





German Research School  
for Simulation Sciences

# Using Automated Performance Modeling to Find Scalability Bugs in Complex Codes




A. Calotoiu<sup>1</sup>, T. Hoefler<sup>2</sup>, M. Poke<sup>1</sup>, F. Wolf<sup>1</sup>

1) German Research School for Simulation Sciences

2) ETH Zurich


September 13, 2013



# Analytical performance modeling



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



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

## Disadvantages

- Time consuming
- Danger of overlooking unscalable code



Generate an empirical model for each part of the program automatically

- Run a manageable number of small-scale performance experiments
- Launch our tool
- Compare extrapolated performance to expectations

## Key ideas

- Exploit that space of function classes underlying such model is small enough to be searched by a computer program
- Abandon model accuracy as the primary success metric and rather focus on the binary notion of **scalability bugs**
- Create requirements models alongside execution models

- German Research School for Simulation Sciences, Laboratory for Parallel Programming (Prof. Dr. Felix Wolf)
- Technische Universität Darmstadt, Institute for Scientific Computing (Prof. Dr. Christian Bischof)
- Swiss Federal Institute of Technology Zurich, Institute of Computer Systems, (Prof. Dr. Torsten Hoefler)
- Forschungszentrum Jülich, Jülich Supercomputing Centre (Dr.-Ing. Bernd Mohr)
- Goethe University Frankfurt, Goethe Center for Scientific Computing (Prof. Dr. Gabriel Wittum)



Overview  
Detailed approach  
Evaluation  
Conclusion

# Scalability bug detector

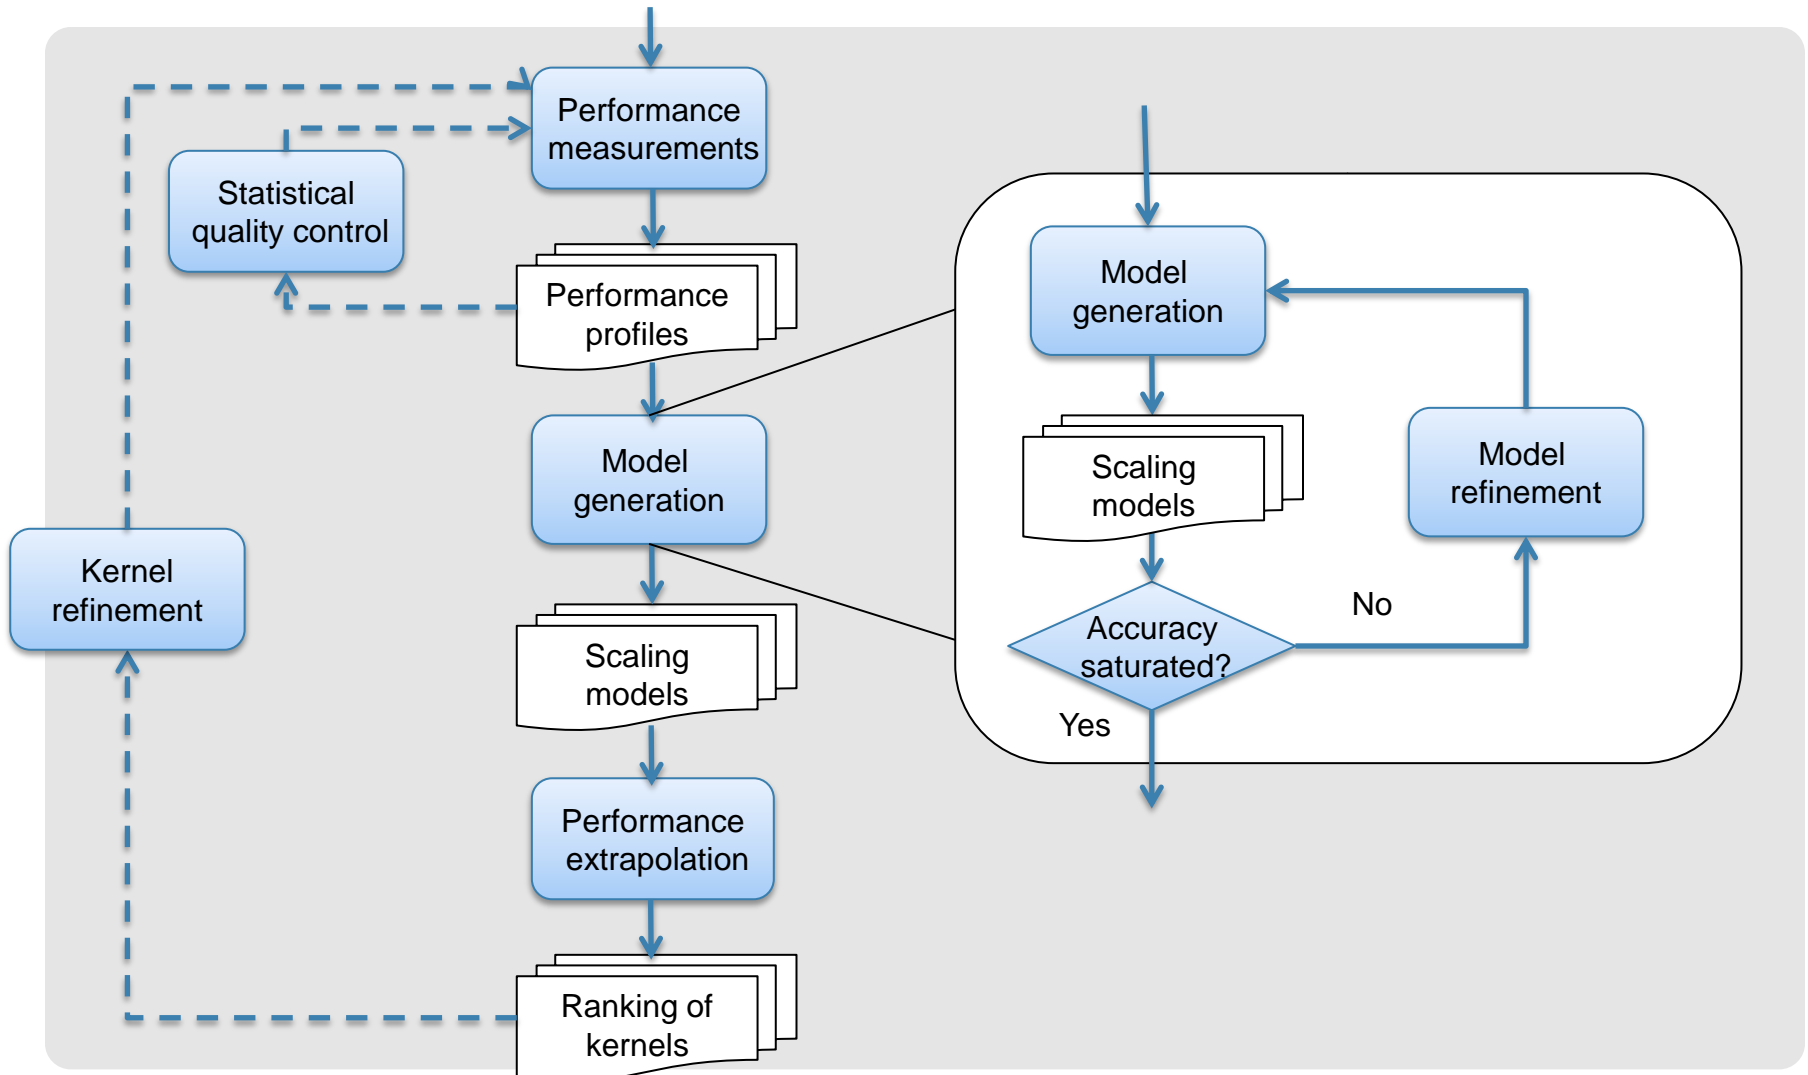
## Input

- Set of performance measurements (profiles) on different processor counts  $\{p_1, \dots, p_{\max}\}$  w/ weak scaling
- Individual measurement broken down by program region (call path)

## Output

- List of program regions (kernels) ranked by their predicted execution time at target scale  $p_t > p_{\max}$
- Or ranked by growth function ( $p_t \rightarrow \infty$ )

- Not 100% accurate but good enough to draw attention to right kernels
- False negatives when phenomenon at scale is not captured in data
- False positives possible but unlikely
- Can also model parameters other than p





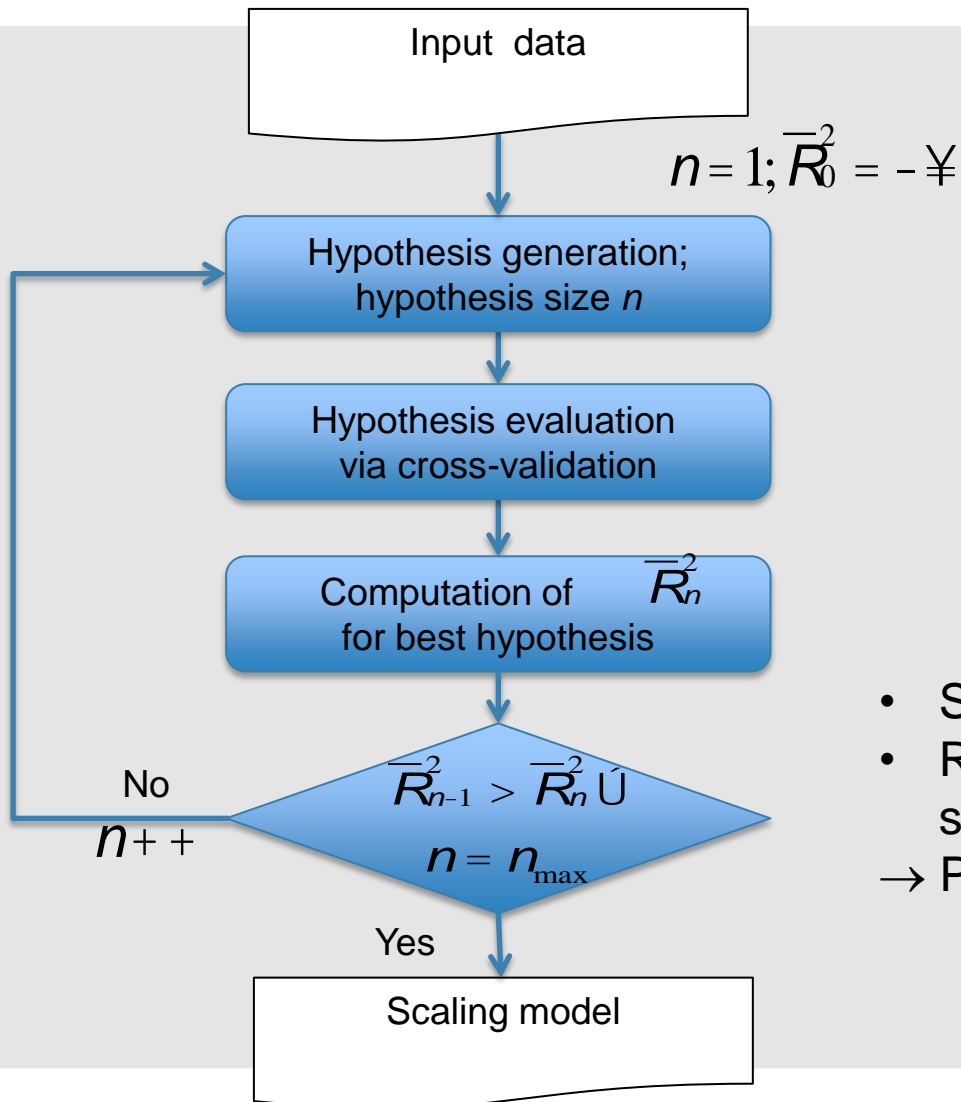
## Performance Model Normal Form (PMNF)

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

- Not exhaustive but works in most practical scenarios
- An assignment of  $n$ ,  $i_k$  and  $j_k$  is called **model hypothesis**
- $i_k$  and  $j_k$  are chosen from sets  $I, J \subset \mathbb{Q}$
- $n$ ,  $|I|$ ,  $|J|$  don't have to be arbitrarily large to achieve good fit

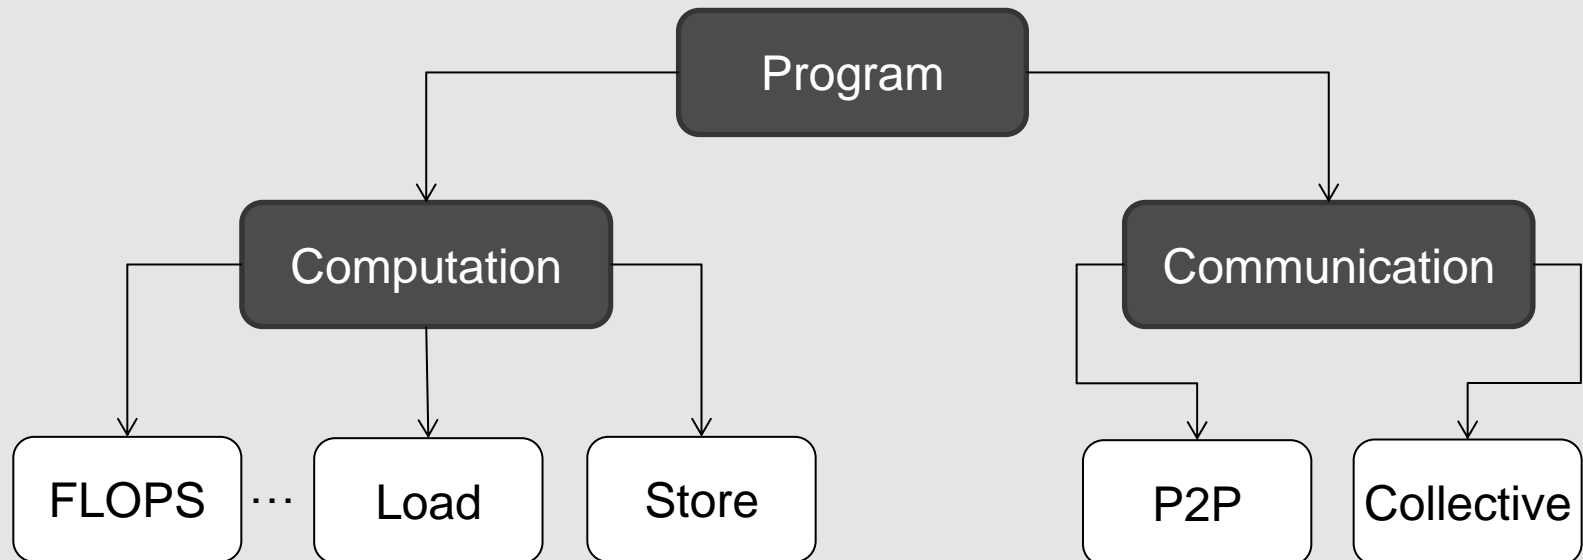
Instead of deriving model through reasoning, make reasonable choices for  $n$ ,  $I$ ,  $J$  and try all assignment options one by one

- Select winner through **cross-validation**



- Start with coarse approximation
  - Refine to the point of statistical shrinkage
- Protection against over-fitting

# Requirements modeling



Disagreement may be indicative of wait states

Time

We demonstrate that our tool

- identifies a scalability issue in a code that is known to have one
- does not identify a scalability issue in a code that is known to have none
- identifies two scalability issues in a code that was thought to have only one

Test platform:  
IBM Blue Gene/Q  
Juqueen in Jülich

$$I = \left\{ \frac{0}{2}, \frac{1}{2}, \frac{2}{2}, \frac{3}{2}, \frac{4}{2}, \frac{5}{2}, \frac{6}{2} \right\}$$

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

$$n = 5$$


Solves neutron transport problem

- 3D domain mapped onto 2D process grid
- Parallelism achieved through pipelined wave-front process

LogGP model for communication developed by Hoisie et al.

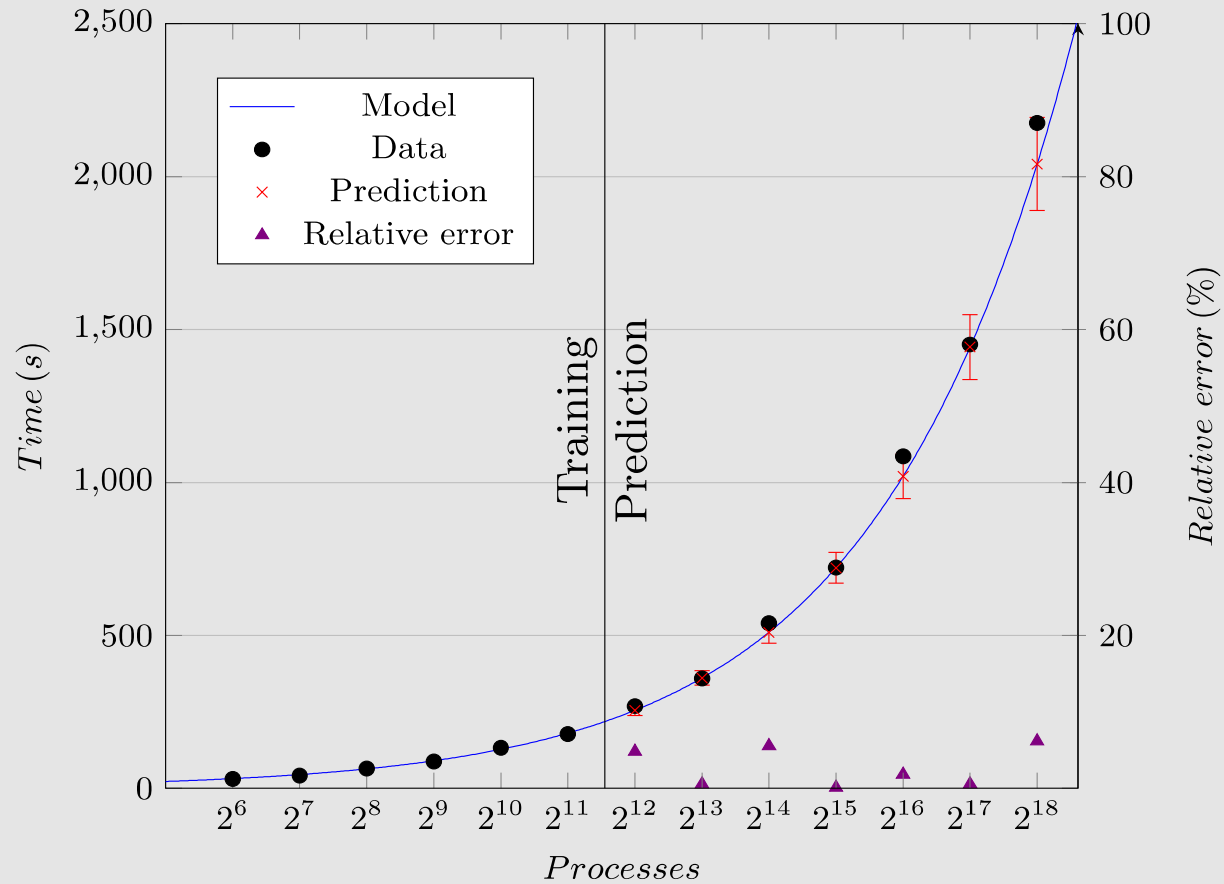
$$t^{comm} = [2(p_x + p_y - 2) + 4(n_{sweep} - 1)] \times t_{msg}$$

$$t^{comm} = c \times \sqrt{p}$$



Kernel	Runtime[%] $p_t=262k$	Increase $\frac{t(p=262k)}{t(p=64)}$	Model [s] $t = f(p)$	Predictive error [%] $p_t=262k$
sweep->MPI_Recv	65.35	16.54	$4.03\sqrt{p}$	5.10
sweep	20.87	0.23	582.19	0.01
global_int_sum-> MPI_Allreduce	12.89	18.68	$1.06\sqrt{p} + 0.03\sqrt{p} \log(p)$	13.60
sweep->MPI_Send	0.40	0.23	$11.49 + 0.09\sqrt{p} \log(p)$	15.40
source	0.25	0.04	$6.86 + 9.13 \cdot 10^{-5} \log(p)$	0.01

$$p_i \leq 8k$$



MILC/su3\_rmd – code from MILC suite of QCD codes with performance model manually created by Hoefler et al.

- Time per process should remain constant except for a rather small logarithmic term caused by global convergence checks

Kernel	Model [s] $t=f(p)$	Predictive Error [%] $p_t=64k$
compute_gen_staple_field	$2.40 \cdot 10^{-2}$	0.43
g_vecdoublesum>MPI_Allreduce	$6.30 \cdot 10^{-6} \cdot \log^2 p$	0.01
mult_adj_su3_fieldlink_lathwec	$3.80 \cdot 10^{-3}$	0.04

$$p_i \leq 16k$$






# MILC – Varying grid points per process

Kernel	Flops	$\overline{R^2}$	Visits	$\overline{R^2}$	Flops/Visit	
	Model[s] flops=f(V)	[*10 <sup>-3</sup> ]	Model visits=f(V)	[*10 <sup>-3</sup> ]	Model	$\overline{R^2}$ [*10 <sup>-3</sup> ]
load_lnglinks	5.64*10 <sup>4</sup> *V	0.030	2.31*10 <sup>3</sup>	0.000	24.42*V	0.030
load_fatlinks_cpu	1.95*10 <sup>6</sup> *V	0.210	7.14*10 <sup>4</sup>	0.000	27.36*V	0.210
ks_congrad	1.16*10 <sup>8</sup> + 3.24*10 <sup>5</sup> *V <sup>5/4</sup>	0.292	5.11*10 <sup>4</sup> + 1.38*10 <sup>4</sup> *V <sup>1/4</sup>	4.000	15.94*V	0.143
imp_gauge_force_cpu	1.65*10 <sup>6</sup> *V	0.015	7.40*10 <sup>4</sup>	0.000	22.28*V	0.015
eo_fermion_force_two_terms_site	4.02*10 <sup>6</sup> *V	0.002	1.27*10 <sup>5</sup>	0.000	31.61*V	0.002

Test platform: Juropa in Jülich (Intel Nehalem cluster)  
p = 32= constant

## Core of the Community Atmospheric Model (CAM)

- Spectral element dynamical core on a cubed sphere grid



Kernel	$p_i \leq 15k$		$p_i \leq 43k$	
	Model [s] $t = f(p)$	Predictive error [%] $p_t = 130k$	Model [s] $t = f(p)$	Predictive error [%] $p_t = 130k$
Box_rearrange->MPI_Reduce	$0.03 + 2.53 \cdot 10^{-6} p^* \sqrt{p} + 1.24 \cdot 10^{-12} p^3$	57.02	$3.63 \cdot 10^{-6} p^* \sqrt{p} + 7.21 \cdot 10^{-13} p^3$	30.34
Vlaplace_sphere_vk	49.53	99.32	$24.44 + 2.26 \cdot 10^{-7} p^2$	4.28
...				
Compute_and_apply_rhs	48.68	1.65	49.09	0.83

## Two issues

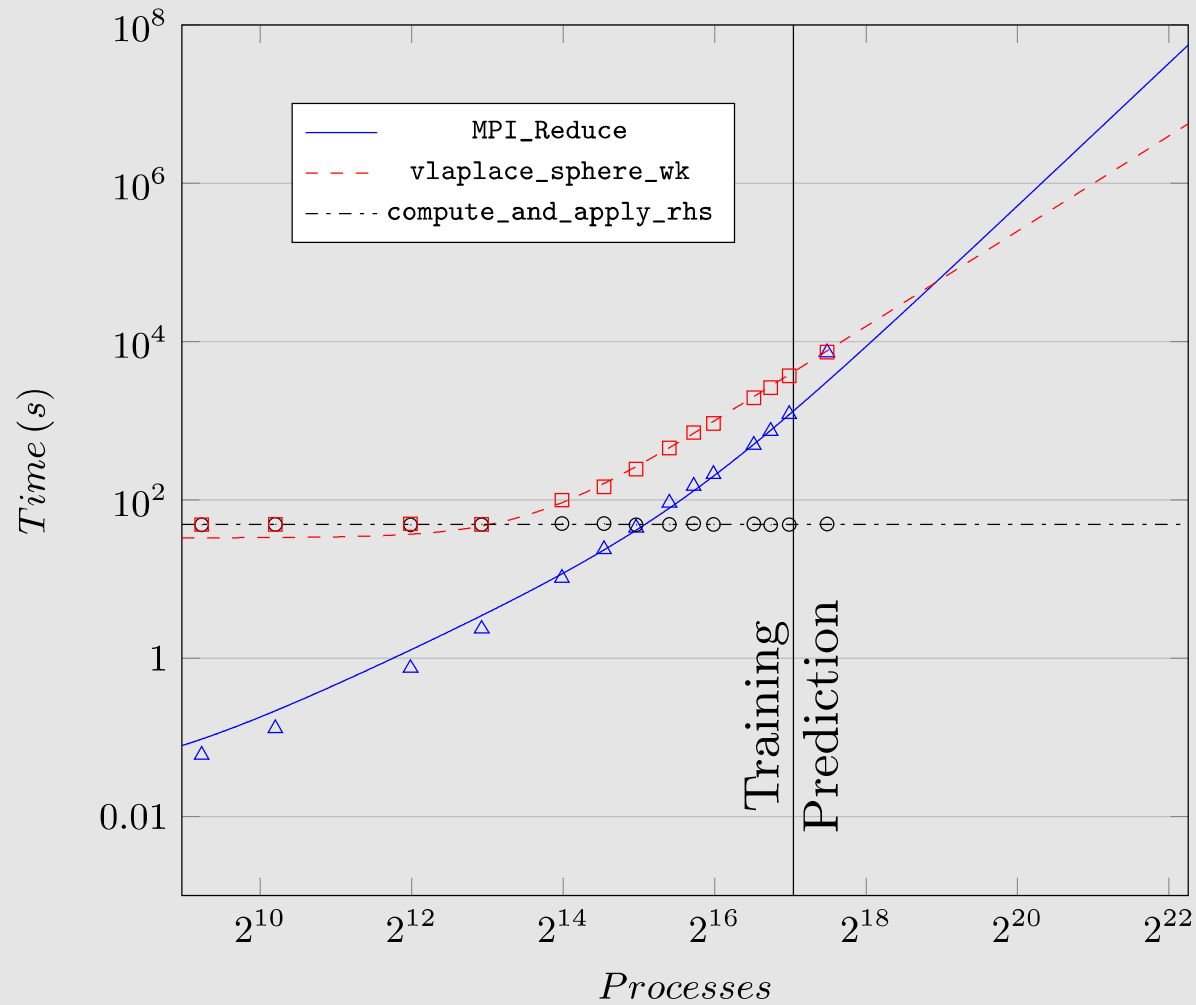
Number of iterations inside a subroutine grows with  $p^2$

- Ceiling for up to and including 15k
- Developers were aware of this issue and had developed work-around

Growth of time spent in reduce function grows with  $p^3$

- Previously unknown
- Function invoked during initialization to funnel data to dedicated I/O processes
- Execution time at 183k ~ 2h, predictive error ~40%

The G8 Research Councils Initiative on Multilateral Research Funding  
Interdisciplinary Program on Application Software towards Exascale Computing for Global Scale Issues



Automated performance modeling is feasible

Generated models accurate enough to identify scalability bugs and in good agreement with hand-crafted models

Advantages of mass production also performance models

- Approximate models are acceptable as long as the effort to create them is low and they do not mislead the user
- Code coverage is as important as model accuracy

Future work

- Study influence of further hardware parameters
- More efficient traversal of search space (allows increase of modeling parameters)
- Integration into Scalasca