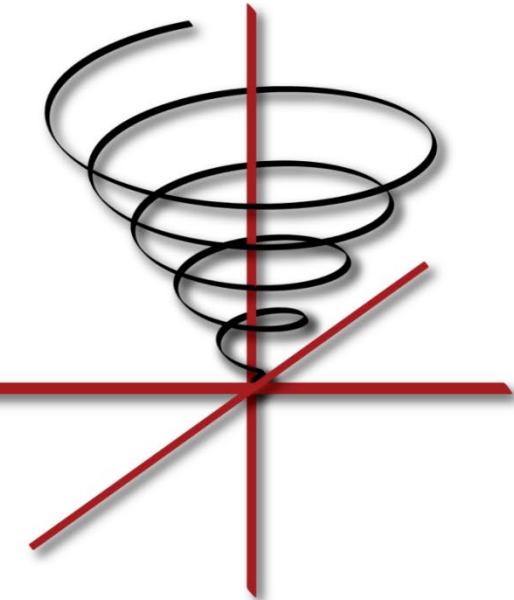


The Design of FFTX



Franz Franchetti

Carnegie Mellon University

in collaboration with

Daniele G. Spampinato, Anuva Kulkarni, Tze Meng Low

Carnegie Mellon University

Doru Thom Popovici, Andrew Canning, Peter McCorquodale,

Brian Van Straalen, Phillip Colella

Lawrence Berkeley National Laboratory

Mike Franusich

SpiralGen, Inc.

This work was supported by DOE ECP project 2.3.3.07

Have You Ever Wondered About This?

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

FFTX and SpectralPACK

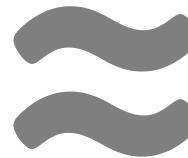
Numerical Linear Algebra

LAPACK

LU factorization
Eigensolves
SVD
...

BLAS

BLAS-1
BLAS-2
BLAS-3



Spectral Algorithms

SpectralPACK

Convolution
Correlation
Upsampling
Poisson solver
...

FFTX

DFT, RDFT
1D, 2D, 3D,...
batch

Define the LAPACK equivalent for spectral algorithms

- **Define FFTX as the BLAS equivalent**
provide user FFT functionality as well as algorithm building blocks
- **Define class of numerical algorithms to be supported by SpectralPACK**
PDE solver classes (Green's function, sparse in normal/k space,...), signal processing,...
- **Library front-end, code generation and vendor library back-end**
mirror concepts from FFTX layer

FFTX and SpectralPACK solve the “spectral motif” long term

Example: Poisson's Equation in Free Space

Partial differential equation (PDE)

$$\Delta(\Phi) = \rho$$

$$\rho : \mathbb{R}^3 \rightarrow \mathbb{R}$$

$$D = \text{supp}(\rho) \subset \mathbb{R}^3$$

Poisson's equation. Δ is the Laplace operator

Solution characterization

$$\Phi : \mathbb{R}^3 \rightarrow \mathbb{R}$$

$$\Phi(\vec{x}) = \frac{Q}{4\pi||\vec{x}||} + o\left(\frac{1}{||\vec{x}||}\right) \text{ as } ||\vec{x}|| \rightarrow \infty$$

$$Q = \int_D \rho d\vec{x}$$

Approach: Green's function

$$\Phi(\vec{x}) = \int_D G(\vec{x} - \vec{y})\rho(\vec{y})d\vec{y} \equiv (G * \rho)(\vec{x}), \quad G(\vec{x}) = \frac{1}{4\pi||\vec{x}||_2}$$

Solution: $\phi(\cdot)$ = convolution of RHS $\rho(\cdot)$ with Green's function $G(\cdot)$. Efficient through FFTs (frequency domain)

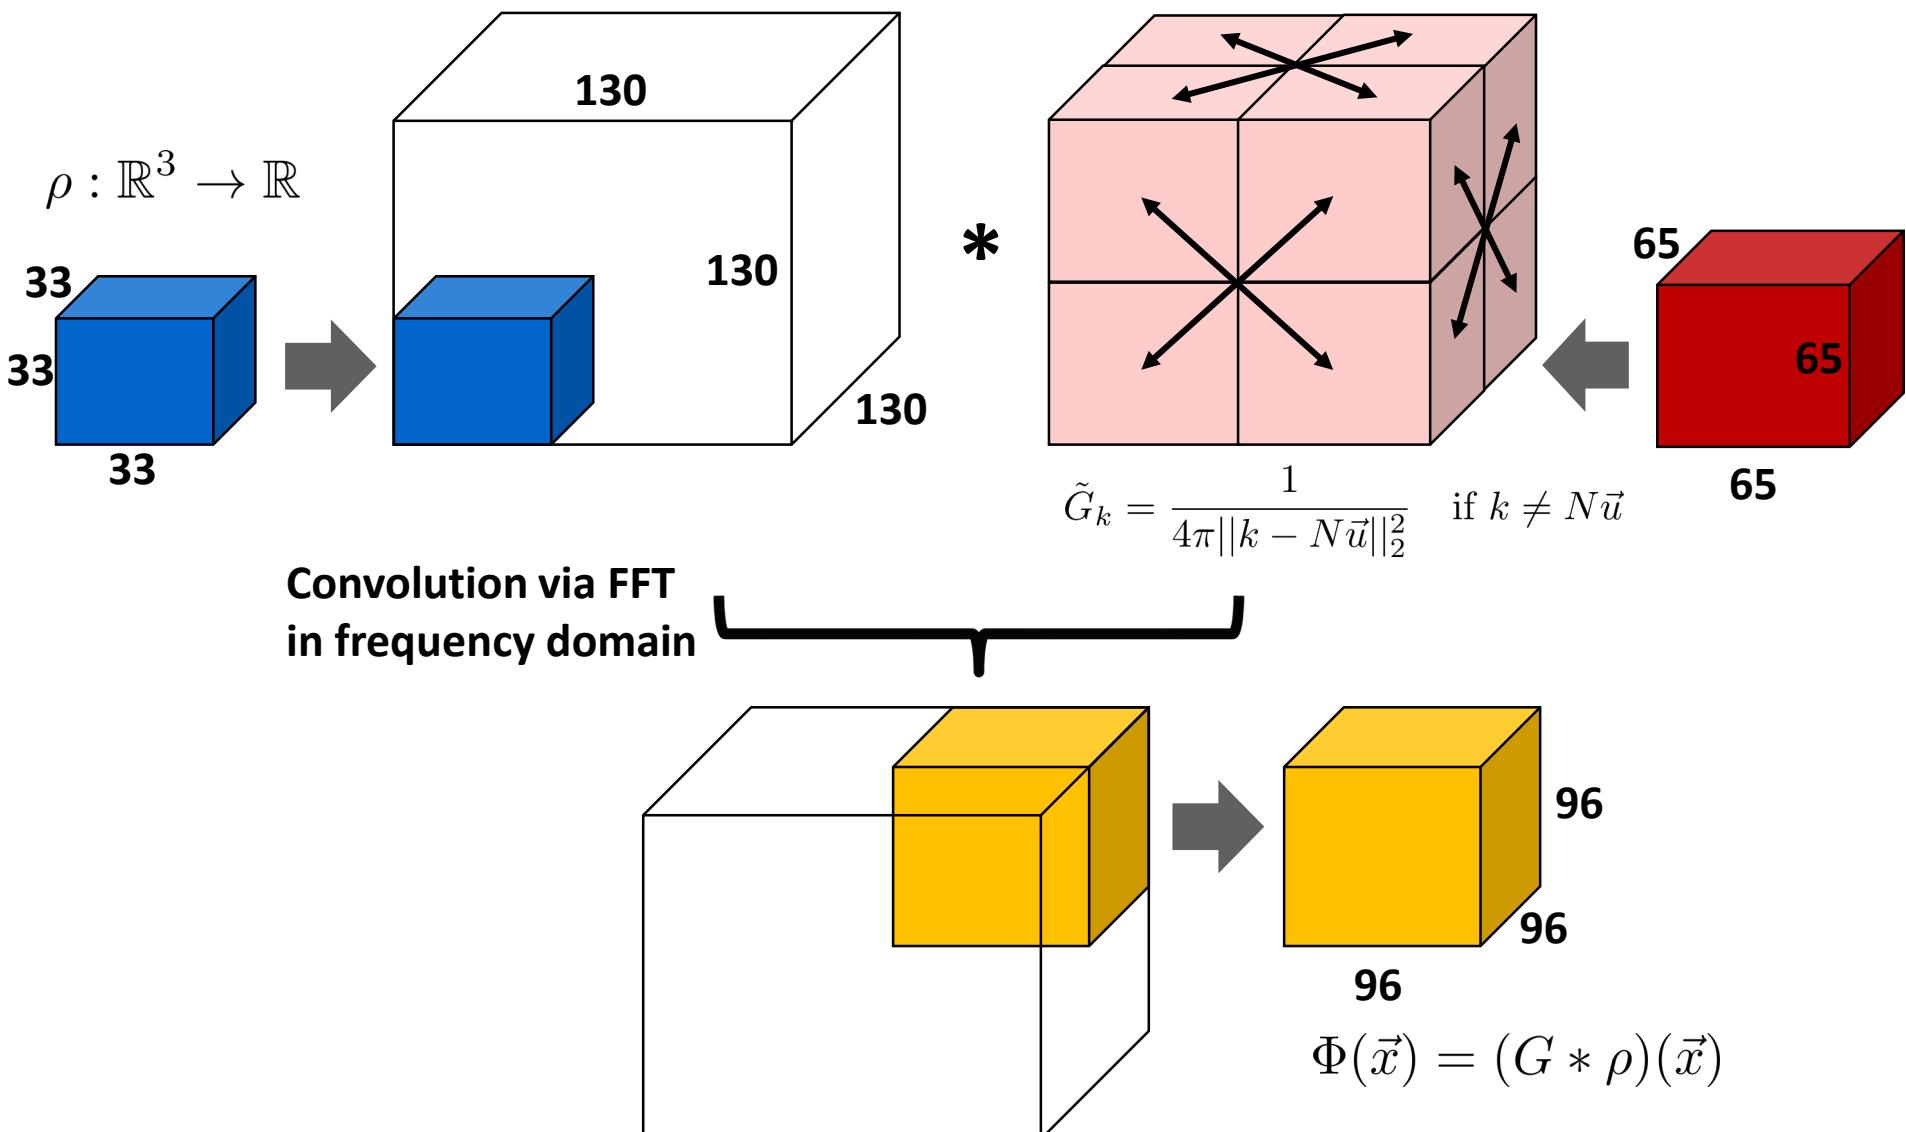
Method of Local Corrections (MLC)

$$\tilde{G}_k = \frac{1}{4\pi||k - N\vec{u}||_2^2} \quad \text{if } k \neq N\vec{u} \quad \text{Green's function kernel in frequency domain}$$

P. McCorquodale, P. Colella, G. T. Banks, and S. B. Baden: **A Local Corrections Algorithm for Solving Poisson's Equation in Three Dimensions.** Communications in Applied Mathematics and Computational Science Vol. 2, No. 1 (2007), pp. 57-81., 2007.

C. R. Anderson: **A method of local corrections for computing the velocity field due to a distribution of vortex blobs.** Journal of Computational Physics, vol. 62, no. 1, pp. 111–123, 1986.

Algorithm: Hockney Free Space Convolution



Hockney: Convolution + problem specific zero padding and output subset

FFTX C++ Code: Hockney Free Space Convolution

```
box_t<3> inputBox(point_t<3>({{0,0,0}}),point_t<3>({32,32,32}));  
array_t<3, double> rho(inputBox);  
// ... set input values.  
  
box_t<3> transformBox(point_t<3>({{0,0,0}}),point_t<3>({{129,129,129}}));  
box_t<3> outputBox(point_t<3>({33,33,33}),point_t<3>({129,129,129}));  
  
point_t<3> kindF({{DFT,DFT,DFT}});  
  
size_t  
auto fo  
pla  
  
auto ke  
  
point_t  
auto in  
  
auto so  
  
context  
context
```

This is a specification dressed as a program

- Needs to be clean and concise
- No code level optimizations and tricks
- Don't think "performance" but "correctness"
- *For production code and software engineering*

```
std::ofstream splFile("hockney.spl");  
export_spl(context, solver, splFile, "hockney33_97_130");  
splFile.close();  
// Offline codegen.  
auto fptr = import_spl<3, double, double>("hockney33_97_130");  
array_t<3, double> Phi(inputBox);  
fptr(&rho, &Phi, 1);
```

FFTX C Code: Hockney Free Space Convolution

```
fftx_plan pruned_real_convolution_plan(fftx_real *in, fftx_real *out, fftx_complex *symbol,
    int n, int n_in, int n_out, int n_freq) {
    int rank = 3,
    batch_rank = 0,
    ...
    fftx_plan plans[5];
    fftx_plan_guru_dft_c2r(rank, &padded_dims, batch_rank,
    &batch_dims, tmp3, tmp4, MY_FFTX_MODE_SUB);

    plans[4] = fftx_plan_guru_copy_real(rank, &out_dimx, tmp4, out, MY_FFTX_MODE_SUB);

    p = fftx_plan_compose(numsubplans, plans, MY_FFTX_MODE_TOP);

    return p;
}
```

This is a specification dressed as a program

- Needs to be clean and concise
- No code level optimizations and tricks
- Don't think "performance" but "correctness"
- *For small and in-development platforms*

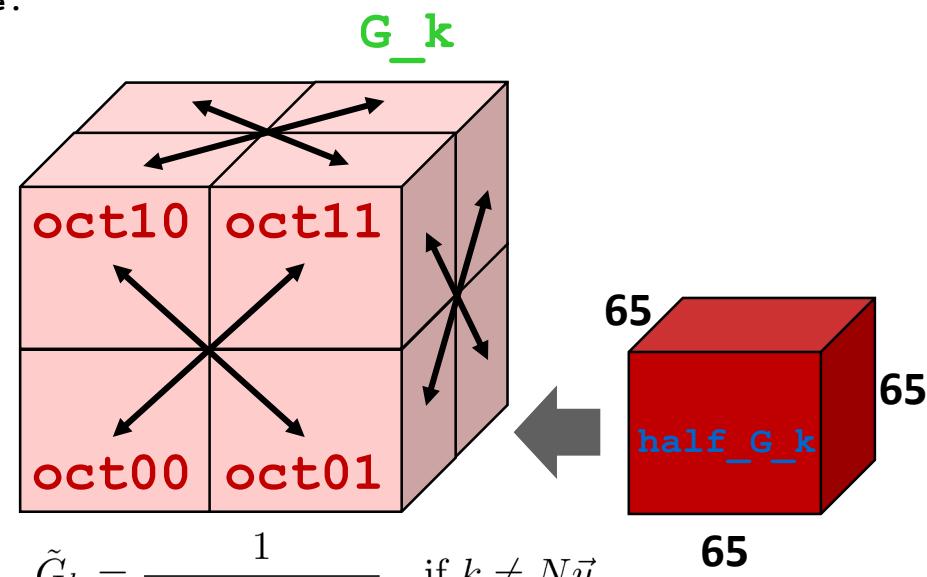
Looks like FFTW calls, but is a specification for SPIRAL

FFTX C Code: Describing The Hockney Symmetry

```
// FFTX data access descriptors.
// Access is to four octants of a symmetric cube.
// Cube size is N^3 and M = N/2.
```

```
fftx_ioidimx oct00[] = {
{ M+1, 0, 0, 0, 1, 1, 1 },
{ M+1, 0, 0, 0, M+1, 2*M, 1 },
{ M+1, 0, 0, 0, (M+1)*(M+1), 4*M*M, 1 } },
oct01[] = {
{ M-1, M-1, M+1, 0, -1, 1, 1 },
{ M+1, 0, 0, 0, M+1, 2*M, 1 },
{ M+1, 0, 0, 0, (M+1)*(M+1), 4*M*M, 1 } },
oct10[] = {
{ M+1, 0, 0, 0, 1, 1, 1 },
{ M-1, M-1, M+1, 0, -(M+1), 2*M, 1 },
{ M+1, 0, 0, 0, (M+1)*(M+1), 4*M*M, 1 } },
oct11[] = {
{ M-1, M-1, M+1, 0, -1, 1, 1 },
{ M-1, M-1, M+1, 0, -(M+1), 2*M, 1 },
{ M+1, 0, 0, 0, (M+1)*(M+1), 4*M*M, 1 } };
...
```

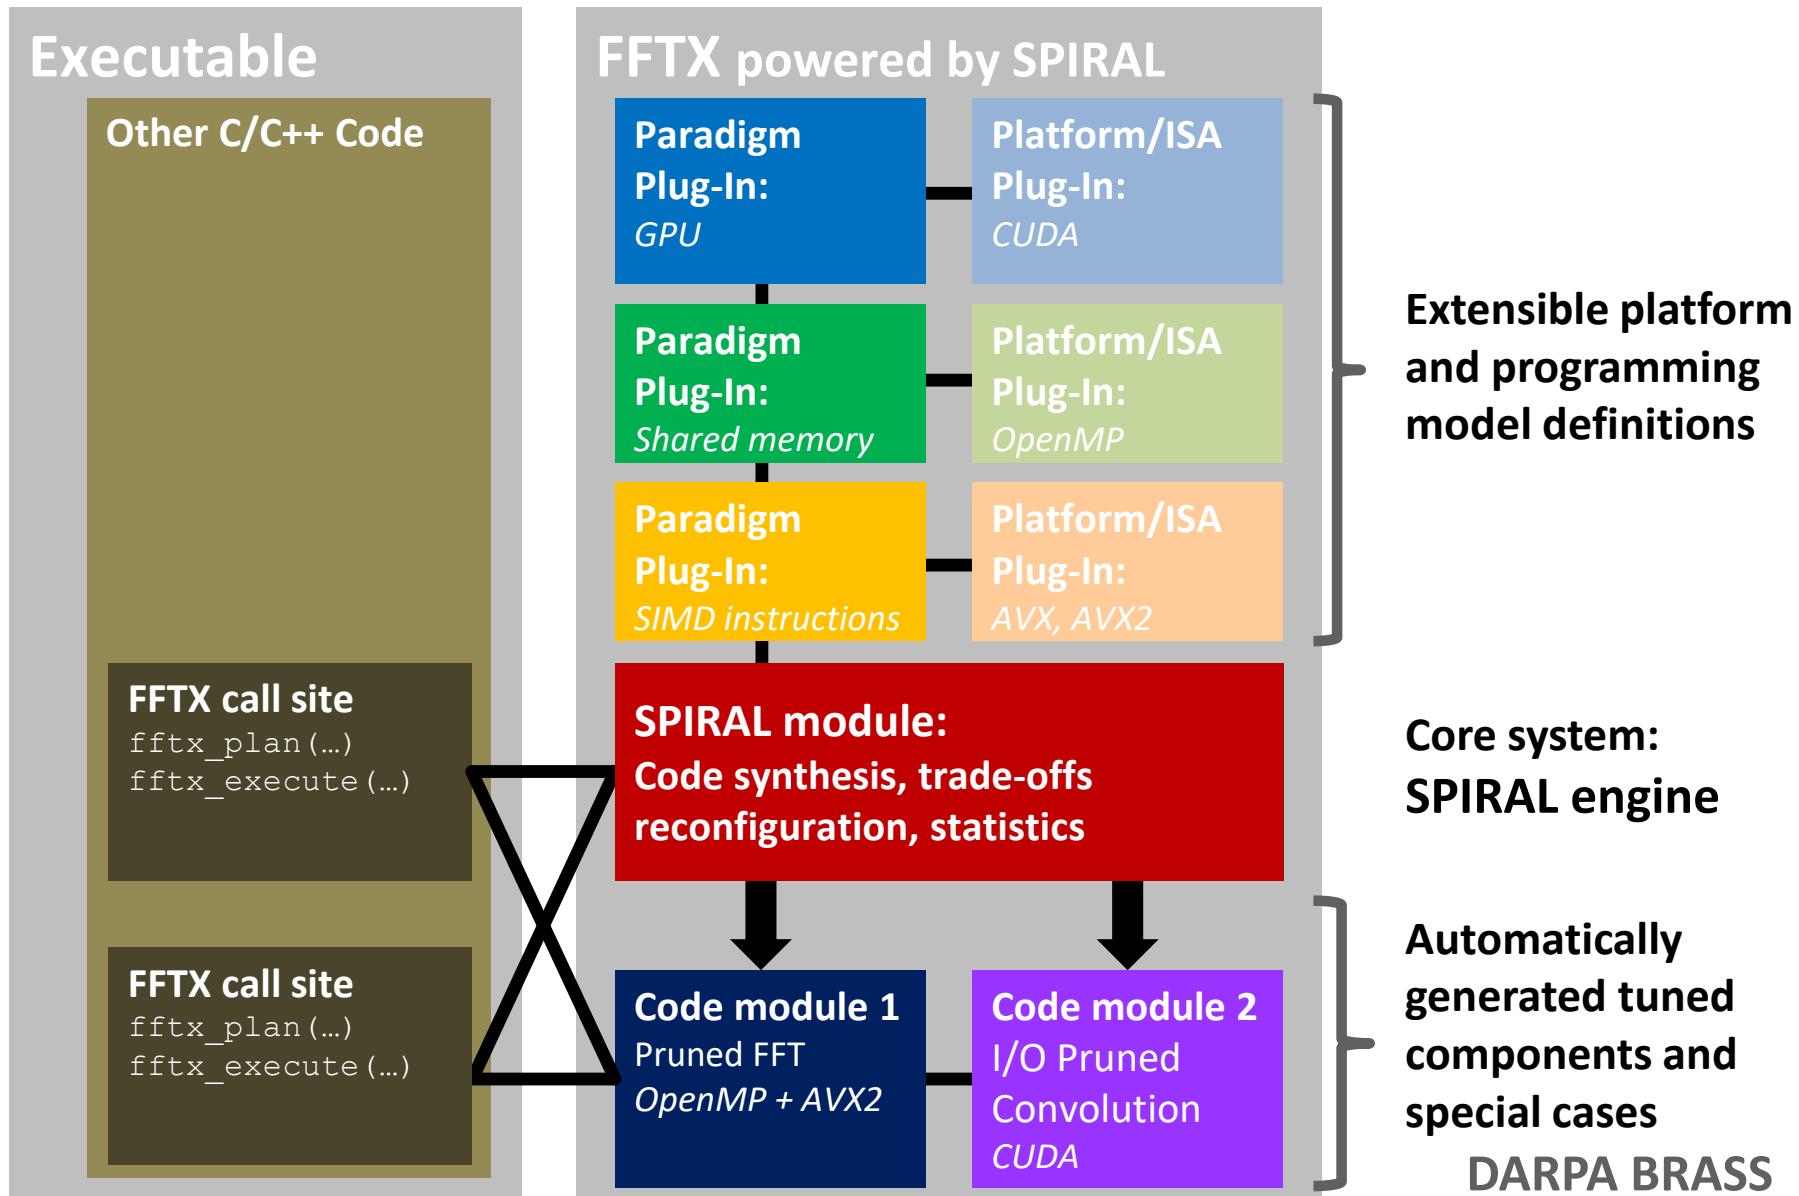
```
fftx_temp_complex half_G_k = fftx_create_zero_temp_complex(rk, f_d);
plans[2] = fftx_plan_guru_copy_complex(rk, oct00, G_k, half_G_k, FFTX_MODE_SUB);
plans[3] = fftx_plan_guru_copy_complex(rk, oct01, G_k, half_G_k, MY_FFTX_MODE_SUB);
plans[4] = fftx_plan_guru_copy_complex(rk, oct10, G_k, half_G_k, MY_FFTX_MODE_SUB);
plans[5] = fftx_plan_guru_copy_complex(rk, oct11, G_k, half_G_k, MY_FFTX_MODE_SUB);
...
```



$$\tilde{G}_k = \frac{1}{4\pi||k - N\vec{u}||_2^2} \quad \text{if } k \neq N\vec{u}$$

We are developing higher-level more natural geometric API

FFTX Backend: SPIRAL



C/C++ FFTX Program Trace

```
fftx_session := [
    rec(op := "fftx_init", flags := IntHexString("8000000")),
    rec(op := "fftx_create_data_real", rank := 1, dims := [ rec(n := 4, is := 1, os := 1) ],
        ptr := IntHexString("000000D236DD2460")),
    ...
    rec(op := "fftx_create_zero_temp_real", rank := 1,
```

```
    re
```

The whole convolution kernel is captured

- DAG with all dependencies
- User-defined call-backs
- Captures pruning, zero-padding and symmetries
- *Lifts sequence of C/C++ library calls to a specification*

```
callback := [
    rec(op := "call", inp := IntHexString("A000000000000001"),
        outp := IntHexString("A000000000000002"), data := IntHexString("A000000000000003")),
    rec(op := "FFTX_COMPLEX_VAR", var := IntHexString("000000D236A0FA30"),
        re := 0.000000e+00, im := 0.000000e+00),
    rec(op := "FFTX_COMPLEX_MOV", target := IntHexString("000000D236A0FA30"),
        source := IntHexString("A000000000000001")),
    rec(op := "FFTX_COMPLEX_MUL", target := IntHexString("000000D236A0FA30"),
        source := IntHexString("A000000000000003")),
    ...
```

SPIRAL Script Captures Performance Engineering

```
# Pruned 3D Real Convolution Pattern
Import(realdft);
Import(filtering);

# set up algorithms needed for multi-dimensional pruned real convolution
opts = S[1].D[1];
opts
PR
opts
IP
opts
opts
opts
```

Recognizes pattern and applies code generation

- Developed by performance engineer + application specialist
- Casts FFTX call sequence as SPIRAL non-terminal
- Does code generation and autotuning
- *Clear separation of concerns frontend/backend*

```
# generate code and autotune
rt := DP(t, opts)[1].ruletree;
c := CodeRuleTree(rt, opts);

# create files
PrintTo(name::".c", PrintCode(name, c, opts));
```

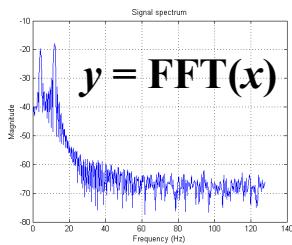
SPIRAL: Go from Mathematics to Software

Given:

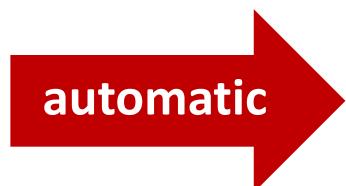
- Mathematical problem specification
core mathematics does not change
- Target computer platform
varies greatly, new platforms introduced often

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));
    ...
}

```



Inspiration: Symbolic Integration

- Rule based AI system
basic functions, substitution
- May not succeed
not all expressions can be symbolically integrated
- Arbitrarily extensible
define new functions as integrals
 $\Gamma(\cdot)$, distributions, Lebesgue integral
- Semantics preserving
rule chain = formal proof
- Automation
Mathematica, Maple

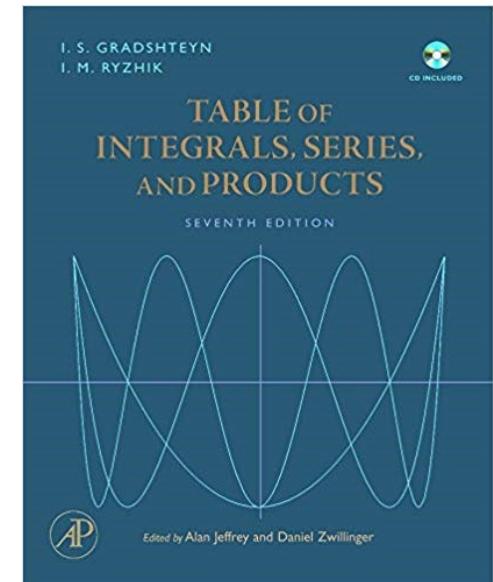
Table of Integrals

BASIC FORMS

- (1) $\int x^n dx = \frac{1}{n+1} x^{n+1}$
- (2) $\int \frac{1}{x} dx = \ln x$
- (3) $\int u dv = uv - \int v du$
- (4) $\int u(x)v'(x)dx = u(x)v(x) - \int v(x)u'(x)dx$

RATIONAL FUNCTIONS

- (5) $\int \frac{1}{ax+b} dx = \frac{1}{a} \ln(ax+b)$
- (6) $\int \frac{1}{(x+a)^2} dx = \frac{-1}{x+a}$
- (7) $\int (x+a)^n dx = (x+a)^n \left(\frac{a}{1+n} + \frac{x}{1+n} \right), n \neq -1$
- (8) $\int x(x+a)^n dx = \frac{(x+a)^{1+n}(nx+x-a)}{(n+2)(n+1)}$

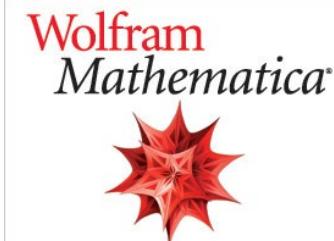


In[31]:=
$$\int_0^{2\pi} \frac{1}{a^2 \cos^2 t + b^2 \sin^2 t} dt$$

 Out[31]:=
$$\frac{2 \sqrt{\frac{b^2}{a^2}} \pi}{b^2}$$

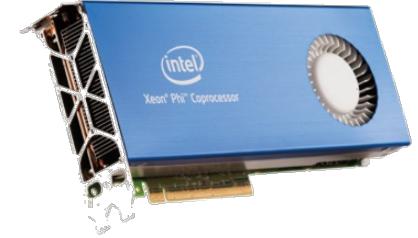
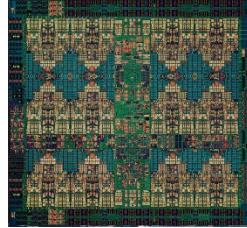
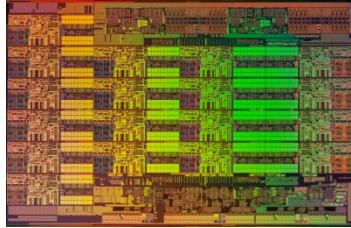
In[33]:=
$$\int_0^{2\pi} \frac{1}{a^2 \left(\frac{e^{it} + e^{-it}}{2} \right)^2 + b^2 \left(\frac{e^{it} - e^{-it}}{2i} \right)^2} dt$$

 Out[33]:= 0



SPIRAL's Target Computing Landscape

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



Intel Xeon 8180M
2.25 Tflop/s, 205 W
28 cores, 2.5–3.8 GHz
2-way–16-way AVX-512

IBM POWER9
768 Gflop/s, 300 W
24 cores, 4 GHz
4-way VSX-3

Nvidia Tesla V100
7.8 Tflop/s, 300 W
5120 cores, 1.2 GHz
32-way SIMT

Intel Xeon Phi 7290F
1.7 Tflop/s, 260 W
72 cores, 1.5 GHz
8-way/16-way LRBni



Snapdragon 835
15 Gflop/s, 2 W
8 cores, 2.3 GHz
A540 GPU, 682 DSP, NEON



Intel Atom C3858
32 Gflop/s, 25 W
16 cores, 2.0 GHz
2-way/4-way SSSE3



Dell PowerEdge R940
3.2 Tflop/s, 6 TB, 850 W
4x 24 cores, 2.1 GHz
4-way/8-way AVX



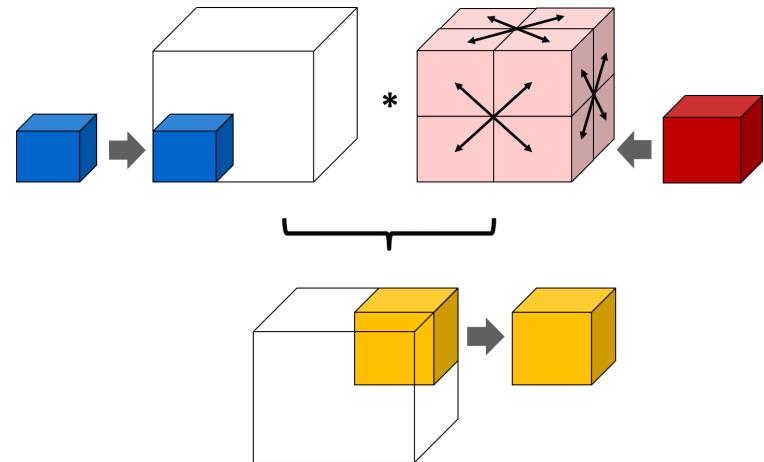
Summit
187.7 Pflop/s, 13 MW
9,216 x 22 cores POWER9
+ 27,648 V100 GPUs

Rules in Internal Domain Specific Language

Linear Transforms

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

Spectral Domain Algorithms



Hardware

- **Multithreading (Multicore)**
- **Vector SIMD (SSE, VMX/Altivec,...)**
- **Message Passing (Clusters, MPP)**
- **Streaming/multibuffering (Cell)**
- **Graphics Processors (GPUs)**
- **Gate-level parallelism (FPGA)**
- **HW/SW partitioning (CPU + FPGA)**

$$\begin{aligned}
 \text{I}_p \otimes \parallel A_{\mu n}, \quad \text{L}_m^{mn} \bar{\otimes} \text{I}_\mu \\
 A \bar{\otimes} \text{I}_\nu &\quad \underbrace{\text{L}_\nu^{2\nu}}_{\text{isa}}, \quad \underbrace{\text{L}_\nu^{2\nu}}_{\text{isa}}, \quad \underbrace{\text{L}_\nu^{2\nu}}_{\text{isa}} \\
 \text{I}_p \otimes \parallel A_n, \quad \underbrace{\text{L}_p^{2^2}}_{\text{all-to-all}} \bar{\otimes} \text{I}_{n/p^2} \\
 \text{I}_n \otimes \text{I}_{2A_{\mu n}}, \quad \text{L}_m^{mn} \bar{\otimes} \text{I}_\mu \\
 \prod_{i=0}^{n-1} A_i, \quad A_n \bar{\otimes} \text{I}_w, \quad P_n \otimes Q_w \\
 \prod_{i=0}^{n-1} A_i, \quad \text{I}_s \bar{\otimes} A, \quad \underbrace{\text{L}_\nu^m}_{\text{bram}} \\
 \underbrace{\text{A}_1}_{\text{fpga}}, \quad \underbrace{\text{A}_2}_{\text{fpga}}, \quad \underbrace{\text{A}_3}_{\text{fpga}}, \quad \underbrace{\text{A}_4}_{\text{fpga}}
 \end{aligned}$$

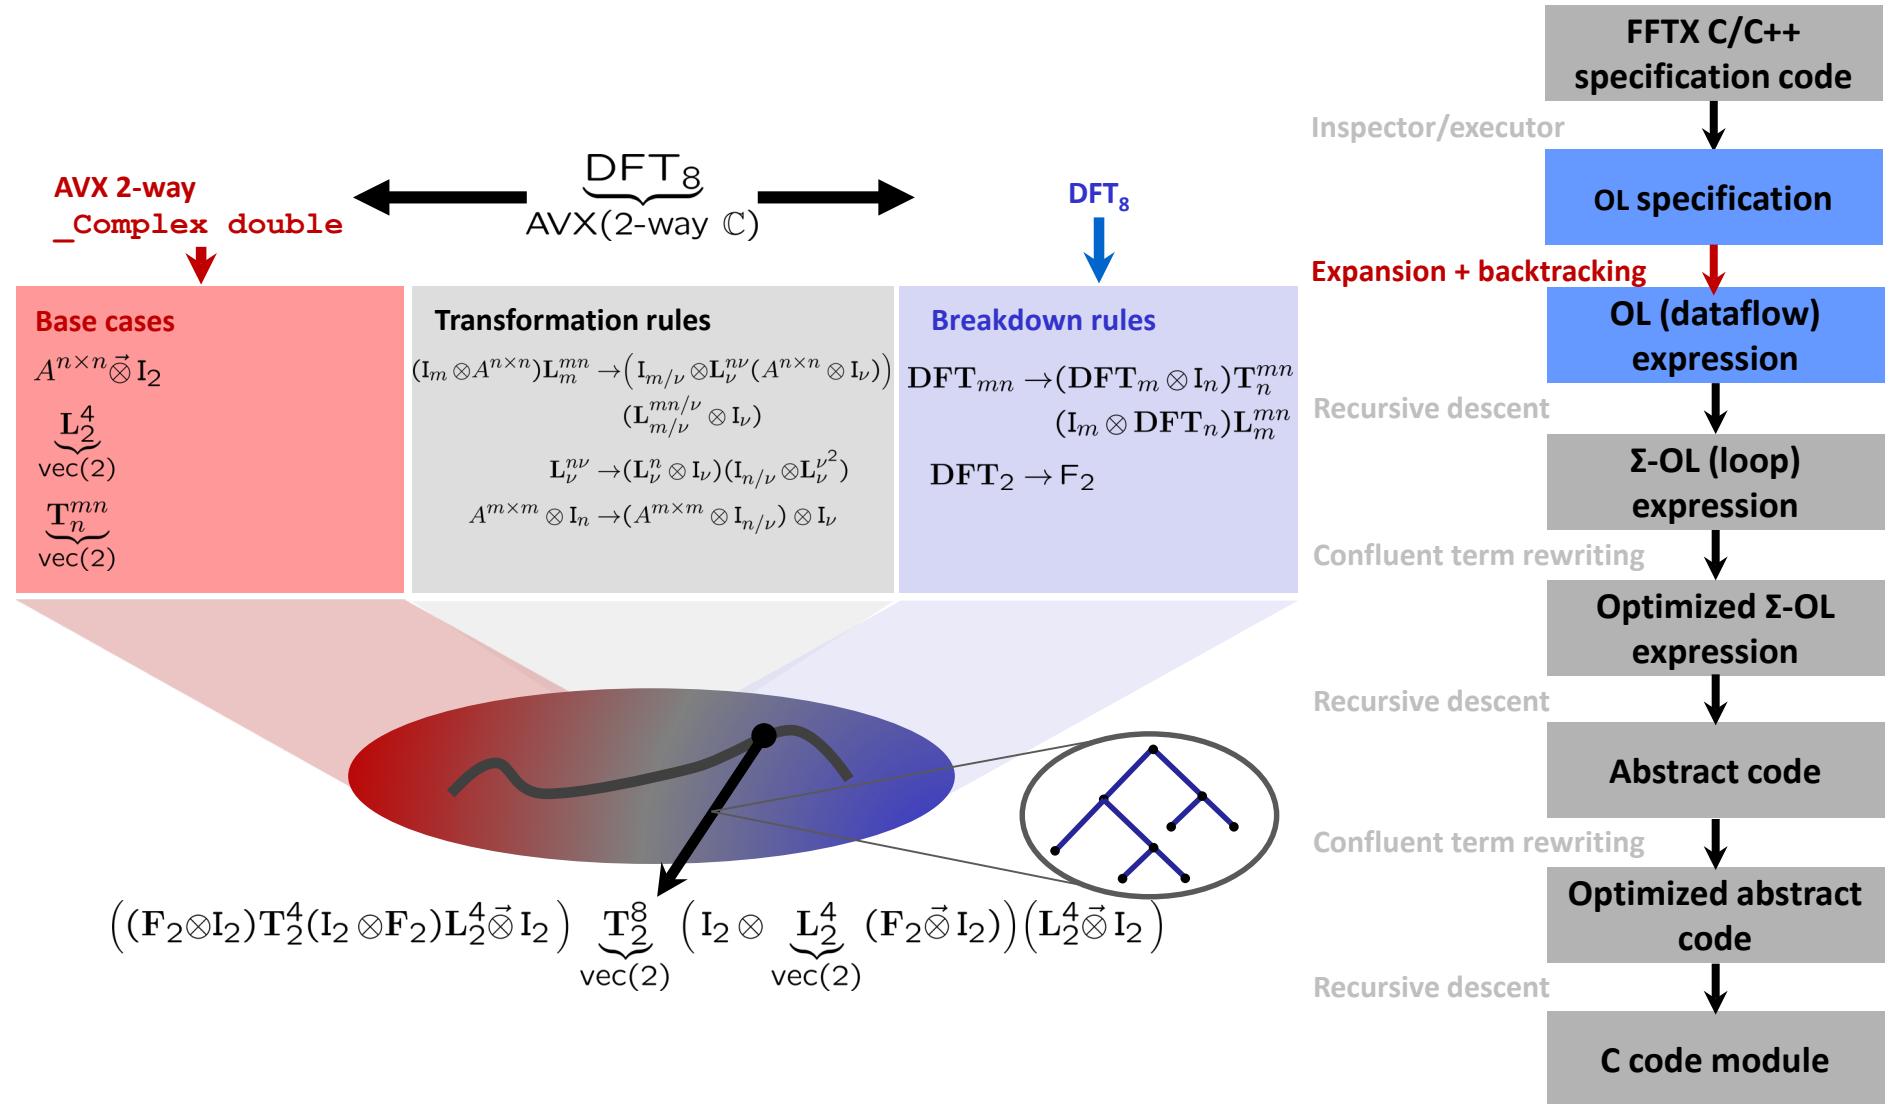
Program Transformations

$$\begin{aligned}
 \underbrace{\text{smp}(p, \mu)}_{AB} &\rightarrow \underbrace{\text{smp}(p, \mu)}_A \underbrace{\text{smp}(p, \mu)}_B \\
 \underbrace{\text{smp}(p, \mu)}_{\text{A}_m \otimes \text{I}_n} &\rightarrow \underbrace{\left(\text{L}_m^{mp} \otimes \text{I}_{n/p} \right) \left(\text{I}_p \otimes (A_m \otimes \text{I}_{n/p}) \right) \left(\text{L}_p^{mp} \otimes \text{I}_{n/p} \right)}_{\text{smp}(p, \mu)} \\
 \underbrace{\text{smp}(p, \mu)}_{\text{L}_m^{mn}} &\rightarrow \begin{cases} \left(\text{I}_p \otimes \text{L}_{m/p}^{mn/p} \right) \left(\text{L}_p^{pn} \otimes \text{I}_{m/p} \right) \\ \text{smp}(p, \mu) \end{cases} \\
 &\quad \begin{cases} \left(\text{L}_m^{pm} \otimes \text{I}_{n/p} \right) \left(\text{I}_p \otimes \text{L}_m^{mn/p} \right) \\ \text{smp}(p, \mu) \end{cases} \\
 \end{cases} \quad \text{Recursive rules}
 \end{aligned}$$

$$\begin{aligned}
 \underbrace{\text{smp}(p, \mu)}_{\text{I}_m \otimes \text{A}_n} &\rightarrow \text{I}_p \otimes \parallel \left(\text{I}_{m/p} \otimes \text{A}_n \right) \\
 \underbrace{\text{smp}(p, \mu)}_{(P \otimes \text{I}_n)} &\rightarrow \left(P \otimes \text{I}_{n/\mu} \right) \bar{\otimes} \text{I}_\mu
 \end{aligned}$$

Base case rules

Autotuning in Constraint Solution Space



Translating an OL Expression Into Code

Constraint Solver Input:

$\underbrace{\text{DFT}_8}_{\text{AVX(2-way)}} \mathbb{C}$

Output =

Ruletrees, expanded into

OL Expression:

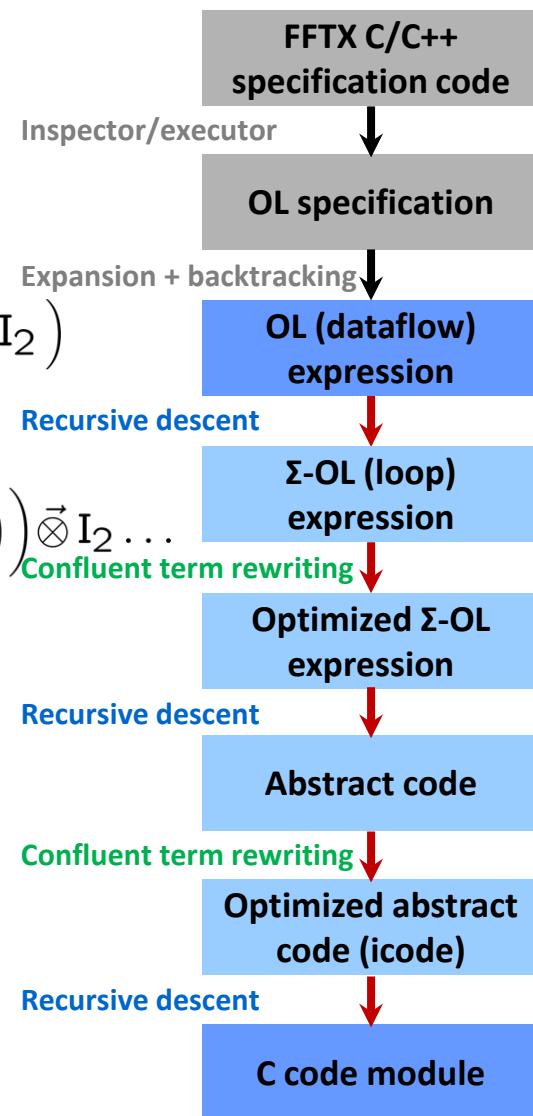
$$\left((F_2 \otimes I_2) T_2^4 (I_2 \otimes F_2) L_2^4 \vec{\otimes} I_2 \right) \underbrace{T_2^8}_{\text{vec}(2)} \left(I_2 \otimes \underbrace{L_2^4}_{\text{vec}(2)} (F_2 \vec{\otimes} I_2) \right) (L_2^4 \vec{\otimes} I_2)$$

Σ -OL:

$$\left(\sum_{j=0}^1 \left(S_{i_2 \otimes (j)_2} F_2 \text{Map}_{x \mapsto \omega_4^{2i+j}}^2 G_{i_2 \otimes (j)_2} \right) \sum_{j=0}^1 \left(S_{(j)_2 \otimes i_2} F_2 G_{i_2 \otimes (j)_2} \right) \right) \vec{\otimes} I_2 \dots$$

C Code:

```
void dft8(_Complex double *Y, _Complex double *X) {
    __m256d s38, s39, s40, s41, ...
    __m256d *a17, *a18;
    a17 = ((__m256d *) X);
    s38 = *(a17);
    s39 = *((a17 + 2));
    t38 = _mm256_add_pd(s38, s39);
    t39 = _mm256_sub_pd(s38, s39);
    ...
    s52 = _mm256_sub_pd(s45, s50);
    *((a18 + 3)) = s52;
}
```



Get Going Quickly: Hockney with OpenACC

```

void ioprunedconv_130_0_62_72_130(double *Y, double *X, double * S) {
    static double D84[260] = {65.5, 0.0, (-0.5000000000001132), (-20.686114762237267),
    (-0.500000000000081), (-10.337014680426078), (-0.5000000000000455),
    ...
for(int i18899 = 0; i18899 <= 1; i18899++) {
    for(int i18912 = 0; i18912 <= 4; i18912++) {
        a9807 = ((2*i18899) + (4*i18912));
        a9808 = (a9807 + 1);
        a9809 = (a9807 + 52);
        a9810 = (a9807 + 53);
        a9811 = (a9807 + 104);
        a9812 = (a9807 + 105);
        s3295 = (*((X + a9807)) + *((X + a9809))
            + *((X + a9811)));
        s3296 = (*((X + a9808)) + *((X + a9810))
            + *((X + a9812)));
        s3297 = (((0.3090169943749474**((X + a9809)))
            - (0.80901699437494745**((X + a9811))))
            + *((X + a9807)));
        ...
        *((104 + Y + a12569)) = ((s3983 - s3987)
            + (0.80901699437494745*t6537)
            + (0.58778525229247314*t6538));
        *((105 + Y + a12569)) = (((s3984 - s3988)
            + (0.80901699437494745*t6538))
            - (0.58778525229247314*t6537));
    }
}

```

FFTX/SPIRAL with OpenACC backend
Compared to cuFFT expert interface



15% faster
on TITAN V



Same speed
on Tesla V100

1,000s of lines of special GPU code, transparently used

Backend: SPIRAL Code Generation

```

__global__ void ker_code0(int *D48, double *D49, double *D50, double *D51, int *D52, double *X) {
    __shared__ double T235[260];
    ...
    if (((threadIdx.x < 13))) {
        for(int i96 = 0; i96 <= 4; i96++) {
            int a31, a32, a33, a34;
            a31 = (2*i96);
            a32 = (threadIdx.x + (13*a31));
            a33 = (threadIdx.x + (13*((a31 + 5) % 10)));
            a34 = (4*i96);
            *((((T235 + 0) + a34) + (20*threadIdx.x))) = (*((T6 + a32)) + *((T6 + a33)));
            *((((1 + (T235 + 0) + a34) + (20*threadIdx.x))) = 0.0;
            *((((2 + (T235 + 0) + a34) + (20*threadIdx.x))) = (*((T6 + a32)) - *((T6 + a33)));
            *((((3 + (T235 + 0) + a34) + (20*threadIdx.x))) = 0.0;
        }
        double t261, t262, t263, t264, t265, t266, t267, t268;
        int a129;
        t263 = (*(((T235+0)+12)+(20*threadIdx.x)))+*(((((T235+0)+8)+(20*threadIdx.x))));
        t264 = (*(((T235+0)+12)+(20*threadIdx.x)))-*(((((T235+0)+8)+(20*threadIdx.x)));
        ...
        *((3 + T5 + a129)) = ((0.58778525229247314*t268) - (0.95105651629515353*t266));
    }
    __syncwarp();
    if (((threadIdx.x < 1))) {
        double t305, t306, t307, t308, t309, t310, t311, t312, t313, t314, t315, t316;
        int a387;
        t305 = (*((T5 + 12)) + *((T5 + 144)));
        ...
    }
}

```

FFTX/SPIRAL with
CUDA backend



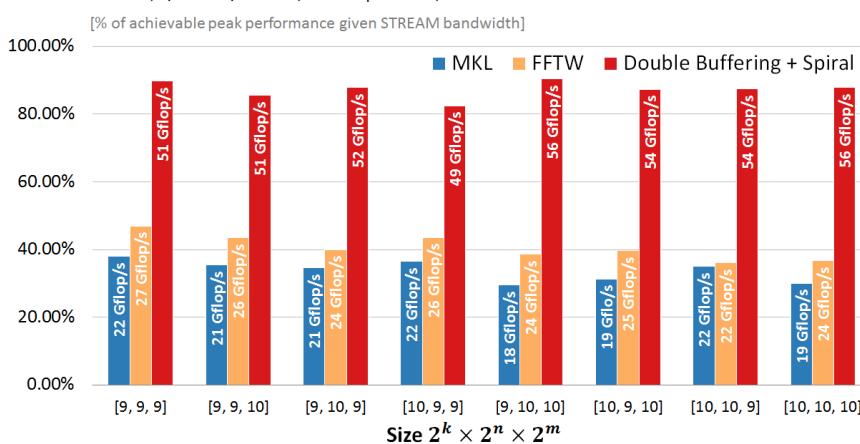
Early result:
220 Gflop/s
70% faster than cuFFT

3,000 lines of code, kernel fusion, cross call data layout transforms

Selected Results: FFTs and Spectral Algorithms

3D FFT performance on Intel Kaby Lake 7700K

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

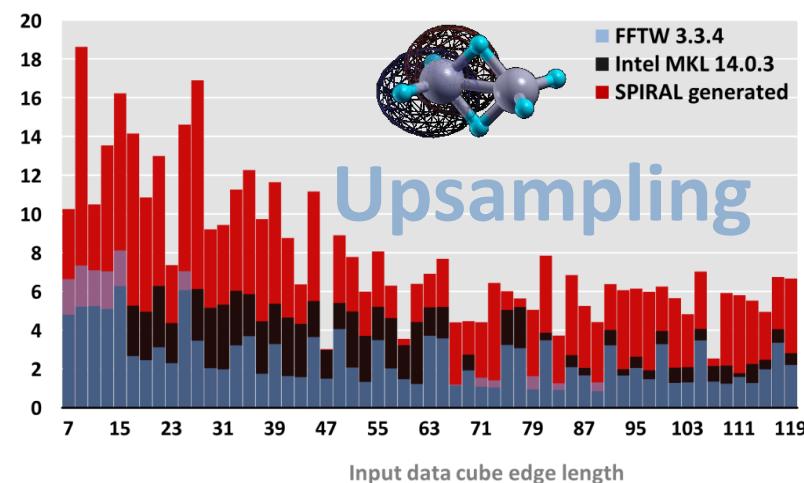


FFT on Multicore

Performance of 2x2x2 Upsampling on Haswell

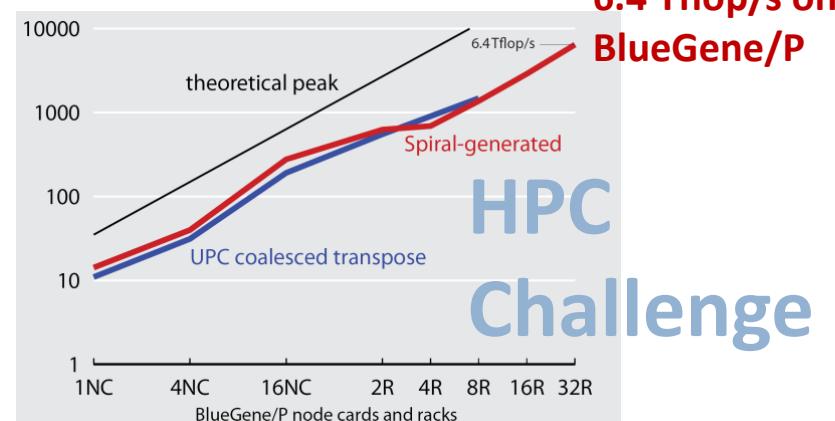
3.5 GHz, AVX, double precision, interleaved input, single core

Performance [Pseudo Gflop/s]



Global FFT (1D FFT, HPC Challenge)

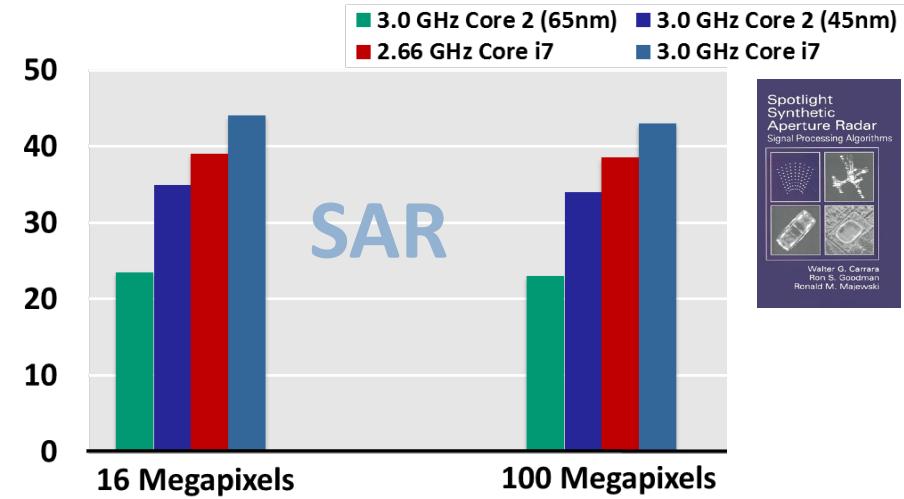
performance [Gflop/s]



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

PFA SAR Image Formation on Intel platforms

performance [Gflop/s]



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

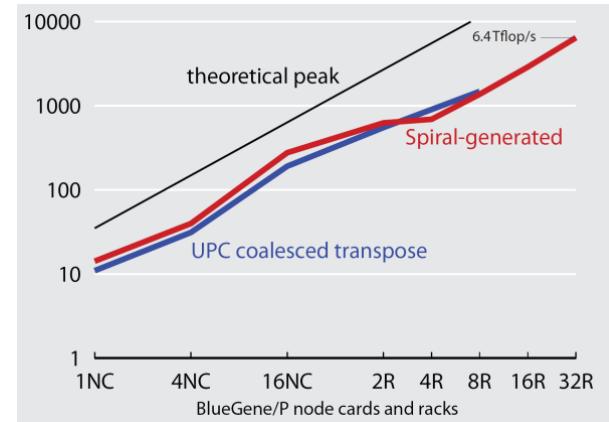
FFTW for BlueGene/L's Double FPU

■ ANL BlueGene/Q Mira

Early Science Program, FFTW for BGQ QPX



Global FFT (1D FFT, HPC Challenge) performance [Gflop/s]



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

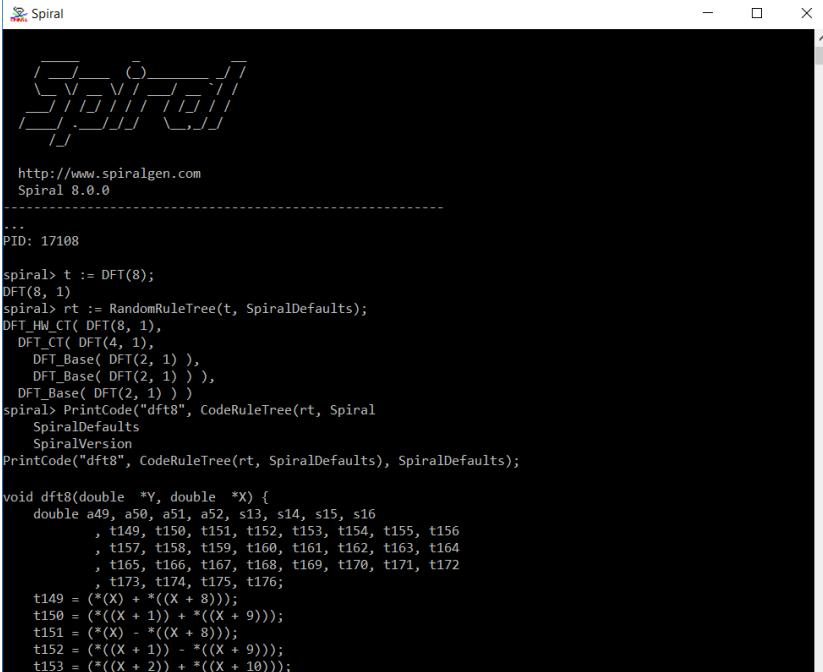


2006 Gordon Bell Prize (Peak Performance Award) with LLNL and IBM

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

SPIRAL 8.1.0: Available Under Open Source

- **Open Source SPIRAL available**
 - non-viral license (BSD)
 - Initial version, effort ongoing to open source whole system
 - Commercial support via SpiralGen, Inc.
- **Developed over 20 years**
 - Funding: DARPA (OPAL, DESA, HACMS, PERFECT, BRASS), NSF, ONR, DoD HPC, JPL, DOE, CMU SEI, Intel, Nvidia, Mercury
- **Open sourced under DARPA PERFECT, continuing under DOE ECP**
- **Tutorial material available online**
[**www.spiral.net**](http://www.spiral.net)



The screenshot shows the Spiral software interface. At the top, there's a logo and the URL <http://www.spiralgen.com>. Below that is the title "Spiral 8.0.0". A tree diagram is displayed above some code. The code is as follows:

```

http://www.spiralgen.com
Spiral 8.0.0
...
PID: 17108

spiral> t := DFT(8);
DFT(8, 1)
spiral> rt := RandomRuleTree(t, SpiralDefaults);
DFT_HW_CTC(DFT(8, 1),
DFT_CTC(DFT(4, 1),
DFT_Base(DFT(2, 1)),
DFT_Base(DFT(2, 1))),
DFT_Base(DFT(2, 1)))
spiral> PrintCode("dft8", CodeRuleTree(rt, Spiral
  SpiralDefaults
  SpiralVersion
PrintCode("dft8", CodeRuleTree(rt, SpiralDefaults), SpiralDefaults);

void dft8(double *Y, double *X) {
    double a49, a50, a51, a52, s13, s14, s15, s16
    , t149, t150, t151, t152, t153, t154, t155, t156
    , t157, t158, t159, t160, t161, t162, t163, t164
    , t165, t166, t167, t168, t169, t170, t171, t172
    , t173, t174, t175, t176;
    t149 = (*X) + *(X + 8));
    t150 = (*((X + 1)) + *((X + 9)));
    t151 = (*X) - *((X + 8));
    t152 = (*((X + 1)) - *((X + 9)));
    t153 = (*((X + 2)) + *((X + 10)));
}

```

