

GraphBLAS: A Programming Specification for Graph Analysis

Scott McMillan, PhD



Copyright 2016 Carnegie Mellon University

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Department of Defense.

References herein to any specific commercial product, process, or service by trade name, trade mark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by Carnegie Mellon University or its Software Engineering Institute.

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.

The Government of the United States has a royalty-free government-purpose license to use, duplicate, or disclose the work, in whole or in part and in any manner, and to have or permit others to do so, for government purposes pursuant to the copyright license under the clause at 252.227-7013 and 252.227-7013 Alternate I.

This material was prepared for the exclusive use of SEI Research Review and may not be used for any other purpose without the written consent of permission@sei.cmu.edu.

DM-0004100

Template updated 12.8.2016

GraphBLAS: A Programming Specification for Graph Analysis

Scott McMillan, PhD

In collaboration with:

MIT/LL, LBNL, Intel, IBM, UC Davis, Indiana, and *many* others.

Summary

Graph analysis is *important* and *pervasive* in the DoD community.

GraphBLAS Forum (a world-wide consortium of researchers)

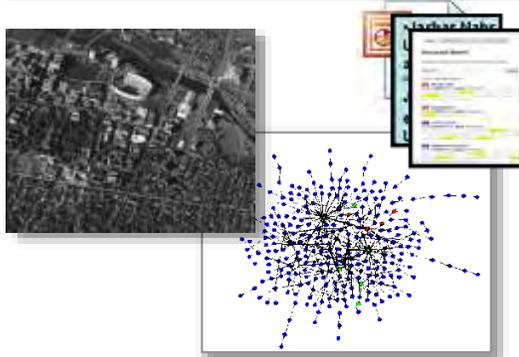
- Government/FFRDCs, academia, and industry
- Goal: application programming specification (API) for graph analysis
- GraphBLAS: Graph Basic Linear Algebra Subprograms

SEI contributions

- Member of the Graph BLAS Forum
- Member of the C API Specification committee
- Early implementation of a C++ library: GraphBLAS Template Library (open-source)

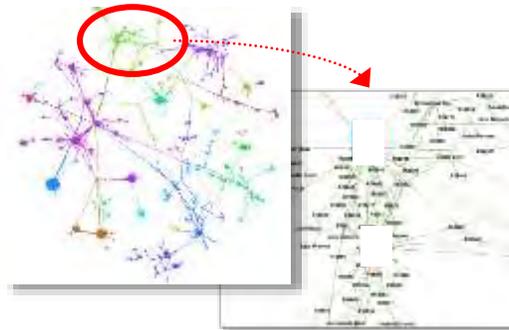
Graph Analysis is *Important* and *Pervasive*

ISR



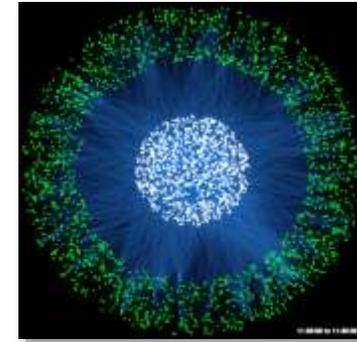
- Graphs represent entities and relationships detected through multi-INT sources
- 1,000s – 1,000,000s tracks and locations
- GOAL: Identify anomalous patterns of life

Social



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

Cyber



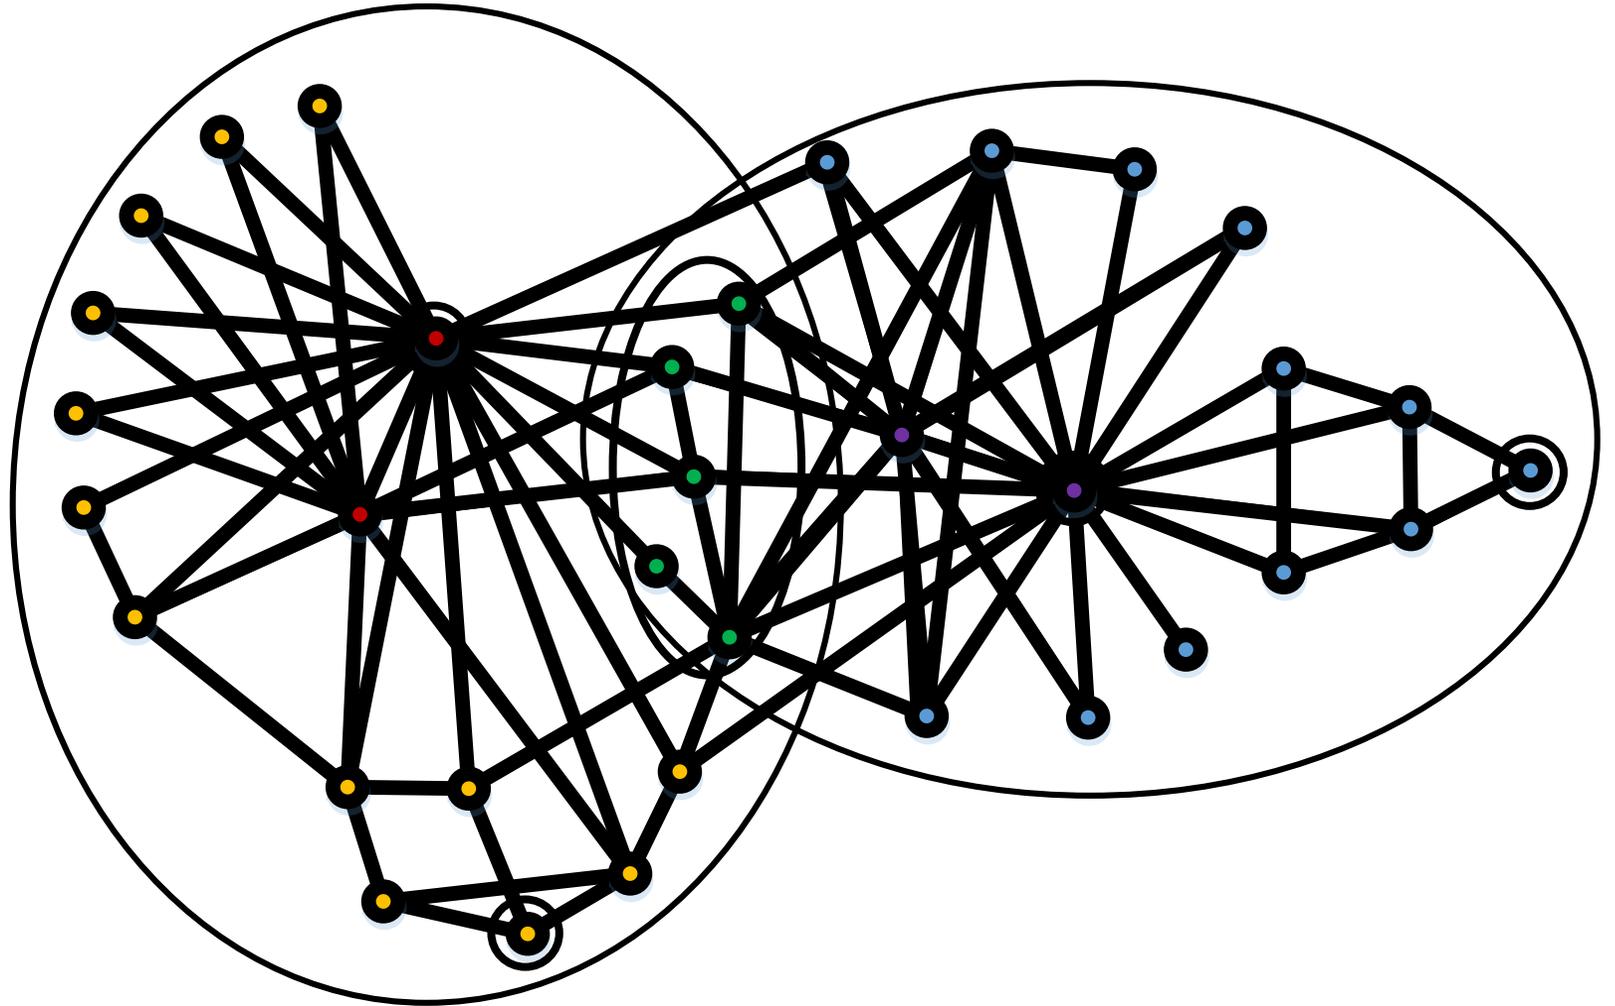
- Graphs represent communication patterns of computers on a network
- 1,000,000s – 1,000,000,000s network events
- GOAL: Identify cyber attacks or malicious software

Common Goal: Detection of subtle patterns in massive graphs

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

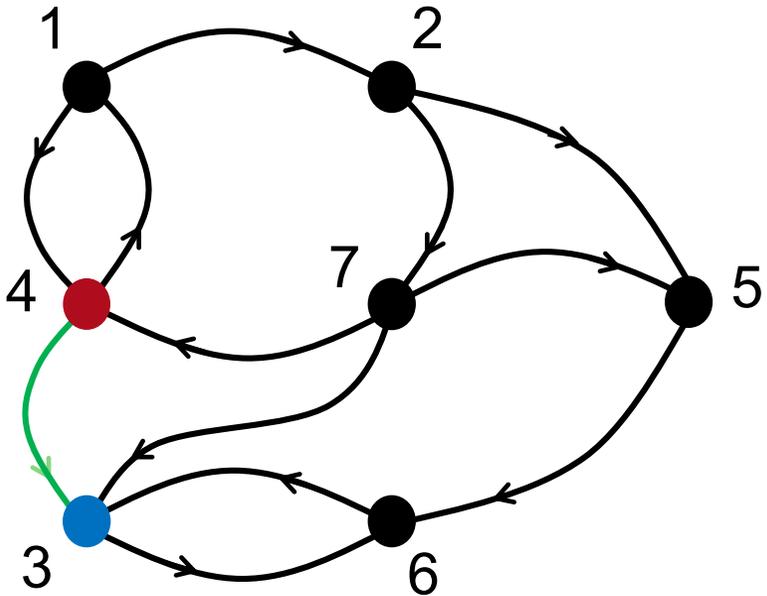
Common Operations

- Finding Neighbors
- Shortest Paths
- Clustering
- “Important” nodes
 - Influencers
 - Bottlenecks
 - Outliers



W. W. Zachary, An information flow model for conflict and fission in small groups,
Journal of Anthropological Research 33, 452-473 (1977).

Graphs as Matrices



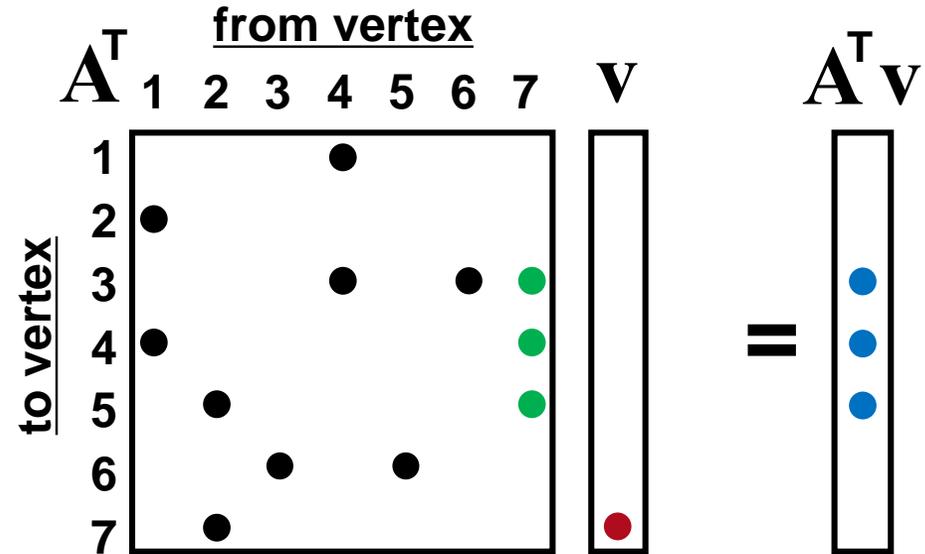
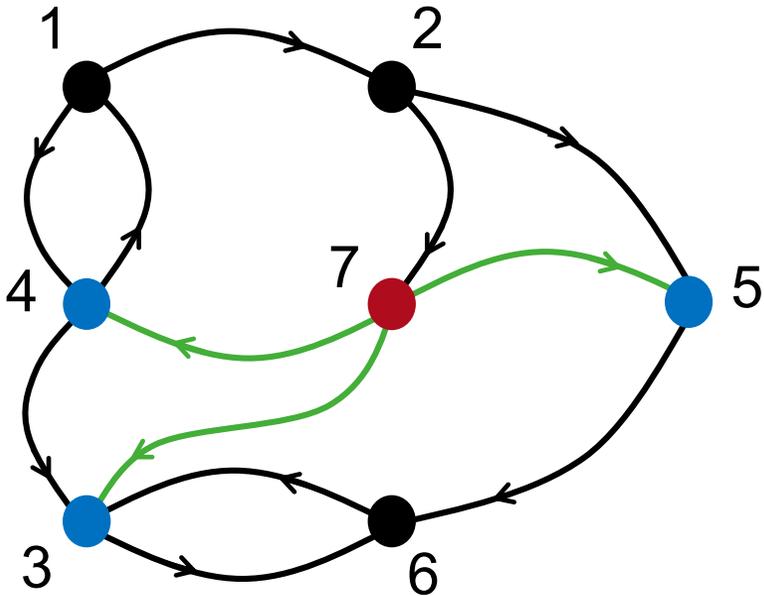
A

	'to' vertex						
	1	2	3	4	5	6	7
1		●		●			
2					●		●
3						●	
4	●		●				
5						●	
6			●				
7		●	●	●			

'from' vertex

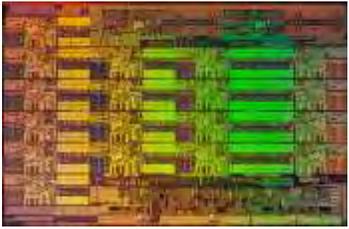
Graphs are represented as adjacency matrices that have *irregular* and *sparse* structure.

Graph Operations as Matrix Operations

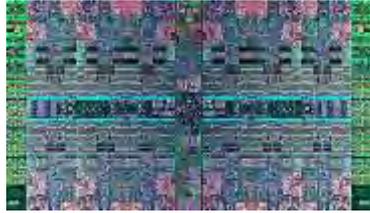


- Matrix multiply \rightarrow find neighbors (most important primitive)
- Used in breadth-first traversal, shortest paths, and many others
- Sparsity and irregularity of matrix structure is a barrier to high performance

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

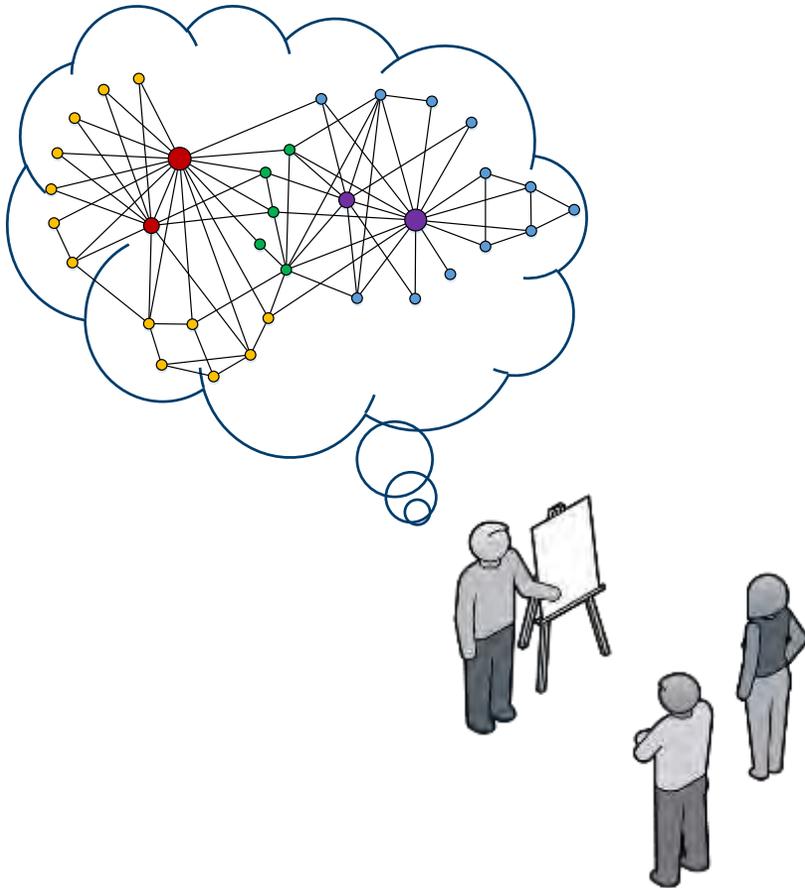


IBM BlueGene/Q
10 Pflop/s, 8 MW
 48k x 16 cores, 1.6 GHz
 4-way QPX

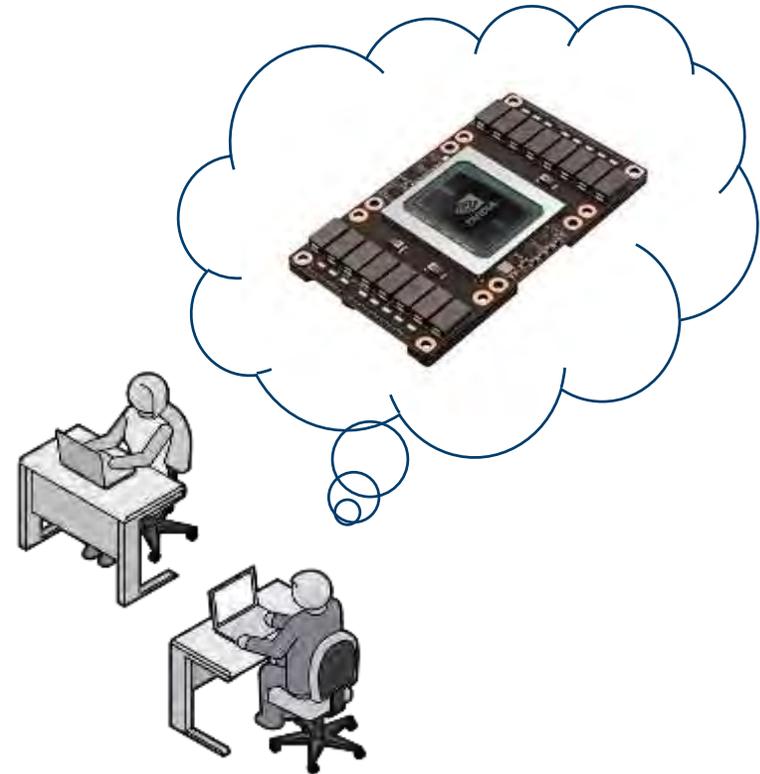
Slide credit: Franz Franchetti, "SPIRAL: Automated Code Generation of Performance Libraries," 2016.

Separation of Concerns

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



Separation of Concerns



Graph Primitives Using Linear Algebra

Standards for Graph Algorithm Primitives

Tim Mattson (Intel Corporation), David Bader (Georgia Institute of Technology), Jon Berry (Sandia National Laboratory), Aydin Buluc (Lawrence Berkeley National Laboratory), Jack Dongarra (University of Tennessee), Christos Faloutsos (Carnegie Mellon University), John Feo (Pacific Northwest National Laboratory), John Gilbert (University of California at Santa Barbara), Joseph Gonzalez (University of California at Berkeley), Bruce Hendrickson (Sandia National Laboratory), Jeremy Kepner (Massachusetts Institute of Technology), Charles Leiserson (Massachusetts Institute of Technology), Andrew Lumsdaine (Indiana University), David Padua (University of Illinois at Urbana-Champaign), Stephen Poole (Oak Ridge National Laboratory), Steve Reinhardt (Cray Corporation), Mike Stonebraker (Massachusetts Institute of Technology), Steve Wallach (Convey Corporation), Andrew Yoo (Lawrence Livermore National Laboratory)

“It is our view that the state of the art in constructing a large collection of graph algorithms in terms of linear algebraic operations is mature enough to support the emergence of a **standard set of primitive building blocks**. This paper is a position paper defining the problem and **announcing our intention to launch an open effort to define this standard.**”

Presented at the IEEE High Performance Extreme Computing Conference. Waltham, MA, Sept. 2013.

GraphBLAS Forum



FFRDCs

- MIT/Lincoln Labs*
- Lawrence Berkeley NL* †
- CMU/Software Engineering Institute †
- Pacific Northwest NL
- Sandia NL

Industry

- Intel* †
- IBM †
- NVIDIA
- Hauwei
- Reservoir Labs
- Galois
- Mellanox
- and others

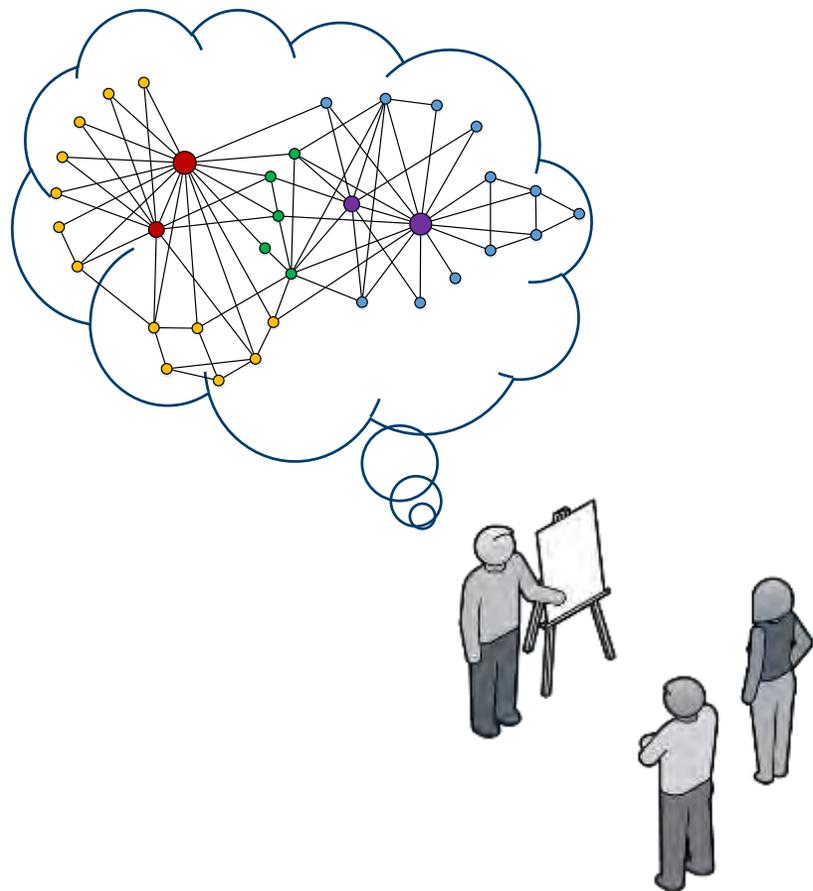
Academia

- UC Santa Barbara*, Davis †, Berkeley
- Georgia Tech*
- Karlsruhe (KIT)*
- CMU
- Indiana U.
- MIT
- U. Washington
- and others

* - steering committee, † - C API subcommittee

Separation of Concerns

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



GraphBLAS Application Programming Interface (API)



GraphBLAS Primitives

Operation	Description
mxm, mxv, vxm	Perform matrix <i>multiplication</i> (e.g., breadth-first traversal)
eWiseAdd, eWiseMult	Element-wise <i>addition</i> and <i>multiplication</i> 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 <i>unary function</i> to each element of matrix (e.g., edge weight modification)
reduce	<i>Reduce</i> along columns or rows of matrices (vertex degree)
transpose	Swaps the rows and columns of a sparse matrix (e.g., reverse directed edges)
buildMatrix	Build an matrix representation from row, column, value tuples
extractTuples	Extract the row, column, value tuples from a matrix representation

GraphBLAS Primitives: The Math



Operation	Mathematical Description	Outputs	Inputs
mxm	$C(\neg M) \oplus = A^T \oplus \cdot \otimes B^T$	C	$\neg, \mathbf{M}, \oplus, \mathbf{A}, \mathbf{T}, \oplus \cdot \otimes, \mathbf{B}, \mathbf{T}$
mxv (vxm)	$c(\neg m) \oplus = A^T \oplus \cdot \otimes b$	c	$\neg, \mathbf{m}, \oplus, \mathbf{A}, \mathbf{T}, \oplus \cdot \otimes, \mathbf{b}$
eWiseMult	$C(\neg M) \oplus = A^T \otimes B^T$	C	$\neg, \mathbf{M}, \oplus, \mathbf{A}, \mathbf{T}, \otimes, \mathbf{B}, \mathbf{T}$
eWiseAdd	$C(\neg M) \oplus = A^T \oplus B^T$	C	$\neg, \mathbf{M}, \oplus, \mathbf{A}, \mathbf{T}, \oplus, \mathbf{B}, \mathbf{T}$
reduce (row)	$c(\neg m) \oplus = \oplus_j A^T(:,j)$	c	$\neg, \mathbf{m}, \oplus, \mathbf{A}, \mathbf{T}, \oplus$
apply	$C(\neg M) \oplus = f(A^T)$	C	$\neg, \mathbf{M}, \oplus, \mathbf{A}, \mathbf{T}, f$
transpose	$C(\neg M) \oplus = A^T$	C	$\neg, \mathbf{M}, \oplus, \mathbf{A} (\mathbf{T})$
extract	$C(\neg M) \oplus = A^T(i,j)$	C	$\neg, \mathbf{M}, \oplus, \mathbf{A}, \mathbf{T}, i, j$
assign	$C(\neg M) (i,j) \oplus = A^T$	C	$\neg, \mathbf{M}, \oplus, \mathbf{A}, \mathbf{T}, i, j$
buildMatrix	$C(\neg M) \oplus = \mathbb{S}^{m \times n}(i,j,v,\oplus)$	C	$\neg, \mathbf{M}, \oplus, \oplus, m, n, i, j, v$
extractTuples	$(i,j,v) = A(\neg M)$	i,j,v	$\neg, \mathbf{M}, \mathbf{A}$

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, \oplus monoid/binary function, $\oplus \cdot \otimes$ semiring, **blue** – optional parameters, **red** – optional modifiers

GraphBLAS Primitives: The Code

$$C(\neg M) \oplus = A^T \oplus \otimes B^T$$

```

GrB_Info GrB_mxm(GrB_Matrix *C, // output
                 const GrB_Matrix Mask,
                 const GrB_BinaryFunction accum,
                 const GrB_Semiring op,
                 const GrB_Matrix A, // input matrix
                 const GrB_Matrix B // input matrix
                 [, const Descriptor desc]);

```

Common Elements

- Matrices (C, Mask, A, and B) are opaque data structures defined by the library implementers.
- Destination (C) is first.
- Output mask (M) specifies which output elements can assigned (optional).
- Accumulation function (accum) allows operation to combine with existing values (optional).
- Descriptor can specify transpose of input matrices or logical complement of the output mask.



The GraphBLAS Forum

- Meeting since 2014
- Teleconferences once per month (C API group more often)
- Forum meetings open to the public:
 - May: IEEE IPDPS (Symposium) – Graph Algorithms Building Blocks Workshop
 - Sept.: IEEE HPEC (Conference) – GraphBLAS BoF
 - Nov.: IEEE/ACM Supercomputing (Conference) – GraphBLAS Working Group
- Status and next steps:
 - Primitives established
 - C Application Programming Interface (API) – First draft review in August
 - Targeted release: November 2016 at Supercomputing
- For more information: <http://graphblas.org>

Current and Future Work

- Continue work on C++ API from the open-source GraphBLAS Template Library (GBTL) [collaboration with Indiana/PNNL]
- Line Project (FY17-18) on automated code generation for high performance graph libraries [collaboration with CMU ECE department]
 - Starting with the GraphBLAS primitives
 - Targeting COTS hardware (CPUs and GPUs)
- DARPA's HIVE (Hierarchical Identify Verify Exploit) BAA released in August
 - References the work of the GraphBLAS Forum
 - Goal: develop a special-purpose graph processing chip

Contact Information

Presenter / Point of Contact

Scott McMillan, PhD

Senior Member of Technical Staff

Telephone: +1 412.268.5156

Email: smcmillan@sei.cmu.edu