

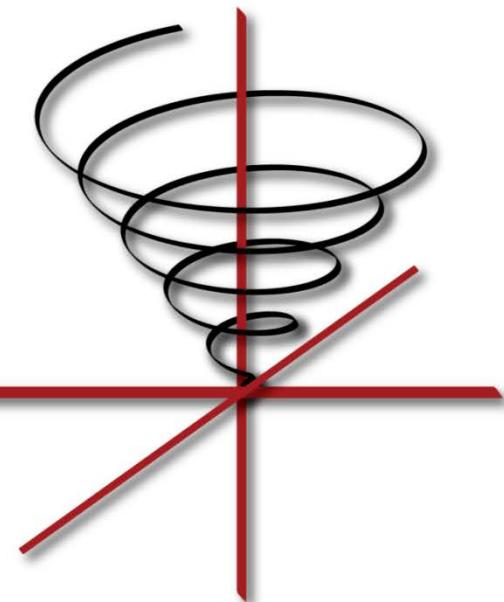
# SPIRAL

# FFT Library Generation and Autotuning

---

Franz Franchetti

Carnegie Mellon University



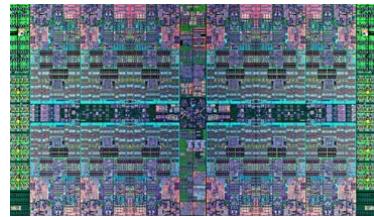
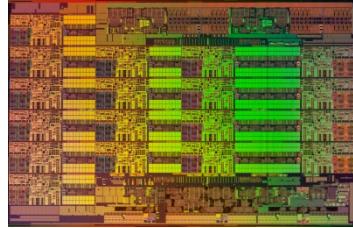
[www.spiral.net](http://www.spiral.net)

In collaboration with the SPIRAL team

This work was supported by DARPA, ONR, DOE NSF, Intel, Mercury

# Today's Computing Landscape

1 Gflop/s = one billion floating-point operations (additions or multiplications) per second



## Intel Xeon E5-2699v3

**662 Gflop/s, 145 W**

18 cores, 2.3 GHz

4-way/8-way AVX2

## IBM POWER8

**384 Gflop/s, 200 W**

12 cores, 4 GHz

2-way/4-way VMX/VSX

## Nvidia Tesla M2090

**666 Gflop/s, 225 W**

512 cores, 1.33 GHz

32-way SIMD

## Intel Xeon Phi

**1.2 Tflop/s, 300 W**

61 cores, 1.24 GHz

8-way/16-way LRBni



## Snapdragon 810

**10 Gflop/s, 2 W**

4 cores, 2.5 GHz

A330 GPU, V50 DSP, NEON



## Intel Atom C2750

**29 Gflop/s, 20 W**

8 cores, 2.4 GHz

2-way/4-way SSSE3

## Dell PowerEdge R920

**1.34 Tflop/s, 850 W**

4x 15 cores, 2.8 GHz

4-way/8-way AVX

## BlueGene/Q

**10 Pflop/s, 8 MW**

48k x 16 cores, 1.6 GHz

4-way QPX

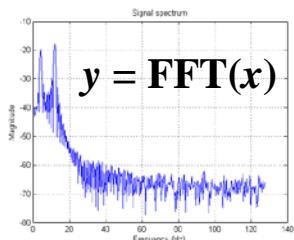
# Idea: Go from Mathematics to Software

## Given:

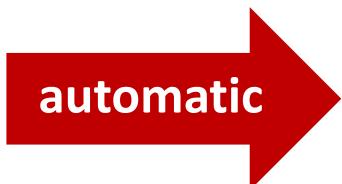
- Mathematical problem specification  
*does not change*
- Computer platform  
*changes often*

## Wanted:

- Very good implementation of specification on platform
- Proof of correctness



on

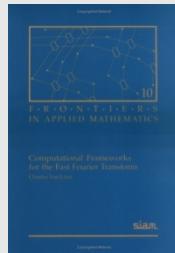
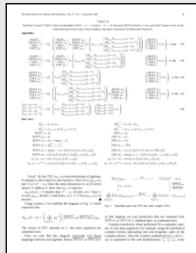
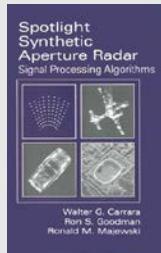


void fft64(double \*Y, double \*X) {  
 ...  
 s5674 = \_mm256\_permute2f128\_pd(s5672, s5673, (0) | ((2) << 4));  
 s5675 = \_mm256\_permute2f128\_pd(s5672, s5673, (1) | ((3) << 4));  
 s5676 = \_mm256\_unpacklo\_pd(s5674, s5675);  
 s5677 = \_mm256\_unpackhi\_pd(s5674, s5675);  
 s5678 = \*(a3738 + 16);  
 s5679 = \*(a3738 + 17));  
 s5680 = \_mm256\_permute2f128\_pd(s5678, s5679, (0) | ((2) << 4));  
 s5681 = \_mm256\_permute2f128\_pd(s5678, s5679, (1) | ((3) << 4));  
 s5682 = \_mm256\_unpacklo\_pd(s5680, s5681);  
 s5683 = \_mm256\_unpackhi\_pd(s5680, s5681);  
 t5735 = \_mm256\_add\_pd(s5676, s5682);  
 t5736 = \_mm256\_add\_pd(s5677, s5683);  
 t5737 = \_mm256\_add\_pd(s5670, t5735);  
 t5738 = \_mm256\_add\_pd(s5671, t5736);  
 t5739 = \_mm256\_sub\_pd(s5670, \_mm256\_mul\_pd(\_mm\_vbroadcast\_sd(&(C22)), t5735));  
 t5740 = \_mm256\_sub\_pd(s5671, \_mm256\_mul\_pd(\_mm\_vbroadcast\_sd(&(C22)), t5736));  
 t5741 = \_mm256\_mul\_pd(\_mm\_vbroadcast\_sd(&(C23)), \_mm256\_sub\_pd(s5677, s5683));  
 t5742 = \_mm256\_mul\_pd(\_mm\_vbroadcast\_sd(&(C23)), \_mm256\_sub\_pd(s5676, s5682));  
 ...  
}



# What is Spiral?

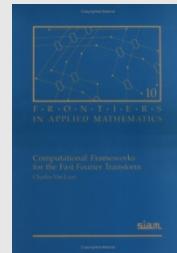
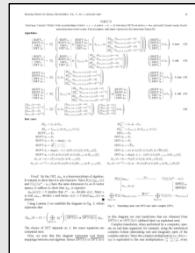
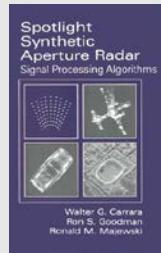
*Traditionally*



High performance library  
optimized for given platform

*Comparable  
performance*

*Spiral Approach*

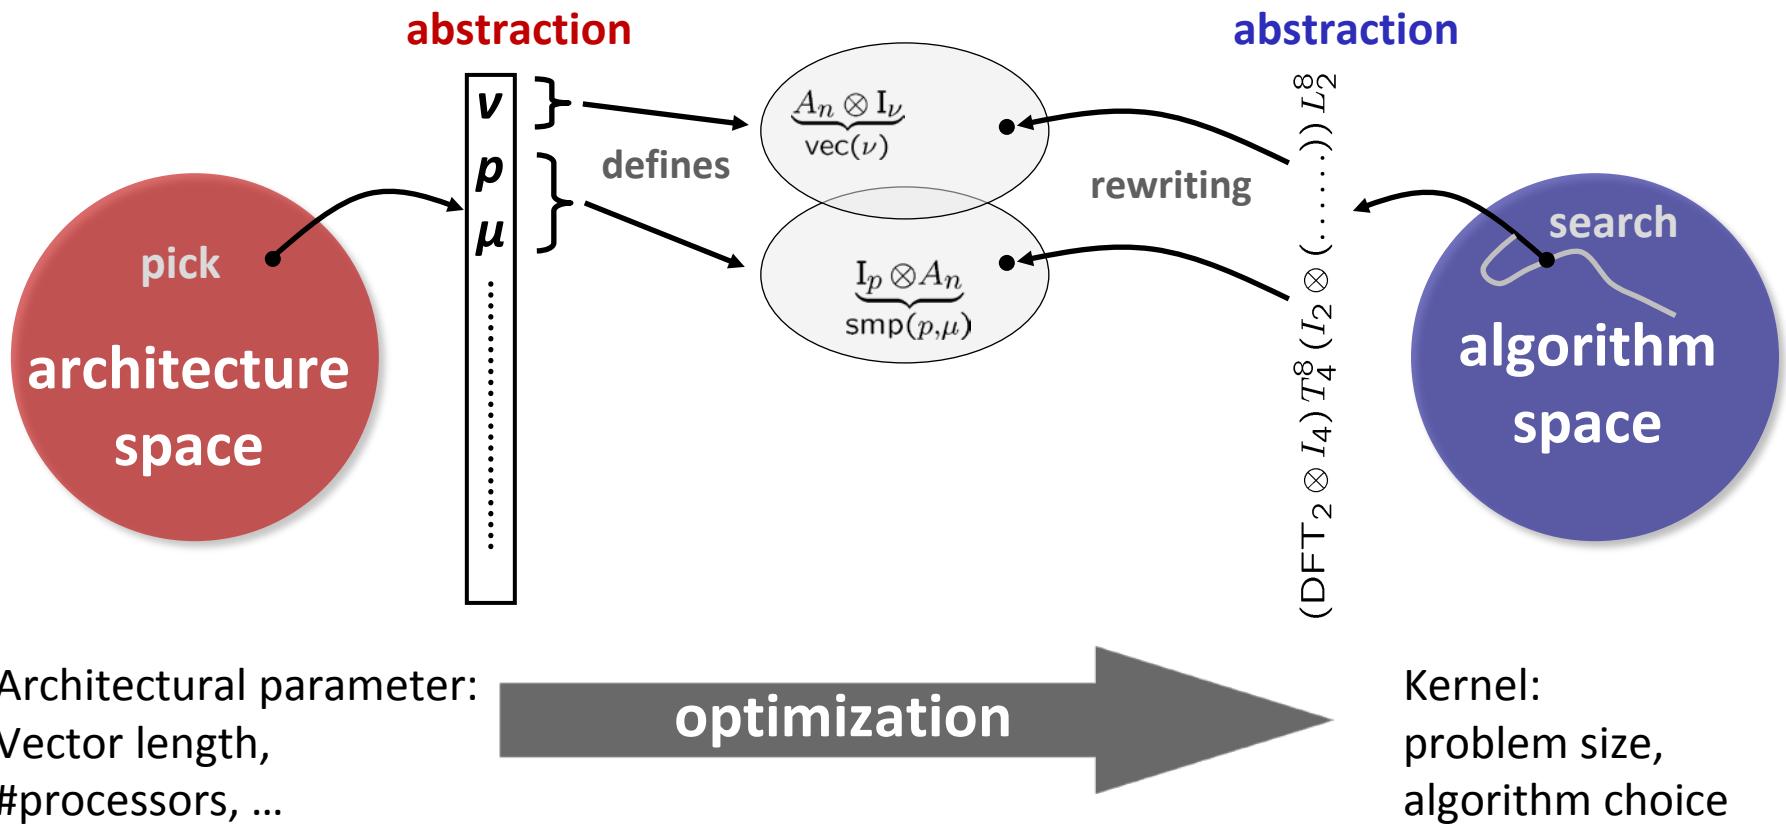


*Spiral*

High performance library  
optimized for given platform

# Platform-Aware Formal Program Synthesis

**Model:** common abstraction  
= spaces of matching formulas

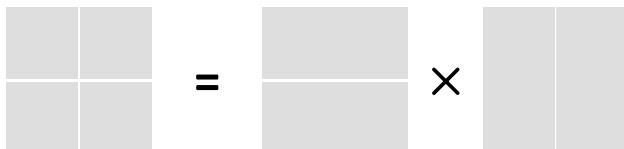


# Algorithms: Rules in Domain Specific Language

## Linear Transforms

$$\begin{aligned}
 \text{DFT}_n &\rightarrow (\text{DFT}_k \otimes \text{I}_m) \text{T}_m^n (\text{I}_k \otimes \text{DFT}_m) \text{L}_k^n, \quad n = km \\
 \text{DFT}_n &\rightarrow P_n (\text{DFT}_k \otimes \text{DFT}_m) Q_n, \quad n = km, \quad \gcd(k, m) = 1 \\
 \text{DFT}_p &\rightarrow R_p^T (\text{I}_1 \oplus \text{DFT}_{p-1}) D_p (\text{I}_1 \oplus \text{DFT}_{p-1}) R_p, \quad p \text{ prime} \\
 \text{DCT-3}_n &\rightarrow (\text{I}_m \oplus \text{J}_m) \text{L}_m^n (\text{DCT-3}_m(1/4) \oplus \text{DCT-3}_m(3/4)) \\
 &\quad \cdot (\mathcal{F}_2 \otimes \text{I}_m) \begin{bmatrix} \text{I}_m & 0 \oplus -\text{J}_{m-1} \\ 0 & \frac{1}{\sqrt{2}}(\text{I}_1 \oplus 2\text{I}_m) \end{bmatrix}, \quad n = 2m \\
 \text{DCT-4}_n &\rightarrow S_n \text{DCT-2}_n \text{diag}_{0 \leq k < n} (1/(2 \cos((2k+1)\pi/4n))) \\
 \text{IMDCT}_{2m} &\rightarrow (\text{J}_m \oplus \text{I}_m \oplus \text{I}_m \oplus \text{J}_m) \left( \left( \begin{bmatrix} 1 \\ -1 \end{bmatrix} \otimes \text{I}_m \right) \oplus \left( \begin{bmatrix} -1 \\ -1 \end{bmatrix} \otimes \text{I}_m \right) \right) \text{J}_{2m} \text{DCT-4}_{2m} \\
 \text{WHT}_{2^k} &\rightarrow \prod_{i=1}^t (\text{I}_{2^{k_1+\dots+k_{i-1}}} \otimes \text{WHT}_{2^{k_i}} \otimes \text{I}_{2^{k_{i+1}+\dots+k_t}}), \quad k = k_1 + \dots + k_t \\
 \text{DFT}_2 &\rightarrow \mathcal{F}_2 \\
 \text{DCT-2}_2 &\rightarrow \text{diag}(1, 1/\sqrt{2}) \mathcal{F}_2 \\
 \text{DCT-4}_2 &\rightarrow \text{J}_2 \mathcal{R}_{13\pi/8}
 \end{aligned}$$

## Matrix-Matrix Multiplication



$$\text{MMM}_{1,1,1} \rightarrow (\cdot)_1$$

$$\text{MMM}_{m,n,k} \rightarrow (\otimes)_{m/m_b \times 1} \otimes \text{MMM}_{m_b, n, k}$$

$$\text{MMM}_{m,n,k} \rightarrow \text{MMM}_{m,nb,k} \otimes (\otimes)_{1 \times n/nb}$$

$$\text{MMM}_{m,n,k} \rightarrow ((\Sigma_{k/k_b} \circ (\cdot)_{k/k_b}) \otimes \text{MMM}_{m,n,k_b}) \circ ((L_{k/k_b}^{mk/k_b} \otimes I_{k_b}) \times I_{kn})$$

$$\text{MMM}_{m,n,k} \rightarrow (L_m^{mn/n_b} \otimes I_{n_b}) \circ ((\otimes)_{1 \times n/n_b} \otimes \text{MMM}_{m,n_b,k}) \circ (I_{km} \times (L_{n/n_b}^{kn/n_b} \otimes I_{n_b}))$$

## Viterbi Decoding



$$\begin{aligned}
 \text{Vit}_{\text{vec}(v)} &\rightarrow \underbrace{\left( \prod (L \times I) \circ (I \otimes C) \right) \circ Id}_{\text{vec}(v)} \\
 &\rightarrow \left( \prod \underbrace{(L \times I) \circ (I \otimes C)}_{\text{vec}(v)} \right) \circ Id \\
 \times &\rightarrow \left( \prod (L \otimes I_v \times I) \circ (I \otimes C \otimes I_v) \circ (\bar{L} \times I) \right) \circ Id \\
 &\rightarrow \prod (L \otimes I_v \times I) \circ (I \otimes (B \otimes I_v)) \circ (\bar{L} \times I)
 \end{aligned}$$

## Synthetic Aperture Radar (SAR)



$$\begin{aligned}
 \text{SAR}_{k \times m \rightarrow n \times n} &\rightarrow \text{DFT}_{n \times n} \circ \text{Interp}_{k \times m \rightarrow n \times n} \\
 \text{DFT}_{n \times n} &\rightarrow (\text{DFT}_n \otimes \text{I}_n) \circ (\text{I}_n \otimes \text{DFT}_n) \\
 \text{Interp}_{k \times m \rightarrow n \times n} &\rightarrow (\text{Interp}_{k \rightarrow n} \otimes_i \text{I}_n) \circ (\text{I}_k \otimes_i \text{Interp}_{m \rightarrow n}) \\
 \text{Interp}_{r \rightarrow s} &\rightarrow \left( \bigoplus_{i=0}^{n-2} \text{InterpSeg}_k \right) \oplus \text{InterpSegPruned}_{k,\ell} \\
 \text{InterpSeg}_k &\rightarrow G_f^{u \cdot n \rightarrow k} \circ \text{iPrunedDFT}_{n \rightarrow u \cdot n} \circ \left( \frac{1}{n} \right) \circ \text{DFT}_n
 \end{aligned}$$

# Formal Approach for all Types of Parallelism

- **Multithreading (Multicore)**

$$\mathbf{I}_p \otimes_{\parallel} A_{\mu n}, \quad \mathbf{L}_m^{mn} \bar{\otimes} \mathbf{I}_{\mu}$$

- **Vector SIMD (SSE, VMX/Altivec,...)**

$$A \hat{\otimes} \mathbf{I}_{\nu} \quad \underbrace{\mathbf{L}_2^{2\nu}}_{\text{isa}}, \quad \underbrace{\mathbf{L}_{\nu}^{2\nu}}_{\text{isa}}, \quad \underbrace{\mathbf{L}_{\nu}^{\nu^2}}_{\text{isa}}$$

- **Message Passing (Clusters, MPP)**

$$\mathbf{I}_p \otimes_{\parallel} A_n, \quad \underbrace{\mathbf{L}_p^{p^2} \bar{\otimes} \mathbf{I}_{n/p^2}}_{\text{all-to-all}}$$

- **Streaming/multibuffering (Cell)**

$$\mathbf{I}_n \otimes_2 A_{\mu n}, \quad \mathbf{L}_m^{mn} \bar{\otimes} \mathbf{I}_{\mu}$$

- **Graphics Processors (GPUs)**

$$\prod_{i=0}^{n-1} A_i, \quad A_n \hat{\otimes} \mathbf{I}_w, \quad P_n \otimes Q_w$$

- **Gate-level parallelism (FPGA)**

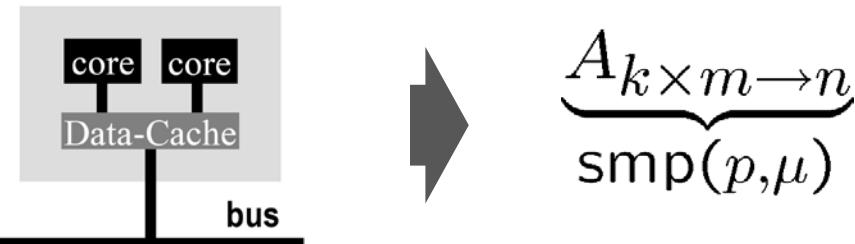
$$\prod_{i=0}^{n-1} \overset{\text{ir}}{A}, \quad \mathbf{I}_s \tilde{\otimes} A, \quad \underbrace{\mathbf{L}_n^m}_{\text{bram}}$$

- **HW/SW partitioning (CPU + FPGA)**

$$\underbrace{A_1}_{\text{fpga}}, \quad \underbrace{A_2}_{\text{fpga}}, \quad \underbrace{A_3}_{\text{fpga}}, \quad \underbrace{A_4}_{\text{fpga}}$$

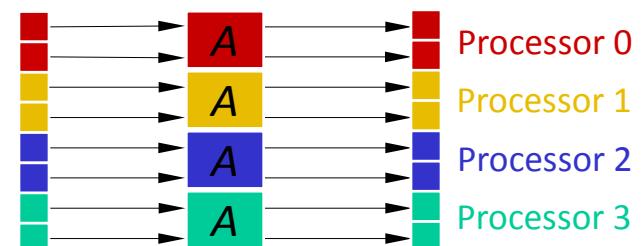
# Modeling Hardware: Base Cases

- **Hardware abstraction: shared cache with cache lines**



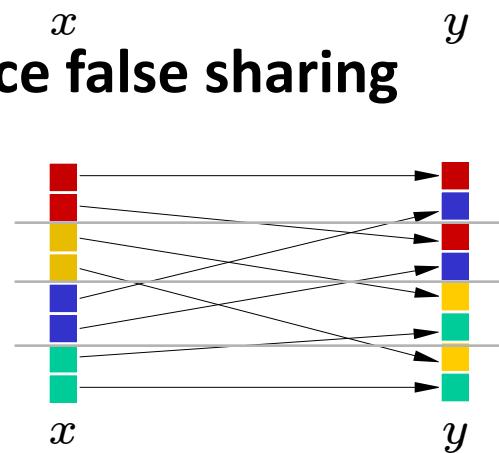
- **Tensor product: embarrassingly parallel operator**

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



- **Permutation: problematic; may produce false sharing**

$$y = L_4^8(x)$$



# Example Program Transformation Rule Set

$$\underbrace{AB}_{\text{smp}(p,\mu)} \rightarrow \underbrace{A}_{\text{smp}(p,\mu)} \underbrace{B}_{\text{smp}(p,\mu)}$$

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

$$\underbrace{L_m^{mn}}_{\text{smp}(p,\mu)} \rightarrow \begin{cases} \underbrace{\left( I_p \otimes L_{m/p}^{mn/p} \right)}_{\text{smp}(p,\mu)} \underbrace{\left( L_p^{pn} \otimes I_{m/p} \right)}_{\text{smp}(p,\mu)} \\ \underbrace{\left( L_m^{pm} \otimes I_{n/p} \right)}_{\text{smp}(p,\mu)} \underbrace{\left( I_p \otimes L_m^{mn/p} \right)}_{\text{smp}(p,\mu)} \end{cases} \quad \text{Recursive rules}$$

$$\underbrace{I_m \otimes A_n}_{\text{smp}(p,\mu)} \rightarrow I_p \otimes \| \left( I_{m/p} \otimes A_n \right)$$

$$\underbrace{(P \otimes I_n)}_{\text{smp}(p,\mu)} \rightarrow \left( P \otimes I_{n/\mu} \right) \overline{\otimes} I_\mu$$

Base case rules

# Autotuning in Constraint Solution Space

Intel Core i7 (2<sup>nd</sup> Gen)



## Base cases

$$\begin{aligned} I_4 \otimes \| A_{16n} \\ L_m^{mn} \bar{\otimes} I_{16} \\ A \hat{\otimes} I_4 \\ L_2^8, L_4^8, L_4^{16}, L_2^4 \otimes I_2 \\ \text{SSE SSE SSE SSE} \end{aligned}$$

## Transformation rules

$$\begin{aligned} \frac{AB}{\text{smp}(p,\mu)} &\rightarrow \frac{A}{\text{smp}(p,\mu)} \frac{B}{\text{smp}(p,\mu)} \\ L_m^{mn} &\rightarrow \begin{cases} (I_p \otimes L_{m/p}^{mn/p}) (L_p^{pn} \otimes I_{m/p}) \\ \text{smp}(p,\mu) \quad \text{smp}(p,\mu) \end{cases} \\ \frac{L_m^{mn}}{\text{smp}(p,\mu)} &\rightarrow \begin{cases} (L_m^{pn} \otimes I_{n/p}) (I_p \otimes L_m^{mn/p}) \\ \text{smp}(p,\mu) \quad \text{smp}(p,\mu) \end{cases} \\ \frac{L_m \otimes A_n}{\text{smp}(p,\mu)} &\rightarrow I_p \otimes \| (I_{m/p} \otimes A_n) \end{aligned}$$

DFT<sub>256</sub>



## Breakdown rules

$$\begin{aligned} \text{DFT}_n &\rightarrow (\text{DFT}_k \otimes I_m) T_m^n \\ &\quad \cdot (I_k \otimes \text{DFT}_m) L_k^n \\ \text{DFT}_n &\rightarrow P_n (\text{DFT}_k \otimes \text{DFT}_m) Q_n \\ \text{DFT}_p &\rightarrow R_p^T (I_1 \oplus \text{DFT}_{p-1}) D_p \\ &\quad \cdot (I_1 \oplus \text{DFT}_{p-1}) R_p \\ \text{DFT}_2 &\rightarrow \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \end{aligned}$$

Expansion + backtracking

## OL specification

**OL (dataflow) expression**

Recursive descent

## $\Sigma$ -OL (loop) expression

Confluent term rewriting

## Optimized $\Sigma$ -OL expression

Recursive descent

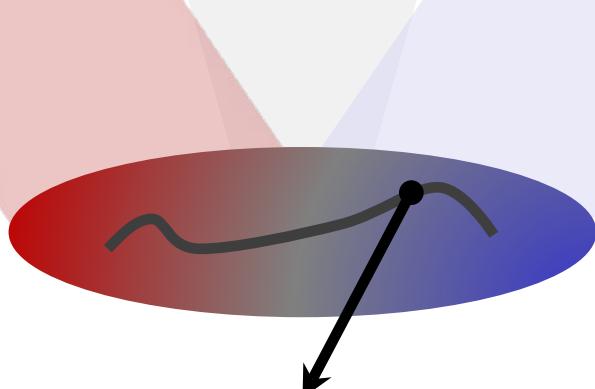
## Abstract code

Confluent term rewriting

## Optimized abstract code

Recursive descent

## C code



$$\left( (L_m^{mp} \otimes I_{n/p\mu}) \bar{\otimes} I_\mu \right) \left( I_p \otimes \| (\text{DFT}_m \otimes I_{n/p}) \right) \left( (L_p^{mp} \otimes I_{n/p\mu}) \bar{\otimes} I_\mu \right) \left( \bigoplus_{i=0}^{p-1} \| T_n^{mn,i} \right) \left( I_p \otimes \| (I_{m/p} \otimes \text{DFT}_n) \right) \left( I_p \otimes \| L_{m/p}^{mn/p} \right) \left( (L_p^{pn} \otimes I_{m/p\mu}) \bar{\otimes} I_\mu \right)$$

# Translating an OL Expression Into Code

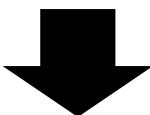
Constraint Solver Input:

$\underbrace{\text{DFT}_8}_{\text{double}}$



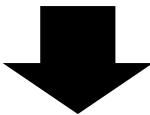
Output =

**OL Expression:**  $(\text{DFT}_2 \otimes I_4) T_4^8 \left( I_2 \otimes \left( (\text{DFT}_2 \otimes I_2) T_2^4 (I_2 \otimes \text{DFT}_2) L_2^4 \right) \right) L_2^8$



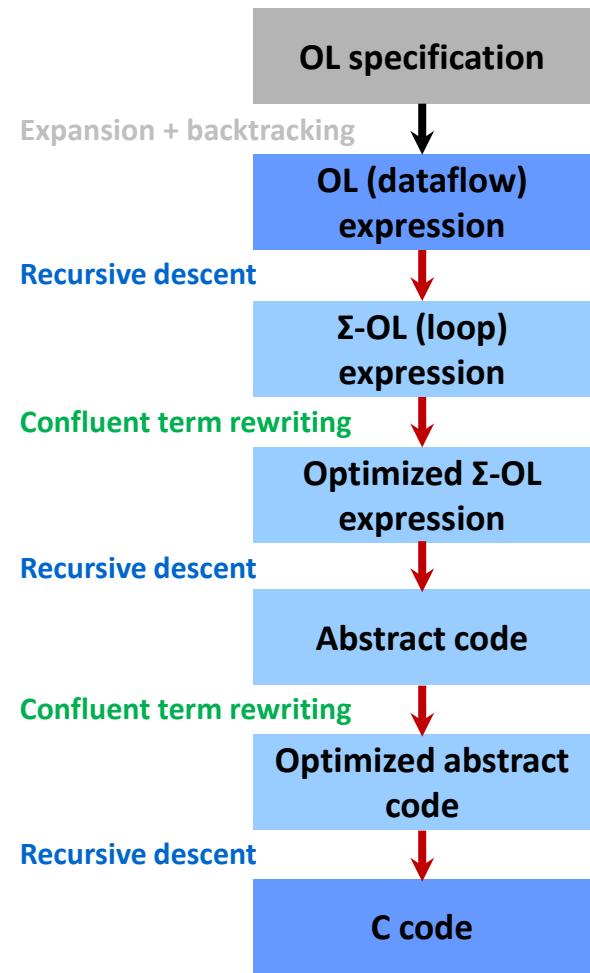
**Σ-OL:**

$$\sum_{j=0}^3 (S_j \text{DFT}_2 G_j) \sum_{k=0}^1 \left( \sum_{l=0}^1 (S_{k,l} \text{diag}(t_{k,l}) \text{DFT}_2 G_l) \sum_{m=0}^1 (S_m \text{diag}(t_m) \text{DFT}_2 G_{k,m}) \right)$$



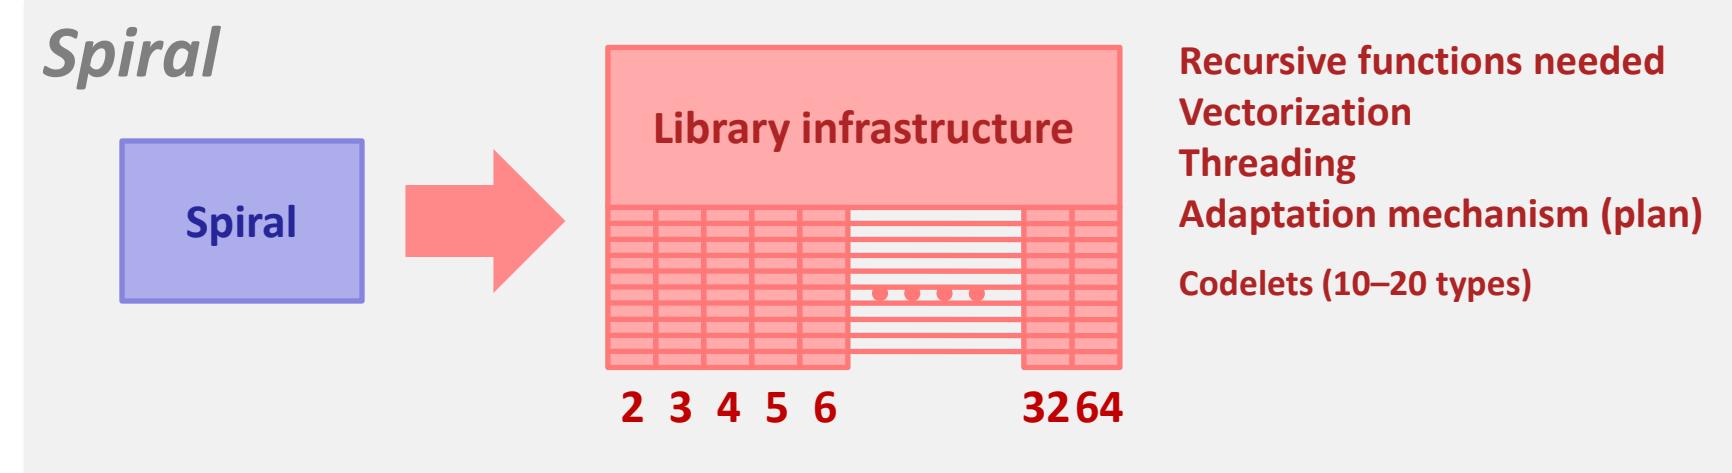
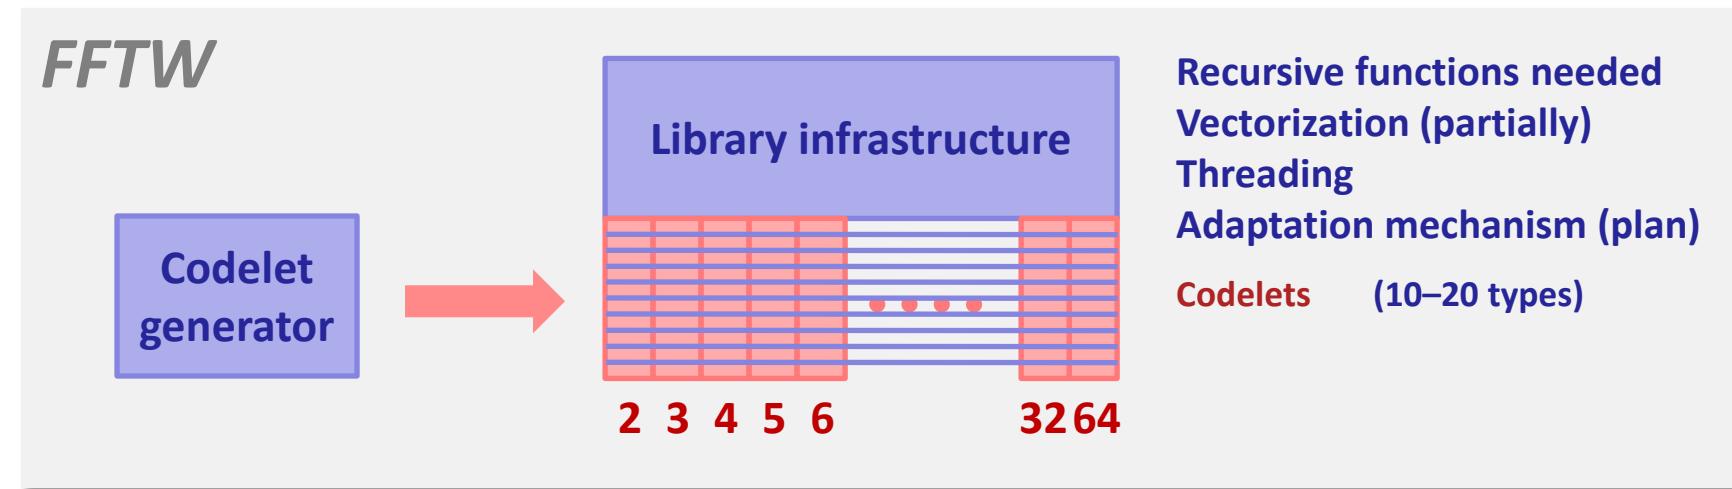
**C Code:**

```
void sub(double *y, double *x) {
    double f0, f1, f2, f3, f4, f7, f8, f10, f11;
    f0 = x[0] - x[3];
    f1 = x[0] + x[3];
    f2 = x[1] - x[2];
    f3 = x[1] + x[2];
    f4 = f1 - f3;
    y[0] = f1 + f3;
    y[2] = 0.7071067811865476 * f4;
    f7 = 0.9238795325112867 * f0;
    f8 = 0.3826834323650898 * f2;
    y[1] = f7 + f8;
    f10 = 0.3826834323650898 * f0;
    f11 = (-0.9238795325112867) * f2;
    y[3] = f10 + f11;
}
```



# SPIRAL: Meta-Tool to FFTW

blue = hand-written, red = generated



# Production-Quality Code

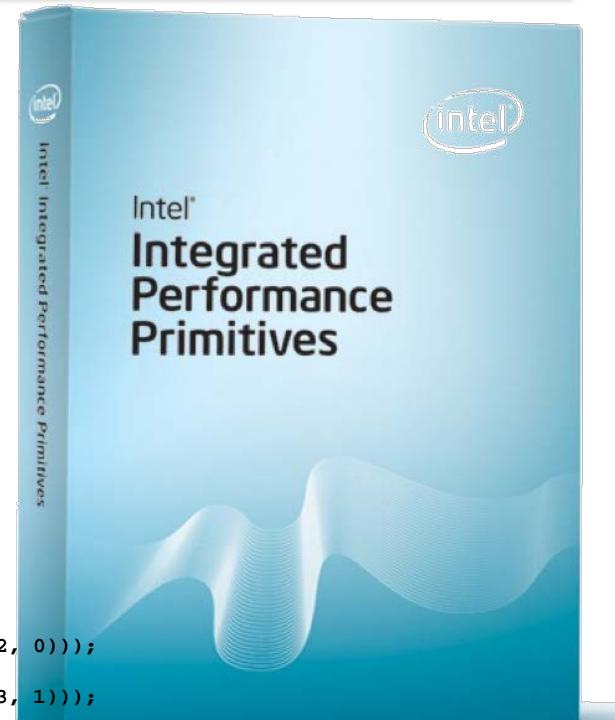
```
s3013 = _mm_loadl_pi(a772, ((float *) X));
s3014 = _mm_loadh_pi(_mm_loadl_pi(a772, ((float *) (Y + 2))), ((float *) (Y + 6)));
```

## *Spiral-Synthesized code in Intel's Library IPP 6 and 7*

- IPP = Intel's performance primitives, part of Intel C++ Compiler suite
  - Generated: 3984 C functions (signal processing) = 1M lines of code
  - Full parallelism support
  - Computer-generated code: Faster than what was achievable by hand

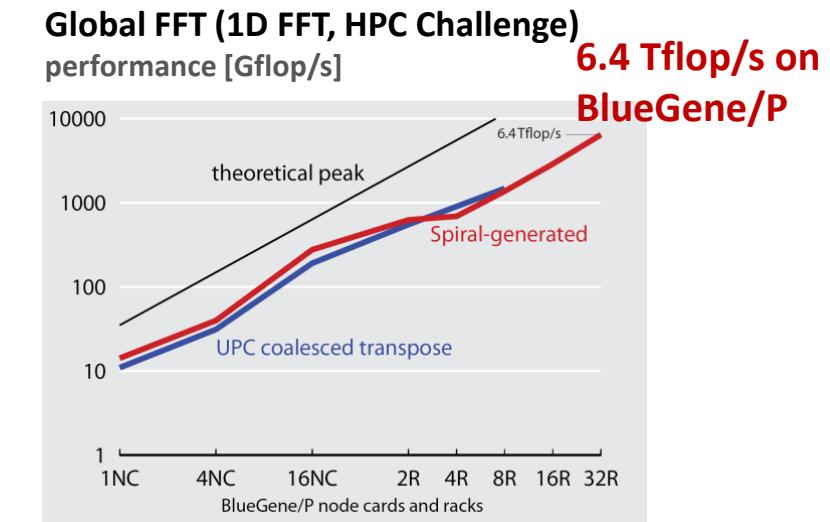
intel integrated Performance Primitives

```
s3017 = _mm_loadl_pi(a772, ((float *) (X + 14)));
s3018 = _mm_loadh_pi(_mm_loadl_pi(a772, ((float *) (X + 24))), ((float *) (X + 20)));
s3019 = _mm_loadl_pi(a772, ((float *) (X + 8)));
s3020 = _mm_loadh_pi(_mm_loadl_pi(a772, ((float *) (X + 10))), ((float *) (X + 4)));
s3021 = _mm_loadl_pi(a772, ((float *) (X + 12)));
s3022 = _mm_shuffle_ps(s3014, s3015, _MM_SHUFFLE(2, 0, 2, 0));
s3023 = _mm_shuffle_ps(s3014, s3015, _MM_SHUFFLE(3, 1, 3, 1));
s3024 = _mm_shuffle_ps(s3016, s3017, _MM_SHUFFLE(2, 0, 2, 0));
s3025 = _mm_shuffle_ps(s3016, s3017, _MM_SHUFFLE(3, 1, 3, 1));
s3026 = _mm_shuffle_ps(s3018, s3019, _MM_SHUFFLE(2, 0, 2, 0));
s3027 = _mm_shuffle_ps(s3018, s3019, _MM_SHUFFLE(3, 1, 3, 1));
s3028 = _mm_shuffle_ps(s3020, s3021, _MM_SHUFFLE(2, 0, 2, 0));
s3029 = _mm_shuffle_ps(s3020, s3021, _MM_SHUFFLE(3, 1, 3, 1));
...
t3794 = _mm_add_ps(s3042, s3043);
t3795 = _mm_add_ps(s3038, t3793);
t3796 = _mm_add_ps(s3041, t3794);
t3797 = _mm_sub_ps(s3038, _mm_mul_ps(_mm_set1_ps(0.5), t3793));
t3798 = _mm_sub_ps(s3041, _mm_mul_ps(_mm_set1_ps(0.5), t3794));
s3044 = _mm_mul_ps(_mm_set1_ps(0.8660254037844386), _mm_sub_ps(s3042, s3043));
s3045 = _mm_mul_ps(_mm_set1_ps(0.8660254037844386), _mm_sub_ps(s3039, s3040));
t3799 = _mm_add_ps(t3797, s3044);
t3800 = _mm_sub_ps(t3798, s3045);
t3801 = _mm_sub_ps(t3797, s3044);
t3802 = _mm_add_ps(t3798, s3045);
a773 = _mm_mul_ps(_mm_set_ps(0, 0, 0, 1), _mm_shuffle_ps(s3013, a772, _MM_SHUFFLE(2, 0, 2, 0)));
t3803 = _mm_add_ps(a773, _mm_mul_ps(_mm_set_ps(0, 0, 0, 1), t3795));
a774 = _mm_mul_ps(_mm_set_ps(0, 0, 0, 1), _mm_shuffle_ps(s3013, a772, _MM_SHUFFLE(3, 1, 3, 1)));
t3804 = _mm_add_ps(a774, _mm_mul_ps(_mm_set_ps(0, 0, 0, 1), t3796));
t3805 = _mm_add_ps(a773, _mm_add_ps(_mm_mul_ps(_mm_set_ps(0.28757036473700154, 0.300462606288665,
t3806 = _mm_add_ps(a774, _mm_add_ps(_mm_mul_ps(_mm_set_ps(0.08706930057606789, 0, 0, 0.087069300576
s3046 = _mm_sub_ps(_mm_mul_ps(_mm_set_ps((-0.25624767158293649), 0.25826039031174486, (-0.300238
s3047 = _mm_add_ps(_mm_mul_ps(_mm_set_ps(0.15689139105158462, (-0.15355568557954136), (-0.011599
```



# SPIRAL: Success in HPC/Supercomputing

- **NCSA Blue Waters**  
PAID Program, FFTs for Blue Waters
- **RIKEN K computer**  
FFTs for the HPC-ACE ISA
- **LANL RoadRunner**  
FFTs for the Cell processor
- **PSC/XSEDE Bridges**  
Large size FFTs
- **LLNL BlueGene/L and P**  
FFTW for BlueGene/L's Double FPU
- **ANL BlueGene/Q Mira**  
Early Science Program, FFTW for BGQ QPX



**BlueGene/P at Argonne National Laboratory**  
128k cores (quad-core CPUs) at 850 MHz

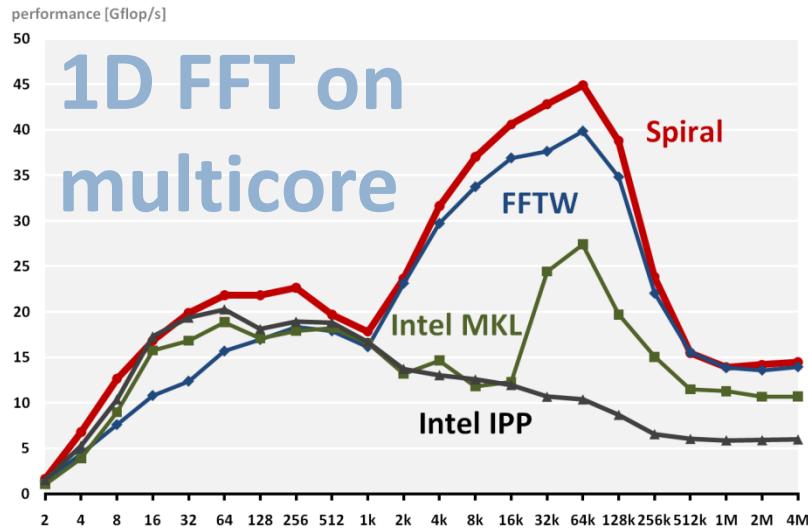


**2006 Gordon Bell Prize (Peak Performance Award) with LLNL and IBM**

**2010 HPC Challenge Class II Award (Most Productive System) with ANL and IBM**

# SPIRAL: Performance Across Platforms

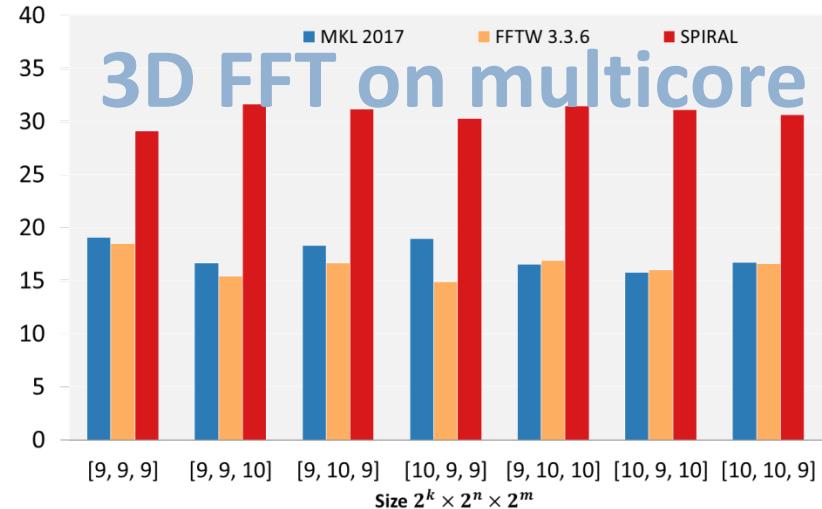
1D DFT on 3.3 GHz Sandy Bridge (4 Cores, AVX)



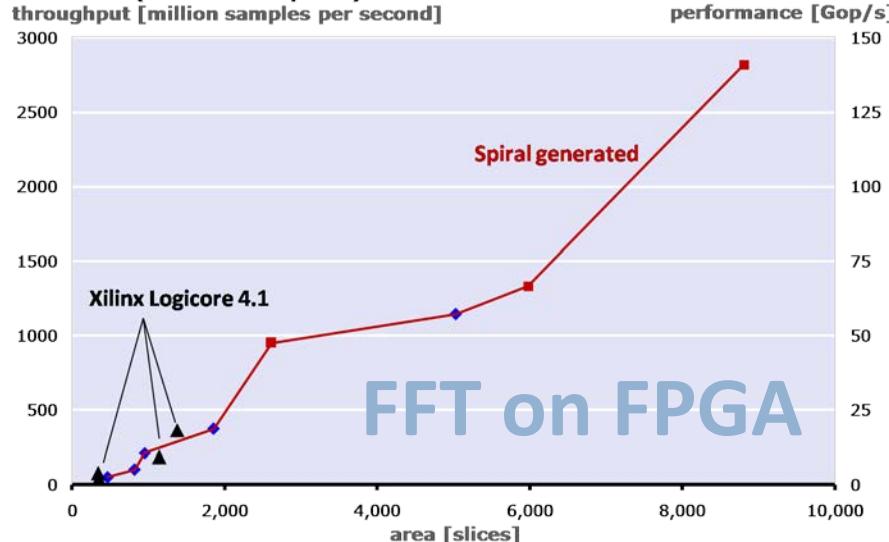
3D FFT Performance on Intel Haswell 4770K

3.5 GHz, 4/8 cores/threads, double-precision, AVX

Performance [Gflop/s]

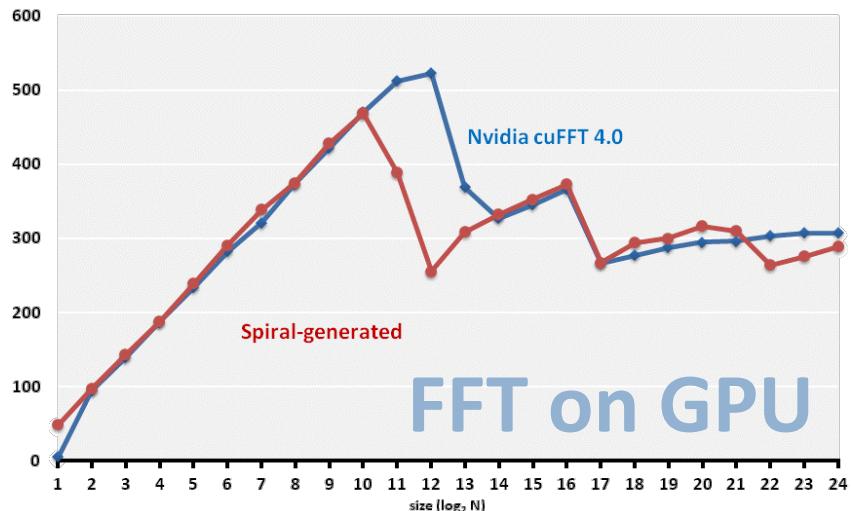


DFT 1024 (16 bit fixed point) on Xilinx Virtex-5 FPGA



1D Batch DFT (Nvidia GTX 480)

performance [GFlop/s], single precision



# Thrust 1: Repeat The Success of LAPACK for FFTs

## Numerical Linear Algebra

LAPACK

ScaLAPACK

LU factorization

Eigensolves

SVD

BLAS, BLACS

BLAS-1

BLAS-2

BLAS-3

## Spectral Algorithms

Convolution

Correlation

Upsampling

Poisson solver

...



FFTW

DFT, RDFT

1D, 2D, 3D,...

batch

## No LAPACK equivalent for spectral methods

- **Medium size 1D FFT (1k—10k data points) is most common library call**  
applications break down 3D problems themselves and then call the 1D FFT library
- **Higher level FFT calls rarely used**  
FFTW *guru* interface is powerful but hard to use, leading to performance loss
- **Low arithmetic intensity and variation of FFT use make library approach hard**  
Algorithm specific decompositions and FFT calls intertwined with non-FFT code

# Thrust 2: Updated FFTW Interface for Node FFTs

## Front End: application facing

- **Backwards compatible to FFTW 2.X and 3.X**  
old code runs unmodified and gains substantially but not fully
- **Delayed execution/futures**  
asynchronous execution, offloading, triggering of execution
- **Data location and layout control**  
specify location of data (host/device, NUMA, NVRAM,...)
- **Performance metrics**  
user control of resource/performance/energy requirements and adaptation

## Backend: developer and hardware facing

- **Extensible to novel and yet-unknown hardware and system architectures**  
future-proof platform extension API: transparently enable accelerators, etc.
- **Plug-in mechanism for special case acceleration packs**  
transparent access to application/architecture specialized FFT kernels

## More Information:

[www.spiral.net](http://www.spiral.net)

[www.spiralgen.com](http://www.spiralgen.com)

*We are about to make SPIRAL available as  
open source under a BSD style license*