

$$A \otimes I_n$$

Automatic Performance Programming?

Markus Püschel
Computer Science

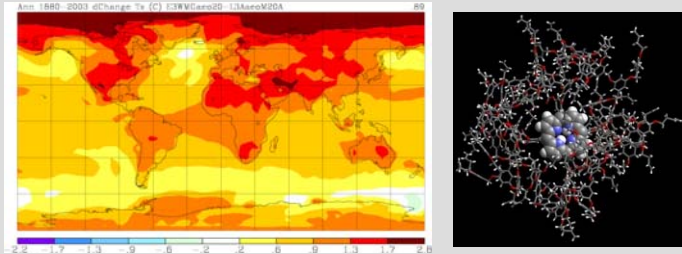
ETH

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

SPIRAL
www.spiral.net

```
__m128i t3 = _mm_unpacklo_epi16(X[0], X[1]);  
__m128i t4 = _mm_unpackhi_epi16(X[0], X[1]);  
__m128i t7 = _mm_unpacklo_epi16(X[2], X[3]);  
__m128i t8 = _mm_unpackhi_epi16(X[2], X[3]);
```

Scientific Computing



Simulations: Physics, Biology, ...

Mainstream Computing



Audio/Image/Video processing, ...

Embedded Computing



Signal processing, communication, ...

Unlimited need for performance

Many applications, but relatively few (~100 to 1000) components:

- Matrix multiplication
- Filters
- Fourier transform
- Coding/decoding
- Geometric transformations
- Graph algorithms
- ...

Fast components → fast applications

Software Performance: Traditional Approach

Algorithms

optimal op count

Software

Compilers

performance optimization

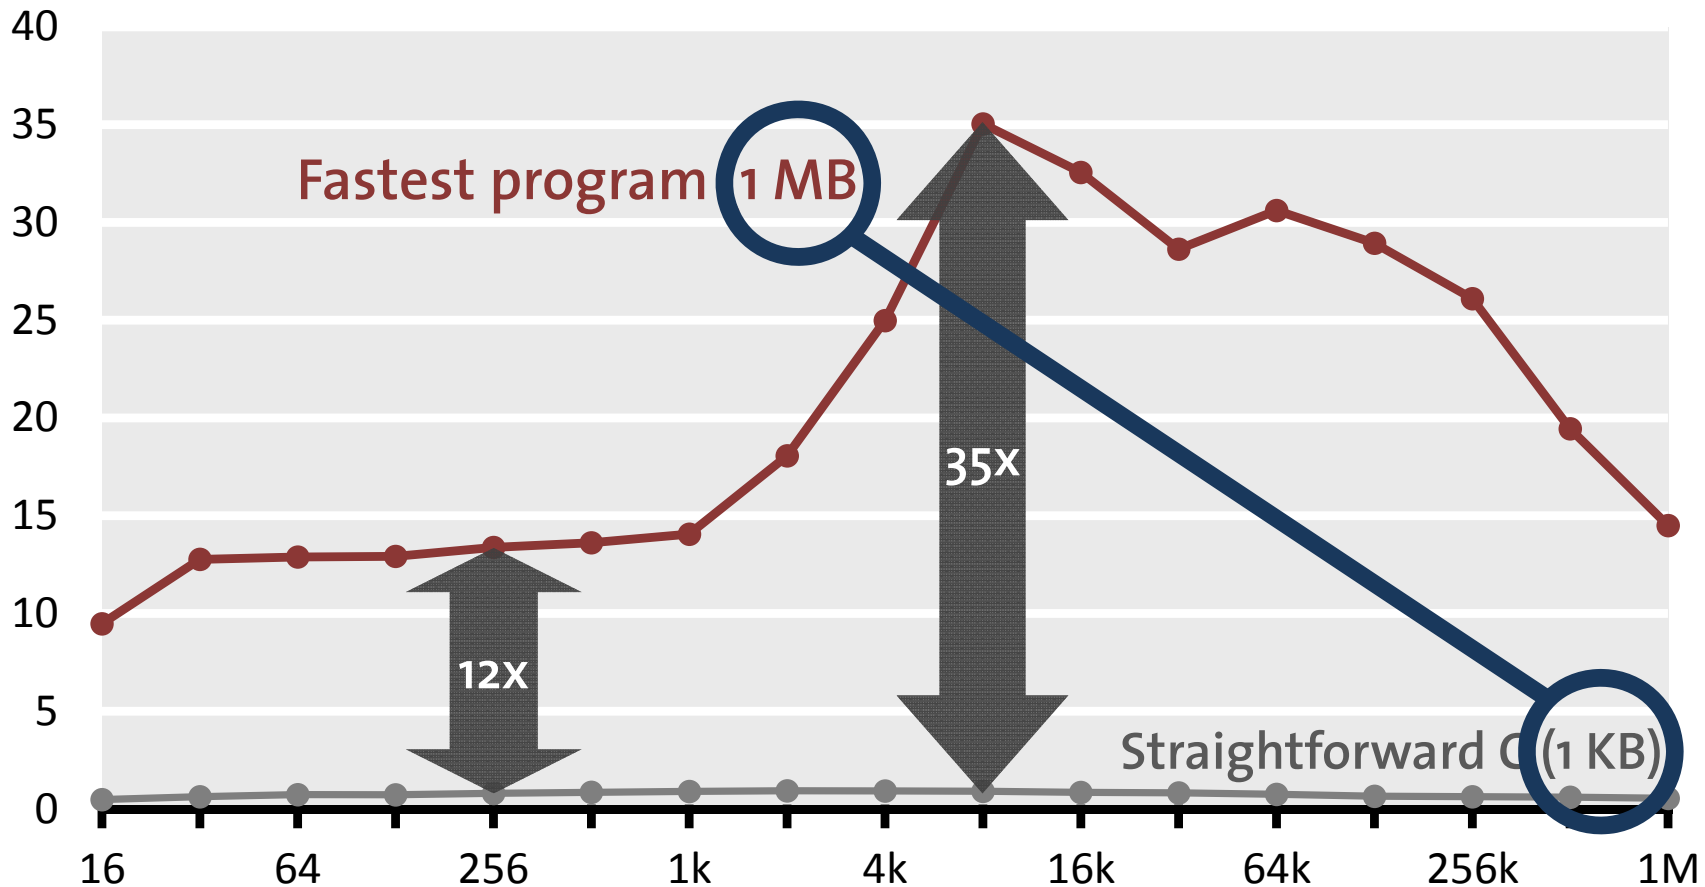
Microarchitecture

How well does that work?

The Problem: Example Fourier Transform

Discrete Fourier transform (single precision) on Intel Core i7 (4 cores)

Performance [Gflop/s]

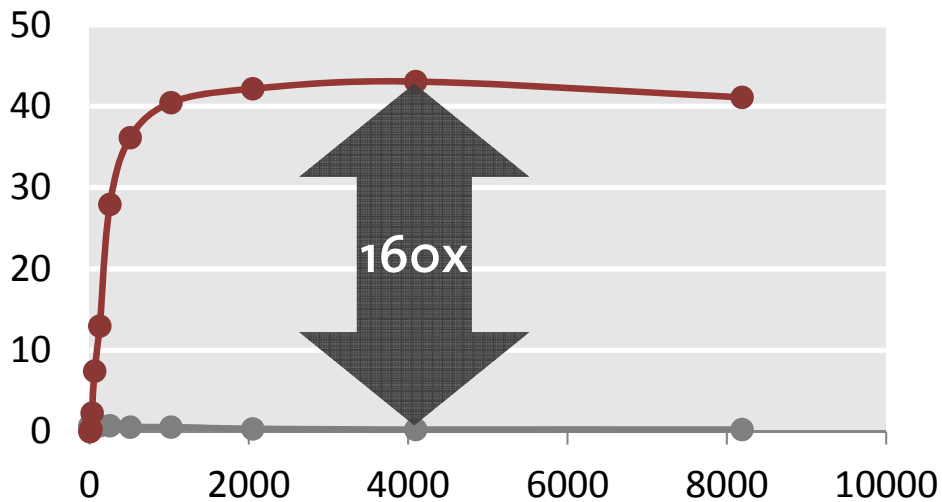


- Same opcount
- Best compiler

The Problem Is Everywhere

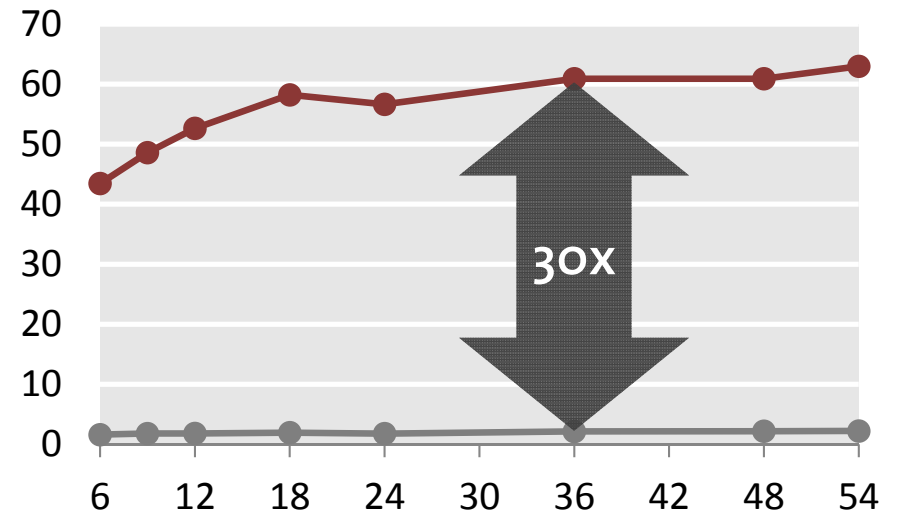
Matrix multiplication

Performance [Gflop/s]



WiFi Receiver

Performance [Mbit/s]



Model predictive control
Eigenvalues
LU factorization
Optimal binary search organization
Image color conversions
Image geometry transformations
Enclosing ball of points
Metropolis algorithm, Monte Carlo
Seam carving
SURF feature detection
Submodular function optimization
Graph cuts, Edmond-Karps Algorithm
Gaussian filter
Black Scholes option pricing
Disparity map refinement

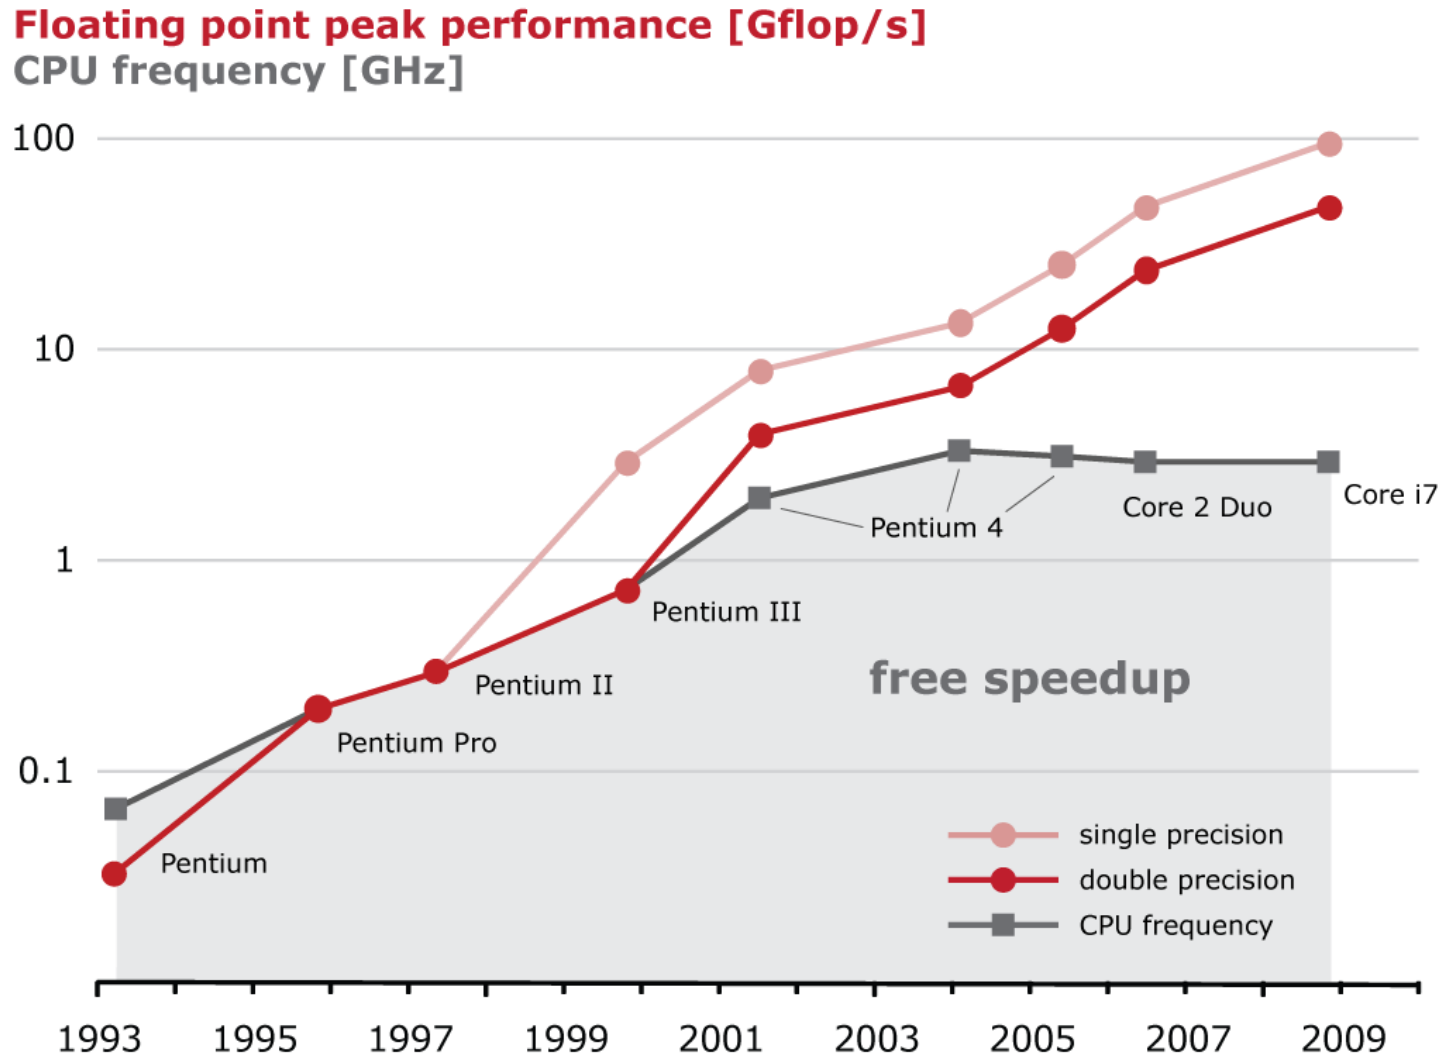
Singular-value decomposition
Mean shift algorithm for segmentation
Stencil computations
Displacement based algorithms
Motion estimation
Multiresolution classifier
Kalman filter
Object detection
IIR filters
Arithmetic for large numbers
Optimal binary search organization
Software defined radio
Shortest path problem
Feature set for biomedical imaging
Biometrics identification

“Theorem:”

Let f be a mathematical function to be implemented on a state-of-the-art processor. Then

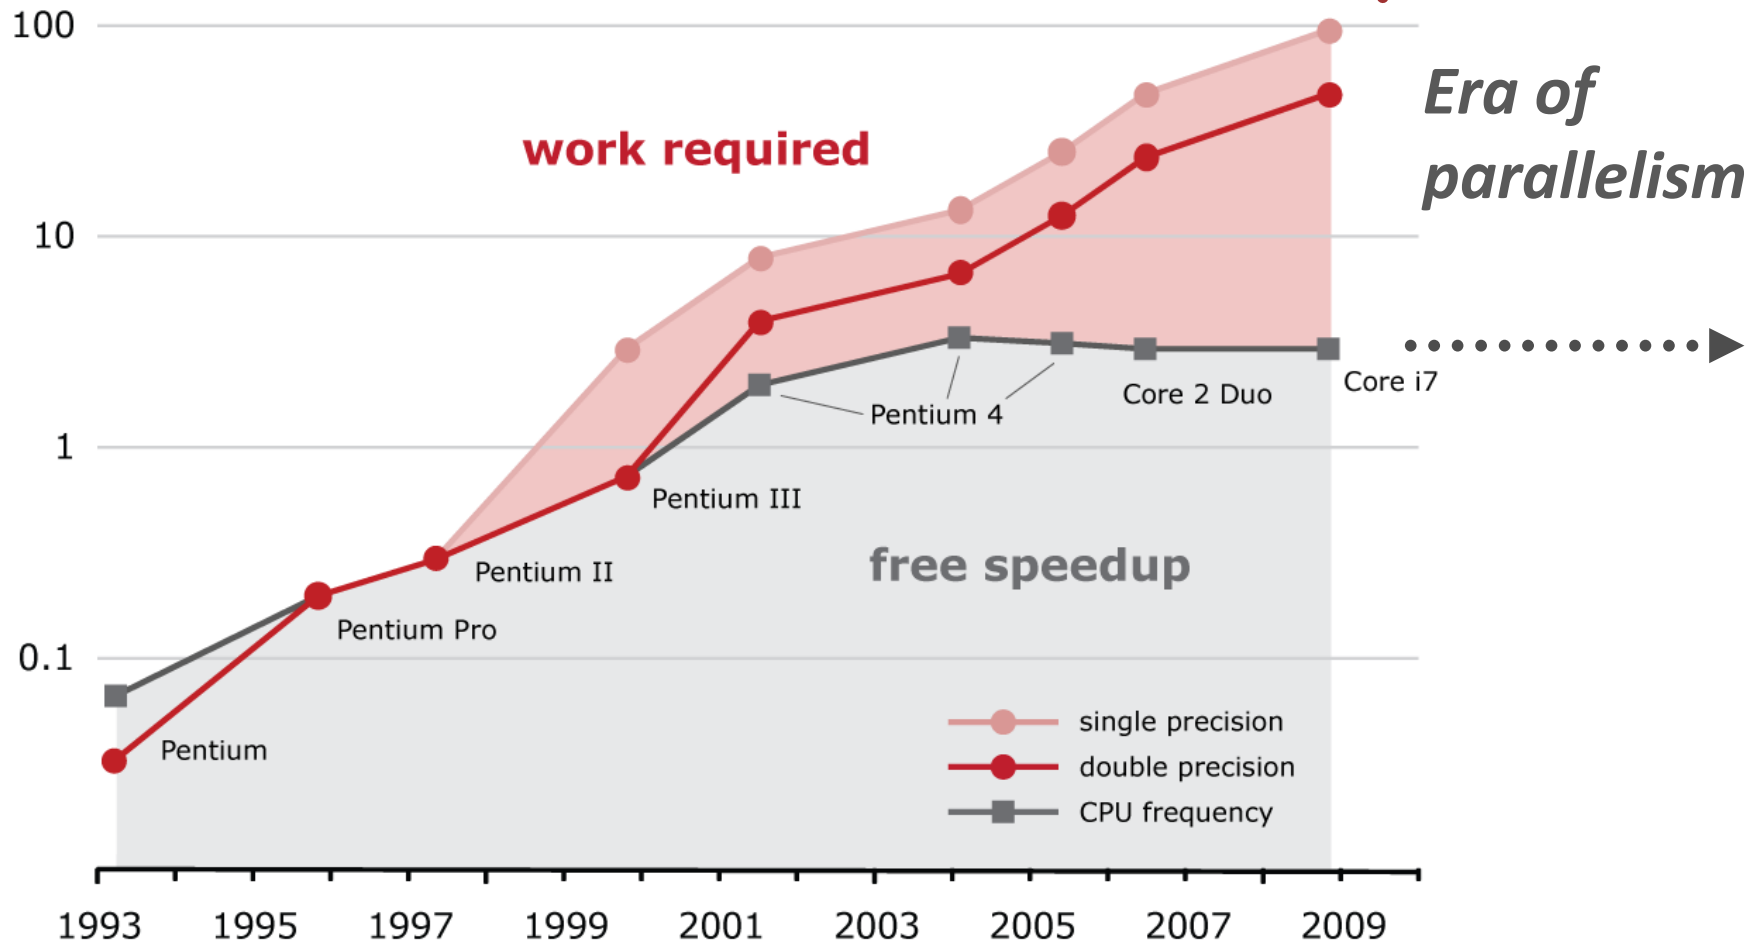
$$\frac{\text{Performance of optimal implementation of } f}{\text{Performance of straightforward implementation of } f} \approx 10-100$$

Evolution of Processors (Intel)



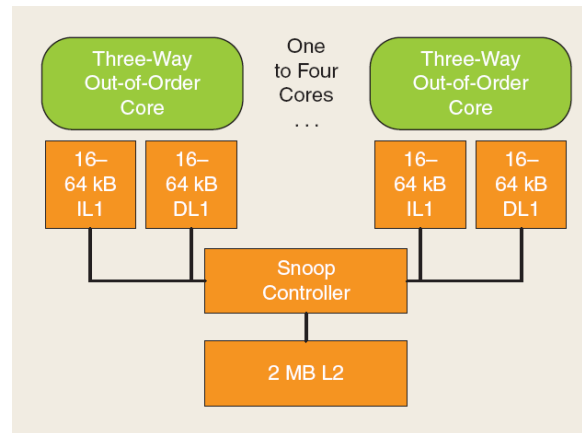
The End of Automatic Speedup

Floating point peak performance [Gflop/s]
CPU frequency [GHz]

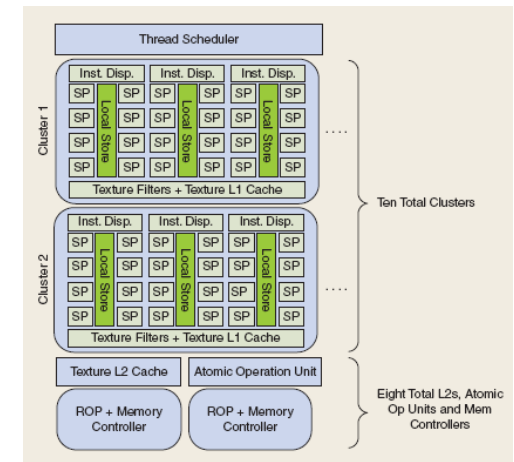


And There Will Be Variety ...

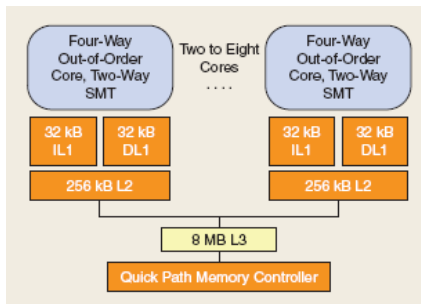
Arm Cortex A9



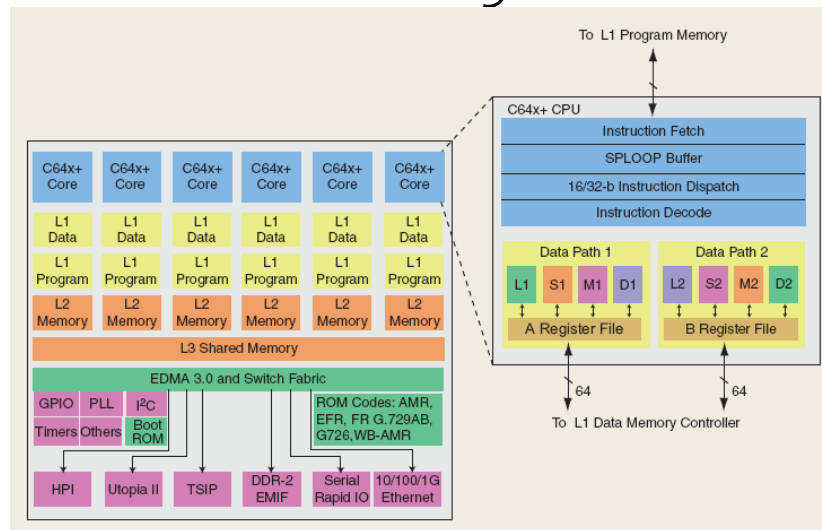
Nvidia G200



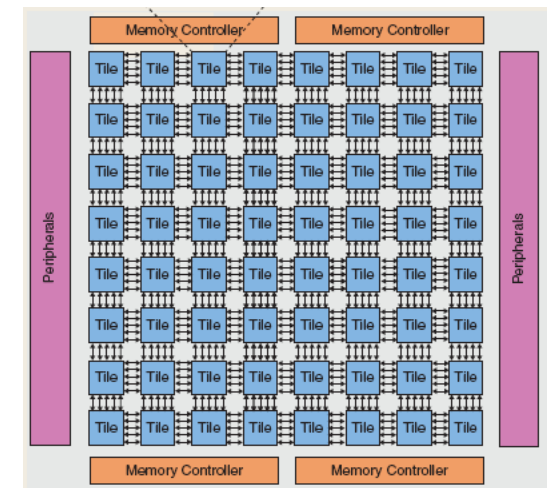
Core i7



TI TNETV3020

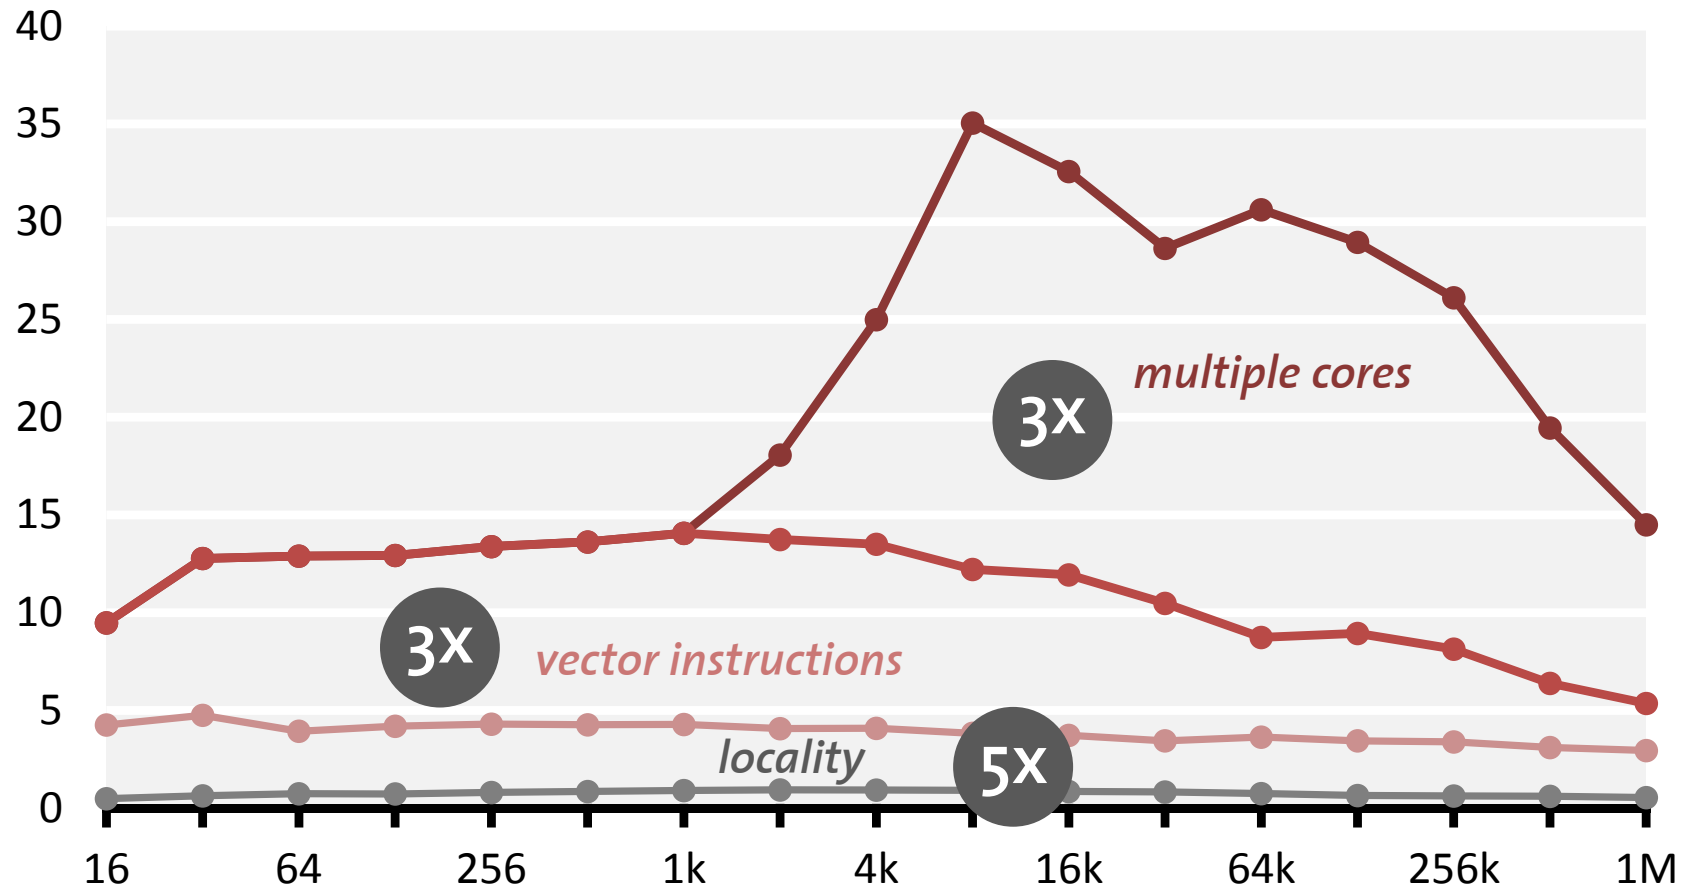


Tilera Tile64



DFT: Analysis

Discrete Fourier transform (single precision) on Intel Core i7 (4 cores)
Performance [Gflop/s]



- Compiler doesn't do the job (why?)
- Doing it by hand: requires guru knowledge

Optimization for Register Locality and ILP

```
// straightforward code
for(i = 0; i < N; i += 1)
  for(j = 0; j < N; j += 1)
    for(k = 0; k < N; k += 1)
      c[i][j] += a[i][k]*b[k][j];
```

```
// unrolling + scalar replacement
for(i = 0; i < N; i += MU) {
  for(j = 0; j < N; j += NU) {
    for(k = 0; k < N; k += KU) {
      t1 = A[i*N + k];
      t2 = A[i*N + k + 1];
      t3 = A[i*N + k + 2];
      t4 = A[i*N + k + 3];
      t5 = A[(i + 1)*N + k];
      <more copies>

      t10 = t1 * t9;
      t17 = t17 + t10;
      t21 = t1 * t8;
      t18 = t18 + t21;
      t12 = t5 * t9;
      t19 = t19 + t12;
      t13 = t5 * t8;
      t20 = t20 + t13;
      <more ops>

      C[i*N + j]           = t17;
      C[i*N + j + 1]       = t18;
      C[(i+1)*N + j]       = t19;
      C[(i+1)*N + j + 1] = t20;
    }
  }
}
```

load

compute

store

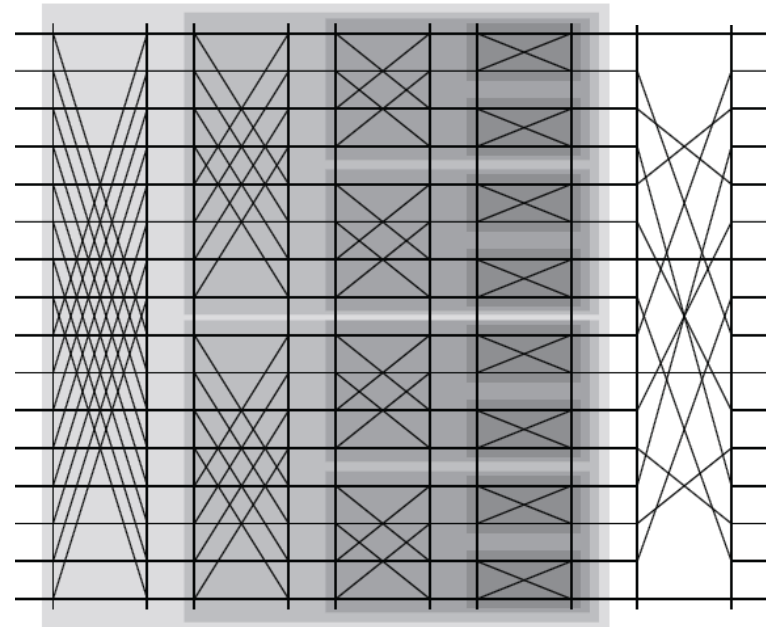
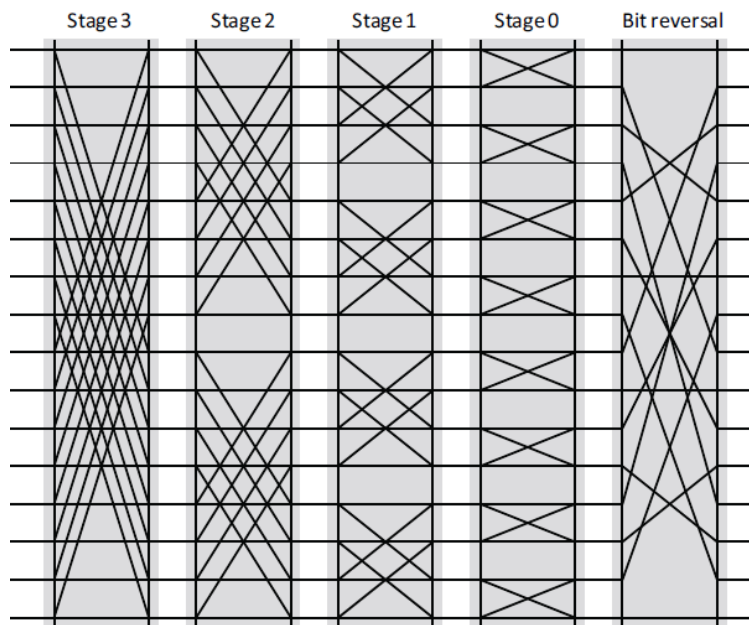
Removes aliasing

Enables register allocation and instruction scheduling

Compiler typically does not do:

- often illegal
- many choices

Optimization for Cache Locality



Many algorithms are formulated iteratively:

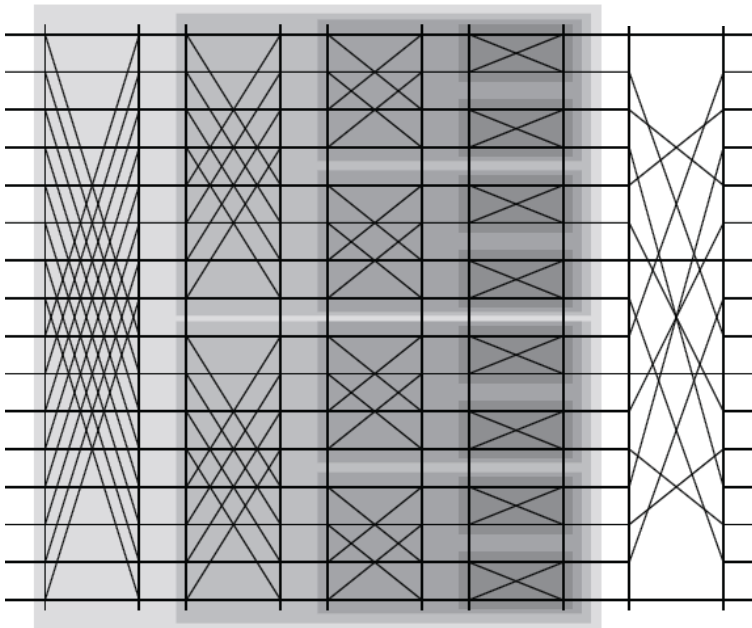
- many passes through data
- poor locality
- poor performance

Restructured for locality

Compiler usually does not do

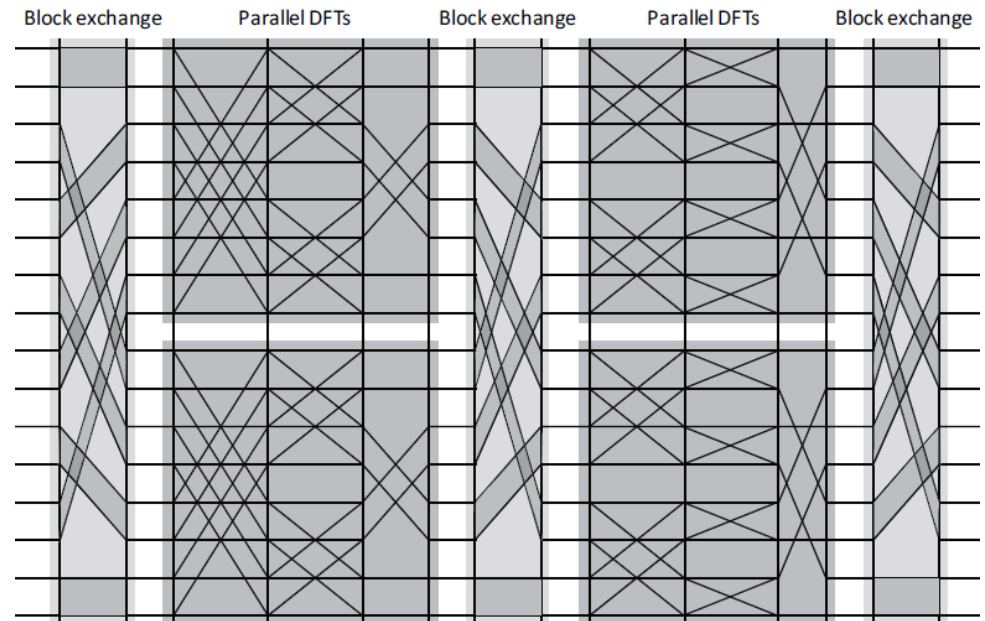
- analysis may be unfeasible
- possibly many choices
- may require algorithm changes
- may require domain knowledge
- may require cache parameters

Optimization for Parallelism (Threads)



Restructured for locality

Parallelism is present, but is not in the “right shape”



Restructured for locality and parallelism (shared memory, 2 cores, 2 elements per cache line)

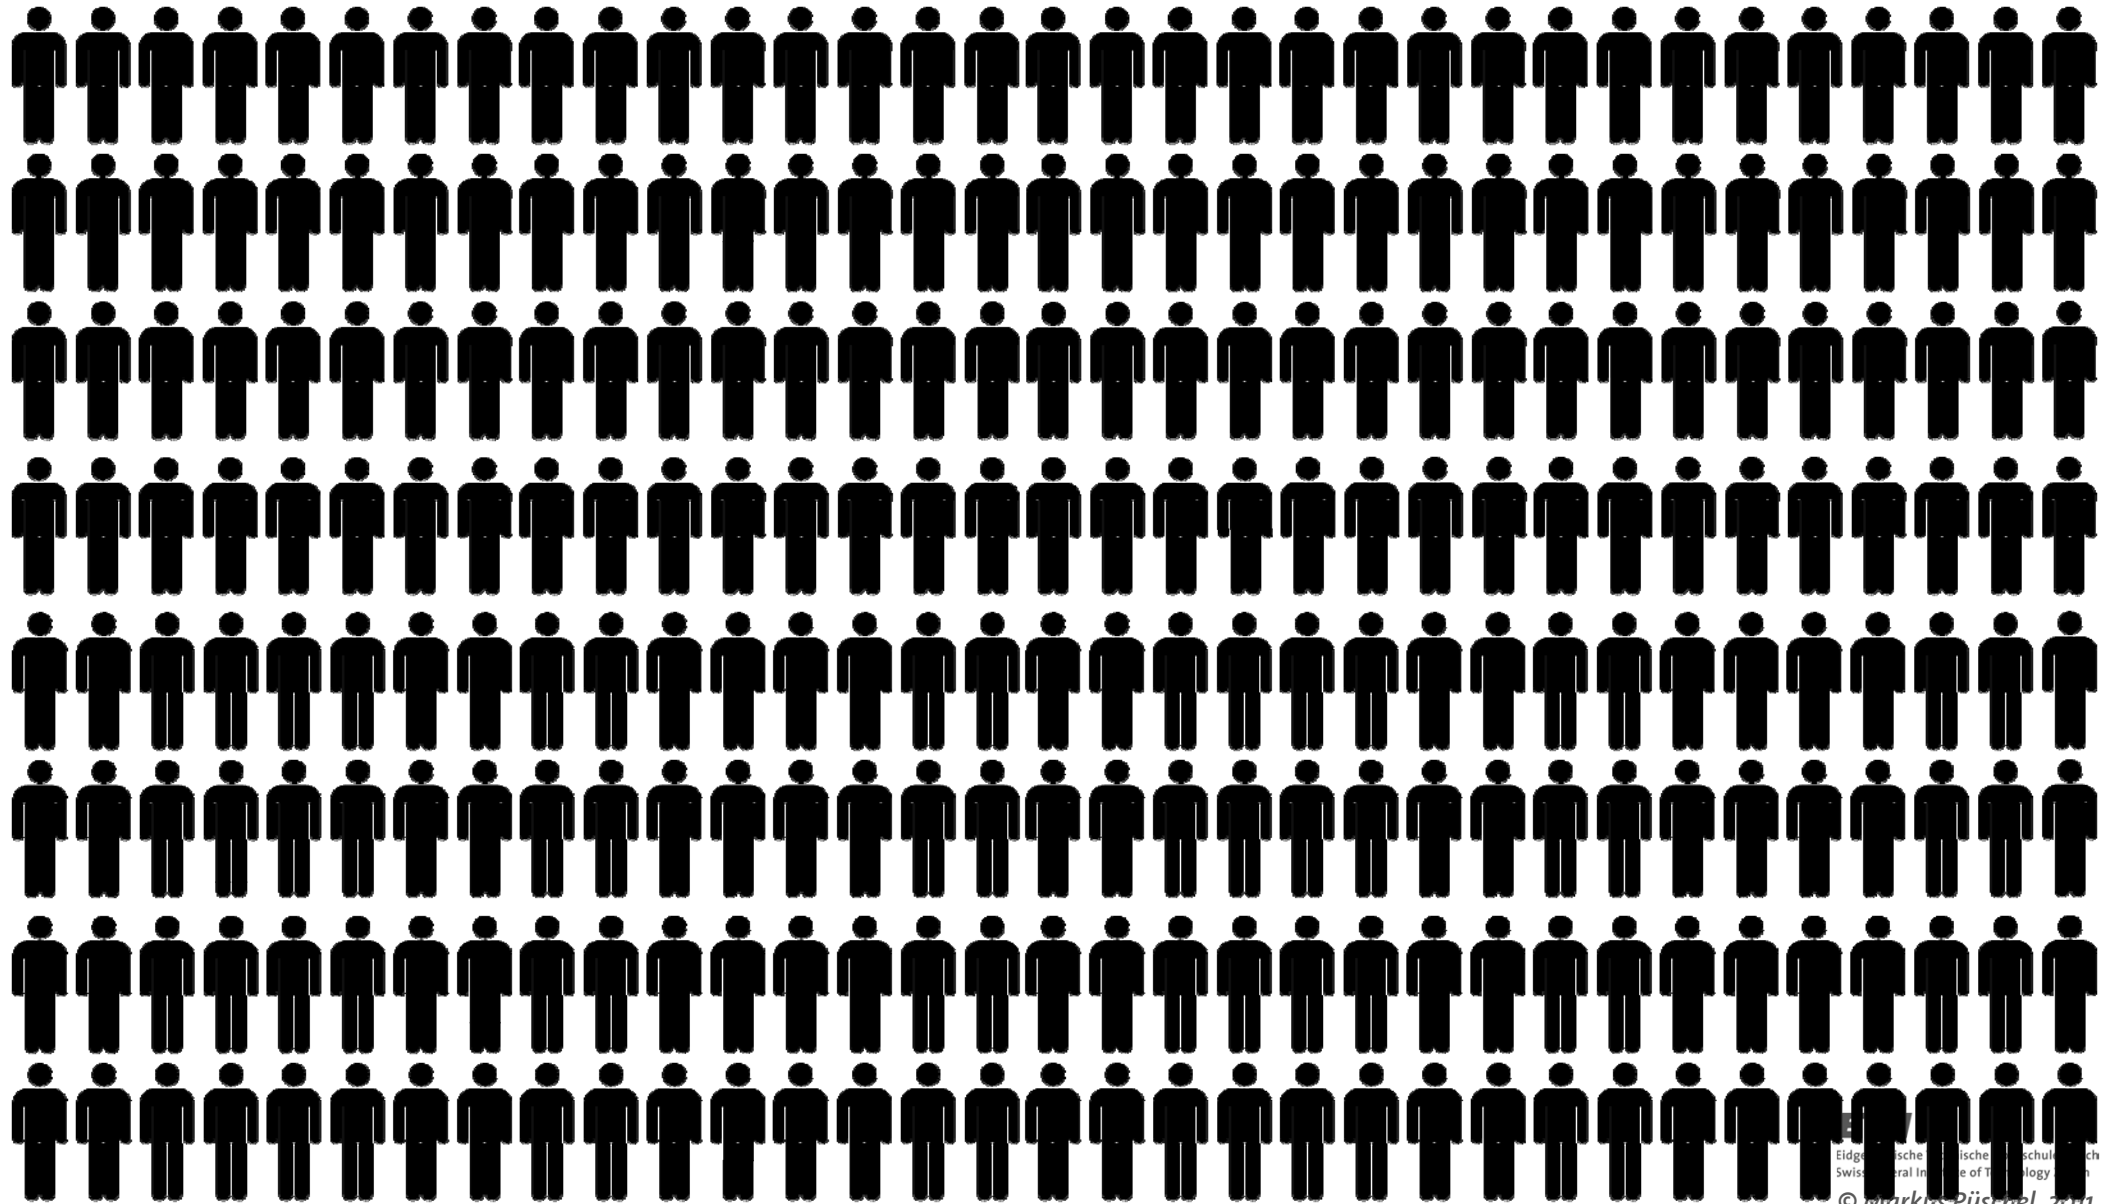
Compiler usually does not do

- analysis may be unfeasible
- may require algorithm changes
- may require domain knowledge
- may require processor parameters

Software Performance: Facts

- Straightforward code is often slow
- Inherent compiler limitations
- End of free speedup for legacy code
- Performance optimization is very hard
- Fast code violates good software engineering practices
- Performance optimization is “vertical”
 - *algorithm changes*
 - *code style*
 - *considers microarchitectural parameters*
- Highest performance is generally non-portable

Current practice: Thousands of programmers re-implement and re-optimize the same functionality for every new processor and for every new processor generation



Algorithms

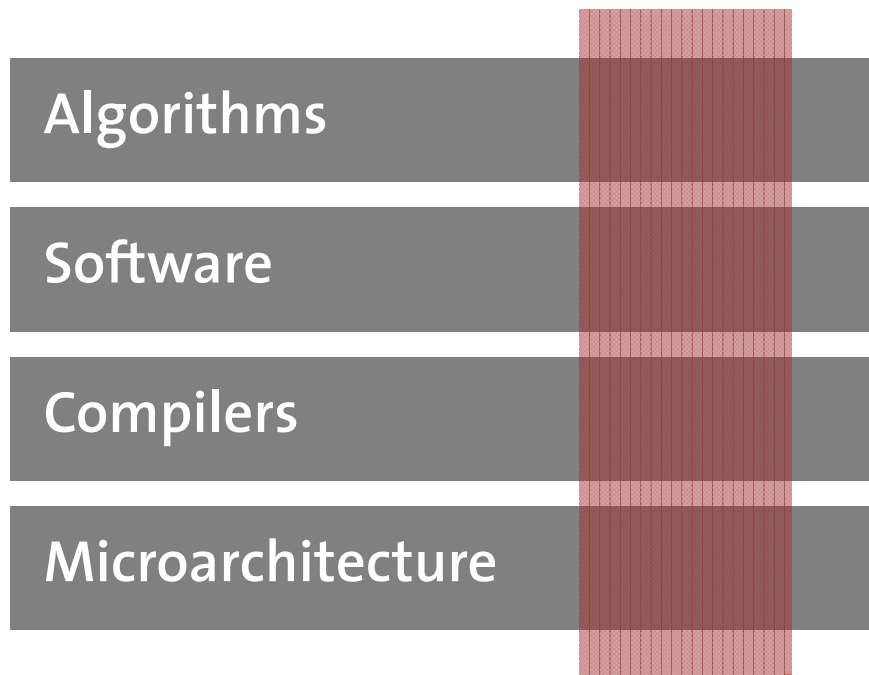
optimal op count

Software

Compilers

performance optimization

Microarchitecture



performance optimization

Performance is different than other software quality features

We need

- Courses on software performance
- Microarchitecture-cognizant algorithm analysis (e.g., Kung's work, cache oblivious algorithms, ...)
- Smarter compilers (e.g., iterative compilation, machine learning)
- *“Vertical” software engineering and programming languages techniques for performance optimization*

Organization

- Software performance issues
- *Automatic performance tuning (autotuning)*
- Automatic performance programming
- Conclusions

ATLAS/PhiPAC: MMM Generator

Bilmes et al. 97, Whaley et al. 98

```
// 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 , ...

Offline tuning: tuning at installation

FFTW: Adaptive DFT Library

Frigo, Johnson 98, 05

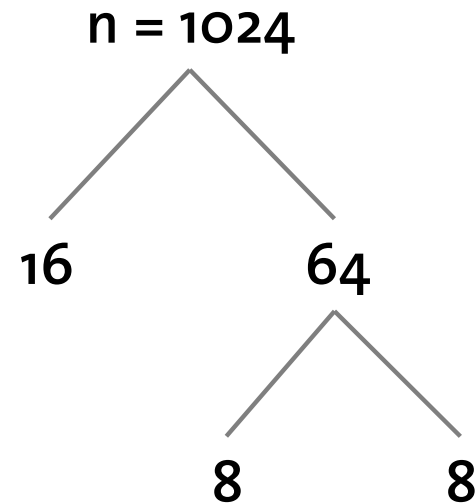
Installation

configure/make

Usage

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

Search for fastest
computation strategy



Online tuning: tuning at time of use

Autotuning

■ Use of models

- OSKI (Vuduc, Demmel, Yelick et al.)
- Adaptive sorting (Li, Padua et al.)

■ Tools

- Iterative compilation (Knijnenburg, O'Boyle, Cooper, ...)
- Machine learning in compilation (O'Boyle, Cavazos, ...)
- Source code tools (POET, Rose)

Most common form of autotuning: Search over alternative implementations to find the fastest for a given platform

Problem: usually parameter based

- not portable to new types of platforms (e.g., parallel)
- not portable to different functions

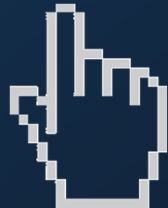
Organization

- Software performance issues
- Automatic performance tuning (autotuning)
- *Automatic performance programming*
- Conclusions

Goal:

Computer generation of high performance code for ubiquitous performance-critical components

Generate Code



“click”

Possible Approach:

Capturing algorithm knowledge:
Domain-specific languages (DSLs)

Structural optimization:
Rewriting systems

High performance code style:
Compiler

Decision making for choices:
Machine learning

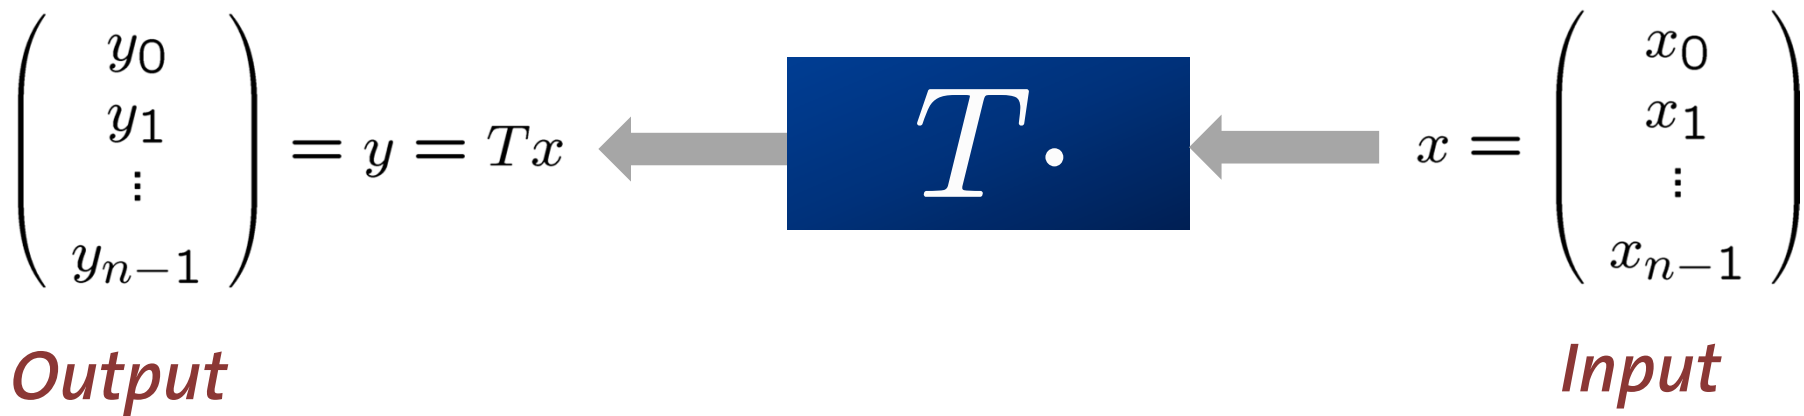
Spiral: Program Generation for Performance (www.spiral.net)



Franz Franchetti
Yevgen Voronenko
Jianxin Xiong
Bryan Singer
Srinivas Chellappa
Frédéric de Mesmay
Peter Milder
José Moura
David Padua
Jeremy Johnson
James Hoe
<many more>

funding: DARPA, NSF, ONR, Intel

Linear Transforms



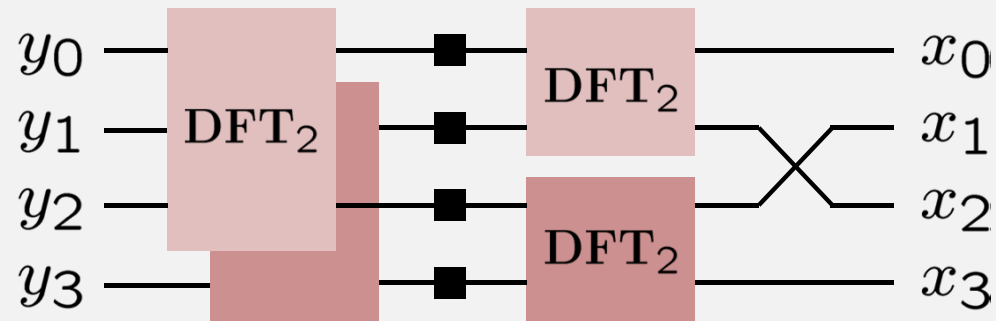
Example: $T = \text{DFT}_n = [e^{-2k\ell\pi i/n}]_{0 \leq k, \ell < n}$

Algorithms: Example FFT, n = 4

Fast Fourier transform (FFT):

$$\begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & i & -1 & -i \\ 1 & -1 & 1 & -1 \\ 1 & -i & -1 & i \end{bmatrix} x = \begin{bmatrix} 1 & \cdot & 1 & \cdot \\ \cdot & 1 & \cdot & 1 \\ 1 & \cdot & -1 & \cdot \\ \cdot & 1 & \cdot & -1 \end{bmatrix} \begin{bmatrix} 1 & \cdot & \cdot & \cdot \\ \cdot & 1 & \cdot & \cdot \\ \cdot & \cdot & 1 & \cdot \\ \cdot & \cdot & \cdot & i \end{bmatrix} \begin{bmatrix} 1 & 1 & \cdot & \cdot \\ 1 & -1 & \cdot & \cdot \\ \cdot & \cdot & 1 & 1 \\ \cdot & \cdot & 1 & -1 \end{bmatrix} \begin{bmatrix} 1 & \cdot & \cdot & \cdot \\ \cdot & \cdot & 1 & \cdot \\ \cdot & 1 & \cdot & \cdot \\ \cdot & \cdot & \cdot & 1 \end{bmatrix} x$$

Data flow graph



Description with matrix algebra (SPL)

$$\text{DFT}_4 = (\text{DFT}_2 \otimes \text{I}_2) \text{T}_2^4 (\text{I}_2 \otimes \text{DFT}_2) \text{L}_2^4$$

Decomposition Rules (>200 for >40 Transforms)

$$\begin{aligned} \mathbf{DFT}_n &\rightarrow P_{k/2,2m}^\top (\mathbf{DFT}_{2m} \oplus (I_{k/2-1} \otimes_i C_{2m} \mathbf{rDFT}_{2m}(i/k))) (\mathbf{RDFT}'_k \otimes I_m), \quad k \text{ even,} \\ \begin{pmatrix} \mathbf{RDFT}_n \\ \mathbf{RDFT}'_n \\ \mathbf{DHT}_n \\ \mathbf{DHT}'_n \end{pmatrix} &\rightarrow (P_{k/2,m}^\top \otimes I_2) \left(\begin{pmatrix} \mathbf{RDFT}_{2m} \\ \mathbf{RDFT}'_{2m} \\ \mathbf{DHT}_{2m} \\ \mathbf{DHT}'_{2m} \end{pmatrix} \oplus \left(I_{k/2-1} \otimes_i D_{2m} \begin{pmatrix} \mathbf{rDFT}_{2m}(i/k) \\ \mathbf{rDHT}_{2m}(i/k) \end{pmatrix} \right) \right) \begin{pmatrix} \mathbf{RDFT}'_k \\ \mathbf{RDFT}'_k \\ \mathbf{DHT}'_k \\ \mathbf{DHT}'_k \end{pmatrix} \otimes I_m, \quad k \text{ even,} \\ \begin{pmatrix} \mathbf{rDFT}_{2n}(u) \\ \mathbf{rDHT}_{2n}(u) \end{pmatrix} &\rightarrow L_m^{2n} \left(I_k \otimes_i \begin{pmatrix} \mathbf{rDFT}_{2m}((i+u)/k) \\ \mathbf{rDHT}_{2m}((i+u)/k) \end{pmatrix} \right) \left(\begin{pmatrix} \mathbf{rDFT}_{2k}(u) \\ \mathbf{rDHT}_{2k}(u) \end{pmatrix} \otimes I_m \right), \\ \mathbf{RDFT-3}_n &\rightarrow (Q_{k/2,m}^\top \otimes I_2) (I_k \otimes_i \mathbf{rDFT}_{2m})(i+1/2)/k) (\mathbf{RDFT-3}_k \otimes I_m), \quad k \text{ even,} \\ \mathbf{DCT-3}_n &\rightarrow P_{k/2,m}^\top (\mathbf{DCT-3}_{2m} \oplus (I_{k/2-1} \otimes_i N_{2m} \mathbf{RDFT-3}_{2m}^\top)) P_{k/2,m} (I_{n/2} \otimes I) (I \otimes \mathbf{RDFT}'_k) \otimes I_m \end{aligned}$$

Rules = algorithm knowledge

(≈100 journal papers)

$$\begin{aligned} \mathbf{DCT-3}_n &\rightarrow (I_m \oplus J_m) L_m^n (\mathbf{DCT-3}_m(1/4) \oplus \mathbf{DCT-3}_m(3/4)) \\ &\quad \cdot (F_2 \otimes I_m) \begin{bmatrix} I_m & 0 \oplus -J_{m-1} \\ \frac{1}{\sqrt{2}}(I_1 \oplus 2I_m) \end{bmatrix}, \quad n = 2m \\ \mathbf{DCT-4}_n &\rightarrow S_n \mathbf{DCT-2}_n \text{diag}_{0 \leq k < n} (1/(2 \cos((2k+1)\pi/4n))) \\ \mathbf{IMDCT}_{2m} &\rightarrow (J_m \oplus I_m \oplus I_m \oplus J_m) \left(\left(\begin{bmatrix} 1 \\ -1 \end{bmatrix} \otimes I_m \right) \oplus \left(\begin{bmatrix} -1 \\ -1 \end{bmatrix} \otimes I_m \right) \right) J_{2m} \mathbf{DCT-4}_{2m} \\ \mathbf{WHT}_{2^k} &\rightarrow \prod_{i=1}^t (I_{2^{k_1+\dots+k_{i-1}}} \otimes \mathbf{WHT}_{2^{k_i}} \otimes I_{2^{k_{i+1}+\dots+k_t}}), \quad k = k_1 + \dots + k_t \\ \mathbf{DFT}_2 &\rightarrow F_2 \\ \mathbf{DCT-2}_2 &\rightarrow \text{diag}(1, 1/\sqrt{2}) F_2 \\ \mathbf{DCT-4}_2 &\rightarrow J_2 R_{13\pi/8} \end{aligned}$$

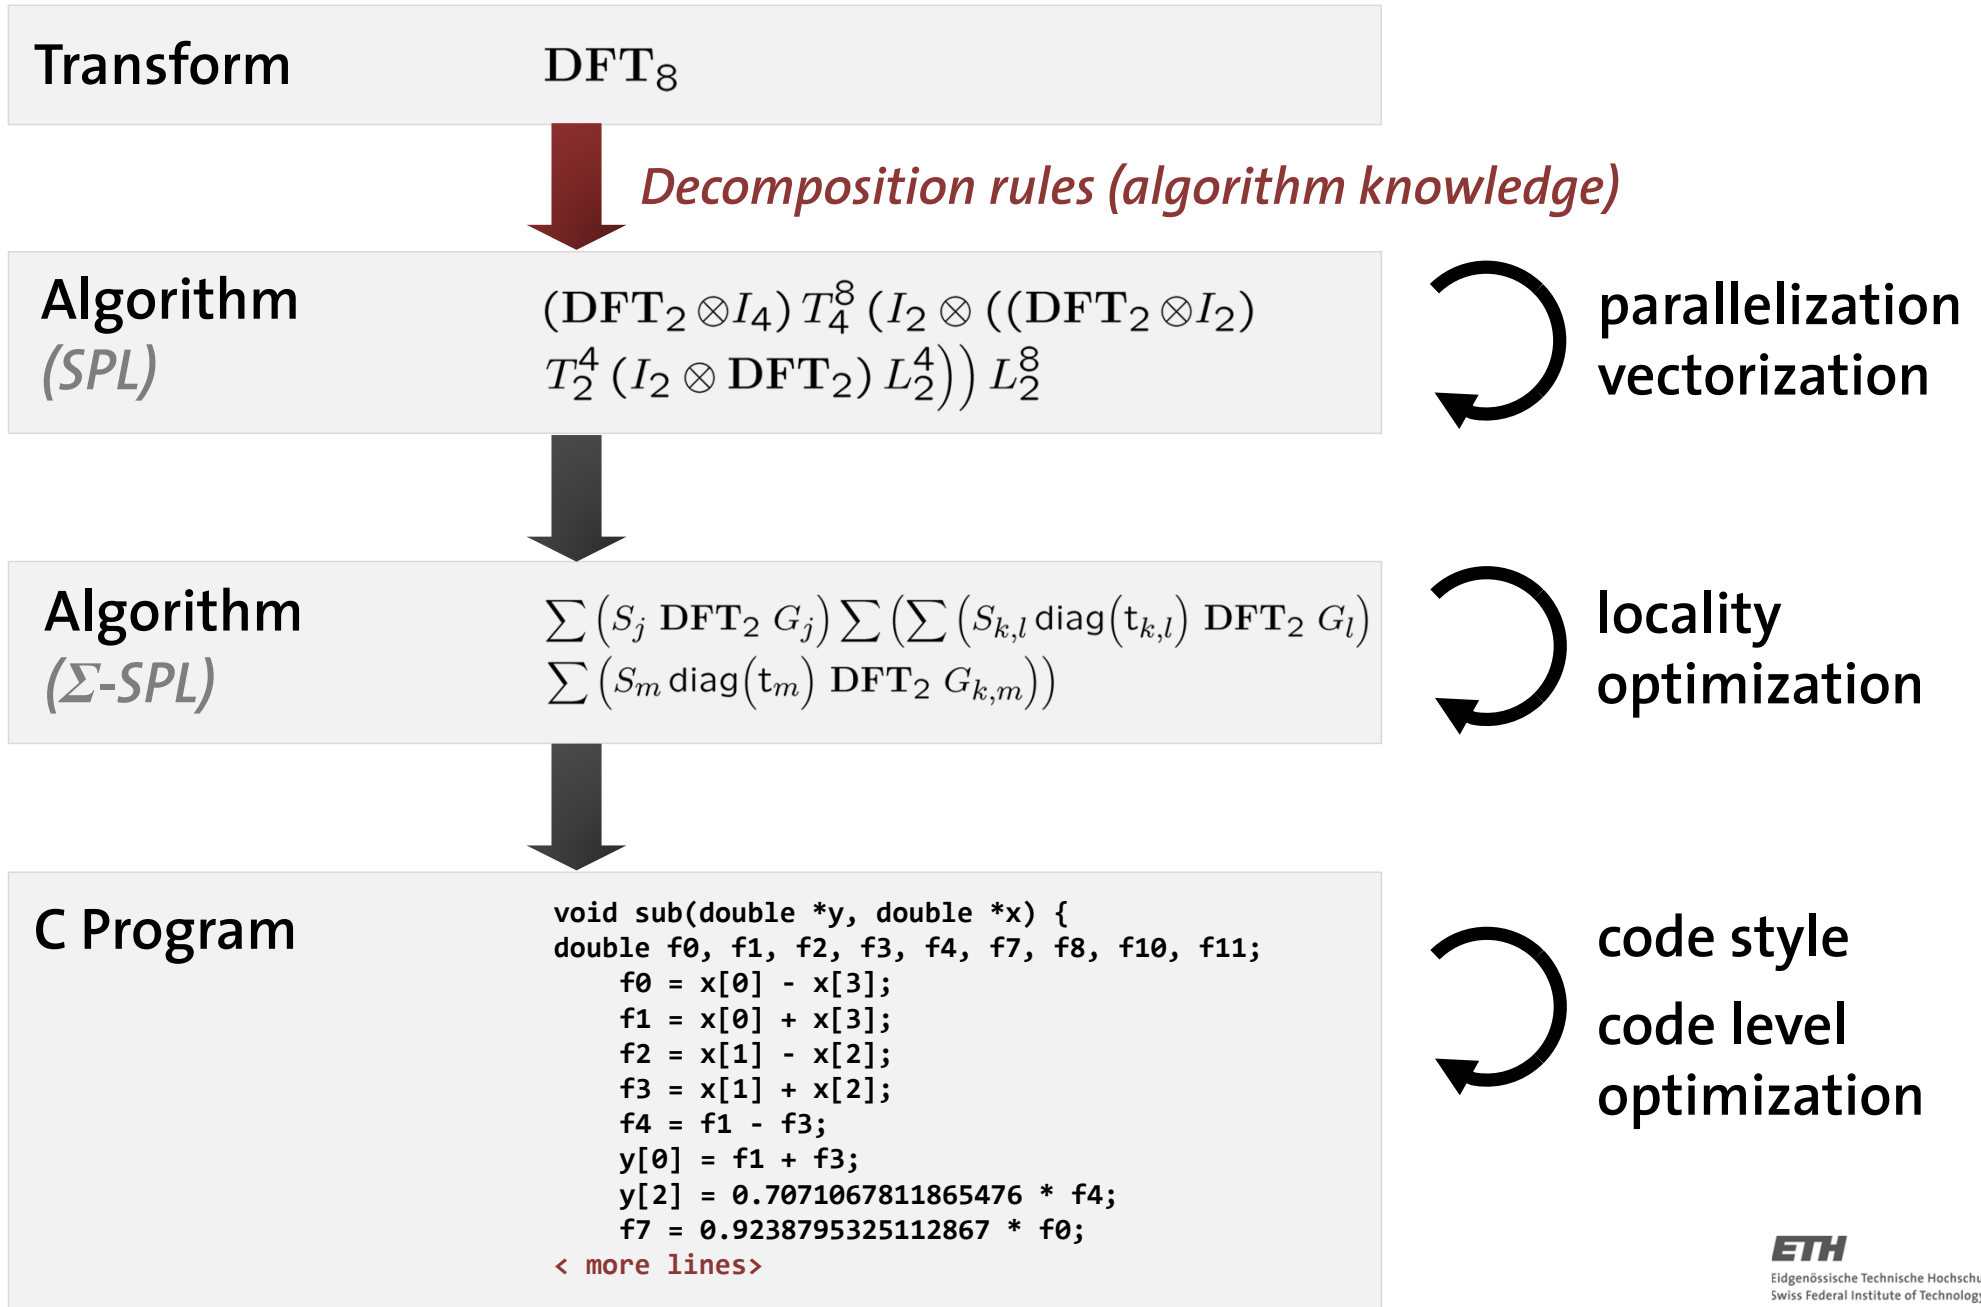
SPL to Code

SPL S	Pseudo code for $y = Sx$
$A_n B_n$	<pre><code for: t = Bx> <code for: y = At></pre>
$I_m \otimes A_n$	<pre>for (i=0; i<m; i++) <code for: y[i*n:1:i*n+n-1] = A(x[i*n:1:i*n+n-1])></pre>
$A_m \otimes I_n$	<pre>for (i=0; i<n; i++) <code for: y[i:n:i+m*n-n] = A(x[i:n:i+m*n-n])></pre>
D_n	<pre>for (i=0; i<n; i++) y[i] = D[i]*x[i];</pre>
L_k^{km}	<pre>for (i=0; i<k; i++) for (j=0; j<m; j++) y[i*m+j] = x[j*k+i];</pre>
F_2	<pre>y[0] = x[0] + x[1]; y[1] = x[0] - x[1];</pre>

$$I_m \otimes A_n = \begin{bmatrix} A_n & & \\ & \dots & \\ & & A_n \end{bmatrix}$$

Gives reasonable, straightforward code

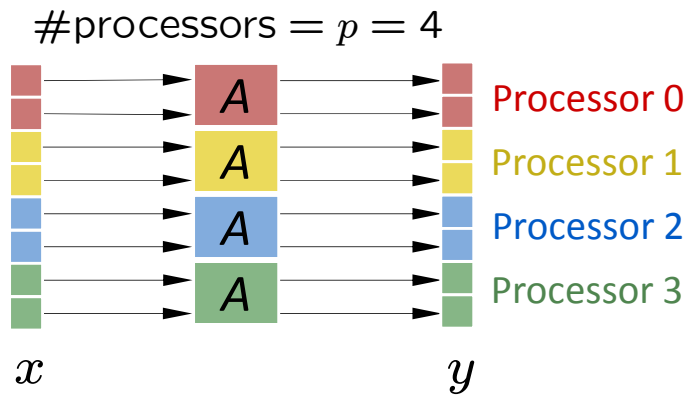
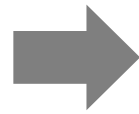
Program Generation in Spiral



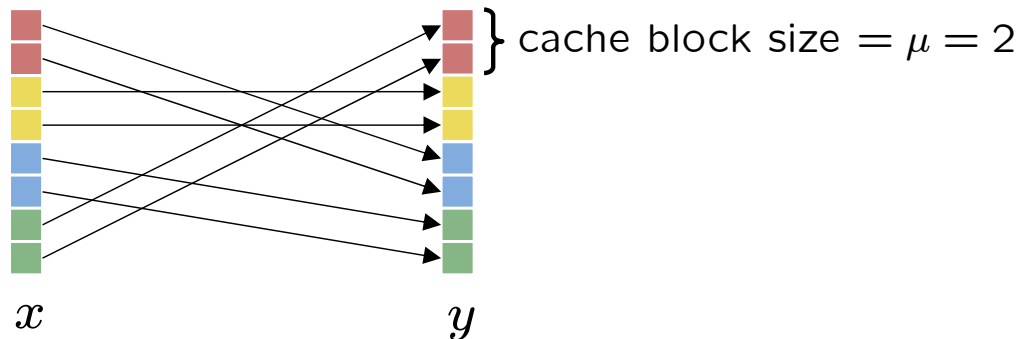
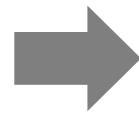
SPL to Shared Memory Code: Basic Idea

“Good” SPL structures

$$y = \left(I_p \otimes A \right) x$$



$$y = \left(P \otimes I_\mu \right) x$$



Rewriting: Bad structures \longrightarrow good structures

Example: SMP Parallelization

$$\underbrace{\text{DFT}_{mn}}_{\text{smp}(p,\mu)} \rightarrow \underbrace{\left((\text{DFT}_m \otimes \text{I}_n) \text{T}_n^{mn} (\text{I}_m \otimes \text{DFT}_n) \text{L}_m^{mn} \right)}_{\text{smp}(p,\mu)}$$

...

$$\rightarrow \underbrace{\left(\text{DFT}_m \otimes \text{I}_n \right)}_{\text{smp}(p,\mu)} \underbrace{\text{T}_n^{mn}}_{\text{smp}(p,\mu)} \underbrace{\left(\text{I}_m \otimes \text{DFT}_n \right)}_{\text{smp}(p,\mu)} \underbrace{\text{L}_m^{nm}}_{\text{smp}(p,\mu)}$$

...

$$\rightarrow \underbrace{\left((\text{L}_m^{mp} \otimes \text{I}_{n/p\mu}) \otimes \text{I}_\mu \right)}_{\text{red}} \underbrace{\left(\text{I}_p \otimes (\text{DFT}_m \otimes \text{I}_{n/p}) \right)}_{\text{blue}} \underbrace{\left((\text{L}_p^{mp} \otimes \text{I}_{n/p\mu}) \otimes \text{I}_\mu \right)}_{\text{red}}$$

$$\underbrace{\left(\bigoplus_{i=0}^{p-1} \text{T}_n^{mn,i} \right)}_{\text{blue}} \underbrace{\left(\text{I}_p \otimes (\text{I}_{m/p} \otimes \text{DFT}_n) \right)}_{\text{blue}} \underbrace{\left(\text{I}_p \otimes \text{L}_{m/p}^{mn/p} \right)}_{\text{blue}} \underbrace{\left((\text{L}_p^{pn} \otimes \text{I}_{m/p\mu}) \otimes \text{I}_\mu \right)}_{\text{red}}$$

load-balanced, no false sharing

One rewriting system for every platform paradigm:
SIMD, distributed memory parallelism, FPGA, ...

Challenge: General Size Libraries

So far:

Code specialized to fixed input size

```
DFT_384(x, y) {  
  ...  
  for(i = ...) {  
    t[2i]    = x[2i] + x[2i+1]  
    t[2i+1] = x[2i] - x[2i+1]  
  }  
  ...  
}
```

- Algorithm fixed (offline tuning)
- Nonrecursive code

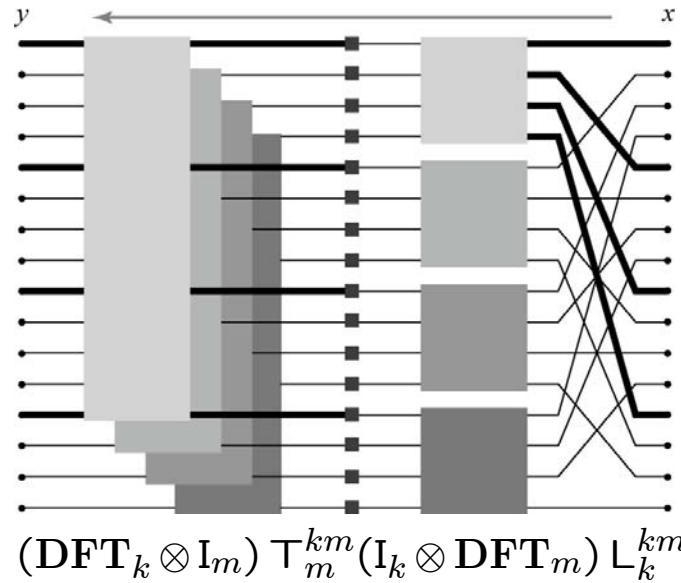
Challenge:

Library for general input size

```
DFT(n, x, y) {  
  ...  
  for(i = ...) {  
    DFT_strided(m, x+mi, y+i, 1, k)  
  }  
  ...  
}
```

- Algorithm cannot be fixed (online tuning)
- Recursive code
- Creates many challenges

Challenge: Recursive Composition



$$16 = 4 \times 4$$

```
void dft(int n, cpx *y, cpx *x) {
```

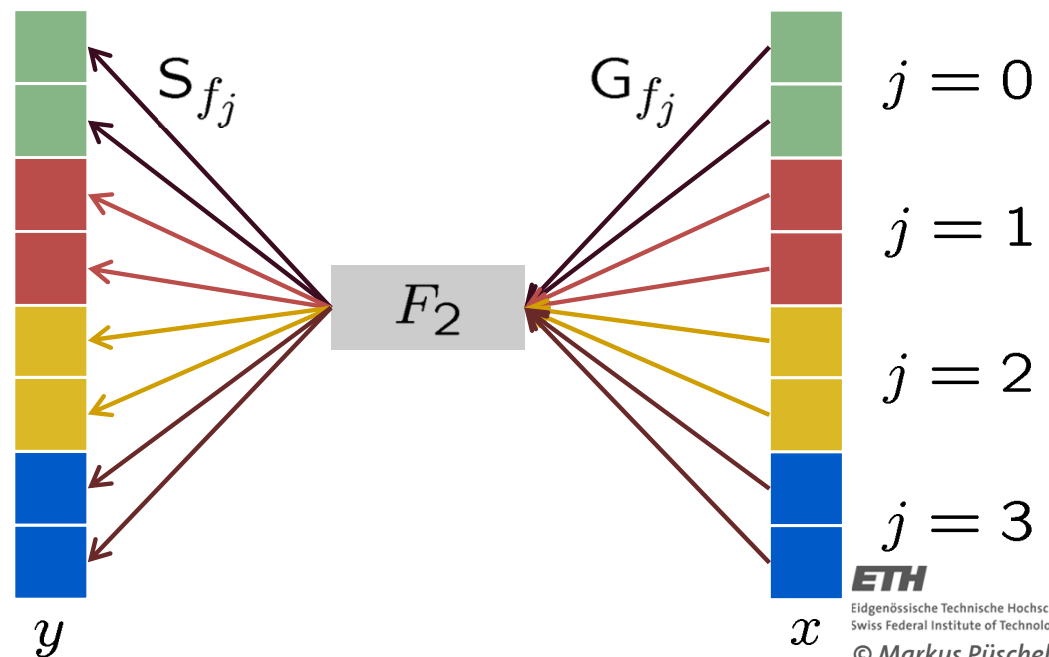
Σ -SPL : Basic Idea

Four additional matrix constructs: Σ , G , S , Perm

- Σ (sum) explicit loop
- G_f (gather) load data with index mapping f
- S_f (scatter) store data with index mapping f
- Perm_f permute data with the index mapping f

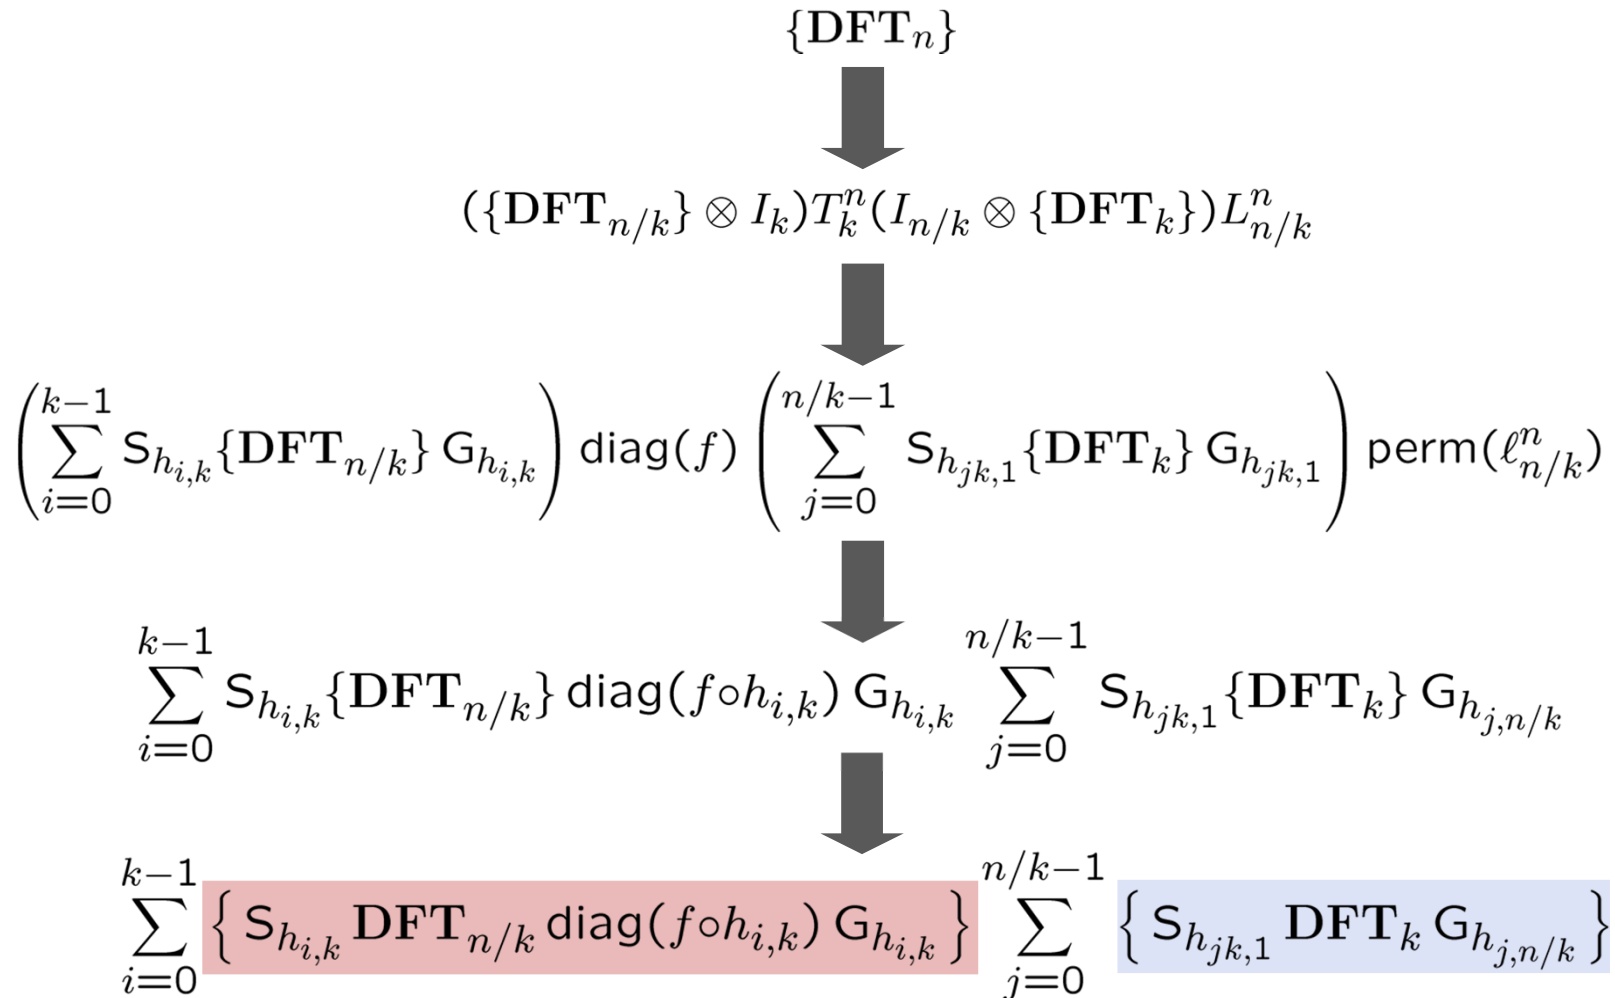
Example: $y = (I_4 \otimes F_2)x \rightarrow y = \sum_{j=0}^3 S_{f_j} F_2 G_{f_j} x$

$$y = \begin{bmatrix} F_2 & & & \\ & F_2 & & \\ & & F_2 & \\ & & & F_2 \end{bmatrix} x$$



Find Recursion Step Closure

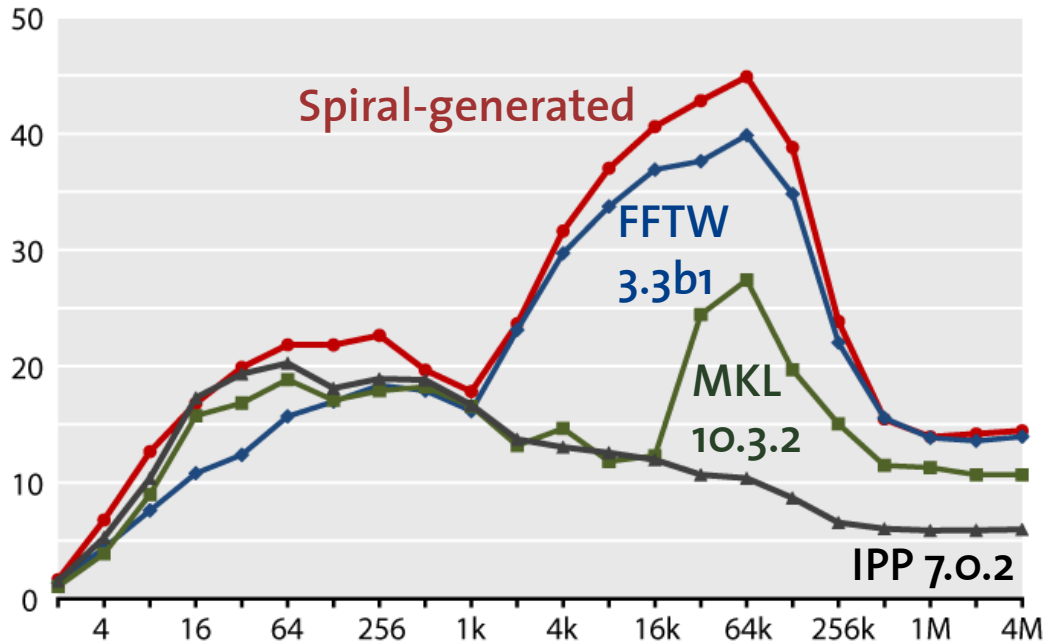
Voronenko, 2008



Repeat until closure

It Really Works

DFT on Sandybridge (3.3 GHz, 4 Cores, AVX)
Performance [Gflop/s]



$$\begin{aligned} \text{DFT}_n &\rightarrow (\text{DFT}_k \otimes I_m) T_m^n (I_k \otimes \text{DFT}_m) L_k^n \\ \text{DFT}_n &\rightarrow P_{k/2, 2m}^\top (\text{DFT}_{2m} \oplus (I_{k/2-1} \otimes_i C_{2m} \text{rDFT}_{2m}(i/k))) (\text{RDFT}_k \otimes I_m) \\ \text{RDFT}_n &\rightarrow (P_{k/2, m}^\top \otimes I_2) (\text{RDFT}_{2m} \oplus (I_{k/2-1} \otimes_i D_{2m} \text{rDFT}_{2m}(i/k))) (\text{RDFT}_k \otimes I_m) \\ \text{rDFT}_{2n}(u) &\rightarrow L_m^{2n} (I_k \otimes_i \text{rDFT}_{2m}((i+u)/k)) (\text{rDFT}_{2k}(u) \otimes I_m) \end{aligned}$$



5MB vectorized, threaded,
general-size, adaptive library

- Many transforms, often the generated code is best
- Vector, shared/distributed memory parallel, FPGAs

Online tuning

Installation

configure/make

Use

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

Search for fastest
computation strategy

Offline tuning

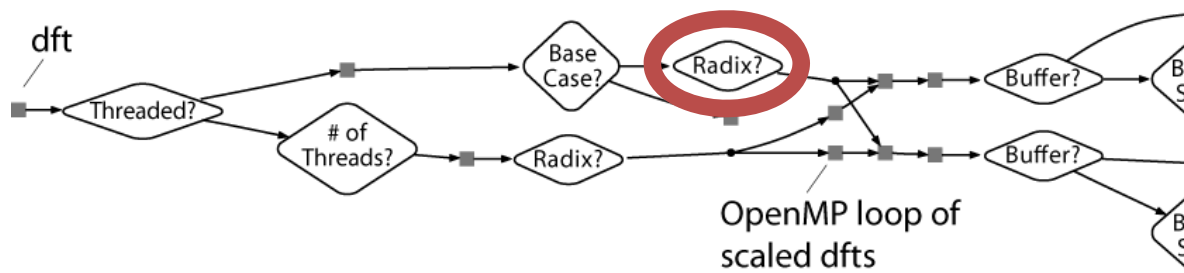
Installation

configure/make
for a few n : search
learn decision trees

Use

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

Machine learning



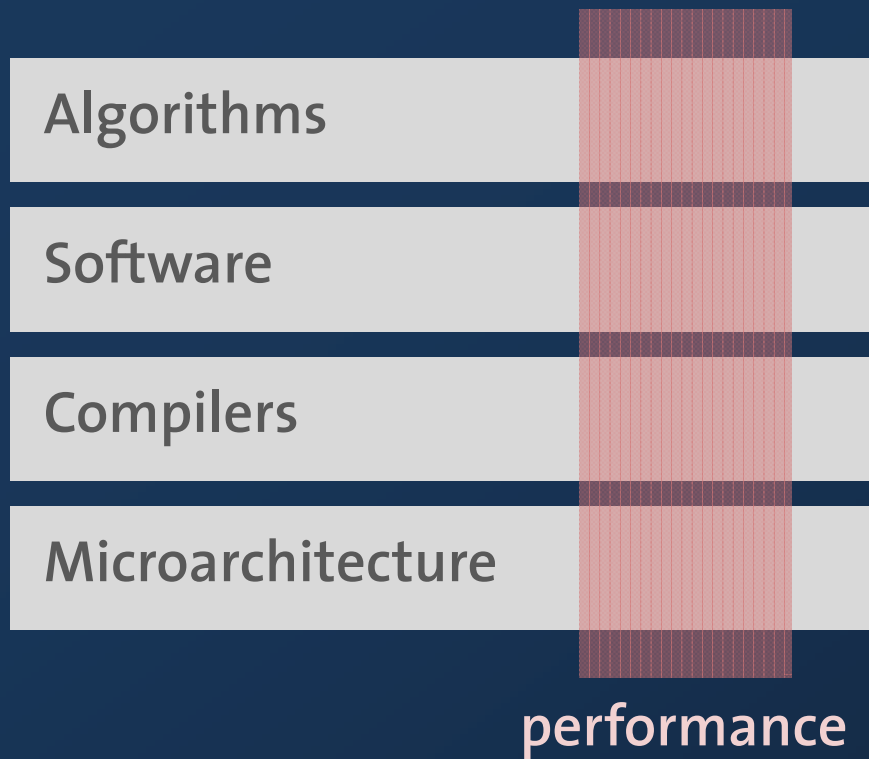
```
if ( n <= 65536 ) {  
  if ( n <= 32 ) {  
    if ( n <= 4 ) {return 2;}  
    else {return 4;}  
  }  
  else {  
    if ( n <= 1024 ) {  
      if ( n <= 256 ) {return 8;}  
      else {return 32;}  
    }  
    else {
```

Program Generators: Related Work

- FFTW codelet generator (Frigo)
- Flame (van de Geijn, Quintana-Orti, Bientinesi, ...)
- Tensor contraction engine (Baumgartner, Sadayappan, Ramanujan, ...)
- cvxgen (Mattingley, Boyd)
- PetaBricks (Ansel, Amarasinghe, ...)
- Spiral

Organization

- Software performance issues
- Automatic performance tuning (autotuning)
- Automatic performance programming
- *Conclusions*



- End of free speedup
- Straightforward code often 10-100x suboptimal
- Performance \neq efficient parallelization
- Likely inherent compiler limitations
- Performance optimization violates good programming practices

So What Do We Do?

- Teaching
- Better autotuning

Search over parameterized alternatives



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

- Example: Program generation for performance
 - Maybe Spiral can be generalized*
 - *DSLs*
 - *Rewriting for platform paradigms*
 - *Search/learning*