

第10回先進スーパーコンピューティング環境研究会（ASE研究会）実施報告

東京大学情報基盤センター 特任准教授 片桐孝洋

2011年10月13日（木）15時00分～18時30分、東京大学情報基盤センター4階遠隔講義室にて、第10回先進スーパーコンピューティング環境研究会（ASE研究会）が開催されました。

国内の大学・研究機関および企業からの参加者が13名あり、活発な議論がなされました。

招待講演として、ETH ZurichからMarkus Püschel教授をお呼びし、自動チューニングと機械学習に関する講演を行いました。また、工学院大学の藤井昭宏博士、筑波大学の山崎育朗様による、数値計算アルゴリズムにおける自動チューニングの講演を行いました。

Püschel教授の講演は、「Automatic Performance Tuning and Machine Learning」と題し、自動チューニングの重要性などの背景に始まり、自動チューニングへ機械学習を適用した事例紹介に至る大変興味深いご講演でした。特に、自動チューニングを行うことは「探索」処理に位置づけられることが多いのですが、Püschel教授の講演では、自動チューニングはコンパイラやプログラマを助けるためのツールが本来の意味であり探索はコスト高である、ですから機械学習を適用する解を示すという主張は明快であり、わかりやすいものでした。Püschel教授らが行ってきた信号処理(FFT)の自動チューニングに関するプロジェクトSpiralにおける自動チューニングの位置づけも大変興味深いものでした。参加者間で、たいへん活発な議論がなされました。

以上の講演内容について、本原稿の後に当日の発表資料を掲載します。ご興味のある方はご覧ください。当日のプログラムを以下に載せます。

Program

- 15:00-15:05 Introduction by Takahiro Katagiri (The University of Tokyo)
- 15:05 - 16:05

Invited Speaker: Professor Markus Püschel (Computer Science, ETH Zurich, Switzerland)

Title: Automatic Performance Tuning and Machine Learning

Abstract: Automatic performance tuning has emerged as a paradigm complementing traditional compilers to port software and performance between platforms. Several techniques have proven useful including adaptive libraries, program generation, domain-specific languages, and models. However, one technique is shared by almost all approaches: search for the fastest among a set of alternative implementations. Typically the search space is huge and hence the search is costly. This may be bearable in offline tuning (e.g., ATLAS) that is performed during installation but becomes cumbersome in online tuning (e.g., FFTW) that is performed at runtime since the input

size is required. We argue that machine learning, which has already been studied and used in the compiler community, can solve this problem and should be added to the portfolio of performance tuning tools. As example we show a successful approach to automatically convert Spiral-generated online-tunable transform libraries into offline-tunable ones.

- 16:05–16:15 Break
- 16:15 – 16:45

Speaker: Akihiro Fujii (Kogakuin University), Osamu Nakamura (Sumitomo Metal Industries)

Title: Automatic tuning for Algebraic Multigrid solver for Fluid analysis

Abstract: This talk presents an online automatic tuning method of Algebraic Multigrid (AMG) solver for fluid analysis based on SMAC method. This type of fluid analysis requires to solve Pressure Poisson equation every time step. As time steps forward, problem matrix does not change, but the right hand side vector changes gradually in many cases. To shorten the total time the solver uses, we try to optimize AMG solver by selecting its parameter setting for each time step. In this talk, we consider automatic tuning technique of AMG solver parameters such as smoothers, multigrid cycles, acceleration coefficient of the smoother and others. We studied the auto tuning method which narrows the search domain by determining parameters in a step-by-step manner based on typical property of the AMG solver parameters. In addition, optimizing parameters is done every constant number of time steps such as 100 time steps to reduce the overhead of measuring the efficiency of various parameter settings. Our AMG library with online auto-tuning mechanism which determines appropriate solver parameter set among 900 parameter settings improved the performance of AMG solver with default setting up to 20 percent in our numerical tests.

- 16:45 – 17:15

Speaker: Ikuro Yamazaki, Hiroto Tadano, Tetsuya Sakurai (Graduate School of Systems and Information Engineering, University of Tsukuba)

Title: A parameter selection for a preconditioner using a cutoff for Krylov subspace methods

Abstract: Large linear systems involving a semi-sparse matrix which has relatively large number of nonzero elements appear in nano-science simulations. Preconditioners with a cutoff is proposed as suitable methods for such semi-sparse linear systems. The performance of these preconditioners is highly dependent on a cutoff parameter. A smaller value of cutoff increases a computational cost in preconditioner construction. Meanwhile, a larger value of cutoff leads to a less effective

preconditioning matrix with a large number of iterations. Hence, the selection of an appropriate cutoff parameter is important. In this talk, we present a strategy to find an efficient cutoff parameter for the preconditioners using a cutoff. We estimate an appropriate cutoff parameter using residual norms before applying an iterative solver, and verify the validity of our strategy by numerical experiments.

- 17:15 – 17:30 Break
- 17:30 – 18: 00

Speaker: Satoshi Ohshima (The University of Tokyo)

Title: Implementation of 3D FEM program on GPU

Abstract: GPU is now utilizing for several scientific applications because GPU has high calculation performance and memory transfer performance. FEM (Finite Element Method) is one of the applications expected to accelerate with GPU. This talk shows the implementation of 3D FEM program on CUDA GPU. The main targets of acceleration are sparse matrix solver (CG method) and matrix assembly. Our GPU implementation has obtained better performance than multi-core CPU. Also we are now trying multi GPUs implementation.

- 18:00 – 18:30

Speaker: Masae Hayashi (The University of Tokyo)

Title: OpenMP/MPI Hybrid Parallel FEM Based on Extended Hierarchical Interface Decomposition for Multi-core Clusters.

Abstract: ILU preconditioner is a powerful and popular preconditioning method for Krylov iterative solvers on sparse matrices derived from FEM applications. Block-Jacobi-type localized ILU preconditioner is generally employed in parallel computation basing on domain decomposition. The localization of ILU process is good for parallel efficiency but decreases the effectiveness of preconditioning since it neglects the effect from out side the domain. Hierarchical Interface Decomposition (HID) and proposed extended version of HID, where additional layers of separators are introduced for robust and efficient computation is a robust and efficient parallel preconditioning method. We are developing preconditioning methods using OpenMP/MPI hybrid parallel programming models on multicore/multisocket clusters. HID and extended version of HID (ExHID) is applied to both of inter-node parallelization with message-passing (e.g. MPI) part and intra-node parallelization with multi-threading (e.g. OpenMP). We implement the OpenMP/MPI hybrid parallel programing model to FEM application solving 3- dimensional linear elasticity problem. The developed code has been tested on the T2K Open Super Computer (T2K/Tokyo) using up to 8 nodes, 16 cores. We report the scalability and the robustness resulted from numerical experiments of

parallel ILU preconditioner with fill-ins and iterative solver and we show the developed code provides better performance compared to that of multicoloring method used for intra-node parallelization.

- 18: 30 Closing Remarks by Takahiro Katagiri (The University of Tokyo)
 - 19:00– Banquet near Nedu station
-



図1 当日の会場の様子

ASE 研究会の開催情報はメーリングリストで発信をしております。研究会メーリングリストに参加ご希望の方は、ASE 研究会幹事の片桐 (katagiri@cc.u-tokyo.ac.jp) までお知らせください。

以上

Automatic Performance Tuning and Machine Learning

Markus Püschel
Computer Science, ETH Zürich

with:

Frédéric de Mesmay
PhD, Electrical and Computer Engineering, Carnegie Mellon

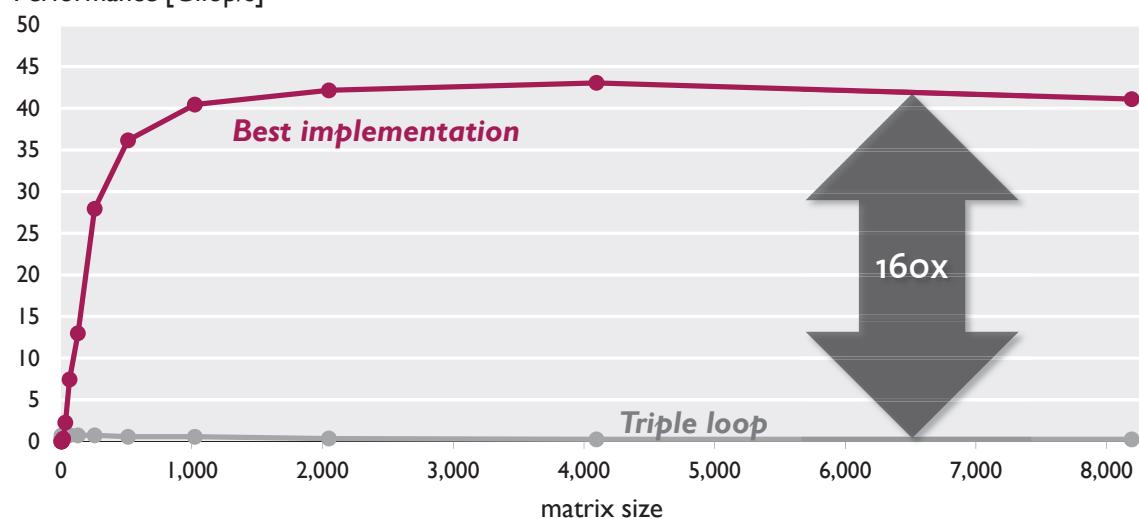


Markus Püschel, ETH Zürich, 2011

Why Autotuning?

Matrix-Matrix Multiplication (MMM) on quadcore Intel platform

Performance [Gflop/s]

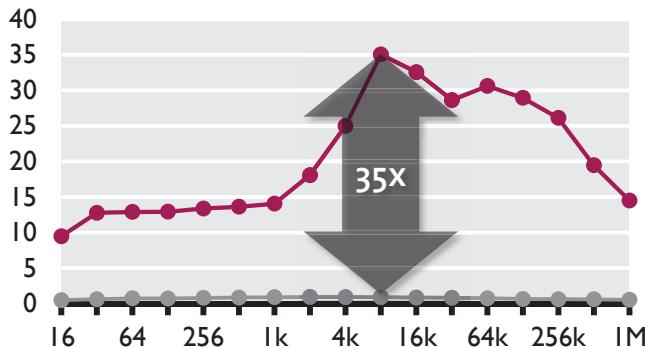


- Same (mathematical) operation count ($2n^3$)
- Compiler underperforms by 160x

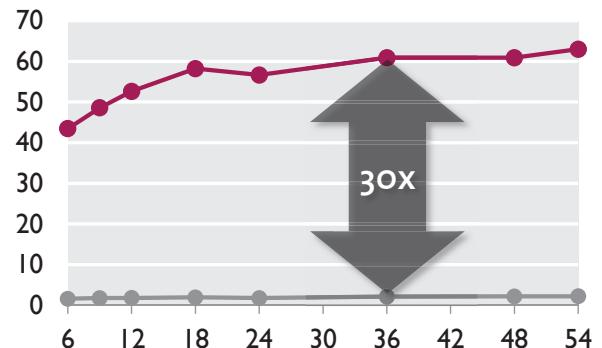
Same for All Critical Compute Functions

DFT auf Intel Core i7

Leistung [Gflop/s]

**WiFi Empfänger (1 Intel Core)**

Leistung [Mbit/s]



Solution: Autotuning

~~**Definition:** Search over alternative implementations or parameters to find the fastest~~

Definition: Automating performance optimization with tools that complement/aid the compiler or programmer

However: Search is an important tool. But expensive.

Solution: Machine learning

Organization

- Autotuning examples
- An example use of machine learning

No search

Markus Püschel, ETH Zürich, 2011

Search

No search

No search

No search

Search

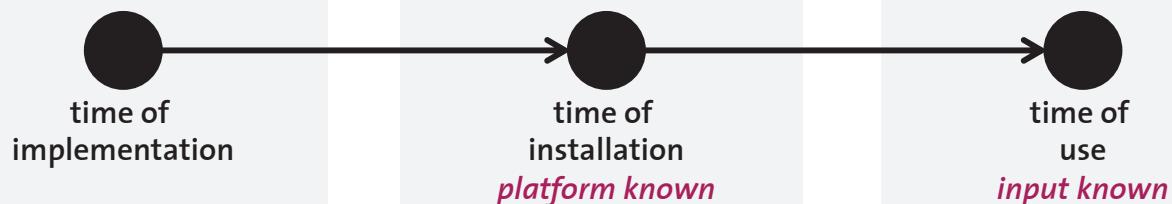
Search

Search

time of
implementation

time of
installation
platform known

time of
use
input known



PhiPac/ATLAS: MMM Generator

Whaley, Bilmes, Demmel, Dongarra, ...

```
// MMM loop-nest
for i = 0:NB:N-1
    for j = 0:NB:M-1
        for k = 0:NB:K-1
```

- *ijk or jik depending on N and M*
- *Blocking for cache*

```
// mini-MMM loop nest
for i' = i:MU:i+NB-1
    for j' = j:NU:j+NB-1
        for k' = k:KU:k+NB-1
```

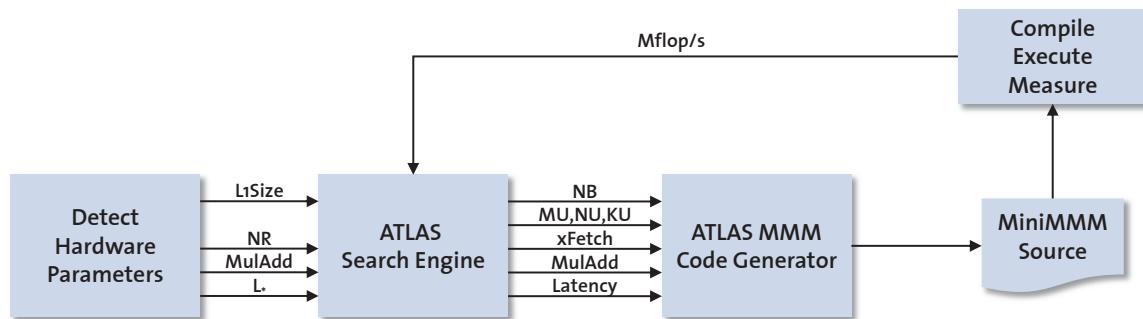
- *Blocking for registers*

```
// micro-MMM loop nest
for k'' = k':1:k'+KU-1
    for i'' = i':1:i'+MU-1
        for j'' = j':1:j'+NU-1
```

- *Unrolling*
- *Scalar replacement*
- *Add/mult interleaving*
- *Skewing*

Search parameters: N_B, M_U, N_U, K_U, L_S, ...

PhiPac/ATLAS: MMM Generator



source: Pingali, Yotov, Cornell U.

No search

Markus Püschel, ETH Zürich, 2011

Search

ATLAS
MMM generator

time of
implementation

time of
installation
platform known

time of
use
input known

FFTW: Discrete Fourier Transform (DFT)

Frigo, Johnson

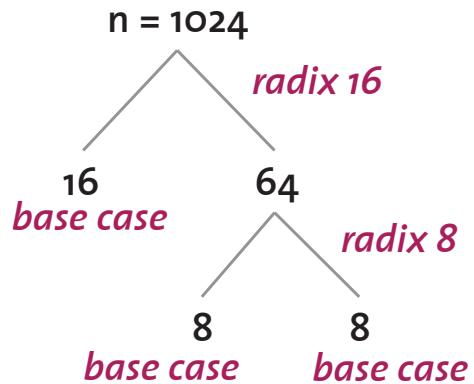
Installation

configure/make

Usage

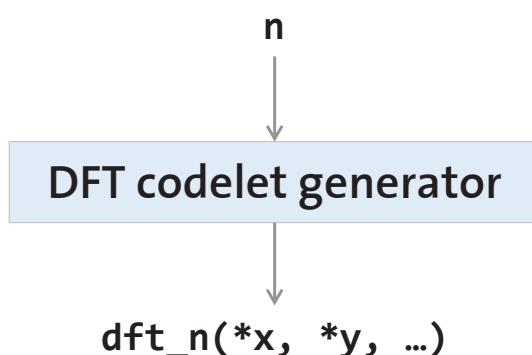
```
d = dft(n)
d(x,y)
```

Twiddles
Search for fastest computation strategy



FFTW: Codelet Generator

Frigo



*fixed size DFT function
straightline code*

FFTW codelet
generator

ATLAS
MMM generator

FFTW adaptive
library

time of
implementation

time of
installation
platform known

time of
use
input known

OSKI: Sparse Matrix-Vector Multiplication

Vuduc, Im, Yelick, Demmel

$$\begin{array}{c|c|c} & = & \\ \hline & \text{Sparse Matrix} & * \\ \hline & = & \end{array}$$

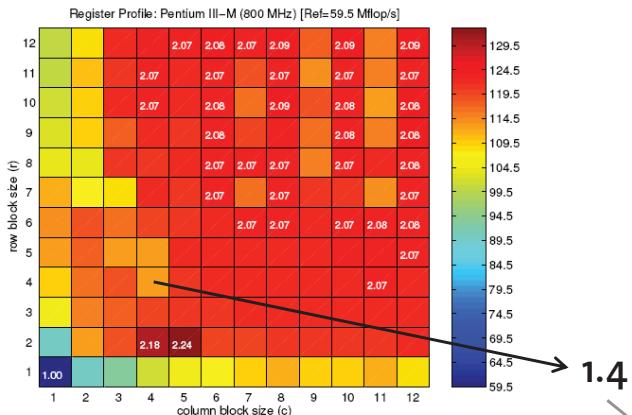
■ Blocking for registers:

- Improves locality (reuse of input vector)
- But creates overhead (zeros in block)

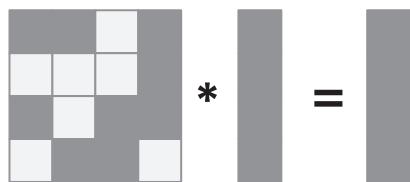
$$\begin{array}{c|c|c} \text{Block-diagonal Matrix} & * & \text{Vector} \\ \hline & = & \end{array}$$

OSKI: Sparse Matrix-Vector Multiplication

Gain by blocking (dense MVM)



Overhead by blocking



$$16/9 = 1.77$$

$$1.4/1.77 = 0.79 \text{ (no gain)}$$

No search

Markus Püschel, ETH Zürich, 2011

Search

FFTW codelet generator

OSKI
sparse MVM

OSKI
sparse MVM

FFTW adaptive library

time of implementation

time of installation
platform known

time of use
input known

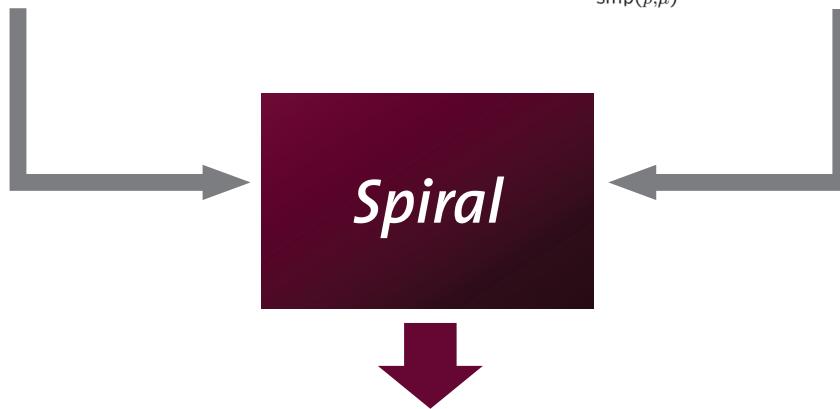
Spiral: Linear Transforms & More

Algorithm knowledge

$$\begin{aligned} \text{DFT}_n &\rightarrow P_{k/2,2m}^\top (\text{DFT}_{2m} \oplus (I_{k/2-1} \otimes_i C_{2m} \text{rDFT}_{2m}(i/k))) (\text{RDFT}'_k \otimes I_m) \\ |\text{rDFT}_{2n}(u)| &\rightarrow L_m^{2n} \left(I_k \otimes_i \left| \begin{array}{l} \text{rDFT}_{2m}((i+u)/k) \\ \text{rDHT}_{2m}((i+u)/k) \end{array} \right| \right) \left(\left| \begin{array}{l} \text{rDFT}_{2k}(u) \\ \text{rDHT}_{2k}(u) \end{array} \right| \otimes I_m \right) \\ \text{RDFT-3}_n &\rightarrow (Q_{k/2,m}^\top \otimes I_2) (I_k \otimes_i \text{rDFT}_{2m})(i+1/2)/k) (\text{RDFT-3}_k \otimes I_m) \end{aligned}$$

Platform description

$$\begin{aligned} \underbrace{A_m \otimes I_n}_{\text{smp}(p,\mu)} &\rightarrow \underbrace{\left(L_m^{mp} \otimes I_{n/p} \right) \left(I_p \otimes (A_m \otimes I_{n/p}) \right) \left(L_p^{mp} \otimes I_{n/p} \right)}_{\text{smp}(p,\mu)} \\ \underbrace{I_m \otimes A_n}_{\text{smp}(p,\mu)} &\rightarrow I_p \otimes \parallel \left(I_{m/p} \otimes A_n \right) \\ \underbrace{(P \otimes I_n)}_{\text{smp}(p,\mu)} &\rightarrow (P \otimes I_{n/\mu}) \overline{\otimes} I_\mu \end{aligned}$$



Optimized implementation
regenerated for every new platform

Program Generation in Spiral (Sketched)

Transform
user specified

Fast algorithm
in SPL
many choices

DFT₈

Algorithm rules



$$(DFT_2 \otimes I_4) T_4^8 (I_2 \otimes ((DFT_2 \otimes I_2) \cdot T_2^4 (I_2 \otimes DFT_2) L_2^4)) L_2^8$$

Σ -SPL

$$\begin{aligned} \sum (S_j DFT_2 G_j) \sum &\left(\sum (S_{k,l} \text{diag}(t_{k,l}) DFT_2 G_l) \right. \\ &\left. \sum (S_m \text{diag}(t_m) DFT_2 G_{k,m}) \right) \end{aligned}$$



Optimized implementation

Optimization at all abstraction levels



parallelization
vectorization



loop
optimizations



constant folding
scheduling
.....

+ search

Machine learning

Machine learning

Spiral: transforms
general input size

Spiral: transforms
fixed input size

Spiral: transforms
general input size

OSKI
sparse MVM

OSKI
sparse MVM

FFTW codelet
generator

ATLAS
MMM generator

FFTW adaptive
library

time of
implementation

time of
installation
platform known

time of
use
input known

Organization

- Autotuning examples
- An example use of machine learning

Online tuning (time of use)

Installation

configure/make

Use

$d = \text{dft}(n)$
 $d(x, y)$

Twiddles

Search for fastest computation strategy

Offline tuning (time of installation)

Installation

configure/make

for a few n : search
learn decision trees

Use

$d = \text{dft}(n)$
 $d(x, y)$

Goal

Integration with Spiral-Generated Libraries

Voronenko 2008

$(\text{DFT}_k \otimes \text{I}_m) \text{T}_m^n (\text{I}_k \otimes \text{DFT}_m) \text{L}_k^n$
+ some platform information → **Spiral** → Online tunable library

$$\begin{aligned}
& \text{DFT}_n \rightarrow P_{k/2,2m}^\top (\text{DFT}_{2m} \oplus (I_{k/2-1} \text{C}_{2m} \text{rDFT}_{2m}(i/k))) (\text{RDFT}'_k \text{I}_m), \quad k \text{ even}, \\
& \begin{cases} \text{RDFT}'_n \\ \text{DHT}'_n \\ \text{DHT}_n \end{cases} \rightarrow (P_{k/2,2m}^\top \text{I}_2) \left(\begin{cases} \text{RDFT}'_{2m} \\ \text{DHT}'_{2m} \\ \text{DHT}_{2m} \end{cases} \oplus \left(I_{k/2-1} \text{D}_{2m} \begin{cases} \text{rDFT}_{2m}(i/k) \\ \text{rDFT}_{2m}(i/k) \\ \text{rDHT}_{2m}(i/k) \\ \text{rDHT}_{2m}(i/k) \end{cases} \right) \right) \left(\begin{cases} \text{RDFT}'_k \\ \text{RDFT}'_k \\ \text{DHT}'_k \\ \text{DHT}_k \end{cases} \text{I}_m \right), \quad k \text{ even}, \\
& \begin{cases} \text{rDFT}_{2n}(u) \\ \text{rDHT}_{2n}(u) \end{cases} \rightarrow L_m^{2n} \left(I_k \begin{cases} \text{rDFT}_{2m}(i+u/k) \\ \text{rDHT}_{2m}(i+u/k) \end{cases} \right) \left(\begin{cases} \text{rDFT}_{2k}(u) \\ \text{rDHT}_{2k}(u) \end{cases} \text{I}_m \right), \\
& \text{RDFT-3}_n \rightarrow (Q_{k/2,m}^\top \text{I}_2) (I_{k-1} \text{rDFT}_{2m}(i+1/2/k)) (\text{RDFT-3}_k \text{I}_m), \quad k \text{ even}, \\
& \text{DCT-2}_n \rightarrow P_{k/2,2m}^\top (\text{DCT-2}_m K_2^{2m} \oplus (I_{k/2-1} \text{N}_{2m} \text{RDFT-3}_{2m}^\top)) B_n (L_{k/2}^{n/2} \text{I}_2) (I_m \text{RDFT}'_k) Q_{m/2,k}, \\
& \text{DCT-3}_n \rightarrow \text{DCT-2}_n^\top, \\
& \text{DCT-4}_n \rightarrow Q_{k/2,2m}^\top (I_{k/2} \text{N}_{2m} \text{RDFT-3}_m^\top) B'_n (L_{k/2}^{n/2} \text{I}_2) (I_m \text{RDFT-3}_k) Q_{m/2,k}, \\
& \text{DFT}_n \rightarrow (\text{DFT}_k \text{I}_m) \text{T}_m^n (I_k \text{DFT}_m) \text{L}_k^n, \quad n = km \\
& \text{DFT}_n \rightarrow P_n (\text{DFT}_k \text{DFT}_m) Q_n, \quad n = km, \quad \gcd(k, m) = 1 \\
& \text{DFT}_p \rightarrow R_p^\top (I_1 \oplus \text{DFT}_{p-1}) R_p (I_1 \oplus \text{DFT}_{p-1}) R_p, \quad p \text{ prime} \\
& \text{DCT-3}_n \rightarrow (I_m \oplus J_m) \text{L}_m^n (\text{DCT-3}_m(1/4) \oplus \text{DCT-3}_m(3/4)) \\
& \quad \cdot (F_2 \text{I}_m) \begin{bmatrix} I_m & 0 \oplus -J_{m-1} \\ 0 \oplus J_{m-1} & \frac{1}{\sqrt{2}}(I_1 \oplus 2I_m) \end{bmatrix}, \quad n = 2m \\
& \text{DCT-4}_n \rightarrow S_n \text{DCT-2}_n \text{diag}_{0 \leq k < n} (1/(2 \cos((2k+1)\pi/4n))) \\
& \text{IMDCT}_{2m} \rightarrow (J_m \oplus I_m \oplus I_m \oplus J_m) \left(\begin{pmatrix} 1 \\ -1 \end{pmatrix} \text{I}_m \right) \oplus \left(\begin{pmatrix} -1 \\ 1 \end{pmatrix} \text{I}_m \right) J_{2m} \text{DCT-4}_{2m} \\
& \text{WHT}_{2^k} \rightarrow \prod_{i=1}^t (I_{2^{k_1+\dots+k_{i-1}}} \text{WHT}_{2^{k_i}} \text{I}_{2^{k_{i+1}+\dots+k_t}}), \quad k = k_1 + \dots + k_t \\
& \text{DFT}_2 \rightarrow F_2 \\
& \text{DCT-2}_2 \rightarrow \text{diag}(1, 1/\sqrt{2}) F_2 \\
& \text{DCT-4}_2 \rightarrow J_2 R_{13\pi/8}
\end{aligned}$$

Organization

- Autotuning examples
- An example use of machine learning
 - *Anatomy of an adaptive discrete Fourier transform library*
 - *Decision tree generation using C4.5*
 - *Results*

Markus Püschel, ETH Zürich, 2011

Discrete/Fast Fourier Transform

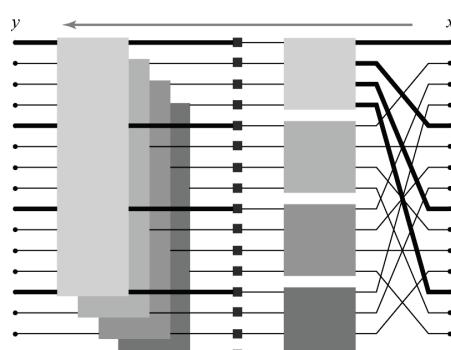
- Discrete Fourier transform (DFT):

$$y = \text{DFT}_n x, \quad \text{DFT}_n = [e^{-2k\ell\pi i/n}]_{0 \leq k, \ell < n}$$

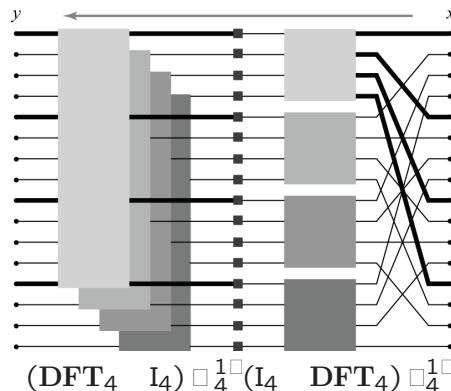
- Cooley/Tukey fast Fourier transform (FFT):

$$\text{DFT}_n = (\text{DFT}_k \otimes \text{I}_m) \text{T}_m^n (\text{I}_k \otimes \text{DFT}_m) \text{L}_k^n, \quad n = km$$

- Dataflow (right to left): $16 = 4 \times 4$



Fast Fourier Transform



$$16 = 4 \times 4$$

```

void dft(int n, cpx *y, cpx *x) {
    if (use_dft_base_case(n))
        dft_bc(n, y, x);
    else {
        int k = choose_dft_radix(n);
        for (int i=0; i < k; ++i)
            dft_strided(m, k, t + m*i, x + m*i);
        for (int i=0; i < m; ++i)
            dft_scaled(k, m, precomp_d[i], y + i, t + i);
    }
}
void dft_strided(int n, int istr, cpx *y, cpx *x) { ... }
void dft_scaled(int n, int str, cpx *d, cpx *y, cpx *x) { ... }

```

Choices used for adaptation

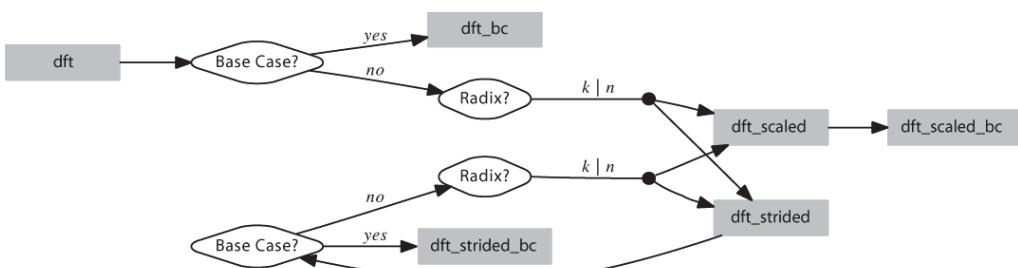
Decision Graph of Library

```

void dft(int n, cpx *y, cpx *x) {
    if (use_dft_base_case(n))
        dft_bc(n, y, x);
    else {
        int k = choose_dft_radix(n);
        for (int i=0; i < k; ++i)
            dft_strided(m, k, t + m*i, x + m*i);
        for (int i=0; i < m; ++i)
            dft_scaled(k, m, precomp_d[i], y + i, t + i);
    }
}
void dft_strided(int n, int istr, cpx *y, cpx *x) { ... }
void dft_scaled(int n, int str, cpx *d, cpx *y, cpx *x) { ... }

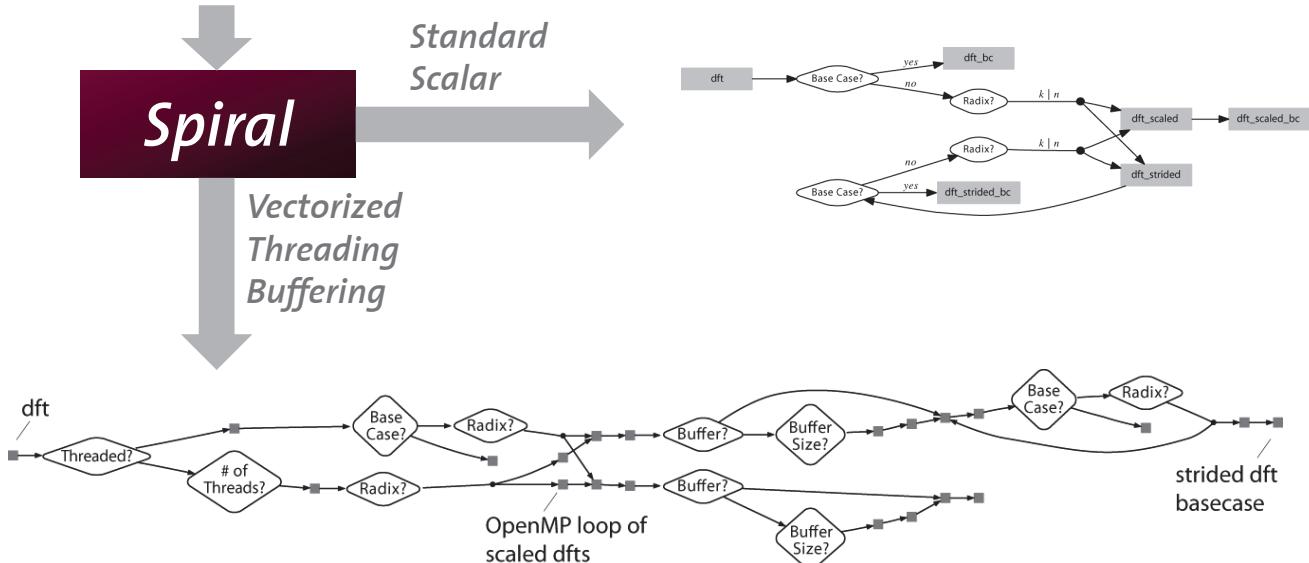
```

Choices used for adaptation



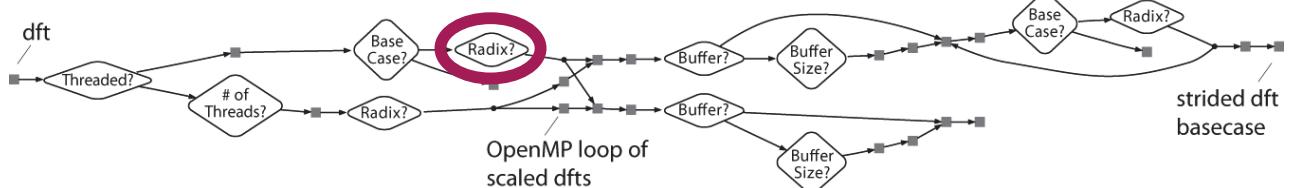
Spiral-Generated Libraries

$$(\text{DFT}_k \otimes \text{I}_m) \top_m^n (\text{I}_k \otimes \text{DFT}_m) \text{L}_k^n$$



- 20 mutually recursive functions
- 10 different choices (occurring recursively)
- Choices are heterogeneous (radix, threading, buffering, ...)

Our Work



Upon installation, generate decision trees for each choice

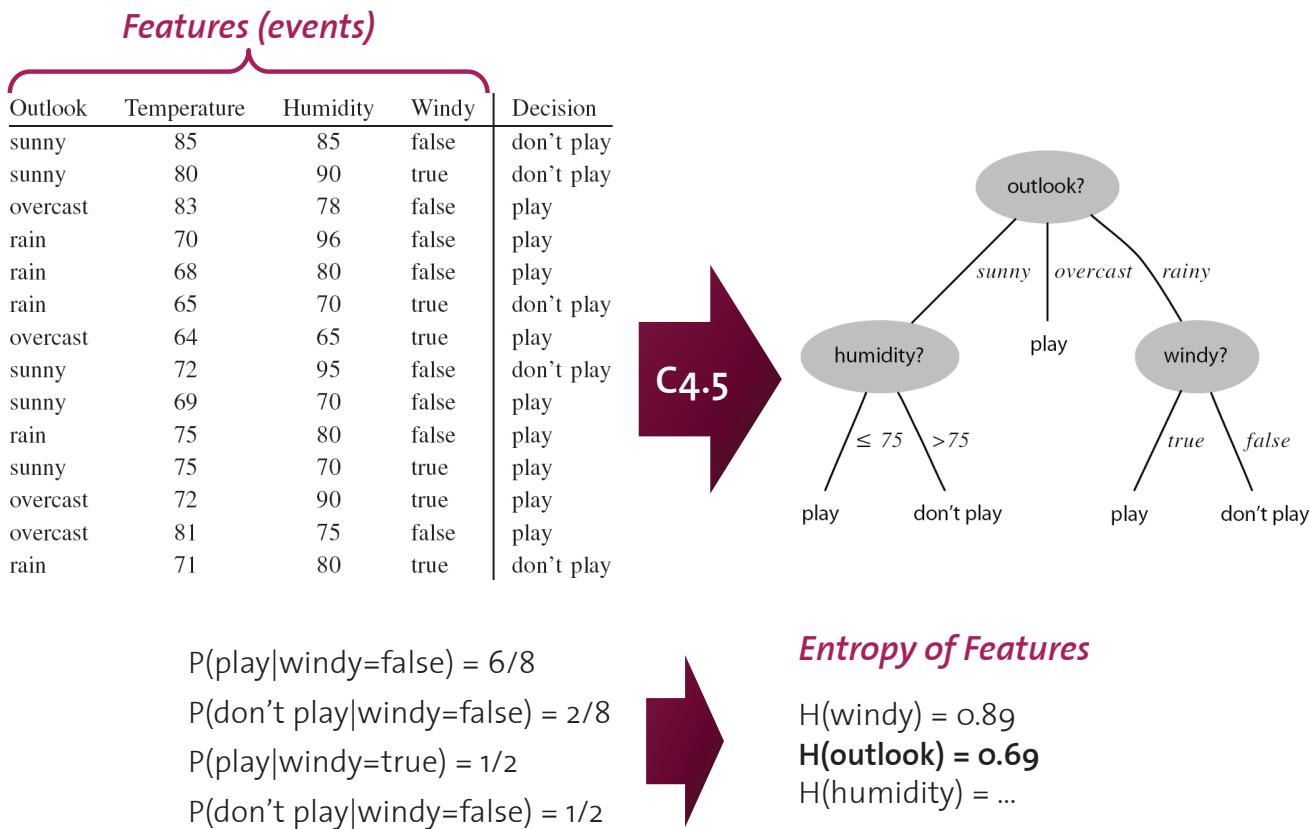
Example:

```

if ( n <= 65536 ) {
    if ( n <= 32 ) {
        if ( n <= 4 ) {return 2;}
        else {return 4;}
    }
    else {
        if ( n <= 1024 ) {
            if ( n <= 256 ) {return 8;}
            else {return 32;}
        }
        else {
            ...
        }
    }
}

```

Statistical Classification: C4.5



Application to Libraries

- Features = arguments of functions (except variable pointers)

```

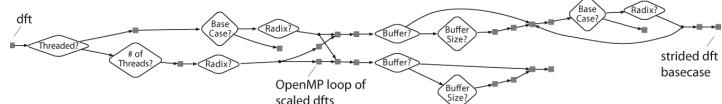
dft(int n, cpx *y, cpx *x)
dft_strided(int n, int istr, cpx *y, cpx *x)
dft_scaled(int n, int str, cpx *d, cpx *y, cpx *x)
  
```

- At installation time:

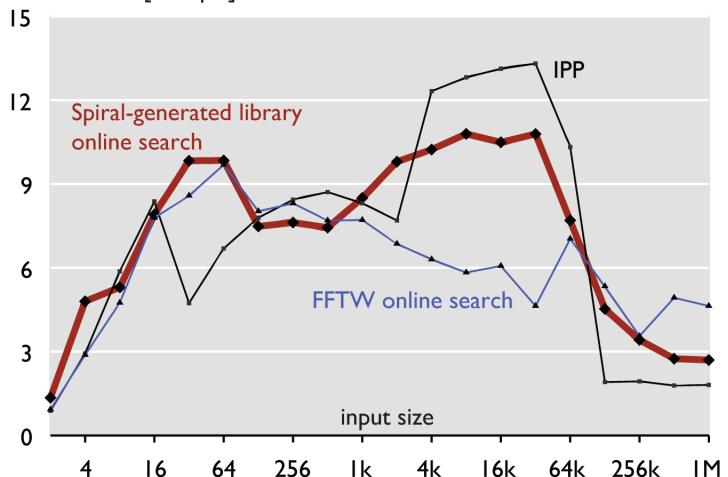
- Run search for a few input sizes n
- Yields training set: features and associated decisions (several for each size)
- Generate decision trees using C4.5 and insert into library
- Some issues resolved (de Mesmay et al. 2010)

Experimental Setup

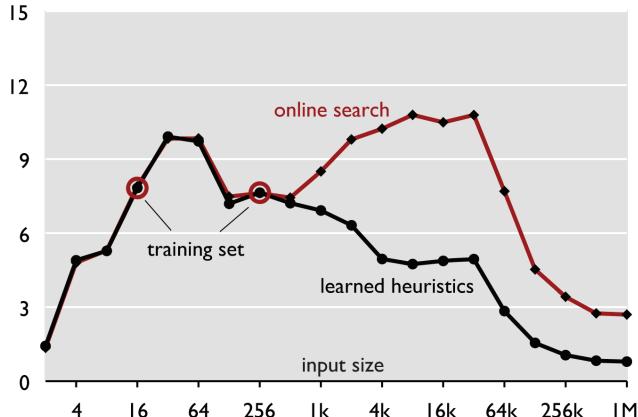
- 3GHz Intel Xeon 5160 (2 Core 2 Duos = 4 cores), icc 10.1
- Spiral-generated library:



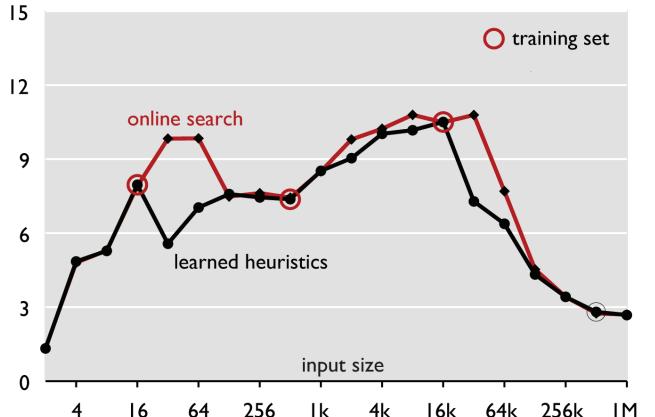
Complex DFT, double precision, up to 4 threads
Performance [GFlop/s]



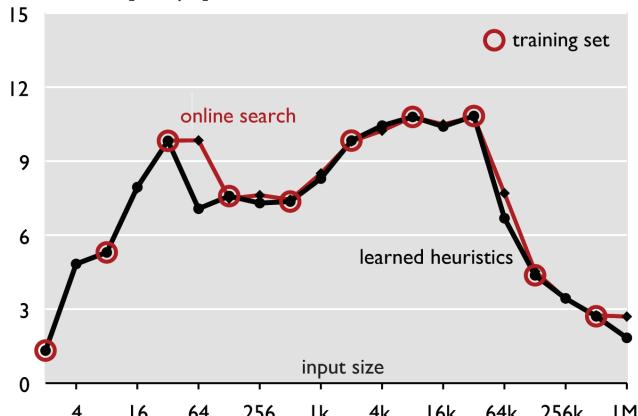
Complex DFT, double precision, up to 4 threads
Performance [GFlop/s]



Complex DFT, double precision, up to 4 threads
Performance [GFlop/s]



Complex DFT, double precision, up to 4 threads
Performance [GFlop/s]



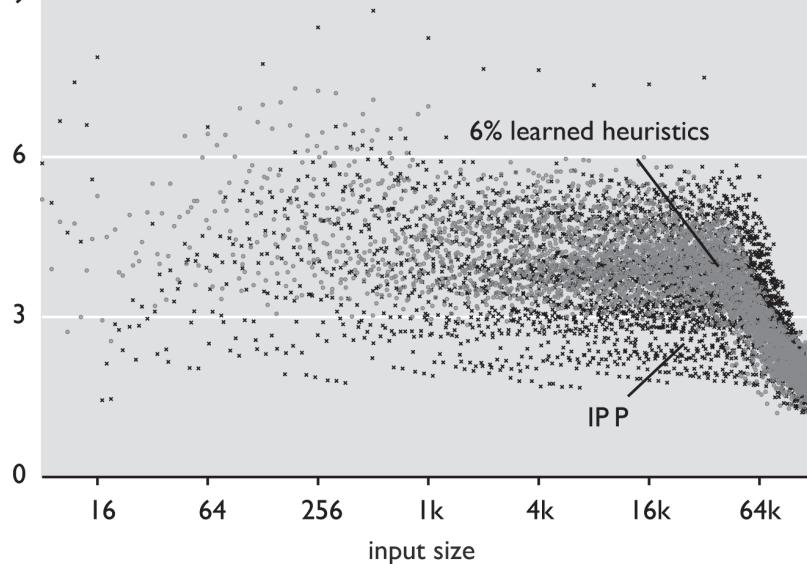
Learning works as expected

“All” Sizes

Complex DFT, double precision, mixed sizes

Performance [GFlop/s]

9



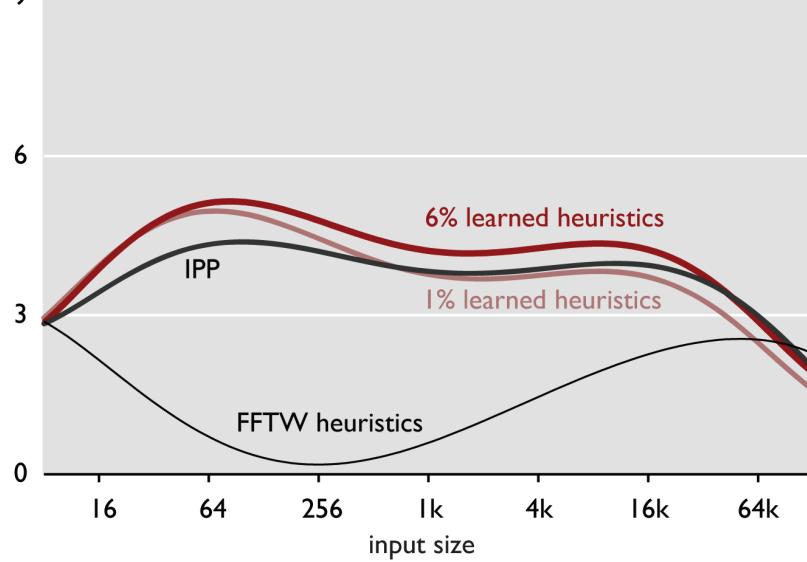
- All sizes $n \leq 2^{18}$, with prime factors ≤ 19

“All” Sizes

Complex DFT, double precision, mixed sizes

Performance [GFlop/s]

9



- All sizes $n \leq 2^{18}$, with prime factors ≤ 19
- Higher order fit of all sizes

Related Work

■ Machine learning in autotuning

- Linear regression for stencil/sorting (Brewer 1995)
- **Pioneering work:**
Learning DFT recursions in Spiral (Singer/Veloso 2000)
- Linear regression/SVM in PhiPAC (Vuduc/Demmel/Bilmes 2001)

■ Machine learning in compilation

- Scheduling (Moss et al. 1997, Cavazos/Moss 2004)
- Branch prediction (Calder et al. 1997)
- Heuristics generation (Monsifrot/Bodin/Quiniou 2002)
- Feature generation (Leather/Bonilla/O'Boyle 2009)

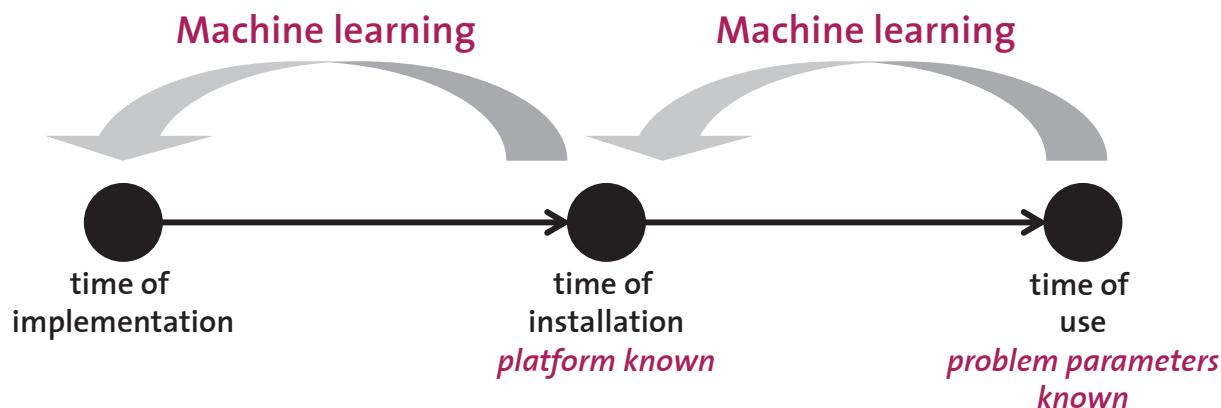
This Talk

- Frédéric de Mesmay, Yevgen Voronenko and Markus Püschel
Offline Library Adaptation Using Automatically Generated Heuristics
 Proc. International Parallel and Distributed Processing Symposium (IPDPS),
 pp. 1-10, 2010

Other Recent ML work in Spiral

- Frédéric de Mesmay, Arpad Rimmel, Yevgen Voronenko and Markus Püschel
Bandit-Based Optimization on Graphs with Application to Library Performance Tuning
 Proc. International Conference on Machine Learning (ICML), pp. 729-736,
 2009

Message of Talk



- **Machine learning should be used in autotuning**
 - Overcomes the problem of expensive searches
 - Relatively easy to do
 - Applicable to any search-based approach
- **Don't forget: autotuning ≠ search**