

A GraphBLAS solution to the SIGMOD 2014 Programming Contest using multi-source BFS


High-Performance Extreme Computing (HPEC) 2020

M. Elekes, A. Nagy, D. Sándor, J.B. Antal, T.A. Davis, G. Szárnyas



MOTIVATION

- Graph algorithms are challenging to program
 - irregular access patterns → poor locality
 - caching and parallelization are difficult
- Optimizations often limit portability
- **GraphBLAS** introduces an abstraction layer using the language of linear algebra
 - graph \equiv sparse matrix
 - navigation step \equiv matvec on semirings




\equiv

A	①	②	③	④	⑤	⑥	⑦
①		●		●			
②					●		●
③						●	
④	●		●				
⑤						●	
⑥			●				
⑦			●	●	●		

MOTIVATION

- Graph algorithms are challenging to program
 - irregular access patterns → poor locality
 - caching and parallelization are difficult
- Optimizations often limit portability
- **GraphBLAS** introduces an abstraction layer using the language of linear algebra
 - graph \equiv sparse matrix
 - navigation step \equiv matvec on semirings



\equiv

A

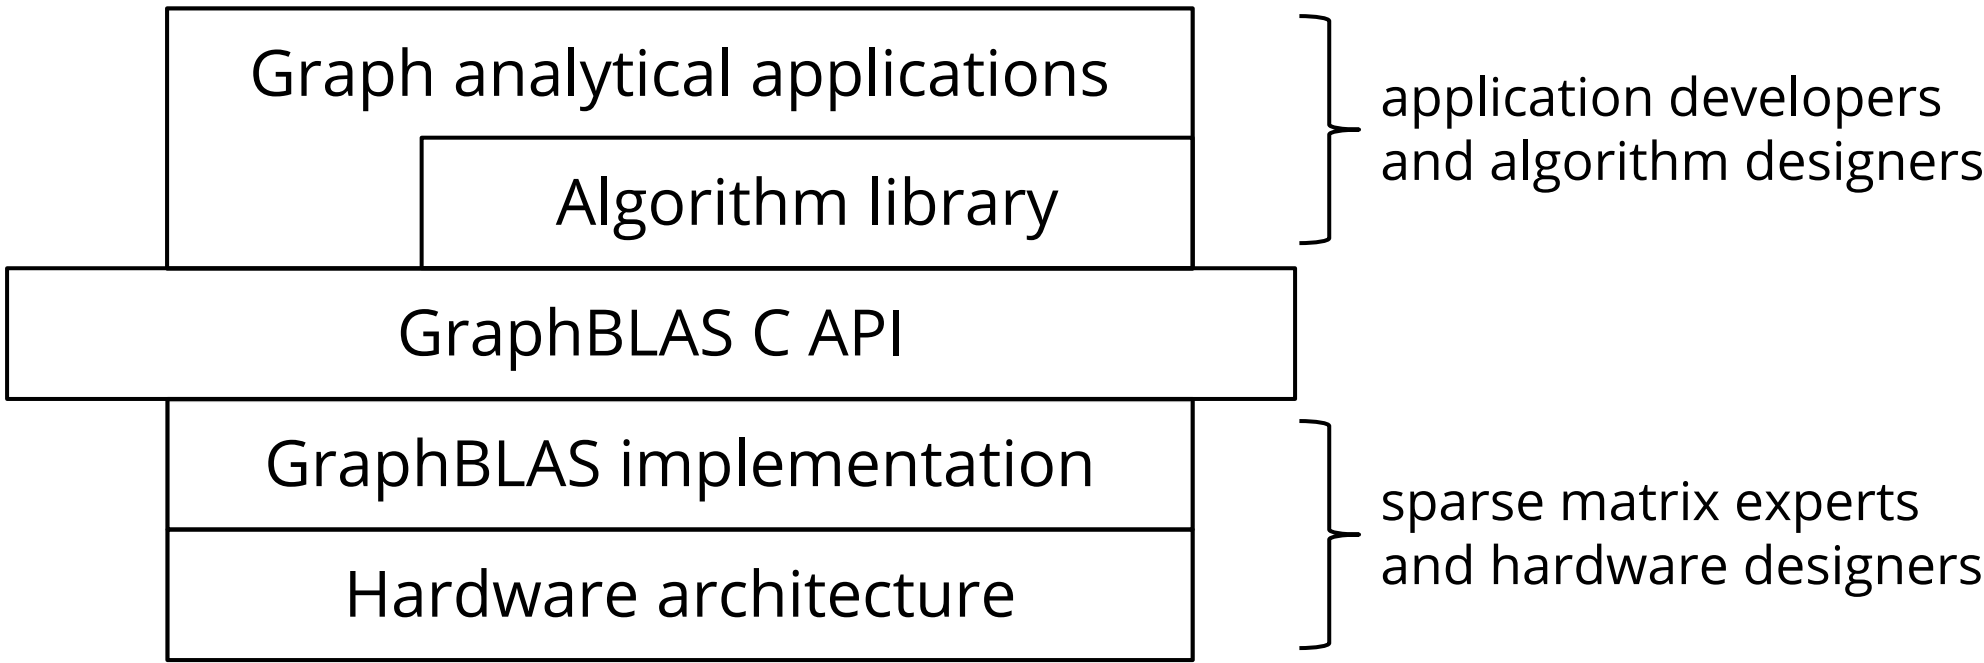
	①	②	③	④	⑤	⑥	⑦
①		●		●			
②					●		●
③						●	
④	●		●				
⑤						●	
⑥			●				
⑦			●	●	●		

v

	①	②	③	④	⑤	⑥	⑦
							●

	①	②	③	④	⑤	⑥	⑦
			●	●	●		

GRAPHBLAS STACK



Textbook algos: BFS, PageRank, triangle count untyped graphs

GraphBLAS-based Cypher engine: RedisGraph property graphs

RQ: How to formulate mixed workloads on property graphs?

SIGMOD 2014 PROGRAMMING CONTEST

Annual contest

- Teams compete on database-related programming tasks
- Highly-optimized C++ implementations


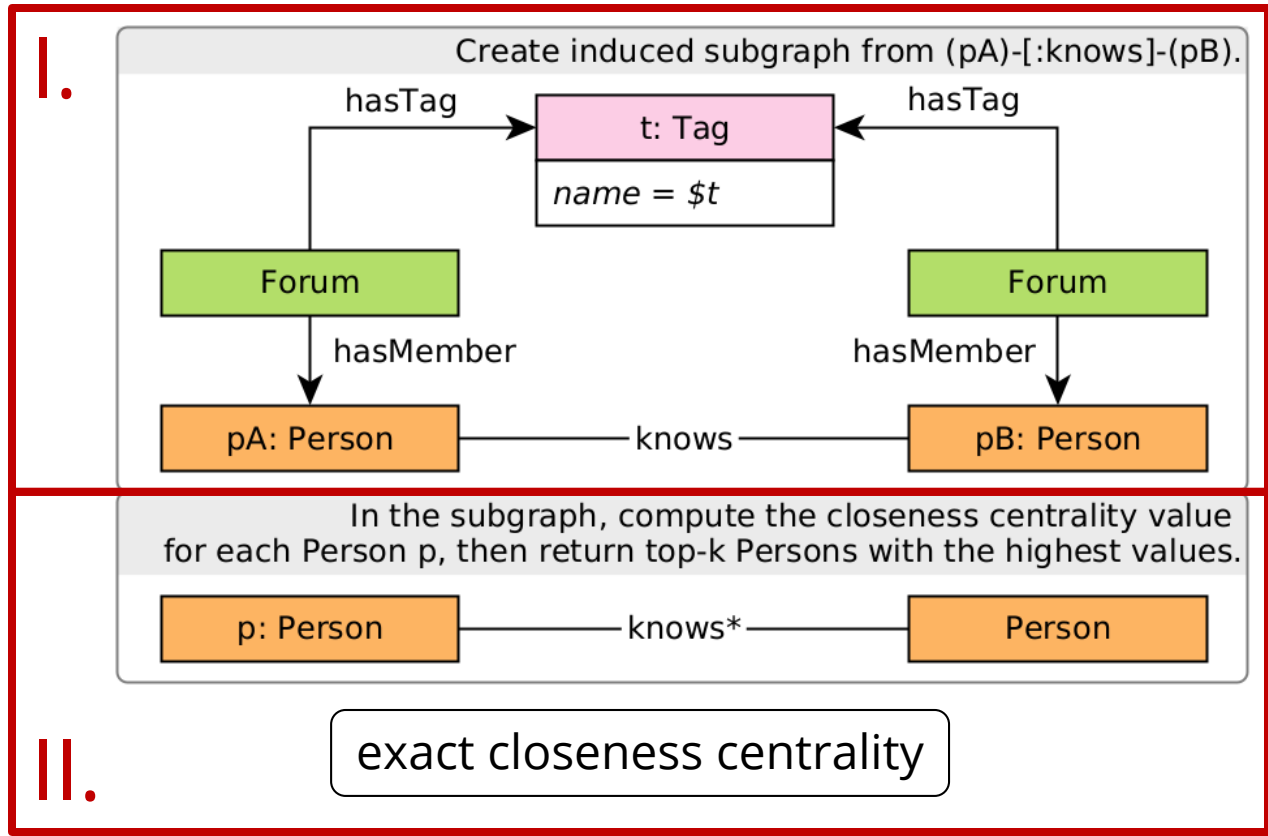
2014 event

- Tasks on the LDBC social network graph
 - Benchmark data set for property graphs
 - People, forums, comments, hashtags, etc.
- 4 queries
 - Mix of filtering operations and graph algorithms



QUERY TEMPLATE


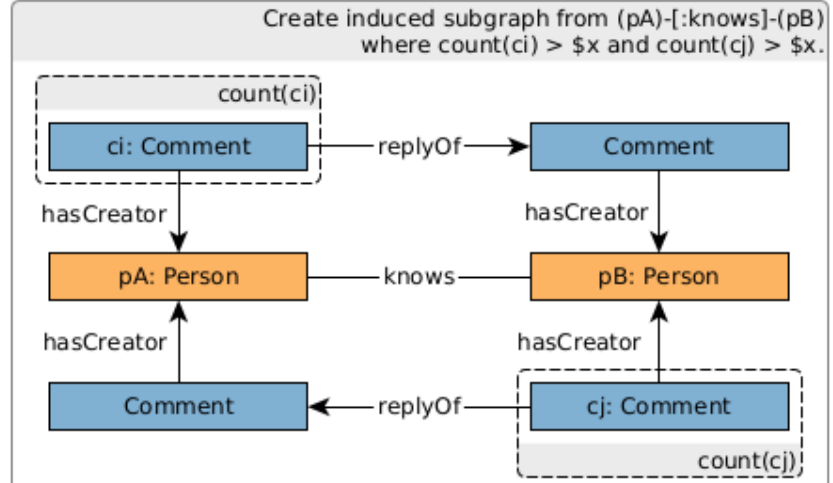
- I. Compute an induced subgraph over Person-knows-Person
- II. Run the graph algorithm on the subgraph




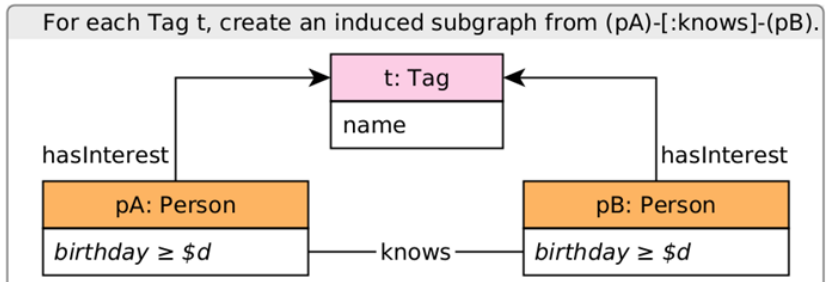
key kernel: all-source BFS

OVERVIEW OF QUERIES 1, 2, 3

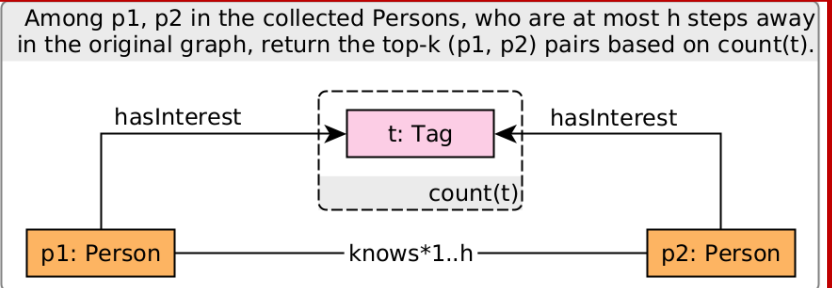
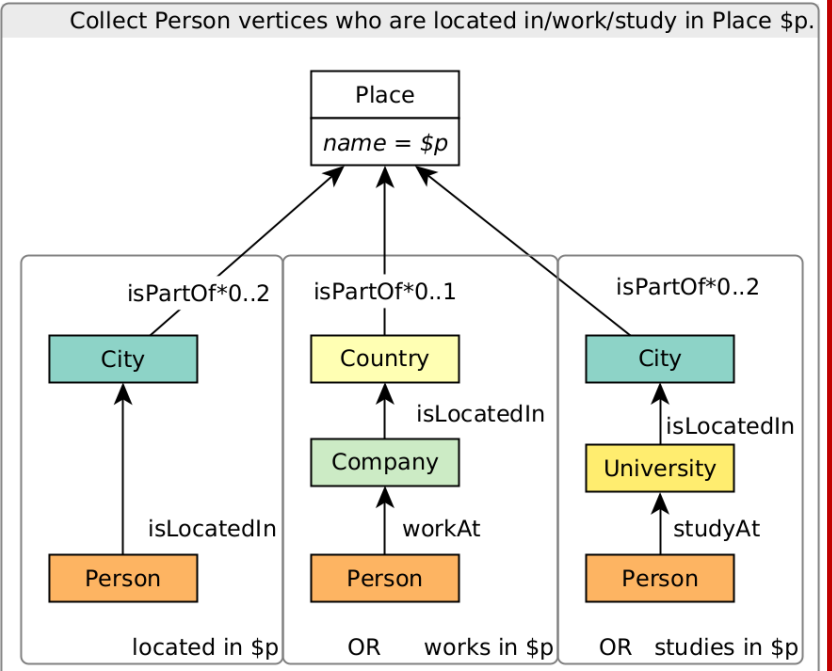
I. Filter the induced subgraph(s)



unweighted shortest path



connected components



pairwise reachability


II. Run the graph algorithm

GRAPHBLAS SOLUTION OF THE QUERIES

- **Loading** includes relabelling UIN64 vertex IDs to a contiguous sequence $0 \dots N - 1$.
- **Filtering the induced subgraph** from the property graph is mostly straightforward and composable with the algorithms.
- **The algorithms** can be concisely expressed in GraphBLAS:
 - Connected components ✓ → FastSV [[Zhang et al., PPSC'20](#)]
 - BFS ✓
 - Bidirectional BFS
 - All-source BFS + bitwise optimization
 - Multi-source bidirectional BFS


BFS

BFS: BREADTH-FIRST SEARCH




Boolean matrices and vectors



frontier



seen




A ① ② ③ ④ ⑤



$\langle \neg \text{seen} \rangle$ mask


seen'






⊕: logical OR
⊗: logical AND

$\text{next} \langle \neg \text{seen} \rangle = A \text{ LOR } \text{LAND frontier}$


$\text{seen}' = \text{seen LOR next}$





BFS: BREADTH-FIRST SEARCH



frontier



mask prevents redundant computations



$\text{next}(\neg \text{seen}) =$
A LOR . LAND frontier

$\text{seen}' =$
seen LOR next

All-source BFS

Q4: CLOSENESS CENTRALITY VALUES

Q4 computes the top- k Person vertices based on their exact *closeness centrality values*:


$$CCV(p) = \frac{(C(p) - 1)^2}{(n - 1) \cdot s(p)}$$

where

- $C(p)$ is the size of the connected component of vertex p ,
- n is the number of vertices in the induced graph,
- $s(p)$ is the sum of geodesic distances to all other reachable persons from p .

$s(p)$ is challenging: needs unweighted all-pairs shortest paths.

BOOLEAN ALL-SOURCE BFS ALGORITHM



Frontier t1 t2 t3 t4 t5

traversals

①	●				
②		●			
③			●		
④				●	
⑤					●

Seen t1 t2 t3 t4 t5

①	●				
②		●			
③			●		
④				●	
⑤					●

A ① ② ③ ④ ⑤

①		●		●	
②	●		●		●
③		●		●	●
④	●		●		
⑤		●	●		

	X	●		●	
●	X		●		●
	●	X		●	●
●		●	X		
	●	●		X	

Seen' t1 t2 t3 t4 t5


①	●	●		●	
②	●	●	●		●
③		●	●	●	●
④	●		●	●	
⑤		●	●		●



Next(\neg Seen) =
A LOR . LAND Frontier

Seen' =
Seen LOR Next

BOOLEAN ALL-SOURCE BFS ALGORITHM



Frontier t1 t2 t3 t4 t5

①		●		●	
②	●		●		●
③		●		●	●
④	●		●		
⑤		●	●		

Seen t1 t2 t3 t4 t5

①	●	●		●	
②	●	●	●		●
③		●	●	●	●
④	●		●	●	
⑤		●	●		●


A ① ② ③ ④ ⑤

①		●		●	
②	●		●		●
③		●		●	●
④	●		●		
⑤		●	●		

①	⊗	⊗	⊗	⊗	⊗
②	⊗	⊗	⊗	⊗	⊗
③	⊗	⊗	⊗	⊗	⊗
④	⊗	⊗	⊗	⊗	⊗
⑤	⊗	⊗	⊗	⊗	⊗

Seen' t1 t2 t3 t4 t5

①	●	●	●	●	●
②	●	●	●	●	●
③	●	●	●	●	●
④	●	●	●	●	●
⑤	●	●	●	●	●



Next(\neg Seen) =
A LOR . LAND Frontier


Seen' =
Seen LOR Next

Bitwise all-source BFS

BITWISE ALL-SOURCE BFS ALGORITHM

- For large graphs, the all-source BFS algorithm might need to run 500k+ traversals
- Two top-ranking teams used bitwise operations to process traversals in batches of 64 [[Then et al., VLDB'15](#)]
- This idea can be adopted in the GraphBLAS algorithm by
 - using UINT64 values
 - performing the multiplication on the BOR .SECOND semiring, where BOR is “bitwise or” and $\text{SECOND}(x, y) = y$
- 5-10x speedup compared to the Boolean all-source BFS

BITWISE ALL-SOURCE BFS ALGORITHM



Using UINT4s here

Frontier

	t1-t4	t5
①	1000	0000
②	0100	0000
③	0010	0000
④	0001	0000
⑤	0000	1000

Seen

	t1-t4	t5
①	1000	0000
②	0100	0000
③	0010	0000
④	0001	0000
⑤	0000	1000

A

	①	②	③	④	⑤
①		●		●	
②	●		●		●
③		●		●	●
④	●		●		
⑤		●	●		

Next =

	t1-t4	t5
①	0101	0000
②	1010	1000
③	0101	1000
④	1010	0000
⑤	0110	0000




Seen'

	t1-t4	t5
①	0101	0000
②	1010	1000
③	0101	1000
④	1010	0000
⑤	0110	0000

Next =
A BOR . SECOND Frontier

Seen' =
Seen BOR Next

BITWISE ALL-SOURCE BFS ALGORITHM



Frontier

	t1-t4	t5
①	0101	0000
②	1010	1000
③	0101	1000
④	1010	0000
⑤	0110	0000

Full VLDB paper
on this algorithm

vs.

9 GrB operations

Seen

	t1-t4	t5
①	0101	0000
②	1010	1000
③	0101	1000
④	1010	0000
⑤	0110	0000

A

	①	②	③	④	⑤
①		●		●	
②	●		●		●
③		●		●	●
④	●		●		
⑤		●	●		

	t1-t4	t5
①	0010	1000
②	0001	0000
③	1000	0000
④	0100	1000
⑤	1001	0000



Seen'

	t1-t4	t5
①	1111	1000
②	1111	1000
③	1111	1000
④	1111	1000
⑤	1111	1000

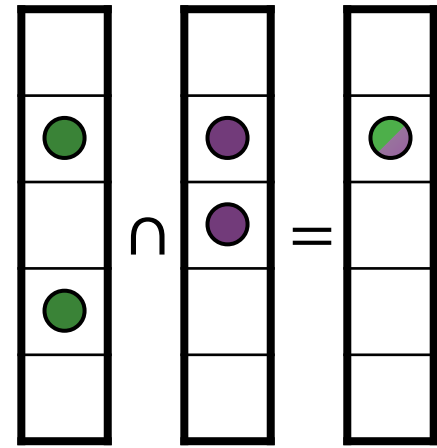
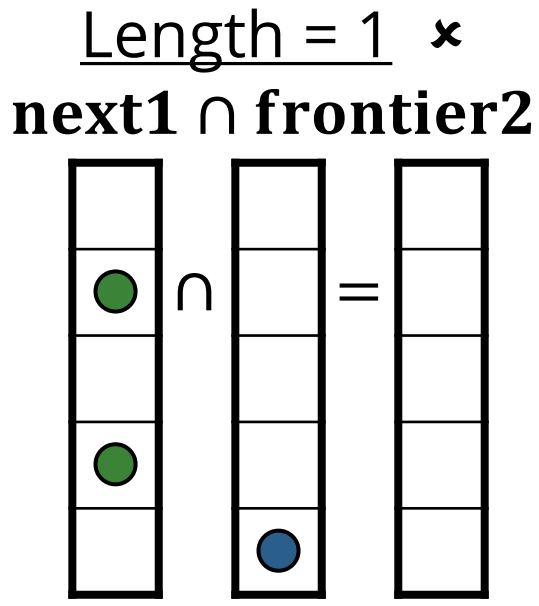




Next =
A BOR . SECOND Frontier

Seen' =
Seen BOR Next

Bidirectional BFS

BIDIRECTIONAL BFS

Advance frontiers alternately and intersect them




Length = 2 ✓
next1 ∩ next2

Bidirectional MSBFS




BIDIRECTIONAL MSBFS ALGORITHM

- **Pairwise reachability problem:**
From a given set of k vertices, which pairs of vertices are reachable from each other with at most h hops?
- **Naïve solution:**
Run a k -source MSBFS for h steps and check reachability. The frontiers get large as they grow exponentially.
- **Better solution:**
Advance all frontiers simultaneously for $\lceil h/2 \rceil$ iterations.



BIDIRECTIONAL MSBFS

Seen[1]: reachability with ≤ 1 hops




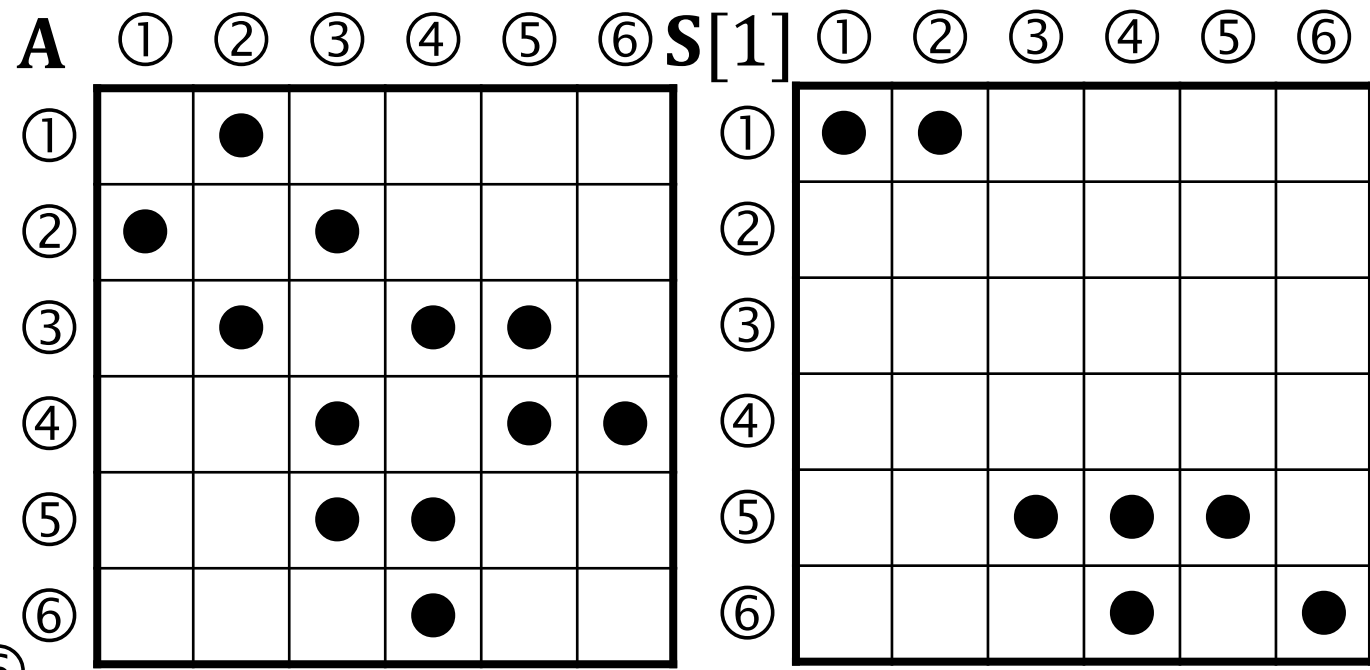



Next[1]

Seen[1]

BIDIRECTIONAL MSBFS


Seen[2]: reachability with ≤ 2 hops



Next[2] < Seen[1] >

Seen[2]

BIDIRECTIONAL MSBFS: PATHS OF LENGTH ≤ 4



Seen[2]^T

	①	②	③	④	⑤	⑥
①	●					
②	●					
③	●				●	●
④					●	●
⑤					●	
⑥						●

From vertex 1, we could get to these vertices with ≤ 2 hops

Seen[2]

	①	②	③	④	⑤	⑥
①	●	●	●			
②						
③						
④						
⑤			●	●	●	
⑥			●	●		●

From vertex 5, we could get to these vertices with ≤ 2 hops

To get paths of at most 4 hops, we compute


Seen[2] LOR.LAND Seen[2]^T

⊗				●	●
⊗	⊗				
⊗	⊗	⊗			
⊗	⊗	⊗	⊗		
⊗	⊗	⊗		⊗	
⊗	⊗	⊗		●	⊗

Here, we found paths between all pairs:

- from ① to ⑤,
- from ① to ⑥,
- from ⑤ to ⑥.

BIDIRECTIONAL MSBFS: PATHS OF LENGTH 3



Next[2]^T

	①	②	③	④	⑤	⑥
①						
②						
③	●					●
④						
⑤						
⑥						

From vertex 1, we could get to this vertex with 2 hops

To get exactly 3-length paths we compute

Next[1]

	①	②	③	④	⑤	⑥
①		●				
②						
③						
④						
⑤			●	●		
⑥				●		

From vertex 5, we could get to these vertices with 1 hop

X						
X						
X						
X						
●						●
						X

Next[1] LOR.LAND Next[2]^T


We found two 3-length paths:

- from ⑤ to ①
- from ⑤ to ⑥.

Benchmark results


BENCHMARK RESULTS

- The top solution of AWFY vs. SuiteSparse:GraphBLAS v3.3.3
- AWFY's solution uses SIMD instructions → difficult to port
- GraphBLAS load times are slow (see details in paper)



DIRECTION-OPTIMIZING TRAVERSAL

- This optimization is subject to future work.
- For low diameter graphs, it is worth to use push/pull phases
- Push/pull is simple to express in GraphBLAS
 - See [[Yang et al., ICPP'18](#)]
- But deciding *when* to change is non-trivial



SUMMARY

- An interesting case study, see [technical report](#)
- GraphBLAS can capture mixed workloads
 - Induced subgraph computations are simple to express
 - Algorithms are concise, bitwise optimizations can be adopted
 - Performance is *sometimes* on par with specialized solutions
- Future optimizations
 - Bitmap-based matrix/vector compression is WIP by Tim Davis
→ **≈5× speedup** without using bitwise operators in our code

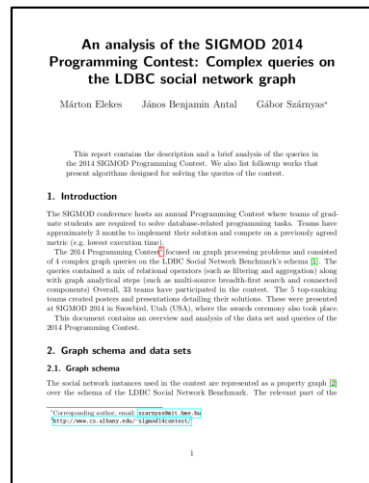


[GraphBLAS](#)

Extended
slide deck



[sigmod2014-contest-graphblas](#)



Ω