



# Adaptive multi-tier intelligent data manager for Exascale




admire-eurohpc.eu

## ADMIRE User Day

# Extra-P With I/O Modeling Capabilities

**Ahmad Tarraf**  
**Technical University Darmstadt**

**December 12<sup>th</sup> 2023**  
**Barcelona Supercomputing Center**

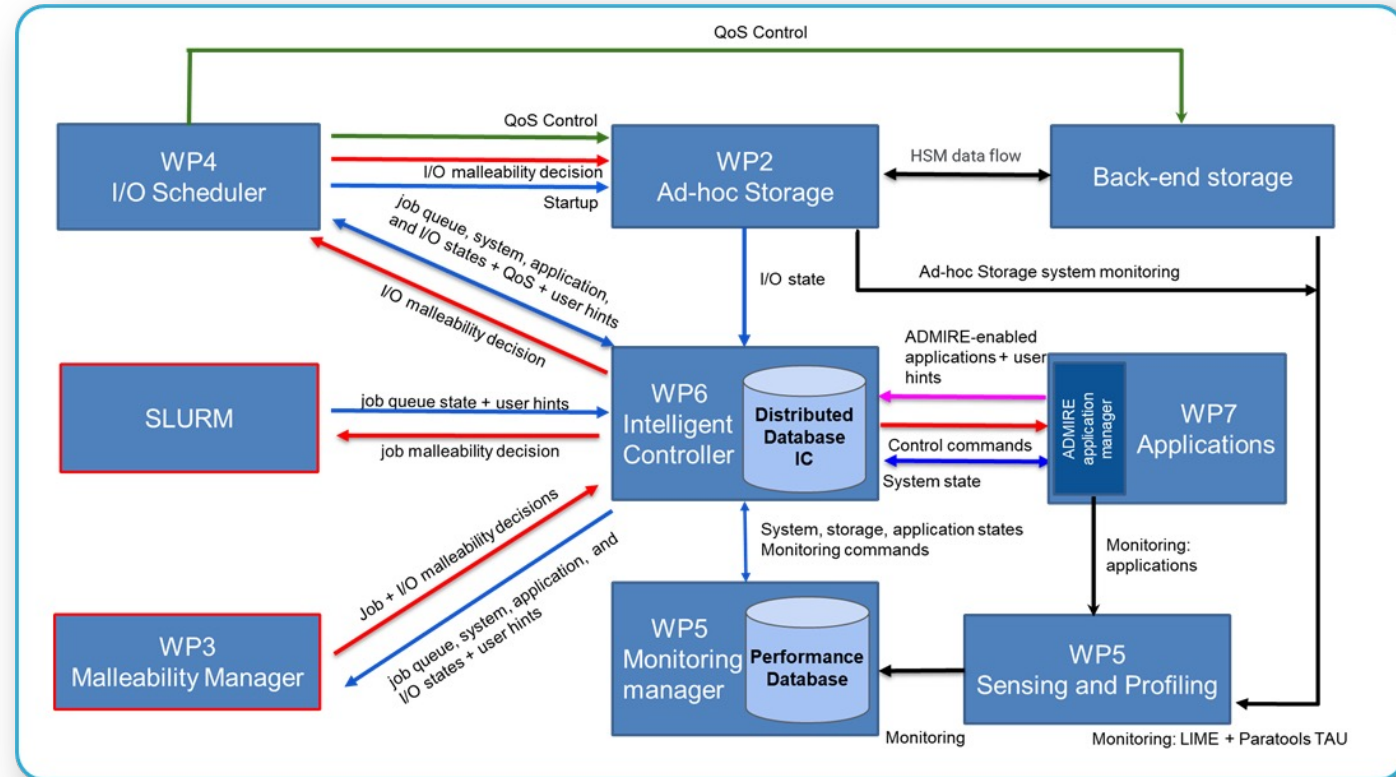


Grant Agreement number: 956748 — ADMIRE — H2020-JTI-EuroHPC-2019-1



## ADMIRE global architecture:

- Intelligent controller acts as the main component forwarding and processing data
- Key aspects examined in the ADMIRE project is **job malleability**
- Malleability manager needs:
  - Current state of the system
  - Jobs in the queue
  - User hints
  - **Scaling performance of an application**



Objective: balance both compute and I/O resources

→ Performance models that cover the scaling behavior applications in both regards

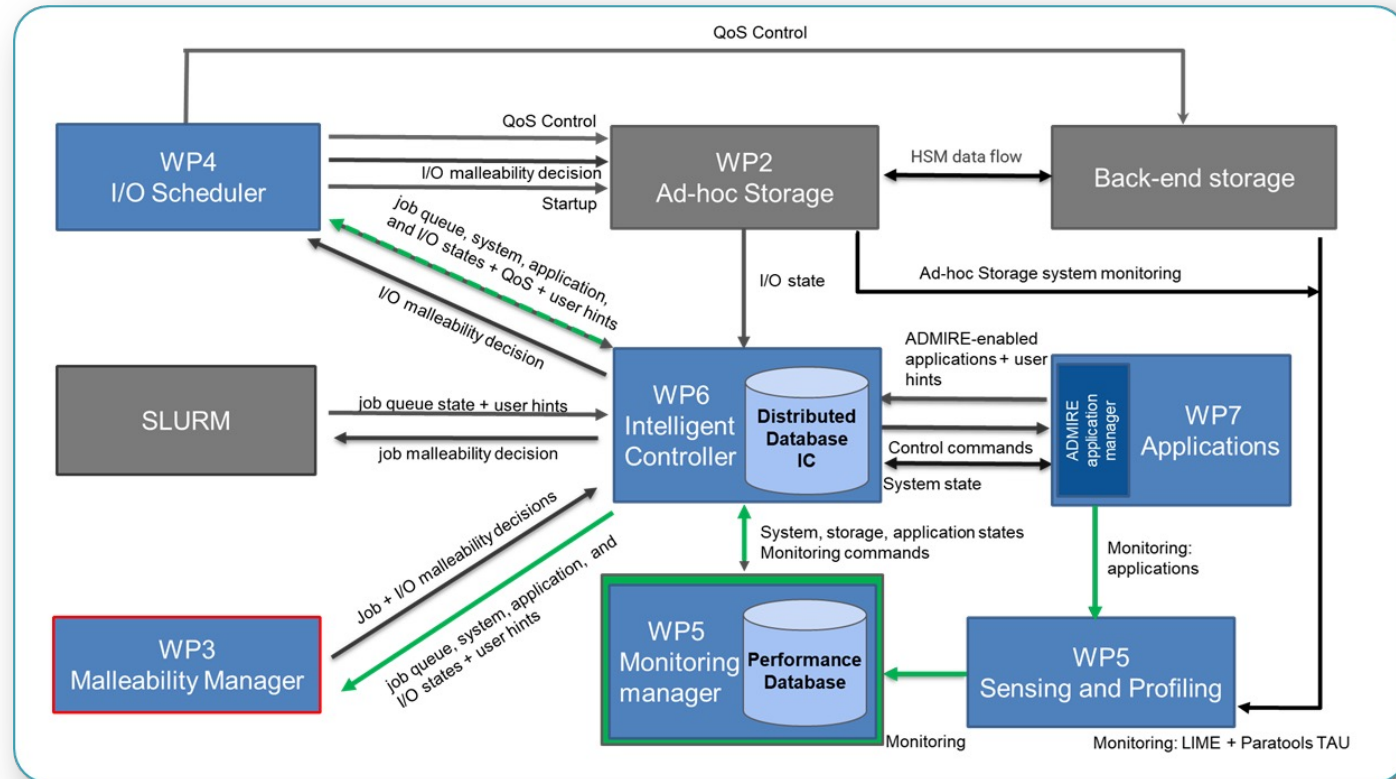
## Generate performance models with Extra-P

- Components that directly influence or interact with the performance models from Extra-P are colored (in blue)
- Data flow of the models is colored in green

So far, Extra-P focuses on modeling the scaling behavior of computational and communicational parts of an applications

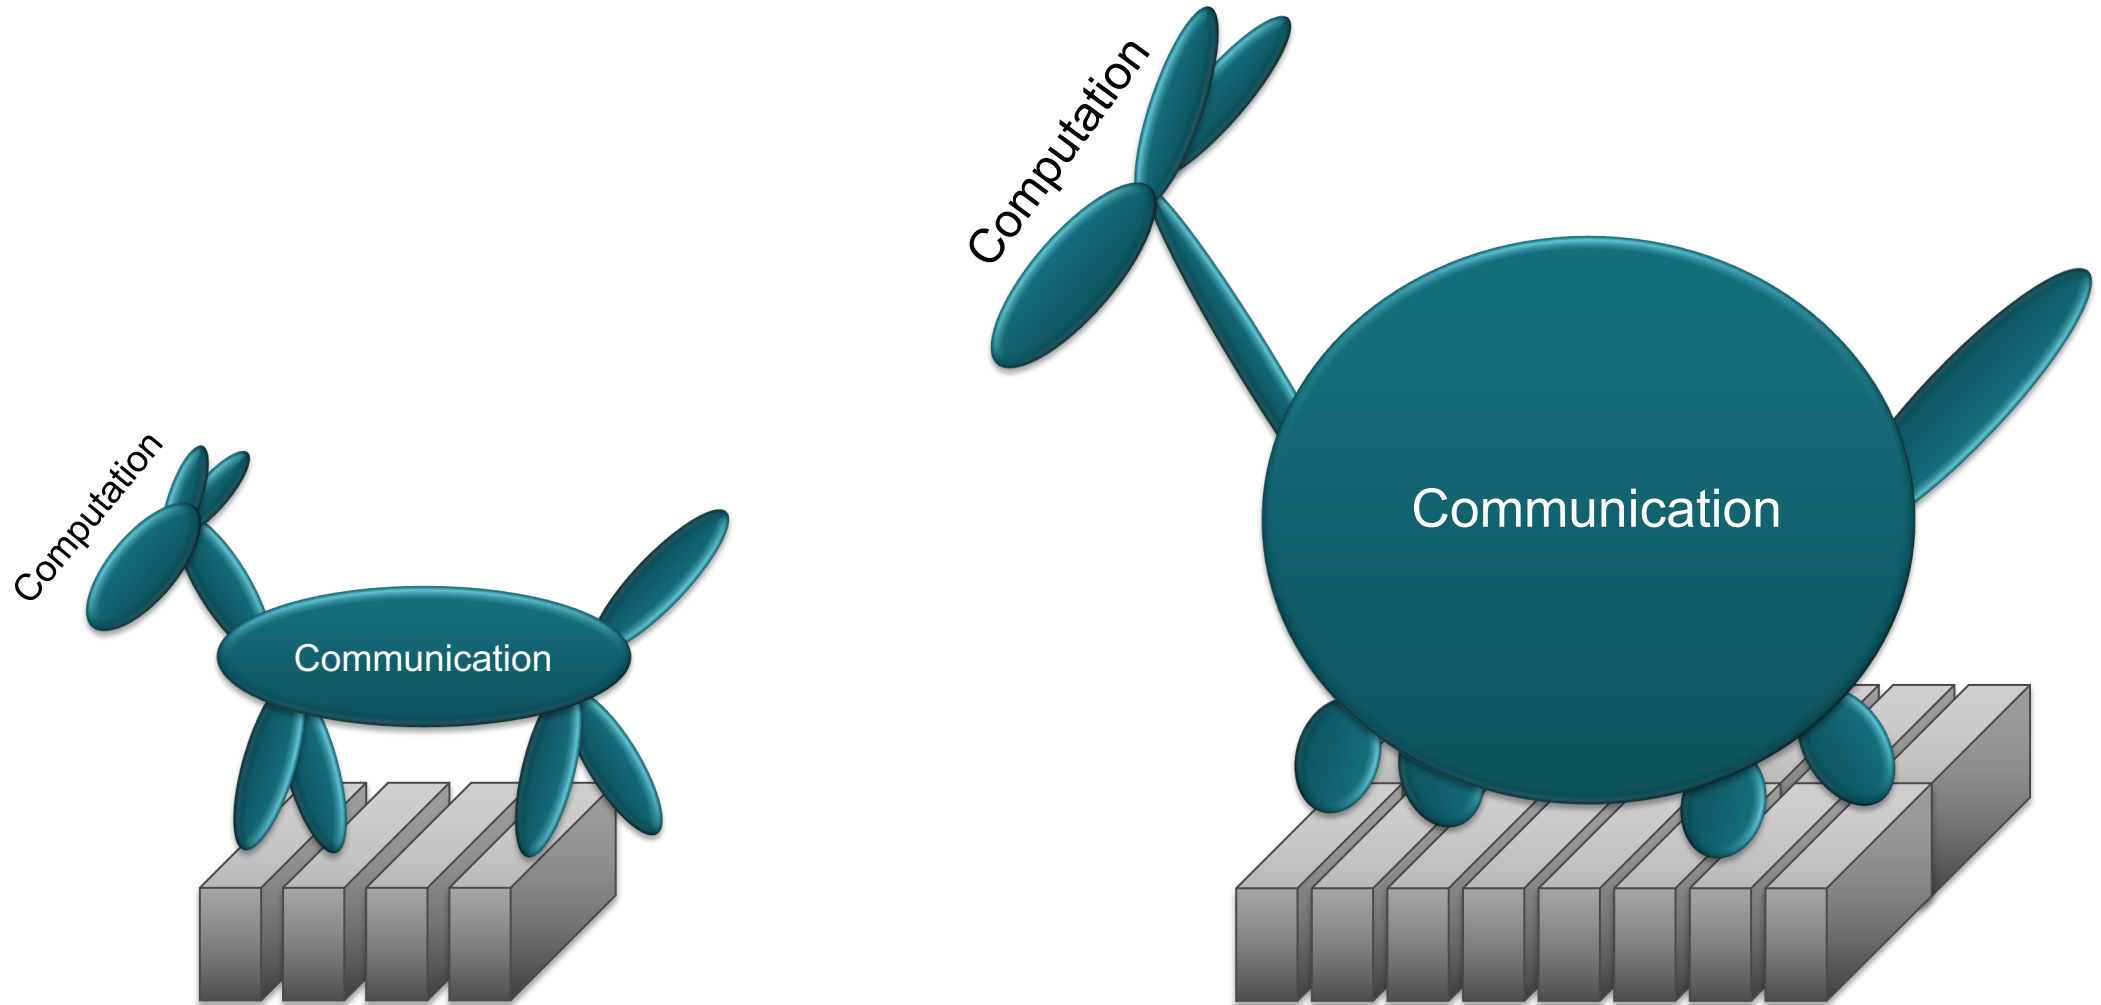
Extend Extra-P to model I/O!

But before, let's have a quick look at Extra-P and performance models in general ...




# Scaling your code can harbor performance surprises\* ...

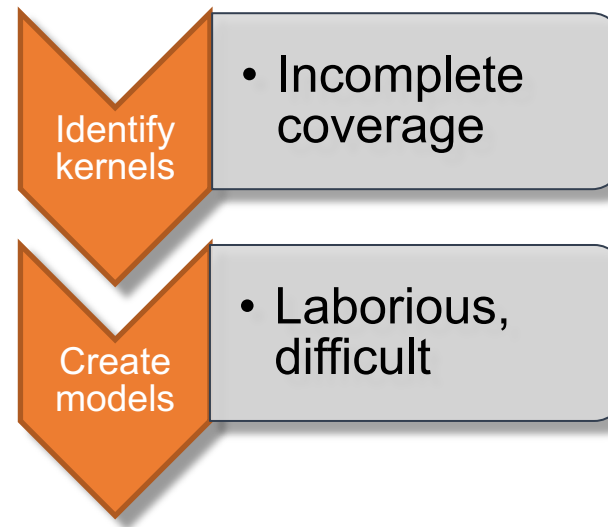
\*Goldsmith et al., 2007



**Performance Model:** Formula that expresses a relevant performance metric as a function of one or more execution parameters



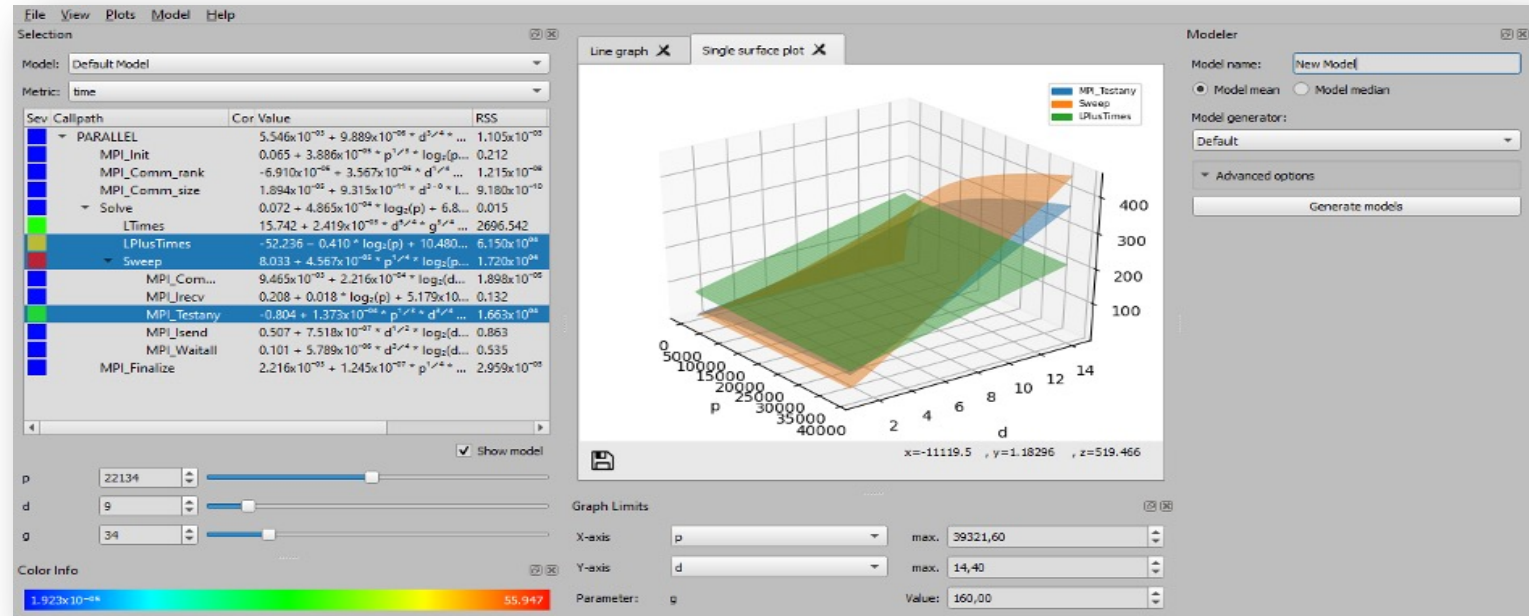
Analytical (i.e., manual) creation  
→ challenging for entire programs



## Extra-P in a nutshell:

- Automatic performance-modelling tool
- Generates **empirical performance models** that predict the scaling behavior of the different parts of an application
- Supports the user in identifying **scalability bugs**

Part of the program whose scaling behavior is unintentionally poor

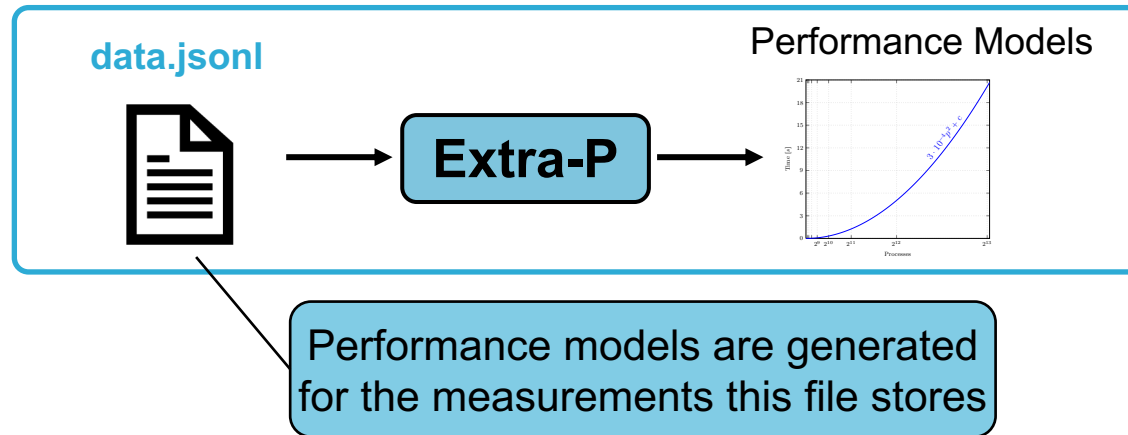


<https://github.com/extra-p/extrap>



<https://youtu.be/Cv2YRCMWqBM>

- In ADMIRE: No human in the loop  
→ Use Extra-P **programming interface**
- Whenever **new traces** are available, just **reinvoke** Extra-P



- **Stores fine graded I/O data in the JSONL file**  
→ new profiles can be simply appended
- **Model fine graded I/O function up till metrics like total I/O time**

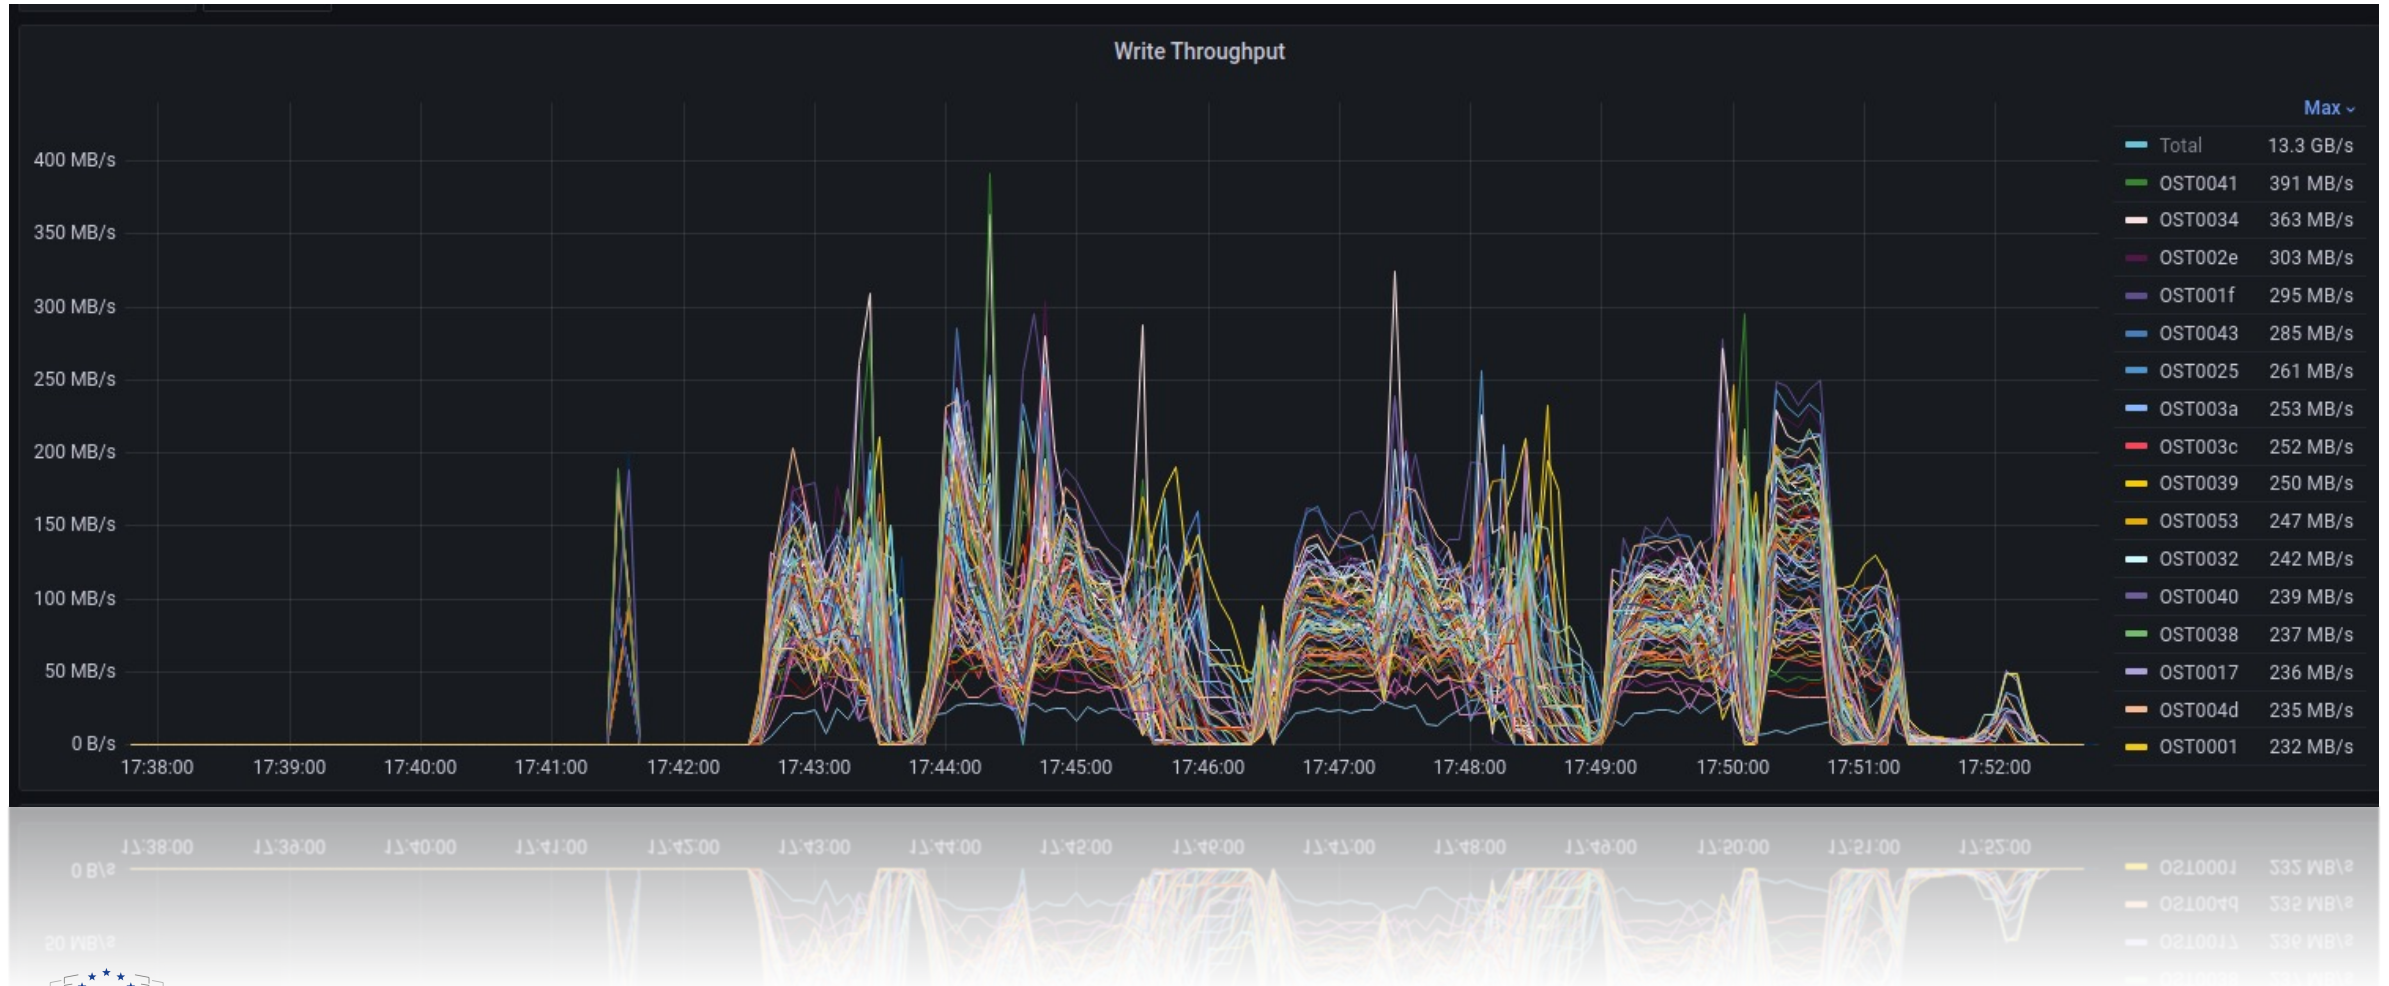
```
python - modeler.py
1 # Imports
2 from extrap.entities.callpath import Callpath
3 from extrap.entities.metric import Metric
4 from extrap.fileio.file_reader.jsonlines_file_reader import read_jsonlines_file
5 from extrap.modelers.model_generator import ModelGenerator
6 import numpy as np
7
8 # Load the JSONL data. For different file formats, there are
9 # different read libraries in extrap.fileio
10 experiment = read_jsonlines_file("./data.jsonl")
11
12 # Initialize model generator
13 model_generator = ModelGenerator(experiment)
14
15 # Create models from data
16 model_generator.model_all()
17
18 # ----- At this point, the models are generated
19 # Next, let's see how to evaluate them
20
21 # 1) Specify the metric and call path, e.g., default and root, respectively
22 cp0 = Callpath("<root>"), Metric("<default>")
23 pm = model_generator.models[cp0].hypothesis.function
24 print(f"Model evaluated at x = 5 and y = 5:\n{pm.evaluate([5,5])}\n")
25
26 # 2) Or iterate over all of them:
27 for model in model_generator.models.values():
28     # Measurement points
29     pts = model.measurements
30     print(f"Measurement points are:\n{pts}\n")
31
32     # Evaluated function @ measurement points
33     pred = model.predictions
34     print(f"Prediction points are:\n{pred}\n")
35
36     # The actual model
37     print(f"Model is:\n{model.hypothesis.function}\n")
38
39     # Evaluate at specific points
40     # Here, data.jsonl is a 2D data set, hence, a 2d array must be provided
41     m = model.hypothesis.function.evaluate(
42         np.array([[1, 24, 58], [1, 29, 30]]))
43     print(f"Model evaluated at x = [1 24 58] and y = [1 29 30] is:\n{m}\n")
44
45     # Or evaluate just at a single point [x,y]
46     pm = model.hypothesis.function.evaluate([2, 3])
47     print(f"Model evaluated at x = 2 and y = 3:\n{pm}\n\n")
```

Monitoring Proxy Captures traces at different levels for each application (example: IMB-IO)

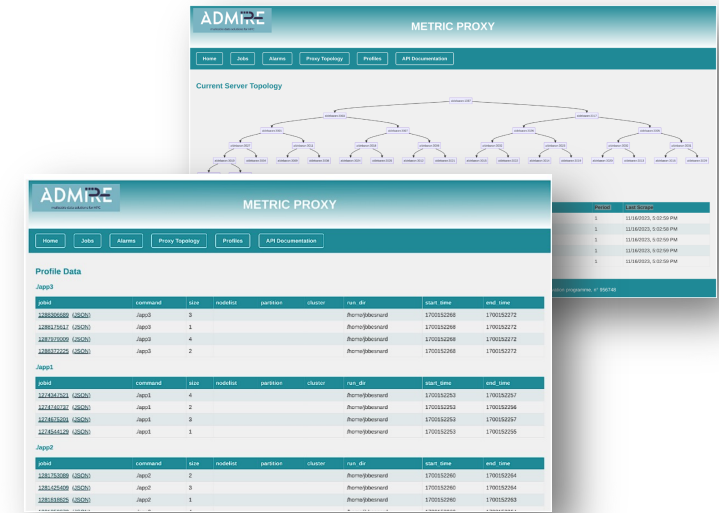




File system side is also available, but currently not used for performance modeling (Ex: HACC-IO)




- Support for four new tools has been added
  - TMIO (Tracing MPI-I/O)
  - Darshan
  - Tau metric proxy
  - FTIO



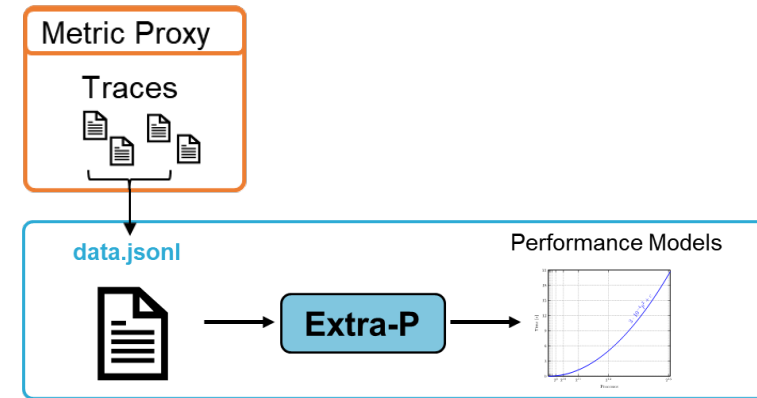
Metric proxy: Graphical user interface

| Source       | Key Advantage                                    |
|--------------|--|
| TMIO         | Async I/O, I/O requirements, flexibility         |
| Darshan      | Universal support                                |
| Metric proxy | Very detailed traces, <b>on-the-fly modeling</b> |
| FTIO         | Modeling I/O phases                              |

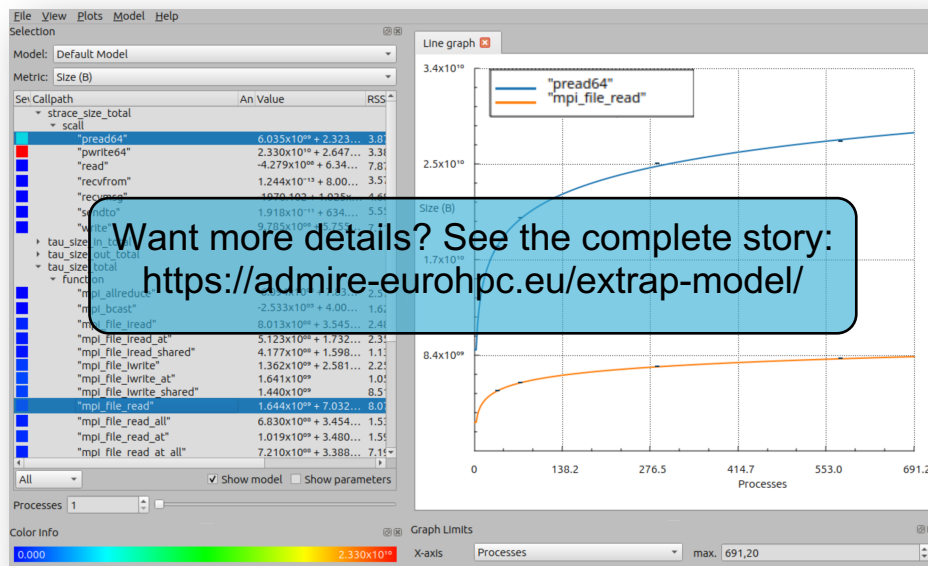


FTIO: Frequency techniques for I/O

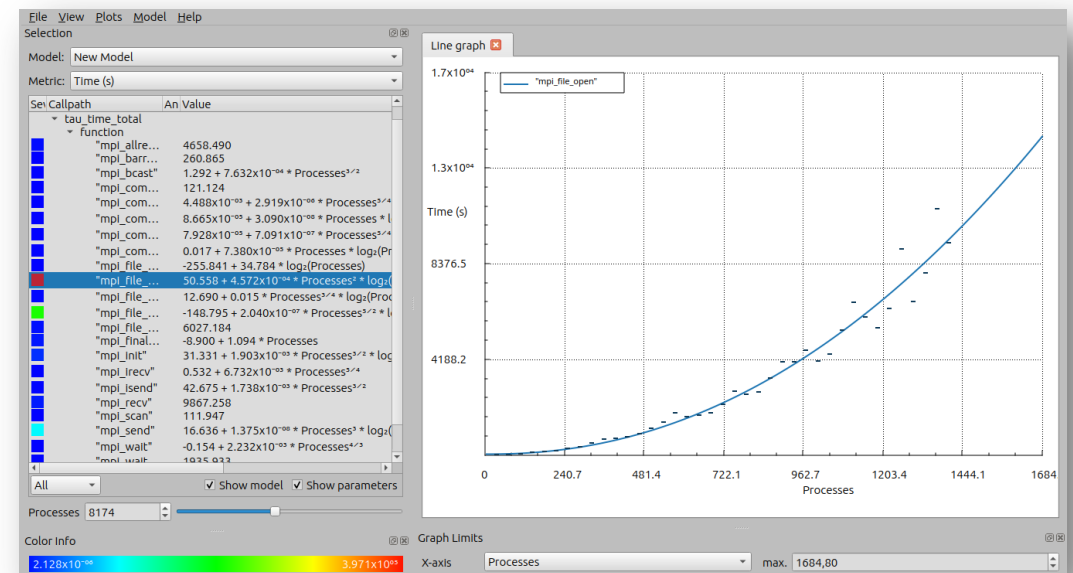
- Metric proxy:
  - Provides traces at different levels
  - Groups traces into a single profile
  - Different methods for grouping exist



IMB-IO



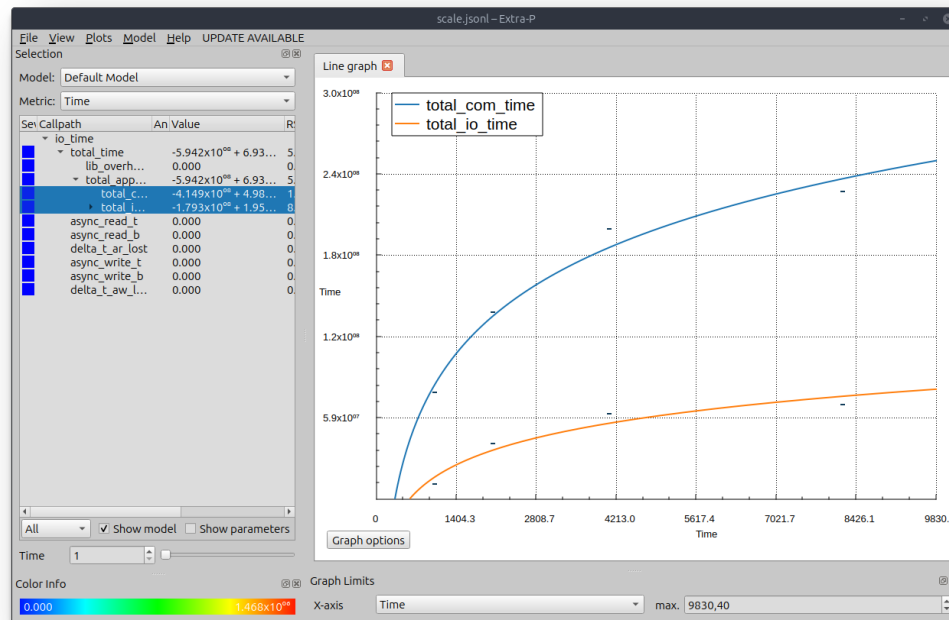
Nek5000



- Nek5000 simulation from Mainz trace Website
- Generate a performance model from any Darshan File (with and without heatmap)

HPC IO ANALYSIS

| Application                     | Workload Family              | Institute                              | Cluster/TOP500 | IO Library | Checkpointing | MPI Ranks | Details | Compare |
|---------------------------------|------------------------------|--|----------------|------------|---------------|-----------|---------|---------|
| Nek5000 (Turbulence Simulation) | Computational Fluid Dynamics | Johannes Gutenberg University of Mainz | Mogon II       | MPI-IO     | False         | 8192      | Details | Compare |
| Nek5000 (Turbulence Simulation) | Computational Fluid Dynamics | Johannes Gutenberg University of Mainz | Mogon II       | MPI-IO     | False         | 4096      | Details | Compare |
| Nek5000 (Turbulence Simulation) | Computational Fluid Dynamics | Johannes Gutenberg University of Mainz | Mogon II       | MPI-IO     | False         | 2048      | Details | Compare |
| Nek5000 (Turbulence Simulation) | Computational Fluid Dynamics | Johannes Gutenberg University of Mainz | Mogon II       | MPI-IO     | False         | 1024      | Details | Compare |



- We acknowledge the support of the European Commission and the German Federal Ministry of Education and Research (BMBF) under the EuroHPC Programmes DEEP-SEA (GA no. 955606, BMBF funding no. 16HPC015) and ADMIRE (GA no. 956748, BMBF funding no. 16HPC006K), which receive support from the European Union's Horizon 2020 programme and DE, FR, ES, GR, BE, SE, GB, CH (DEEP-SEA) or DE, FR, ES, IT, PL, SE (ADMIRE).
- We gratefully acknowledge the computing time provided on the high-performance computer Lichtenberg at the NHR Centers NHR4CES at TU Darmstadt. This is funded by the Federal Ministry of Education and Research, and the state governments participating on the basis of the resolutions of the GWK for national high performance computing at universities ([www.nhr-verein.de/unsere-partner](http://www.nhr-verein.de/unsere-partner)). Moreover, we acknowledge the computing time provided on the high-performance computer at the University of Turin from the laboratory on High-Performance Computing for Artificial Intelligence.





Thank you for your attention!

Questions?