TECHNISCHE UNIVERSITÄT DARMSTADT

Parallel Programming

Validating the Performance of GPU-Offloading with Differential Performance Models

Motivation

- When offloading computation to the GPU we expect
  - Things to get faster
  - More energy efficient

- Are these expectations for a specific code fulfilled?
- What about individual application parts?
- What happens when scaling up?

Workflow

1. Obtain hardware capabilities
2. Collect baseline for the CPU-only version
3. Benchmark the GPU port
4. Generate performance models and explore them

1. Obtaining Hardware Capabilities

TECHNISCHE UNIVERSITÄT DARMSTADT

Parallel Programming

- Once per system/hardware
- Determining hardware capabilities
  - We use roofline models
    - Easy to create & understand
  - Generated with empirical roofline toolkit

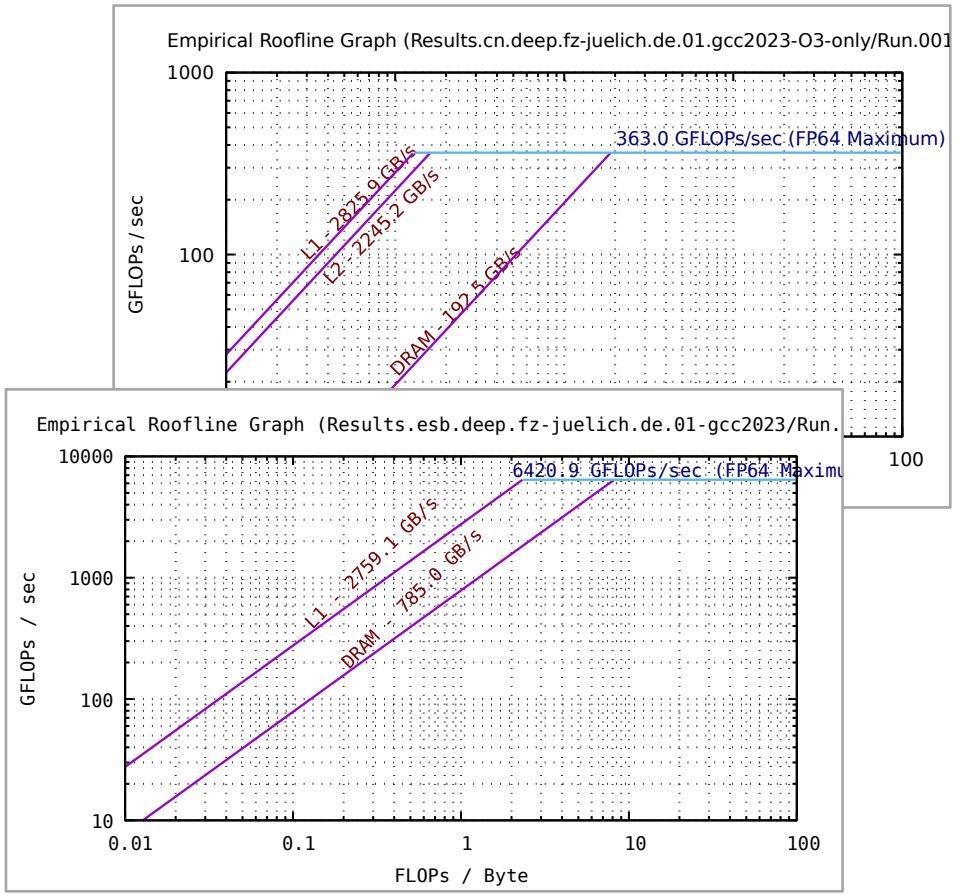

2. Baseline Measurements

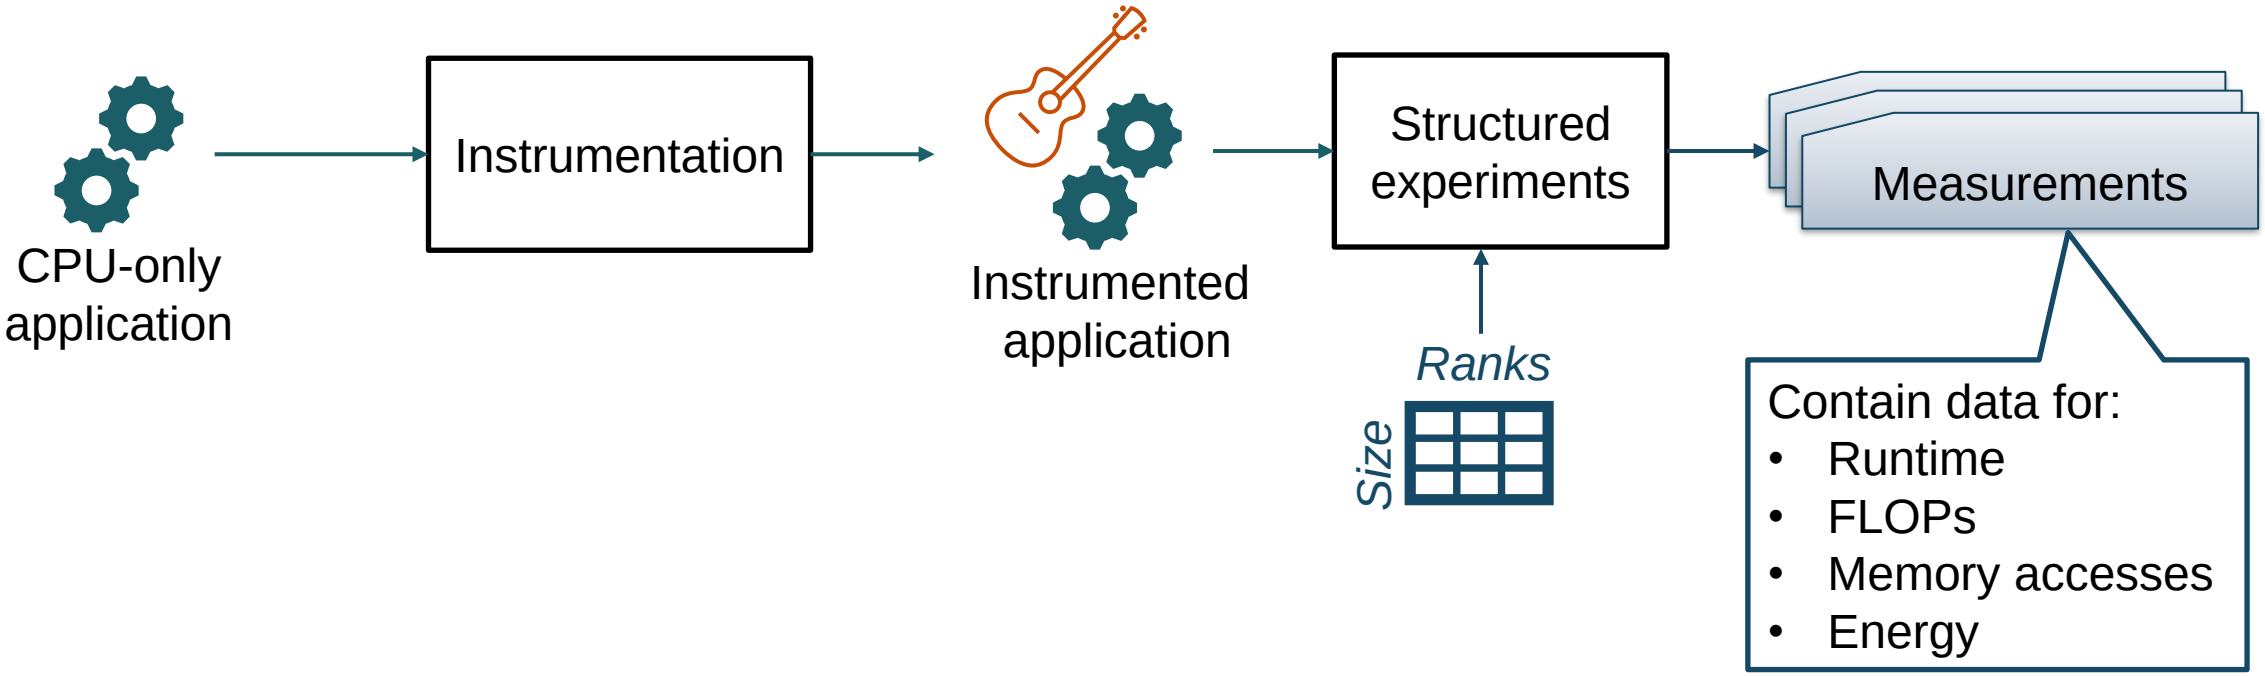

3. Benchmark the GPU port

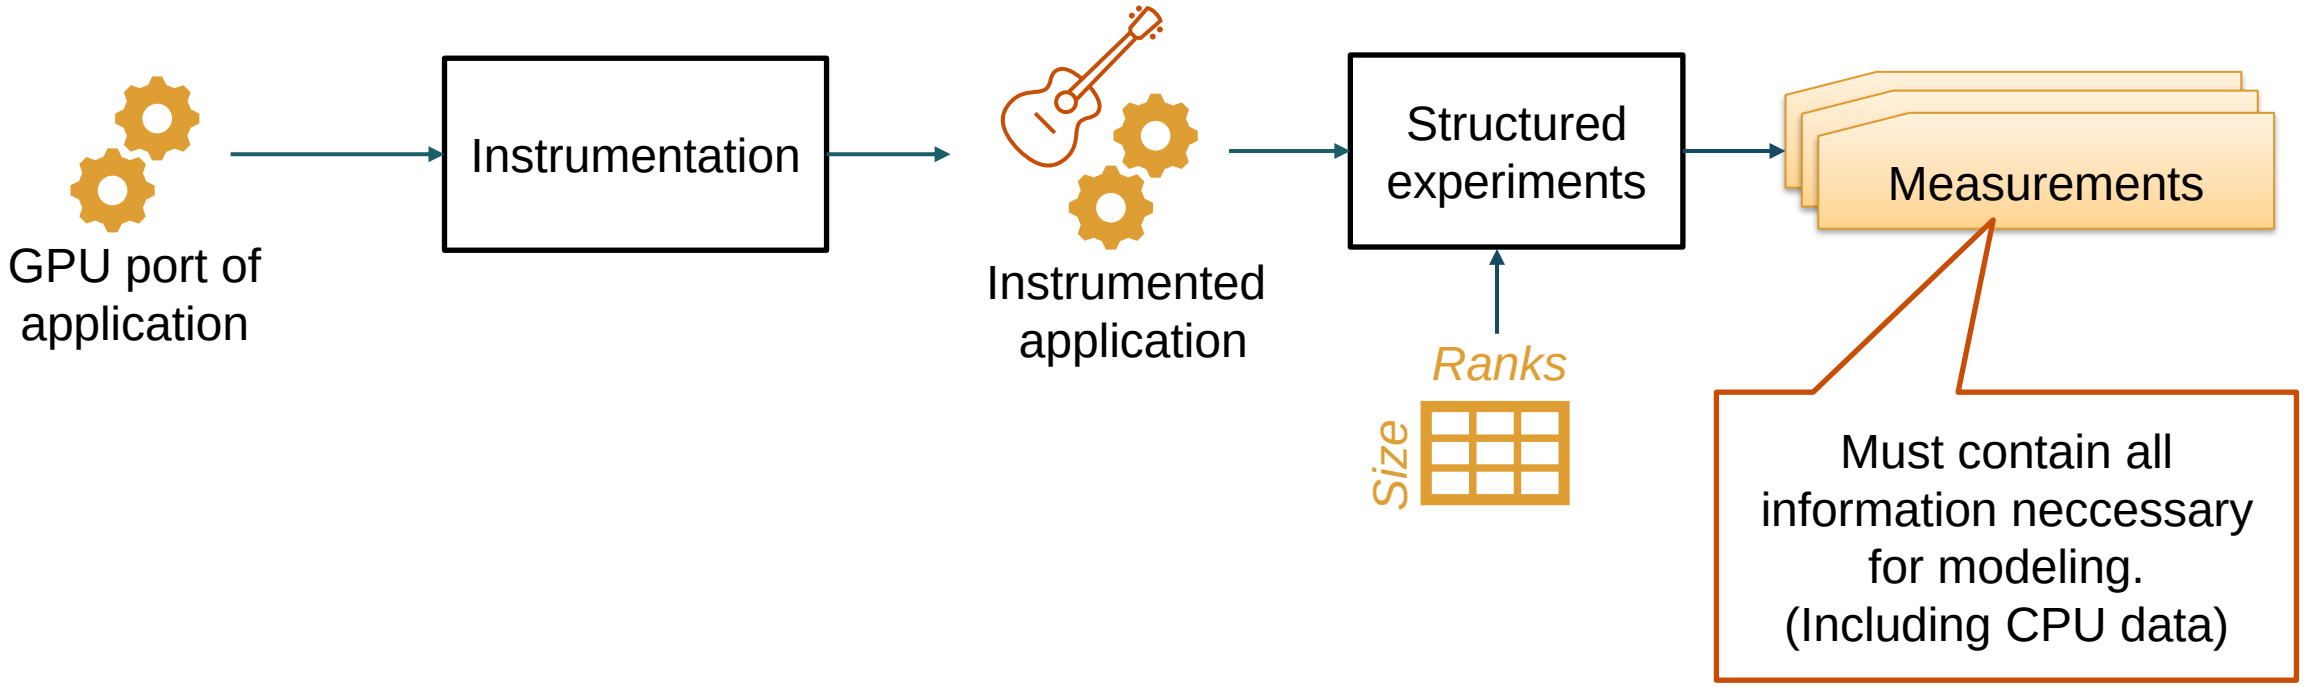

Joint CPU–GPU-Profiling

```
cudaStream_t stream1, stream2;
void foo(double* a_gpu, ...) { ...
  cudaMemcpyAsync(a_gpu, ..., H2D, stream2);
  gpu_foo<<<..., stream2>>>(a_gpu, ...); ...
}

void bar(...) { ...
  gpu_bar<<<..., stream1>>>(...); ...
}

int main() { ...
  foo(a_gpu, ...); ...
  bar(...); ...
  qux(...); ...
  cudaDeviceSynchronize(); ...
}
```

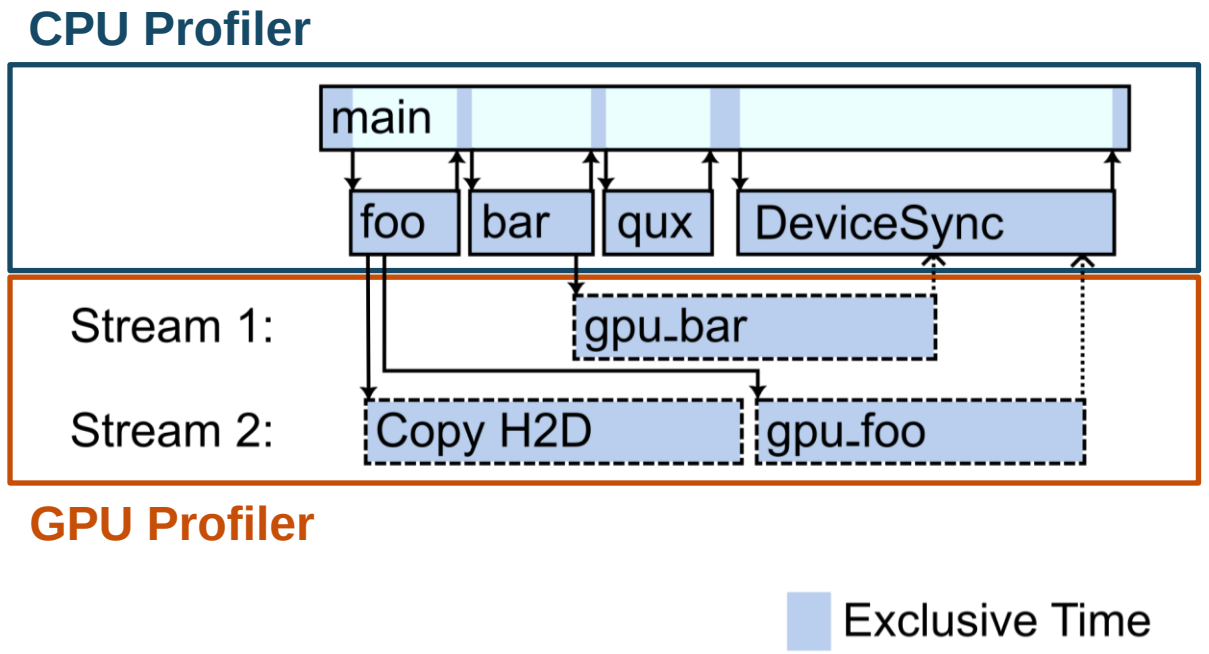

Joint CPU–GPU-Profiling II

- We collect timestamps on start and end of each function and kernel
  - Using compiler instrumentation on CPU
  - CUPTI Callback API for GPU
- We profile the CPU and trace the GPU events
  - We store only limited trace data
  - Calling context is tracked
  - During post-processing we convert to a unified profile

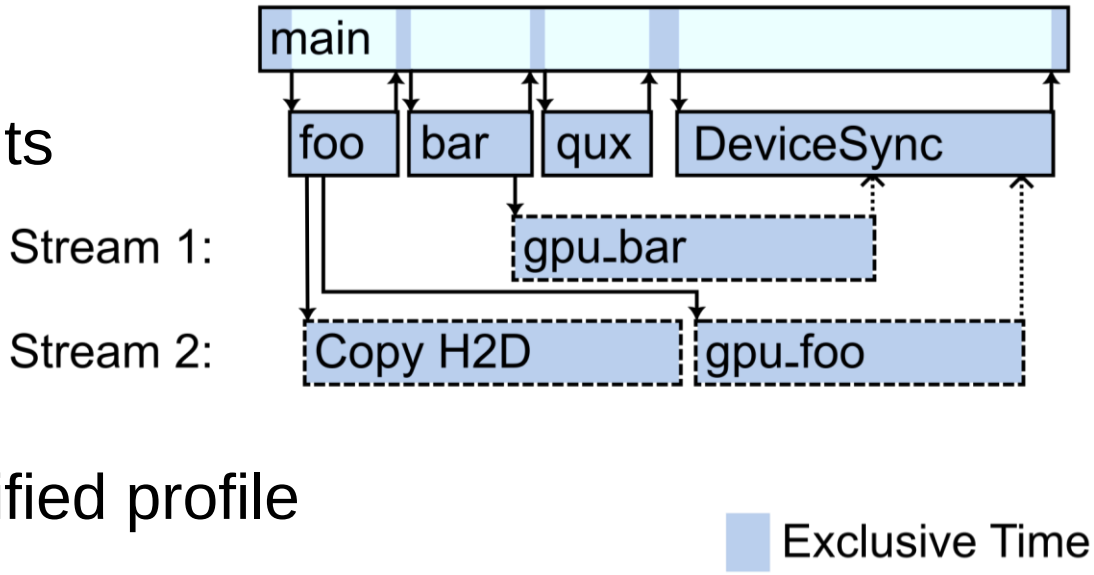

Joint CPU–GPU-Profiling III

- This allows recording of
  - Synchronization of GPU actions with CPU functions
  - Overlap of concurrent GPU activities
- Small result files with data necessary for modeling
- Full call path up to kernel

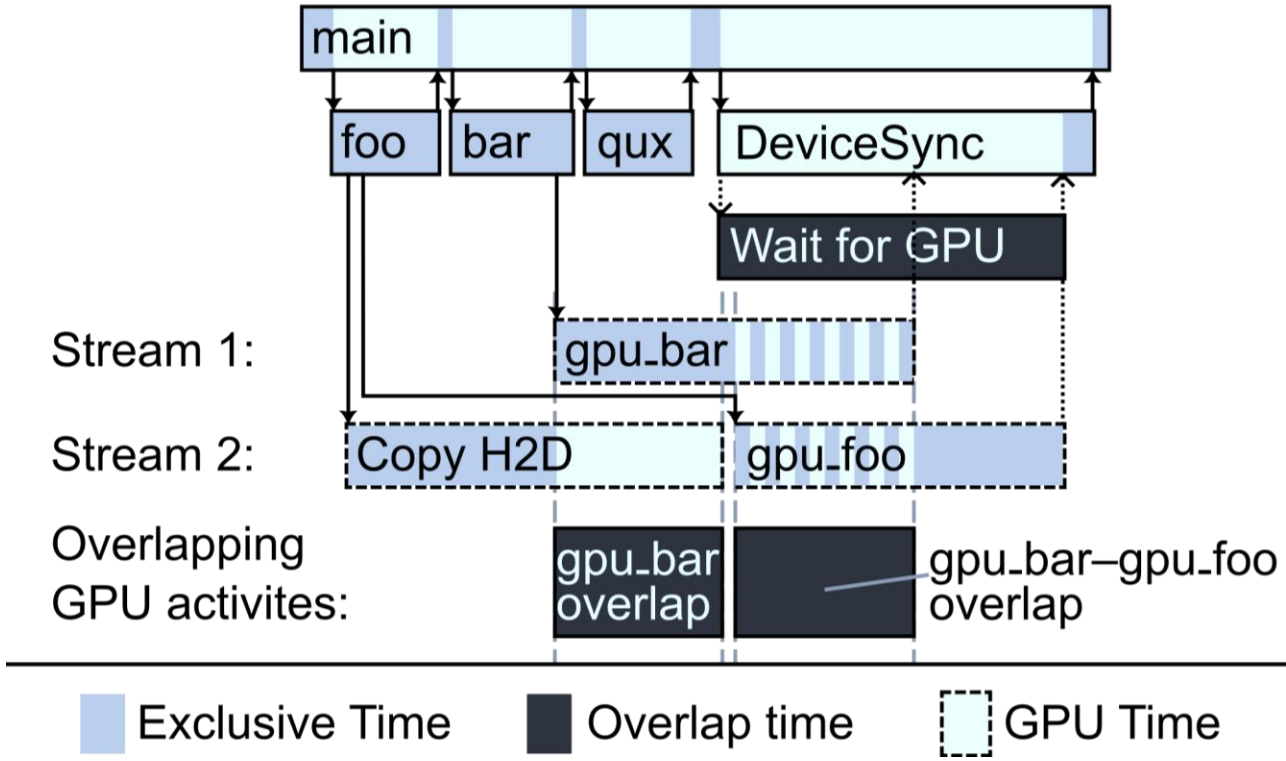

4. Differential Performance Modeling

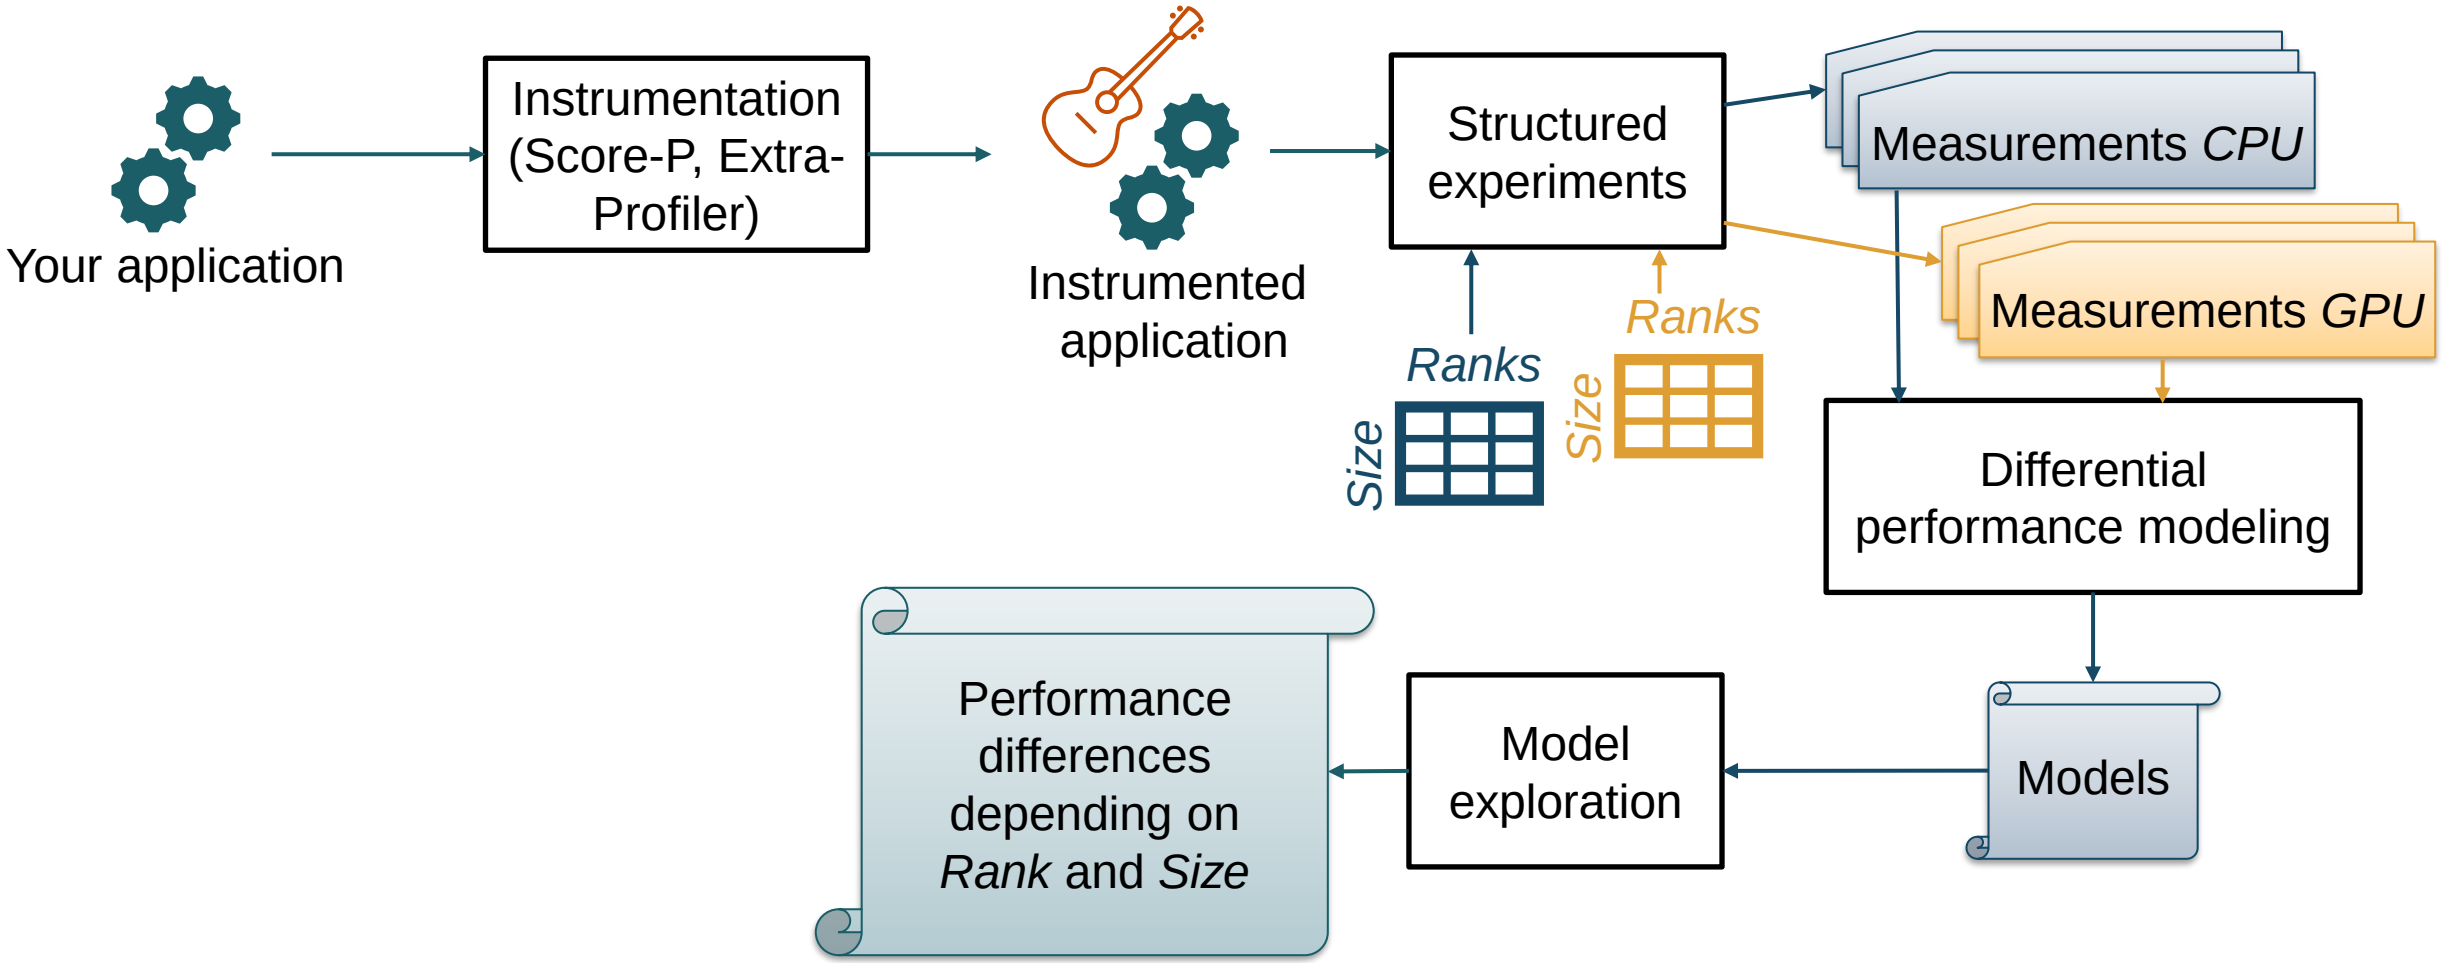

4. Differential performance modeling

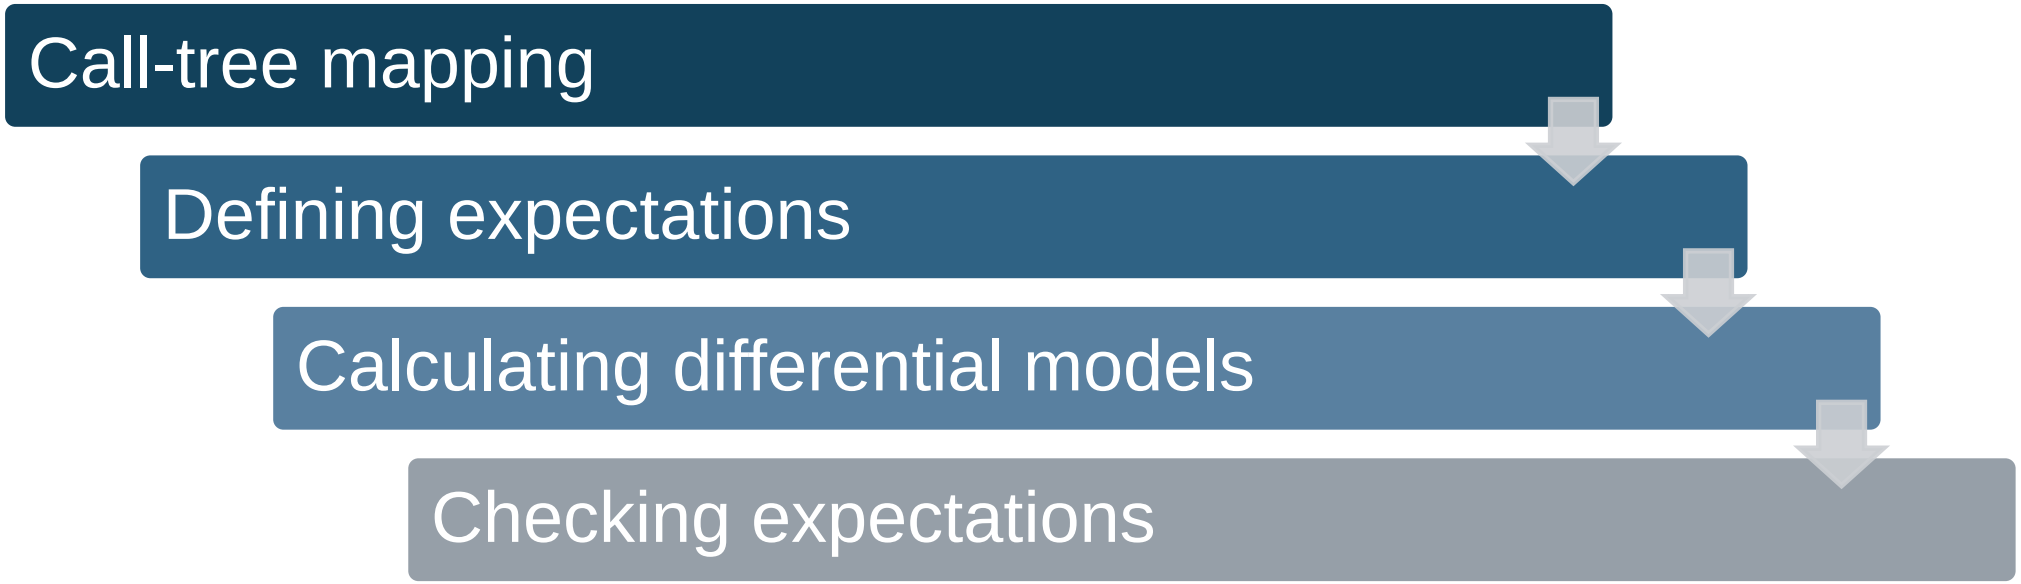

Call-Tree Mapping

- Prefix matching with aggregation

Remember we are using exclusive values

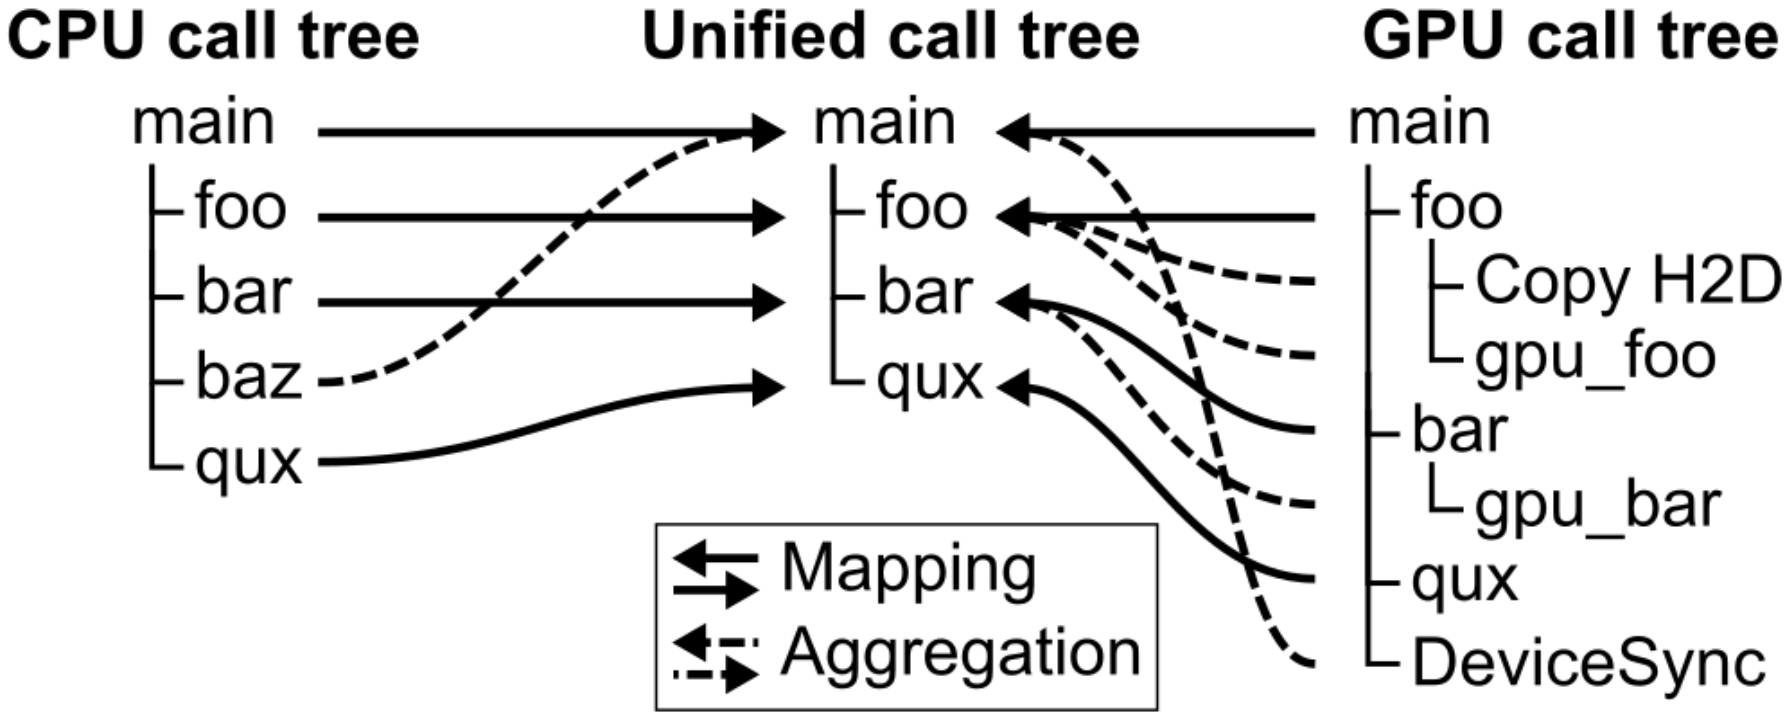

Defining Expectations

Faster than CPU-only implementation	• Runtime of GPU port is lower than runtime of CPU-only version
Uses less energy	• Energy usage of GPU port is lower than runtime of CPU-only version
Uses the hardware well	• Achieves same or higher hardware efficiency

Performance Modeling with Extra-P

TECHNISCHE UNIVERSITÄT DARMSTADT

Parallel Programming

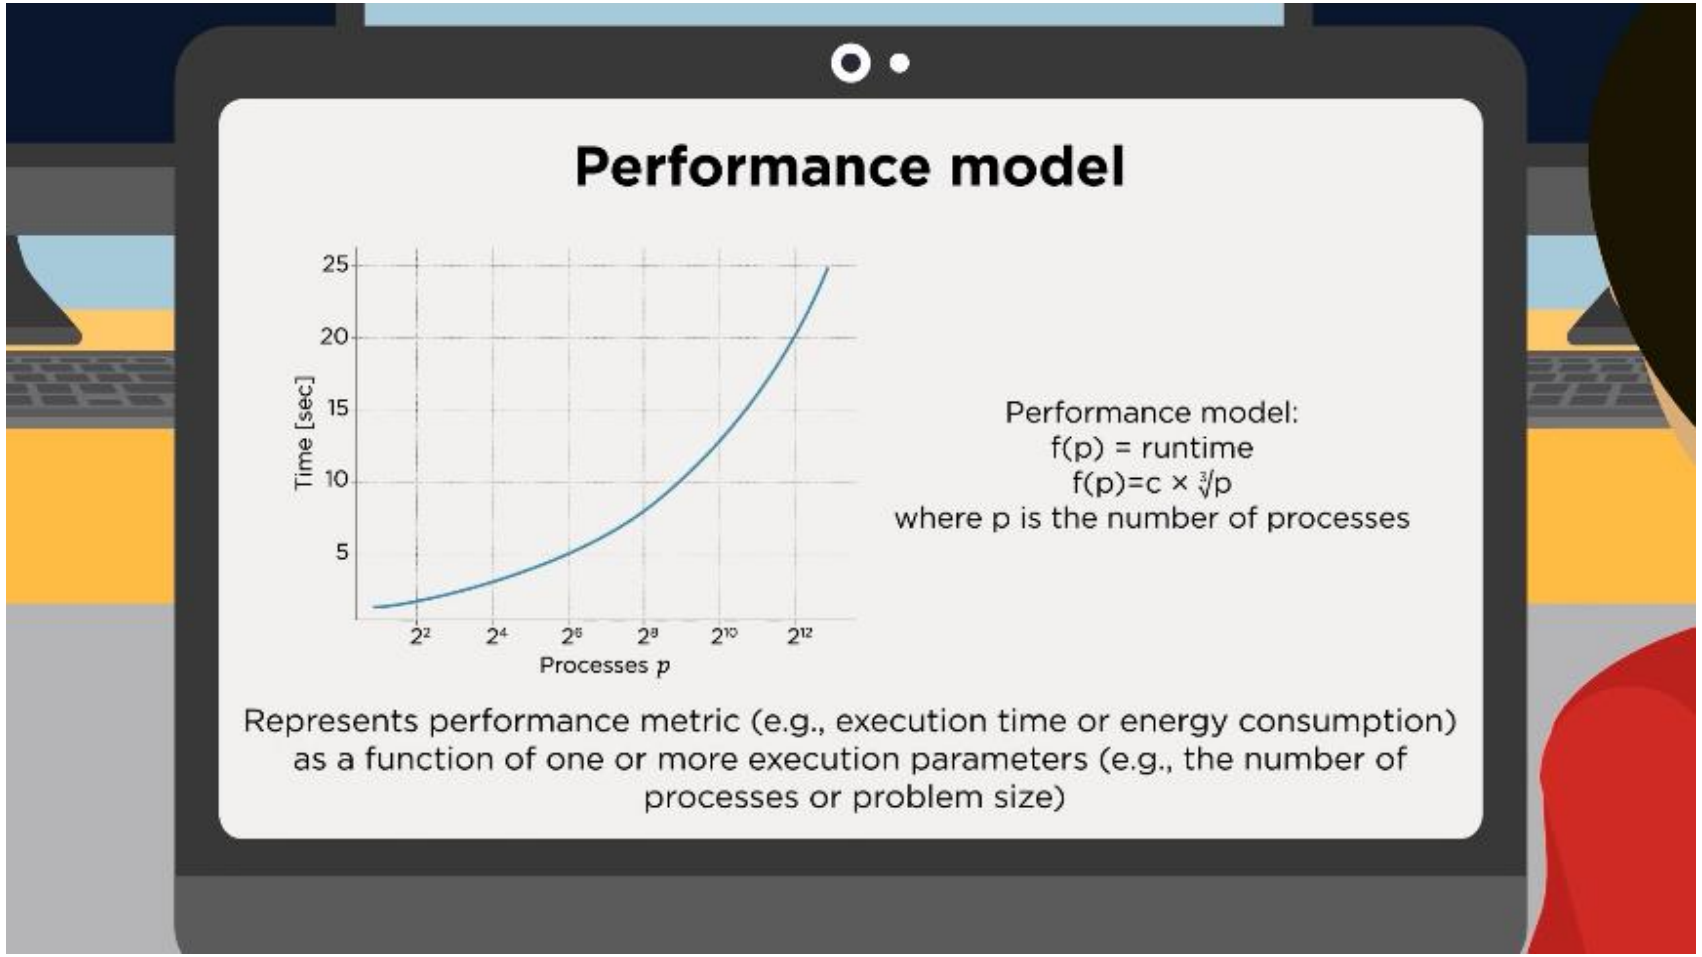

Watch Extra-P overview video

https://www.youtube.com/watch?v=Cv2YRCMWqBM

Automatic Performance Modeling

TECHNISCHE UNIVERSITÄT DARMSTADT
Parallel Programming

Input

```
main() {
    foo()
    bar()
    compute()
}
```

Instrumentation

- All functions

Performance measurements

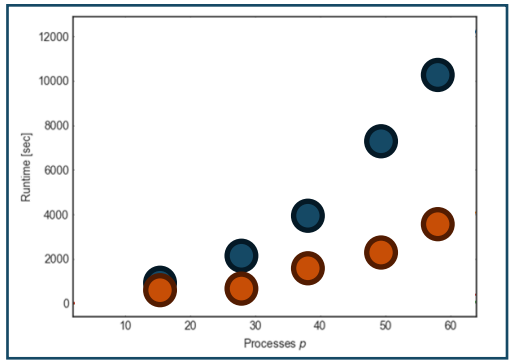

Extra-P

Output

Human-readable performance models of all functions (e.g., $t(p) = c_1 \cdot \log(p) + c_2$)

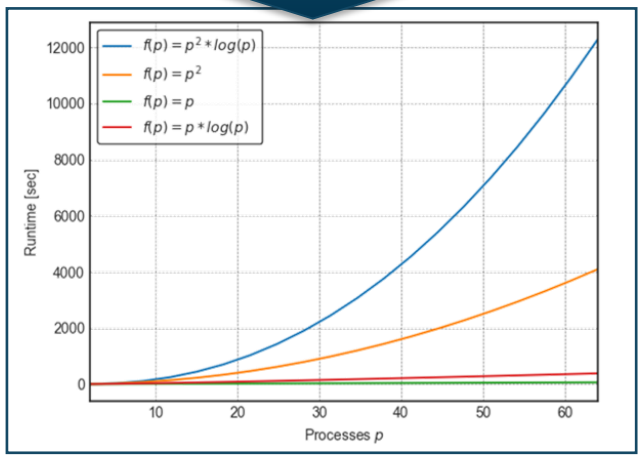

Performance Model Normal Form

$$f(x) = \sum_{k=1}^{n} c_k \cdot x^{i_k} \cdot log_2^{j_k}(x)$$

$$n \in \mathbb{N}$$
$$i_k \in I$$
$$j_k \in J$$
$$I, J \subset \mathbb{Q}$$

$$n = 1$$
$$I = \{0, 1, 2\}$$
$$J = \{0, 1\}$$

c_1 $\quad$ $c_1 \cdot \log x$
$c_1 \cdot x$ $\quad$ $c_1 \cdot x \cdot \log x$
$c_1 \cdot x^2$ $\quad$ $c_1 \cdot x^2 \cdot \log x$

Calculating Differential Models

- All models f_c^{GPU} and f_c^{CPU} are mathematical expressions
  - We can calculate with them

Differential models express the difference

- $\Delta_c(p_1, \dots) = f_c^{\mathrm{GPU}}(p_1, \dots) - f_c^{\mathrm{CPU}}(p_1, \dots)$
  - For a specific call tree entry c

Hardware Efficiency

$$\text{Hardware efficiency} = \frac{\text{achieved FLOPs per second}}{\text{achievable FLOPs per second}}$$

- Modeled with Extra-P
  - We build models for FLOPs, time, and memory accesses from measurements
    - All models are mathematical expressions
- We present this as a hardware adjusted runtime model to the user

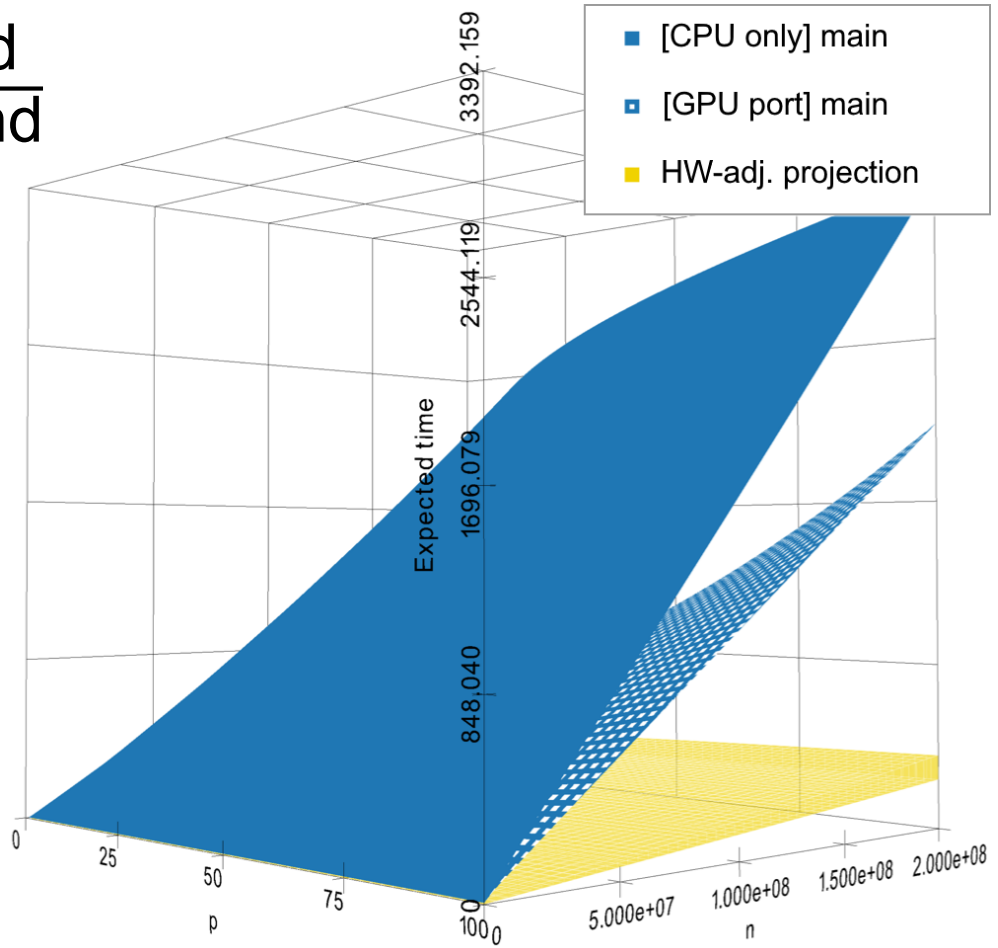

Interactive Exploration

TECHNISCHE UNIVERSITÄT DARMSTADT

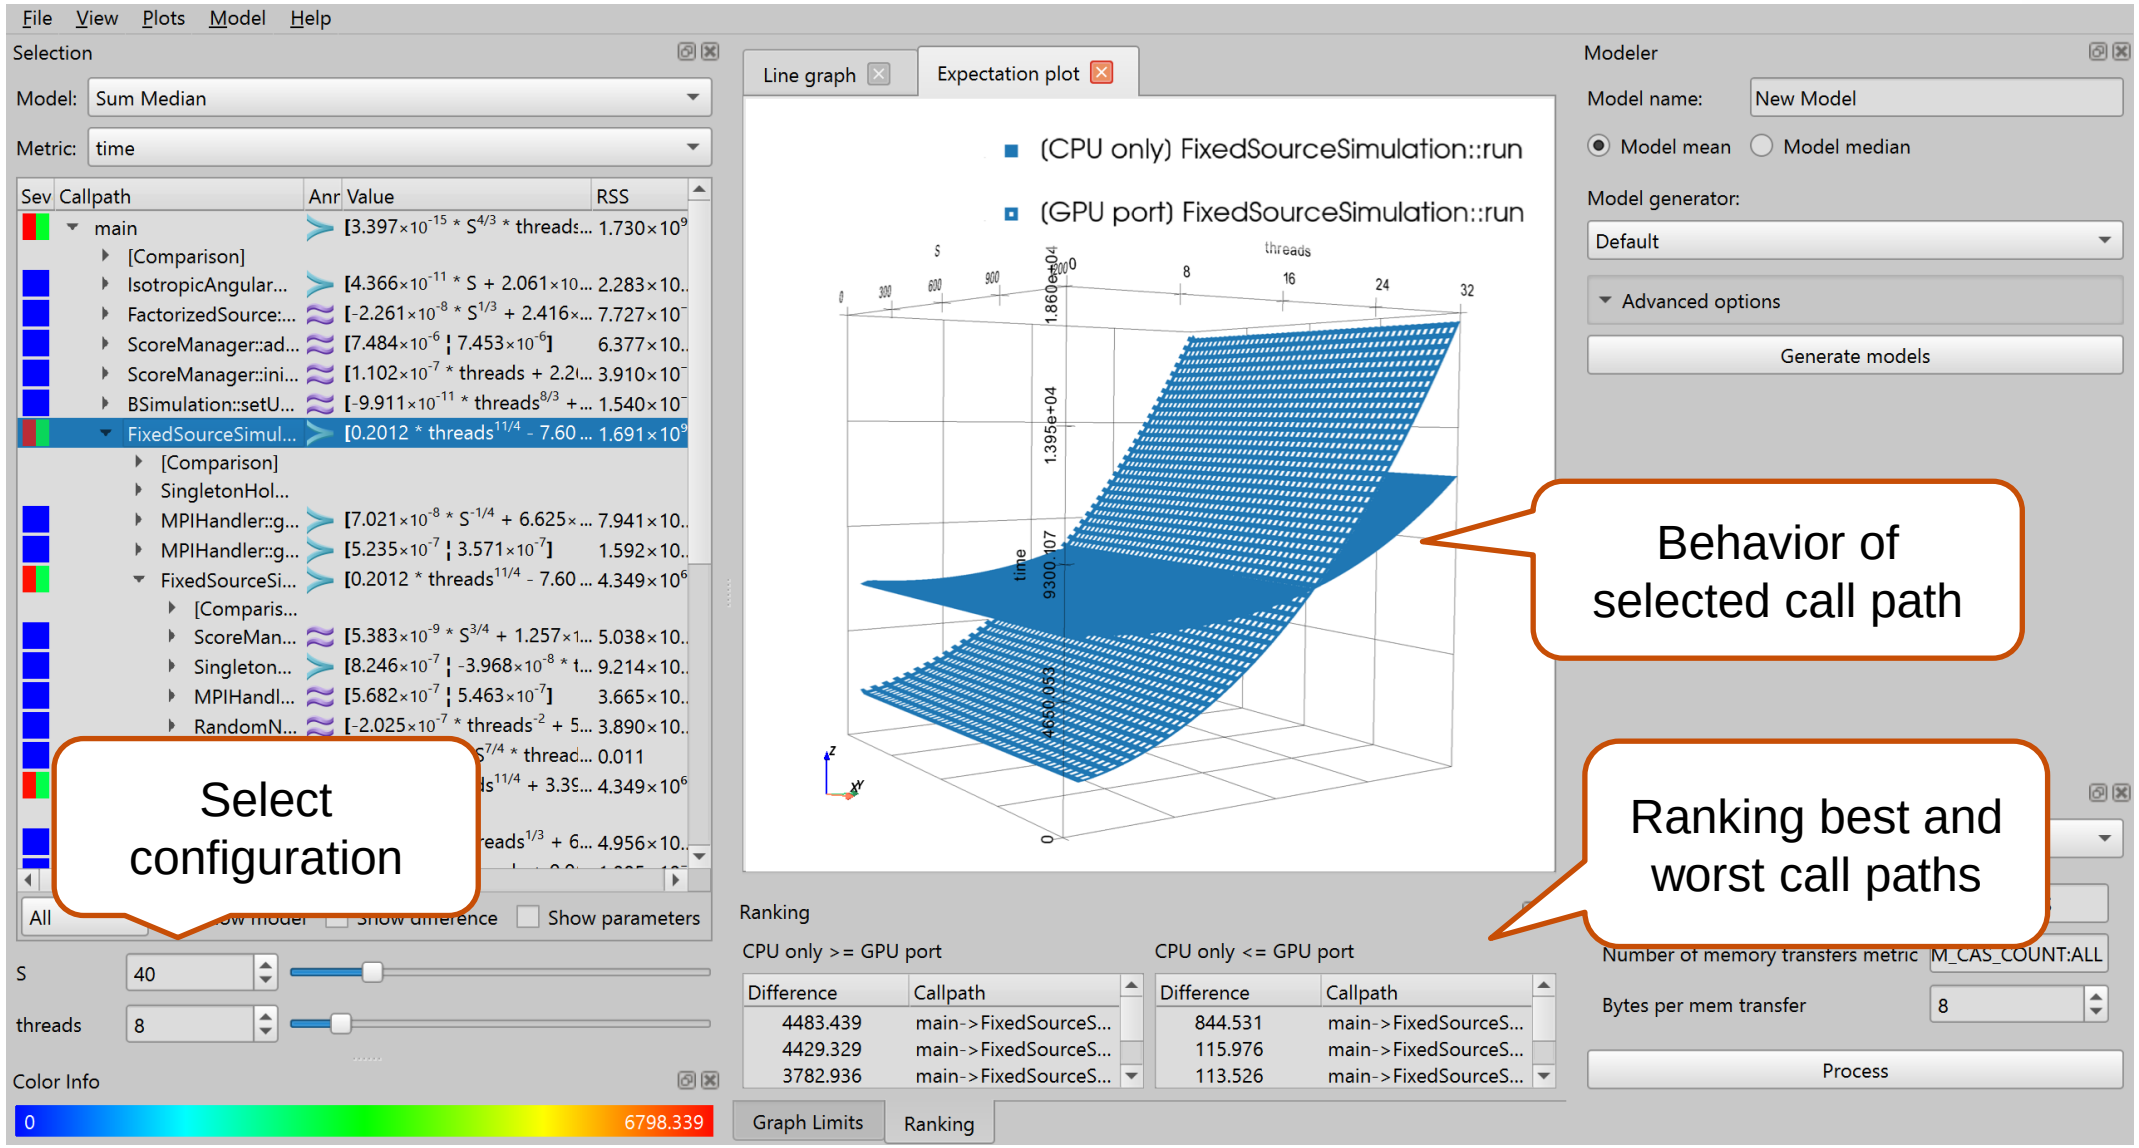

Interactive Exploration

TECHNISCHE UNIVERSITÄT DARMSTADT

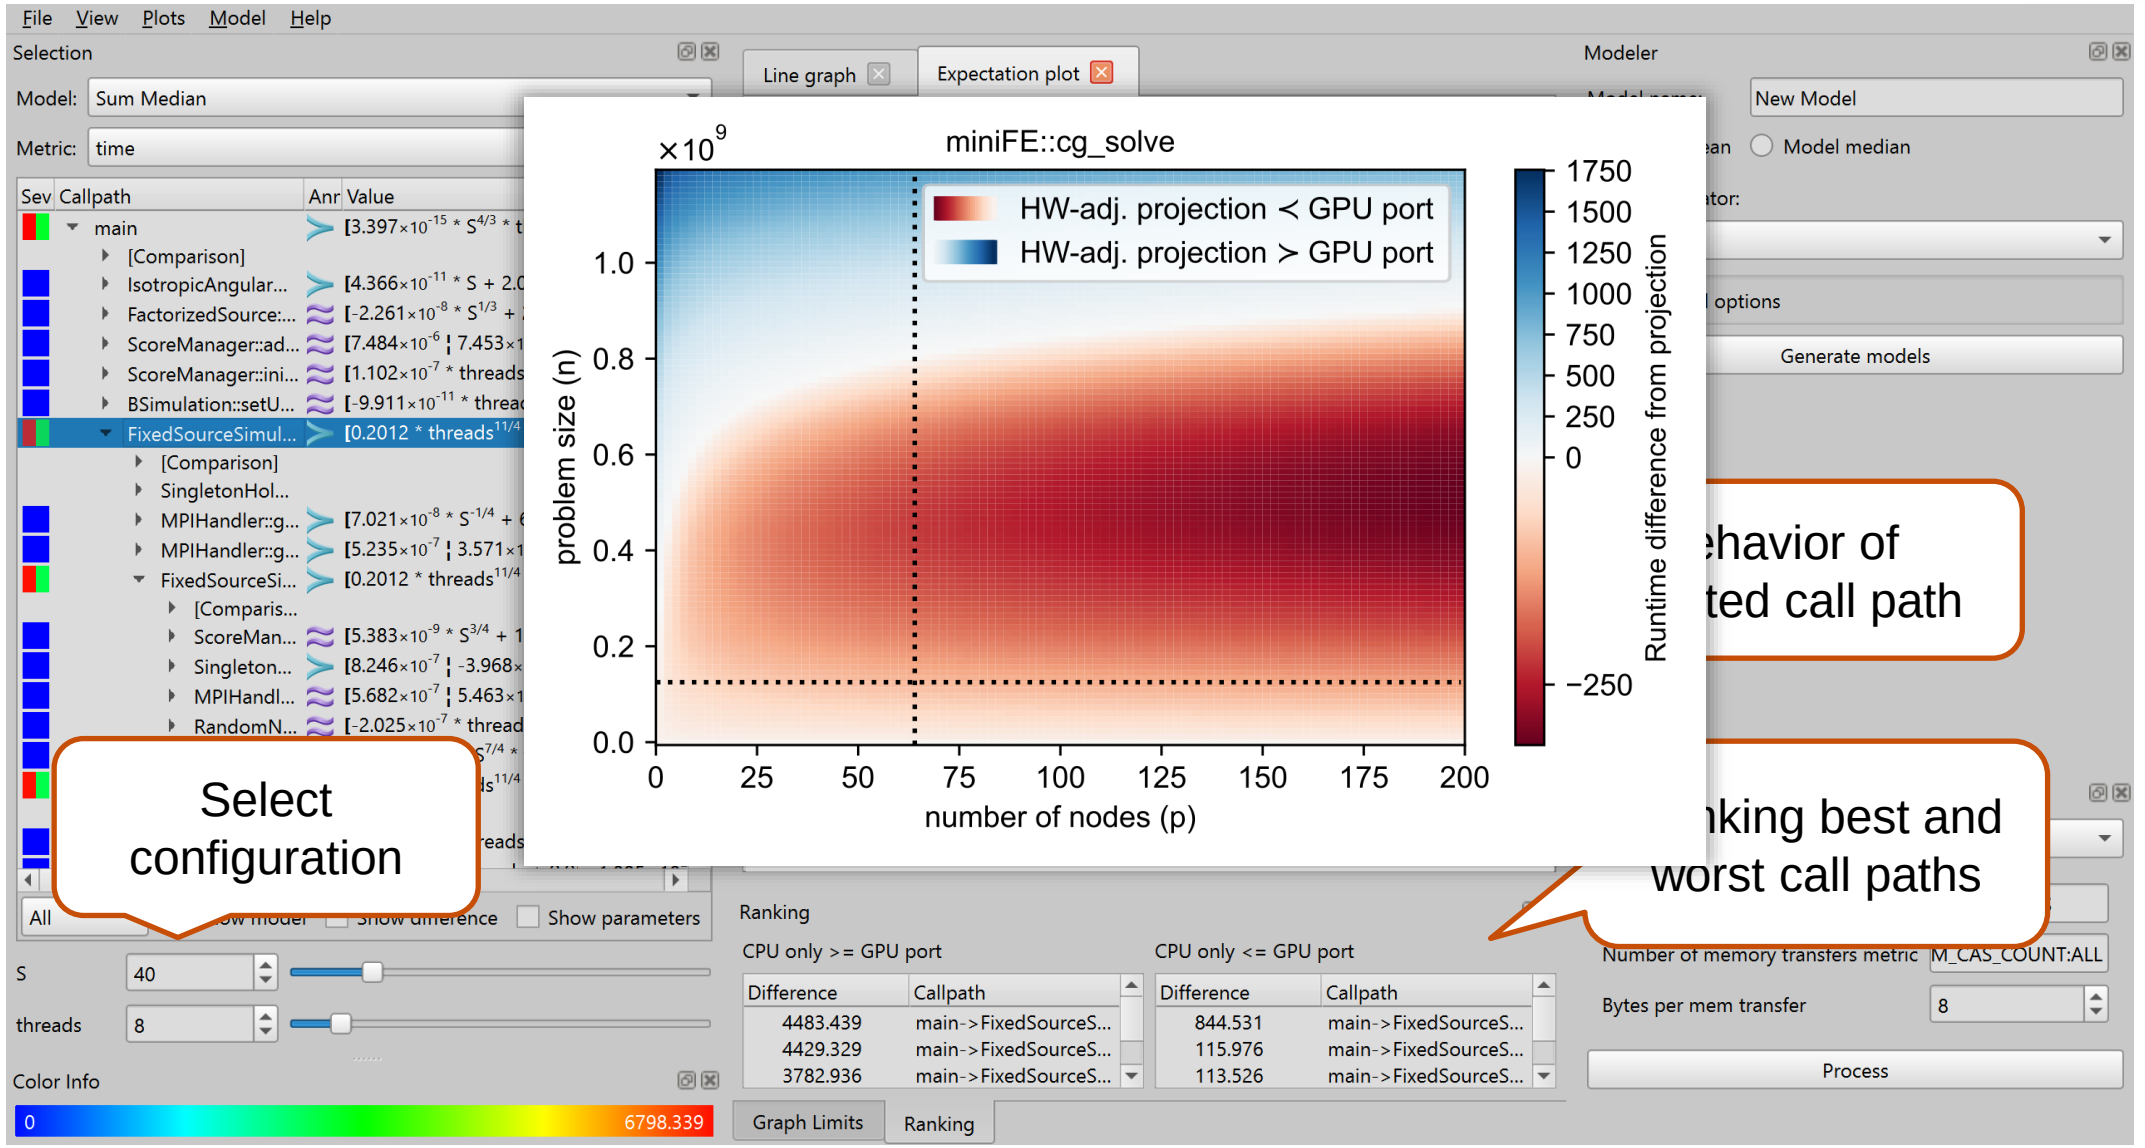

Interactive Exploration

TECHNISCHE UNIVERSITÄT DARMSTADT

Parallel Programming

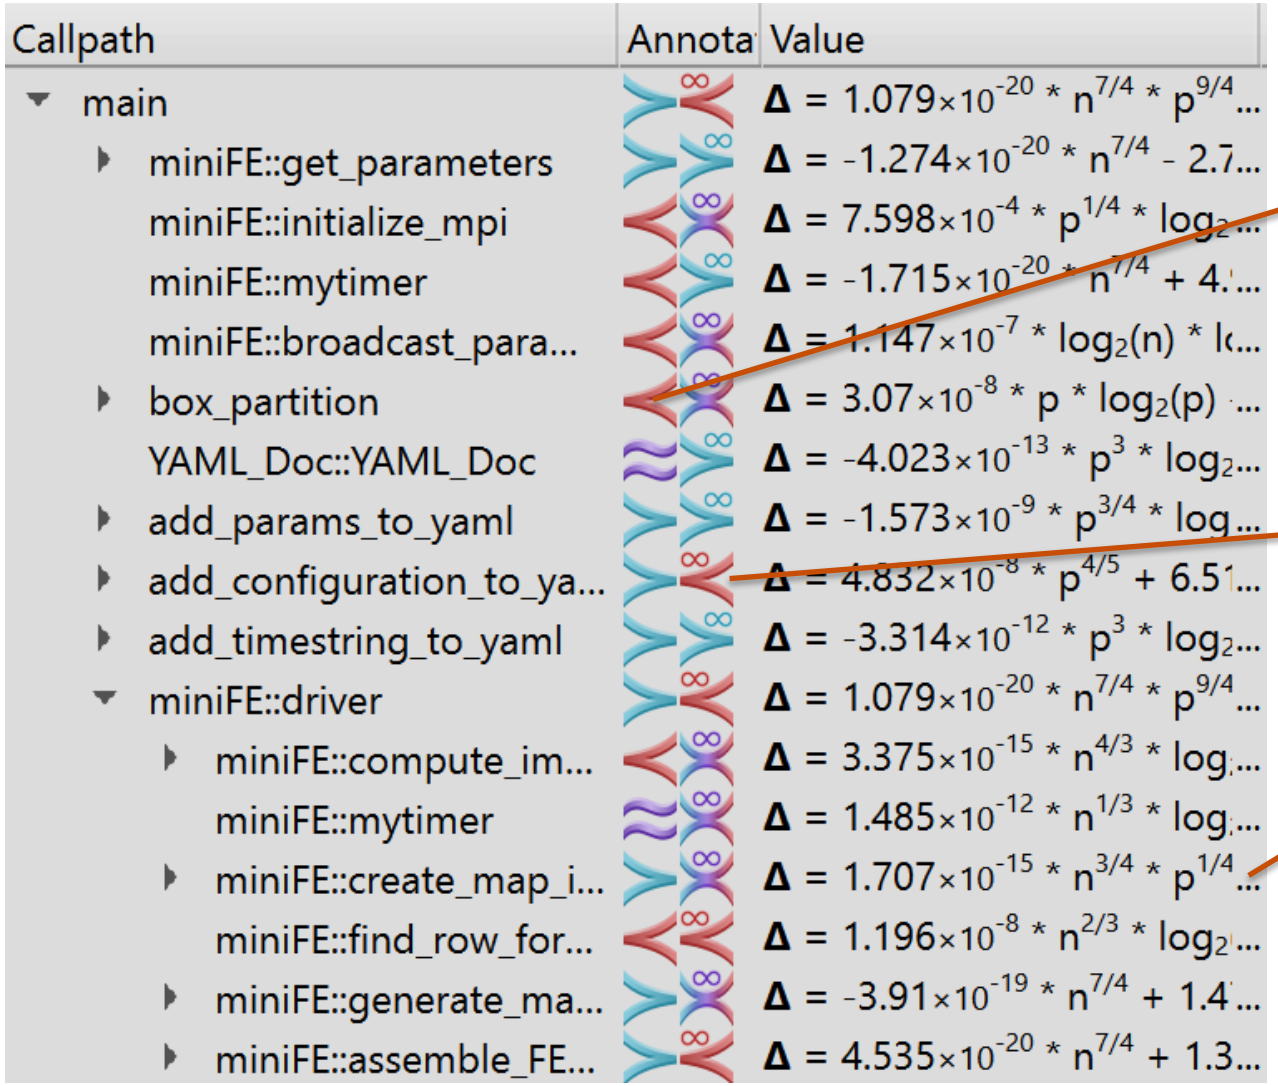

Result: GPU port...
... is slower than expected
... is faster than expected
... meets expectation

Result: GPU port is asymptotically...
... slower than expected
... faster than expected
... faster and slower than expected depending on the parameter

Differential performance models

Summary

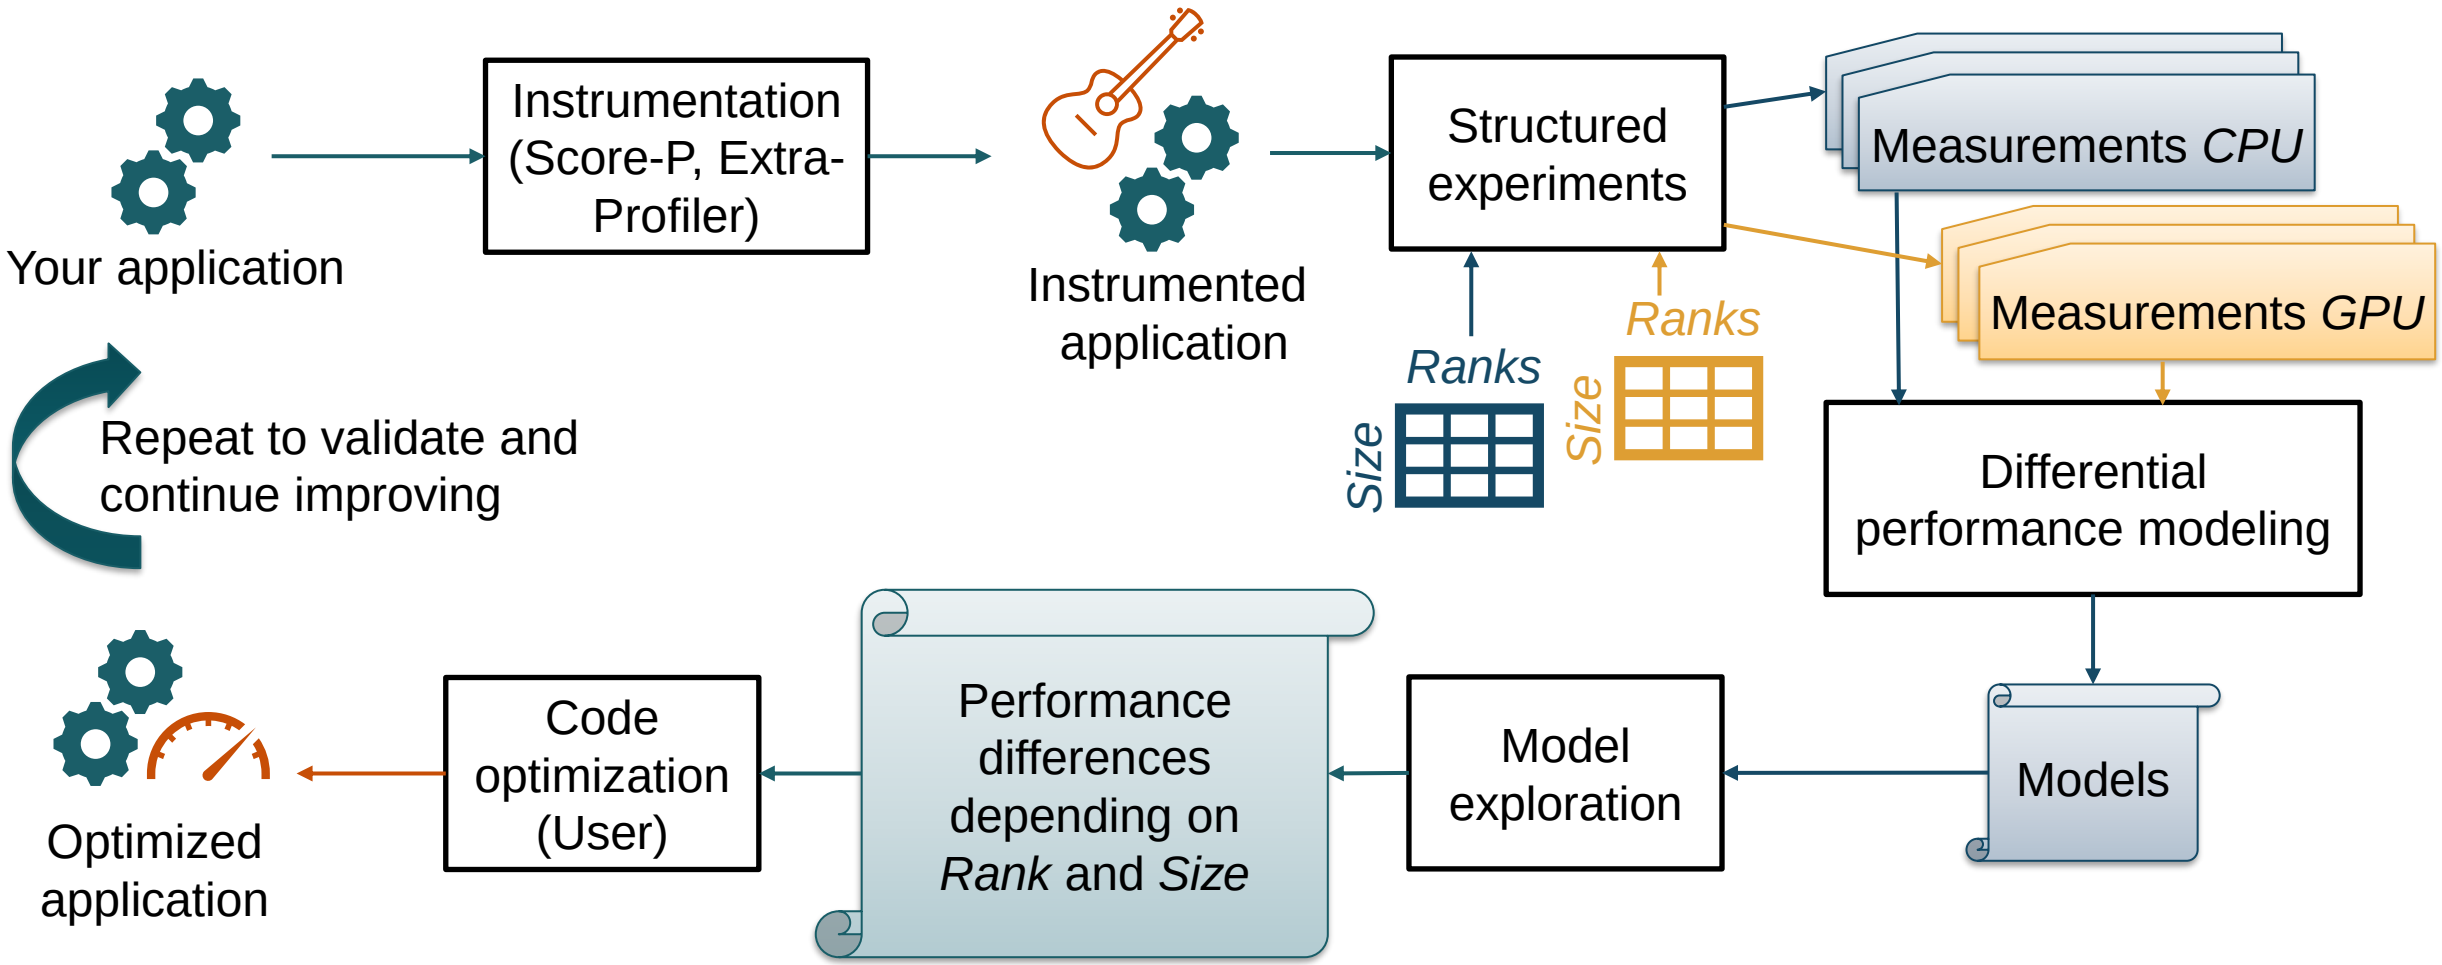
Your application
Instrumentation (Score-P, Extra-Profiler)
Instrumented application
Structured experiments
Measurements CPU
Measurements GPU
Ranks
Size
Ranks
Size
Differential performance modeling
Models
Model exploration
Performance differences depending on Rank and Size
Code optimization (User)
Optimized application
Repeat to validate and continue improving

TECHNISCHE UNIVERSITÄT DARMSTADT

Parallel Programming

Acknowledgement

We acknowledge the support of the European Commission and the German Federal Ministry of Education and Research (BMBF) under the EuroHPC programme DEEP-SEA (GA 955606, BMBF funding No. 16HPC015), which receives support from the European Union's Horizon 2020 programme and DE, FR, ES, GR, BE, SE, GB, CH. This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project No. 449683531 (ExtraNoise).

SPONSORED BY THE

Federal Ministry of Education and Research

Funded by

DFG Deutsche Forschungsgemeinschaft

German Research Foundation