

Formal Loop Merging for Signal Transforms

Franz Franchetti

Yevgen S. Voronenko

Markus Püschel

Department of Electrical & Computer Engineering
Carnegie Mellon University

This work was supported by NSF through awards 0234293, and 0325687.
Franz Franchetti was supported by the Austrian Science Fund FWF,
Erwin Schrödinger Fellowship J2322.



<http://www.spiral.net>

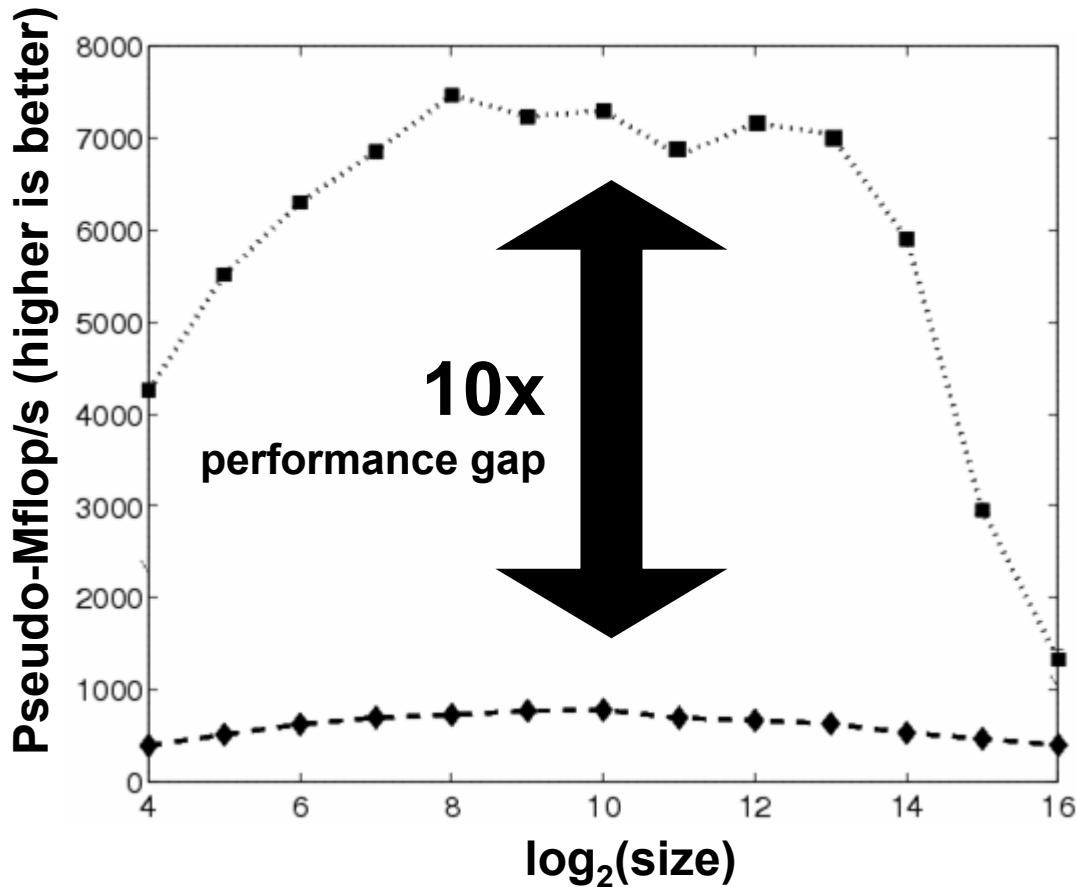
Problem

- **Runtime of (uniprocessor) numerical applications typically dominated by few compute-intensive kernels**
 - Examples: discrete Fourier transform, matrix-matrix multiplication
- **These kernels are hand-written for every architecture (open-source and commercial libraries)**
- **Writing fast numerical code is becoming increasingly **difficult**, **expensive**, and **platform dependent**, due to:**
 - Complicated memory hierarchies
 - Special purpose instructions
(short vector extensions, fused multiply-add)
 - Other microarchitectural features
(deep pipelines, superscalar execution)



Example: Discrete Fourier Transform (DFT)

Performance on Pentium 4 @ 3 GHz



vendor library



roughly the same
operations count



GNU Scientific Library

Writing fast code is hard. Are there alternatives?



Automatic Code Generation and Adaptation

- **ATLAS:** Code generator for basic linear algebra subroutines (BLAS)
[Whaley, et. al., 1998] [Yotov, et al., 2005]
- **FFTW:** Adaptive library for computing the discrete Fourier transform (DFT) and its variants
[Frigo and Johnson, 1998]
- **SPIRAL:** Code generator for linear signal transforms (including DFT)
[Püschel, et al., 2004]
- **See also:** Proceedings of the IEEE special issue on “Program Generation, Optimization, and Adaptation,” Feb. 2005.

- **Focus of this talk:**

A new approach to automatic loop merging in SPIRAL



Talk Organization

- **SPIRAL Background**
- **Automatic loop merging in SPIRAL**
- **Experimental Results**
- **Conclusions**



Talk Organization

- **SPIRAL Background**
- Automatic loop merging in SPIRAL
- Experimental Results
- Conclusions

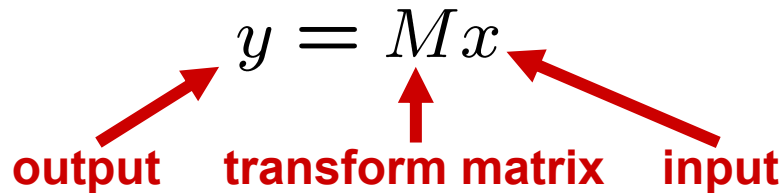


SPIRAL: DSP Transforms

- SPIRAL generates optimized code for linear signal transforms, such as discrete Fourier transform (DFT), discrete cosine transforms, FIR filters, wavelets, and many others.
- Linear transform = **matrix-vector product**:

$$y = Mx$$

output transform matrix input



- Example: DFT of input vector x

$$y = \text{DFT}_n x$$

$$\text{DFT}_n = \left[\omega_n^{kl} \right]_{0 \leq k, l < n}, \quad \omega_n = e^{-2\pi\sqrt{-1}/n}$$



SPIRAL: Fast Transform Algorithms

- Reduce computation cost from $O(n^2)$ to $O(n \log n)$
- For every transform there are **many** fast algorithms
- Algorithm = **sparse matrix factorization**

$$\begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & i & -1 & -i \\ 1 & -1 & 1 & -1 \\ 1 & -i & -1 & i \end{pmatrix} x \rightarrow \begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & -1 & 0 \\ 0 & 1 & 0 & -1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & i \end{pmatrix} \begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & -1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & -1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} x$$

12 adds
4 mults

4 adds

1 mult

(when multiplied with input vector x)

4 adds

$$\text{DFT}_4 \rightarrow (\text{DFT}_2 \otimes \text{I}_2) D (\text{I}_2 \otimes \text{DFT}_2) P$$

- SPIRAL generates the space of algorithms using **breakdown rules** in the domain-specific **Signal Processing Language (SPL)**

$$\text{DFT}_{mn} \rightarrow (\text{DFT}_m \otimes \text{I}_n) D (\text{I}_m \otimes \text{DFT}_n) P$$



SPL (Signal Processing Language)

- SPL expresses transform algorithms as structured sparse matrix factorization

- Examples:

$$F_2 = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \quad A \otimes B = [a_{k,l} B]_{k,l}$$
$$A \oplus B = \begin{bmatrix} A & \\ & B \end{bmatrix} \quad I \otimes B = \begin{bmatrix} B & & \\ & \dots & \\ & & B \end{bmatrix}$$

- SPL grammar in Backus-Naur form

$$\begin{aligned} \langle \text{spl} \rangle & ::= \langle \text{generic} \rangle \mid \langle \text{symbol} \rangle \mid \langle \text{transform} \rangle \mid \\ & \quad \langle \text{spl} \rangle \cdots \langle \text{spl} \rangle \mid \\ & \quad \langle \text{spl} \rangle \oplus \cdots \oplus \langle \text{spl} \rangle \mid \\ & \quad \langle \text{spl} \rangle \otimes \cdots \otimes \langle \text{spl} \rangle \mid \\ & \quad \dots \\ \langle \text{generic} \rangle & ::= \text{diag}(a_0, \dots, a_{n-1}) \mid \dots \\ \langle \text{symbol} \rangle & ::= I_n \mid J_n \mid L_k^n \mid R_\alpha \mid F_2 \mid \dots \\ \langle \text{transform} \rangle & ::= \mathbf{DFT}_n \mid \mathbf{WHT}_n \mid \mathbf{DCT-2}_n \mid \mathbf{Filt}_n(h[z]) \mid \dots \end{aligned}$$



Compiling SPL to Code Using Templates

$$y = \mathcal{L}_n^{mn} x$$

```
for i=0..n-1
  for j=0..m-1
    y[i+n*j]=x[m*i+j]
```

$$y = (A_n \oplus B_m)x$$

```
y[0:1:n-1] = call A(x[0:1:n-1])
y[n:1:n+m-1] = call B(x[n:1:n+m-1])
```

$$y = (I_n \otimes B_m)x$$


```
for i=0..n-1
  y[im:1:im+m-1] = call B(x[im:1:im+m-1])
```

$$y = (I_n \otimes B_m) \mathcal{L}_n^{mn} x$$

```
for i=0..n-1
  y[im:1:im+m-1] = call B(x[i:n:i+m-1])
```



Some Transforms and Breakdown Rules in SPIRAL

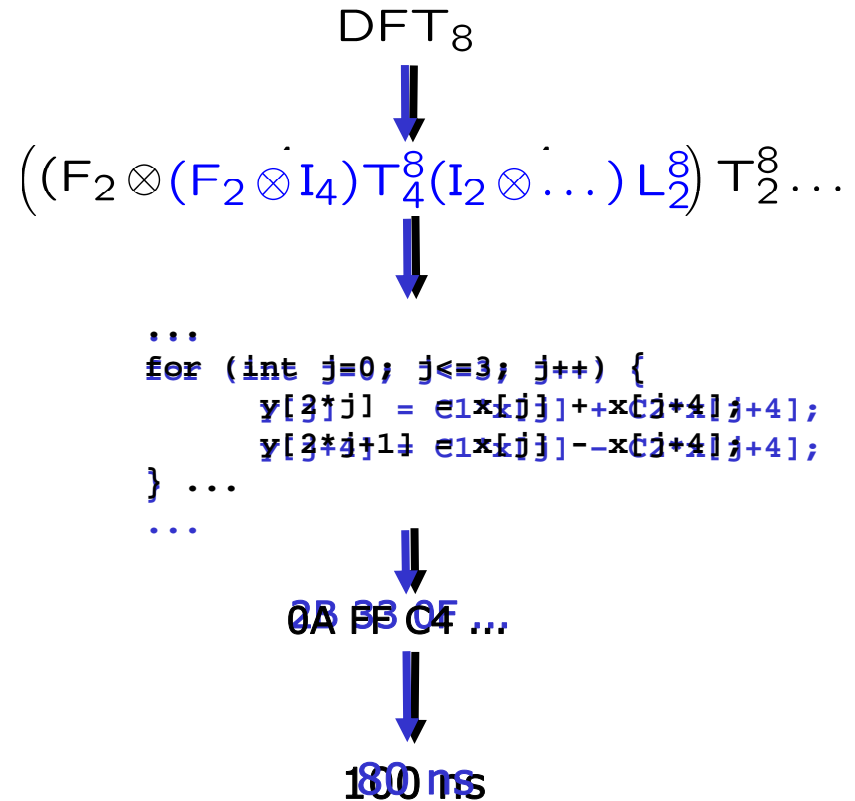
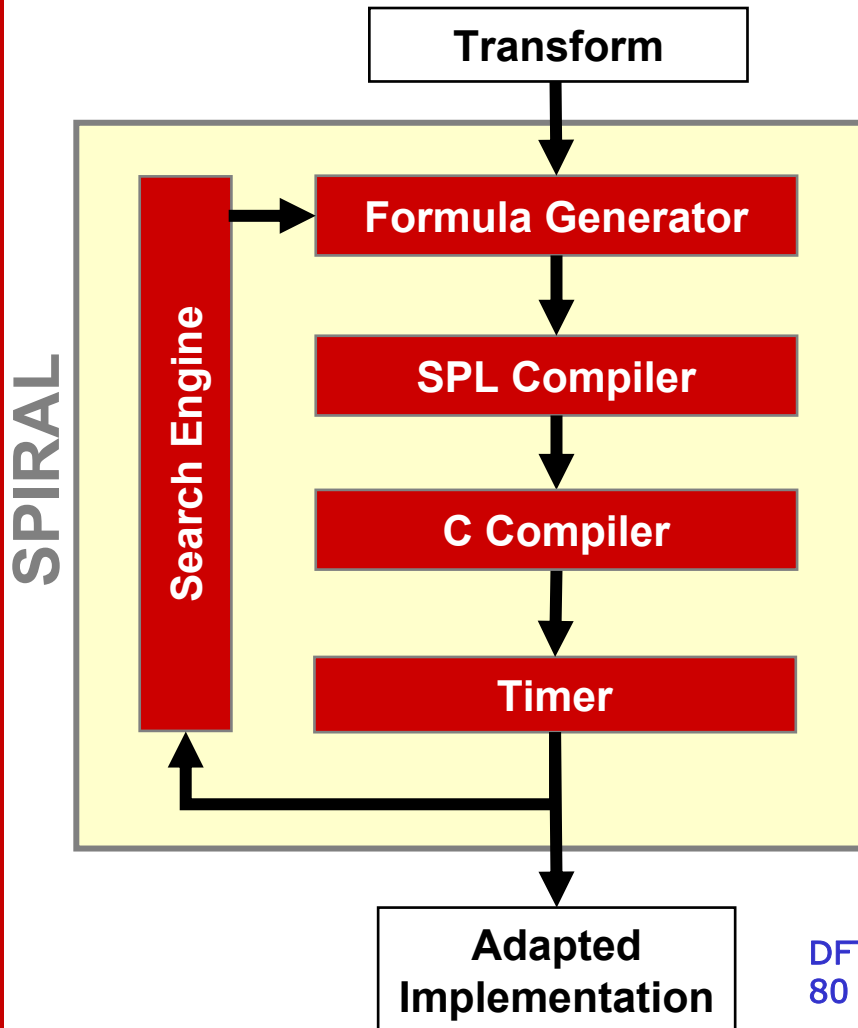


Spiral contains **30+ transforms** and **100+ rules**



SPIRAL Architecture


Approach: Empirical search over alternative recursive algorithms



DFT_8.c
80 ns



Problem: Fusing Permutations and Loops



Two passes over the working set
Complex index computation

```
void sub(double *y, double *x) {
    double t[8];
    for (int i=0; i<=7; i++)
        t[(i/4)+2*(i%4)] = x[i];
    for (int i=0; i<4; i++){
        y[2*i] = t[2*i] + t[2*i+1];
        y[2*i+1] = t[2*i] - t[2*i+1];
    }
}
```

C compiler cannot do this

One pass over the working set
Simple index computation

```
void sub(double *y, double *x) {
    for (int j=0; j<=3; j++){
        y[2*j] = x[j] + x[j+4];
        y[2*j+1] = x[j] - x[j+4];
    }
}
```

State-of-the-art
SPiRAL: Hardcoded with templates
FFTW: Hardcoded in the infrastructure

How does hardcoding scale?



General Loop Merging Problem

 = permutations

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

$$\text{DFT}_n \rightarrow P_n (\text{DFT}_k \otimes \text{DFT}_m) Q_n, \quad n = km, \text{ gcd}(k, m) = 1$$

$$\text{DFT}_p \rightarrow R_p^I (\text{I}_1 \oplus \text{DFT}_{p-1}) D_p (\text{I}_1 \oplus \text{DFT}_{p-1}) R_p, \quad p \text{ prime}$$

$$\text{DCT-3}_n \rightarrow (\text{I}_m \oplus \text{J}_m) \text{L}_m^n (\text{DCT-3}_m(1/4) \oplus \text{DCT-3}_m(3/4)) \\ \cdot (\text{F}_2 \otimes \text{I}_m) \begin{bmatrix} \text{I}_m & 0 \oplus -\text{J}_{m-1} \\ \frac{1}{\sqrt{2}}(\text{I}_1 \oplus 2\text{I}_m) \end{bmatrix}, \quad n = 2m$$

$$\text{DCT-4}_n \rightarrow S_n \text{DCT-2}_n \text{diag}_{0 \leq k < n} (1/(2 \cos((2k+1)\pi/4n)))$$

$$\text{IMDCT}_{2m} \rightarrow (\text{J}_m \oplus \text{I}_m \oplus \text{I}_m \oplus \text{J}_m) \left(\left(\begin{bmatrix} 1 \\ -1 \end{bmatrix} \otimes \text{I}_m \right) \oplus \left(\begin{bmatrix} -1 \\ -1 \end{bmatrix} \otimes \text{I}_m \right) \right) \text{J}_{2m} \text{DCT-4}_{2m}$$

- Combinatorial explosion: Implementing templates for all rules and all recursive combinations is **unfeasible**
- In many cases even theoretically not understood

Our Solution in SPIRAL

- Loop merging at C code level: **impractical**
- Loop merging at SPL level: **not possible**
- Solution:
 - New language Σ -SPL – an abstraction level between SPL and code
 - Loop merging through Σ -SPL formula manipulation

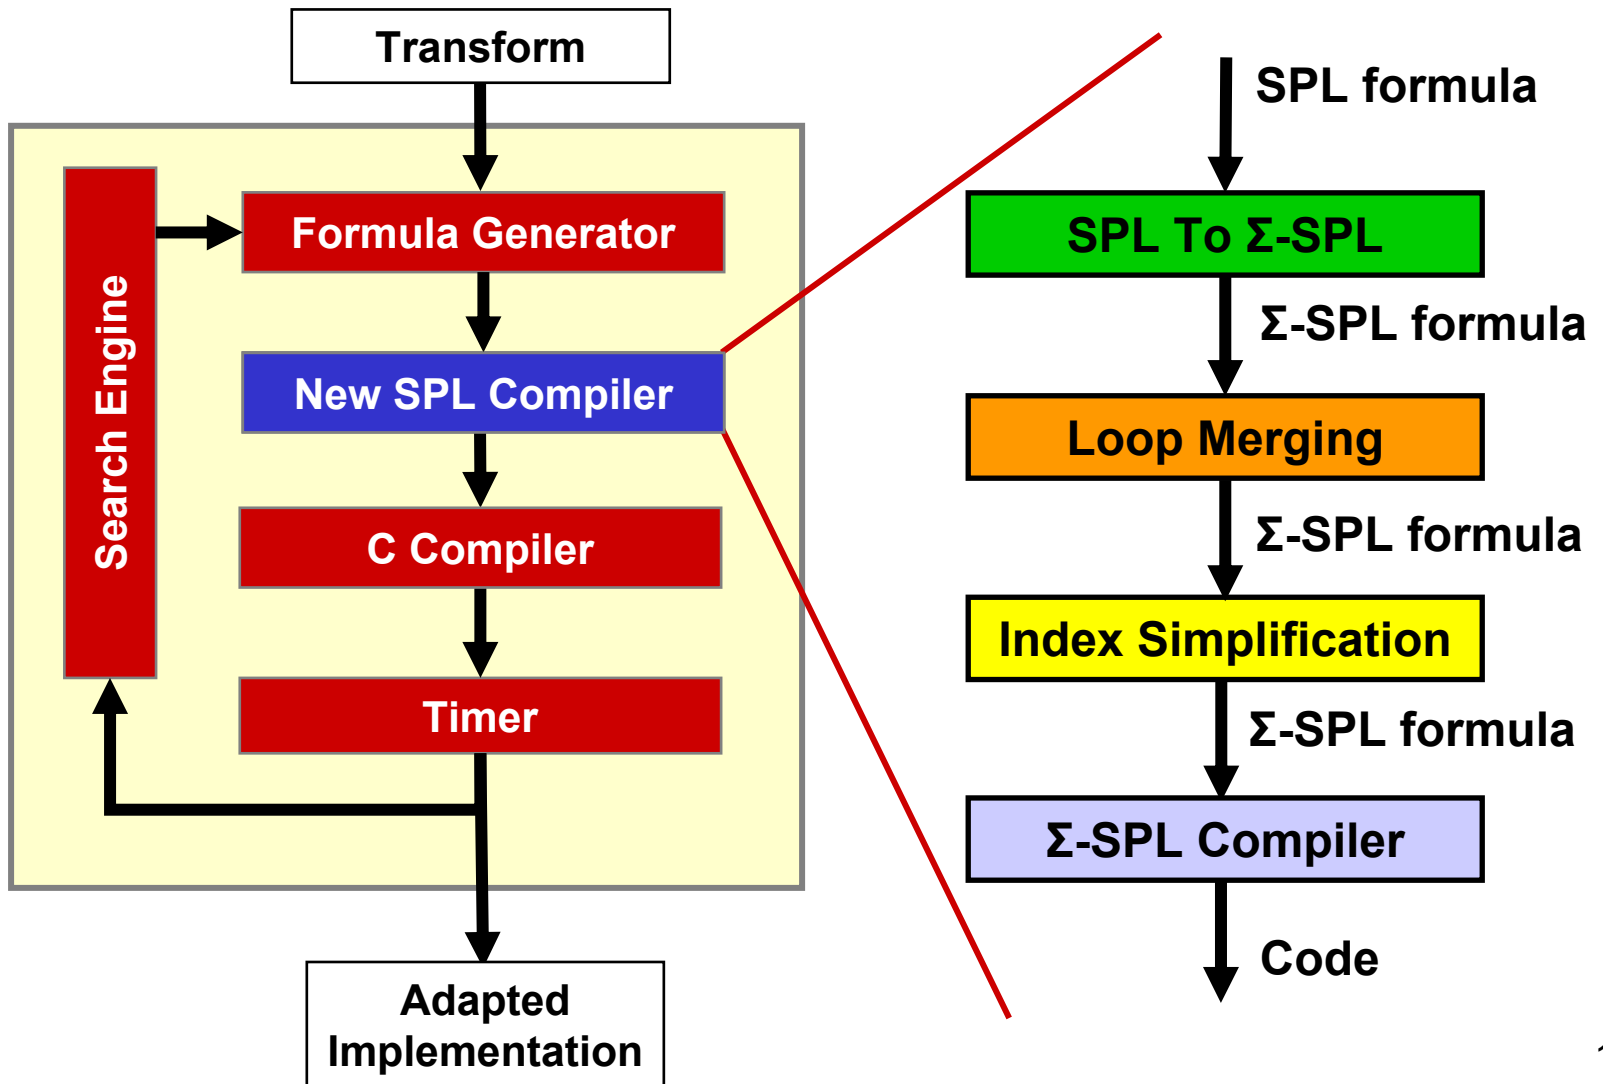


Talk Organization

- SPIRAL Background
- **Automatic loop merging in SPIRAL**
- Experimental Results
- Conclusions




New Approach for Loop Merging




Σ -SPL

- Four central constructs: Σ , G , S , Perm
 - Σ (sum) – makes loops explicit
 - G_f (gather) – reads data using the index mapping f
 - S_f (scatter) – writes data using the index mapping f
 - Perm_f – permutes data using the index mapping f
- Every Σ -SPL formula still represents a matrix factorization

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



Loop Merging With Rewriting Rules



$$y = (I_4 \otimes F_2) L_4^8 x$$

$$\left(\sum_{j=0}^3 S_{(j)_{4 \otimes 2}} F_2 G_{(j)_{4 \otimes 2}} \right) \text{Perm}_{\ell_4^8}$$

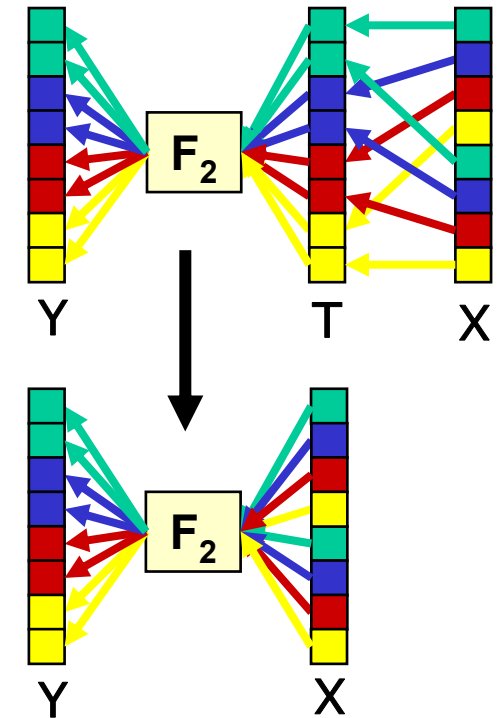
$$\left(\sum_{j=0}^3 S_{(j)_{4 \otimes 2}} F_2 G_{\ell_4^8 \circ ((j)_{4 \otimes 2})} \right)$$

$$\left(\sum_{j=0}^3 S_{(j)_{4 \otimes 2}} F_2 G_{2 \otimes (j)_4} \right)$$

```

for (int j=0; j<=3; j++) {
  y[2*j]   = x[j] + x[j+4];
  y[2*j+1] = x[j] - x[j+4];
}

```



Rules:

$$G_r \text{Perm}_p \rightarrow G_{por}$$

$$\ell_m^{mn} \circ ((j)_m \otimes v_n) \rightarrow v_n \otimes (j)_m$$


Application: Loop Merging For FFTs

DFT breakdown rules:

Cooley-Tukey FFT $\text{DFT}_{km} \rightarrow (\text{DFT}_k \otimes \text{I}_m) \text{T}_m^{km} (\text{I}_k \otimes \text{DFT}_m) \text{L}_k^{km}$

Prime factor FFT $\text{DFT}_{km} \rightarrow \text{V}_{k,m}^T (\text{DFT}_k \otimes \text{I}_m) (\text{I}_k \otimes \text{DFT}_m) \text{V}_{k,m}$
 $\text{gcd}(k, m) = 1$

Rader FFT $\text{DFT}_p \rightarrow \text{W}_p^T (\text{I}_1 \oplus \text{DFT}_{p-1}) D_p (\text{I}_1 \oplus \text{DFT}_{p-1}) \text{W}_p$
 $p - \text{prime}$

Index mapping functions are **non-trivial**:

$\text{L}_k^{km} \rightarrow \text{Perm}_{\ell_k^{km}} \quad \ell_k^{km}(i) = \left\lfloor \frac{i}{m} \right\rfloor + k(i \bmod m)$

$\text{V}_{k,m} \rightarrow \text{Perm}_{v_{k,m}} \quad v_{k,m}(i) = \left(m \left\lfloor \frac{i}{m} \right\rfloor + k(i \bmod m) \right) \bmod km$

$\text{W}_p \rightarrow \text{Perm}_{w_{\phi,g}^p} \quad w_{\phi,g}^p(i) = \begin{cases} 0, & i = 0, \\ \phi g^i \bmod p, & \text{else.} \end{cases}$



Example

Given DFT_{pq}

p – prime

p-1 = rs

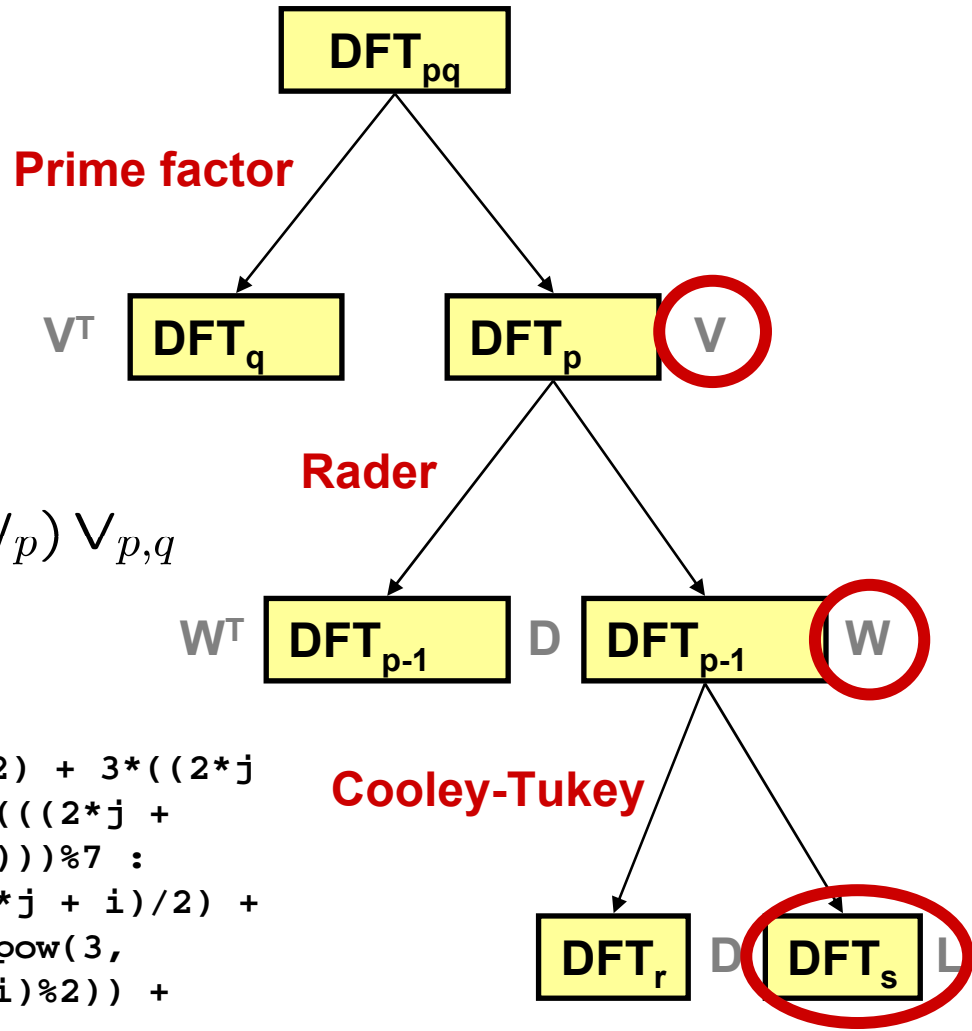
Formula fragment

$$(I_p \otimes (I_1 \oplus (I_r \otimes \text{DFT}_s) L_r^{rs}) W_p) V_{p,q}$$

Code for one memory access

p=7; q=4; r=3; s=2;

```
t=x[((21*((7*k + ((((((2*j + i)/2) + 3*((2*j + i)%2)) + 1)) ? (5*pow(3, (((2*j + i)/2) + 3*((2*j + i)%2)) + 1))))%7 : (0)))/7) + 8*((7*k + ((((((2*j + i)/2) + 3*((2*j + i)%2)) + 1)) ? (5*pow(3, (((2*j + i)/2) + 3*((2*j + i)%2)) + 1))))%7 : (0)))]%7))%28)];
```



Task: Index simplification



Index Simplification: Basic Idea

- **Example:** Identity necessary for fusing successive Rader and prime-factor step

$$\left(\varphi g^{(b+si) \bmod N'} \right) \bmod N = \left((\varphi g^b)(g^s)^i \right) \bmod N$$
$$s|N', N'|N, 0 \leq i < n$$

- Performed at the Σ -SPL level through rewrite rules on function objects:

$$\overline{w}_{\phi, g}^{N' \rightarrow N} \circ \overline{h}_{b, s}^{n \rightarrow N'} \rightarrow \overline{w}_{\phi g^b, g^s}^{n \rightarrow N}$$

- **Advantages:**
 - no analysis necessary
 - efficient (or doable at all)

Index Simplification Rules for FFTs

Cooley-Tukey

$$\begin{aligned} \ell_m^{mn} \circ ((j)_m \otimes f^{k \rightarrow n}) &\rightarrow f^{k \rightarrow n} \otimes (j)_m \\ (f^{1 \rightarrow m} \otimes h) \circ g &\rightarrow f \otimes (h \circ g) \\ (h \otimes g^{1 \rightarrow n}) \circ f &\rightarrow (h \circ f) \otimes g \\ (f_0 \otimes f_1) \circ (g_0 \otimes g_1) &\rightarrow (f_0 \circ g_0) \otimes (f_1 \circ g_1) \end{aligned}$$

Transitional

$$\begin{aligned} v_n &\rightarrow h_{0,1}^{n \rightarrow n} \\ f^{m \rightarrow M} \otimes g^{1 \rightarrow N} &\rightarrow h_{g(0),N}^{M \rightarrow MN} \circ f \\ g^{1 \rightarrow N} \otimes f^{m \rightarrow M} &\rightarrow h_{Mg(0),1}^{M \rightarrow MN} \circ f \end{aligned}$$

Cooley-Tukey + Prime factor

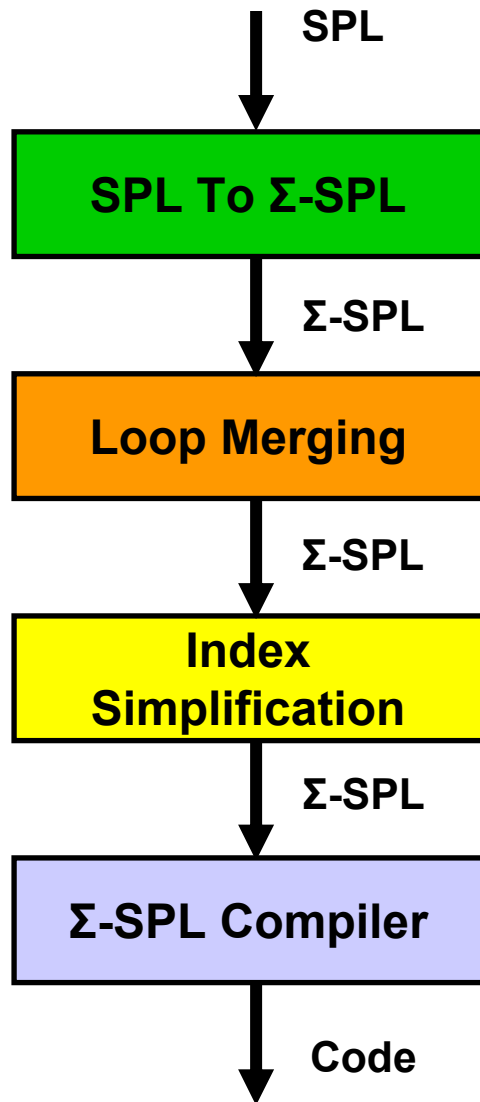
$$\begin{aligned} v_{r,s} \circ h_{b,1}^{s \rightarrow rs} &\rightarrow \bar{h}_{b,r}^{s \rightarrow rs} \\ h_{b_1,s_1}^{nk \rightarrow mnk} \circ h_{b_2,s_2}^{n \rightarrow nk} &\rightarrow h_{b_1+s_1b_2,s_1s_2}^{n \rightarrow mnk} \\ \bar{h}_{b_1,s_1}^{nk \rightarrow mnk} \circ h_{b_2,s_2}^{n \rightarrow nk} &\rightarrow \bar{h}_{b_1+s_1b_2,s_1s_2}^{n \rightarrow mnk} \\ \bar{h}_{b_1,s_1}^{nk \rightarrow mnk} \circ \bar{h}_{b_2,s_2}^{n \rightarrow nk} &\rightarrow \bar{h}_{b_1+s_1b_2,s_1s_2}^{n \rightarrow mnk} \end{aligned}$$

Cooley-Tukey + Prime factor + Rader

$$\begin{aligned} w_{\phi,g}^N \circ (0)_+^{1 \rightarrow N} &\rightarrow (0)_+^{1 \rightarrow N} \\ w_{\phi,g}^N \circ (N-1)_+^{N-1 \rightarrow N} &\rightarrow \bar{w}_{\phi,g}^{N-1 \rightarrow N} \\ \bar{w}_{\phi,g}^{N' \rightarrow N} \circ h_{b,s}^{n \rightarrow N'} &\rightarrow \bar{w}_{\phi g^b, g^s}^{n \rightarrow N} \\ \bar{w}_{\phi,g}^{N' \rightarrow N} \circ \bar{h}_{b,s}^{n \rightarrow N'} &\rightarrow \bar{w}_{\phi g^b, g^s}^{n \rightarrow N} \end{aligned}$$

**These 15 rules cover all combinations.
Some encode novel optimizations.**

Loop Merging For the FFTs : Example (cont'd)



$$(I_p \otimes (I_1 \oplus (I_r \otimes \text{DFT}_s) L_r^s) W_p) V_{p,q}$$

$$\sum_{j_1=0}^{p-1} \left(S_{((j_1)_p \otimes \iota_q) \circ (0)_+^{1 \rightarrow q}} G_{v^p, q \circ ((j_1)_p \otimes \iota_q) \circ w_{1,g}^q \circ (0)_+^{1 \rightarrow q}} \right. \\ \left. + \sum_{j_0=0}^{r-1} S_{((j_1)_p \otimes \iota_q) \circ (1)_+^{q-1 \rightarrow q} \circ ((j_0)_r \otimes \iota_s)} \text{DFT}_s \right. \\ \left. G_{v^p, q \circ ((j_1)_p \otimes \iota_q) \circ w_{1,g}^q \circ (1)_+^{q-1 \rightarrow q} \circ l_r^s \circ ((j_0)_r \otimes \iota_s)} \right)$$

```

// I
in
for
    b1 = qj1
    y1[ j1+1 ] = x1[ ( j1+1 ) % 28 ];
    p1 = 1; b1 = 7*j1;
    for(int j0 = 0; j0 <= 2; j0++) {
        y[b1 + 2*j0 + 1] = x[(b1 + 4*p1)%28] + x[(b1 + 24*p1)%28];
        y[b1 + 2*j0 + 2] = x[(b1 + 4*p1)%28] - x[(b1 + 24*p1)%28];
        p1 = (p1*3%7);
    }
}
  
```




```

// Input: _Complex double x[28], output: y[28]
double t1[28];
for(int i5 = 0; i5 <= 27; i5++)
    t1[i5] = x[(7*3*(i5/7) + 4*2*(i5%7))%28];
for(int i1 = 0; i1 <= 3; i1++) {
    double t3[7], t4[7], t5[7];
    for(int i6 = 0; i6 <= 6; i6++)
        t5[i6] = t1[7*i1 + i6];
    for(int i8 = 0; i8 <= 6; i8++)
        t4[i8] = t5[i8 ? (5*pow(3, i8))%7 : 0];
    {
        double t7[1], t8[1];
        t8[0] = t4[0];
        t7[0] = t8[0];
        t3[0] = t7[0];
    }
    {
        double t10[6], t11[6], t12[6];
        for(int i13 = 0; i13 <= 5; i13++)
            t12[i13] = t4[i13 + 1];
        for(int i14 = 0; i14 <= 5; i14++)
            t11[i14] = t12[(i14/2) + 3*(i14%2)];
        for(int i3 = 0; i3 <= 2; i3++) {
            double t14[2], t15[2];
            for(int i15 = 0; i15 <= 1; i15++)
                t15[i15] = t11[2*i3 + i15];
            t14[0] = (t15[0] + t15[1]);
            t14[1] = (t15[0] - t15[1]);
            for(int i17 = 0; i17 <= 1; i17++)
                t10[2*i3 + i17] = t14[i17];
        }
        for(int i19 = 0; i19 <= 5; i19++)
            t3[i19 + 1] = t10[i19];
    }
}
for(int i20 = 0; i20 <= 6; i20++)
    y[7*i1 + i20] = t3[i20];
}

```

```

// Input: _Complex double x[28], output: y[28]
int p1, b1;
for(int j1 = 0; j1 <= 3; j1++) {
    y[7*j1] = x[(7*j1%28)];
    p1 = 1; b1 = 7*j1;
    for(int j0 = 0; j0 <= 2; j0++) {
        y[b1 + 2*j0 + 1] = x[(b1 + 4*p1)%28] +
                               x[(b1 + 24*p1)%28];
        y[b1 + 2*j0 + 2] = x[(b1 + 4*p1)%28] -
                               x[(b1 + 24*p1)%28];
        p1 = (p1*3%7);
    }
}

```

After, 2 Loops.

Before, 11 Loops.



Talk Organization

- SPIRAL Background
- Automatic loop merging in SPIRAL
- **Experimental Results**
- Conclusions




Benchmarks Setup

- **Comparison against FFTW 3.0.1**
- **Pentium 4 3.6 GHz**
- **We consider sizes requiring at least one Rader step (sizes with large prime factor)**
- **We divide sizes into levels depending on number of Rader steps needed (Rader FFT has most expensive index mapping)**




One Rader Step

Average SPIRAL speedup: factor of 2.7




Two Rader Steps

Average SPIRAL speedup: factor of 3.3



Three Rader Steps **Average SPIRAL speedup: factor of 3.4**



Talk Organization

- SPIRAL Background
- Automatic loop merging in SPIRAL
- Experimental Results
- **Conclusions**



Conclusion

- **General loop optimization framework for linear DSP transforms in SPIRAL**
- **Loop optimization at the “right” abstraction level: Σ -SPL**
- **Application to FFT: Speedups of a factor of 2-5 over FFTW**
- **Future work: Other Σ -SPL optimizations**
 - **Loop merging for other transforms**
 - **Loop elimination, interchange, peeling**

<http://www.spiral.net>

