Extra-P: Application Performance Modelling and Application Mapping

Alexander Geiß, Technical University of Darmstadt

TECHNISCHE UNIVERSITÄT DARMSTADT

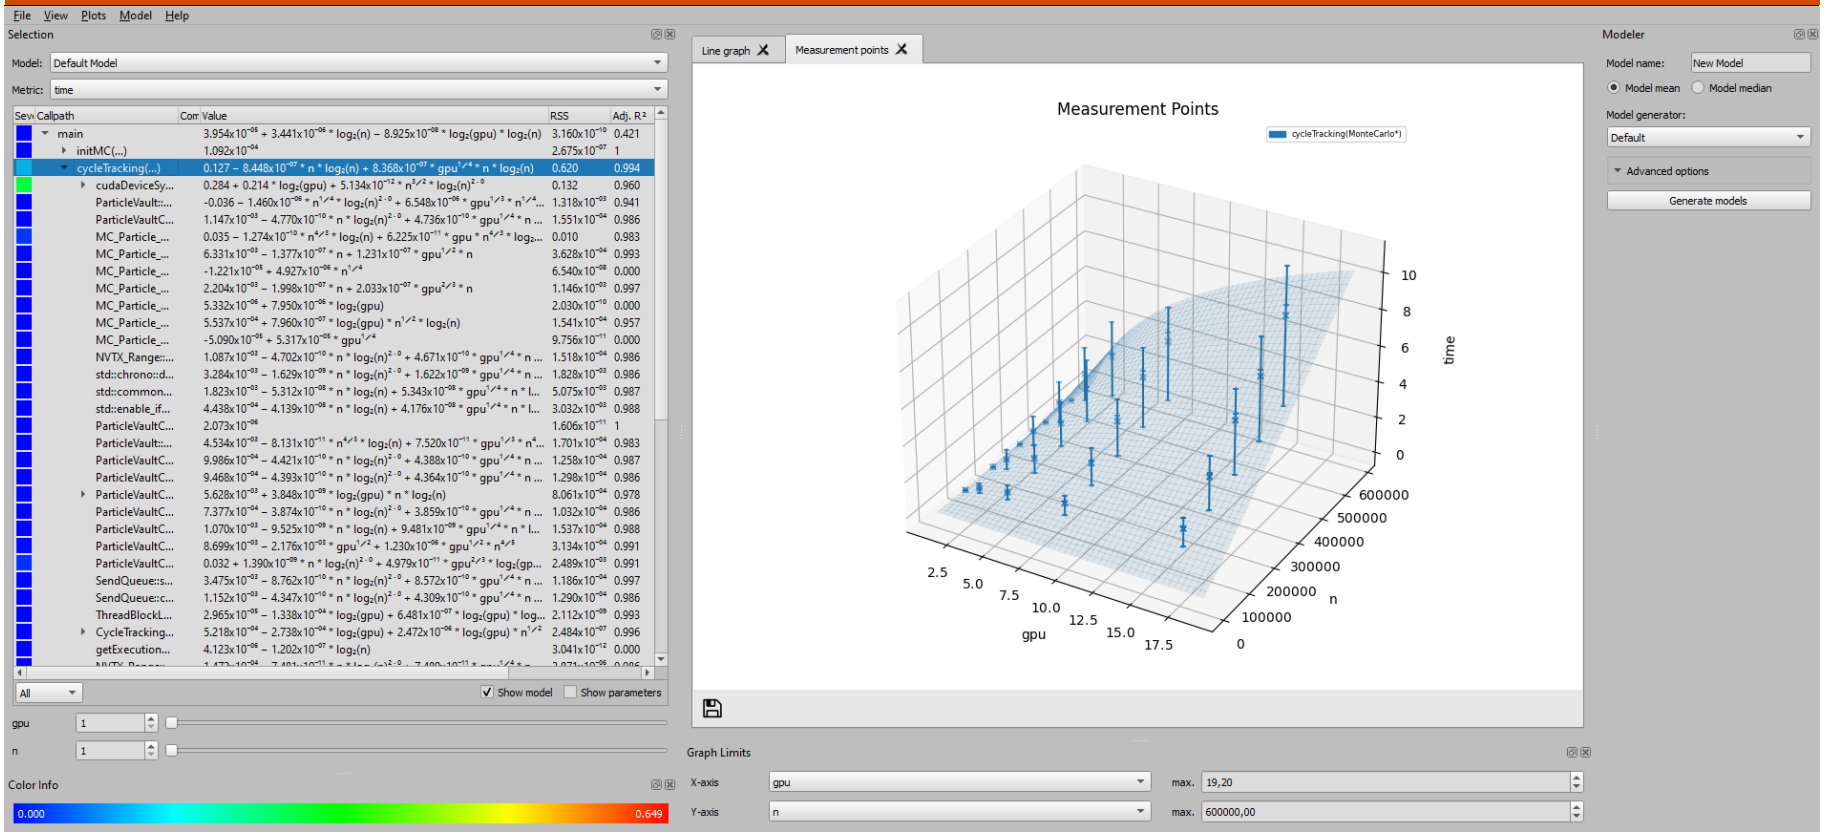

Scaling code to a bigger machine can unveil unpleasant surprises

TECHNISCHE UNIVERSITÄT DARMSTADT

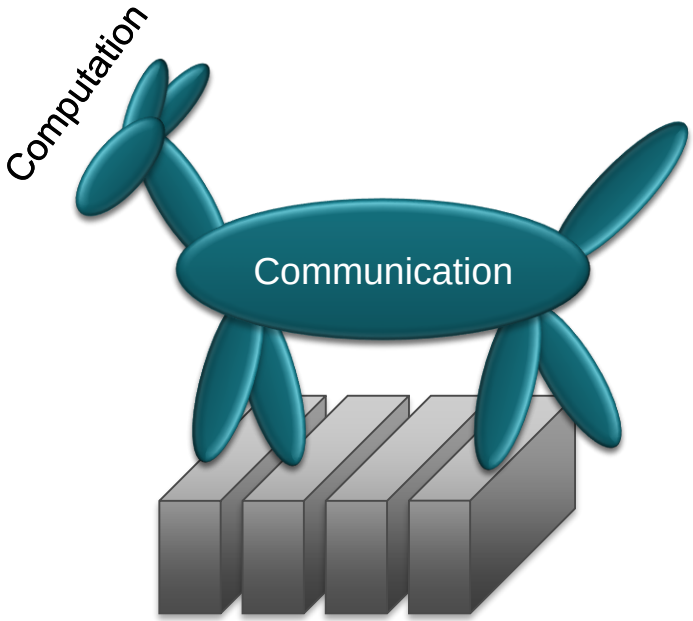

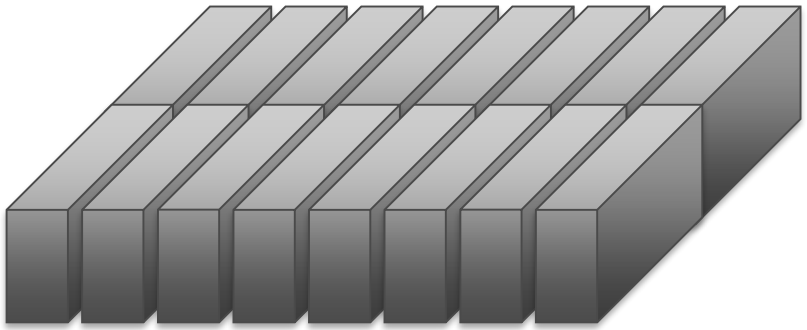

Scaling code to a bigger machine can unveil unpleasant surprises

TECHNISCHE UNIVERSITÄT DARMSTADT

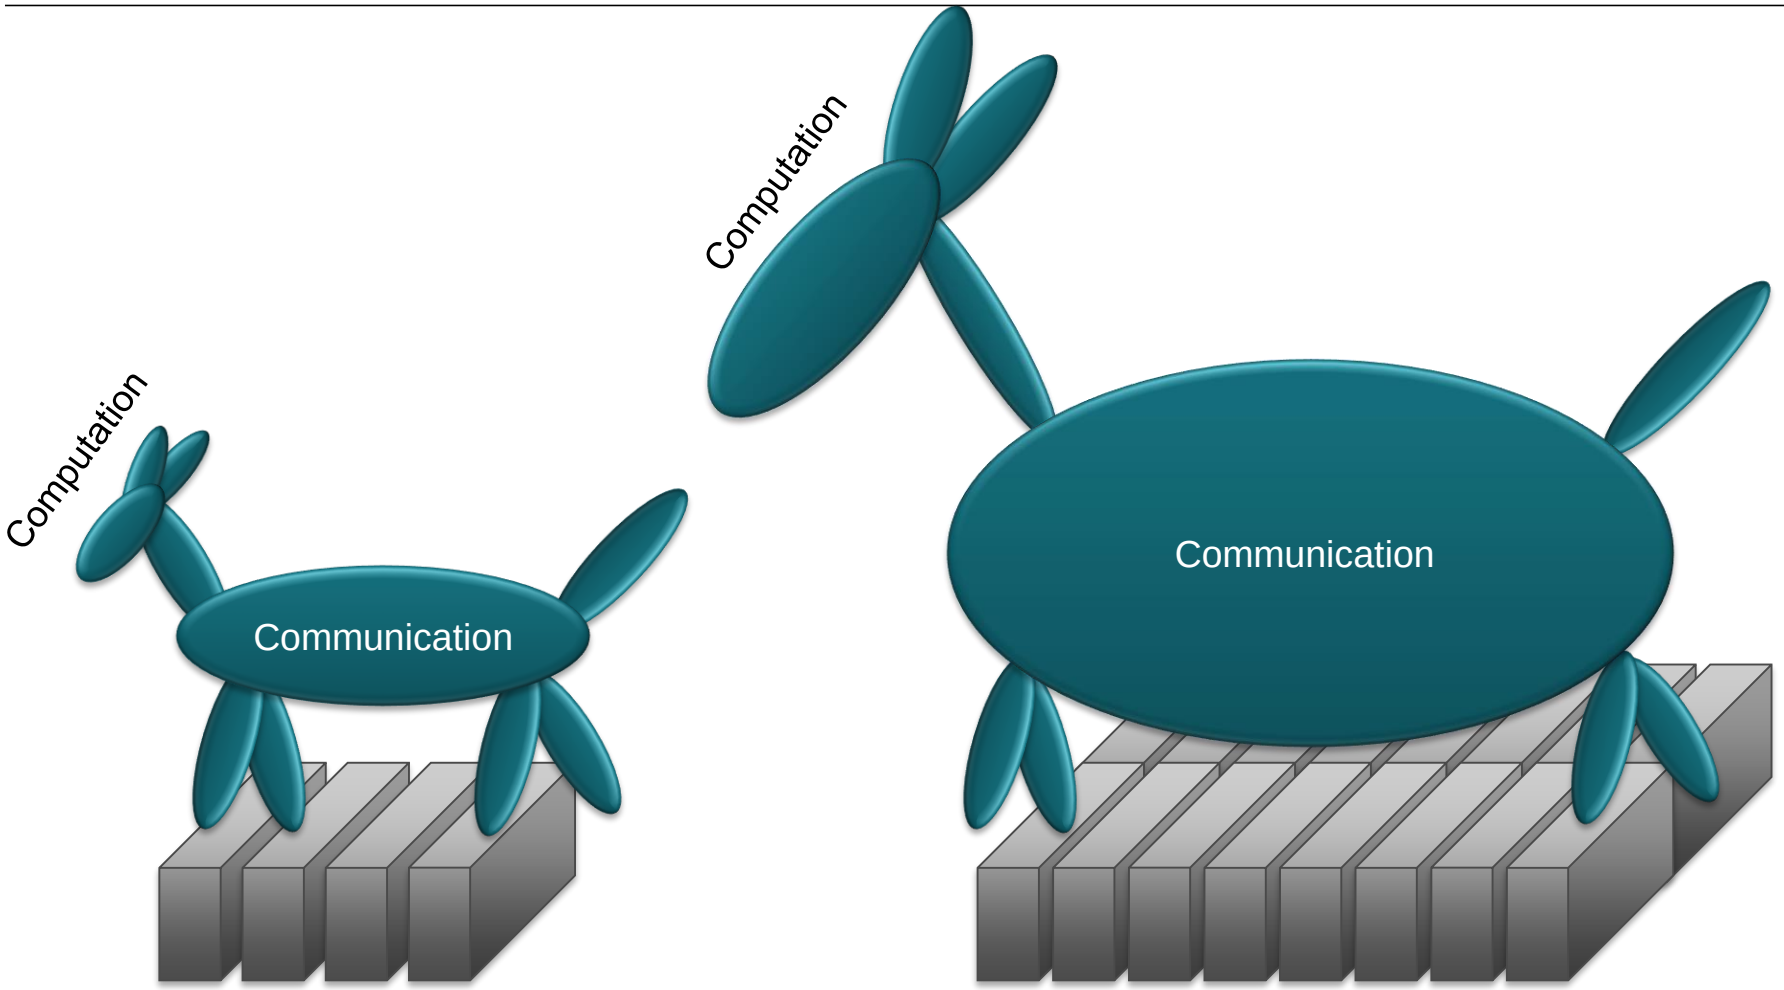

Scaling code to a bigger machine can unveil unpleasant surprises

TECHNISCHE UNIVERSITÄT DARMSTADT

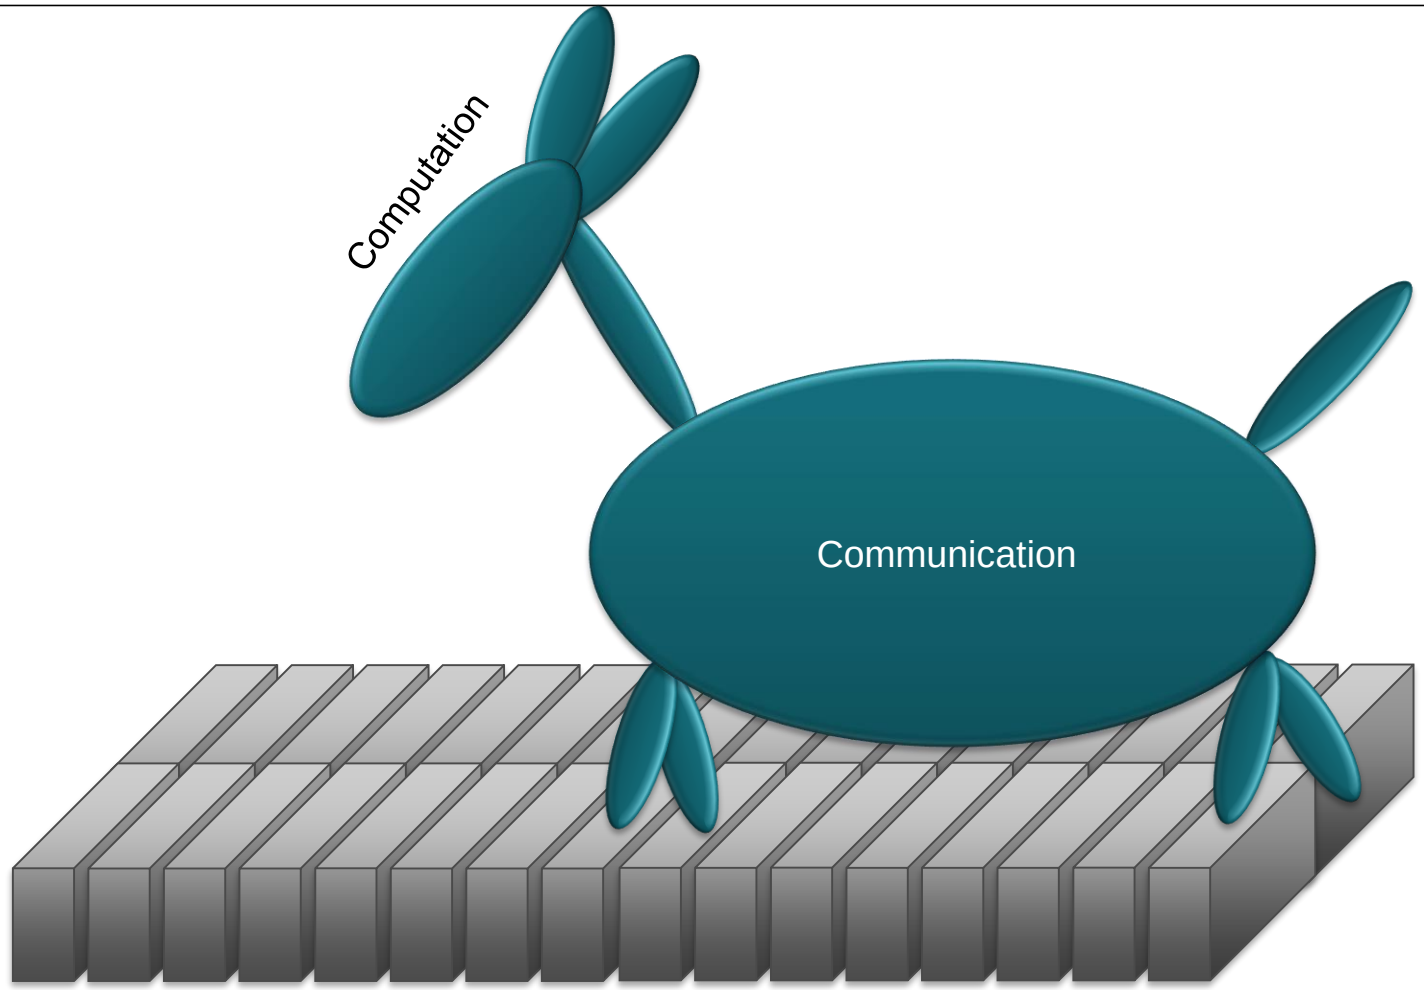

Communication

Scaling code to a bigger machine can unveil unpleasant surprises

TECHNISCHE UNIVERSITÄT DARMSTADT

We need to find scaling issues before they occur

We need:
a model for performance behavior

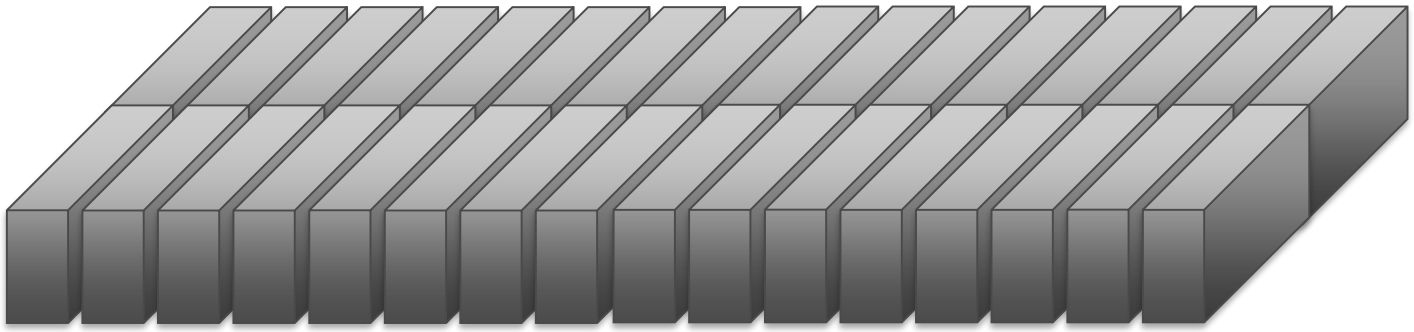

Performance model

TECHNISCHE UNIVERSITÄT DARMSTADT

Formula that expresses a relevant performance metric as a function of one or more execution parameters

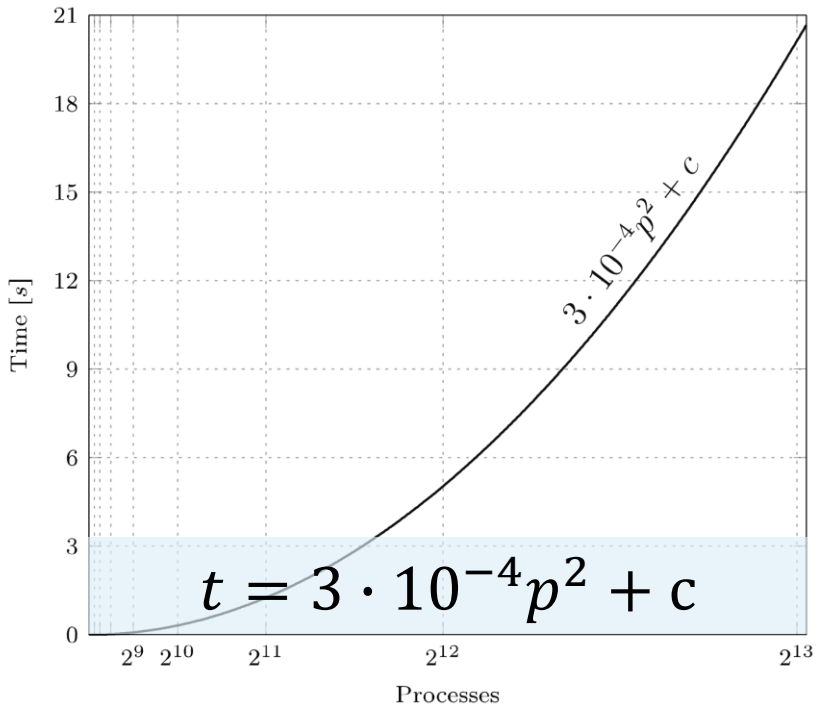

Manual creation challenging for entire programs

Identify kernels
- Incomplete coverage

Create models
- Laborious, difficult

Automatic empirical performance modelling with Extra-P

TECHNISCHE UNIVERSITÄT DARMSTADT

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

Input

e.g. Score-P

Instrumentation

Performance measurements

M_i M_j

Extra-P

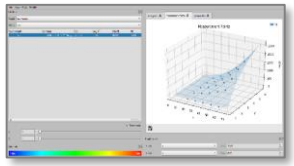

Output

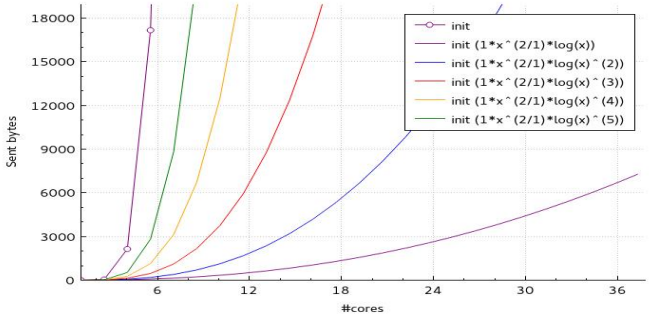

Human-readable, multi-parameter performance models

$$f(x_1, \ldots, x_m) = \sum_{k=1}^{n} c_k \prod_{l=1}^{m} x_l^{i_{kl}} \cdot log_2^{j_{kl}}(x_l)$$

Modelling process

TECHNISCHE UNIVERSITÄT DARMSTADT

Performance model normal form

$$f(x_1, \ldots, x_m) = \sum_{k=1}^{n} c_k \prod_{l=1}^{m} x_l^{i_{kl}} \cdot log_2^{j_{kl}}(x_l)$$

$$c_1$$
$$c_1 + c_2 \cdot x_1$$
$$c_1 + c_2 \cdot x_1 + c_3 \cdot x_2$$
$$c_1 + c_2 \cdot x_1 \cdot x_2^2$$
$$c_1 + c_2 \cdot x_1 + c_3 \cdot x_1 \cdot x_2^3 \cdot \log_2^2 x_2$$

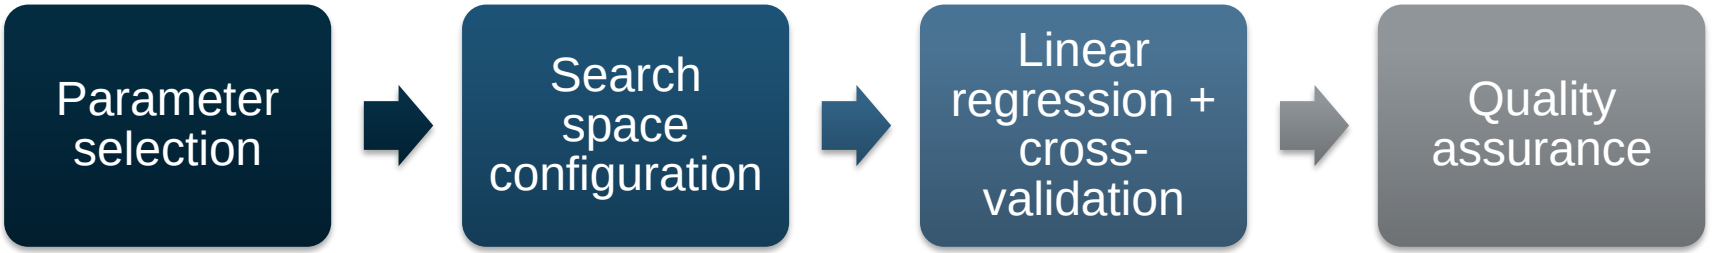

Assumptions & limitations

TECHNISCHE UNIVERSITÄT DARMSTADT

- Scaling behaviour expressible with performance model normal form
- Only one scaling behaviour for all the measurements; no jumps
- Some MPI collective operations switch their algorithm
  - results in bad models
- Example: red model tries to model measurements of different algorithms
  - First 4 points – one function
  - Last 4 points – another function (linear)

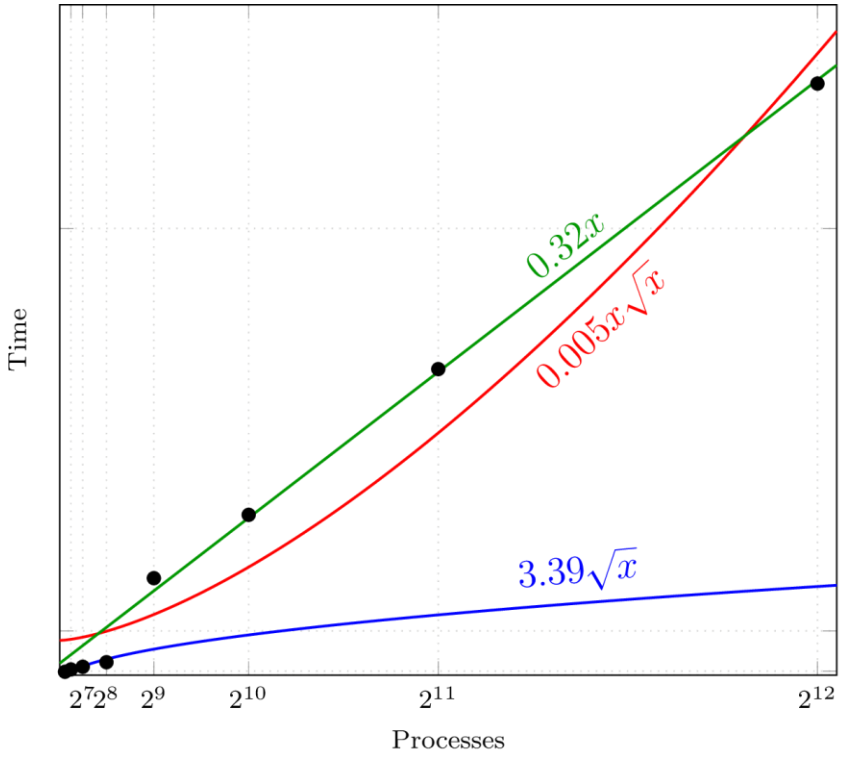

Modelling application requirements

TECHNISCHE UNIVERSITÄT DARMSTADT

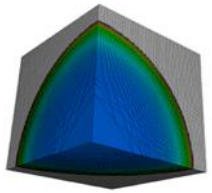

Lulesh

Models represent **per-process** effects

p – number of processes
n – problem size per process

Requirement	Metric	Model
Computation	#FLOPs	$10^5 \cdot n \cdot \log(n) \cdot p^{0.25} \cdot \log(p)$
Communication	#Bytes sent & received	$10^3 \cdot n \cdot p^{0.25} \cdot \log(p)$
Memory access	#Loads & stores	$10^5 \cdot n \cdot \log(n) \cdot \log(p)$
Memory footprint	#Bytes used	$10^5 \cdot n \cdot \log(n)$
Memory locality	Stack distance	$\mathrm{Constant}$

TECHNISCHE UNIVERSITÄT DARMSTADT

DEEP-SEA

Extensions for heterogeneous systems

UPCOMING FEATURES

GPU Applications

TECHNISCHE UNIVERSITÄT DARMSTADT

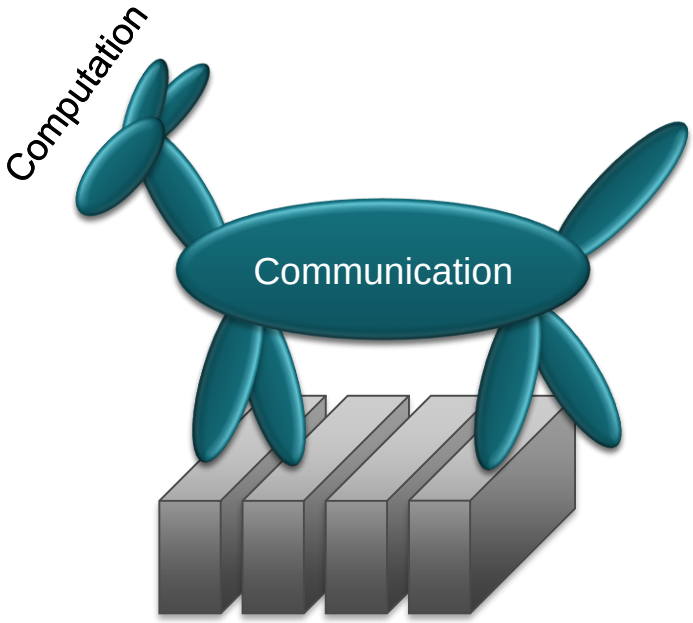

GPU Applications

TECHNISCHE UNIVERSITÄT DARMSTADT

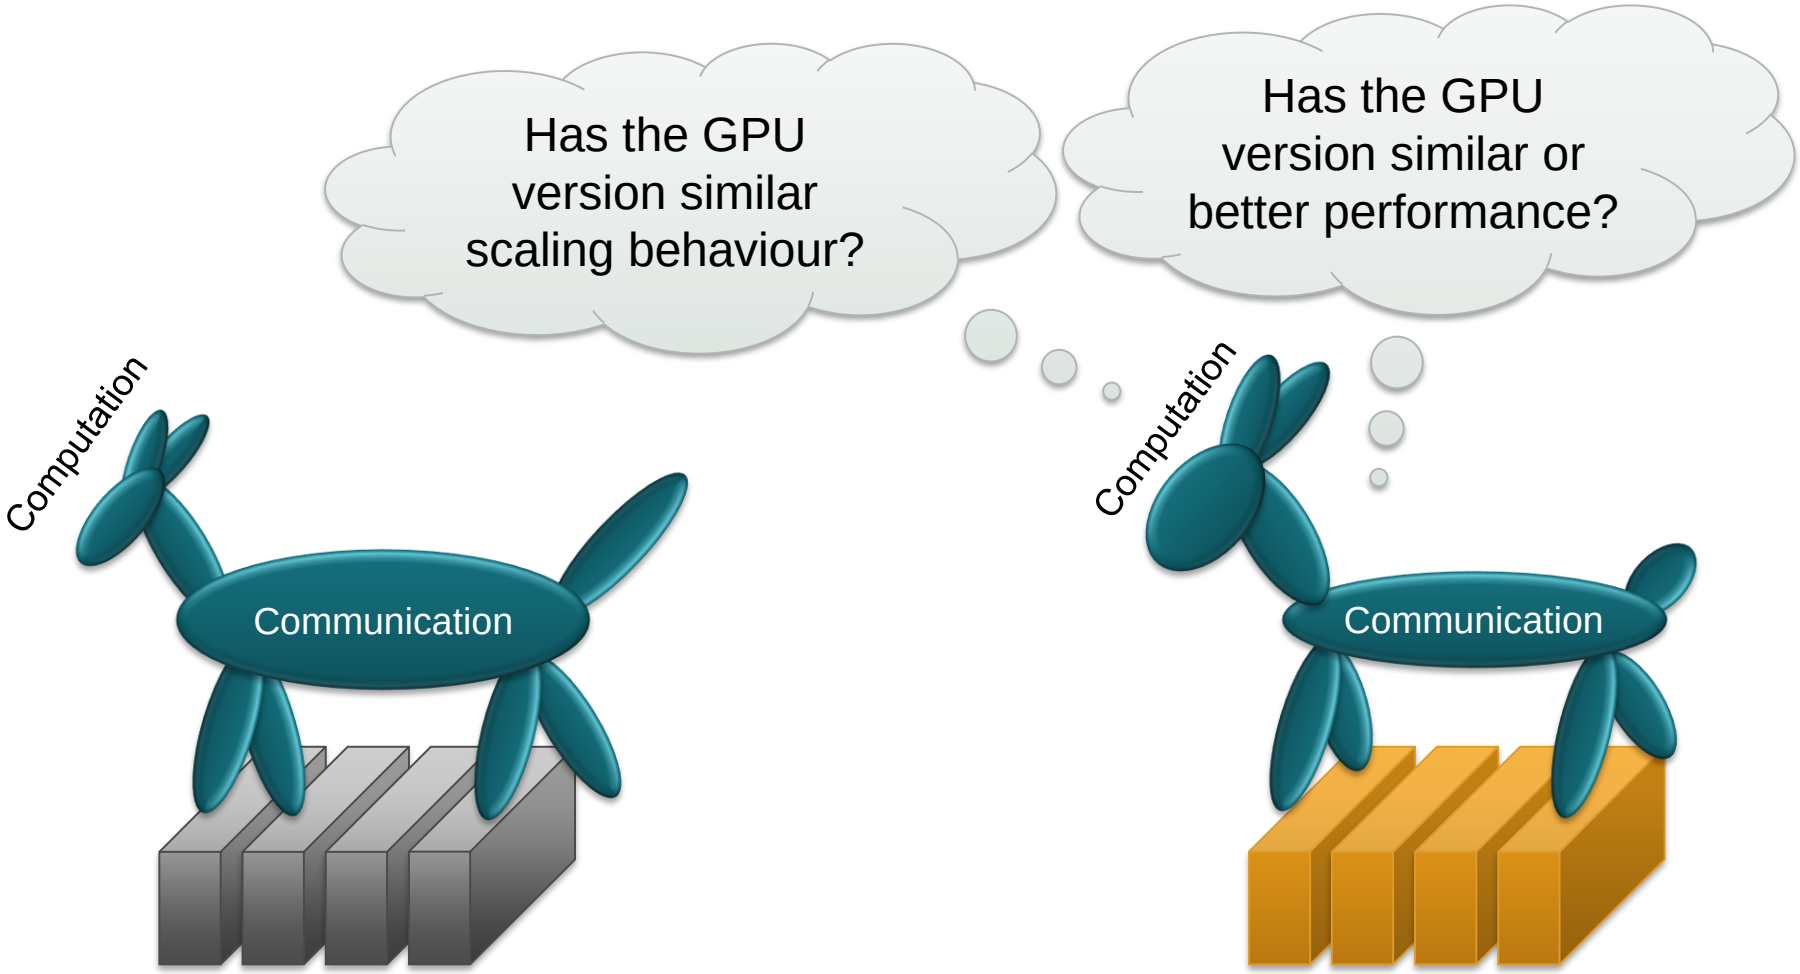

Generating GPU models

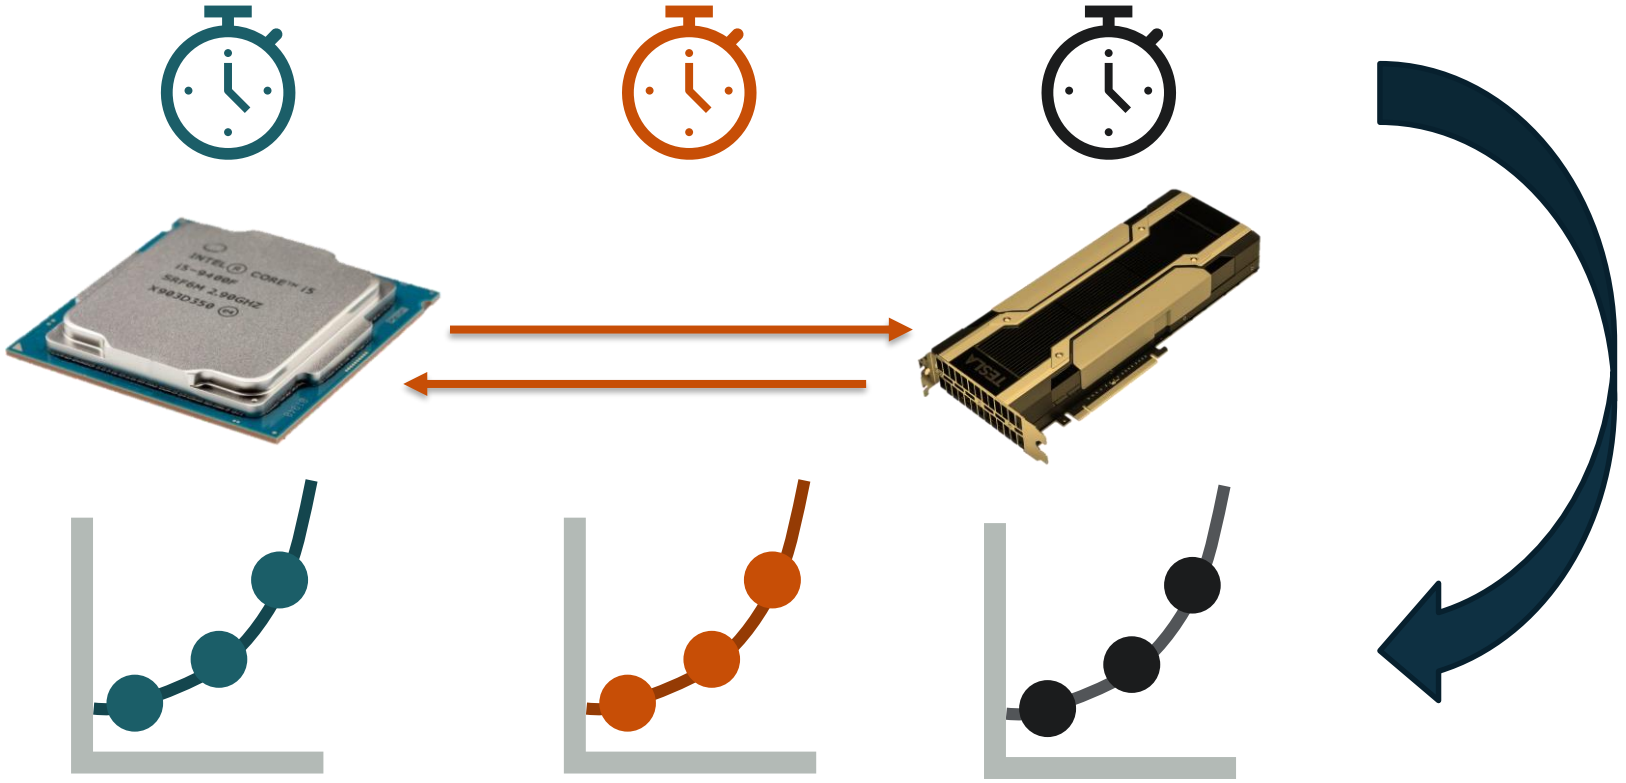

Host Computation
Synchronization with GPU

Data Transfer

Kernel
Idle

Usage

Has the GPU version similar scaling behaviour?

Has the GPU version similar or better performance?

TECHNISCHE UNIVERSITÄT DARMSTADT

- Is runtime of CPU only > host computation + data transfer time + runtime on accelerator?
  - Comparison of CPU app model with models for host computation, data transfer and kernels

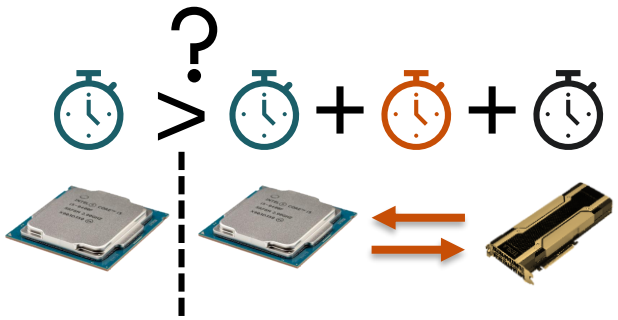

- How much work can the CPU do, while the GPU is doing the offloaded work?
  - Synchronization model

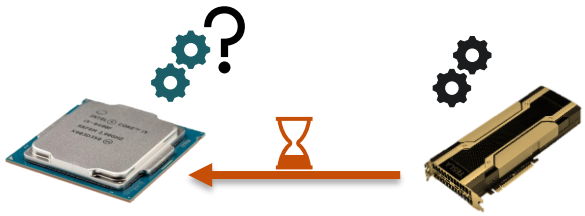

- Is the GPU well utilized?
  - Idle time model

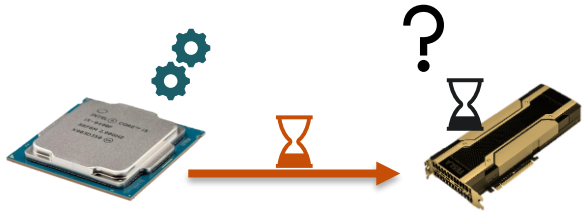

How to map an app onto MSA systems?

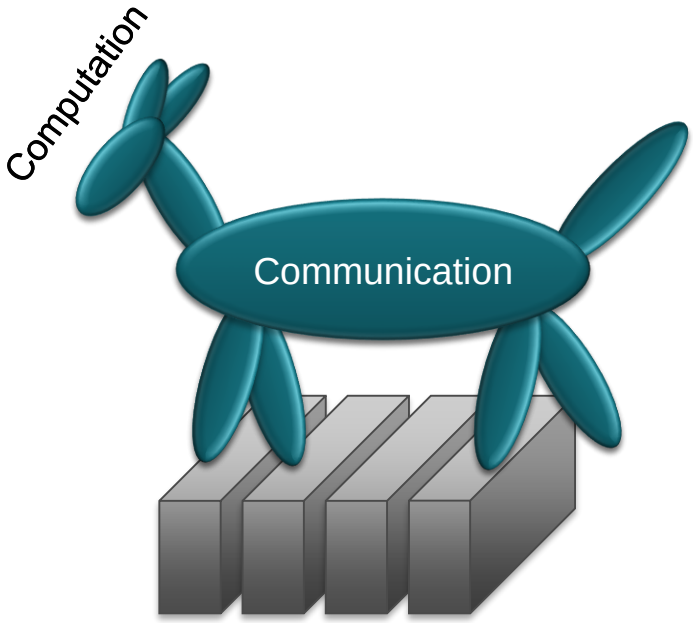

How to map an app onto MSA systems?

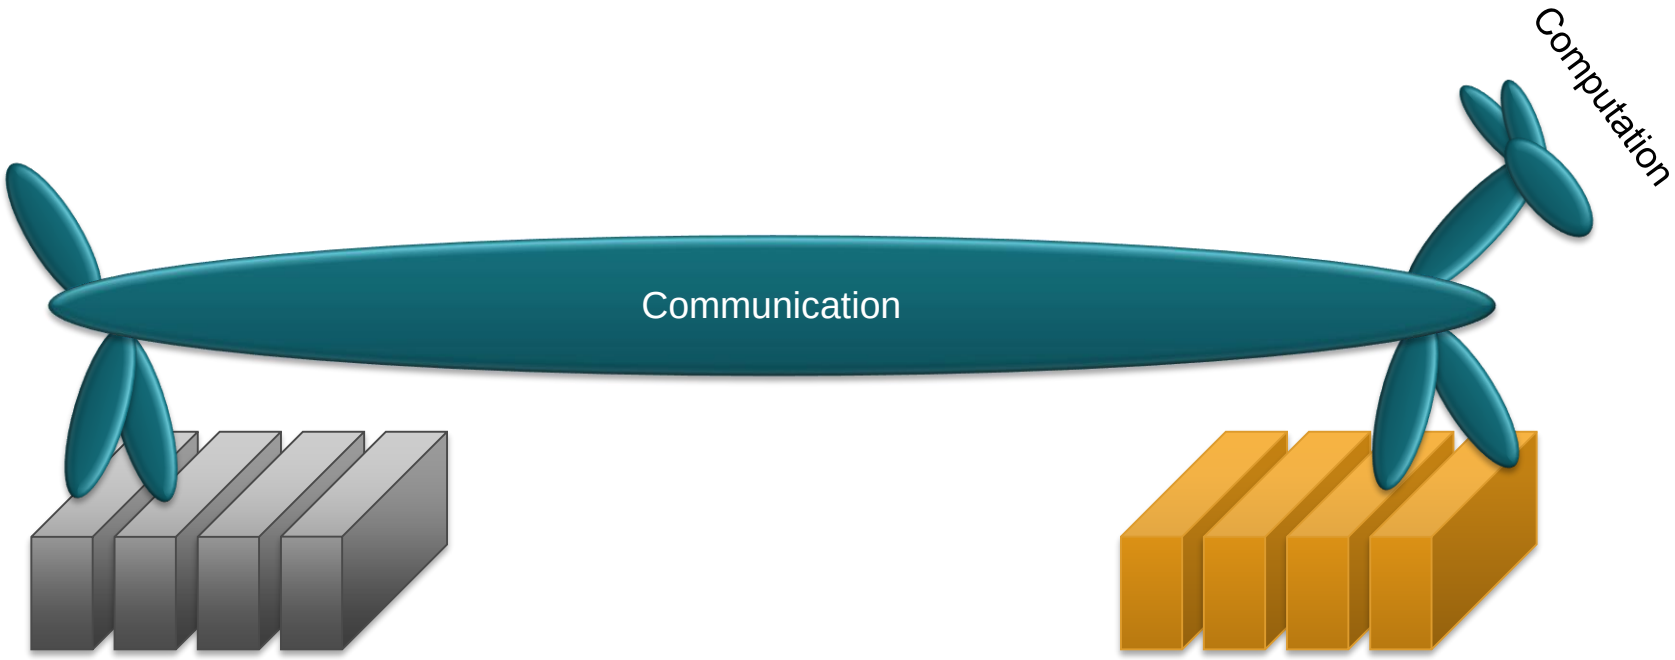

How to map an app onto MSA systems?

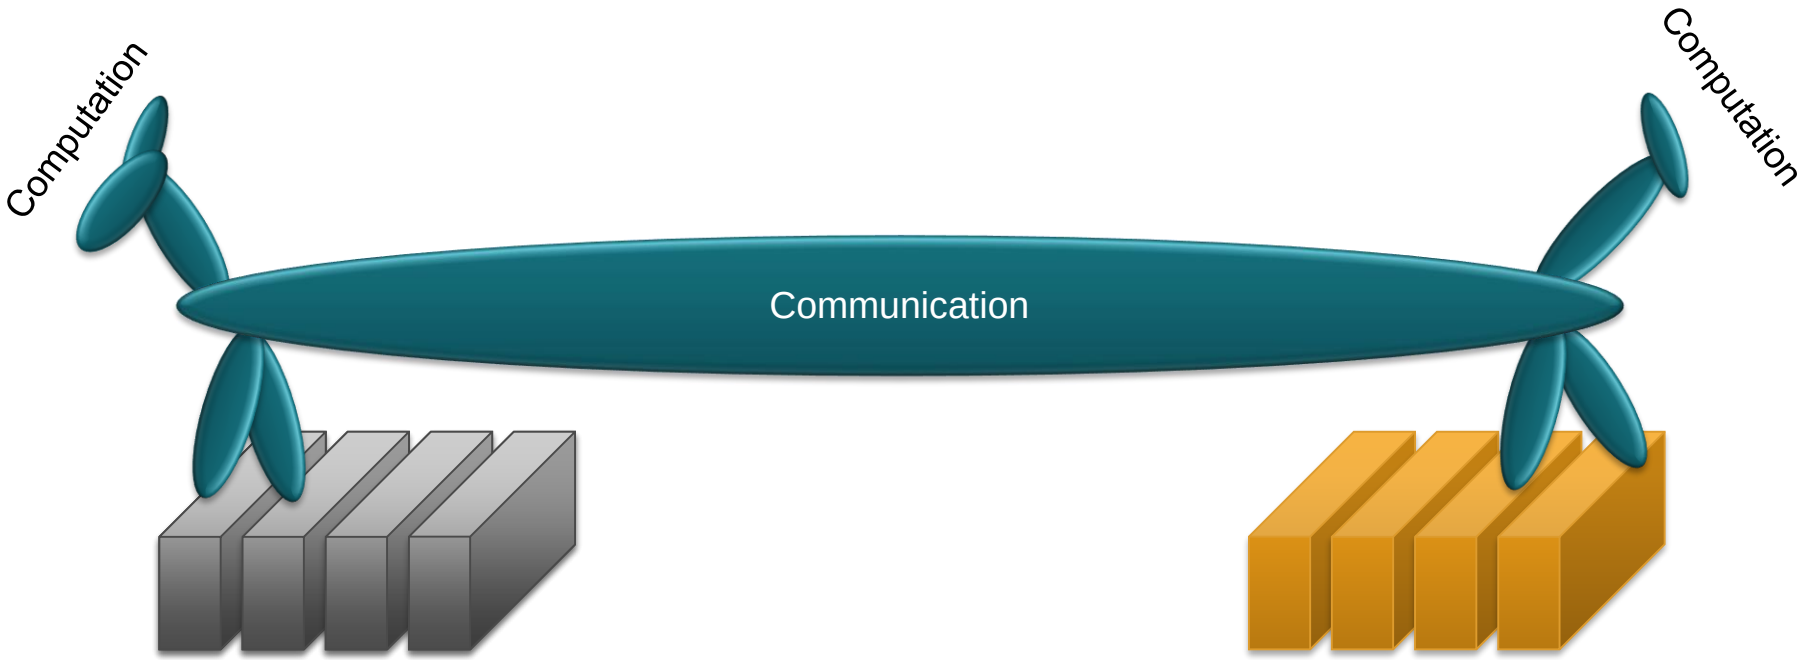

How to map an app onto MSA systems?

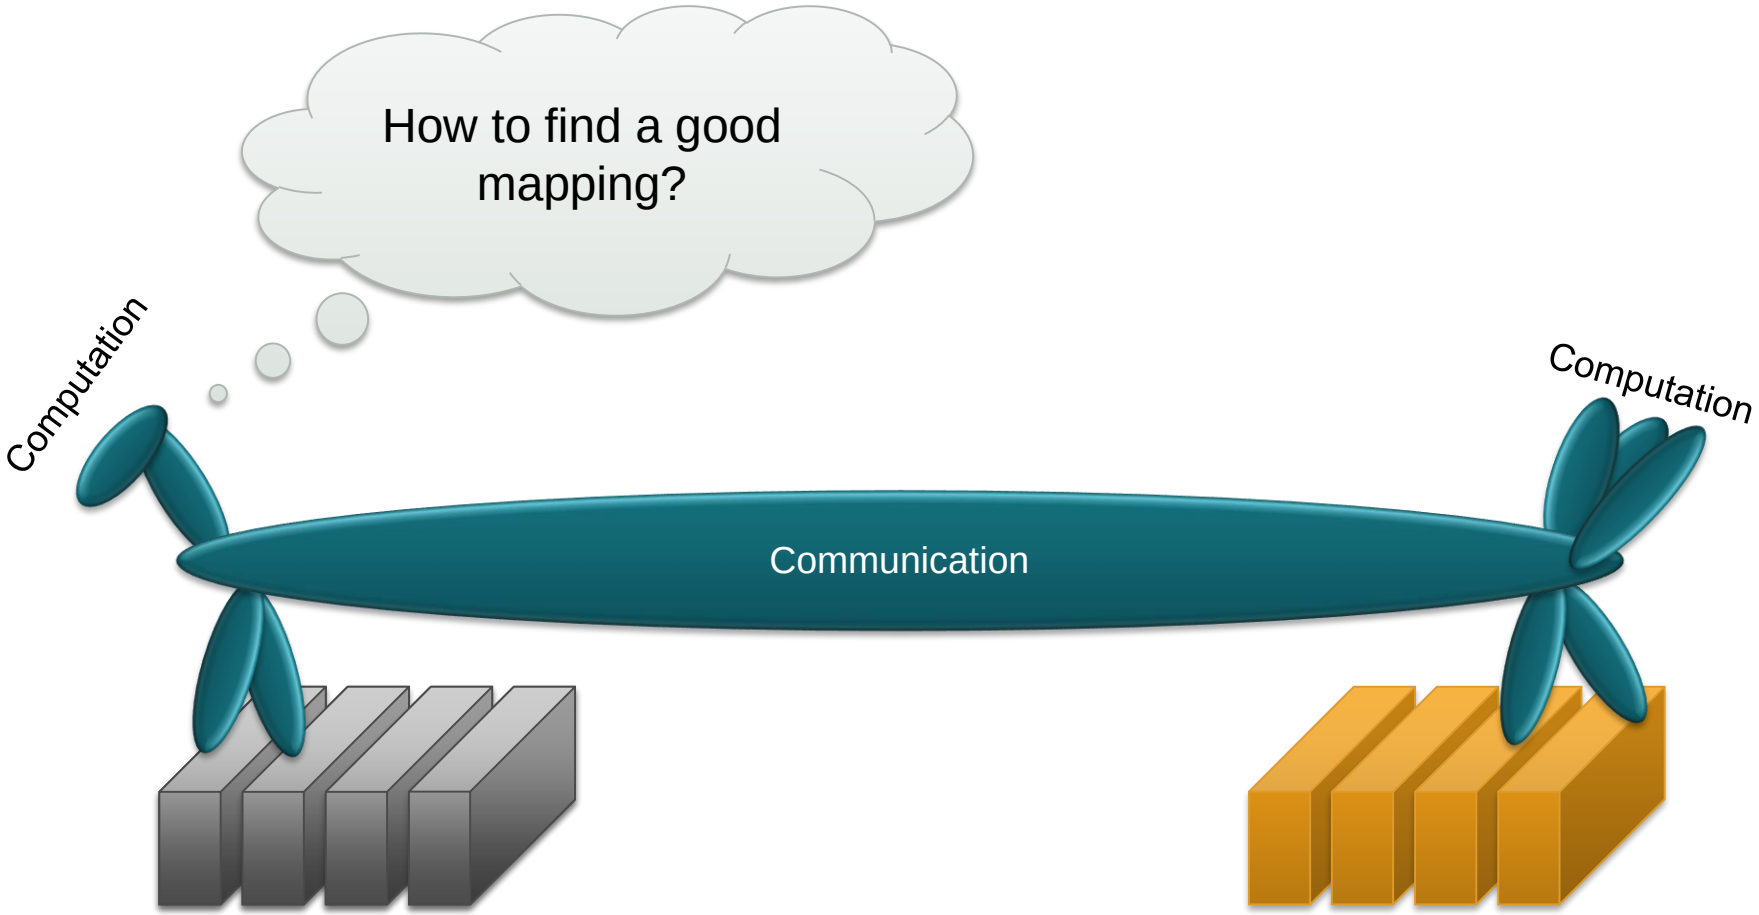

Objectives

TECHNISCHE UNIVERSITÄT DARMSTADT

Design a strategy for mapping application parts to the MSA modules

- Create support for model comparison
- Create portable performance models
- Use these models to determine best target modules
- Reduce user involvement, if possible

Summary

TECHNISCHE UNIVERSITÄT DARMSTADT

- Applications can exhibit unwanted performance behaviour when scaling up
- Performance models help to find issues before they occur
  - Laborious to do by hand
  - Extra-P automates this step
- Extra-P will also support GPUs
  - Helps checking for optimization opportunities and unwanted behaviour
- Extra-P will assist in mapping of applications onto MSA systems

Check your app with Extra-P before scaling up!

Find it on GitHub:
https://github.com/extra-p/extrap