


FFT and Solvers for Exascale: FFTX and SpectralPACK a first look

F. Franchetti, D. G. Spampinato, A. Kulkarni, T. M. Low (Carnegie Mellon University), M. Franusich (SpiralGen, Inc.), D. T. Popovici, A. Canning, P. McCorquodale, B. Van Straalen, P. Colella (Lawrence Berkeley National Laboratory)

Our approach

Have you ever wondered about this?



- No analogue to LAPACK for spectral methods
- Medium-size 3D 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
 - FFTX guru interface is powerful but hard to use, leading to performance loss
 - Low arithmetic intensity and variation of FFT use make library approach hard

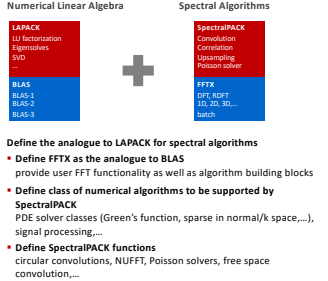
FFTW is de-facto standard interface for FFT

- FFTW 3.X is the high-performance reference implementation: supports multicore/SMP and MPI, and Cell processor
- Vendor libraries support the FFTW 3.X interface: Intel MKL, IBM ESSL, AMD ACML (end-of-life), Nvidia cuFFT, Cray LibSCL/CRAFFT
- Issue 1: 1D FFTW call is standard kernel for many applications
- Parallel libraries and applications reduce to 1D FFTW call: P3DFFT, Qbox, PS/DNS, CPMD, HACC, ...
- Supported by modern languages and environments: Python, Matlab, ...
- Issue 2: FFTW is slowly becoming obsolete
- FFTW 2.1.5 (still in use, 1997), FFTW 3 (2004) minor updates since then
- Risk: loss of high-performance FFT standard library
- Development currently dormant, except for small bug fixes
- No native support for accelerators (GPUs, Xeon Phi, FGAs) and SIMD
- Parallel/MPI version does not scale beyond 32 nodes

FFTX: FFTW revamped for Exascale

- Modernized FFTW-style interface
- Backwards compatible to FFTW 2.X and 3.X
- Old code runs unmodified and gains substantially but not fully
- Small number of new features: futures/delayed execution, offloading, data placement, callback kernels
- Code generation backend using SPIRAL
- Library/application kernels are interpreted as specifications in DSL
- extract semantics from source code and known library semantics
- Compilation and advanced performance optimization: cross-call and cross library optimization, accelerator off-loading...
- Fine control over resource expenditure of optimization: compile-time, initialization-time, invocation time, optimization resources
- Reference library implementation and bindings to vendor libraries
- library-only reference implementation for ease of development

FFTX and SpectralPACK: long-term vision



- Define the analogue to LAPACK for spectral algorithms
- Define FFTX as the analogue to BLAS
 - 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, ...
 - Define SpectralPACK functions: circular convolutions, NUFFT, Poisson solvers, free space convolution, ...

Front end

Poisson's equation in free space

Partial differential equation (PDE) Solution characterization

$\Delta(\Phi) = \rho$

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

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

Poisson's equation is the Laplace operator $\Delta = \nabla \cdot \nabla$

Approach: Green's function


$\Phi(\vec{x}) = \int_D G(\vec{x} - \vec{y}) \rho(\vec{y}) d\vec{y} + G(\vec{x} \cdot \vec{y}) \vec{y}$, $G(\vec{x}) = \frac{1}{4\pi|\vec{x}|^2}$

Solution: $\Phi(\vec{x})$ is convolution of RHS $\rho(\vec{y})$ with Green's function $G(\vec{x})$. Efficient through FFTs (frequency domain)

$\hat{\Phi}_k = \frac{1}{4\pi|k - N\vec{v}|^2}$ if $k \neq N\vec{v}$

Green's function kernel in frequency domain

Hockney free-space convolution



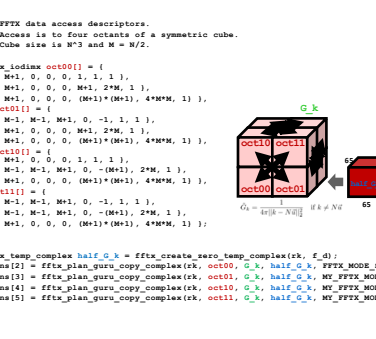
Hockney: Convolution = problem-specific zero padding and output subset

FFTX source code for Hockney

```
fftw_plan pruned_cgal_convolution_plan(fftw_real *in, fftw_real *out,
    int rank = 3,
    batch_rank = 0,
    ...
    fftw_plan plans[5];
    fftw_plan p;
    tmp1 = fftw_create_szero_temp_real(rank, spadded_dims);
    plans[0] = fftw_plan_guru_copy_real(rank, sin_dim, in, tmp1, MY_FFTX_MODE_SUB);
    tmp2 = fftw_create_temp_complex(rank, sfreq_dims);
    plans[1] = fftw_plan_guru_dft_r2c(rank, spadded_dims, batch_rank,
        sbatch_dims, tmp1, tmp2, MY_FFTX_MODE_SUB);
    tmp3 = fftw_create_temp_complex(rank, sfreq_dims);
    plans[2] = fftw_plan_guru_pointwise_c2c(rank, sfreq_dim, batch_rank,
        sbatch_dim, tmp2, tmp3, symbol, (fftw_callback)complex_scaling,
        MY_FFTX_MODE_SUB | FFTX_PW_POINTWISE);
    tmp4 = fftw_create_temp_real(rank, spadded_dims);
    plans[3] = fftw_plan_guru_dft_c2r(rank, spadded_dims, batch_rank,
        sbatch_dims, tmp3, tmp4, MY_FFTX_MODE_SUB);
    plans[4] = fftw_plan_guru_copy_real(rank, sout_dim, tmp4,
        out, MY_FFTX_MODE_SUB);
    p = fftw_plan_compose(sumsubplans, plans, MY_FFTX_MODE_TOY);
    return p;
```

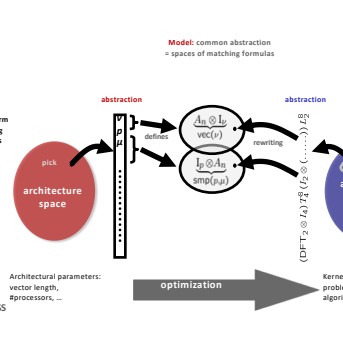
Looks like FFTW calls, but is a specification for SPIRAL

Describing the Green's function symmetry




Back end

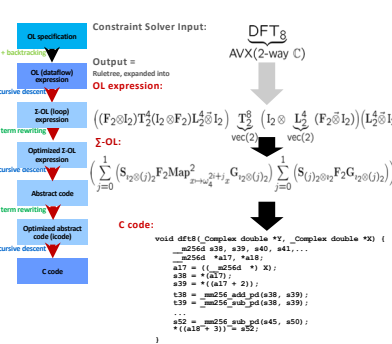
Platform-aware formal program synthesis



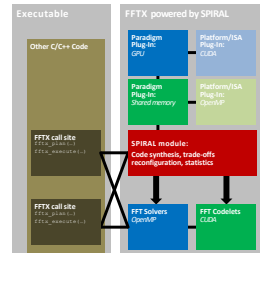
Autotuning in constraint solution space



Translating an OL expression into code



FFTX backend: SPIRAL



Technology + Results

Generated code For Hockney convolution

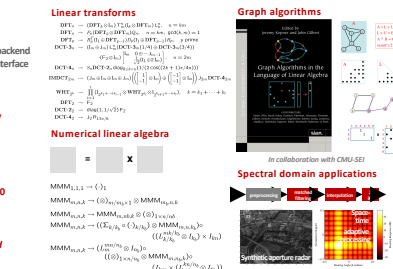
```
#include <math.h>
#define N 1024
#define N3 N*N*N
void hockney_convolution(int *in, int *out, double *rho) {
    // ...
}
// ...
```

15% faster on TITAN v1

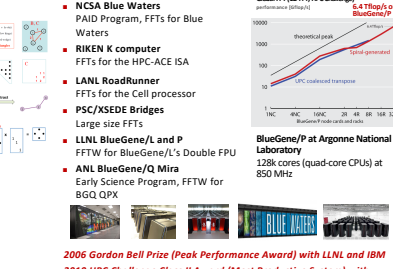
Same speed on Tesla V100

1,000s of lines of code, cross-call optimization, etc., transparently used

Algorithms: rules in domain-specific language



SPIRAL: success in HPC/supercomputing



SPIRAL 8.0: available under open source

- Open-source SPIRAL available
 - non-viral license (BSD)
 - initial version, effort going 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 Sel, Intel, Nvidia, Mercury
 - Open source under DARPA
 - PERFECT
- www.spiral.net

