# Energy-Aware Scheduling for Serverless Scientific Workflows: A Machine Learning Approach

Meeting with Prof. Viktor Prasanna

**Lucas de Sousa Rosa**
roses.lucas@usp.br

**Alfredo Goldman**
gold@ime.usp.br

# My PhD Network


**This is me**

PhD student

**Alfredo Goldman**

My Advisor

**Danilo Santos**

UGA

**Raphael Camargo**

UFABC

**Pedro Bruel**

HP Labs

**Anderson Andrei**

HP Labs

**Marcos Amarís**

UFPA


# Frontier

1.1 EFlop/s - 23 MW

Rank 1


# Aurora

0.6 EFlop/s - 24 MW

Rank 2

# Fugaku

0.4 EFlop/s - 30 MW

Rank 4

# Motivation

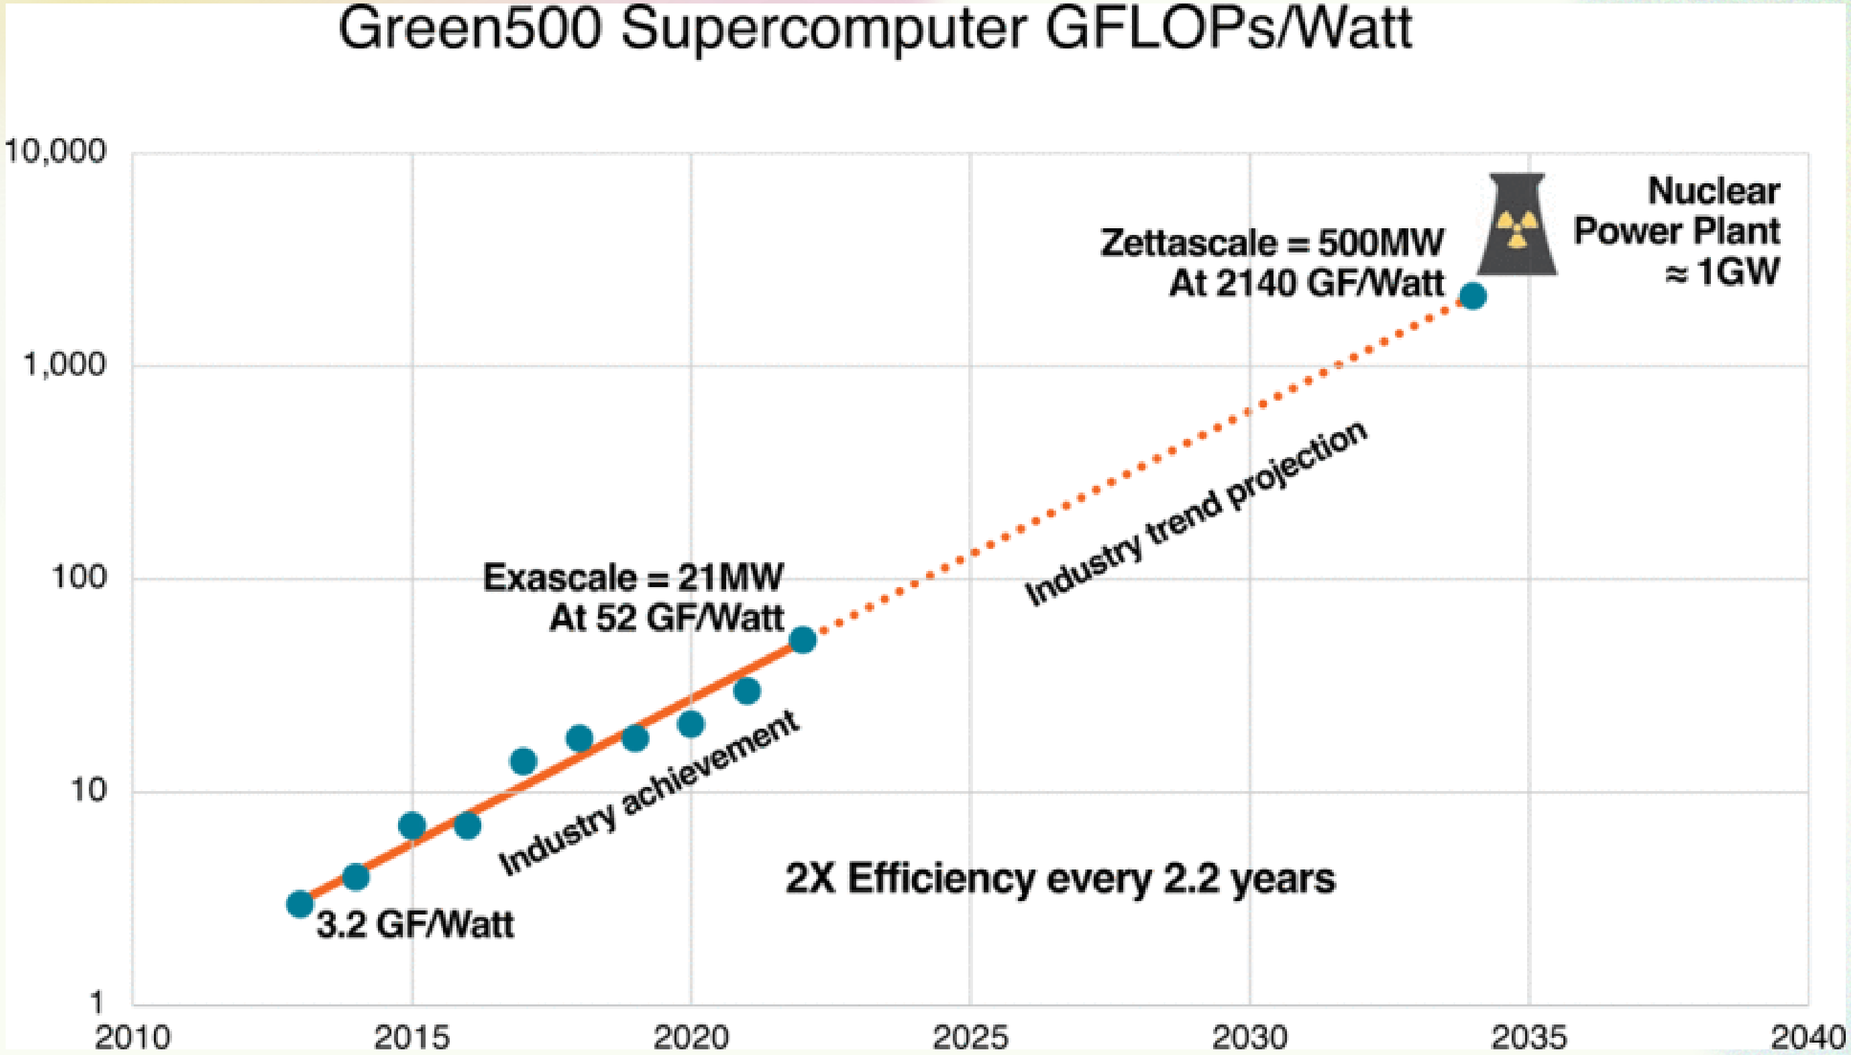


**Lisa Su**

CEO of AMD

Innovation for the Next Decade of High-Performance Computing - ISSCC ’23 [2]

# Main Objective

**Specialized, power-efficient hardware [2,7]**

Low-power devices and components

**Efficient cooling technologies [7]**

Air/liquid cooling or mixed cooling

Hardware Level

**Objective**

Match increasing computing demands with reduced power consumption

Software Level

**Enhance resources management and system scheduling [7]**

Improving energy-awareness of workflow management systems (WMSs)

# Scientific Workflows

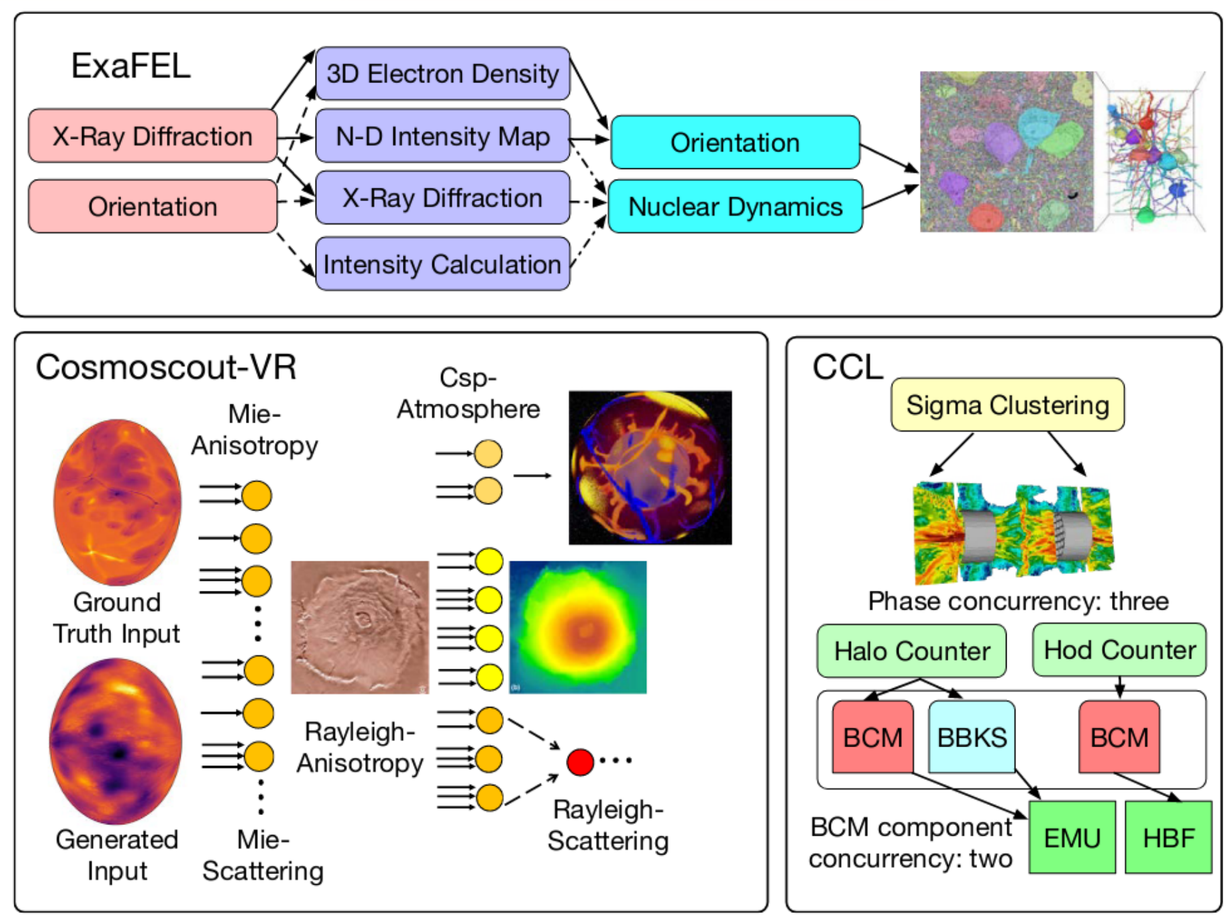


DayDream: Executing Dynamic Scientific Workflows on Serverless Platforms with Hot Starts - SC '22 [5]

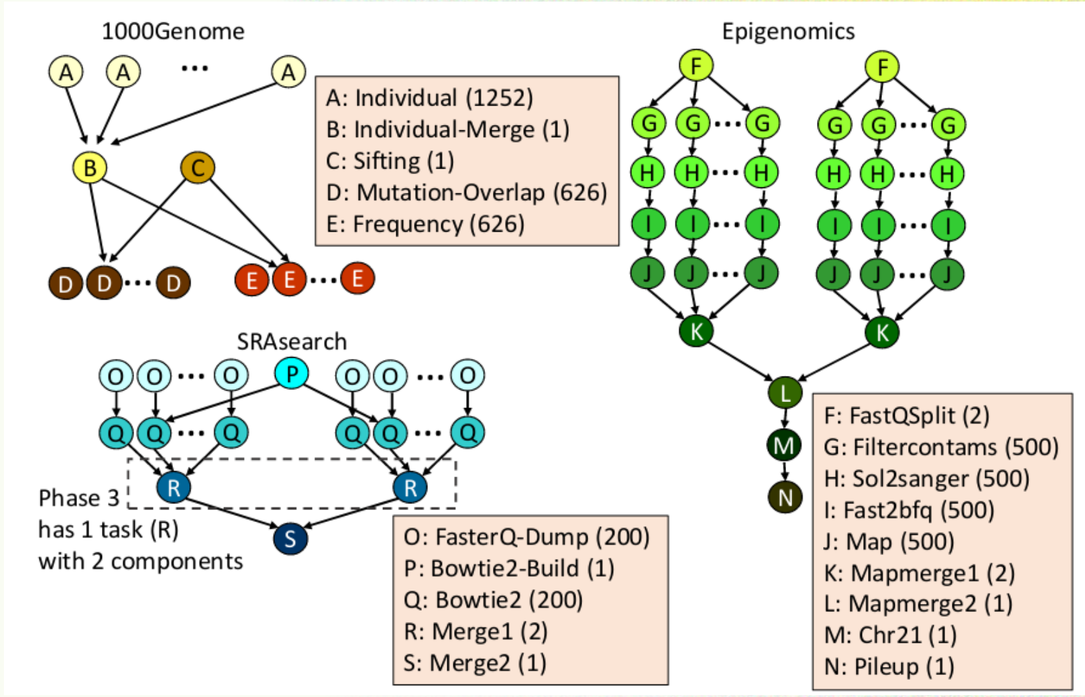


Mashup: making serverless computing useful for HPC workflows via hybrid execution - PPoPP '22 [6]

Apache Airflow

argo

pegasus

# A New Trend

## Traditional deployment

Monolithic applications running on on-premises clusters and IaaS

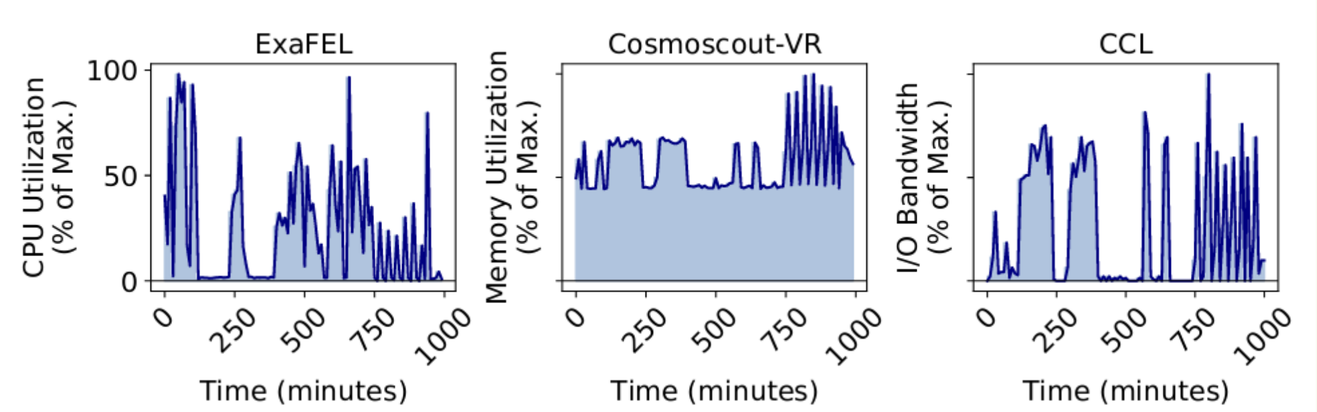


Roy et al. *

*Resources under-utilization, over-provisioning, and high expenses [3-5]*

## Emerging trend

Serverless computing can alleviate these challenges

fission

kubernetes

*Data locality, cold start, and resource volatility [3-5]*

*DayDream: Executing Dynamic Scientific Workflows on Serverless Platforms with Hot Starts - SC '22 [5]

# Workflows Scheduling

## The scheduling problem

- Distributing computational resources among different tasks accodring to specific **constraints**
- Balancing energy efficiency with other objectives presents a challenge [4]

→ Minimize energy consumption
Minimize workflow makespan
Maximize resources utilization

## Solving the scheduling problem

- Workflows scheduling in HPC belongs to the NP-hard problem class
- Usually, we see in the literature heuristics with manually-tuned parameters

In recent years, the application of ML has significantly increased in task scheduling [1]

# Machine Learning and Scheduling

## Advantages of ML

- ML techniques can handle complex scenarios with multiple states of the computing environment (high heterogeneity)
- Have the ability to self-adapt and self-learn

Some studies applied RL and DRL for task scheduling [1,9], while others utilized regression methods to create new policies [1]

Only a small number of recent research papers [9] have implemented such techniques in the context of **serverless computing**

# Bringing Everything Together

**Energy-aware scheduling**

Introduce energy-saving innovations at the software level, focusing on scheduling

**Serverless scientific workflows**

Make serverless more attractive for scientific workflows

**Machine Learning techniques**

Develop models to address multiple complex challenges

# Thank you!

Questions?!

**Lucas de Sousa Rosa**
roses.lucas@usp.br

**Alfredo Goldman**
gold@ime.usp.br

# References

[1] G. P. Koslovski, K. Pereira, and P. R. Albuquerque, “DAG-based workflows scheduling using Actor–Critic Deep Reinforcement Learning,” Future Generation Computer Systems, vol. 150, pp. 354–363, Jan. 2024, doi: 10.1016/j.future.2023.09.018.

[2] L. Su and S. Naffziger, “1.1 Innovation For the Next Decade of Compute Efficiency,” in 2023 IEEE International Solid-State Circuits Conference (ISSCC), San Francisco, CA, USA: IEEE, Feb. 2023, pp. 8–12. doi: 10.1109/ISSCC42615.2023.10067810.

[3] A. Elshamy, A. Alquraan, and S. Al-Kiswany, “A Study of Orchestration Approaches for Scientific Workflows in Serverless Computing,” in Proceedings of the 1st Workshop on SErverless Systems, Applications and MEthodologies, Rome Italy: ACM, May 2023, pp. 34–40. doi: 10.1145/3592533.3592809.

# References

[4] L. Versluis and A. Iosup, “TaskFlow: An Energy- and Makespan-Aware Task Placement Policy for Workflow Scheduling through Delay Management,” in Companion of the 2022 ACM/SPEC International Conference on Performance Engineering, Bejing China: ACM, Jul. 2022, pp. 81–88. doi: 10.1145/3491204.3527466.

[5] R. B. Roy, T. Patel, and D. Tiwari, “DayDream: Executing Dynamic Scientific Workflows on Serverless Platforms with Hot Starts,” in SC22: International Conference for High Performance Computing, Networking, Storage and Analysis, Dallas, TX, USA: IEEE, Nov. 2022, pp. 1–18. doi: 10.1109/SC41404.2022.00027.

[6] R. B. Roy, T. Patel, V. Gadepally, and D. Tiwari, “Mashup: making serverless computing useful for HPC workflows via hybrid execution,” in Proceedings of the 27th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, Seoul Republic of Korea: ACM, Apr. 2022, pp. 46–60. doi: 10.1145/3503221.3508407.

# References

[7] X. Liao et al., “Moving from exascale to zettascale computing: challenges and techniques,” Frontiers Inf Technol Electronic Eng, vol. 19, no. 10, pp. 1236–1244, Oct. 2018, doi: 10.1631/FITEE.1800494.

[8] “TOP500.” Accessed: Mar. 05, 2024. [Online]. Available: https://www.top500.org/

[9] A. Mampage, S. Karunasekera, and R. Buyya, “Deep reinforcement learning for application scheduling in resource-constrained, multi-tenant serverless computing environments,” Future Generation Computer Systems, vol. 143, pp. 277–292, Jun. 2023, doi: 10.1016/j.future.2023.02.006.