SPEEDING UP SCIENCE THROUGH PARAMETRIC OPTIMIZATION ON HPC CLUSTERS

Speaker: Jonas Weßner

Co-authors:
Rüdiger Berlich,
Matthias F. M. Lutz,
and Kilian Schwarz

h_da
HOCHSCHULE DARMSTADT
UNIVERSITY OF APPLIED SCIENCES

GSI

TABLE OF CONTENTS

OUR RESEARCH AT GSI

Parametric Optimization for Quantum Physics

Problem:
Lattice Quantum Chromodynamics (QCD) simulations are computationally expensive

Goal:
Effective quantum field theories for QCD

Method:
Transformation to effective degrees of freedom in terms of unknown parameters (LEC)

Challenge:
How to determine the LECs efficiently through optimization?

OUR RESEARCH AT GSI

THE GENEVA OPTIMIZATION LIBRARY

Large-scale Parametric Optimization in Distributed Environments

PARAMETRIC OPTIMIZATION WITH GENEVA

github.com/gemfony/geneva

gemfony.eu

- Open-source framework in C++ for parametric optimization
- Focus on distributed execution: grids, clouds, **clusters**
- Used at GSI for QCD research
- **Recently:**
  MPI Consumer for distributed execution on HPC clusters:
  https://doi.org/10.1007/s41781-023-00098-6
- **Today's Presentation:**
  Performance of MPI Consumer

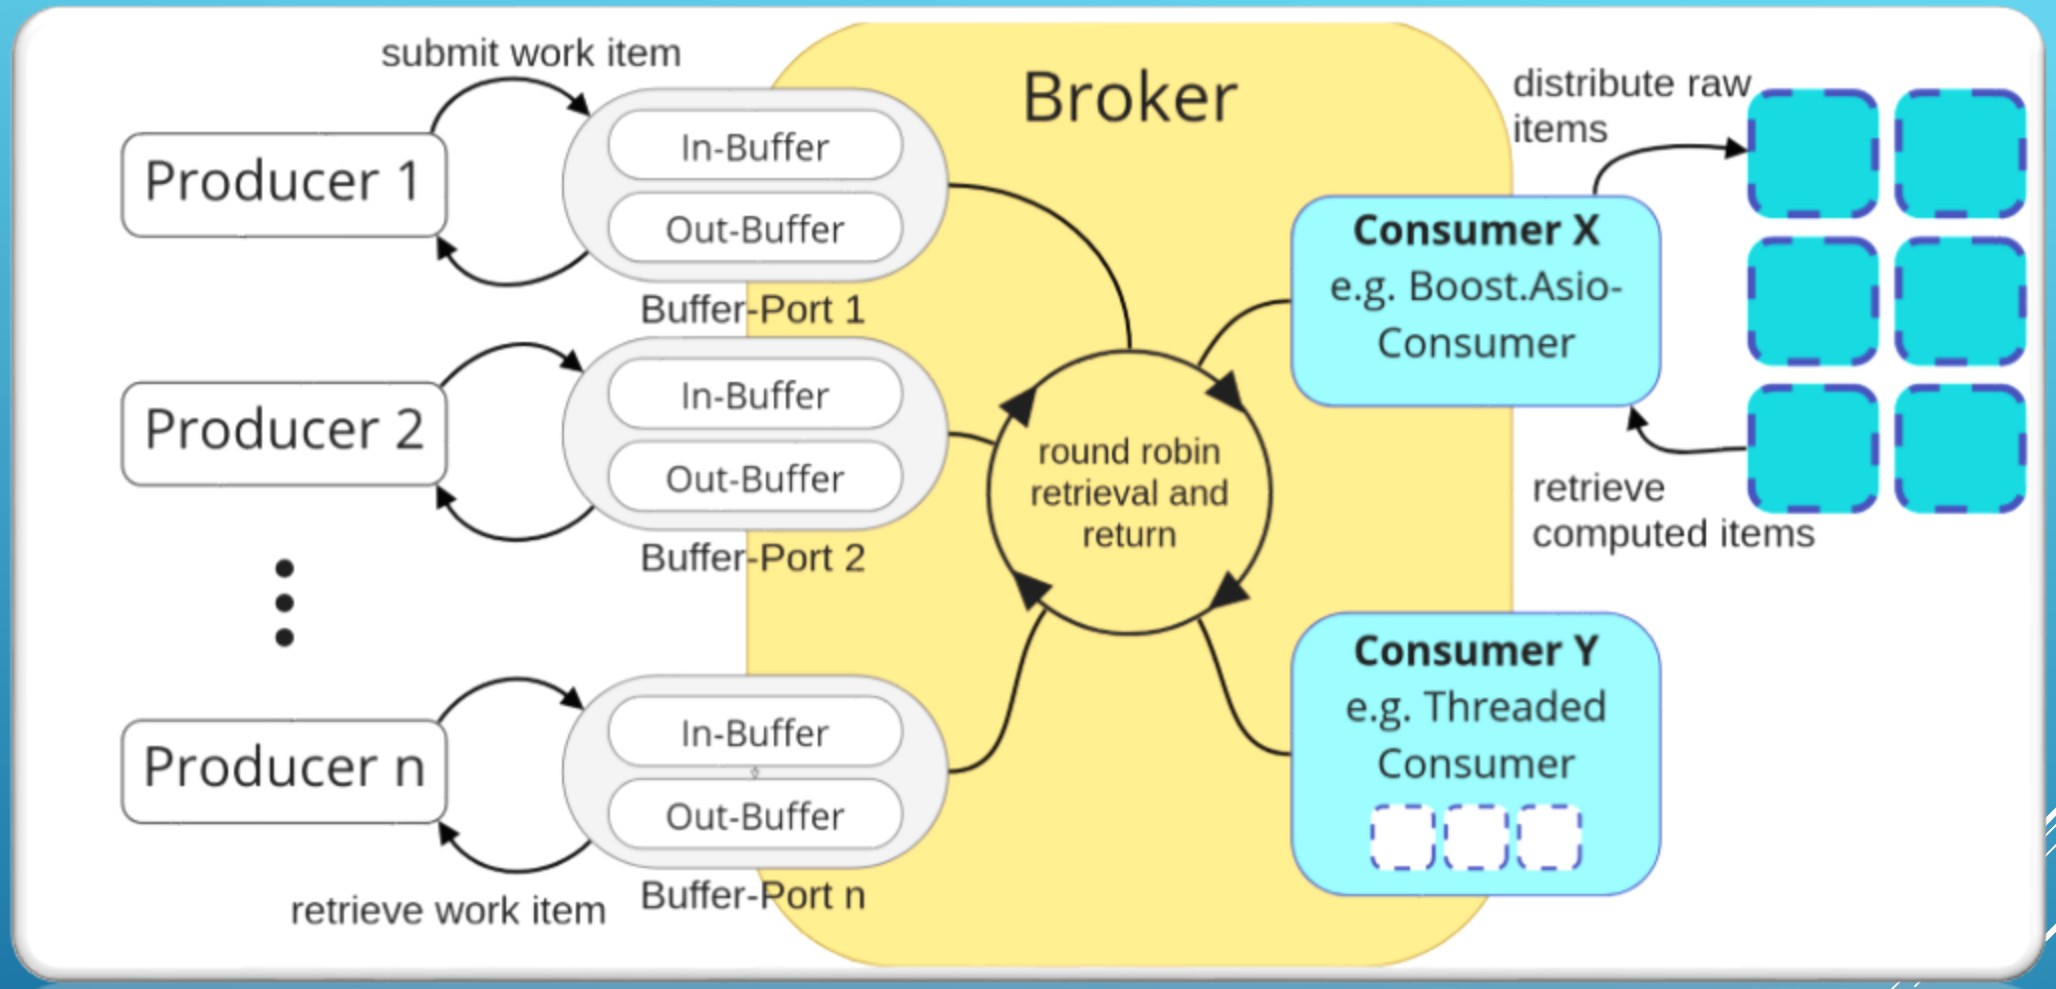

PARALLELIZATION IN GENEVA

Graphic taken from our paper: *https://doi.org/10.1007/s41781-023-00098-6*

PERFORMANCE OF THE MPI CONSUMER

Comparison with Boost.Asio and Boost.Beast Consumers

On a Single 128-core Machine

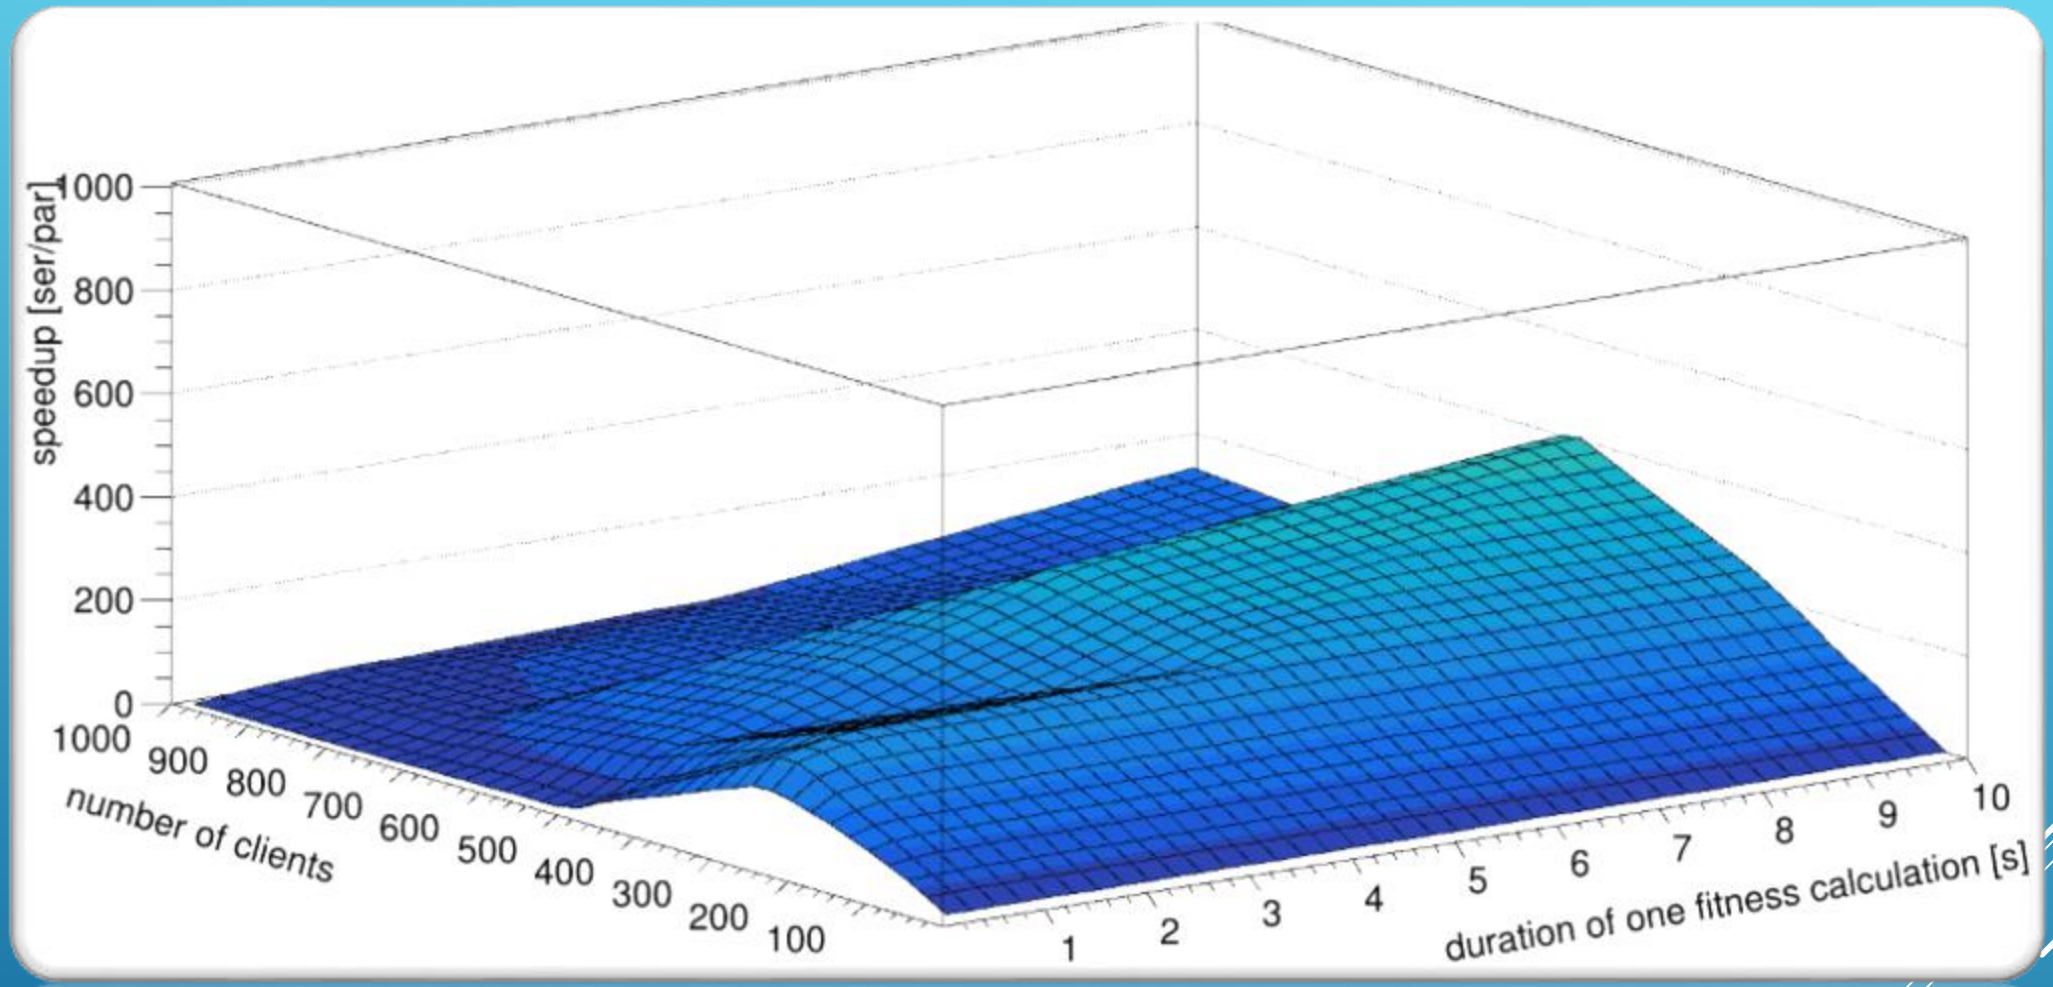

BOOST.ASIO CONSUMER

- LIMITED SCALABILITY
- TYPICALLY AT GSI: 5+ SEC EVALUATION TIME

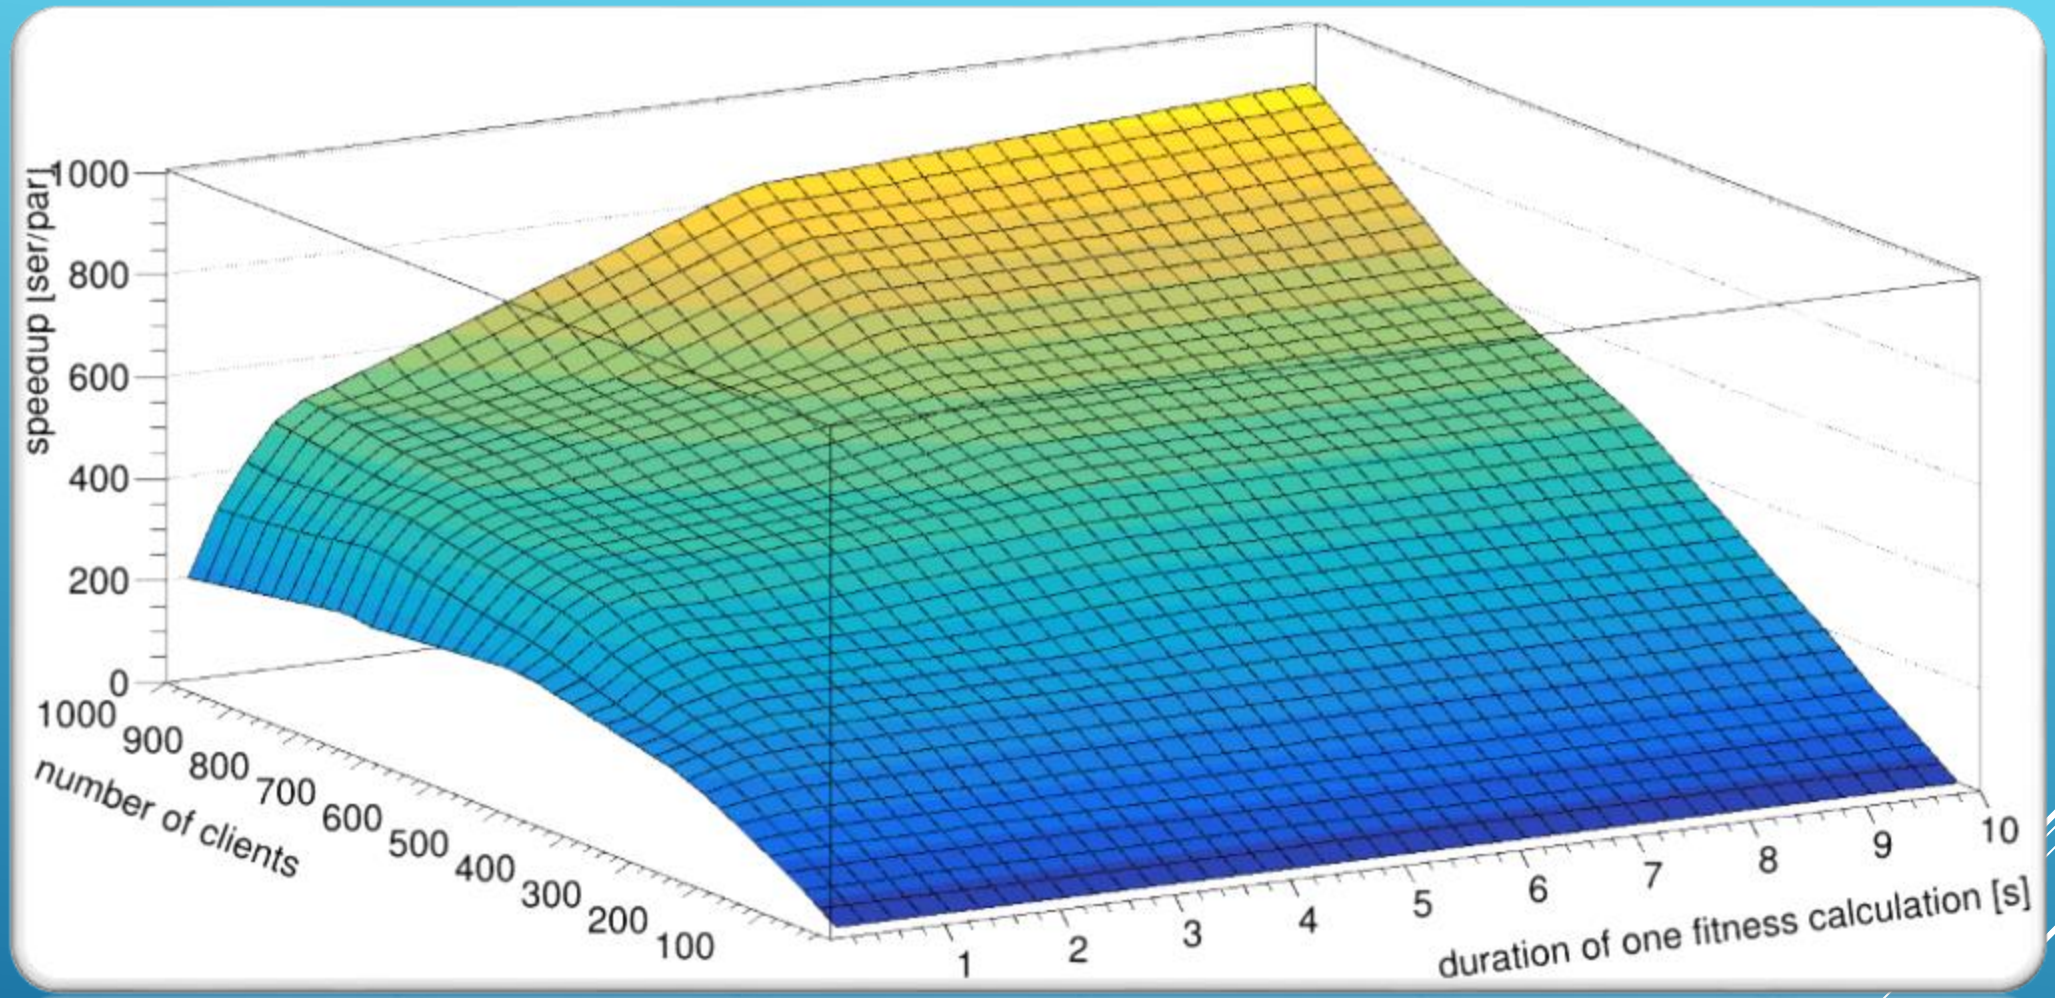

BOOST.BEAST CONSUMER WITH 64 THREADS

IMPROVED SCALABILITY

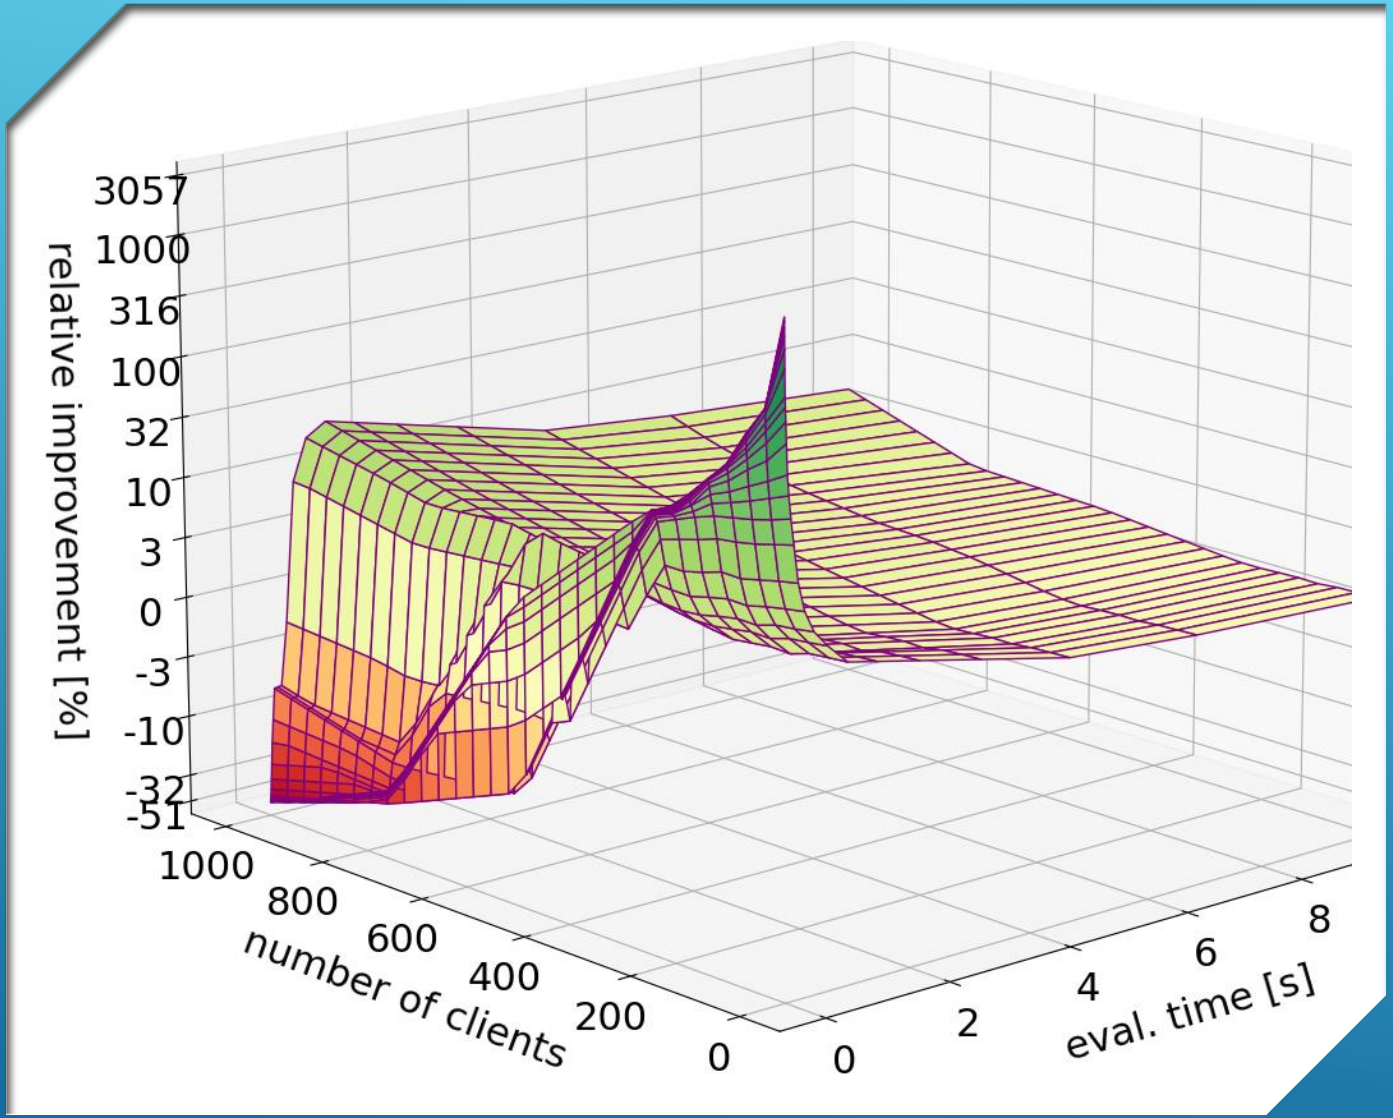

MPI VS. BEAST

- IMPROVEMENTS FOR ALL RELEVANT WORKLOADS

MPI CONSUMER PERFORMANCE

On GSI's *Green IT Cube* HPC Cluster

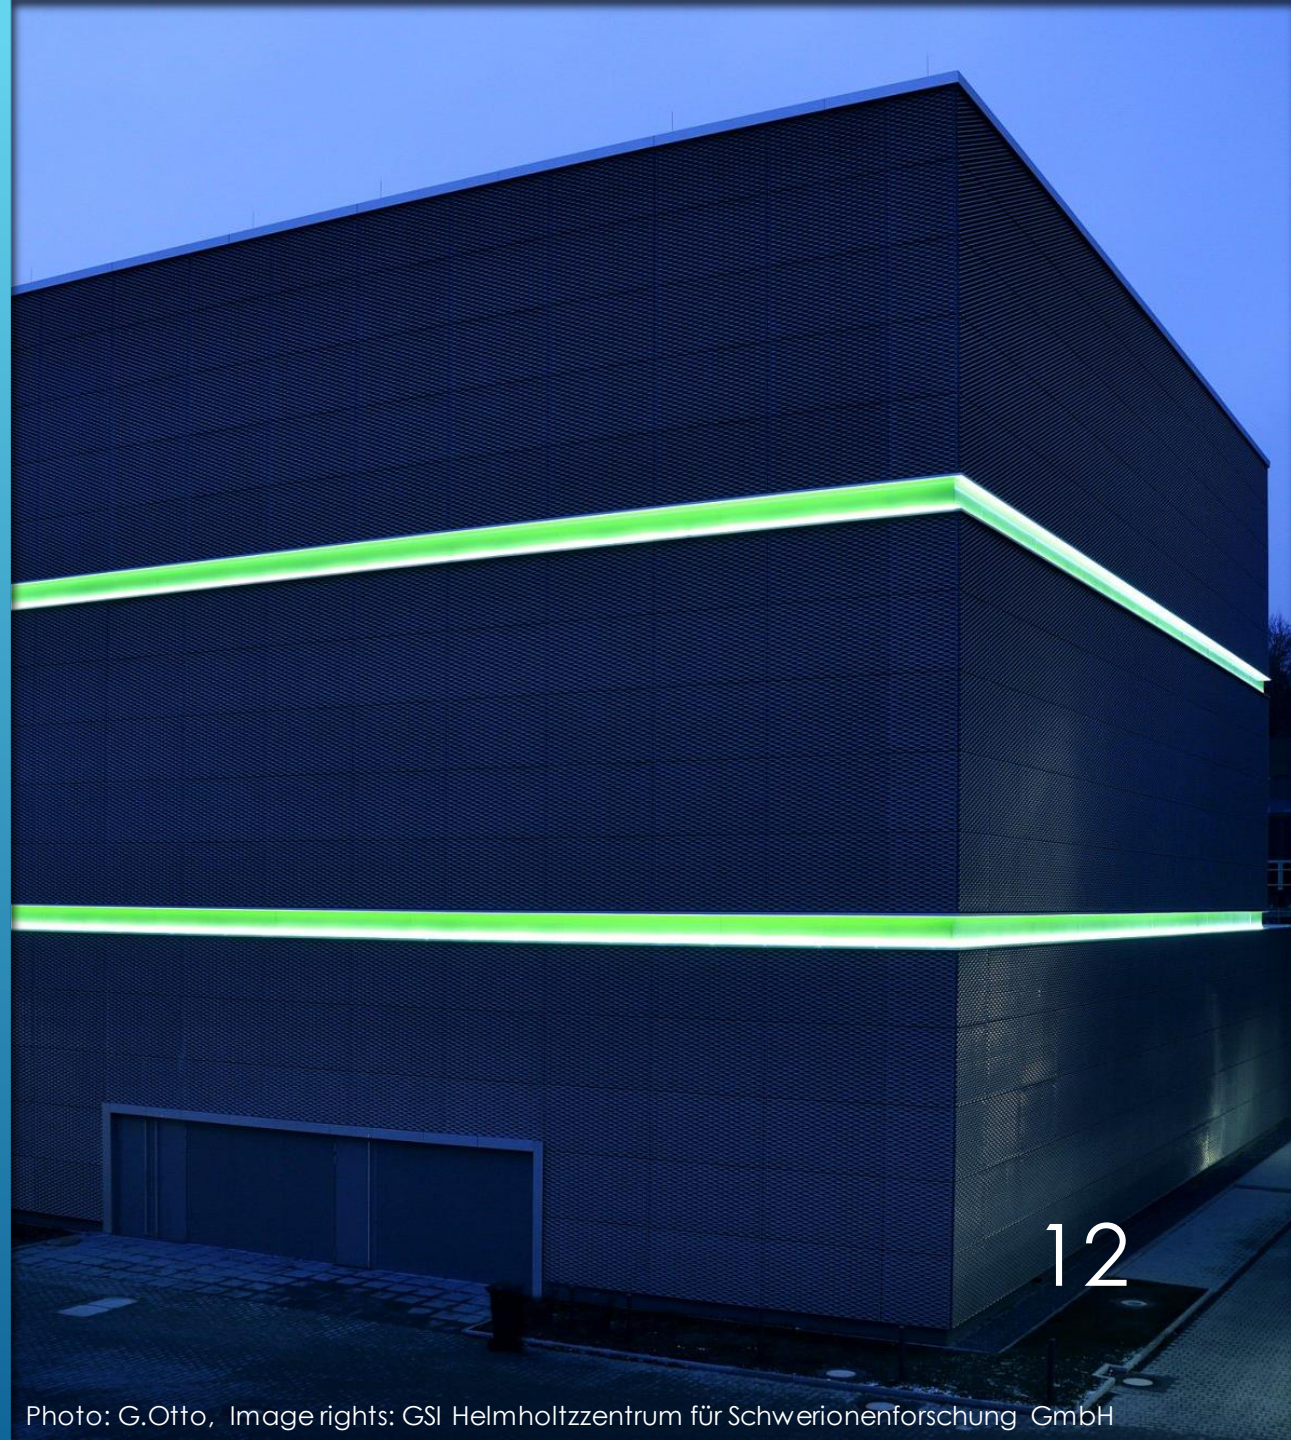

Photo: T. Ernsting, Image rights: Hessen schafft Wissen

TESTS ON HPC CLUSTER

- GSI's green IT cube:
  - >400 nodes
  - 56-256 cores and 128GB-1TB RAM each
  - Slurm workload manager
- Test setup:
  - 400 distributed clients
  - Mean over 10 measurements
  - -> Compare results with results on multicore machine

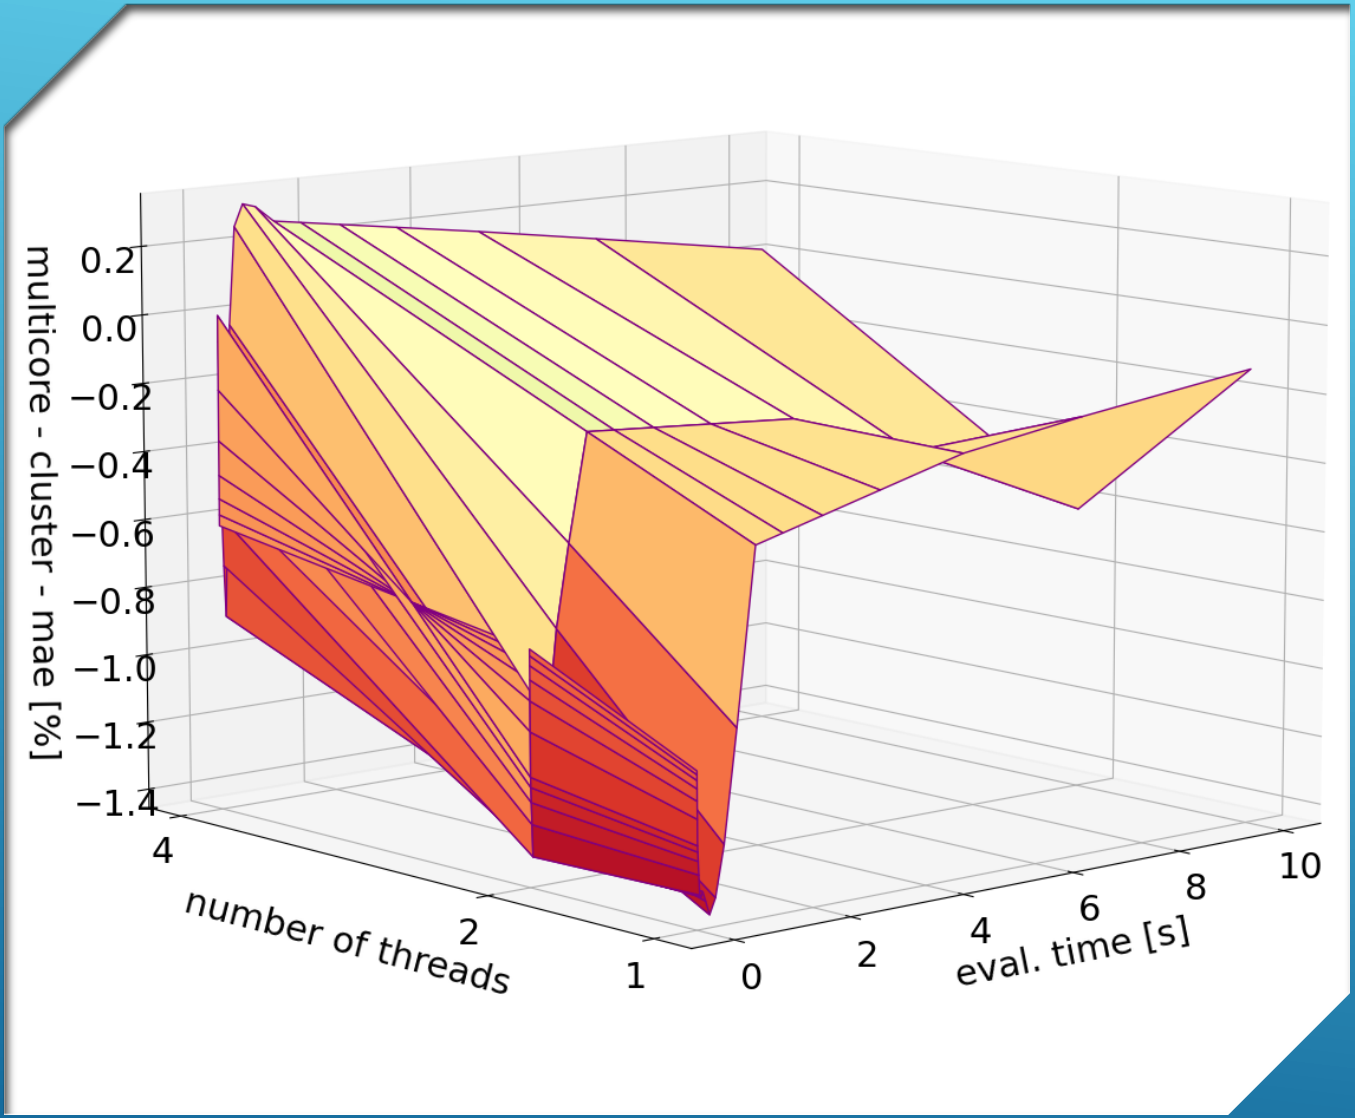

DIFFERENCE OF 128-CORE MACHINE AND CLUSTER

MEASUREMENTS ARE REPRODUCIBLE ON CLUSTER WITH MINIMAL ERROR

- Geneva optimization library handles large-scale problems
- Better user experience for HPC because of MPI Consumer
- MPI Consumer system design to be published in *Computing and Software for Big Science*: https://doi.org/10.1007/s41781-023-00098-6
- Improved Scalability:
  Even more efficient than ASIO/BEAST Consumers

SUMMARY

FUTURE WORK

- MPI Consumer to be brought in production on GSI's Green IT Cube cluster with 1000+ clients
- More optimization algorithms and metrics for Geneva optimization library
- Paper about performance evaluation in *Proceedings of Science* (PoS)

Paper about design details:

BACKUP SLIDES

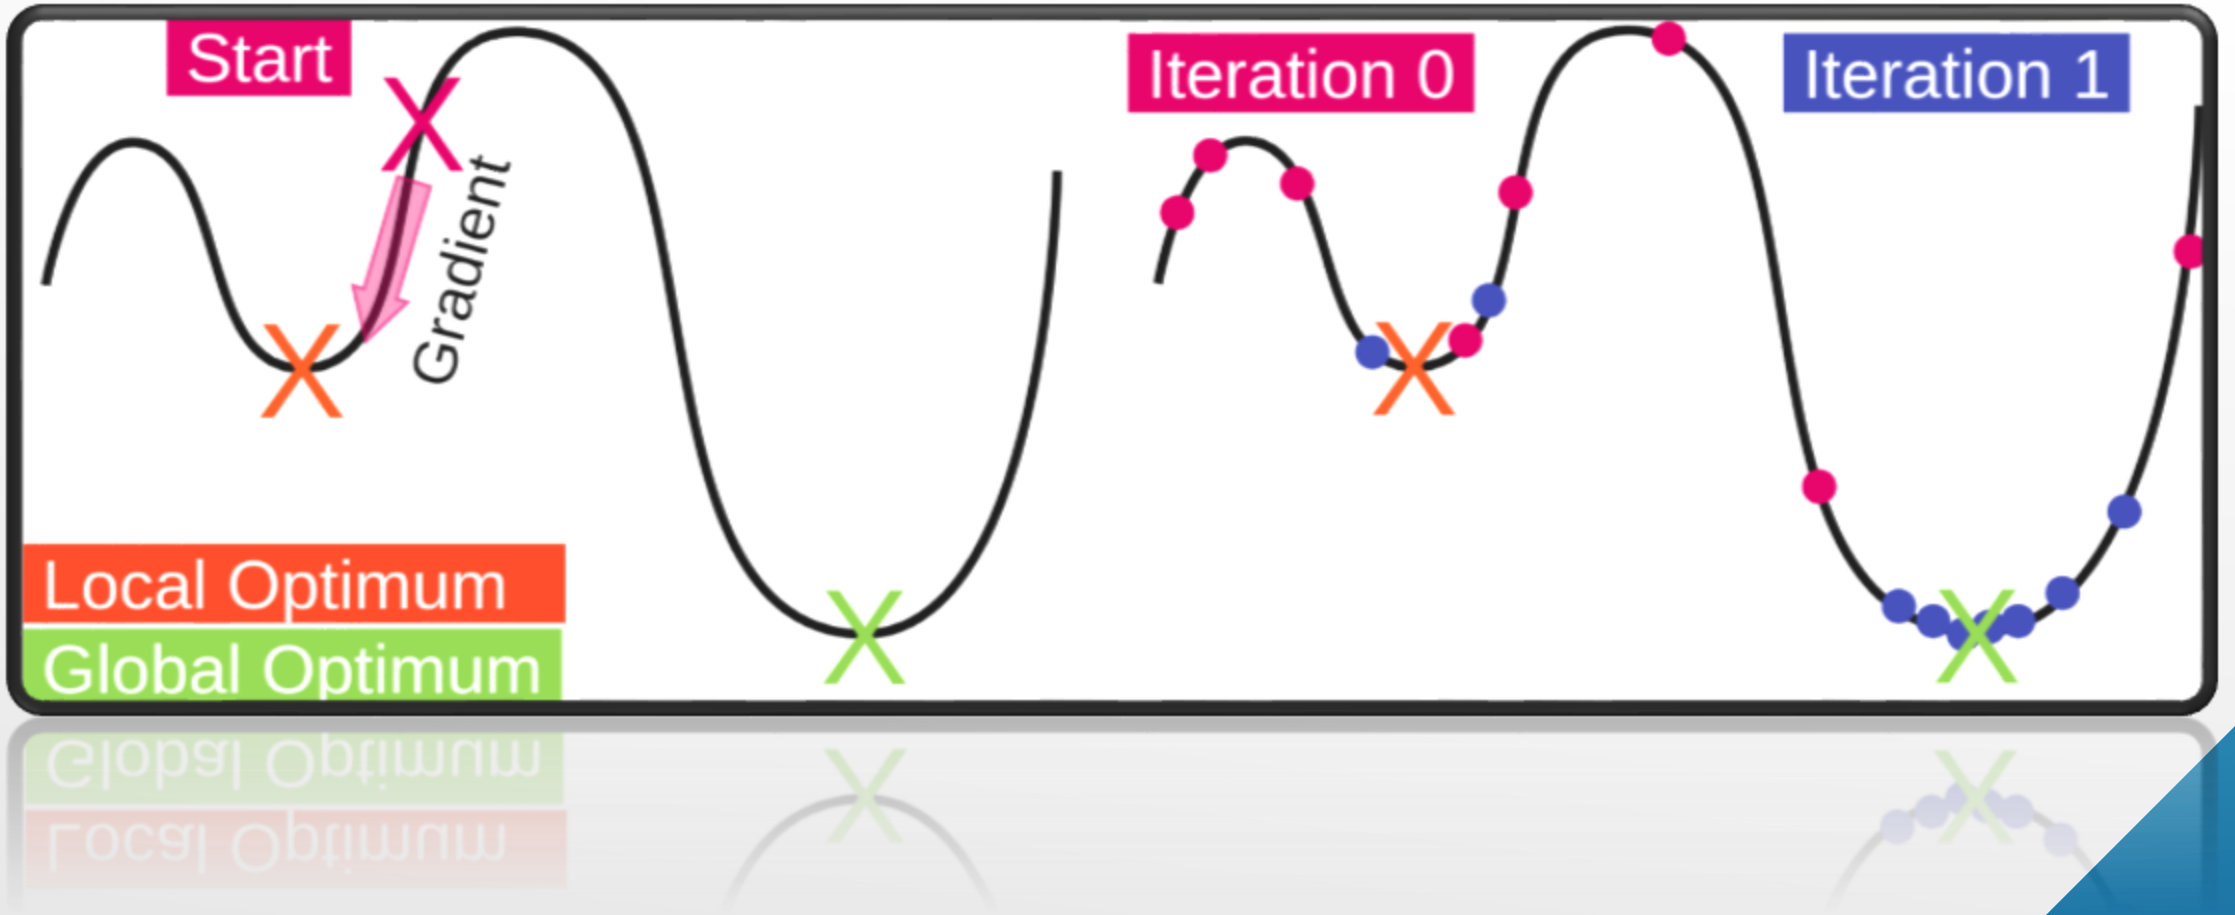

PARAMETRIC OPTIMIZATION

FROM SIMPLE TO ROUGH SURFACES

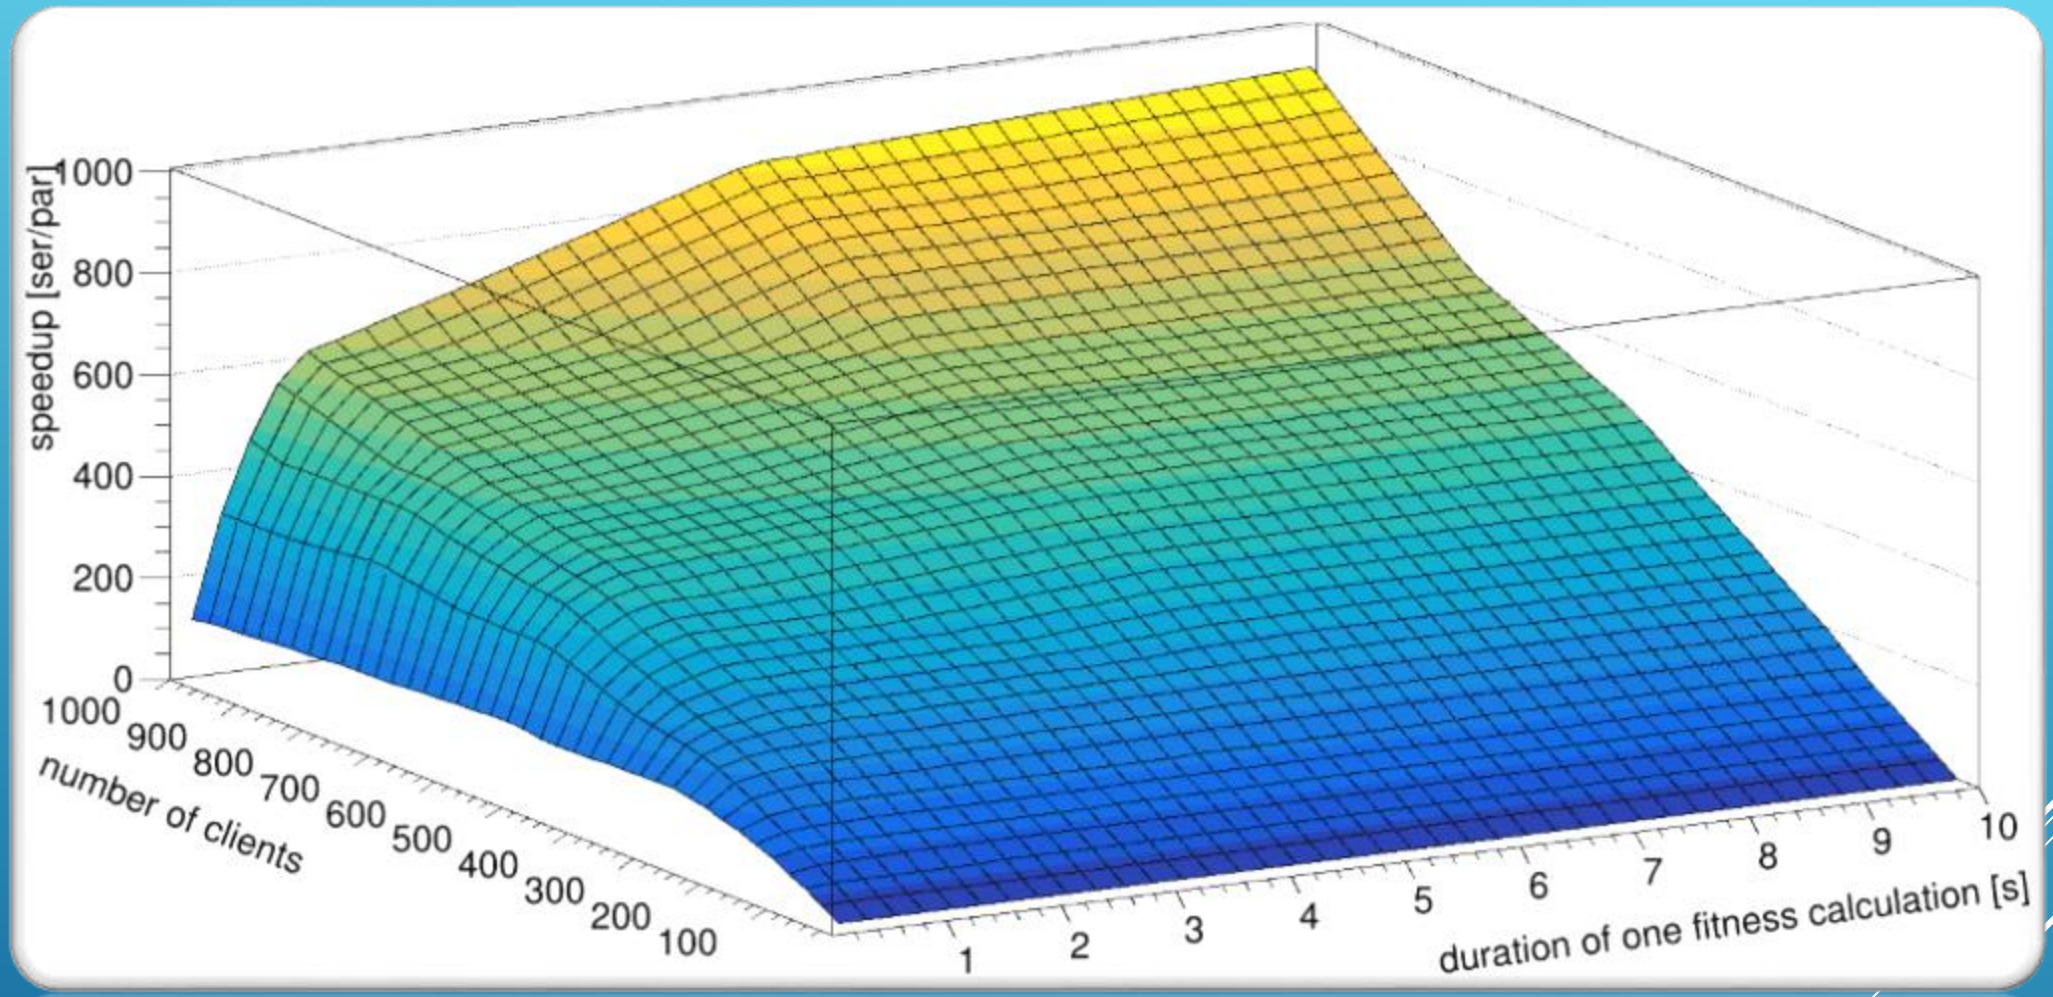

MPI CONSUMER SCALABILITY

SERVER WITH 64 THREADS RUNNING ON 128-CORE MACHINE

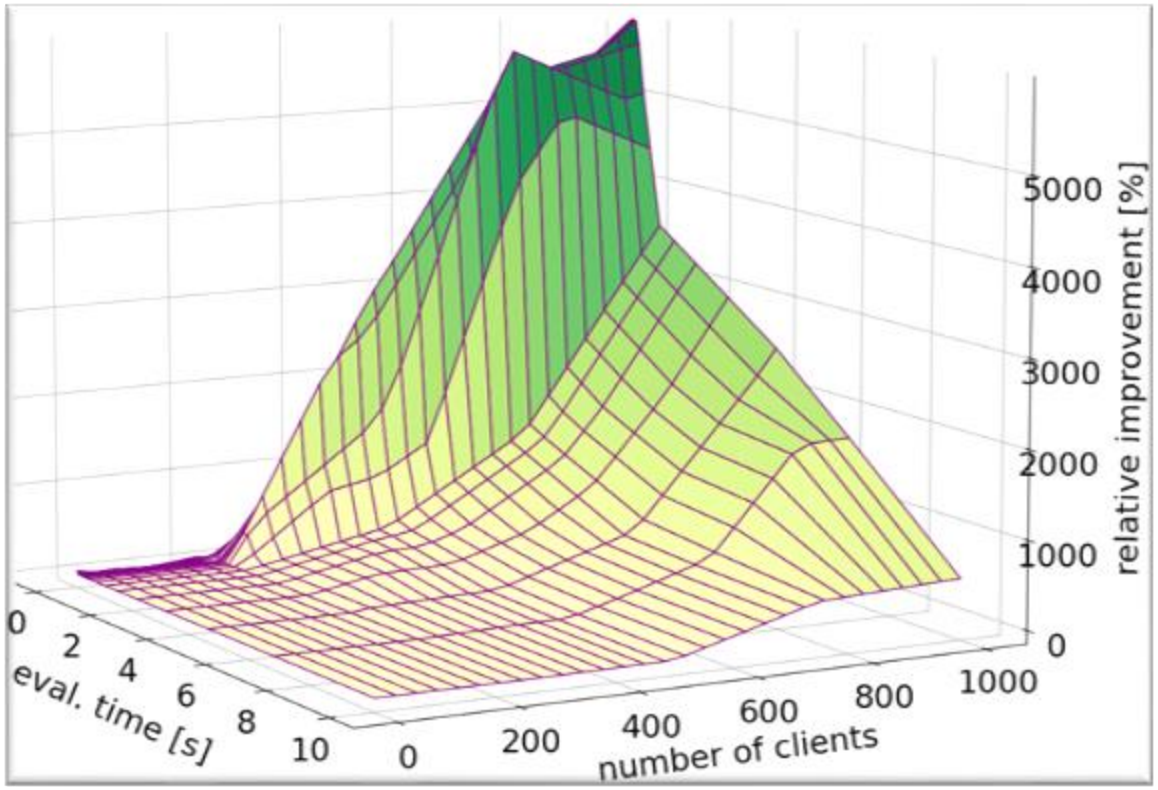

MPI COMPARED WITH BOOST.ASIO CONSUMER

RELATIVE IMPROVEMENT

BENEFIT OF ASYNC CLIENT REQUESTS

RELATIVE IMPROVEMENT

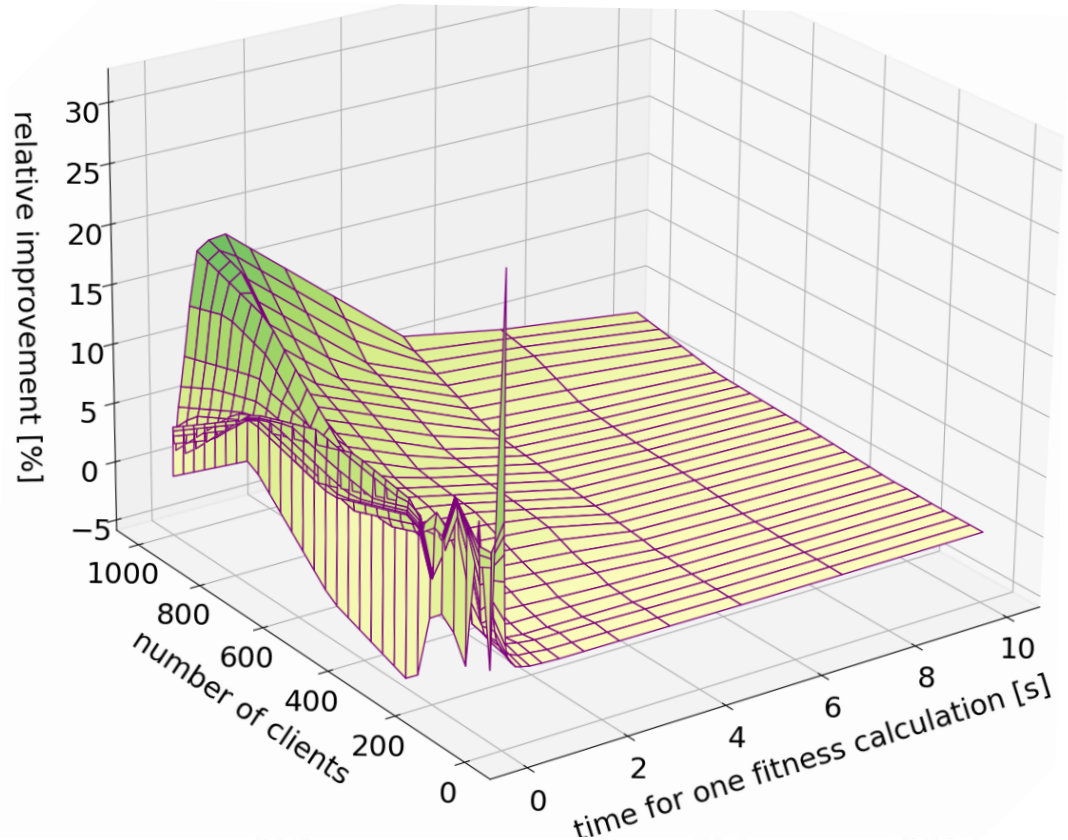

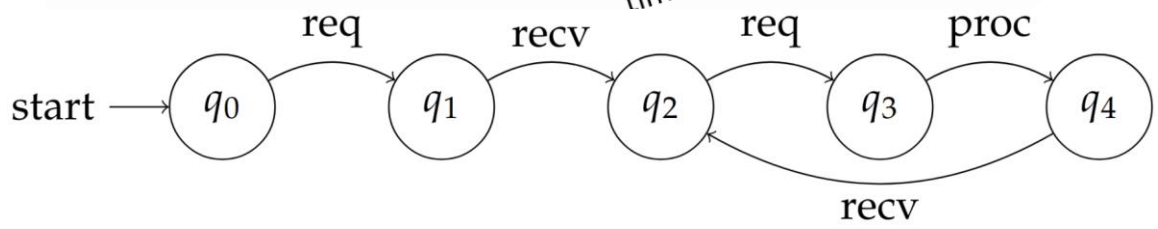

Detailed explanation in our previous publication:
https://doi.org/10.1007/s41781-023-00098-6

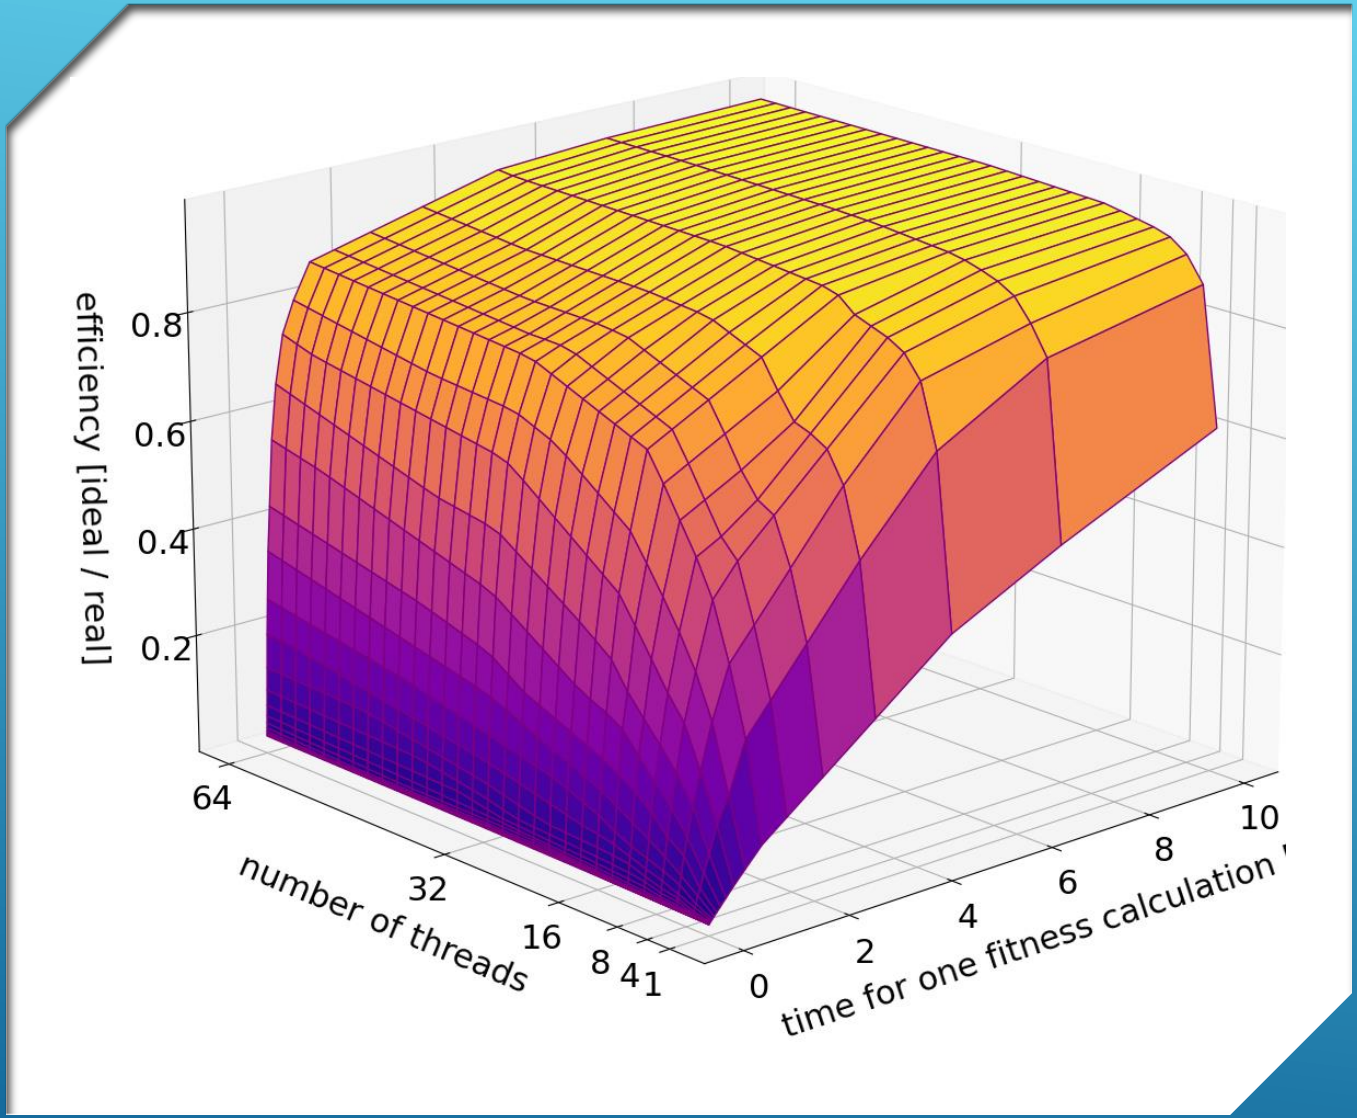

IMPACT OF THREAD POOL SIZE

SYSTEM EFFICIENCY

Detailed explanation in our publication:
https://doi.org/10.1007/s41781-023-00098-6

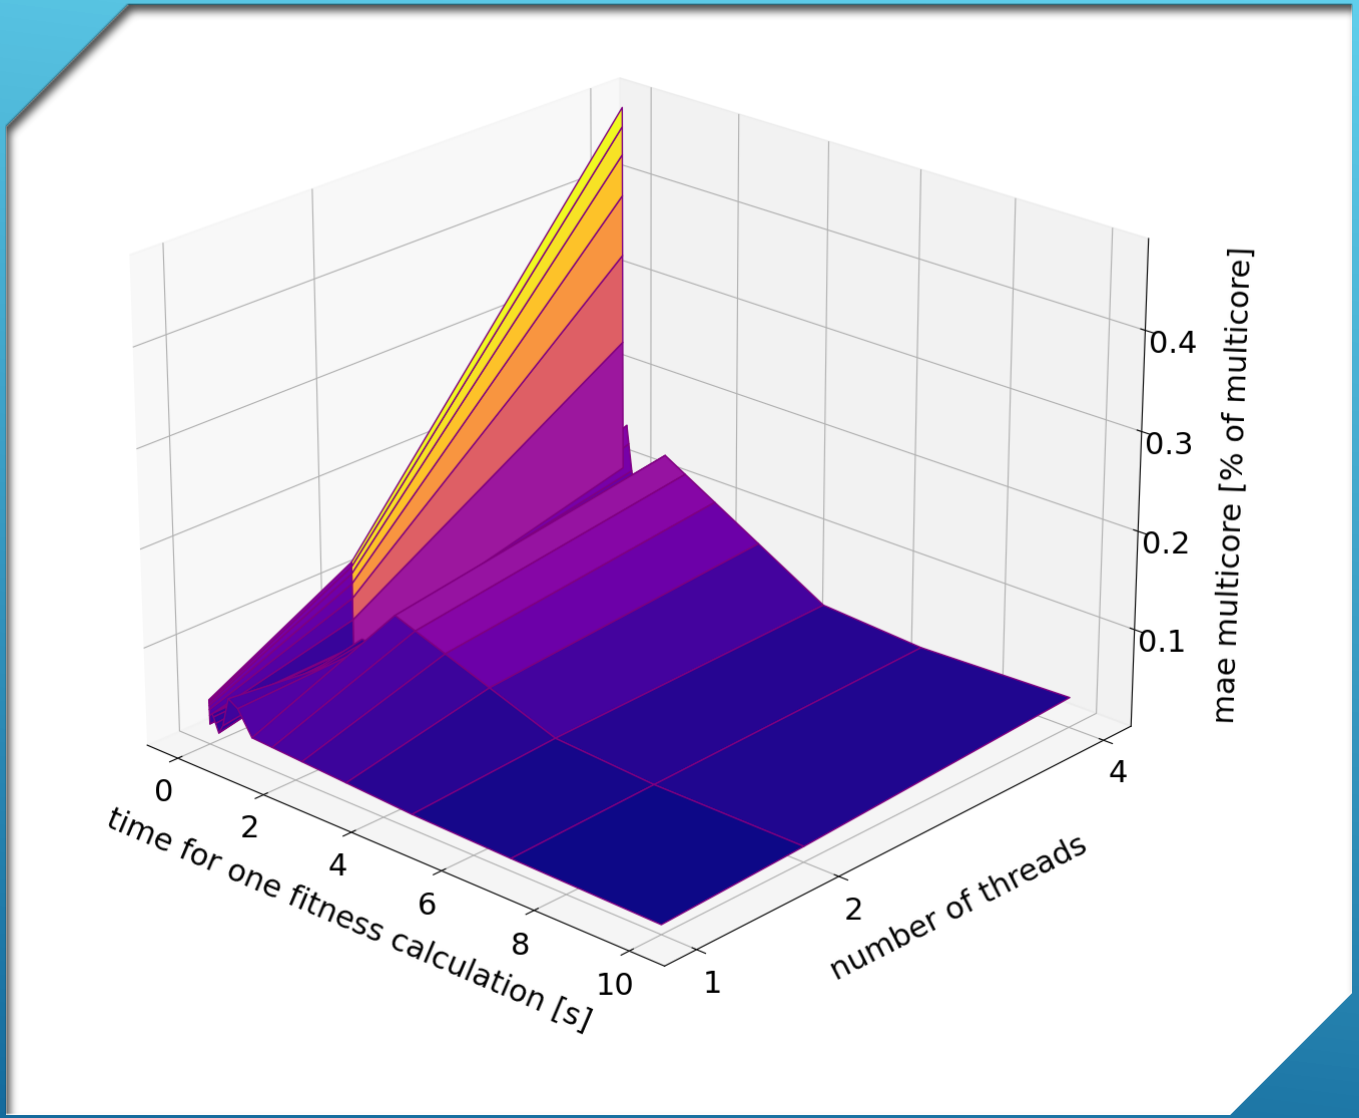

MEAN ABSOLUTE ERROR ON 128-CORE MACHINE

- 10 MEASUREMENTS PER DATA POINT
- 400 CLIENTS

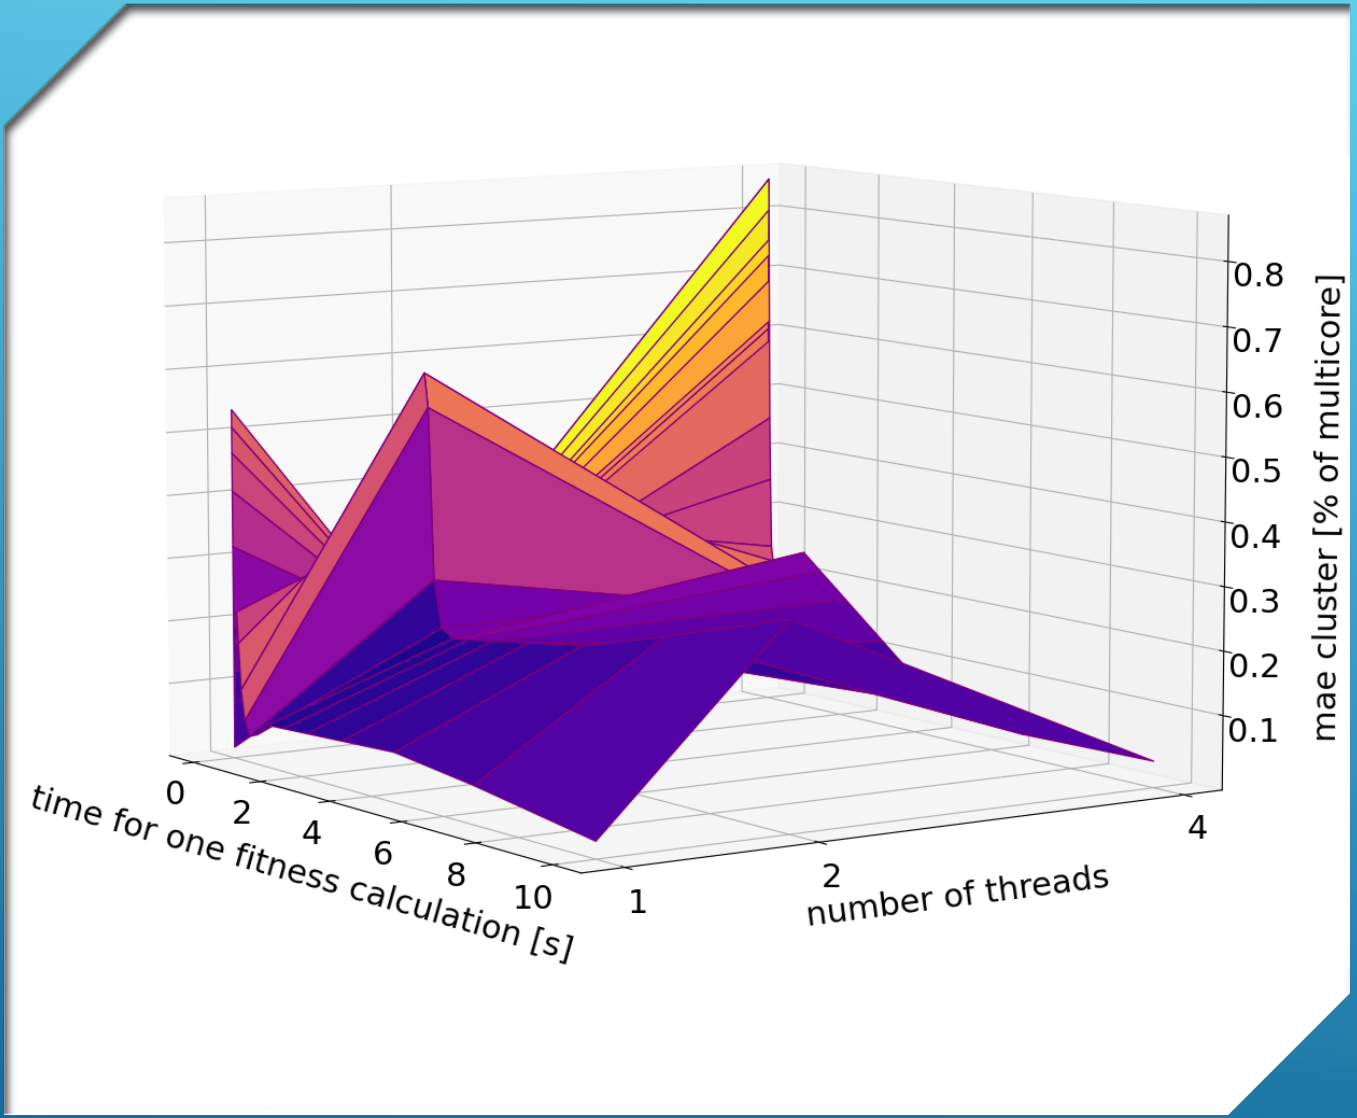

MEAN ABSOLUTE ERROR HPC CLUSTER

- 10 MEASUREMENTS PER DATA POINT
- 400 CLIENTS