

$$A \otimes I_n$$


Program Synthesis For Performance

Markus Püschel
Computer Science

ETH

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

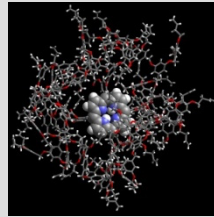
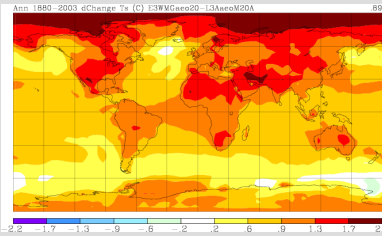
SPIRAL

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



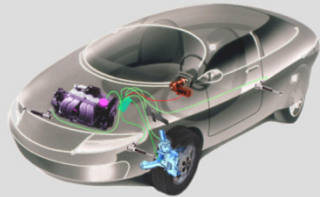
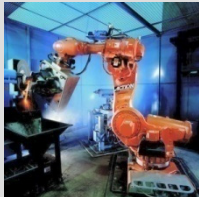
Simulations: Physics, Biology, ...

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

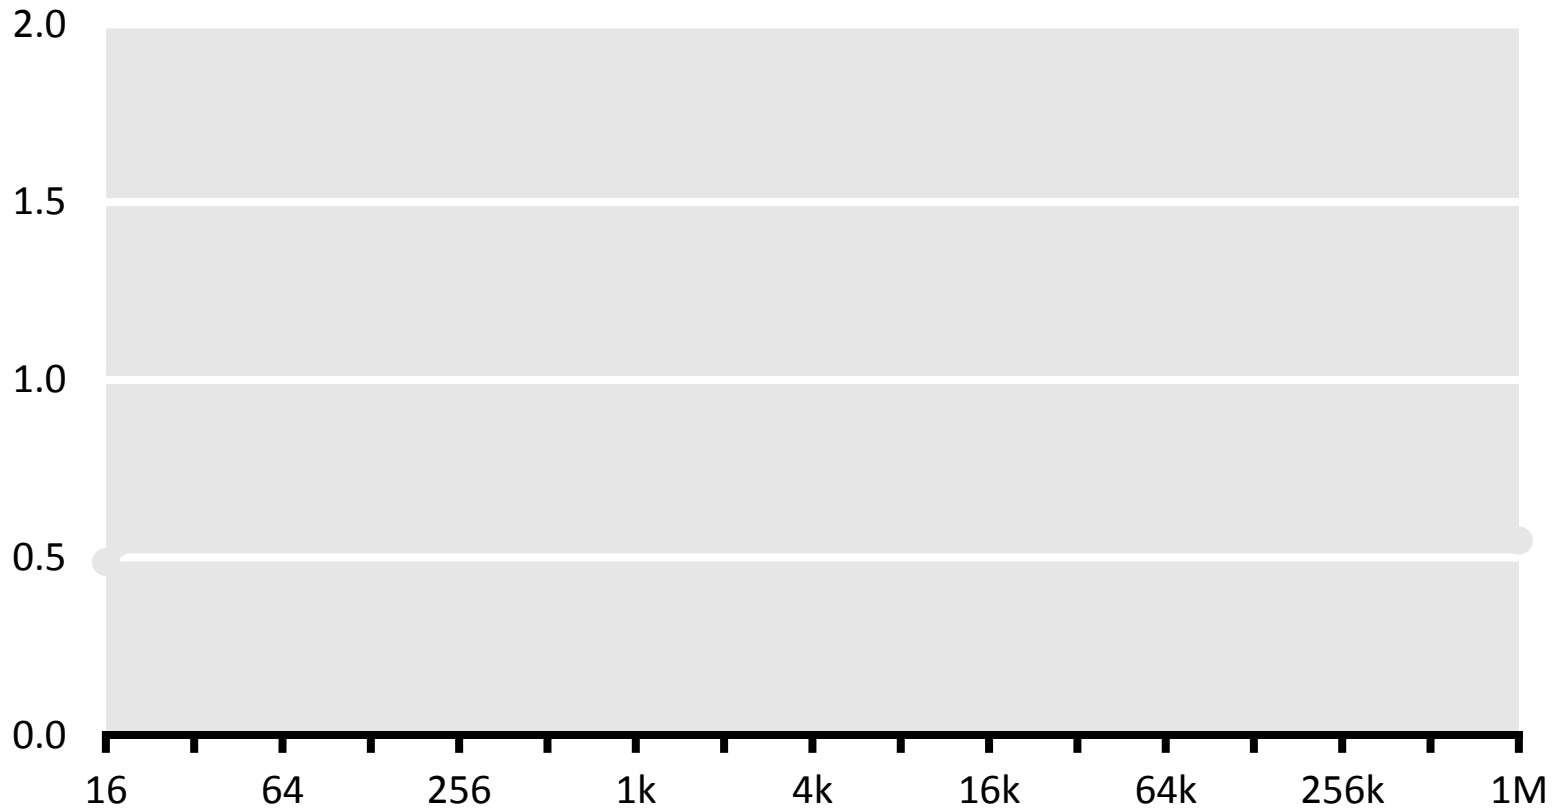
performance optimization

Microarchitecture

How well does that work?

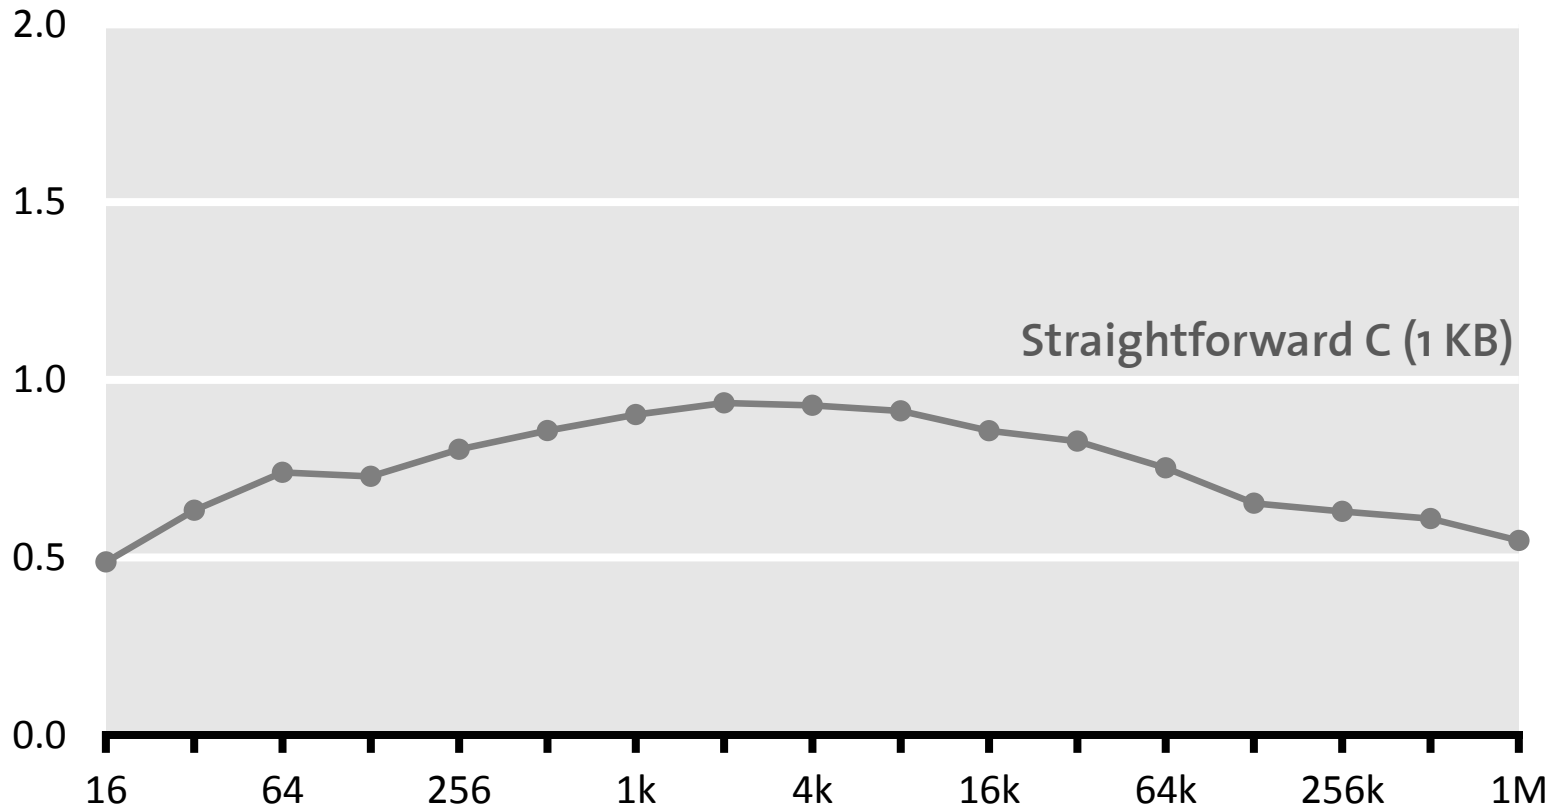
Discrete Fourier transform (single precision) on Intel Core i7 (4 cores)

Performance [Gflop/s]



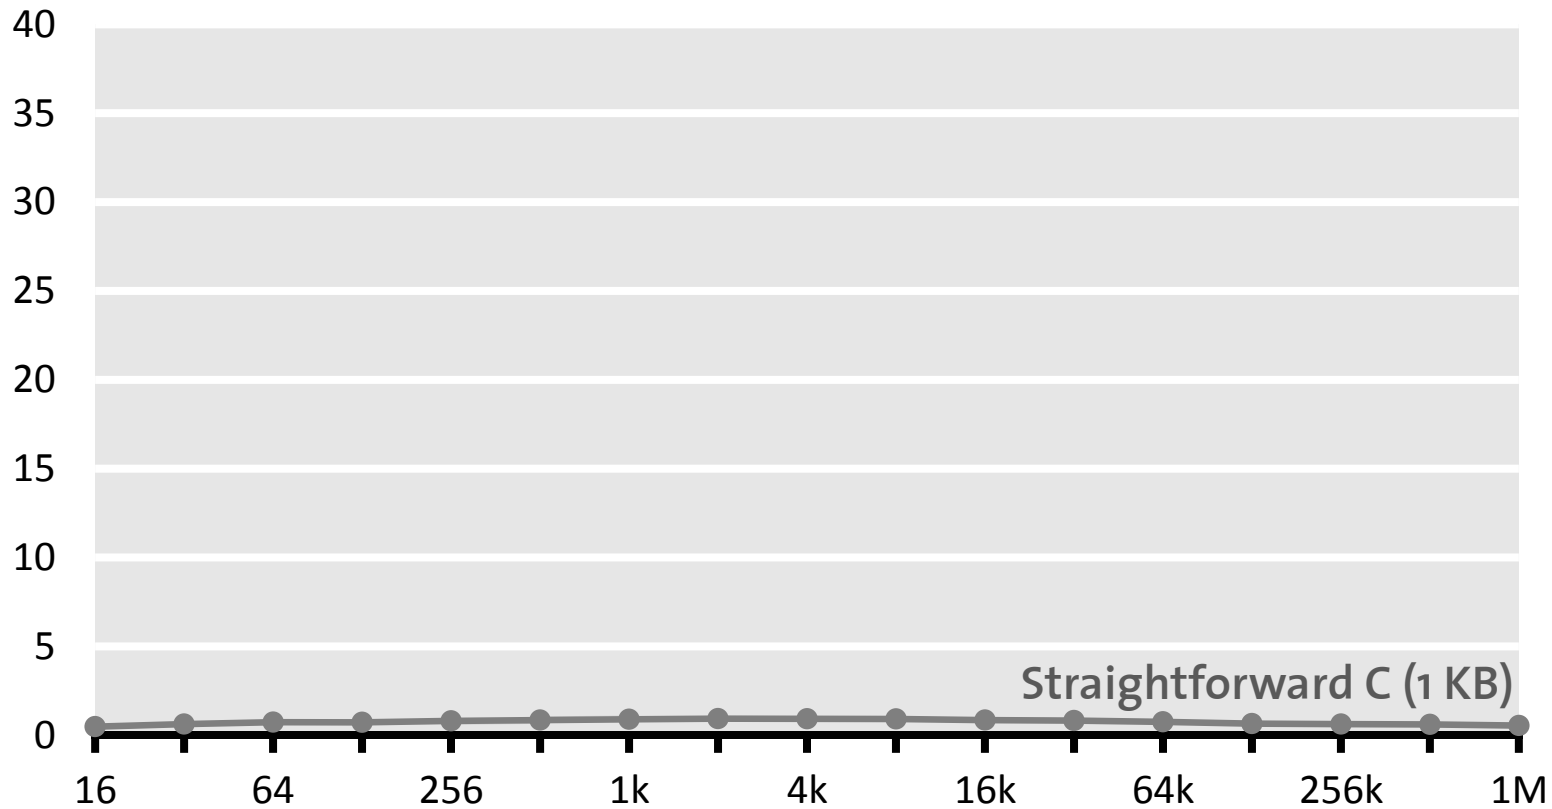
Discrete Fourier transform (single precision) on Intel Core i7 (4 cores)

Performance [Gflop/s]



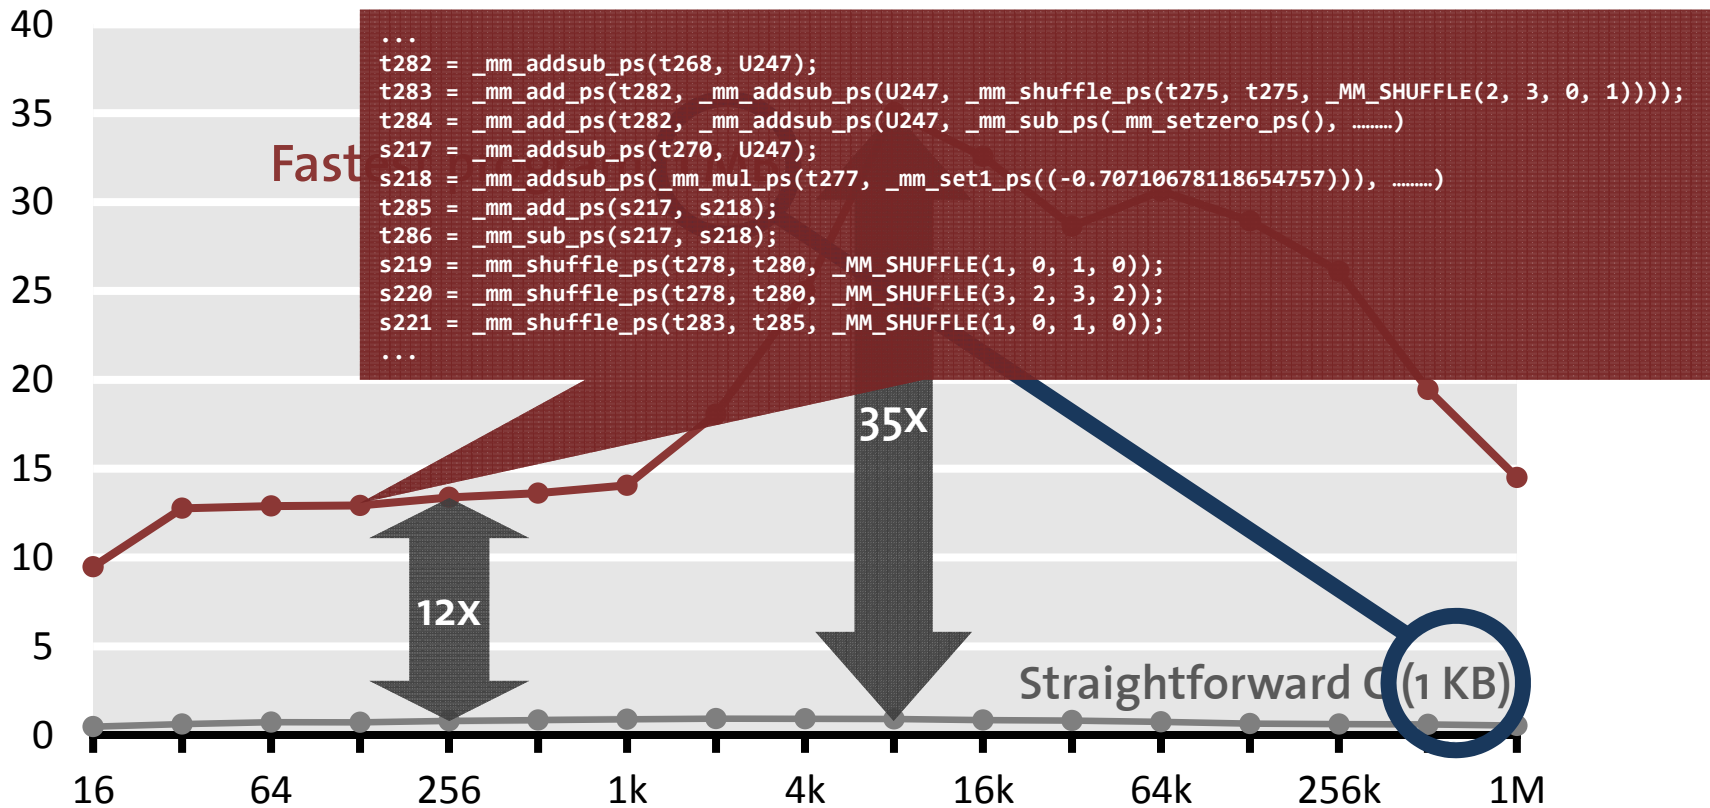
Discrete Fourier transform (single precision) on Intel Core i7 (4 cores)

Performance [Gflop/s]



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

Performance [Gflop/s]

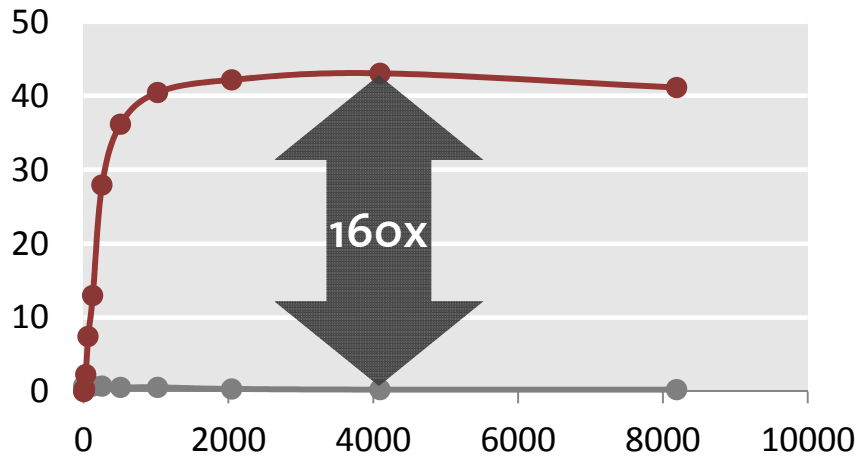


- Same opcount
- Best compiler

The Problem Is Everywhere

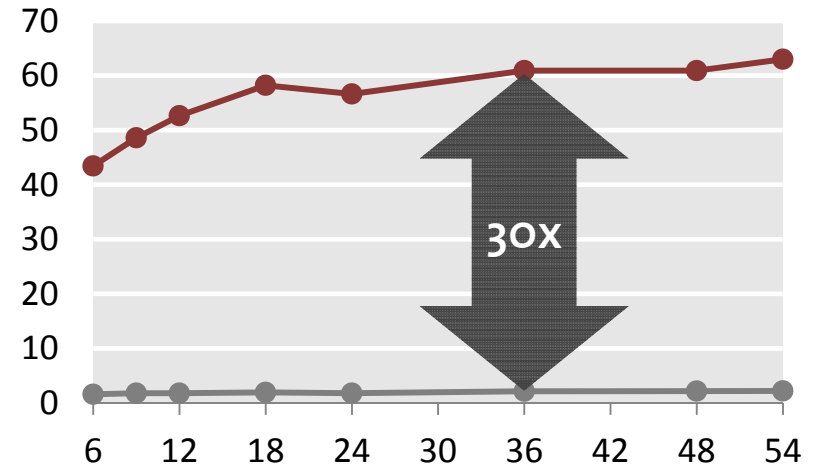
Matrix multiplication

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:”

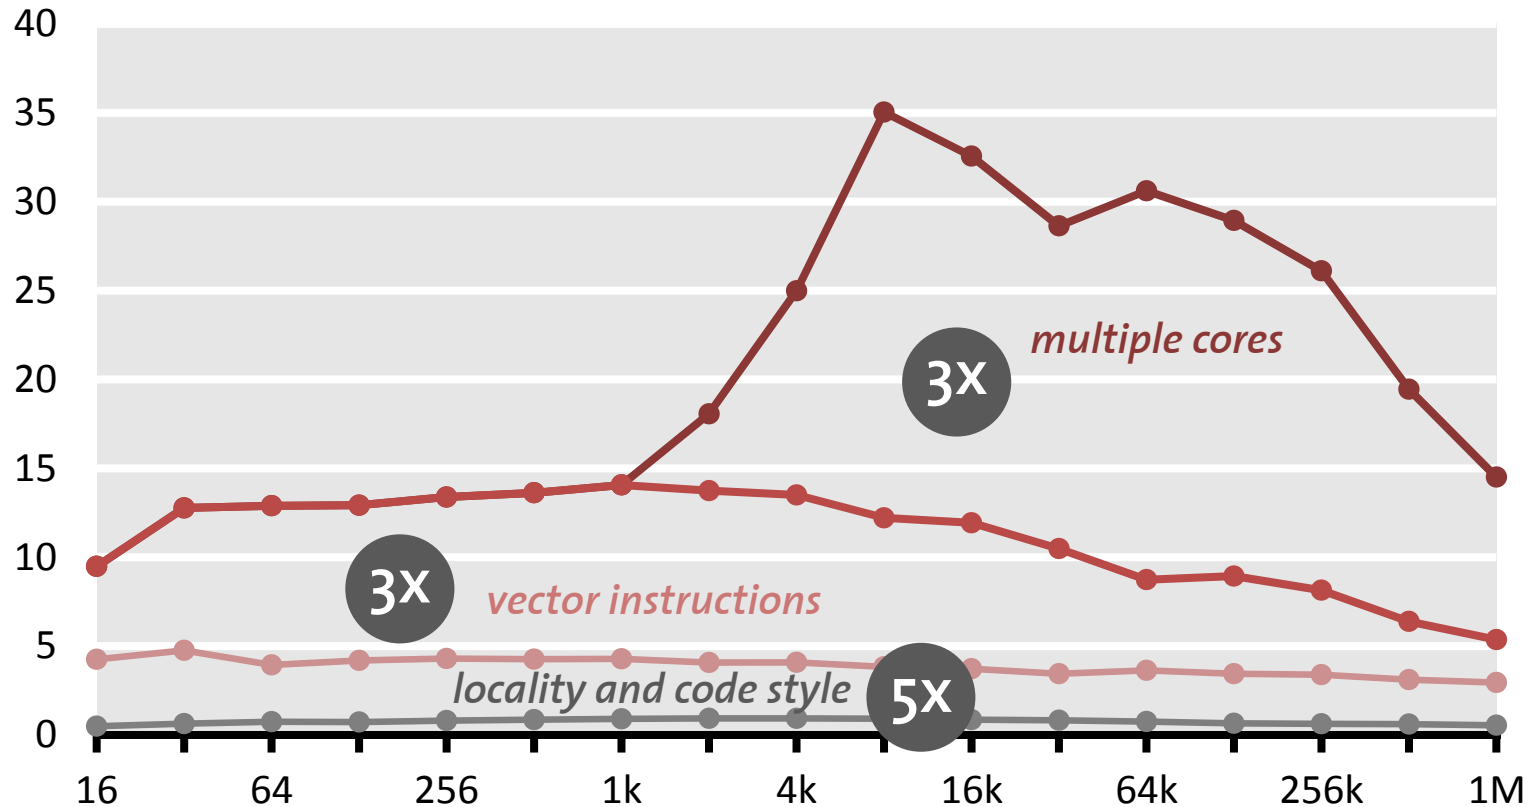
Let f be a mathematical function to be implemented on a state-of-the-art processor. Then

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

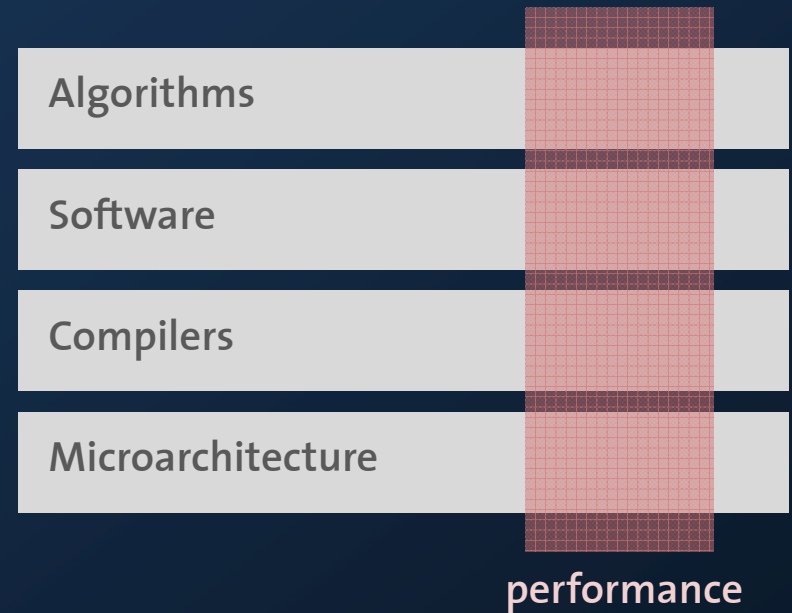
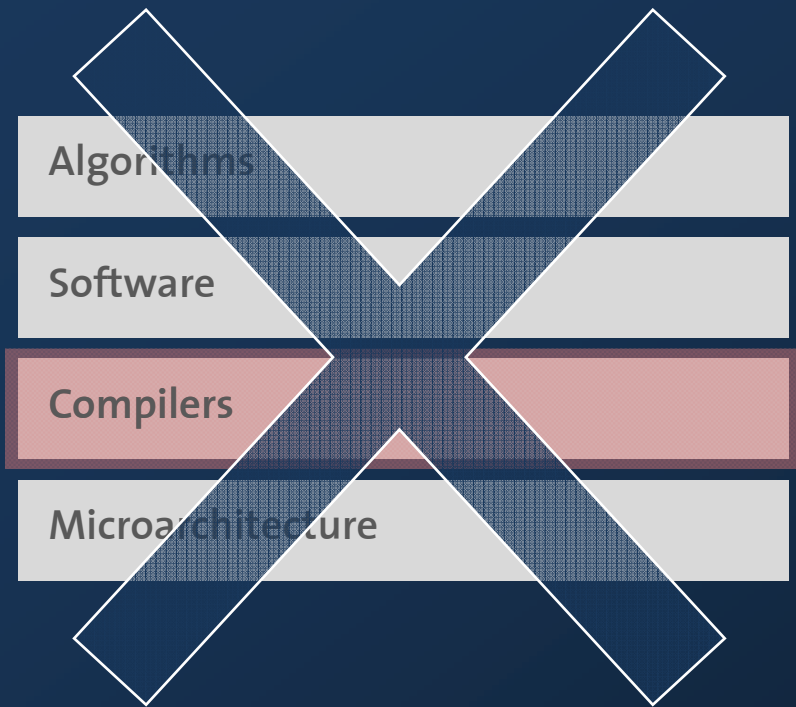
DFT: Analysis

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

Performance [Gflop/s]

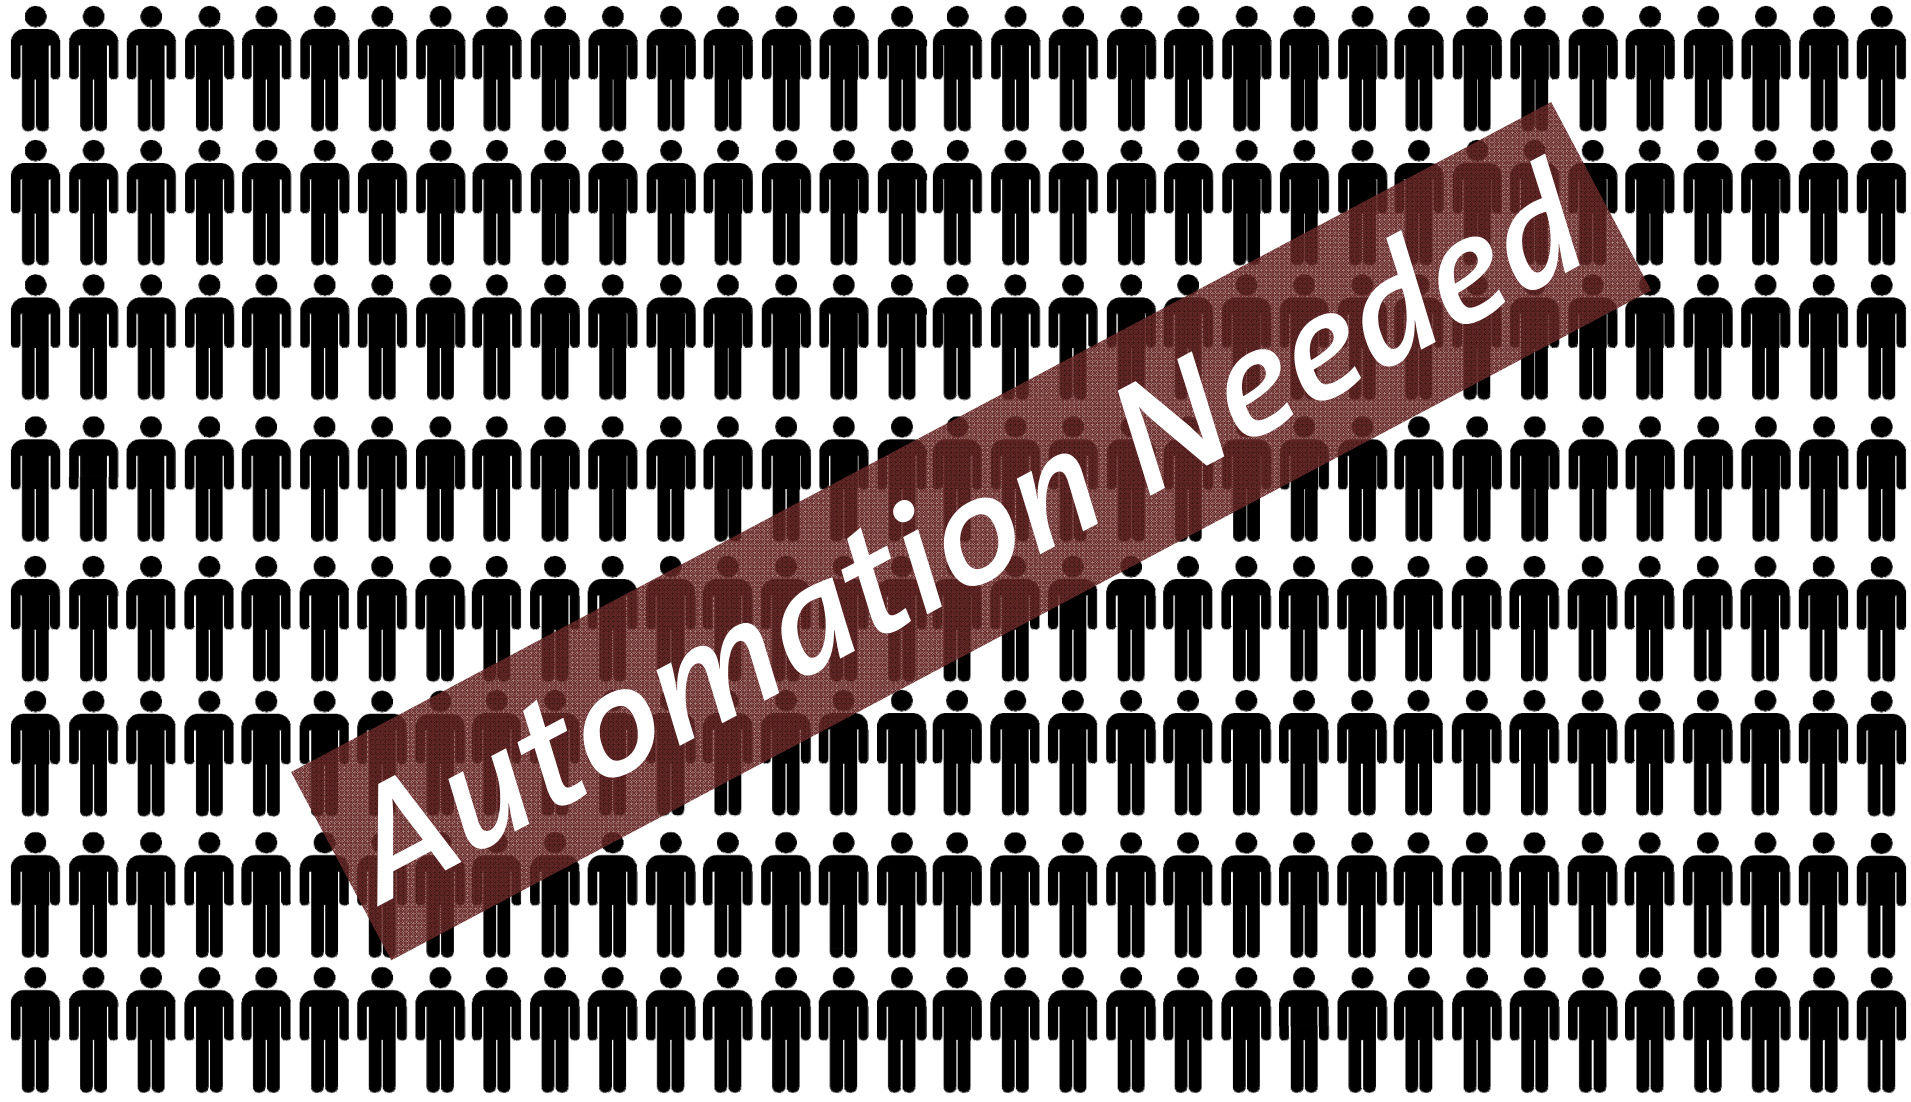


- Compiler doesn't do the job
- Doing it by hand: requires guru knowledge



Performance is different than other software quality features

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



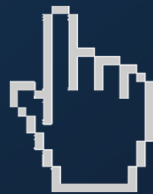
Organization

- Software performance issues
- *Synthesis of fast mathematical libraries*
- Some benchmarks
- Conclusions

Goal:

Program synthesis of high performance components

Generate Code



“click”

Select convolutional code

Select a preset code or customize parameters

- custom*
- Voyager

rate	1 / <input type="text" value="2"/>	code rate (?)
K	<input type="text" value="7"/>	constraint length (?)
polynomials	<input type="text" value="109"/>	polynomials for the code in decimal notation (?)
	<input type="text" value="79"/>	
- NASA-DSN
- CCSDS/NASA-GSFC
- WiMax
- CDMA IS-95A
- LTE (3GPP - Long Term Evolution)
- UWB (802.15)
- CDMA 2000
- Cassini
- Mars Pathfinder & Stereo

Select implementation options

frame length unpadded frame length

Vectorization level type of code [\(?\)](#)

DFT IP Cores

Viterbi Decoder

parameter	value	range	explanation
-----------	-------	-------	-------------

Problem specification

transform size	<input type="text" value="64"/>	4–32768	Number of samples (?)
direction	<input type="text" value="forward"/>		forward or inverse DFT (?)
data type	<input type="text" value="fixed point"/>		fixed or floating point (?)
	<input type="text" value="16"/> bits	4–32 bits	fixed point precision (?)
	<input type="text" value="unscaled"/>		scaling mode (?)

Parameters controlling implementation

architecture	<input type="text" value="fully streaming"/>		iterative or fully streaming (?)
radix	<input type="text" value="2"/>	2, 4, 8, 16, 32, 64	size of DFT basic block (?)
streaming width	<input type="text" value="2"/>	2–64	number of complex words per cycle (?)
data ordering	<input type="text" value="natural in / natural out"/>		natural or digit-reversed data order (?)
BRAM budget	<input type="text" value="1000"/>		maximum # of BRAMs to utilize (-1 for no limit) (?)

Possible Approach:

Capturing algorithm knowledge:

Domain-specific languages (DSLs)

Structural optimization:

Rewriting systems

High performance code style:

Compiler

Decision making for choices:

Machine learning

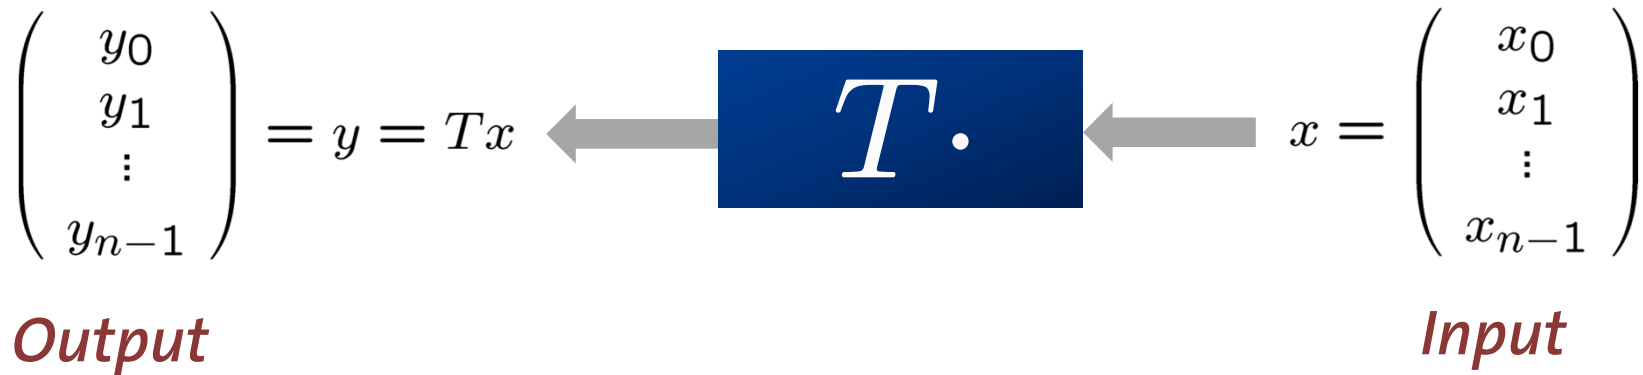
Spiral: Program Generation for Performance (www.spiral.net)



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



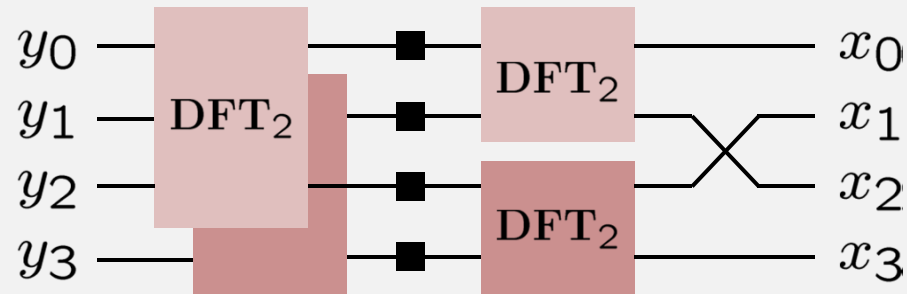
Example: $T = \text{DFT}_n = [e^{-2kl\pi i/n}]_{0 \leq k, l < n}$

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)

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

Rules = algorithm knowledge
 (≈100 journal papers)

$\text{RDFT}'_k) Q_{m/2,k}$

$$\begin{aligned} \text{DFT}_n &\rightarrow P_n (\text{DFT}'_k \otimes \text{DFT}'_m) Q_n, \quad n = km, \text{ gcd}(k, m) = 1 \\ \text{DFT}_p &\rightarrow R_p^\top (I_1 \oplus \text{DFT}_{p-1}) D_p (I_1 \oplus \text{DFT}_{p-1}) R_p, \quad p \text{ prime} \\ \text{DCT-3}_n &\rightarrow (I_m \oplus J_m) L_m^n (\text{DCT-3}_m(1/4) \oplus \text{DCT-3}_m(3/4)) \\ &\quad \cdot (F_2 \otimes I_m) \begin{bmatrix} I_m & 0 \oplus -J_{m-1} \\ \frac{1}{\sqrt{2}}(I_1 \oplus 2I_m) \end{bmatrix}, \quad n = 2m \\ \text{DCT-4}_n &\rightarrow S_n \text{DCT-2}_n \text{diag}_{0 \leq k < n} (1/(2 \cos((2k+1)\pi/4n))) \\ \text{IMDCT}_{2m} &\rightarrow (J_m \oplus I_m \oplus I_m \oplus J_m) \left(\left(\begin{bmatrix} 1 \\ -1 \end{bmatrix} \otimes I_m \right) \oplus \left(\begin{bmatrix} -1 \\ -1 \end{bmatrix} \otimes I_m \right) \right) J_{2m} \text{DCT-4}_{2m} \\ \text{WHT}_{2k} &\rightarrow \prod_{i=1}^t (I_{2^{k_1+\dots+k_{i-1}}} \otimes \text{WHT}_{2^{k_i}} \otimes I_{2^{k_{i+1}+\dots+k_t}}), \quad k = k_1 + \dots + k_t \\ \text{DFT}_2 &\rightarrow F_2 \\ \text{DCT-2}_2 &\rightarrow \text{diag}(1, 1/\sqrt{2}) F_2 \\ \text{DCT-4}_2 &\rightarrow J_2 R_{13\pi/8} \end{aligned}$$

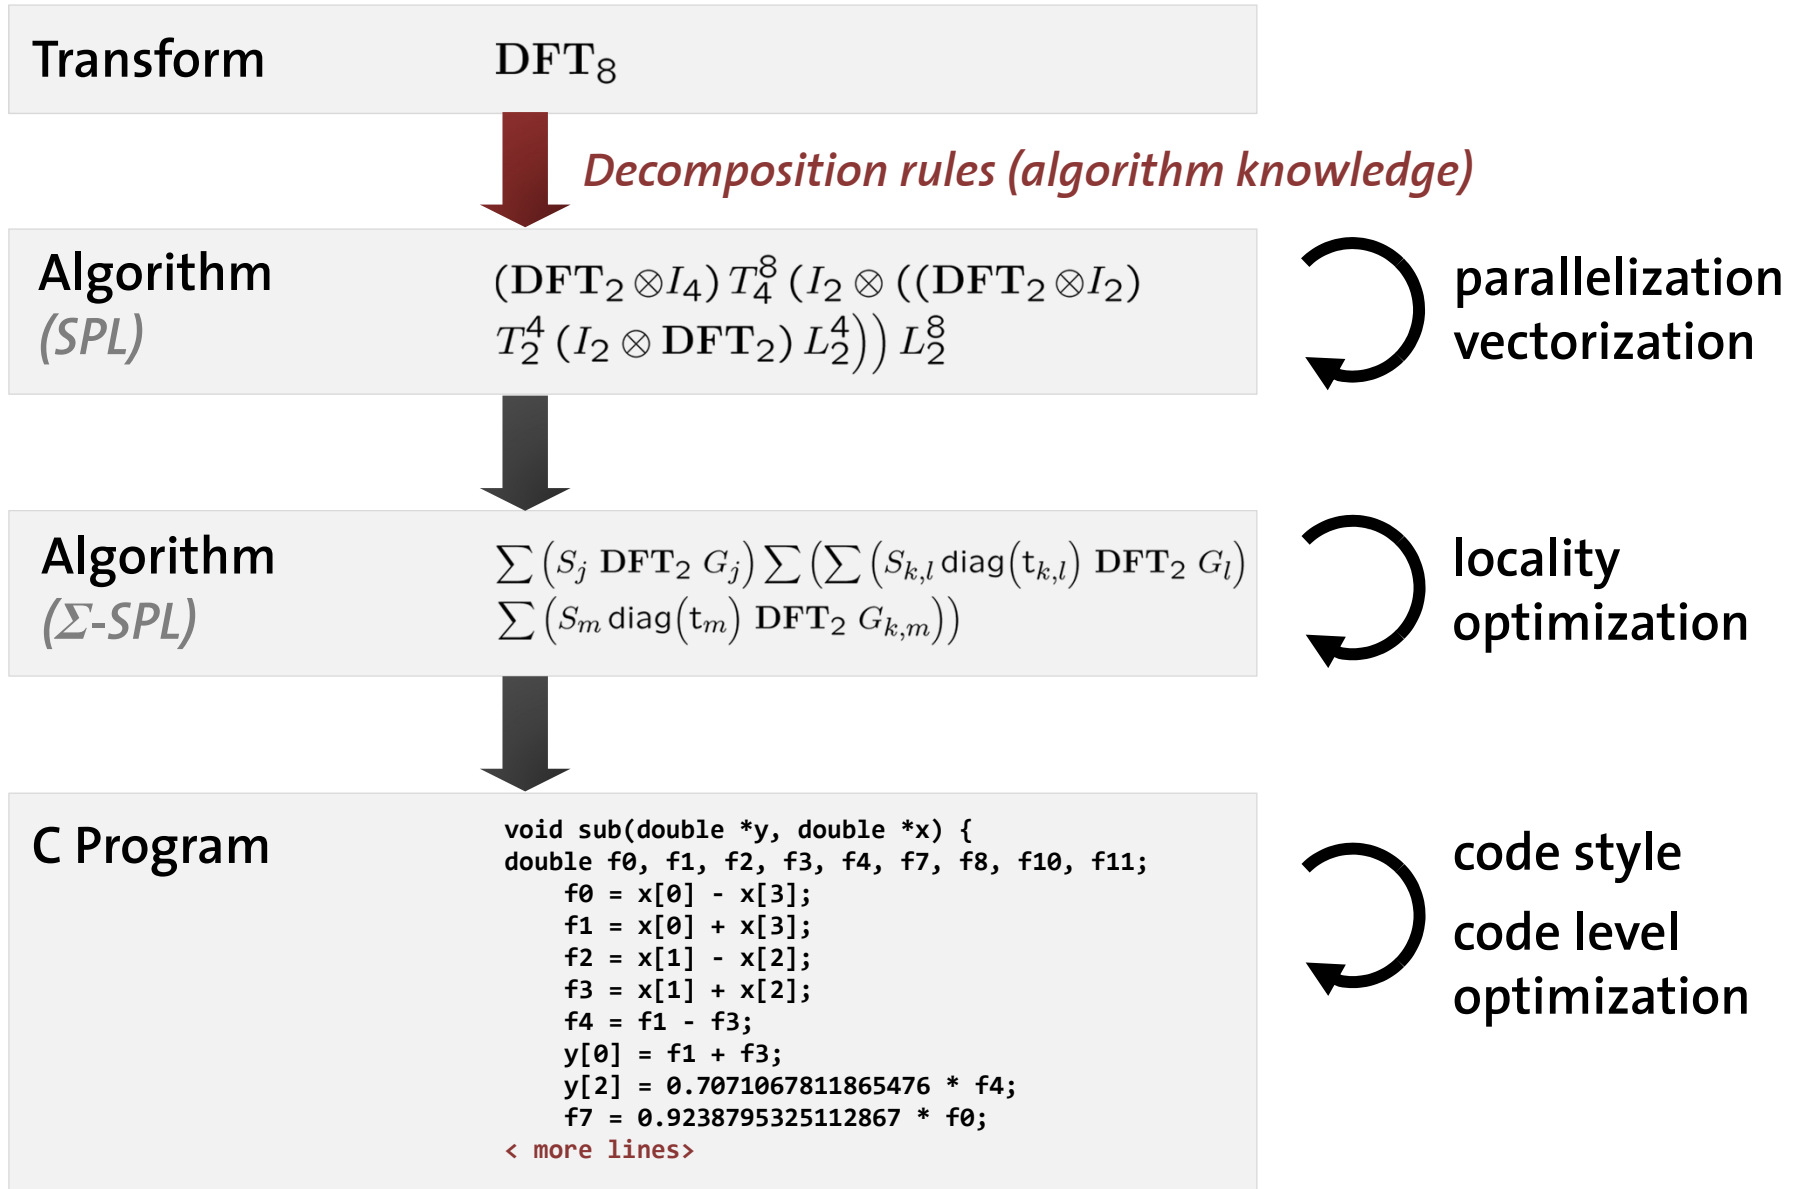
SPL to Code

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

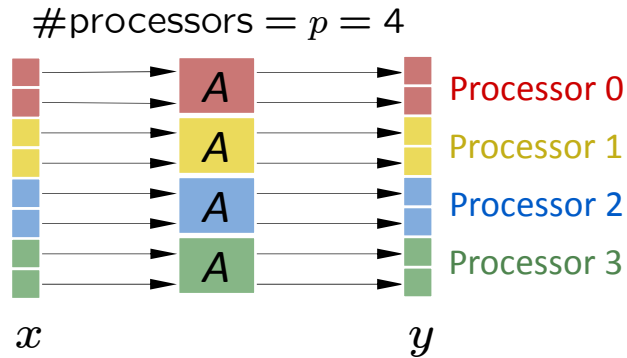
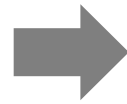
Program Generation in Spiral



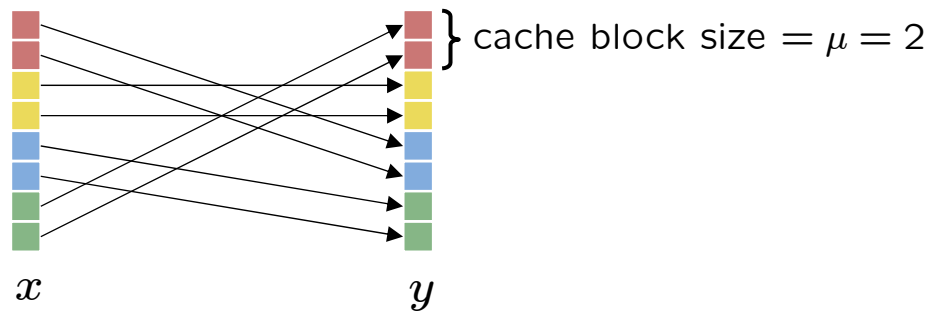
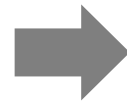
SPL to Shared Memory Code: Basic Idea

“Good” SPL structures

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



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



Rewriting: Bad structures  good structures

Example: SMP Parallelization

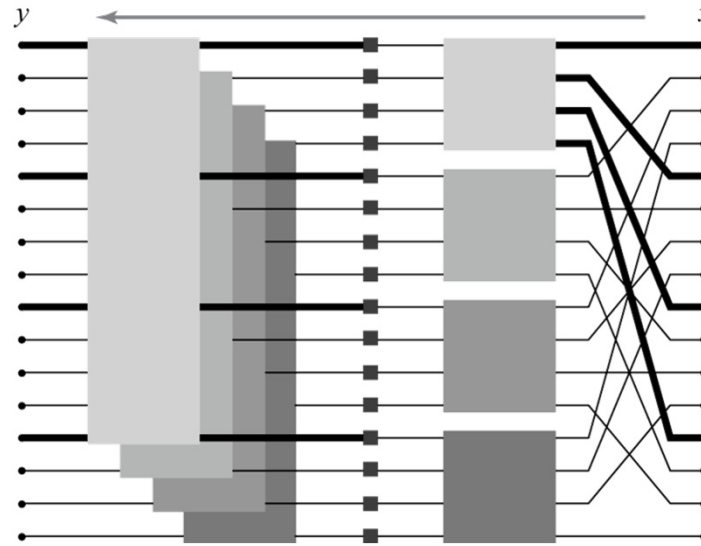
Franchetti et al., 2006

$$\begin{aligned}
 \underbrace{\text{DFT}_{mn}}_{\text{smp}(p,\mu)} &\rightarrow \underbrace{\left((\text{DFT}_m \otimes \text{I}_n) \text{T}_n^{mn} (\text{I}_m \otimes \text{DFT}_n) \text{L}_m^{mn} \right)}_{\text{smp}(p,\mu)} \\
 \dots & \\
 &\rightarrow \underbrace{\left(\text{DFT}_m \otimes \text{I}_n \right)}_{\text{smp}(p,\mu)} \underbrace{\text{T}_n^{mn}}_{\text{smp}(p,\mu)} \underbrace{\left(\text{I}_m \otimes \text{DFT}_n \right)}_{\text{smp}(p,\mu)} \underbrace{\text{L}_m^{nm}}_{\text{smp}(p,\mu)} \\
 \dots & \\
 &\rightarrow \underbrace{\left((\text{L}_m^{mp} \otimes \text{I}_{n/p\mu}) \otimes \text{I}_\mu \right)}_{\text{red}} \underbrace{\left(\text{I}_p \otimes (\text{DFT}_m \otimes \text{I}_{n/p}) \right)}_{\text{blue}} \underbrace{\left((\text{L}_p^{mp} \otimes \text{I}_{n/p\mu}) \otimes \text{I}_\mu \right)}_{\text{red}} \\
 &\quad \underbrace{\left(\bigoplus_{i=0}^{p-1} \text{T}_n^{mn,i} \right)}_{\text{blue}} \underbrace{\left(\text{I}_p \otimes (\text{I}_{m/p} \otimes \text{DFT}_n) \right)}_{\text{blue}} \underbrace{\left(\text{I}_p \otimes \text{L}_{m/p}^{mn/p} \right)}_{\text{blue}} \underbrace{\left((\text{L}_p^{pn} \otimes \text{I}_{m/p\mu}) \otimes \text{I}_\mu \right)}_{\text{red}}
 \end{aligned}$$

load-balanced, no false sharing

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

Challenge: Recursive Composition

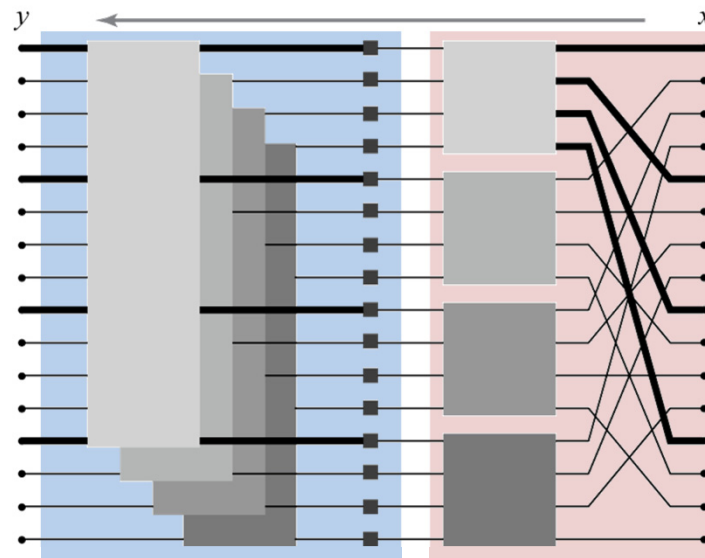


$$16 = 4 \times 4$$

$$(\text{DFT}_k \otimes \text{I}_m) \text{T}_m^{km} (\text{I}_k \otimes \text{DFT}_m) \text{L}_k^{km}$$

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

Challenge: Recursive Composition



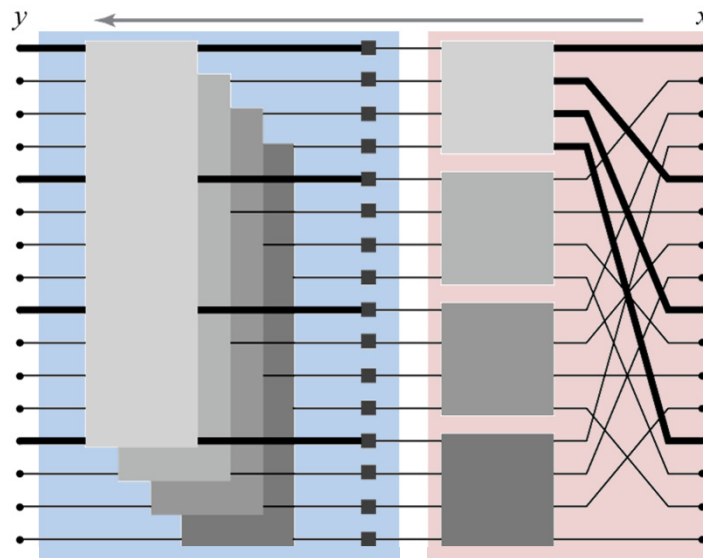
$$16 = 4 \times 4$$

$$(\text{DFT}_k \otimes \text{I}_m) \text{T}_m^{km} (\text{I}_k \otimes \text{DFT}_m) \text{L}_k^{km}$$

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

```
    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);
```

Challenge: Recursive Composition



$$16 = 4 \times 4$$

$$(DFT_k \otimes I_m) T_m^{km} (I_k \otimes DFT_m) L_k^{km}$$

```
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);
  }
}
```

How to discover these DFT variants?

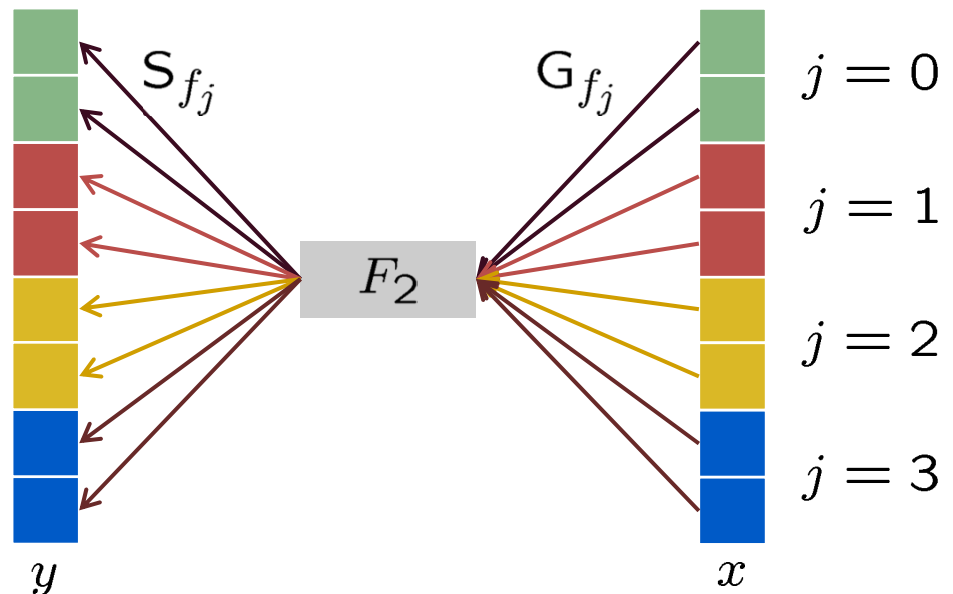
Σ -SPL : Basic Idea

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

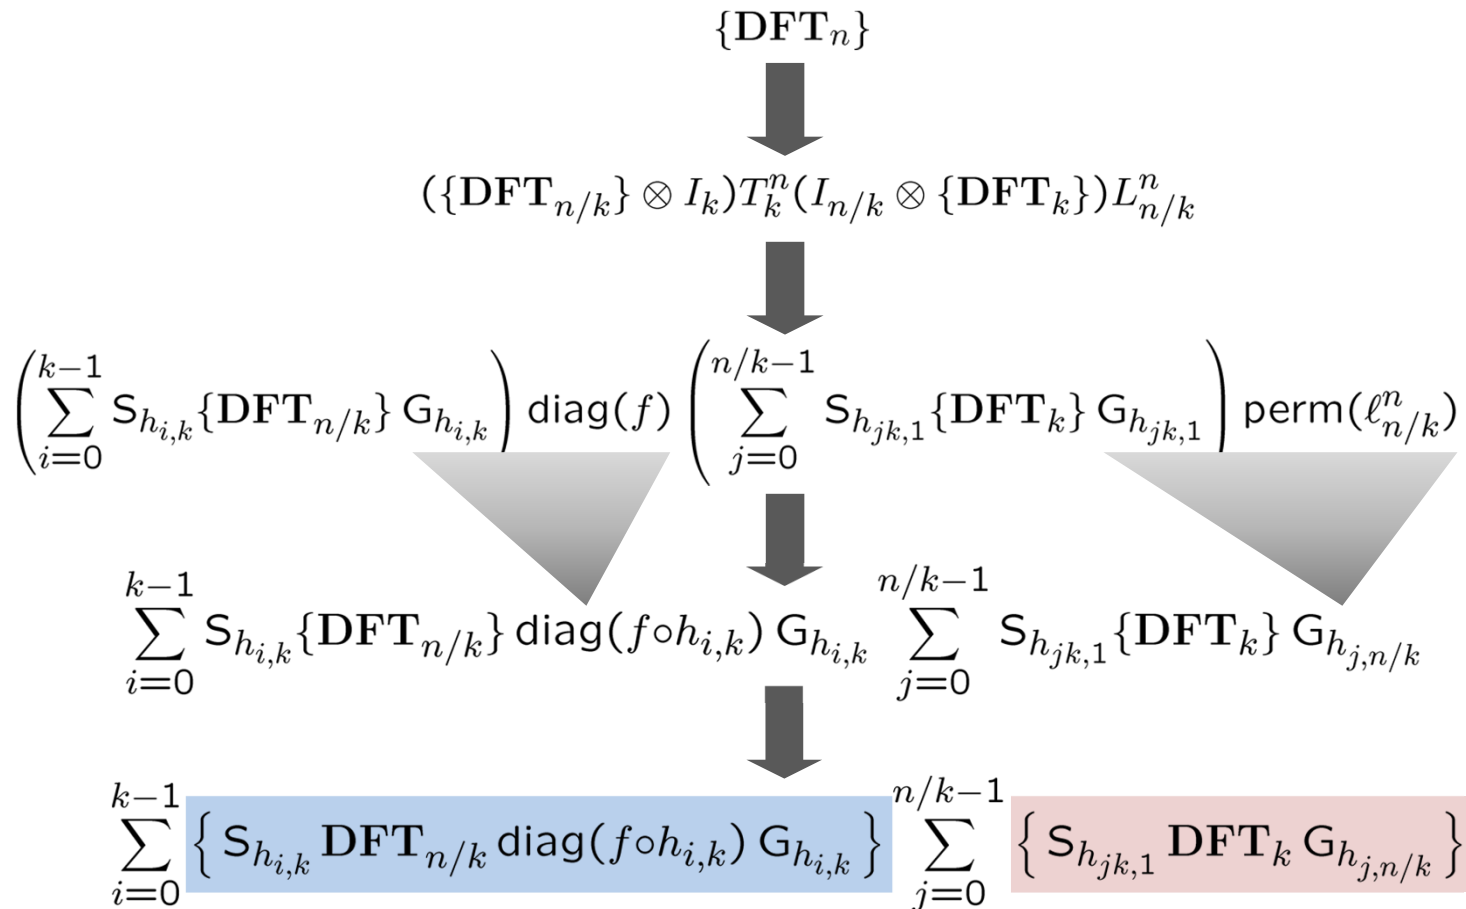
Example: $y = (I_4 \otimes F_2)x \rightarrow y = \sum_{j=0}^3 S_{f_j} F_2 G_{f_j} x$

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



Find Recursion Step Closure

Voronenko, 2008



Repeat until closure

Online tuning

Installation

configure/make

Use

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

Search for fastest recursion

Offline tuning

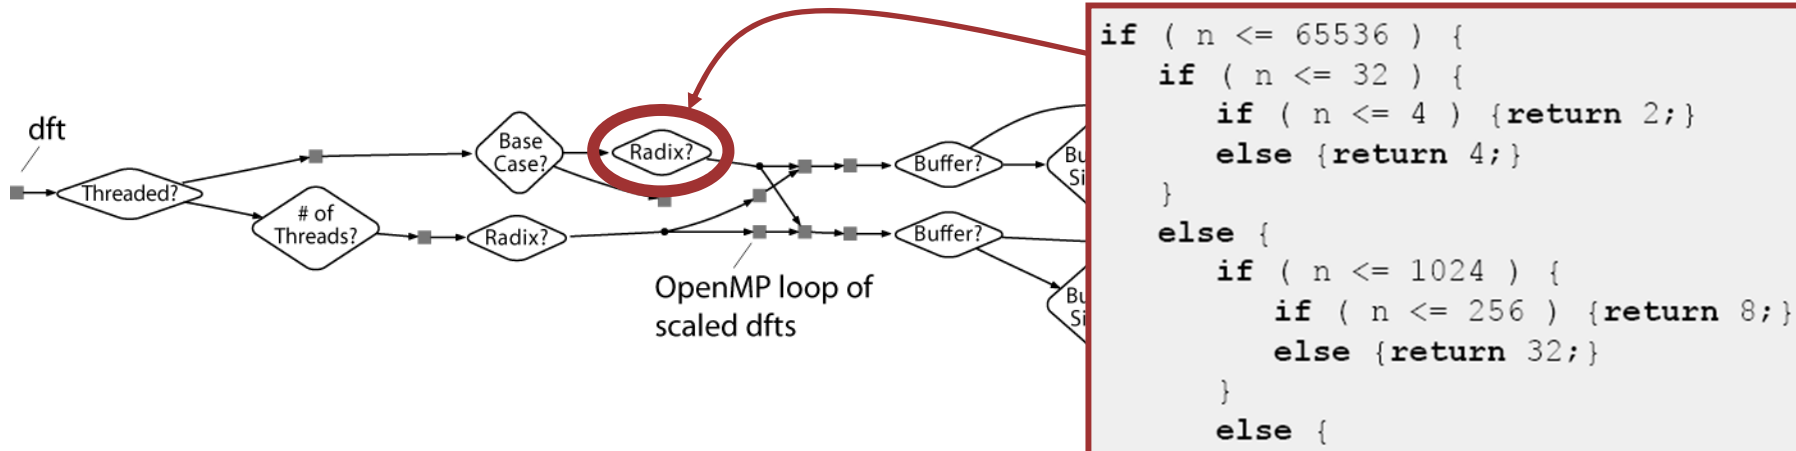
Installation

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

Use

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

Machine learning (e.g., C4.5)

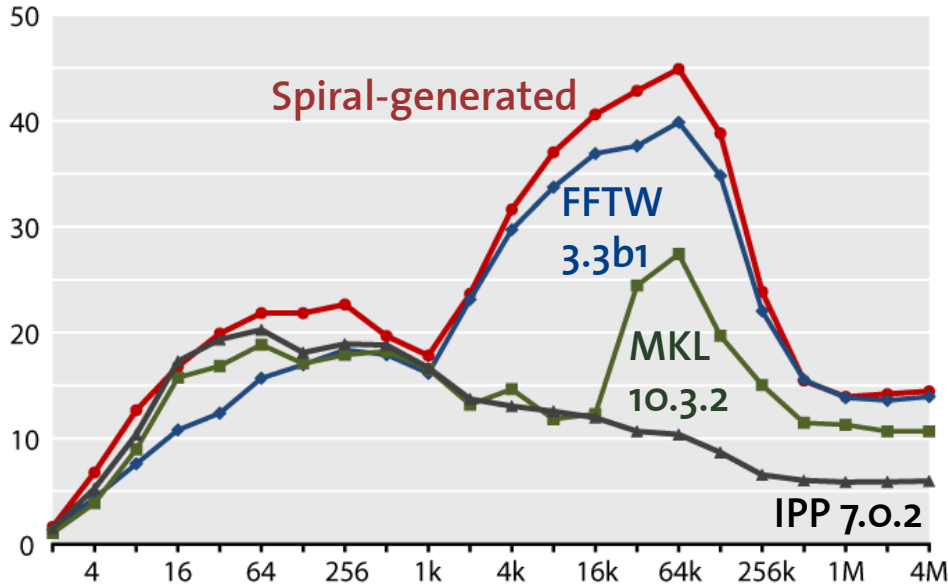


Organization

- Software performance issues
- Synthesis of fast mathematical libraries
- *Some benchmarks*
- Conclusions

It Really Works

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



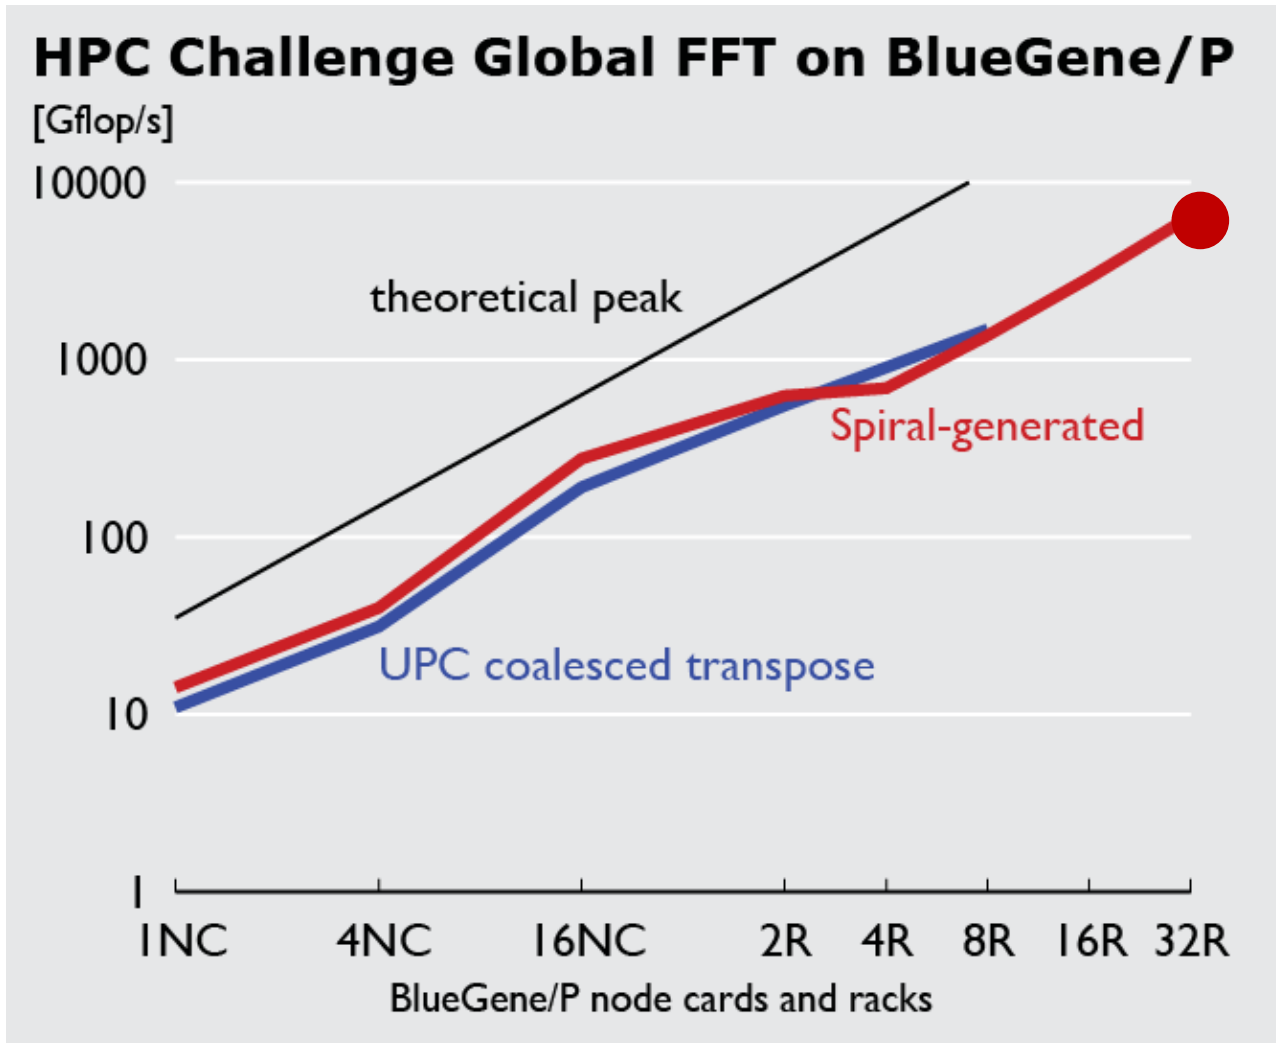
$$\begin{aligned}
 \text{DFT}_n &\rightarrow (\text{DFT}_k \otimes I_m) T_m^m (I_k \otimes \text{DFT}_m) L_k^n \\
 \text{DFT}_n &\rightarrow P_{k/2, 2m}^\top (\text{DFT}_{2m} \oplus (I_{k/2-1} \otimes_i C_{2m} \text{rDFT}_{2m}(i/k))) (\text{RDFT}_k \otimes I_m) \\
 \text{RDFT}_n &\rightarrow (P_{k/2,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



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

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

Very Large Scale: BG/P



6.4 Tflop/s

32 racks
= 32K node cards
= 128K cores

Related Work

■ Program generators for performance

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

■ Metaprogramming

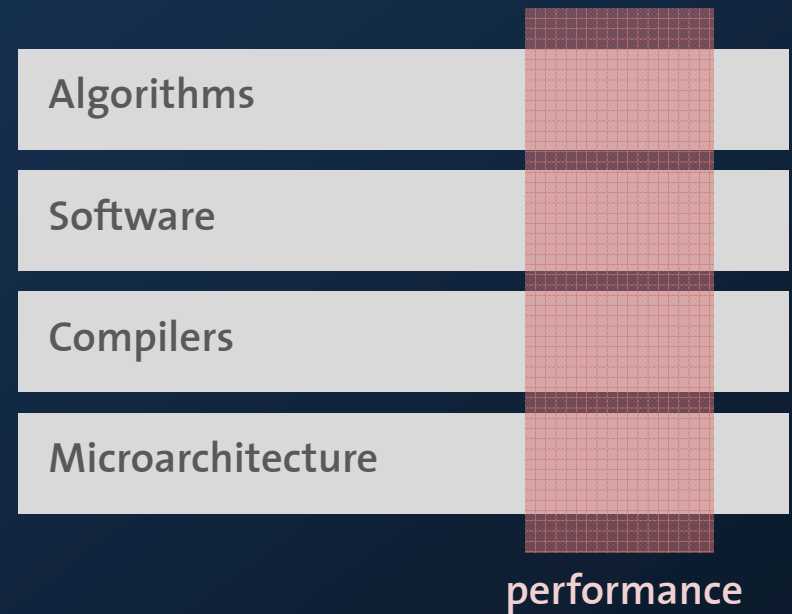
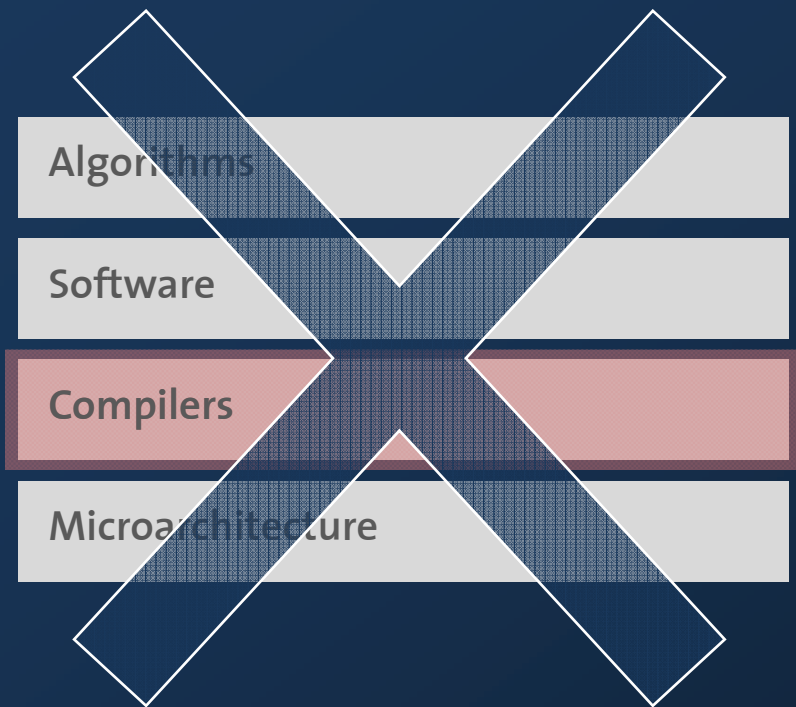
- [Eigen](#)
- [Pochoir](#)

■ Autotuning

- ATLAS/PhiPAC (Whaley, Bilmes, Demmel, Dongarra, ...)
- FFTW adaptive library (Frigo, Johnson)
- OSKI (Vuduc et al.)
- Adaptive sorting (Li et al.)

Organization

- Software performance issues
- Synthesis of fast mathematical libraries
- Some benchmarks
- *Conclusions*



*We need tools aiding the programmer or compiler
in achieving performance*

Spiral

Generate Code

Program synthesis for performance



“click”

Principles

Capturing algorithm knowledge:
DSLs

Structural optimization:
Rewriting

Decision making:
Search and learning

$$\begin{aligned}\text{DFT}_n &\rightarrow (\text{DFT}_k \otimes \text{I}_m) \text{T}_m^n (\text{I}_k \otimes \text{DFT}_m) \text{L}_k^n, \quad n = km \\ \text{DFT}_n &\rightarrow P_n (\text{DFT}_k \otimes \text{DFT}_m) Q_n, \quad n = km, \text{gcd}(k, m) = 1 \\ \text{DCT-4}_n &\rightarrow S_n \text{DCT-2}_n \text{diag}_{0 \leq k < n} (1 / (2 \cos((2k + 1)\pi / 4n))) \\ \text{IMDCT}_{2m} &\rightarrow (\text{J}_m \oplus \text{I}_m \oplus \text{I}_m \oplus \text{J}_m) \left(\left(\begin{bmatrix} 1 \\ -1 \end{bmatrix} \otimes \text{I}_m \right) \oplus \left(\begin{bmatrix} -1 \\ -1 \end{bmatrix} \otimes \text{I}_m \right) \right) \text{J}_{2m} \text{DCT-4}_{2m}\end{aligned}$$

$$\underbrace{A_m \otimes I_n}_{\text{smp}(p, \mu)} \rightarrow \underbrace{L_m^{mn}}_{\text{smp}(p, \mu)} \left(I_p \otimes_{\parallel} (I_{n/p} \otimes A_m) \right) \underbrace{L_n^{mn}}_{\text{smp}(p, \mu)}$$

More information: www.spiral.net