

Tripoline: Generalized Incremental Graph Processing via Graph Triangle Inequality

Xiaolin Jiang*, Chengshuo Xu*, Xizhe Yin*, Zhijia Zhao, Rajiv Gupta Computer Science and Engineering University of California Riverside

*First three authors contributed equally to this research.

Background

<u>Streaming graph processing</u>

- A stream of updates (e.g., new edges) are continuously applied to the graph
- E.g., social network, online shopping





Existing Solutions to Streaming Graph Processing

• Incremental evaluation

- Kineograph [EuroSys'12], Naiad [SOSP'13], Tornado [SIGMOD'16], Kickstarter [ASPLOS'17]
- Start the evaluation from results of a prior evaluation, rather than from scratch



Limitation in Existing Incremental Graph Processing

• <u>Require a priori knowledge of (standing) query</u>



Limitation in Existing Incremental Graph Processing

- <u>Require a priori knowledge of (standing) query</u>
 - A serious issue for vertex-specific queries (e.g., BFS, SSSP, SSWP, etc.)



Limitation in Existing Incremental Graph Processing

- <u>Require a priori knowledge of (standing) query</u>
 - A serious issue for vertex-specific queries (e.g., BFS, SSSP, SSWP, etc.)



Goal of This Work

- Require a prior knowledge of (standing) query

Incremental evaluation of an arbitrary user query (of the same type)



- Properties regarding three different vertices form a triangle
 - Inspired the classic triangle inequality



Triangle Inequality in Euclidean Space



Distance-based Triangle Inequality in Graphs* *A sketch-based distance oracle for web-scale graphs [Sarma'10]

Graph triangle inequalities exist for many vertex-specific properties



Graph triangle inequalities exist for many vertex-specific properties



• Abstraction with a generalized "Addition" and "Greater Than"



for any u, r, and x property $(u, r) \oplus property(r, x) \ge property(u, x)$

Key Idea: Connect Queries with Graph Triangle Inequality

Build constraints between Standing Query and User Query



Graph Triangle Inequality-based Incremental Evaluation



triangle inequality-based inc. eval. of q(u)

Assumptions and Correctness

• Assumptions: Monotonicity and Async-Safety

- from GraphLab [UAI'10], GRAPE [VLDB'17], Subway [Eurosys'20], Kickstarter [ASPLOS'17], etc.



Complexity: Handling Directed Graphs



inc. eval. of q(u)

Complexity: Handling Directed Graphs



Complexity: Handling Directed Graphs



Complexity: Standing Query Selection

• Triangle-based Selection

- select the query $q(r^*)$ that can minimize the **initial values** for all possible user queries

$$r^* = \underset{r \in V}{\arg\min} \sum_{\forall x \in V} property(u, r) \oplus property(r, x)$$

 $property(u, r) \bigoplus property(r, x) \ge property(u, x)$

- Topology-based Selection
 - select the query $q(r^*)$ that can maximize the **reachability** to all possible user queries

$$r^* = \underset{r \in V}{\operatorname{arg max}} \operatorname{degree}(r)$$

Complexity: Standing Query Selection

- <u>Two-Step Standing Query Selection</u>
 - Step-1: maintaining K high-degree standing queries continunously and incrementally

 $Standing_K = \{q(r_1), q(r_2), \cdots, q(r_K)\}$

- Step-2: given a user query q(u), select one standing query to apply triangle inequality

$$r^* = \underset{r \in Standing_K}{\arg \min} property(u, r)$$

Implementation

- Tripoline: implemented on top of Aspen [PLDI'19]
 - graph streaming: based on compressed tree
 - *parallel processing*: based on Ligra [PPoPP'13]





Evaluation Setup

• Machine environment

- Intel Xeon CPU E5-2683 v4 CPU (32 cores) and 512GB memory
- CentOS 7.9 and g++ 8.3.

• Graph queries

- BFS, SSSP, SSWP, Viterbi, SSNP, SSNSP, SSR, Radii
- source vertex: 256 randomly selected vertices

• Real-world large graphs

- LiveJournal (LJ), Twitter (TW), Orkut (OR), Friendster (FR)

• Streaming scenario

- starting ratio of edges: 50%, 60%, 70%
- graph update batch size: 10K
- num. of standing queries: 1-64 (16 by default)

• Speedups of Tripoline over the non-incremental evaluation

Graph	SSSP	SSWP	Viterbi	BFS	SSNP	SSR	Radii	SSNSP
OR-60	2.42x	33.91x	37.94x	1.25x	29.06x	10.86x	1.22x	1.09x
	[0.17s]	[0.01s]	[0.01s]	[0.13s]	[0.01s]	[0.01s]	[2.43s]	[0.27s]
FR-60	1.34x	35.23x	41.48x	1.02x	18.77x	10.44x	1.18x	1.00x
	[12.26s]	[0.38s]	[0.38s]	[7.16s]	[0.45s]	[0.45s]	[56.43s]	[9.58s]
LJ-60	1.81x	11.56x	26.88x	1.12x	11.53x	5.50x	1.16x	1.03x
	[0.13s]	[0.01s]	[0.02s]	[0.07s]	[0.02s]	[0.02s]	[1.31s]	[0.20s]
TW-60	1.95x	15.97x	19.14x	1.56x	13.42x	8.32x	1.15x	1.18x
	[1.45s]	[0.13s]	[0.13s]	[0.94s]	[0.14s]	[0.15s]	[11.85s]	[2.27s]
Avg.	1.89x	23.28x	30.52x	1.24x	18.24x	8.83x	1.18x	1.08x

- Reduction in vertex function activations
 - higher activation ratio \rightarrow lower speedup

$$R_{act} = \frac{Num_{act} \ triangle - based \ inc. eval.}{Num_{act} \ non - inc. eval.}$$

Table: Vertex Function Activation Ratio

	OR-60	FR-60	LJ-60	TW-60	
SSSP	44.4%	61.7%	56%	52.8%	
SSWP	1.9e-9	1.3e-8	0.79%	4.0e-8	
Viterbi	^{3.5} min	-max natu	re of the S	SWP -7	
BFS	82.2%	98%	89.4%	65.8%	
SSNP	1.9e-7	1.4e-8	0.78%	3.6e-8	
SSR	3.3e-7	1.7e-8	0.78%	3.2e-8	
Radii	98.9%	91.9%	92.21% 93.9%		
SSNSP	98.9%	99.97%	98.58%	94.9%	

• Speedup distribution across queries



Figure: Speedups of 256 Queries (16 for Radii) on LiveJournal

 $r^* = \underset{r \in Standing_K}{\operatorname{arg min}} property(u, r)$

• Confirmation of standing query selection heuristics



Figure: Correlations between Speedups and *property(u,r)*

• Integeration into Differential Dataflow (with Shared Arrangements) [VLDB'20]

Graph	Method	BFS	SSSP	SSWP
LJ-100	DD-SA	1.10s	8.41s	4.63s
	DD-SA-Tri	1.11s	3.24s	0.52s
	Speedup	[0.99x]	[2.60x]	[8.90x]
	DD-SA	10.69s	58.63s	32.68s
TW-100	DD-SA-Tri	10.71s	14.72s	7.74s
	Speedup	[1.00x]	[3.98x]	[4.22x]

Table: DD with Triangle Inequality Optimization

Table: Reduction of Reduce Operations

Graph	Method	BFS	SSSP	SSWP
	DD-SA	9156594	30418846	20622003
LJ-100	DD-SA-Tri	8956638	17570555	6292821
	Reduction	[1.02x]	[1.73x]	[3.28x]

DD-SA: DD with shared arrangements enabled (baseline) DD-SA-Tri: DD with triangle inequality opt.



Conclusion

- Graph triangle inequality is common for vertex-specific graph queries
- It is key to enable **generalized incremental evaluation**
- Prototyped <u>**Tripoline**</u> system
- Observed significant speedups for multiple types of graph queries

Thank You!