

Extra-P:

Insightful Automatic Performance Modeling

Alexandru Calotoiu¹, Torsten Hoefler², Martin Schulz³
Sergei Shudler¹ and Felix Wolf¹

¹ TU Darmstadt , ² ETH Zürich , ³ TU München



TECHNISCHE
UNIVERSITÄT
DARMSTADT

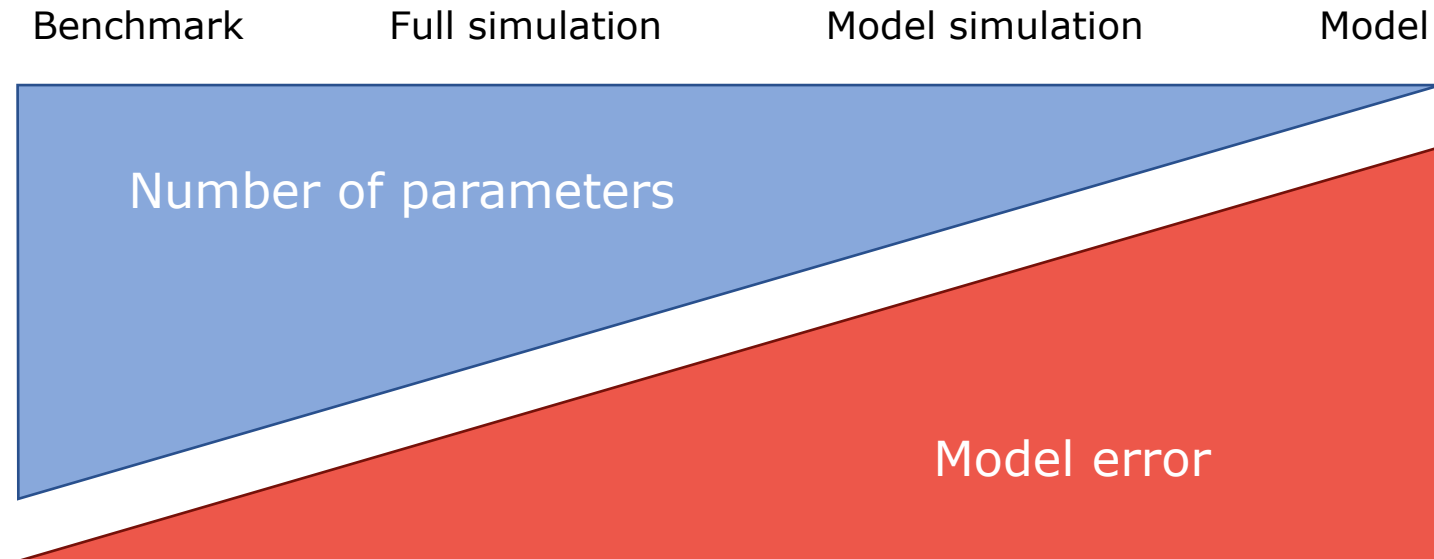
ETH zürich

 Lawrence Livermore
National Laboratory

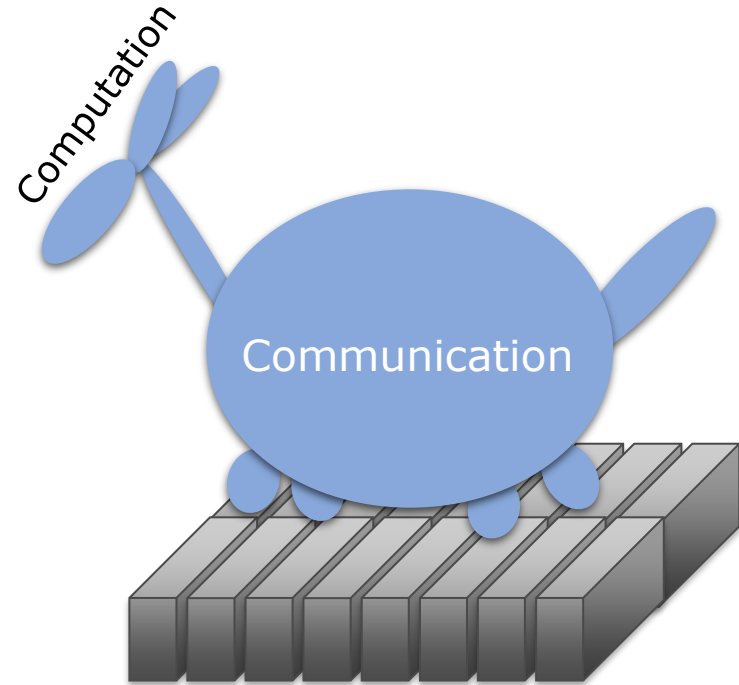
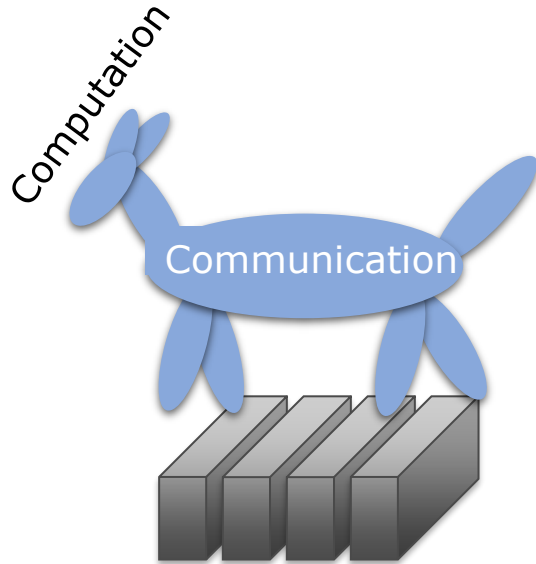


Technische Universität München


Spectrum of performance analysis methods



Motivation - latent scalability bugs

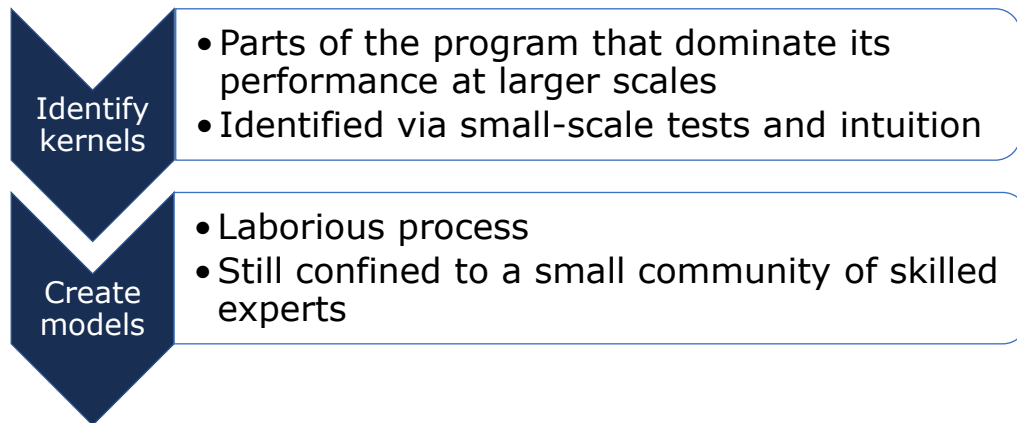


Scaling model



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

Analytical performance modeling



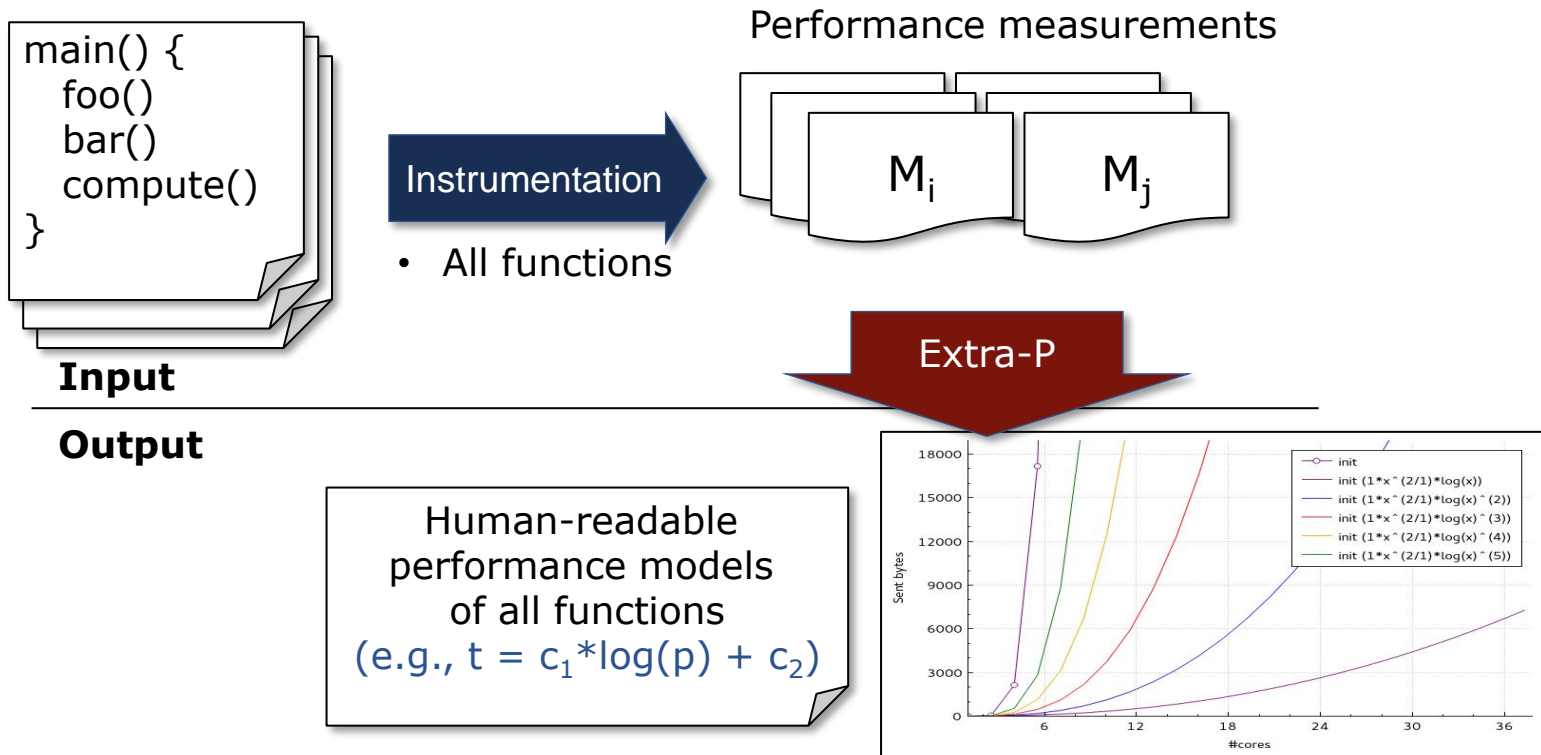
Disadvantages:

- Time consuming
- Danger of overlooking unscalable code

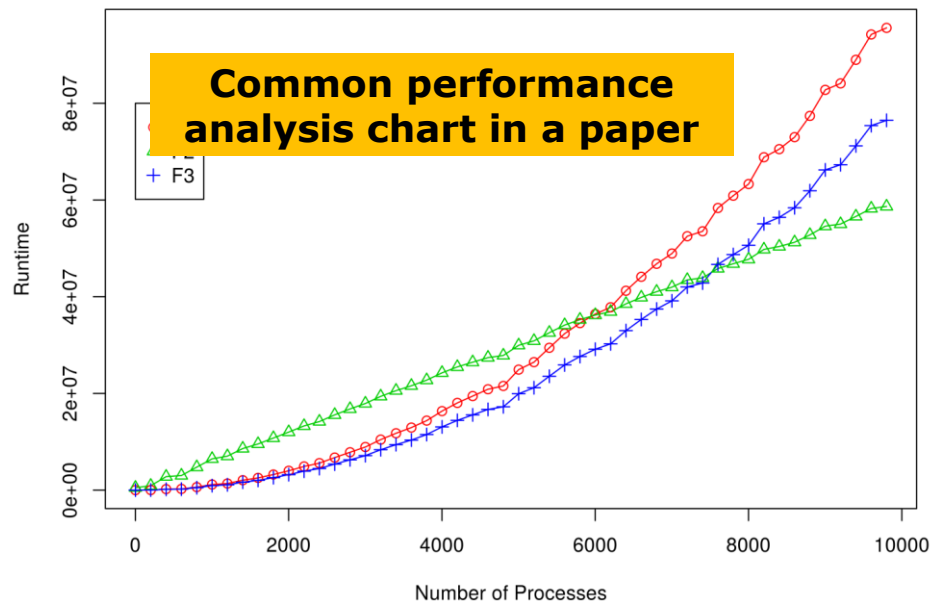
Hoisie et al.: *Performance and scalability analysis of teraflop-scale parallel architectures using multi-dimensional wavefront applications*. International Journal of High Performance Computing Applications, 2000

Bauer et al.: *Analysis of the MILC Lattice QCD Application su3_rmd*. CCGrid, 2012

Automatic performance modeling



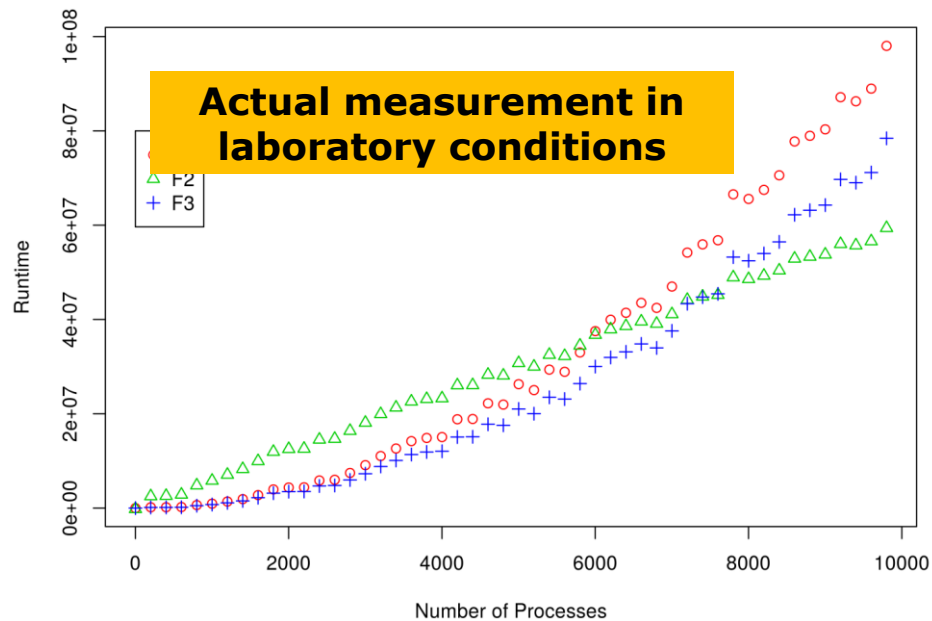
Primary focus on scaling trend



Ranking

1. F_2
2. F_1
3. F_3


Primary focus on scaling trend



Ranking

1. F_2
2. F_1
3. F_3

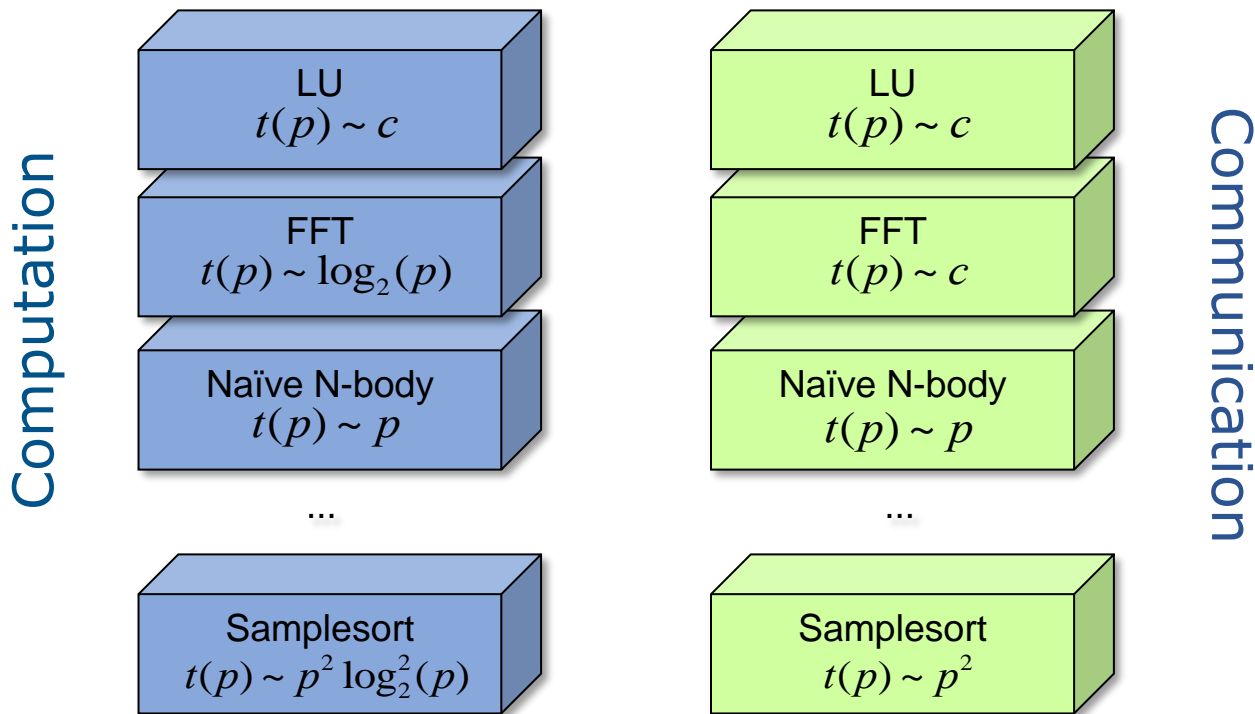
Primary focus on scaling trend



Ranking

1. F_2
2. F_1
3. F_3

Model building blocks



Performance model normal form

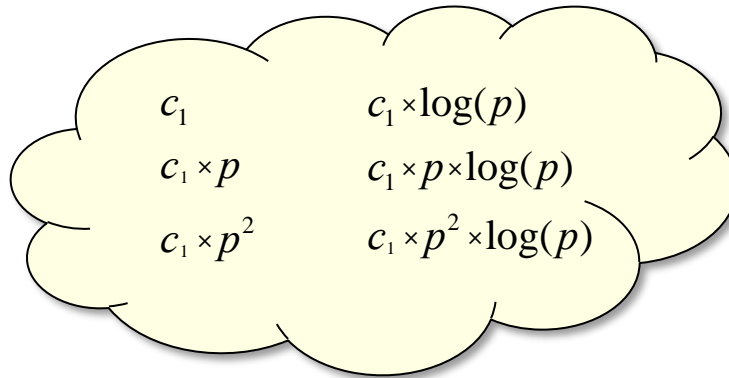
$$f(p) = \prod_{k=1}^n c_k \times p^{i_k} \times \log_2^{j_k}(p)$$

$$\begin{array}{l} n \in \mathbb{N} \\ i_k \in I \\ j_k \in J \\ I, J \in \mathbb{Q} \end{array}$$

$$n = 1$$

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

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



Performance model normal form

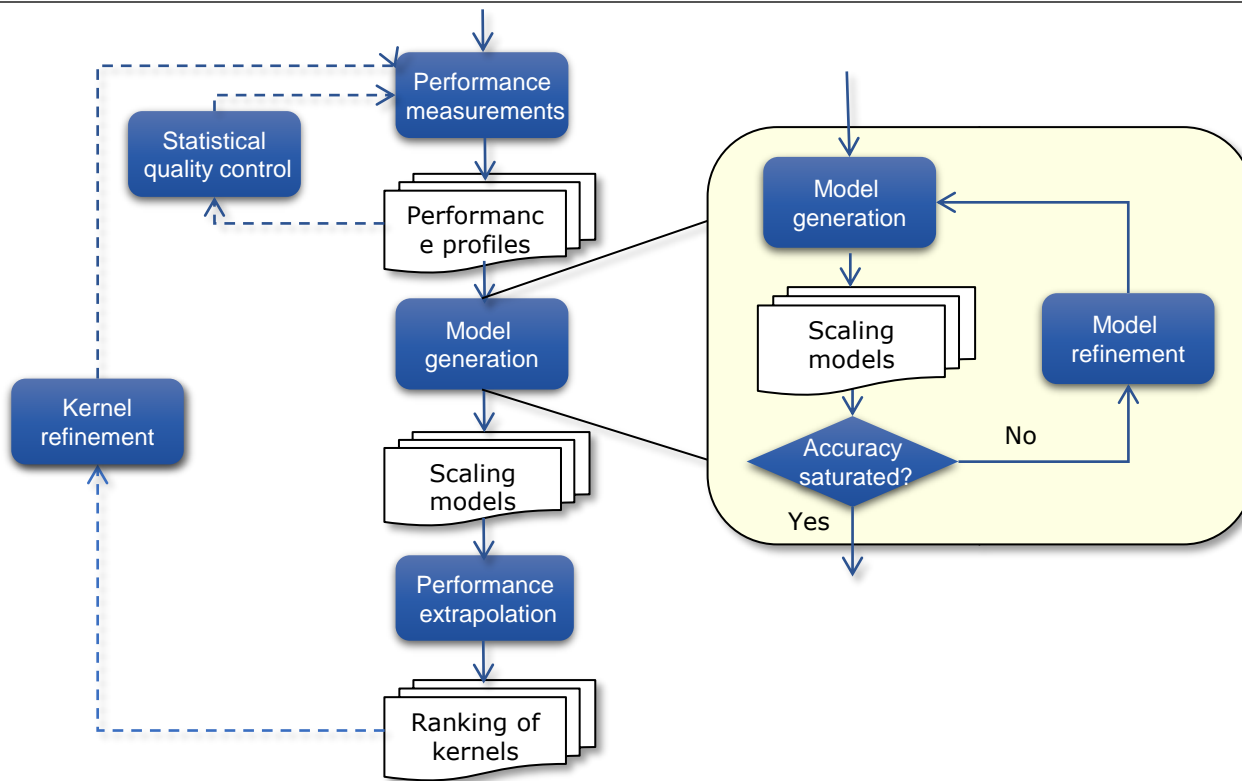
$$f(p) = \sum_{n=0}^{\infty} \sum_{I \in \{0,1,2\}^n} \sum_{J \in \{0,1\}^n} c_{I,J} \cdot p^{I_0} (\log(p))^{I_1} p^{I_2}$$

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

$c_1 + c_2 \times p$
 $c_1 + c_2 \times p^2$
 $c_1 + c_2 \times \log(p)$
 $c_1 + c_2 \times p \times \log(p)$
 $c_1 + c_2 \times p^2 \times \log(p)$


$c_1 \times \log(p) + c_2 \times p$
 $c_1 \times \log(p) + c_2 \times p \times \log(p)$
 $c_1 \times \log(p) + c_2 \times p^2$
 $c_1 \times \log(p) + c_2 \times p^2 \times \log(p)$
 $c_1 \times p + c_2 \times p \times \log(p)$
 $c_1 \times p + c_2 \times p^2$
 $c_1 \times p + c_2 \times p^2 \times \log(p)$
 $c_1 \times p \times \log(p) + c_2 \times p^2$
 $c_1 \times p \times \log(p) + c_2 \times p^2 \times \log(p)$
 $c_1 \times p^2 + c_2 \times p^2 \times \log(p)$

Workflow



Assumptions & limitations


- 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 (!)





Changing growth trends

- Ranking according to growth rate difficult:

$$\log^2(p) ? \sqrt{p}$$



Changing growth trends (2)



Ranking of kernels

- Kernels are ranked according the leading-order terms in the models
- Leading-order term \rightarrow big-O notation
- For example: $O(x)$ comes before $O(x^2)$

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 required



Performance measurements (profiles)



Performance measurements (3)

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

Performance measurements (profiles)

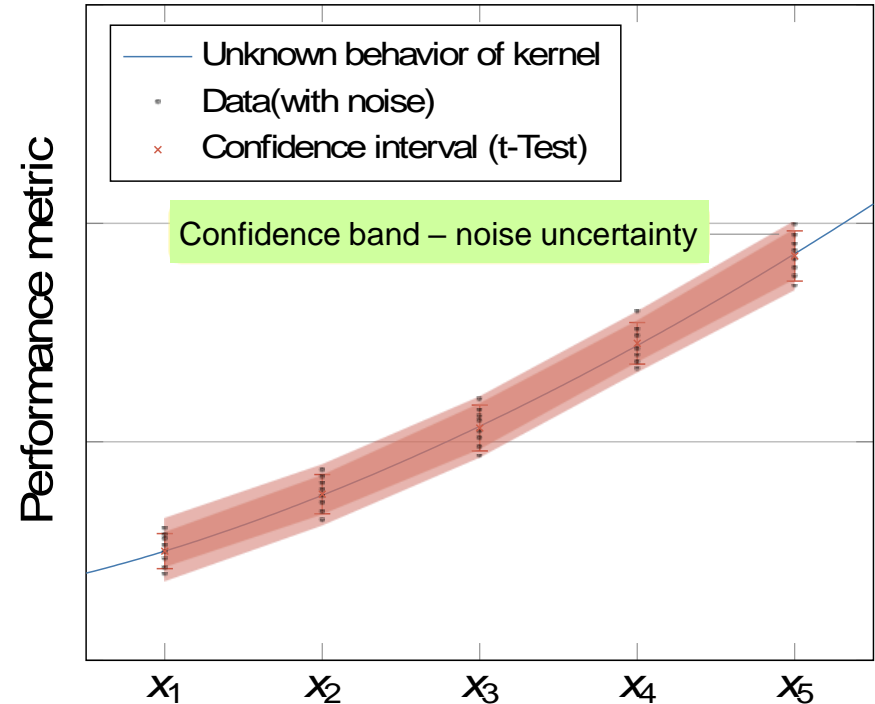


Statistical quality control

- 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



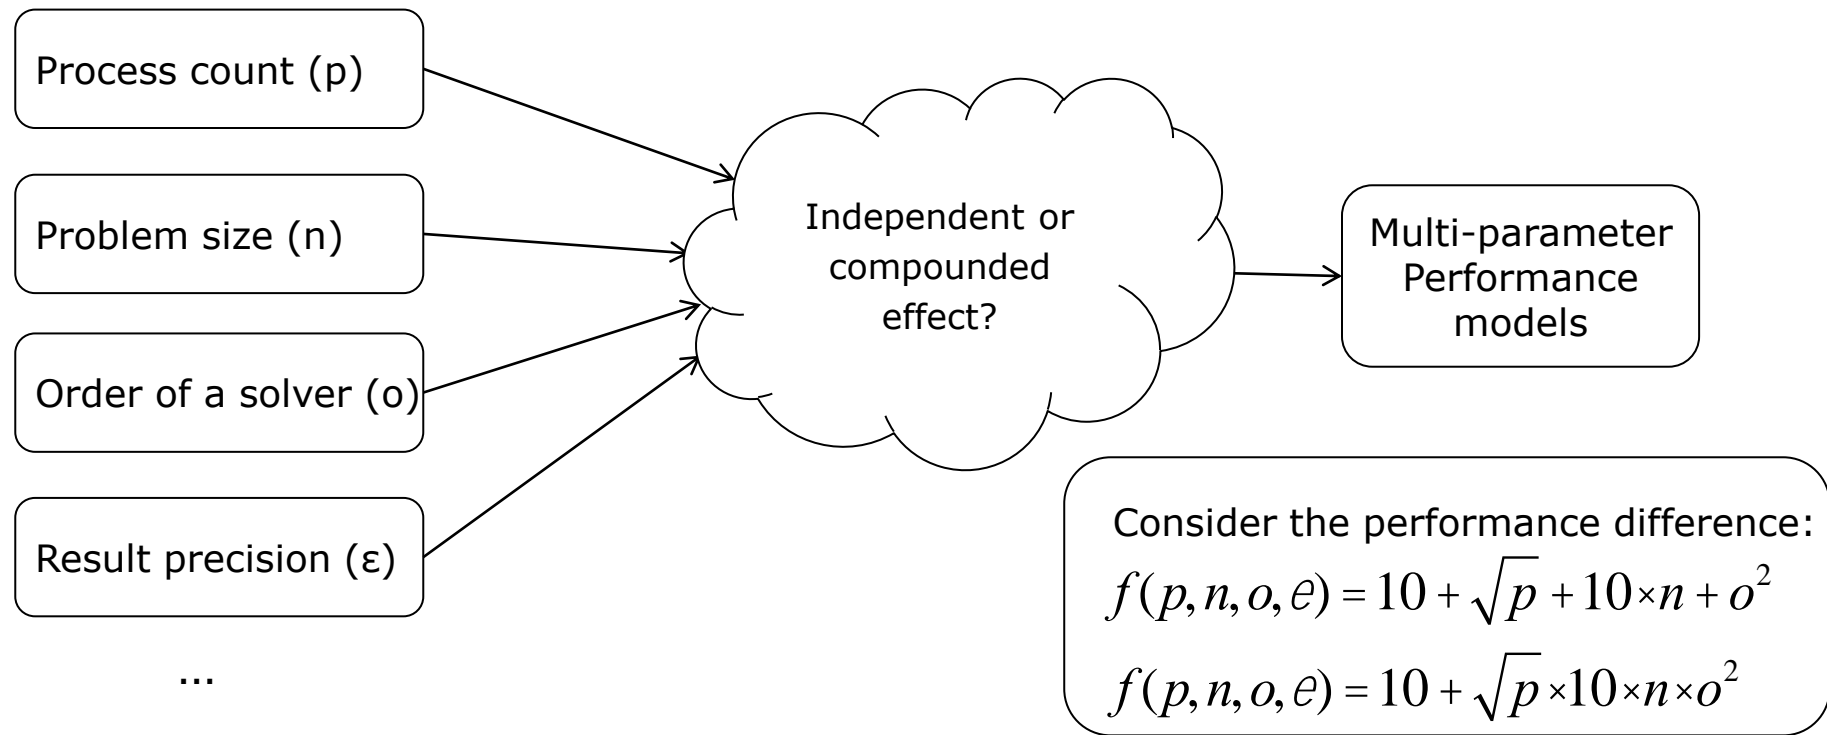
Adjusted R²

- R² represents how well the determined function fits the M available measurements
- Adjusted R² adjusts for N, the number of terms used
 - Adj. R² decreases → more useless variables
 - Adj. R² increases → more useful variables
- Rule of thumb: adj. R² > 0.95

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

$$\bar{R}^2 = 1 - (1 - R^2) \times \frac{M - 1}{M - N - 2}$$

Extra-P 3.0: Fast multi-parameter performance modeling



Extra-P 3.0: Fast multi-parameter performance modeling

Expanded performance model normal form

$$f(p) = \mathring{a} \prod_{k=1}^n c_k \prod_{l=1}^m p_l^{i_{kl}} \times \log_2^{j_{kl}}(p_l)$$

$$\begin{aligned} n &\hat{=} \mathbb{N} \\ m &\hat{=} \mathbb{N} \\ i_{kl} &\hat{=} I \\ j_{kl} &\hat{=} J \\ I, J &\hat{=} \mathbb{Q} \end{aligned}$$

$$\begin{aligned} n &= 2 \\ m &= 2 \\ I &= \left\{ \frac{0}{4}, \frac{1}{4}, \dots, \frac{12}{4} \right\} \\ J &= \{0, 1, 2\} \end{aligned}$$

Model candidates	Constant	
	• Single parameter	$c_1 + c_2 \times p_1$
	• Multiple parameters ...	
	• Additive	$c_1 + c_2 \times p_1 + c_3 \times p_2$
	• Multiplicative	$c_1 + c_2 \times p_1 \times p_2$
	• Complex	$c_1 + c_2 \times p_1 + c_3 \times p_1^2 \times p_2 \times \log_2(p_2)$

Extra-P 3.0: Fast multi-parameter performance modeling

- **Hierarchical search** – Reduces search from all possible models to combinations of the best single parameter models of each parameter
- **Modified golden section search** – Orders single parameter search space and applies a modified binary search
- Assuming 300.000 models searched per second* – 3 parameter models

Exhaustive search	$C(59.319, 3) = 34.786.300.841.019$	~3.5 years/model
Hierarchical search	$3 \times 9.139 + 512 = 27.929$	~11 models/second
Hierarchical search + Modified golden section search	$3 \times 26 + 512 = 590$	~508 models/second

*This is a simplification, multi-parameter models take much longer to be evaluated

Using Extra-P 2.0

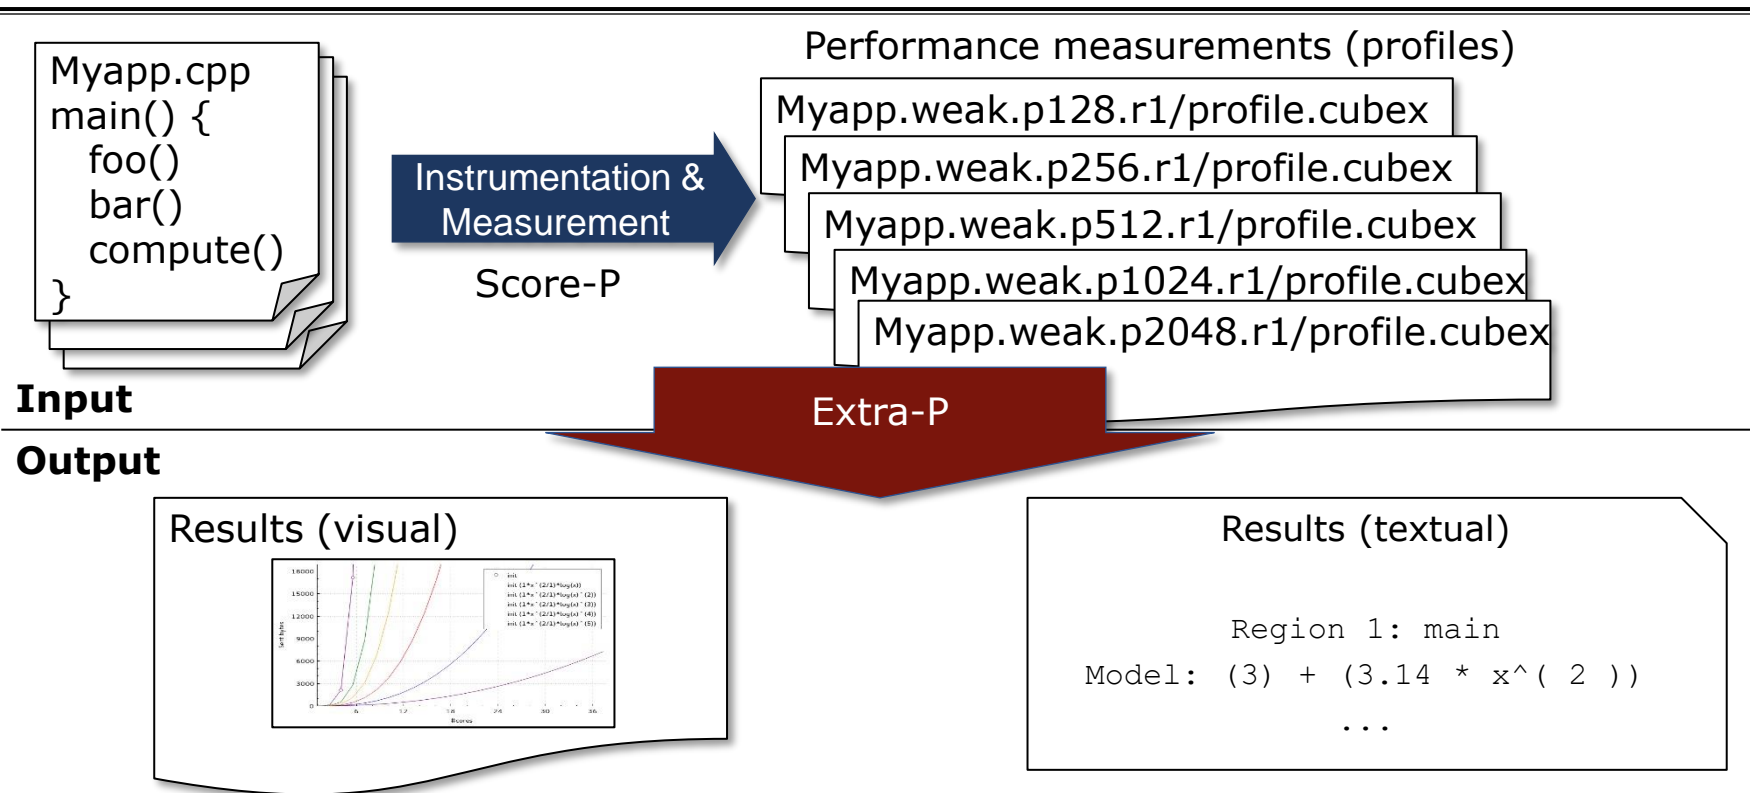
Extra-P in the tuning workshop


- Installed in: `/home/hpc/a2c06/lu23voj/...`
- Run: `source /home/hpc/a2c06/lu23voj/load-extrap`
- Extra-P executables: `extrap, extrap-print`

Installing Extra-P

- Download and install Qt4 or Qt5
- Download and install Python3+ and PyQt5 (Extra-P 3.0 requires matplotlib)
- Download Extra-P: <http://www.scalasca.org/software/extra-p/download.html>
- Unpack & install Extra-P
- ```
./configure --prefix=<extra-p-install-dir> CPPFLAGS=-I<python.h path>
PYTHON=<Python interpreter>; make; make install
```

# Automatic performance modeling with Extra-P





## Modeling sets of Cube experiments

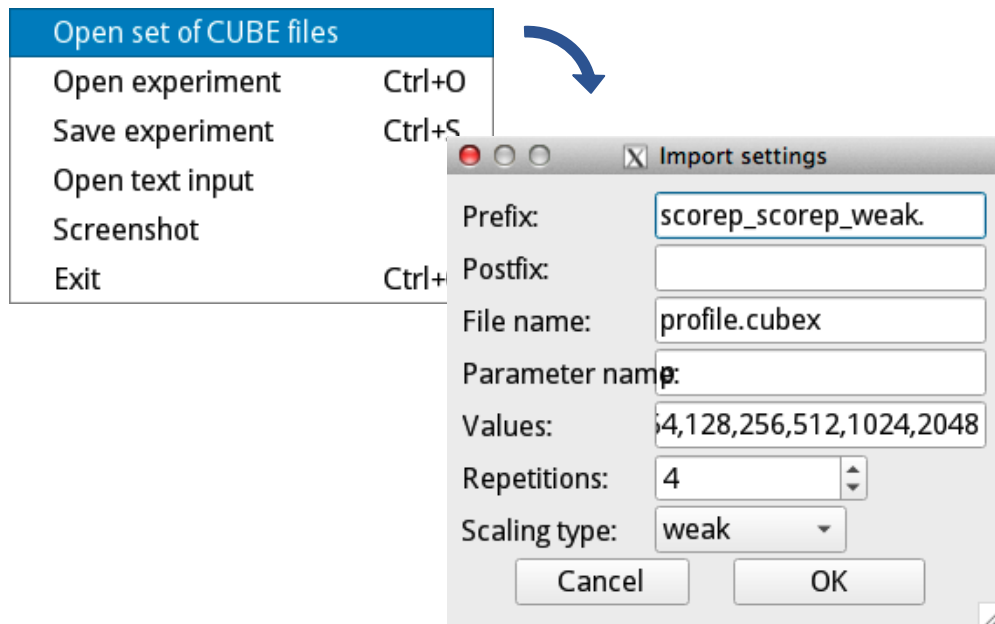
## Extra-P Cube input description

---

- Modeling tool expects Cube files in the following format:  
`<DIR>/<PREFIX><X><POSTFIX>.r<{1,..,REPS}>/<FILENAME>`
- DIR, PREFIX, X, POSTFIX, REPS and FILENAME must all be defined.
  - X – value of varied parameter e.g. number of processes
  - REPS – number of repeated experiments with same parameter value

## Extra-P Cube input description

<DIR>/<PREFIX><X><POSTFIX>.r<{1,..,REPS}>/<FILENAME>

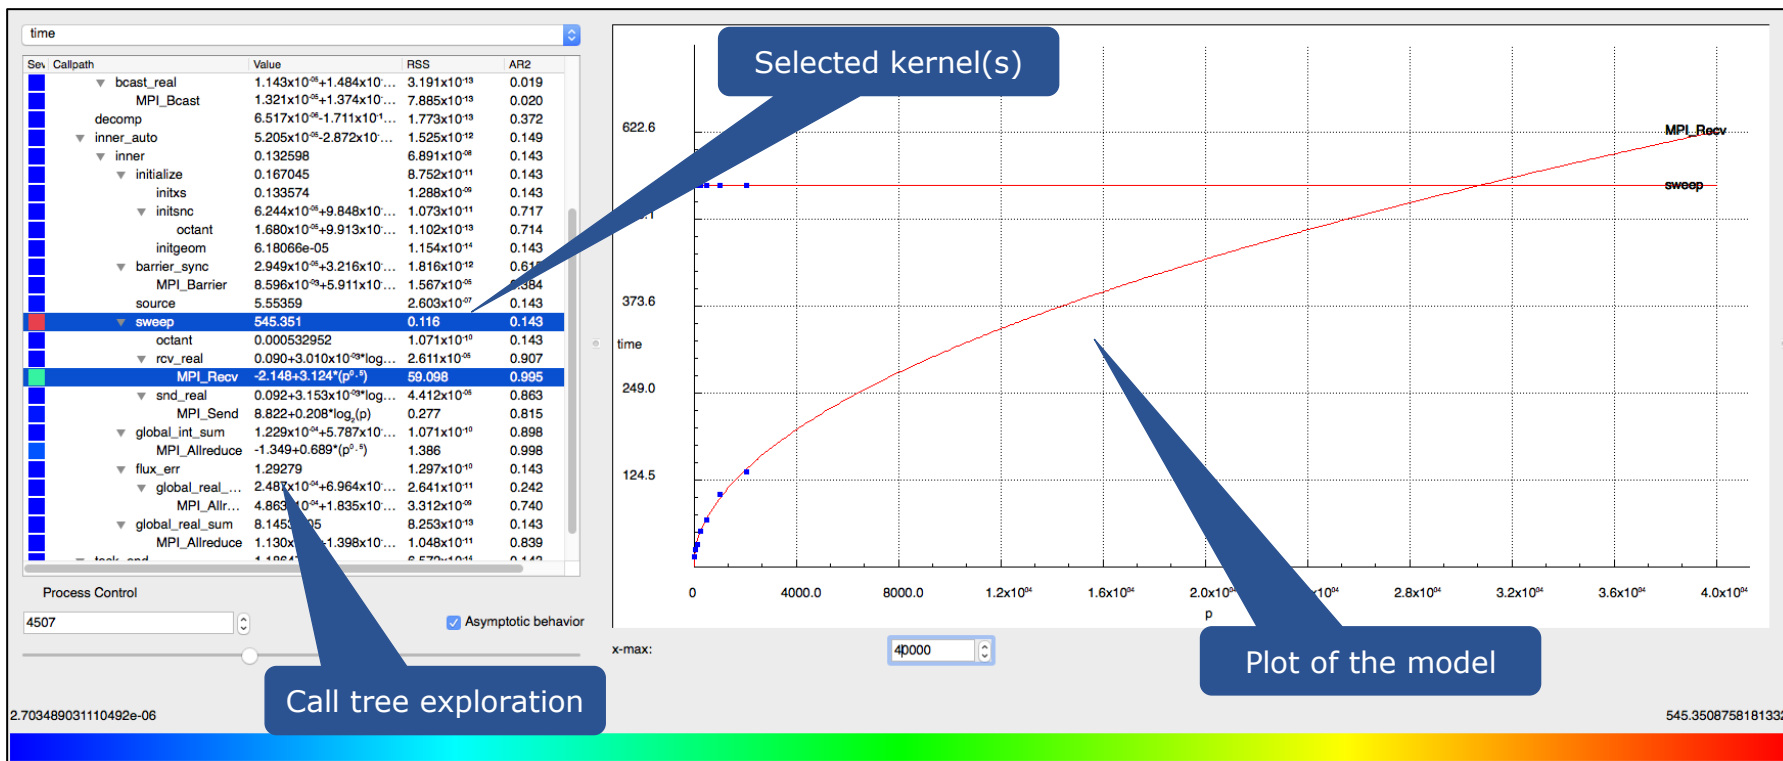


Best effort approach to identify and populate the fields automatically based on selected directory contents.



## Visualization with Extra-P

# Extra-P user interface



## Extra-P call tree view

Metric selection

Call tree exploration

Model

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

Asymptotic view of  
model functions vs.  
value at given value

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

time

| Sev | Callpath        | Value                                                   | RSS                     | AR2   |
|-----|-----------------|---------------------------------------------------------|-------------------------|-------|
|     | MPI_Bcast       | $1.321 \times 10^{-06} + 1.374 \times 10^{-06} \log(p)$ | $7.885 \times 10^{-13}$ | 0.020 |
|     | decomp          | $6.517 \times 10^{-06} - 1.711 \times 10^{-06} \log(p)$ | $1.773 \times 10^{-13}$ | 0.372 |
|     | inner_auto      | $5.205 \times 10^{-05} - 2.872 \times 10^{-05} \log(p)$ | $1.525 \times 10^{-12}$ | 0.149 |
|     | inner           | 0.132598                                                | $6.891 \times 10^{-08}$ | 0.143 |
|     | initialize      | 0.167045                                                | $8.752 \times 10^{-11}$ | 0.143 |
|     | initxs          | 0.133574                                                | $1.288 \times 10^{-09}$ | 0.143 |
|     | initsnc         | $6.244 \times 10^{-05} + 9.848 \times 10^{-05} \log(p)$ | $1.073 \times 10^{-11}$ | 0.717 |
|     | octant          | $1.680 \times 10^{-05} + 9.913 \times 10^{-05} \log(p)$ | $1.102 \times 10^{-13}$ | 0.714 |
|     | initgeom        | 6.18066e-05                                             | $1.154 \times 10^{-14}$ | 0.143 |
|     | barrier_sync    | $2.949 \times 10^{-05} + 3.216 \times 10^{-05} \log(p)$ | $1.816 \times 10^{-12}$ | 0.615 |
|     | MPI_Barrier     | $8.596 \times 10^{-03} + 5.911 \times 10^{-03} \log(p)$ | $1.567 \times 10^{-05}$ | 0.384 |
|     | source          | 5.55359                                                 | $2.603 \times 10^{-07}$ | 0.143 |
|     | sweep           | 545.351                                                 | 0.116                   | 0.143 |
|     | octant          | $8.800000000000002$                                     | $1.071 \times 10^{-10}$ | 0.143 |
|     | rcv_real        | $0.090 + 3.010 \times 10^{-03} \log(p)$                 | $2.611 \times 10^{-05}$ | 0.907 |
|     | MPI_Recv        | $-2.148 + 3.124 \cdot (p^{0.5})$                        | 59.098                  | 0.995 |
|     | snd_real        | $0.092 + 3.153 \times 10^{-03} \log(p)$                 | $4.412 \times 10^{-05}$ | 0.863 |
|     | MPI_Send        | $8.822 + 0.208 \cdot \log_2(p)$                         | 0.277                   | 0.815 |
|     | global_int_sum  | $1.229 \times 10^{-04} + 5.787 \times 10^{-04} \log(p)$ | $1.071 \times 10^{-10}$ | 0.898 |
|     | MPI_Allreduce   | $-1.349 + 0.689 \cdot (p^{0.5})$                        | 1.386                   | 0.998 |
|     | flux_err        | 1.29279                                                 | $1.297 \times 10^{-10}$ | 0.143 |
|     | global_real_max | $2.487 \times 10^{-04} + 6.964 \times 10^{-04} \log(p)$ | $2.641 \times 10^{-11}$ | 0.242 |
|     | MPI_Allreduce   | $4.863 \times 10^{-04} + 1.835 \times 10^{-04} \log(p)$ | $3.312 \times 10^{-09}$ | 0.740 |
|     | global_real_sum | 8.1453e-05                                              | $8.253 \times 10^{-13}$ | 0.143 |
|     | MPI_Allreduce   | $1.130 \times 10^{-04} + 1.398 \times 10^{-04} \log(p)$ | $1.048 \times 10^{-11}$ | 0.839 |
|     | task_end        | 1.18647e-05                                             | $6.572 \times 10^{-14}$ | 0.143 |
|     | MPI_Finalize    | $2.768 \times 10^{-03} + 2.188 \times 10^{-03} \log(p)$ | $1.074 \times 10^{-05}$ | 0.143 |

Process Control

27421


Asymptotic behavior

## 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



## 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 a small data set must be modeled
- Example provided in </home/hpc/a2c06/lu23voj/tutorial/experiment2.txt>

```
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



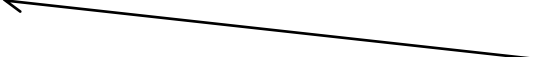
## Extra-P input in text form

---

- Useful when no CUBE files are available or when a small data set must be modeled
- Example provided in `/home/hpc/a2c06/lu23voj/tutorial/experiment2.txt`

```
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

---

- Useful when no CUBE files are available or when a small data set must be modeled
- Example provided in </home/hpc/a2c06/lu23voj/tutorial/experiment2.txt>

```
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

---

- Useful when no CUBE files are available or when a small data set must be modeled
- Example provided in </home/hpc/a2c06/lu23voj/tutorial/experiment2.txt>

```
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

## Extra-P input in text form

---

- Useful when no CUBE files are available or when a small data set must be modeled
- Example provided in </home/hpc/a2c06/lu23voj/tutorial/experiment2.txt>

```
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

## Extra-P input in text form

---

- Useful when no CUBE files are available or when a small data set must be modeled
- Example provided in </home/hpc/a2c06/lu23voj/tutorial/experiment2.txt>

```
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

## Extra-P input in text form

---

- Useful when no CUBE files are available or when a small data set must be modeled
- Example provided in </home/hpc/a2c06/lu23voj/tutorial/experiment2.txt>

```
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




## Results in text form

## Produce text output

---

- Step 1: Save results

|                        |        |
|------------------------|--------|
| Open set of CUBE files |        |
| Open experiment        | Ctrl+O |
| Save experiment        | Ctrl+S |
| Open text input        |        |
| Screenshot             |        |
| Exit                   | Ctrl+Q |



Save the experiment currently open in Extra-P to an output file

- Step 2: Run `extrap-print <Extra-P experiment file>`

## Text output format

---

Measurements and model data for each experiment and metric:

```
Callpath: driver->inner_auto->inner->sweep
Metric: max_time
```

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

```
Data:
(32, 47.4735) 95% CI [0, 0]
(64, 48.2489) 95% CI [0, 0]
(128, 48.9725) 95% CI [0, 0]
(256, 50.53) 95% CI [0, 0]
(512, 51.8811) 95% CI [0, 0]
```

Measurements for each input element (e.g., #processes)

```
Model:
46.1285+0.261952* (p^0.5)
```

Best-fit model

```
RSS: 0.419493
```

RSS: Residual sum of squares

```
Adjusted R^2: 0.994
```

Adjusted R<sup>2</sup> (explained previously)



## Hands-on exercises

## Extra-P exercises

---

- Run: `extrap`
- Examples: `/home/hpc/a2c06/lu23voj/tutorial/{blast,sweep3D}`
- Open the examples in the GUI
- Open the textual input example
- Produce textual output and inspect it

## Feedback

---

- What additional features would you like to see?
- What additional capabilities would you like to see?