

Research Review 2017

Automated Code Generation for High-Performance, Future-Compatible Graph Libraries

Dr. Scott McMillan, Senior Research Scientist

CMU PI: Prof. Franz Franchetti, ECE

Copyright 2017 Carnegie Mellon University. All Rights Reserved.

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

The view, opinions, and/or findings contained in this material are those of the author(s) and should not be construed as an official Government position, policy, or decision, unless designated by other documentation.

NO WARRANTY. THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

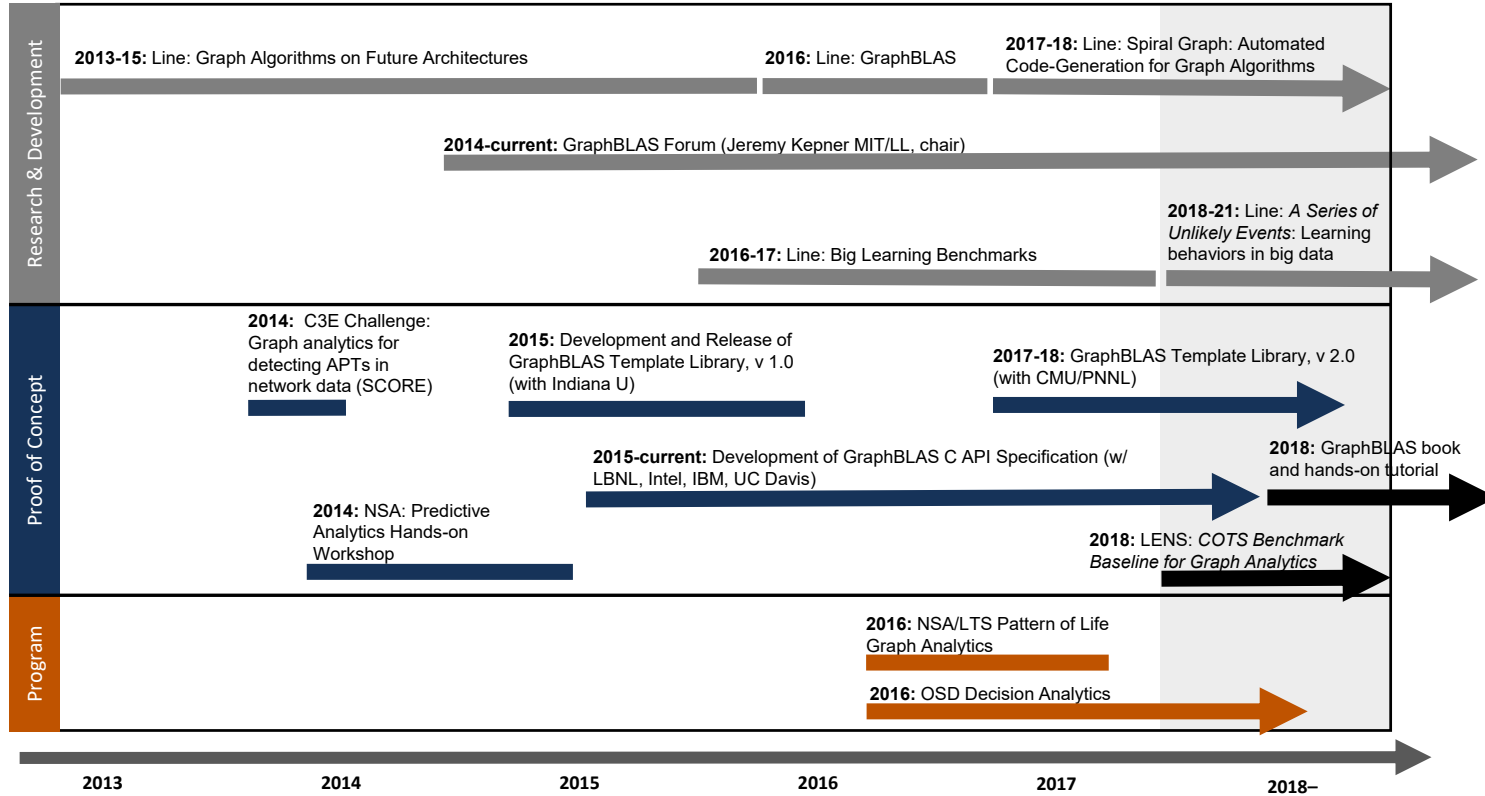
[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution. Please see Copyright notice for non-US Government use and distribution.

This material may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

Carnegie Mellon® is registered in the U.S. Patent and Trademark Office by Carnegie Mellon University.

DM17-0774

Data-Intensive Computing Efforts at the SEI



SpiralGraph: Automated Code Generation for *Future-Compatible*, High-Performance Graph Libraries

Problem:

- Heterogeneous high-performance computing (HHPC) architectures are becoming more complex (the NSCI push to exascale).
- Graph algorithms are difficult to program efficiently even on today's hardware architectures.
- Exascale trend: Programming these systems will be much more difficult.¹

Solution:

- Create an automated code generation tool that produces high-performance graph algorithm implementations for specified hardware.
- **Make graph algorithms performance-portable and future-compatible.**

Approach:

- Create formal abstractions of graph algorithms and primitives (build on GraphBLAS).
- Extend formal abstractions of chosen hardware architectures (build on Spiral and DARPA HACMS, DESA, PERFECT, BRASS).
- Create tool for mapping graph algorithms to hardware architectures for efficient code generation of data-intensive applications.

¹FACT SHEET: National Strategic Computing Initiative, 29 July 2015.

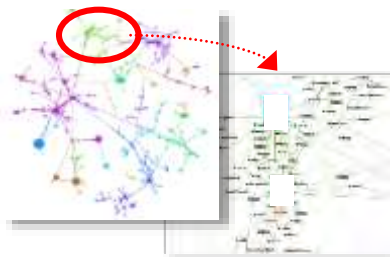
Graph Analysis is *Important and Pervasive*

ISR



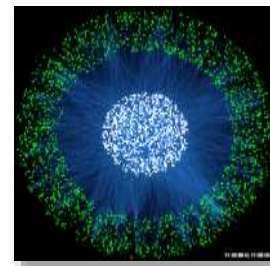
- Graphs represent entities and relationships detected through multi-INT sources
- 1,000s – 1,000,000s tracks and locations
- GOAL: Identify anomalous patterns of life

Social



- Graphs represent relationships between individuals or documents
- 10,000s – 10,000,000s individual and interactions
- GOAL: Identify hidden social networks

Cyber

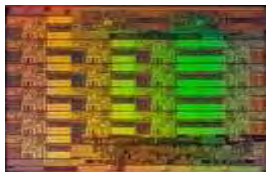


- Graphs represent communication patterns of computers on a network
- 1,000,000s – 1,000,000,000s network events
- GOAL: Identify cyber attacks or malicious software

Common Goal: Detection of subtle patterns in massive graphs

Slide credit: Jeremy Kepner, et al. "Mathematical Foundations of the GraphBLAS", IEEE HPEC, Sept. 2016.

Today's Computing Landscape



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 P100
10.6 Tflop/s, 250 W
3584 cores, 1.48 GHz
64-way SIMT



Intel Xeon Phi
1.2 Tflop/s, 300 W
61 cores, 1.24 GHz
8-way/16-way LRBni



Qualcomm 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



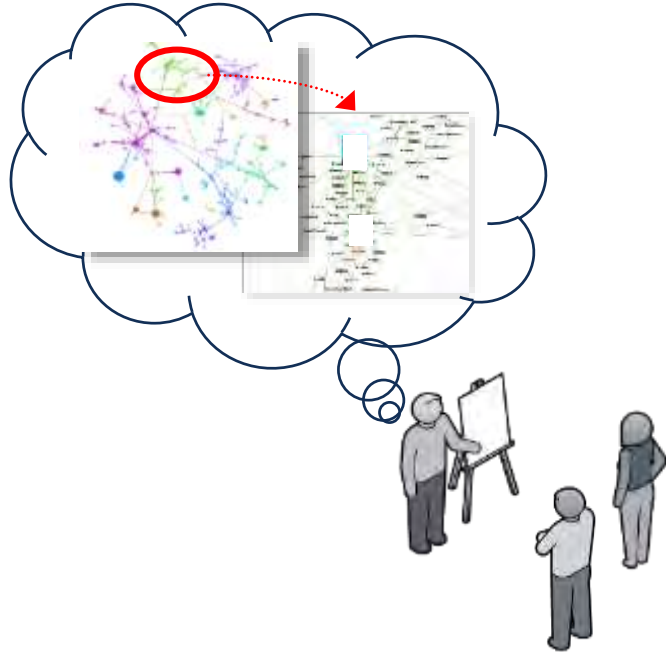
IBM BlueGene/Q
10 Pflop/s, 8 MW
48k x 16 cores, 1.6 GHz
4-way QPX

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

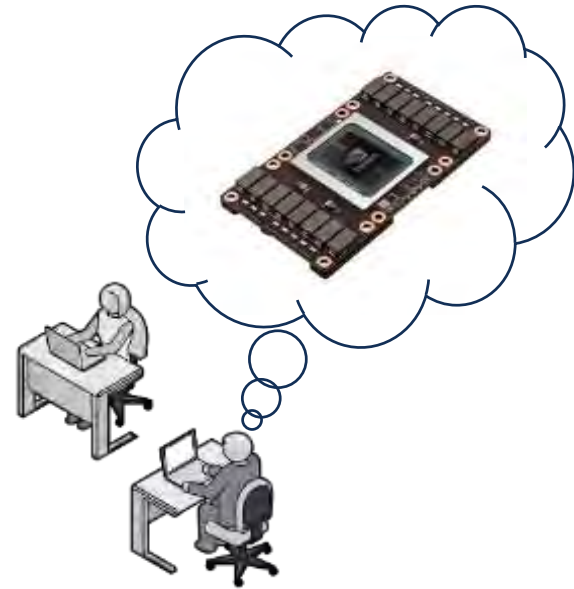
Slide credit: Franz Franchetti, "SPIRAL: Automated Code Generation of Performance Libraries," 2016.

Separation of Concerns

Separate the complexity of graph analysis from the complexity of hardware systems:



Separation of Concerns



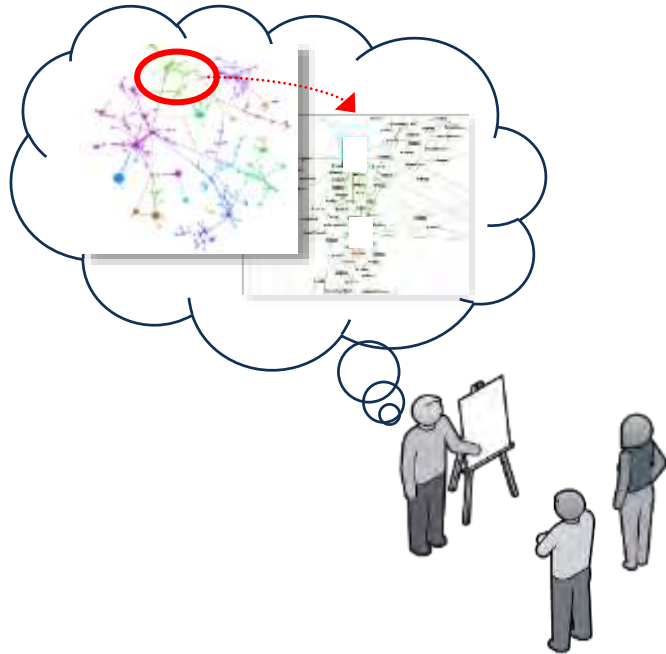
GraphBLAS Primitives

Operation	Description
mxm, mxv, vxm	Perform matrix multiplication (e.g., breadth-first traversal)
eWiseAdd, eWiseMult	Element-wise addition and multiplication of matrices (e.g., graph union, intersection)
extract	Extract a sub-matrix from a larger matrix (e.g., sub-graph selection)
assign	Assign to a sub-matrix of a larger matrix (e.g., sub-graph assignment)
apply	Apply unary function to each element of matrix (e.g., edge weight modification)
reduce	Reduce along columns or rows of matrices (vertex degree)
transpose	Swaps the rows and columns of a sparse matrix (e.g., reverse directed edges)
build	Build a matrix representation from row, column, value tuples
extractTuples	Extract the row, column, value tuples from a matrix representation

<http://graphblas.org> "A. Buluc, T. Mattson, S. McMillan, J. Moreira, C. Yang, "The GraphBLAS C API Specification, v 1.0.2," updated August 2017.

Separation of Concerns

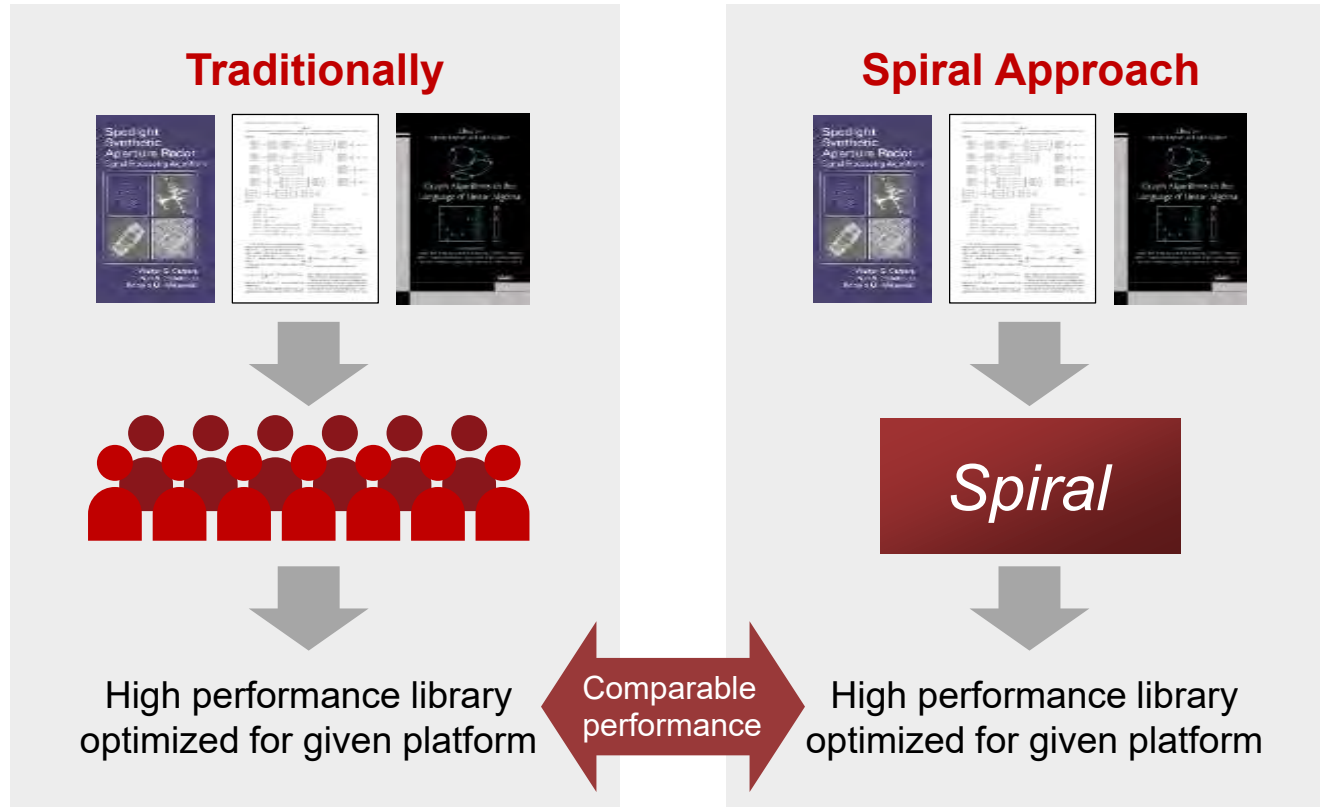
GOAL: write once, run everywhere...fast (**with** help from hardware experts).



GraphBLAS Application Programming Interface (API)

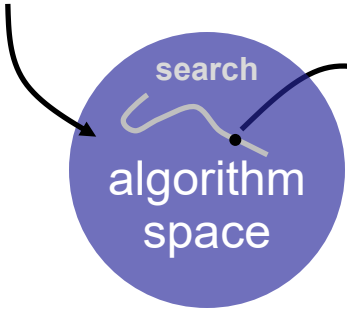


What is Spiral?



Spiral: Platform-Aware Formal Program Synthesis

GraphBLAS Math:
 $C(L, z) = (L \oplus \cdot \otimes L^T)$
 count = $\bigoplus_{i,j} C(i,j)$



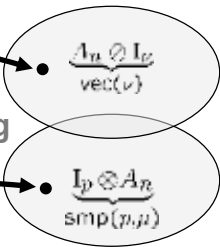
Kernel:
 problem size,
 algorithm choice

Model: common abstraction
 = spaces of matching formulas

abstraction

$(DFT_2 \otimes I_4) T_4^8 (I_2 \otimes (\dots)) L_2^8$

rewriting



defines

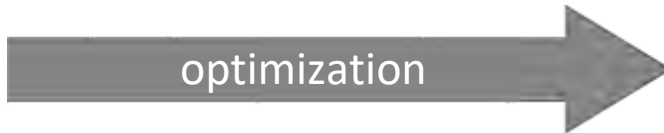
abstraction



pick

architecture
 space

Architectural parameters:
 Vector length,
 #processors, ...



GraphBLAS Primitives: The Math

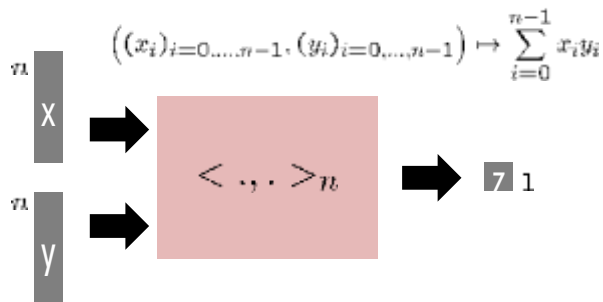
Operation	Mathematical Description	Output	Inputs
mxm	$C \langle \neg M, z \rangle = C \odot (A^T \oplus \otimes B^T)$	C	$\neg, M, z, \odot, A, T, \oplus \otimes, B, T$
mxv, (vxm)	$c \langle \neg m, z \rangle = c \odot (A^T \oplus \otimes b)$	c	$\neg, m, z, \odot, A, T, \oplus \otimes, b$
eWiseMult	$C \langle \neg M, z \rangle = C \odot (A^T \otimes B^T)$	C	$\neg, M, z, \odot, A, T, \otimes, B, T$
eWiseAdd	$C \langle \neg M, z \rangle = C \odot (A^T \oplus B^T)$	C	$\neg, M, z, \odot, A, T, \oplus, B, T$
reduce (row)	$c \langle \neg m, z \rangle = c \odot [\oplus_j A^T(:,j)]$	c	$\neg, m, z, \odot, A, T, \oplus$
apply	$C \langle \neg M, z \rangle = C \odot f(A^T)$	C	$\neg, M, z, \odot, A, T, f$
transpose	$C \langle \neg M, z \rangle = C \odot A^T$	C	$\neg, M, z, \odot, A (T)$
extract	$C \langle \neg M, z \rangle = C \odot A^T(i,j)$	C	$\neg, M, z, \odot, A, T, i, j$
assign	$C \langle \neg M, z \rangle (i,j) = C(i,j) \odot A^T$	C	$\neg, M, z, \odot, A, T, i, j$
build (meth.)	$C = \mathbb{S}^{m \times n}(i,j,v,\odot)$	C	\odot, m, n, i, j, v
extractTuples (meth.)	$(i,j,v) = A$	i,j,v	A

Notation: **i,j** – index arrays, **v** – scalar array, **m** – 1D mask, **other bold-lower** – vector (column), **M** – 2D mask, **other bold-caps** – matrix, **T** – transpose, \neg - structural complement, **z** – clear output, \oplus monoid/binary function, $\oplus \otimes$ semiring, **blue** – optional parameters, **red** – optional modifiers

SPIRAL's Math Framework

High Level Operators

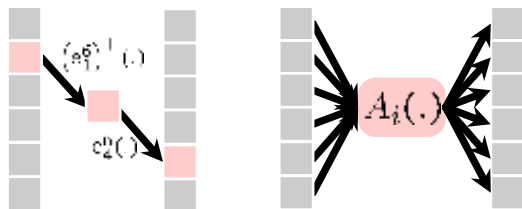
$$\langle \dots \rangle_n: \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$$



Loop Abstraction

$$\prod_{i=0}^{n-1} A_i: (D \rightarrow R)^n \rightarrow (D \rightarrow R)$$

$$A_i \mapsto (x \mapsto A_0(x) \sqcup \dots \sqcup A_{n-1}(x))$$



Basic Operators

$$\text{Pointwise}_{n, f_i}: \mathbb{R}^n \rightarrow \mathbb{R}^n$$

$$(x_i)_i \mapsto f_0(x_0) \oplus \dots \oplus f_{n-1}(x_{n-1})$$

$$\text{Atomic}_{f(\cdot, \cdot)}: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$$

$$(x, y) \mapsto f(x, y)$$

$$\text{Pointwise}_{n \times n, f_i}: \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^n$$

$$((x_i)_i, (y_i)_i) \mapsto f_0(x_0, y_0) \odot \dots \odot f_{n-1}(x_{n-1}, y_{n-1})$$

$$\text{Reduction}_{n, f_i}: \mathbb{R}^n \rightarrow \mathbb{R}$$

$$(x_i)_i \mapsto f_{n-1}(x_{n-1}, f_{n-2}(x_{n-2}, f_{n-3}(\dots f_0(x_0, \text{Id}())) \dots))$$

Rule Based Compiler

$$\text{Code}(y = (A \circ B)(x)) \rightarrow \{\text{decl}(t), \text{Code}(t = B(x)), \text{Code}(y = A(t))\}$$

$$\text{Code}\left(y = \left(\sum_{i=0}^{n-1} A_i\right)(x)\right) \rightarrow \{y := \vec{0}, \text{for}(i = 0..n-1) \text{Code}(y += A_i(x))\}$$

$$\text{Code}(y = (e_i^n)^T(x)) \rightarrow y[0] := x[i]$$

$$\text{Code}(y = e_i^n(x)) \rightarrow \{y = \vec{0}, y[i] := x[0]\}$$

$$\text{Code}(y = \text{Atomic}_f(x)) \rightarrow y[0] := f(x[i])$$

Leverages DARPA HACMS

Autotuning in Constraint Solution Space

Intel Core i7 (2nd Gen)

TriangleCount

Base cases

$$L \otimes_{16} A_{16 \times 16}$$

$$L_{16 \times 16} \otimes_{16} I_{16}$$

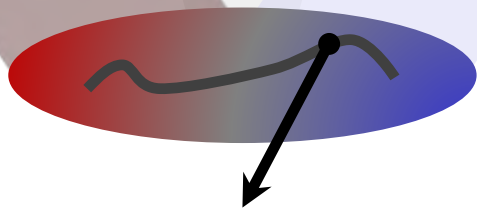
$$A_{16 \times 16}$$

$$\underbrace{L_{16 \times 16}}_{SSE} \cdot \underbrace{L_{16 \times 16}}_{SSE} \cdot \underbrace{L_{16 \times 16}}_{SSE} \cdot \underbrace{L_{16 \times 16}}_{SSE}$$

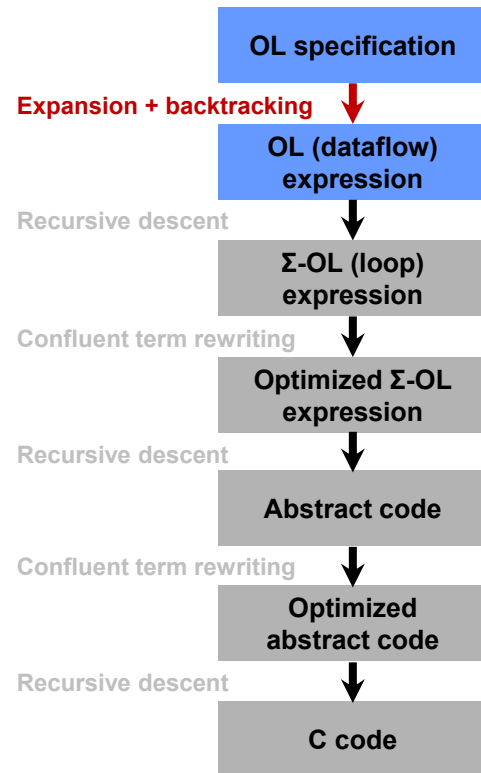
Transformation rules

Breakdown rules

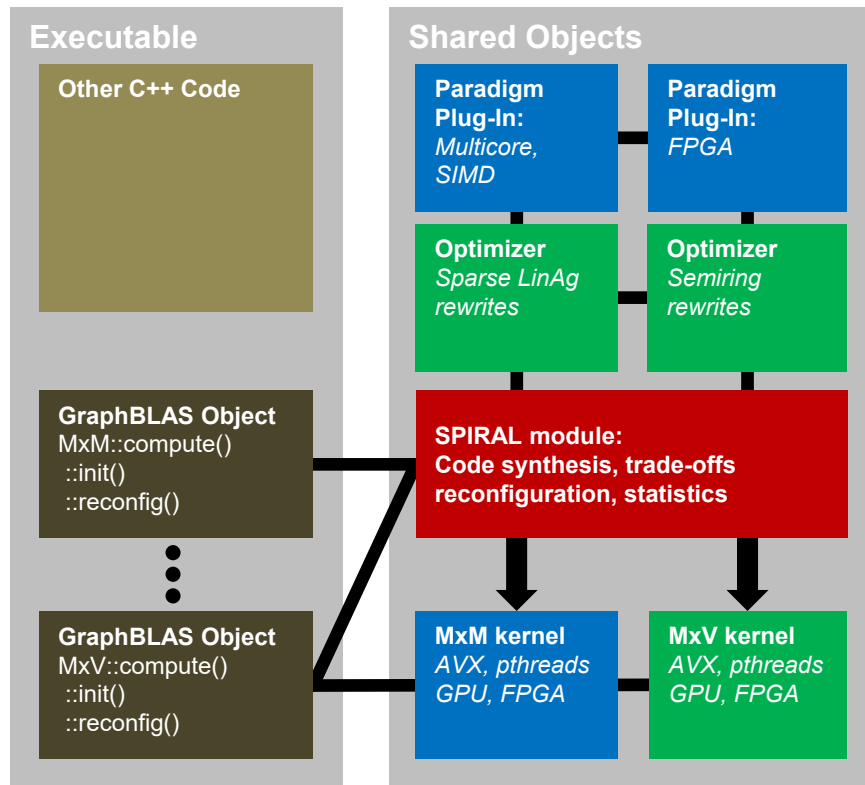
triangles = $||L \otimes (L \oplus \cdot \otimes L^T)||_1$



$$\left((L_{16 \times 16}^{\otimes 4} \otimes L_{16 \times 16}) \otimes I_{16} \right) \left(I_{16} \otimes (DFT_{16 \times 16} \otimes L_{16 \times 16}) \right) \left((L_{16 \times 16}^{\otimes 2} \otimes L_{16 \times 16}) \otimes I_{16} \right) \left(\bigoplus_{i=0}^{15} T_{16 \times 16}^{m_{i,5}} \right) \left(I_{16} \otimes (L_{16 \times 16} \otimes DFT_{16 \times 16}) \right) \left(I_{16} \otimes L_{16 \times 16}^{\otimes 2} \right) \left((L_{16 \times 16}^{\otimes 2} \otimes L_{16 \times 16}) \otimes I_{16} \right)$$



Long-Range Goal: SPIRAL as JIT and GraphBLAS Optimizer



Source Code

- C++, GraphBLAS calls, other supported libraries
- Code = **specification**, not program

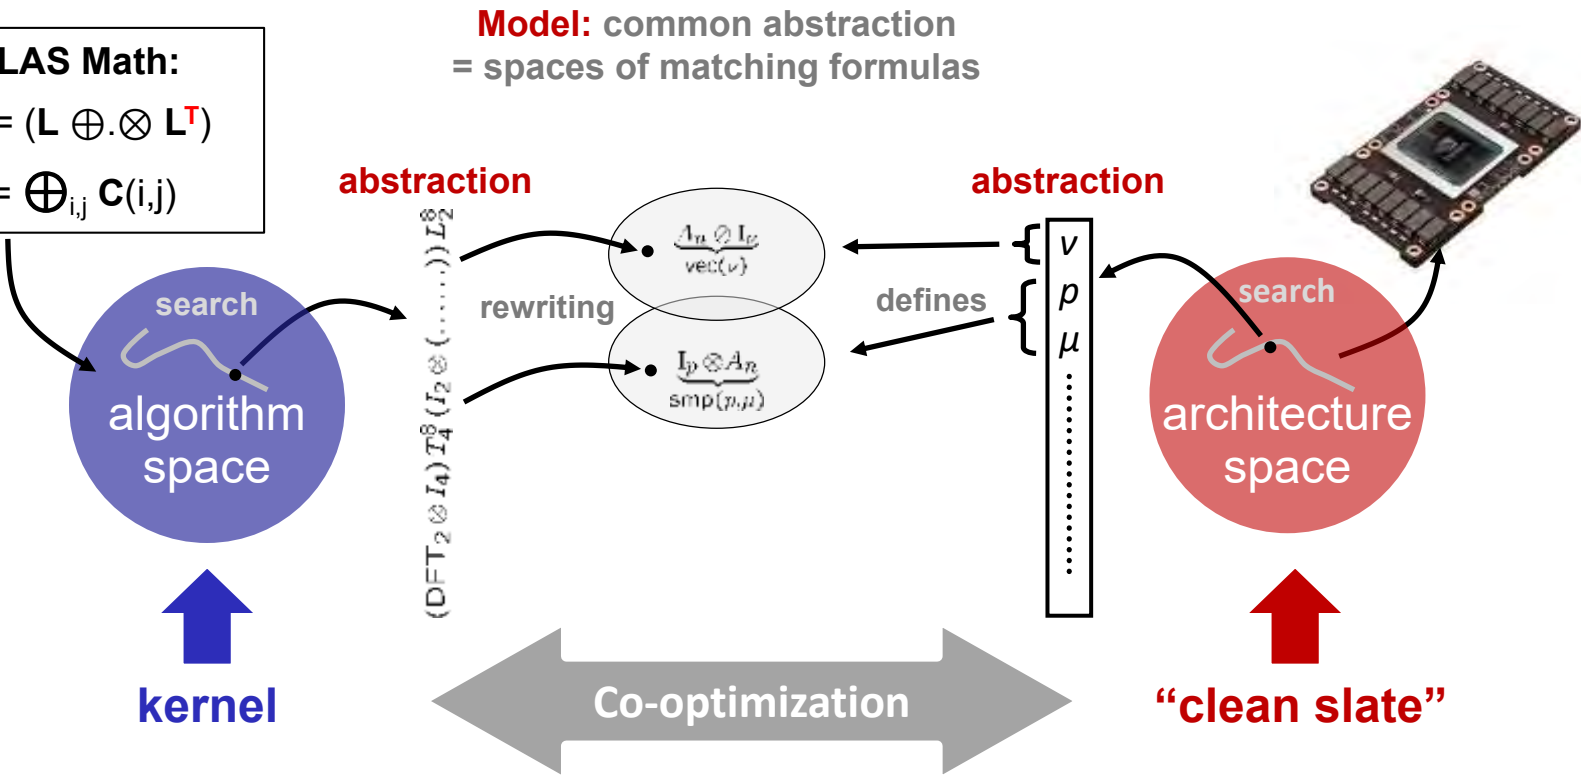
SPIRAL Module

- Acts as JIT, delayed execution engine, Inspector/executor
- Implements telescoping language ideas
- Rewrites code into better algorithms
- Compiles to range of platforms CPU, GPU, FPGA
- Plug-in mechanism for post deployment reconfiguration and update

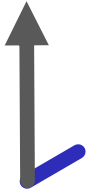
Leverages DARPA BRASS

Formal Approach To Co-Optimization

GraphBLAS Math:
 $C(L, z) = (L \oplus \cdot \otimes L^T)$
 count = $\bigoplus_{i,j} C(i,j)$



$$(\hat{\mathcal{A}}, \hat{\mathcal{M}}) = \operatorname{argmin}_{\theta, \xi} C_m(\mathcal{A}(\theta), \mathcal{M}(\xi))$$



M (»)

“What is the right architecture for my application?”
“What architecture features are good for my application?”

Summary and Future Work

- GraphBLAS C API Specification is complete
 - Mathematical descriptions of primitive operations are complete
 - Algorithm development using the API is in progress (dozens completed)
- Exploration of performant code generation and data structures has begun
- Goals for FY18:
 - Integrate primitives and necessary “knowledge” into Spiral code generation technology
 - Target different hardware platforms
 - Multi-core CPUs
 - Accelerators: FPGAs, Graphics Processing Units (GPUs)
 - Generate code for algorithms in benchmark and Challenge problems
 - Graph 500: breadth-first search, shortest paths
 - DARPA HIVE Graph Challenge: subgraph isomorphism
- Long-Range Goal: Co-synthesis of hardware and software

Contact Information

Presenter / SEI PI

Dr. Scott McMillan
Senior Research Scientist

Email: smcmillan@sei.cmu.edu

Presenter / CMU PI

Prof. Franz Franchetti
ECE Department

Email: franzf@ece.cmu.edu