Shaping the Digital Future of ErUM Research: Sustainability & Ethics

From Data to Duty: A Simulation Framework for Sustainable Distributed Computing

Maximilian Horzela, Henri Casanova, Frédéric Suter, and others — Aachen, 29. July 2025

The WLCG

- World's largest computing grid
- Provides computing resources to store, distribute & analyse LHC data
  - >170 computing centers in 42 countries
  - >1 Million computer cores
  - >2 Exabyte of storage
  - ~1TB/s average transfer rate

The WLCG

- World's largest computing grid
- Provides computing resources to store, distribute & analyse LHC data
  - >170 computing centers in 42 countries
  - >1 Million computer cores
  - >2 Exabyte of storage
  - ~1TB/s average transfer rate

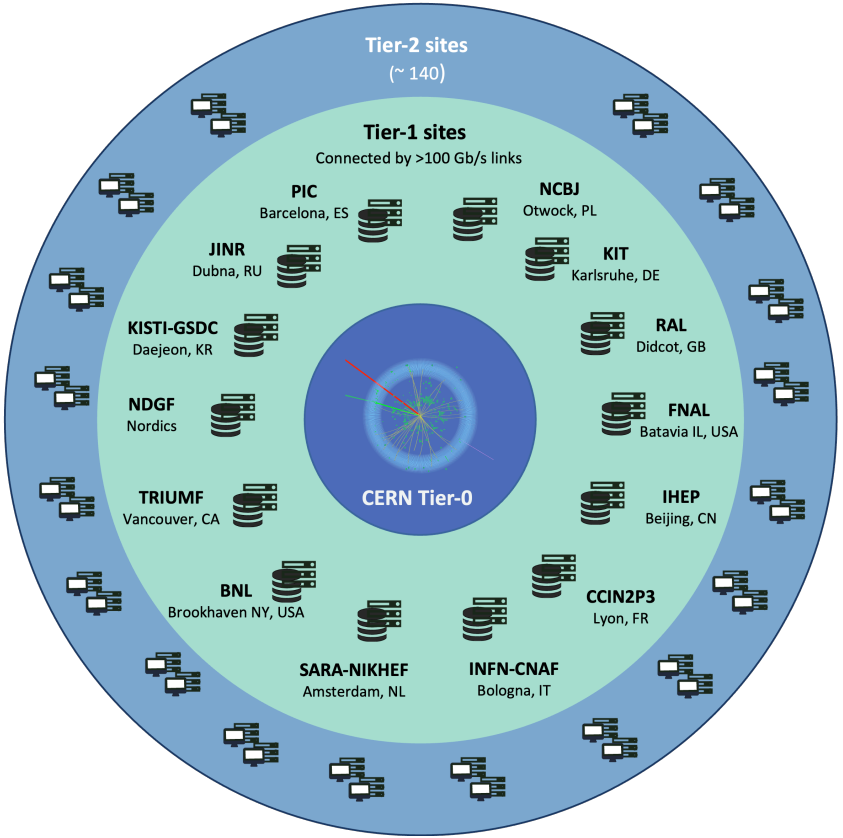

The WLCG — Germany

- Substantial contribution by German sites with high reliability
  - ∼10% of total compute
  - ∼10% of network traffic

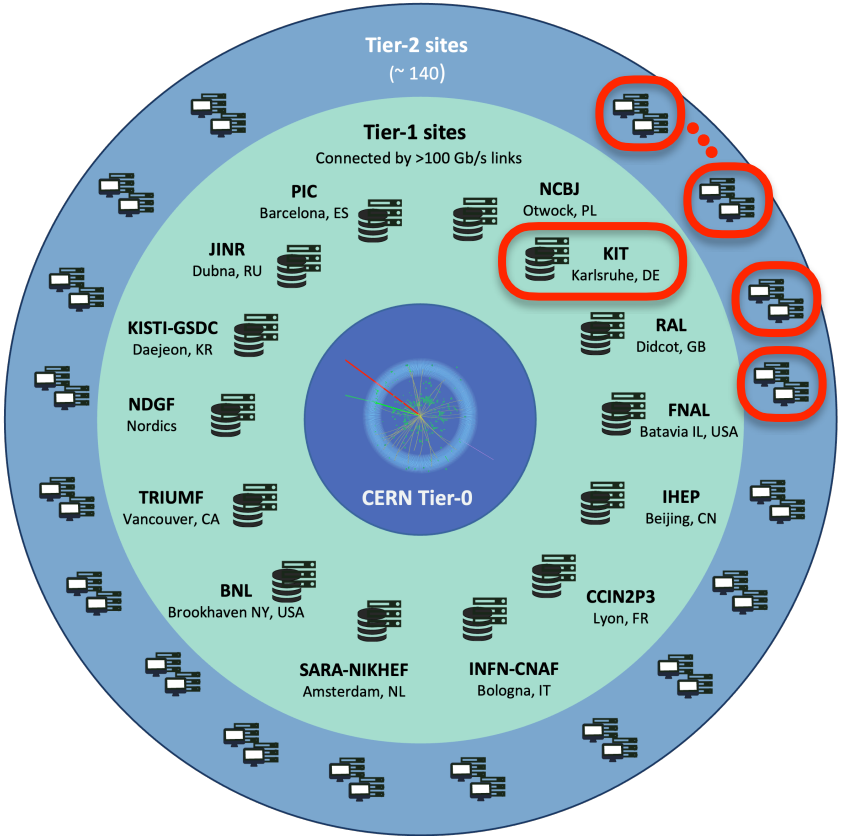

The WLCG — Germany

- Substantial contribution by German sites with high reliability
  - ∼10% of total compute
  - ∼10% of network traffic

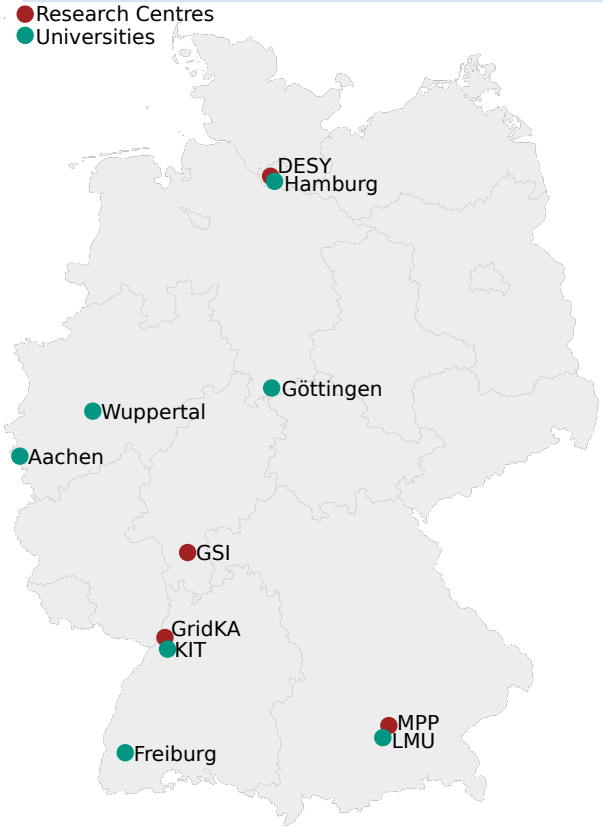

The WLCG — Germany

- Substantial contribution by German sites with high reliability
  - ∼10% of total compute
  - ∼10% of network traffic
- Decision to move from scattered network of sites to more consolidated infrastructure
  - NHR Supercomputers for compute

The WLCG — Germany

- Substantial contribution by German sites with high reliability
  - ∼10% of total compute
  - ∼10% of network traffic
- Decision to move from scattered network of sites to more consolidated infrastructure
  - NHR Supercomputers for compute
  - Data provisioning by DESY and KIT

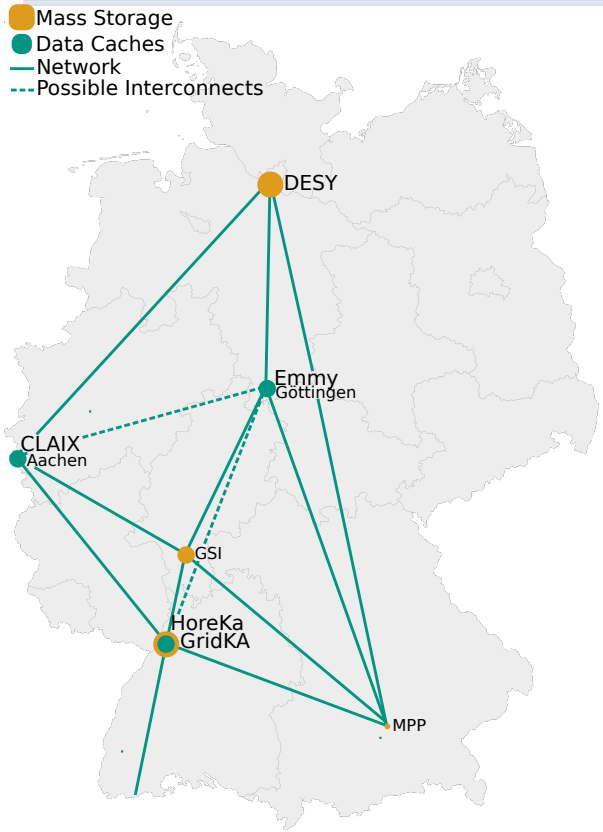

The WLCG — Germany

- Substantial contribution by German sites with high reliability
  - ∼10% of total compute
  - ∼10% of network traffic
- Decision to move from scattered network of sites to more consolidated infrastructure
  - NHR Supercomputers for compute
  - Data provisioning by DESY and KIT
- Remain a reliable partner in the WLCG

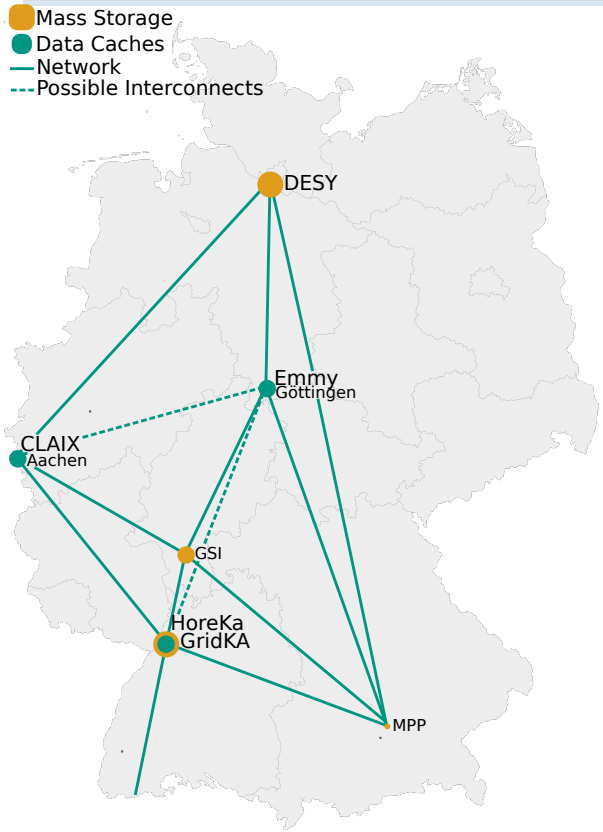

Challenges in Distributed Computing

- New paradigms entering the game require adjustments to the computing model

Challenges in Distributed Computing

- New paradigms entering the game require adjustments to the computing model
  - More heterogeneous hardware

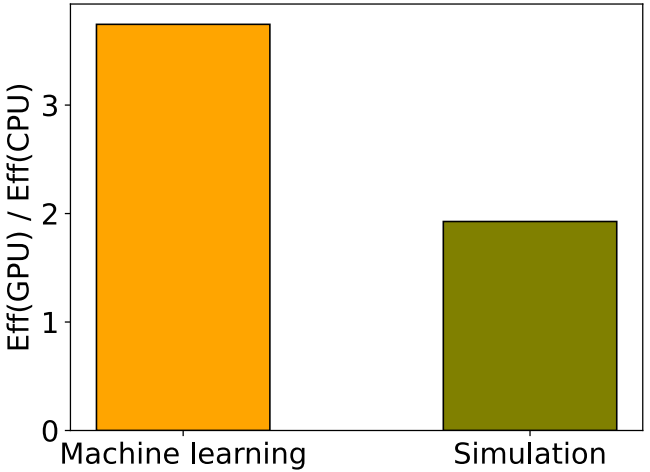

Challenges in Distributed Computing

- New paradigms entering the game require adjustments to the computing model
  - More heterogeneous hardware
  - Interactive & opportunistic resources

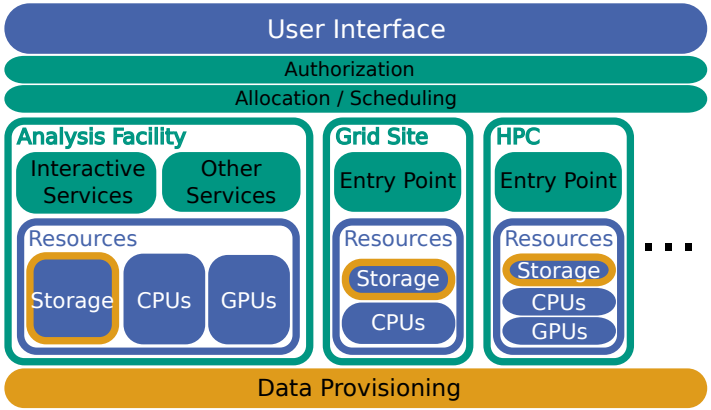

Challenges in Distributed Computing

- New paradigms entering the game require adjustments to the computing model
  - More heterogeneous hardware
  - Interactive & opportunistic resources
  - Federated data infrastructures

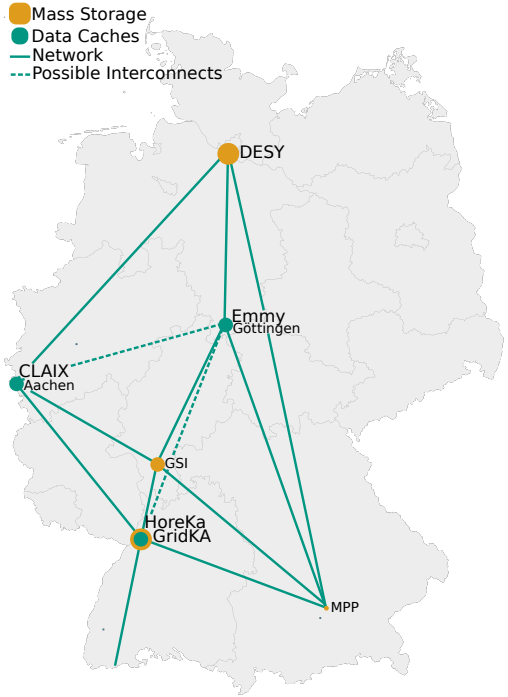

Challenges in Distributed Computing

- New paradigms entering the game require adjustments to the computing model
  - More heterogeneous hardware
  - Interactive & opportunistic resources
  - Federated data infrastructures
  - Energy efficiency

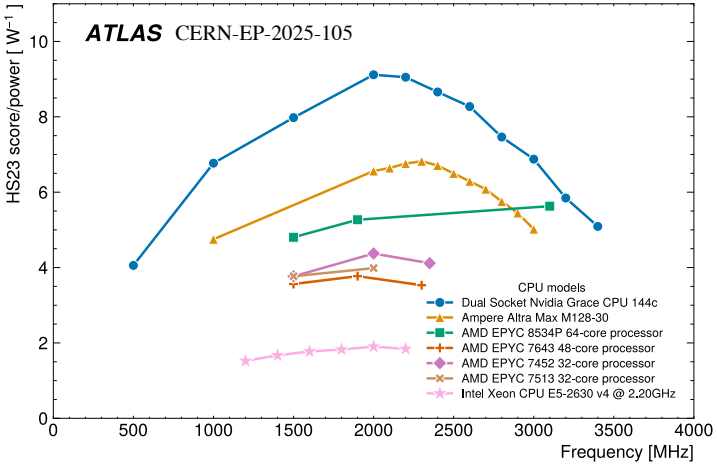

Challenges in Distributed Computing

- New paradigms entering the game require adjustments to the computing model
  - More heterogeneous hardware
  - Interactive & opportunistic resources
  - Federated data infrastructures
  - Energy efficiency
  - Energy source heterogeneity

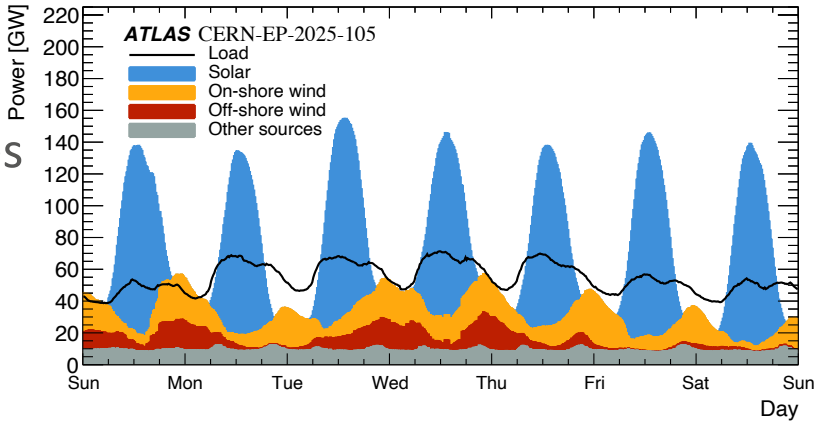

Challenges in Distributed Computing

- New paradigms entering the game require adjustments to the computing model
  - More heterogeneous hardware
  - Interactive & opportunistic resources
  - Federated data infrastructures
  - Energy efficiency
  - Energy source heterogeneity

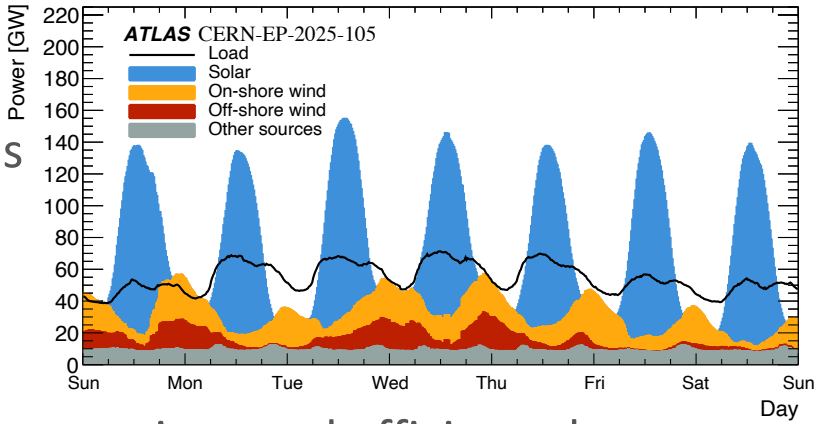

- Continue ***reliable operation*** without interruptions and efficiency losses with future ***sustainable infrastructure***

Solving the Challenges with Modelling

- How can we obtain successful computing models prior to building the infrastructures?
  - Historically → experience and gut feeling
  - Evidence-based → models calibrated on real-world data (digital twins)
- There is plenty of activity and examples (for single centres)

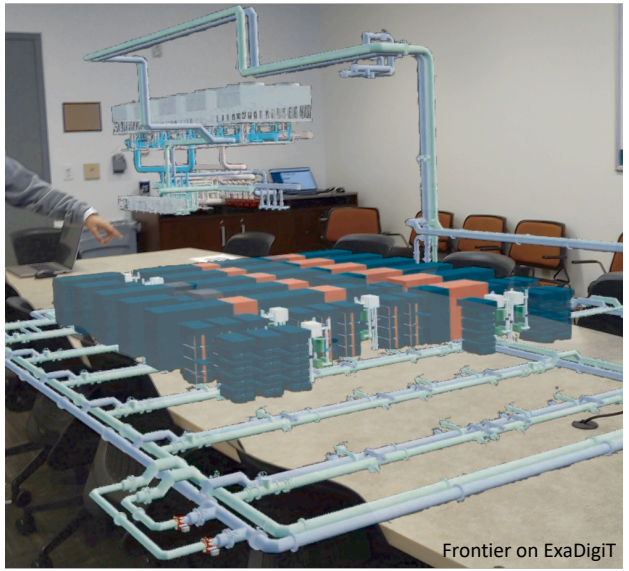
Frontier on ExaDigiT

Modelling Execution Traces — SimGrid

- Execution of computation, communication & data retrieval/storage activities on hosts, links & disks
- SimGrid toolkit
  - Implementation of flow models & resource representations describing activity progression & resource sharing
  - Demonstrated accuracy, scalability, and expressiveness

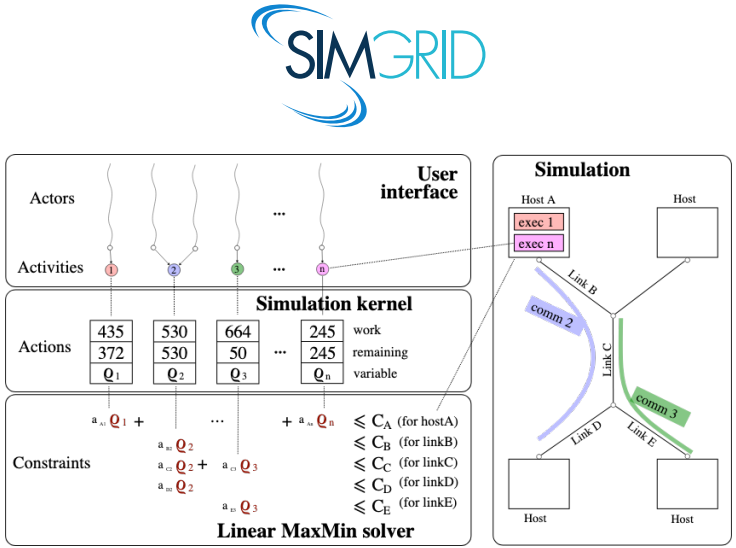

Modelling Energy Consumption — SimGrid

- Simulating execution traces
  → Prediction of dissipated energy
  - Worker machine computations (DVFS, utilised cores) ✓

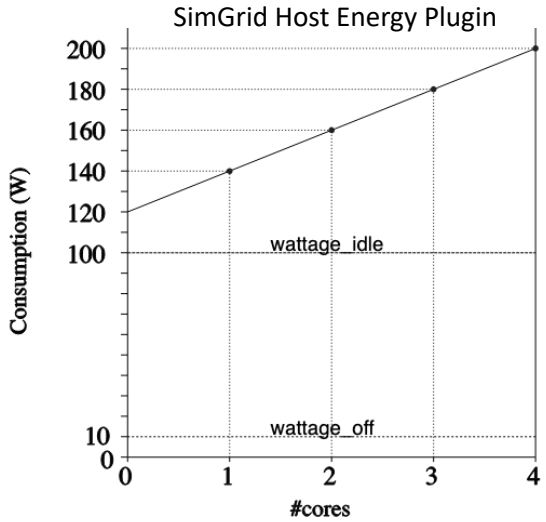

10.1109/CLUSTER.2017.66

Modelling Energy Consumption — SimGrid

- Simulating execution traces
  → Prediction of dissipated energy
  - Worker machine computations (DVFS, utilised cores) ✓
  - Network traffic ✓

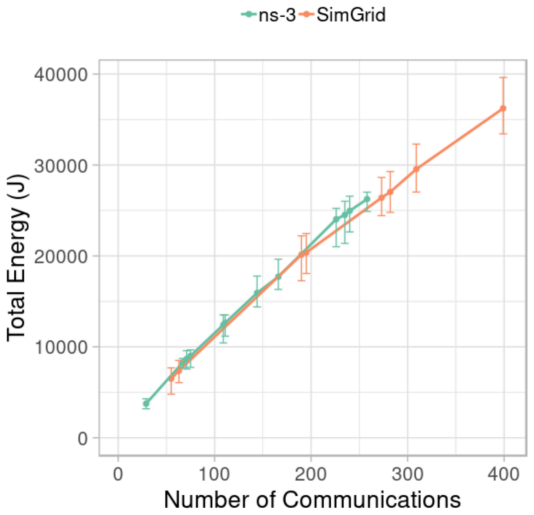

10.1007/978-3-030-15032-7_88

Modelling Energy Consumption — SimGrid

- Simulating execution traces
  → Prediction of dissipated energy
  - Worker machine computations (DVFS, utilised cores) ✓
  - Network traffic ✓
  - Retrieving and storing data ❌

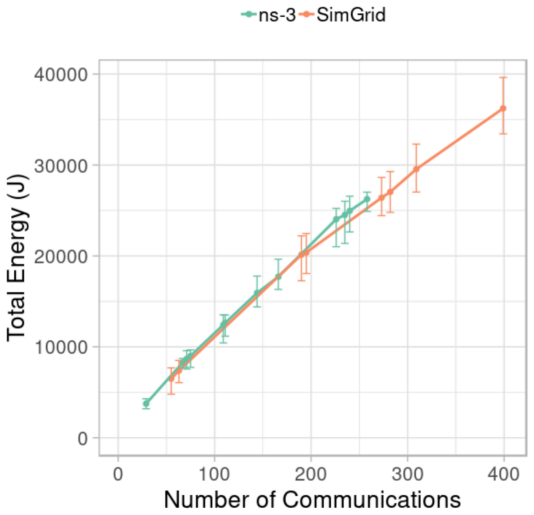

10.1007/978-3-030-15032-7_88

Modelling Energy Consumption — SimGrid

- Simulating execution traces
  → Prediction of dissipated energy
  - Worker machine computations (DVFS, utilised cores) ✓
  - Network traffic ✓
  - Retrieving and storing data ✗
  - Infrastructure, i.e. Cooling ✓
- Powering with solar panels and batteries ✓

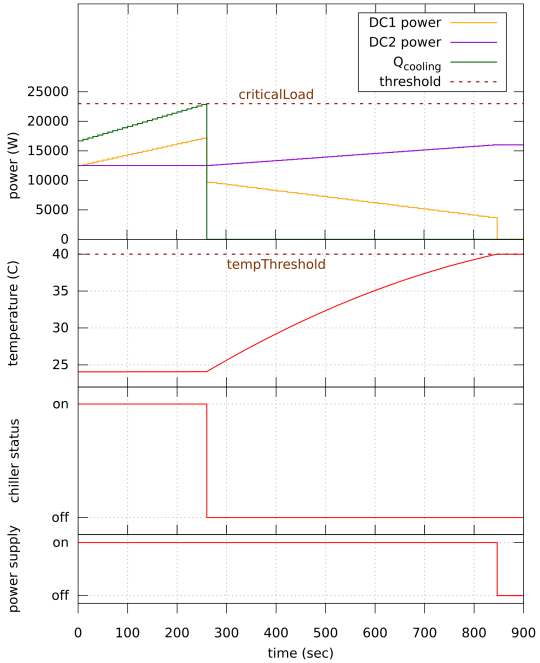

10.1145/3200921.3200932

Modelling Services — WRENCH

- Concept of jobs that bundle a collection of activities that need to be executed
- Scheduler services distribute jobs on hardware
- Activities on hardware are managed by services & execution controllers
  - Computing services manage core usage
  - Storage services manage read/write actions
  - Transfer services manage data communications
  - Execution controllers start and stop processes

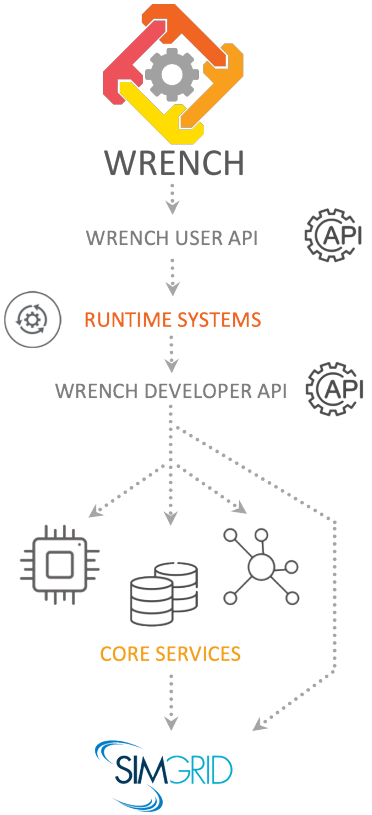

Modelling the Real-world — DCSim

- Definition and implementation of computing model
  - Job definition, mix, & scheduling
  - Dataset definition
  - Data & data transfer management
    - Location, Streaming, Caching
  - Workflows
  - Job Monitoring
- Because built with SimGrid & WRENCH can simulate energy consumption

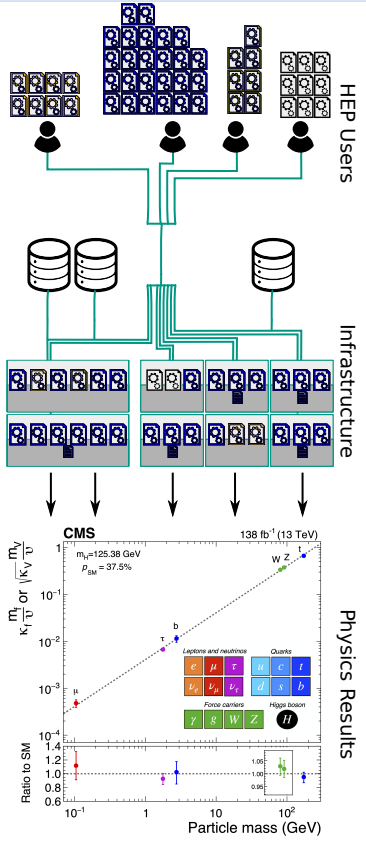

Calibration of Models

- Models are only abstractions with parameters
- Limited knowledge about real-world systems
- Don’t expect out-of-the box agreement, but learn what is missing from real-world data
  - Tuning of simulation models & validation with independent data split
  - If tune generalises, also works for other platforms

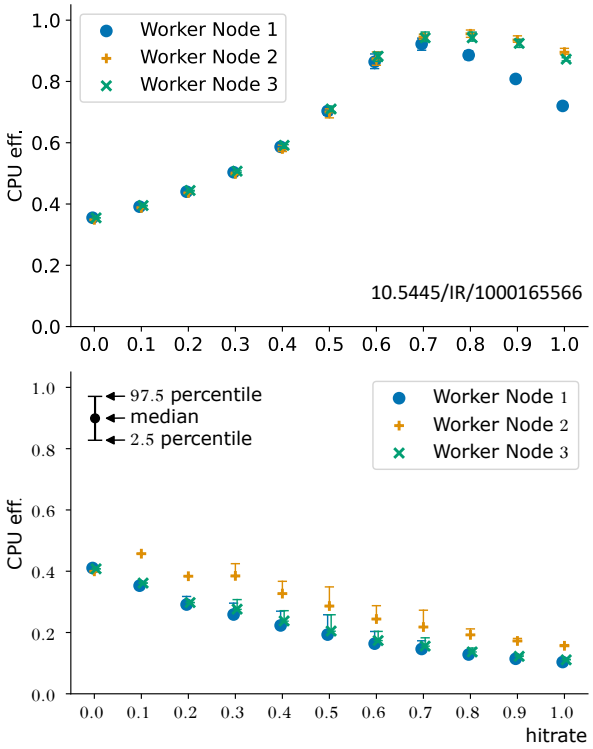

Real-world Data

- No freely accessible repository of real-world data
- But HEP experiments store job monitoring data for jobs on WLCG sites
  - CPU time & efficiency, Bytes read & written, …
  - Estimates of energy consumption per job

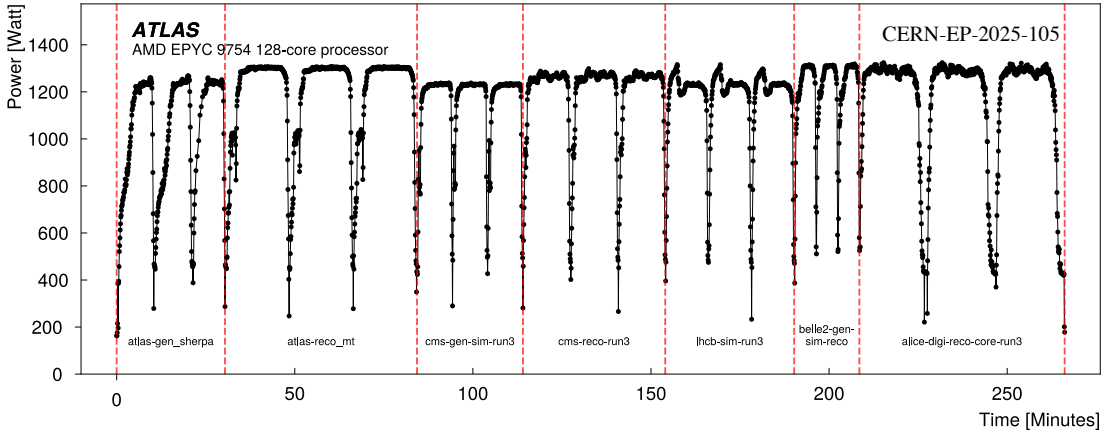

Conclusions

- We must optimise computing infrastructures (and software) for maximum compute efficiency and minimum energy consumption or CO2 emissions!
- Distributed computing systems are too complex for gut feeling or perceived experience
- We have the data and software to construct credible simulation models
- Let's experiment with future infrastructures already now!