A Hierarchical Synchronous Parallel Model for Wide-Area Graph Analytics

<u>Shuhao Liu*</u>, Li Chen, Baochun Li, Aiden Carnegie University of Toronto April 17, 2018

What is Graph Analytics?

What is Graph Analytics?



What is Graph Analytics?

Shortest Path



What is Graph Analytics?

Shortest Path





What is Graph Analytics?

Shortest Path



PageRank



What kind of graph are we interested in?



Large

Large Geo-Distributed

Process it in *parallel* (Spark, Hadoop…)

- Process it in *parallel* (Spark, Hadoop…)
- Parallel graph analytics frameworks:

- Process it in *parallel* (Spark, Hadoop…)
- Parallel graph analytics frameworks:
 - ► State-of-the-art: Gemini [OSDI 16'], etc.

- Process it in *parallel* (Spark, Hadoop…)
- Parallel graph analytics frameworks:
 - ► State-of-the-art: Gemini [OSDI 16'], etc.
 - Assume:

- Process it in *parallel* (Spark, Hadoop…)
- Parallel graph analytics frameworks:
 - ► State-of-the-art: Gemini [OSDI 16'], etc.
 - Assume:
 - Data is accessible via fast network

- Process it in *parallel* (Spark, Hadoop…)
- Parallel graph analytics frameworks:
 - ► State-of-the-art: Gemini [OSDI 16'], etc.
 - Assume:
 - Data is accessible via fast network
 - A cluster of *well-connected* workers

Analyze Geo-Distributed Data

Analyze Geo-Distributed Data

Reduce WAN traffic

Analyze Geo-Distributed Data

- Reduce WAN traffic
- Wide-area data analytics
 - ► Placement: Iridium [SIGCOMM 15'], etc.
 - ► Workload: Clarinet [OSDI 16'], Gaia [NSDI 17']

Analyze Geo-Distributed Graph

- Reduce WAN traffic
- Wide-area graph analytics
 - ► Placement: Mayer et al. [ICDCS 16'], Zhou et al. [ICDCS 17']
 - Workload: ?

Can we optimize the workload with the awareness of WAN transfers?

Connected Components: BSP in two regions



WAN transfer:









Connected Components: sync locally -> globally



WAN transfer:

Connected Components: sync locally -> globally



WAN transfer:

Connected Components: sync locally -> globally



Connected Components: sync locally -> globally



HSP: Design

- Two modes of synchronization: local + global
- Local synchronization
 - Keep mirrored vertices static, compute locally without WAN communication
- Global synchronization

HSP: Analysis

- Convergence
 - We have proven HSP has the same convergence guarantee as BSP
- WAN traffic \propto # Global Synchronizations
 - We have proven HSP has a much higher rate of convergence with the same amount of WAN traffic (if BSP converges linearly or super-linearly)





Graph++

GraphLab









Proof-of-Concept

PageRank example



Implementation in GraphX



Dataset	# Vertices	# Edges	Harmonic Diameter
enwiki-2013	4,206,785	101,355,853	5.24
uk-2014-host	4,769,354	50,829,923	21.48

Workle	oad	# HSP Global Sync.	# BSP Super- step	HSP Usage (GB)	BSP Usage (GB)	Reduction (%)
enwiki-	PR	46	74	18.39	27.14	32.2
2013	CC	5	7	0.69	0.91	23.4
2015	SP	7	10	0.59	0.84	30.6
uk-	PR	35	52	21.48	31.47	31.7
2014-	CC	12	20	0.71	0.95	25.4
host	SP	15	23	0.50	0.64	22.4

Dataset	# Vertices	# Edges	Harmonic Diameter
enwiki-2013	4,206,785	101,355,853	5.24
uk-2014-host	4,769,354	50,829,923	21.48

Workle	oad	# HSP Global Sync.	# BSP Super- step	HSP Usage (GB)	BSP Usage (GB)	Reduction (%)
enwiki_	PR	46	74	18.39	27.14	32.2
2013	CC	5	7	0.69	0.91	23.4
2015	SP	7	10	0.59	0.84	30.6
uk-	PR	35	52	21.48	31.47	31.7
2014-	CC	12	20	0.71	0.95	25.4
host	SP	15	23	0.50	0.64	22.4

Dataset	# Vertices	# Edges	Harmonic Diameter
enwiki-2013	4,206,785	101,355,853	5.24
uk-2014-host	4,769,354	50,829,923	21.48

Workle	oad	# HSP Global Sync.	# BSP Super- step	HSP Usage (GB)	BSP Usage (GB)	Reduction (%)
enwiki-	PR	46	74	18.39	27.14	32.2
2013	CC	5	7	0.69	0.91	23.4
2015	SP	7	10	0.59	0.84	30.6
uk-	PR	35	52	21.48	31.47	31.7
2014-	CC	12	20	0.71	0.95	25.4
host	SP	15	23	0.50	0.64	22.4

Dataset	# Vertices	# Edges	Harmonic Diameter	
enwiki-2013	4,206,785	101,355,853	5.24	
uk-2014-host	4,769,354	50,829,923	21.48	

Workle	oad	# HSP Global Sync.	# BSP Super- step	HSP Usage (GB)	BSP Usage (GB)	Reduction (%)
enwiki	PR	46	74	18.39	27.14	32.2
2013	CC	5	7	0.69	0.91	23.4
2015	SP	7	10	0.59	0.84	30.6
uk-	PR	35	52	21.48	31.47	31.7
2014-	CC	12	20	0.71	0.95	25.4
host	SP	15	23	0.50	0.64	22.4

Dataset	# Vertices	# Edges	Harmonic Diameter
enwiki-2013	4,206,785	101,355,853	5.24
uk-2014-host	4,769,354	50,829,923	21.48

Workle	oad	# HSP Global Sync.	# BSP Super- step	HSP Usage (GB)	BSP Usage (GB)	Reduction (%)
enwiki-	PR	46	74	18.39	27.14	32.2
2013	CC	5	7	0.69	0.91	23.4
2015	SP	7	10	0.59	0.84	30.6
uk-	PR	35	52	21.48	31.47	31.7
2014-	CC	12	20	0.71	0.95	25.4
host	SP	15	23	0.50	0.64	22.4

Monetary cost



HSP: Takeaways

- ► HSP = BSP + local mode
 - WAN efficiency: faster && cheaper
 - **Correctness:** strong convergence guarantee
 - **Transparency:** independent from apps

Thanks! Q & A

HSP: Design

- Switch local -> global
 - When all local partitions have run d iterations
 - When a local partition that has already converged
- Run one global iteration and switch back to local mode

Runtime



Rate of convergence

