


Open Source SPIRAL 8.4 Summary

Franz Franchetti
Carnegie Mellon University



Tutorial based on joint work with the Spiral team at CMU, UIUC, and Drexel


SPIRAL: AI for Performance Engineering

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));
    ...
}
```




Some of SPIRAL's Targets

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




Intel Xeon 8380HL
2.5 Tflop/s, 205 W
28 cores, 2.9—4.3 GHz
2-way—16-way AVX-512



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



Nvidia A100
9.7 Tflop/s, 400 W
6912 cores, 1.41 GHz
312 Tflop/s tensor cores




Cerebras WSE2
5.8 Pflop/s 20kW
850,000 cores




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



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





Fugaku
537 Pflop/s, 30 MW
7,630,848 cores A64FX



Google Bristlecone
72 qubits

Current Research Directions

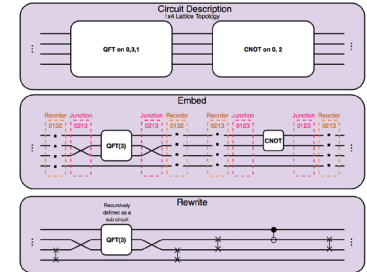
SPIRAL for Graphs



```
# OL Algorithm Specification
Accum(i4, 1, X.N-1,
  Accum_X(i6, [ i4, 0 ], i4,
    Dot([ i6, add(i4, V(1)) ],
      [ i4, add(i4, V(1)) ],
        sub(sub(X.N, i4), V(1)))
  )
)
```



$$\Delta = \Delta + \frac{1}{2} \alpha_{10} A_{00} \alpha_{01}$$

Spiral for Quantum Computing

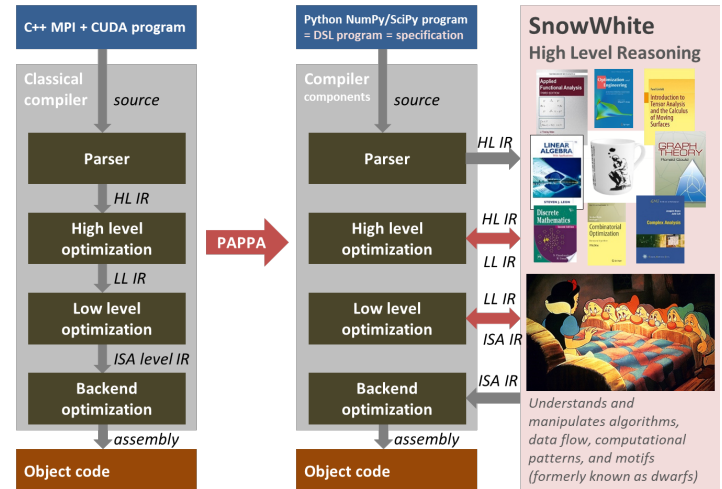


```
[[0,1,0],
 [1,1,0],
 [0,1,1],
 [0,0,1]]
Apply a 3-qubit Fourier
transform to qubits {0,1,2}
Apply a CNOT from qubit 0 -> 2
qCirc(arch, [ ([0,3,1], qFT ), ([0,2], qCNOT) ]);
qEmbed([0,3,1], arch, qFT) * qEmbed([0,2], arch, qCNOT)
Reord([0,1,3,2], arch, F)*Junc([0,2,1,3], F)*Tensor(qFT(3), I(2))*Junc([0,2,1,3], B)*Reord([0,1,3,2], arch, B)
Swap([3,2], 4)
Tensor(I(4), (CNOT(0->1)*CNOT(1->0)*CNOT(0->1)))
qCirc(subarch, [ ([0], qFT), ... ])
Tensor(I(4), (CNOT(1->0)*CNOT(0->1)*CNOT(0->1)))
Tensor(I(4), (CNOT(0->1)*CNOT(1->0)*CNOT(0->1)))
Tensor(I(4), (CNOT(1->0)*CNOT(0->1)*CNOT(1->0)))
```

HW/SW Co-Design



SPiRAL as AI in Compilers



SPIRAL 8.4.0: 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 (ECP, XStack, SciDAC), CMU SEI, Intel, Nvidia, Mercury
 - Open sourced under DARPA PERFECT

```

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

spiral> t := DFT(8);
DFT(8, 1)
spiral> rt := RandomRuleTree(t, SpiralDefaults);
DFT HW_CT( DFT(8, 1),
  DFT_CT( 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));

```

www.spiral.net



SPIRAL's History



DARPA/Defense Sciences Office
Applied & Computational Mathematics Program
Origins




Louis Auslander (1928-1997)

Prepared for DARPA/DSO by Anna Tsao¹

¹ This retrospective was made possible because of the contributions of many individuals who gave generously of their time and energy in providing input, feedback, and writing, with special thanks to William Barker, Gregory Beylkin, Dennis Braunreiter, Russel Caffisch, Douglas Cochran, Ronald Colfman, James Crowley, Benjamin Dembart, Jon Ebert, Abbas Emami-Naeini, Fariba Fahroo, George Fann, Franz Franchetti, Kristen Fuller, Leslie Greengard, Mark Gyure, Robert Harrison, Peter Heller, Jeremy Johnson, Robert Kosut, José Moura, Arje Nachman, Stanley Osher, Mark Oxtley, Vladimir Rokhlin, Leonid Rudin, Randall Sands, Carey Schwartz, Mark Stalzer, Scott Stewart, Bruce Suter, Haydn Wadley, Steven Wax, and Stuart Wolf.


SPIRAL As AI System

Traditionally



High performance library
optimized for given platform

Spiral Approach

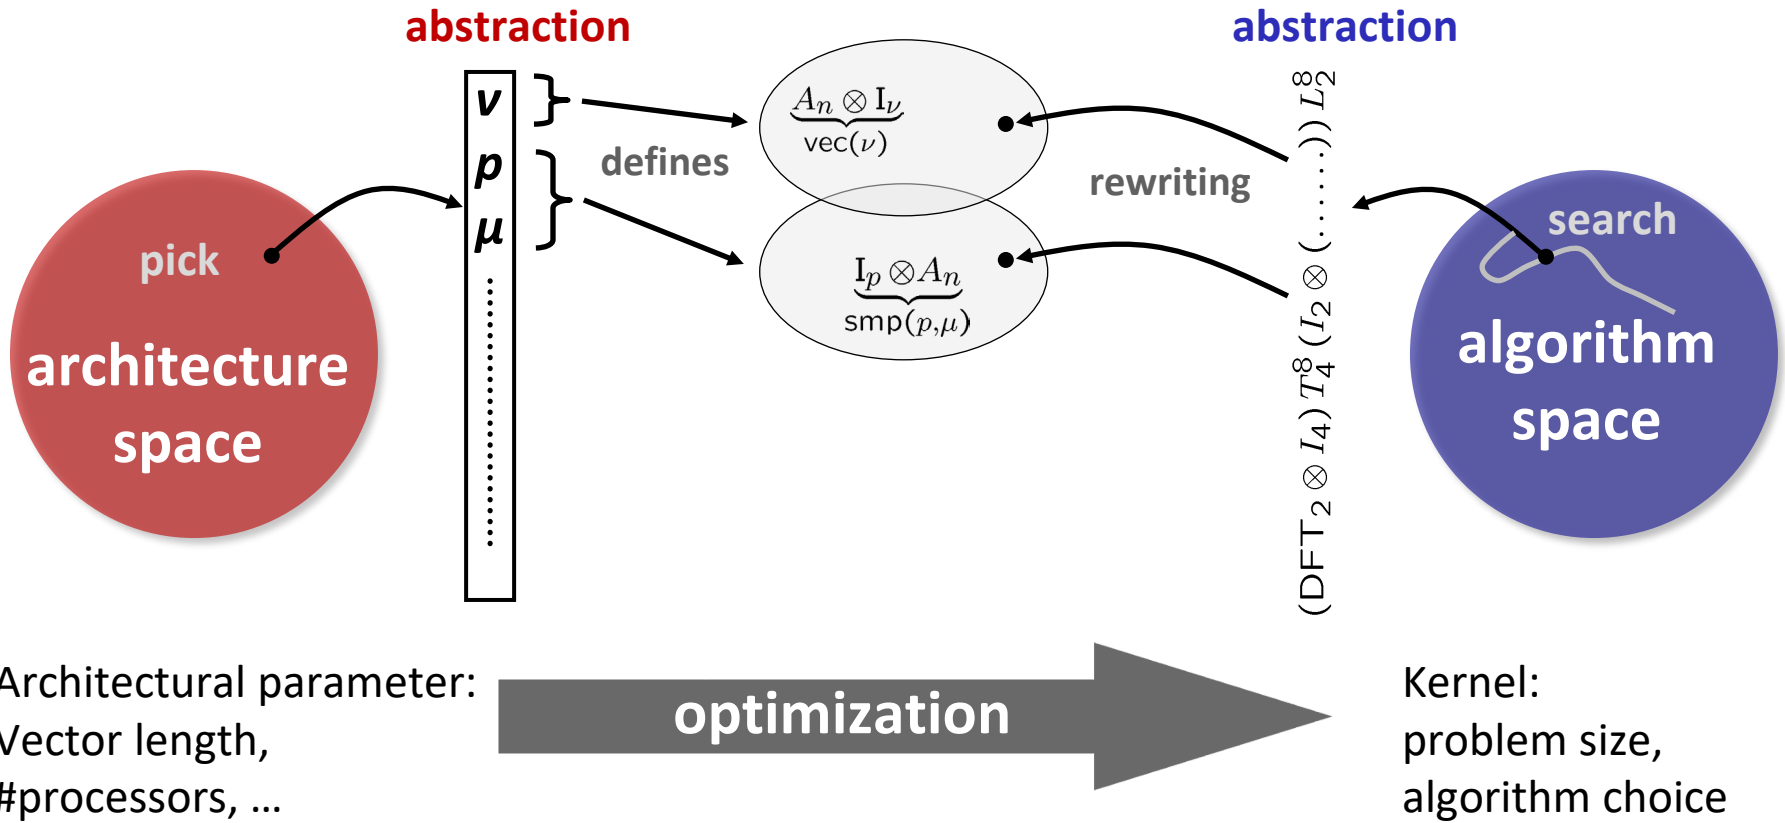


High performance library
optimized for given platform

Comparable performance

Platform-Aware Formal Program Synthesis

Model: common abstraction
= spaces of matching formulas

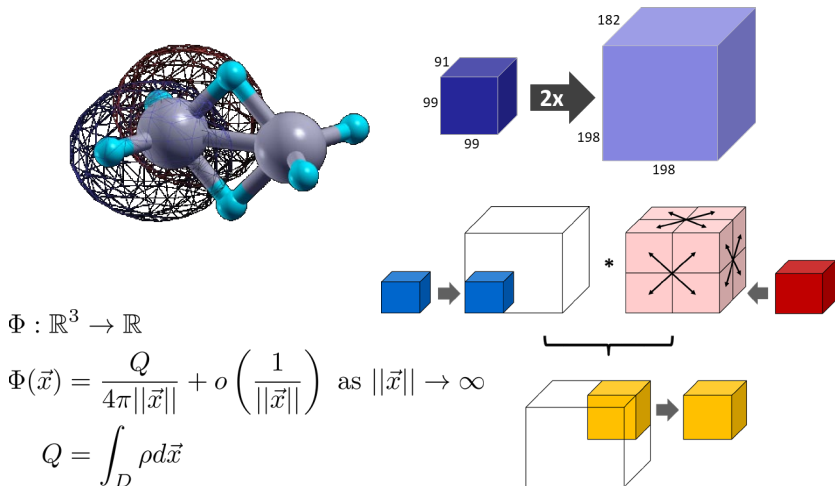


Some Application Domains in SPL


Linear Transforms

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

PDEs/HPC Simulations




Software Defined Radio



$$\begin{aligned} \mathbf{F}_{K,F} &\rightarrow \prod_{i=1}^F \left((\text{I}_{2^{K-2}} \otimes_j B_{F-i,j}) \text{L}_{2^{K-2}}^{2^{K-1}} \right) \\ \mathbf{F}_{K,F} \nu &\rightarrow \prod_{i=1}^F \left((\text{I}_{2^{K-2}/\nu} \otimes_{j_1} \text{L}_{\nu}^{-2\nu} \tilde{B}_{F-i,j_1}^{\nu}) (\text{L}_{2^{K-2}/\nu}^{2^{K-1}/\nu} \otimes \text{I}_{\nu}) \right) \\ B_{i,j} &: \begin{cases} \pi_U = \min_{d_U} (\pi_A + \beta_{A \rightarrow U}, \pi_B + \beta_{B \rightarrow U}) \\ \pi_V = \min_{d_V} (\pi_A + \beta_{A \rightarrow V}, \pi_B + \beta_{B \rightarrow V}) \end{cases} \end{aligned}$$

Radar Image Formation (SAR, STAP)



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



Formal Approach for all Types of Parallelism

- **Multithreading** (Multicore)

$$I_p \otimes_{\parallel} A_{\mu n}, \quad L_m^{mn} \bar{\otimes} I_{\mu}$$

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

$$A \hat{\otimes} I_{\nu} \quad \underbrace{L_2^{2\nu}}_{isa}, \quad \underbrace{L_{\nu}^{2\nu}}_{isa}, \quad \underbrace{L_{\nu}^{\nu^2}}_{isa}$$

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

$$I_p \otimes_{\parallel} A_n, \quad \underbrace{L_p^{p^2} \bar{\otimes} I_{n/p^2}}_{\text{all-to-all}}$$

- **Streaming/multibuffering** (Cell)

$$I_n \otimes_2 A_{\mu n}, \quad L_m^{mn} \bar{\otimes} I_{\mu}$$

- **Graphics Processors** (GPUs)

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

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

$$\prod_{i=0}^{n-1} A_i^{ir}, \quad I_s \tilde{\otimes} A, \quad \underbrace{L_n^m}_{\text{bram}}$$

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

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

Autotuning in Constraint Solution Space

AVX 2-way
_Complex double

$\overbrace{\text{DFT}_8}^{\text{DFT}_8}$
AVX(2-way C)

DFT₈

Base cases

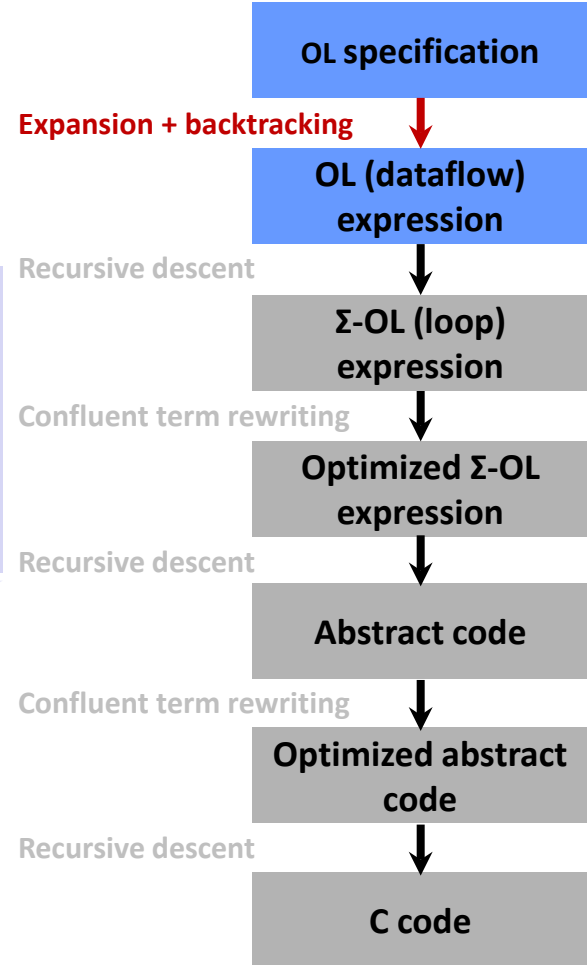

$A^{n \times n} \otimes \vec{I}_2$
 $\underbrace{L_2^4}_{\text{vec}(2)}$
 $\underbrace{T_n^{mn}}_{\text{vec}(2)}$

Transformation rules

$(I_m \otimes A^{n \times n}) L_m^{mn} \rightarrow (I_{m/\nu} \otimes L_{\nu}^{n\nu} (A^{n \times n} \otimes I_{\nu})) (L_{m/\nu}^{mn/\nu} \otimes I_{\nu})$
 $L_{\nu}^{n\nu} \rightarrow (L_{\nu}^n \otimes I_{\nu}) (I_{n/\nu} \otimes L_{\nu}^{\nu^2})$
 $A^{m \times m} \otimes I_n \rightarrow (A^{m \times m} \otimes I_{n/\nu}) \otimes I_{\nu}$

Breakdown rules

$\text{DFT}_{mn} \rightarrow (\text{DFT}_m \otimes I_n) T_n^{mn}$
 $(I_m \otimes \text{DFT}_n) L_m^{mn}$
 $\text{DFT}_2 \rightarrow F_2$



Translating an SPL Expression Into Code

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

Output =
Ruletree, expanded into

SPL Expression:

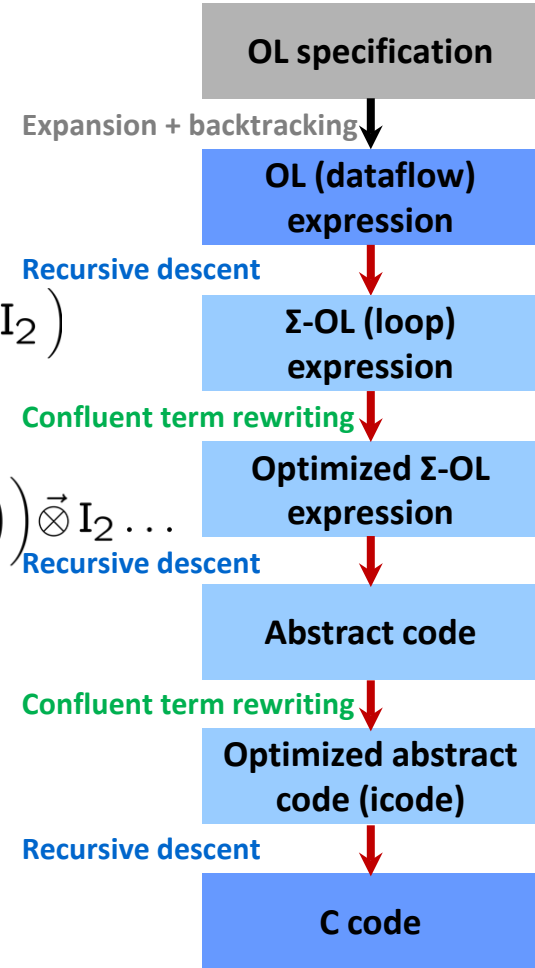
$$\left((F_2 \otimes I_2) T_2^4 (I_2 \otimes F_2) L_2^4 \vec{\otimes} I_2 \right) \underbrace{T_2^8}_{\text{vec}(2)} \left(I_2 \otimes \underbrace{L_2^4}_{\text{vec}(2)} (F_2 \vec{\otimes} I_2) \right) \left(L_2^4 \vec{\otimes} I_2 \right)$$

Σ -SPL:

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

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;
}
```





Some SPIRAL Generated C/AVX Code




```

int dwmonitor(float *X, double *D) {
    __m128d u1, u2, u3, u4, u5, u6, u7, u8 , x1, x10, x13, x14, x17, x18, x19, x2, x3, x4, x6, x7, x8, x9;
    int w1;
    unsigned _xm = __mm_getcsr();
    __mm_setcsr(_xm & 0xffff0000 | 0x0000dfc0);
    u5 = __mm_set1_pd(0.0);
    u2 = __mm_cvtps_pd(__mm_addsub_ps(__mm_set1_ps FLT_MIN, __mm_set1_ps(X[0])));
    u1 = __mm_set_pd(1.0, (-1.0));
    for(int i5 = 0; i5 <= 2; i5++) {
        x6 = __mm_addsub_pd(__mm_set1_pd((DBL_MIN + DBL_MIN)), __mm_loaddup_pd(&(D[i5])));
        x1 = __mm_addsub_pd(__mm_set1_pd(0.0), u1);
        x2 = __mm_mul_pd(x1, x6);
        x3 = __mm_mul_pd(__mm_shuffle_pd(x1, x1, MM_SHUFFLE2(0, 1)), x6);
        x4 = __mm_sub_pd(__mm_set1_pd(0.0), __mm_min_pd(x3, x2));
        u3 = __mm_add_pd(__mm_max_pd(__mm_shuffle_pd(x4, x4, MM_SHUFFLE2(0, 1)), __mm_max_pd(x3, x2)), __mm_set1_pd(DBL_MIN));
        u5 = __mm_add_pd(u5, u3);
        x7 = __mm_addsub_pd(__mm_set1_pd(0.0), u1);
        x8 = __mm_mul_pd(x7, u2);
        x9 = __mm_mul_pd(__mm_shuffle_pd(x7, x7, MM_SHUFFLE2(0, 1)), u2);
        x10 = __mm_sub_pd(__mm_set1_pd(0.0), __mm_min_pd(x9, x8));
        u1 = __mm_add_pd(__mm_max_pd(__mm_shuffle_pd(x10, x10, MM_SHUFFLE2(0, 1)), __mm_max_pd(x9, x8)), __mm_set1_pd(DBL_MIN));
    }
    u6 = __mm_set1_pd(0.0);
    for(int i3 = 0; i3 <= 1; i3++) {
        u8 = __mm_cvtps_pd(__mm_addsub_ps(__mm_set1_ps FLT_MIN, __mm_set1_ps(X[(i3 + 1)])));
        u7 = __mm_cvtps_pd(__mm_addsub_ps(__mm_set1_ps FLT_MIN, __mm_set1_ps(X[(3 + i3)])));
        x14 = __mm_add_pd(u8, __mm_shuffle_pd(u7, u7, MM_SHUFFLE2(0, 1)));
        x13 = __mm_shuffle_pd(x14, x14, MM_SHUFFLE2(0, 1));
        u4 = __mm_shuffle_pd(__mm_min_pd(x14, x13), __mm_max_pd(x14, x13), MM_SHUFFLE2(1, 0));
        u6 = __mm_shuffle_pd(__mm_min_pd(u6, u4), __mm_max_pd(u6, u4), MM_SHUFFLE2(1, 0));
    }
    x17 = __mm_addsub_pd(__mm_set1_pd(0.0), u6);
    x18 = __mm_addsub_pd(__mm_set1_pd(0.0), u5);
    x19 = __mm_cmpge_pd(x17, __mm_shuffle_pd(x18, x18, MM_SHUFFLE2(0, 1)));
    w1 = (__mm_testc_si128(__mm_castpd_si128(x19), __mm_set_epi32(0xffffffff, 0xffffffff, 0xffffffff, 0xffffffff)) -
        (__mm_testnzc_si128(__mm_castpd_si128(x19), __mm_set_epi32(0xffffffff, 0xffffffff, 0xffffffff, 0xffffffff))));
    __asm nop;
    if (__mm_getcsr() & 0x0d) {
        __mm_setcsr(_xm);
        return -1;
    }
    __mm_setcsr(_xm);
    return w1;
}

```

Dynamic Range: Single Node/Shared Memory

Architecture ↑



- Dynamic range:**
- <1W – 10 kW
 - 1 to 500/41k cores (CPU/GPU)
 - 500 MB – 6 TB RAM
 - 1 Gflop/s – 200 Tflop/s
 - 20 years of release dates



...one source code, one tool, always highest performance...



2000: SSE2



2007: SSSE3



2011: AVX



2013: AVX2



2019: AVX512



Symbolic Verification

- Transform = Matrix-vector multiplication
matrix fully defines the operation

$$\text{DFT}_4 = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & j & -1 & -j \\ 1 & -1 & 1 & -1 \\ 1 & -j & -1 & j \end{bmatrix}$$

= ?

- Algorithm = Formula
represents a matrix expression, can be evaluated to a matrix

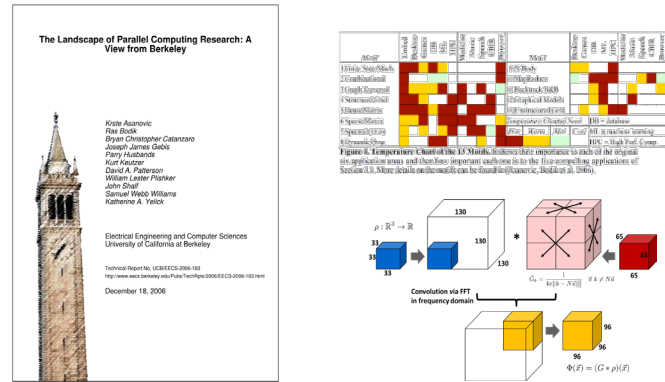
$$(\text{DFT}_2 \otimes \text{I}_2) T_2^4 (\text{I}_2 \otimes \text{DFT}_2) L_2^4 = \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & -1 & 0 \\ 0 & 1 & 0 & -1 \end{bmatrix} \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & -1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & -1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & j \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

SPIRAL as DSL Frontend and Compiler Plugin

Multi-language, Multi target

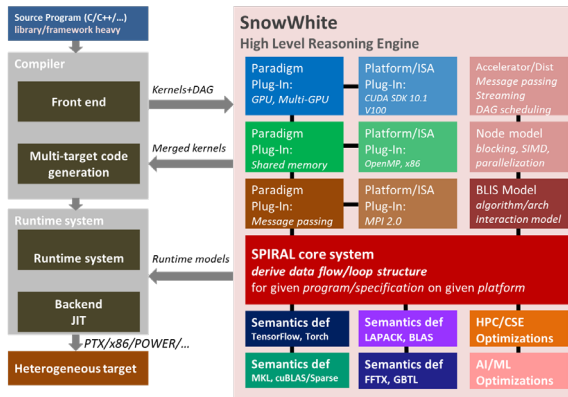


Cross-motif optimization



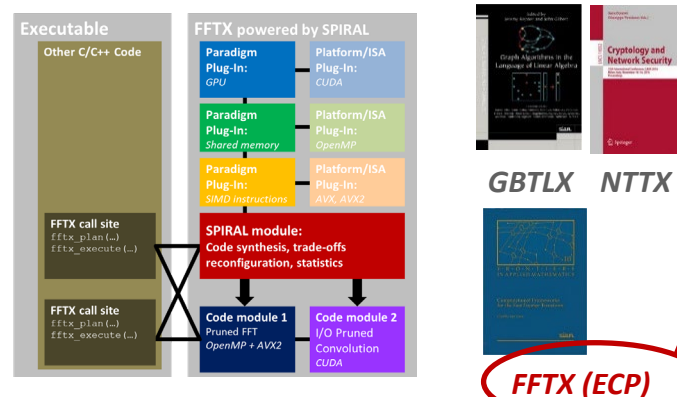
Cross-call, cross-library, cross-motif

SnowWhite: SPIRAL inside compilers



DARPA PAPP, X-Stack Bluestone

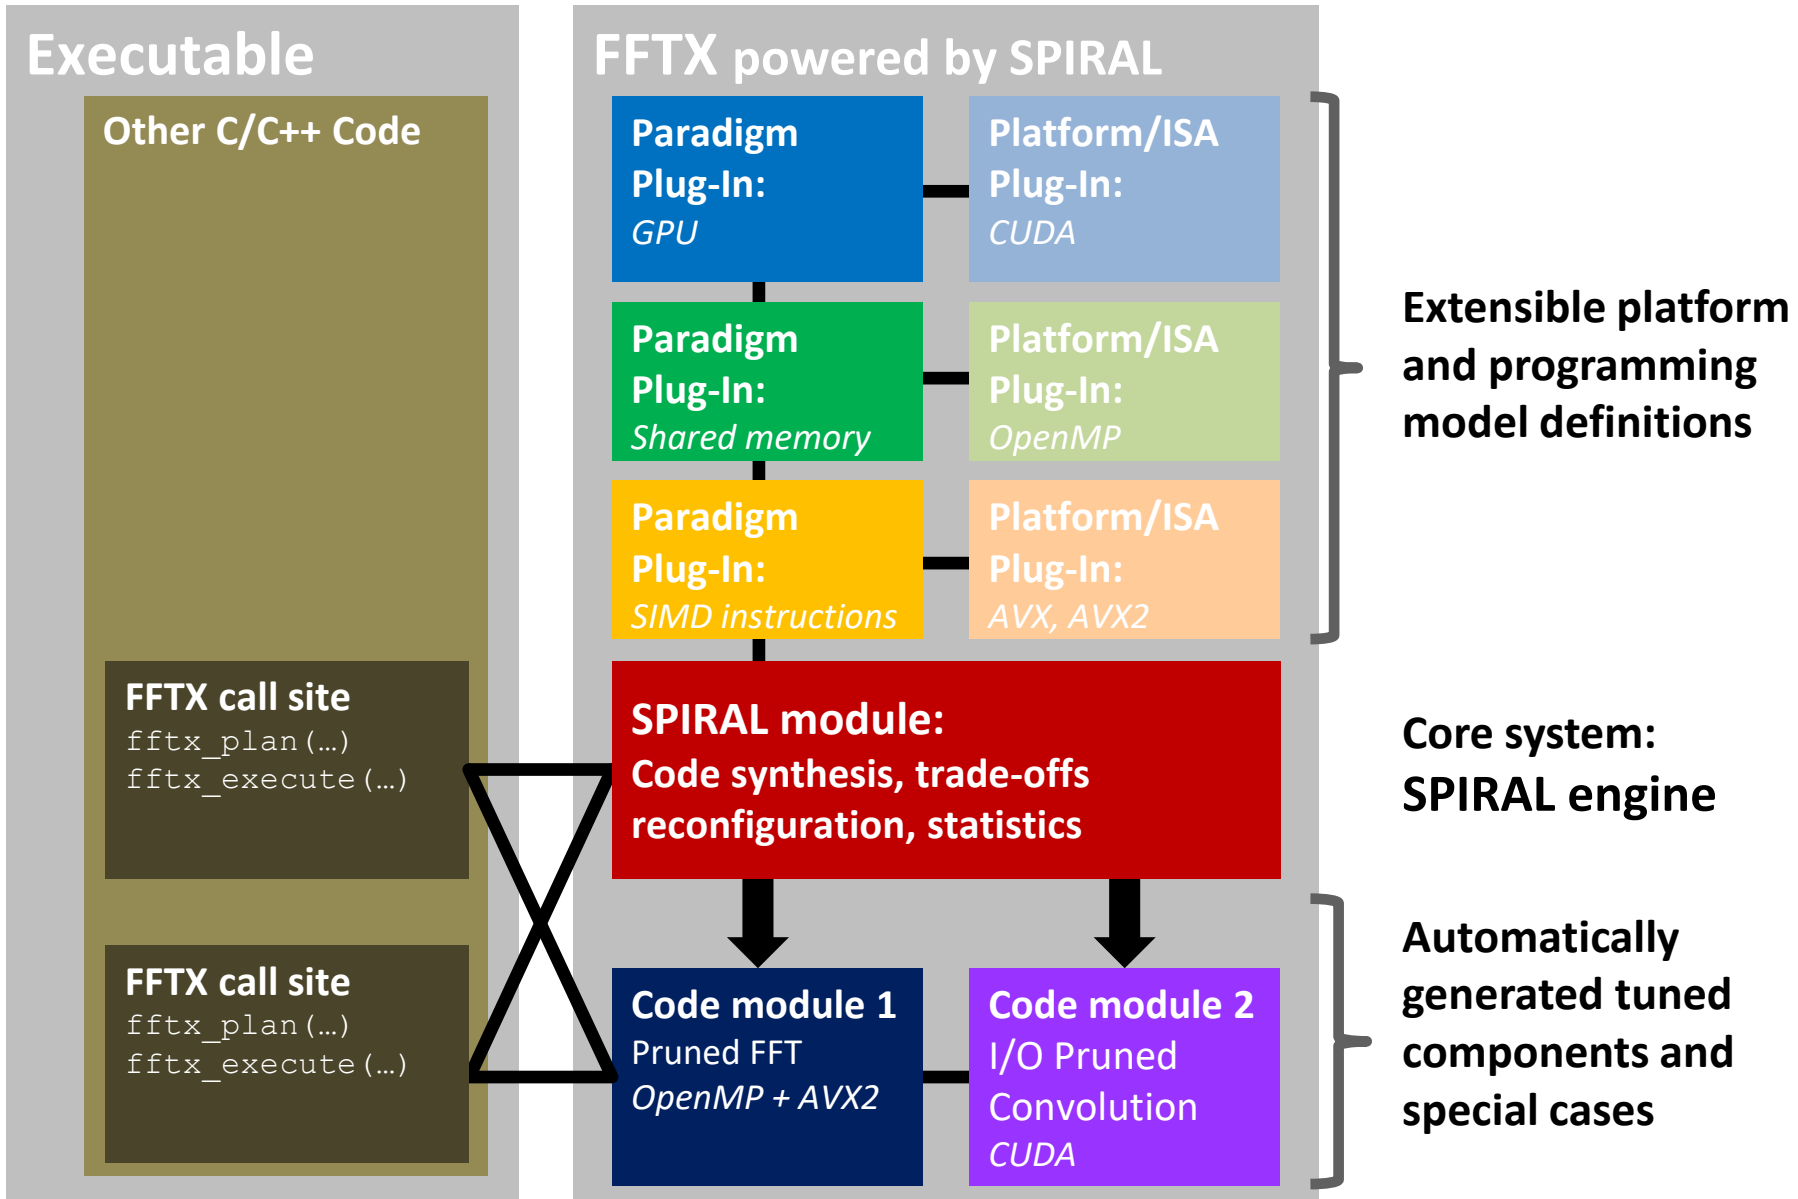
LibraryX, powered by SPIRAL



Multiple active libraries, one infrastructure

One instance of a family of active libraries

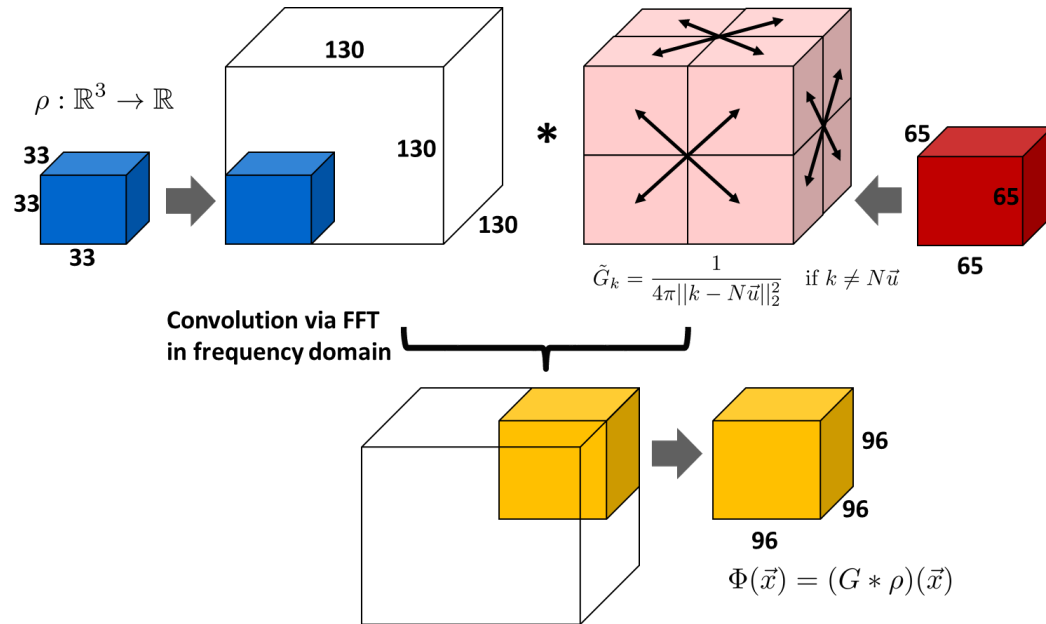
SPIRAL as Core Engine in LibraryX





In Development: Python API and Beyond

```
embIn = numpy.zeros((N, N, N))
embIn[0:Ns, 0:Ns, 0:Ns] = In
FI = numpy.fft.rfftn(embIn)
C = FI * S
embOut = numpy.fft.irfftn(C)
Out = embOut[N-Nd:N, N-Nd:N, N-Nd:N]
```





Beyond Linear: Operator Language

Definition


- **Operator: Multiple vectors -> Multiple vectors**
- **Stateless**
- **Higher-dimensional data is linearized**
- **Operators are potentially nonlinear**

$$M : \begin{cases} \mathbb{C}^{n_0} \times \dots \times \mathbb{C}^{n_{k-1}} \rightarrow \mathbb{C}^{N_0} \times \dots \times \mathbb{C}^{N_{\ell-1}} \\ (\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_{k-1}) \mapsto M(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_{k-1}) \end{cases}$$

Example: Scalar product

$$\langle \cdot, \cdot \rangle_n: \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$$

$$\left((x_i)_{i=0, \dots, n-1}, (y_i)_{i=0, \dots, n-1} \right) \mapsto \sum_{i=0}^{n-1} x_i y_i$$



OL: Linear Algebra + Functional Programming

■ Application specific: Safety Distance as Rewrite Rule

$$\text{SafeDist}_{V,A,b,\varepsilon}(\cdot, \cdot, \cdot) \rightarrow \left(P[x, (a_0, a_1, a_2)](\cdot) < d_{\infty}^2(\cdot, \cdot) \right) (\cdot, \cdot, \cdot)$$

$$\text{with } a_0 = \frac{1}{2b}, a_1 = \frac{V}{b} + \varepsilon \left(\frac{A}{b} + 1 \right), a_2 = \left(\frac{A}{b} + 1 \right) \left(\frac{A}{2} \varepsilon^2 + \varepsilon V \right)$$

Problem specification: hand-developed or automatically produced

■ One-time effort: mathematical library

$$d_{\infty}^n(\cdot, \cdot) \rightarrow \|\cdot\|_{\infty}^n \circ (-)_n$$



$$(\diamond)_n \rightarrow \text{Pointwise}_{n \times n, (a,b) \mapsto a \diamond b}, \quad \diamond \in \{+, -, \cdot, \wedge, \vee, \dots\}$$

$$\|\cdot\|_{\infty}^n \rightarrow \text{Reduction}_{n, (a,b) \mapsto \max(|a|, |b|)}$$

$$\langle \cdot, \cdot \rangle_n \rightarrow \text{Reduction}_{n, (a,b) \mapsto a+b} \circ \text{Pointwise}_{n \times n, (a,b) \mapsto ab}$$

$$P[x, (a_0, \dots, a_n)] \rightarrow \langle (a_0, \dots, a_n), \cdot \rangle \circ (x^i)_n$$

$$(x^i)_n \rightarrow \text{Induction}_{n, (a,b) \mapsto ab, 1}$$



Library of well-known identities expressed in OL

More Information:

www.spiral.net

www.spiralgen.com