

# Extra-P: Insightful Automatic Performance Modeling



TECHNISCHE  
UNIVERSITÄT  
DARMSTADT

---

Alexander Geiß<sup>1</sup>, Marcus Ritter<sup>1</sup>, Benedikt Naumann<sup>1</sup>,  
Alexandru Calotoiu<sup>2</sup>, Torsten Hoefler<sup>2</sup>, and Felix Wolf<sup>1</sup>

---

<sup>1</sup> TU Darmstadt , <sup>2</sup> ETH Zürich

**ETH** zürich

# Spectrum of performance analysis methods

---

Benchmark

Full simulation

Model simulation

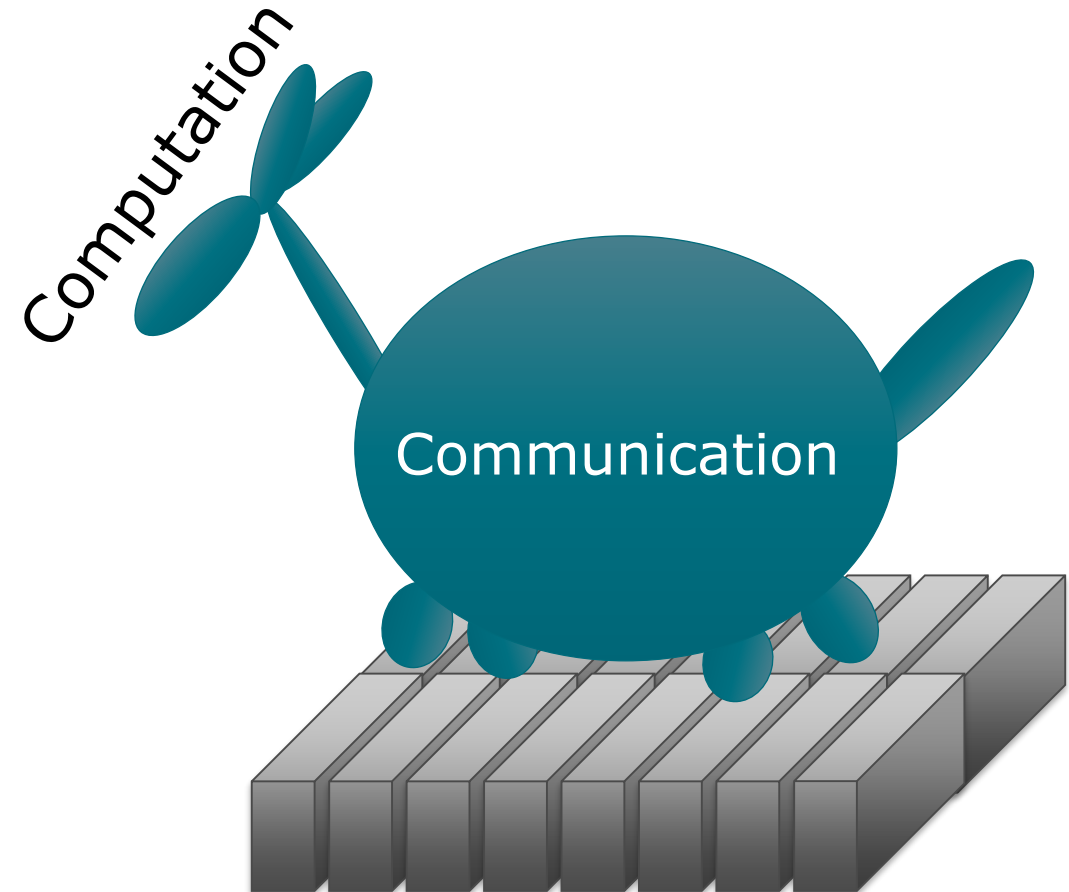
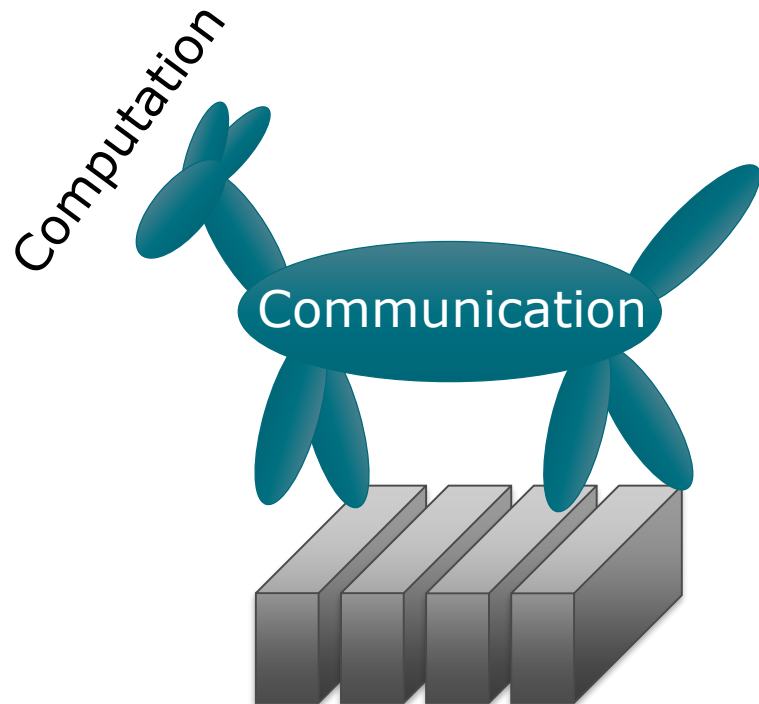
Model



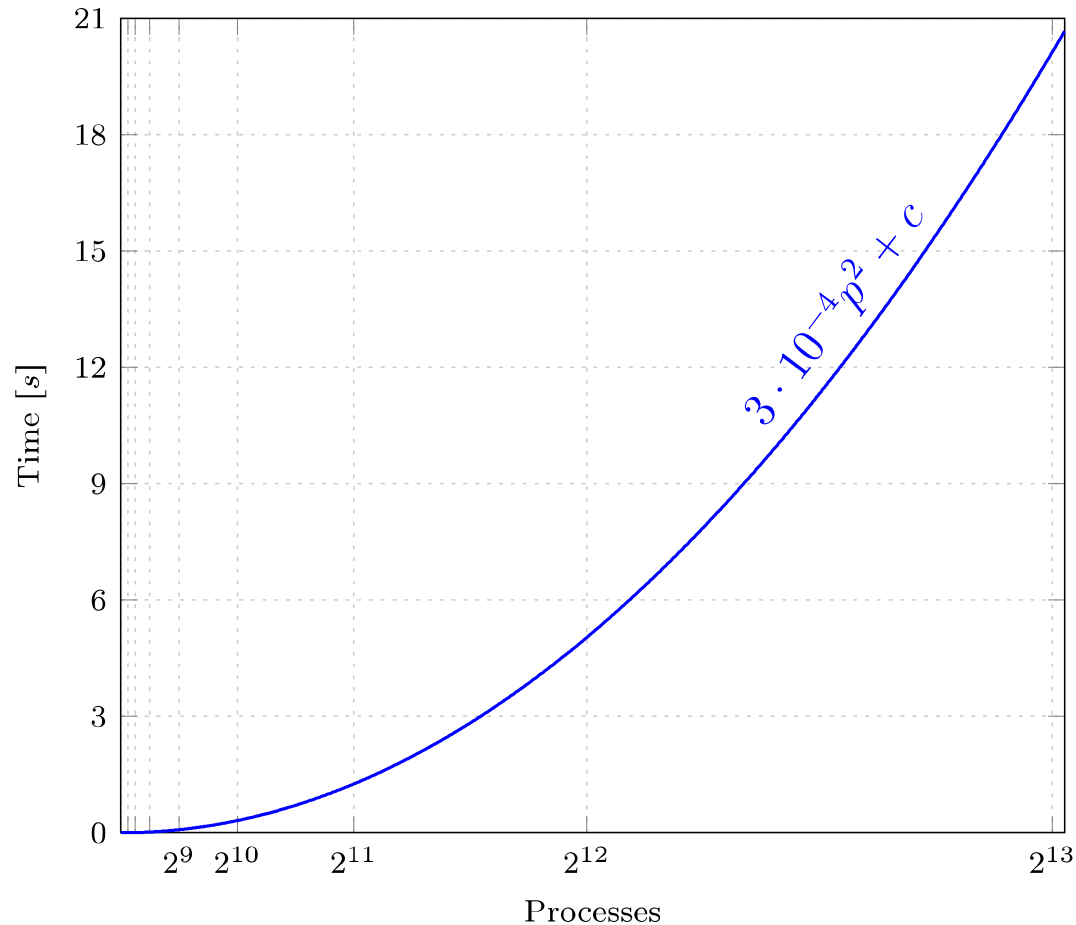
© 2011 IEEE. Reprinted, with permission, from T. Hoefler, W. Gropp, W. Kramer and M. Snir, "Performance modeling for systematic performance tuning," SC '11.

## Motivation - latent scalability bugs

---



# Scaling model



- Represents performance metric as a function of the number of processes
- Provides insight into the program behavior at scale

# Analytical performance modeling

## Identify kernels

- Parts of the program that dominate its performance at larger scales
- Identified via small-scale tests and intuition

## Create models

- Laborious process
- Still confined to a small community of skilled experts

## Disadvantages:

- Time consuming
- Danger of overlooking unscalable code



Hoisie et al.: *Performance and scalability analysis of teraflop-scale parallel architectures using multi-dimensional wavefront applications*. International Journal of High Performance Computing Applications, 2000

Bauer et al.: *Analysis of the MILC Lattice QCD Application su3\_rmd*. CCGrid, 2012

# Automatic performance modeling

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

**Input**

**Instrumentation**

- All functions

Performance measurements

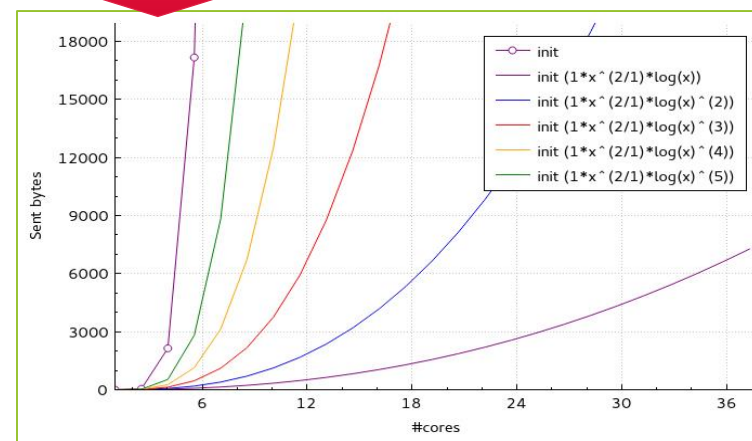
$M_i$

$M_j$

**Extra-P**

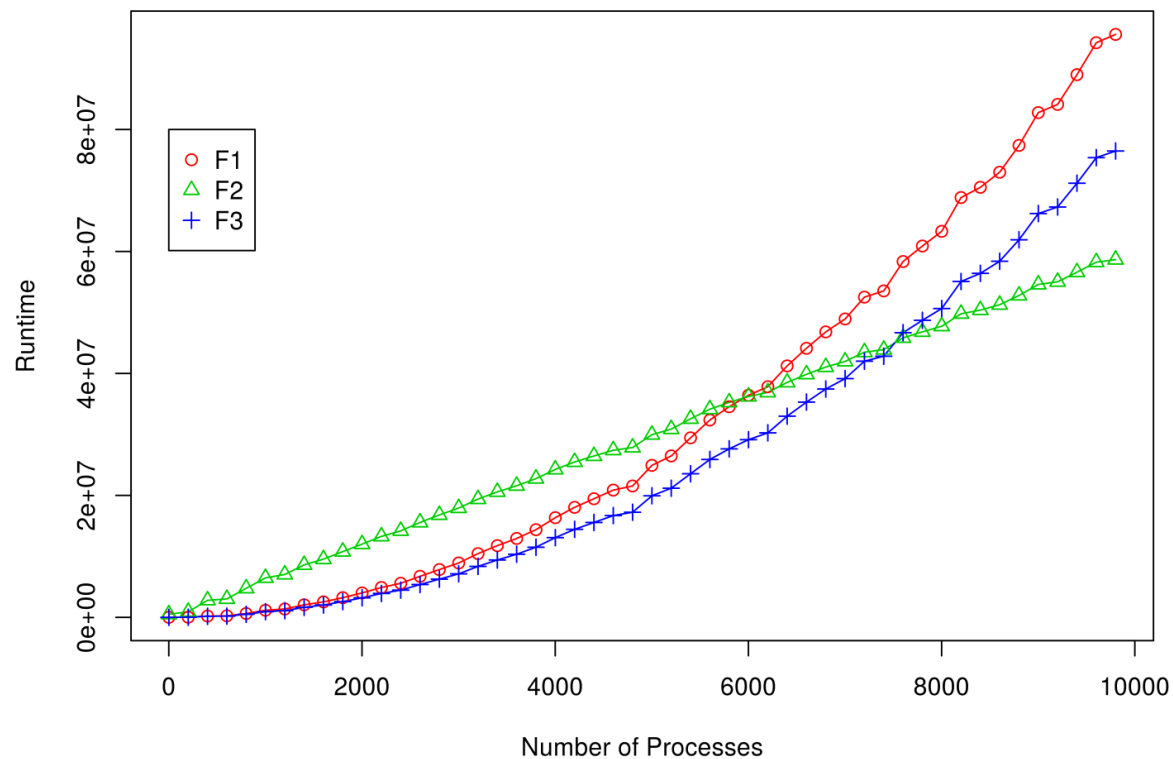
**Output**

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

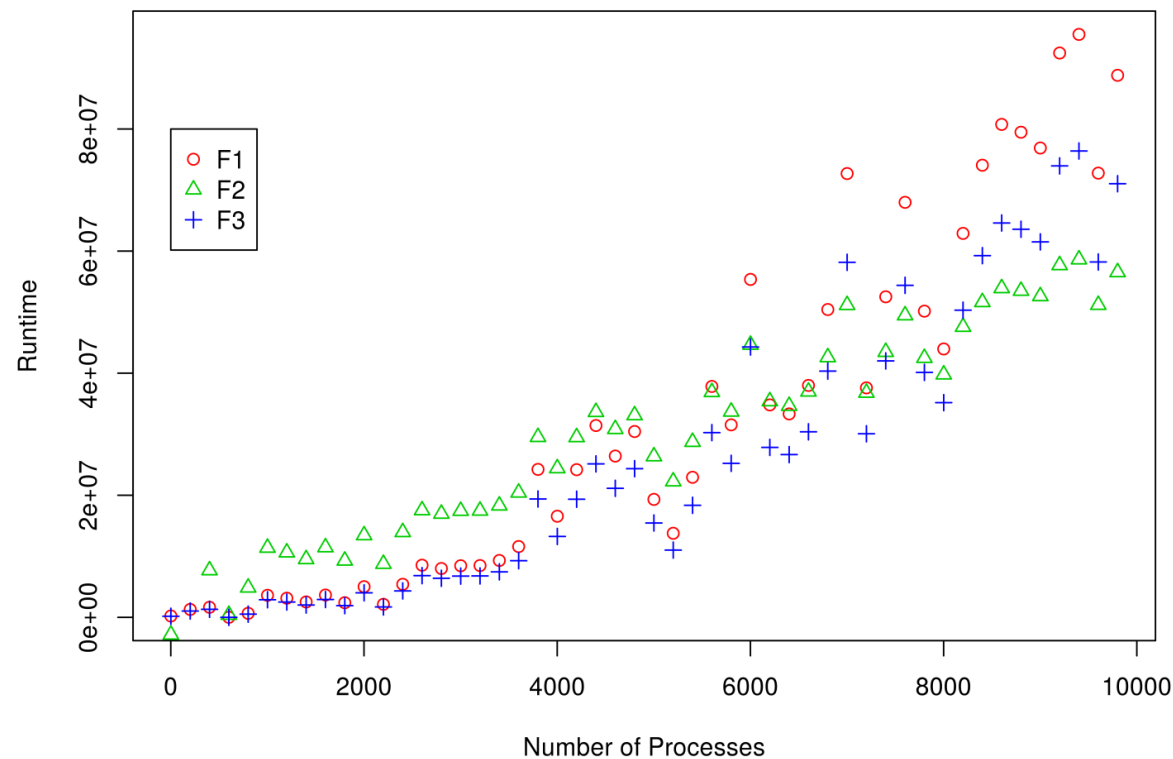


# Primary focus on scaling trend

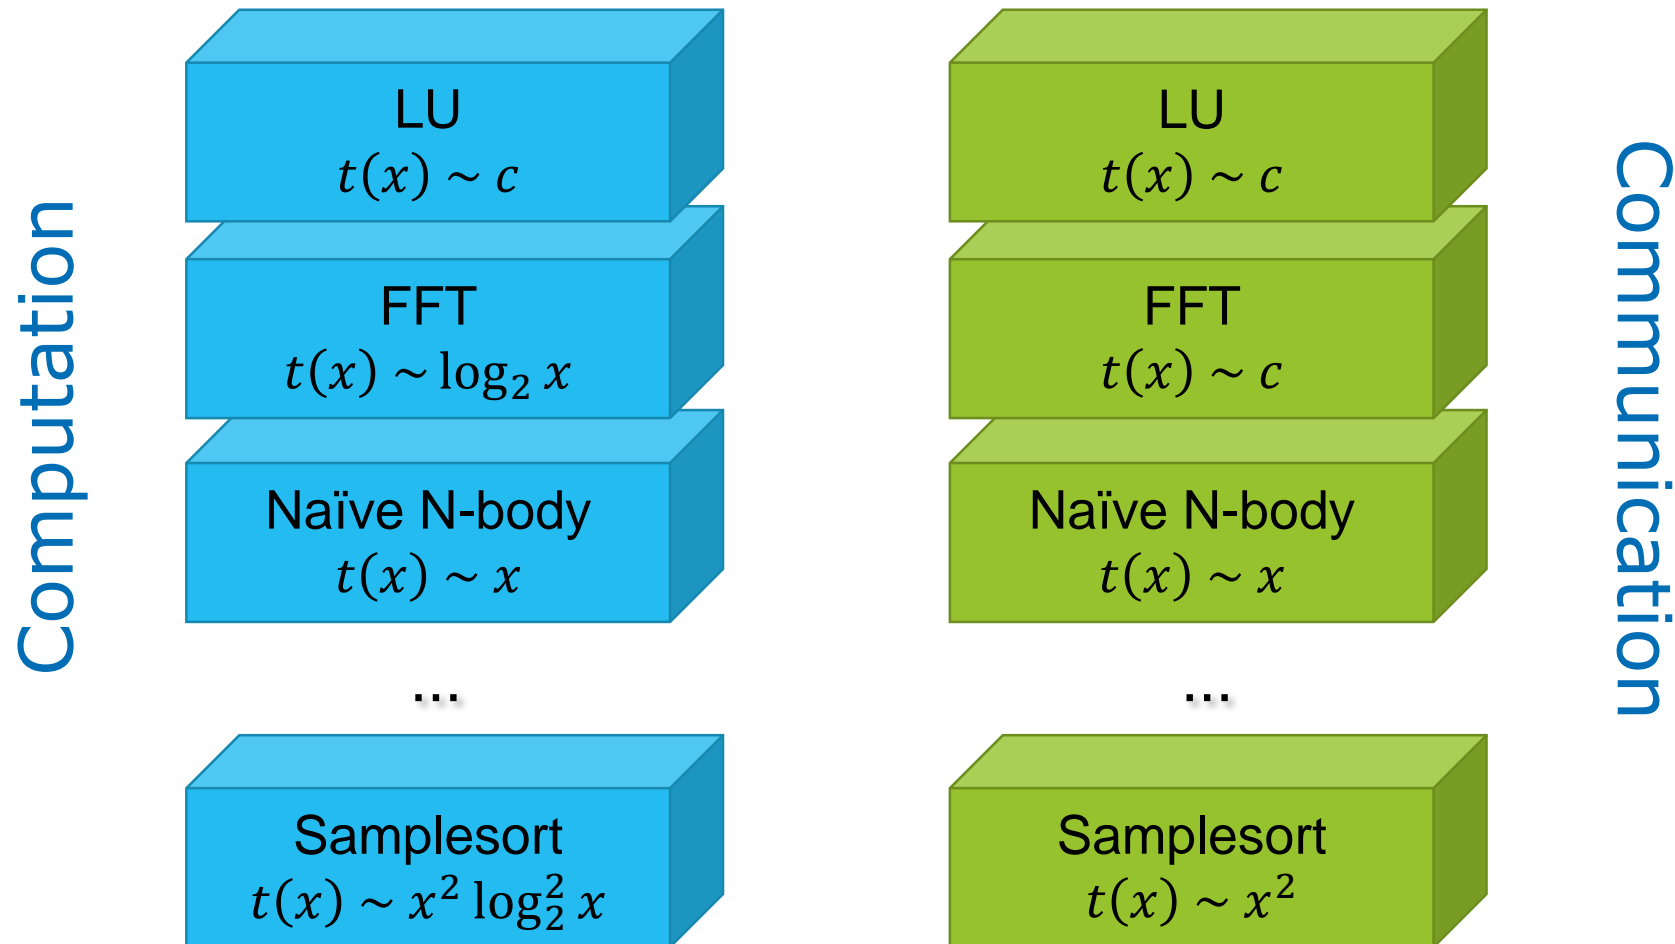
Common performance analysis chart in a paper



Production Reality



# Model building blocks





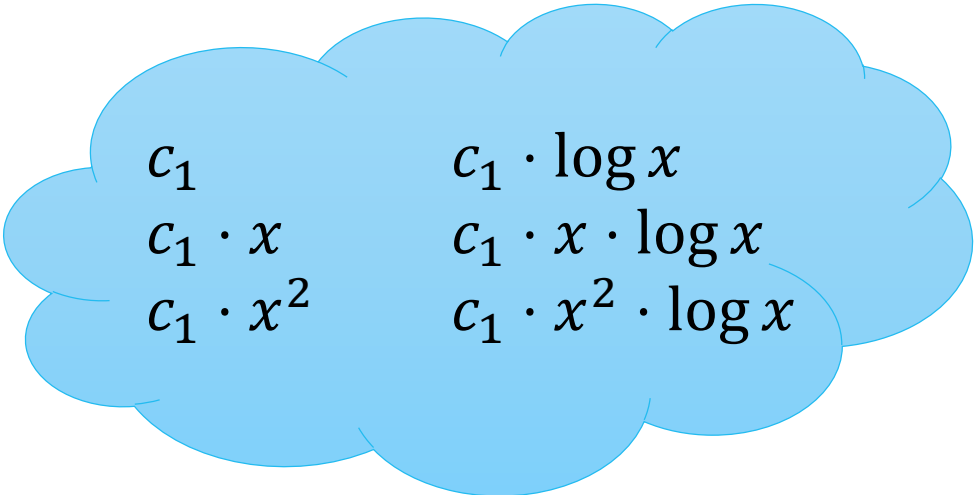
# Performance model normal form

---

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

$$\begin{aligned} n &\in \mathbb{N} \\ i_k &\in I \\ j_k &\in J \\ I, J &\subset \mathbb{Q} \end{aligned}$$

$$\begin{aligned} n &= 1 \\ I &= \{0, 1, 2\} \\ J &= \{0, 1\} \end{aligned}$$


$$\begin{array}{ll} c_1 & c_1 \cdot \log x \\ c_1 \cdot x & c_1 \cdot x \cdot \log x \\ c_1 \cdot x^2 & c_1 \cdot x^2 \cdot \log x \end{array}$$

## Performance model normal form

$$n = 2$$

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

$$J = \{0, 1\}$$

$$c_1$$

$$c_1 \cdot x$$

$$c_1 \cdot x^2$$

$$c_1 \cdot \log x$$

$$c_1 \cdot x \cdot \log x$$

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

$$c_1 \cdot \log x + c_2 \cdot x$$

$$c_1 \cdot \log x + c_2 \cdot x \cdot \log x$$

$$c_1 \cdot \log x + c_2 \cdot x^2$$

$$c_1 \cdot \log x + c_2 \cdot x^2 \cdot \log x$$

$$c_1 \cdot x + c_2 \cdot x \cdot \log x$$

$$c_1 \cdot x + c_2 \cdot x^2$$

$$c_1 \cdot x + c_2 \cdot x^2 \cdot \log x$$

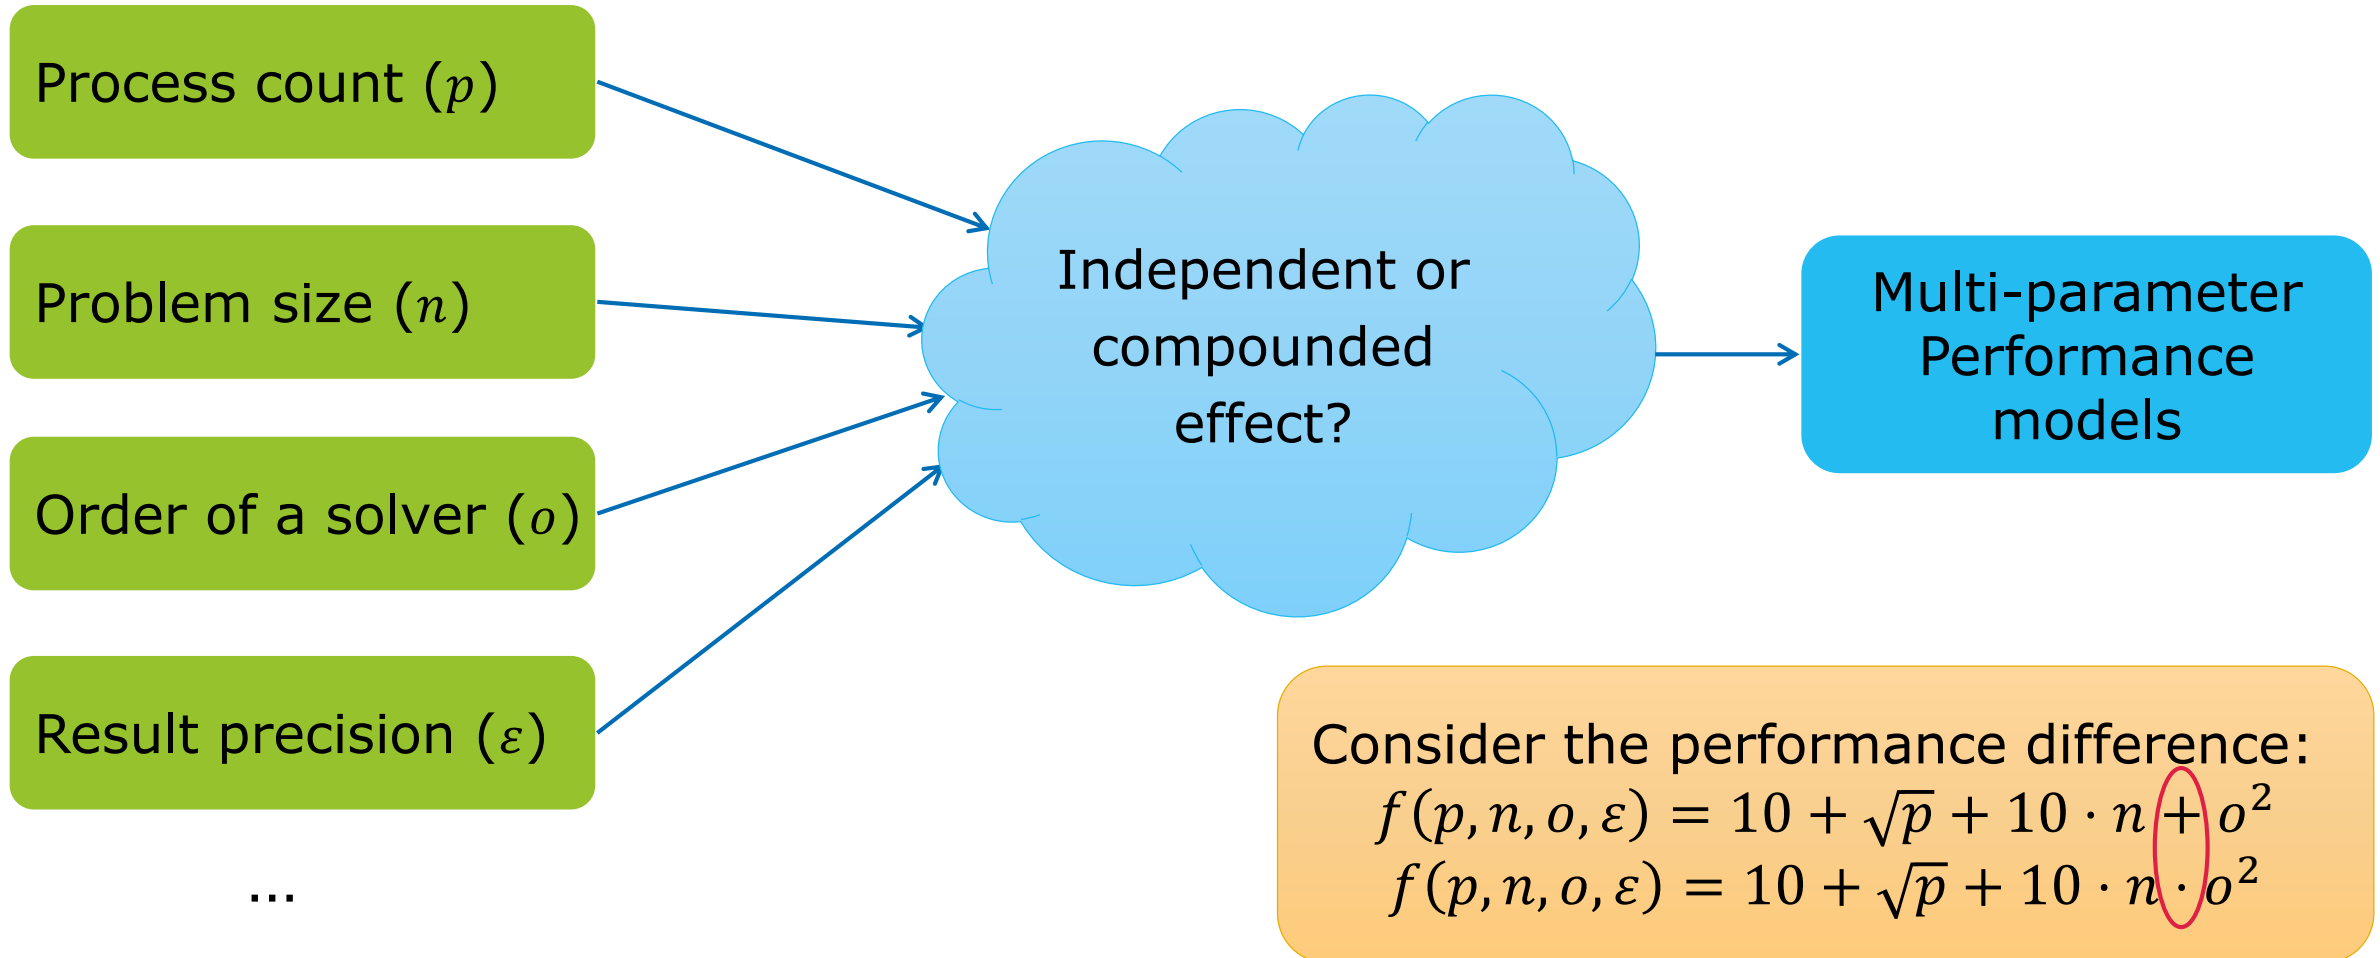
$$c_1 \cdot x \cdot \log x + c_2 \cdot x^2$$

$$c_1 \cdot x \cdot \log x + c_2 \cdot x^2 \cdot \log x$$

$$c_1 \cdot x^2 + c_2 \cdot x^2 \cdot \log x$$

$$\in \mathbb{N}$$

# Fast multi-parameter performance modeling



# Fast multi-parameter performance modeling

- Expanded performance model normal form

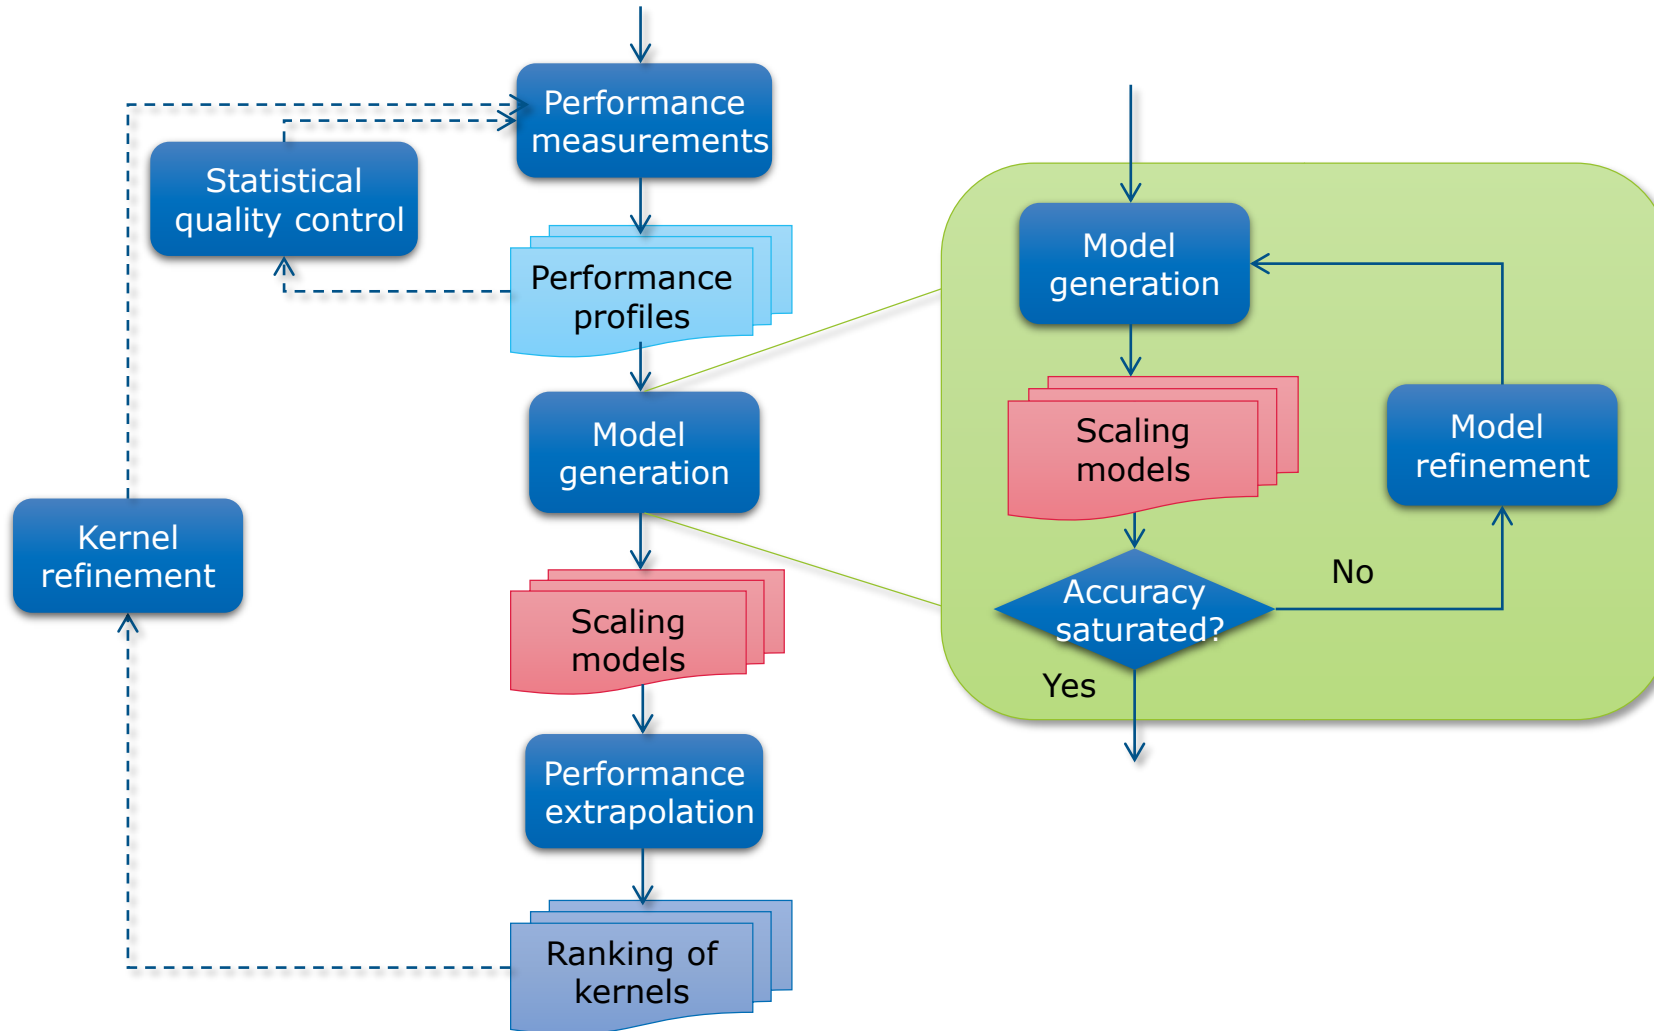
$$f(x_1, \dots, x_m) = \sum_{k=1}^n c_k \prod_{l=1}^m x_l^{i_{kl}} \cdot \log_2^{j_{kl}}(x_l)$$

$$\begin{aligned} m, n &\in \mathbb{N} \\ i_k &\in I \\ j_k &\in J \\ I, J &\subset \mathbb{Q} \end{aligned}$$

## Model candidates

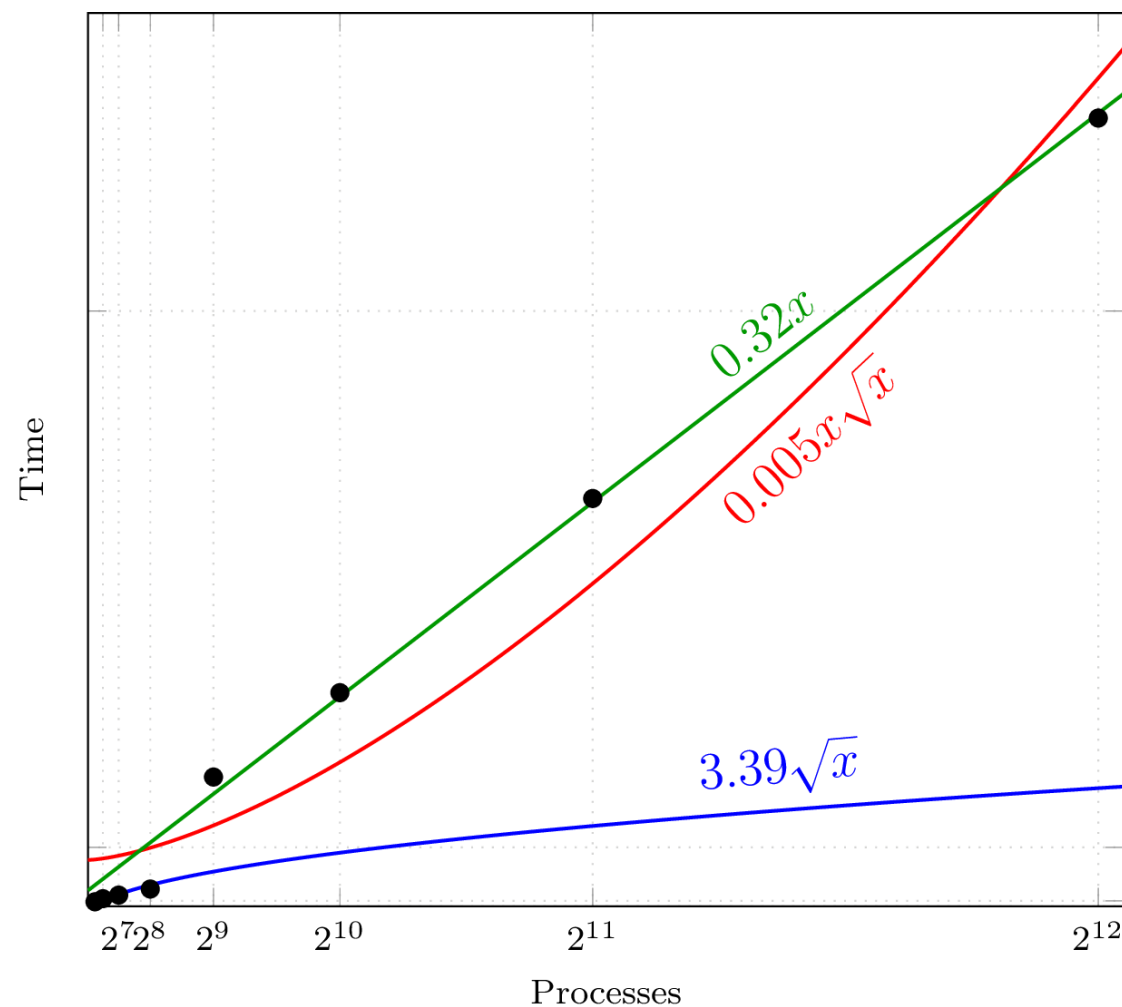
- Constant  $c_1$
- Single parameter  $c_1 + c_2 \cdot x_1$
- Multiple parameters  $\dots$ 
  - Additive  $c_1 + c_2 \cdot x_1 + c_3 \cdot x_2$
  - Multiplicative  $c_1 + c_2 \cdot x_1 \cdot x_2$
  - Complex  $c_1 + c_2 \cdot x_1 \cdot x_2 + c_3 \cdot \log x_2 \cdot x_2^3$

# Workflow



## Assumptions & limitations

- Scaling behavior expressible with performance model normal form
- Only one scaling behavior 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)
  - Adj. R2 = 0.95085 (!)



# Performance measurements

---

- Different ways of collecting measurements
- Score-P (<http://www.vi-hps.org/projects/score-p/>)
- Other profiling tools, e.g. HPCToolkit
- Manual ad-hoc measurements



## Performance measurements (2)

- At least 5 different measurements recommended

Performance measurements (profiles)

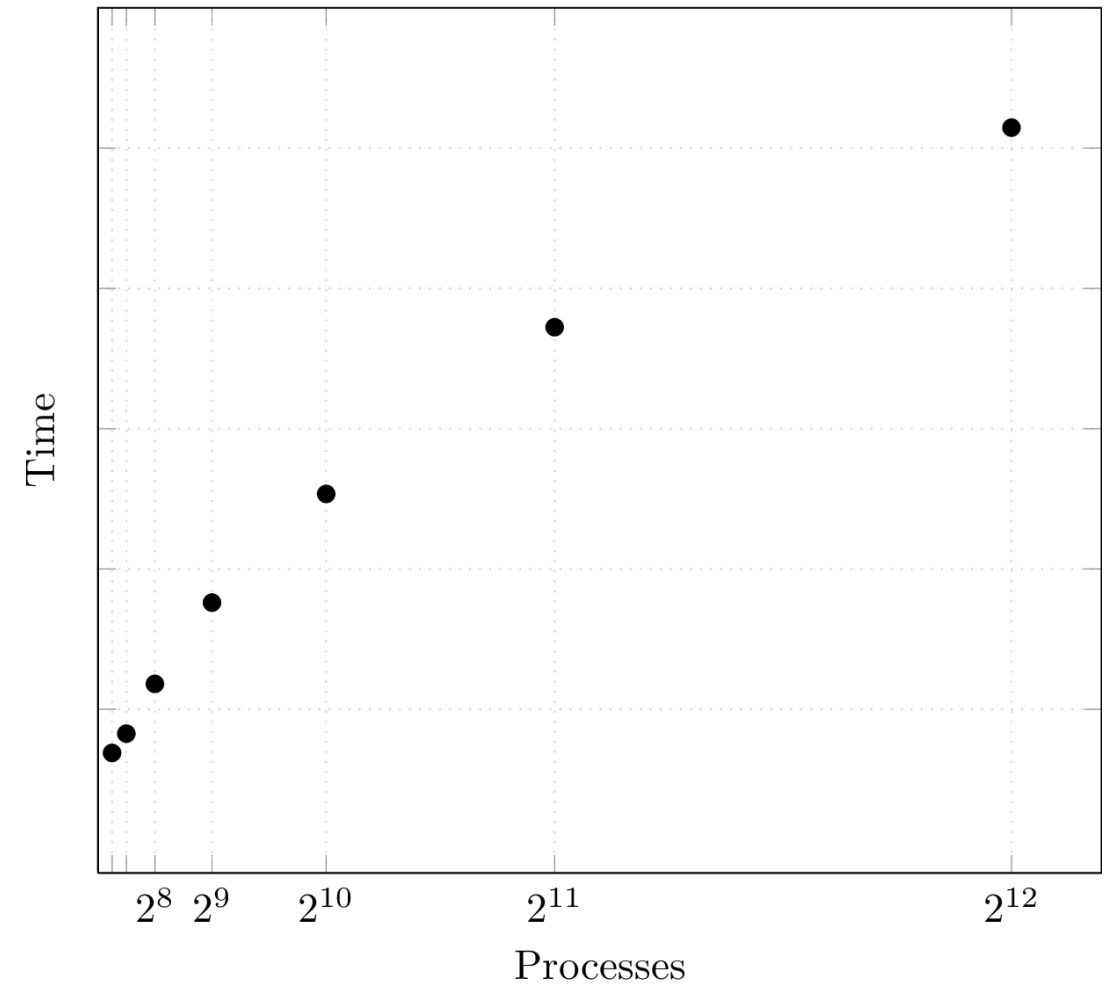
$$p_1 = 256$$

$$p_2 = 512$$

$$p_3 = 1024$$

$$p_4 = 2048$$

$$p_5 = 4096$$

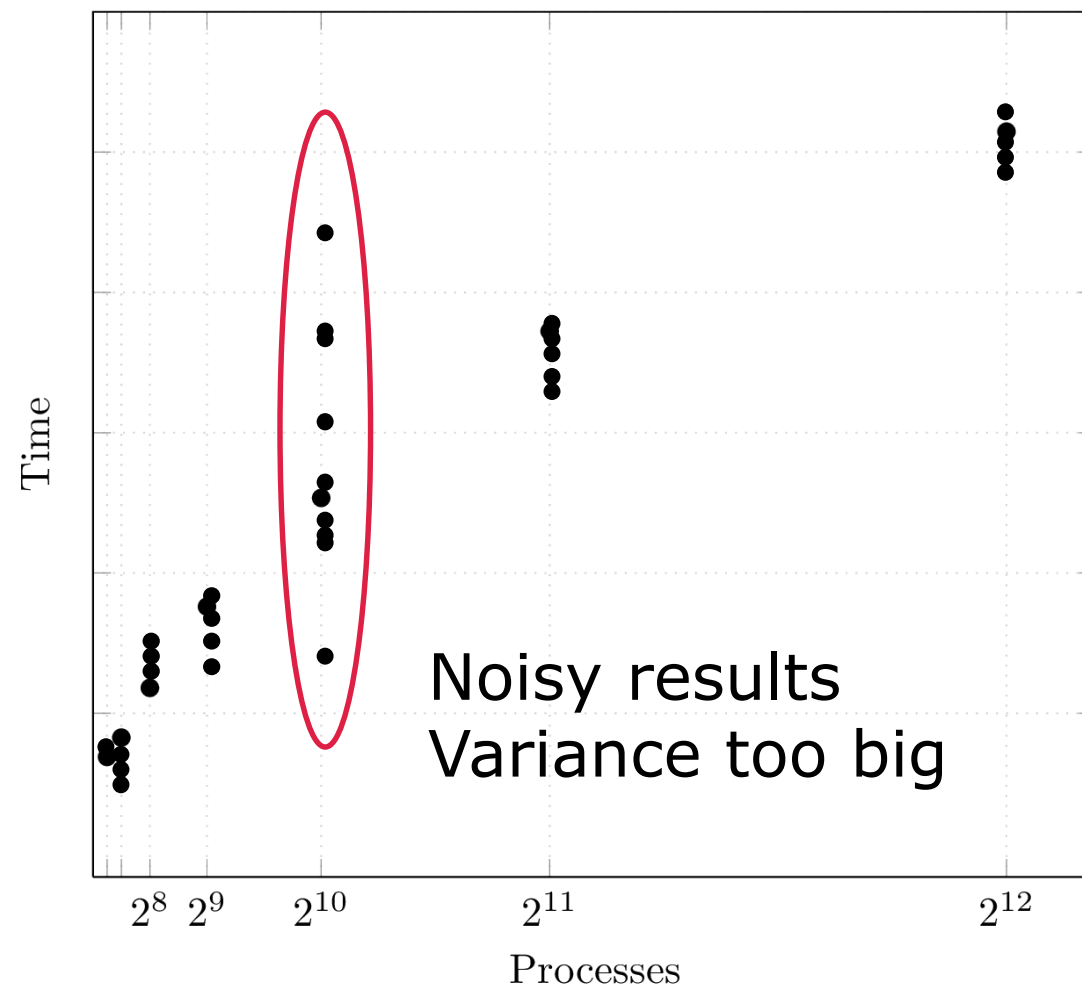
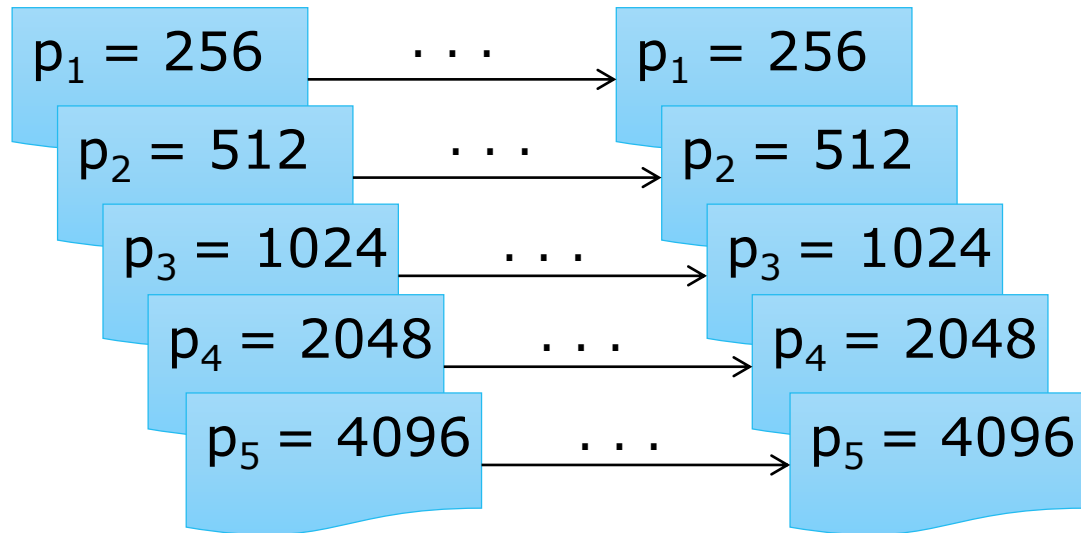




## Performance measurements (3)

- At least 5 different measurements recommended
- Each measurement repeated multiple times

Performance measurements (profiles)

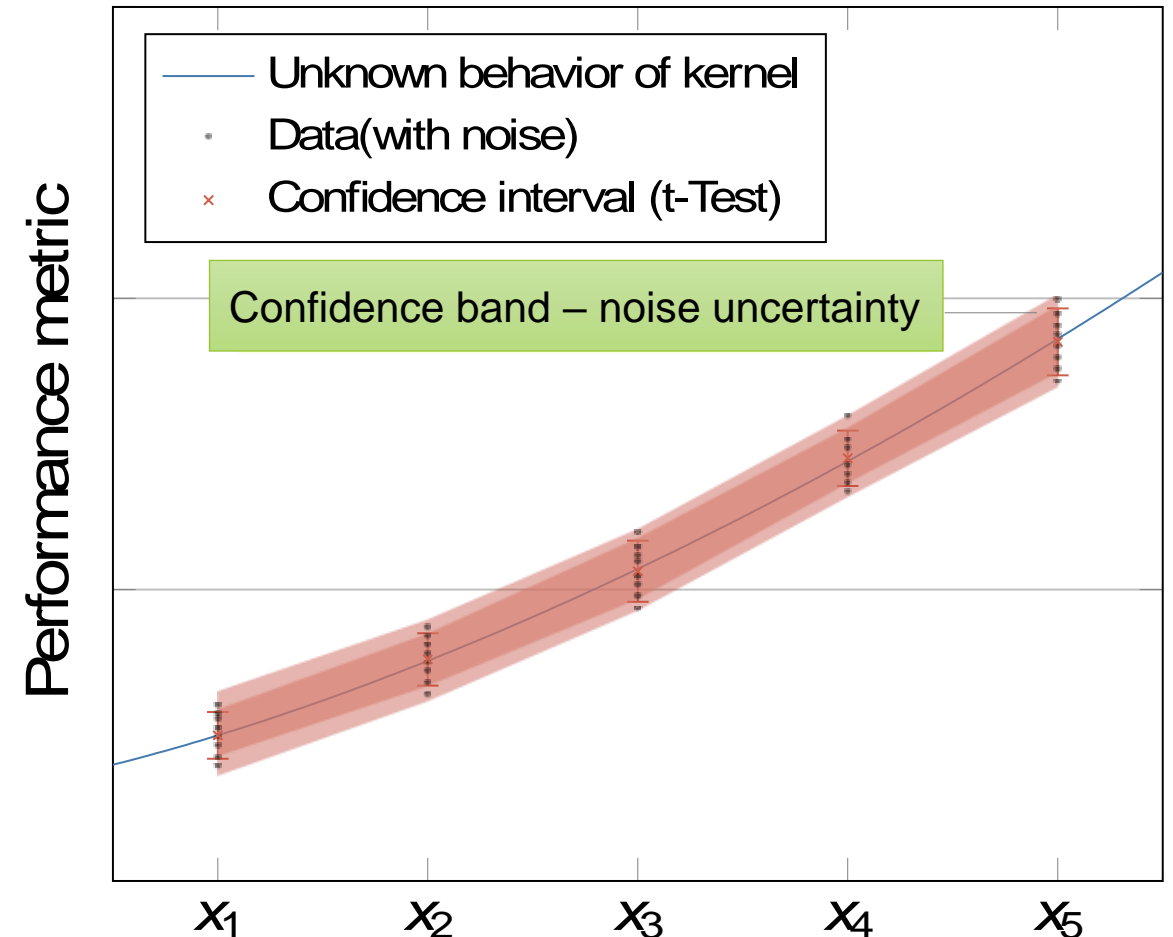


## Statistical quality control

- If the **confidence interval** is too wide, the fit will not be optimal, or overfitting might occur

$$CI = f(\text{mean}, \text{stddev})$$

- To improve CI – increase repetitions, include different configurations



## Adjusted $R^2$

---

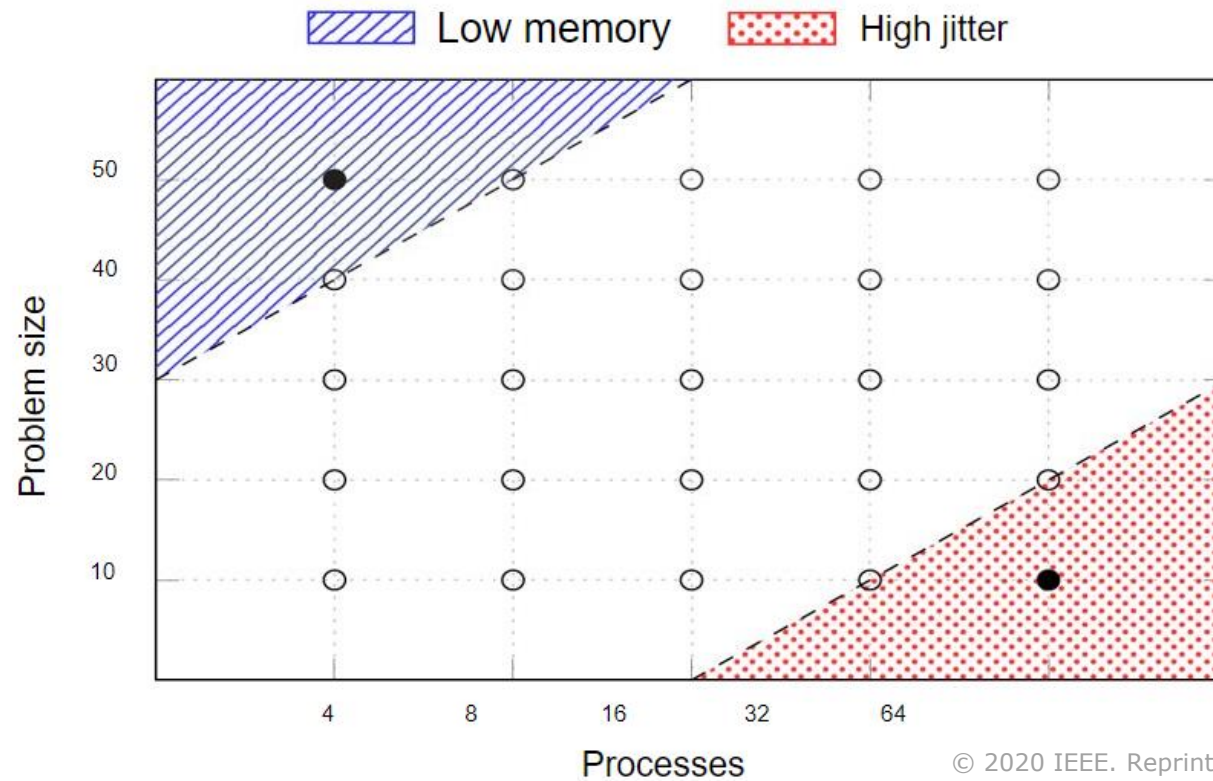
- $R^2$  represents how well the determined function fits the  $M$  available measurements
- Adjusted  $R^2$  adjusts for  $N$ , the number of terms used
  - Adj.  $R^2$  decreases  $\rightarrow$  more useless variables
  - Adj.  $R^2$  increases  $\rightarrow$  more useful variables
- Rule of thumb:  $\text{adj. } R^2 > 0.95$

$$R^2 = 1 - \frac{\text{residualSumSquares}}{\text{totalSumSquares}}$$

$$\overline{R^2} = 1 - (1 - R^2) \cdot \frac{M - 1}{M - N - 2}$$

# Sparse Modeling

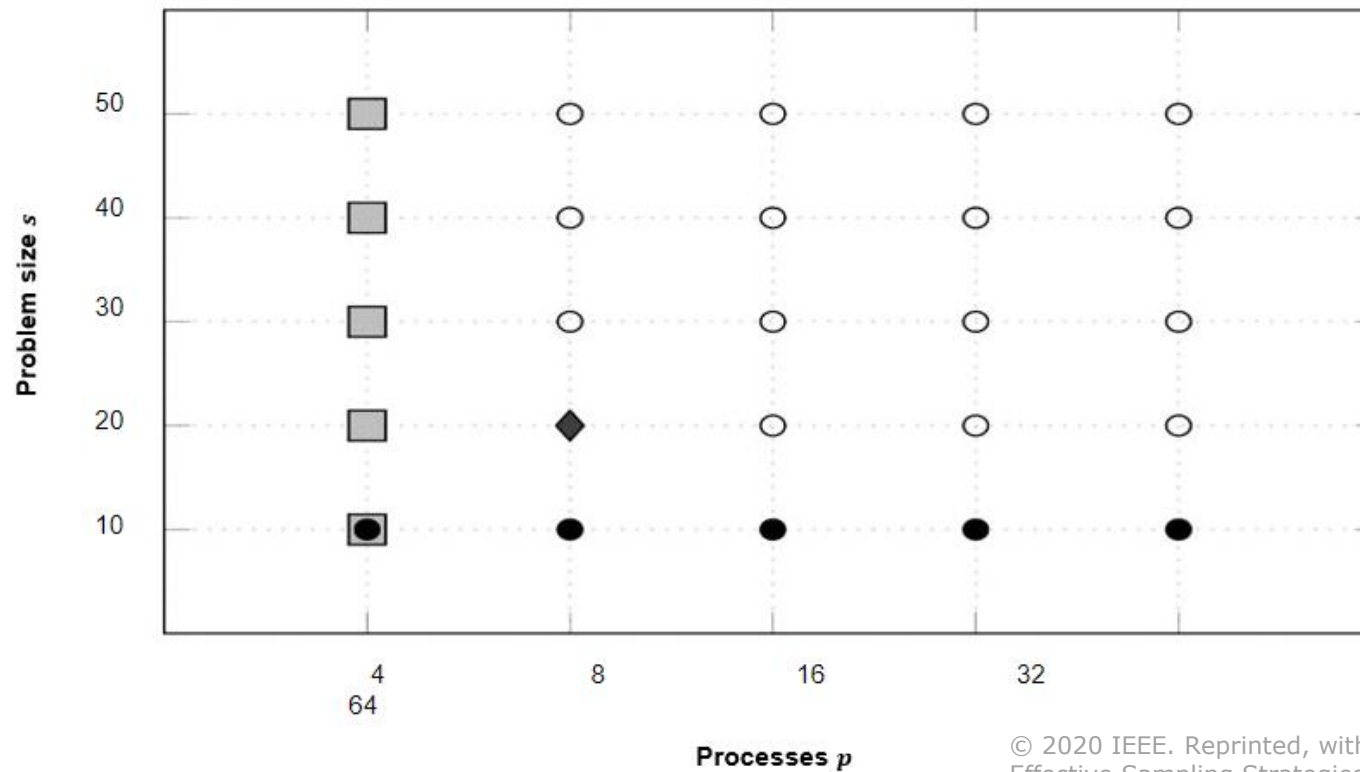
- Experiments can be expensive
- So far we needed  $5^{m+1}$  experiments,  $m$ =number of parameters



© 2020 IEEE. Reprinted, with permission, from M. Ritter et al. "Learning Cost-Effective Sampling Strategies for Empirical Performance Modeling," IPDPS 2020.

## Extra-P 4.0: Sparse Modeling

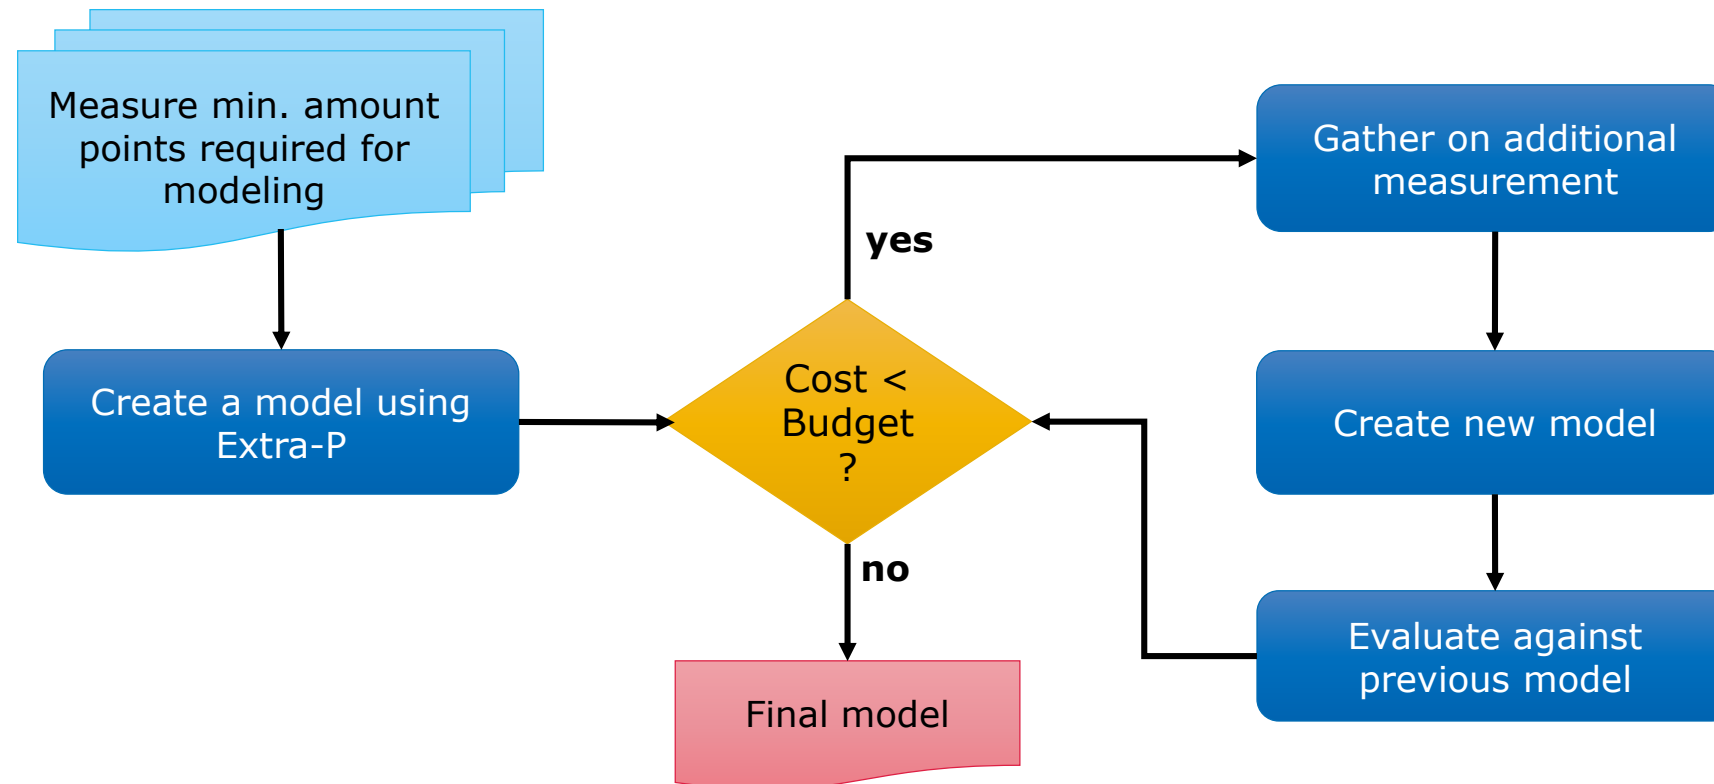
- Using our new sparse modeling approach we can model with less points!
- We only need  $5 \cdot m$  experiments,  $m$ =number of parameters



© 2020 IEEE. Reprinted, with permission, from M. Ritter et al. "Learning Cost-Effective Sampling Strategies for Empirical Performance Modeling," IPDPS 2020.

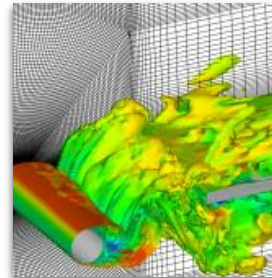
## Extra-P 4.0: Sparse Modeling

- Recommended experiment configuration strategy using our heuristic guideline



## Extra-P 4.0: Sparse Modeling

- Using sparse modeling we can reduce the average modeling cost by  $\sim 85\%$  (on synthetic data)
- We can retain  $\sim 92\%$  of the model accuracy (on synthetic data)
- Allows a more flexible experiment design



### FASTEST

- 70% decrease in cost
- $\sim 2\%$  prediction error

Image by  
Institute for  
Numerical  
Methods in  
Mechanical  
Engineering,  
TU Darmstadt

### Kripke

- 99% decrease in cost
- $\sim 39\%$  prediction error

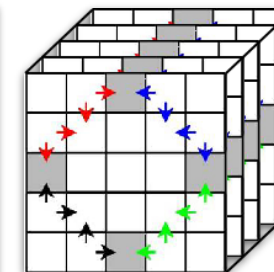
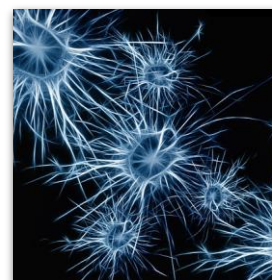
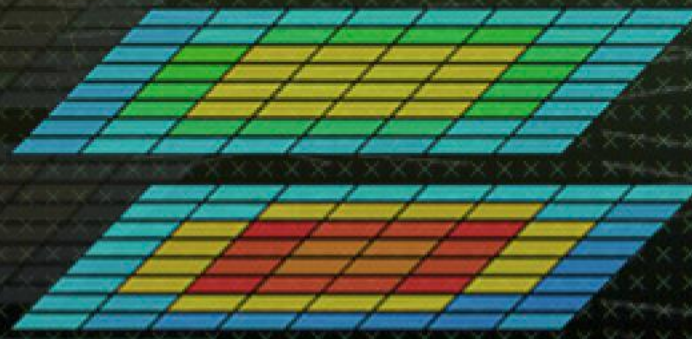


Image by  
Lawrence  
Livermore  
National  
Laboratory  
(CC BY-  
NC-SA 4.0)



### Relearn

- 85% decrease in cost
- $\sim 11\%$  prediction error



## Using Extra-P



## Installing Extra-P

---

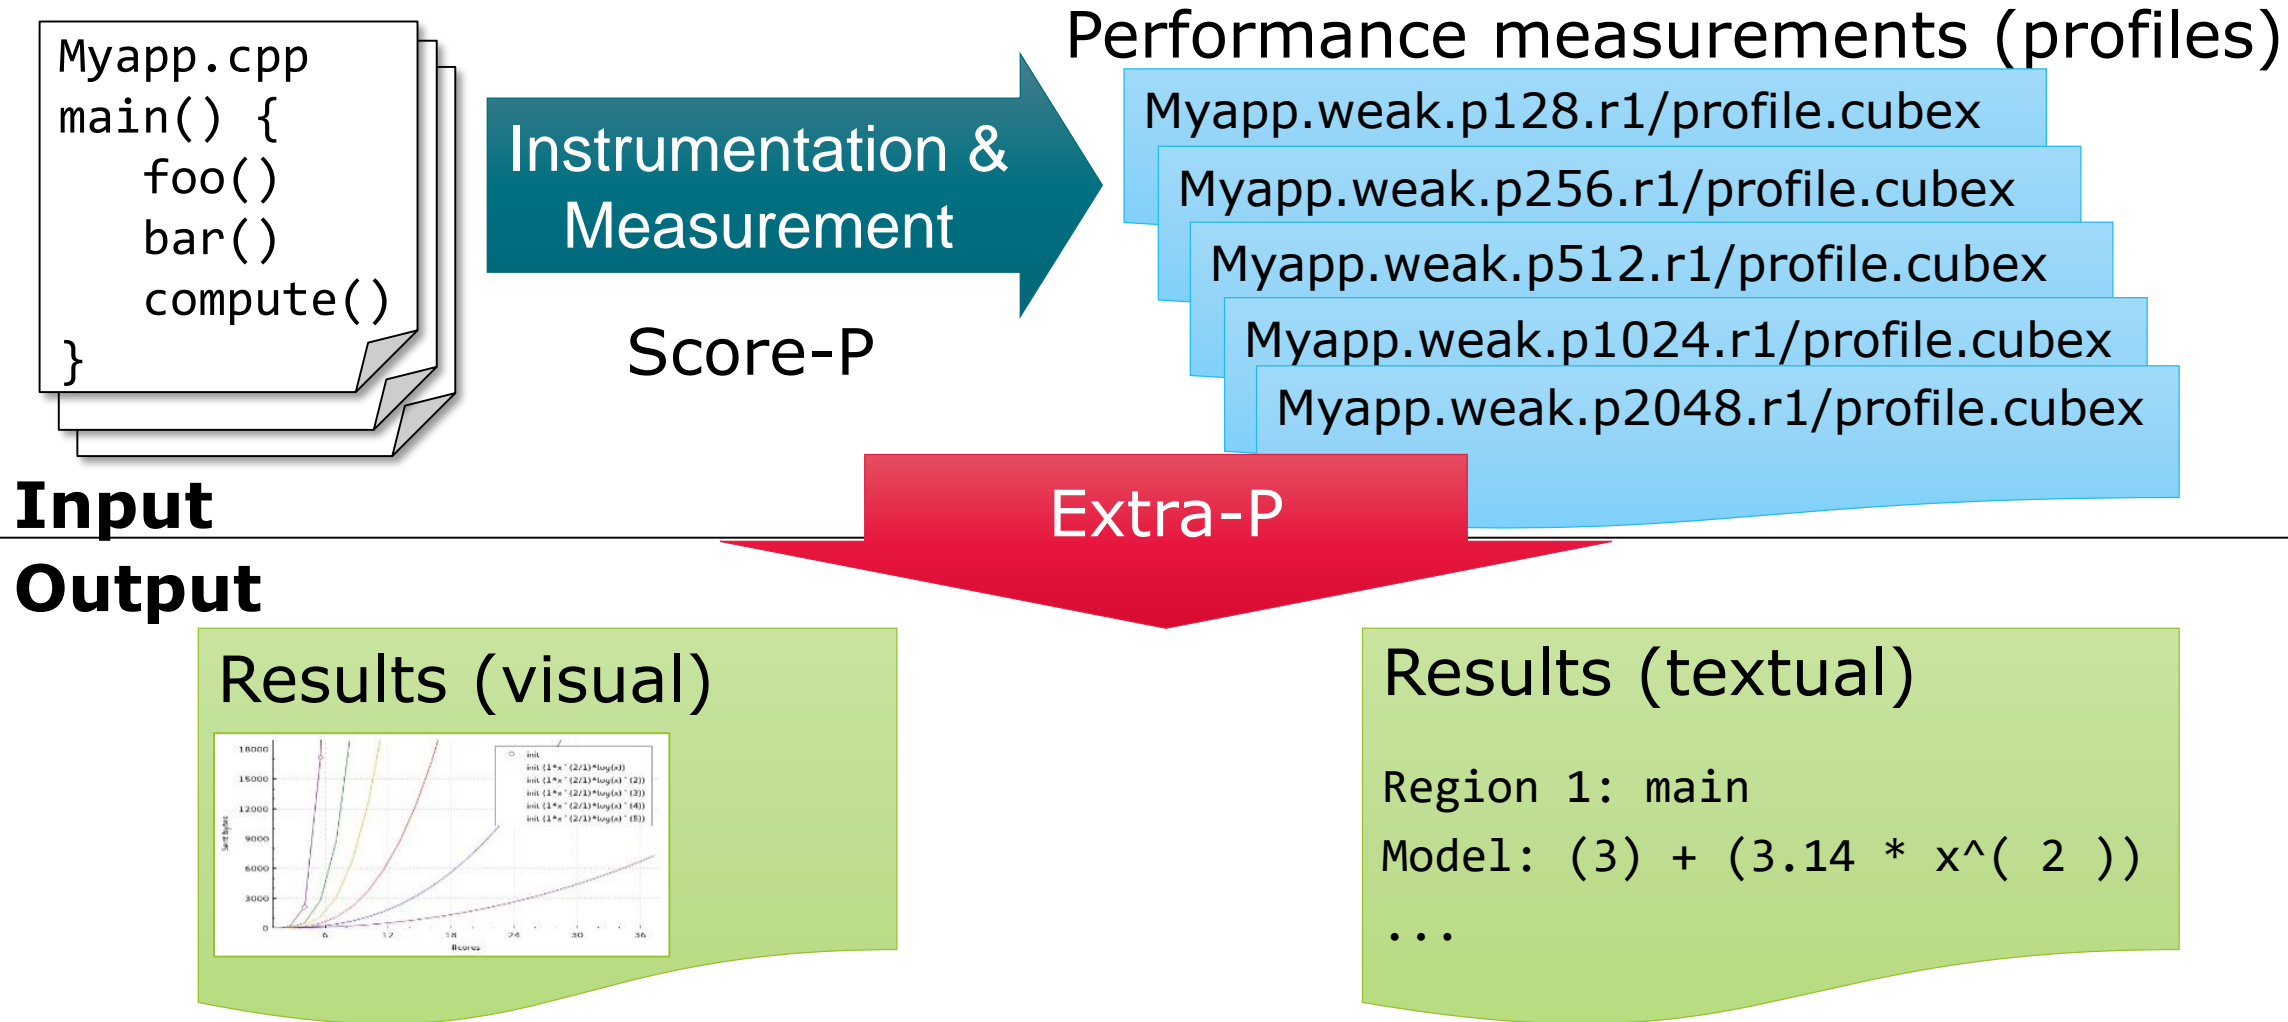
- Easy to install via pip
- Just run: `python -m pip install extrap --upgrade`
  - The `--upgrade` forces the installation of a new version
- All dependencies (packages) will be installed automatically

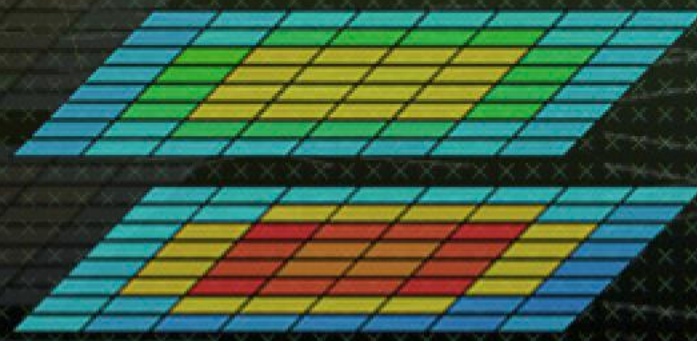
## Extra-P in the tuning workshop

---

- Available at: <https://github.com/extra-p/extrap>
- When installed on the system simply run:
  - `extrap` – for the command line version
  - `extrap-gui` – for the graphical user interface version
- The GUI version is not intended to be used on the cluster

# Automatic performance modeling with Extra-P





## Modeling sets of Cube experiments

## Extra-P Cube input description

---

Modeling tool expects Cube files in the following format:

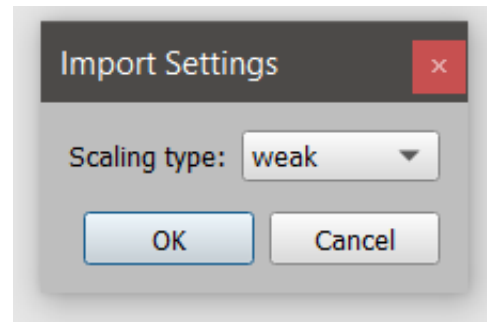
- `<DIR>/<PREFIX>.<PARAMETERS>.r<REPETITION>/<FILENAME>.cubex`
  - DIR, PREFIX, FILENAME – are just names, no meaning for Extra-P
  - REPETITION – number of the repetition of the experiments with same parameter values
- `<PARAMETERS>:=<PARAM1><VALUE1>...<PARAMn><VALUEn>`
  - List of parameter-value-pairs separated by "."
  - PARAM – varied parameter e.g. number of processes
  - VALUE – value of the varied parameter

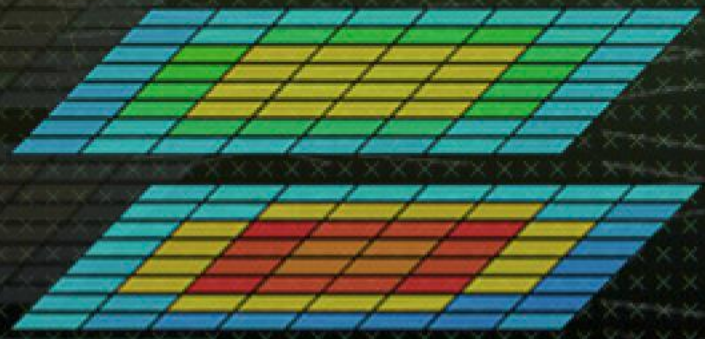
# Extra-P Cube input description

---

## Open set of CUBE files

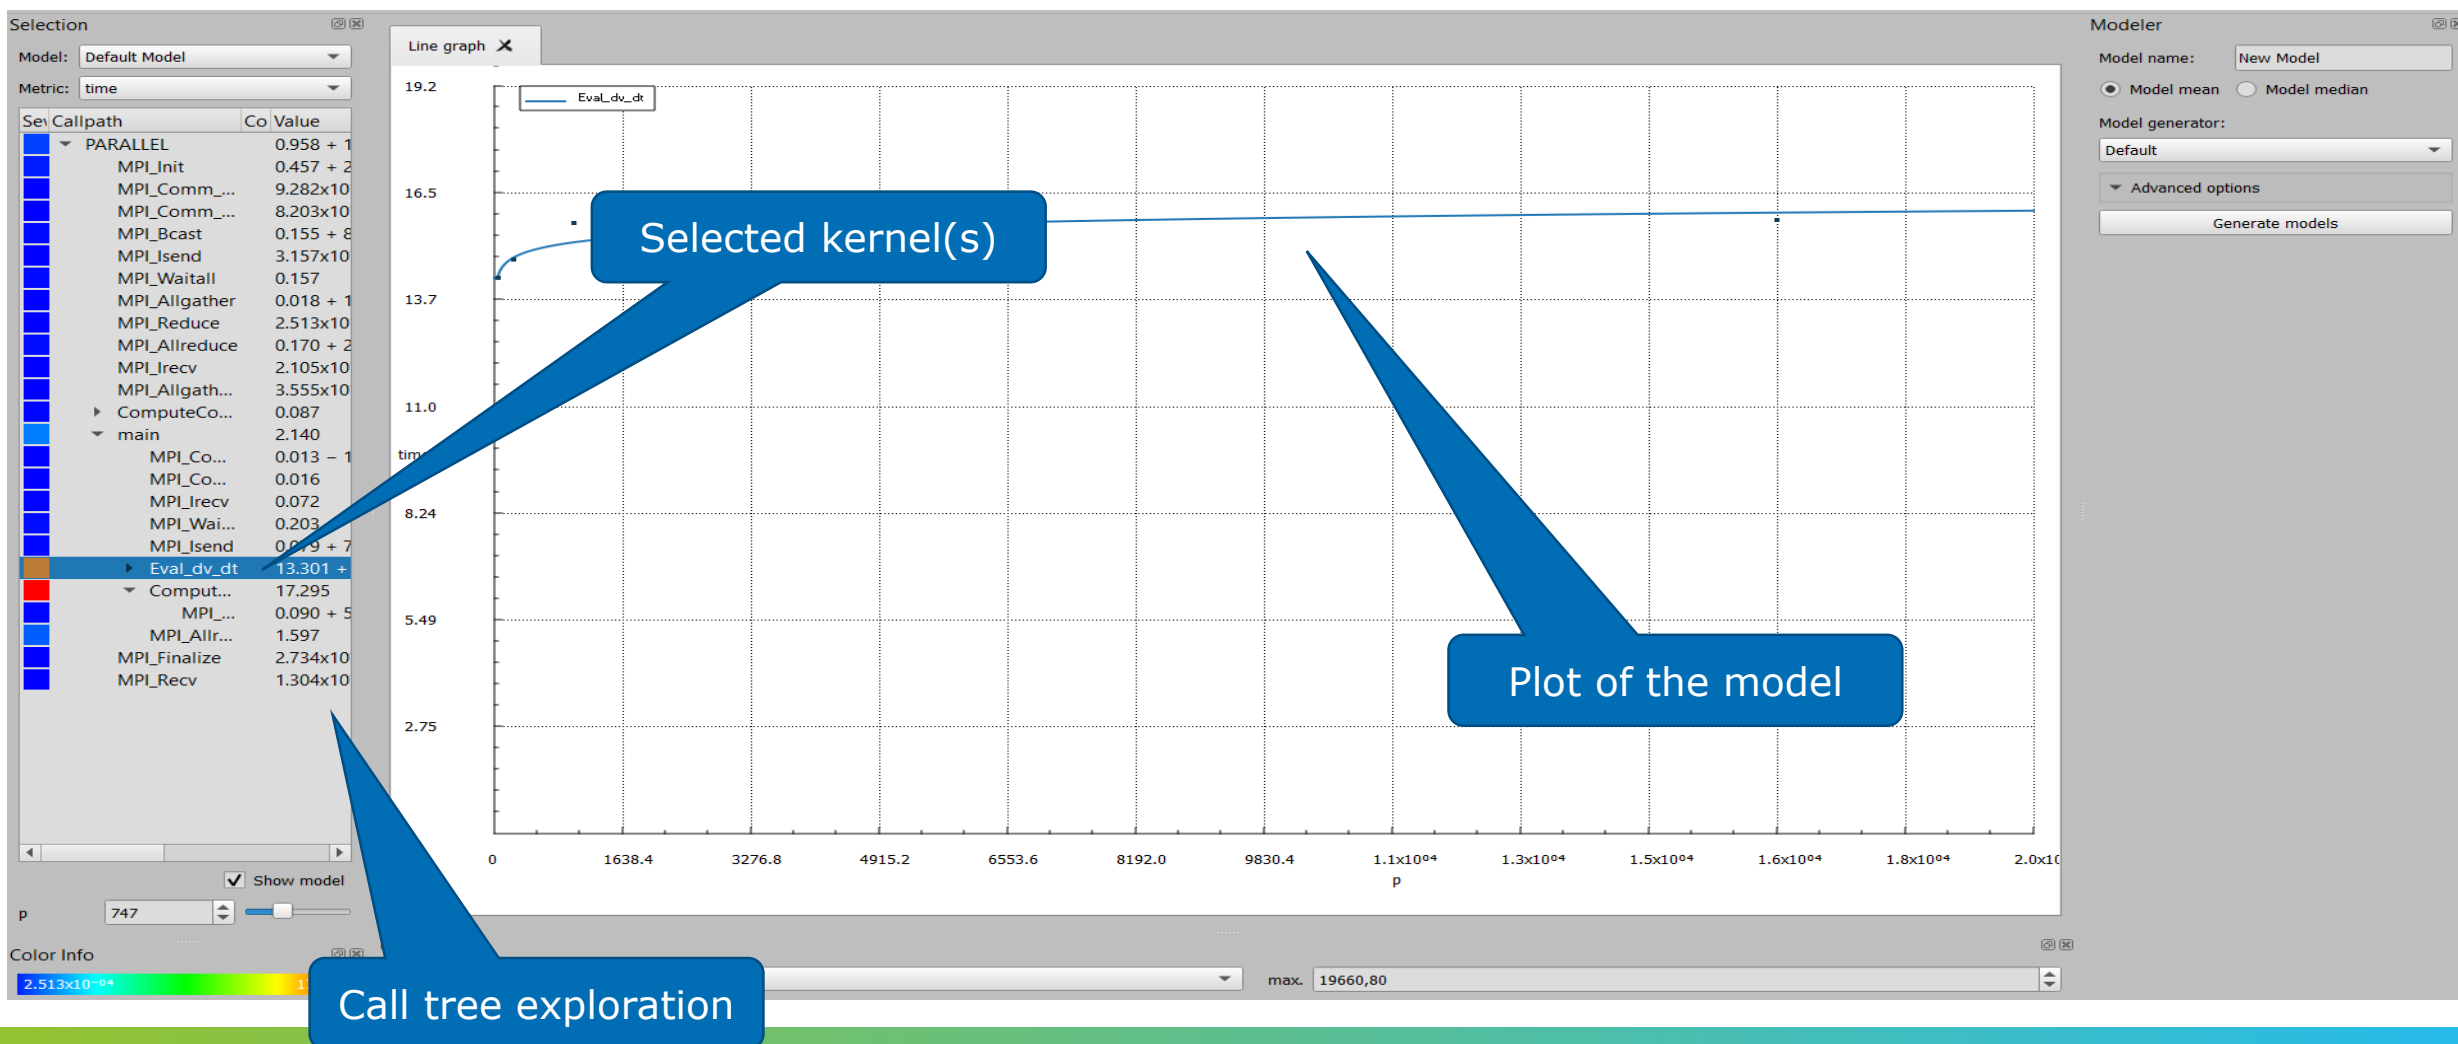
Open experiment	Ctrl+O
Save experiment	Ctrl+S
Open text input	
Screenshot	
Exit	Ctrl+Q





## Visualization with Extra-P

# Extra-P user interface





# Extra-P call tree view

Metric selection

Model selection

Call tree exploration

Model

Quality of fit metrics:  
Residual sum of squares  
and Adjusted R<sup>2</sup>

Impact of each kernel on  
the metric at the  
selected process count  
compared to the other  
kernels

Asymptotic behavior

Model: Default Model  
Metric: time

Seq Call	Comn	Value	RSS	Adj. R <sup>2</sup>	SMAPE	RE
PARALLEL		$0.958 + 1.837 \times 10^{-05} * p * \log_2(p)$	0.092	0.990	3.629	0
MPI_Init		$0.457 + 2.393 \times 10^{-10} * p^{9/4}$	$4.871 \times 10^{-04}$	0.998	1.036	0
MPI_Comm_size		$9.282 \times 10^{-04} + 1.232 \times 10^{-09} * p * \log_2(p)$	$8.406 \times 10^{-10}$	0.981	1.370	0
MPI_Comm_rank		$8.203 \times 10^{-04} + 1.796 \times 10^{-10} * p * \log_2(p)$	$1.241 \times 10^{-11}$	0.987	0.194	0
MPI_Bcast		$0.155 + 8.580 \times 10^{-11} * p^{7/3}$	$6.011 \times 10^{-04}$	0.997	4.769	0
MPI_Isend		$3.157 \times 10^{-03} + 3.410 \times 10^{-05} * \log_2(p)$	$1.641 \times 10^{-08}$	0.568	1.707	0
MPI_Waitall		0.157	$1.788 \times 10^{-04}$	1	2.934	0
MPI_Allgather		$0.018 + 1.218 \times 10^{-09} * p^{7/4}$	$9.760 \times 10^{-06}$	0.980	7.293	0
MPI_Reduce		$2.513 \times 10^{-04}$	$1.119 \times 10^{-08}$	1	13.626	0
MPI_Allreduce		$0.170 + 2.478 \times 10^{-04} * p^{1/3} * \log_2(p)$	$7.463 \times 10^{-04}$	0.790	6.737	0
MPI_Irecv		$2.105 \times 10^{-03} + 5.098 \times 10^{-05} * \log_2(p)$	$2.552 \times 10^{-08}$	0.693	3.059	0
MPI_Allgatherv		$3.555 \times 10^{-04} + 2.418 \times 10^{-07} * p^{5/4}$	$2.974 \times 10^{-06}$	0.997	8.098	0
ComputeCornerForces		0.087	$1.425 \times 10^{-06}$	1	0.557	0
main		2.140	$4.007 \times 10^{-03}$	1	1.264	0
MPI_Comm_rank		$0.013 - 1.971 \times 10^{-05} * \log_2(p)$	$9.469 \times 10^{-09}$	0.319	0.350	0
MPI_Comm_size		0.016	$1.771 \times 10^{-08}$	1	0.341	0
MPI_Irecv		0.072	$9.185 \times 10^{-06}$	1	1.582	0
MPI_Waitall		0.203	0.019	1	22.235	0
MPI_Isend		$0.079 + 7.508 \times 10^{-04} * \log_2(p)$	$5.459 \times 10^{-06}$	0.695	1.291	0
Eval_dv_dt		$13.301 + 0.190 * \log_2(p)$	0.803	0.392	2.763	0
ComputeCornerForces		17.295	0.038	1	0.340	0
MPI_Reduce		$0.090 + 5.990 \times 10^{-09} * p^{5/4} * \log_2(p)$	$4.877 \times 10^{-05}$	0.684	3.088	0
MPI_Allreduce		1.597	0.051	1	4.829	0
MPI_Finalize		$2.734 \times 10^{-04}$	$5.817 \times 10^{-08}$	1	28.617	0
MPI_Recv		$1.304 \times 10^{-04} + 4.542 \times 10^{-07} * p$	$5.099 \times 10^{-06}$	0.829	48.915	0

p 747

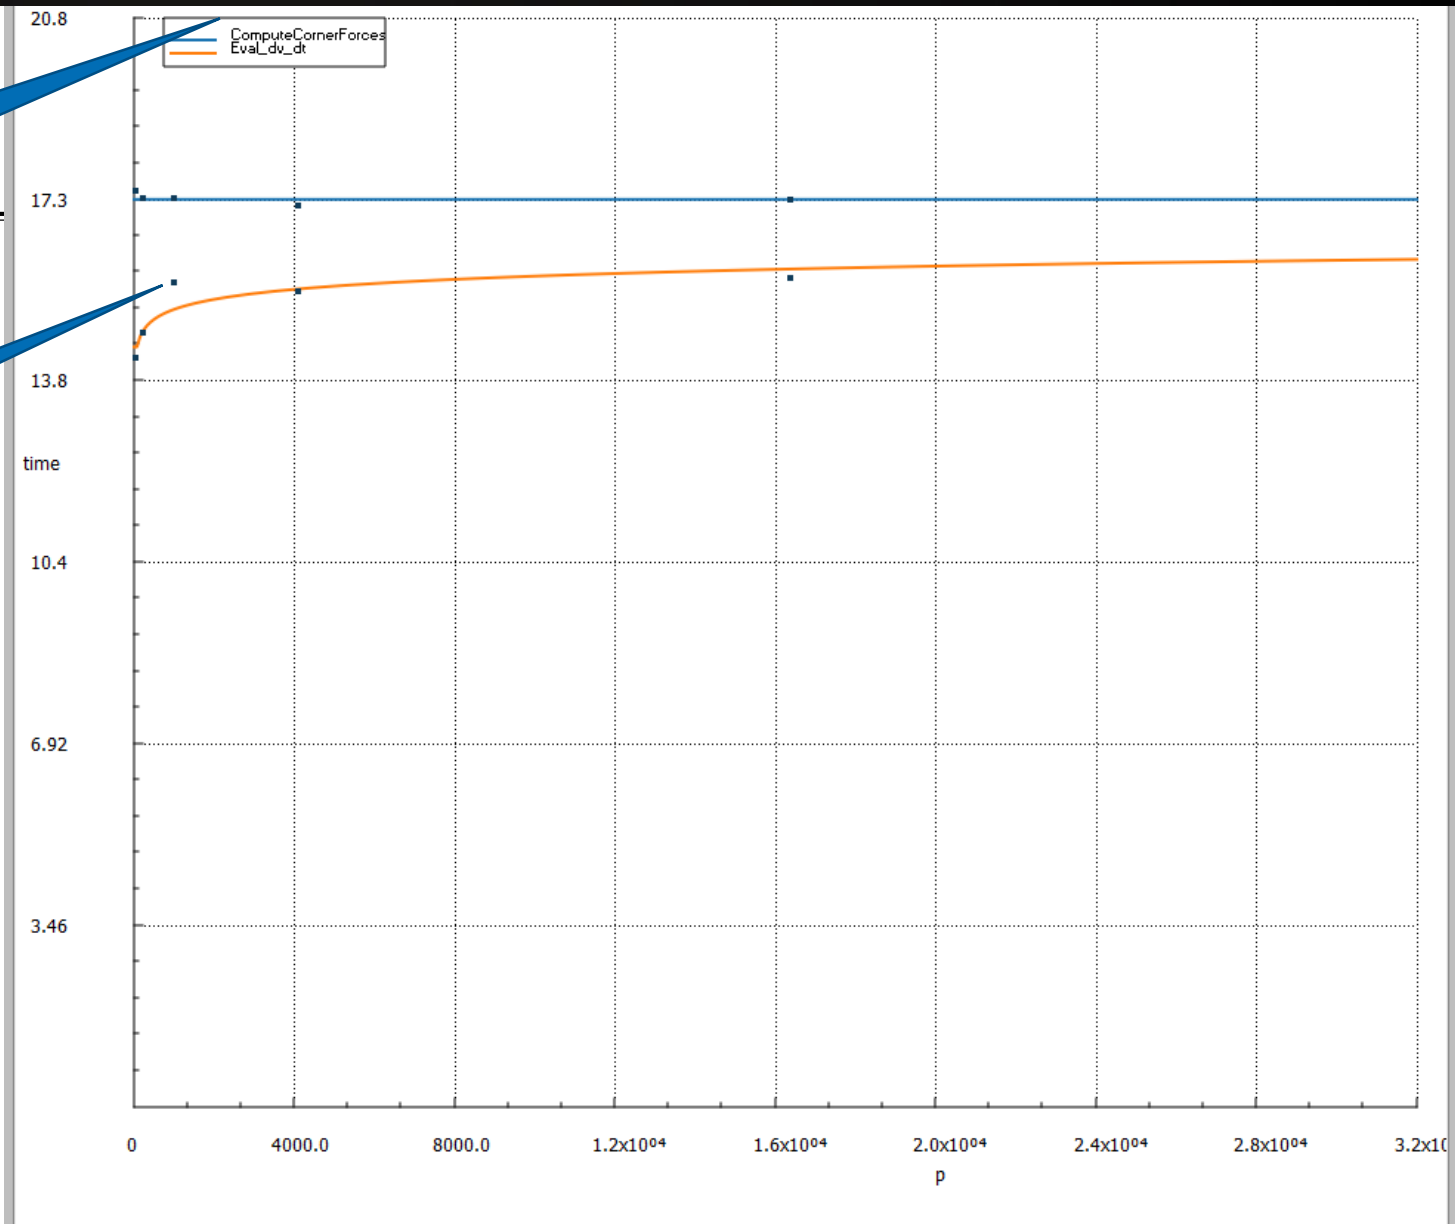
Show model

## Extra-P model view

Models selected in the Call path view

Measurement values

X axis scale control for prediction of behavior at other process counts



Graph Limits

X-axis

p

max.

32000,00

34

# Modeling measurements from a text file

## Choose input file

---

Open set of CUBE files

Open experiment      Ctrl+O

Save experiment      Ctrl+S

Open text input

Screenshot

Exit      Ctrl+Q



Select input file in the  
GUI

## Extra-P input in text form

---

- Useful when no CUBE files are available or when modeling a small data set

```
PARAMETER p  
POINTS 1000 2000 4000 8000 16000  
METRIC metric1  
REGION region1  
DATA 1 1 1 1 1  
DATA 4 4 4 3.99 4.01  
DATA 16 15.999 16.01 16.01 15.99  
DATA 64 64 64 64.01 63.99  
DATA 256.01 255.99 256 256
```

### Parameter name

This name will be used in the GUI as well as in the textual output

## Extra-P input in text form

---

```
PARAMETER p
POINTS 1000 2000 4000 8000 16000
METRIC metric1
REGION region1
DATA 1 1 1 1 1
DATA 4 4 4 3.99 4.01
DATA 16 15.999 16.01 16.01 15.99
DATA 64 64 64 64.01 63.99
DATA 256.01 255.99 256 256
```



**Measurement points**  
Use at least 5, preferably 6,  
but in general the more the better

## Extra-P input in text form

---

```
PARAMETER p
POINTS 1000 2000 4000 8000 16000
METRIC metric1
REGION region1
DATA 1 1 1 1 1
DATA 4 4 4 3.99 4.01
DATA 16 15.999 16.01 16.01 15.99
DATA 64 64 64 64.01 63.99
DATA 256.01 255.99 256 256
```

Metric name

Region name

Both used to determine the output  
Cube file hierarchical structure and  
identify separate data sets

## Extra-P input in text form

---

```
PARAMETER p
POINTS 1000 2000 4000 8000 16000
METRIC metric1
REGION region1
DATA 1 1 1 1 1
DATA 4 4 4 3.99 4.01
DATA 16 15.999 16.01 16.01 15.99
DATA 64 64 64 64.01 63.99
DATA 256.01 255.99 256 256
```

### Data points

Each row corresponds to a point; all values in a row are considered repeat measurements of the same experiment



## Using the command line tool

## Extra-P command line tool

---

- Provides the same functionality, without visualization for use on cluster
- Usage guideline and command can be found at: <https://github.com/extra-p/extrap>
- 1.) Run: `extrap`
- Command Format: `extrap OPTIONS (--cube | --text | --talpas | --json | --extra-p-3) FILEPATH`
- 2.) Select input type: `extrap --text /lrz/sys/courses/vihps/material/extrap_data/input.txt`

## Extra-P command line tool

### 3.) Output:

Callpath: compute ← Callpath, kernel of the application that was measured

Metric: time ← Metric name; either Score-P metrics (time, bytes, etc.) or custom metrics

Measurement point: (2.00E+01)	Mean: 8.19E+01	Median: 8.20E+01	} ← Measurements for each input element (e.g., #processes)
Measurement point: (3.00E+01)	Mean: 1.79E+02	Median: 1.78E+02	
Measurement point: (4.00E+01)	Mean: 3.19E+02	Median: 3.19E+02	
Measurement point: (5.00E+01)	Mean: 5.05E+02	Median: 5.06E+02	
Measurement point: (6.00E+01)	Mean: 7.25E+02	Median: 7.26E+02	

Model:  $-0.8897934098062804 + 0.20168243826499183 * x^{(2)}$  ← Best-fit model

RSS: 3.43E+01 ← RSS: Residual sum of squares

Adjusted R<sup>2</sup>: 1.00E+00 ← Adjusted R<sup>2</sup> (explained previously)

## Hands-on exercises

## Extra-P exercises

---

- Run: `extrap`
- Example data: `tests/data`
- Open the examples in the GUI
- Use the command line tool
- Open the text based, JSON input example
- Produce textual output and inspect it

- What additional features would you like to see?
- Did you find any bugs?

You can contact us via email: [extra-p-support@lists.parallel.informatik.tu-darmstadt.de](mailto:extra-p-support@lists.parallel.informatik.tu-darmstadt.de)

Or on GitHub using the issues tool: <https://github.com/extra-p/extrap>