GraphAlg: Algorithm Support in a Graph Database, Done Right Sciloke

<u>Daan de Graaf</u>, Robert Brijder, Soham Chakraborty^{*}, George Fletcher, Bram van de Wall, Nikolay Yakovets

TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY

* TU Delft others TU Eindhoven

About Me

- 1st year PhD @TU Eindhoven (Database Group)
- Current Research: Algorithm support in GDBs
- Research Interests:
 - Query Processing
 - Hardware Acceleration
 - Compilers









Proposed Approach

Key Features

Types of Graph Queries





Proposed Approach



Limitations of Graph Query Languages

Graph query languages are great for simple queries

... but **lack expressive power** for Graph Analytics

Users waste resources and complicate workflows

by processing in external tools





Proposed Approach

O Key Features

Traditional Approaches to Analytics in Databases

Approach	Key Problems	Available in
Built-in Algorithms Library	- Fixed set of Algorithms	_∩eo4j
Pregel API	Performance issuesNot analysable	ArangoDB
User-defined operators	- Unsafe - Not analysable	UMBRA
Recursive CTE	Difficult to writePerformance issues	
Procedural SQL	OverheadLimited analysis	
Algorithm DSL	ProprietaryNo integration with queries	Oracle Labs



O Key Features

What Should Algorithm Support Look Like?

- **Flexible:** Fully custom algorithms
- User-friendly: Clear and convenient syntax
- Fully integrated: Native support in the database
- **Optimizable:** Efficiently process large graphs

Proposed Approach

Key Features

Introducing GraphAlg

- A language built for graph algorithms
- Fully integrated into **AvantGraph**¹
- Highly optimizable
- Embed algorithms into queries



source: avantgraph.io



Proposed Approach

Key Features

Computational Model

- Vertex-centric (Pregel)
- Vertex/Edge sets
- Linear Algebra
 - High-level operations that are easily parallelized
 - Semantics are well-defined and widely taught
 - Proven efficient for graph analytics, see **<u>GraphBLAS</u>**.



source: GraphBLAS Forum

Edited by Jeremy Kepner and John Gilbert



Graph Algorithms in the Language of Linear Algebra



CONTRIBUTORS

Bader, Bliss, Bond, Buluç, Dunlavy, Edelman, Faloutsos, Fineman, Gilbert, Heitsch, Hendrickson, Kegelmeyer, Kepner, Kolda, Leskovec, Madduri, Mohindra, Nguyen, Rader, Reinhardt, Robinson & Shah

sian.

source: SIAM







Powerful Language, Small Core

- Reducible to core language, without loss of expressivity
- Equivalent to MATLANG¹
- Loop construct **balancing expressivity & optimizability**



Proposed Approach



Cross-Optimization

- Unified IR for query and algorithm
- Eliminate query/algorithm interface boundary
- Holistic optimization & execution



Benchmark: LDBC Graphalytics

- Expressivity: Support **all algorithms** in Graphalytics spec.
- Performance: Use Graphalytics synthetic and real-world datasets, comparing:
 - Reference implementations
 - DuckDB (Python API)
 - Neo4J (Pregel API)



The graph & RDF benchmark reference

Implications

- A significant jump of programmability for graph databases
- Blurring the line with graph analytics frameworks
- Graph databases as a platform for large-scale data analysis

Thanks! Questions?