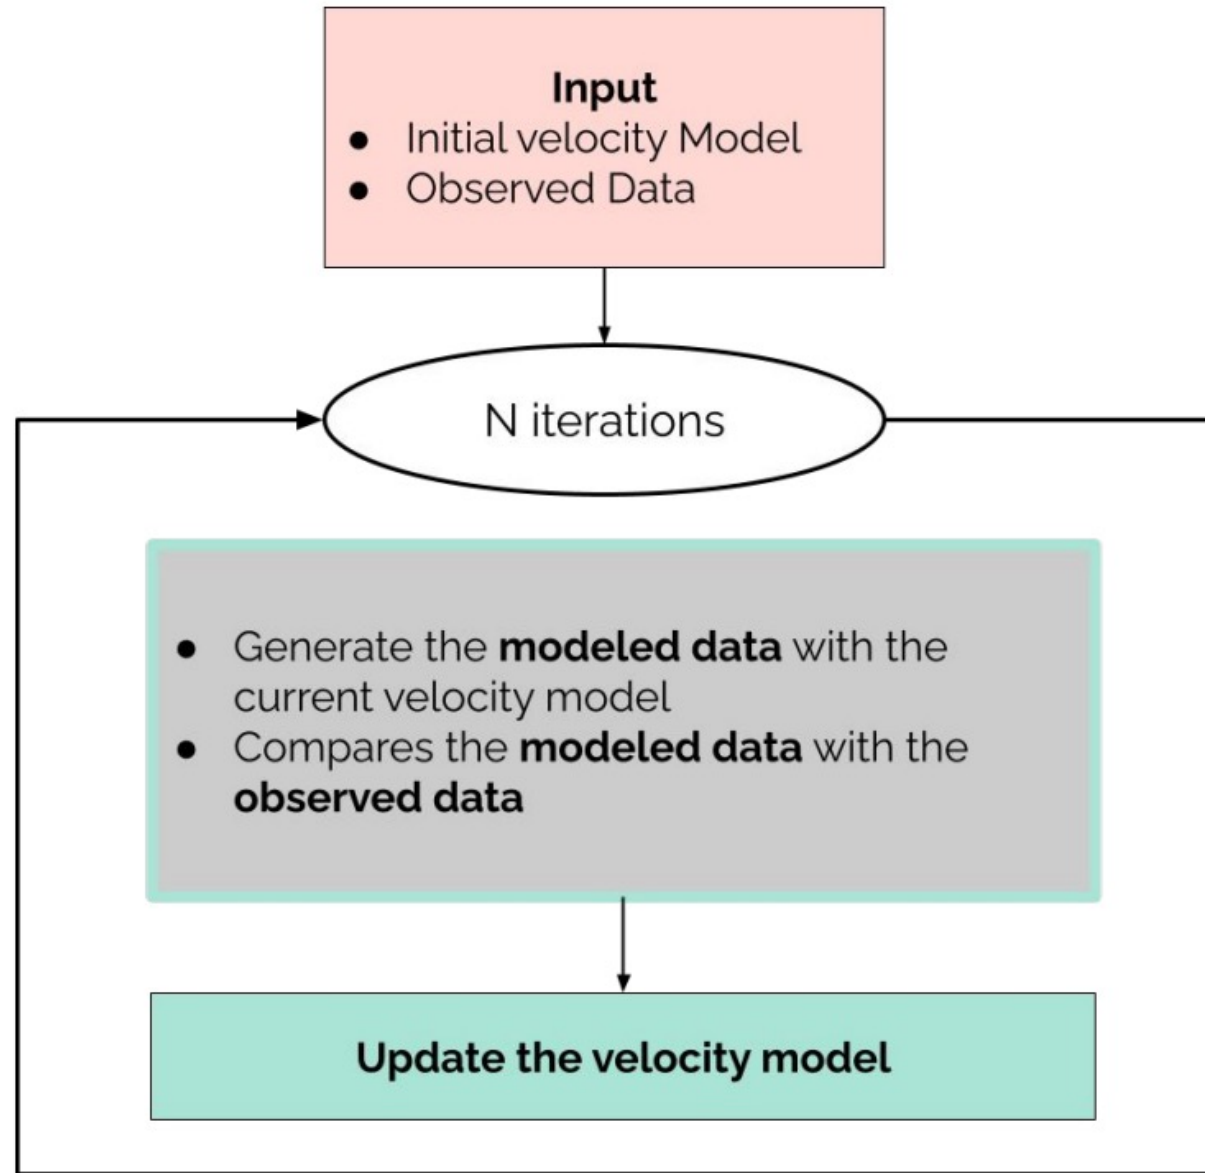



Figure from [7]

# Case Study: 3D Full Waveform Inversion Overview




# DeLIA in 3D FWI

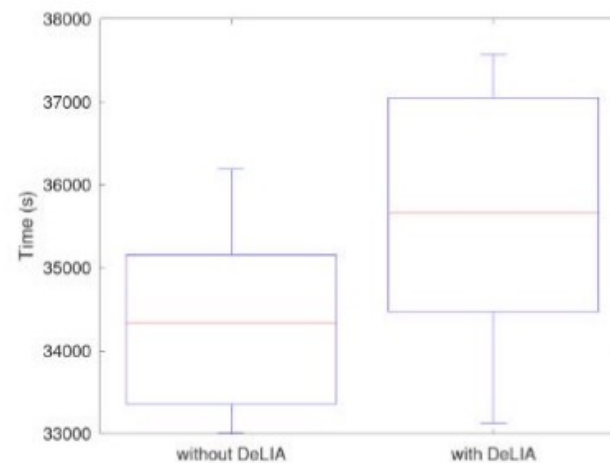
## Experiments



Using four nodes 8 shots (a)



Using eight nodes 8 shots (b)



Using eight nodes 32 shots (b)

	4 nodes 8 shots	8 nodes 8 shots	8 nodes 32 shots
DeLIA overhead	2.54%	8.8%	3.84%
RSD without DeLIA	6.88%	3.37%	3.78%
RSD with DeLIA	17.40%	7.20%	5.14%





# DeLIA Usability

- DeLIA provides an API to programmers to use the features in their software;
- The main parameters for DeLIA are defined by the developer in a JSON file.
- The library and its documentation are available at <https://lappsufrn.gitlab.io/delia>.

```
{
  "FT_FOLDER" : "../projects/fault_tolerance/data",
  "CHECKPOINTING_GLOBAL_ITERATION": 1,
  "TRIGGER_SIGNAL" : true,
  "CHECKPOINTING_TIME_GLOBAL": 60,
  "CHECKPOINTING_TIME_LOCAL" : 10,
  "TRIGGER_HEARTBEAT_MONITORING": {
    "TIME_MAX_WAIT" : 60,
    "SLEEP_THREAD_TIME" : 5
  }
}
```





## Next Steps

- Checkpoint with compressed data.
- The process shares encoded pieces of local data with more than one process. If one piece is lost, the other can reconstruct all data using the Reed-Solomon technique [6].
- Portability to other languages (e.g., python).



**Thank you  
for your time**



# References

- [1] Chetan S, A. Ranganathan, and R. Campbell. Towards fault tolerance pervasive computing. IEEE Technology and Society Magazine, 24(1):38-44, 2005.
- [2] Coulouris, George F., Jean Dollimore, and Tim Kindberg. Distributed systems: concepts and design. pearson education, 2005.
- [3] de Haan, Dick, et al. Potential effects of seismic surveys on harbour porpoises. No. C126/15. IMARES, 2015.
- [4] Herault T. and Robert Y. Fault-tolerance techniques for high-performance computing. Springer, 2015.
- [5] Kalaiselvi S. and Rajaraman V.. A survey of checkpointing algorithms for parallel and distributed computers. Sadhana, 25(5):489-510, 2000.
- [6] Reed, Daniel A., and Celso L. Mendes. "Reliability challenges in large systems." Future Generation Computer Systems 22.3 (2006): 293-302.
- [7] Silva, Suzane Adrielly. Análise qualitativa do método de inversão completa das formas de onda no domínio do tempo. BS thesis. Universidade Federal do Rio Grande do Norte, 2017.
- [8] Stephen B. Wicker and Vijay K. Bhargava. Reed-Solomon codes and their applications. John Wiley & Sons, 1999.

# Fairy WANDS for Fusion

Stefanie Reuter<sup>1</sup>, Jonathan Hollocombe<sup>2</sup>, Toby James<sup>2</sup>

<sup>1</sup>Research Computing Services, UIS, University of Cambridge, UK

<sup>2</sup>UK Atomic Energy Authority

Women in HPC Poster Competition

# Fairy WANDS for Fusion



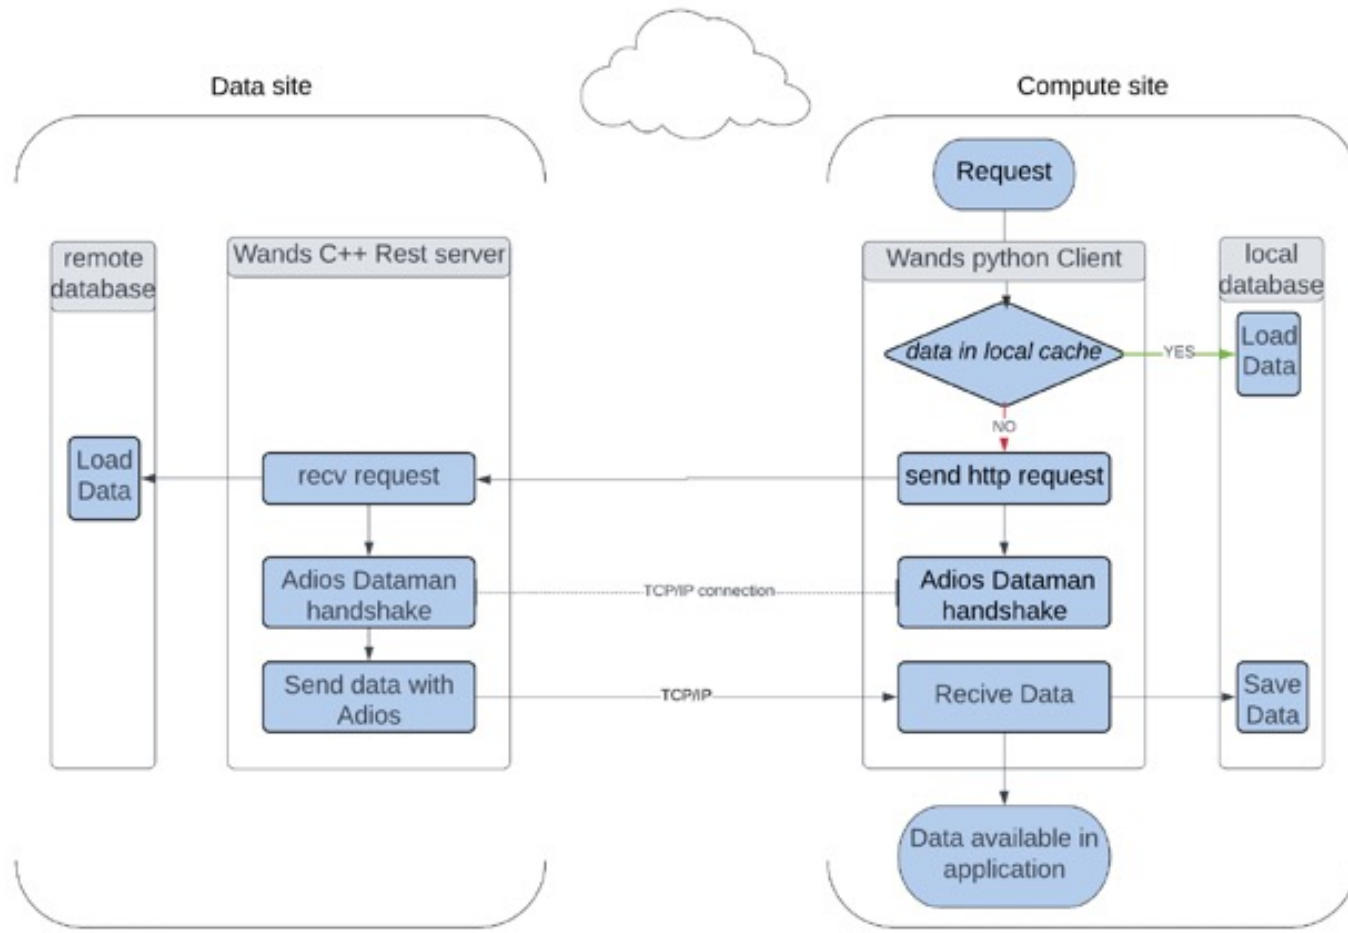
Illustration of JET fusion tokamak reactor  
(Credit: EUROfusion)



Cambridge Service for Data Driven Discovery  
(Credit: Joe Bishop)



# Wide Area Network Data Streaming



## Python example usage

```
from wands import Wands

filename = "exampleFile.h5"
signals = ["dataset_1",
           "dataset_2"]
local_database = "/path/to/local/database"

wo = Wands(local_database, Port="12345")
data_dict = wo.request(filename, signals)

#do whatever you want with the data
```

Note: .h5 in remote database and .bp in local database



# Thank you for your attention

Results?!?



**UNIVERSITY OF  
CAMBRIDGE**

**Fairy WANDS for Fusion**  
Wide Area Network Data Streaming  
Stefanie Reuter<sup>1</sup> Jonathan Hollocombe<sup>2</sup> Toby James<sup>2</sup>  
<sup>1</sup>University of Cambridge - UK <sup>2</sup>UKAEA

### Background

Fusion reactors, such as UKAEA's MAST-U, are pulsed fusion reactors that run individual reactions known as shots.



Figure 1. MAST-U tokamak reactor (Credit: UKAEA)

Next generation fusion reactors are expected to produce data in the order of petabytes per shot, that need to be analysed to adjust for the next shot. Computational resources for analysis are not co-located with the reactor, creating a need for efficient data streaming across wide area networks.

### WANDS Workflow

Data requests are issued from the compute site, triggering a remote request only if data is not available locally. Every variable that is received from the remote site via TCP/IP is stored locally, allowing fast access in future requests. ADIOS2, a unified, high-speed, high-performance I/O framework, underpins the remote data transfer and local caching.



Figure 2. Interaction flowchart of various elements within WANDS, beginning from a data request at the compute site to the availability of data for analysis.

### Example Usage (Python client)

```
from wands import Wands

# Create a Wands object
wands = Wands(
    filename = "example_data.h5",
    hostname = "localhost",
    port = 12345,
    database = "wands_db"
)

# Make a request
data_block = wands.request("example_data")

# Print the data
print(data_block["example_data"])
```

### Results and Discussion

A 7.5 GB sample data HDF5 file from one MAST-U shot is used to benchmark transfer speeds. Transferring the full file via rsync takes the longest, see left in Figure 3a. The middle bar shows a WANDS remote request for a 72.6 MB subset of the original file needed for an analysis script, and compares this to the same request if the data can be loaded from the local database after having been requested previously (right).



Figure 3a. Wide Area Network Streaming Times

In a second experiment, the large data file was intended to only include the required data. Figure 3b shows the time needed to send this reduced file with rsync, and compares this to a remote request via WANDS.



Figure 3b. Wide Area Network Streaming Times

### Conclusion and Limitations

- WANDS is able to match performance for a single data request.
- WANDS offers significant performance benefits if data is reused.
- By initiating the next request while running the analysis on the first data, the intended analysis time can be further improved.
- If all data available on the data site is requested at once, the use of WANDS does not provide any benefit when compared to bulk transfer.
- As the data transferred becomes immediately available in main memory, the available RAM on the compute side puts a maximum limit on the amount of data that can be loaded per script at a time.

### References

1. Hollocombe, J., James, T., & Reuter, S. (2023). The MAST-U tokamak reactor: a framework for high-performance data management. *Plasma Physics and Controlled Fusion*, 65(12), 012001.
2. MAST-U tokamak reactor. (2023). <https://www.ukaea.ac.uk/mastru>.

### Collaboration Partners





# The perfect answer

Elisabeth Ortega, PhD. - R&D&I Manager

22/05/2023

ISC23 Exhibitor Forum - WHPC

[www.hpcnow.com](http://www.hpcnow.com)

... for the question:




[Image by studiogstock - Freepik](#)

which I was not able to answer properly until I had a conversation with a **woman in HPC** at my early career in HPC (one year ago)

**which I was not able to answer properly until I had a conversation with a **woman in HPC** at my early career in HPC (one year ago)**

(which I'll answer with a story)




## The topic:

Improving the durability of planes by adding extra pieces of metal on the critical parts of the plane considering the bullet hits of the returning planes, but where?

- a) In the bullet holes (red dots)
- b) In other places






## The topic:

Improving the durability of planes by adding extra pieces of metal on the critical parts of the plane considering the bullet hits of the returning planes, but where?


- a) In the bullet holes (red dots)
- b) In other places



## It's not wrong data, it's the wrong approach



## It's not wrong data, it's the wrong approach



**The perfect answer is ...**

**The perfect answer is ...**

**don't ask it to me**

**The perfect answer is ...**



**don't ask it to me**





# Thanks for your attention!

(and visit us at booth D404!)



# Space Weather Forecasting using Celery and django

Dr Jenny Wong

Senior Research Software Engineer

University of Birmingham, UK



UNIVERSITY OF  
BIRMINGHAM




WHPC  
WOMEN IN HIGH  
PERFORMANCE COMPUTING



# Space Weather Forecasting

- Forecasting space weather conditions in the Earth's ionosphere is critical to protecting key infrastructure
- Variations in space weather are caused by Coronal Mass Ejections from the Sun's surface
- Produces disturbances in communications and electrical systems, as well as spectacular aurorae










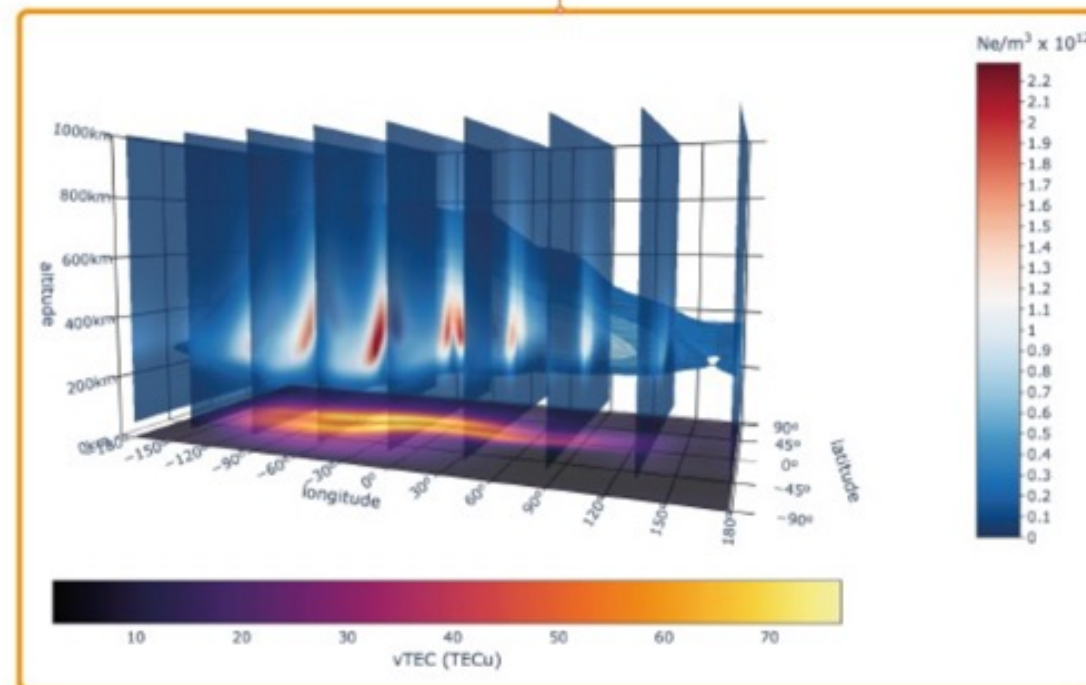
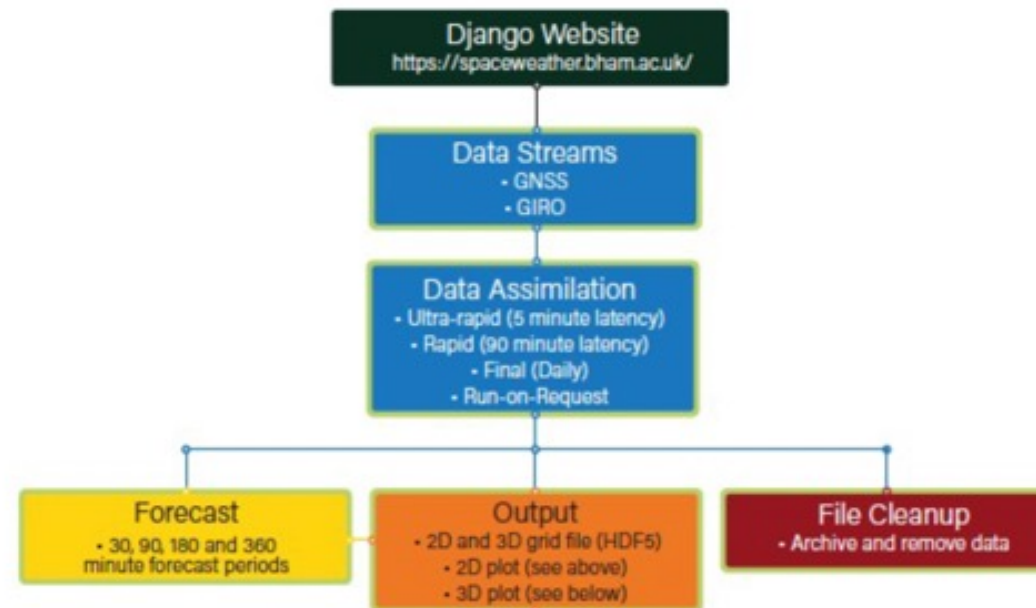


## Making automated, HPC-powered scientific results accessible in near-real time

- We present a system for operationalising HPC tasks for data assimilation in space weather forecasting using Celery and Django
- Celery is a distributed task queue used to execute asynchronous jobs
- Django is popular Python-based web framework where simulation outputs are presented

AIDA Ultra Rapid - 04/28/2023, 13:45:00







# Space Weather Forecasting using Celery and django

Making automated, HPC-powered scientific results accessible in near-real time.

Author Dr Jenny Wong | Senior Research Software Engineer | University of Birmingham, UK

Principal Investigator Dr Sean Dodgson | Head of Space Environment Research (SERENE) | University of Birmingham, UK


## Space Weather Forecasting

Forecasting space weather conditions in the Earth's ionosphere is critical to protecting key infrastructure, such as satellite-based positioning and navigation systems, high frequency radio communications, and the electric power grid.

Variations in space weather are caused by coronal mass ejections from the Sun's surface, travelling up to 5 million kilometres per hour towards the Earth, dragging electrons in the ionosphere to produce disturbances in communications and electrical systems, as well as spectacular auroras.

In conjunction with the University of Birmingham's SERENE group, we present a system for operationalising HPC tasks for data assimilation in space weather forecasting using Celery and Django. Celery is a distributed task queue that allows us to execute tasks asynchronously in a distributed environment, while Django is a popular web framework that provides a high-level view of user interaction. Our system integrates both these tools to automate the process of running space weather simulations on an HPC cluster for data assimilation and presenting outputs on the website.

Figure 22: 2D plot showing ultra-rapid 15 minute (km/s) assimilation output. (Clockwise from top-left) Vertical Total Electron Count, Maximum Usable Frequency, Height and value of peak F2 layer frequency.



## Operationalise with Celery<sup>1</sup> and Django<sup>2</sup>

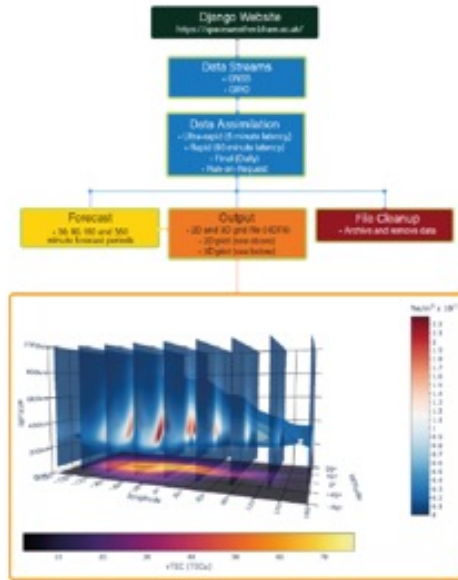
High performance computing plays a significant role in simulating and forecasting space weather, as we perform data assimilation on enormous amounts of data that require a large amount of computing power. However, manually running HPC tasks can be time-consuming and error-prone. Automating HPC tasks can greatly enhance efficiency and generate scientific results accessible in near real-time.

The Django website is designed to be modular and flexible and provides an intuitive web interface for users to submit and monitor their tasks, as well as visualise results. It allows users to access automated assimilation outputs with 5, 10, 30 minute and daily intervals, as well as forecasts. Users can also define and monitor their custom simulations with a custom request function.

Celery is used to download observational data from hundreds of GPS and GNSS stations around the world where the assimilation ingests up to 250 gigabytes of data per hour. Then, Celery sends assimilation and forecast jobs to the Storm scheduler, which enables efficient and resilient resource utilisation on a federated server. Celery is also used to create visualisations and automate file cleanup.

In conclusion, this project demonstrates that Celery and Django are powerful and flexible tools for operationalising HPC tasks for data assimilation in space weather forecasting. It enables users to efficiently utilise HPC resources while providing an easy-to-use website for interacting with jobs and simulation outputs. Our system applies to a wide range of HPC tasks in research software, and we believe it will be a useful framework for researchers to personalise their code.

Figure 10: 3D plot showing ultra-rapid (15 minute (km/s)) assimilation output. Total Electron Count (km/s) as a function of altitude, longitude and latitude with plots in longitude and a VTEC line.



## Acknowledgements

The author wishes to acknowledge funding from European Space Agency: 10 ionospheric modelling - UTR-00009. Project design and execution was delivered by the SERENE group and the Advanced Research Computing group for research software support.

1 Celery (2020). Retrieved from: <https://github.com/celery>

2 Django Software Foundation (2016). Django. Retrieved from: <https://djangoproject.com>

# Space weather forecasting using Celery and Django: Making automated, HPC-powered scientific results accessible in near-real time.

Foyer D-G - 2nd Floor  
Women in HPC Poster



**Dr Jenny Wong**  
Senior Research Software Engineer  
University of Birmingham, UK



@\_jennywong\_



thisisjennywong





# Poster Reception

---

May 23, 2023

Foyer D-G

3:30 - 5:30 PM

# Interested in Sponsoring Us?

---

Talk to us today about how you can fund our global initiatives!





# Deb Goldfarb

Director of Product + Strategy  
for HPC & Batch, AWS

# LESSONS

**Stuff I learned along the way  
which turned out to be important**

Debra Goldfarb  
Director, HPC Products and Strategy, AWS

# Inspired by mistakes...

## ...My own and those of others





# Lesson 1

## ...Laying a Foundation

A photograph showing a wooden frame structure, likely for a foundation or wall, laid out on a dirt surface. The frame is made of light-colored wooden planks and is composed of several interconnected rectangular sections. The background is a rough, uneven dirt ground.

What gives me value? Is it just productivity?

What is the place / priority of my career?

What am I not willing to give to further my career?

What am I not willing to do to keep my job?



# Lesson 2

## ...Most of Us Have Impostor Syndrome



"It's perfectly okay to occasionally feel like a fraud when it comes to your career. I'm just not sure you need to say that on your resume."

# Lesson 3

## Do it anyway

you get in life  
WHAT YOU HAVE THE  
COURAGE  
to ask for.  
- OPRAH WINFREY

# Final Thoughts

Let your values drive your career

Don't crush yourself with false expectations or comparisons

Ask for what you want...and deserve

Celebrate where you are


# Early Career Speakers




**Paula Esposito**  
La Laguna University



**Aleksandra Kowalczyk**  
University of Warsaw



**Ekaterina Zossimova**  
University of Exeter



**Helena Vela Beltran**  
Barcelona  
Supercomputing Center

# Performance Analysis of Parallel Codes with FancyJCL

High Performance Computing Group  
Department of Computer Engineering and Systems, University of La Laguna

Sergio Afonso  
Óscar Gómez-Cárdenes  
Paula Expósito  
Vicente Blanco  
Francisco Almeida







## Problematic of acceleration on mobile devices

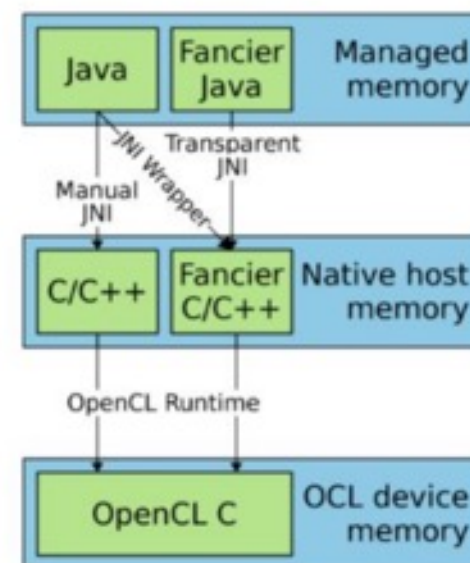
Workflow for execute an app on a mobile device's GPU



The advances on SoC hardware are not up to par with the software's



Communication between memory spaces





# FancyJCL

Through a sequential interface users can accelerate applications

FancyJCL is built on the top of the Fancier library for avoiding memory copies

## Adding a constant to an array with FancyJCL

```
public class Foo {  
    public void run (long[] array, long kConstant) {  
        Stage stage = new Stage();  
        stage.setKernelSource("array[d0] += kConstant;");  
        stage.setInputs(Map.of("array", array, "kConstant", kConstant));  
        stage.setOutputs(Map.of("array", array));  
        stage.setRunConfiguration(new RunConfiguration(new long[]{array.length},  
                                                         new long[]{4}));  
        stage.runSync();  
    }  
}
```

# Results


## Algorithms



x5 Image Processing Algorithms  
x8 Different Image Resolution

## Devices

Device	System-on-Chip	Characteristics
Xiaomi Mi Mix 2	Snapdragon 845	Programmable GPU Unified Memory Architecture
Snapdragon 865 SDK	Snapdragon 865	
Vivo iQOO 7	Snapdragon 888	



In horizontal, 5 groups of kernels with 8 resolutions each. In vertical, achieved speedup.

# Thank you for your attention

Sergio Afonso	<a href="mailto:safonsof@ull.edu.es">safonsof@ull.edu.es</a>
Óscar Gómez-Cárdenes	<a href="mailto:ogomezca@ull.edu.es">ogomezca@ull.edu.es</a>
<b>Paula Expósito</b>	<b><a href="mailto:pexposit@ull.edu.es">pexposit@ull.edu.es</a></b>
Vicente Blanco	<a href="mailto:vblanco@ull.edu.es">vblanco@ull.edu.es</a>
Francisco Almeida	<a href="mailto:falmeida@ull.edu.es">falmeida@ull.edu.es</a>

**Acknowledgements:** This work has been supported by the Spanish Ministry of Science and Innovation with the PID2019-107228RB-I00, TED2021-131019B-I00 and PDC2022-134013-I00 projects; and by the Government of the Canary Islands with the project ProID2021010012.



ISC HIGH PERFORMANCE, HAMBURG 2023

# **APPLICATION OF ADVANCED TEXT ANALYSIS IN THE STUDY OF SCIENTIFIC LITERATURE**

ALEKSANDRA KOWALCZUK  
INTERDISCIPLINARY CENTRE FOR MATHEMATICAL AND  
COMPUTATIONAL MODELLING



UNIVERSITY OF WARSAW  
Interdisciplinary Centre for Mathematical  
and Computational Modelling  
[icm.edu.pl](http://icm.edu.pl)



# OUTLINE

- Background
- State of Art
- Methods
- Preliminary Results
- The Future



# BACKGROUND

In view of the constantly growing amount of scientific literature, there is a real need to make it easier for experts like medical doctors to access resources published around the world and quickly draw conclusions from them.




## What is the Value?

Discovery of novel tools

Identification of promising new paths in science

Better performance of research organisations

Improvement of career paths for researchers





# STATE OF ART

We can observe a very dynamic development of AI natural language models and ready-to-use tools dedicated to text analysis, processing and generation.



## Challenges

Different scientific data formats

Accessing and searching relevant scientific literature

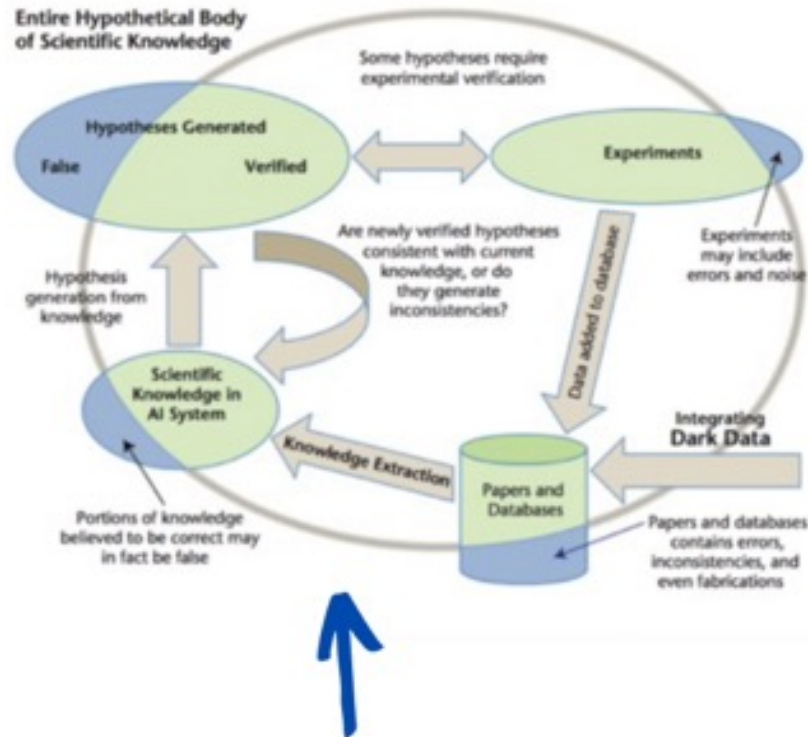
The direction of valuable research work

More and more advanced language models for NLP

Hardware and computing power

Scientific, social and legal impact

# METHODS



Scientific Literature (unstructured data)



Auto-Curation



(structured knowledge)



Engine of scientific discovery:  
Taxila



Intelligent algorithms and scripts:  
language models for NLP

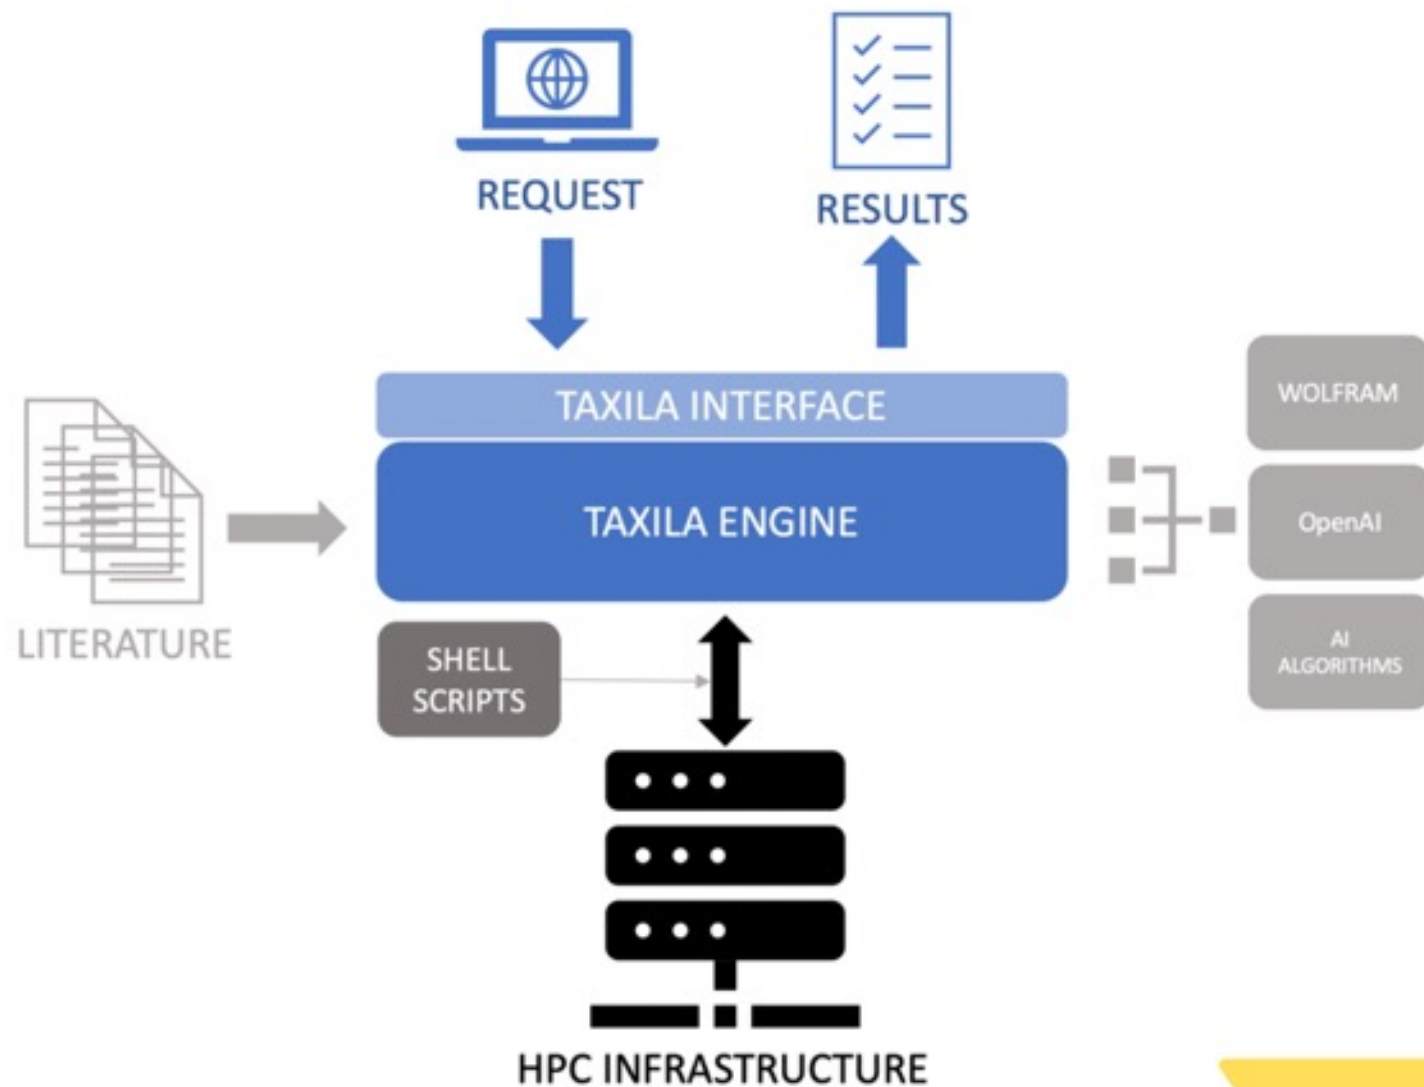


Ready-to-use tools:  
OpenAI, Wolfram



Analyzing text in biological context:  
CAR-T therapy

# ARCHITECTURE



# PRELIMINARY RESULTS



# THE FUTURE

Can a machine write a good scientific article?

Can a machine win the Nobel Prize?

*"Our ultimate goal is to develop an AI System that can make major scientific discoveries that would improve the state of the world and impacting the way we do science."*

Hiraoki Kitano, SBI Tokyo



# SOURCES

## [1] Taxila Curation

<https://curation.taxila.io>

## [2] SBI Publications

<https://medium.com/@sbijapan>

## [3] OpenAI

<https://openai.com>

## [4] Wolfram Alpha Series of Articles

<https://writings.stephenwolfram.com>

## [5] T-PAIR: Temporal Node-pair Embedding for Automatic Biomedical Hypothesis Generation

IEEE XPLORE, Vol. 34 Issue: 6

Uchenna Akujuobi, Michael Spranger, Sucheendra K. Palaniappan, Xiangliang Zhang

## [6] What Is ChatGPT Doing ... and Why Does It Work?

Wolfram Research, Inc., March 2023

Stephen Wolfram

# THANK YOU!

ALEXANDRA.KOWALCZUK@GMAIL.COM

# High Performance Computing for Decoding Biosensor Signals



Ekaterina Zossimova and Michael Walter

University of Exeter, UK & Freiburg Center  
for Interactive Materials and Bioinspired  
Technologies, Germany

22 May 2023




# Opto-plasmonic biosensor



Vincent, S., Subramanian, S. & Vollmer, F. Optoplasmonic characterisation of reversible disulfide interactions at single thiol sites in the attomolar regime. Nat Commun 11, 2043 (2020). Licensed under [Creative Commons CC BY](#)



# Opto-plasmonic biosensor





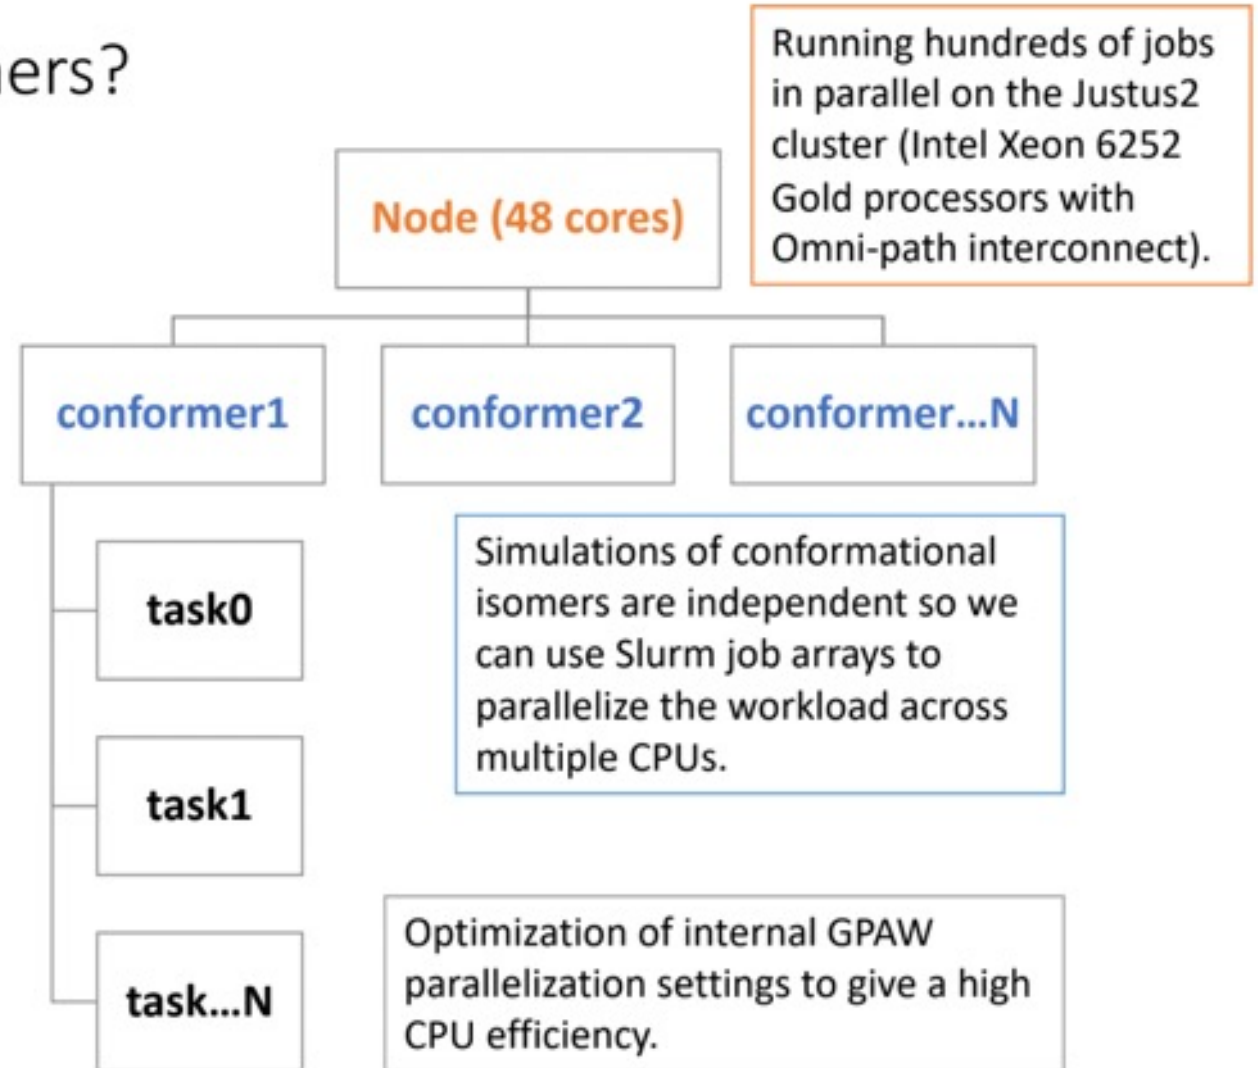


# Excess polarizability

What is the effect of charge/ protonation state?



# Excess polarizability

What about conformational isomers?



# Conclusions and outlook

## Progress

- Decoding the biosensor signals by calculating the static polarizability of molecules and also by quantifying the statistical dispersion due to charge or conformational isomers.
- Using HPC resources to run hundreds of jobs in parallel, reducing the simulation time from several days to less than 1 hour.

## Challenges

- It is not possible to use this method for optically active molecules. This requires a full time-dependent electronic structure calculation which is significantly more computationally intensive.

The authors acknowledge support by the state of Baden-Württemberg through bwHPC and the German Research Foundation (DFG) through grant no INST 40/575-1 FUGG (JUSTUS 2 cluster).

Ekaterina Zossimova  
✉ ez216@exeter.ac.uk



# INTEGRATION OF A PARALLEL EFFICIENCY MONITORING TOOL INTO AN HPC PRODUCTION SYSTEM

Co-director, David Vicente

Co-director, Marta García-Gasulla

Ponent, Víctor López

Author, Helena Vela Beltran



**Barcelona  
Supercomputing  
Center**  
*Centro Nacional de Supercomputación*



**FIB**

ISC  
May 2023



# CONTEXT



The screenshot shows the HPC User Portal interface. At the top, there's a navigation bar with 'ADMINISTRATION' and 'SWITCH'. Below it, a message states: 'Please Note: It is your responsibility to ensure that your account is properly configured for HPC use. For more information, please refer to the HPC User Portal documentation.' The main content area is titled 'All accounts' and 'All machines'. It displays a table with columns: ID, Name, Status, User, Machine, CPU, Submit time, Start, Workload, Nodes, Tasks, CPU%, and Memory%. The table lists several accounts and machines, including '1214222', '1214223', '1214224', '1214225', '1214226', '1214227', '1214228', and '1214229'.

ID	Name	Status	User	Machine	CPU	Submit time	Start	Workload	Nodes	Tasks	CPU%	Memory%
1214222	1214222	Pending	user1	Machine 1	100%	2023-10-26 11:00:00	N/A	10:00:00	1	10	N/A	N/A
1214223	1214223	Pending	user2	Machine 2	100%	2023-10-26 11:00:00	N/A	10:00:00	2	20	N/A	N/A
1214224	1214224	Pending	user3	Machine 3	100%	2023-10-26 11:00:00	N/A	10:00:00	1	10	N/A	N/A
1214225	1214225	Running	user4	Machine 4	100%	2023-10-26 11:00:00	2023-10-26 11:00:00	10:00:00	1	1	100%	100%
1214226	1214226	Pending	user5	Machine 5	100%	2023-10-26 11:00:00	N/A	10:00:00	2	20	N/A	N/A
1214227	1214227	Running	user6	Machine 6	100%	2023-10-26 11:00:00	2023-10-26 11:00:00	10:00:00	2	20	100%	100%
1214228	1214228	Pending	user7	Machine 7	100%	2023-10-26 11:00:00	N/A	10:00:00	1	10	N/A	N/A
1214229	1214229	Pending	user8	Machine 8	100%	2023-10-26 11:00:00	N/A	10:00:00	1	10	N/A	N/A

User Support Team

## Problem 01

HPC resources are not used properly.

## Problem 02

Many scientific application developers do not come from a Computer Science background.

## Problem 03

Performance analysis applications can be confusing and difficult to use for non-experienced users.

## Problem 04

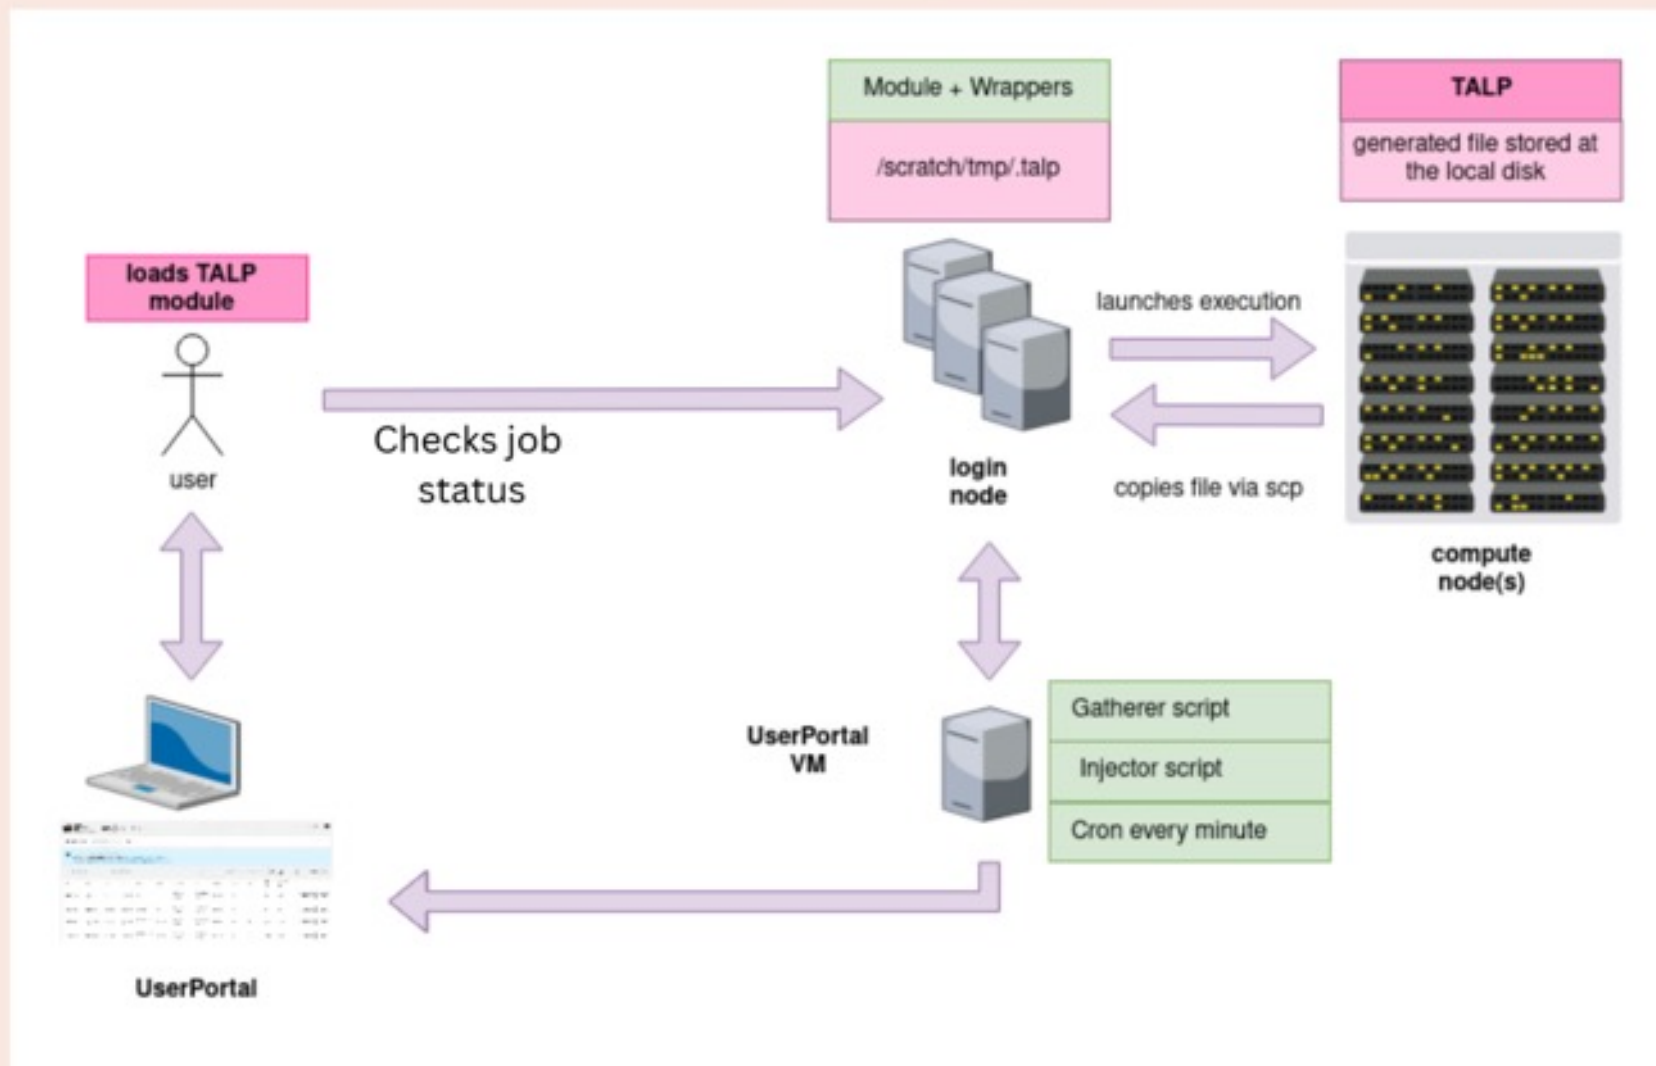
HPC users might not be aware of performance degradation.



Best Practices for Performance and Programmability (BePPP)



# INTEGRATION SCHEMA




Results

TALP report ?

MPI Execution: MPI Execution  
Elapsed Time: 0.00 s  
Parallel Efficiency: 0.84  
↳ Communication Efficiency: 1.00  
↳ Load Balance: 0.84  
↳ lbin: 0.99  
↳ lbinOut: 0.85



# THANK YOU!







# Poster Reception

---

May 23, 2023

Foyer D-G

3:30 - 5:30 PM

# Supporting Institutions - Anchor





# Supporting Institutions - Career




**Hewlett Packard  
Enterprise**



**WHPC**  
WOMEN IN HPC

# Supporting Institutions - Volunteer



# Interested in Sponsoring Us?

---

Talk to us today about how you can fund our global initiatives!



# How to get involved!

## **Join Women in HPC for FREE!**

- Select a Chapter/Affiliate for local events
- Let us know if you would like to get involved:
  - SC23 - Denver, USA
  - ISC 2024

**womeninhpc.org**

