

TORSTEN HOEFLER

An Overview of Static & Dynamic Techniques for Automatic Performance Modeling

in collaboration with Alexandru Calotoiu and Felix Wolf @ RWTH Aachen
with students Arnamoy Bhattacharyya and Grzegorz Kwasniewski @ SPCL
presented at ISC 2016, Frankfurt, July 2016

My sinful youth



- **Original findings:**
 - If carefully tuned, NBC speeds up a 3D solver
 - Full code published*
 - 800^3 domain – 4 GB array
 - 1 process per node, 8-96 nodes*
 - Opteron 246 (old even in 2006, retired now)*
 - Super-linear speedup for 96 nodes
 - ~5% better than linear*

- **9 years later: attempt to reproduce 😊!**

*System A: 28 quad-core nodes,
Xeon E5520*

*System B: 4 nodes, dual
Opteron 6274*

“Neither the experiment in A nor the one in B could reproduce the results presented in the original paper, where the usage of the NBC library resulted in a performance gain for practically all node counts, reaching a superlinear speedup for 96 cores (explained as being due to cache effects in the inner part of the matrix vector product).”



How to report a performance result?

Scientific Benchmarking of Parallel Computing Systems

Twelve ways to tell the masses when reporting performance results

@SC'15

Torsten Hoefler
Dept. of Computer Science
ETH Zurich
Zurich, Switzerland
htor@inf.ethz.ch

Roberto Belli
Dept. of Computer Science
ETH Zurich
Zurich, Switzerland
bellir@inf.ethz.ch

ABSTRACT

Measuring and reporting performance of parallel computers constitutes the basis for scientific advancement of high-performance computing (HPC). Most scientific reports show performance improvements of new techniques and are thus obliged to ensure reproducibility or at least interpretability. Our investigation of a stratified sample of 120 papers across three top conferences in the field shows that the state of the practice is lacking. For example, it is often unclear if reported improvements are deterministic or observed by chance. In addition to distilling best practices from existing work, we propose statistically sound analysis and reporting techniques and simple guidelines for experimental design in parallel computing and codify them in a portable benchmarking library. We aim to improve the standards of reporting research results and initiate a discussion in the HPC field. A wide adoption of our minimal set of rules will lead to better interpretability of performance results and improve the scientific culture in HPC.

Reproducing experiments is one of the main principles of the scientific method. It is well known that the performance of a computer program depends on the application, the input, the compiler, the runtime environment, the machine, and the measurement methodology [20, 43]. If a single one of these aspects of *experimental design* is not appropriately motivated and described, presented results can hardly be reproduced and may even be misleading or incorrect.

The complexity and uniqueness of many supercomputers makes reproducibility a hard task. For example, it is practically impossible to recreate most hero-runs that utilize the world's largest machines because these machines are often unique and their software configurations changes regularly. We introduce the notion of *interpretability*, which is weaker than reproducibility. We call an *experiment interpretable* if it provides enough information to allow scientists to understand the experiment, draw own conclusions, assess their certainty, and possibly generalize results. In other words, interpretable experiments support sound conclusions and convey precise information among scientists. Obviously, every scientific

Analytical application performance modeling

■ Scalability bug prediction

- Find latent scalability bugs early on (before machine deployment)

SC13: A. Calotoiu, TH, M. Poke, F. Wolf: Using Automated Performance Modeling to Find Scalability Bugs in Complex Codes

■ Automated performance testing

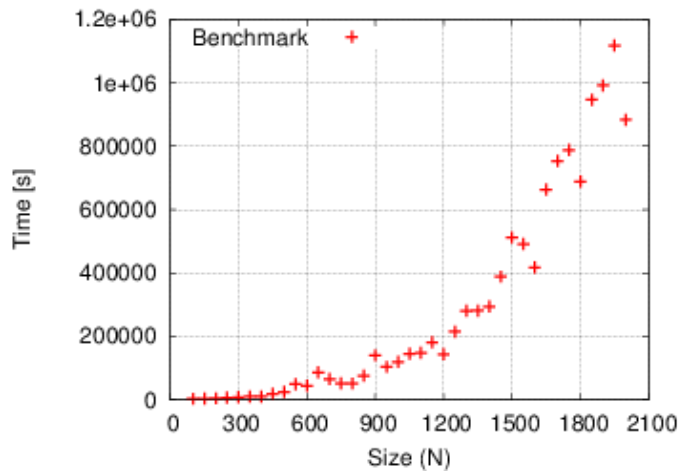
- Performance modeling as part of a software engineering discipline in HPC

ICS'15: S. Shudler, A. Calotoiu, T. Hoefler, A. Strube, F. Wolf: Exascaling Your Library: Will Your Implementation Meet Your Expectations?


■ Hardware/Software co-design

- Decide how to architect systems

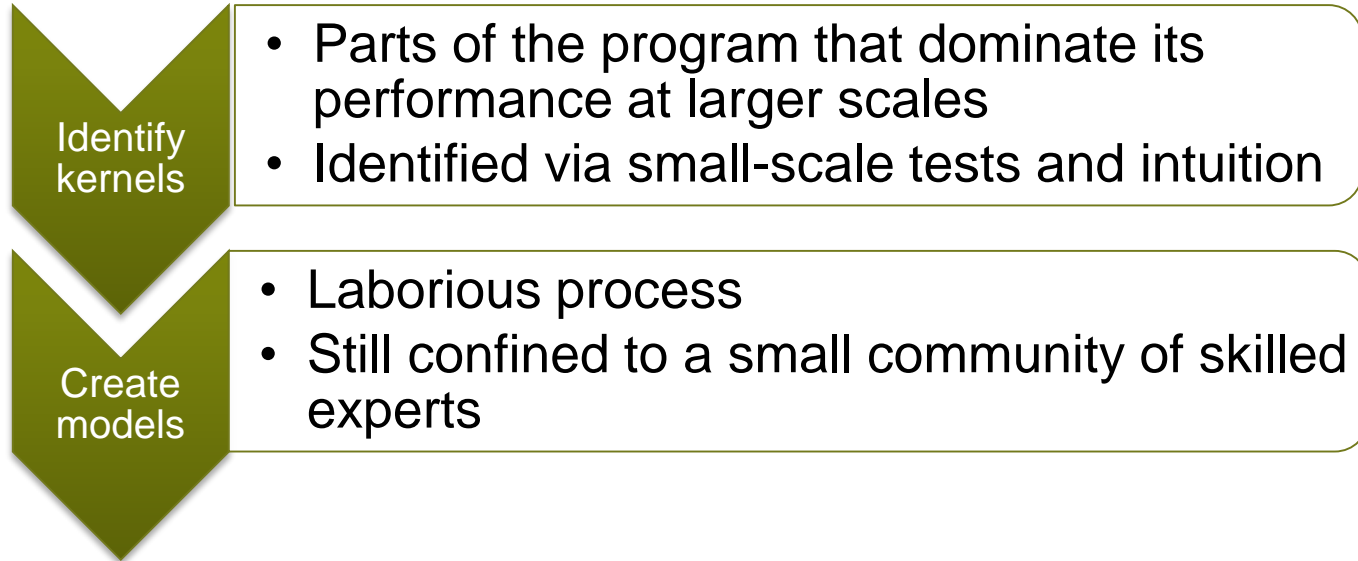
■ Making performance development intuitive



vs.



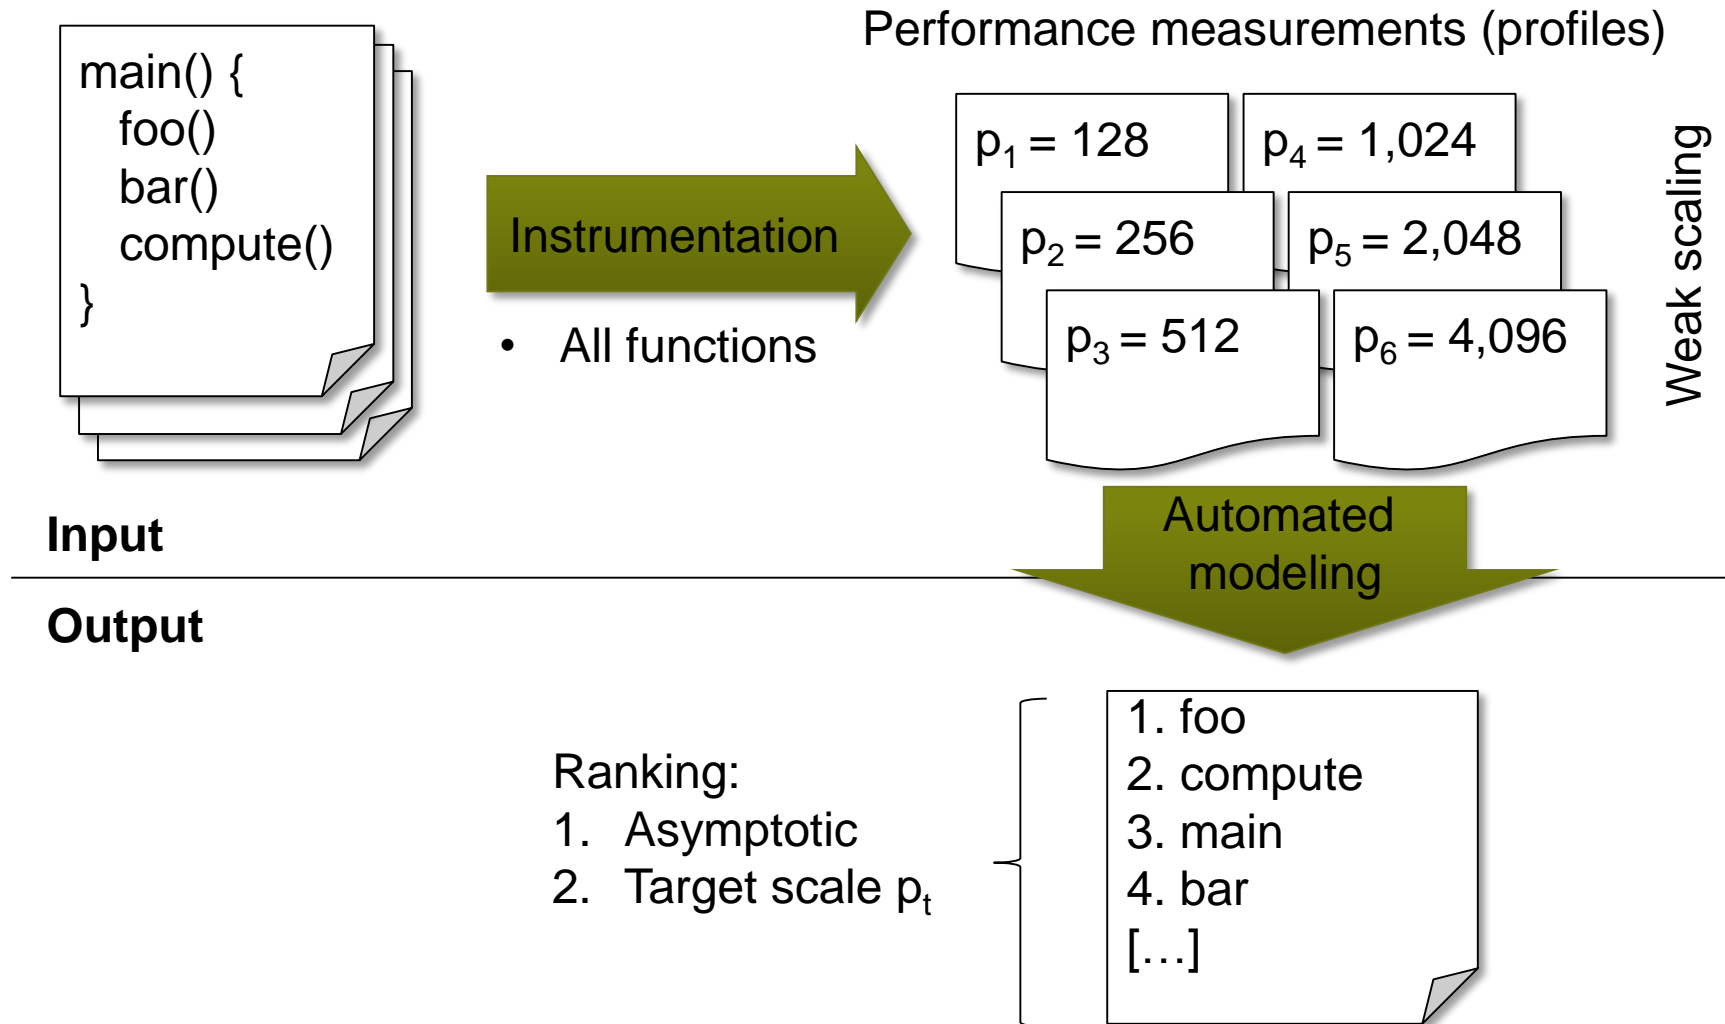
Manual analytical performance modeling




■ Disadvantages

- Time consuming
- Error-prone, may overlook unscalable code

Our first step: scalability bug detector




Primary focus on scaling trend



Our ranking

1. F_1
2. F_3
3. F_2


Primary focus on scaling trend



Our ranking

1. F_1
2. F_3
3. F_2

Primary focus on scaling trend



Our ranking

1. F_1
2. F_3
3. F_2

Survey result: performance model normal form

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

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

$$n = 1$$

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

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

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

Survey result: performance model normal form

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

 $n \in \mathbb{N}$
 $i_k \in I$
 $n = 2$
 $I = \{0, 1, 2\}$
 $J = \{0, 1\}$

$c_1 + c_2 \times p$

$c_1 + c_2 \times p^2$

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

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

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

$c_1 \cdot \log(p) + c_2 \cdot p$

$c_1 \cdot \log(p) + c_2 \cdot p \cdot \log(p)$

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

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

$c_1 \cdot p + c_2 \cdot p \cdot \log(p)$

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

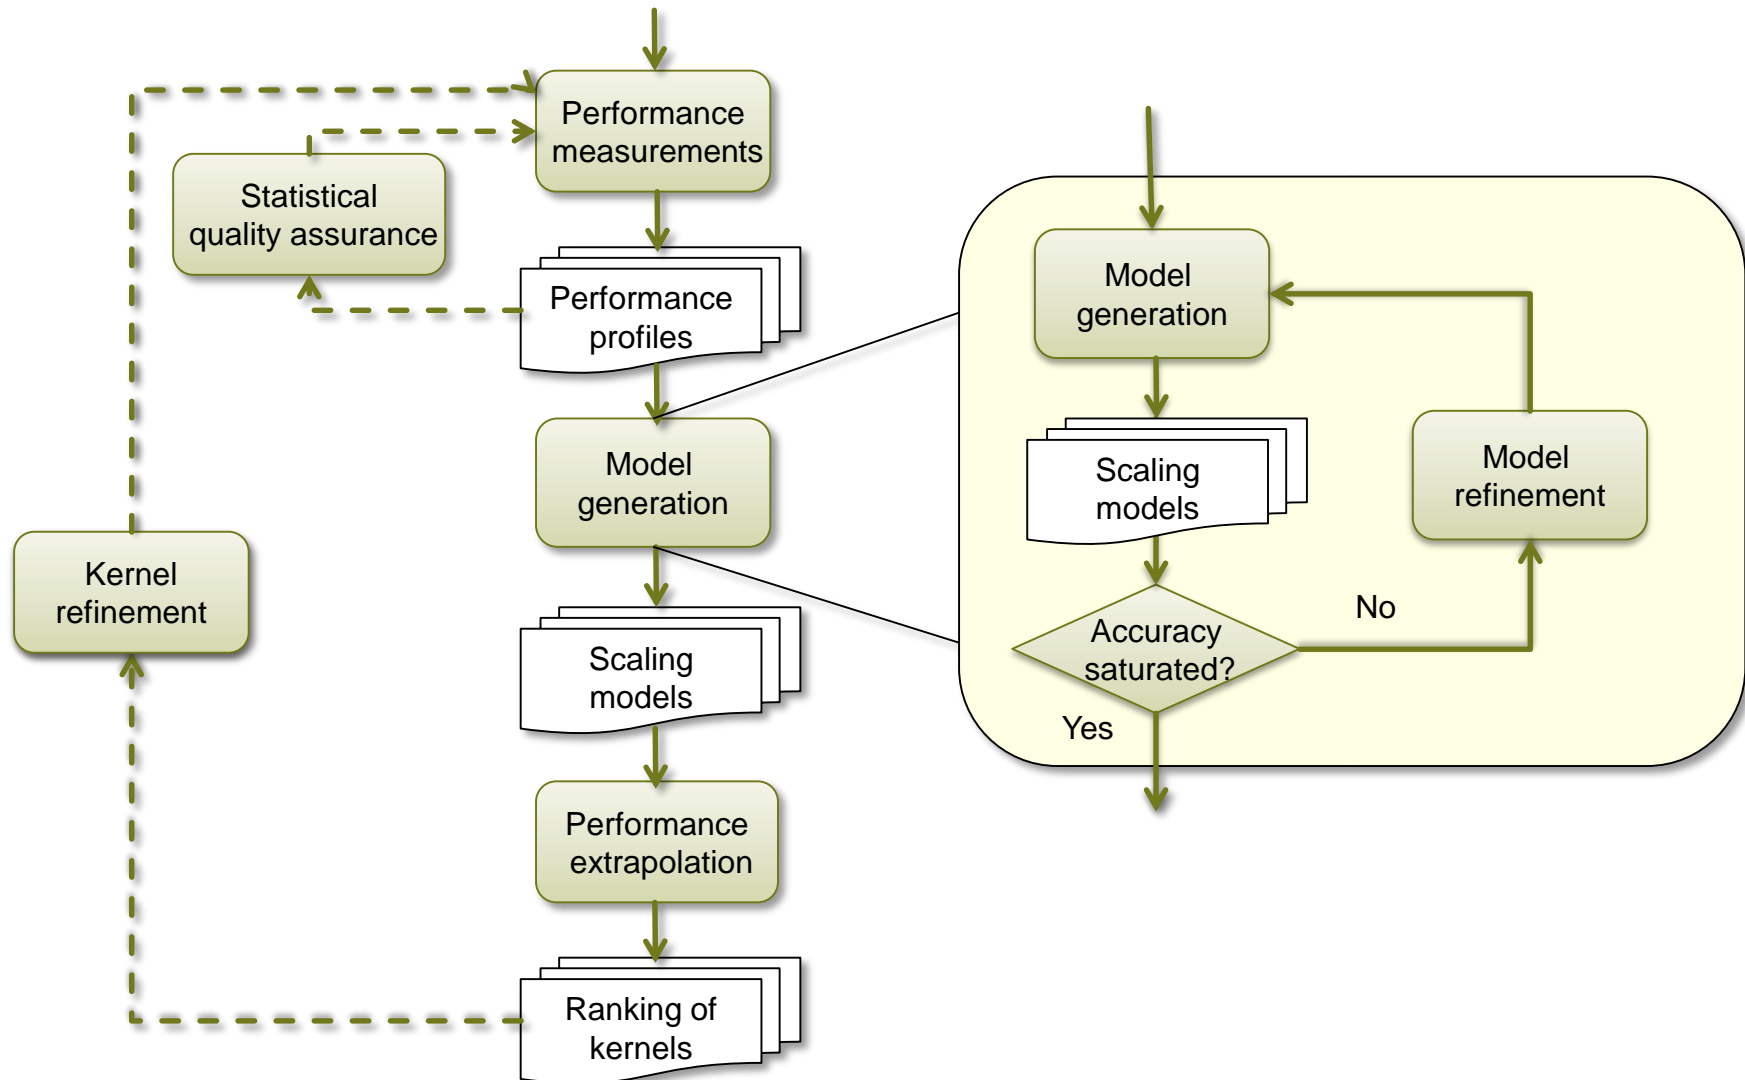
$c_1 \cdot p + c_2 \cdot p^2 \cdot \log(p)$

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

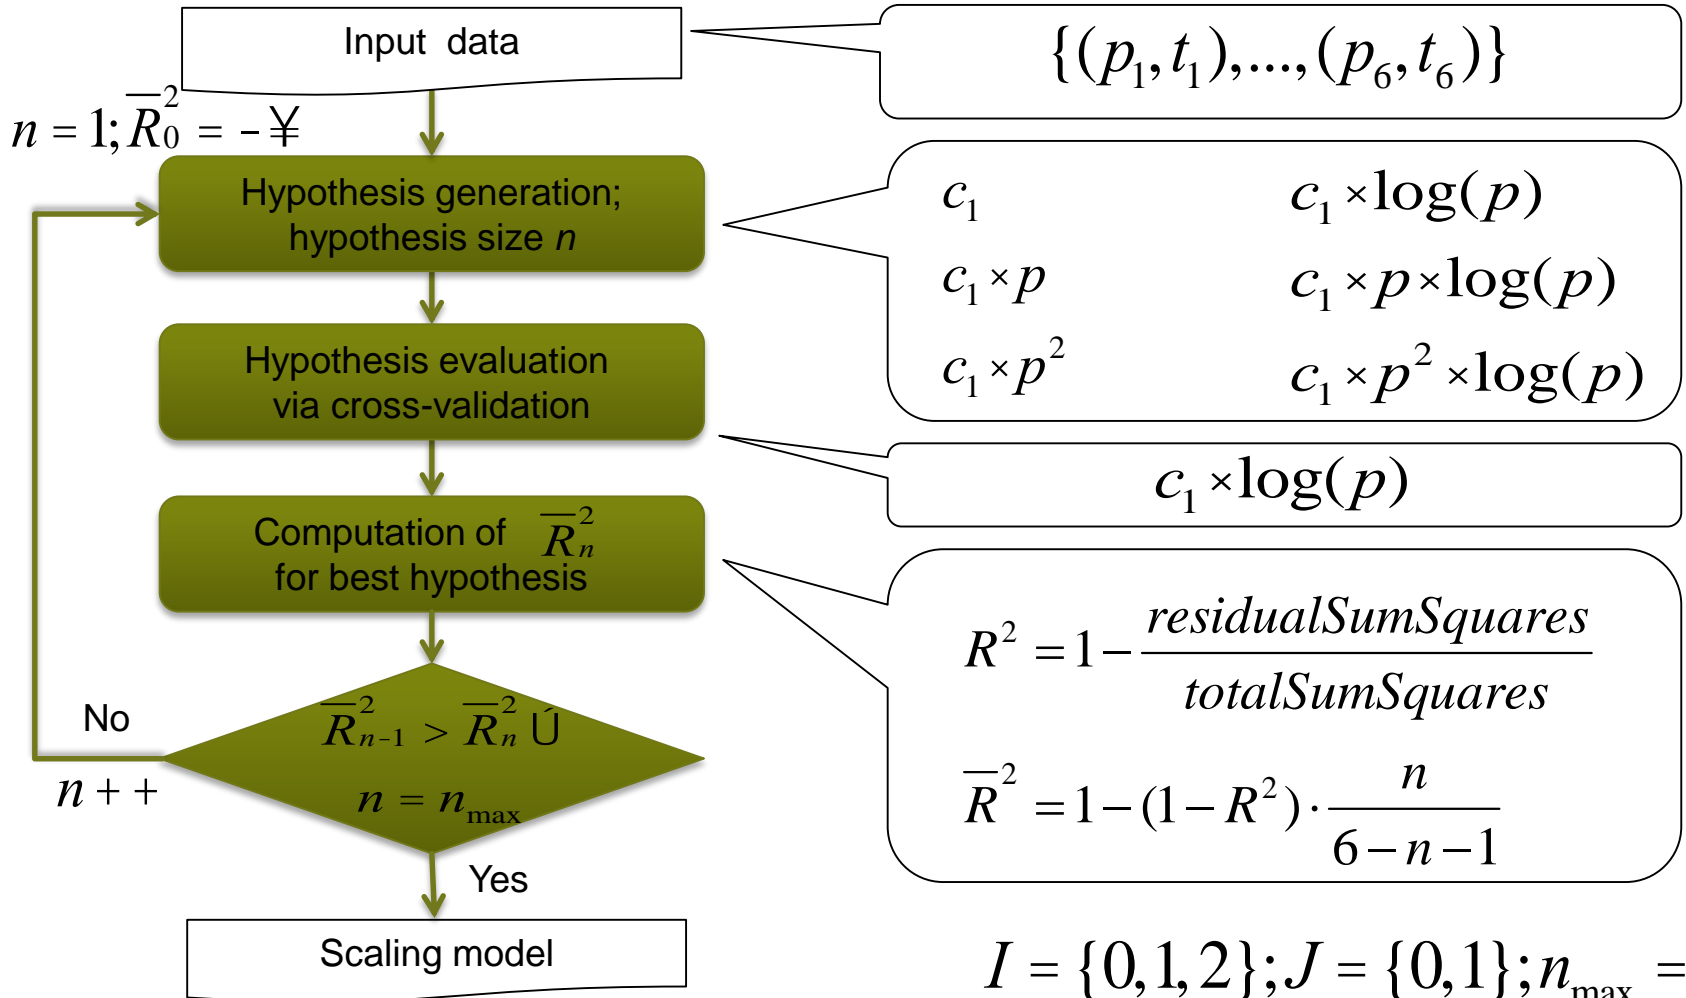
$c_1 \cdot p \cdot \log(p) + c_2 \cdot p^2 \cdot \log(p)$

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

Our automated generation workflow



Model refinement



Is this all? No, it's just the beginning ...

- **We face several problems:**
 - Multiparameter modeling – search space explosion
Interesting instance of the curse of dimensionality
 - Modeling overheads
Cross validation (leave-one-out) is slow and
Our current profiling requires a lot of storage (>TBs)




Static analysis of explicitly parallel programs

- Structures that determine program scalability

LOOPS

- Assumption:**
Other instructions do not influence it
- Example:**

```
for (x=0; x < n/p; x++)  
  for (y=1; y < n; y=2*y )  
    veryComplicatedOperation(x,y) ;
```



Counting arbitrary affine loop nests

■ Affine loops

```

x=x0;           // Initial assignment
while(cTx < g)  // Loop guard
  x=Ax + b;      // Loop update
  
```

■ Perfectly nested affine loops

```

while(c1Tx < g1) {
  x = A1x + b1;
  while(c2Tx < g2) {
    ...
    x = Ak-1x + bk-1;
    while(ckTx < gk) {
      x = Akx + bk;
      while(ck+1Tx < gk+1) {... }
      x = Ukx + vk; }
    x = Uk-1x + vk-1;
    ... }
  x = U1x + v1;}
  
```

$A_k, U_k \in \mathbb{R}^{m \times m}$, $b_k, v_k, c_k \in \mathbb{R}^m$, $g_k \in \mathbb{R}$ and $k = 1 \dots r$.

Counting arbitrary affine loop nests

- **Example**

```
for (j=1; j < n/p + 1; j= j*2)
    for (k=j; k < m; k = k + j )
        veryComplicatedOperation(j,k);
```

Counting arbitrary affine loop nests

■ Example

```

for (j=1; j < n/p + 1; j= j*2)
    for (k=j; k < m; k = k + j )
        veryComplicatedOperation(j,k);
  
```

```

while ( $c_1^T x < g_1$ ) {
   $x = A_1 x + b_1$ ;
  while ( $c_2^T x < g_2$ ) {
    ...
     $x = A_{k-1} x + b_{k-1}$ ;
    while ( $c_k^T x < g_k$ ) {
       $x = A_k x + b_k$ ;
      while ( $c_{k+1}^T x < g_{k+1}$ ) { ... }
       $x = U_k x + v_k$ ; }
     $x = U_{k-1} x + v_{k-1}$ ;
    ... }
   $x = U_1 x + v_1$ ; }
  
```

Counting arbitrary affine loop nests

■ Example

```

for (j=1; j < n/p + 1; j= j*2)
  for (k=j; k < m; k = k + j )
    veryComplicatedOperation(j,k);
  
```

$$\begin{pmatrix} j \\ k \end{pmatrix} = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} j \\ k \end{pmatrix} + \begin{pmatrix} 1 \\ 0 \end{pmatrix};$$

```

while (c1Tx < g1) {
  x = A1x + b1;
  while (c2Tx < g2) {
    ...
    x = Ak-1x + bk-1;
    while (ckTx < gk) {
      x = Akx + bk;
      while (ck+1Tx < gk+1) { ... }
      x = Ukx + vk; }
    x = Uk-1x + vk-1;
    ... }
  x = U1x + v1; }
  
```

Counting arbitrary affine loop nests

■ Example

```

for (j=1; j < n/p + 1; j= j*2)
  for (k=j; k < m; k = k + j )
    veryComplicatedOperation (j ,k) ;

```

$$\begin{pmatrix} j \\ k \end{pmatrix} = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} j \\ k \end{pmatrix} + \begin{pmatrix} 1 \\ 0 \end{pmatrix};$$

$$\text{while}(\mathbf{1} \ 0 \begin{pmatrix} j \\ k \end{pmatrix} < n/p + 1) \{$$

```

while (c1Tx < g1) {
  x = A1x + b1;
  while (c2Tx < g2) {
    ...
    x = Ak-1x + bk-1;
    while (ckTx < gk) {
      x = Akx + bk;
      while (ck+1Tx < gk+1) { ... }
      x = Ukx + vk; }
    x = Uk-1x + vk-1;
    ... }
  x = U1x + v1; }

```

}

Counting arbitrary affine loop nests

■ Example

```

for (j=1; j < n/p + 1; j= j*2)
  for (k=j; k < m; k = k + j )
    veryComplicatedOperation (j ,k) ;

```

$$\begin{pmatrix} j \\ k \end{pmatrix} = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} j \\ k \end{pmatrix} + \begin{pmatrix} 1 \\ 0 \end{pmatrix};$$

$$\text{while}((1 \ 0) \begin{pmatrix} j \\ k \end{pmatrix} < n/p + 1) \{$$

$$\begin{pmatrix} j \\ k \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} j \\ k \end{pmatrix} + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$

$$\text{while}((0 \ 1) \begin{pmatrix} j \\ k \end{pmatrix} < m) \{$$

$$\}$$

$$\}$$

```

while (c1Tx < g1) {
  x = A1x + b1;
  while (c2Tx < g2) {
    ...
    x = Ak-1x + bk-1;
    while (ckTx < gk) {
      x = Akx + bk;
      while (ck+1Tx < gk+1) { ... }
      x = Ukx + vk; }
    x = Uk-1x + vk-1;
    ... }
  x = U1x + v1; }

```

Counting arbitrary affine loop nests

Example

```

for (j=1; j < n/p + 1; j= j*2)
  for (k=j; k < m; k = k + j )
    veryComplicatedOperation (j ,k) ;
  
```

$$\begin{pmatrix} j \\ k \end{pmatrix} = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} j \\ k \end{pmatrix} + \begin{pmatrix} 1 \\ 0 \end{pmatrix};$$

$$\text{while}((1 \ 0) \begin{pmatrix} j \\ k \end{pmatrix} < n/p + 1) \{$$

$$\begin{pmatrix} j \\ k \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} j \\ k \end{pmatrix} + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$

$$\text{while}((0 \ 1) \begin{pmatrix} j \\ k \end{pmatrix} < m) \{$$

$$\begin{pmatrix} j \\ k \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ \mathbf{1} & \mathbf{1} \end{pmatrix} \begin{pmatrix} j \\ k \end{pmatrix} + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$

$$\} \begin{pmatrix} j \\ k \end{pmatrix} = \begin{pmatrix} \mathbf{2} & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} j \\ k \end{pmatrix} + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$

$$\}$$

```

while (c1Tx < g1) {
  x = A1x + b1;
  while (c2Tx < g2) {
    ...
    x = Ak-1x + bk-1;
    while (ckTx < gk) {
      x = Akx + bk;
      while (ck+1Tx < gk+1) { ... }
      x = Ukx + vk; }
    x = Uk-1x + vk-1;
    ... }
  x = U1x + v1; }
  
```

Counting arbitrary affine loop nests

■ Example

```

for (j=1; j < n/p + 1; j= j*2)
  for (k=j; k < m; k = k + j )
    veryComplicatedOperation (j, k) ;

```

$$x = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} x + \begin{pmatrix} 1 \\ 0 \end{pmatrix};$$

while((1 0)x < n/p + 1){

$$x = \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix} x + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$

while((0 1)x < m){

$$x = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} x + \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

$$} x = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix} x + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$

}

where $x = \begin{pmatrix} j \\ k \end{pmatrix}$

```

while (c1Tx < g1) {
  x = A1x + b1;
  while (c2Tx < g2) {
    ...
    x = Ak-1x + bk-1;
    while (ckTx < gk) {
      x = Akx + bk;
      while (ck+1Tx < gk+1) { ... }
      x = Ukx + vk; }
    x = Uk-1x + vk-1;
    ... }
  x = U1x + v1; }

```

Overview of the whole system

Parallel program

```

do i = 1, procCols
  call mpi_irecv( buff, ..., dp_type, reduce_exch_proc(i),
  > i, mpi_comm_world, request, ierr )
  call mpi_send( buff2, ..., dp_type, reduce_exch_proc(i),
  > i, mpi_comm_world, ierr )
  call mpi_wait( request, status, ierr )
enddo

do i = id * n/p, ( id + ) * n/p
  do j = 1, nSize
    call compute
  
```



Closed form representation

$$x(i_1, \dots, i_r) = A_{final}(i_1, \dots, i_r) \cdot x_0 + b_{final}(i_1, \dots, i_r)$$

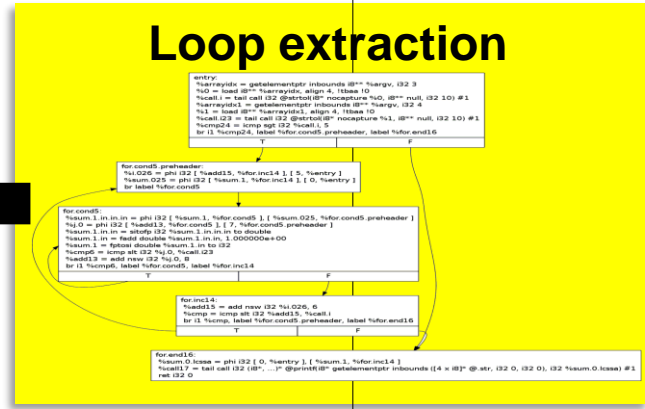
with
 $i_r = 0 \dots n_k(x_{0,k}), k = 1 \dots r$

Affine loop synthesis

```

while (c1^T x < g1) {
  x = A1 x + b1;
  while (c2^T x < g2) {
    ...
    x = Ak-1 x + bk-1;
    while (ck^T x < gk) {
      x = Ak x + bk;
      while (c_{k+1}^T x < g_{k+1}) { ... }
      x = Uk x + vk;
    }
    x = Uk-1 x + vk-1;
  }
  ...
  x = U1 x + v1;
}

```




Number of iterations

$$N = \sum_{i_1=0}^{n_1(x_{0,1})} \sum_{i_2=0}^{n_2(x_{0,2})} \dots \sum_{i_{r-1}=0}^{n_{r-1}(x_{0,r-1})} n_r(x_{0,r})$$

Program analysis

$$W = N|_{p=1}$$

$$D = N|_{p \rightarrow \infty}$$


What problems are remaining?

- **Well, what about non-affine loops?**

- More general abstract interpretation (next step)
- Not solvable → will always have undefined terms

$$N = \frac{na \cdot u}{nprows}$$

- **Back to PMNF?**


- Generalize to multiple input parameters
 - a) *Bigger search-space* 😞
 - b) *Bigger trace files* 😞

- **Ad-hoc (partial) solution: online machine learning – PEMOGEN**



- Replace cross-validation with LASSO (regression with L_1 regularizer)
Much cheaper!
- Replace LASSO with online LASSO [1]
No traces!

PEMOGEN – static analysis

- Also integrated into LLVM compiler
 - Automatic kernel detection and instrumentation (Loop Call Graph)
 - Static dataflow analysis reduces parameter space for each kernel




Quality: NAS UA and Mantevo MiniFE



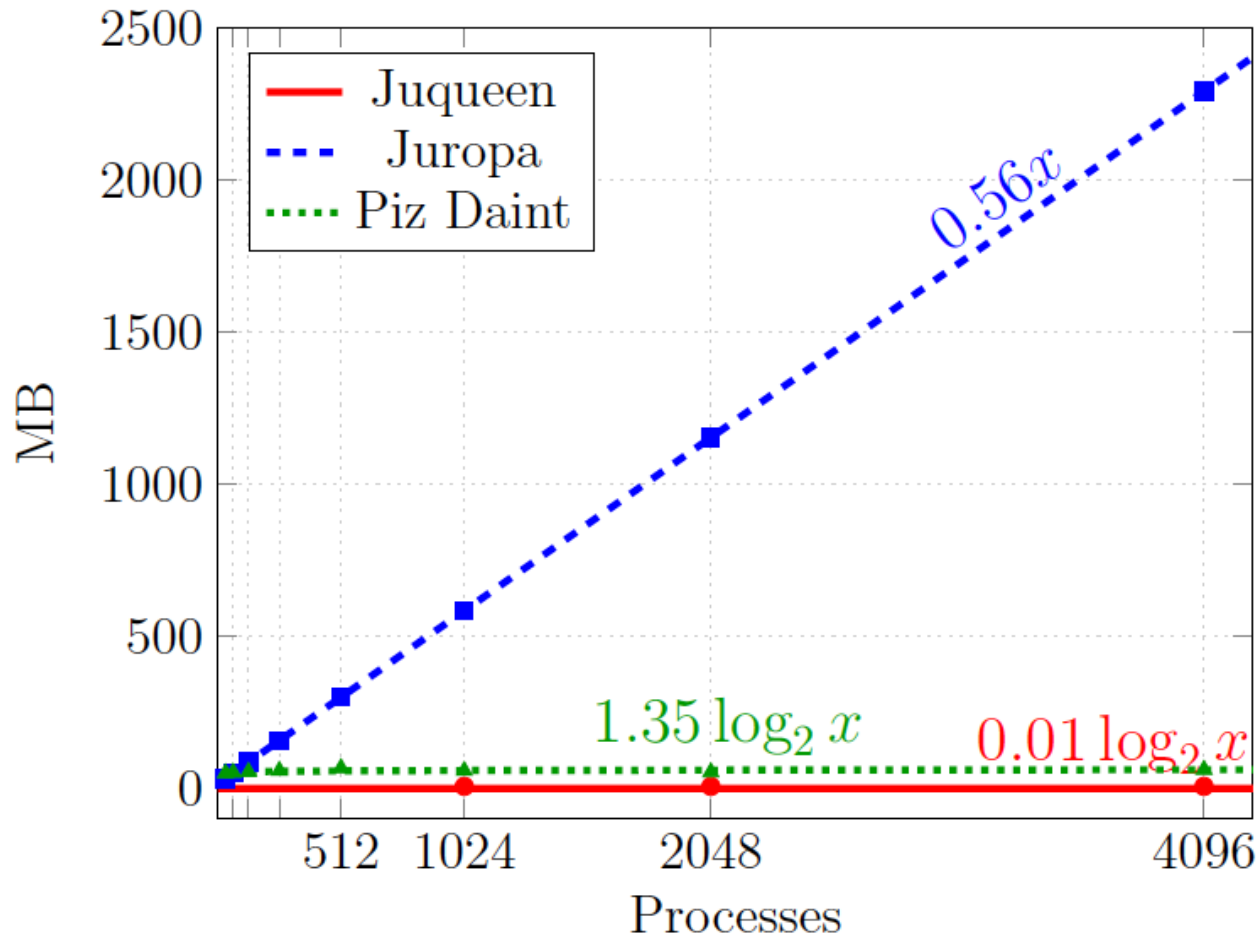
Overhead: Mantevo

Use-case A: automatic testing (Allreduce time)



- Divergence on Piz
- Daint is $O(p^{0.67})$, the highest of all three

Use-case B: automatic testing (MPI memory size)



- Linear memory consumption on Juropa
- ParaStation MPI
- uses RC over IB



Performance Analysis 2.0 – Automatic Models

- Is feasible
Still a long way to go ...
- Offers insight
- Requires low effort
- Improves code coverage



Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich



German Research School
for Simulation Sciences

A. Calotoiu, T. Hoefler, M. Poke, F. Wolf: Using Automated Performance Modeling to Find Scalability Bugs in Complex Codes. *Supercomputing (SC13)*.




T. Hoefler, G. Kwasniewski: Automatic Complexity Analysis of Explicitly Parallel Programs. *SPAA 2014*.

A. Bhattacharyya, T. Hoefler: PEMOGEN: Automatic Adaptive Performance Modeling during Program Runtime, *PACT 2014*

S. Shudler, A. Calotoiu, T. Hoefler, A. Strube, F. Wolf: Exascalating Your Library: Will Your Implementation Meet Your Expectations? *ICS 2015*



SPAA 2014

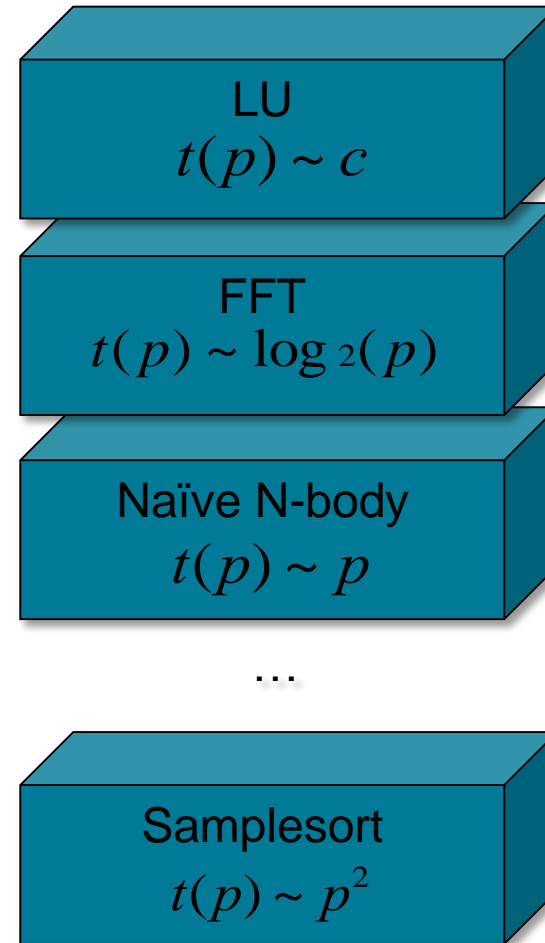
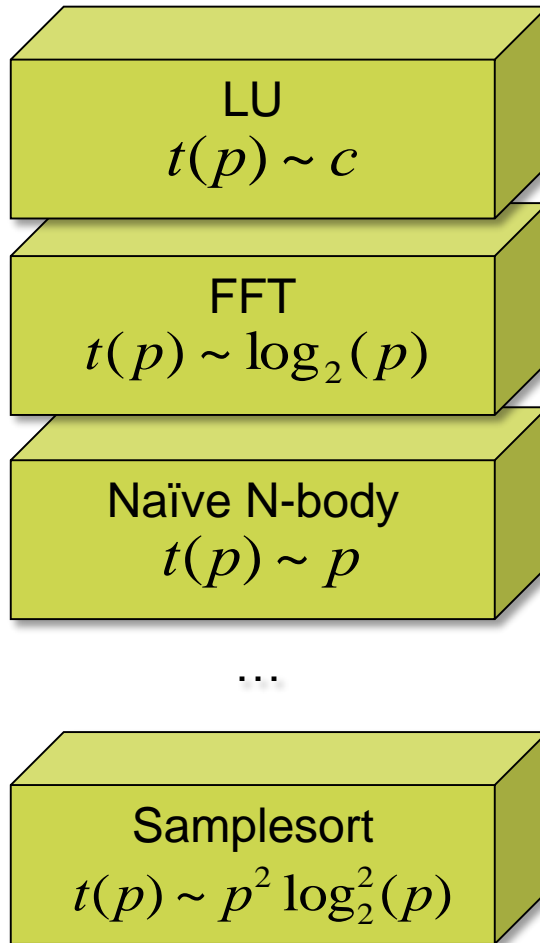


ICS

Backup

How to mechanize the expert? → Survey!

Computation



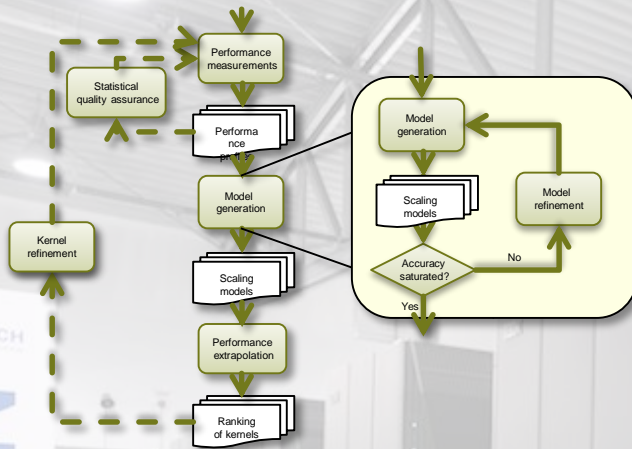
Communication



JÜLICH
JUQUEEN

IBM
supercomputer Blue Gene Q

Evaluation overview




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


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

$$n = 5$$


Sweep3D




MILC



HOMME




XNS



Sweep3D communication performance


- Solves neutron transport problem
- 3D domain mapped onto 2D process grid
- Parallelism achieved through pipelined wave-front process

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



- **LogGP model for communication developed by Hoisie et al.**
 - We assume $p = p_x \cdot p_y \rightarrow$ Equation (6) in [1]

Sweep3D communication performance




Kernel [2 of 40]	Runtime[%] $p_i=262k$	Model [s] $t = f(p)$	Predictive error [%] $p_i=262k$
sweep → MPI_Recv	65.35	$4.03\sqrt{p}$	5.10
sweep	20.87	582.19	01

#bytes = const.
#msg = const.

$p_i \in 8k$

MILC

- MILC/su3_rmd – from MILC suite of QCD codes with performance model manually created
- Time per process should remain constant except for a rather small logarithmic term caused by global convergence checks




Kernel [3 of 479]	Model [s] $t=f(p)$	Predictive Error [%] $p_i=64k$
compute_gen_staple_field	2.40×10^{-2}	0.43
g_vecdoublesum \rightarrow MPI_Allreduce	$6.30 \times 10^{-6} \times \log_2^2(p)$	0.01
mult_adj_su3_fieldlink_lathwec	3.80×10^{-3}	0.04

$$p_i \in 16k$$

HOMME

- Core of the Community Atmospheric Model (CAM)
- Spectral element dynamical core on a cubed sphere grid




Kernel [3 of 194]	Model [s] $t = f(p)$	Predictive error [%] $p_t = 130k$
box_rearrange → MPI_Reduce	$0.026 + 2.53 \times 10^{-6} p \times \sqrt{p} + 1.24 \times 10^{-12} p^3$	57.02
vlaplace_sphere_vk	49.53	99.32
compute_and_apply_rhs	48.68	1.65

$$p_i \text{ } \pounds \text{ } 15k$$

HOMME (2)


- Core of the Community Atmospheric Model (CAM)
- Spectral element dynamical core on a cubed sphere grid



Kernel [3 of 194]	Model [s] $t = f(p)$	Predictive error [%] $p_t = 130k$
box_rearrange → MPI_Reduce	$3.63 \times 10^{-6} p \times \sqrt{p} + 7.21 \times 10^{-13} p^3$	30.34
vlaplace_sphere_vk	$24.44 + 2.26 \times 10^{-7} p^2$	4.28
compute_and_apply_rhs	49.09	0.83

$$p_i \text{ } \underline{\text{£ 43k}}$$

HOMME (3)



What about strong scaling?

- **Wall-clock time not necessarily monotonically increasing – harder to capture model automatically**
 - Different invariants require different reductions across processes


	Weak scaling	Strong scaling
Invariant	Problem size per process	Overall problem size
Model target	Wall-clock time	Accumulated time
Reduction	Maximum / average	Sum

- **Superlinear speedup through cache effects**
 - Measure and model re-use distance?

XNS

- **Finite element flow simulation program with numerous equations represented:**

- Advection diffusion
- Navier-Stokes
- Shallow water




- **Strong scaling analysis**

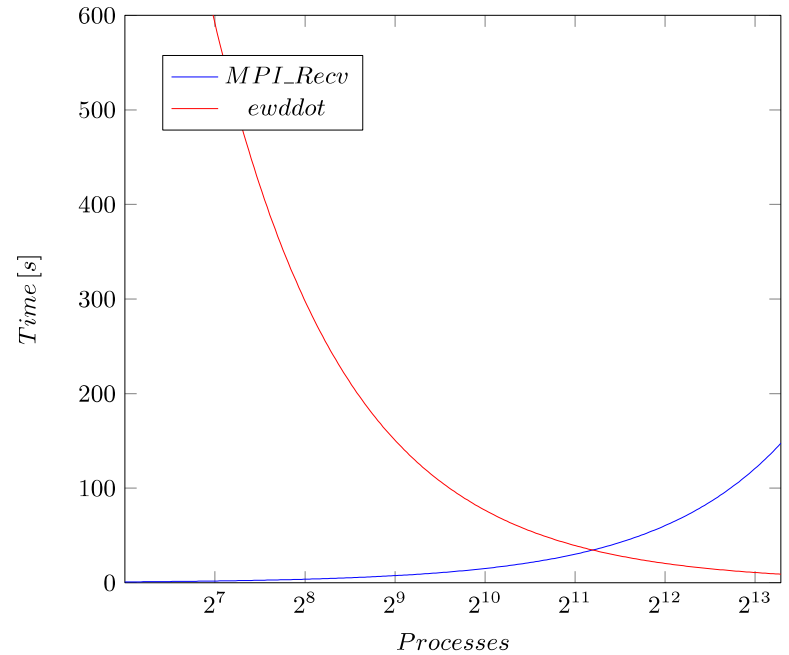
- $P = \{128; \dots; 4,096\}$
- 5 measurements per p_i
- Using accumulated time across processes as metric

XNS (2)

Accumulated time



Wallclock time



Kernel	Runtime[%] p=128	Runtime[%] p=4,096	Model [s] t = f(p)
ewdgennprm->MPI_Recv	0.46	51.46	$0.029 \times p^2$
ewddot	44.78	5.04	$p \times \log(p)$

#bytes = ~p
#msg = ~p

Step back – what do we really care about?

▪ Work


$$W = T_1$$

▪ Depth

$$D = T_\infty$$


▪ Parallel efficiency

$$E_p = \frac{T_1}{pT_p}$$



Related work: counting loop iterations

- Polyhedral model




- piplib
- PolyLib
- PPL
- Polly
- ...

Related work: counting loop iterations

- Polyhedral model

```

for (j = 1; j <= n; j = j*2)
  for (k = j; k <= n; k = k++)
    veryComplicatedOperation(j, k);
  
```



$$j \in [1, n]$$

$$k \in [j, n]$$

$$N = (n+1) \log_2 n - n + 2$$


$$N = \frac{n(n+1)}{2}$$

Related work: counting loop iterations

- When the polyhedral model cannot handle it

```
j=10;  
k=10;  
while (j>0) {  
    j=j+k;  
    k--;  
}
```

?



Algorithm in details

Closed form representation of a loop

- Single affine statement

$$x = Lx + p$$

- Counting function

$$n(x_0)$$

$$x = x_0;$$

$$\text{while}(c^T x < g)$$

$$x = Ax + b;$$

Example
 $x(i, x_0) = A^i \cdot x_0 + p(i)$

$$x = \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix} x + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$

$$n(x_0) = \arg \min_{\{x \in \mathbb{R}^2 \mid c^T \cdot x(i, x_0) \geq g\}}$$

$$x = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} x + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$

}

$$x(i, x_0) = A^i x_0 + \sum_{j=0}^{i-1} A^j \cdot b$$

$$x(i, x_0) = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}^i x_0 + \sum_{j=0}^{i-1} \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}^j \begin{pmatrix} 0 \\ 0 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ i & 1 \end{pmatrix} x_0 + \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

$$n(x_0) = \left\lceil \frac{m - k_0}{j_0} \right\rceil$$

Algorithm in details

Folding the loops

$$x = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} x + \begin{pmatrix} 1 \\ 0 \end{pmatrix};$$
$$\text{while}((1 \ 0)x < n/p) \{$$
$$x = \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix} x + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$
$$\text{while}((0 \ 1)x < m) \{$$
$$x = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} x + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$
$$\} x = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix} x + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$
$$\}$$

Algorithm in details

Folding the loops

$$x = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} x + \begin{pmatrix} 1 \\ 0 \end{pmatrix};$$

while((1 0)x < n/p){

$$x = \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix} x + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$

while((0 1)x < m){

$$x = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} x + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$

$$} x = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix} x + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$

}

$$x = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} x + \begin{pmatrix} 1 \\ 0 \end{pmatrix};$$

while((1 0)x < n/p){

$$x = \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix} x + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$

$$x = \begin{pmatrix} 1 & 0 \\ i & 1 \end{pmatrix} x + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$

$$x = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix} x + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$

}



Algorithm in details

Folding the loops

$$x = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} x + \begin{pmatrix} 1 \\ 0 \end{pmatrix};$$

while((1 0)x < n/p){

$$x = \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix} x + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$

while((0 1)x < m){

$$x = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} x + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$

$$\} x = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix} x + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$

}



$$x = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} x + \begin{pmatrix} 1 \\ 0 \end{pmatrix};$$


while((1 0)x < n/p){

$$x = \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix} x + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$

$$x = \begin{pmatrix} 1 & 0 \\ i & 1 \end{pmatrix} x + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$

$$x = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix} x + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$

}



$$x = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} x + \begin{pmatrix} 1 \\ 0 \end{pmatrix};$$

while((1 0)x < n/p){

$$x = \begin{pmatrix} 2 & 0 \\ i+1 & 0 \end{pmatrix} x + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$

}

Algorithm in details

Starting conditions

```
 $x_{0,1} \longrightarrow x = x_0;$   
     $while(c_1^T x < g_1)\{$   
 $x_{0,2} \longrightarrow x = A_1 x + b_1;$   
     $while(c_2^T x < g_2)\{$   
 $x_{0,3} \longrightarrow x = A_2 x + b_2;$   
     $while(c_3^T x < g_3)\{$   
         $x = A_3 x + b_3;$   
     $\}x = U_2 x + v_2;$   
     $\}x = U_1 x + v_1;$   
     $\}$ 
```

Algorithm in details

Counting the number of iterations

We have:

Algorithm in details

Counting the number of iterations

We have:

- The closed form for each loop:
 - *Single affine statement*
 - *Counting function*
- Starting condition for each loop

Algorithm in details

Counting the number of iterations

We have:

- The closed form for each loop:
 - *Single affine statement*
 - *Counting function*
- Starting condition for each loop

Number of iterations:

$$N = \sum_{i_1=0}^{n_1(x_{0,1})} \sum_{i_2=0}^{n_2(x_{0,2})} \dots \sum_{i_{r-1}=0}^{n_{r-1}(x_{0,r-1})} n_r(x_{0,r}).$$

Algorithm in details

Counting the number of iterations

- The equation computes the precise number of iterations

$$N = \sum_{i_1=0}^{n_1(x_{0,1})} \sum_{i_2=0}^{n_2(x_{0,2})} \dots \sum_{i_{r-1}=0}^{n_{r-1}(x_{0,r-1})} n_r(x_{0,r}).$$


Algorithm in details

Counting the number of iterations

- The equation gives precise number of iterations

$$N = \sum_{i_1=0}^{n_1(x_{0,1})} \sum_{i_2=0}^{n_2(x_{0,2})} \dots \sum_{i_{r-1}=0}^{n_{r-1}(x_{0,r-1})} n_r(x_{0,r})$$

- But simplification may fail → Sum approximation
 - *Approximate sums by integrals*
→ lower and upper bounds



Solving more general problems

Solving more general problems

- Multipath loops

Solving more general problems

- **Multipath loops**
- **Conditional statements**

Solving more general problems

- Multipath loops
- Conditional statements
- Non-affine loops

```
do j=1, lastrow-firstrow+1
  sum = 0.d0
```

$$\text{lastrow-firstrow+1} = \text{row_size} = \frac{\text{na}}{\text{nprows}}$$

```
  do k=rowstr(j), rowstr(j+1)-1
    sum = sum + a(k)*p(colidx(k))
```

$$\text{rowstr}(j+1)-1-\text{rowstr}(j) = u$$

```
  enddo
```

```
  w(j) = sum
```

```
enddo
```

$$N = \frac{\text{na} \cdot u}{\text{nprows}}$$

Case studies

- **NAS Parallel Benchmarks: EP**

$$N(m, p) = \left\lceil \frac{2^{m-16} \cdot (u + 2^{16})}{p} \right\rceil$$

```
u:  do i=1,100
      ik =kk/2
      if (ik .eq. 0) goto 130
      kk=ik
      continue
```

Case studies

■ NAS Parallel Benchmarks: EP

$$N(m, p) = \left\lceil \frac{2^{m-16} \cdot (u + 2^{16})}{p} \right\rceil$$

```

u:  do i=1,100
      ik =kk/2
      if (ik .eq. 0) goto 130
      kk=ik
      continue
  
```


$$W = T_1 \approx 2^m$$

$$D = T_\infty \approx 1$$

$$E_P = \frac{2^m}{p \left\lceil \frac{2^m}{p} \right\rceil}$$

$$E_P \approx 1 \text{ if } p \leq 2^m$$

$$E_P \approx 2^m / p \text{ if } p > 2^m$$



Case studies


CG – conjugate gradient

$$W \approx k_1 \left\lceil \frac{m}{p} \right\rceil + k_2 \sqrt{\left\lceil \frac{m}{p} \right\rceil} + k_3 \log_2(\sqrt{p})$$

IS – integer sort

$$D = T_\infty = \infty \left(3(b+t) + 2 \left\lceil \frac{m}{p} \right\rceil + p + u_1 + u_2 \right)$$

$$E_p = \frac{D = T_\infty = \infty}{k_4} \frac{1}{p \left(k_1 \left\lceil \frac{m}{p} \right\rceil + k_2 \sqrt{\left\lceil \frac{m}{p} \right\rceil} + k_3 \log_2(\sqrt{p}) \right)}$$



Counting Arbitrary Affline Loop Nests

- Why affine loops?
 - Closed form representation of the loop

```
x=x0;           // Initial assignment  
while(cTx < g)  // Loop guard  
  x=Ax + b;     // Loop update
```



$$x(i, x_0) = L(i) \cdot x_0 + p(i)$$

$$n(x_0, c, g) = \arg \min_d (c^T \cdot x(d, x_0) \geq g)$$

Counting Arbitrary Affline Loop Nests

Why affine loops?

- Closed form representation of the loop

```

x=x0;           // Initial assignment
while (cTx < g) // Loop guard
  x=Ax + b;      // Loop update

```



$$x(i, x_0) = L(i) \cdot x_0 + p(i)$$

$$n(x_0, c, g) = \arg \min_d (c^T \cdot x(d, x_0) \geq g)$$

Example

```

for ( k=j; k < m; k = k + j )
  veryComplicatedOperation(j, k);

```

$$x = \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix} x + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$

```
while ((0 1)x < m) {
```

$$x = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} x + \begin{pmatrix} 0 \\ 0 \end{pmatrix};$$

```
}
```



$$x(i, x_0) = \begin{pmatrix} 1 & 0 \\ i & 1 \end{pmatrix} x_0 + \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

$$n(x_0) = \left\lceil \frac{m - k_0}{j_0} \right\rceil$$

where $x_0 = \begin{pmatrix} j_0 \\ k_0 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix} x + \begin{pmatrix} 0 \\ 0 \end{pmatrix}$

Loops

- **Multipath affine loops**

```
x=1;
while(x < n/p + 1) {
  y=x;
  while(y < m) { S1; y=2*y; }
  z=x;
  while(z < m) { S2; z= z + x; }
  x=2*x;
}
```
