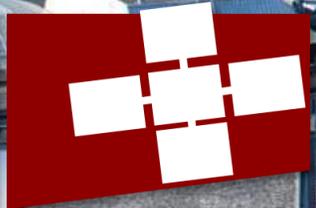


M. BESTA, T. HOEFLER

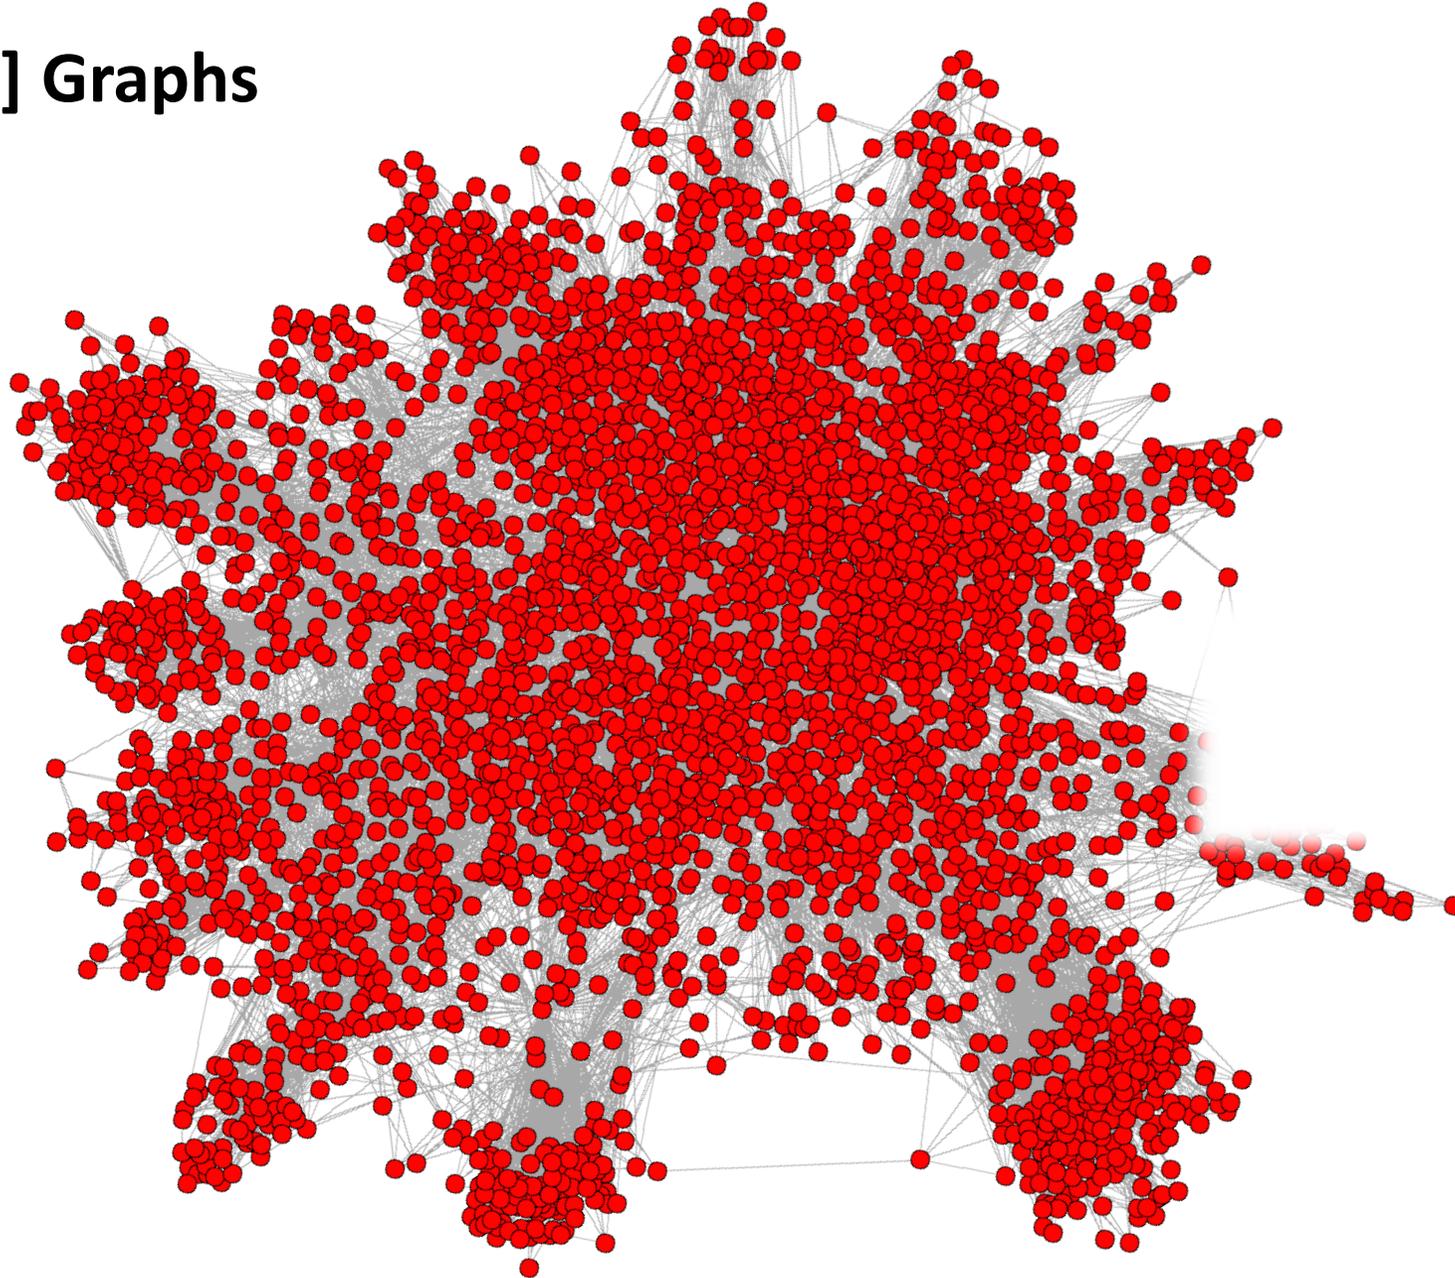
Towards High-Performance Processing, Storage, and Analytics of Extreme-Scale Graphs

With contributions from Dimitri Stanojevic, Simon Weber, Lukas Gianinazzi, Andrey Ivanov, Marc Fischer, Robert Gerstenberger, Heng Lin, Xiaowei Zhu, Bowen Yu, Xiongchao Tang, Wei Xue, Wenguang Chen, Lufei Zhang, Xiaosong Ma, Xin Liu, Weimin Zheng, and Jingfang Xu and others.



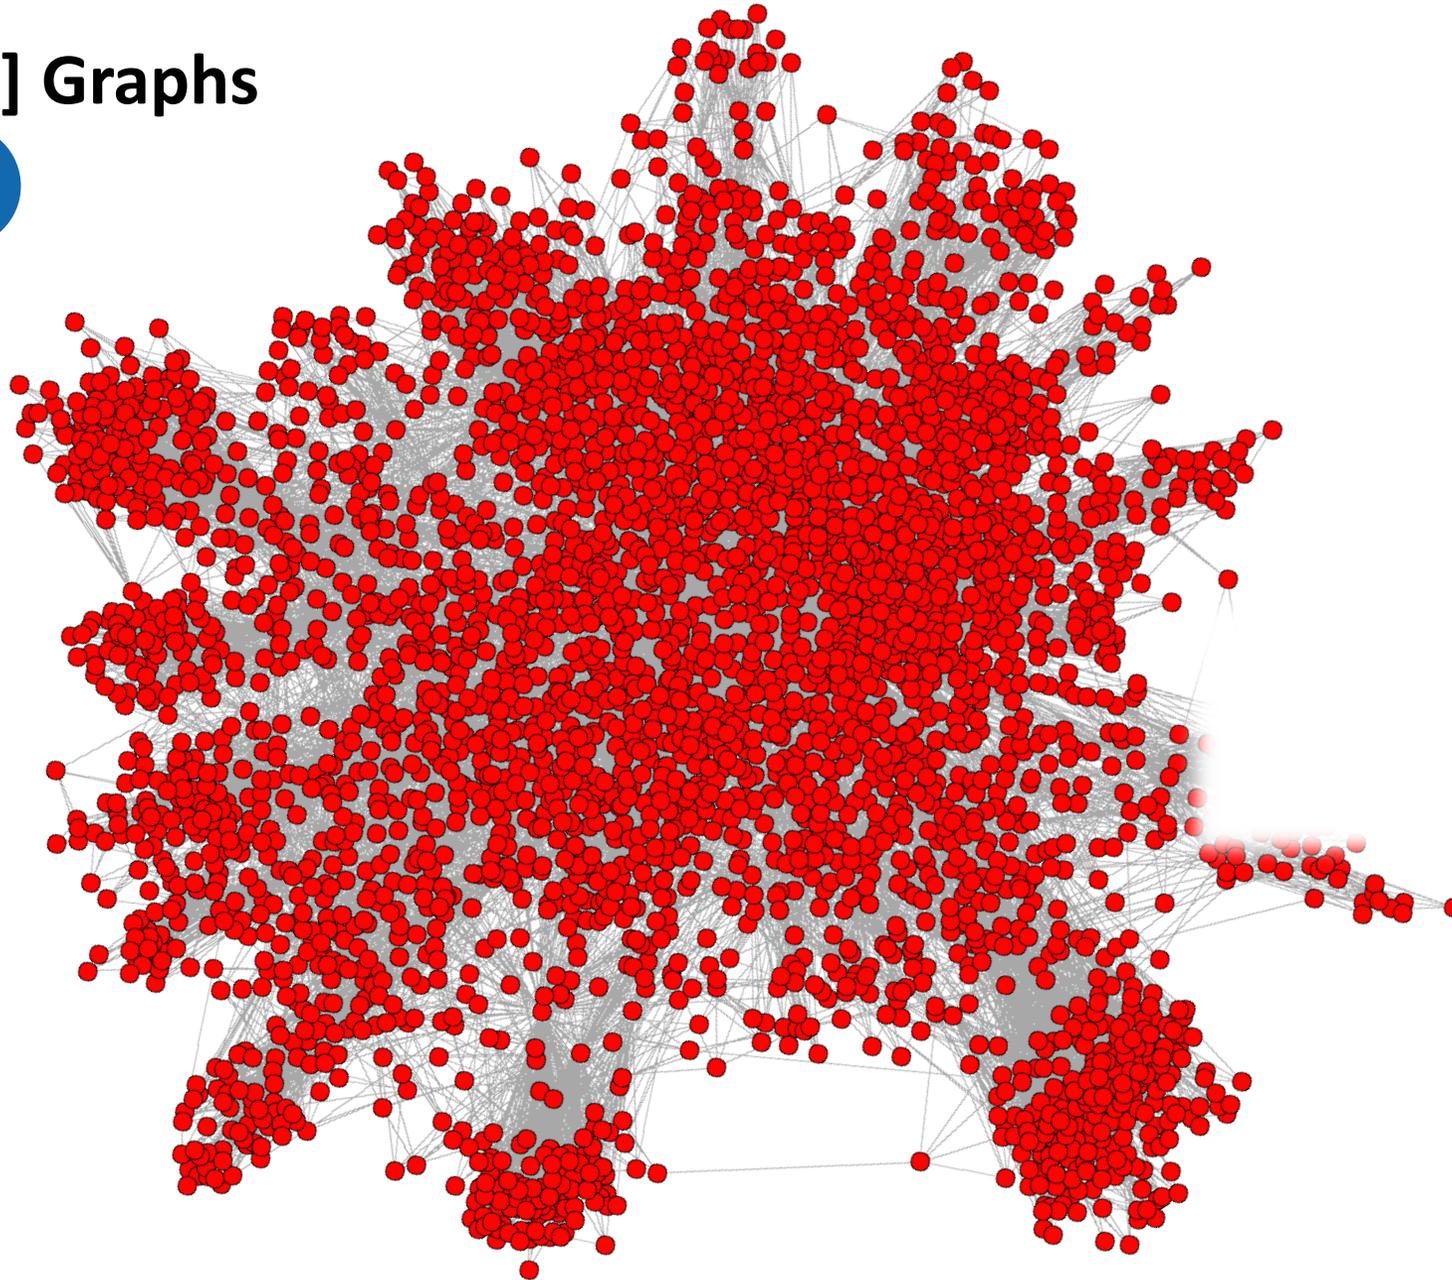
[Extreme-Scale] Graphs

[Extreme-Scale] Graphs



[Extreme-Scale] Graphs

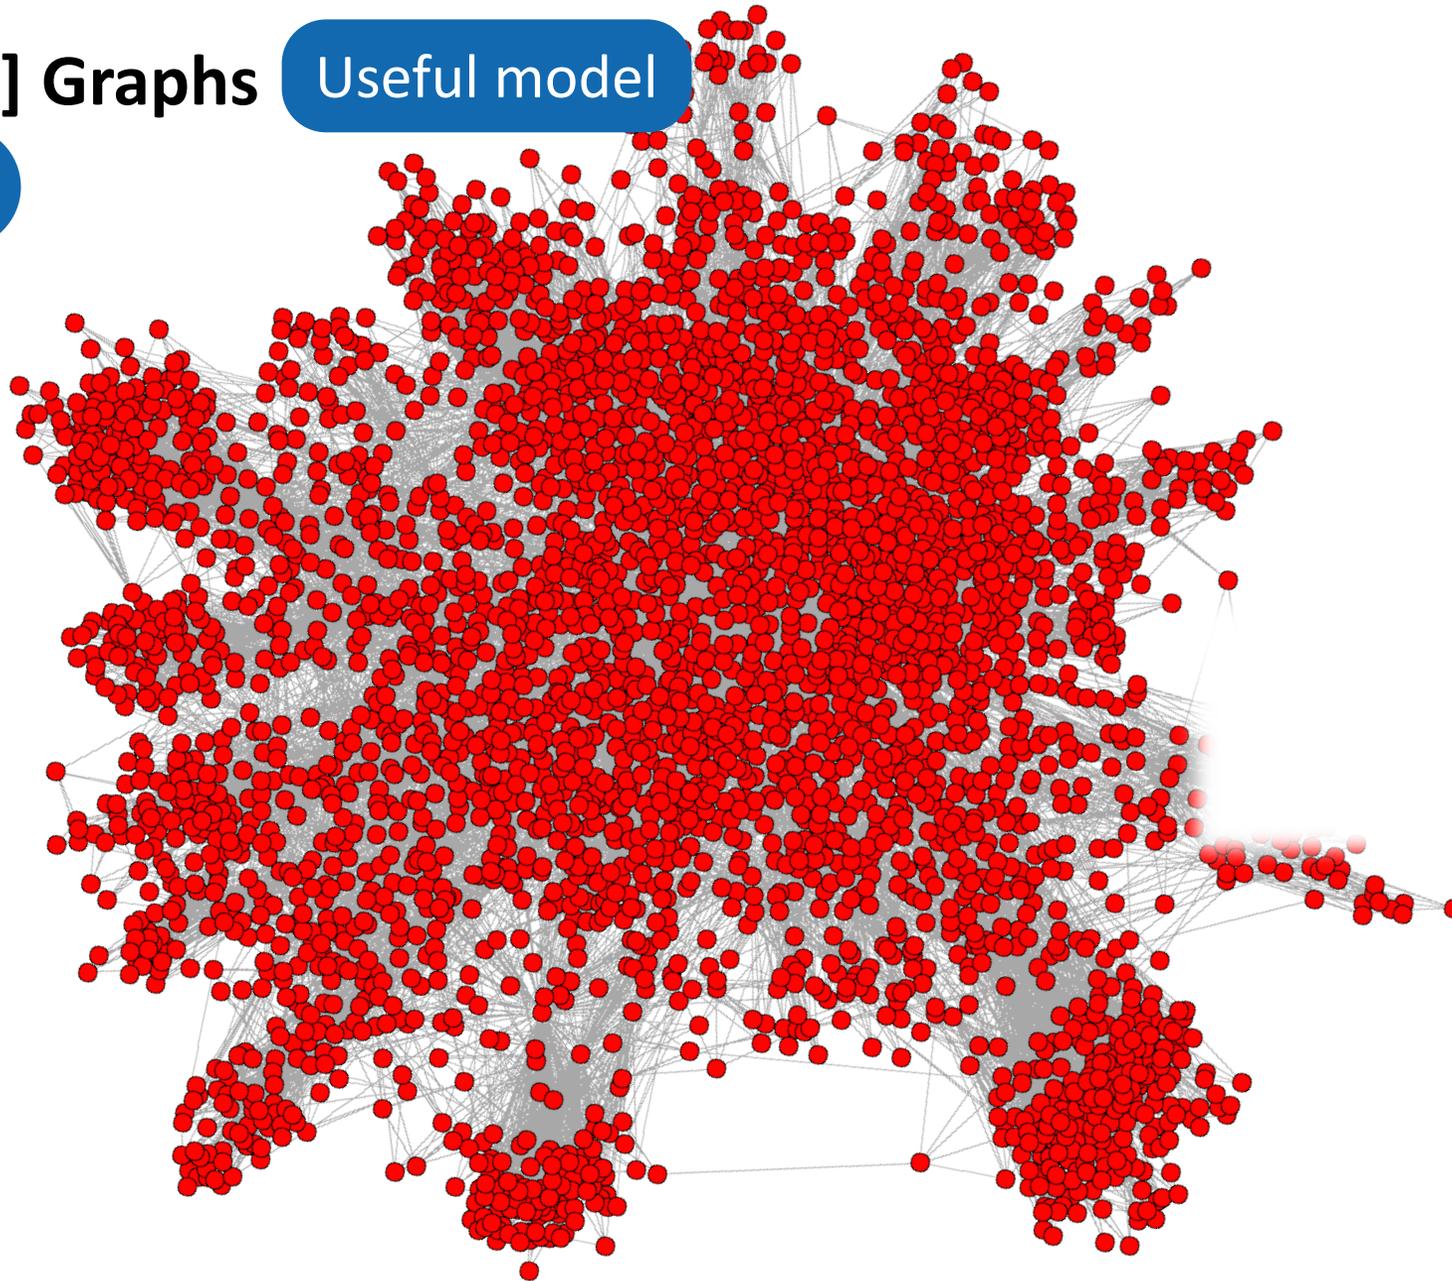
Why do we care?



[Extreme-Scale] Graphs

Useful model

Why do we care?

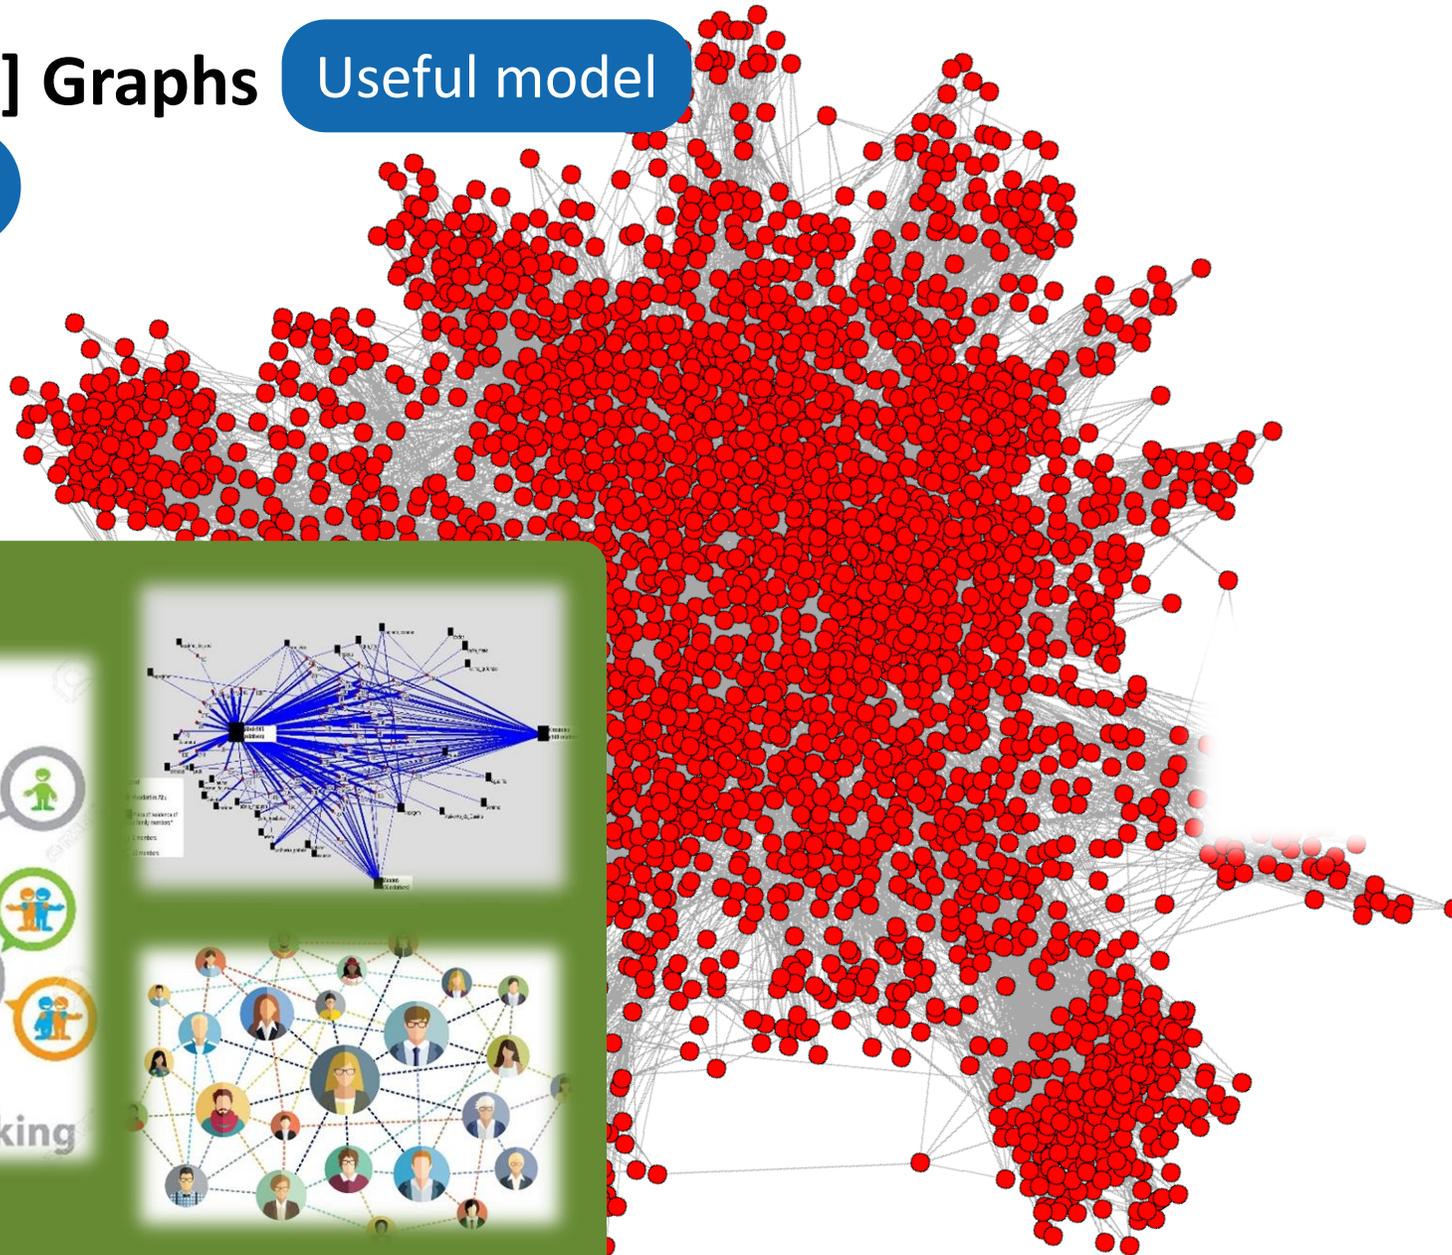
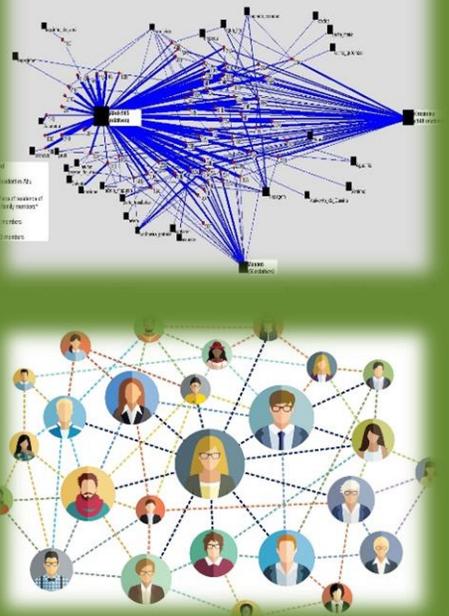


[Extreme-Scale] Graphs

Useful model

Why do we care?

Social networks

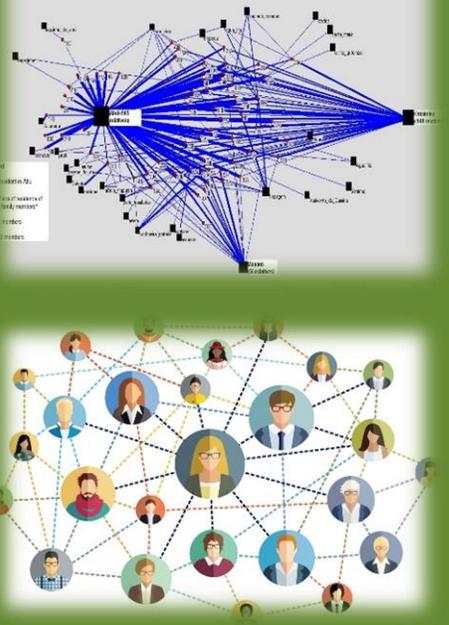


[Extreme-Scale] Graphs

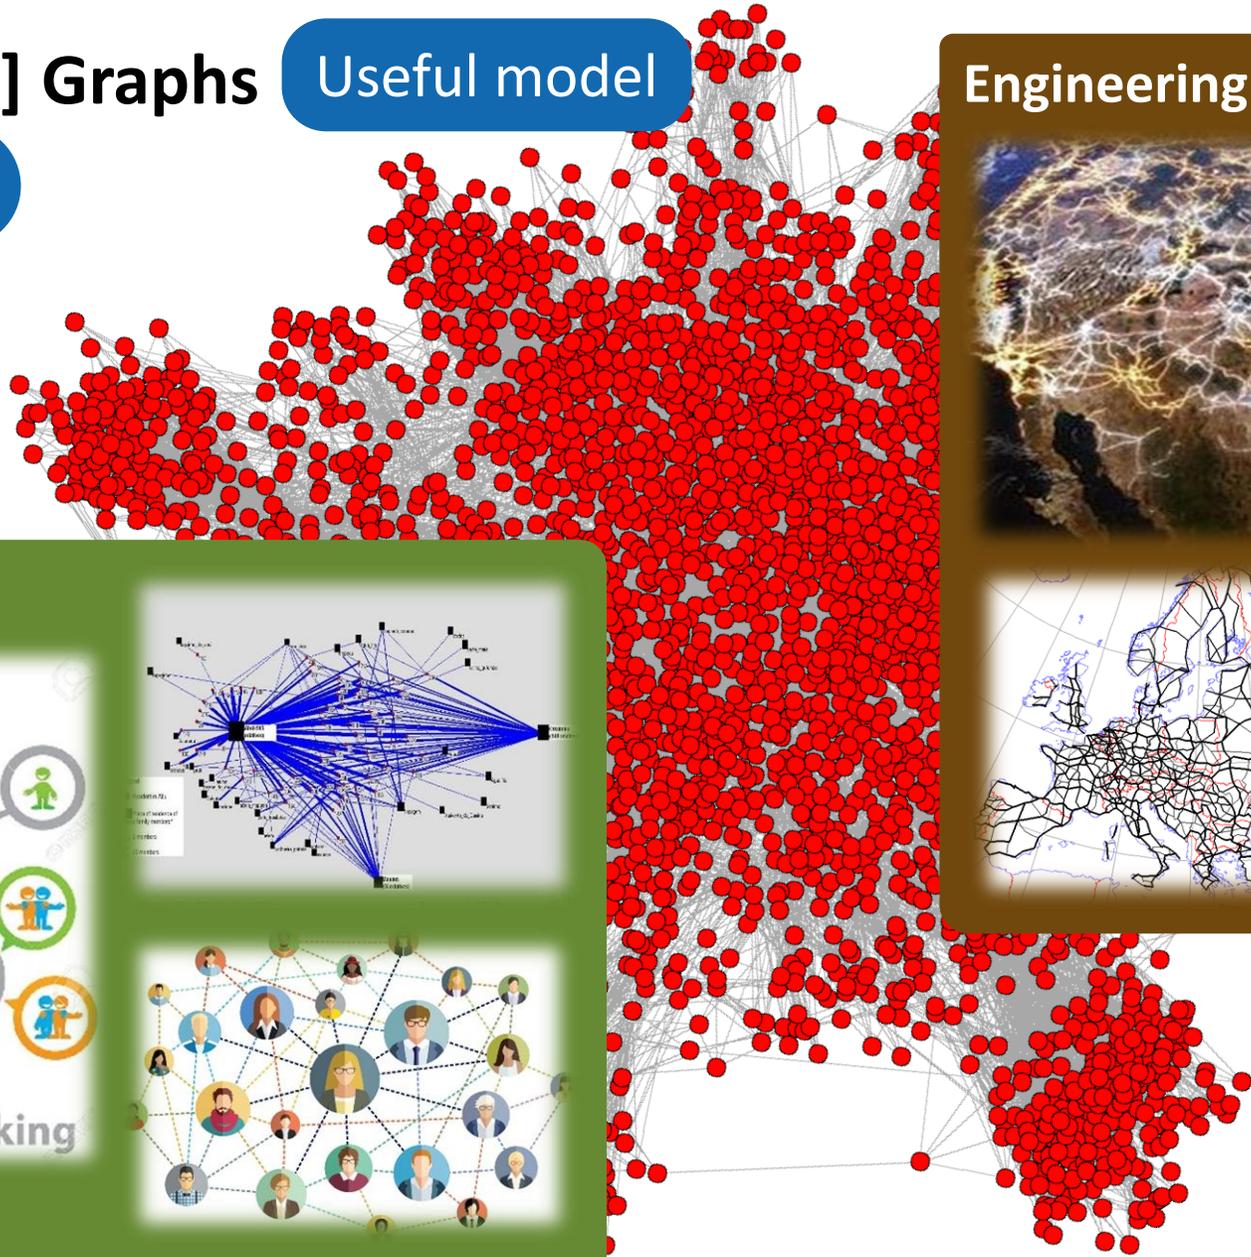
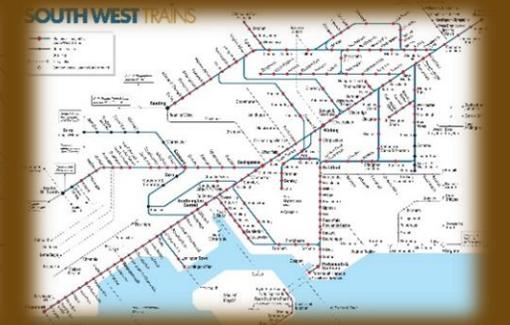
Useful model

Why do we care?

Social networks



Engineering networks



[Extreme-Scale] Graphs

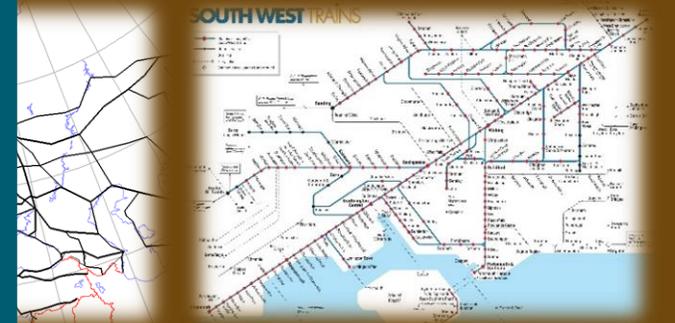
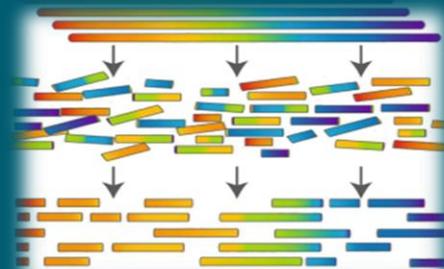
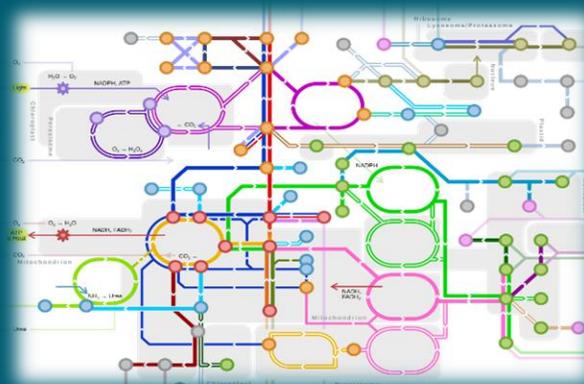
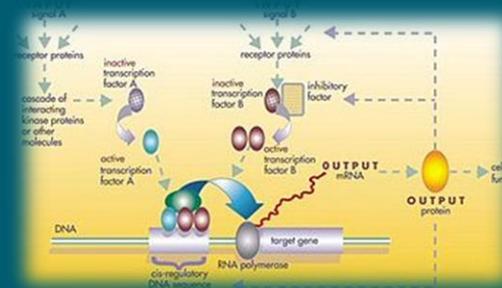
Useful model

Engineering networks

Why do we care?

Social networks

Biological networks



[Extreme-Scale] Graphs

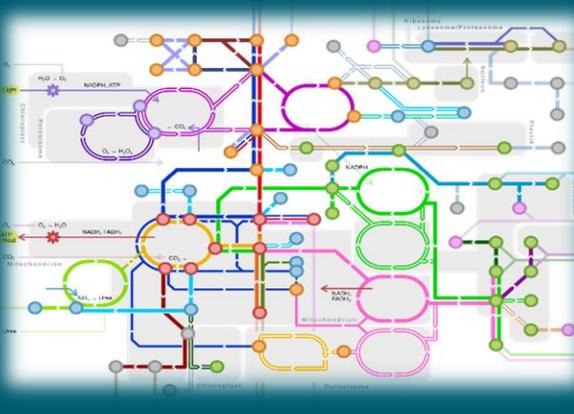
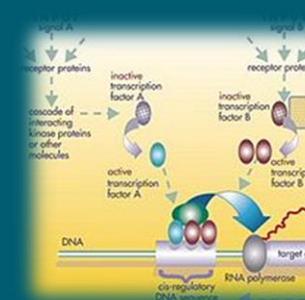
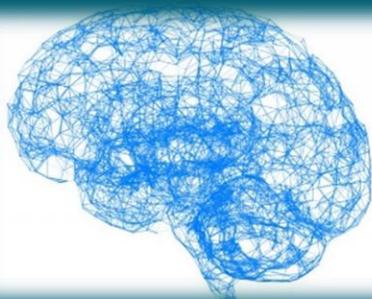
Useful model

Why do we care?

Social networks



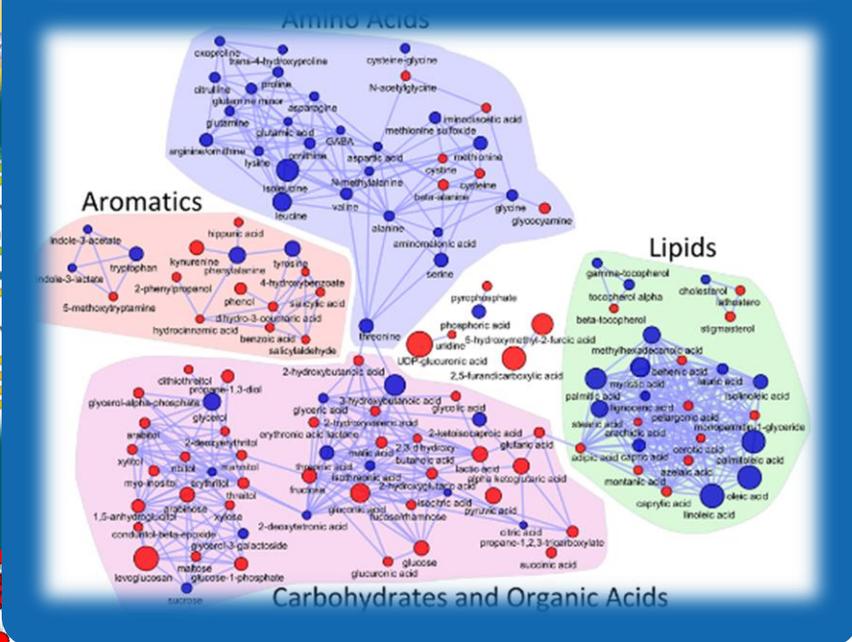
Biological networks



Engineering networks

Physics, chemistry

$$\frac{1}{\sqrt{2}}|\text{cat}\rangle + \frac{1}{\sqrt{2}}|\text{dog}\rangle$$



[Extreme-Scale] Graphs

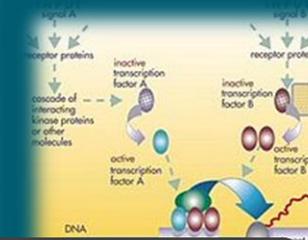
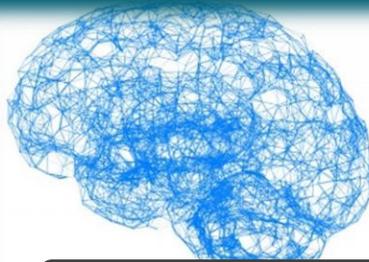
Useful model

Why do we care?

Social networks



Biological networks



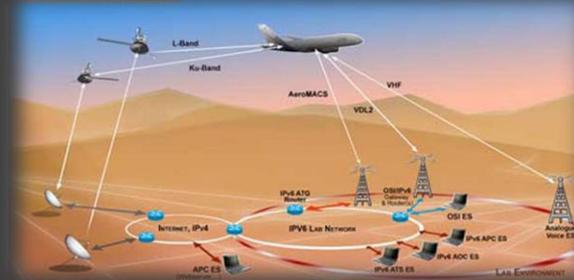
Engineering networks

Physics, chemistry

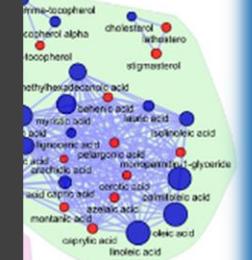
$$\frac{1}{\sqrt{2}}|\text{cat}\rangle + \frac{1}{\sqrt{2}}|\text{dog}\rangle$$



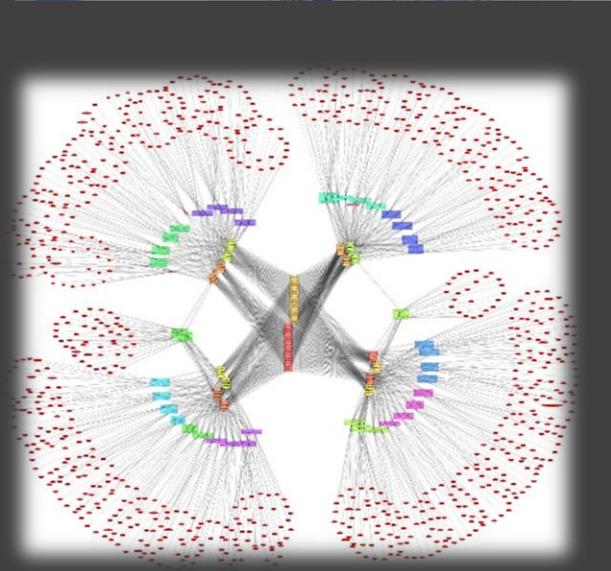
Communication networks



Lipids



Acids



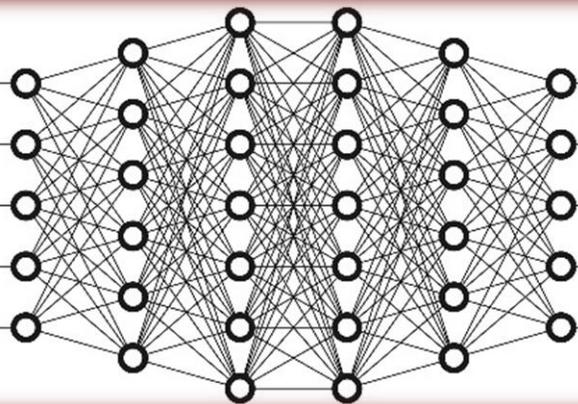
[Extreme-Scale] Graphs

Useful model

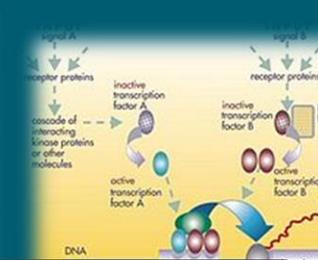
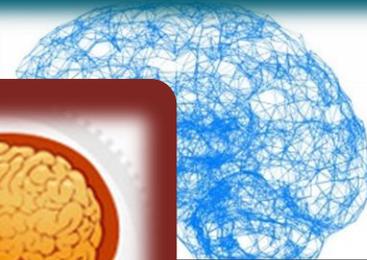
Why do we care?

Social

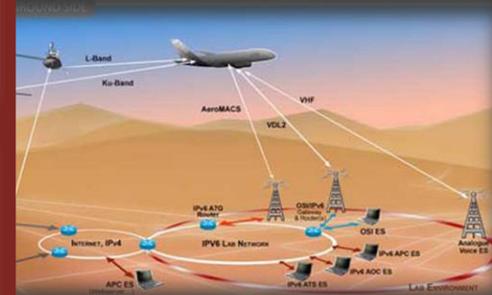
Machine learning



Biological networks



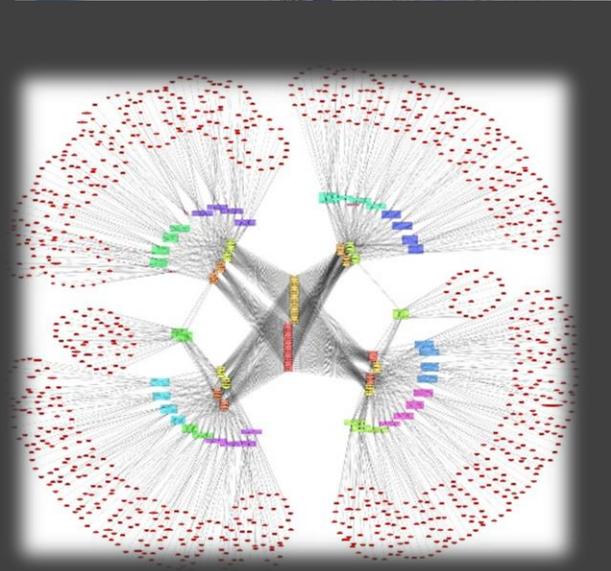
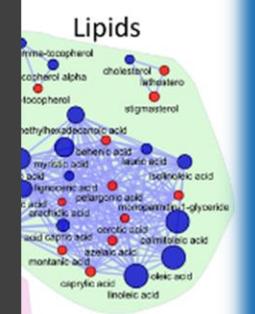
Communication networks



Engineering networks

Physics, chemistry

$$\frac{1}{\sqrt{2}}|\text{cat}\rangle + \frac{1}{\sqrt{2}}|\text{dog}\rangle$$



[Extreme-Scale] Graphs

Useful model

Engineering networks

Why do we care?

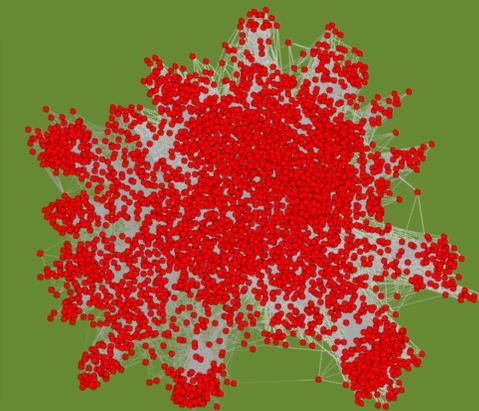
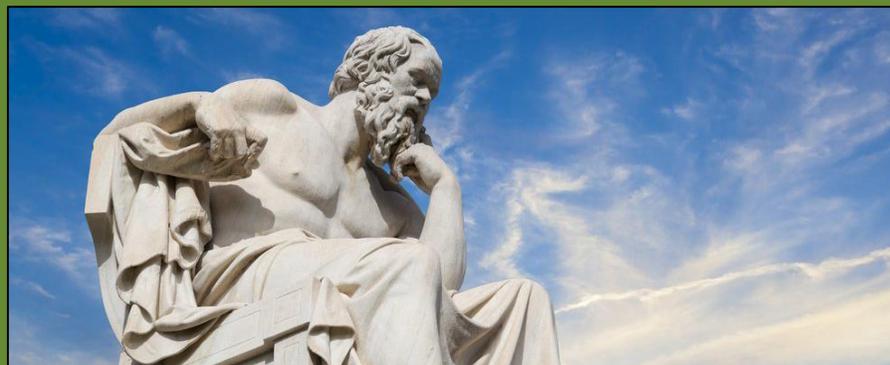
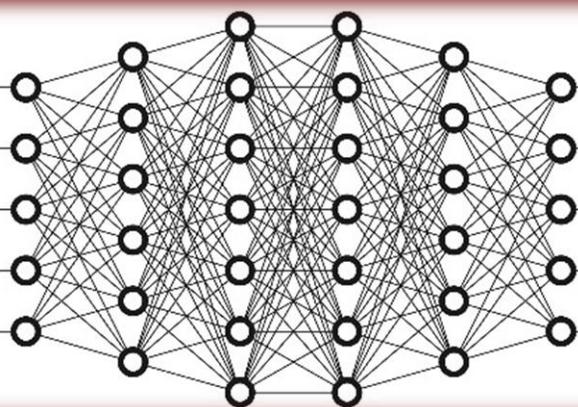
...even philosophy ☺

Physics, chemistry

Biological networks

Machine learning

Social



FOSDEM 2016 / Schedule / Events / Developer rooms / Graph Processing / Modeling a Philosophical Inquiry: from MySQL to a graph database

Modeling a Philosophical Inquiry: from MySQL to a graph database

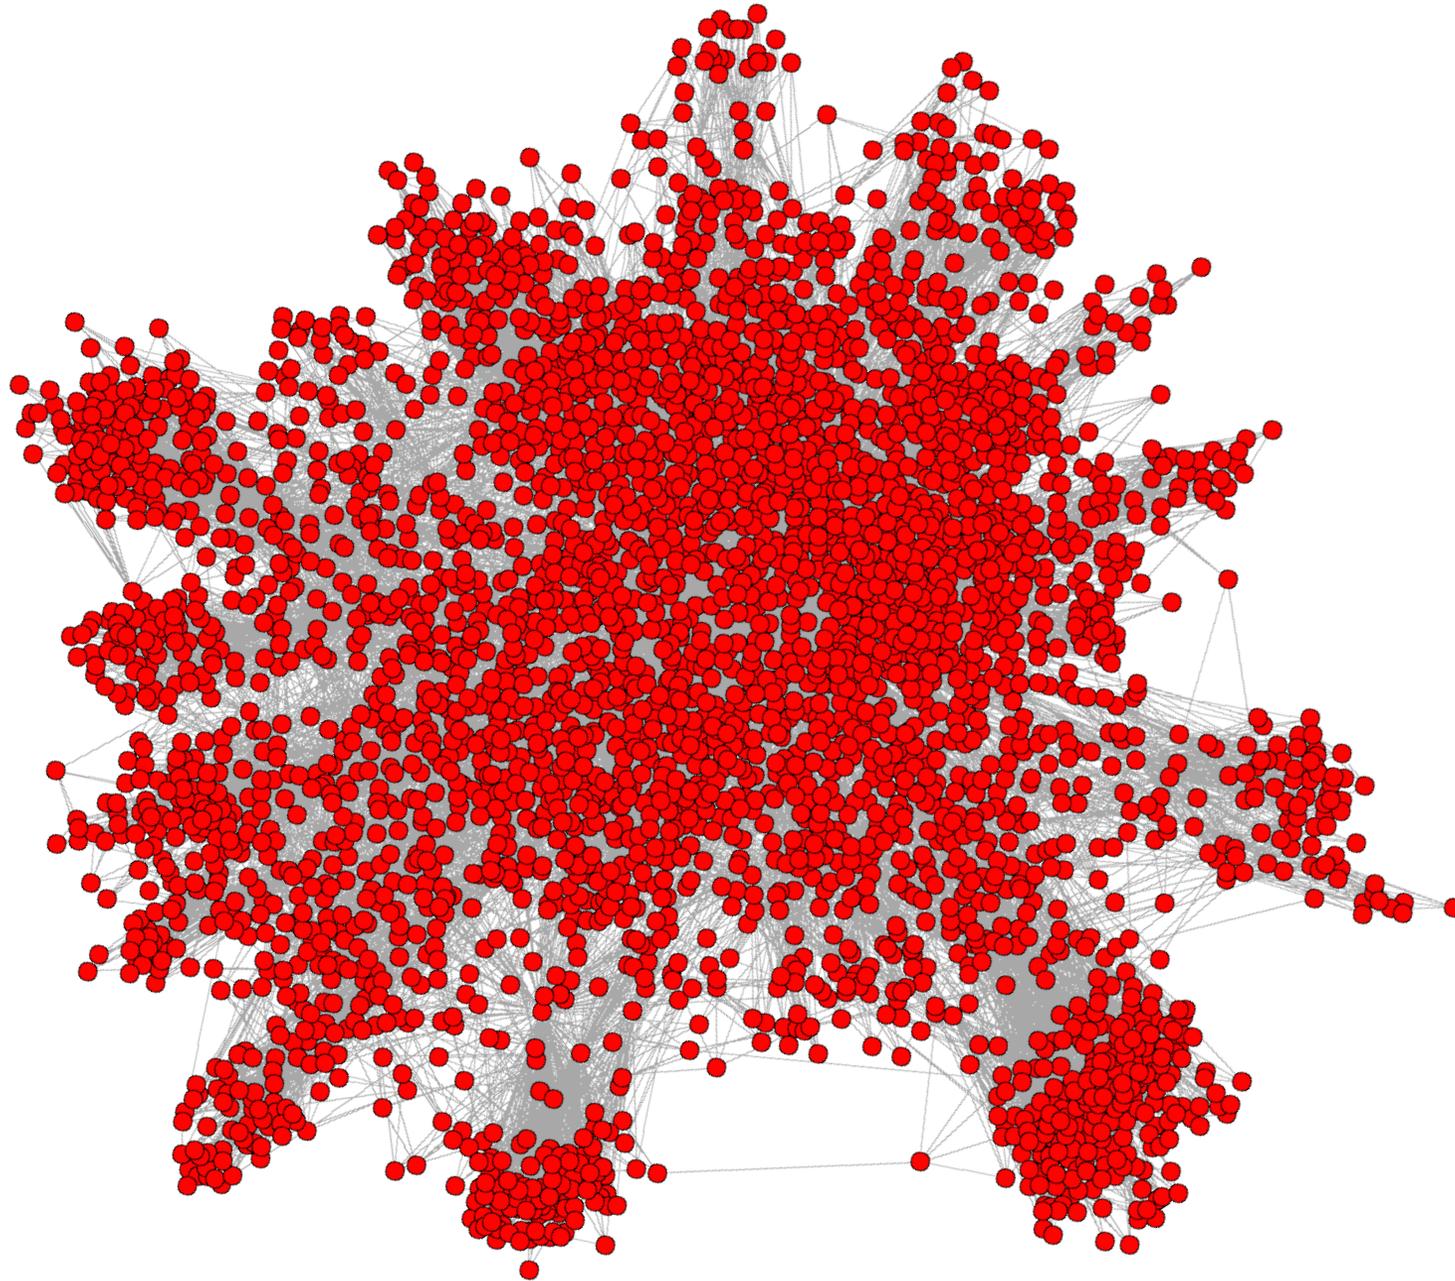
The short story of a long refactoring process

- 📍 Track: Graph Processing devroom
- 📍 Room: AW1.126
- 📅 Day: Saturday
- 🕒 Start: 12:45
- 🕒 End: 13:35

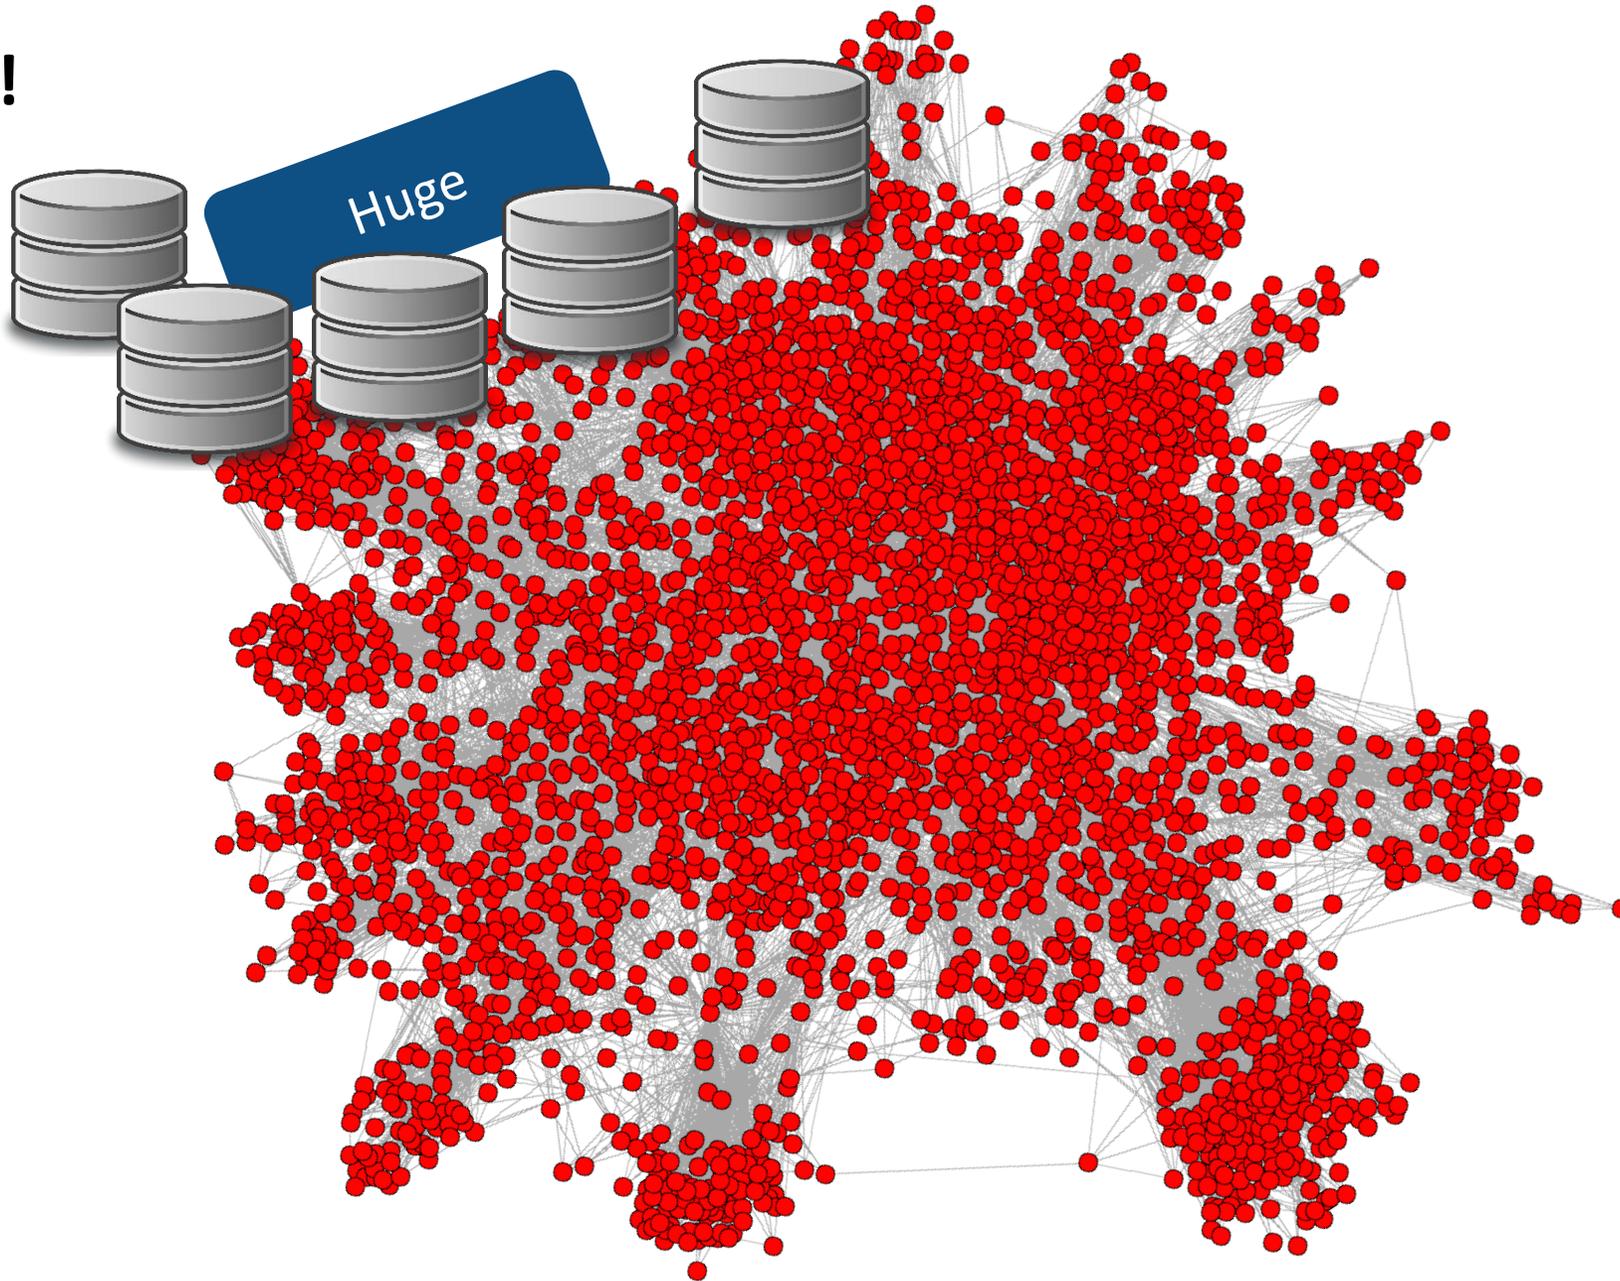


Bruno Latour wrote a book about philosophy (an inquiry into modes of existence). He decided that the paper book was no place for the numerous footnotes, documentation or glossary, instead giving access to all this information surrounding the book through a web application which would present itself as a reading companion. He also offered to the community of readers to submit their contributions to his inquiry by writing new documents to be added to the platform. The first version

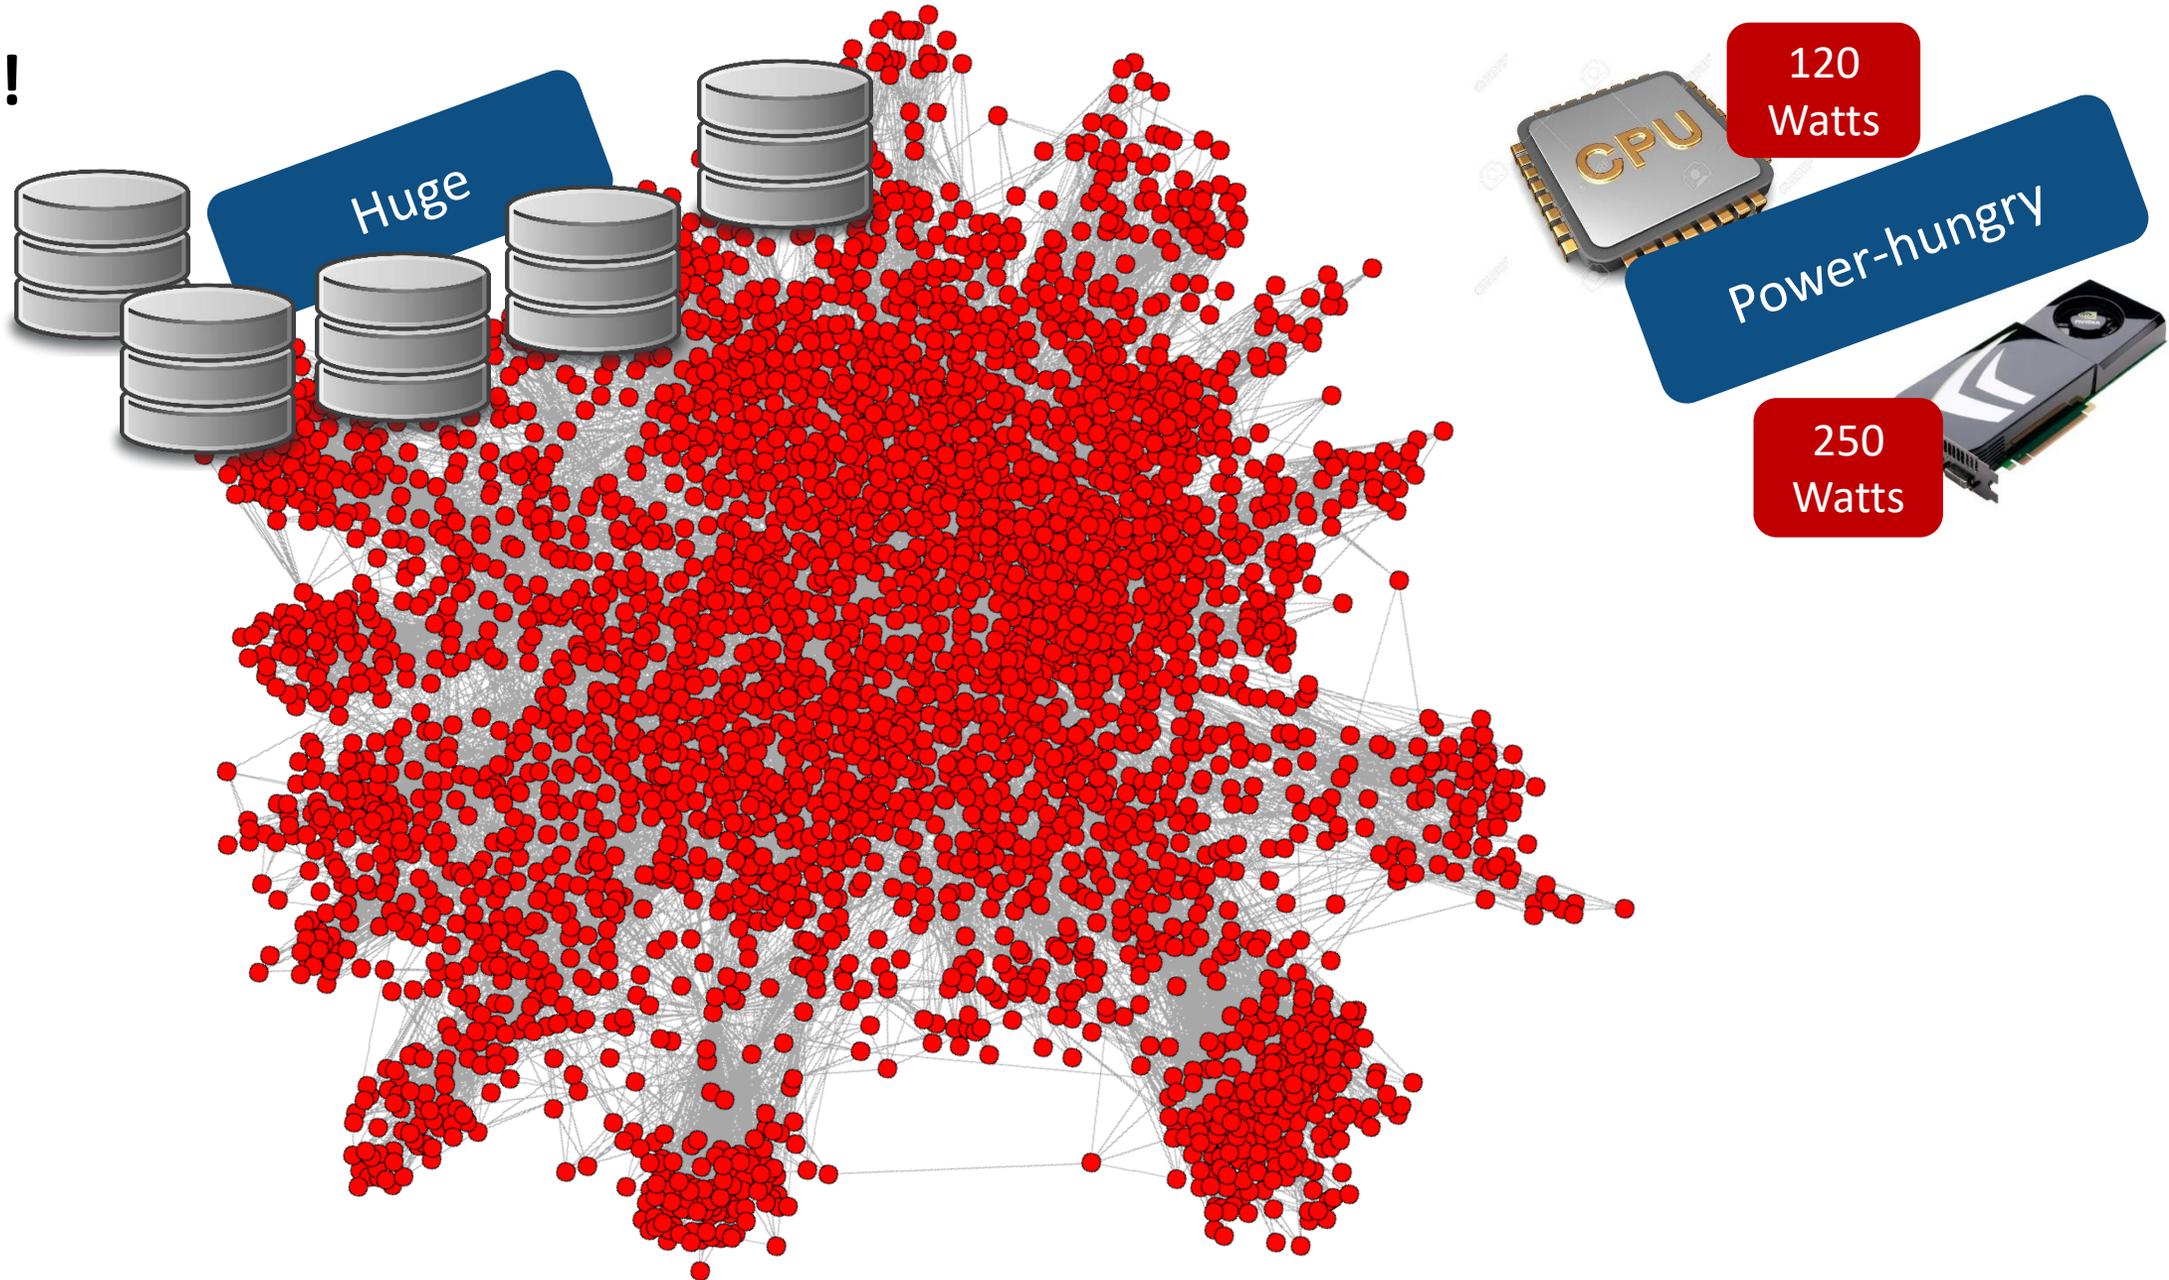
Problems!



Problems!



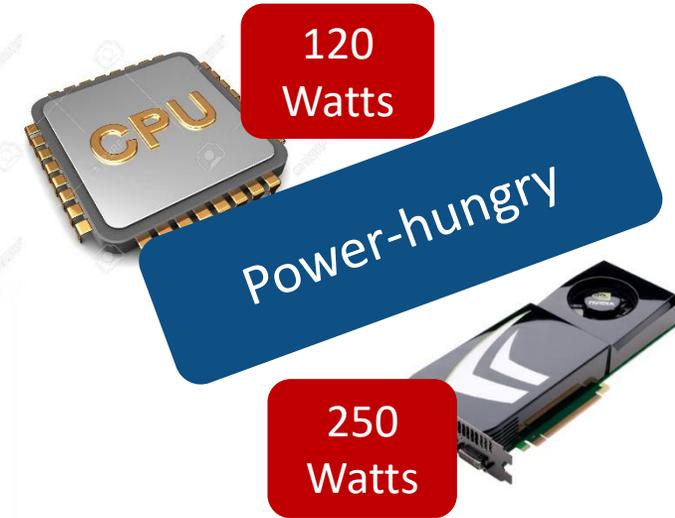
Problems!



Problems!



Problems!



120 Watts

Power-hungry

250 Watts



Problems!



Problems!



Huge

120 Watts

Power-hungry

250 Watts

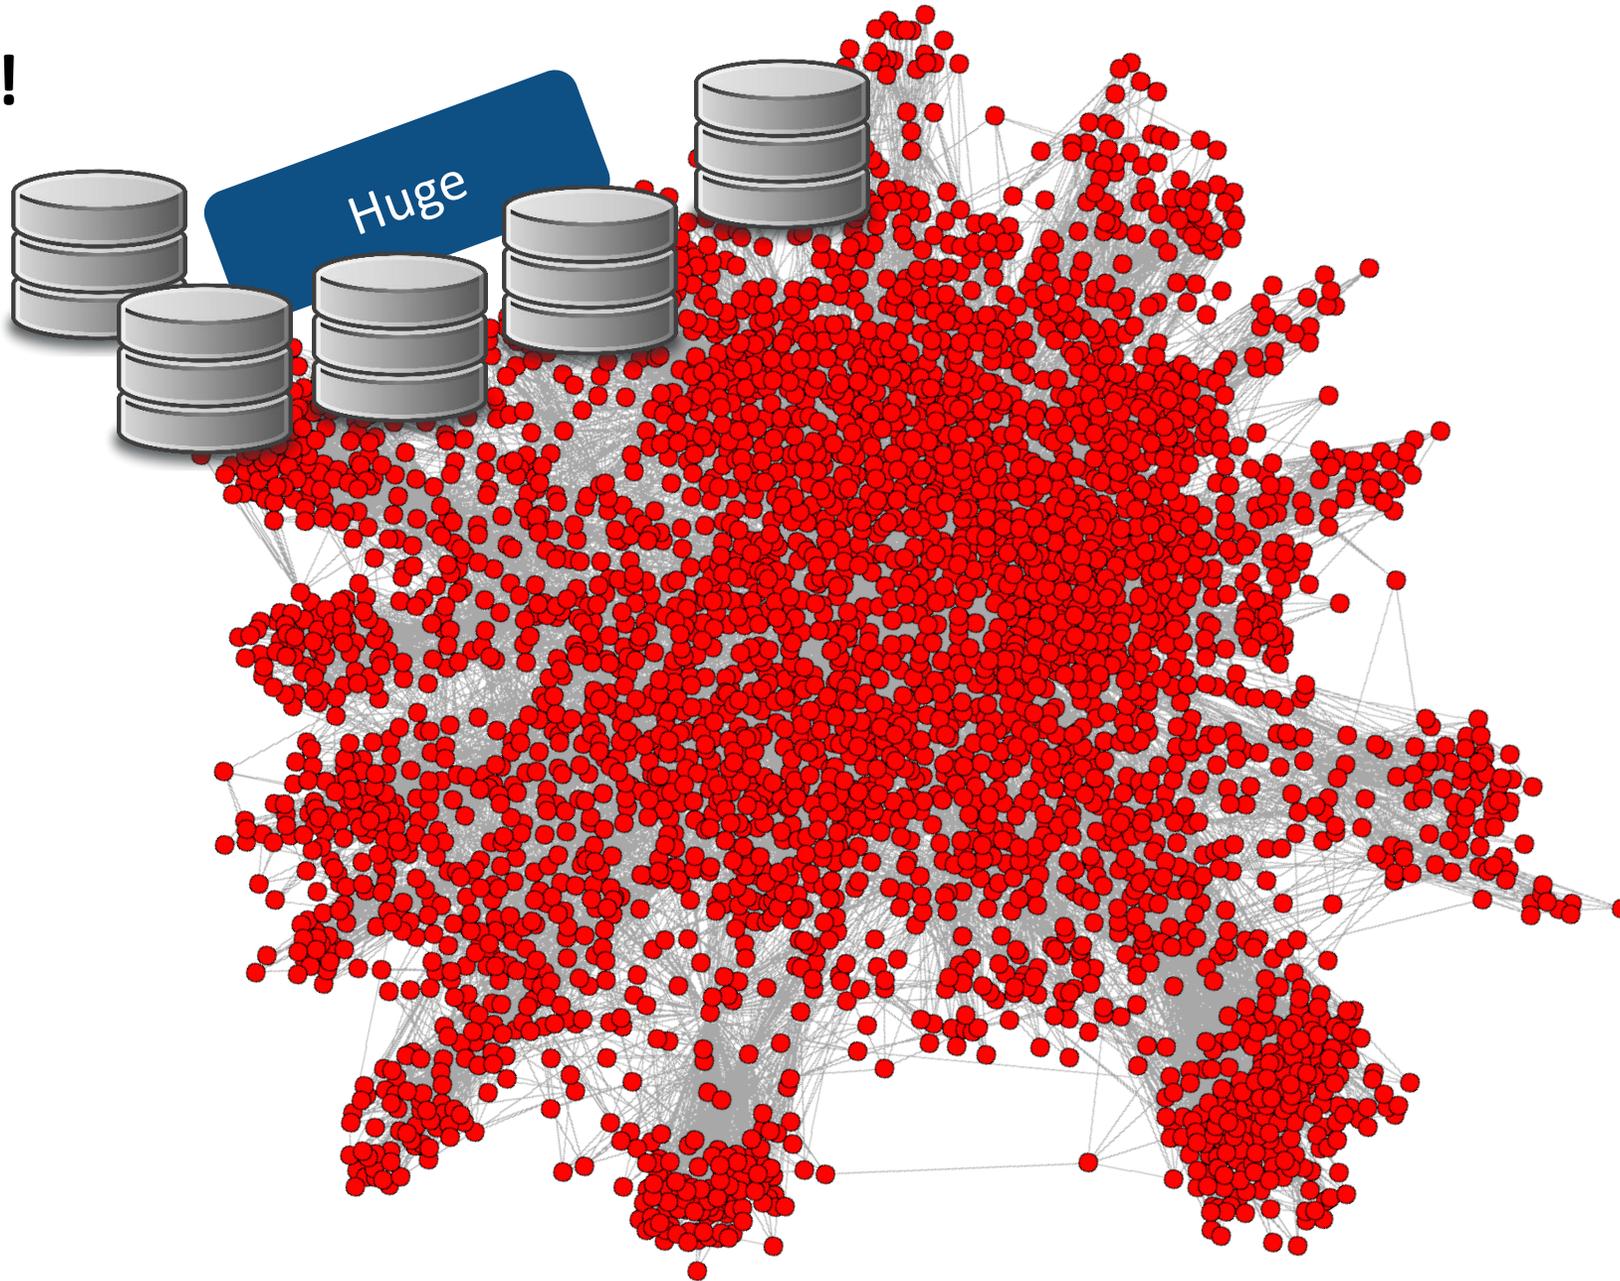
It's Complicated...

Irregular

Communication-heavy

Synchronization-heavy

Problems!

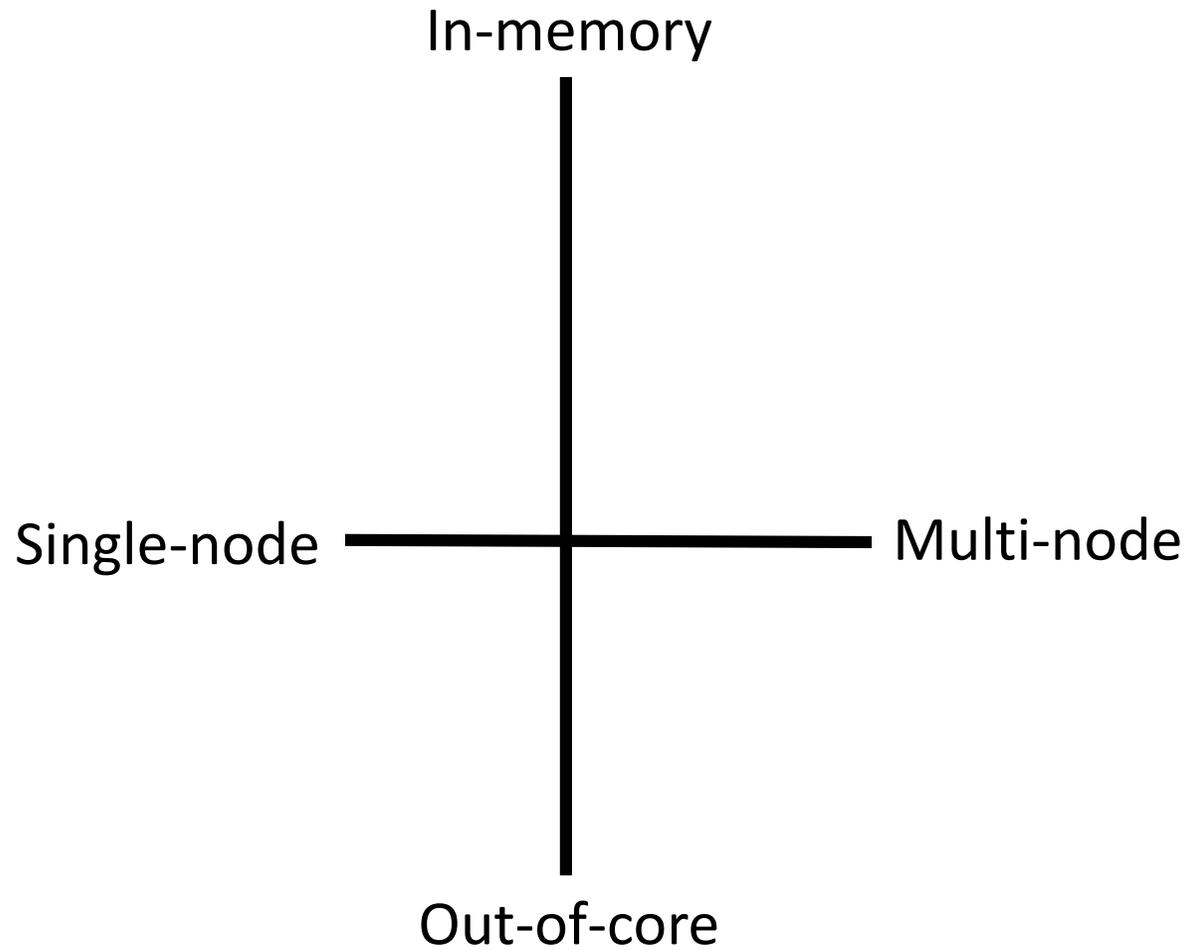


Problems!

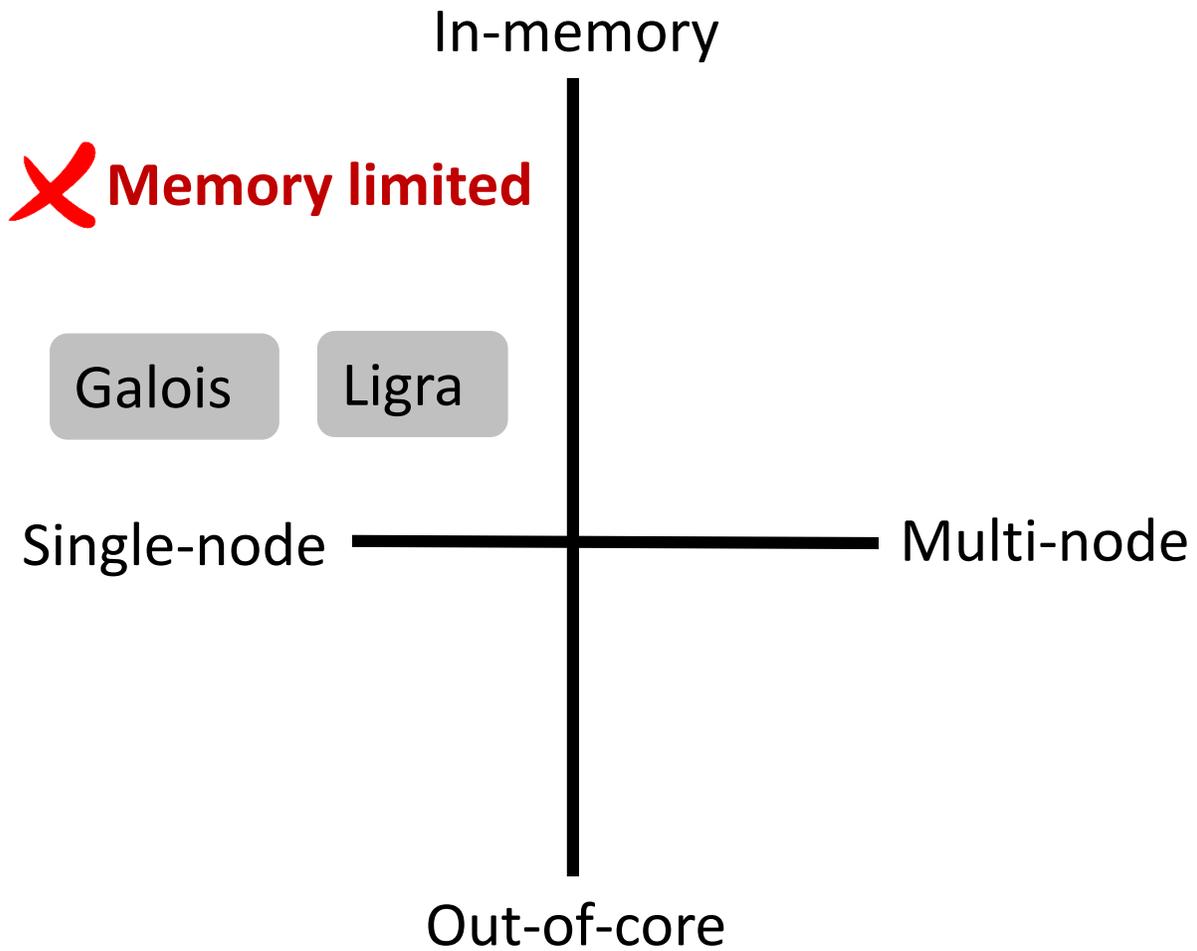


Practice of Extreme-Scale Graph Processing

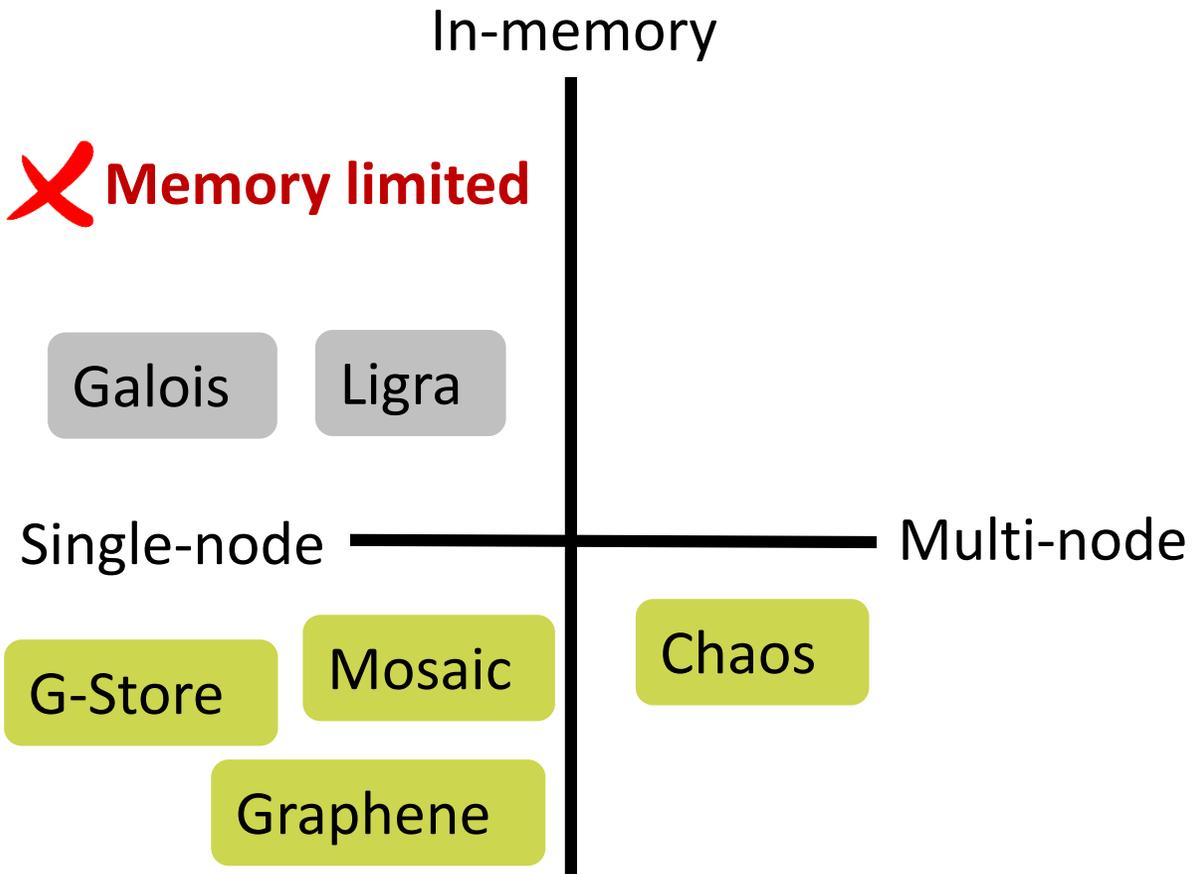
Practice of Extreme-Scale Graph Processing



Practice of Extreme-Scale Graph Processing



Practice of Extreme-Scale Graph Processing

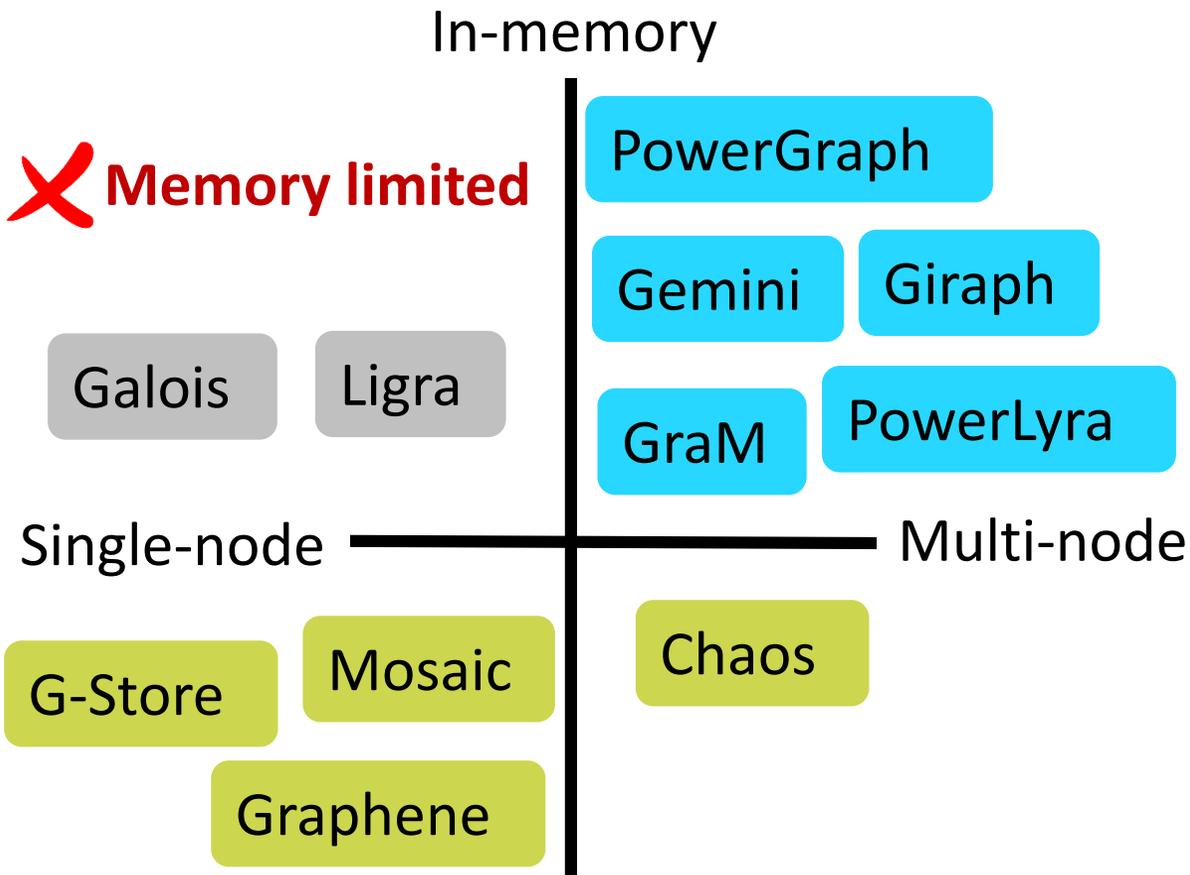


X Memory limited

X Too slow

(> 20 mins per PageRank iteration on 1-trillion-edge graph)

Practice of Extreme-Scale Graph Processing

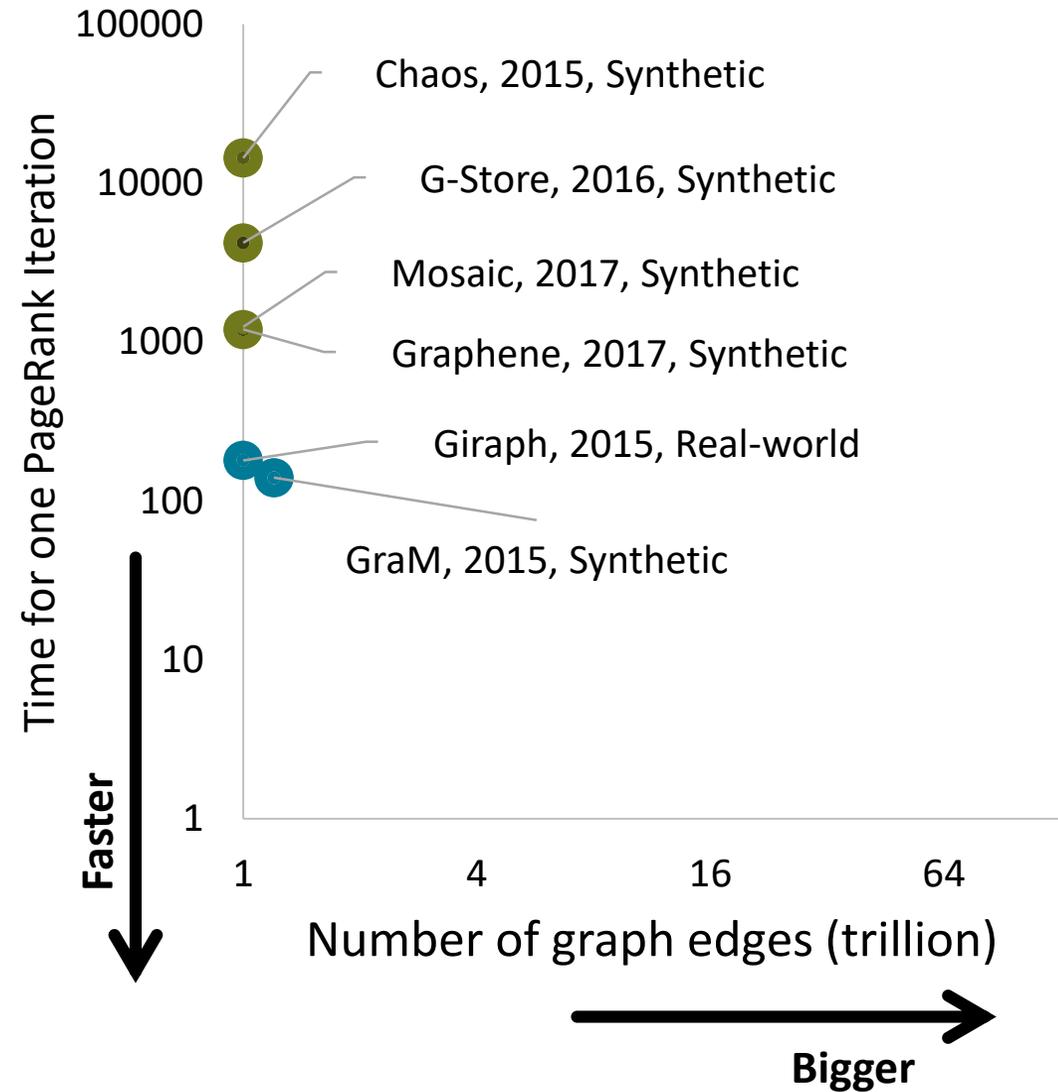
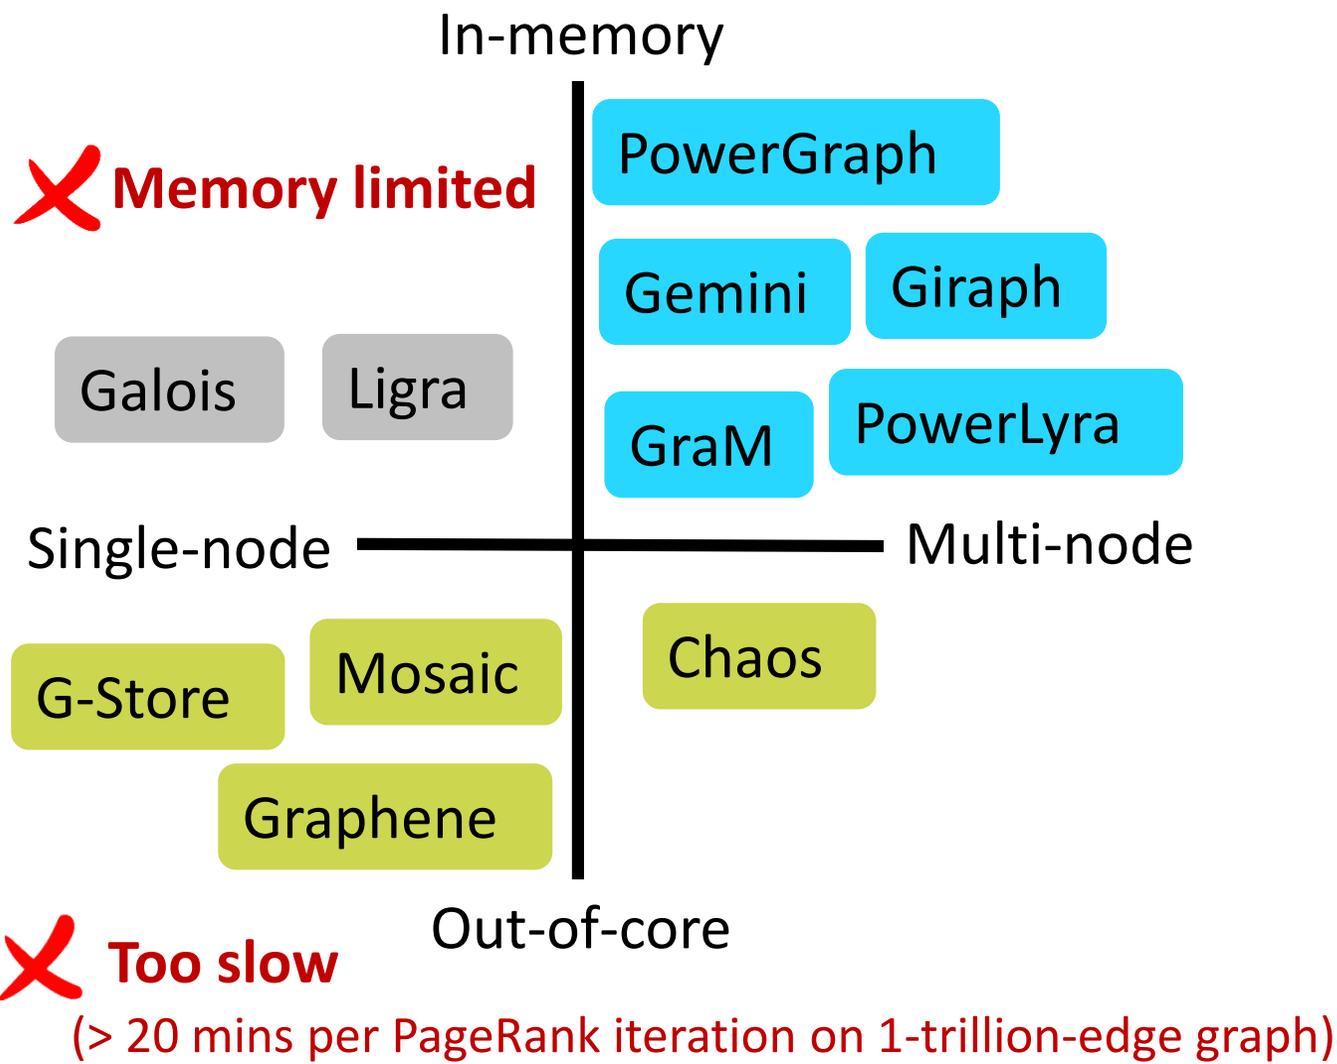


X Memory limited

X Too slow

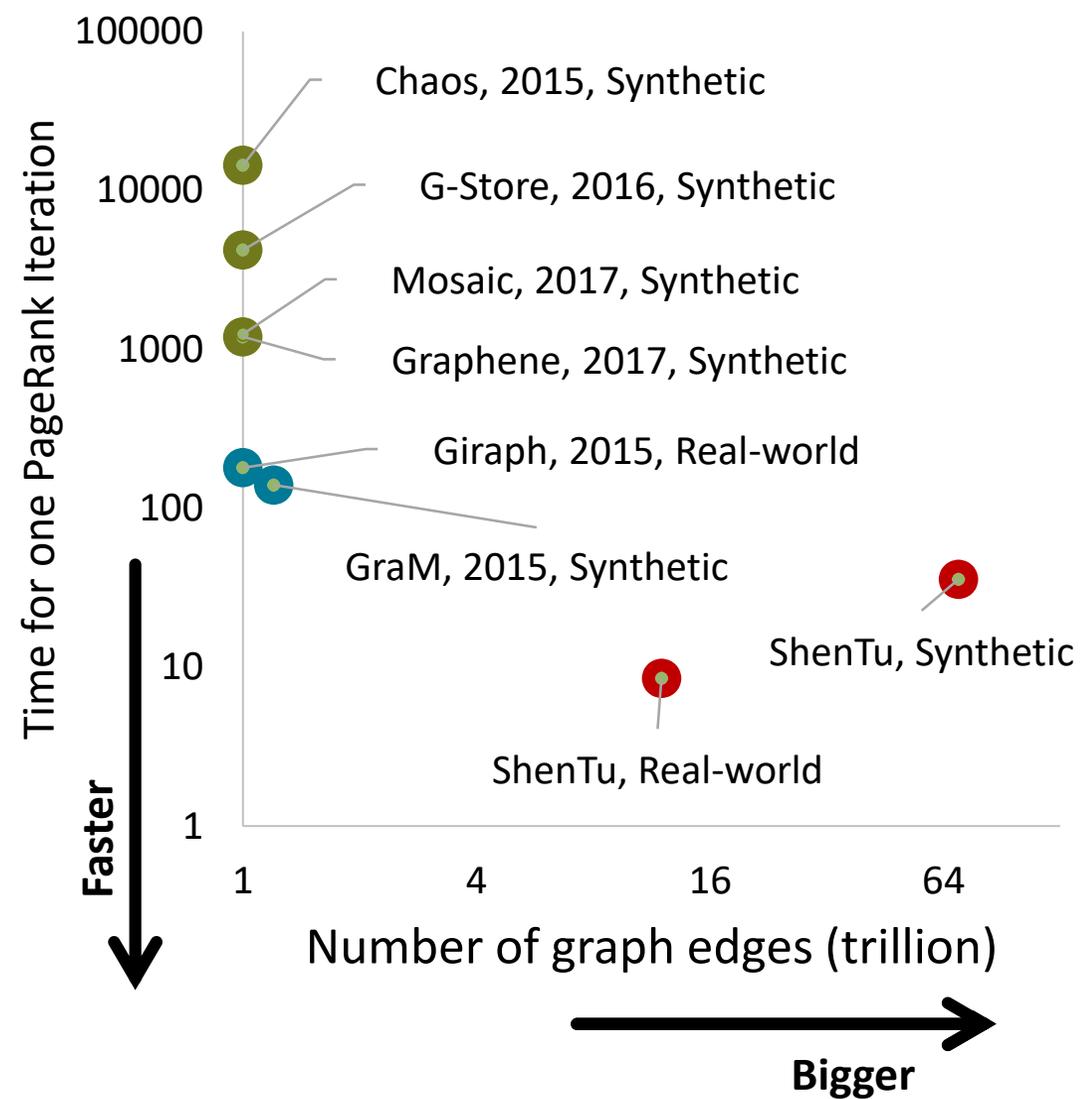
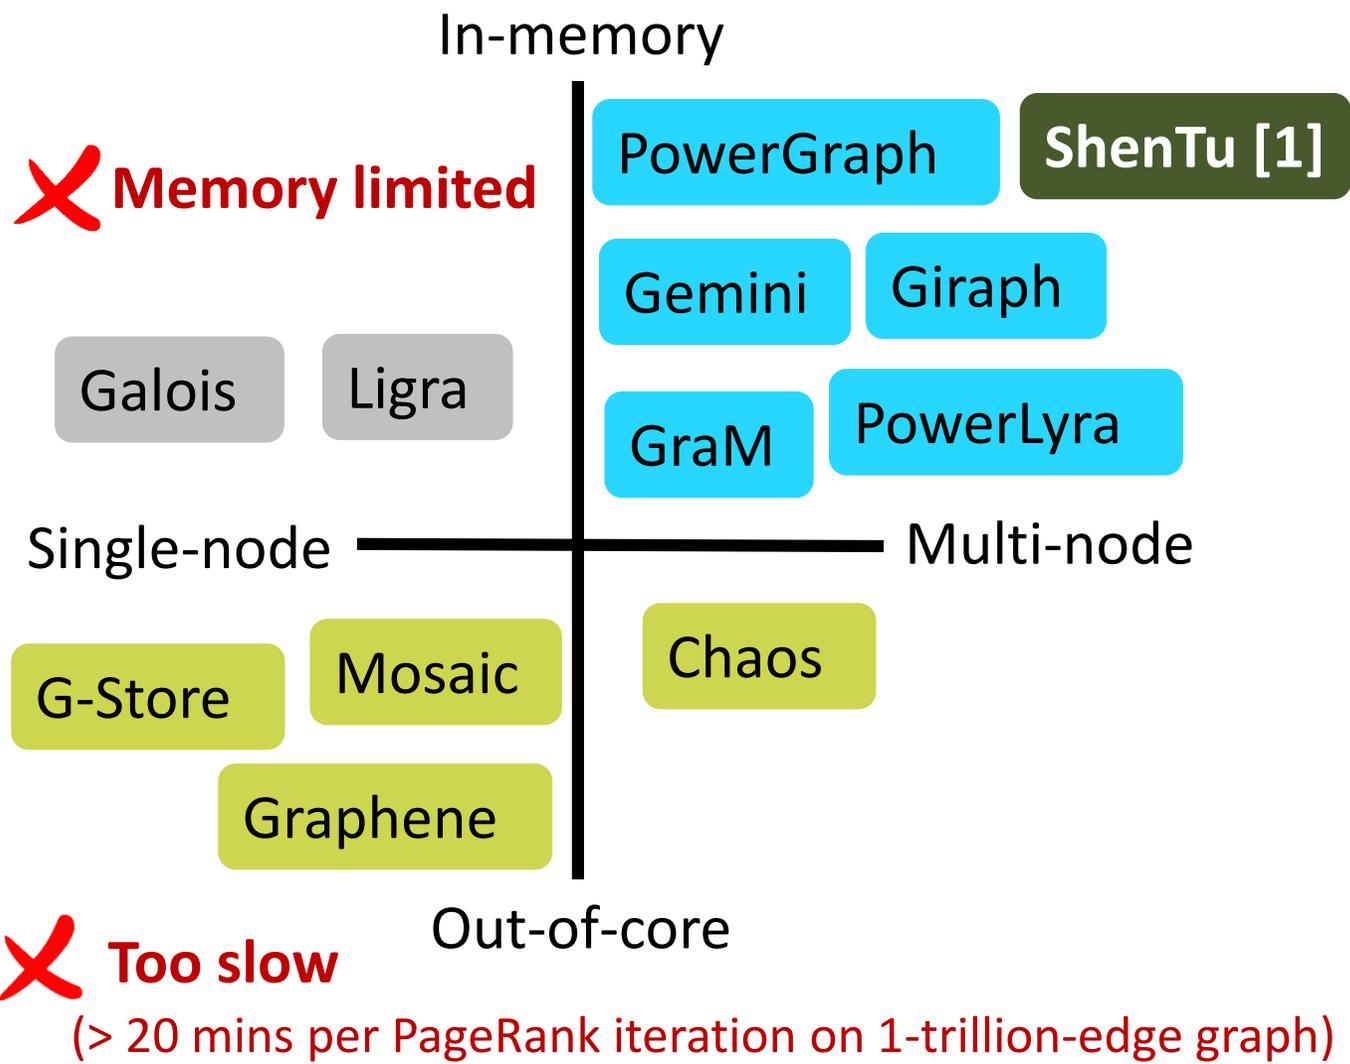
(> 20 mins per PageRank iteration on 1-trillion-edge graph)

Practice of Extreme-Scale Graph Processing



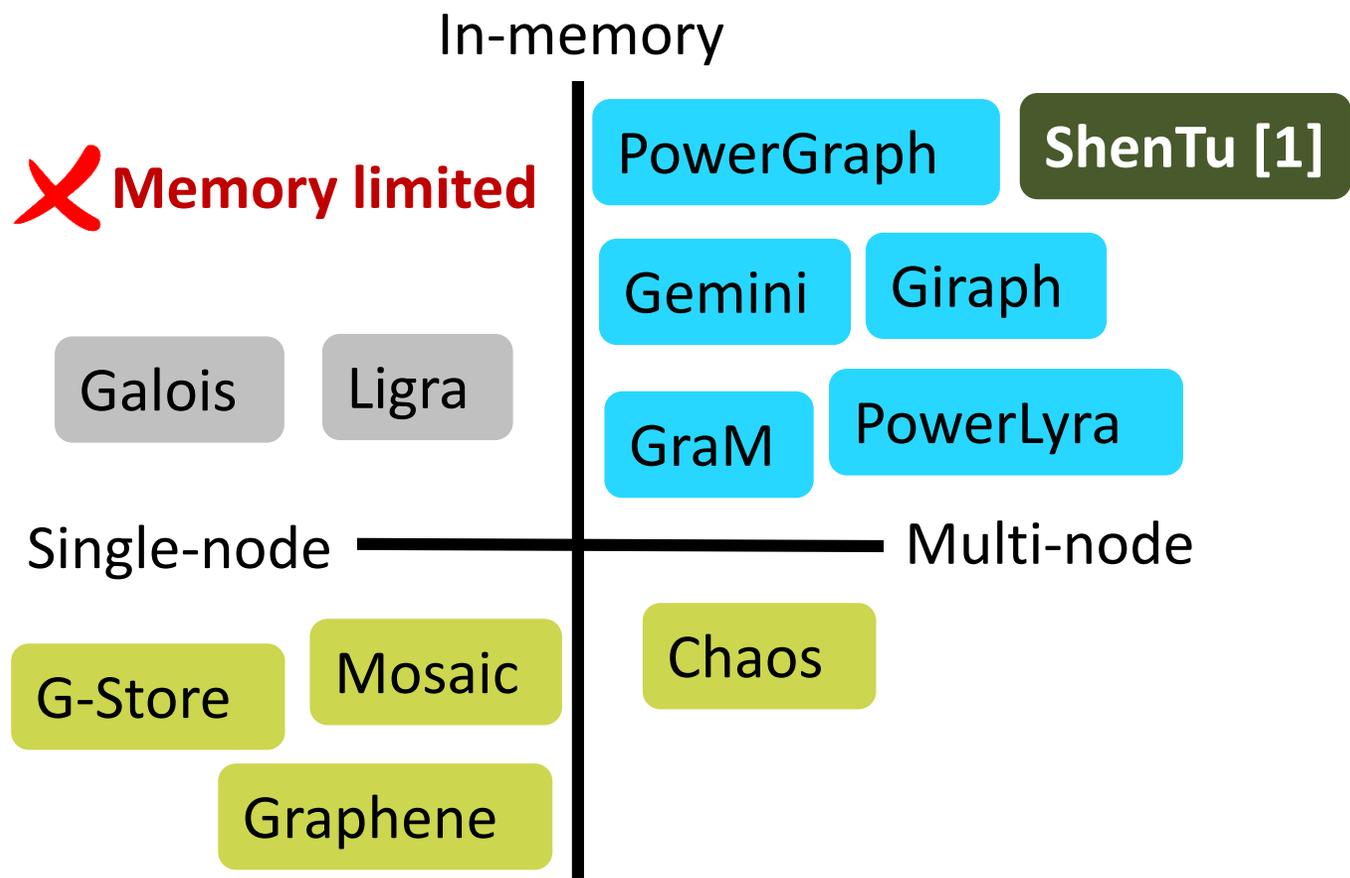
[1] Heng Lin et al.: ShenTu: Processing Multi-Trillion Edge Graphs on Millions of Cores in Seconds, SC18, Gordon Bell Finalist

Practice of Extreme-Scale Graph Processing

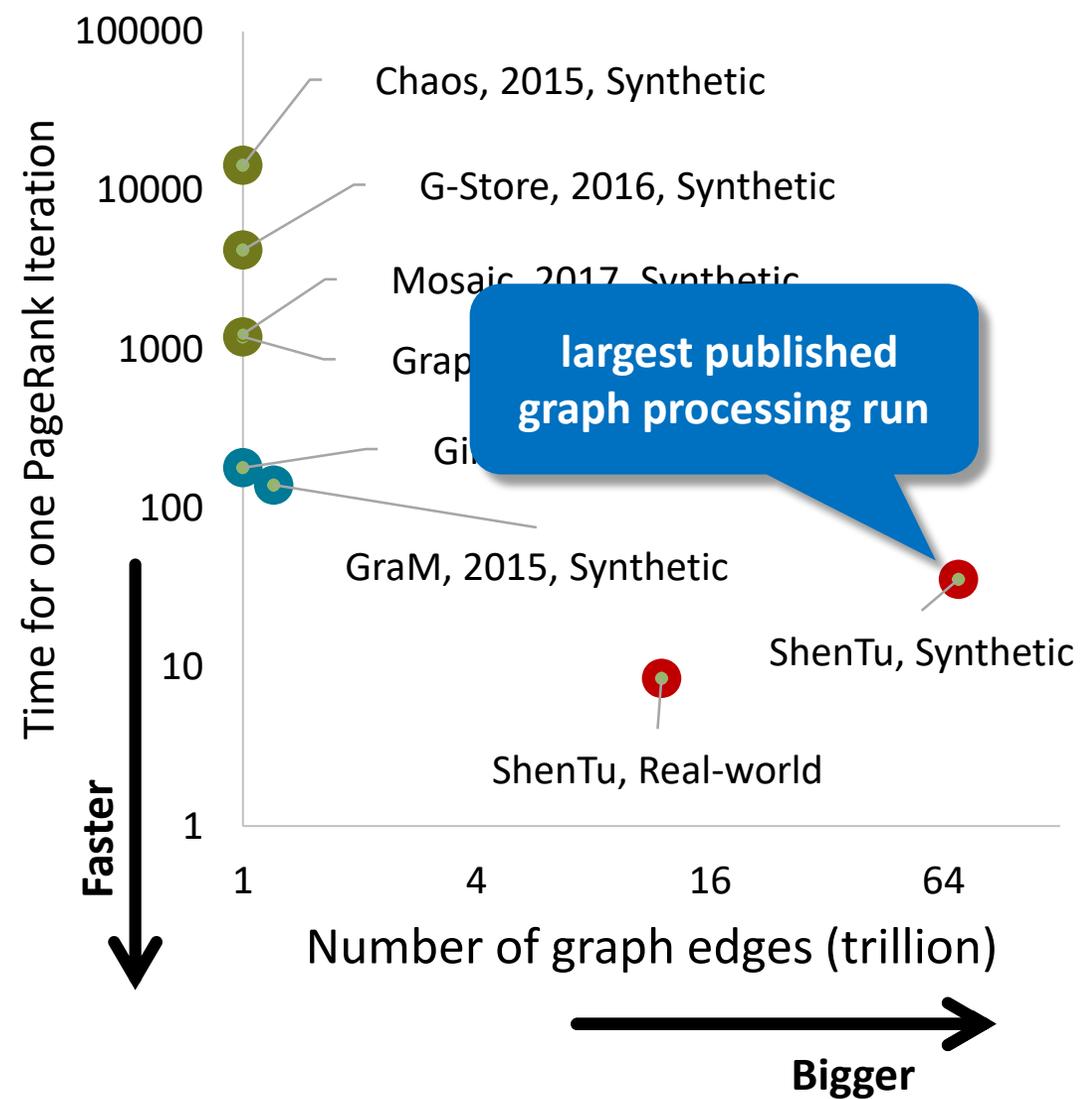


[1] Heng Lin et al.: ShenTu: Processing Multi-Trillion Edge Graphs on Millions of Cores in Seconds, SC18, Gordon Bell Finalist

Practice of Extreme-Scale Graph Processing



Too slow (> 20 mins per PageRank iteration on 1-trillion-edge graph)



[1] Heng Lin et al.: ShenTu: Processing Multi-Trillion Edge Graphs on Millions of Cores in Seconds, SC18, Gordon Bell Finalist

How large are extreme-scale graphs today?

How large are extreme-scale graphs today?

Webgraph datasets

Graph	Crawl date	Nodes	Arcs
uk-2014	2014	787 801 471	47 614 527 250
eu-2015	2015	1 070 557 254	91 792 261 600
gsh-2015	2015	988 490 691	33 877 399 152

Web data commons datasets

Granularity	#Nodes	#Arcs
Page	3,563 million	128,736 million
Host	101 million	2,043 million
Pay-Level-Domain	43 million	623 million

KONECT graph datasets

en	wikipedia edits (en)	<input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	50,757,442	572,591,272
TW	Twitter (WWW)	<input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	41,652,230	1,468,365,182
TF	Twitter (MPI)	<input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	52,579,682	1,963,263,821
FR	Friendster	<input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	68,349,466	2,586,147,869
UL	UK domain (2007)	<input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	105,153,952	3,301,876,564

How large are extreme-scale graphs today?

Largest Published Graph Computation [1]
 Gordon Bell Finalist 2018
 ShenTu on Sunway TaihuLight

Webgraph datasets

Graph	Crawl date	Nodes	Arcs
uk-2014	2014	787 801 471	47 614 527 250
eu-2015	2015	1 070 557 254	91 792 261 600
gsh-2015	2015	988 490 691	33 877 399 152

Web data commons datasets

Granularity	#Nodes	#Arcs
Page	3,563 million	128,736 million
Host	101 million	2,043 million
Pay-Level-Domain	43 million	623 million

KONECT graph datasets

en	wikipedia edits (en)	50,757,442	572,591,272
TW	Twitter (WWW)	41,652,230	1,468,365,182
TF	Twitter (MPI)	52,579,682	1,963,263,821
FR	Friendster	68,349,466	2,586,147,869
UL	UK domain (2007)	105,153,952	3,301,876,564

[1] Heng Lin et al.: ShenTu: Processing Multi-Trillion Edge Graphs on Millions of Cores in Seconds, SC18, Gordon Bell Finalist

How large are extreme-scale graphs today?

Sogou 搜狗 271 billion vertices,
12 trillion edges

Kronecker graph: 4.4 trillion vertices, 70 trillion edges

Webgraph datasets

Graph	Crawl date	Nodes	Arcs
uk-2014	2014	787 801 471	47 614 527 250
eu-2015	2015	1 070 557 254	91 792 261 600
gsh-2015	2015	988 490 691	33 877 399 152

Web data commons datasets

Granularity	#Nodes	#Arcs
Page	3,563 million	128,736 million
Host	101 million	2,043 million
Pay-Level-Domain	43 million	623 million

KONECT graph datasets

en	wikipedia edits (en)	50,757,442	572,591,272
TW	Twitter (WWW)	41,652,230	1,468,365,182
TF	Twitter (MPI)	52,579,682	1,963,263,821
FR	Friendster	68,349,466	2,586,147,869
UL	UK domain (2007)	105,153,952	3,301,876,564

Largest Published Graph Computation [1]
Gordon Bell Finalist 2018
 ShenTu on Sunway TaihuLight

[1] Heng Lin et al.: ShenTu: Processing Multi-Trillion Edge Graphs on Millions of Cores in Seconds, SC18, Gordon Bell Finalist

How large are extreme-scale graphs today?

Sogou 搜狗 271 billion vertices, 12 trillion edges

Kronecker graph: 4.4 trillion vertices, 70 trillion edges

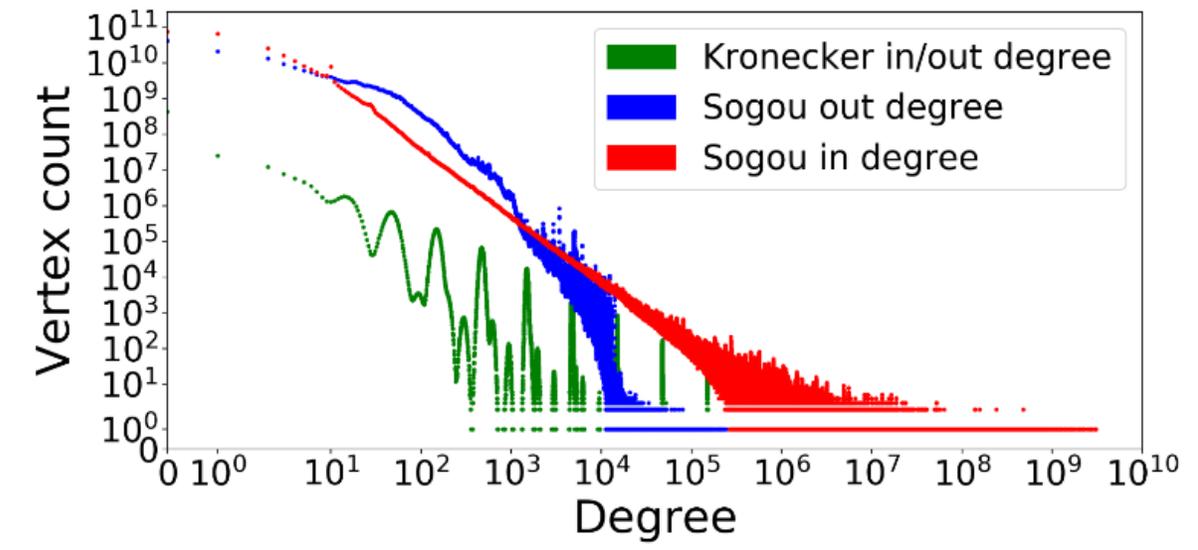
Webgraph datasets

Graph	Crawl date	Nodes	Arcs
uk-2014	2014	787 801 471	47 614 527 250
eu-2015	2015	1 070 557 254	91 792 261 600
gsh-2015	2015	988 490 691	33 877 399 152

KONECT graph datasets

en	wikipedia edits (en)	50,757,442	572,591,272
TW	Twitter (WWW)	41,652,230	1,468,365,182
TF	Twitter (MPI)	52,579,682	1,963,263,821
FR	Friendster	68,349,466	2,586,147,869
UL	UK domain (2007)	105,153,952	3,301,876,564

Largest Published Graph Computation [1]
Gordon Bell Finalist 2018
 ShenTu on Sunway TaihuLight



Web data commons datasets

Granularity	#Nodes	#Arcs
Page	3,563 million	128,736 million
Host	101 million	2,043 million
Pay-Level-Domain	43 million	623 million

[1] Heng Lin et al.: ShenTu: Processing Multi-Trillion Edge Graphs on Millions of Cores in Seconds, SC18, Gordon Bell Finalist

Sunway TaihuLight

Sunway TaihuLight



TaihuLight Top500 ranking: #3 (2018 Nov), #1 (2016, 2017)

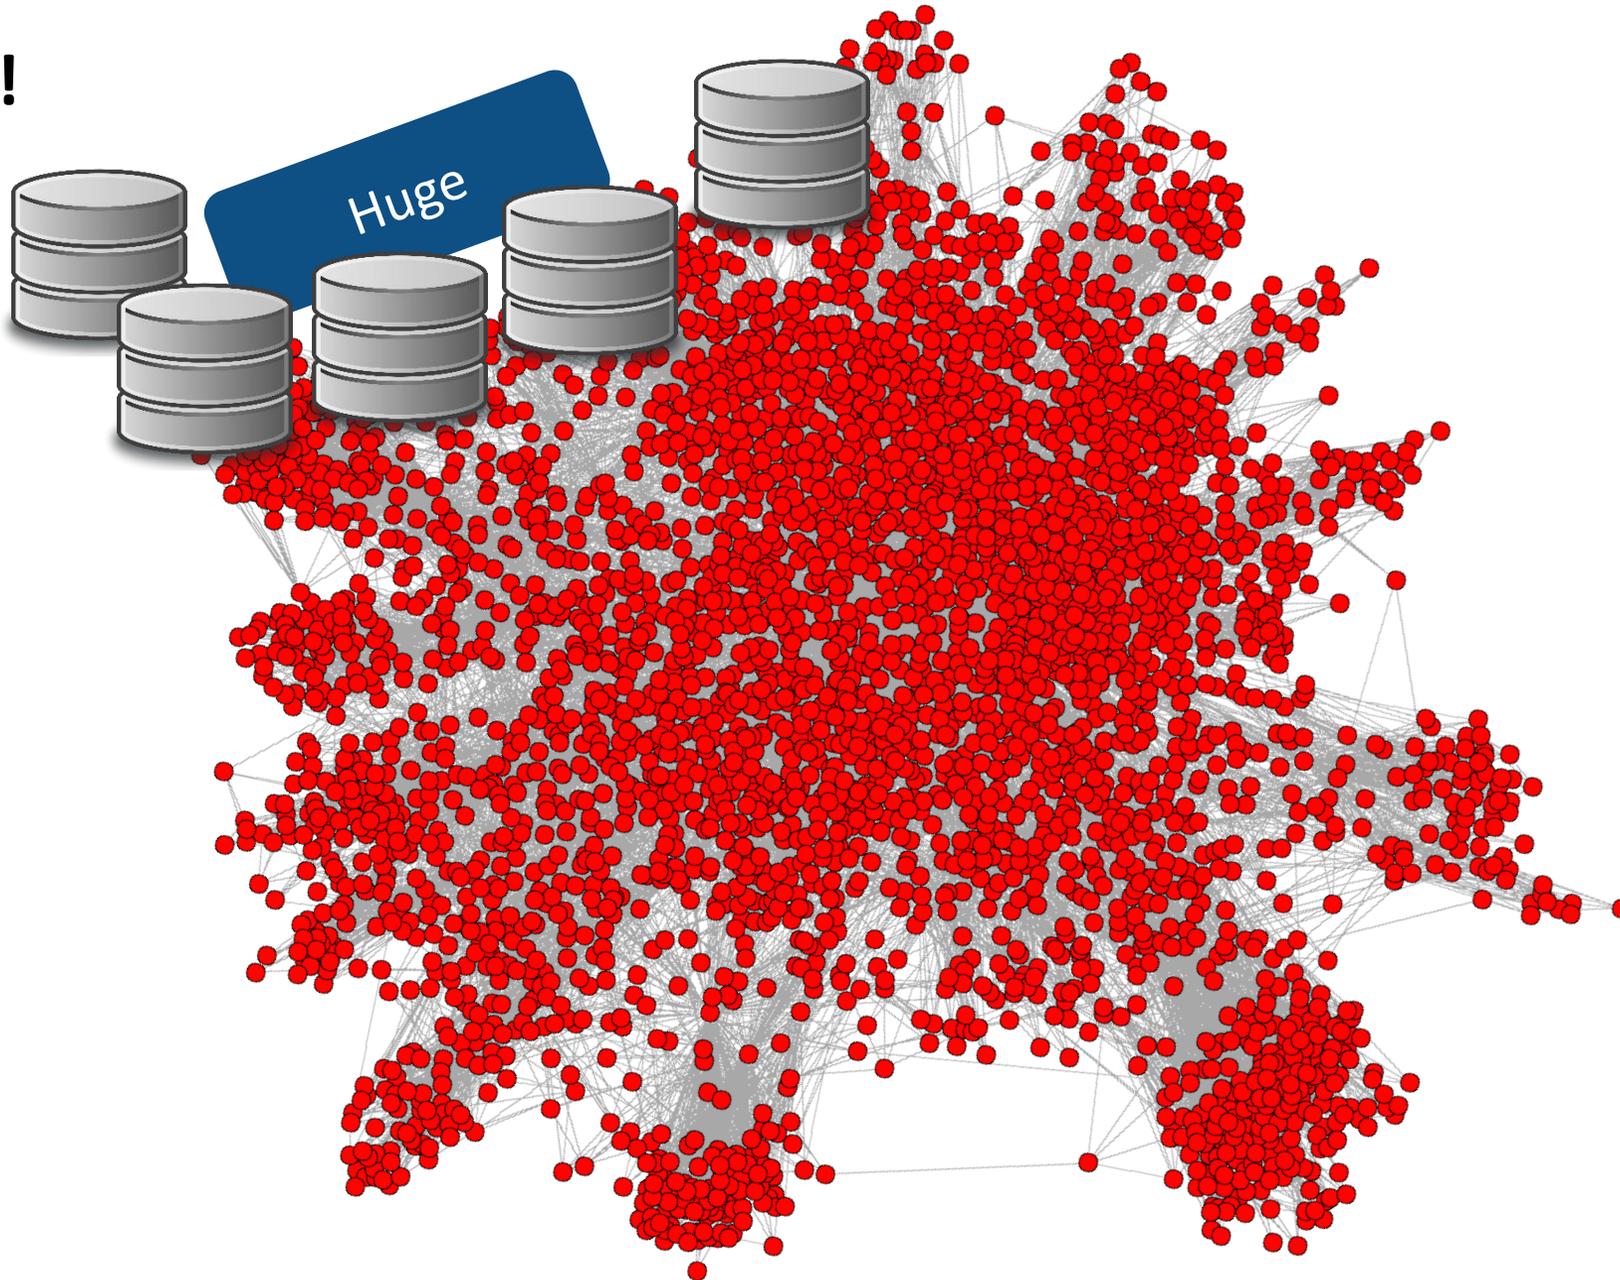
Sunway TaihuLight



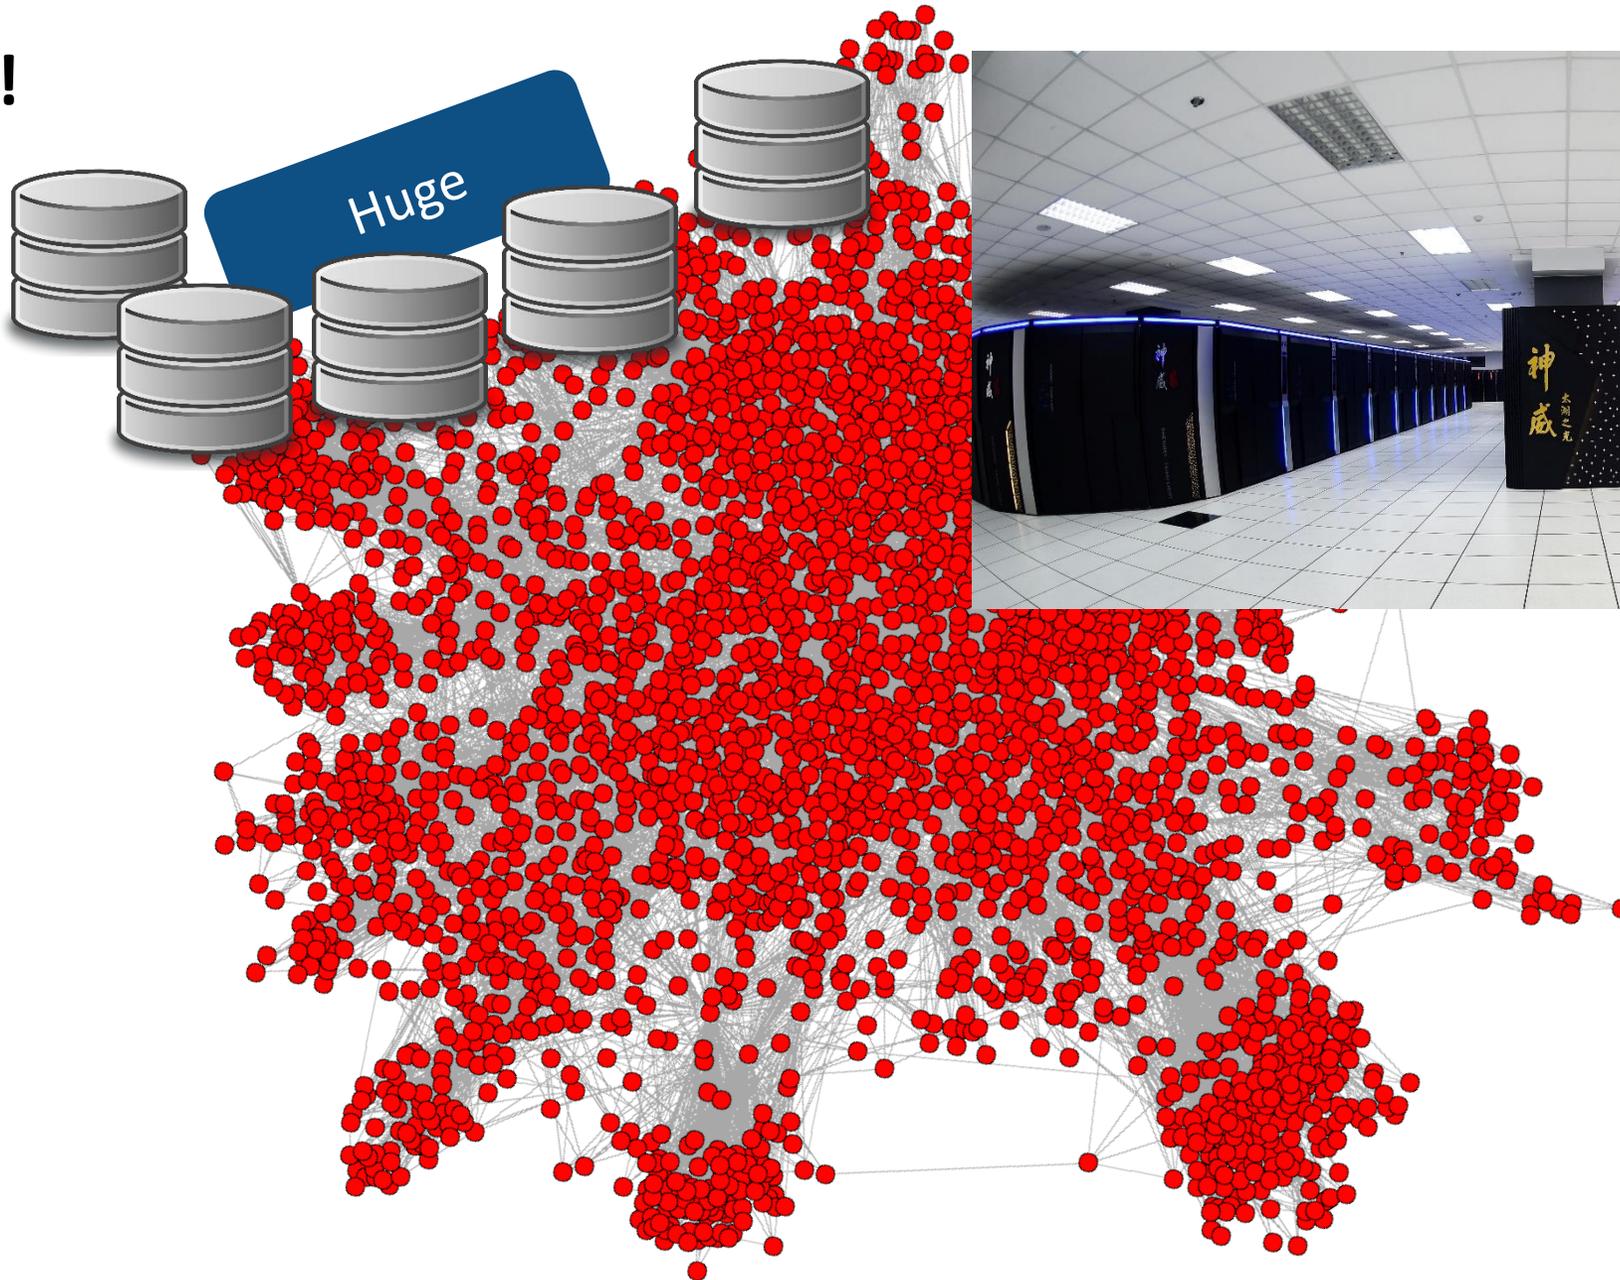
Nearly full
system run
(~41,000 nodes,
~1 million cores)

- PageRank iteration on **12 trillion edges in 8.5s** (1.4 TPEPS)
 - On 70 trillion edges, nearly **2 TPEPS** for PageRank and WCC
 - 774 GPEPS for BFS
- (PEPS = processes edges per second as opposed to TEPS)**

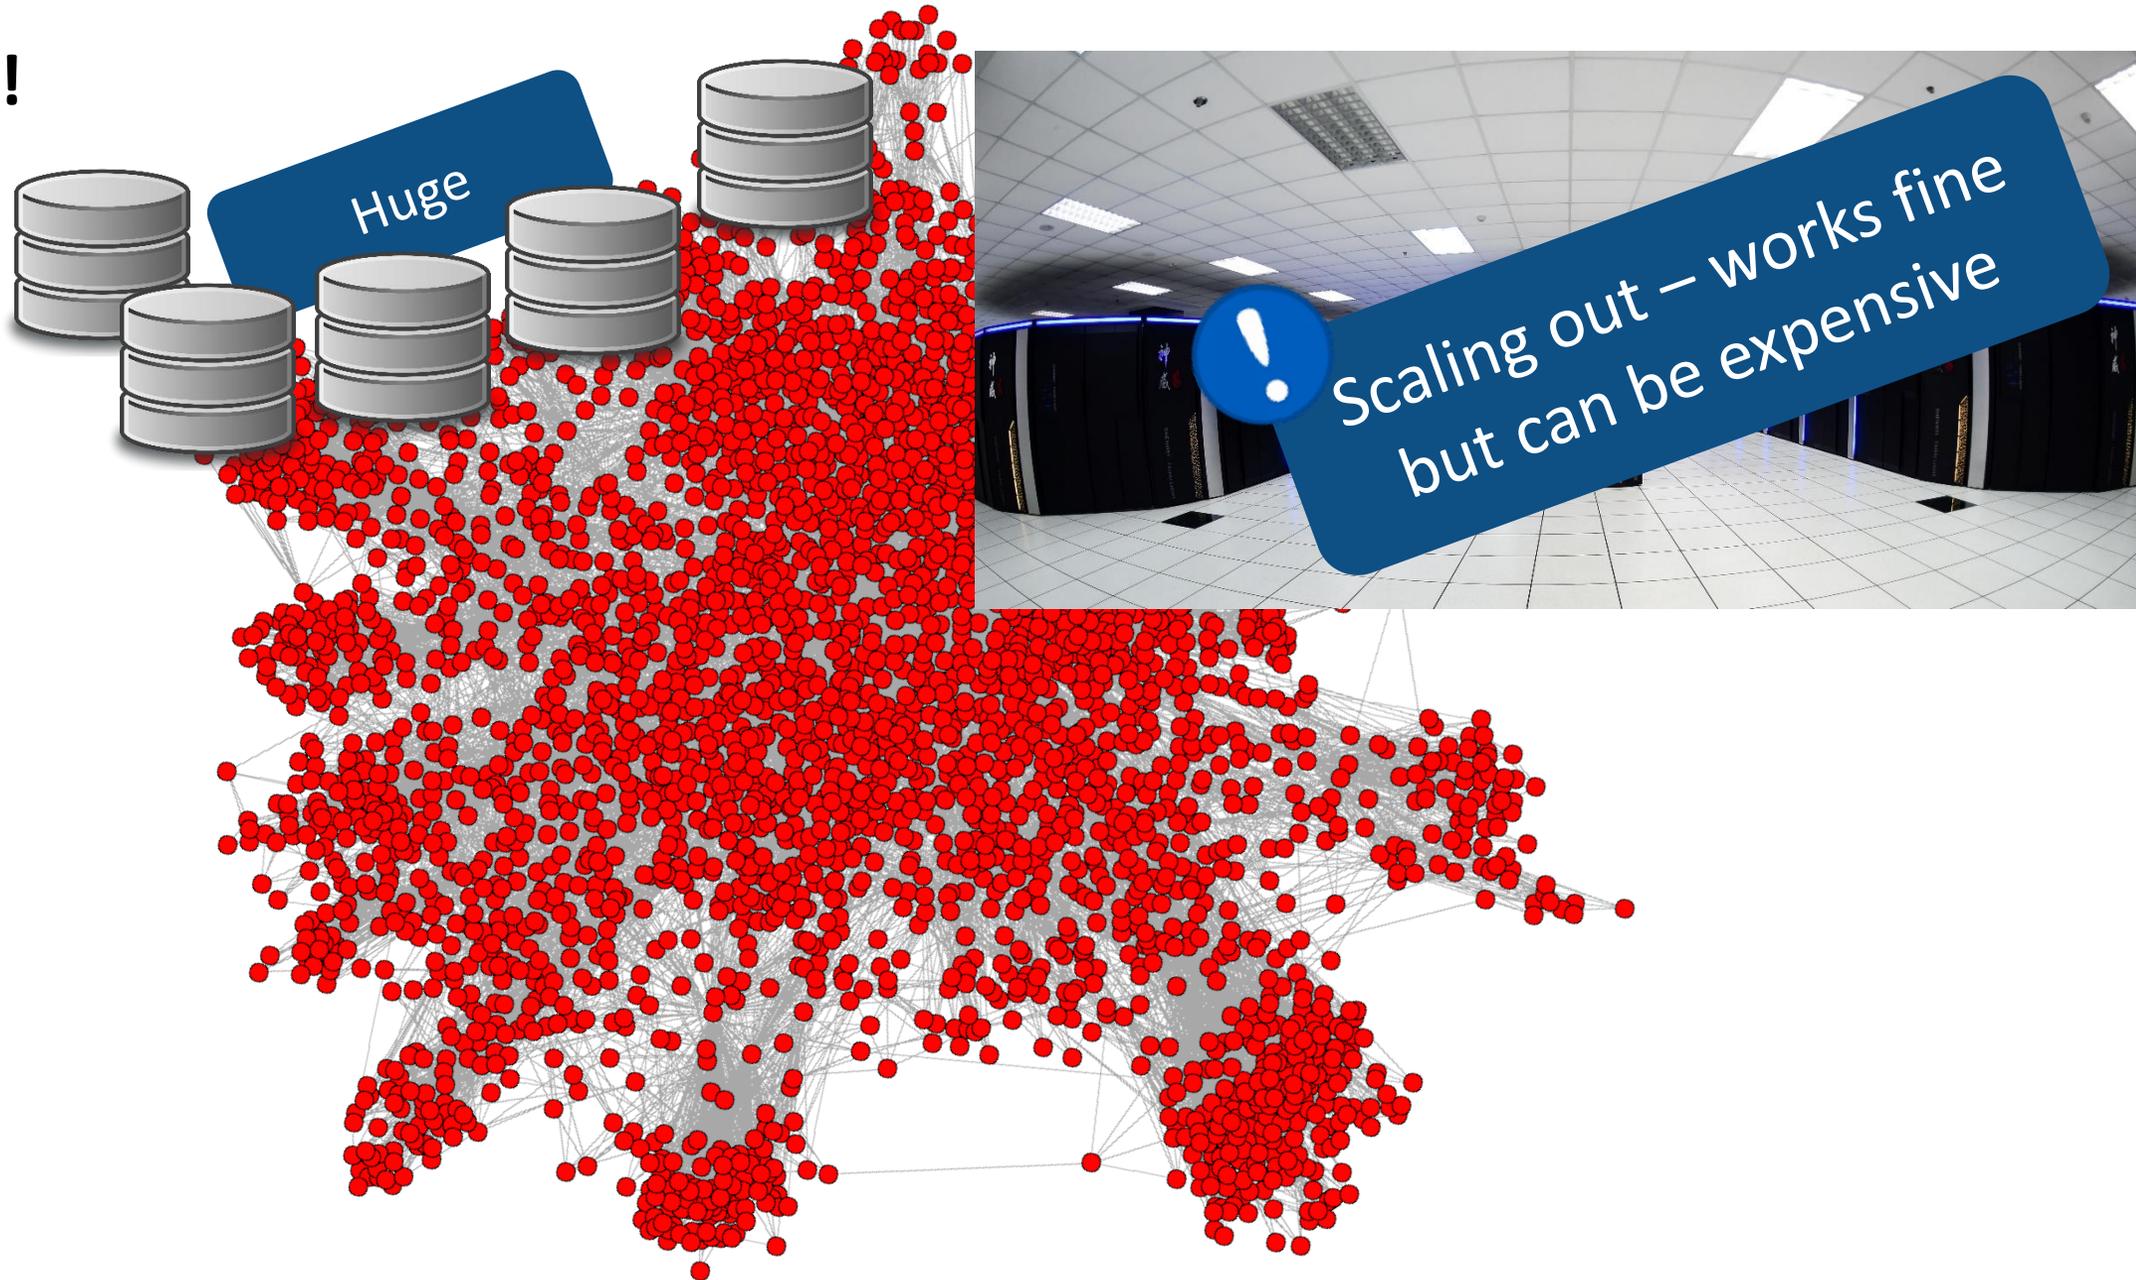
Problems!



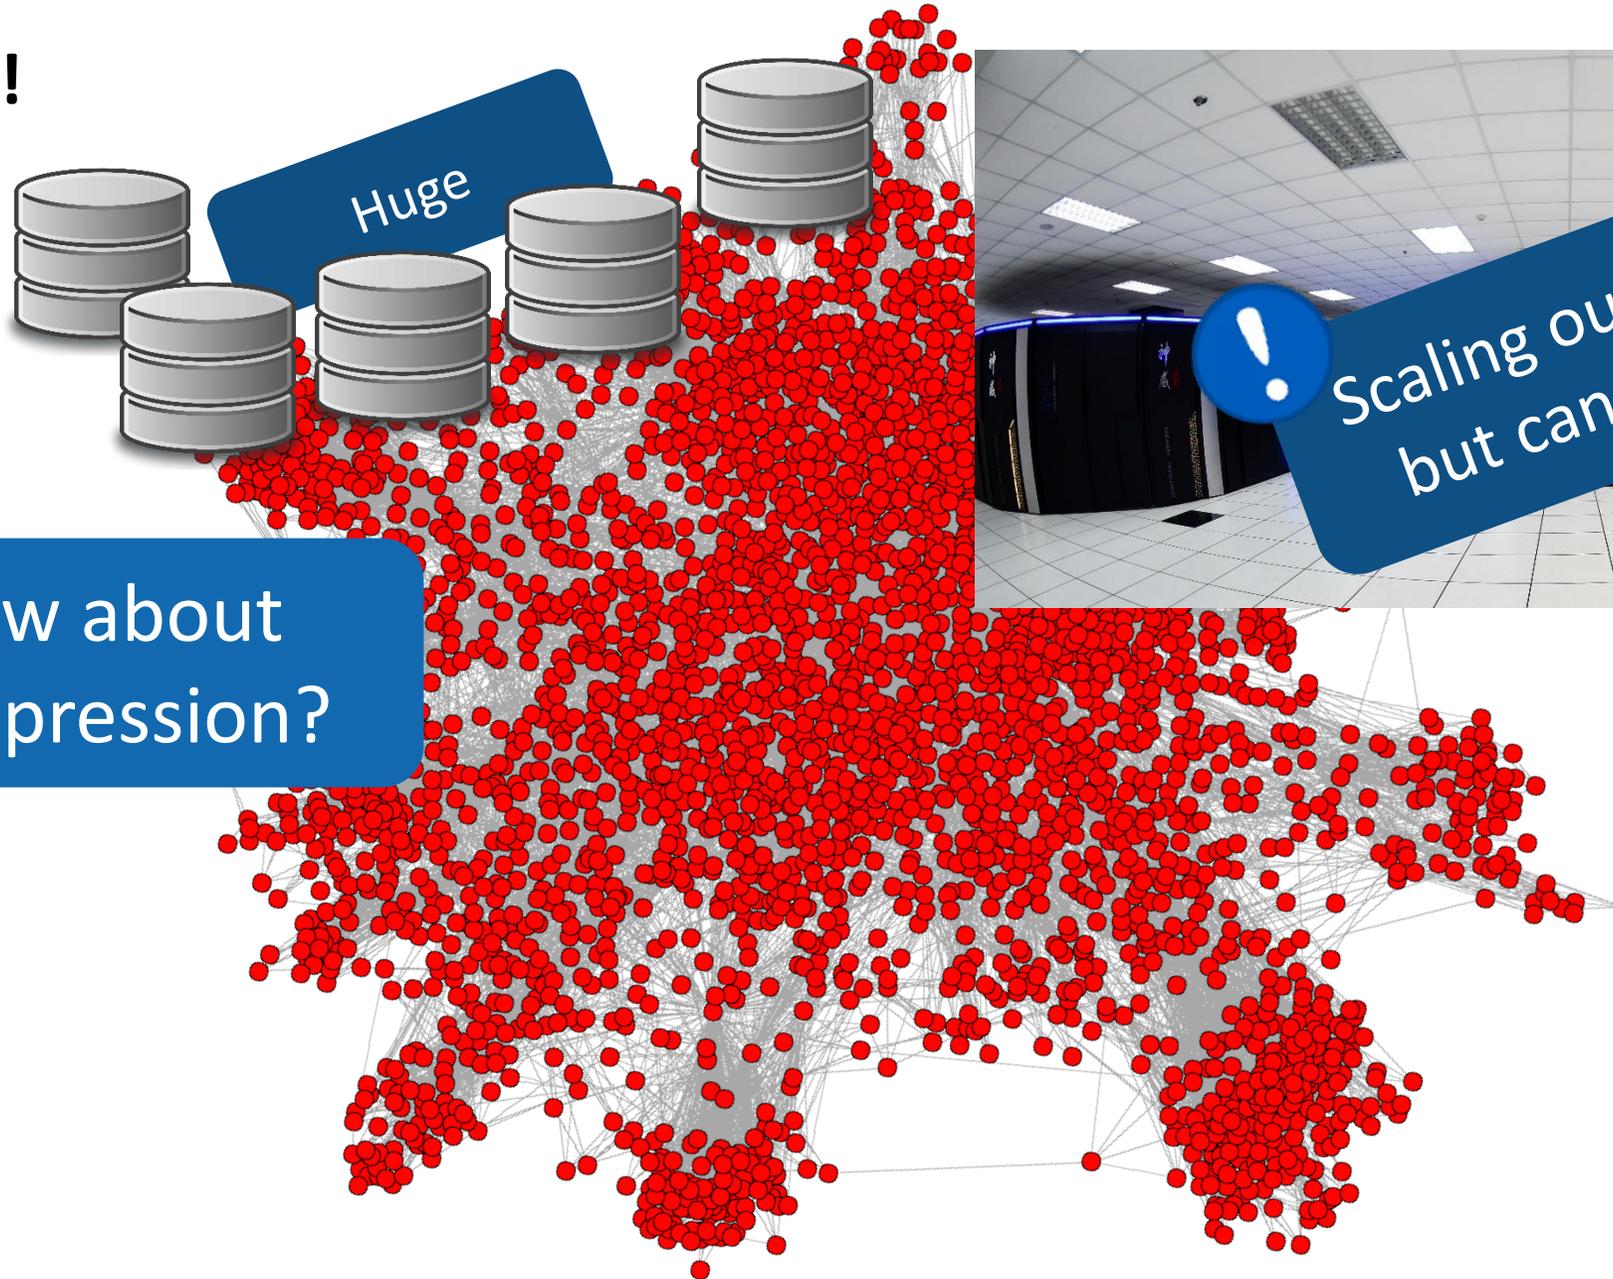
Problems!



Problems!



Problems!



Problems!



[Traditional] Compression
incurs expensive
decompression



Problems!



! Scaling out – works fine but can be expensive

? How about compression?

[Traditional] Compression incurs expensive decompression



Log(Graph): A Near-Optimal High-Performance Graph Representation

Maciej Besta[†], Dimitri Stanojevic, Tijana Zivic, Jagpreet Singh, Maurice Hoerold, Torsten Hoefler[†]
Department of Computer Science, ETH Zurich
[†]Corresponding authors (maciej.best@inf.ethz.ch, htor@inf.ethz.ch)

ABSTRACT

Today’s graphs used in domains such as machine learning or social network analysis may contain hundreds of billions of edges. Yet, they are not necessarily stored efficiently, and standard graph representations such as adjacency lists waste a significant number of bits while graph compression schemes

discovering relationships in graph data. The sheer size of such graphs, up to hundreds of billions of edges, exacerbates the number of needed memory banks, increases the amount of data transferred between CPUs and memory, and may lead to I/O accesses while processing graphs. Thus, reducing the size of such graphs is becoming increasingly important.

Problems!



! Scaling out – works fine but can be expensive

? How about compression?

[Traditional] Compression incurs expensive decompression



@ ACM PACT'18

Log(Graph): A Near-Optimal High-Performance Graph Representation

Maciej Besta[†], Dimitri Stanojevic, Tijana Zivic, Jagpreet Singh, Maurice Hoerold, Torsten Hoefler[†]
 Department of Computer Science, ETH Zurich
[†]Corresponding authors (maciej.besta@inf.ethz.ch, thor@inf.ethz.ch)

! Log(Graph): effective compression with low-overhead decompression!

ABSTRACT
 Today's graphs... social networks... edges. Yet... dard graph... significant number of bits while graph compression schemes... size of such graphs is becoming increasingly important.



What is the **lowest storage** we can
(hope to) use to store a graph?

What is the **lowest storage** we can
(hope to) use to store a graph?

! The storage
lower bound Ω

What is the **lowest storage** we can
(hope to) use to store a graph?

! The storage
lower bound Ω

Which one? 😊

What is the **lowest storage** we can
(hope to) use to store a graph?

The storage
lower bound Ω

Which one? 😊

Shannon's approach
logarithmic
(one needs at least $\log |S|$
bits to store an object
from an arbitrary set S)

What is the **lowest storage** we can
(hope to) use to store a graph?

The storage
lower bound Ω

Which one? 😊

Shannon's approach
logarithmic
(one needs at least $\log |S|$
bits to store an object
from an arbitrary set S)

$$S = \{x_1, x_2, x_3, \dots\}$$

x_1	\rightarrow	0 ... 01
x_2	\rightarrow	0 ... 10
x_3	\rightarrow	0 ... 11
		...

What is the **lowest storage** we can (hope to) use to store a graph?

The storage **lower bound** Ω

Which one? 😊

Shannon's approach **logarithmic**
(one needs at least $\log |S|$ bits to store an object from an arbitrary set S)

 Key idea

$$S = \{x_1, x_2, x_3, \dots\}$$

x_1	\rightarrow	0 ... 01
x_2	\rightarrow	0 ... 10
x_3	\rightarrow	0 ... 11
		...

What is the **lowest storage** we can (hope to) use to store a graph?

The storage **lower bound** Ω

Which one? 😊

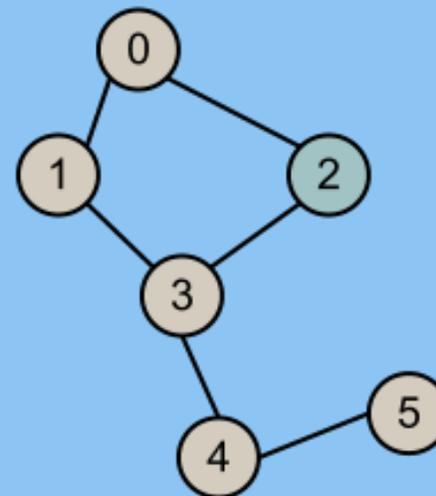
Shannon's approach **logarithmic**
(one needs at least $\log |S|$ bits to store an object from an arbitrary set S)

 Key idea

! $S = \{x_1, x_2, x_3, \dots\}$

$x_1 \rightarrow 0 \dots 01$
$x_2 \rightarrow 0 \dots 10$
$x_3 \rightarrow 0 \dots 11$
...

Encode different parts of a graph representation using (logarithmic) **storage lower bounds**



What is the **lowest storage** we can (hope to) use to store a graph?

The storage **lower bound** Ω

Which one? 😊

Shannon's approach **logarithmic**
(one needs at least $\log |S|$ bits to store an object from an arbitrary set S)

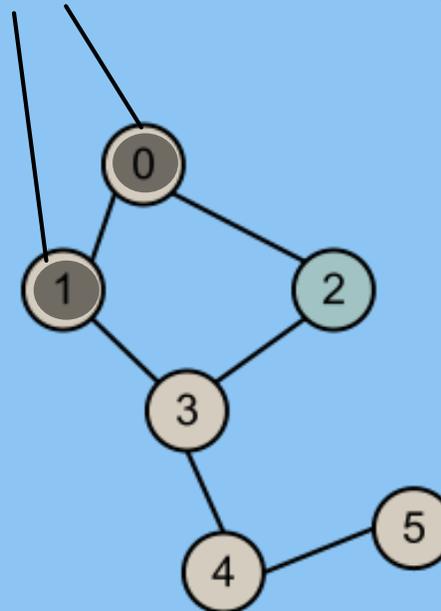
$$S = \{x_1, x_2, x_3, \dots\}$$

$x_1 \rightarrow$	0 ... 01
$x_2 \rightarrow$	0 ... 10
$x_3 \rightarrow$	0 ... 11
	...

 Key idea

Encode different parts of a graph representation using (logarithmic) **storage lower bounds**

Vertex labels



What is the **lowest storage** we can (hope to) use to store a graph?

The storage **lower bound** Ω

Which one? 😊

Shannon's approach **logarithmic**
(one needs at least $\log |S|$ bits to store an object from an arbitrary set S)

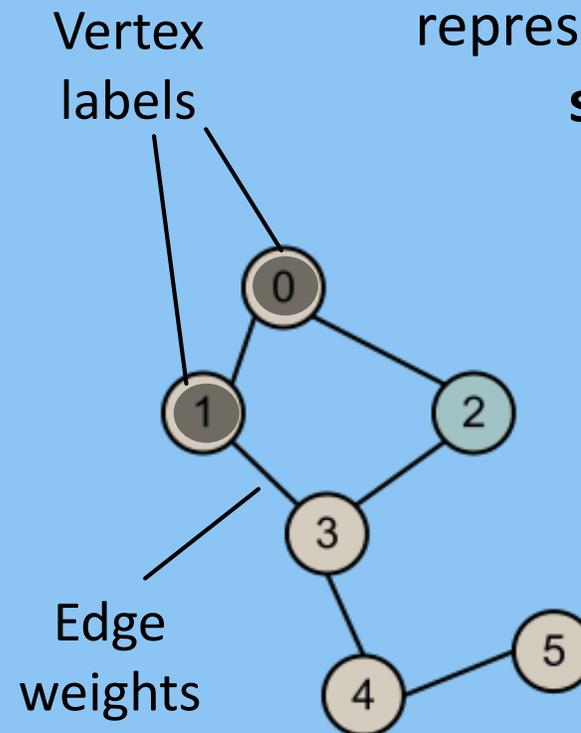
$$S = \{x_1, x_2, x_3, \dots\}$$

$x_1 \rightarrow$	0 ... 01
$x_2 \rightarrow$	0 ... 10
$x_3 \rightarrow$	0 ... 11
	...



Key idea

Encode different parts of a graph representation using (logarithmic) **storage lower bounds**



What is the **lowest storage** we can (hope to) use to store a graph?

The storage **lower bound** Ω

Which one? 😊

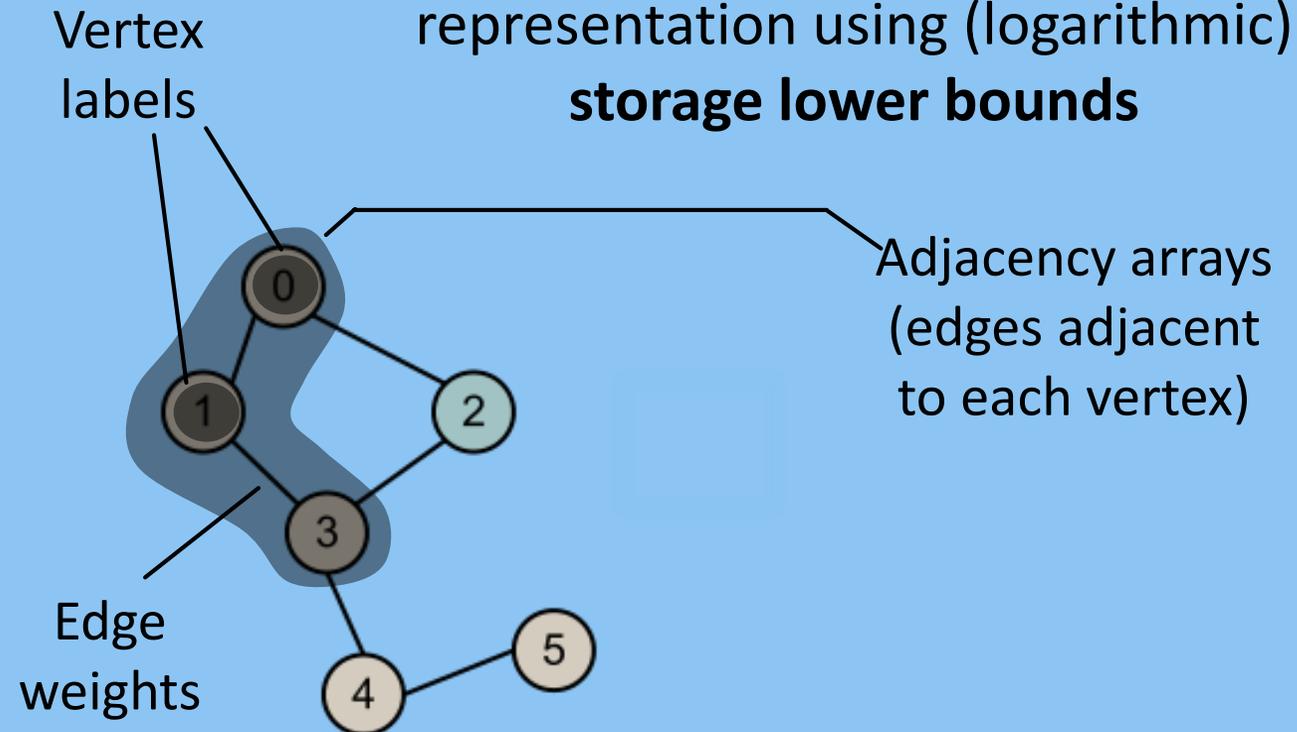
Shannon's approach **logarithmic**
(one needs at least $\log |S|$ bits to store an object from an arbitrary set S)

$$S = \{x_1, x_2, x_3, \dots\}$$

$x_1 \rightarrow 0 \dots 01$
$x_2 \rightarrow 0 \dots 10$
$x_3 \rightarrow 0 \dots 11$
...

Key idea

Encode different parts of a graph representation using (logarithmic) **storage lower bounds**



What is the **lowest storage** we can (hope to) use to store a graph?

The storage **lower bound** Ω

Which one? 😊

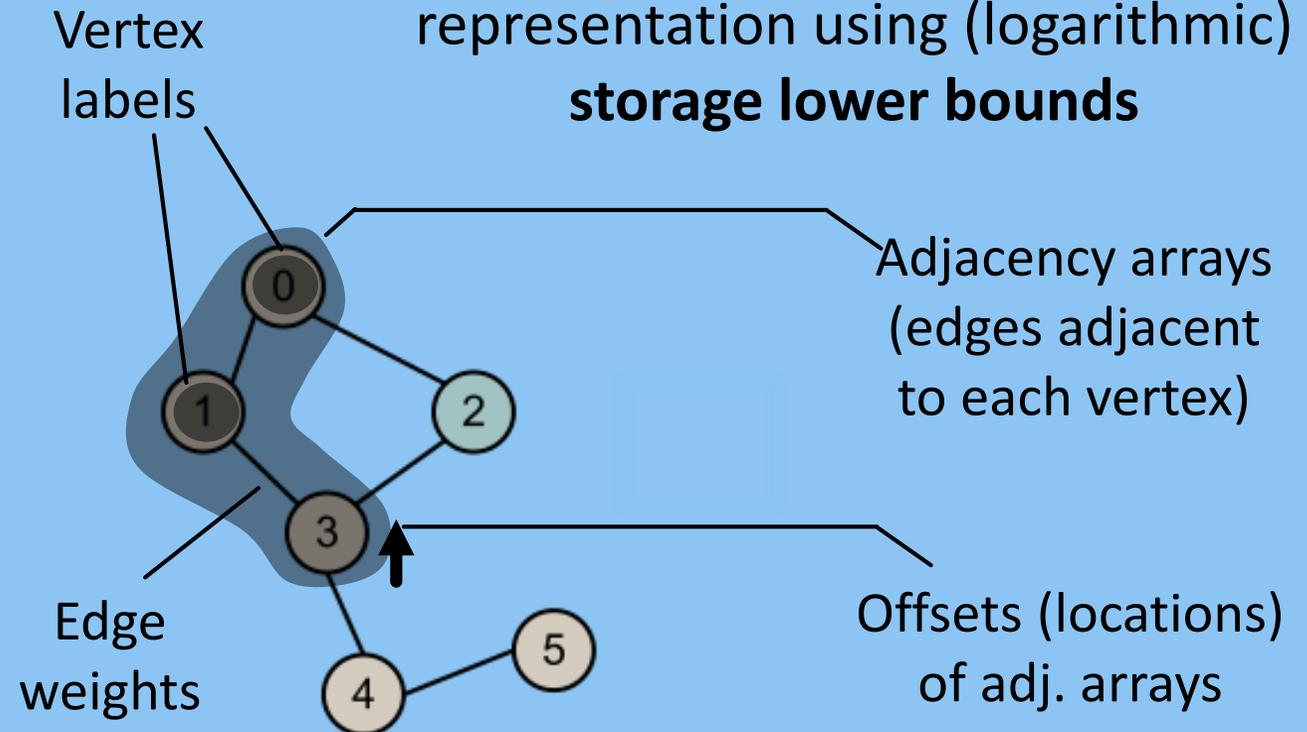
Shannon's approach **logarithmic**
(one needs at least $\log |S|$ bits to store an object from an arbitrary set S)

$$S = \{x_1, x_2, x_3, \dots\}$$

$x_1 \rightarrow 0 \dots 01$
$x_2 \rightarrow 0 \dots 10$
$x_3 \rightarrow 0 \dots 11$
...

Key idea

Encode different parts of a graph representation using (logarithmic) **storage lower bounds**



What is the **lowest storage** we can (hope to) use to store a graph?

The storage **lower bound** Ω

Which one? 😊

Shannon's approach **logarithmic**
(one needs at least $\log |S|$ bits to store an object from an arbitrary set S)

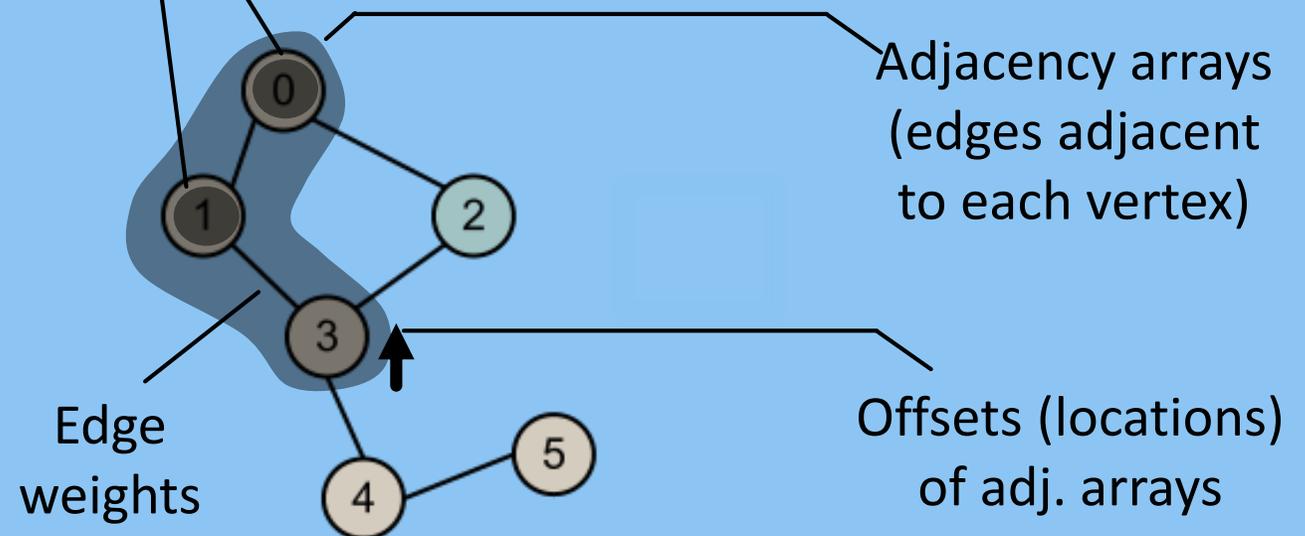
$$S = \{x_1, x_2, x_3, \dots\}$$

$$\begin{aligned} x_1 &\rightarrow 0 \dots 01 \\ x_2 &\rightarrow 0 \dots 10 \\ x_3 &\rightarrow 0 \dots 11 \\ &\dots \end{aligned}$$

Key idea

Log (Vertex labels)

Encode different parts of a graph representation using (logarithmic) **storage lower bounds**



What is the **lowest storage** we can (hope to) use to store a graph?

The storage **lower bound** Ω

Which one? 😊

Shannon's approach **logarithmic**
(one needs at least $\log |S|$ bits to store an object from an arbitrary set S)

$$S = \{x_1, x_2, x_3, \dots\}$$

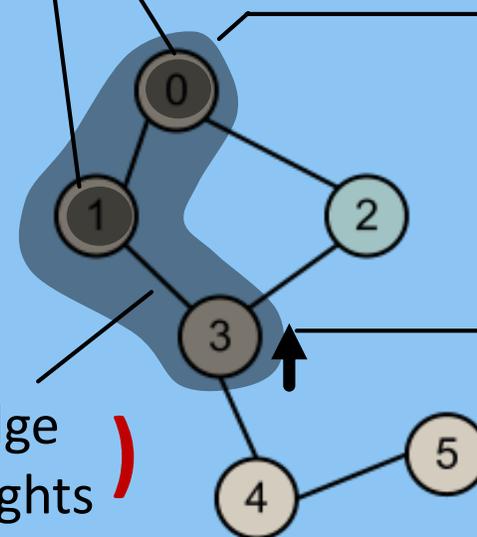
$x_1 \rightarrow 0 \dots 01$
$x_2 \rightarrow 0 \dots 10$
$x_3 \rightarrow 0 \dots 11$
...

Key idea

Log (Vertex labels)

Log (Edge weights)

Encode different parts of a graph representation using (logarithmic) **storage lower bounds**



Adjacency arrays (edges adjacent to each vertex)

Offsets (locations) of adj. arrays

What is the **lowest storage** we can (hope to) use to store a graph?

The storage **lower bound** Ω

Which one? 😊

Shannon's approach **logarithmic**
(one needs at least $\log |S|$ bits to store an object from an arbitrary set S)

$$S = \{x_1, x_2, x_3, \dots\}$$

$x_1 \rightarrow 0 \dots 01$
$x_2 \rightarrow 0 \dots 10$
$x_3 \rightarrow 0 \dots 11$
...

Key idea

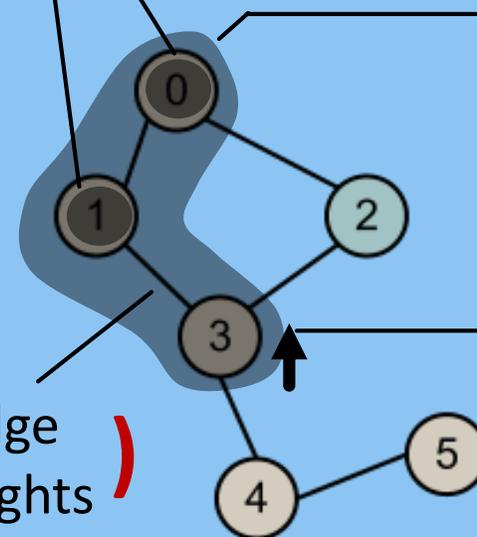
Encode different parts of a graph representation using (logarithmic) **storage lower bounds**

Log (Vertex labels)

Log ((edges adjacent to each vertex))

Log (Edge weights)

Offsets (locations) of adj. arrays



What is the **lowest storage** we can (hope to) use to store a graph?

The storage **lower bound** Ω

Which one? 😊

Shannon's approach **logarithmic**
(one needs at least $\log |S|$ bits to store an object from an arbitrary set S)

$$S = \{x_1, x_2, x_3, \dots\}$$

$x_1 \rightarrow$	0 ... 01
$x_2 \rightarrow$	0 ... 10
$x_3 \rightarrow$	0 ... 11
	...

Key idea

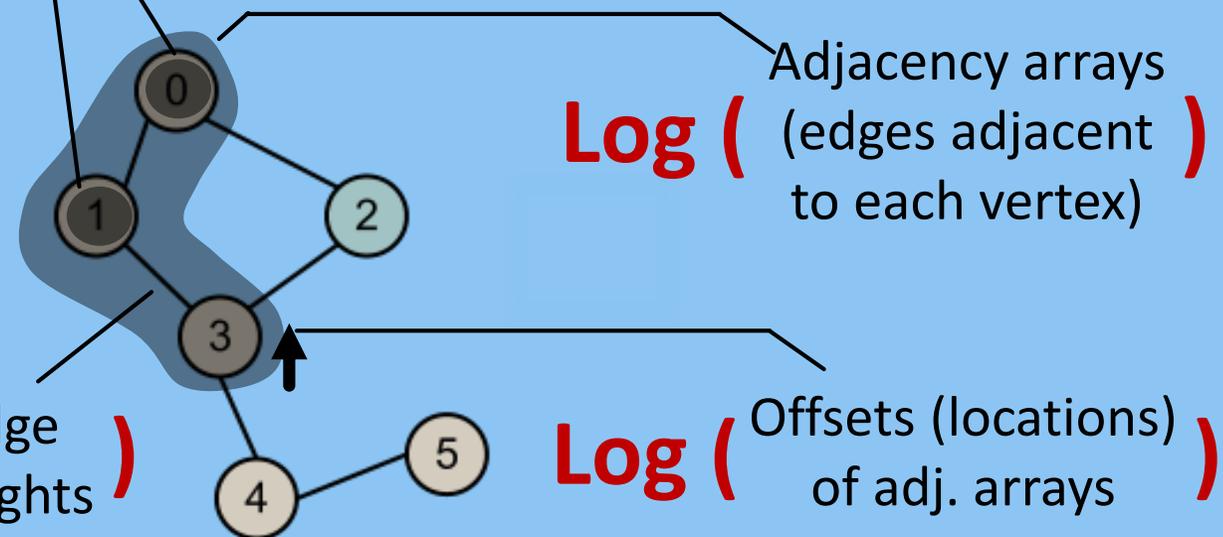
Encode different parts of a graph representation using (logarithmic) **storage lower bounds**

Log (Vertex labels)

Log (Edge weights)

Log ((edges adjacent to each vertex))

Log (Offsets (locations) of adj. arrays)



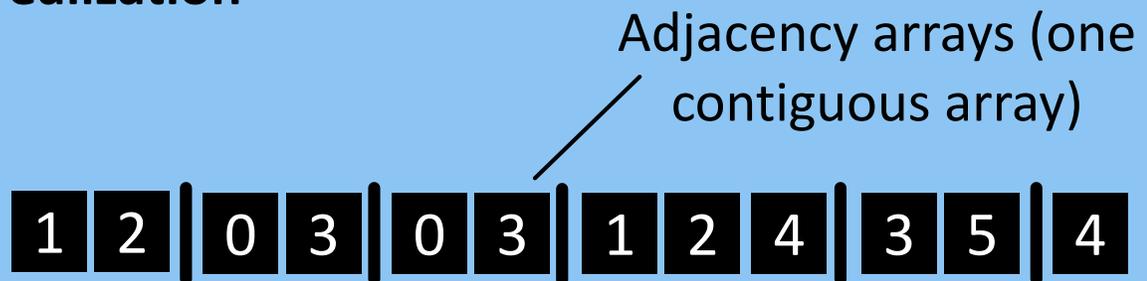
ADJACENCY ARRAY GRAPH REPRESENTATION

ADJACENCY ARRAY GRAPH REPRESENTATION

Physical realization

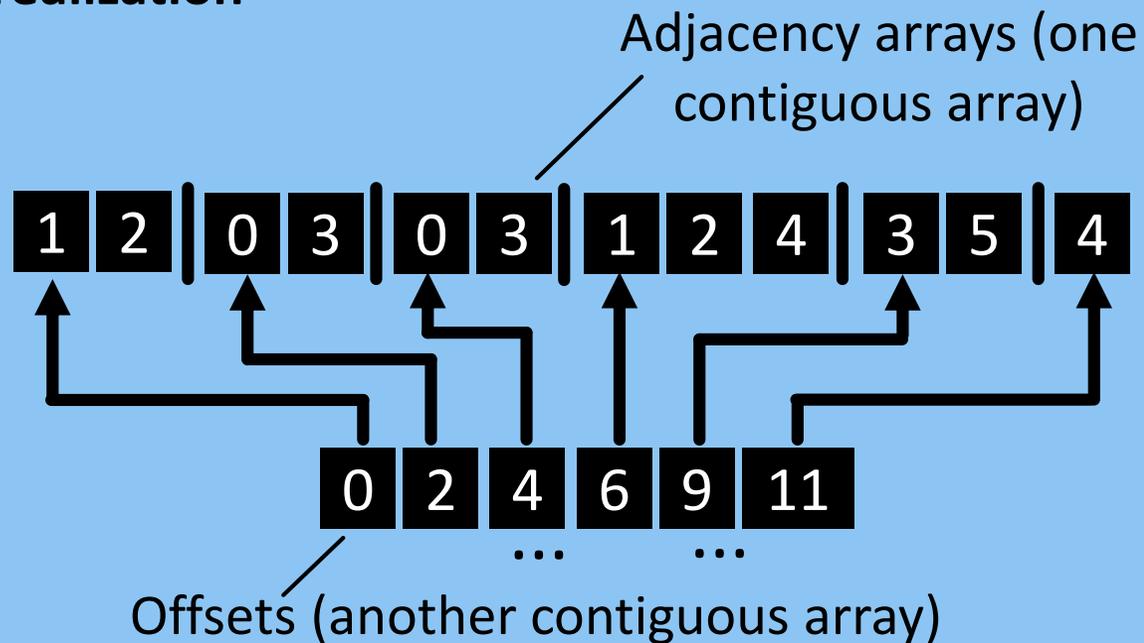
ADJACENCY ARRAY GRAPH REPRESENTATION

Physical realization



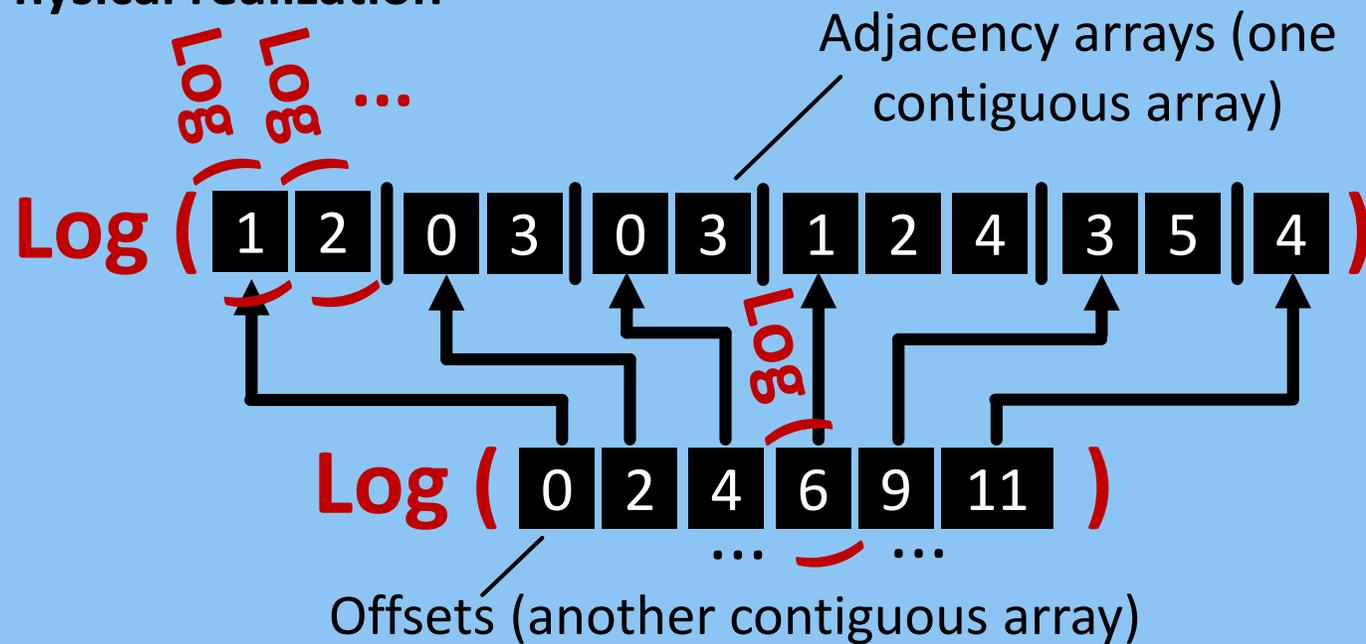
ADJACENCY ARRAY GRAPH REPRESENTATION

Physical realization



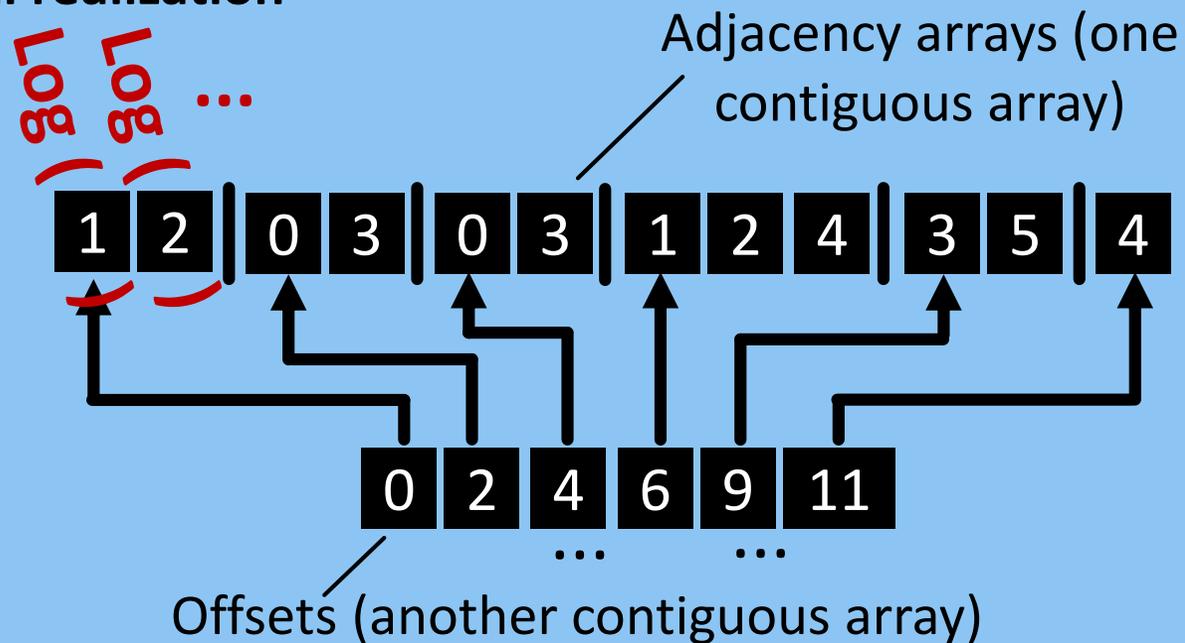
ADJACENCY ARRAY GRAPH REPRESENTATION

Physical realization



ADJACENCY ARRAY GRAPH REPRESENTATION

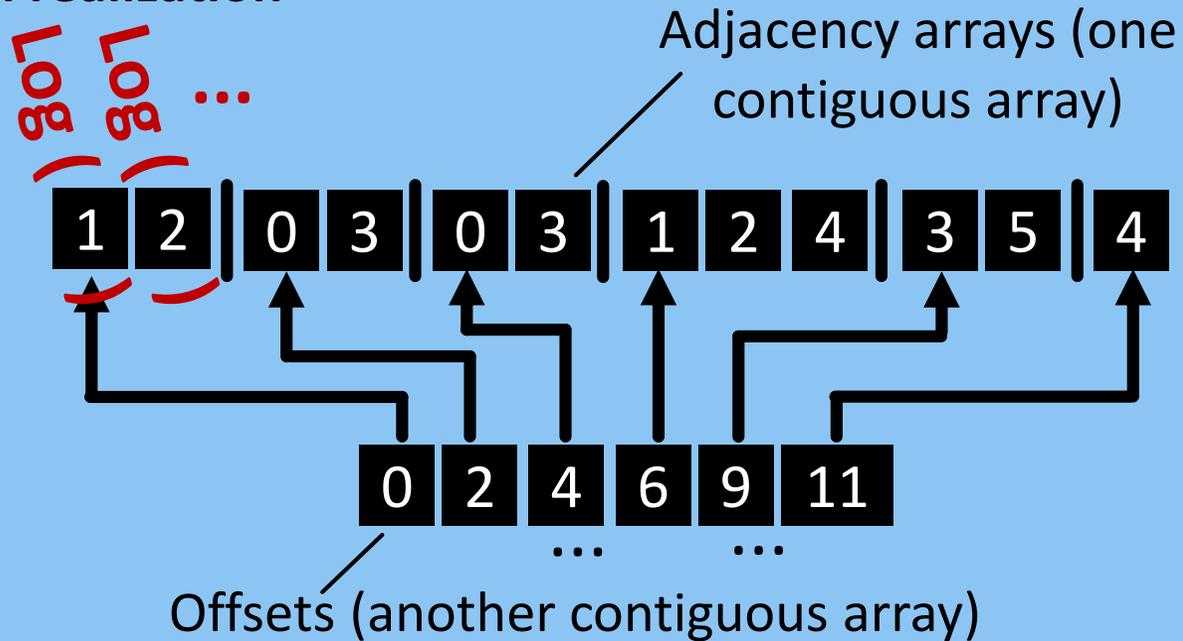
Physical realization



ADJACENCY ARRAY GRAPH REPRESENTATION

1 $\text{Log}(\text{Vertex labels}), \text{Log}(\text{Edge weights})$

Physical realization

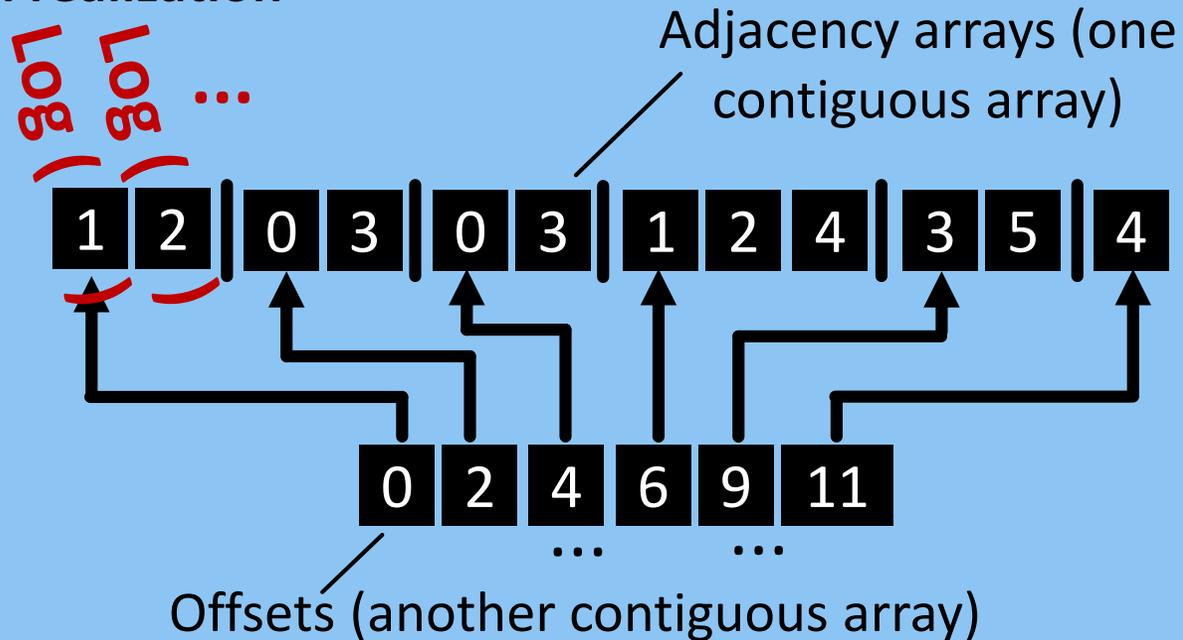


ADJACENCY ARRAY GRAPH REPRESENTATION

1 $\text{Log}(\text{Vertex labels})$, $\text{Log}(\text{Edge weights})$

Global bound $\lceil \log n \rceil$

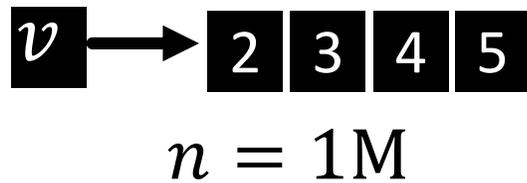
Physical realization



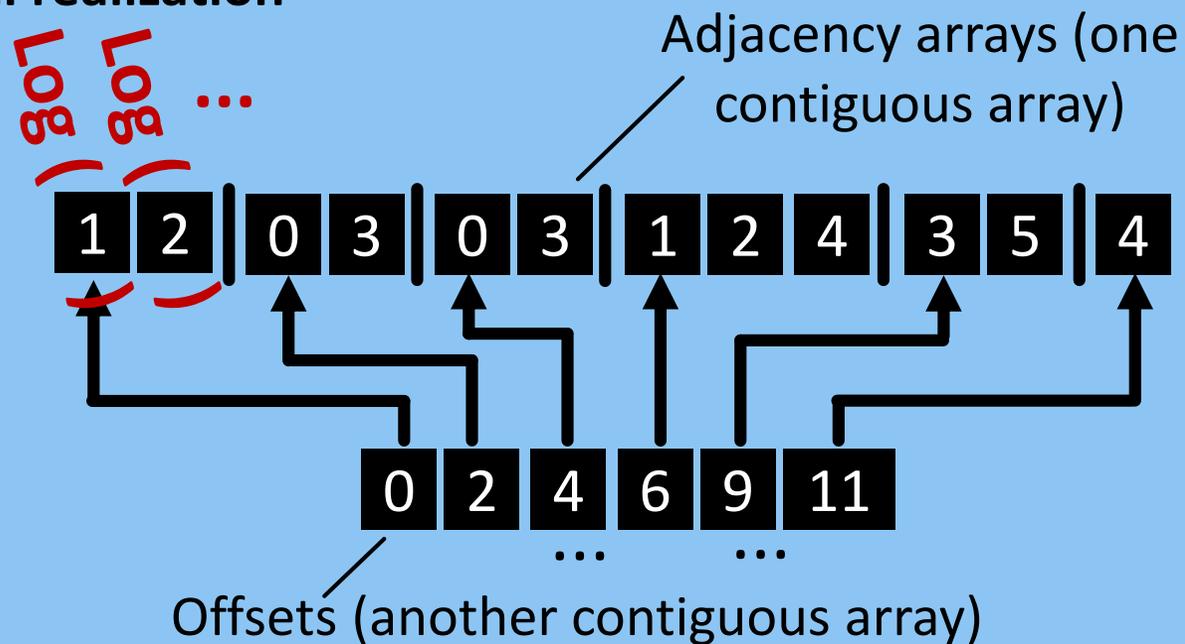
ADJACENCY ARRAY GRAPH REPRESENTATION

1 $\text{Log}(\text{Vertex labels}), \text{Log}(\text{Edge weights})$

Global bound $\lceil \log n \rceil$

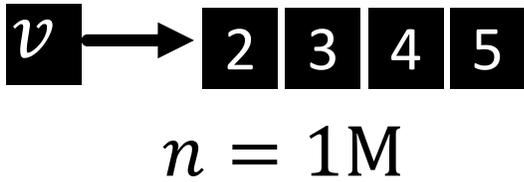


Physical realization



ADJACENCY ARRAY GRAPH REPRESENTATION

1 $\text{Log}(\text{Vertex labels}), \text{Log}(\text{Edge weights})$

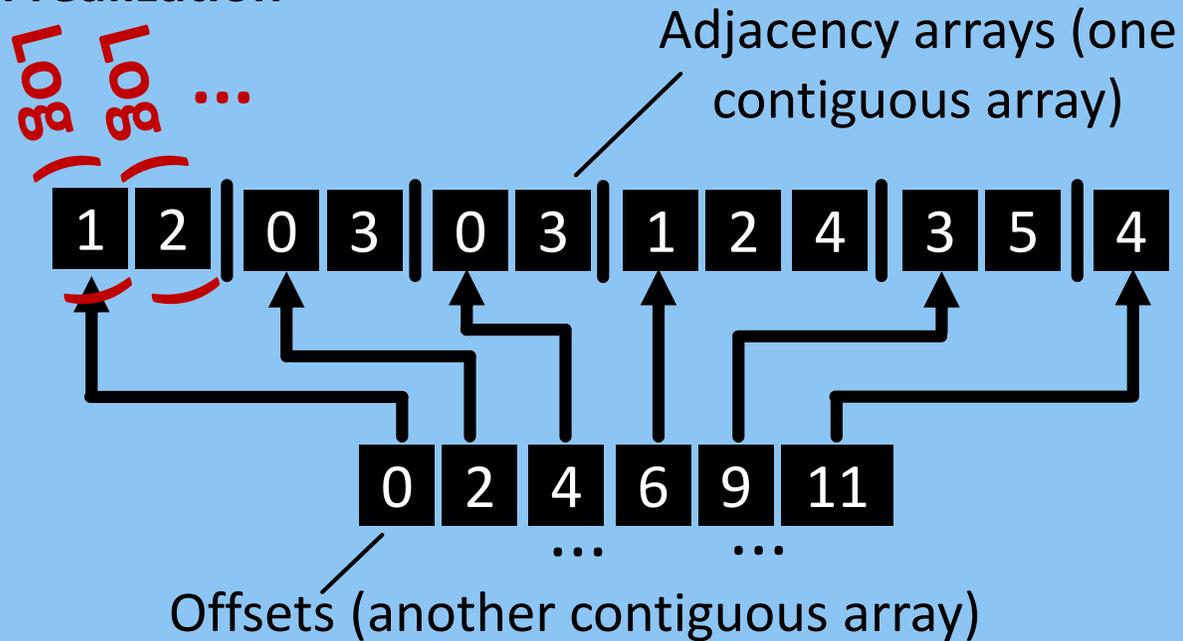


Global bound $\lceil \log n \rceil$



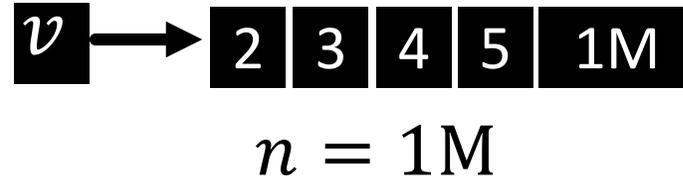
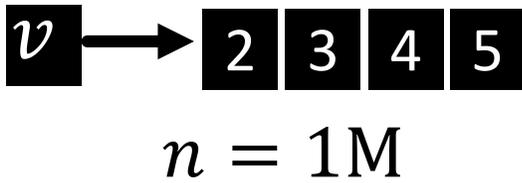
Local bounds $\lceil \log(\max(v_i)) \rceil$

Physical realization



ADJACENCY ARRAY GRAPH REPRESENTATION

1 $\text{Log}(\text{Vertex labels}), \text{Log}(\text{Edge weights})$

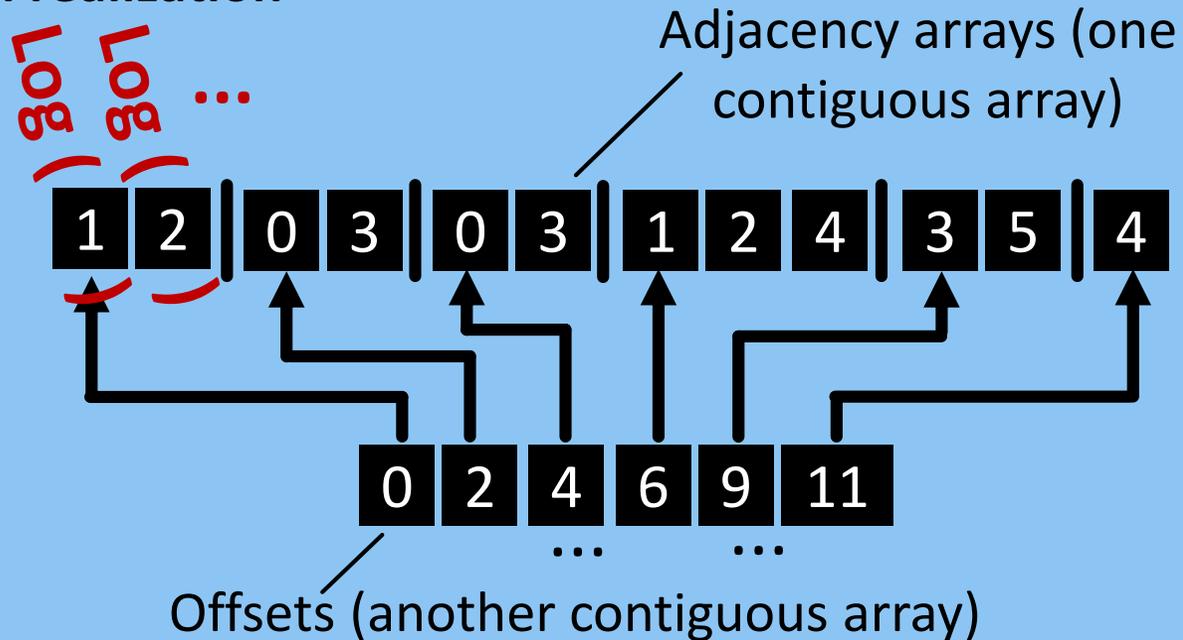


Global bound $\lceil \log n \rceil$



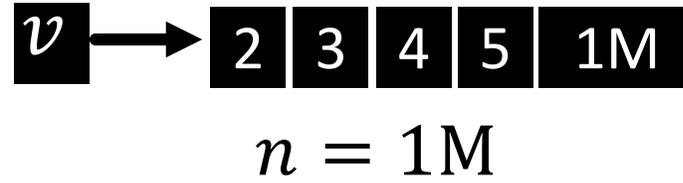
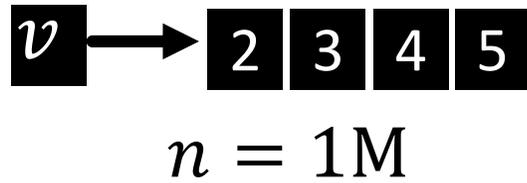
Local bounds $\lceil \log(\max(v_i)) \rceil$

Physical realization

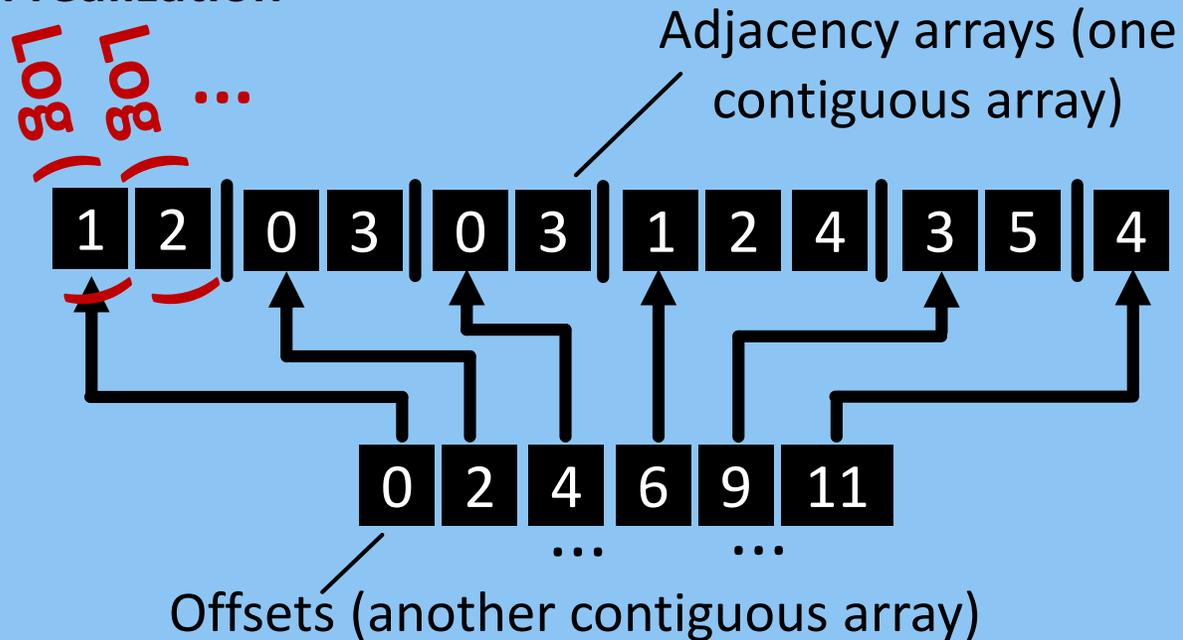


ADJACENCY ARRAY GRAPH REPRESENTATION

1 $\text{Log}(\text{Vertex labels}), \text{Log}(\text{Edge weights})$



Physical realization



Global bound $\lceil \log n \rceil$



Local bounds $\lceil \log(\max(v_i)) \rceil$



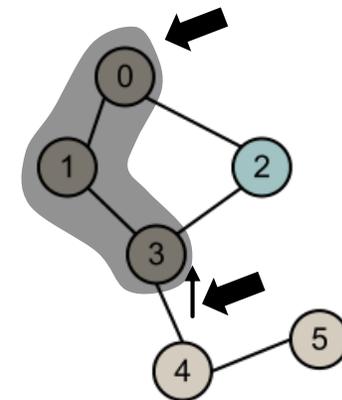
Permute vertex labels to reduce such maximum labels in as many neighborhoods as possible

2 **Log** (Offset structure)

3 **Log** (Adjacency structure)

2 Log (Offset structure)

3 Log (Adjacency structure)



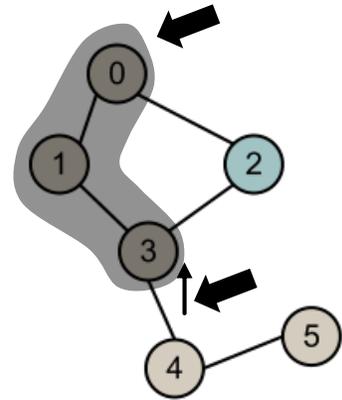
[1] G. J. Jacobson. Succinct Static Data Structures. 1988

M. Besta et al.: "Log(Graph): A Near-Optimal High-Performance Graph Representation", PACT'18

2 **Log** (Offset structure)

3 **Log** (Adjacency structure)

Succinct data structures [1]



[1] G. J. Jacobson. Succinct Static Data Structures. 1988

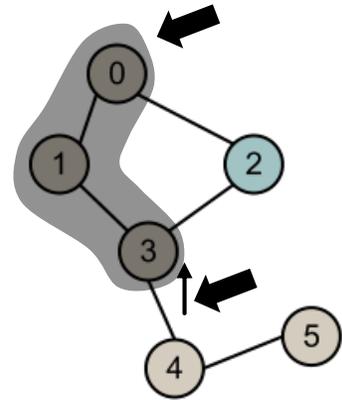
M. Besta et al.: "Log(Graph): A Near-Optimal High-Performance Graph Representation", PACT'18

2 **Log** (Offset structure)

3 **Log** (Adjacency structure)

Succinct data structures [1]

They use $\lceil Q \rceil + o(Q)$ bits ($\lceil Q \rceil$ - lower bound),
they answer various queries in $o(Q)$ time.



[1] G. J. Jacobson. Succinct Static Data Structures. 1988

M. Besta et al.: "Log(Graph): A Near-Optimal High-Performance Graph Representation", PACT'18

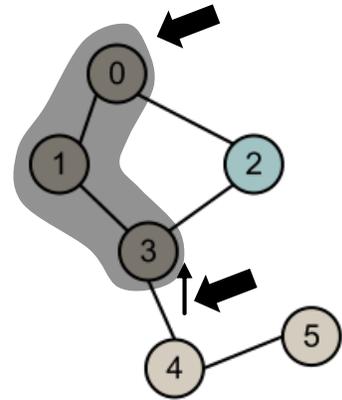
2 **Log** (Offset structure)

3 **Log** (Adjacency structure)

Succinct data structures [1]

They use $\lceil Q \rceil + o(Q)$ bits ($\lceil Q \rceil$ - lower bound), they answer various queries in $o(Q)$ time.

Compact data structures [1]



[1] G. J. Jacobson. Succinct Static Data Structures. 1988

M. Besta et al.: "Log(Graph): A Near-Optimal High-Performance Graph Representation", PACT'18

2 **Log** (Offset structure)

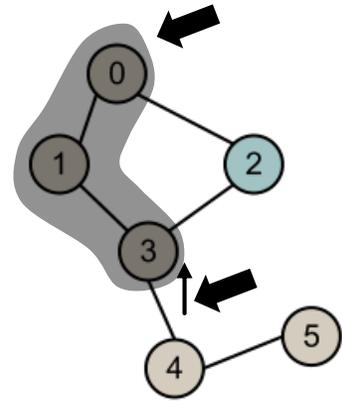
3 **Log** (Adjacency structure)

Succinct data structures [1]

They use $\lceil Q \rceil + o(Q)$ bits ($\lceil Q \rceil$ - lower bound), they answer various queries in $o(Q)$ time.

Compact data structures [1]

They use $O(Q)$ bits ($\lceil Q \rceil$ - lower bound), they answer various queries in $o(Q)$ time.



[1] G. J. Jacobson. Succinct Static Data Structures. 1988

M. Besta et al.: "Log(Graph): A Near-Optimal High-Performance Graph Representation", PACT'18

2 **Log** (Offset structure)

3 **Log** (Adjacency structure)

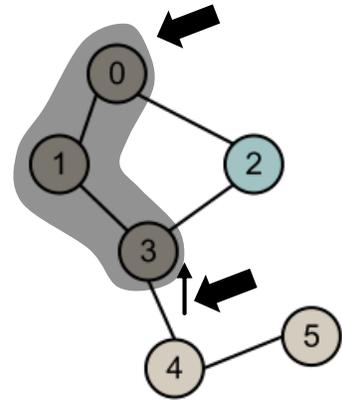
Succinct data structures [1]

They use $\lceil Q \rceil + o(Q)$ bits ($\lceil Q \rceil$ - lower bound), they answer various queries in $o(Q)$ time.

Compact data structures [1]

They use $O(Q)$ bits ($\lceil Q \rceil$ - lower bound), they answer various queries in $o(Q)$ time.

We show that
they are in practice
both small and fast!

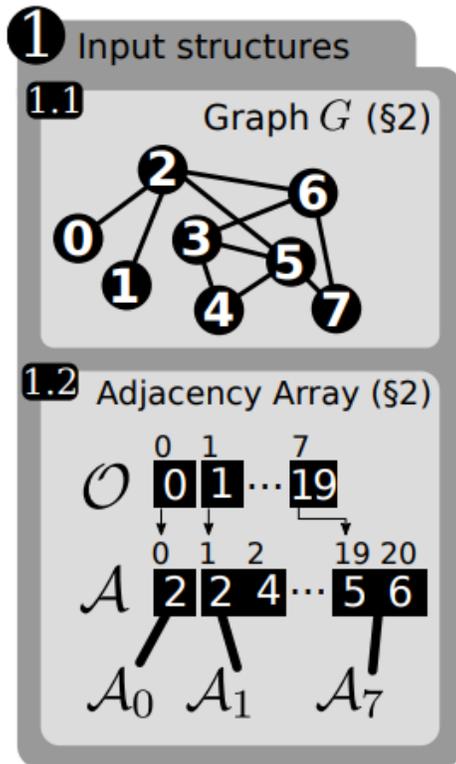


[1] G. J. Jacobson. Succinct Static Data Structures. 1988

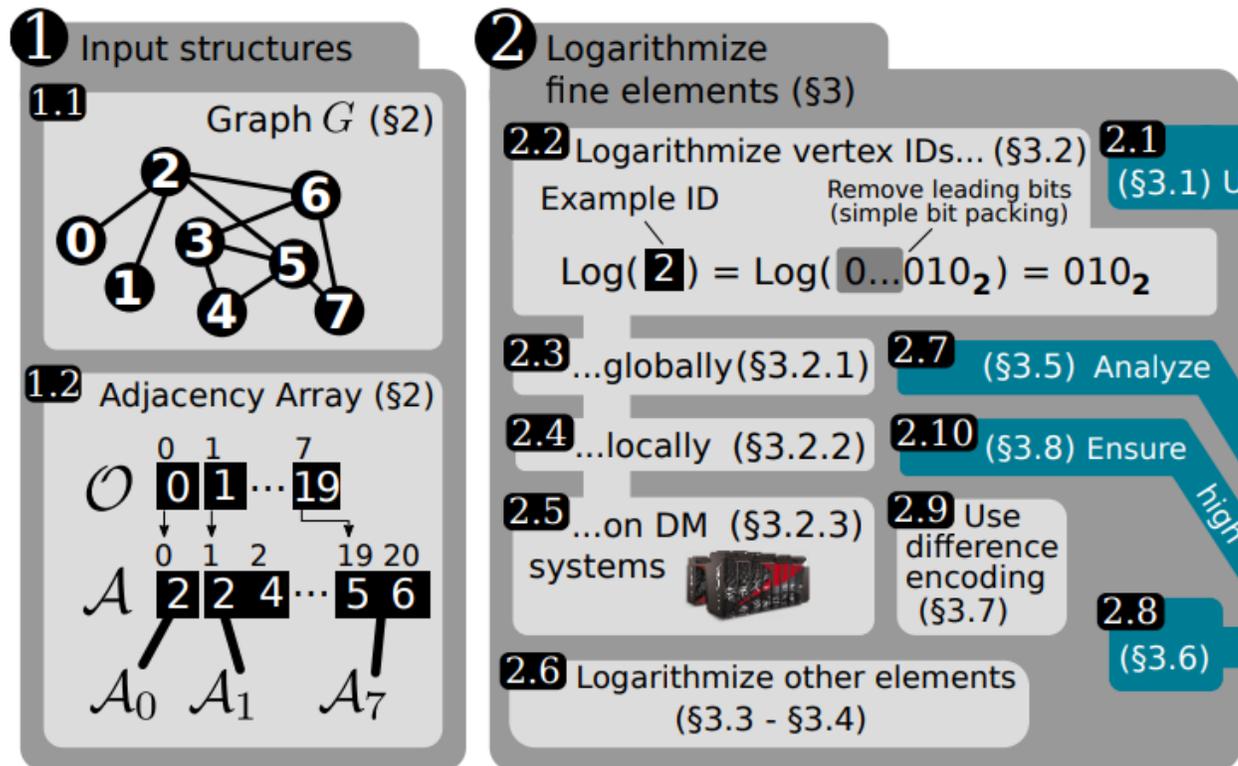
M. Besta et al.: "Log(Graph): A Near-Optimal High-Performance Graph Representation", PACT'18

OVERVIEW OF FULL LOG(GRAPH) DESIGN

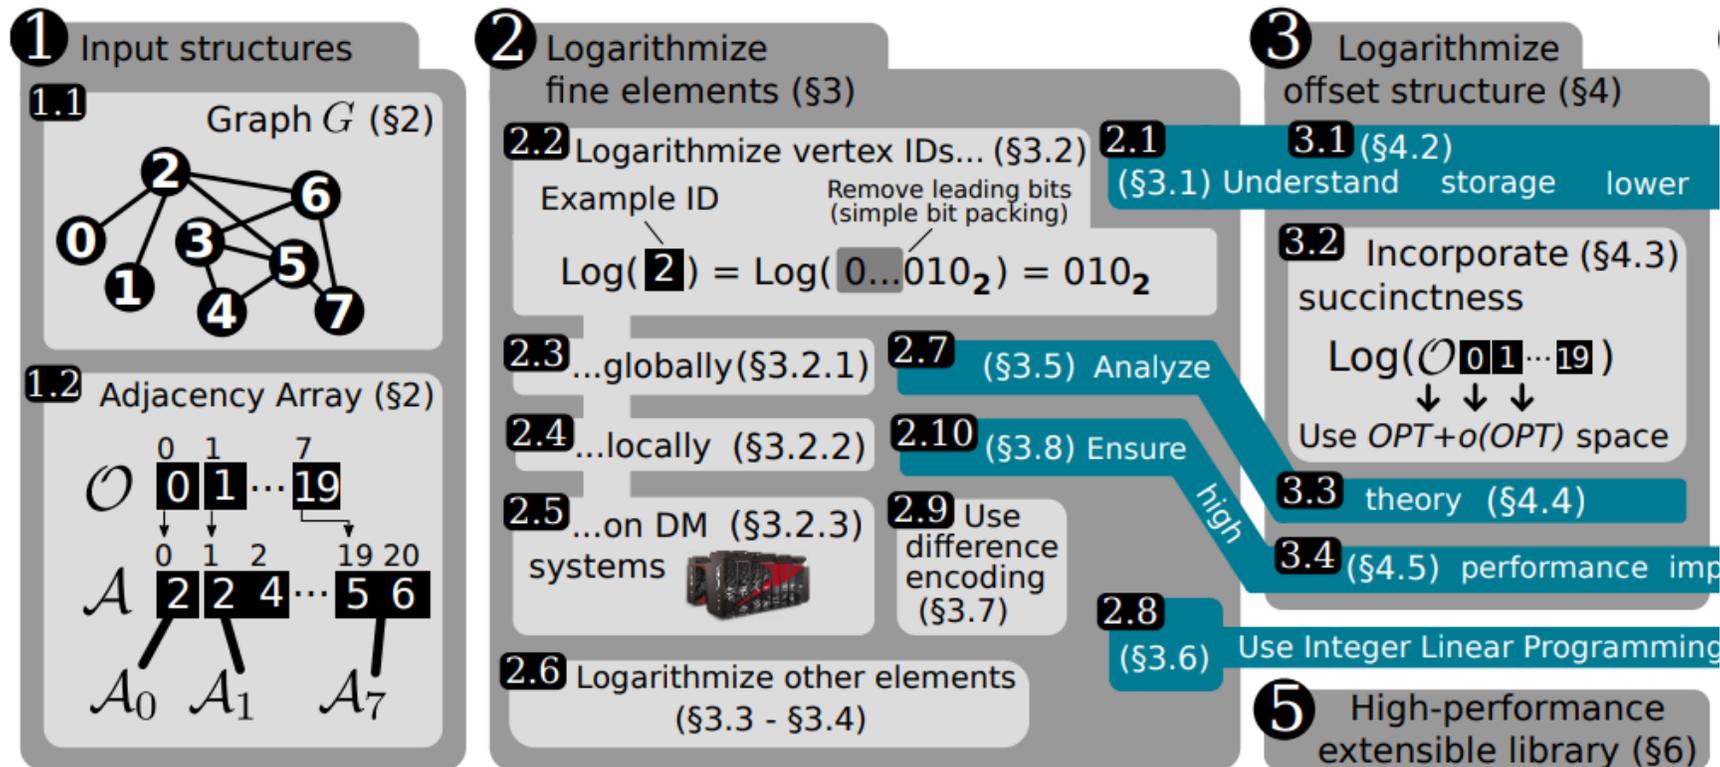
OVERVIEW OF FULL LOG(GRAPH) DESIGN



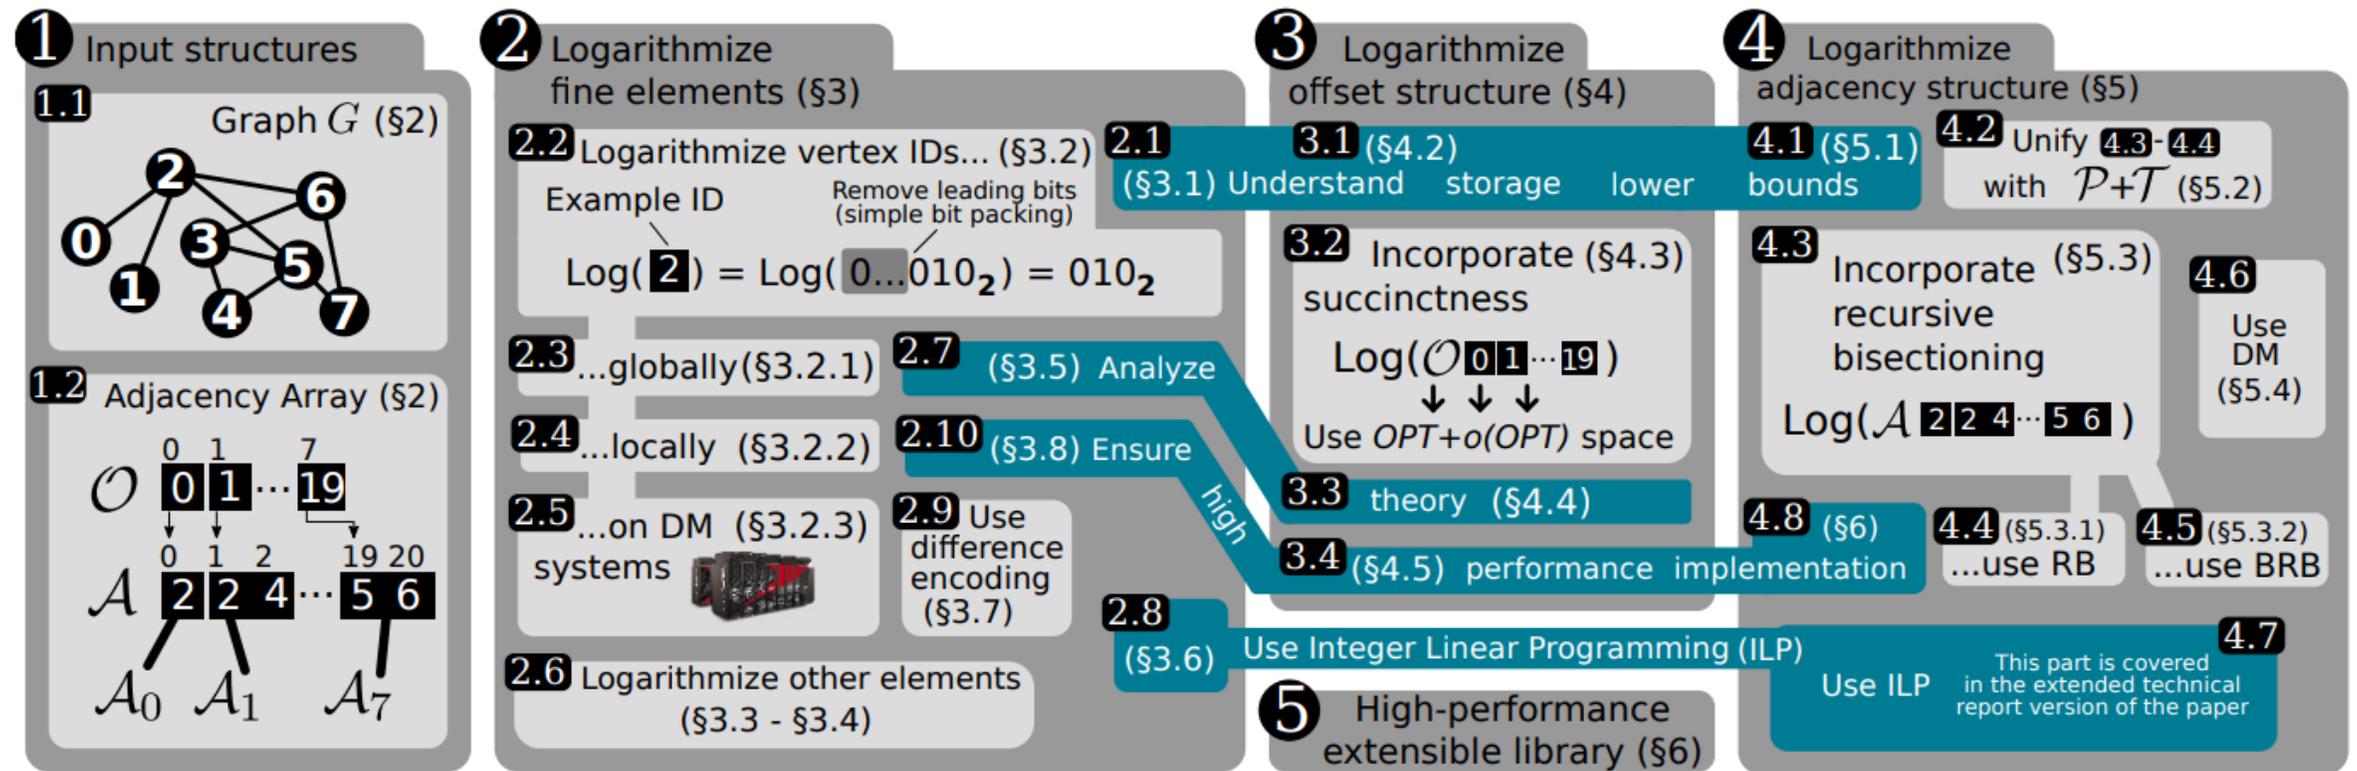
OVERVIEW OF FULL LOG(GRAPH) DESIGN



OVERVIEW OF FULL LOG(GRAPH) DESIGN



OVERVIEW OF FULL LOG(GRAPH) DESIGN



OVERVIEW OF FULL LOG(GRAPH) DESIGN

1 Input structures

1.1 Graph G (§2)

1.2 Adjacency Array (§2)

2 Logarithmize fine elements (§3)

2.1 Logarithmize vertex IDs... (§3.2)

Example ID: $\text{Log}(2) = \text{Log}(0\dots010_2) = 010_2$

Remove leading bits (simple bit packing)

2.2 ...globally (§3.2.1)

2.3 ...locally (§3.2.2)

2.4 ...on DM (§3.2.3) systems

2.5 Logarithmize other elements (§3.3 - §3.4)

2.6 (§3.5) Analyze

2.7 (§3.8) Ensure

2.8 Use difference encoding (§3.7)

2.9 Use Integer Linear Programming (ILP)

2.10 High-performance extensible library (§6)

3 Logarithmize offset structure (§4)

3.1 Incorporate succinctness

$\text{Log}(001\dots19)$

Use $OPT + o(OPT)$ space

3.2 theory (§4.4)

3.3 performance implementation

3.4 Use ILP

4 Logarithmize adjacency structure (§5)

4.1 Unify 4.3-4.4 with $P+T$ (§5.2)

4.2 Incorporate recursive bisectioning

$\text{Log}(A 224\dots56)$

4.3 Use DM (§5.4)

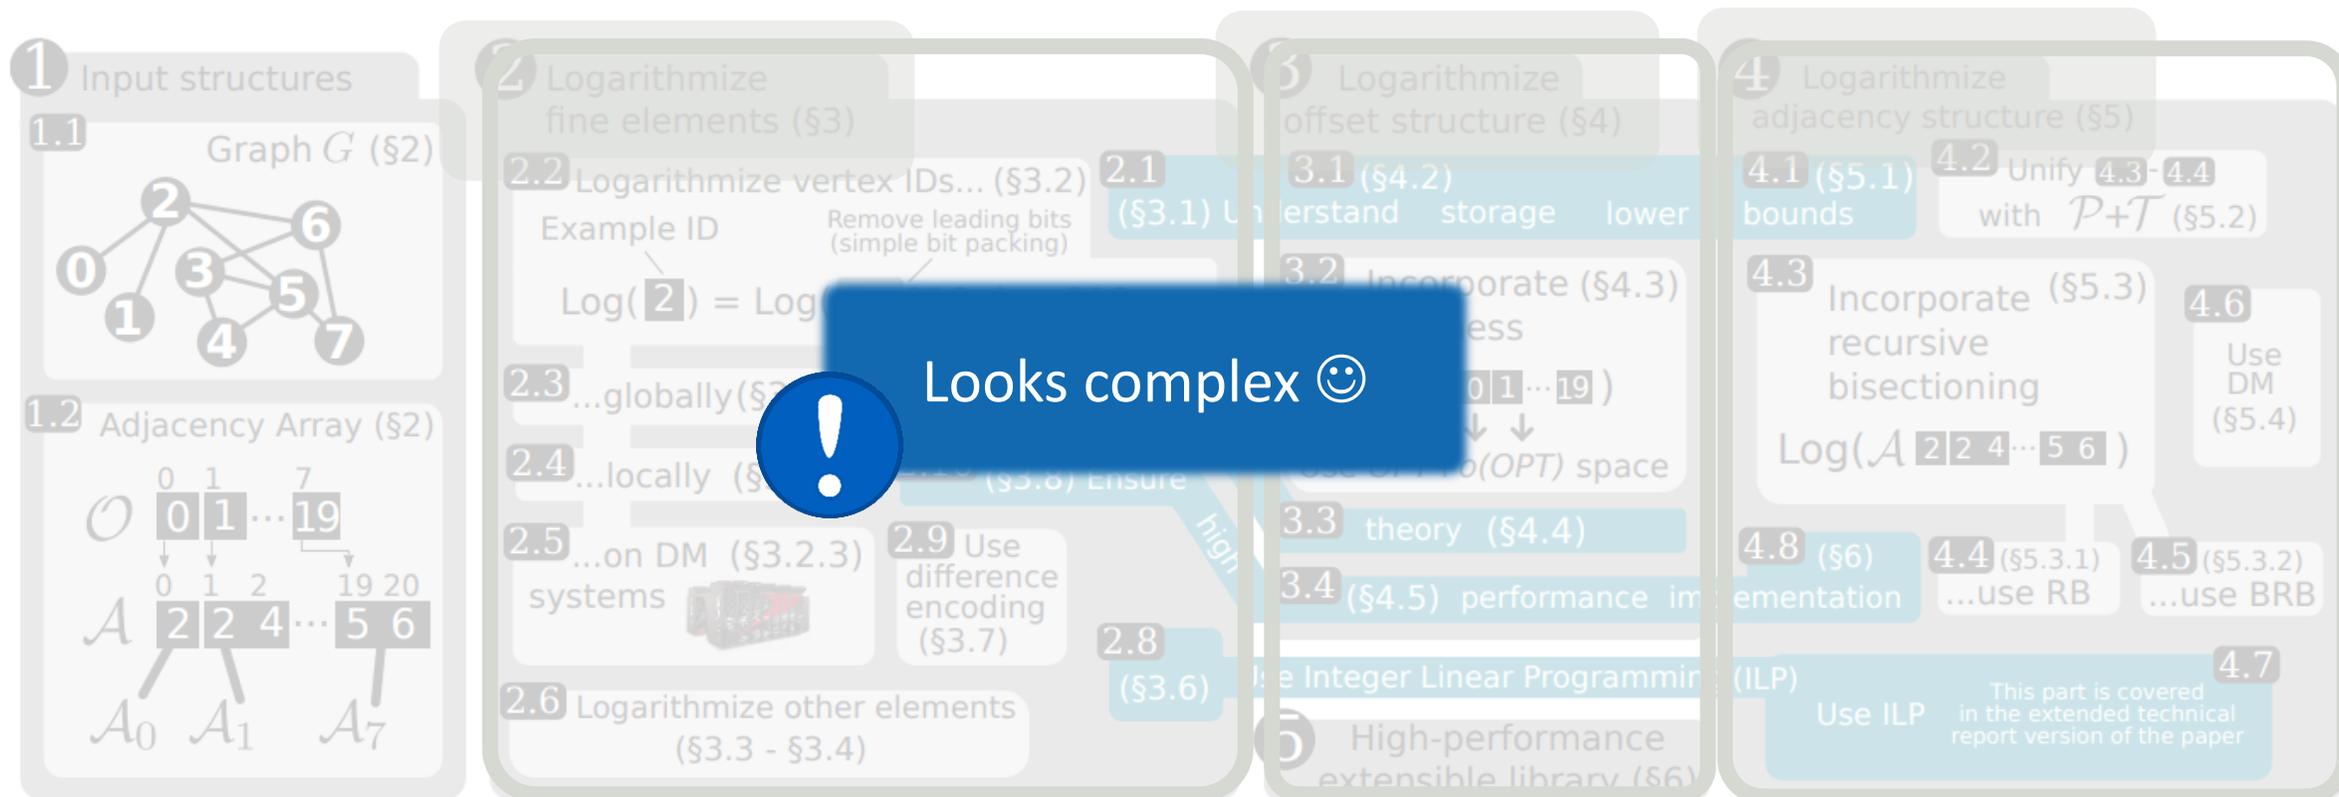
4.4 ...use RB

4.5 ...use BRB

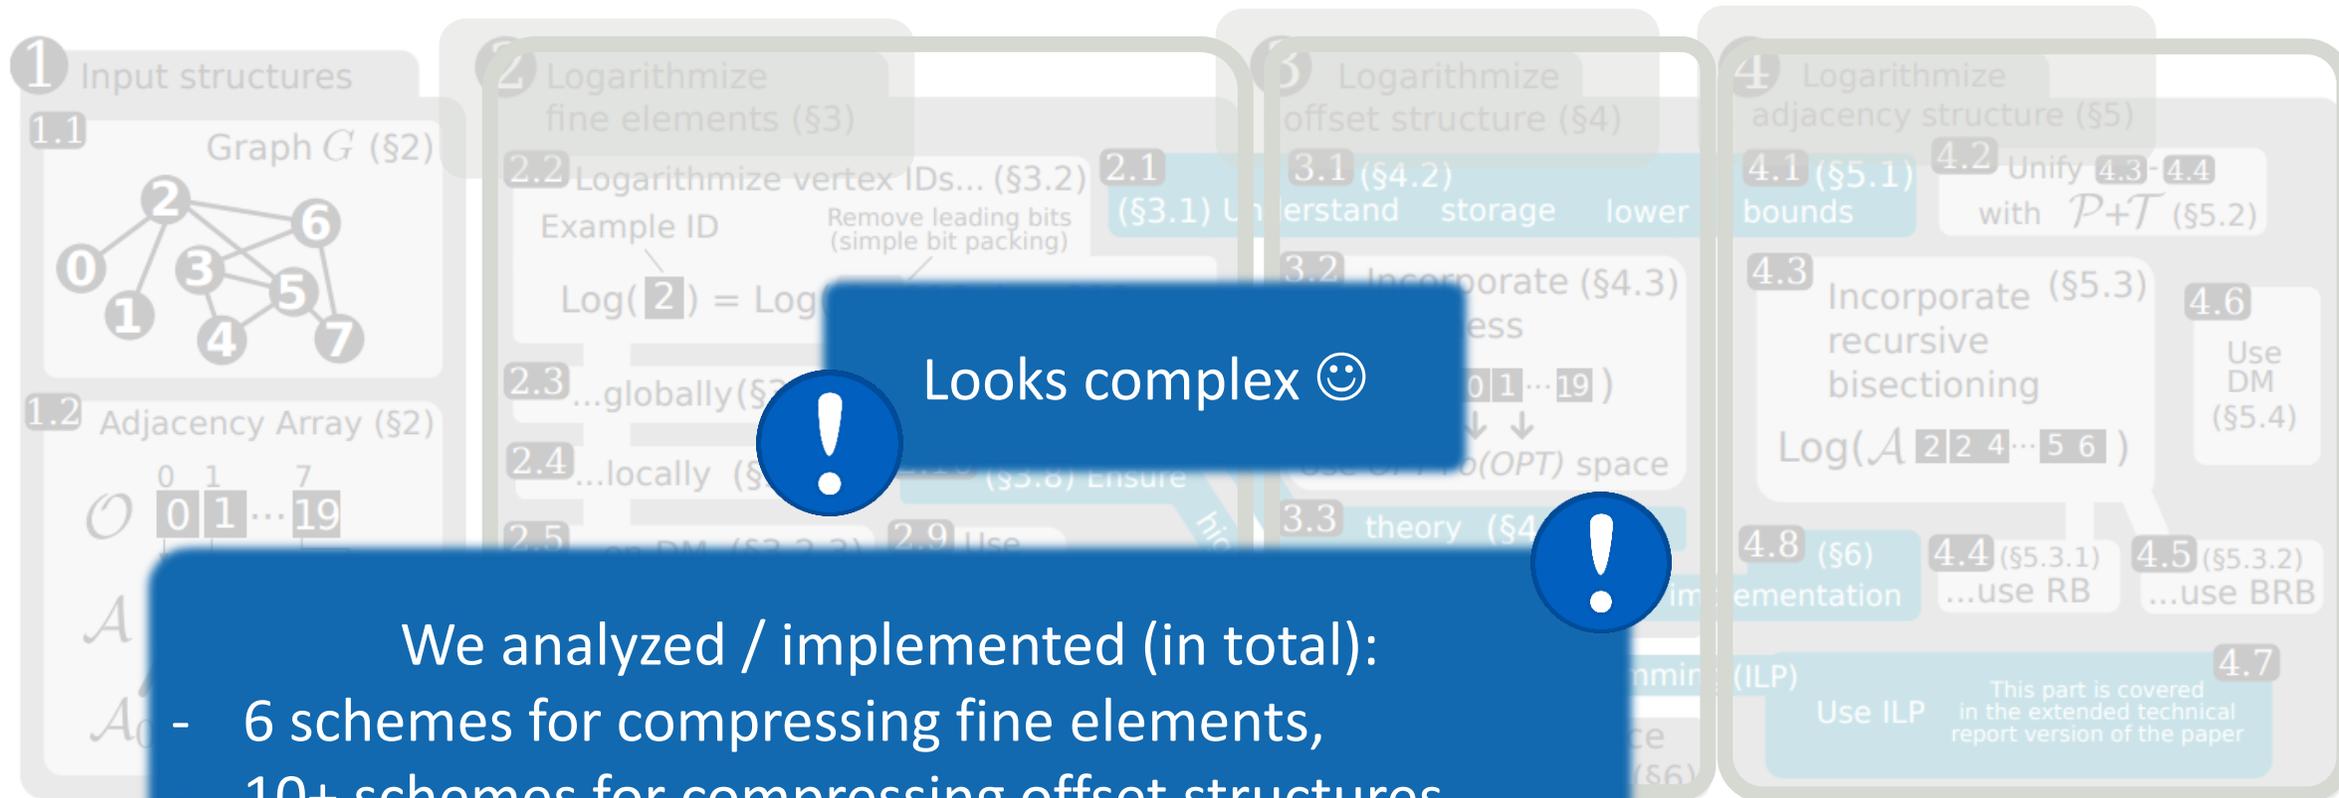
4.6 Use ILP

4.7 This part is covered in the extended technical report version of the paper

OVERVIEW OF FULL LOG(GRAPH) DESIGN



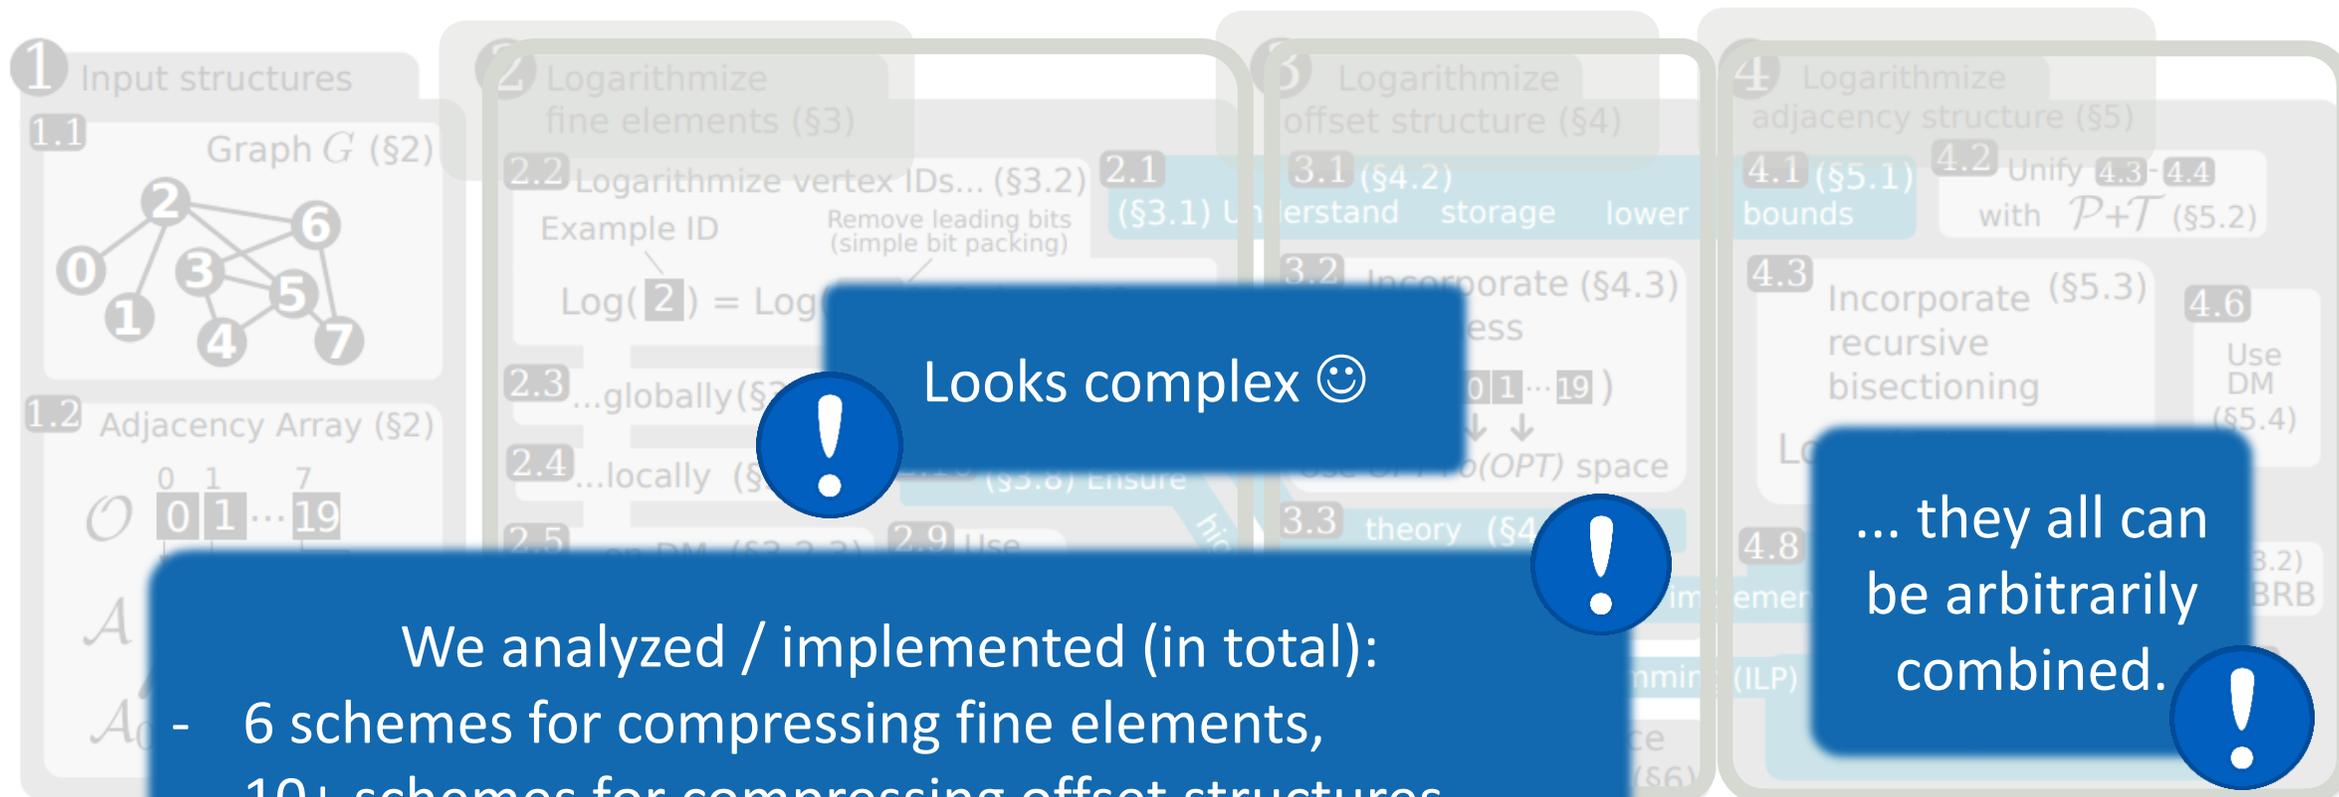
OVERVIEW OF FULL LOG(GRAPH) DESIGN



We analyzed / implemented (in total):

- 6 schemes for compressing fine elements,
- 10+ schemes for compressing offset structures,
- 4+ schemes for compressing adjacency structures

OVERVIEW OF FULL LOG(GRAPH) DESIGN



The image shows a collage of design steps for Log(Graph) with callouts:

- 1 Input structures**
 - 1.1 Graph G (§2) 
 - 1.2 Adjacency Array (§2) 
- 2 Logarithmize fine elements (§3)**
 - 2.1 Logarithmize vertex IDs... (§3.2) 
 - 2.2 ...globally (§3.1)
 - 2.3 ...locally (§3.2)
 - 2.4 ...locally (§3.2)
 - 2.5 ...locally (§3.2)
 - 2.9 Use DM (§3.2)
- 3 Logarithmize offset structure (§4)**
 - 3.1 (§4.2) Understand storage lower
 - 3.2 Incorporate (§4.3) 
 - 3.3 theory (§4.3)
- 4 Logarithmize adjacency structure (§5)**
 - 4.1 (§5.1) bounds
 - 4.2 Unify 4.3-4.4 with $P+T$ (§5.2)
 - 4.3 Incorporate (§5.3) recursive bisectioning
 - 4.6 Use DM (§5.4)
 - 4.8 (ILP)

Callouts:

- Looks complex 😊
- ... they all can be arbitrarily combined.

We analyzed / implemented (in total):

- 6 schemes for compressing fine elements,
- 10+ schemes for compressing offset structures,
- 4+ schemes for compressing adjacency structures

OVERVIEW OF FULL LOG(GRAPH) DESIGN

How to ensure fast, manageable, and extensible implementation of all these schemes?



Looks complex 😊



... they all can be arbitrarily combined.



We analyzed / implemented (in total):

- 6 schemes for compressing fine elements,
- 10+ schemes for compressing offset structures,
- 4+ schemes for compressing adjacency structures

OVERVIEW OF FULL LOG(GRAPH) DESIGN

How to ensure fast, manageable, and extensible implementation of all these schemes?

We use C++ templates to develop a platform that facilitates implementation, benchmarking, analysis, and extending the discussed and many other schemes

Looks complex 😊

We analyzed / implemented (in total):

- 6 schemes for compressing fine elements,
- 10+ schemes for compressing offset structures,
- 4+ schemes for compressing adjacency structures

... they all can be arbitrarily combined.

OVERVIEW OF FULL LOG(GRAPH) DESIGN

How to ensure fast, manageable, and extensible implementation of all these schemes?

We use C++ templates to develop a platform that facilitates implementation, benchmarking, analysis, and extending the discussed and many other schemes

Looks complex 😊

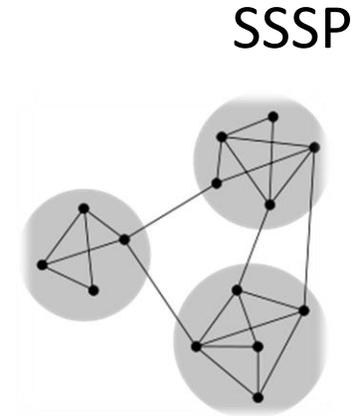
We analyzed / implemented (in total):

- 6 schemes for compressing fine elements,
- 10+ schemes for compressing offset structures,
- 4+ schemes for compressing adjacency structures

... they all can be arbitrarily combined.

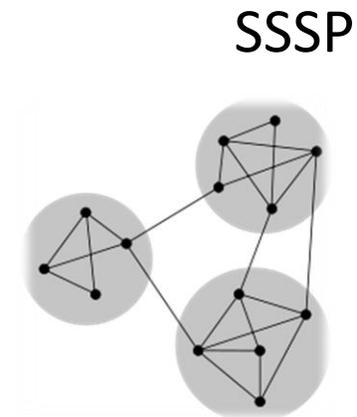
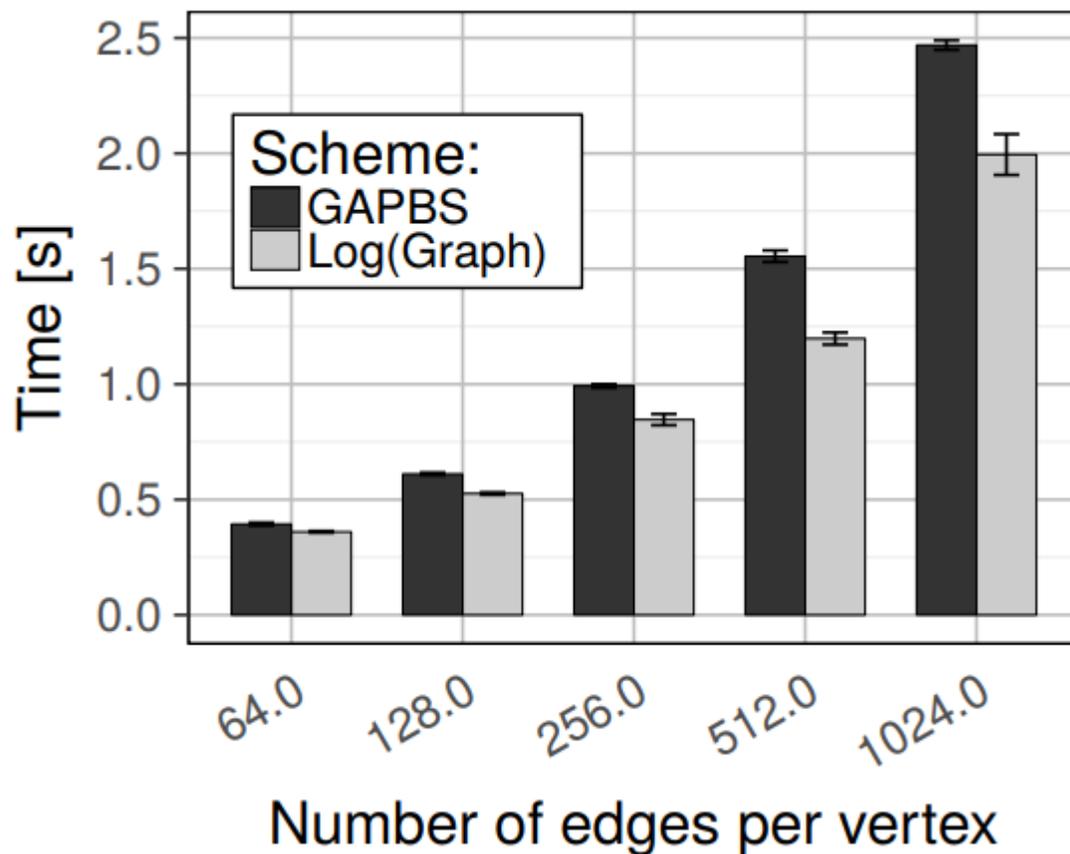
COMING SOON

1 **Log** (Vertex labels), **Log** (Edge weights) **Storage, Performance**



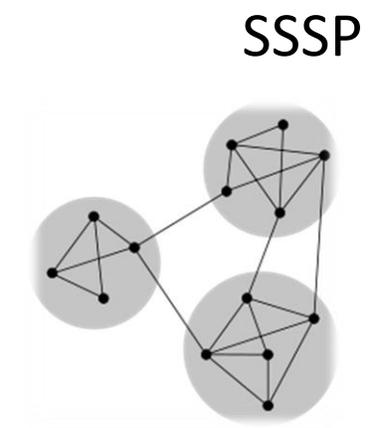
Kronecker graphs
Number of vertices: 4M

1 $\text{Log}(\text{Vertex labels})$, $\text{Log}(\text{Edge weights})$ Storage, Performance

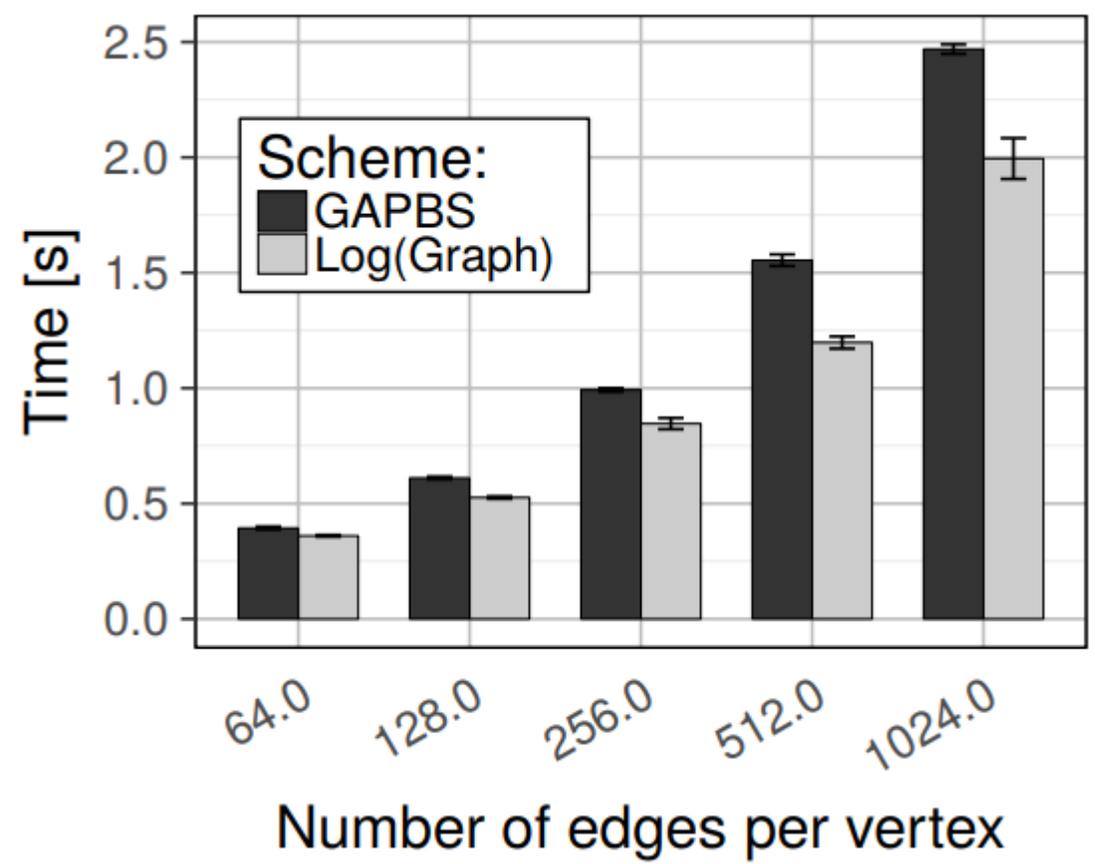


Kronecker graphs
Number of vertices: 4M

1 **Log (Vertex labels), Log (Edge weights)** **Storage, Performance**

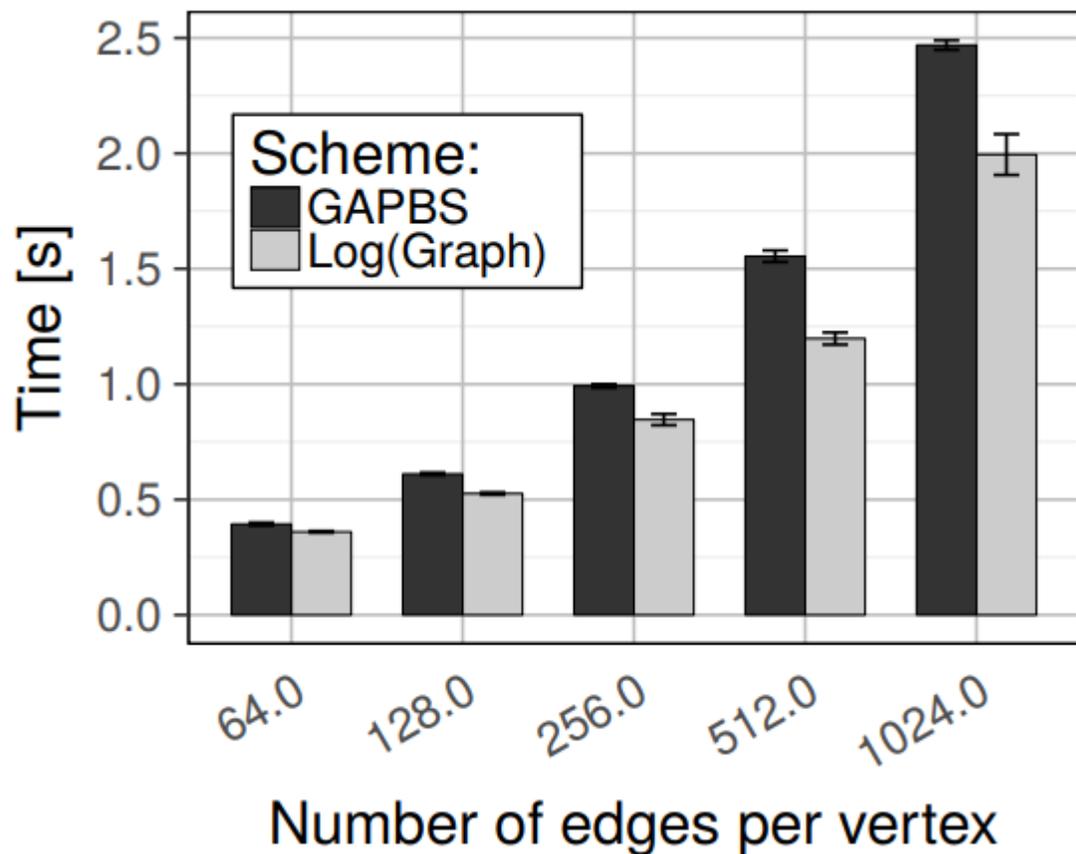


SSSP
Kronecker graphs
Number of vertices: 4M



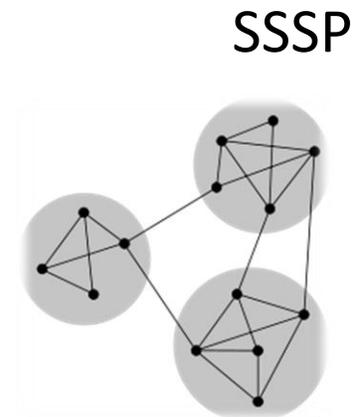
Log(Graph) consistently reduces storage overhead (by 20-35%)

1 $\text{Log}(\text{Vertex labels})$, $\text{Log}(\text{Edge weights})$ **Storage, Performance**



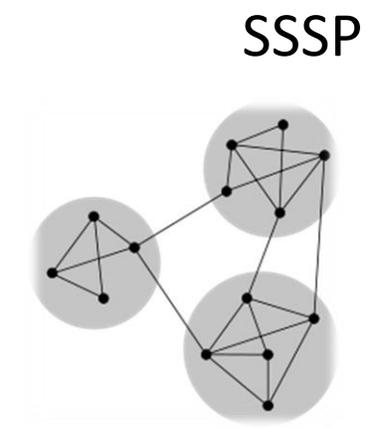
Log(Graph) accelerates GAPBS

Log(Graph) consistently reduces storage overhead (by 20-35%)



Kronecker graphs
Number of vertices: 4M

1 $\text{Log}(\text{Vertex labels})$, $\text{Log}(\text{Edge weights})$ **Storage, Performance**

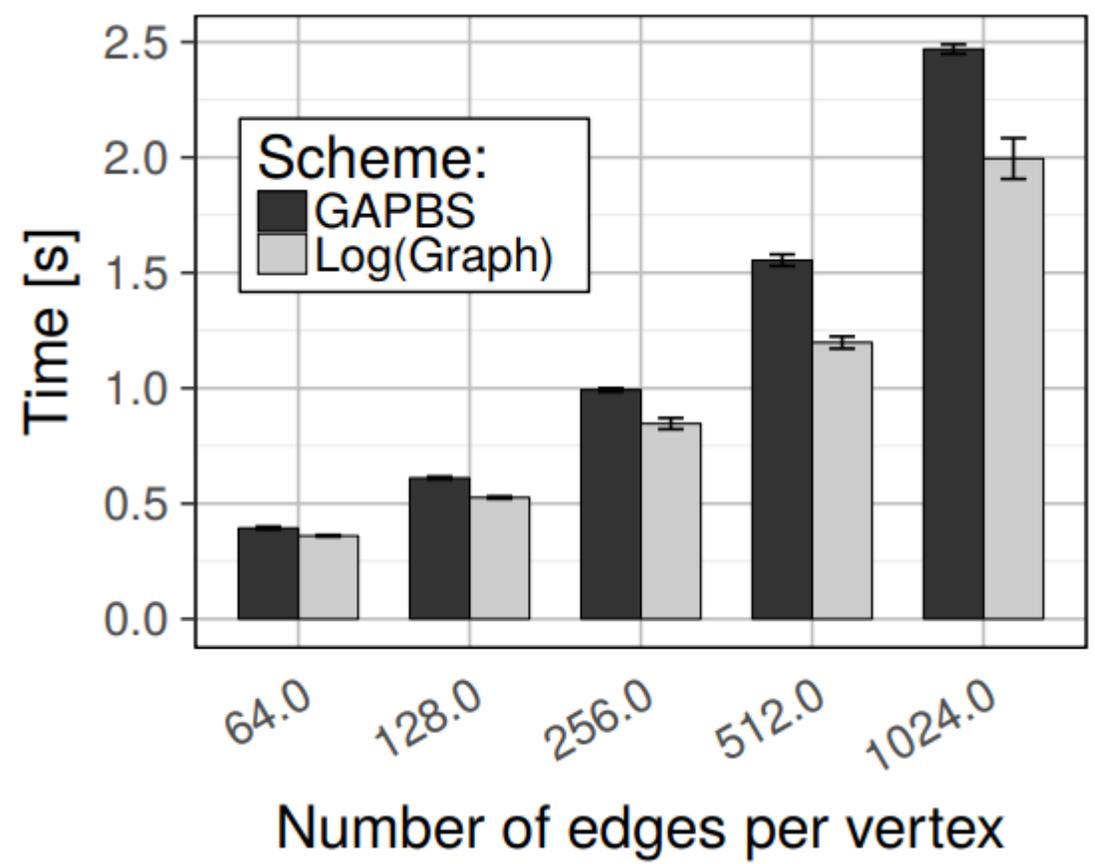


SSSP
Kronecker graphs
Number of vertices: 4M

Log(Graph)
accelerates GAPBS

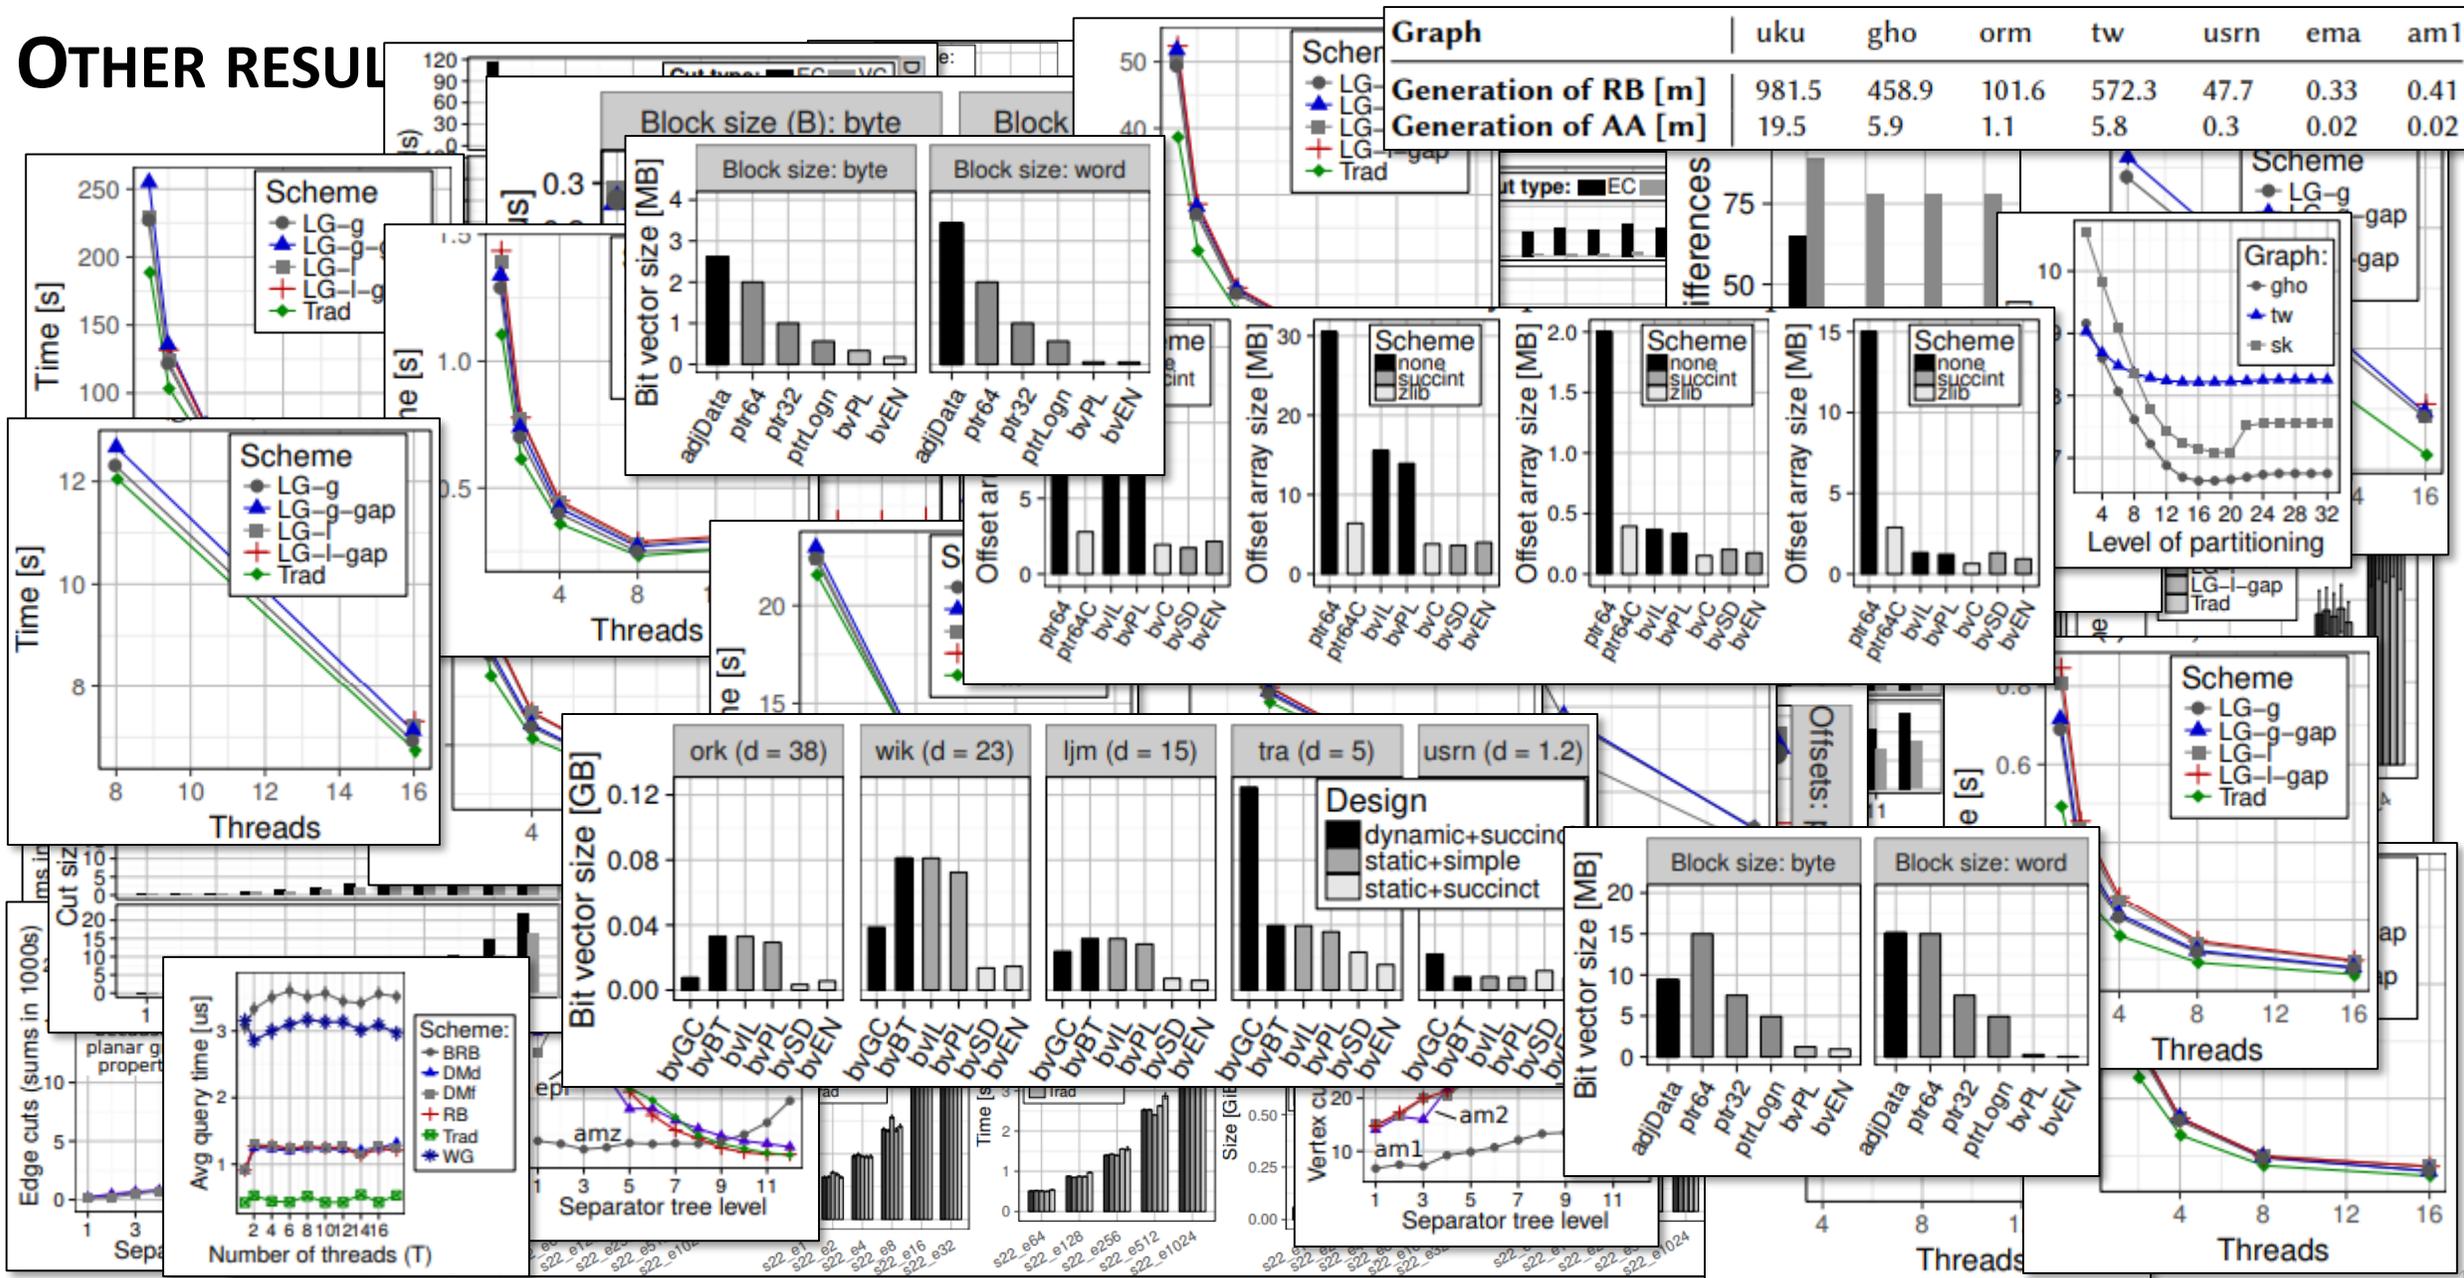
Both storage and performance
are improved **simultaneously**

Log(Graph) consistently
reduces storage overhead
(by 20-35%)

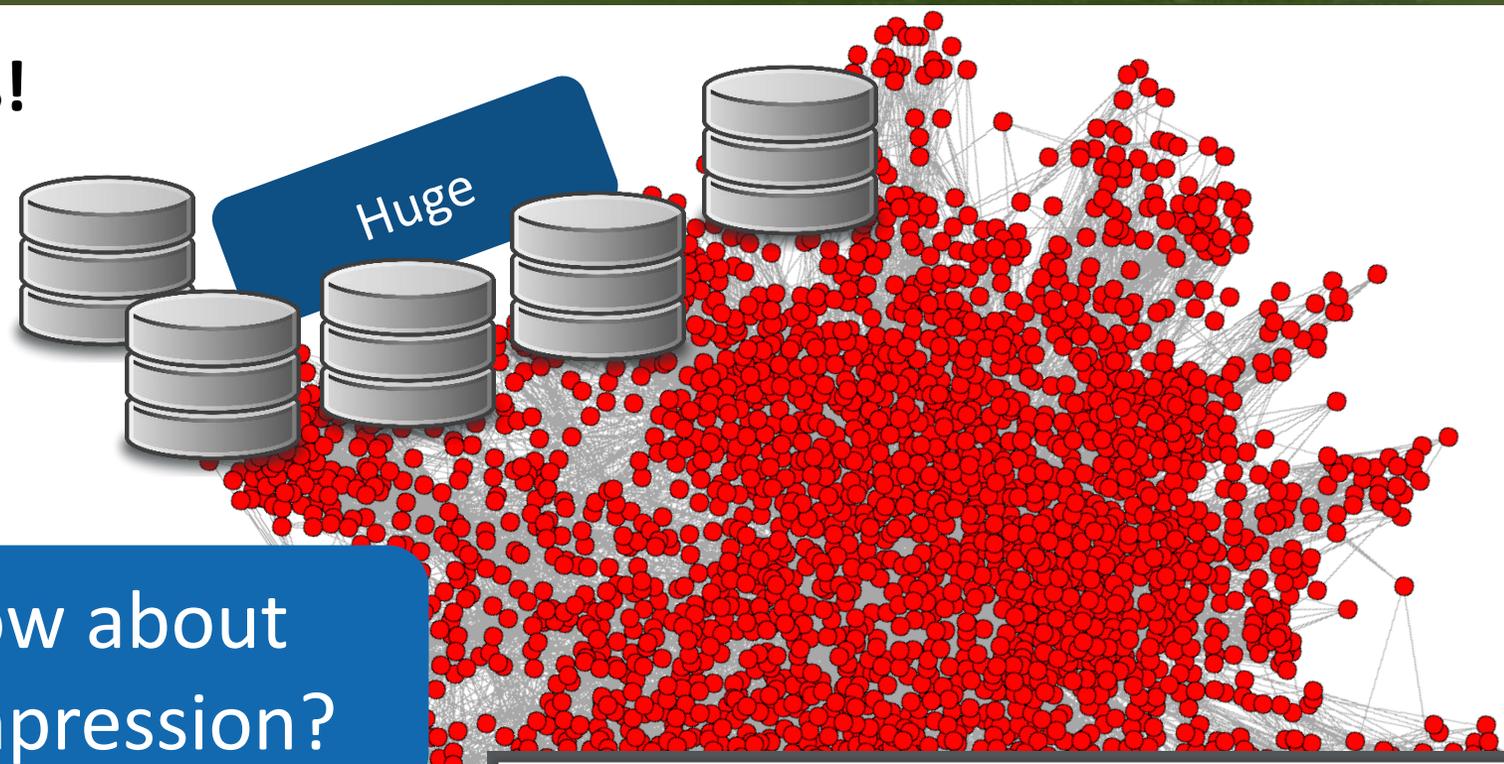


OTHER RESULTS

OTHER RESULT



Problems!



? How about compression?

Compression incurs expensive decompression



Log(Graph): A Near-Optimal High-Performance Graph Representation

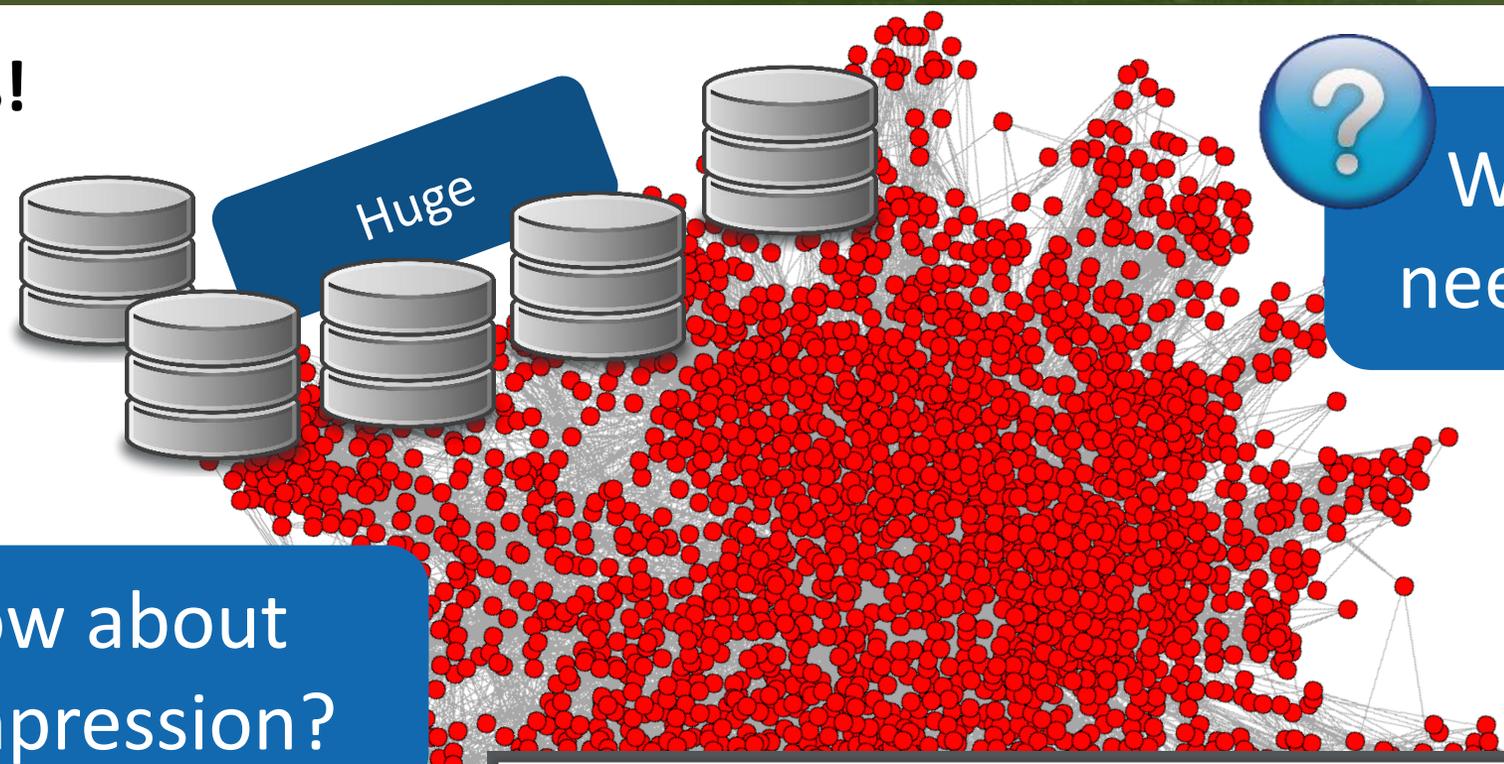
Maciej Besta[†], Dimitri Stanojevic, Tijana Zivic, Jagpreet Singh, Maurice Hoerold, Torsten Hoefler[†]
Department of Computer Science, ETH Zurich
[†]Corresponding authors (maciej.best@inf.ethz.ch, ttor@inf.ethz.ch)

ABSTRACT
Today's graphs, such as social networks, have billions of nodes and edges. Yet, standard graph representations are inefficient, as they store a significant number of bits while graph compression schemes...

size of such graphs is becoming increasingly important.

! Log(Graph): effective compression with low-overhead decompression!

Problems!



? What if we don't need full precision?

? How about compression?

Compression incurs expensive decompression



Log(Graph): A Near-Optimal High-Performance Graph Representation

Maciej Besta[†], Dimitri Stanojevic, Tijana Zivic, Jagpreet Singh, Maurice Hoerold, Torsten Hoefler[†]
 Department of Computer Science, ETH Zurich

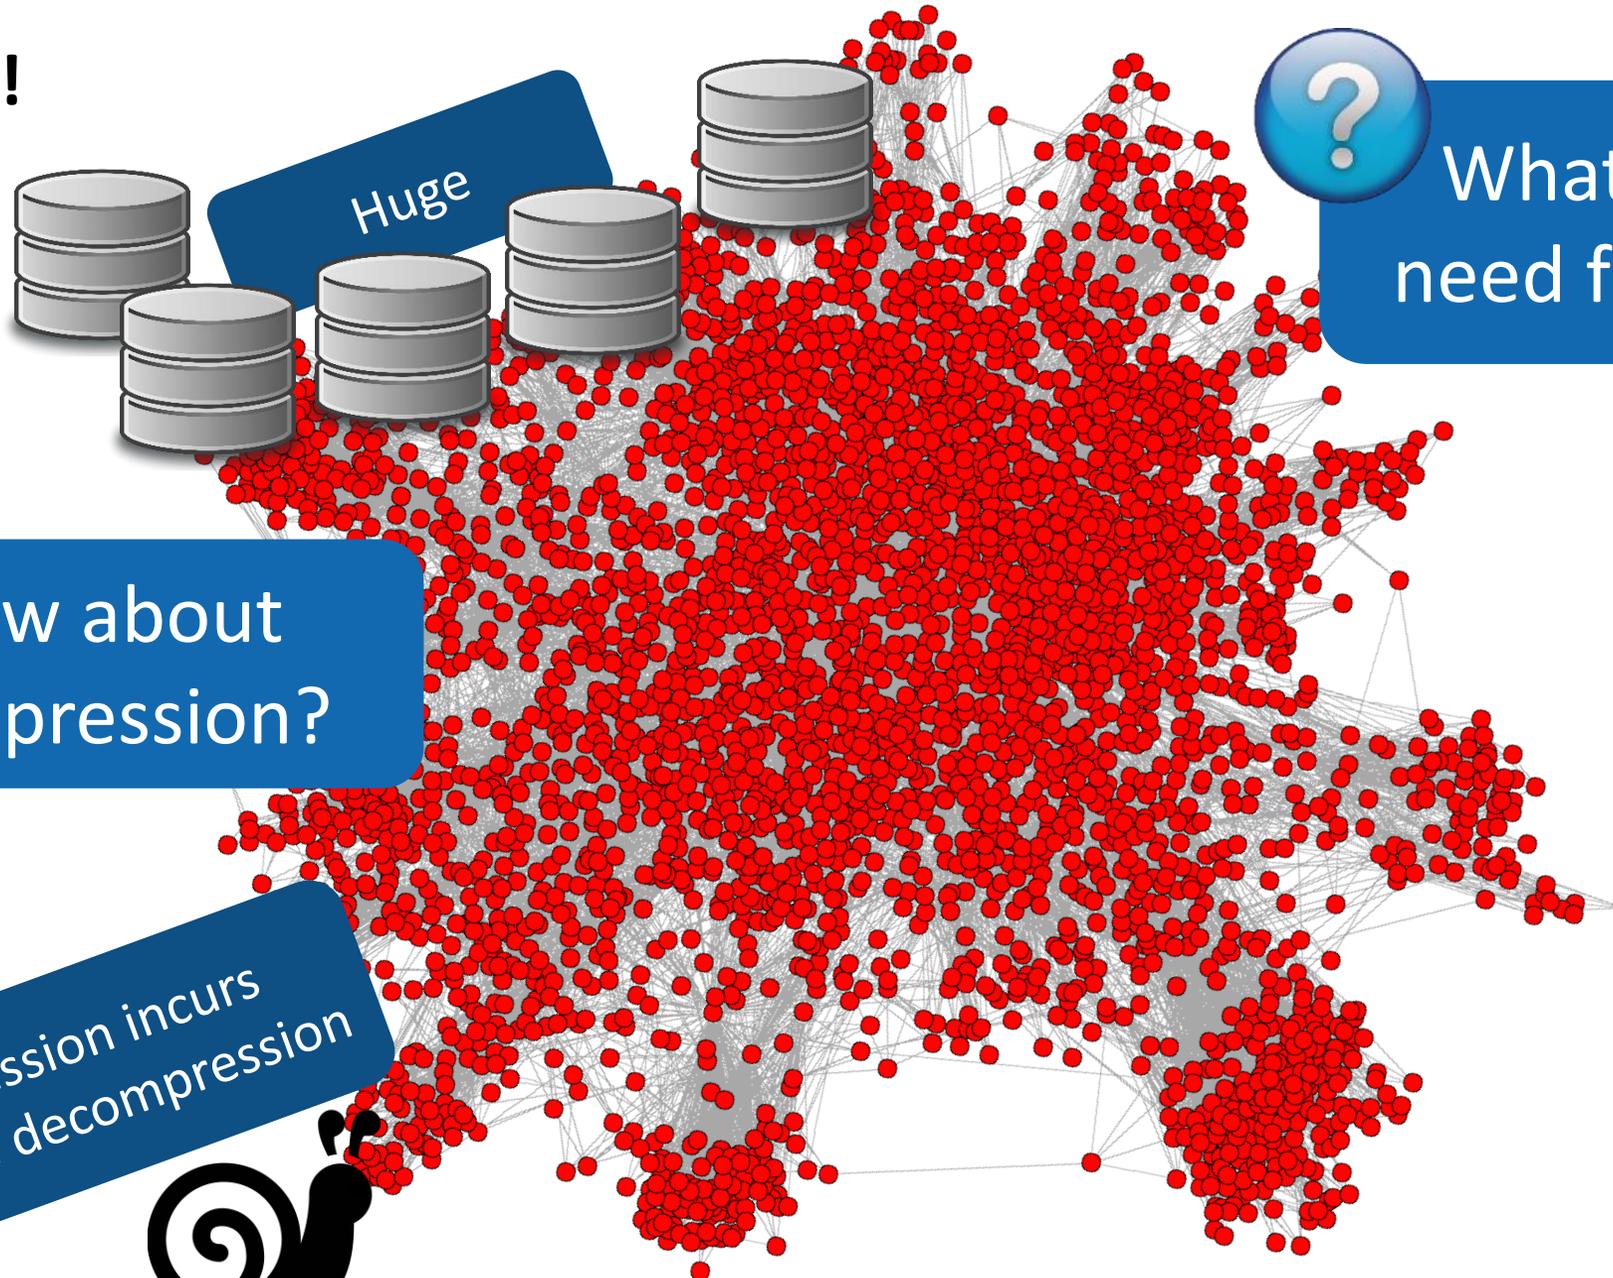
[†]Corresponding authors (maciej.besta@inf.ethz.ch, ttor@inf.ethz.ch)

ABSTRACT
 Today's graphs, such as social networks, have billions of nodes and edges. Yet, standard graph representations require a significant number of bits while graph compression schemes...

size of such graphs is becoming increasingly important.

! **Log(Graph): effective compression with low-overhead decompression!**

Problems!



Problems!



What if we don't need full precision?

How about compression?

Accepted @ ACM/IEEE SC19, Best paper finalist

Compression incurs expensive decompression



Slim Graph: Practical Lossy Graph Compression for Approximate Graph Processing, Storage, and Analytics

Maciej Besta, Simon Weber, Lukas Gianinazzi,
Robert Gerstenberger, Andrey Ivanov, Yishai Oltchik, Torsten Hoefler
Department of Computer Science; ETH Zurich

ABSTRACT

We propose Slim Graph: the first programming model and framework for practical lossy graph compression that facilitates high-performance approximate graph processing, storage, and analytics. Slim Graph enables the developer to

I/O operations, the amount of data communicated over the network, and by storing a larger fraction of data in caches.

There exist many *lossless* schemes for graph compression, including WebGraph [21], k^2 -trees [25], and others [17]. They provide various degrees of storage reductions. Unfortunately,

JPEG compression level: 0%



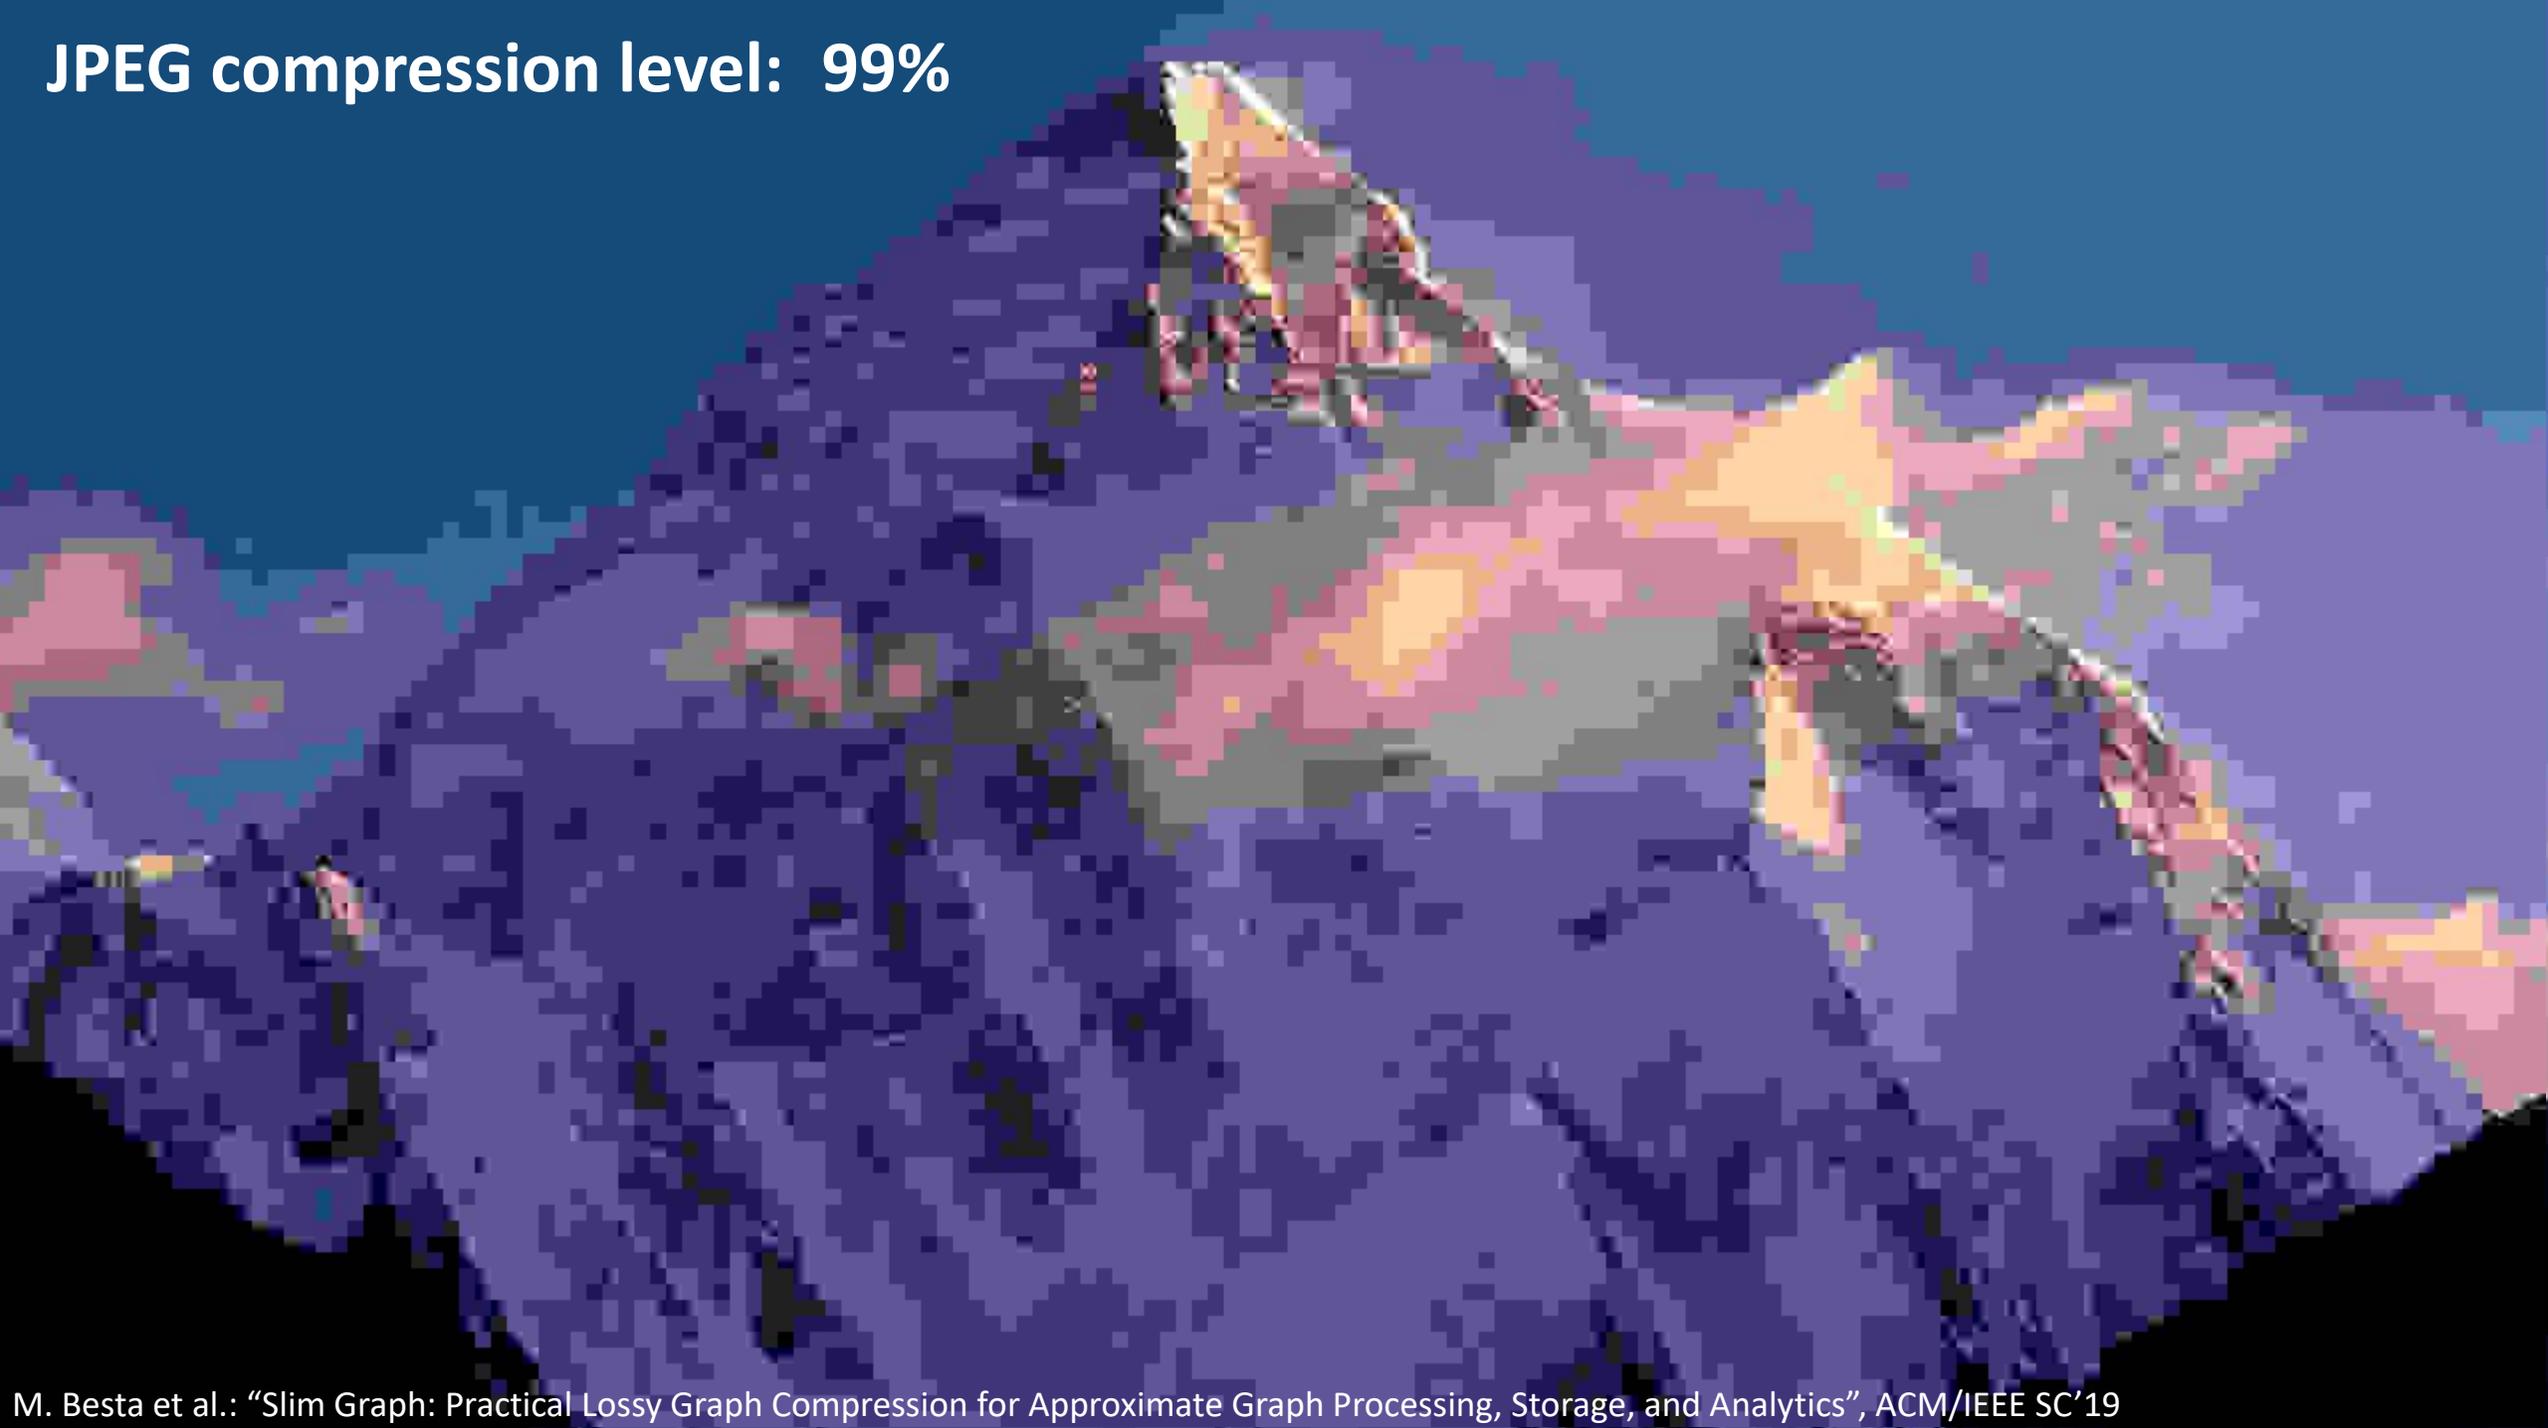
JPEG compression level: 50%



JPEG compression level: 90%



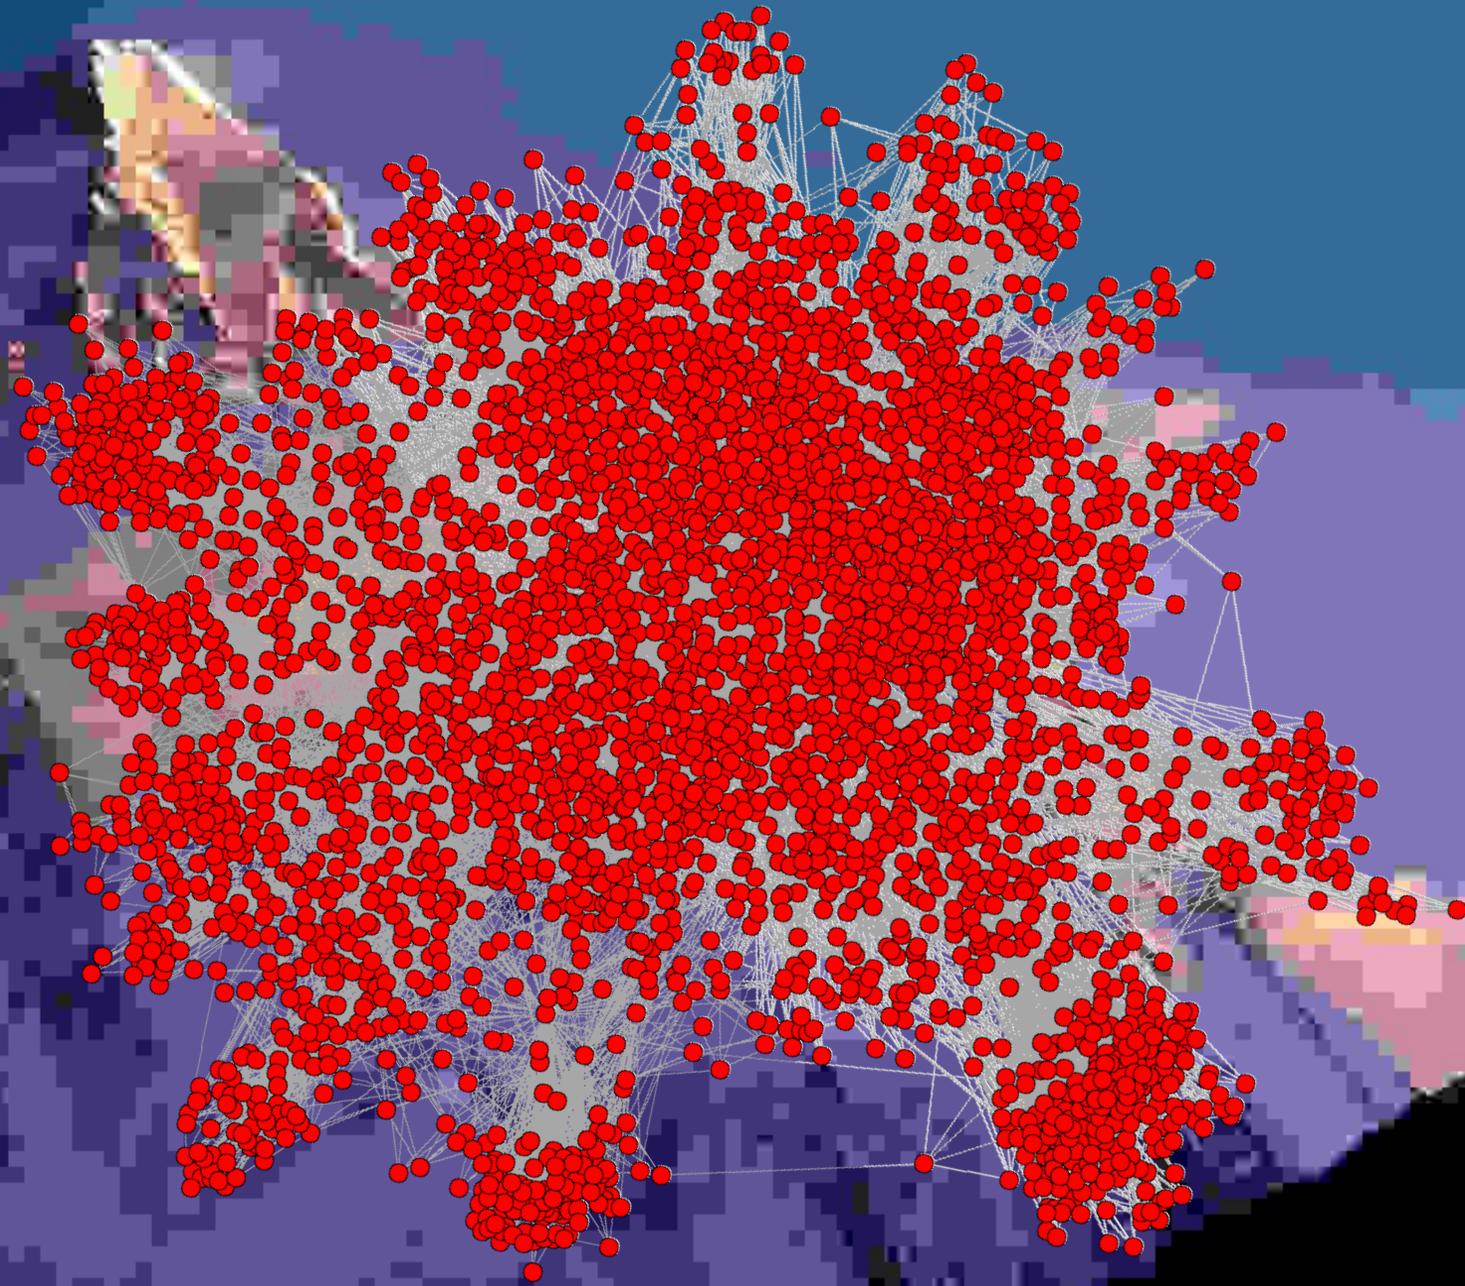
JPEG compression level: 99%



JPEG compression level: 99%



Can we apply a similar reasoning to graphs? How?

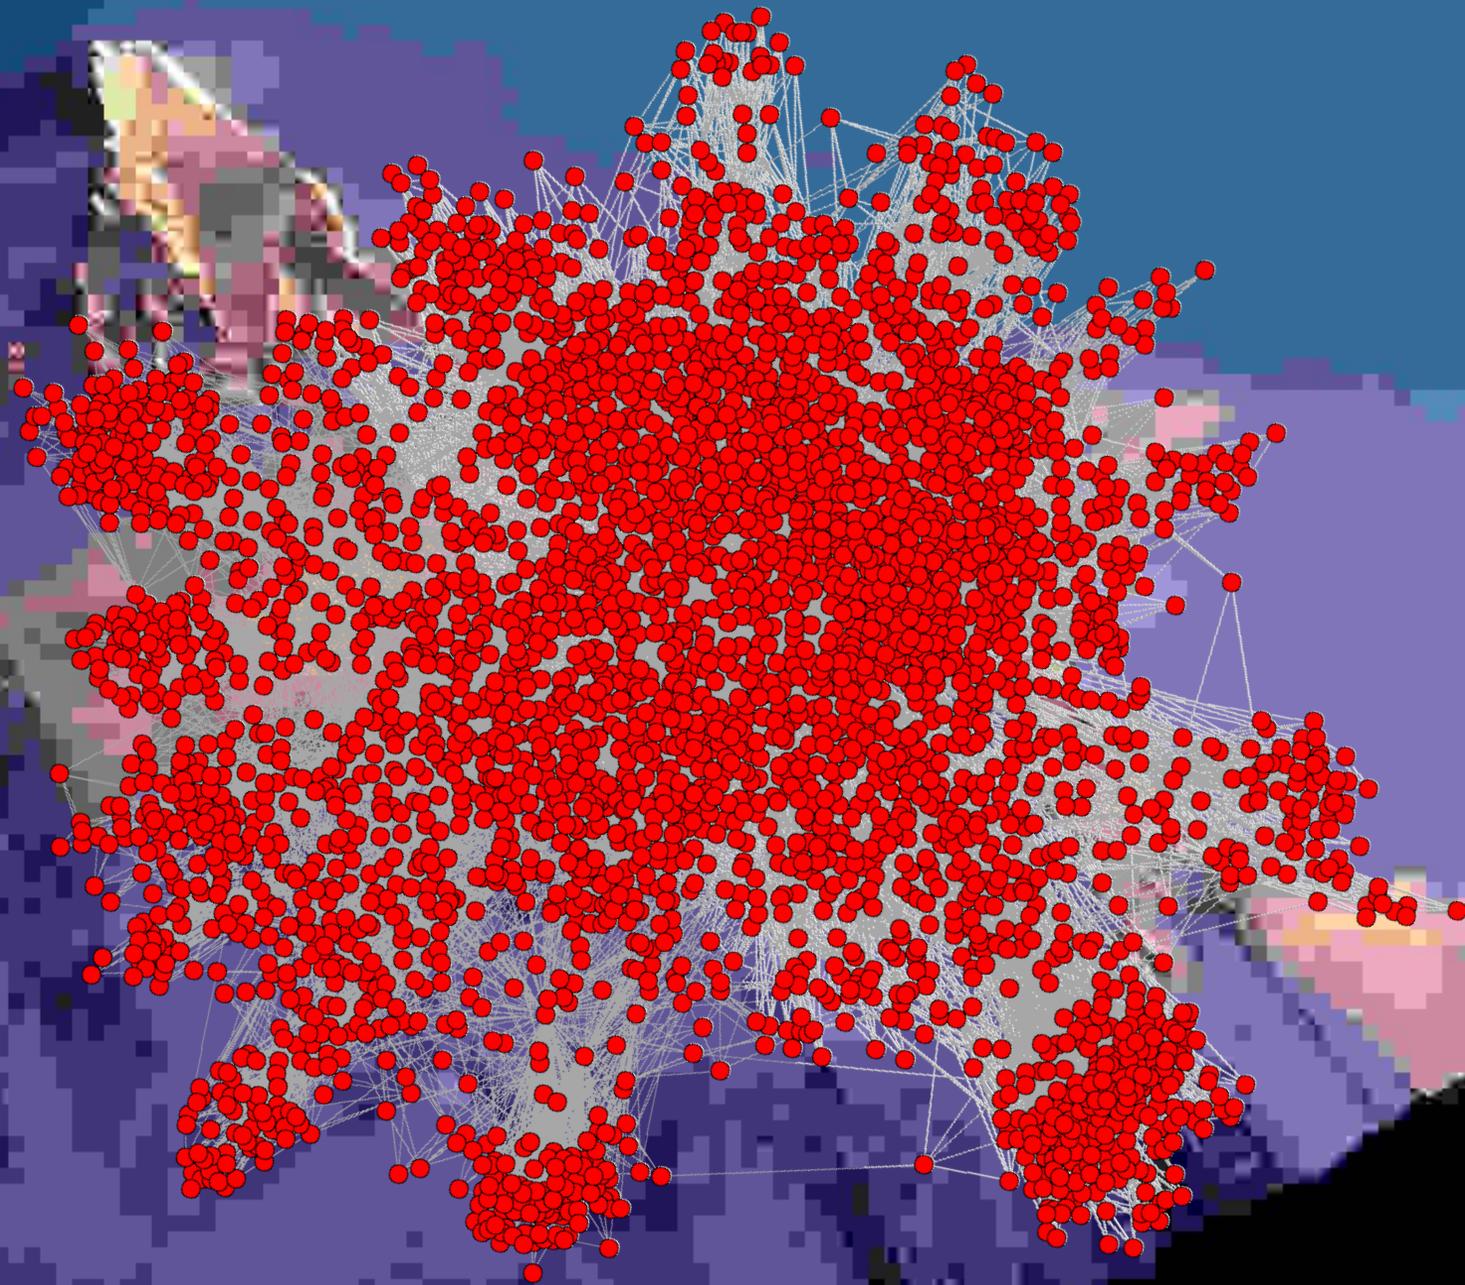


JPEG compression level: 99%



Can we apply a similar reasoning to graphs? How?

What should we pay attention to? (there is no “visual similarity” measure in this case...)



JPEG compression level: 99%

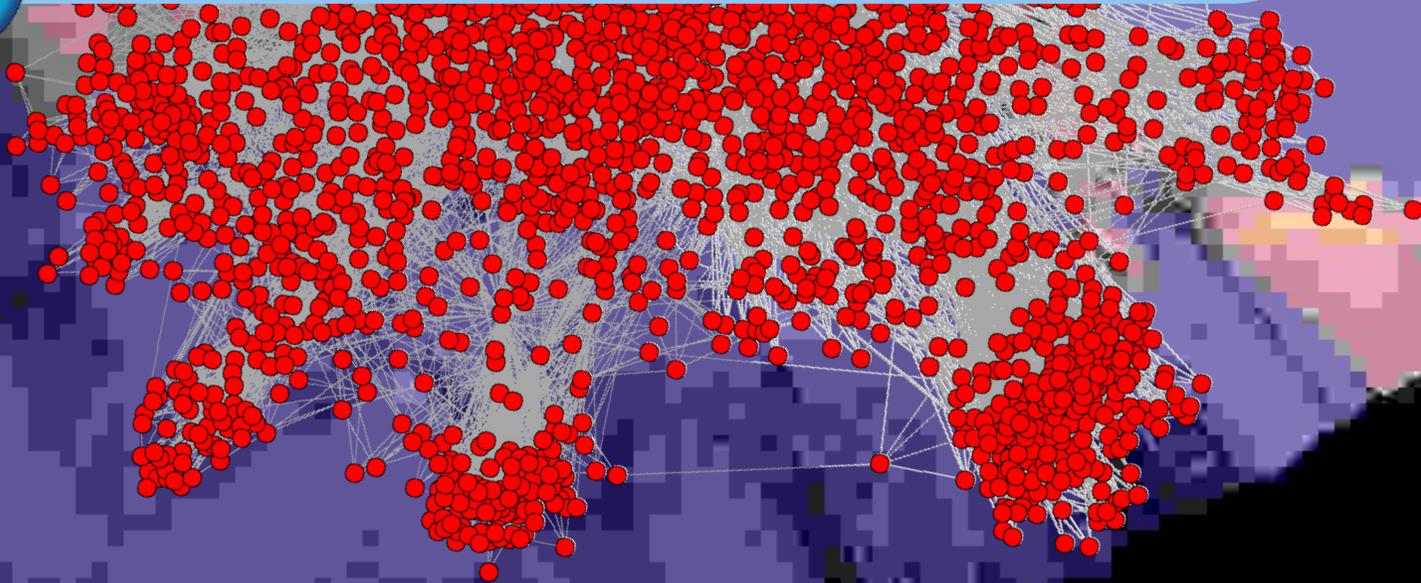


Can we apply a similar reasoning to graphs? How?

What should we pay attention to? (there is no “visual similarity” measure in this case...)



There are many theoretical works into sparsifying graphs (spanners, spectral sparsifiers, cut sparsifiers, ...). How to efficiently develop, use, and compare them, and which ones to select?





Can we apply a similar reasoning to graphs? How?



There are many theoretical works into sparsifying graphs (spanners, spectral sparsifiers, cut sparsifiers, ...). How to efficiently develop, use, and compare them, and which ones to select?



What should we pay attention to? (there is no “visual similarity” measure in this case...)

Slim Graph: A programming model and framework for effective lossy graph compression

Slim Graph Overview

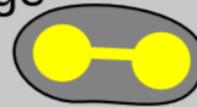
Programming Model

```
// Example kernel:  
atomic reduce_triangle(...) {  
  // Remove an edge from  
  // a triangle with a given  
  // probability  
}
```

A developer specifies
compression kernels
that remove selected
parts of a graph

Kernels focus on:

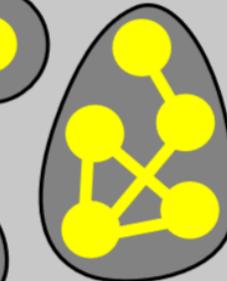
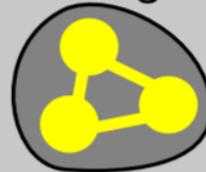
Edge



Vertex



Triangle



Subgraph

Slim Graph Overview

Programming Model

Kernels focus on:

```
// Example kernel:
atomic reduce_triangle(...) {
  // Remove an edge from
  // a triangle with a given
  // probability
}
```

A developer specifies compression kernels that remove selected parts of a graph

Compilation

Processing Engine

In stage 1, compression kernels are executed in parallel to compress graphs

I/O engines (e.g., GraphChi) can be used to compress very large graphs

In stage 2, algorithms compute on compressed graphs

Graphs are processed in-memory; if they do not fit, Slim Graph uses distributed or I/O engines

Slim Graph Overview

Programming Model

Kernels focus on:

```
// Example kernel:
atomic reduce_triangle(...) {
  // Remove an edge from
  // a triangle with a given
  // probability
}
```

A developer specifies compression kernels that remove selected parts of a graph

Compilation

Processing Engine

In stage 1, compression kernels are executed in parallel to compress graphs

I/O engines (e.g., GraphChi) can be used to compress very large graphs

In stage 2, algorithms compute on compressed graphs

Graphs are processed in-memory; if they do not fit, Slim Graph uses distributed or I/O engines

Analytics Subsystem & Accuracy Metrics

Generation of graphs

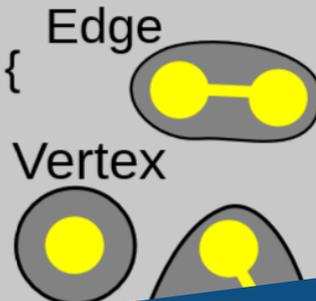
Slim Graph Overview

Programming Model

```
// Example kernel:
atomic reduce_triangle(...) {
  // Remove an edge from
  // a triangle with a given
  // probability
}
```

A developer specifies compression kernels that remove selected parts of a graph

Kernels focus on:



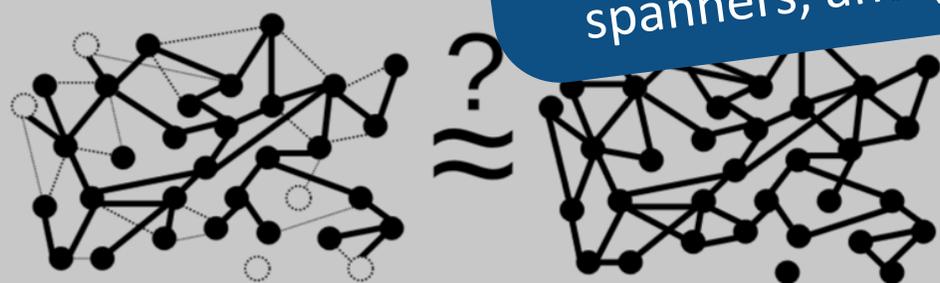
Compilation

In **stage 1**, compression kernels are executed in parallel to compress graphs

I/O engines (e.g., GraphChi) can be used to compress very large graphs

“Compression kernels”: an abstraction that enables expressing fundamental classes of lossy compression: (1) sampling, (2) spectral sparsifiers, (3) spanners, and (4) lossy summaries.

Analytics Subsystem



In **stage 2**, algorithms compute on compressed graphs

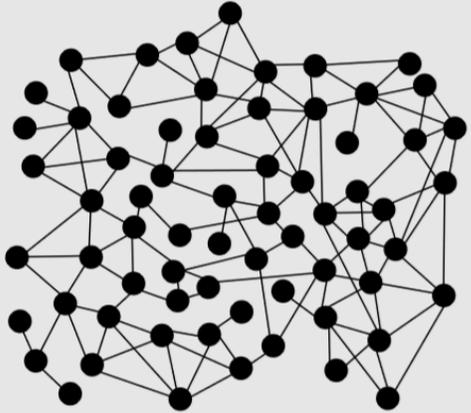
Graphs are processed in-memory; if they do not fit, Slim Graph uses distributed or I/O engines

Generation of graphs

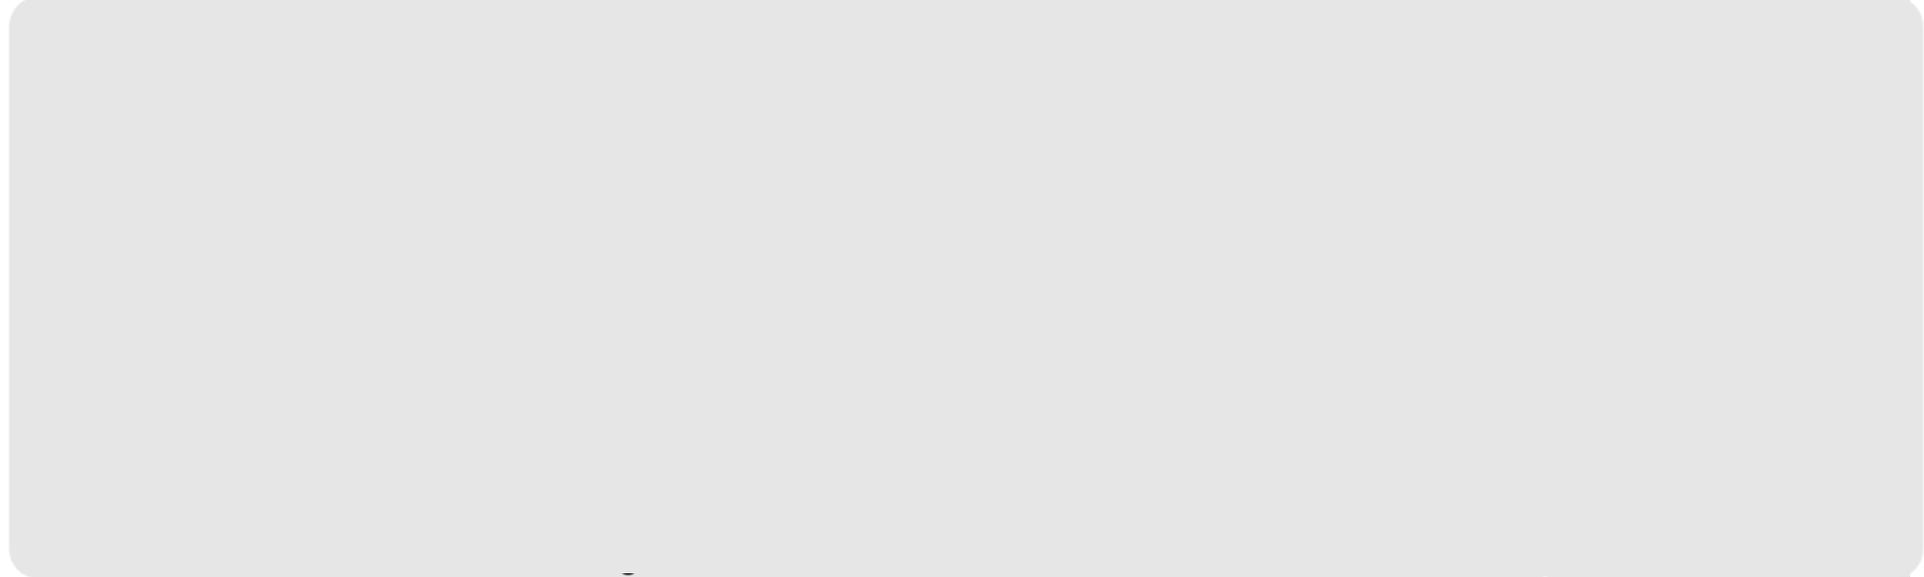
Example Compression Kernel: Triangle Reduction

Input graph:

$n = 67$
 $m = 132$



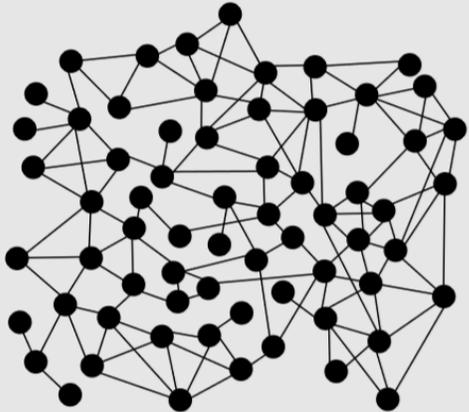
(§ 4.3) Triangle Compression Kernels (implementing Triangle Reduction, a novel graph compression method proposed in this work):



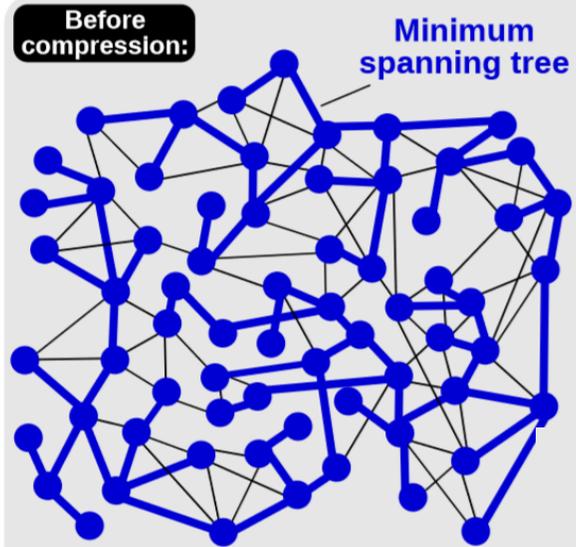
Example Compression Kernel: Triangle Reduction

Input graph:

$n = 67$
 $m = 132$



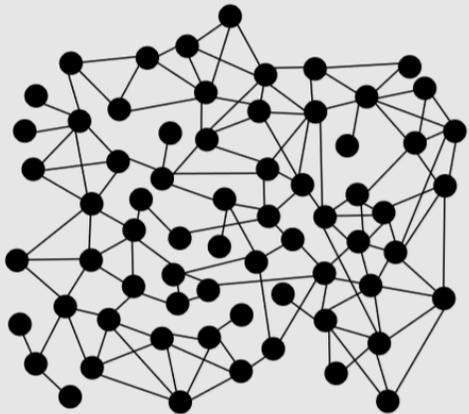
(§ 4.3) Triangle Compression Kernels (implementing Triangle Reduction, a novel graph compression method proposed in this work):



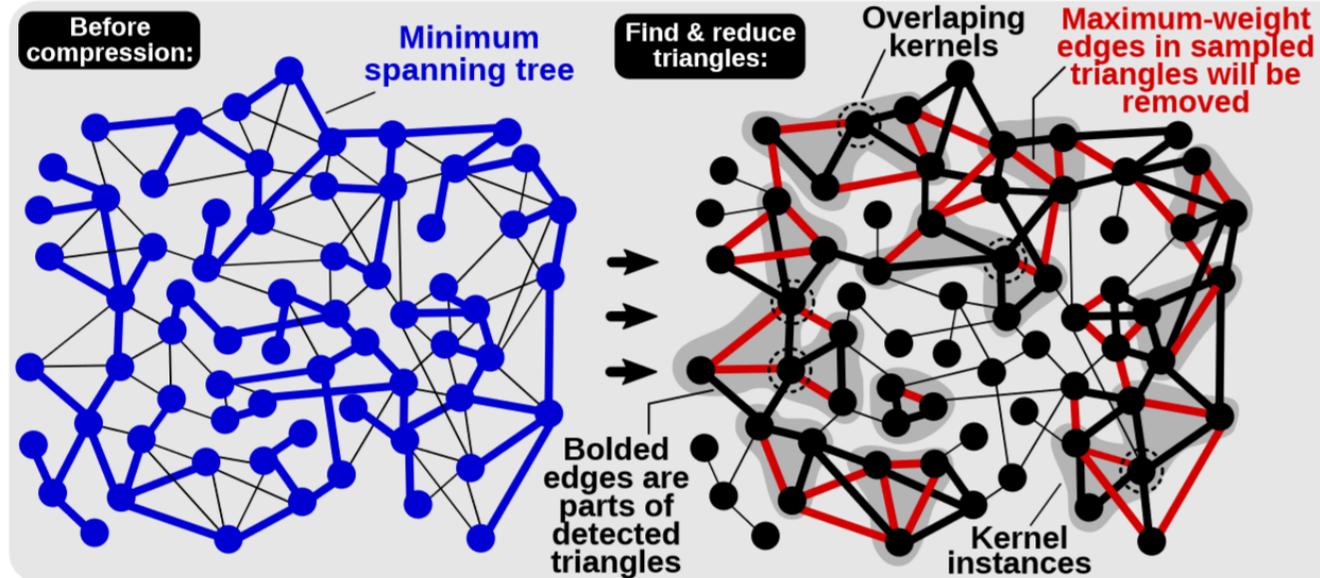
Example Compression Kernel: Triangle Reduction

Input graph:

$n = 67$
 $m = 132$



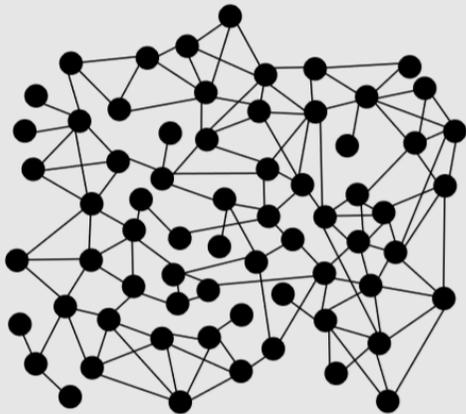
(§ 4.3) Triangle Compression Kernels (implementing Triangle Reduction, a novel graph compression method proposed in this work):



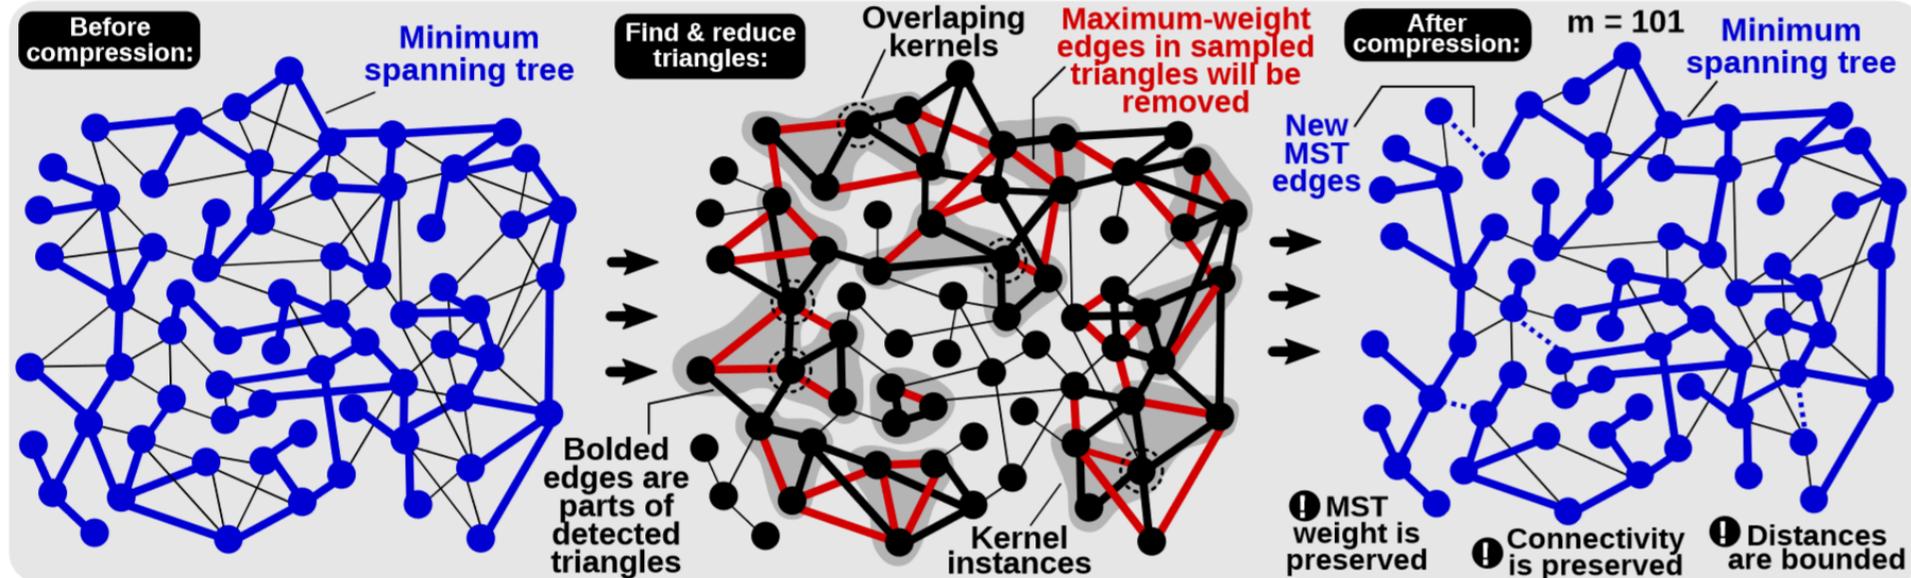
Example Compression Kernel: Triangle Reduction

Input graph:

$n = 67$
 $m = 132$

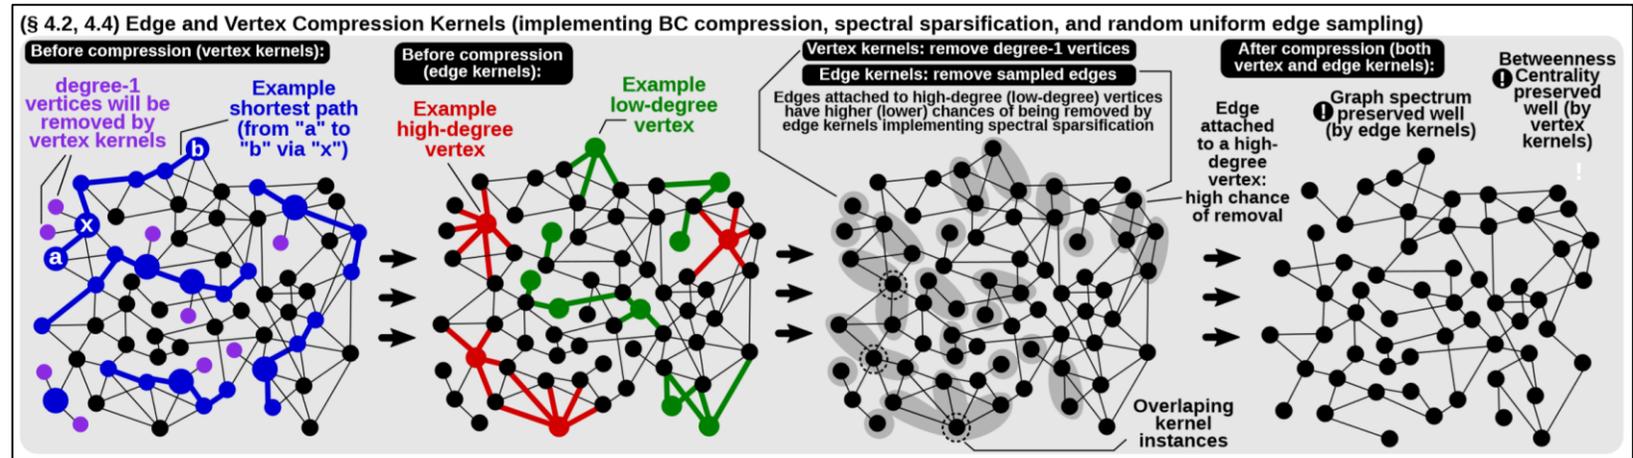


(§ 4.3) Triangle Compression Kernels (implementing Triangle Reduction, a novel graph compression method proposed in this work):

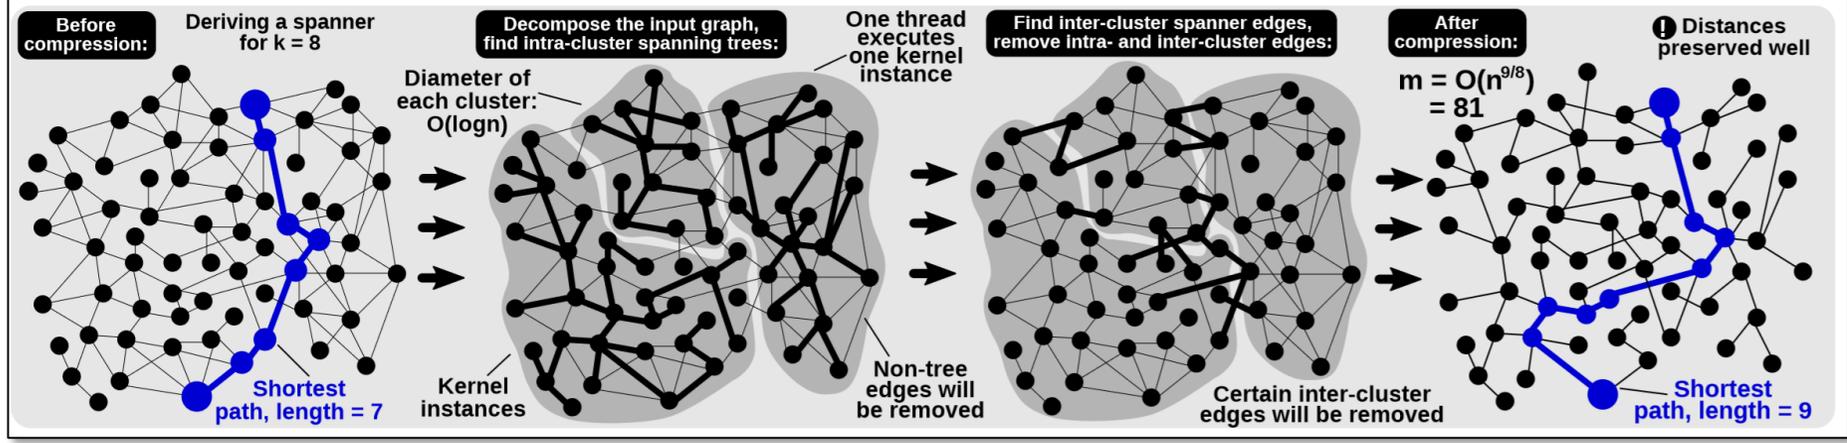
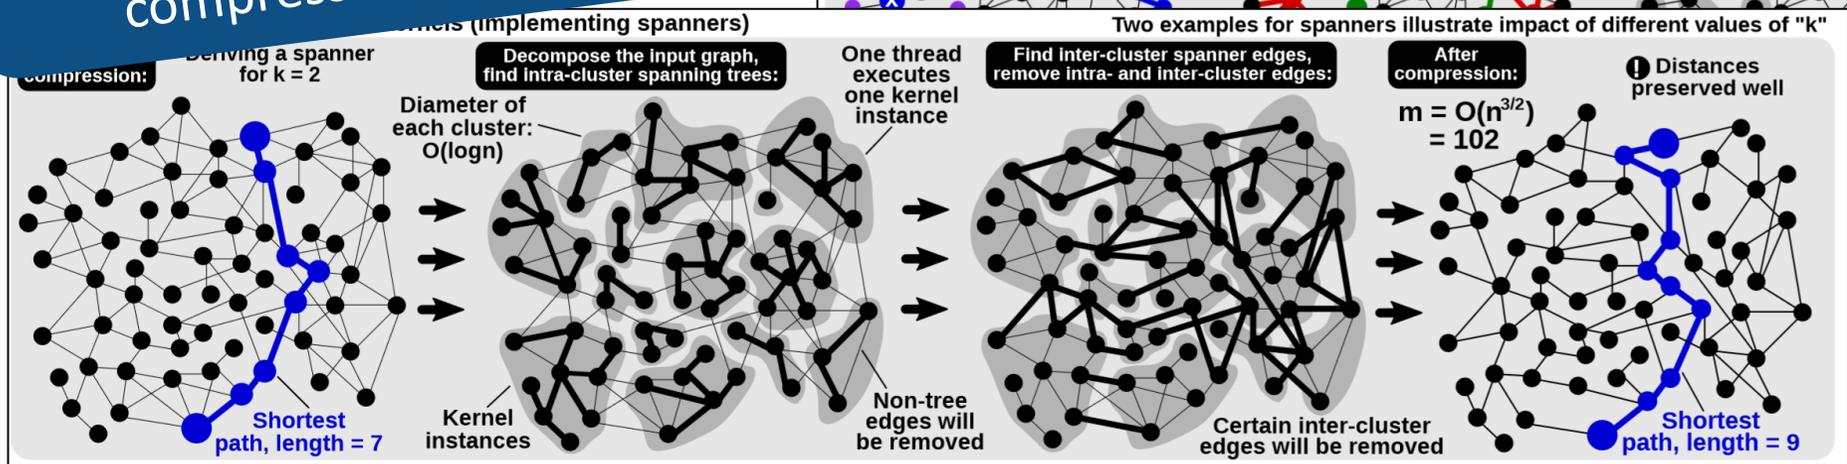
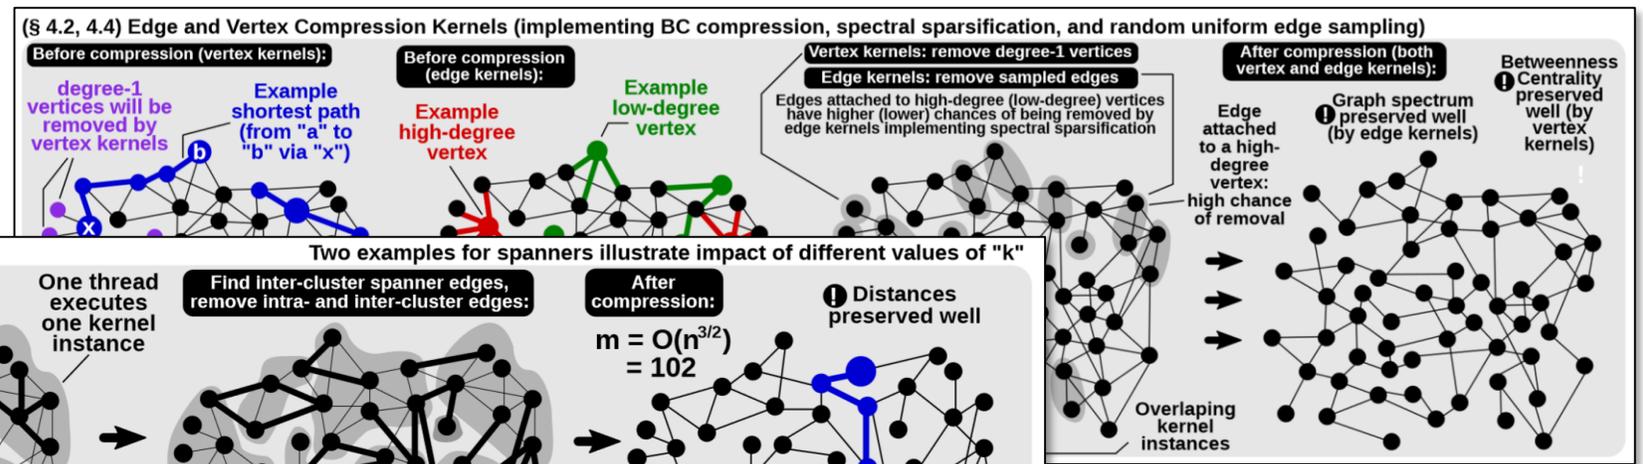


Other types of compression kernels enable expressing and implementing other lossy graph compression schemes

Other types of compression kernels enable expressing and implementing other lossy graph compression schemes

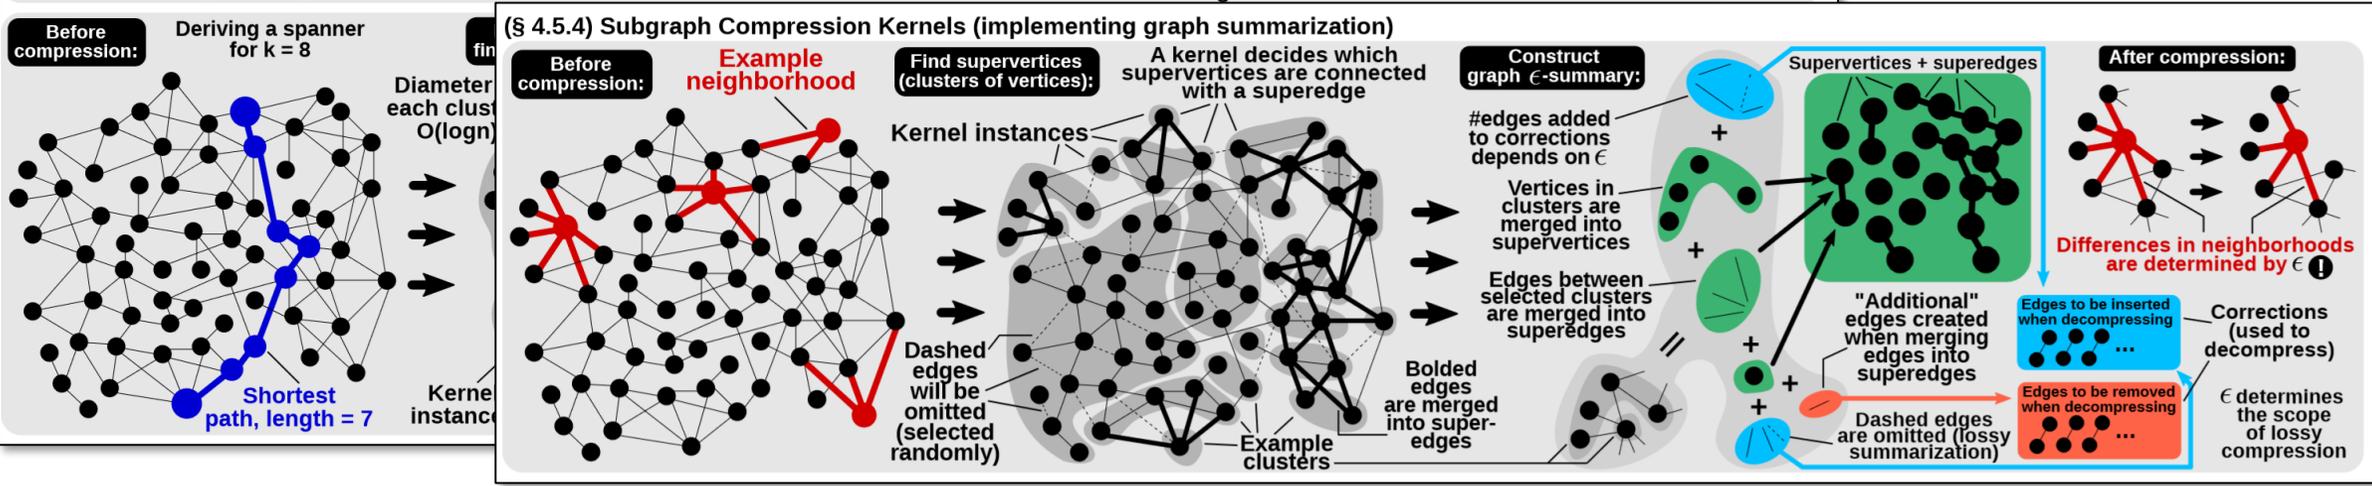
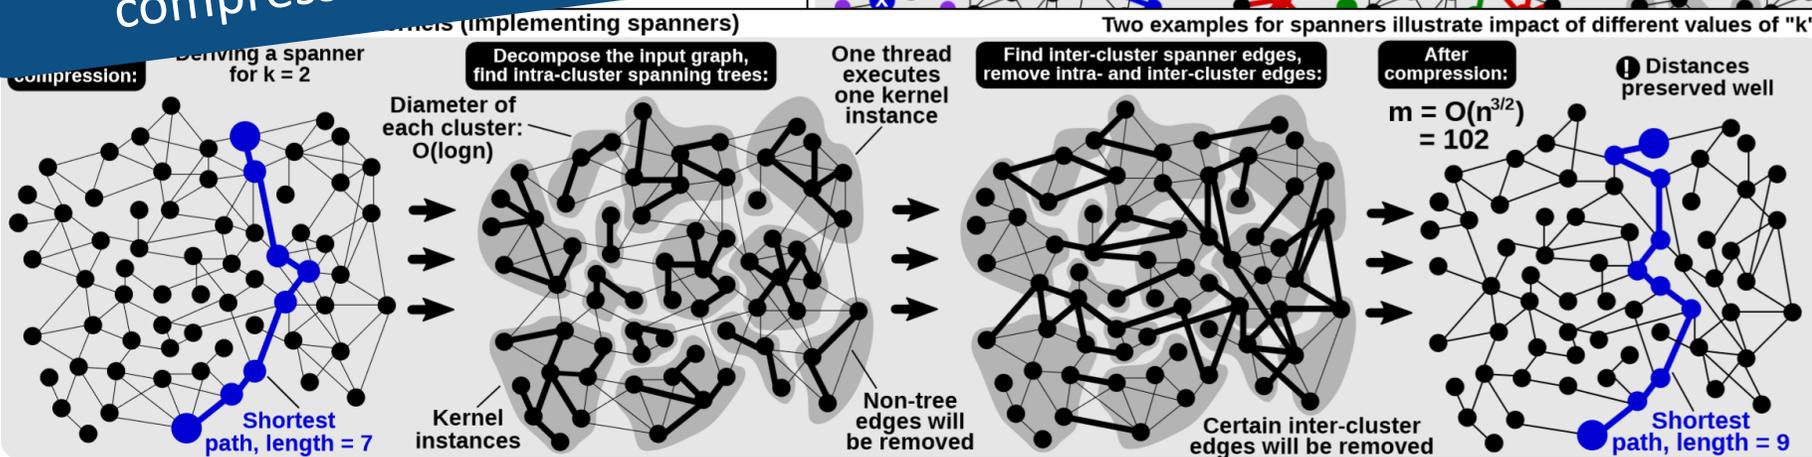
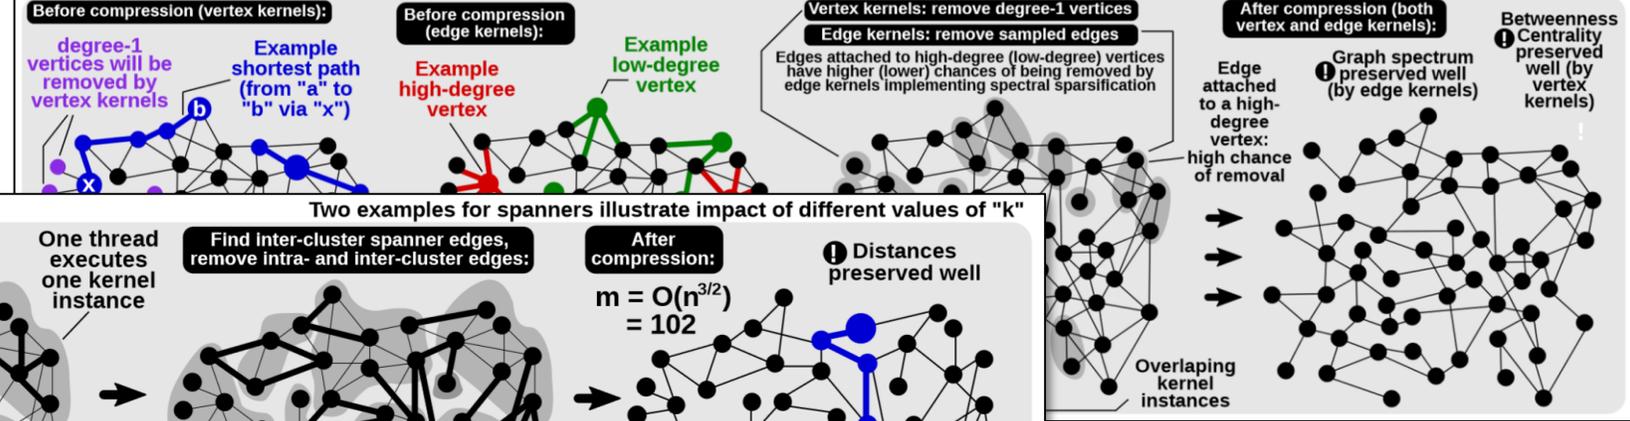


Other types of compression kernels enable expressing and implementing other lossy graph compression schemes

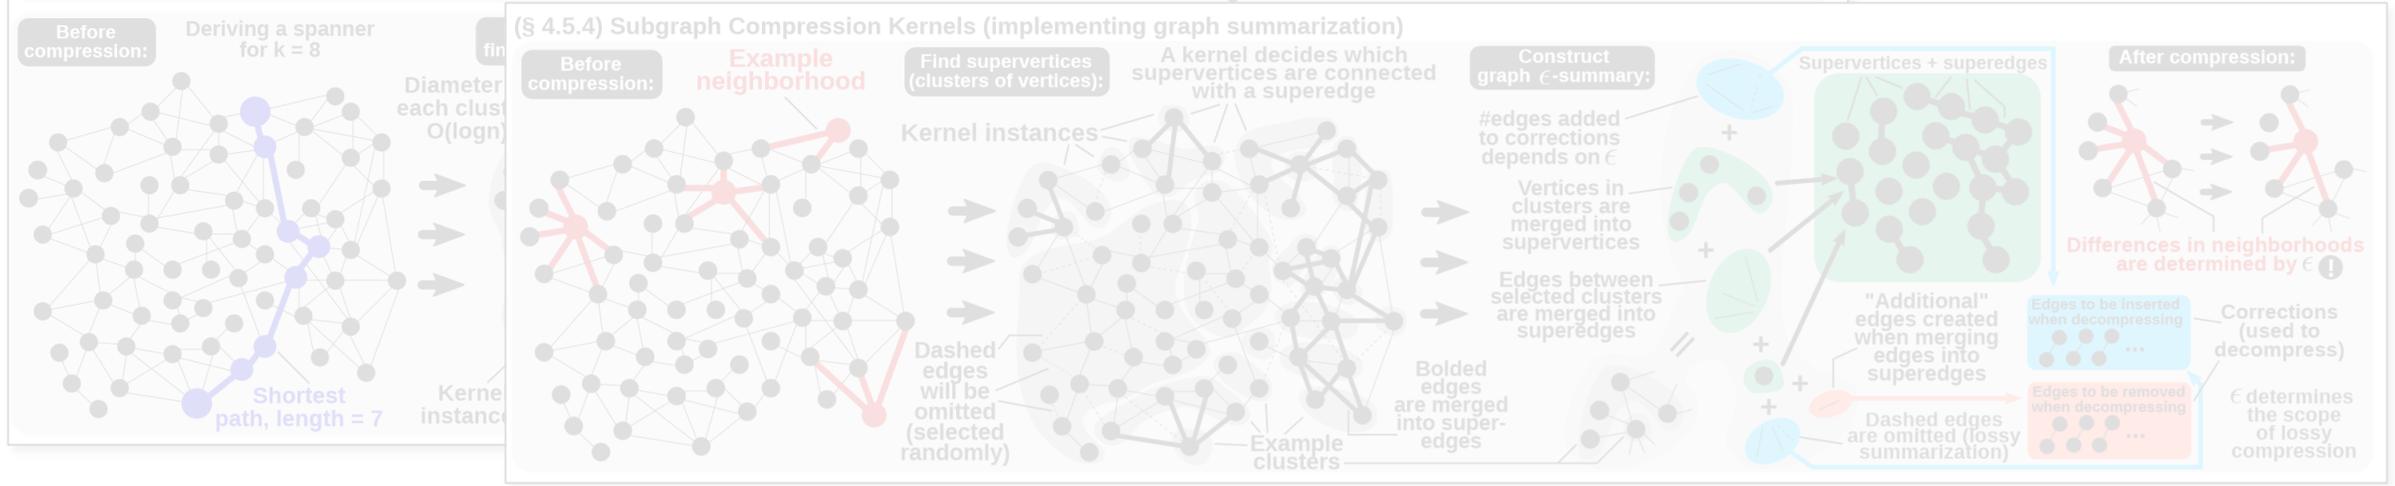
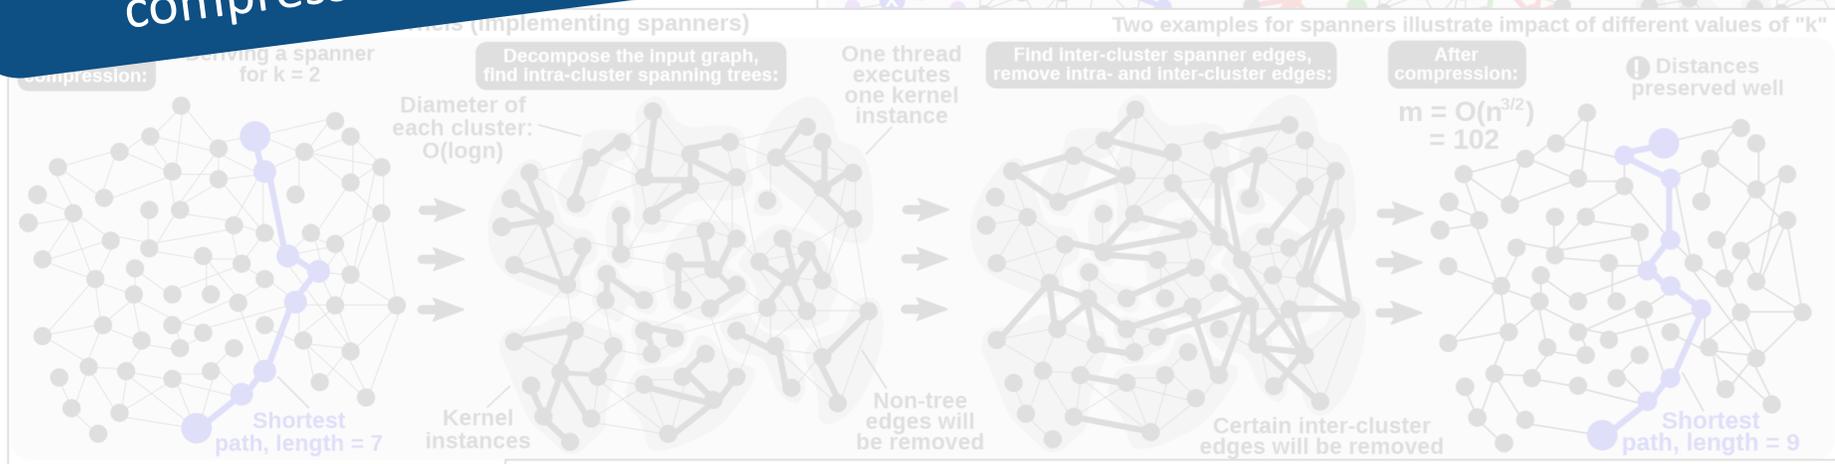
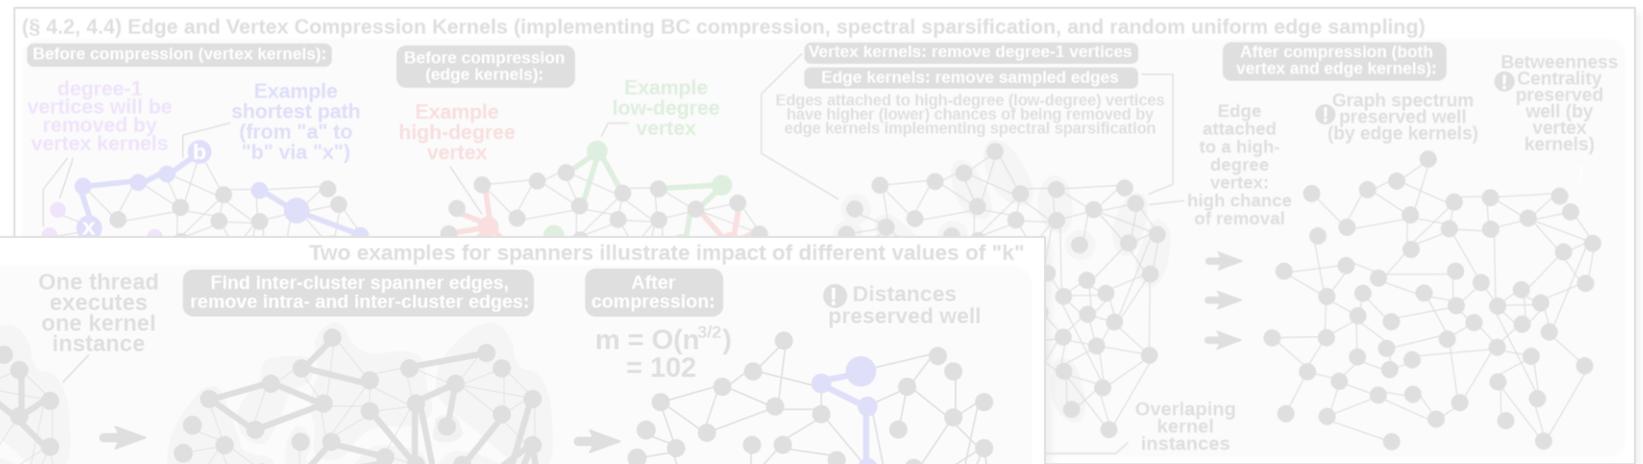


Other types of compression kernels enable expressing and implementing other lossy graph compression schemes

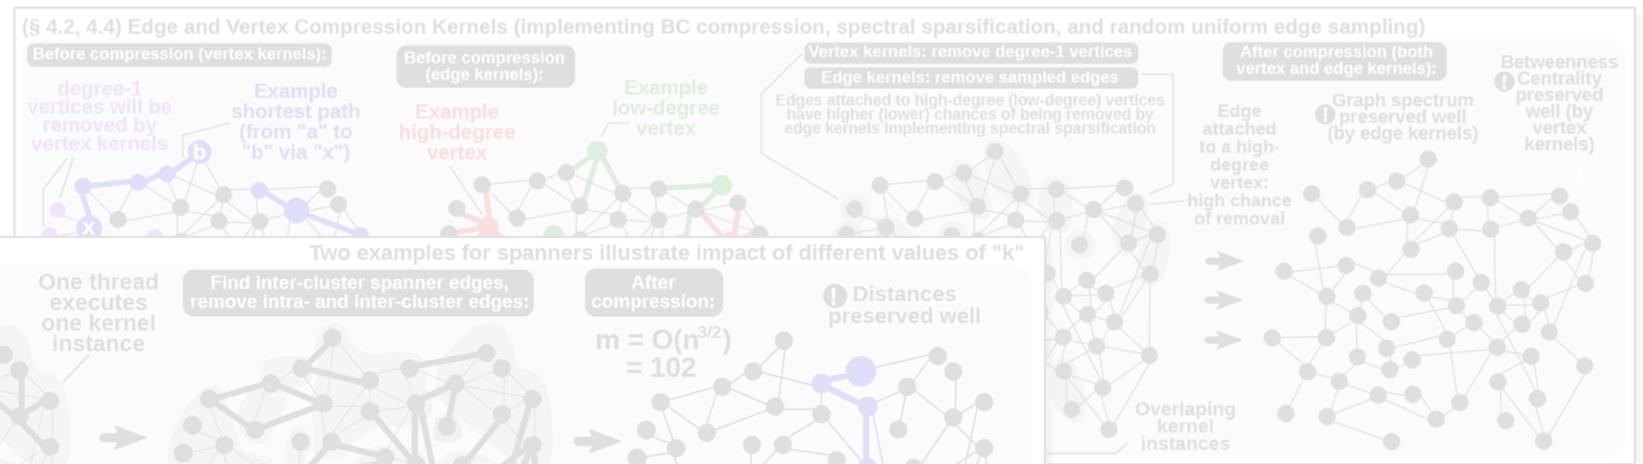
(§ 4.2, 4.4) Edge and Vertex Compression Kernels (implementing BC compression, spectral sparsification, and random uniform edge sampling)



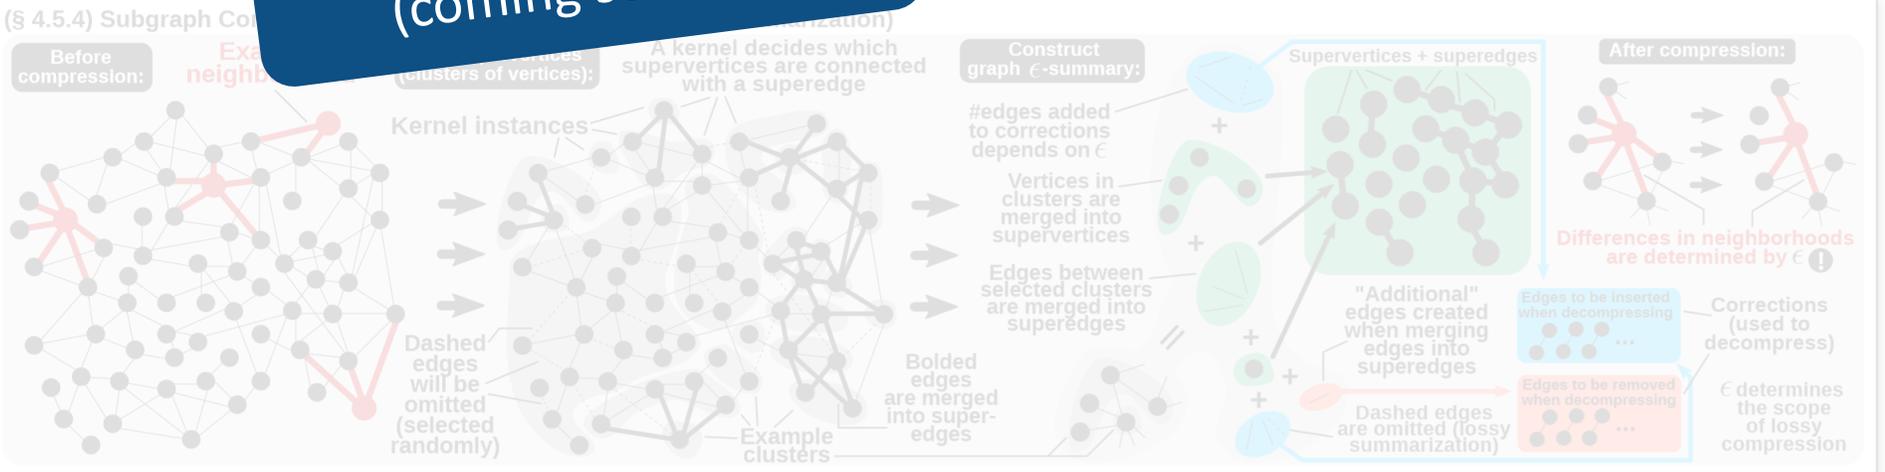
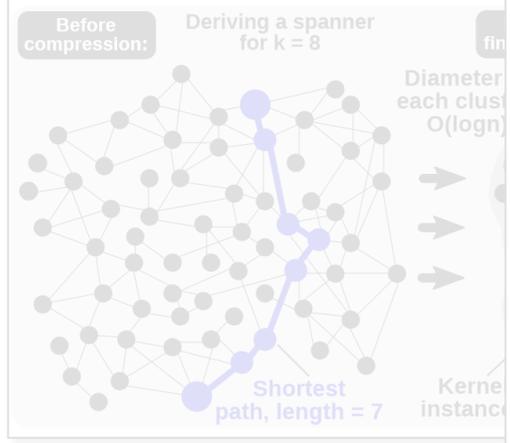
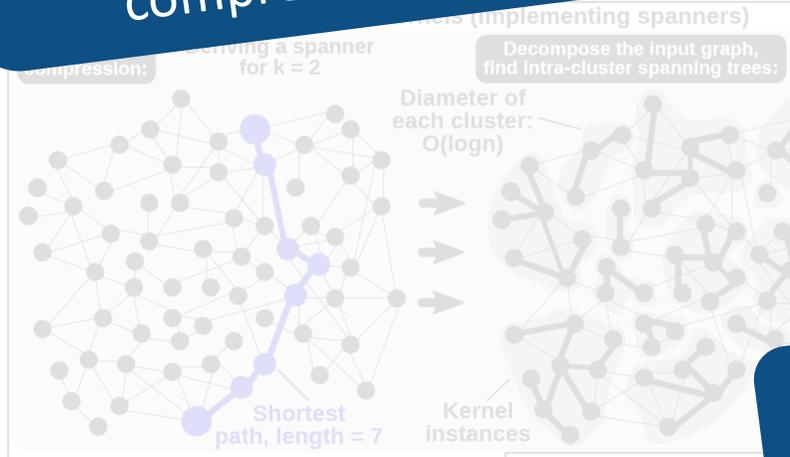
Other types of compression kernels enable expressing and implementing other lossy graph compression schemes



Other types of compression kernels enable expressing and implementing other lossy graph compression schemes

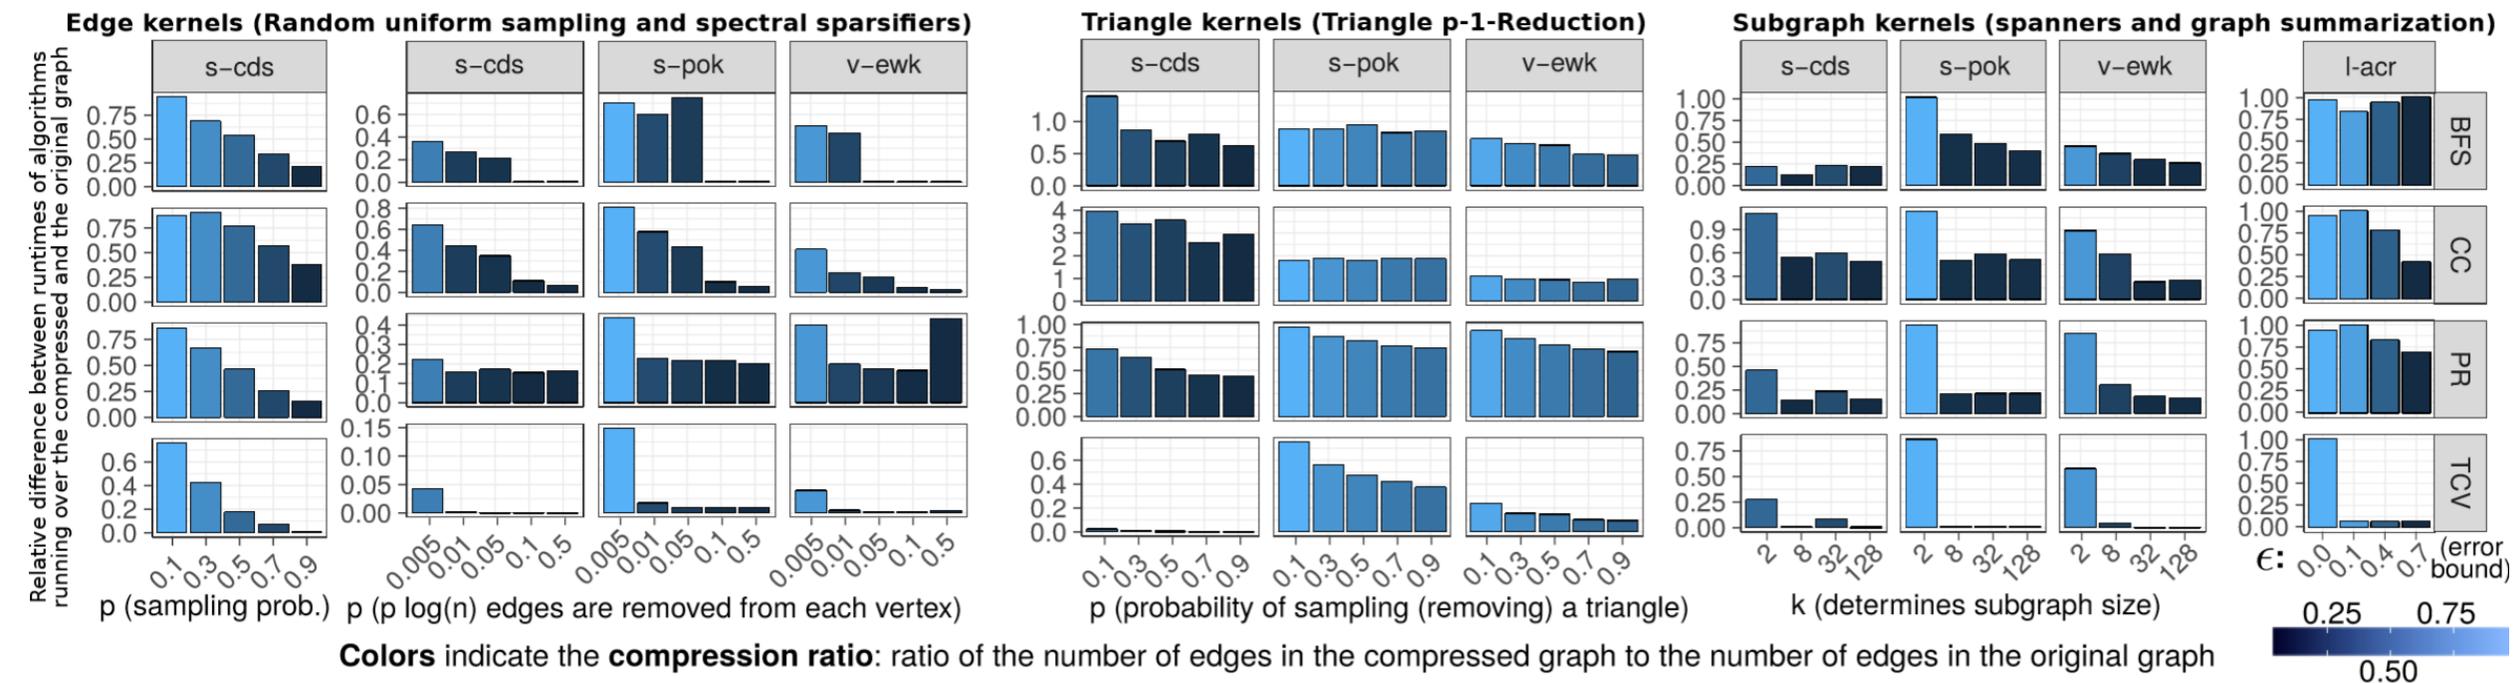


Details in the paper (coming soon) 😊



Storage Reductions vs. Speedups vs. Accuracy Loss [vs. Compression Overhead]

Various real-world graphs are used



Storage Reductions vs. Speedups vs. Accuracy Loss [vs. Compression Overhead]

Various real-world graphs are used

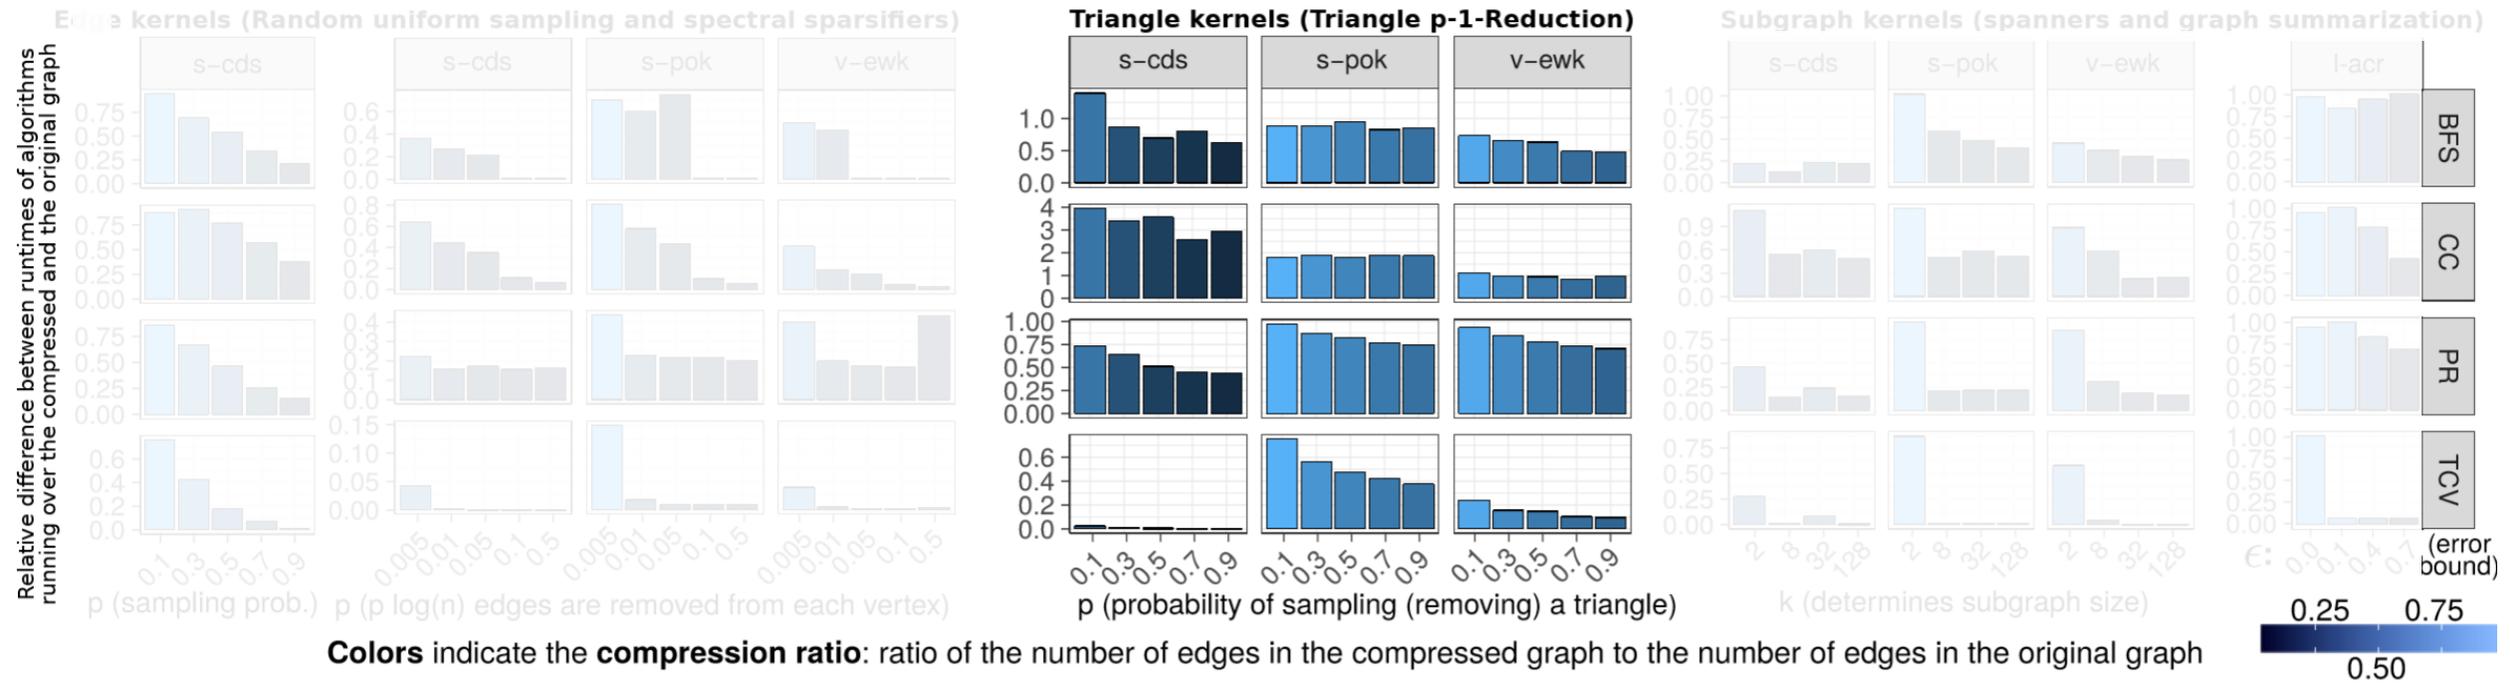


Figure 4: Analysis of **storage and performance** tradeoffs of various lossy compression schemes implemented in Slim Graph (when varying compression parameters).

Storage Reductions vs. Speedups vs. Accuracy Loss [vs. Compression Overhead]

Various real-world graphs are used

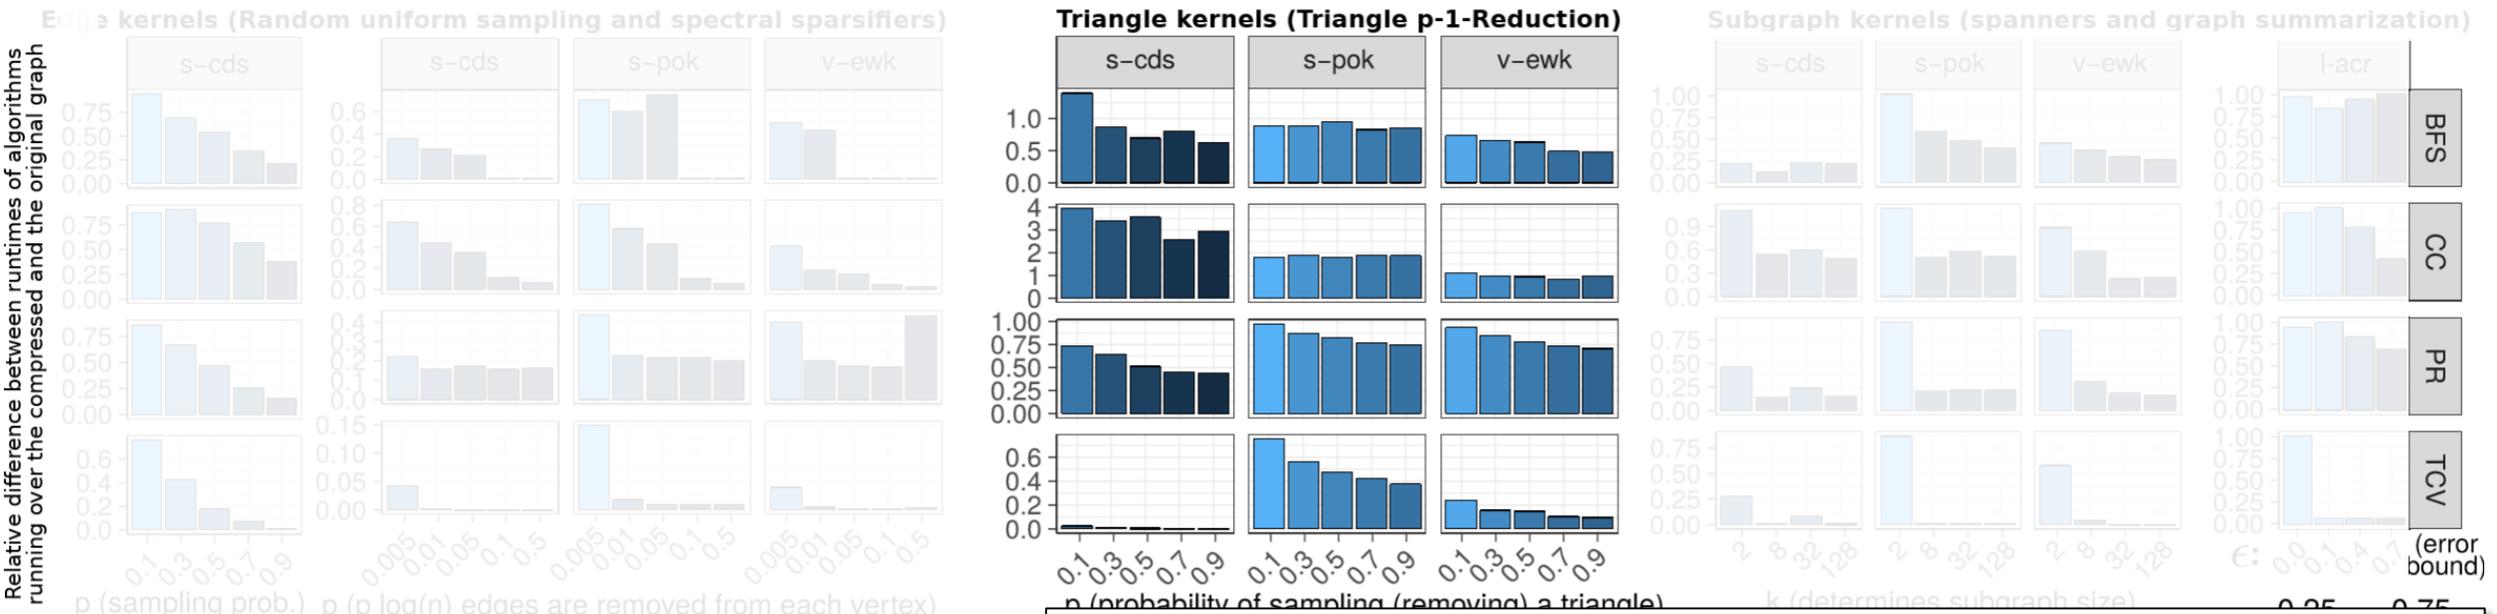
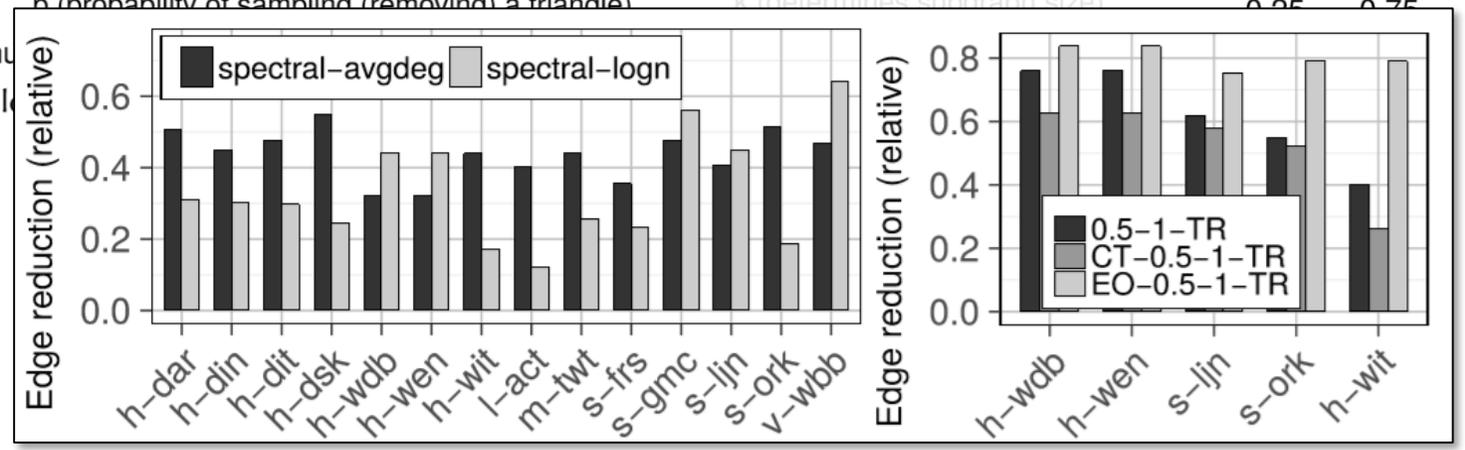


Figure 4: Analysis of storage and performance tradeoffs of various kernels



Storage Reductions vs. Speedups vs. Accuracy Loss [vs. Compression Overhead]

Various real-world graphs are used

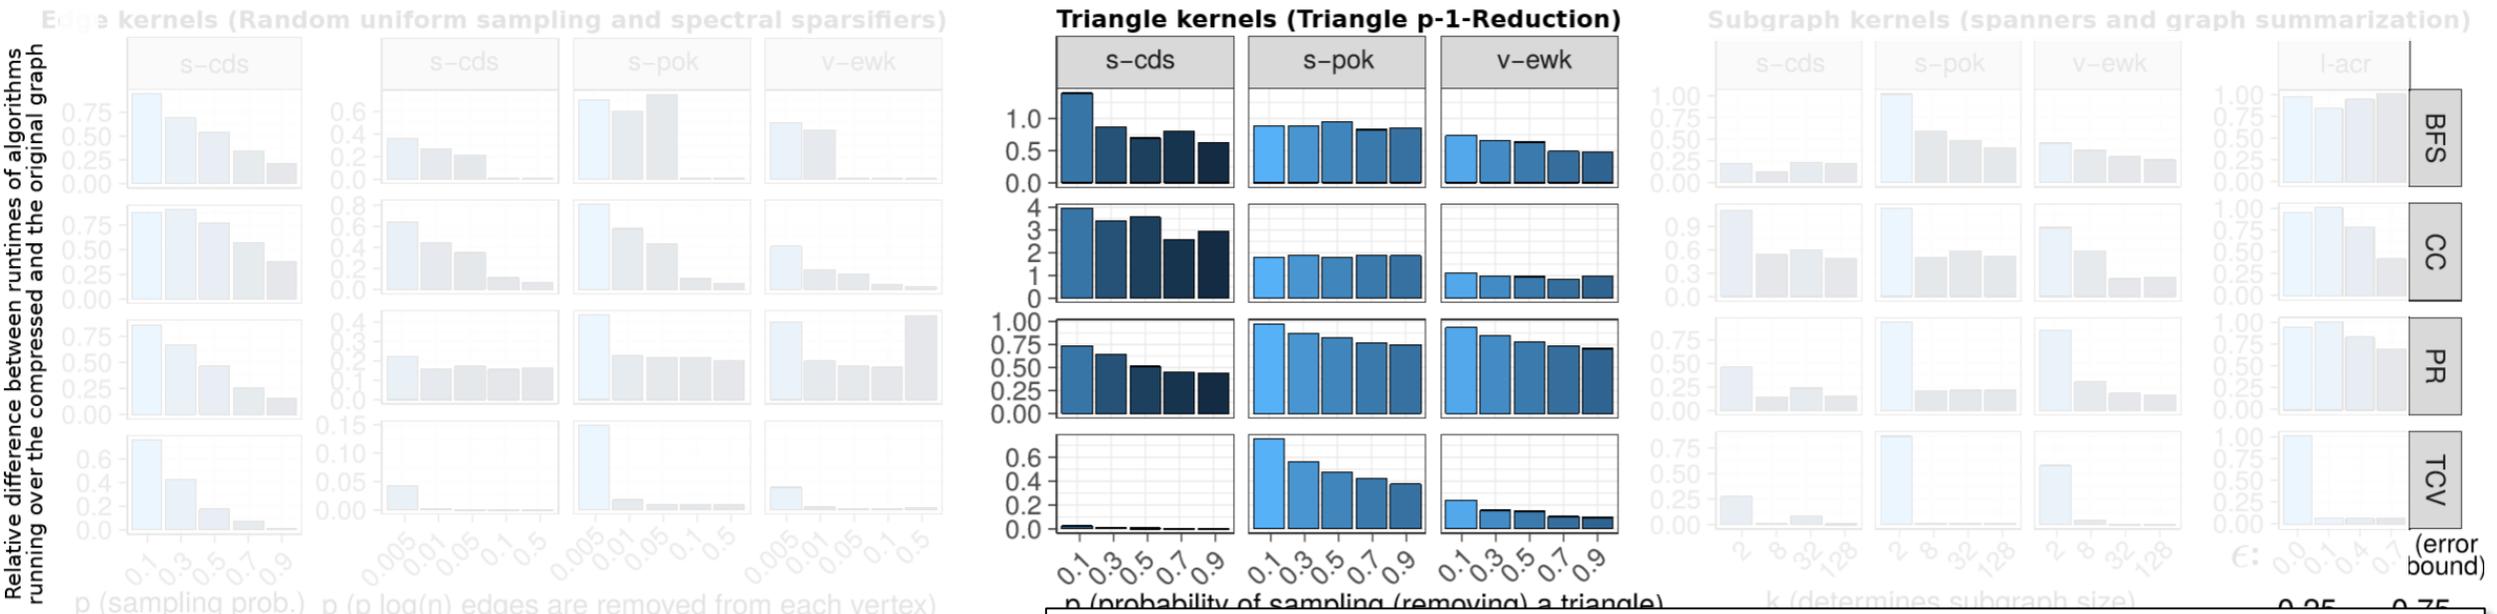
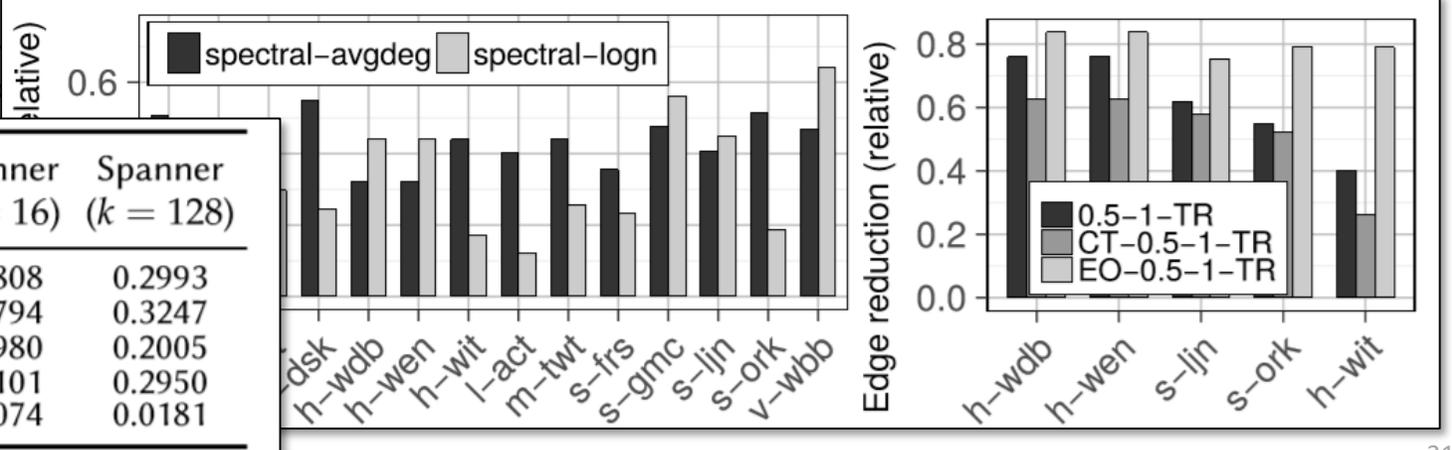


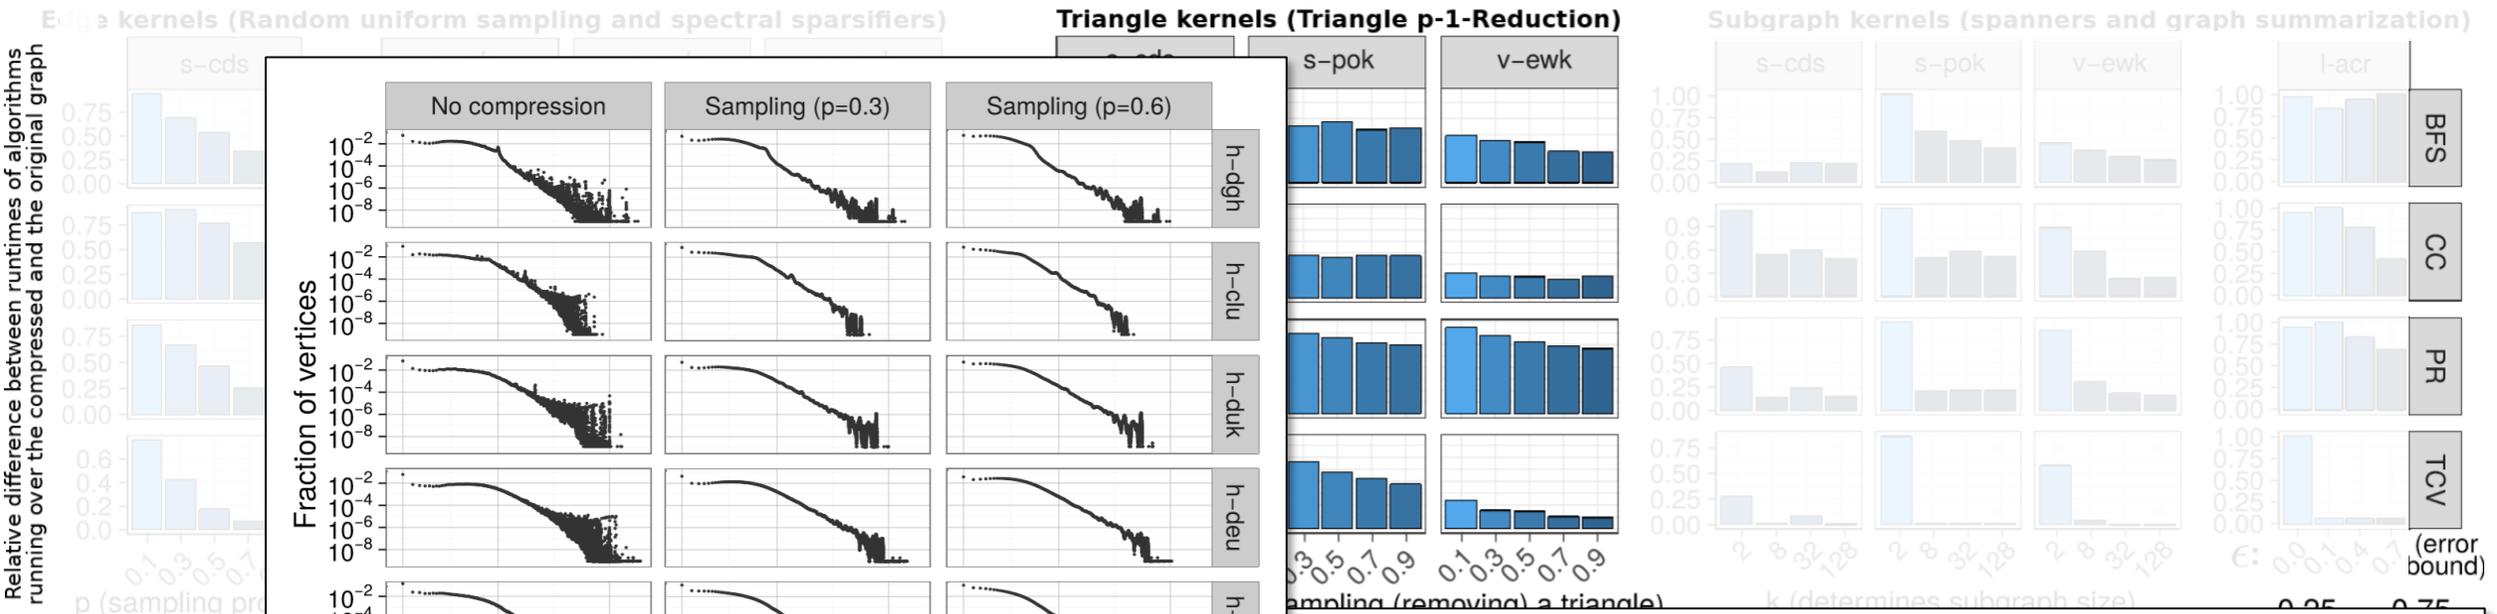
Figure 4: Analysis of storage and performance tradeoffs of various kernels

Graph	EO		Uniform		Spanner		
	0.8-1-TR	1.0-1-TR	($p = 0.2$)	($p = 0.5$)	($k = 2$)	($k = 16$)	($k = 128$)
s-you	0.0121	0.0167	0.1932	0.6019	0.0054	0.2808	0.2993
h-hud	0.0187	0.0271	0.0477	0.1633	0.0340	0.2794	0.3247
il-dbl	0.0459	0.0674	0.0749	0.2929	0.0080	0.1980	0.2005
v-skt	0.0410	0.0643	0.0674	0.2695	0.0311	0.1101	0.2950
v-usa	0.0089	0.0100	0.1392	0.5945	0.0000	0.0074	0.0181



Storage Reductions vs. Speedups vs. Accuracy Loss [vs. Compression Overhead]

Various real-world graphs are used



Relative difference between runtimes of algorithms running over the compressed and the original graph

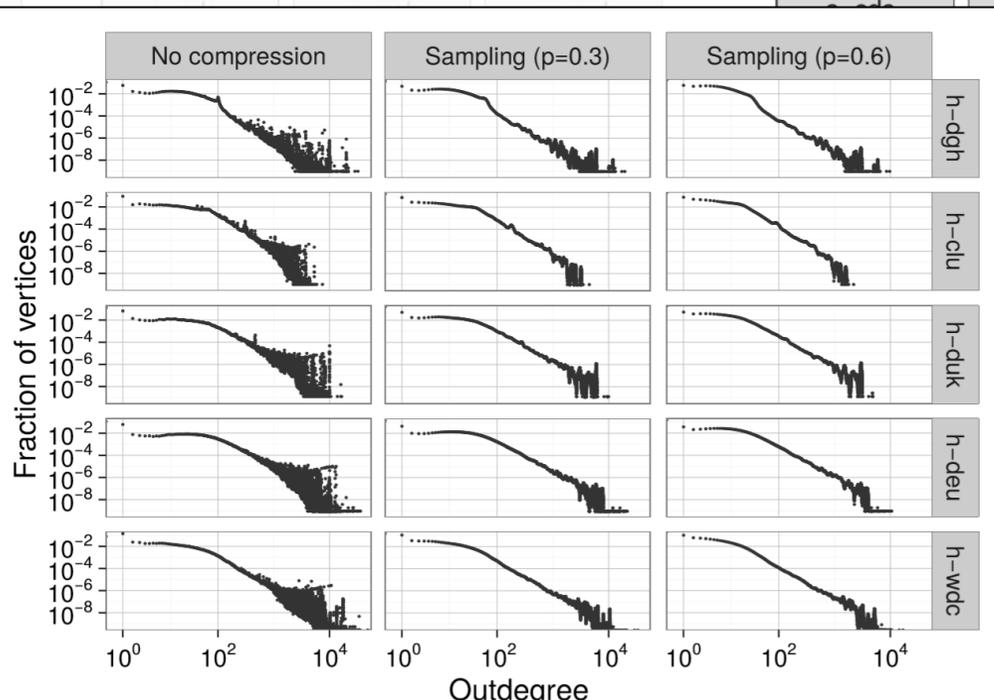
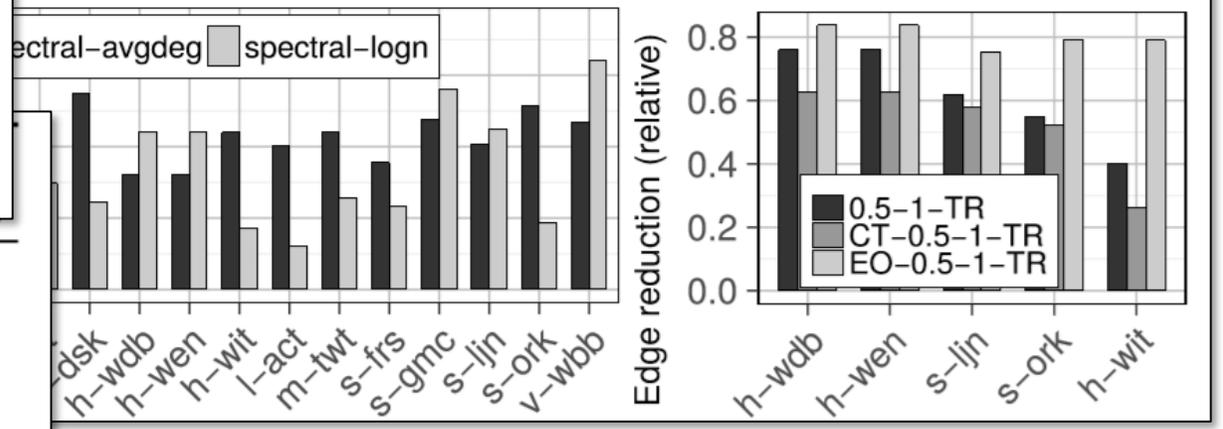


Figure 7: (Accuracy) Impact of random uniform sampling on the degree distribution of large graphs (the largest, h-wdc, has $\approx 128B$ edges).

Figure 4: Analysis of various kernels

Graph	EO	0.1	0.3	0.5	0.7	0.9	1.0
s-you	0.0121	0.0167	0.1932	0.6019	0.0054	0.2808	0.2993
h-hud	0.0187	0.0271	0.0477	0.1633	0.0340	0.2794	0.3247
il-dbl	0.0459	0.0674	0.0749	0.2929	0.0080	0.1980	0.2005
v-skt	0.0410	0.0643	0.0674	0.2695	0.0311	0.1101	0.2950
v-usa	0.0089	0.0100	0.1392	0.5945	0.0000	0.0074	0.0181



Storage Reductions vs. Speedups vs. Accuracy Loss [vs. Compression Overhead]

Various real-world graphs are used

	$ V $	$ E $	Shortest s - t path length	Average path length	Diameter	Average degree	Maximum degree	#Triangles	#Connected components	Chromatic number	Max. indep. set size	Max. cardinal. matching size
	n	m	\mathcal{P}	\bar{P}	D	\bar{d}	d	T	\mathcal{C}	C_R	\hat{I}_S	\hat{M}_C
Original graph	n	m	\mathcal{P}	\bar{P}	D	\bar{d}	d	T	\mathcal{C}	C_R	\hat{I}_S	\hat{M}_C
Lossy ϵ -summary	n	$m \pm 2\epsilon m$	$1, \dots, \infty$	$1, \dots, \infty$	$1, \dots, \infty$	$\bar{d} \pm \epsilon \bar{d}$	$d \pm \epsilon d$	$T \pm 2\epsilon m$	$\mathcal{C} \pm 2\epsilon m$	$C_R \pm 2\epsilon m$	$\hat{I}_S \pm 2\epsilon m$	$\hat{M}_C \pm 2\epsilon m$
Simple p -sampling	n	$(1-p)m$	∞	∞	∞	$(1-p)\bar{d}$	$(1-p)d$	$(1-p^3)T$	$\leq \mathcal{C} + pm$	$\geq C_R - pm$	$\leq \hat{I}_S + pm$	$\geq \hat{M}_C - pm$
Spectral ϵ -sparsifier	n	$\tilde{O}(n/\epsilon^2)$	$\leq n$	$\leq n$	$\leq n$	$\tilde{O}(1/\epsilon^2)$	$\geq d/2(1+\epsilon)$	$\tilde{O}(n^{3/2}/\epsilon^3)$	$\stackrel{w.h.p.}{=} \mathcal{C}$	$\leq d/2(1+\epsilon)$	$\geq 2(1+\epsilon)n/d$	≥ 0
$O(k)$ -spanner	n	$O(n^{1+1/k})$	$O(k\mathcal{P})$	$O(k\bar{P})$	$O(kD)$	$O(n^{1/k})$	$\leq d$	$O(n^{1+2/k})$	\mathcal{C}	$O(n^{1/k} \log n)$	$\Omega\left(\frac{n^{1-1/k}}{\log n}\right)$	≥ 0
EO p -1-Triangle Red.	n	$\leq m - \frac{pT}{3d}$	$\stackrel{w.h.p.}{\leq} \mathcal{P} + p\mathcal{P}$	$\leq \bar{P} + \frac{pT}{n(n-1)}$	$\stackrel{w.h.p.}{\leq} D + pD$	$\leq \bar{d} - \frac{pT}{dn}$	$\geq d/2$	$\leq (1 - \frac{p}{d})T$	\mathcal{C}	$\geq C_R - pT$	$\leq \hat{I}_S + pT$	$\geq \hat{M}_C/2$
remove k deg-1 vertices	$n - k$	$m - k$	\mathcal{P}	$\geq \bar{P} - \frac{kD}{n}$	$\geq D - 2$	$\geq \bar{d} - \frac{k}{n}$	d	T	\mathcal{C}	C_R	$\geq \hat{I}_S - k$	$\geq \hat{M}_C - k$

Table 3: The impact of various compression schemes on the outcome of selected graph algorithms. Bounds that do not include inequalities hold deterministically. If not otherwise stated, the other bounds hold in expectation. Bounds annotated with w.h.p. hold w.h.p. (if the involved quantities are large enough). Note that since the listed compression schemes (except the scheme where we remove the degree 1 vertices) return a subgraph of the original graph, m , C_R , \bar{d} , d , T , and \hat{M}_C never increase. Moreover, \mathcal{P} , \bar{P} , D , \mathcal{C} , and \hat{I}_S never decrease during compression. ϵ is a parameter that controls how well a spectral sparsifier approximates the original graph spectrum.

Figure 4: Analy

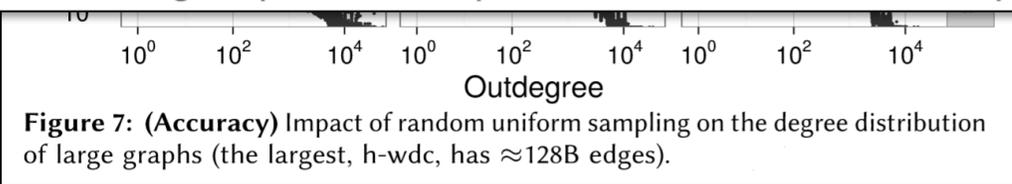
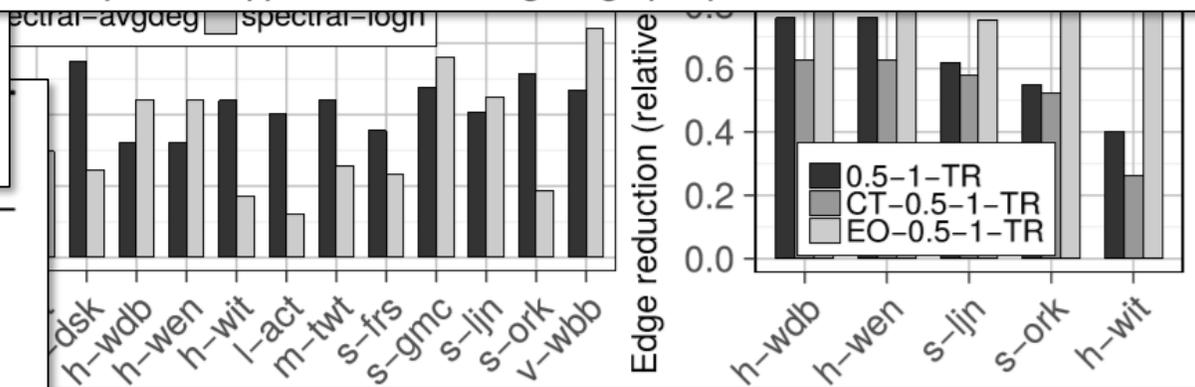
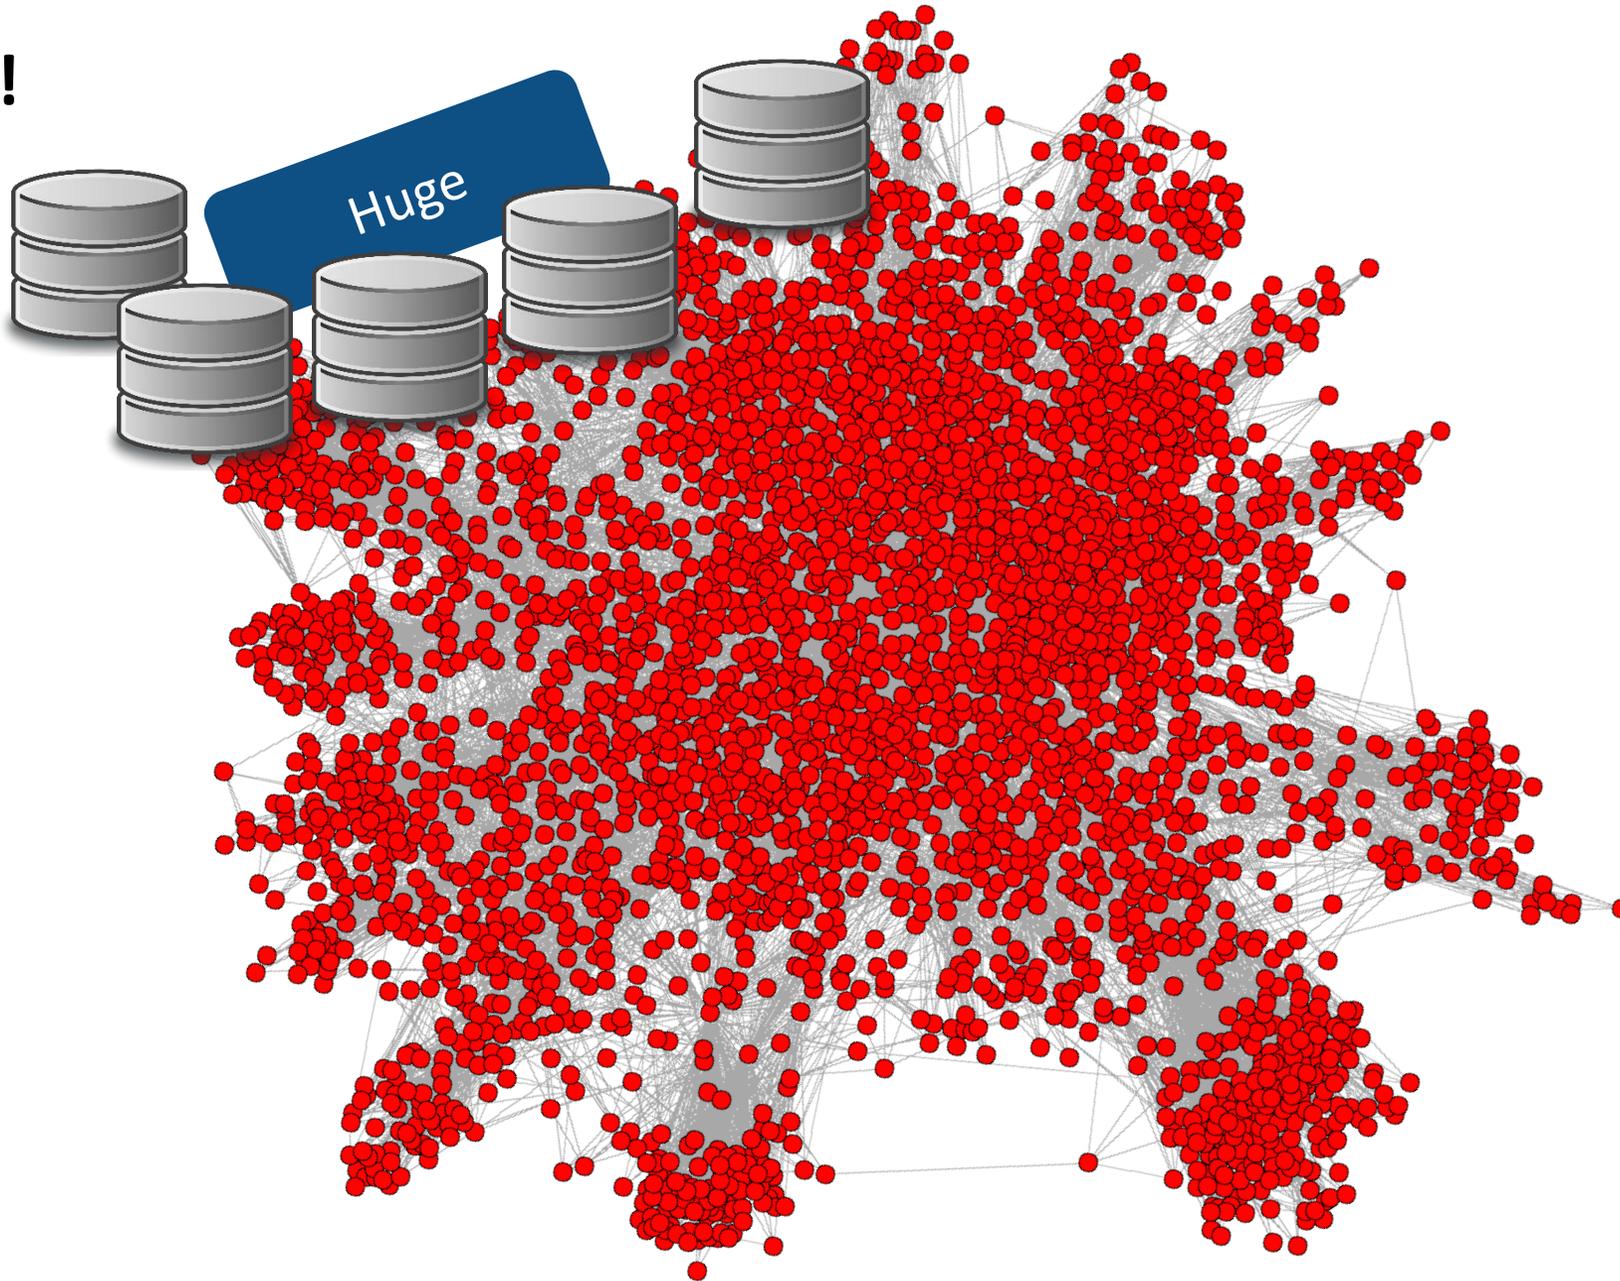


Figure 7: (Accuracy) Impact of random uniform sampling on the degree distribution of large graphs (the largest, h-wdc, has $\approx 128B$ edges).

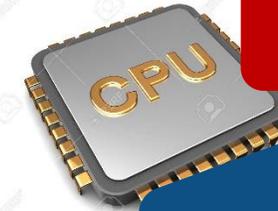
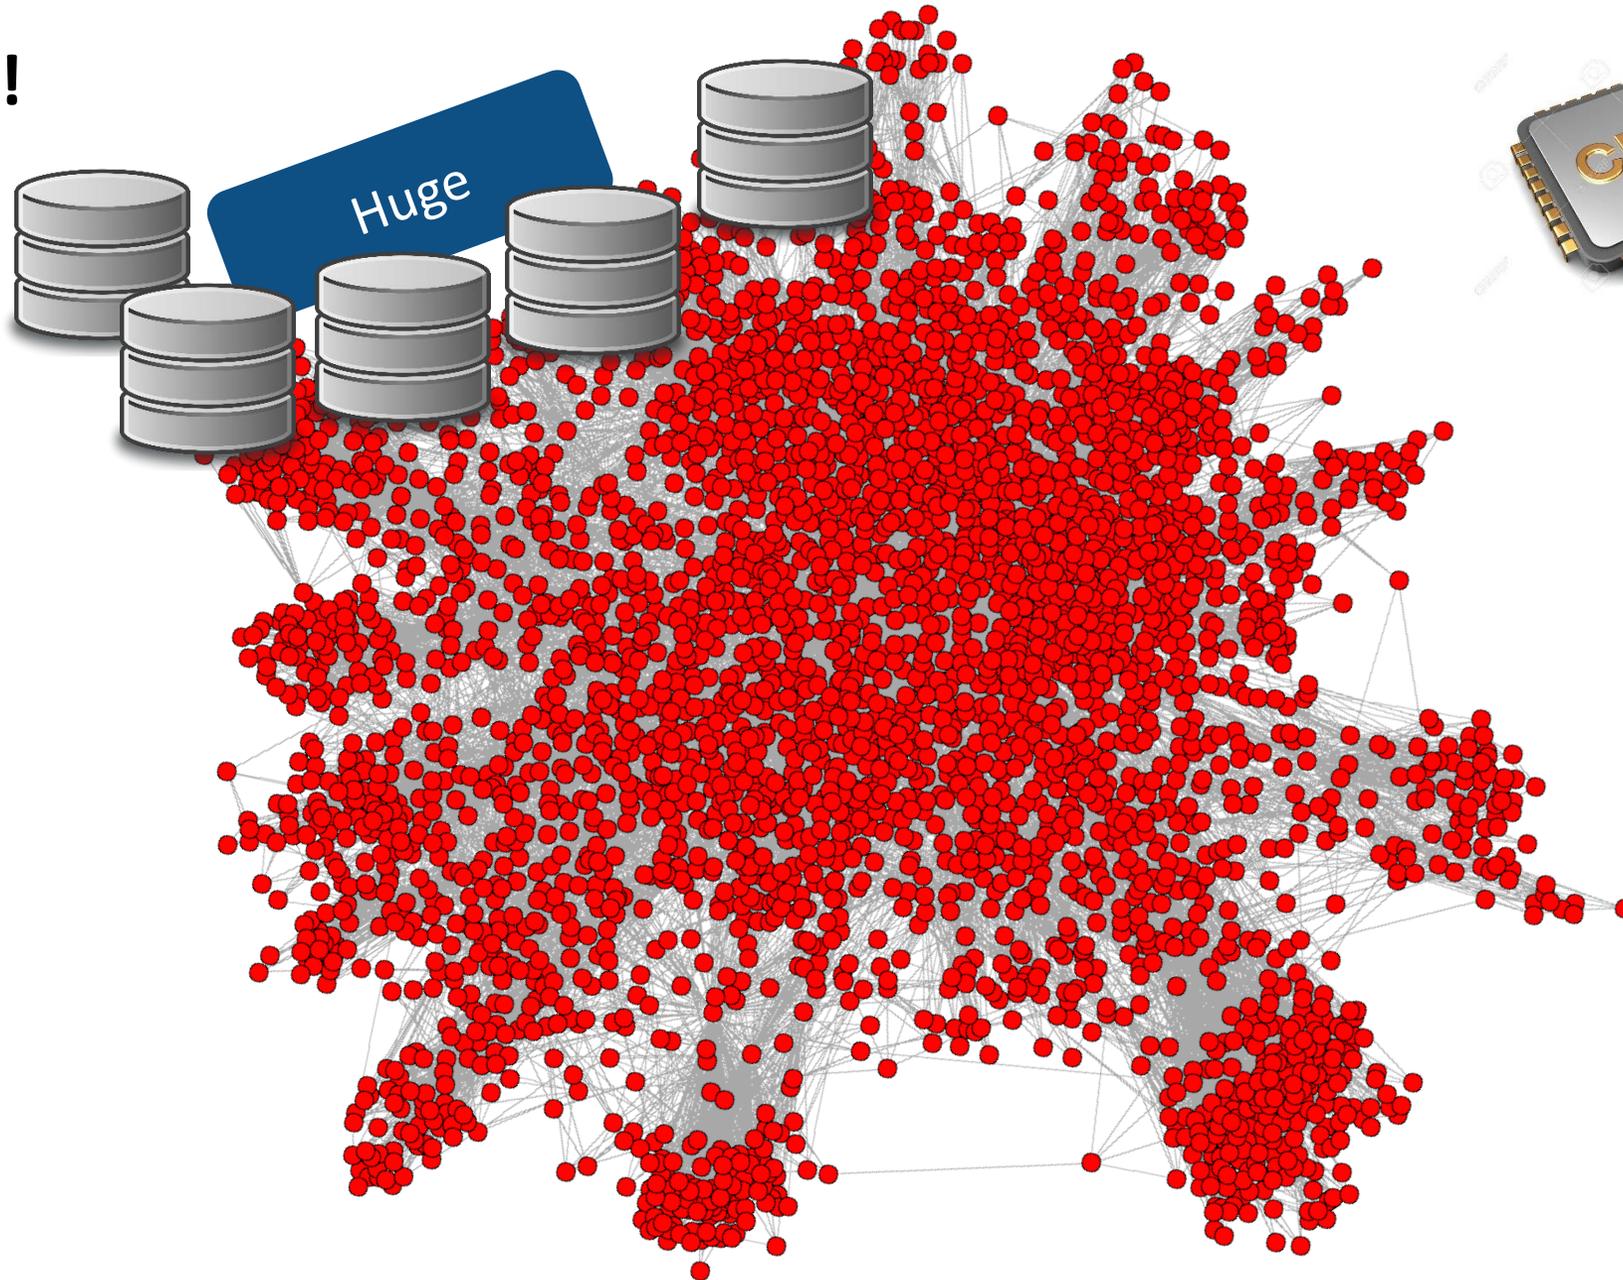
Graph	EO	0.8-1	0.0121	0.0167	0.1932	0.6019	0.0054	0.2808	0.2993
s-you	0.0121	0.0167	0.1932	0.6019	0.0054	0.2808	0.2993		
h-hud	0.0187	0.0271	0.0477	0.1633	0.0340	0.2794	0.3247		
il-dbl	0.0459	0.0674	0.0749	0.2929	0.0080	0.1980	0.2005		
v-skt	0.0410	0.0643	0.0674	0.2695	0.0311	0.1101	0.2950		
v-usa	0.0089	0.0100	0.1392	0.5945	0.0000	0.0074	0.0181		



Problems!



Problems!



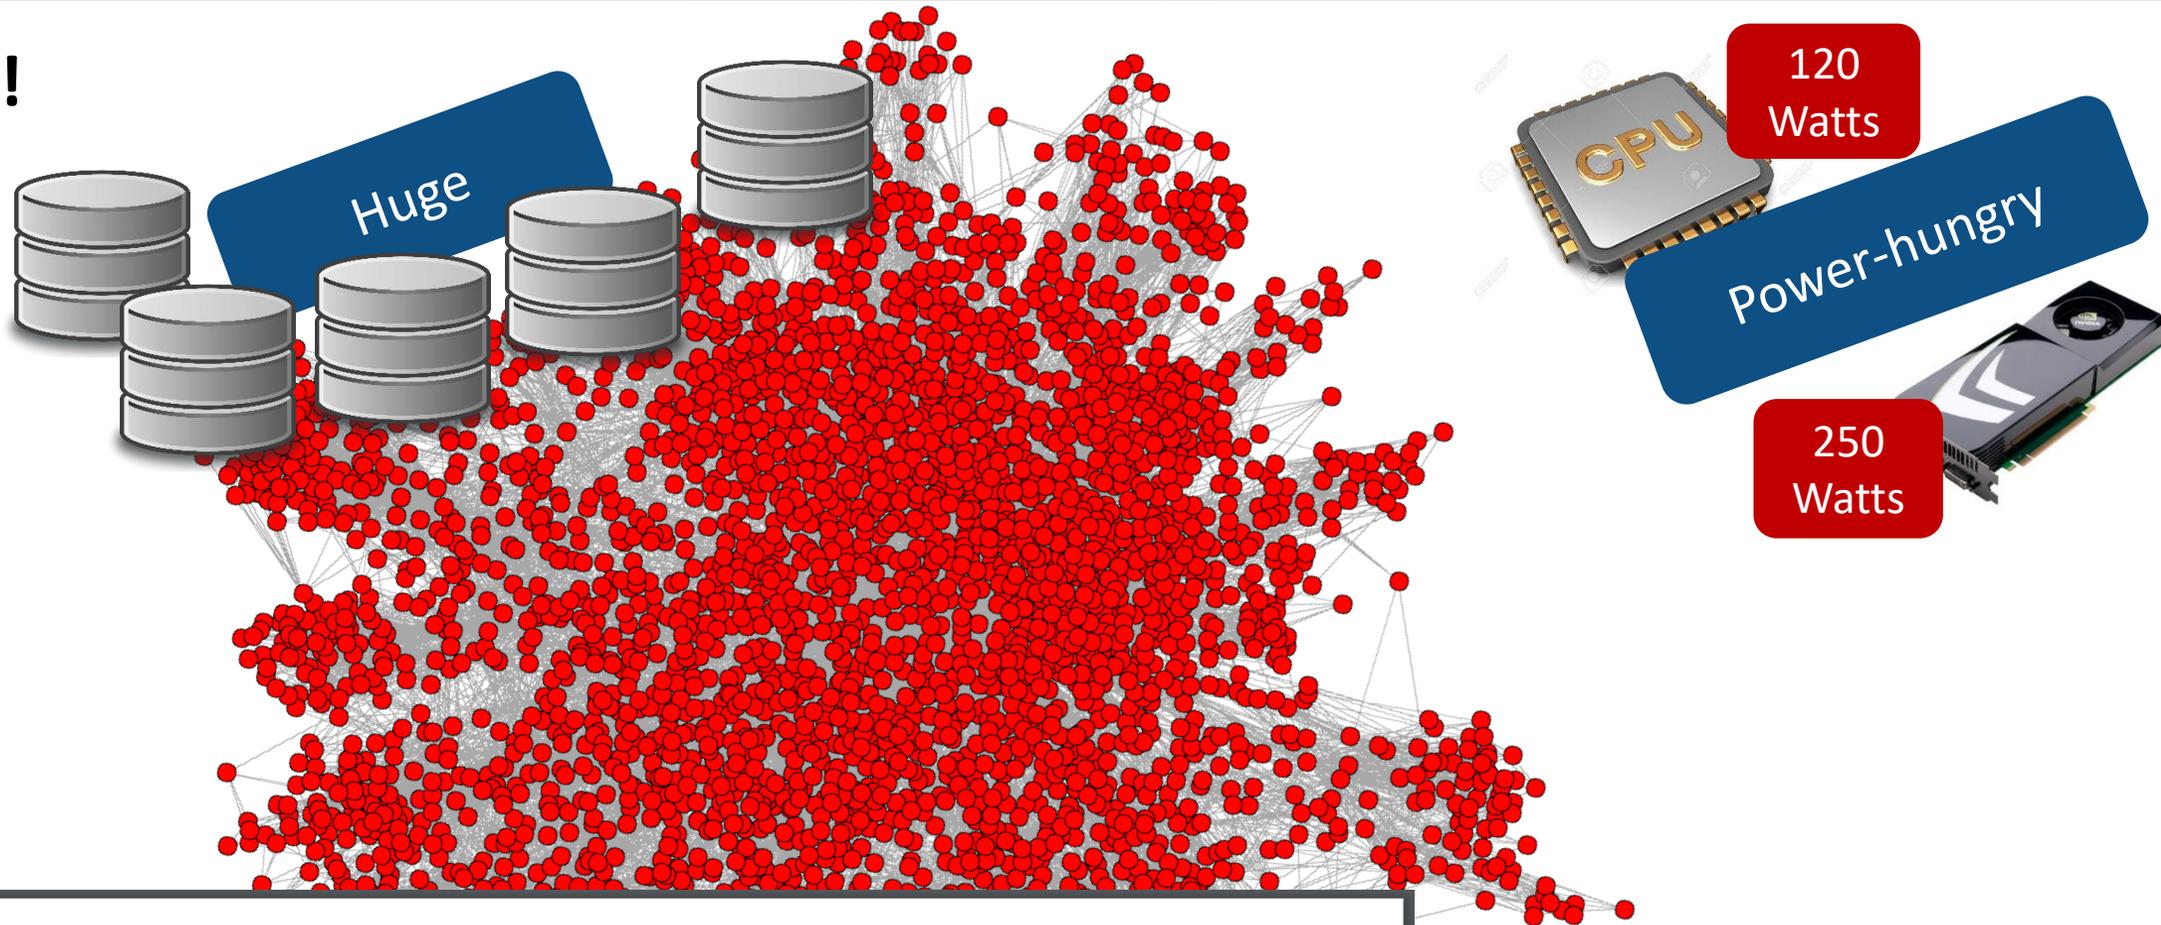
120
Watts

Power-hungry



250
Watts

Problems!



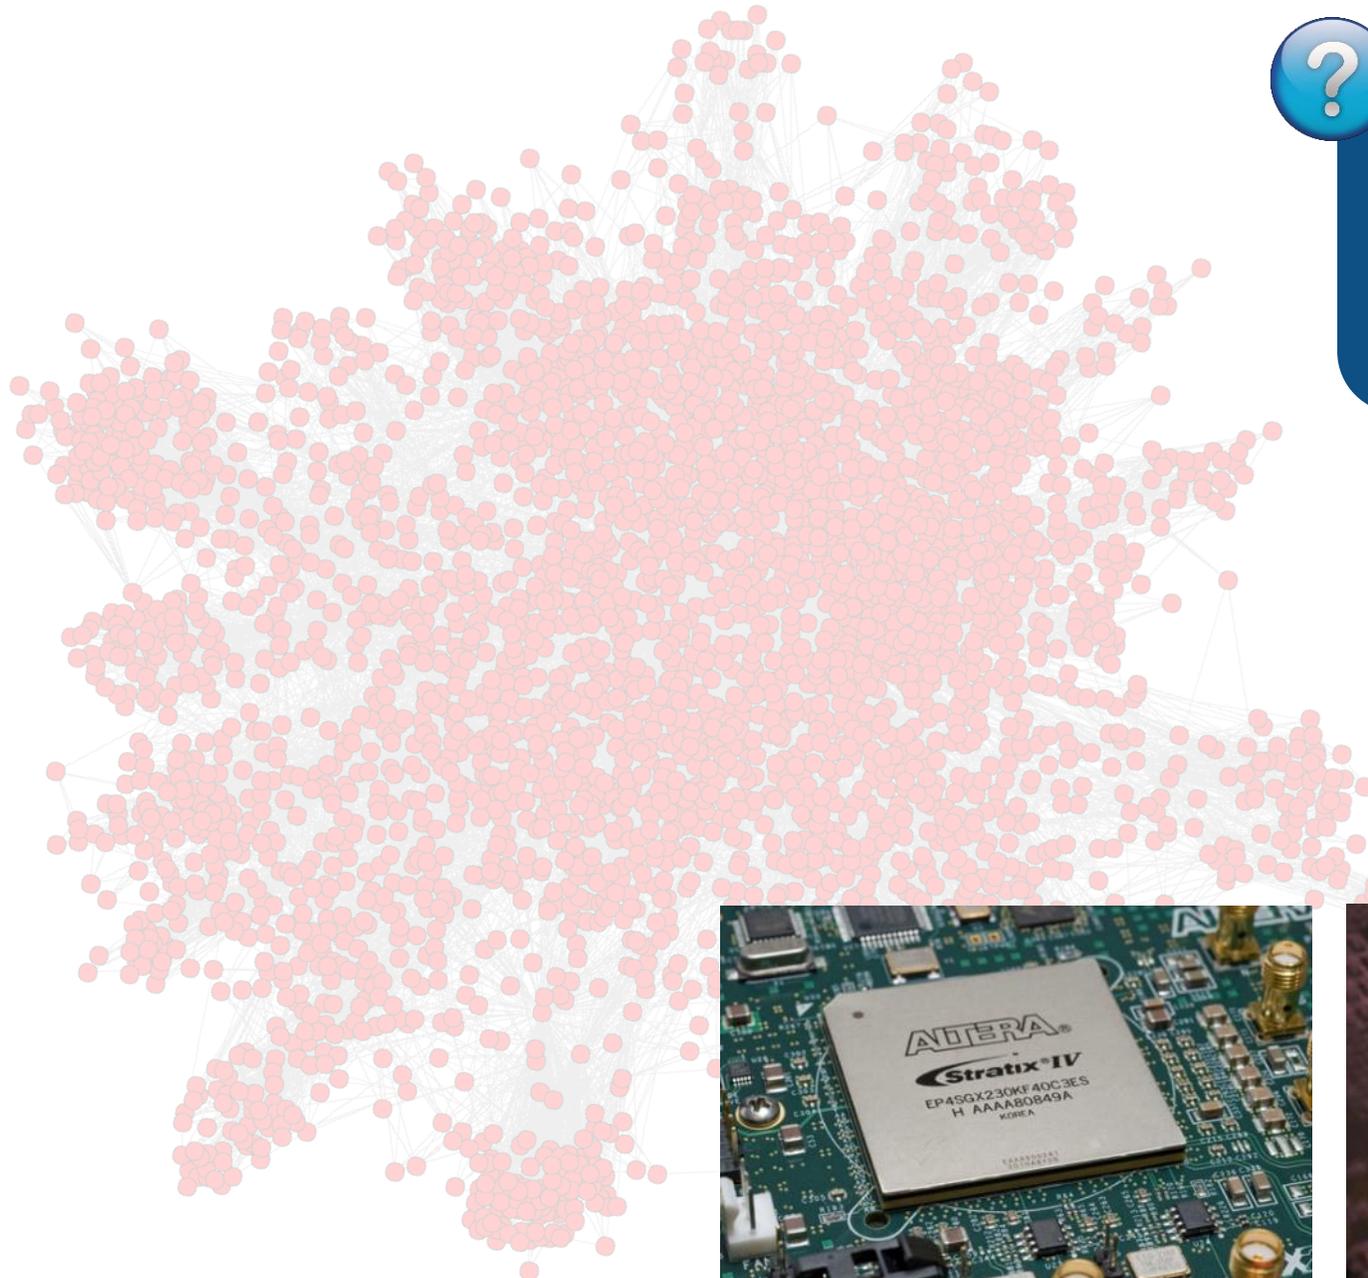
Substream-Centric Maximum Matchings on FPGA

Maciej Besta, Marc Fischer, Tal Ben-Nun, Johannes De Fine Licht, Torsten Hoefler
Department of Computer Science, ETH Zurich

ABSTRACT

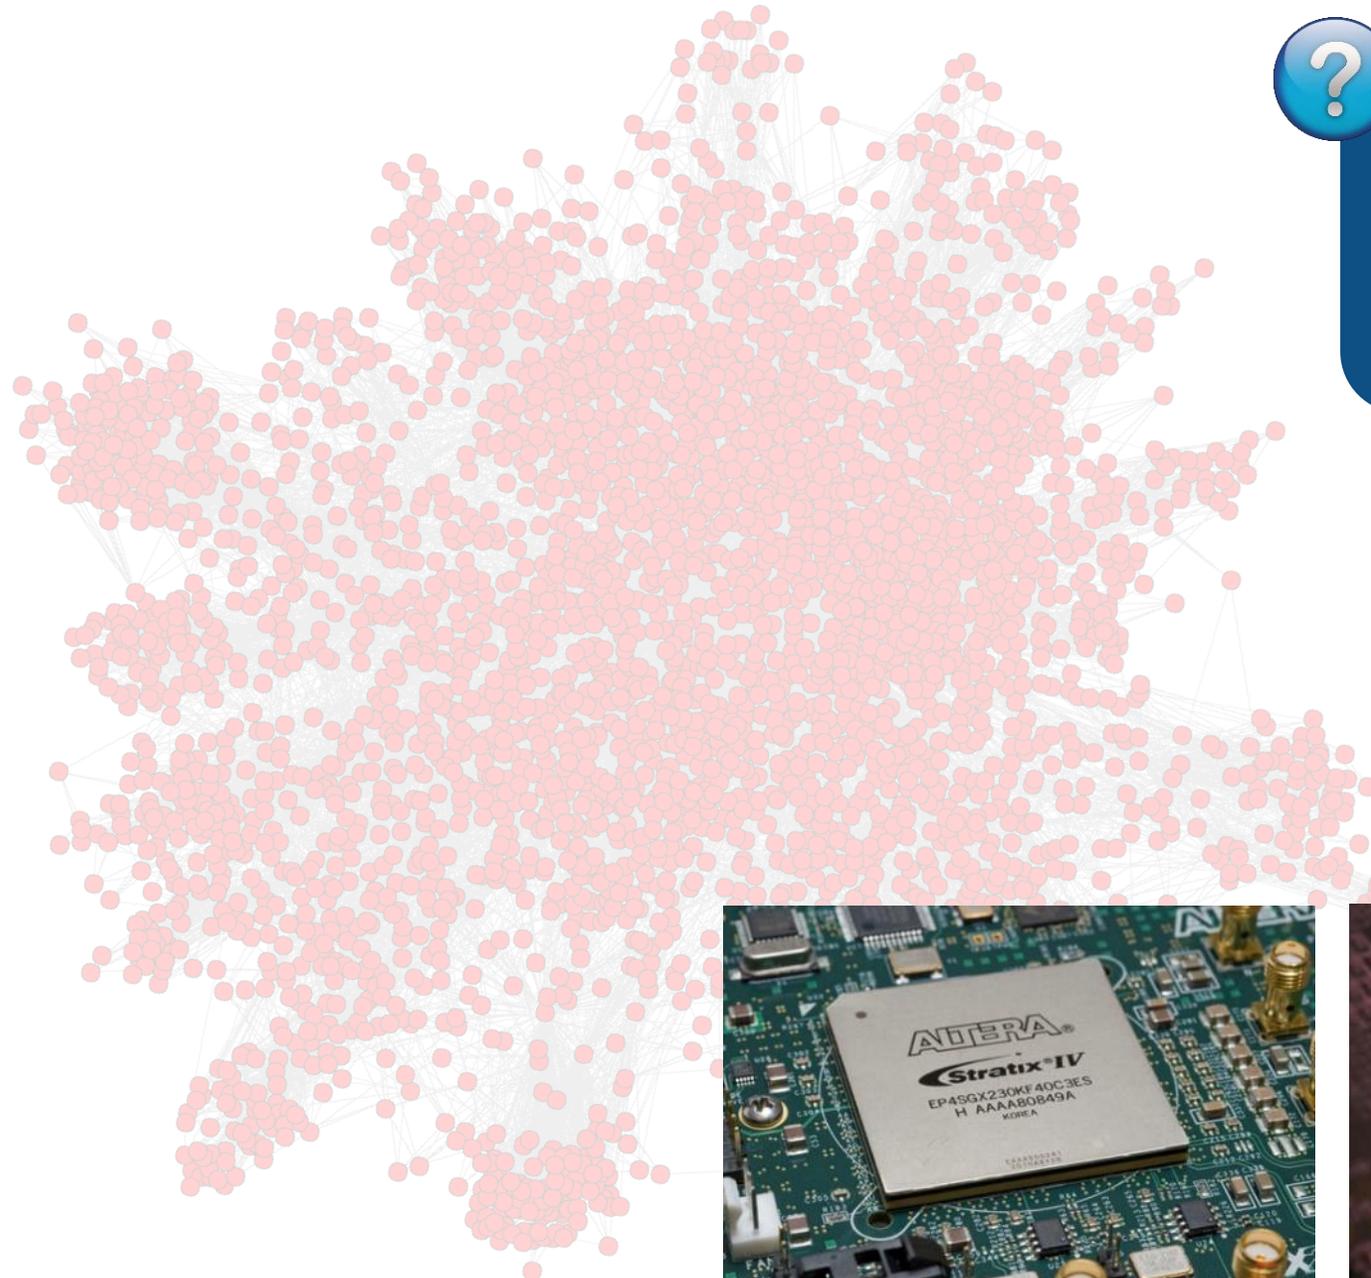
Developing high-performance and energy-efficient algorithms for maximum matchings is becoming increasingly important in social network analysis, computational sciences, schedul-

compound can be used to show the locations of double bonds in the chemical structure [59]. As deriving the exact MM is usually computationally expensive, significant focus has been placed on developing fast approximate solutions [17].



What graph programming paradigm for FPGAs and why?





What graph programming paradigm for FPGAs and why?

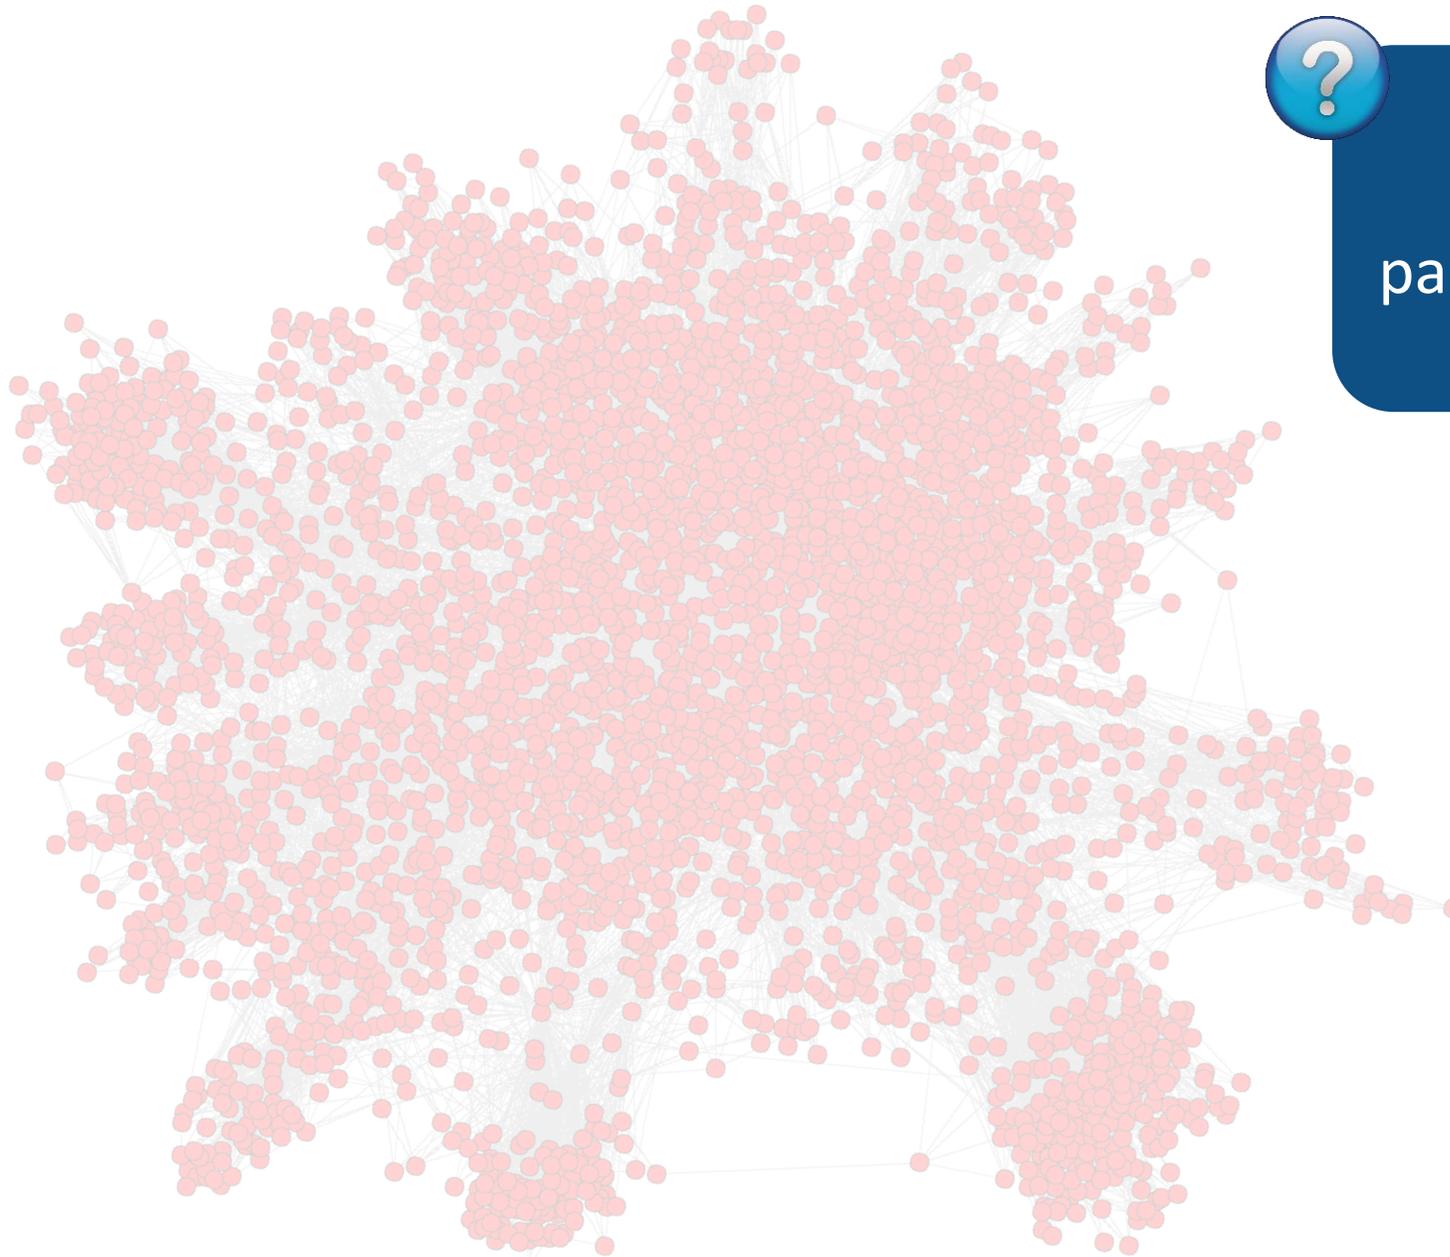




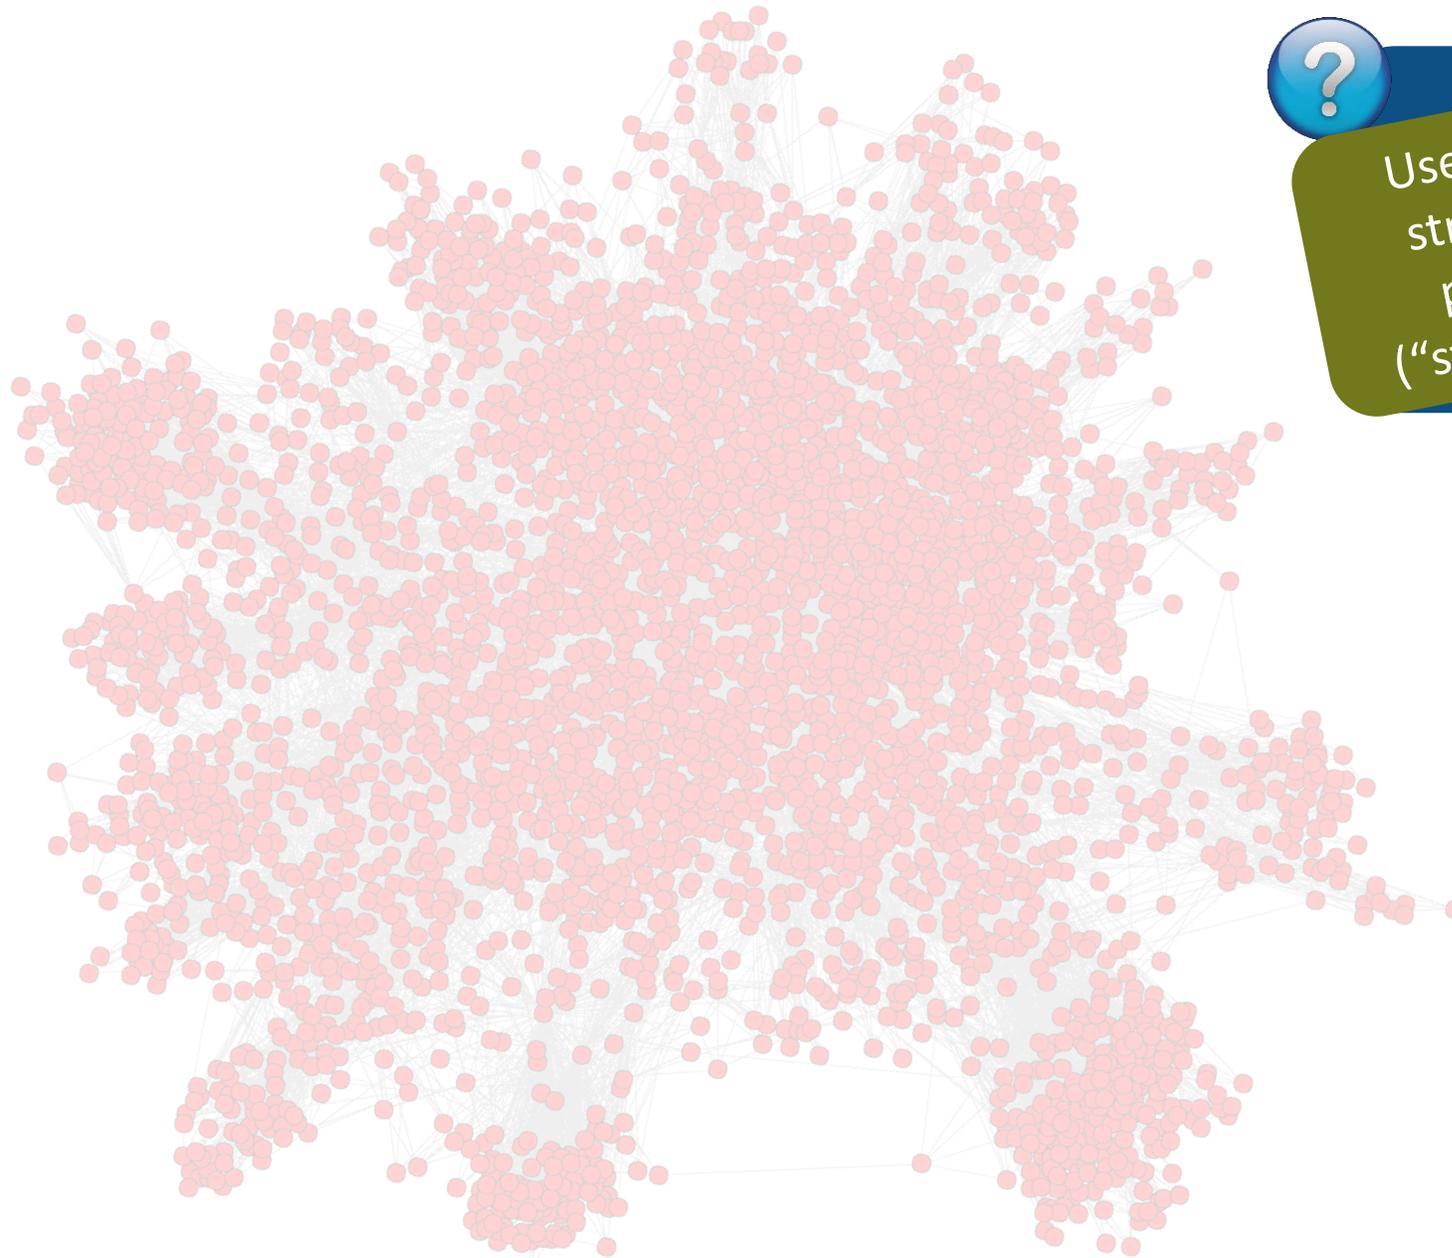
What graph programming paradigm for FPGAs and why?

To be able to utilize pipelining well, we really want to use edge streaming



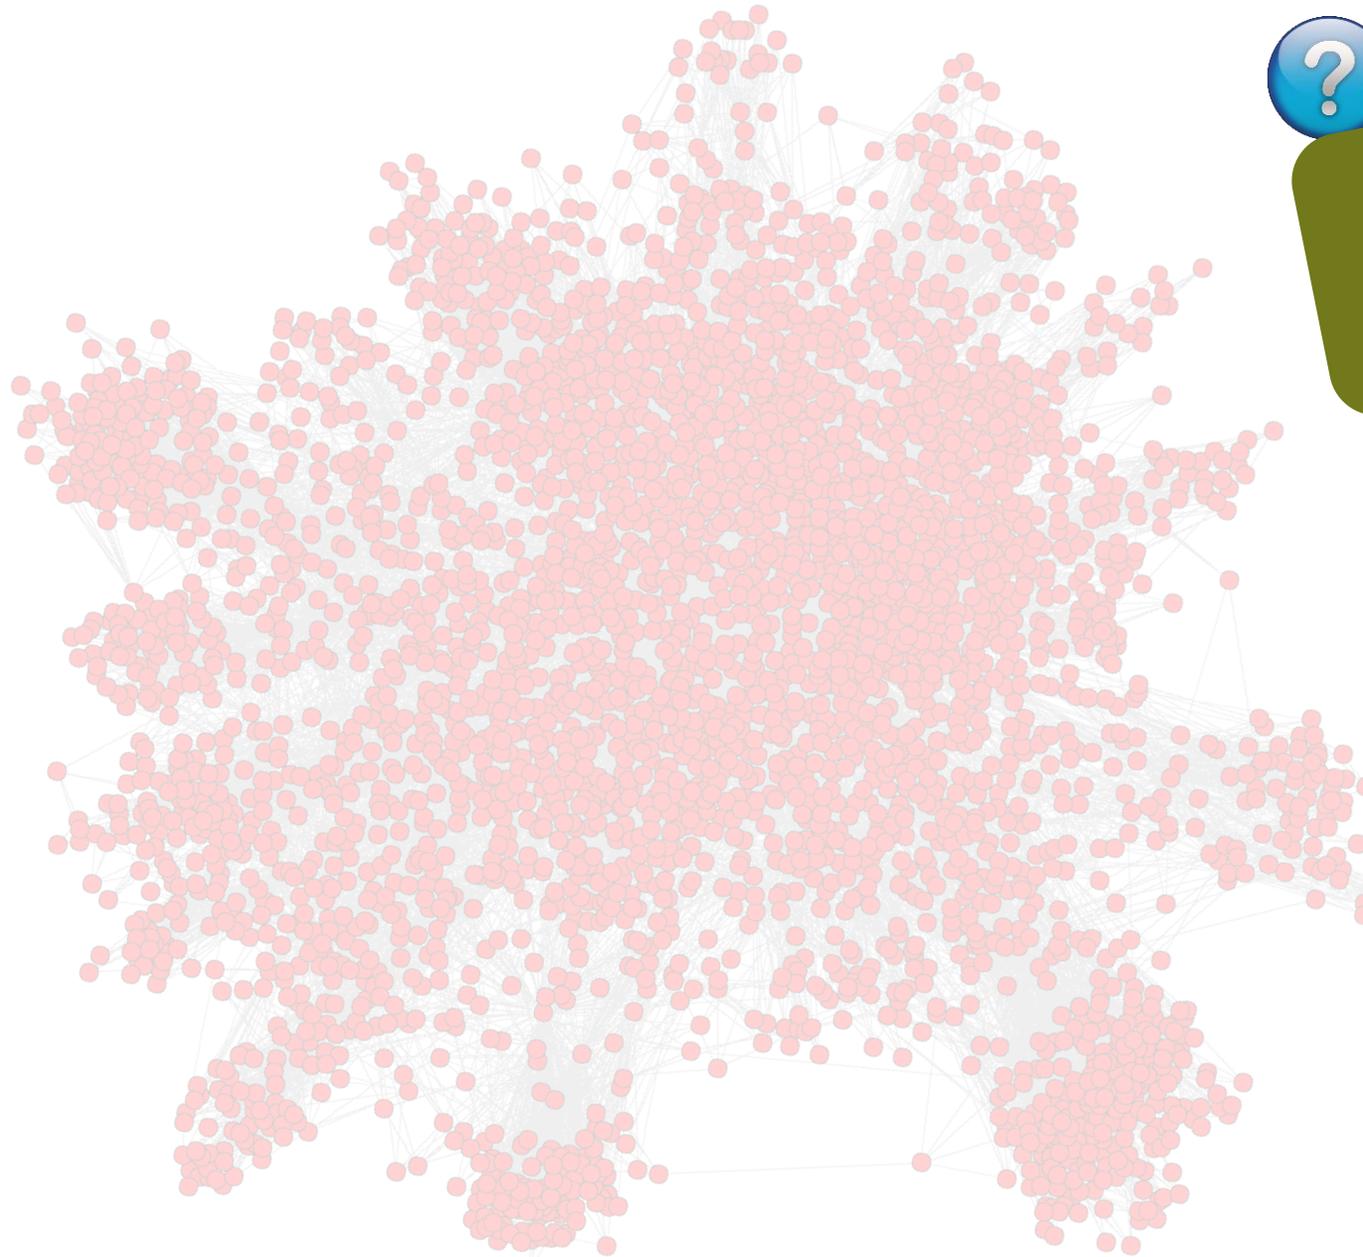


What graph programming paradigm for FPGAs and why?



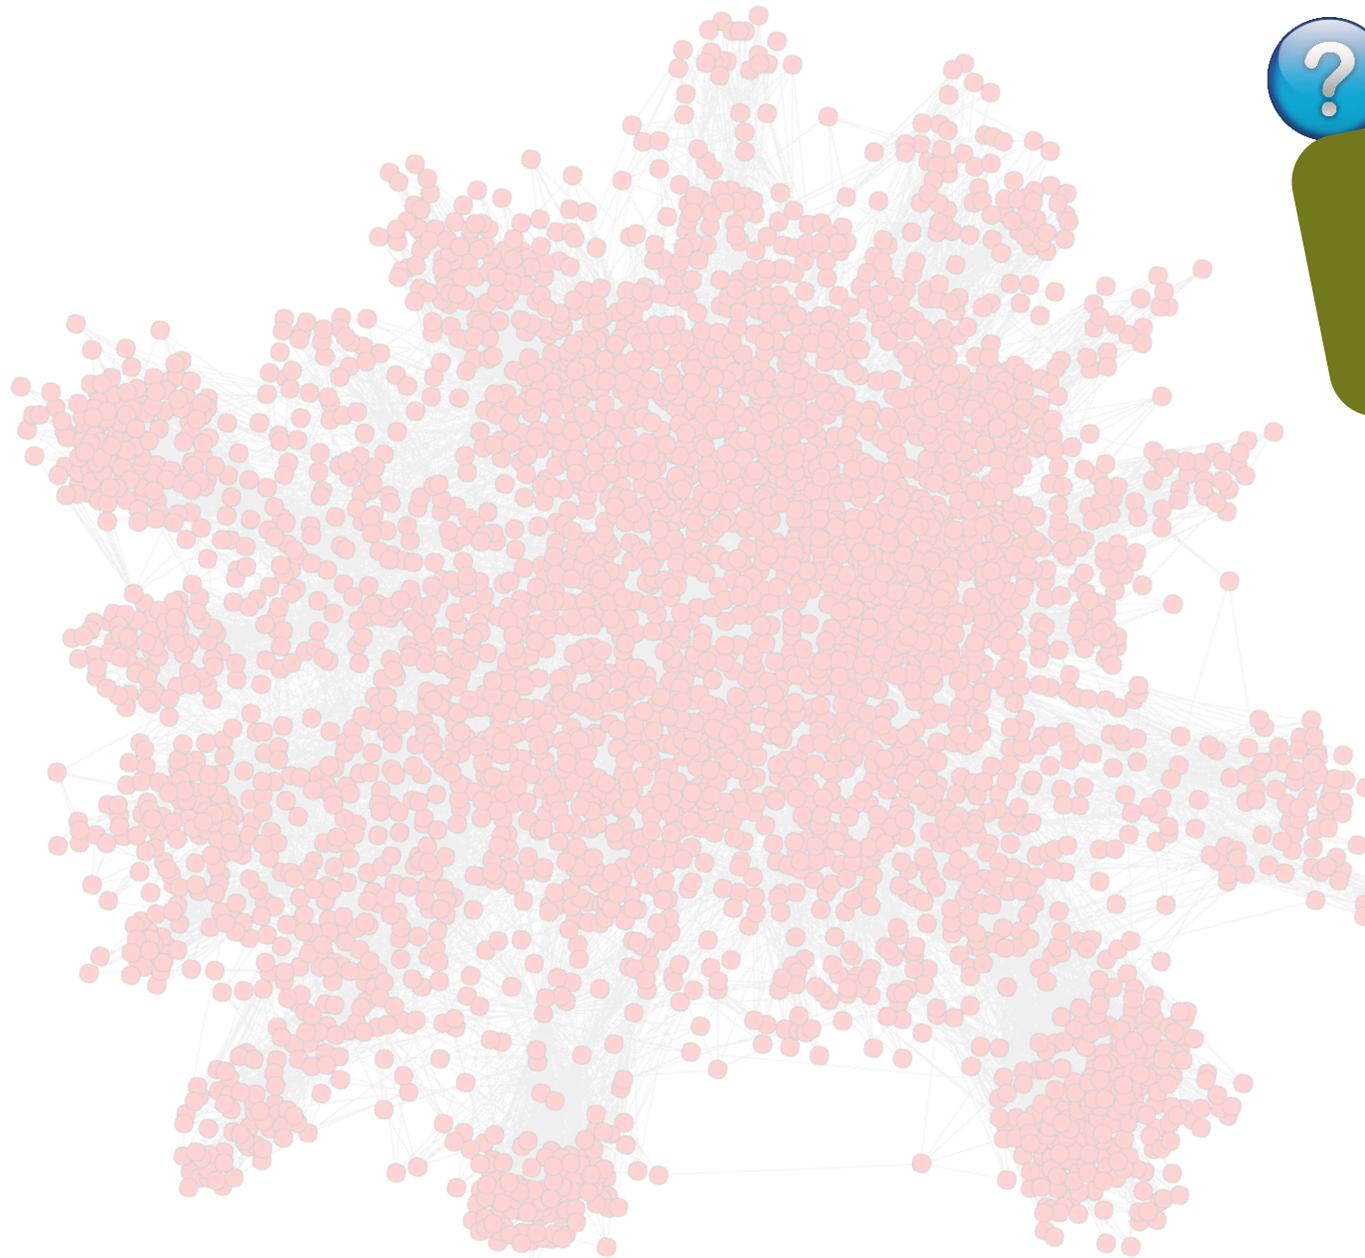
Use some form of edge streaming; we can use pipelining efficiently (“streaming \approx pipelining”)

Why?

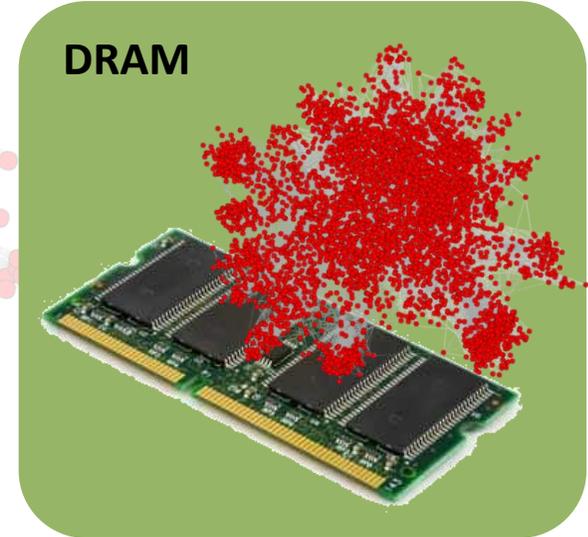


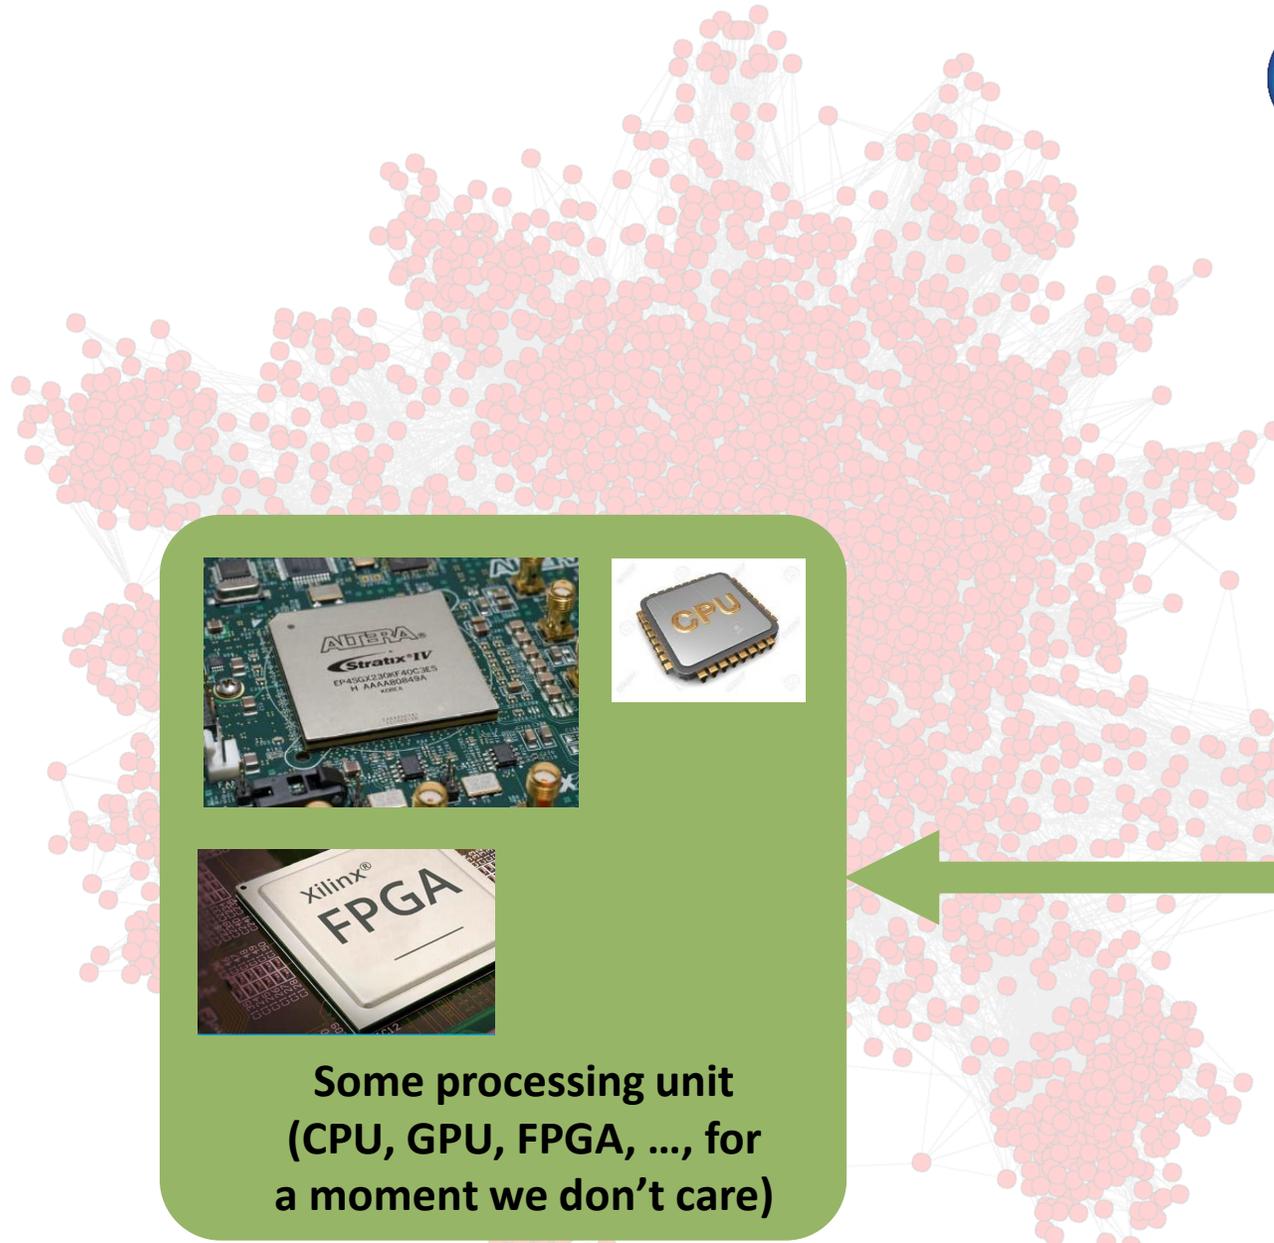
Use some form of edge streaming; we can use pipelining efficiently (“streaming \approx pipelining”)





Use some form of edge streaming; we can use pipelining efficiently (“streaming \approx pipelining”)



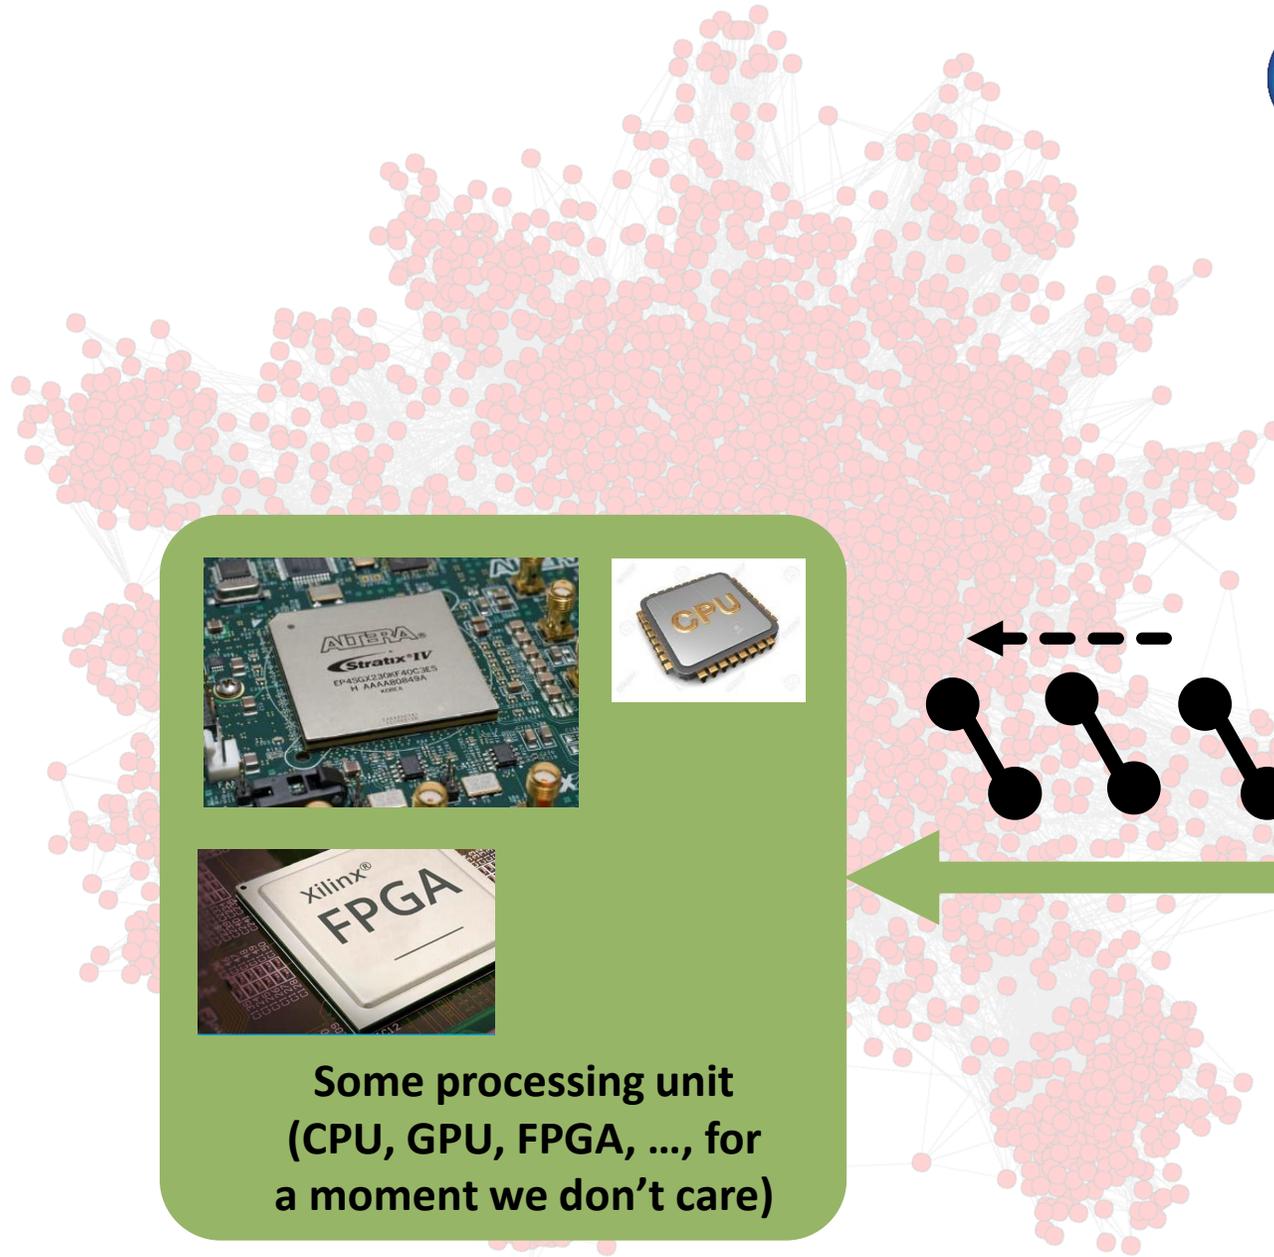


Use some form of edge streaming; we can use pipelining efficiently (“streaming \approx pipelining”)

Some processing unit (CPU, GPU, FPGA, ..., for a moment we don't care)

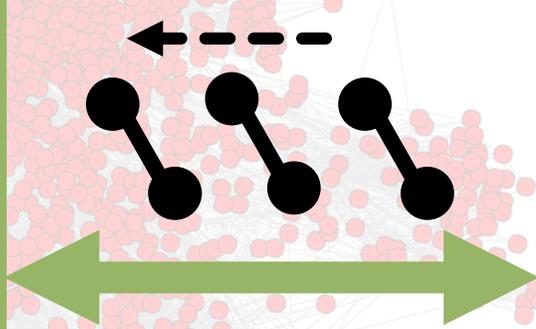
DRAM



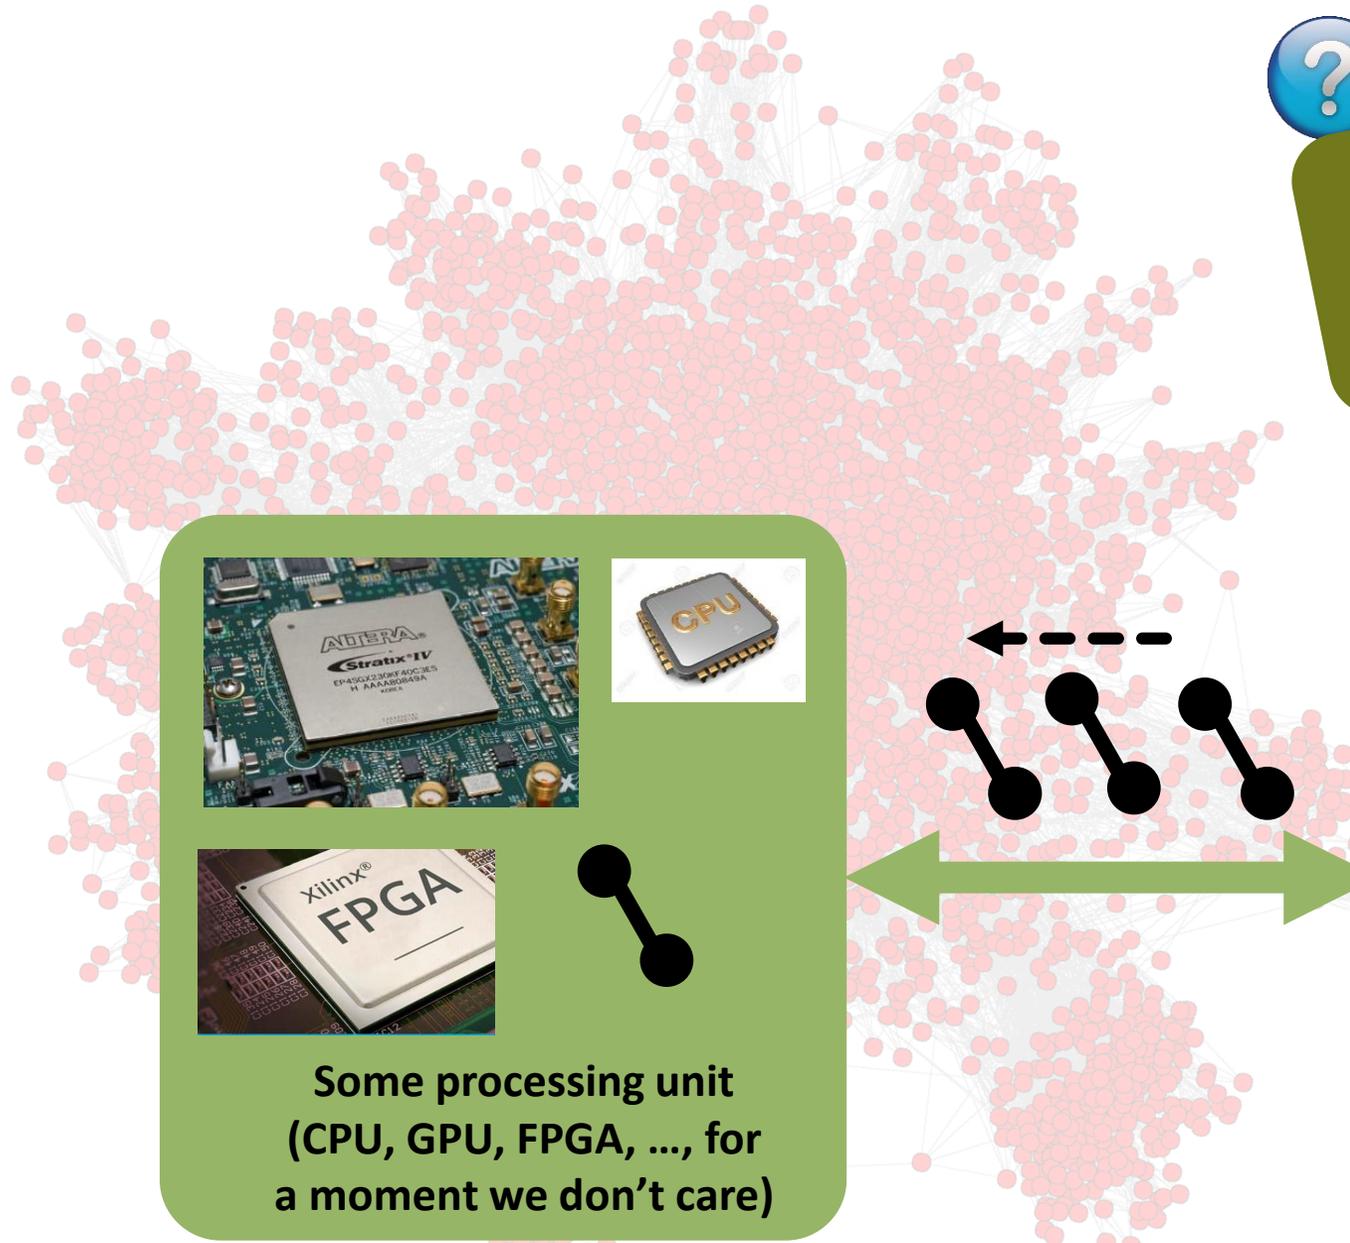


Use some form of edge streaming; we can use pipelining efficiently (“streaming \approx pipelining”)

Some processing unit (CPU, GPU, FPGA, ..., for a moment we don't care)

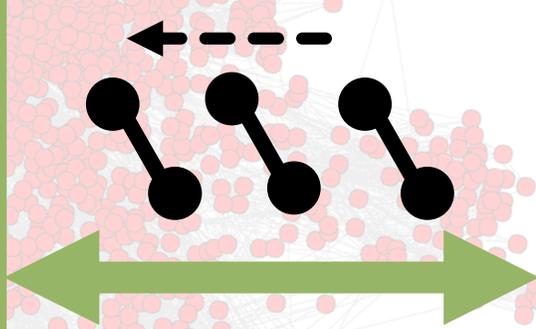


DRAM

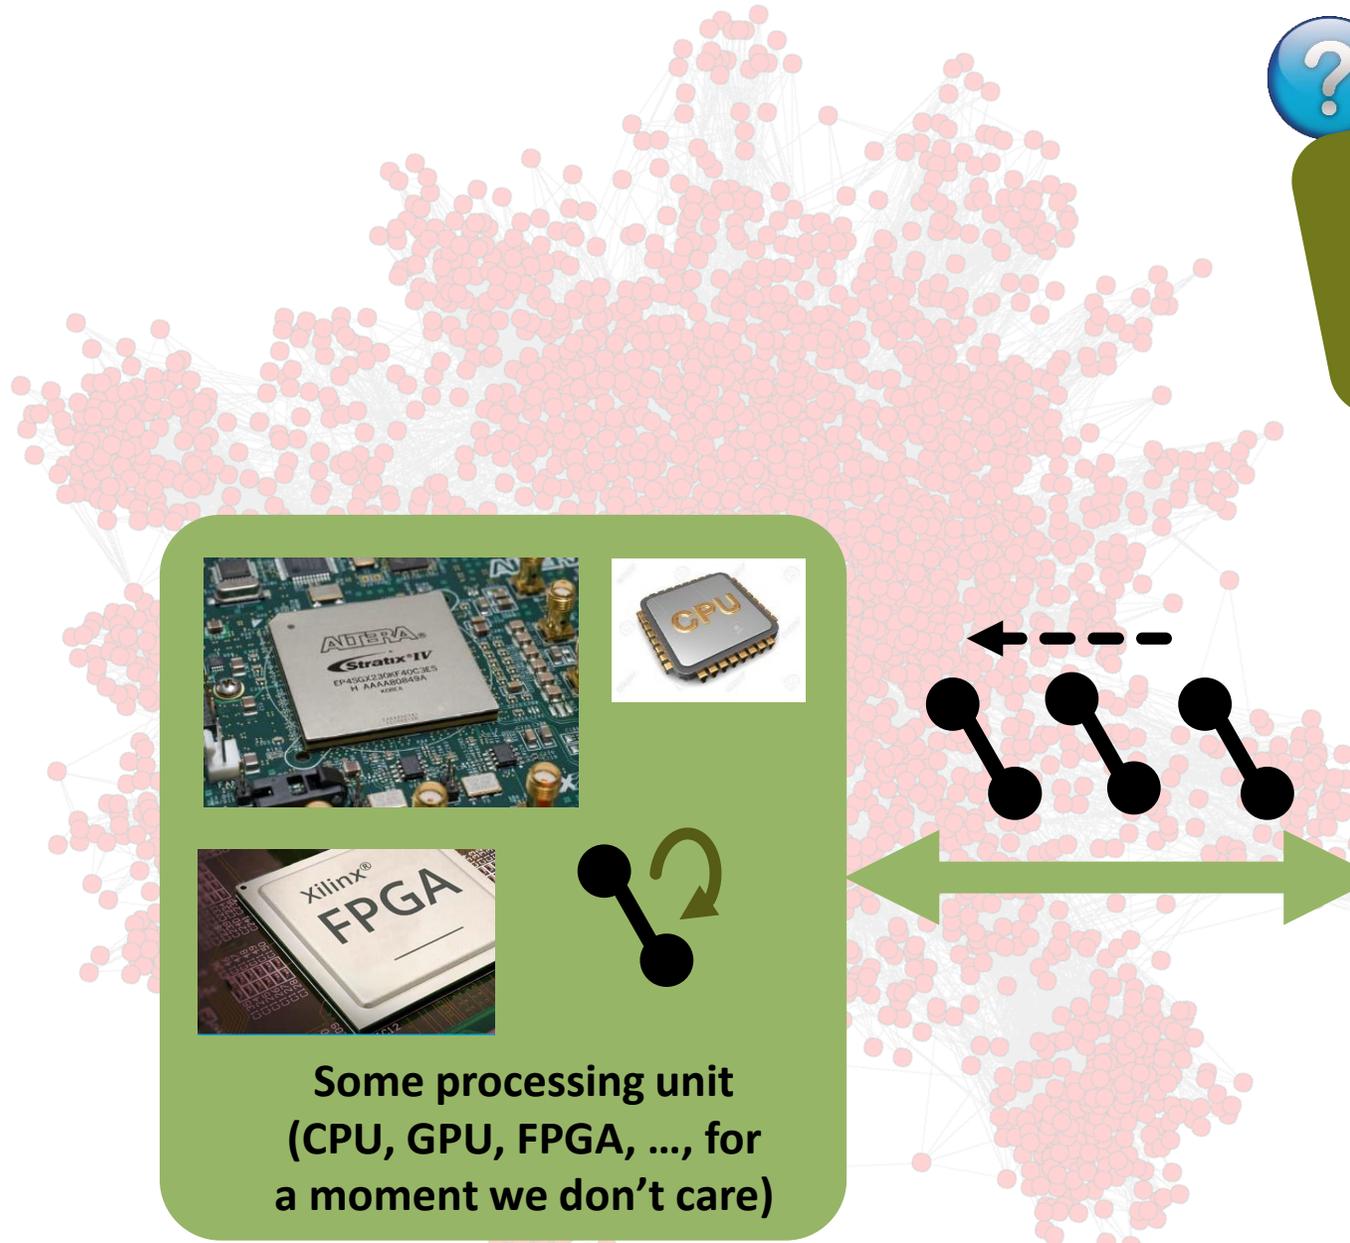


Use some form of edge streaming; we can use pipelining efficiently (“streaming \approx pipelining”)

Some processing unit (CPU, GPU, FPGA, ..., for a moment we don't care)

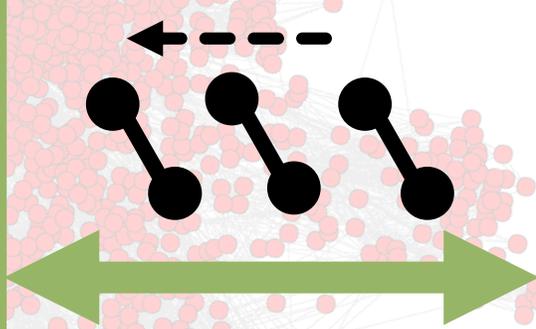


DRAM

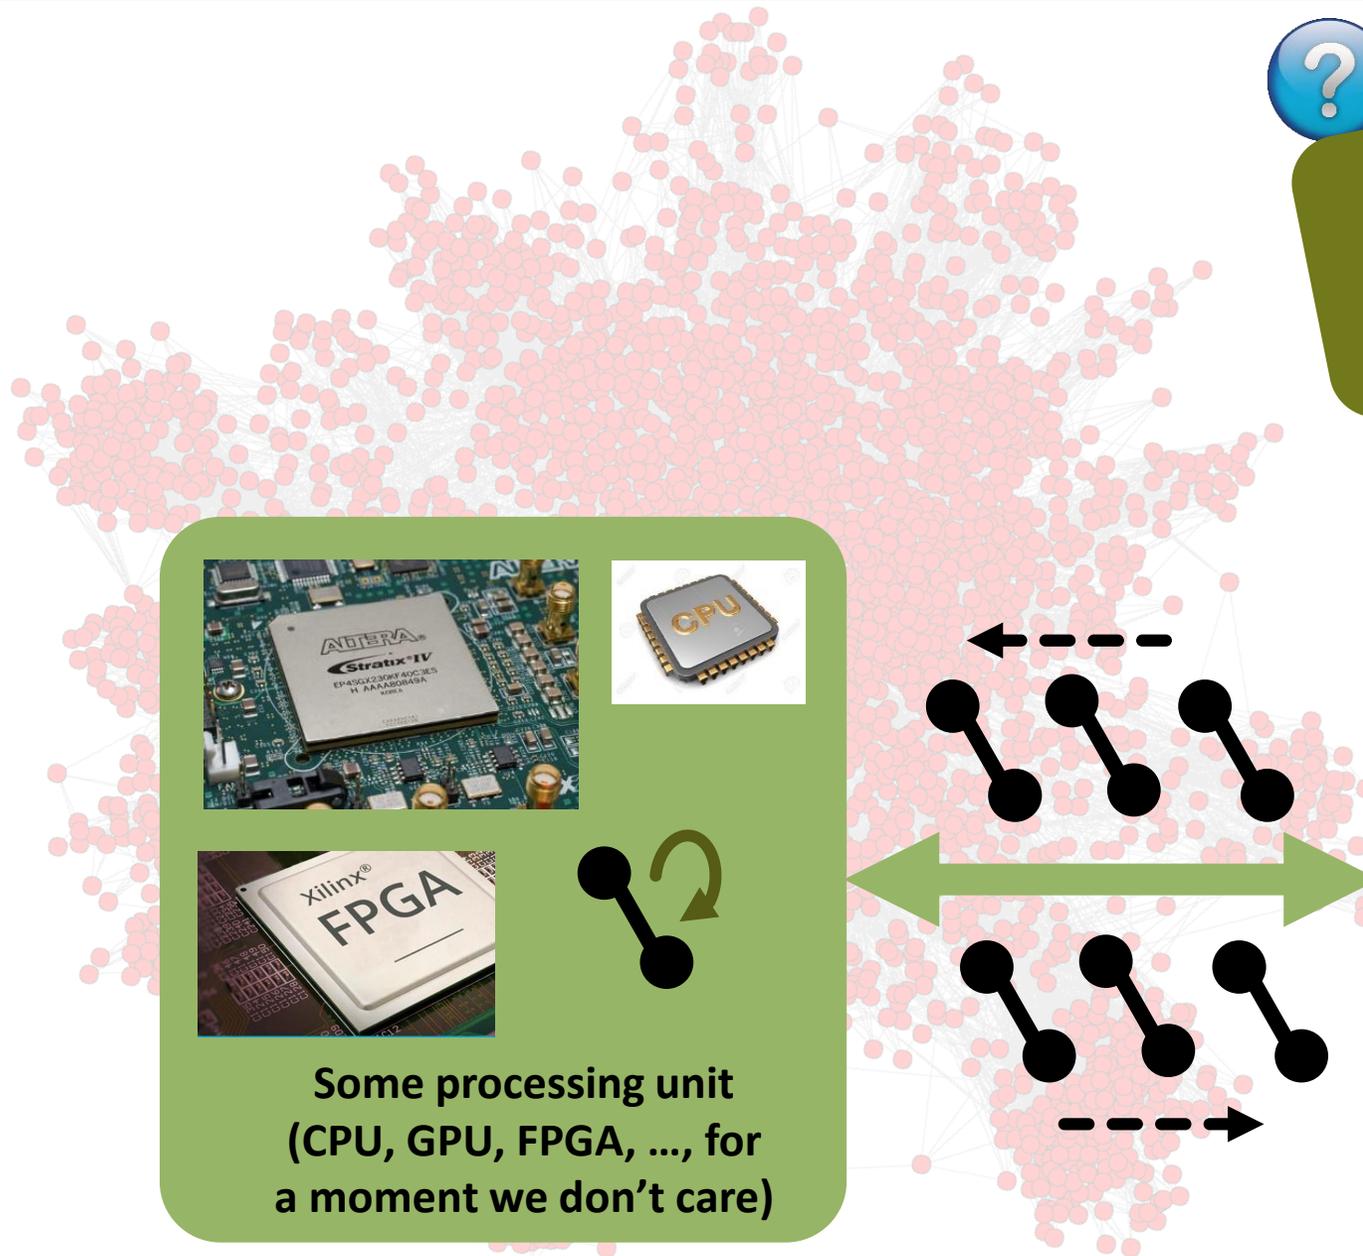


Use some form of edge streaming; we can use pipelining efficiently (“streaming \approx pipelining”)
Why?

Some processing unit (CPU, GPU, FPGA, ..., for a moment we don't care)



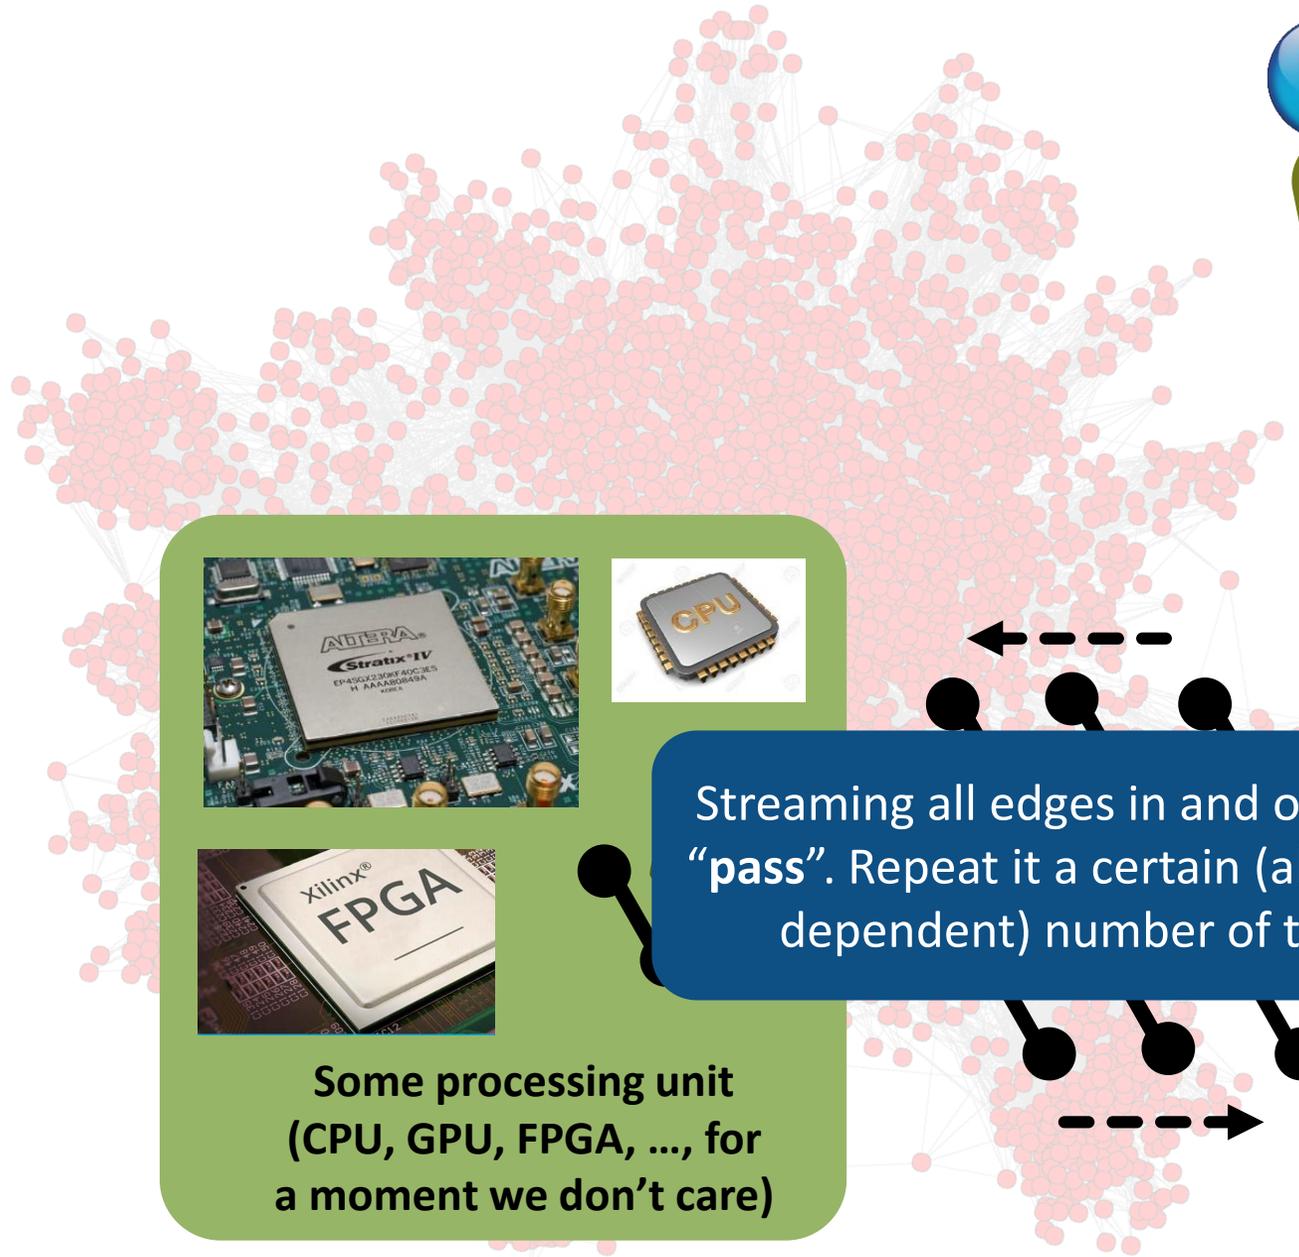
DRAM



Use some form of edge streaming; we can use pipelining efficiently (“streaming \approx pipelining”)

Some processing unit (CPU, GPU, FPGA, ..., for a moment we don't care)

DRAM

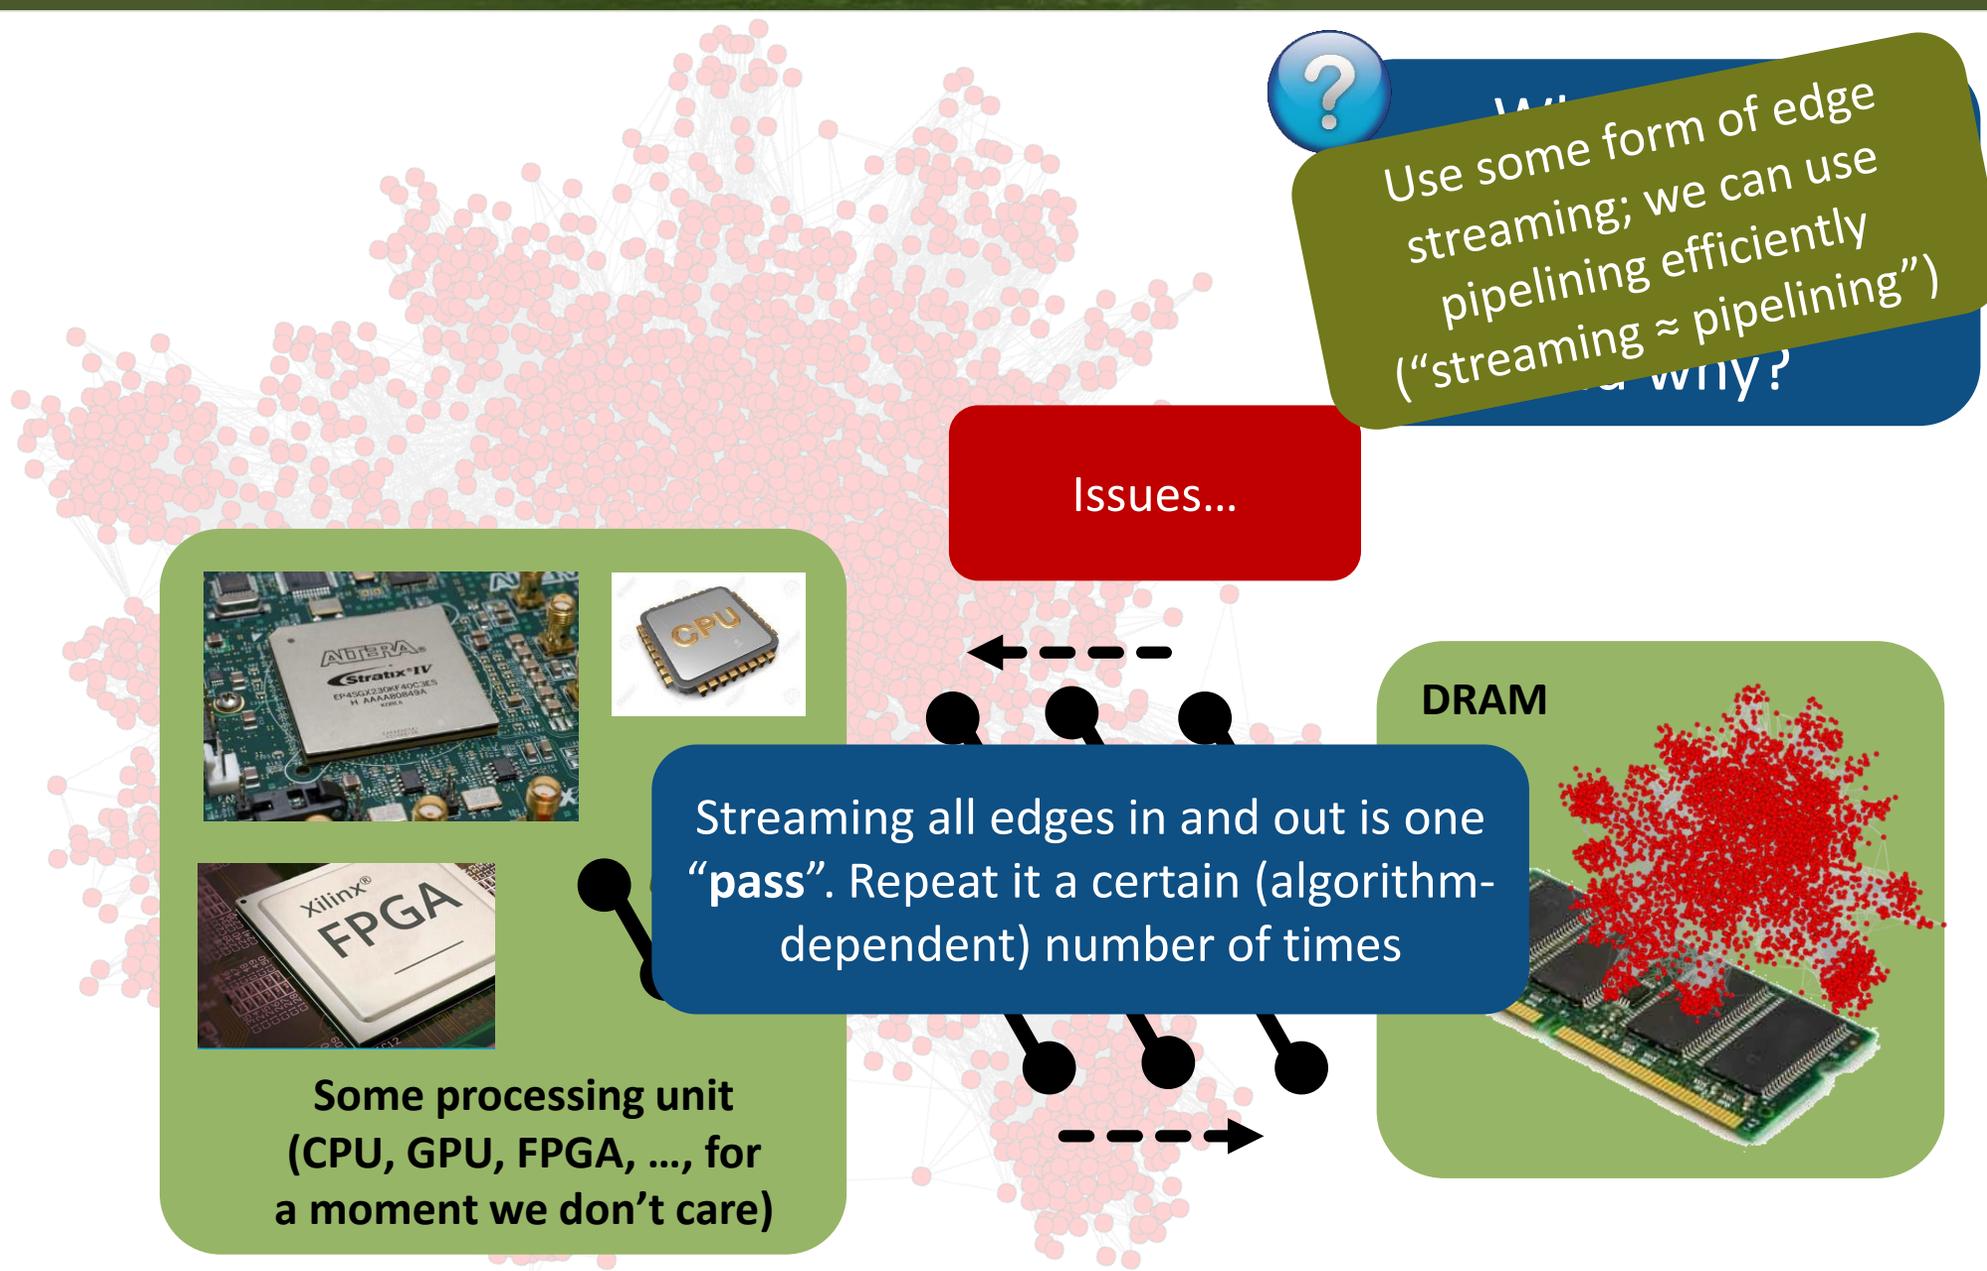


Use some form of edge streaming; we can use pipelining efficiently (“streaming \approx pipelining”)

Some processing unit (CPU, GPU, FPGA, ..., for a moment we don't care)

Streaming all edges in and out is one “pass”. Repeat it a certain (algorithm-dependent) number of times

DRAM



Use some form of edge streaming; we can use pipelining efficiently (“streaming ≈ pipelining”)

Issues...

Some processing unit (CPU, GPU, FPGA, ..., for a moment we don't care)

Streaming all edges in and out is one “pass”. Repeat it a certain (algorithm-dependent) number of times

DRAM

...How to minimize the number of “passes” over edges? This can get really bad in the “traditional” edge-centric approach (e.g., BFS needs D passes; D = diameter [1]).

Use some form of edge streaming; we can use pipelining efficiently (“streaming \approx pipelining”)

...Processing edges is sequential – how to incorporate parallelism?

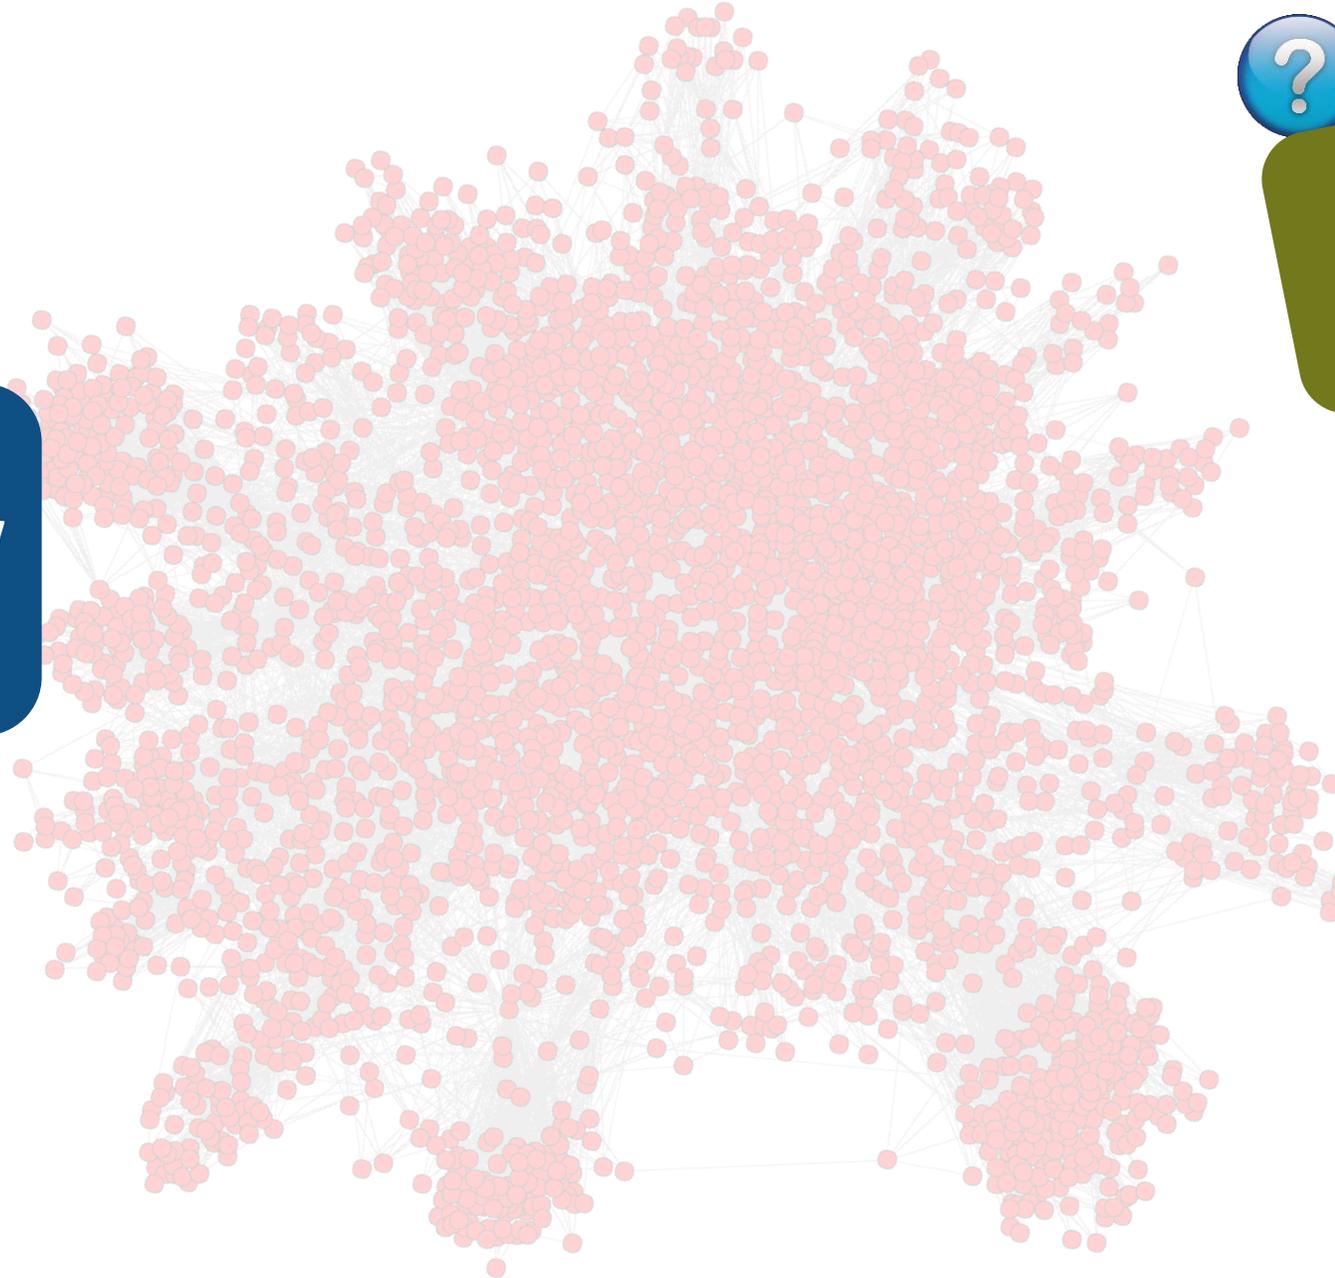
Issues...

Some processing unit (CPU, GPU, FPGA, ..., for a moment we don't care)

Streaming all edges in and out is one “pass”. Repeat it a certain (algorithm-dependent) number of times

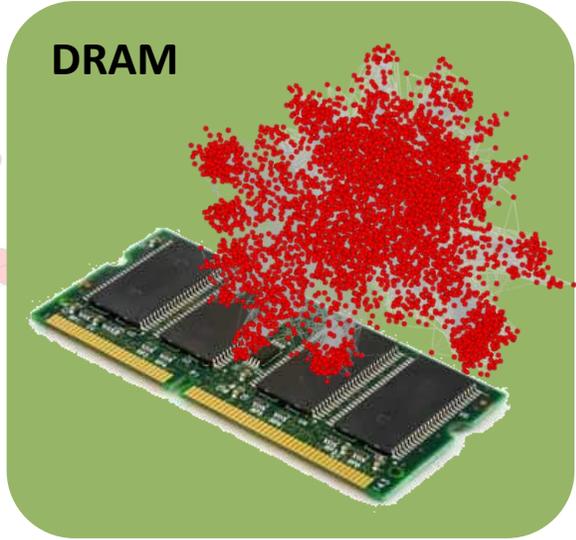
DRAM

[1] A. Roy et al. X-stream: Edge-Centric Graph Processing using Streaming Partitions. SOSP. 2013.



...Processing edges is sequential – how to incorporate parallelism?

Use some form of edge streaming; we can use pipelining efficiently (“streaming \approx pipelining”)



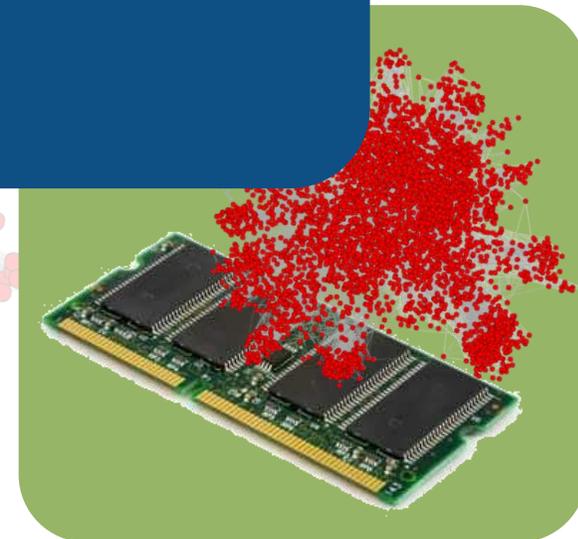
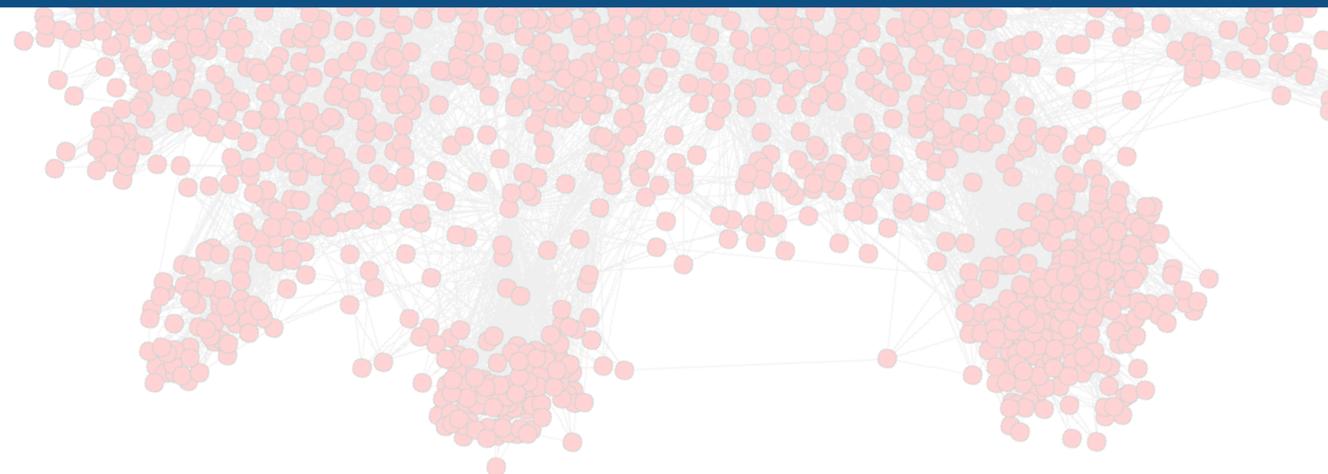


Use some form of edge streaming; we can use pipelining efficiently (streaming \approx pipelining")
... why?



...Processing is sequential to improve parallelism

Substream-Centric: A new paradigm for processing graphs

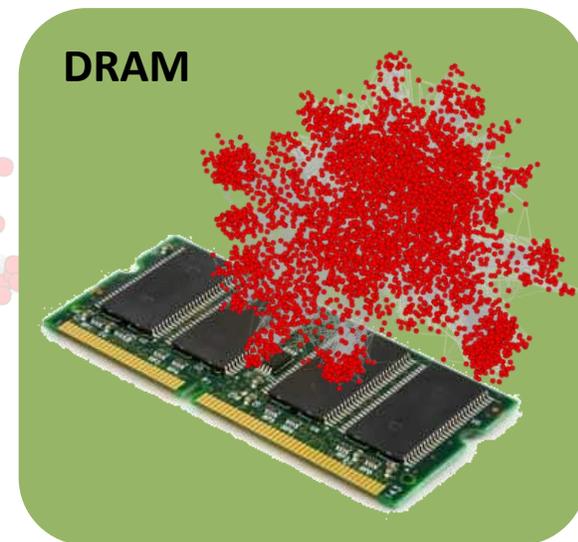
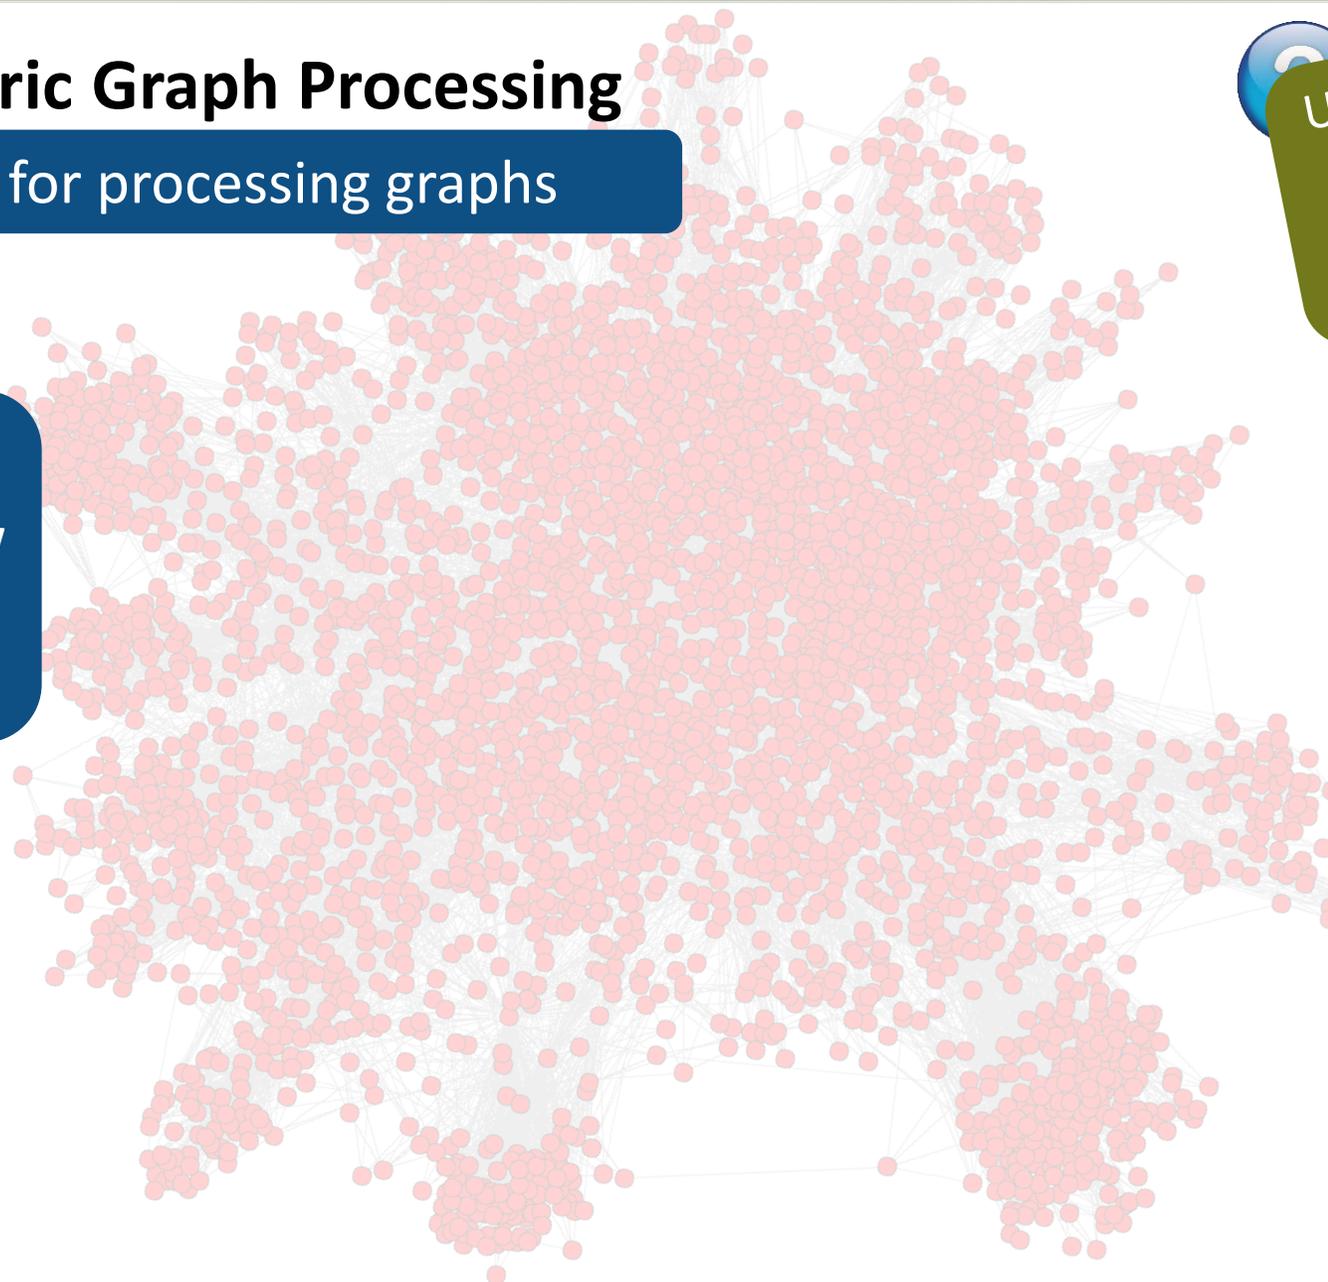


Substream-Centric Graph Processing

A new paradigm for processing graphs

...Processing edges is sequential – how to incorporate parallelism?

Use some form of streaming (aka **edge-centric**); we can use pipelining efficiently (“streaming \approx pipelining”)



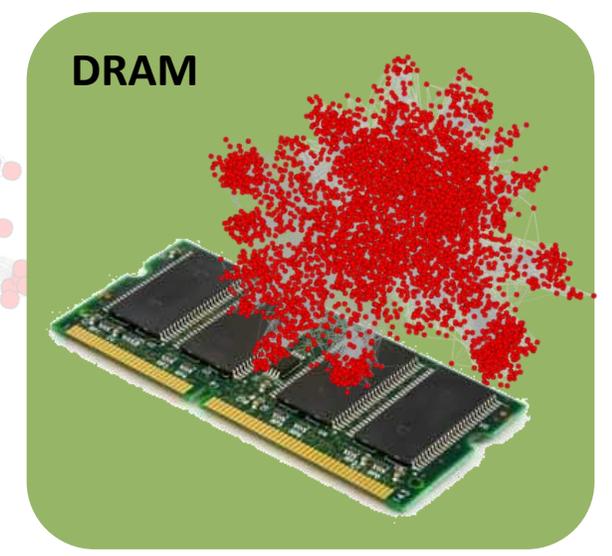
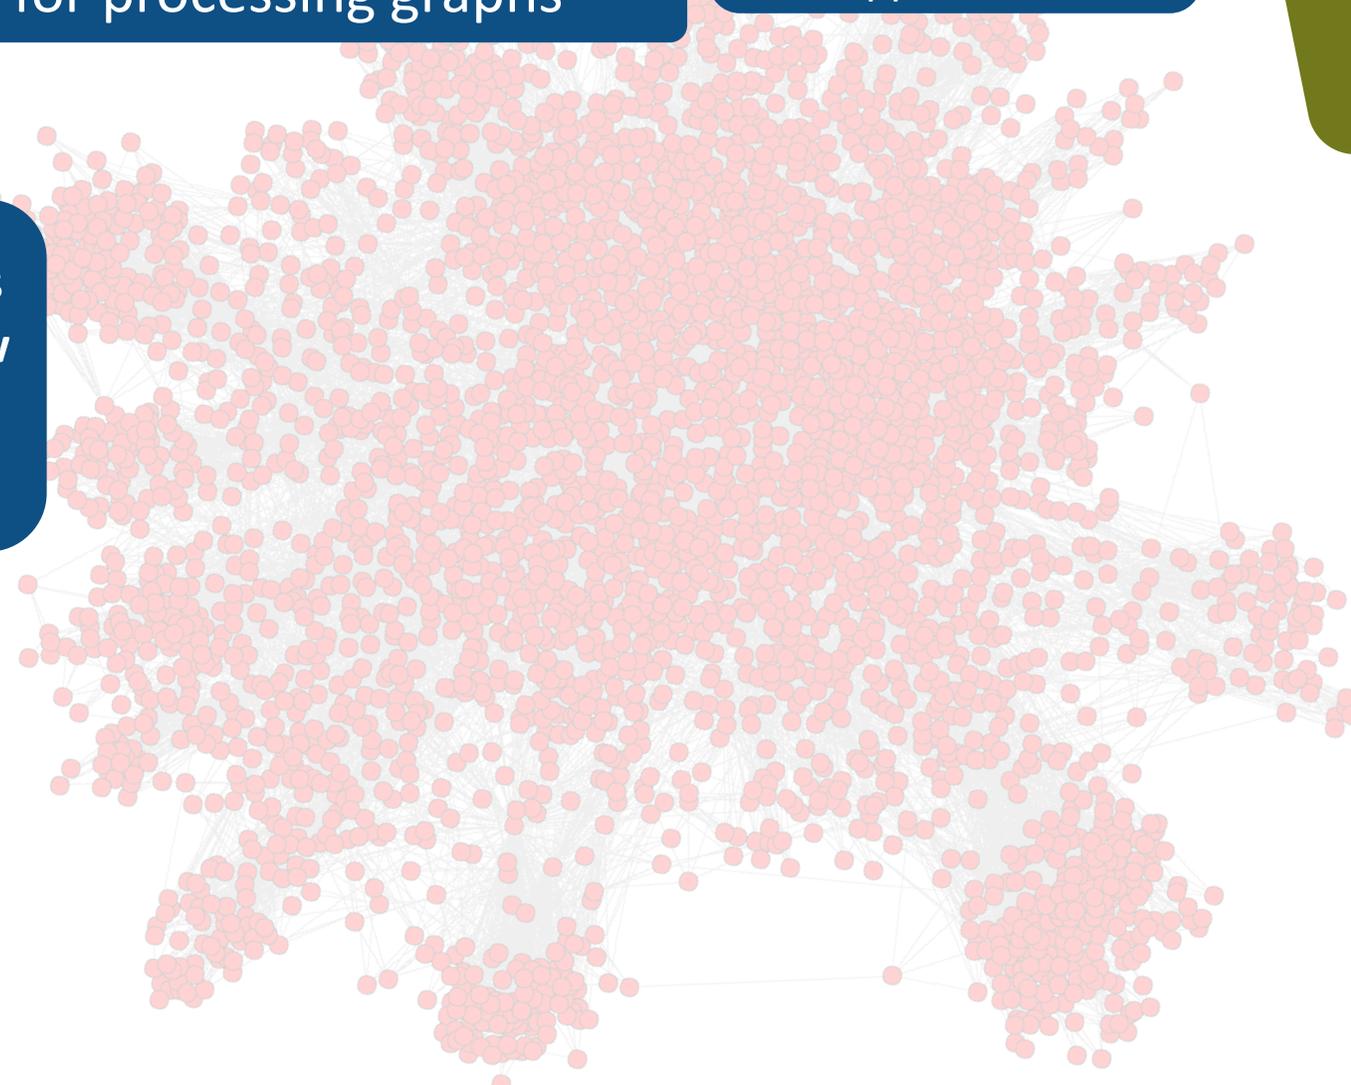
Substream-Centric Graph Processing

A new paradigm for processing graphs

It enhances edge-centric streaming approaches

Use some form of streaming (aka **edge-centric**); we can use pipelining efficiently ("streaming \approx pipelining")

...Processing edges is sequential – how to incorporate parallelism?



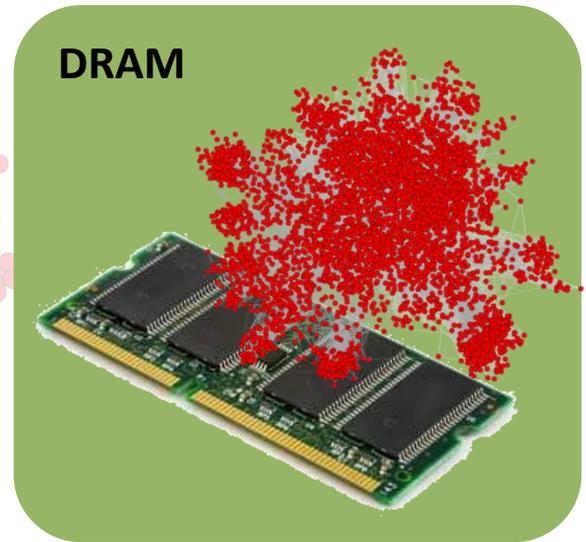
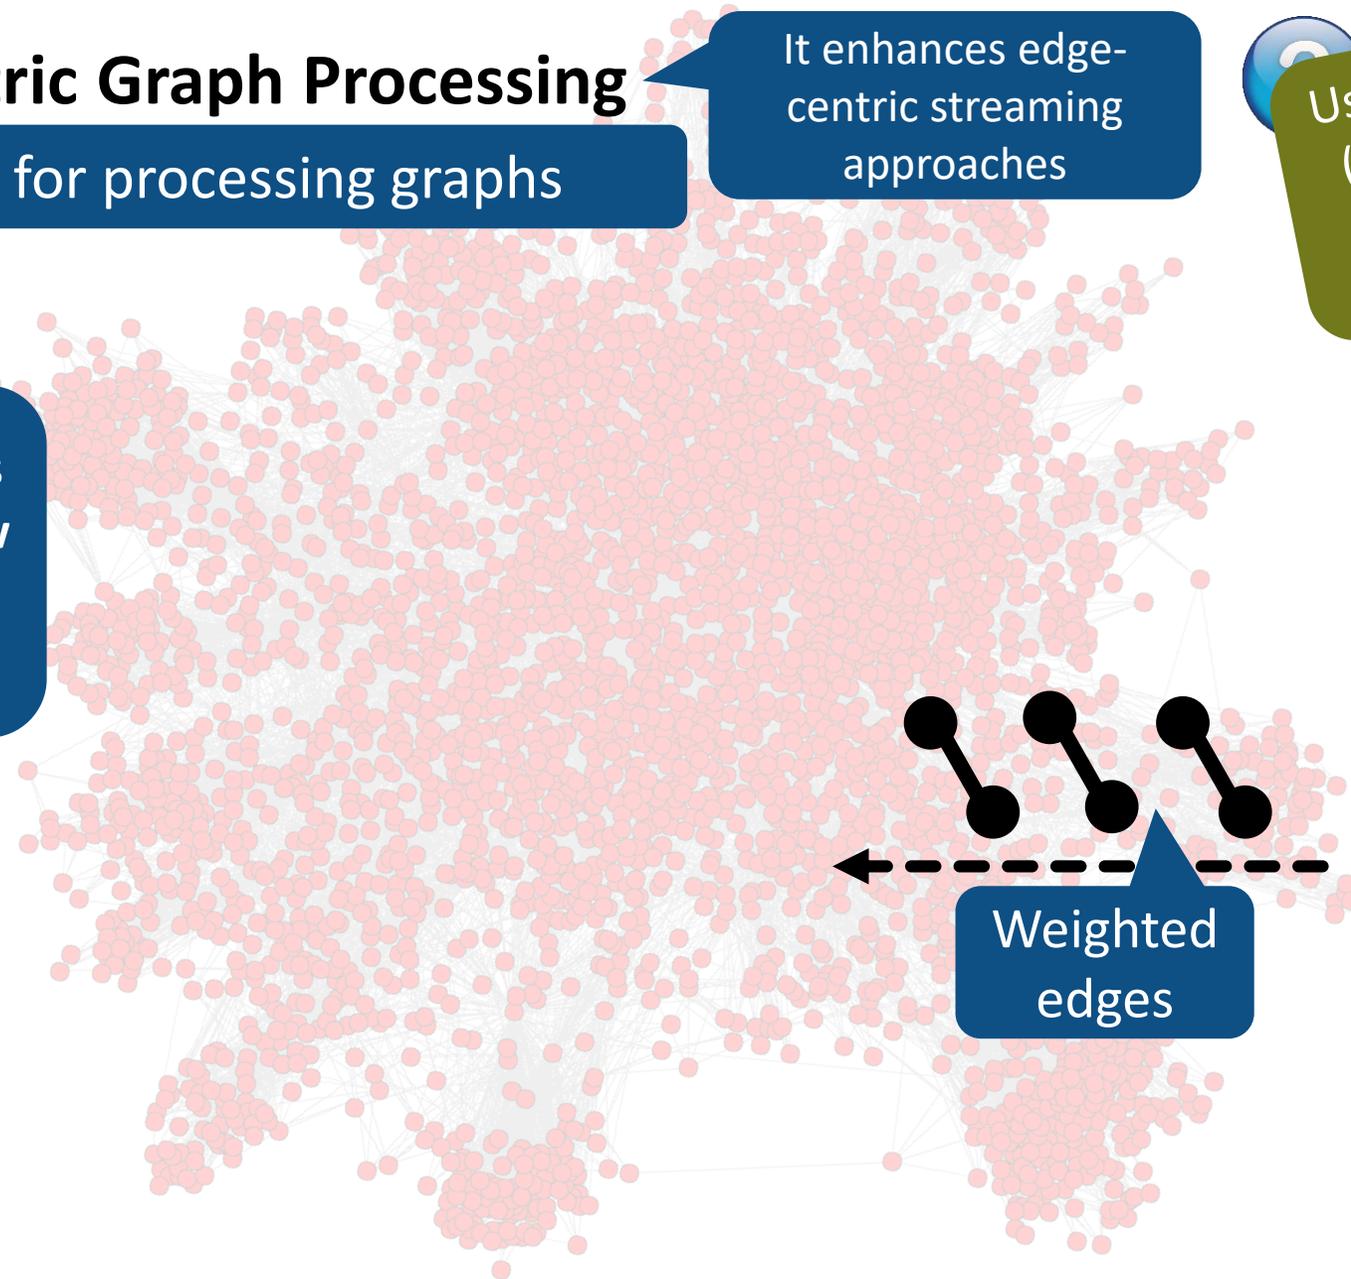
Substream-Centric Graph Processing

A new paradigm for processing graphs

It enhances edge-centric streaming approaches

Use some form of streaming (aka edge-centric); we can use pipelining efficiently ("streaming \approx pipelining")

...Processing edges is sequential – how to incorporate parallelism?



Substream-Centric Graph Processing

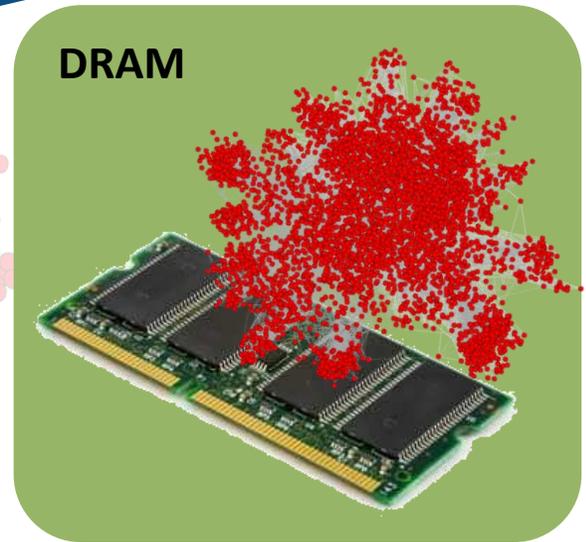
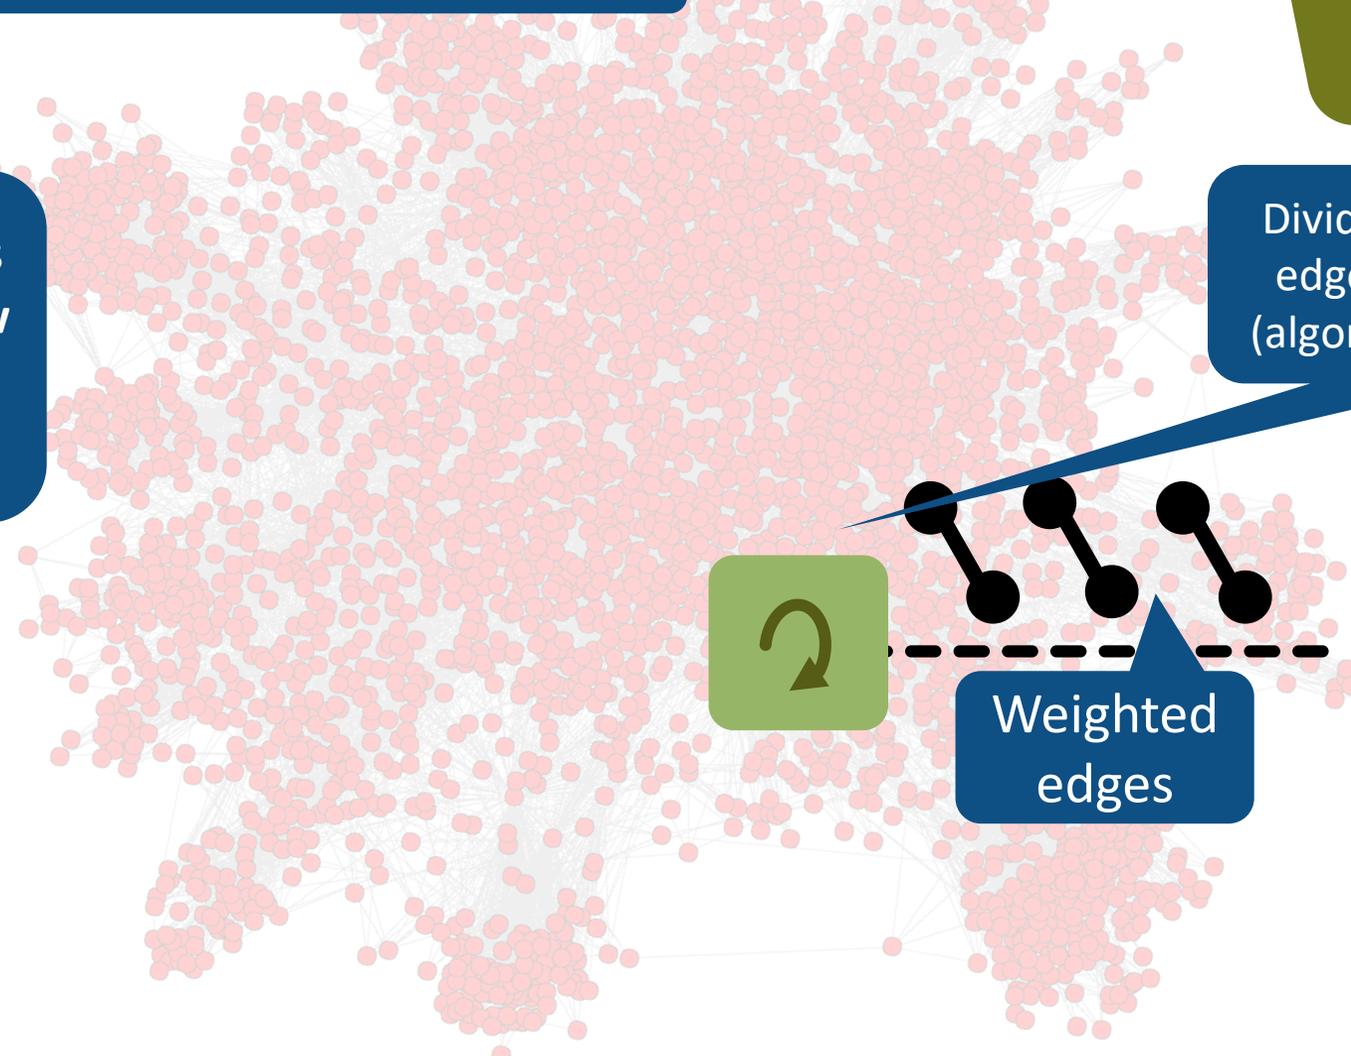
A new paradigm for processing graphs

It enhances edge-centric streaming approaches

Use some form of streaming (aka **edge-centric**); we can use pipelining efficiently ("streaming \approx pipelining")

...Processing edges is sequential – **how to incorporate parallelism?**

Divide the input stream of edges according to some (algorithm-specific) pattern



Substream-Centric Graph Processing

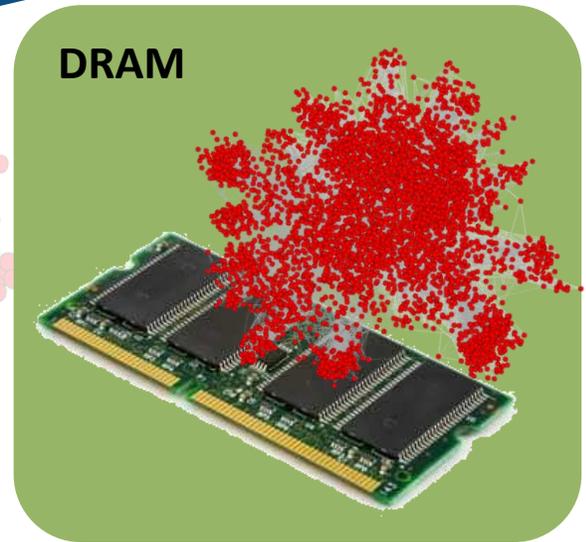
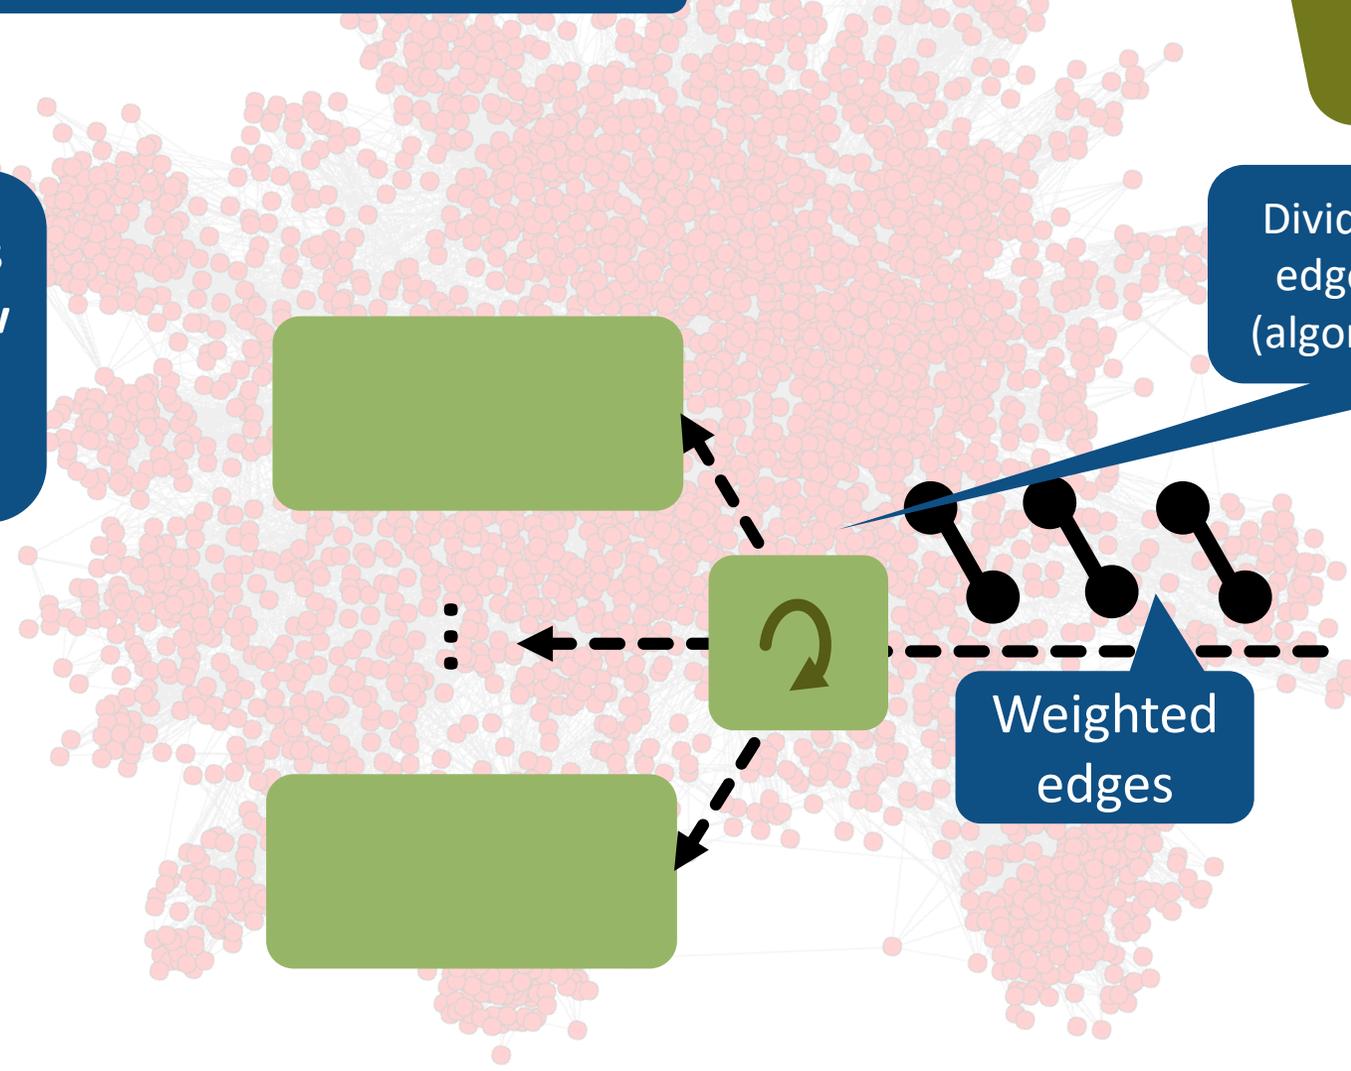
A new paradigm for processing graphs

It enhances edge-centric streaming approaches

Use some form of streaming (aka **edge-centric**); we can use pipelining efficiently ("streaming \approx pipelining")

...Processing edges is sequential – **how to incorporate parallelism?**

Divide the input stream of edges according to some (algorithm-specific) pattern



Substream-Centric Graph Processing

A new paradigm for processing graphs

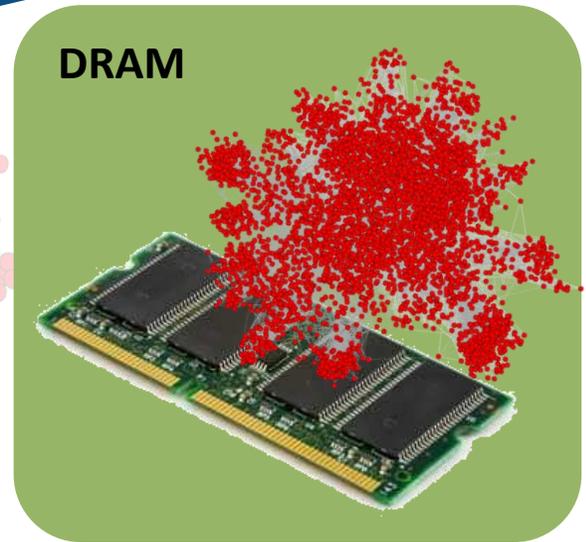
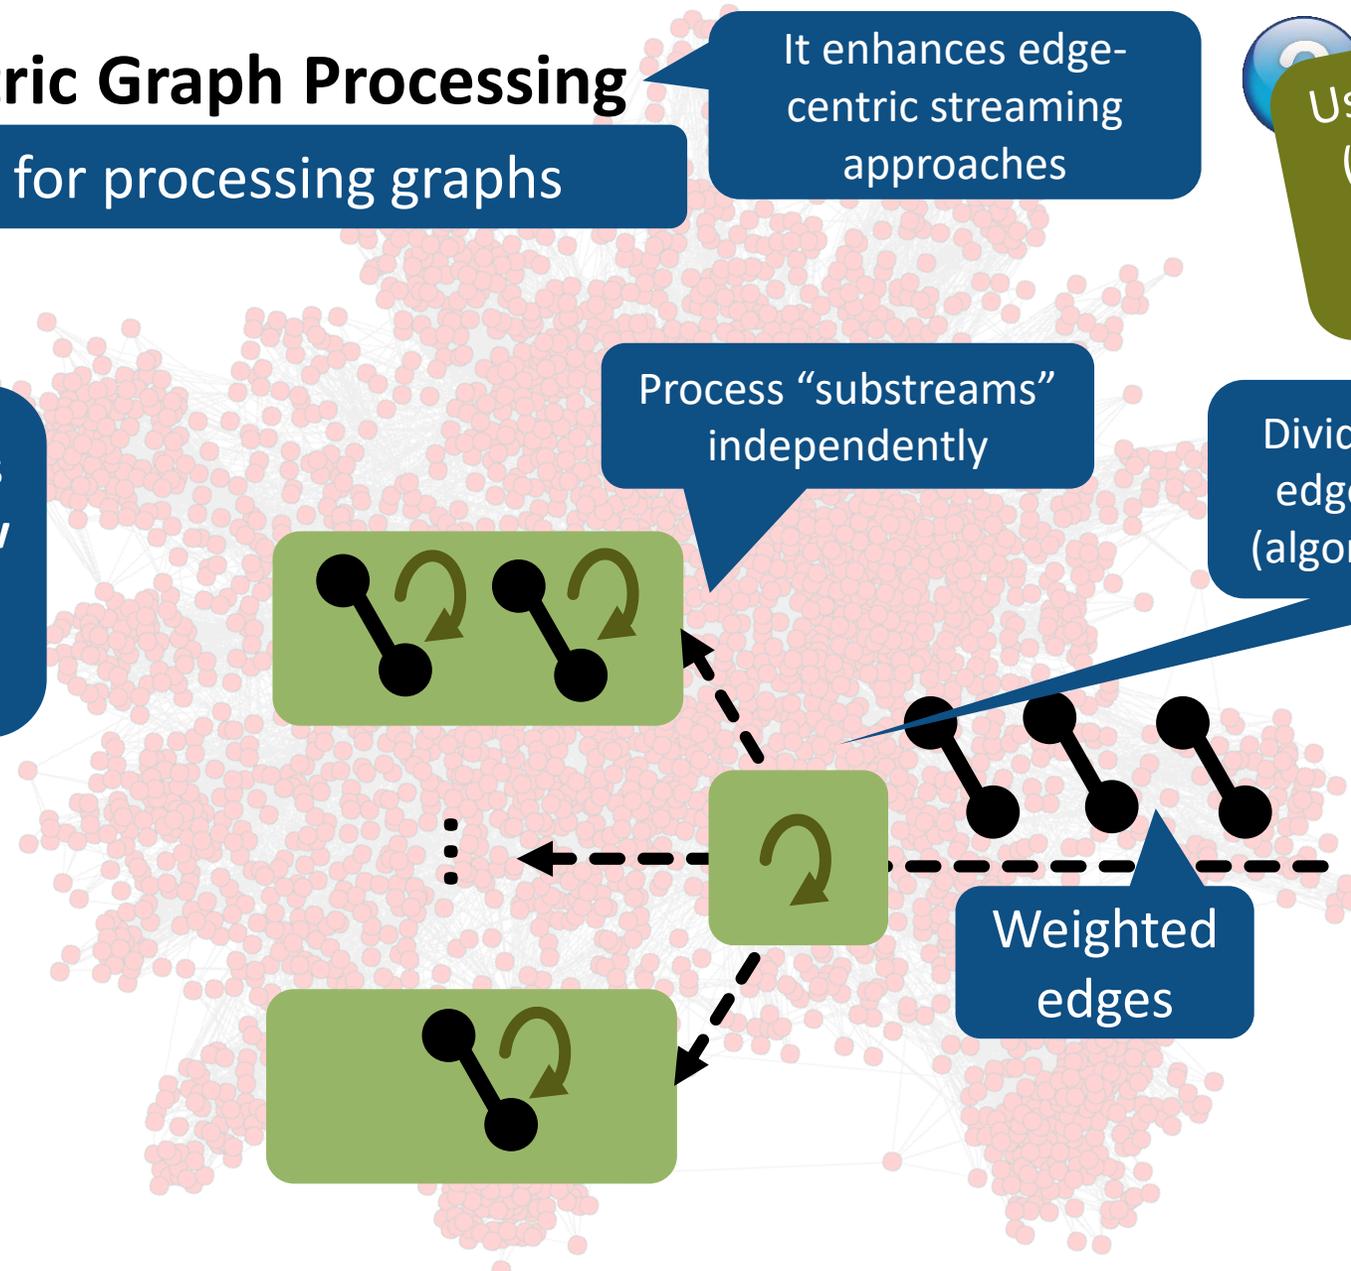
It enhances edge-centric streaming approaches

Use some form of streaming (aka **edge-centric**); we can use pipelining efficiently ("streaming \approx pipelining")

...Processing edges is sequential – **how to incorporate parallelism?**

Process "substreams" independently

Divide the input stream of edges according to some (algorithm-specific) pattern



Substream-Centric Graph Processing

A new paradigm for processing graphs

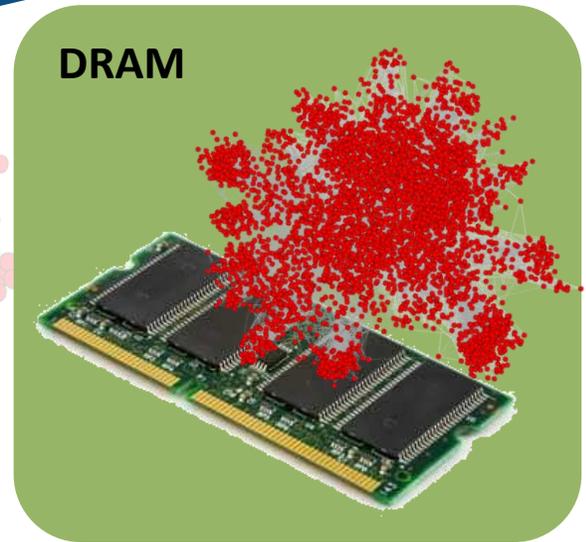
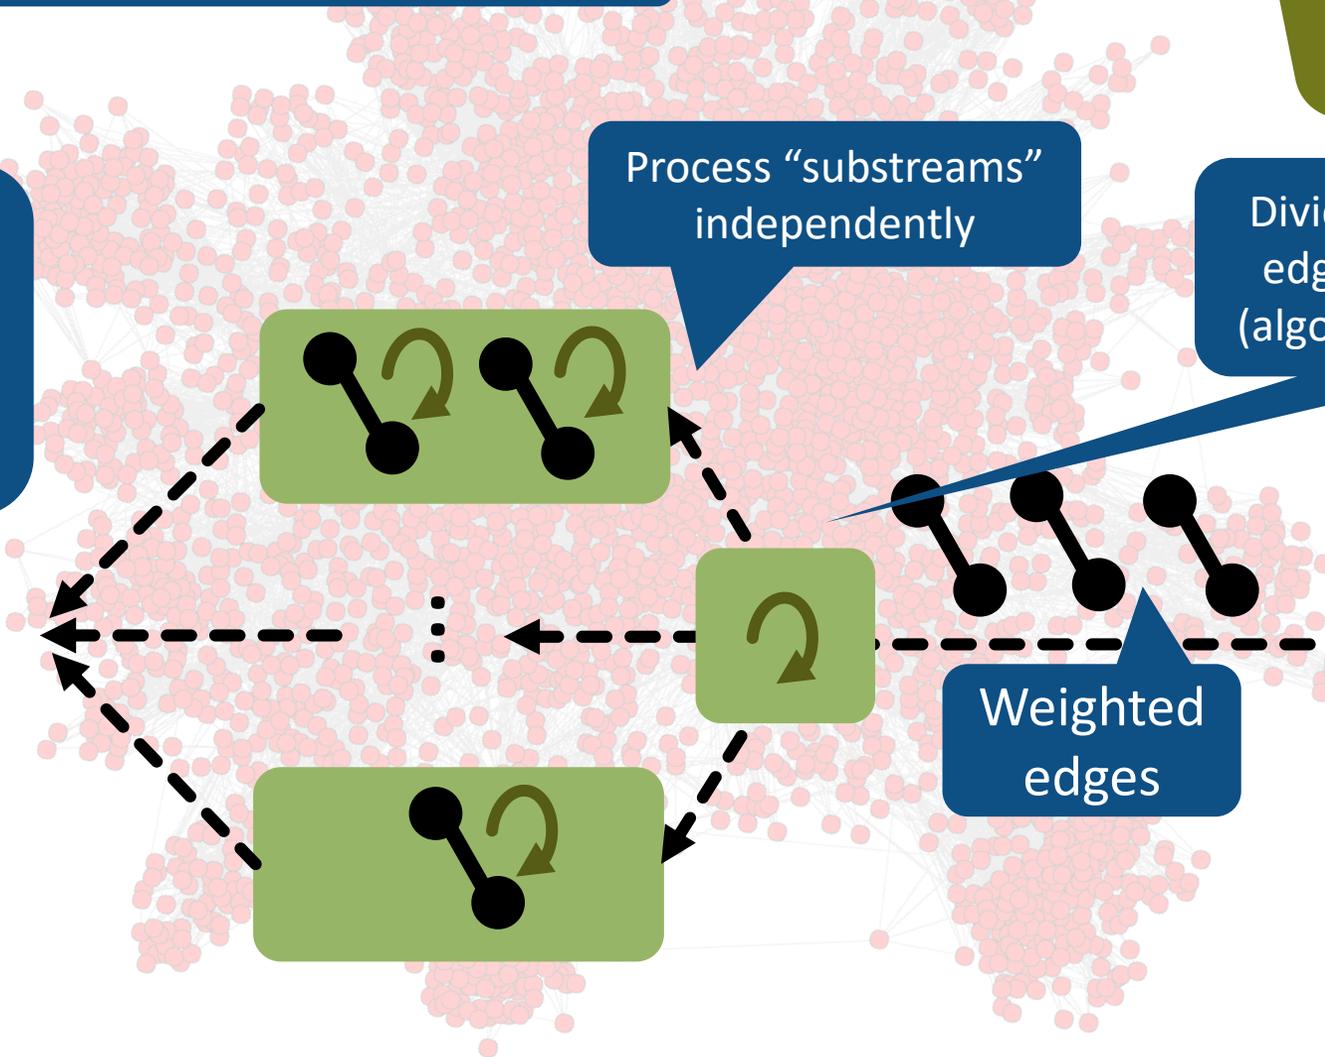
It enhances edge-centric streaming approaches

Use some form of streaming (aka **edge-centric**); we can use pipelining efficiently ("streaming \approx pipelining")

...Processing edges is sequential – **how to incorporate parallelism?**

Process "substreams" independently

Divide the input stream of edges according to some (algorithm-specific) pattern



Substream-Centric Graph Processing

A new paradigm for processing graphs

It enhances edge-centric streaming approaches

Use some form of streaming (aka **edge-centric**); we can use pipelining efficiently ("streaming \approx pipelining")

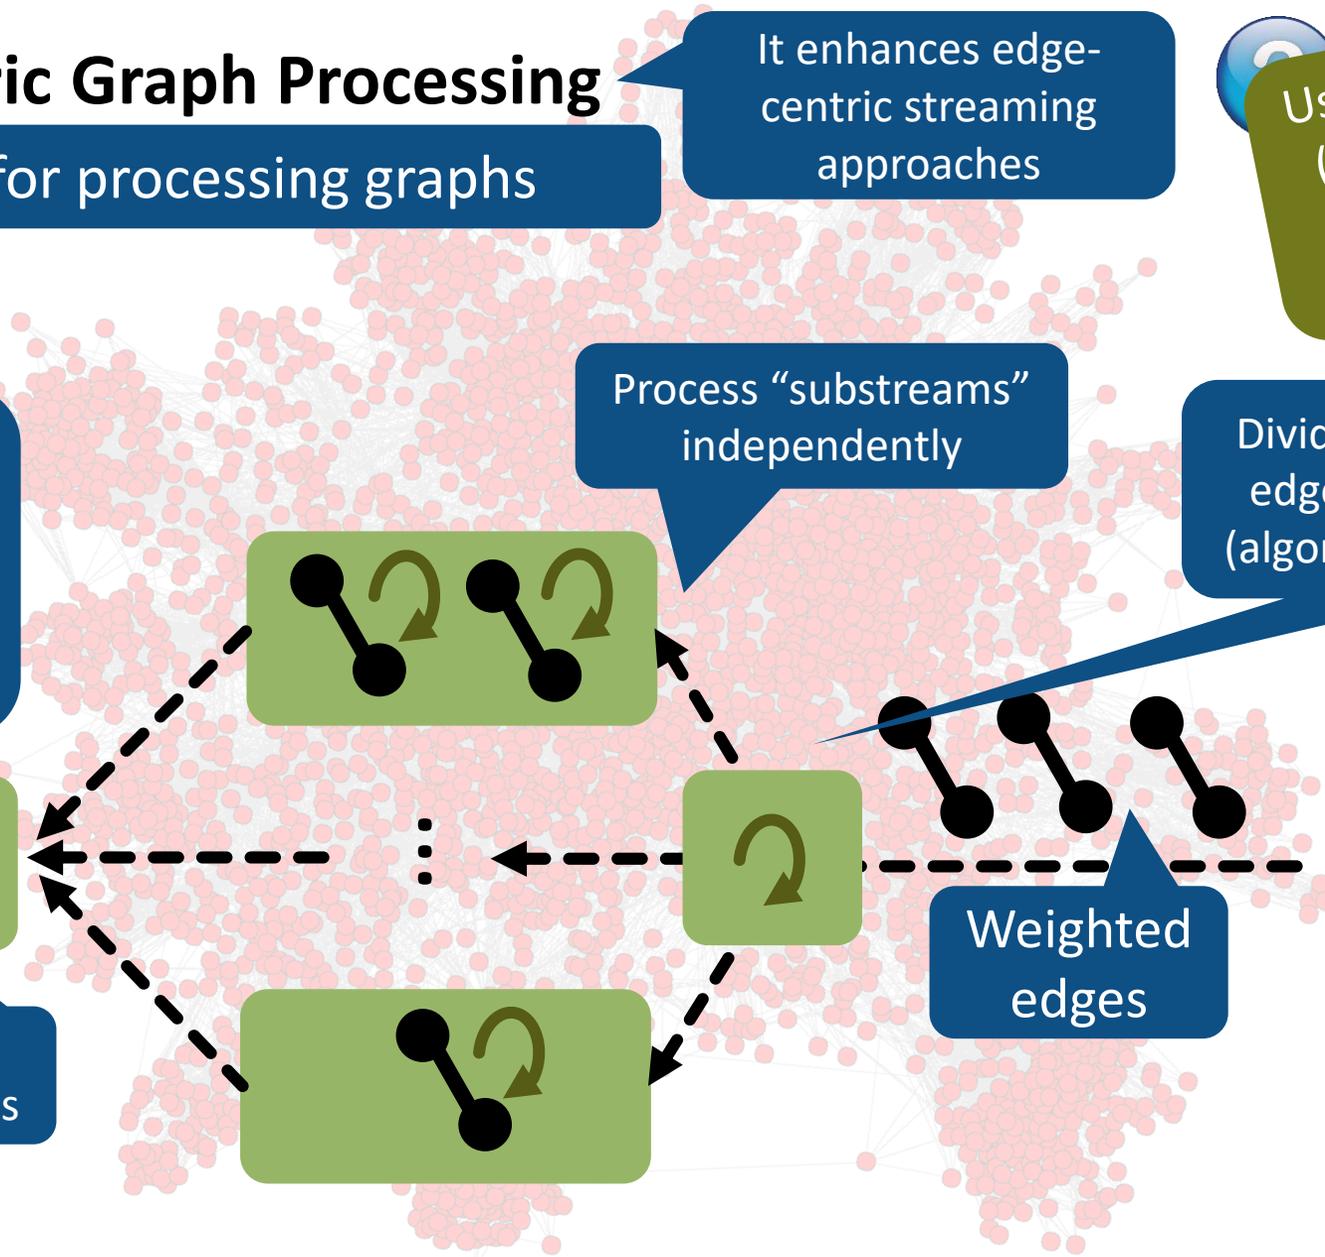
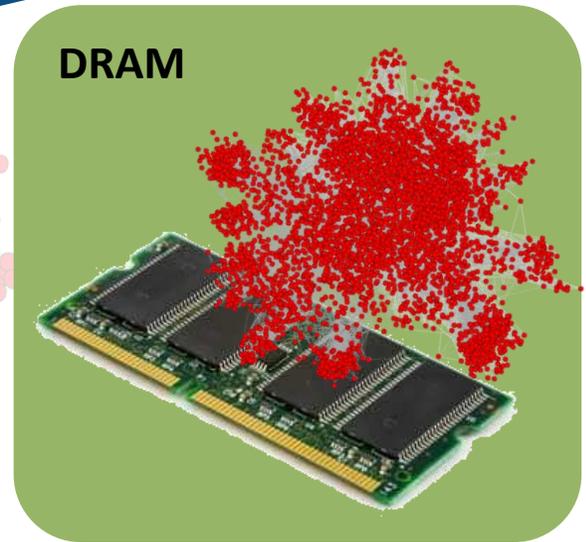
...Processing edges is sequential – **how to incorporate parallelism?**

Process "substreams" independently

Divide the input stream of edges according to some (algorithm-specific) pattern

Merge substreams

Weighted edges



Substream-Centric Graph Processing

A new paradigm for processing graphs

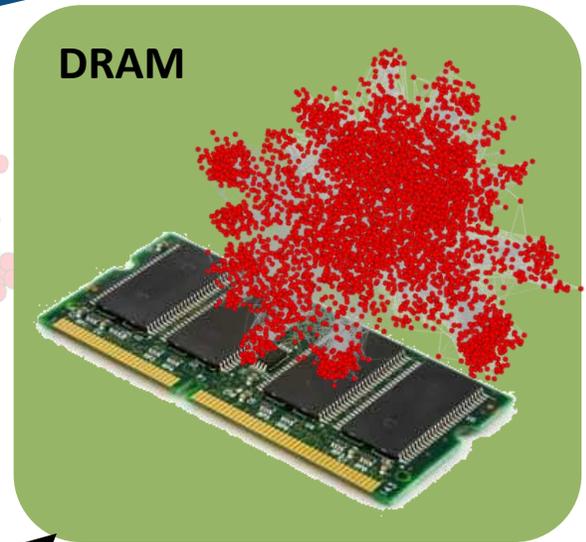
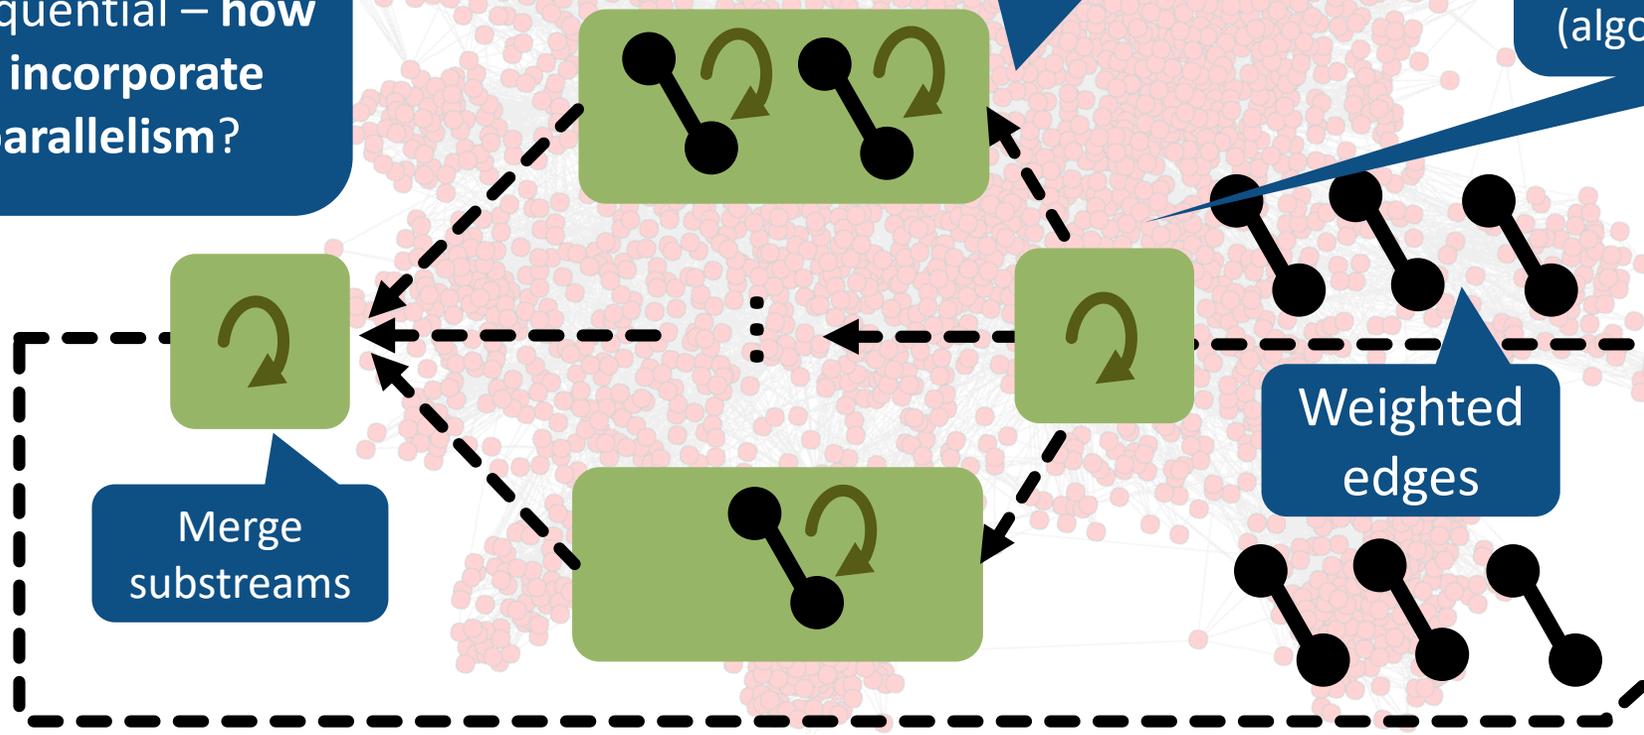
It enhances edge-centric streaming approaches

Use some form of streaming (aka edge-centric); we can use pipelining efficiently ("streaming \approx pipelining")

...Processing edges is sequential – how to incorporate parallelism?

Process "substreams" independently

Divide the input stream of edges according to some (algorithm-specific) pattern



Substream-Centric Graph Processing

A new paradigm for processing graphs

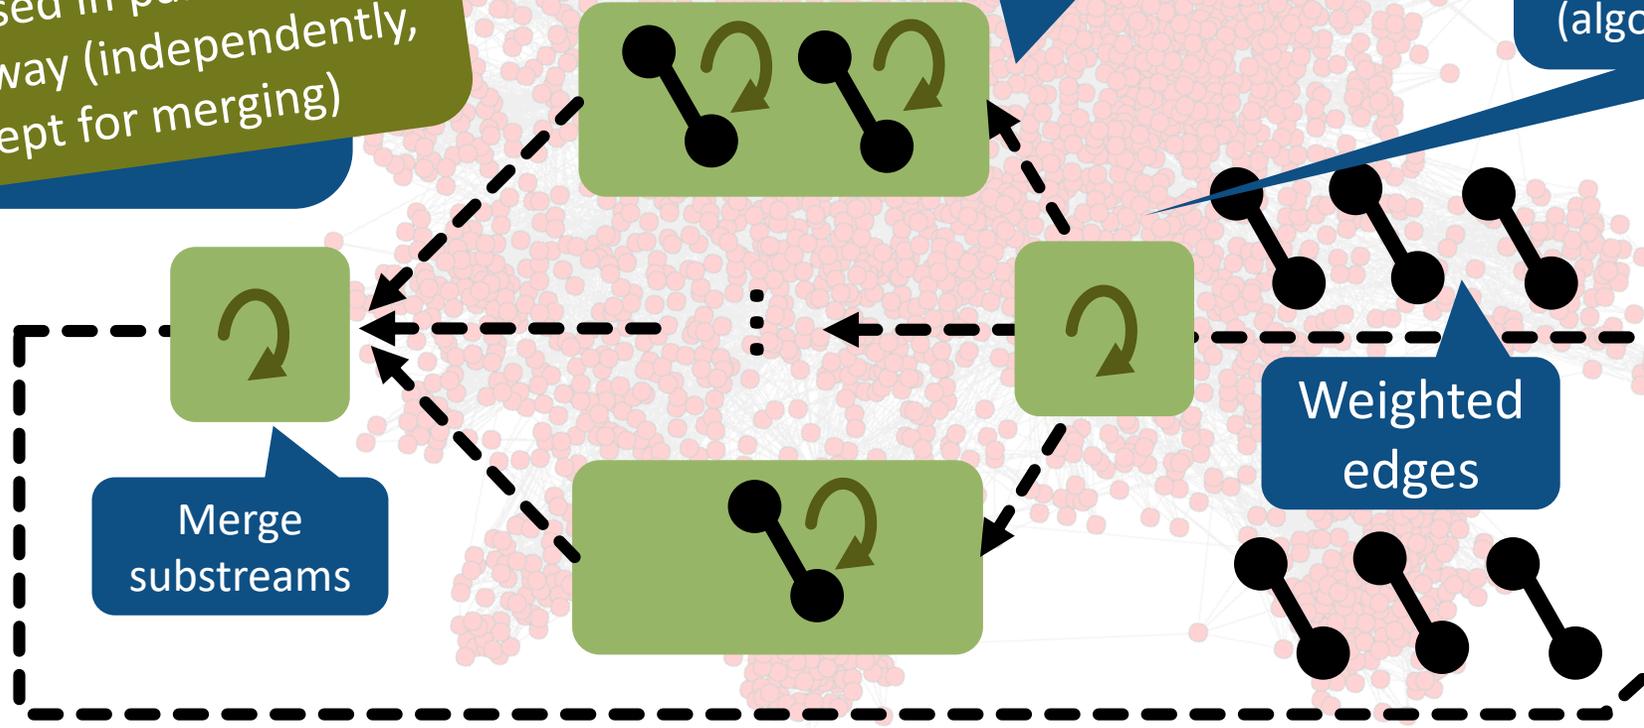
It enhances edge-centric streaming approaches

Use some form of streaming (aka **edge-centric**); we can use pipelining efficiently ("streaming \approx pipelining")

Substreams (pipelines) are processed in parallel, in a simple way (independently, except for merging)

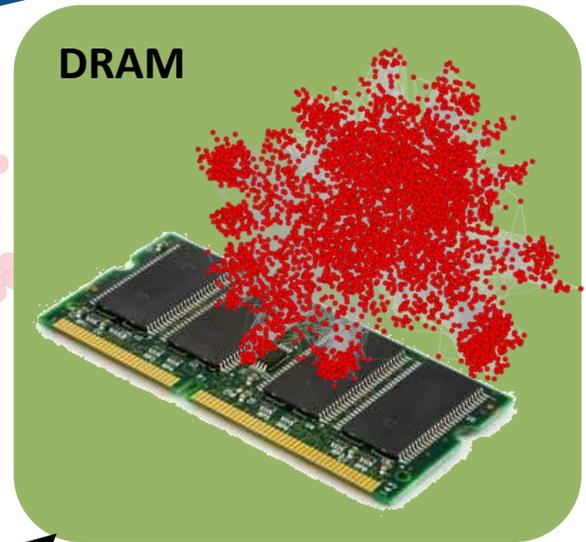
Process "substreams" independently

Divide the input stream of edges according to some (algorithm-specific) pattern



Merge substreams

Weighted edges



Substream-Centric Graph Processing

A new paradigm for processing graphs

It enhances edge-centric streaming approaches

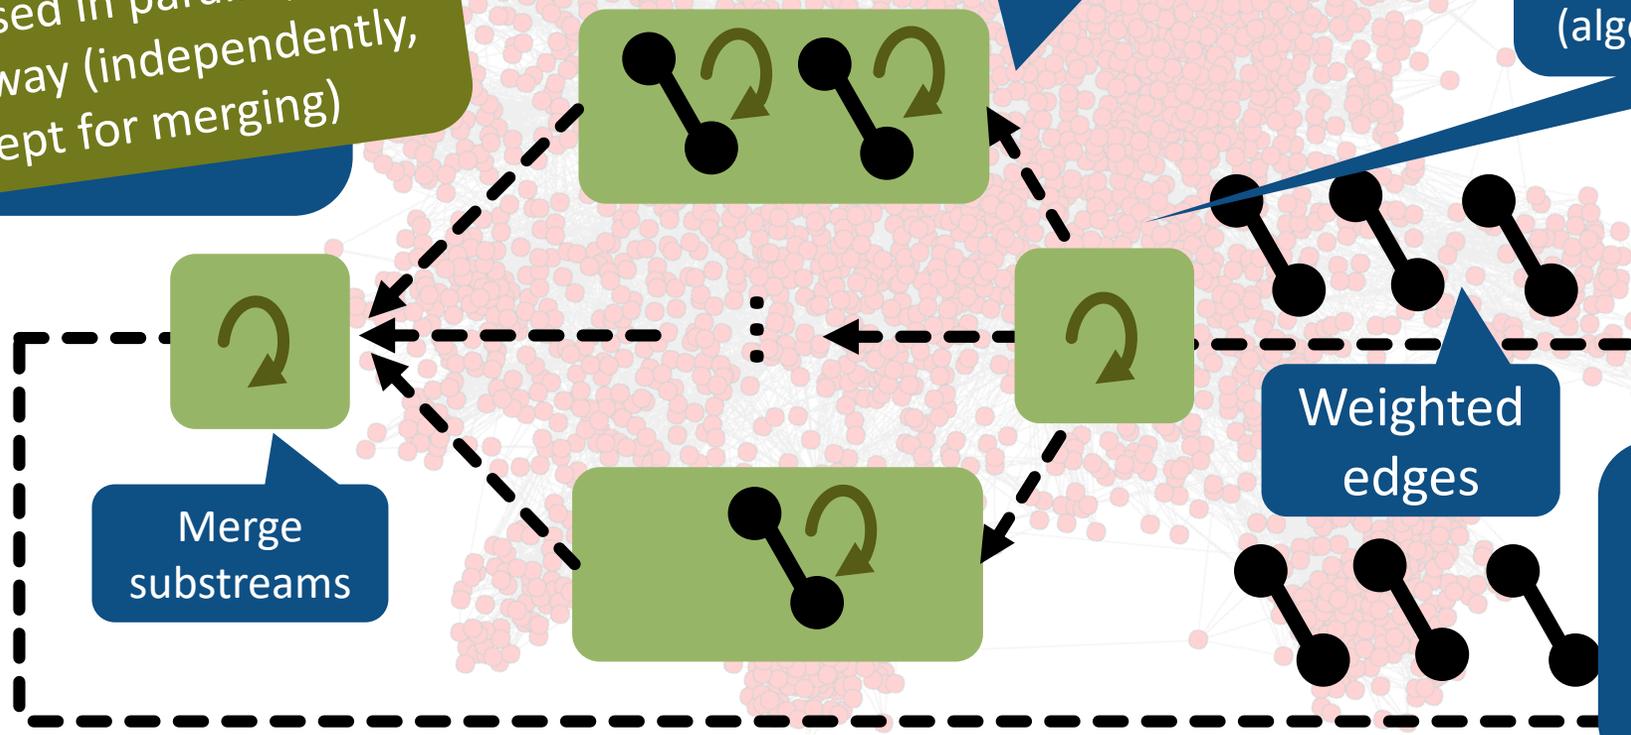
Use some form of streaming (aka **edge-centric**); we can use pipelining efficiently ("streaming \approx pipelining")

Substreams (pipelines) are processed in parallel, in a simple way (independently, except for merging)

Also, it enables (tunable) approximation and a (tunable) number of passes

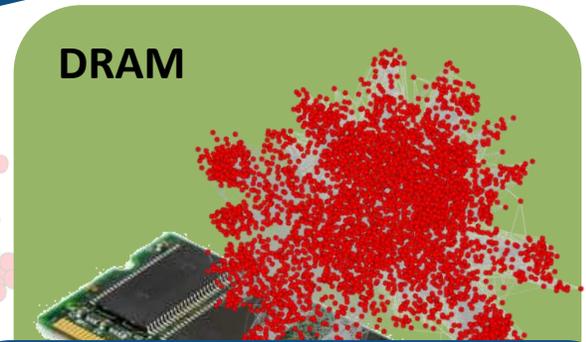
Process substreams independently

Divide the input stream of edges according to some (algorithm-specific) pattern



Merge substreams

Weighted edges



Use case: expressing maximum matchings in this paradigm?

Substream-Centric Graph Processing

A new paradigm for processing graphs

It enhances edge-centric streaming approaches

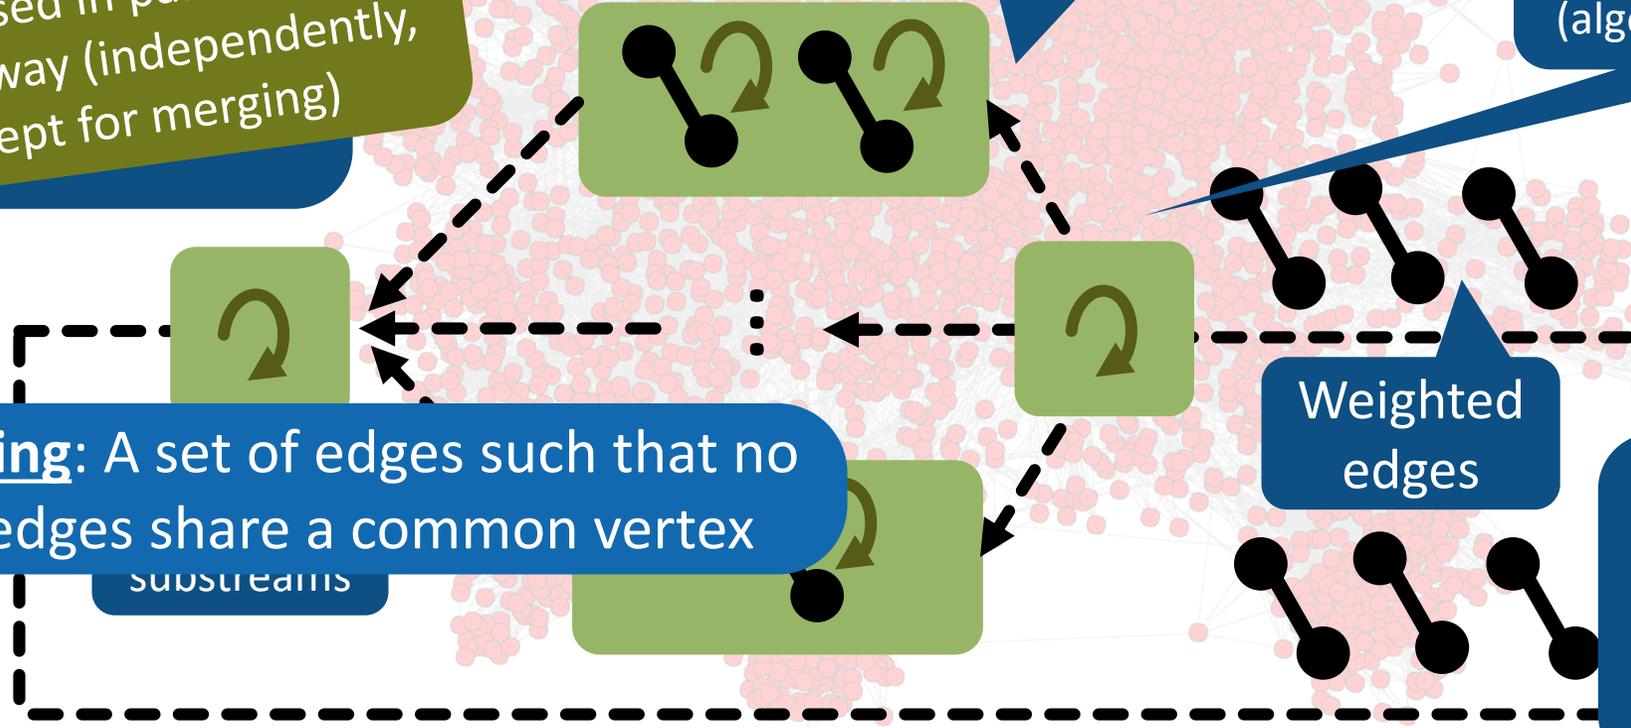
Use some form of streaming (aka **edge-centric**); we can use pipelining efficiently ("streaming \approx pipelining")

Substreams (pipelines) are processed in parallel, in a simple way (independently, except for merging)

Also, it enables (tunable) approximation and a (tunable) number of passes

...streams" independently

Divide the input stream of edges according to some (algorithm-specific) pattern



Matching: A set of edges such that no two edges share a common vertex

Use case: expressing maximum matchings in this paradigm?

Substream-Centric Graph Processing

A new paradigm for processing graphs

It enhances edge-centric streaming approaches

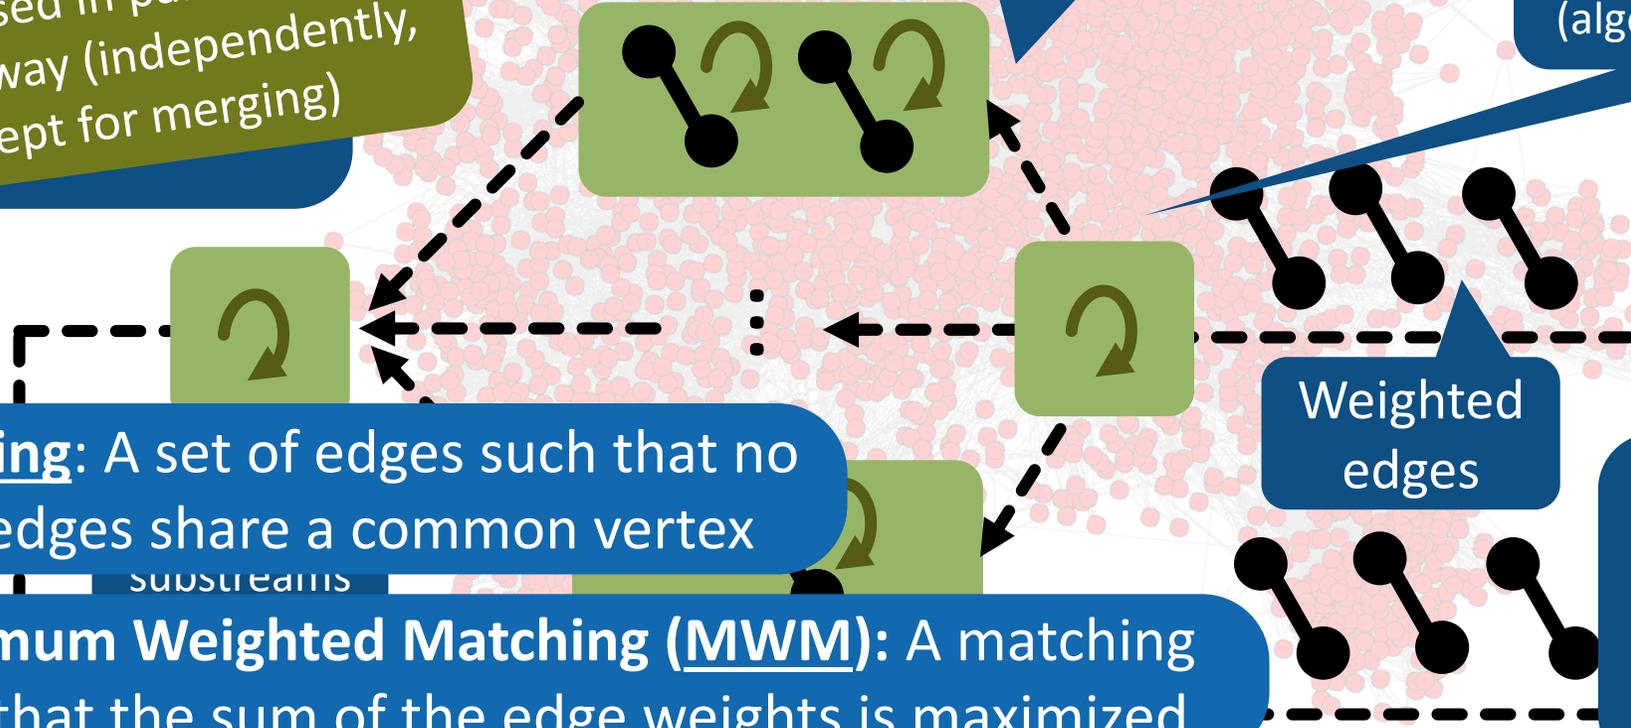
Use some form of streaming (aka **edge-centric**); we can use pipelining efficiently ("streaming \approx pipelining")

Also, it enables (tunable) approximation and a (tunable) number of passes

Substreams independently

Divide the input stream of edges according to some (algorithm-specific) pattern

Substreams (pipelines) are processed in parallel, in a simple way (independently, except for merging)

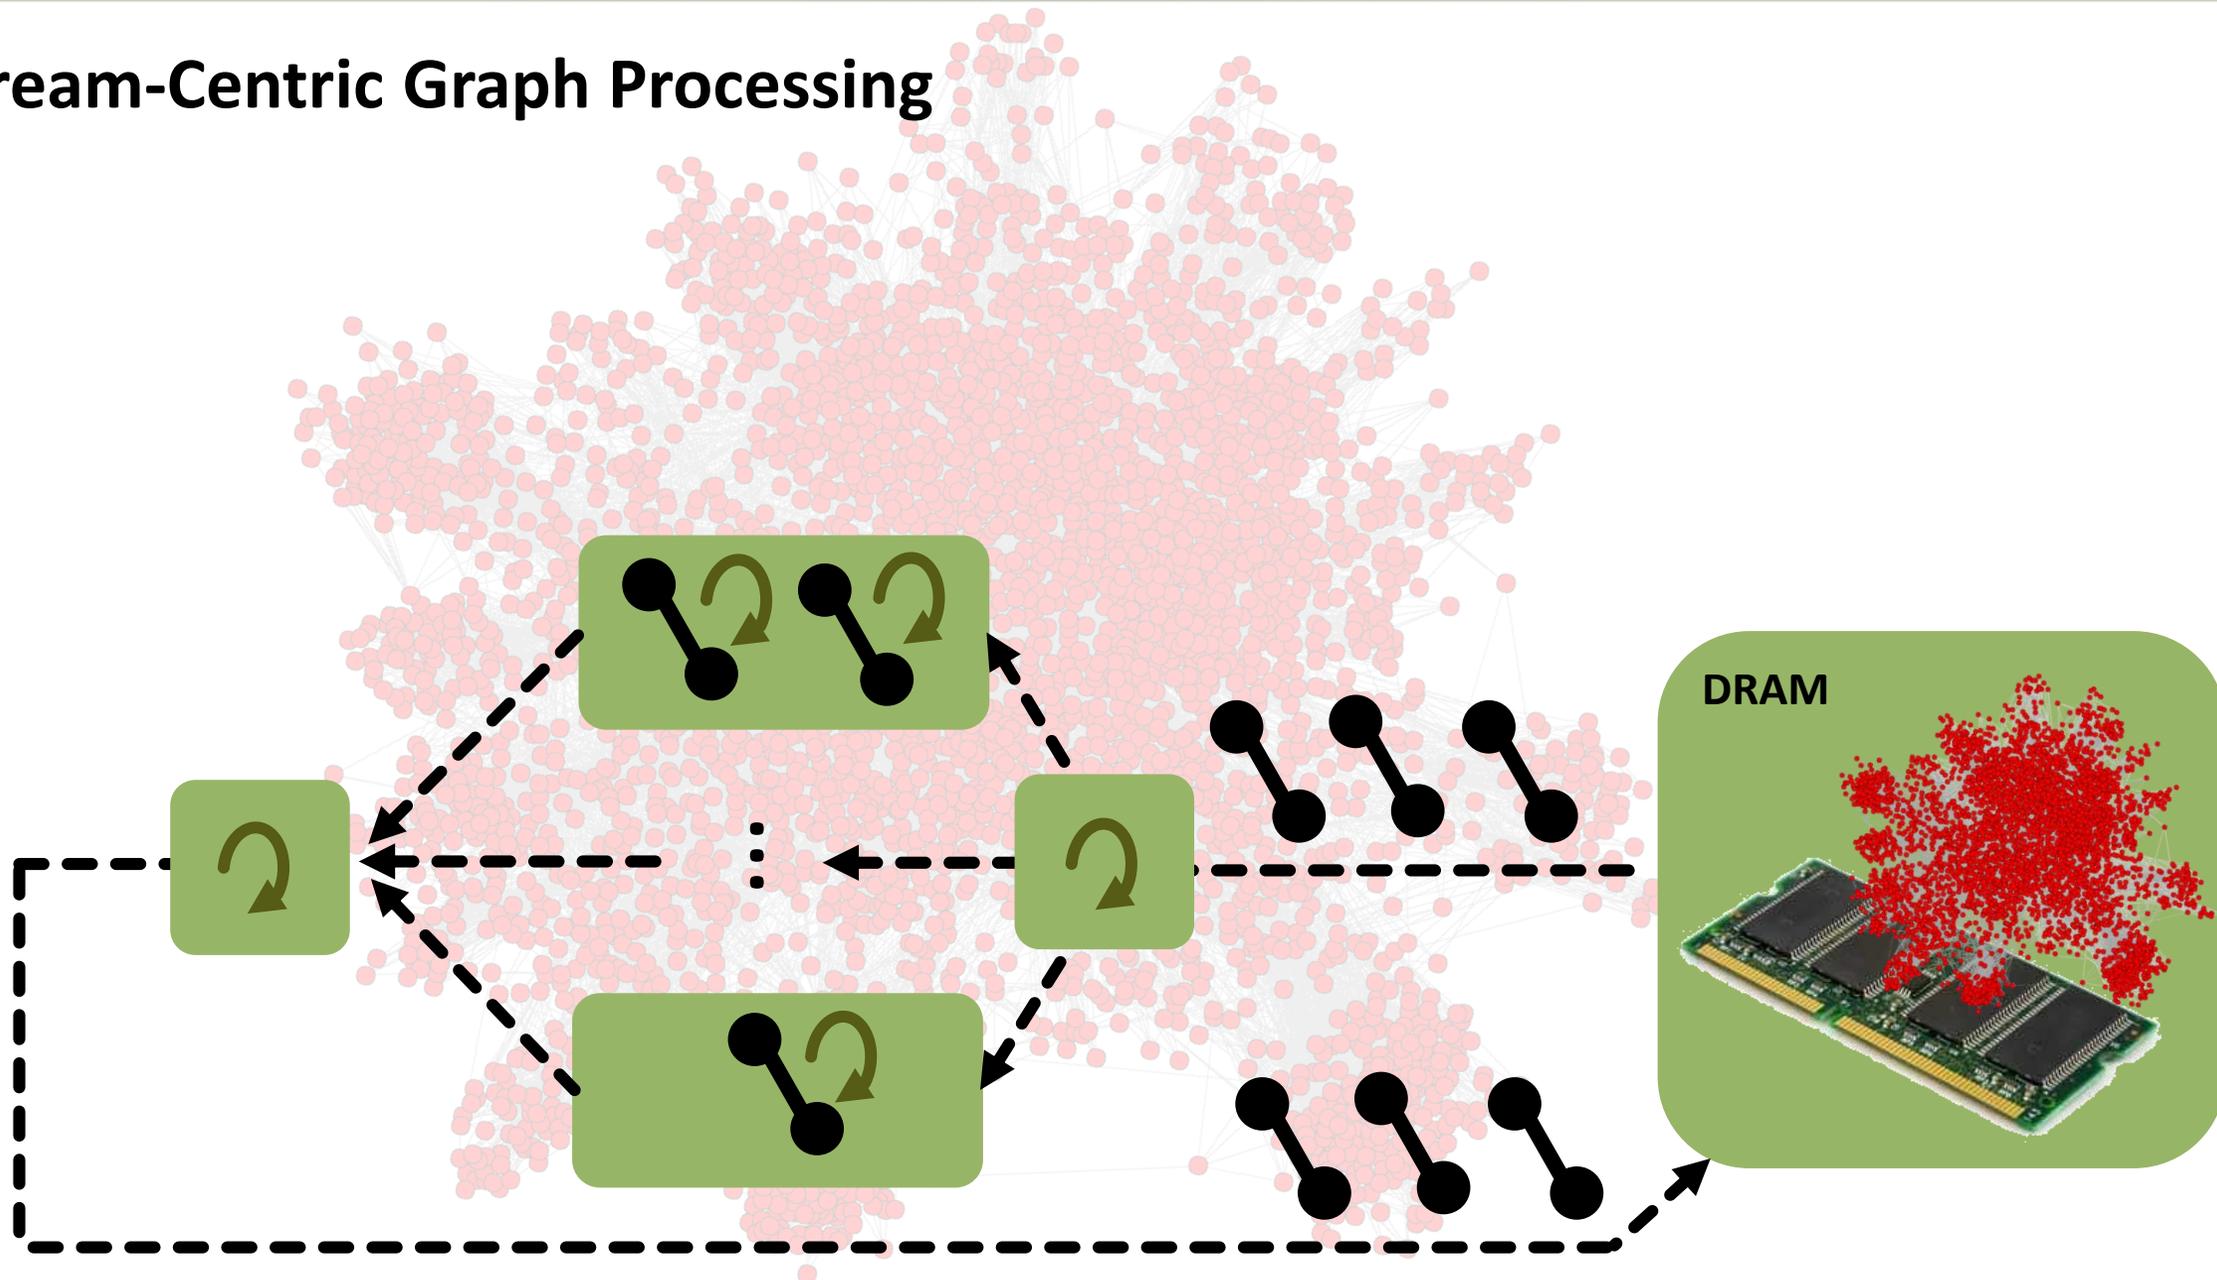


Matching: A set of edges such that no two edges share a common vertex

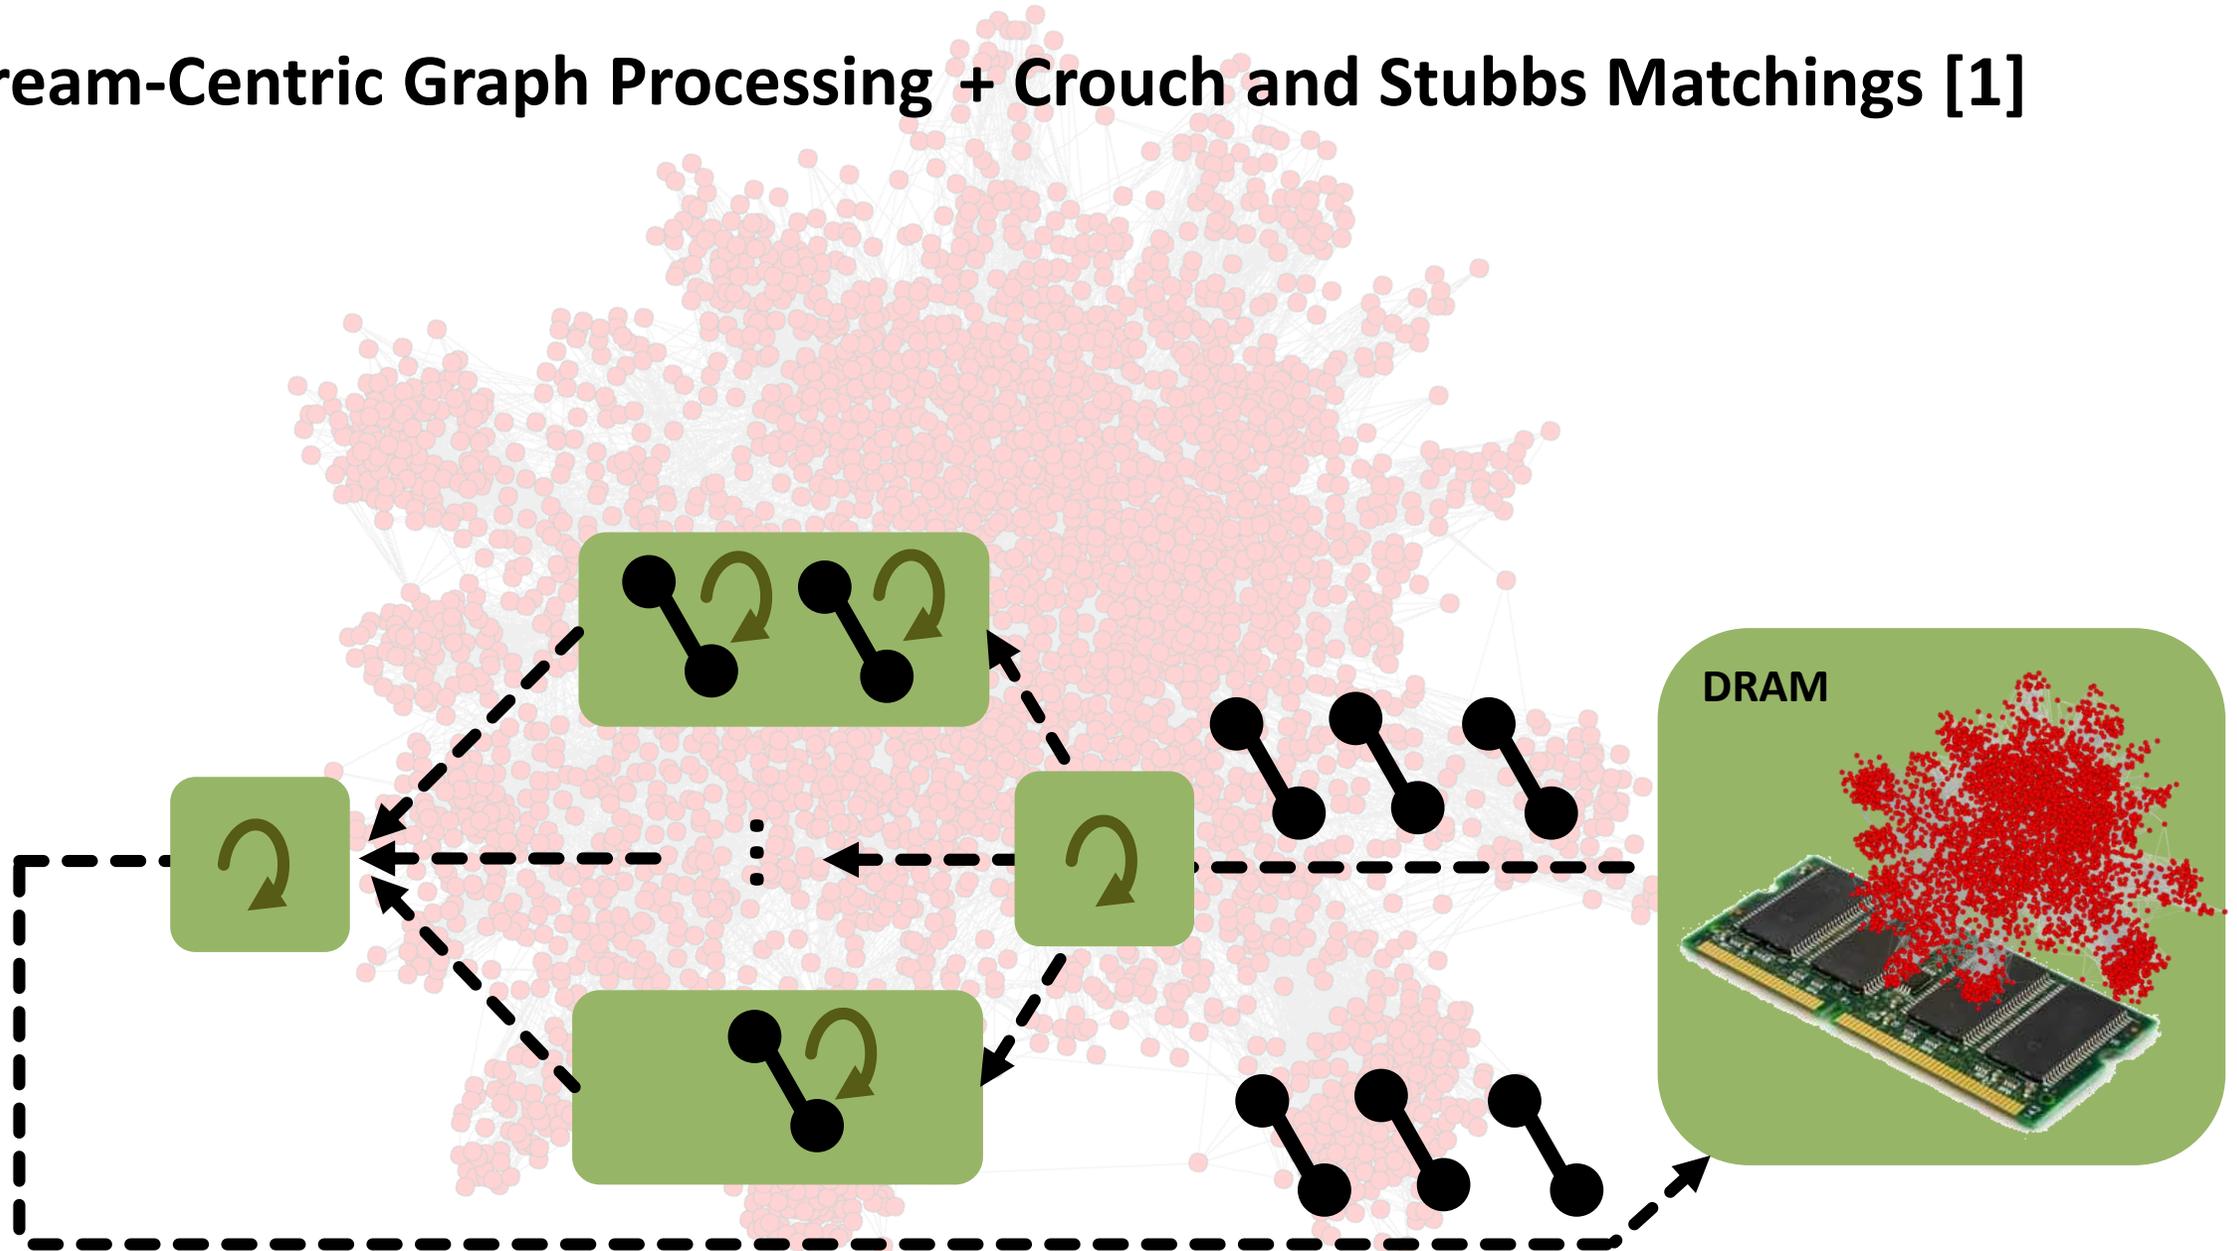
Maximum Weighted Matching (MWM): A matching such that the sum of the edge weights is maximized

Use case: expressing maximum matchings in this paradigm?

Substream-Centric Graph Processing

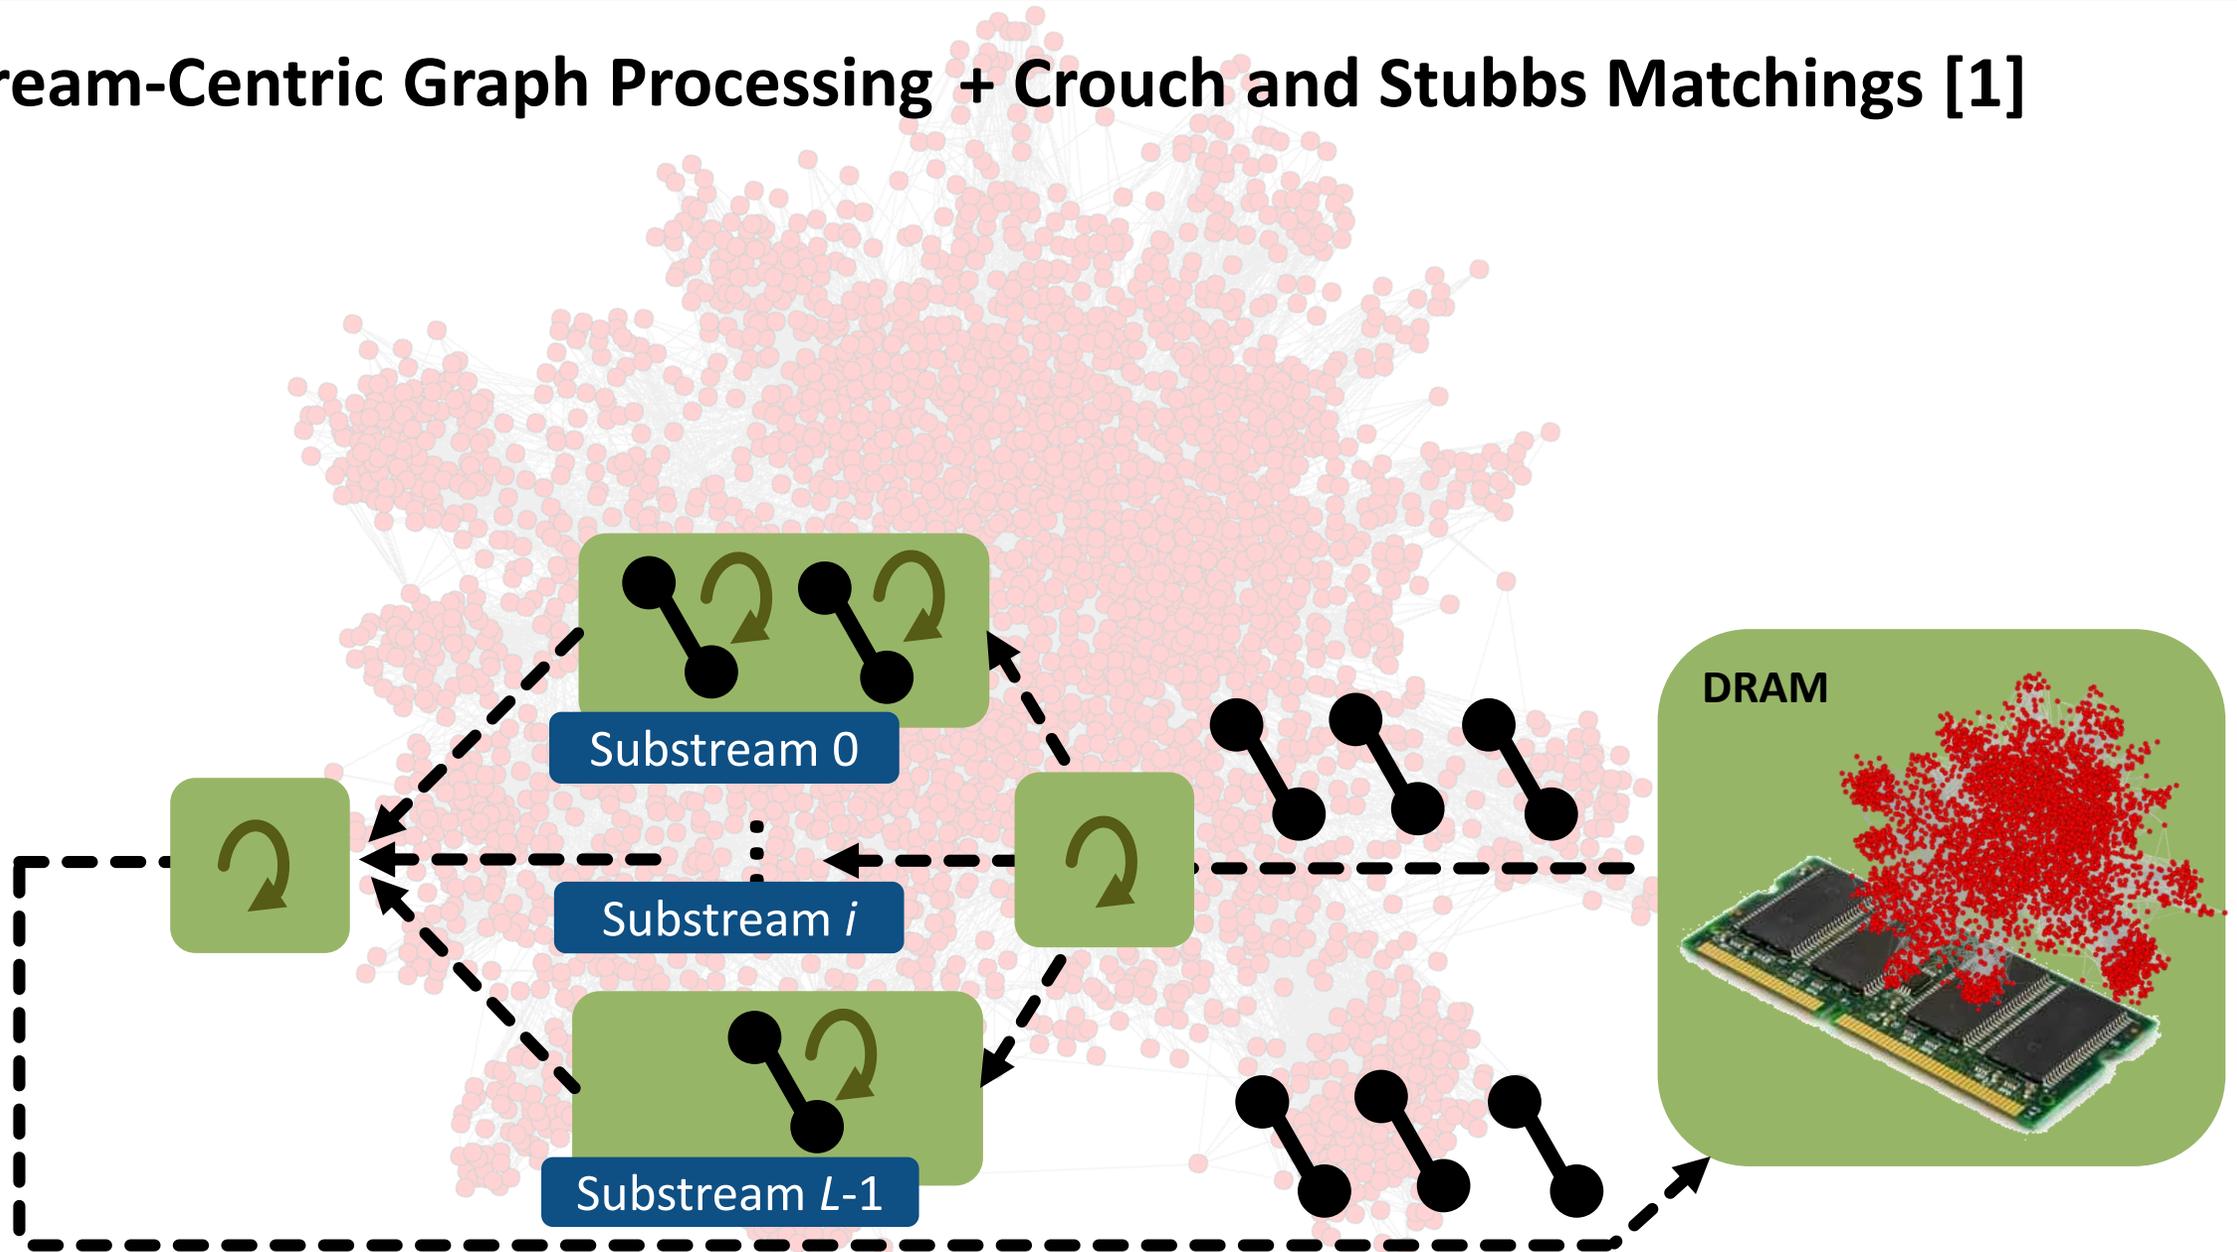


Substream-Centric Graph Processing + Crouch and Stubbs Matchings [1]



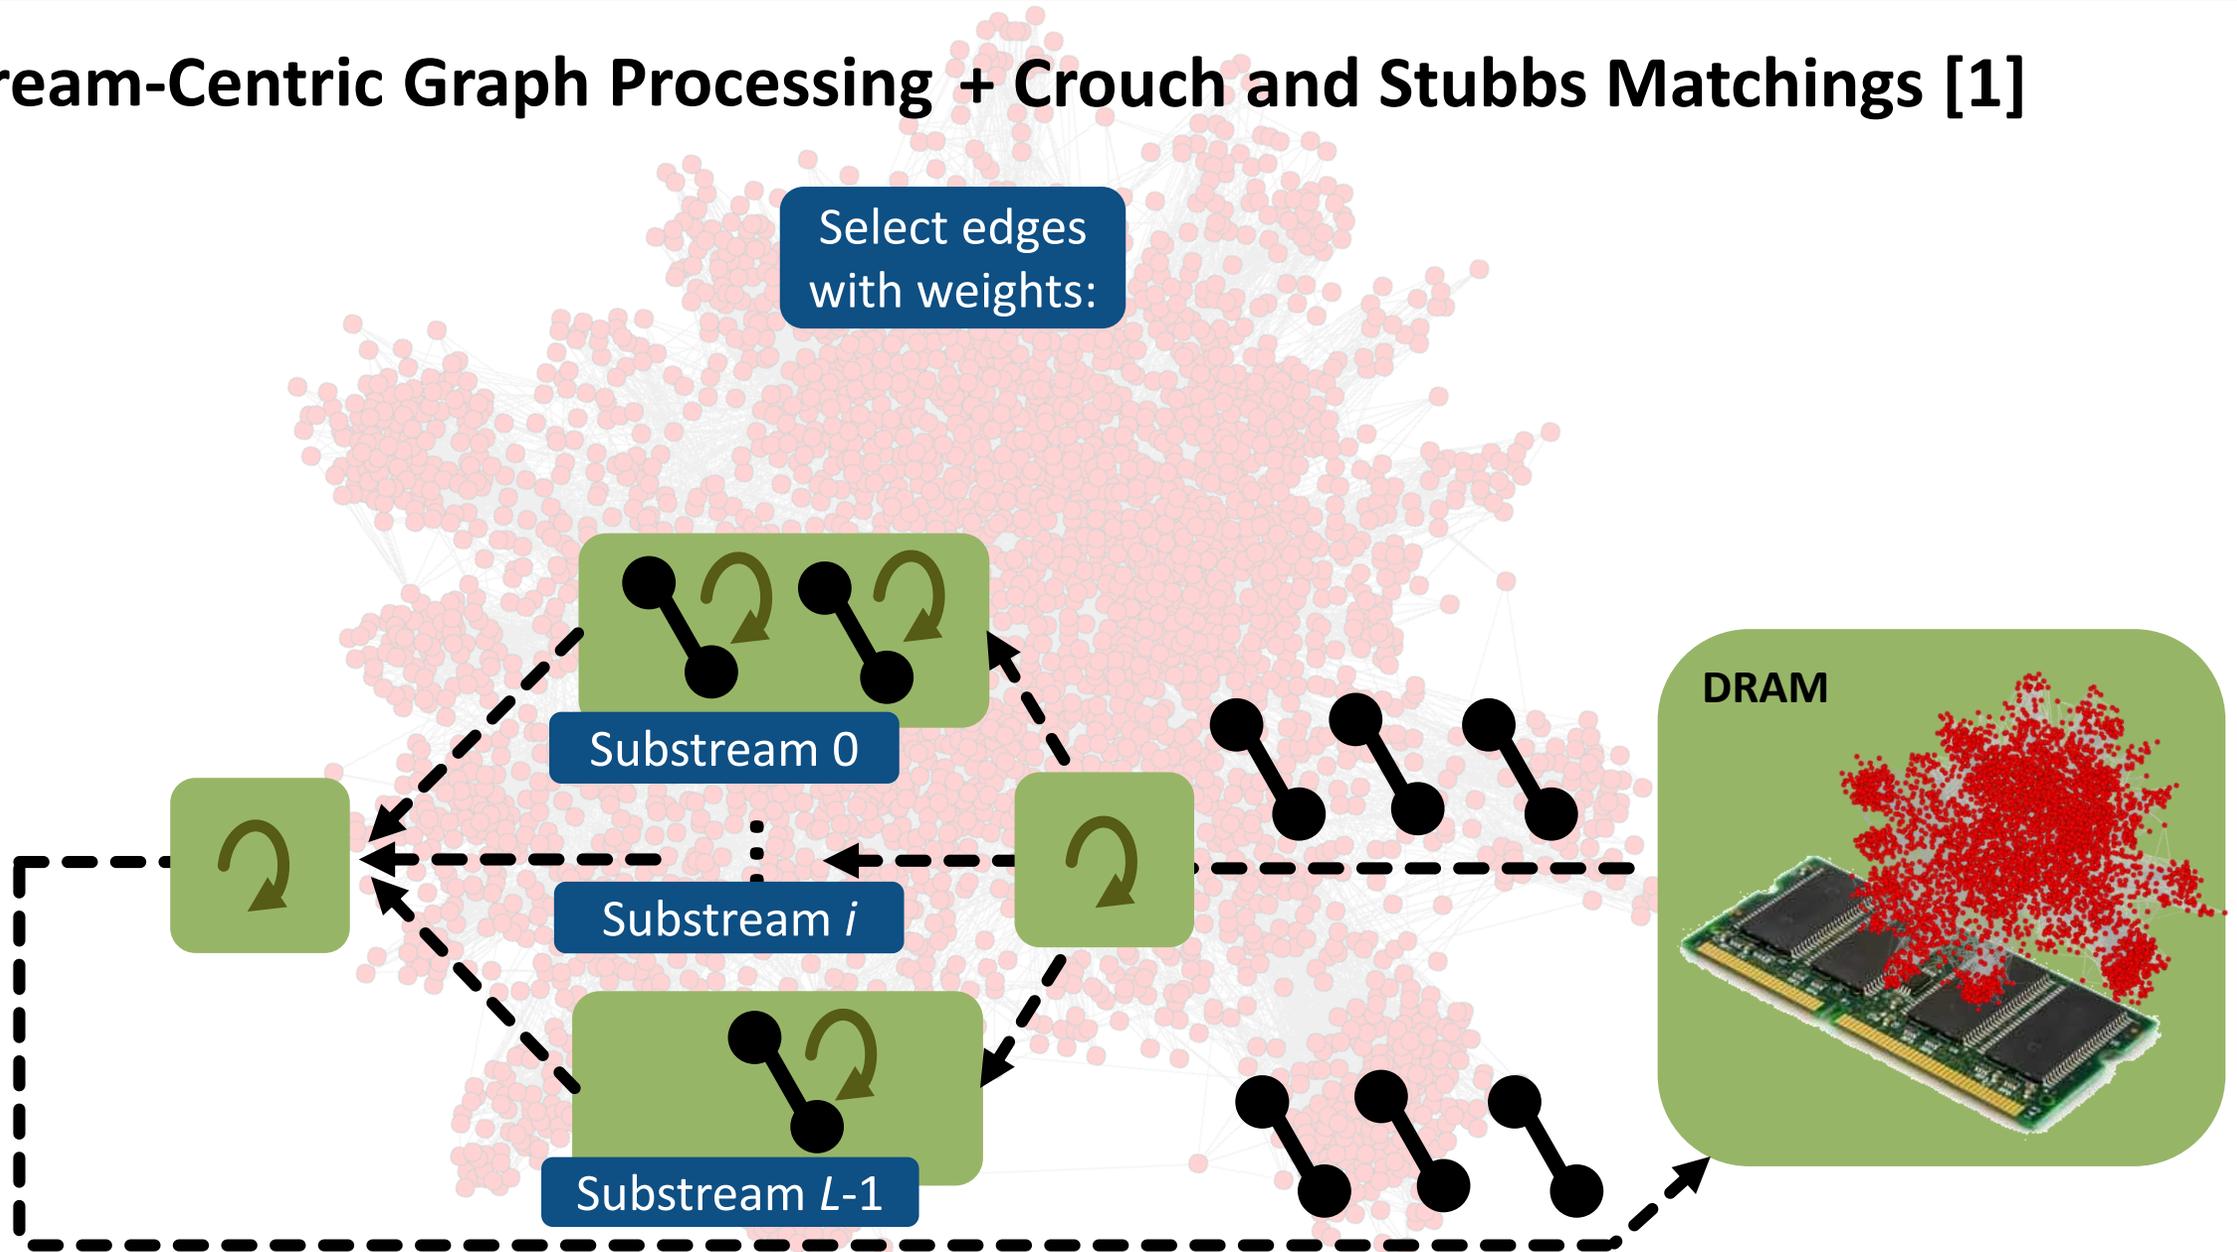
[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

Substream-Centric Graph Processing + Crouch and Stubbs Matchings [1]

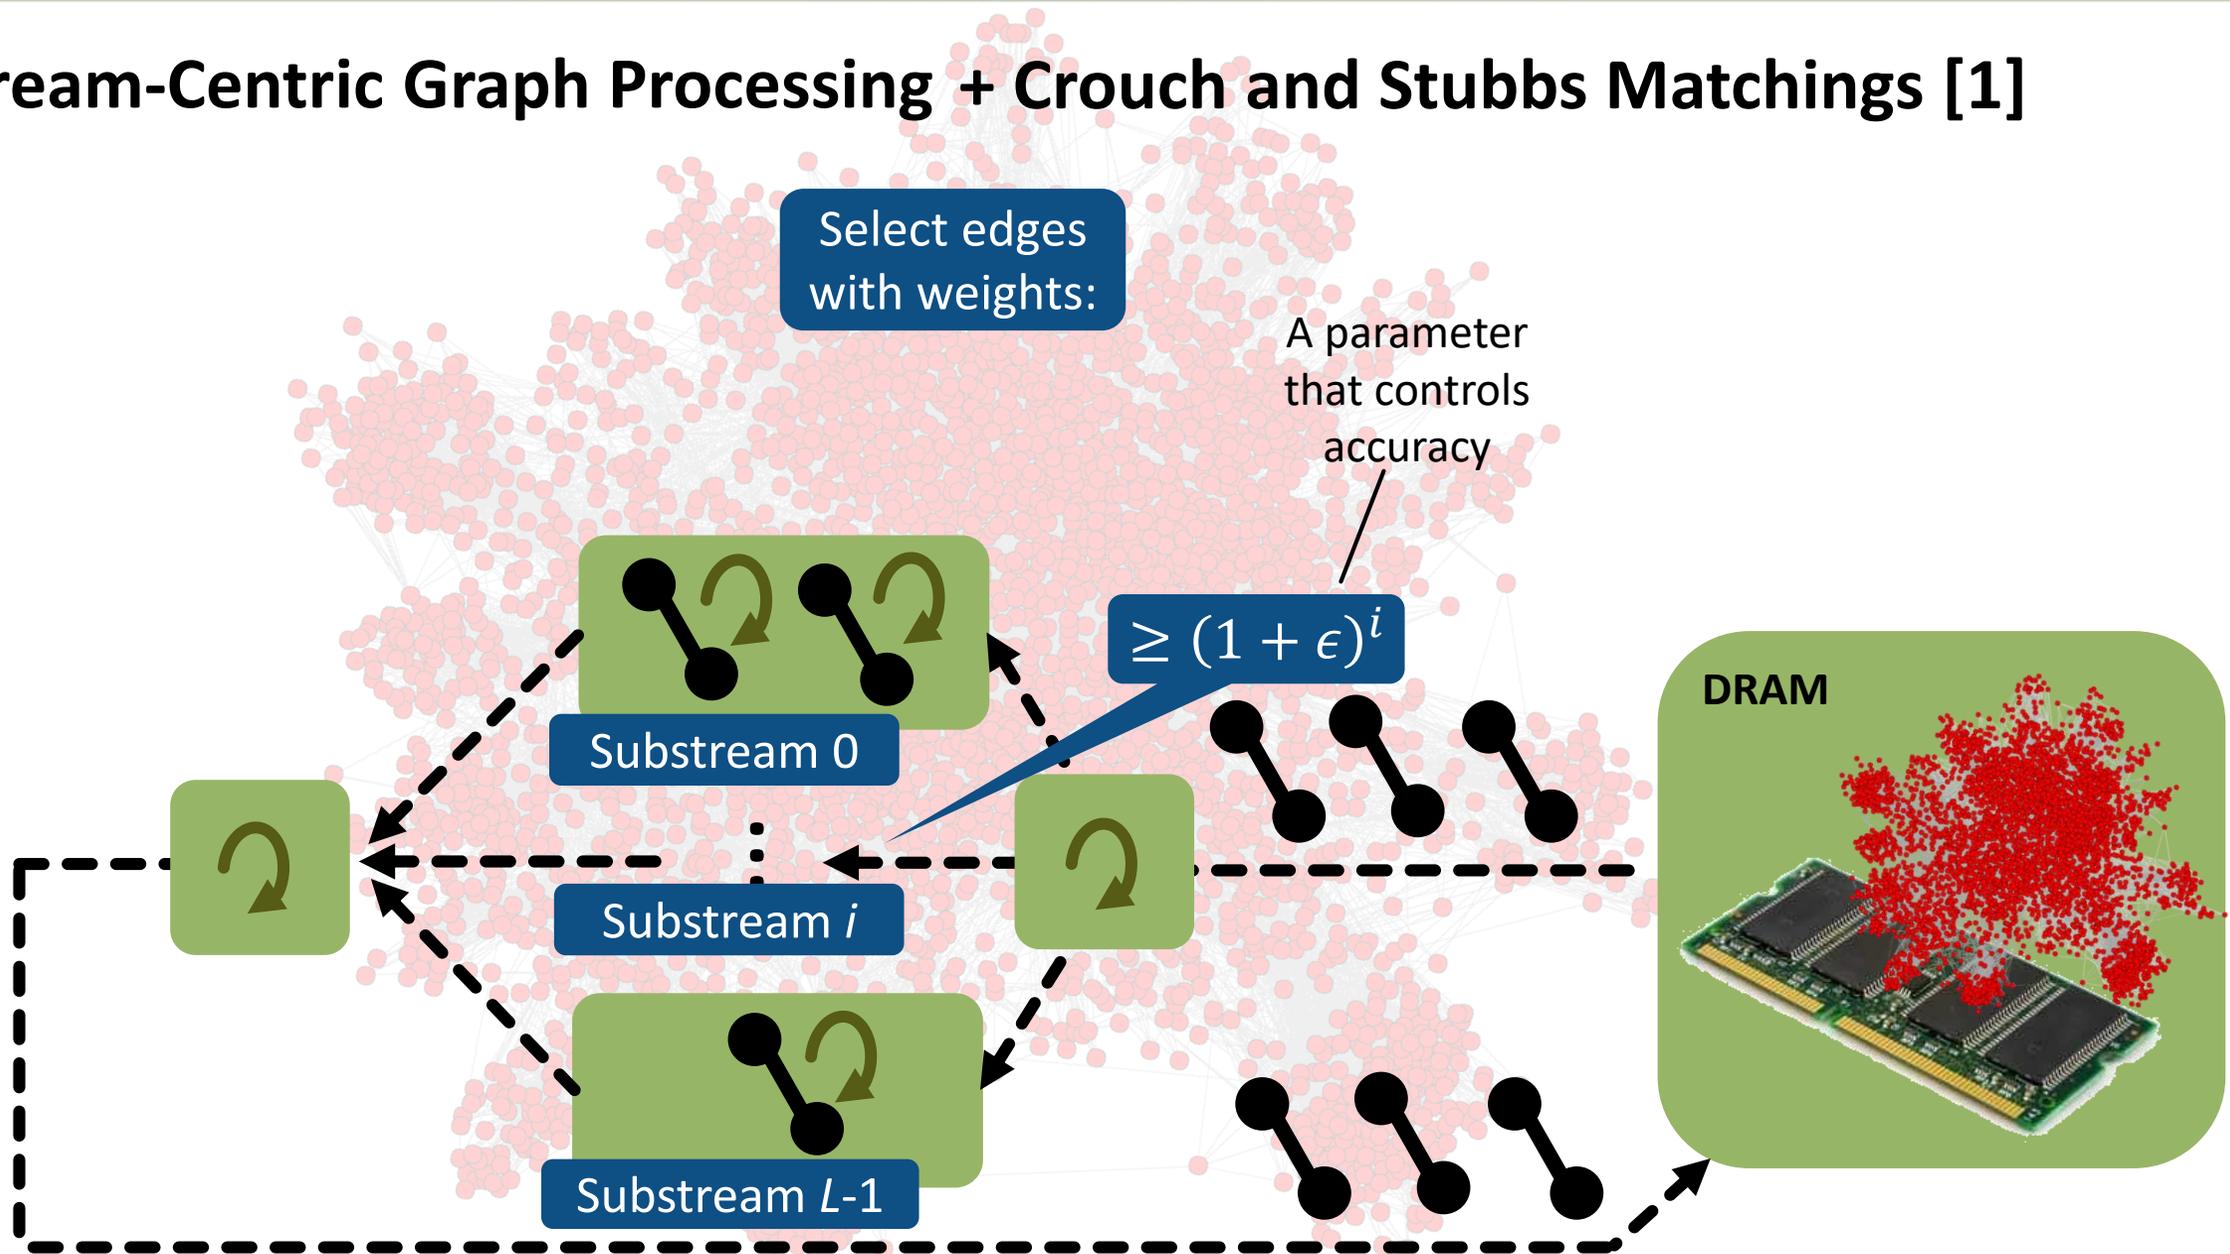


[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

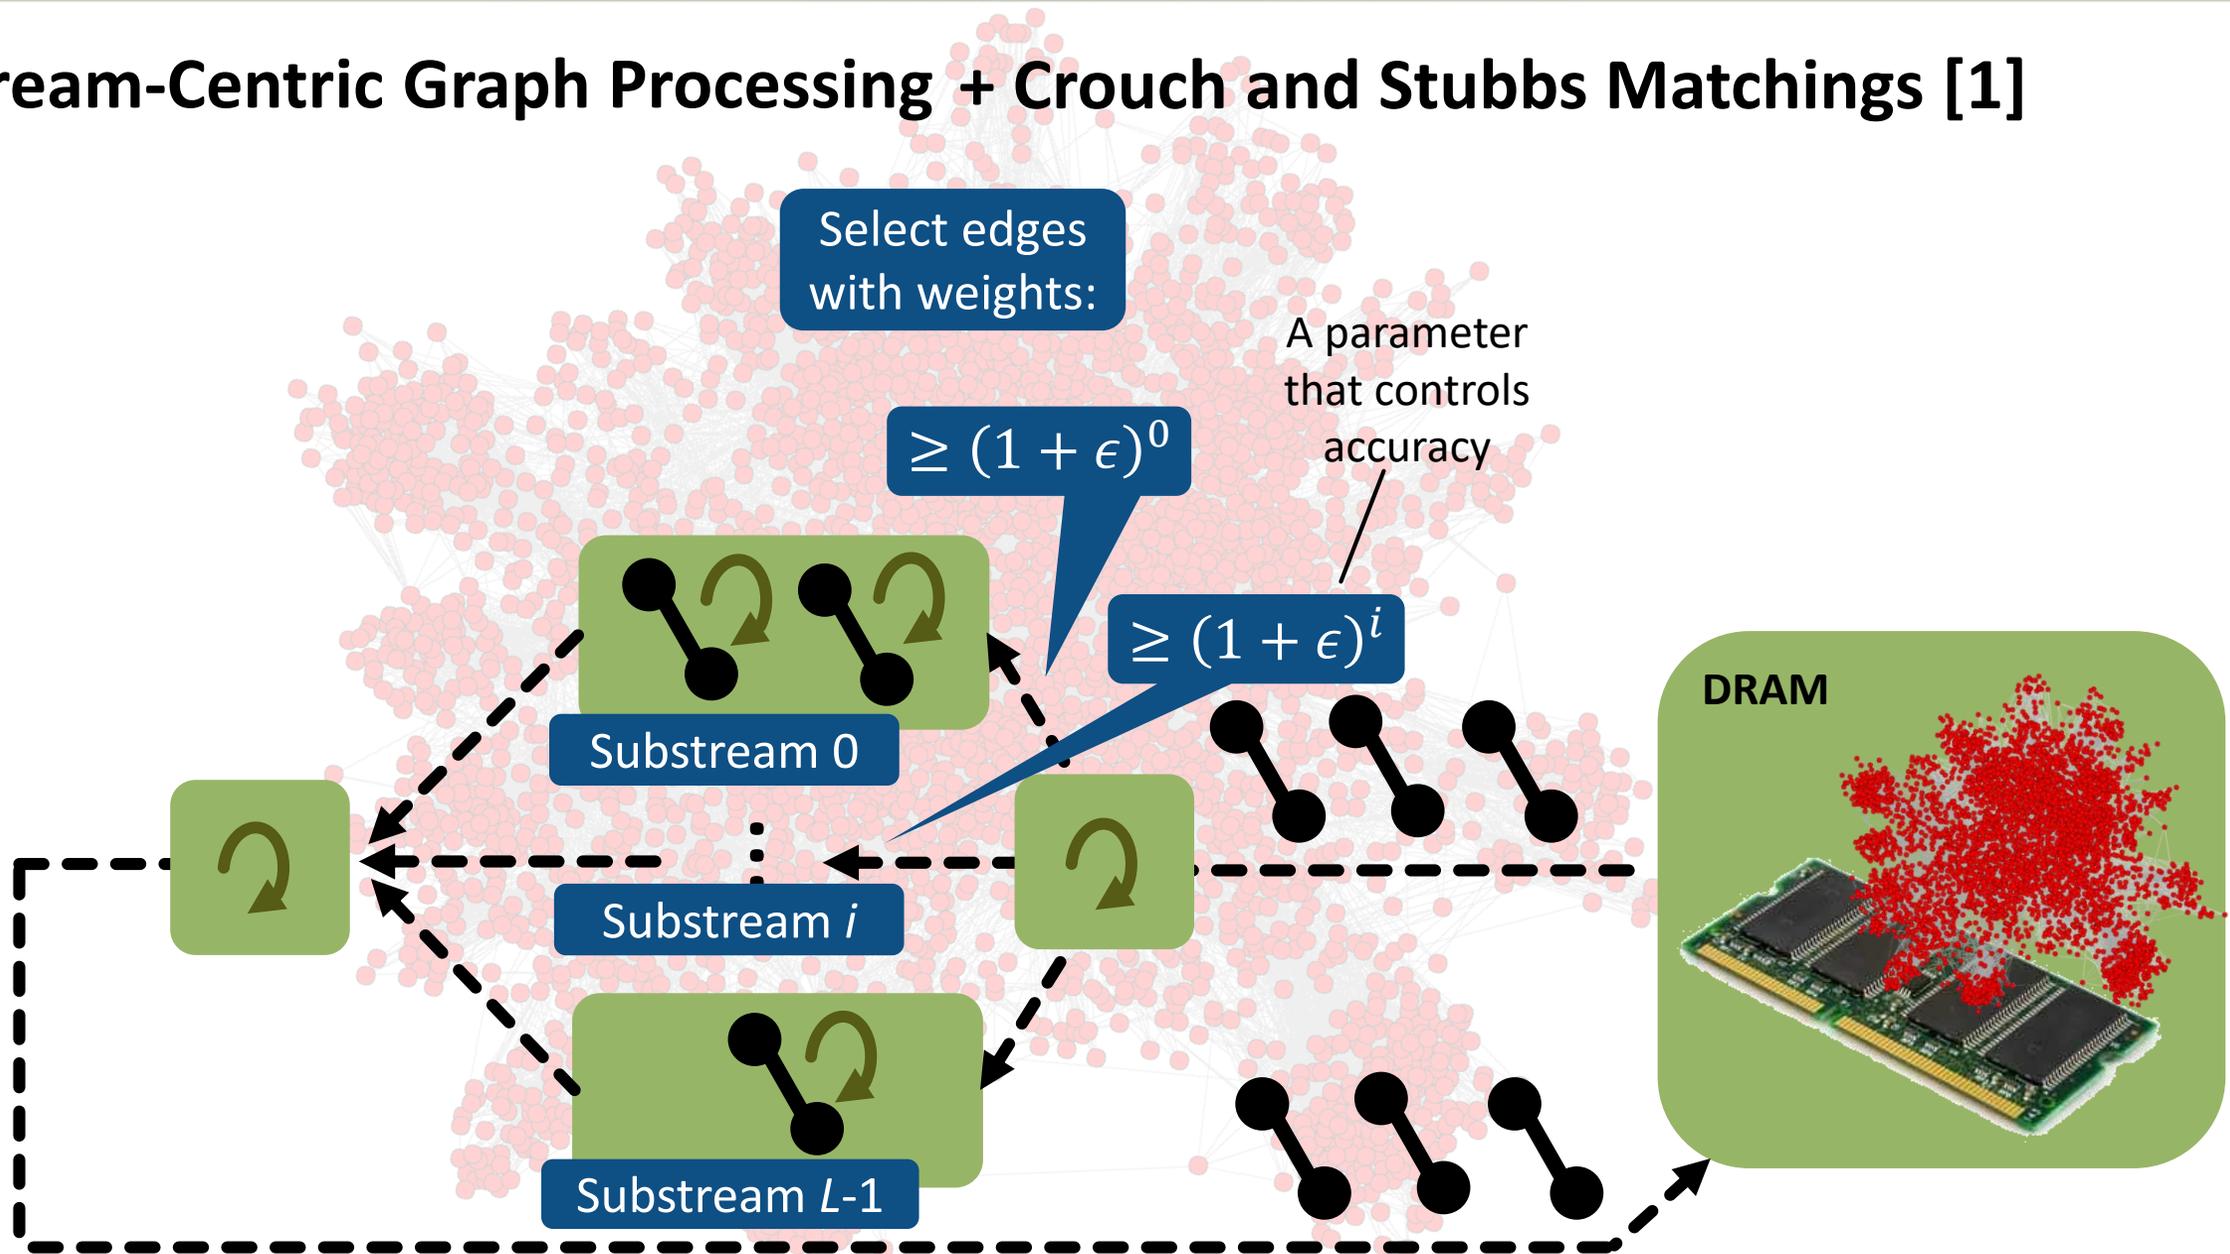
Substream-Centric Graph Processing + Crouch and Stubbs Matchings [1]



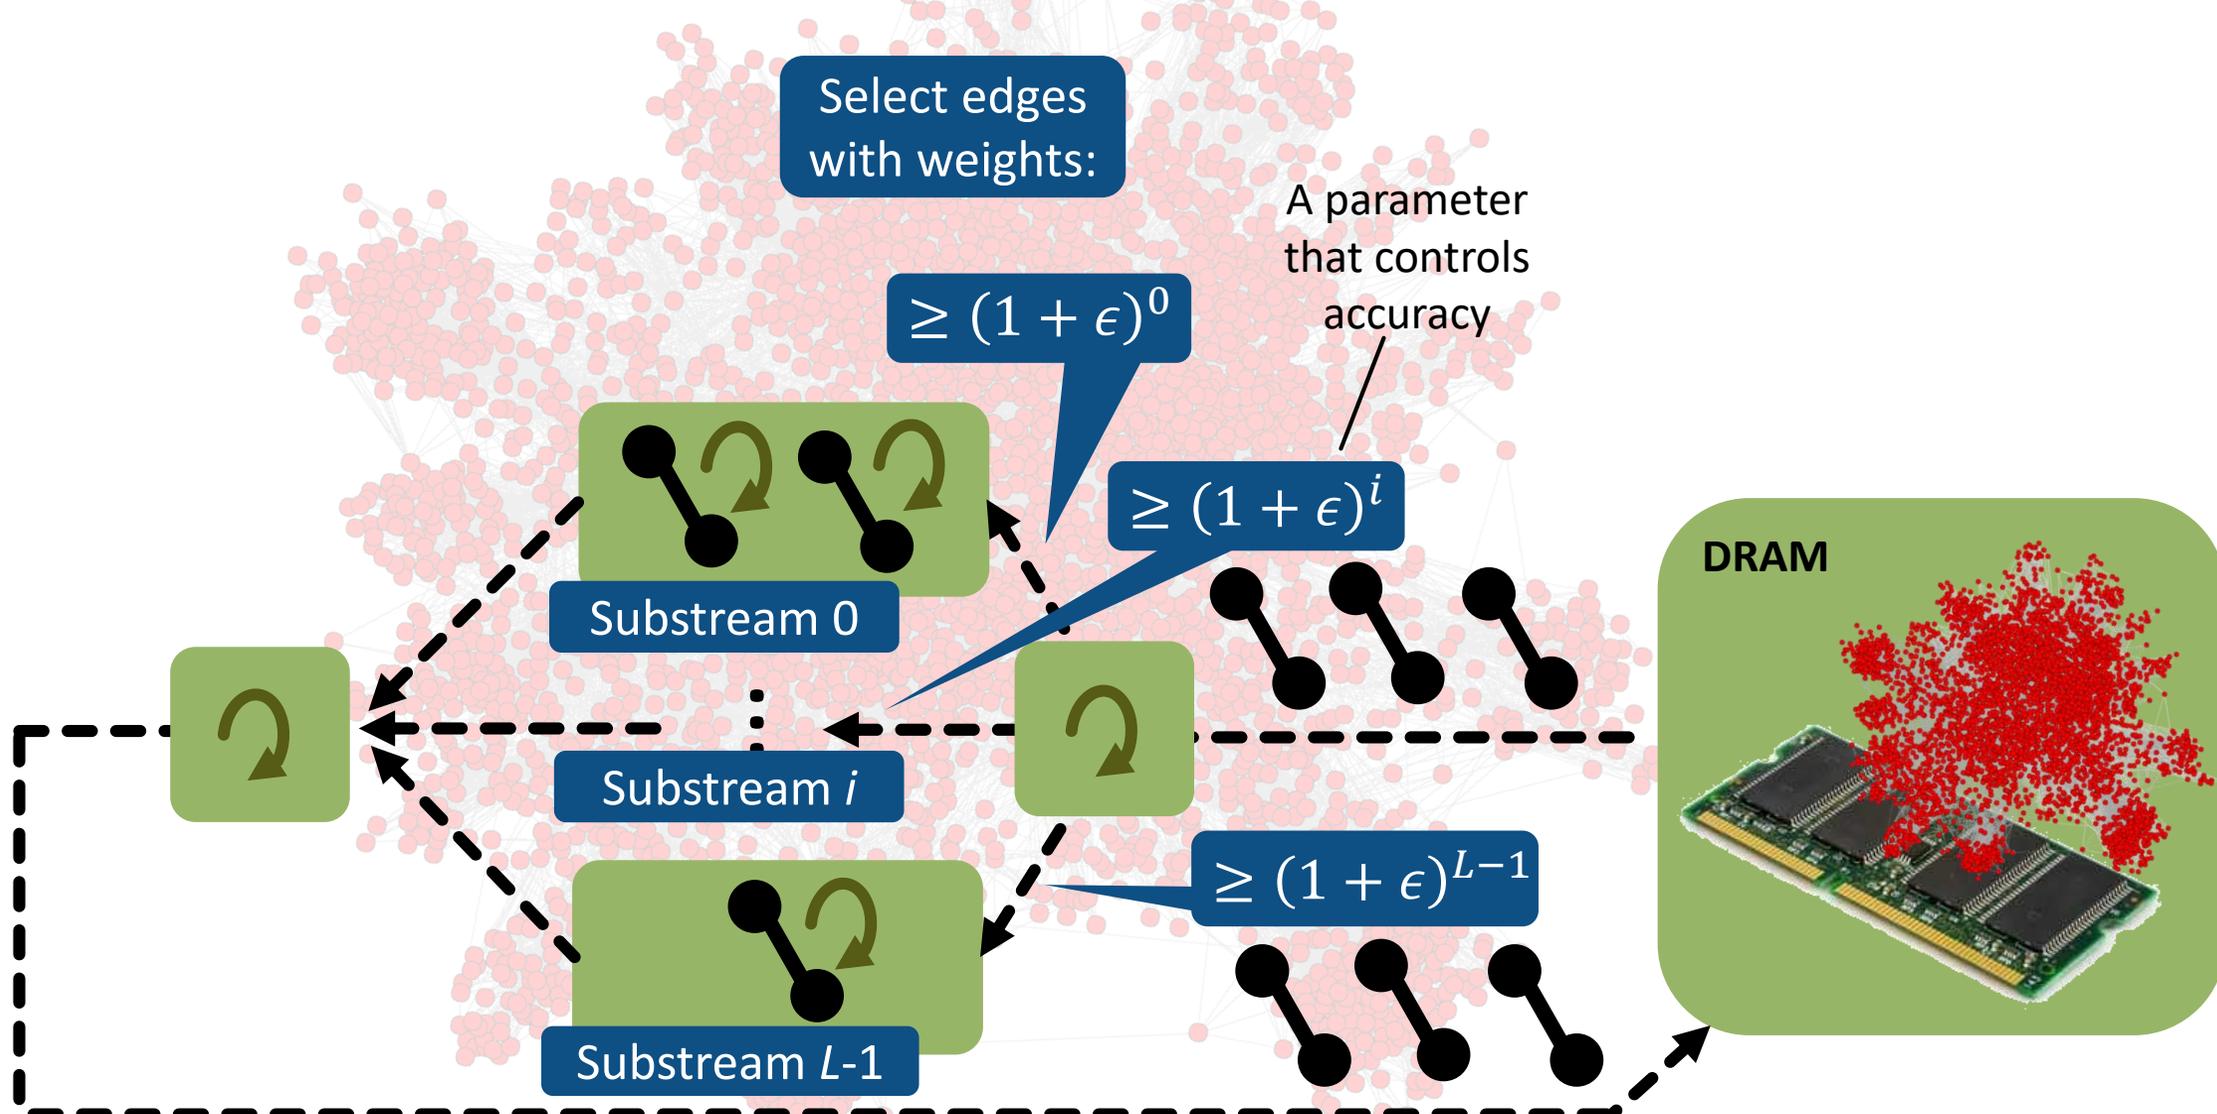
Substream-Centric Graph Processing + Crouch and Stubbs Matchings [1]



Substream-Centric Graph Processing + Crouch and Stubbs Matchings [1]

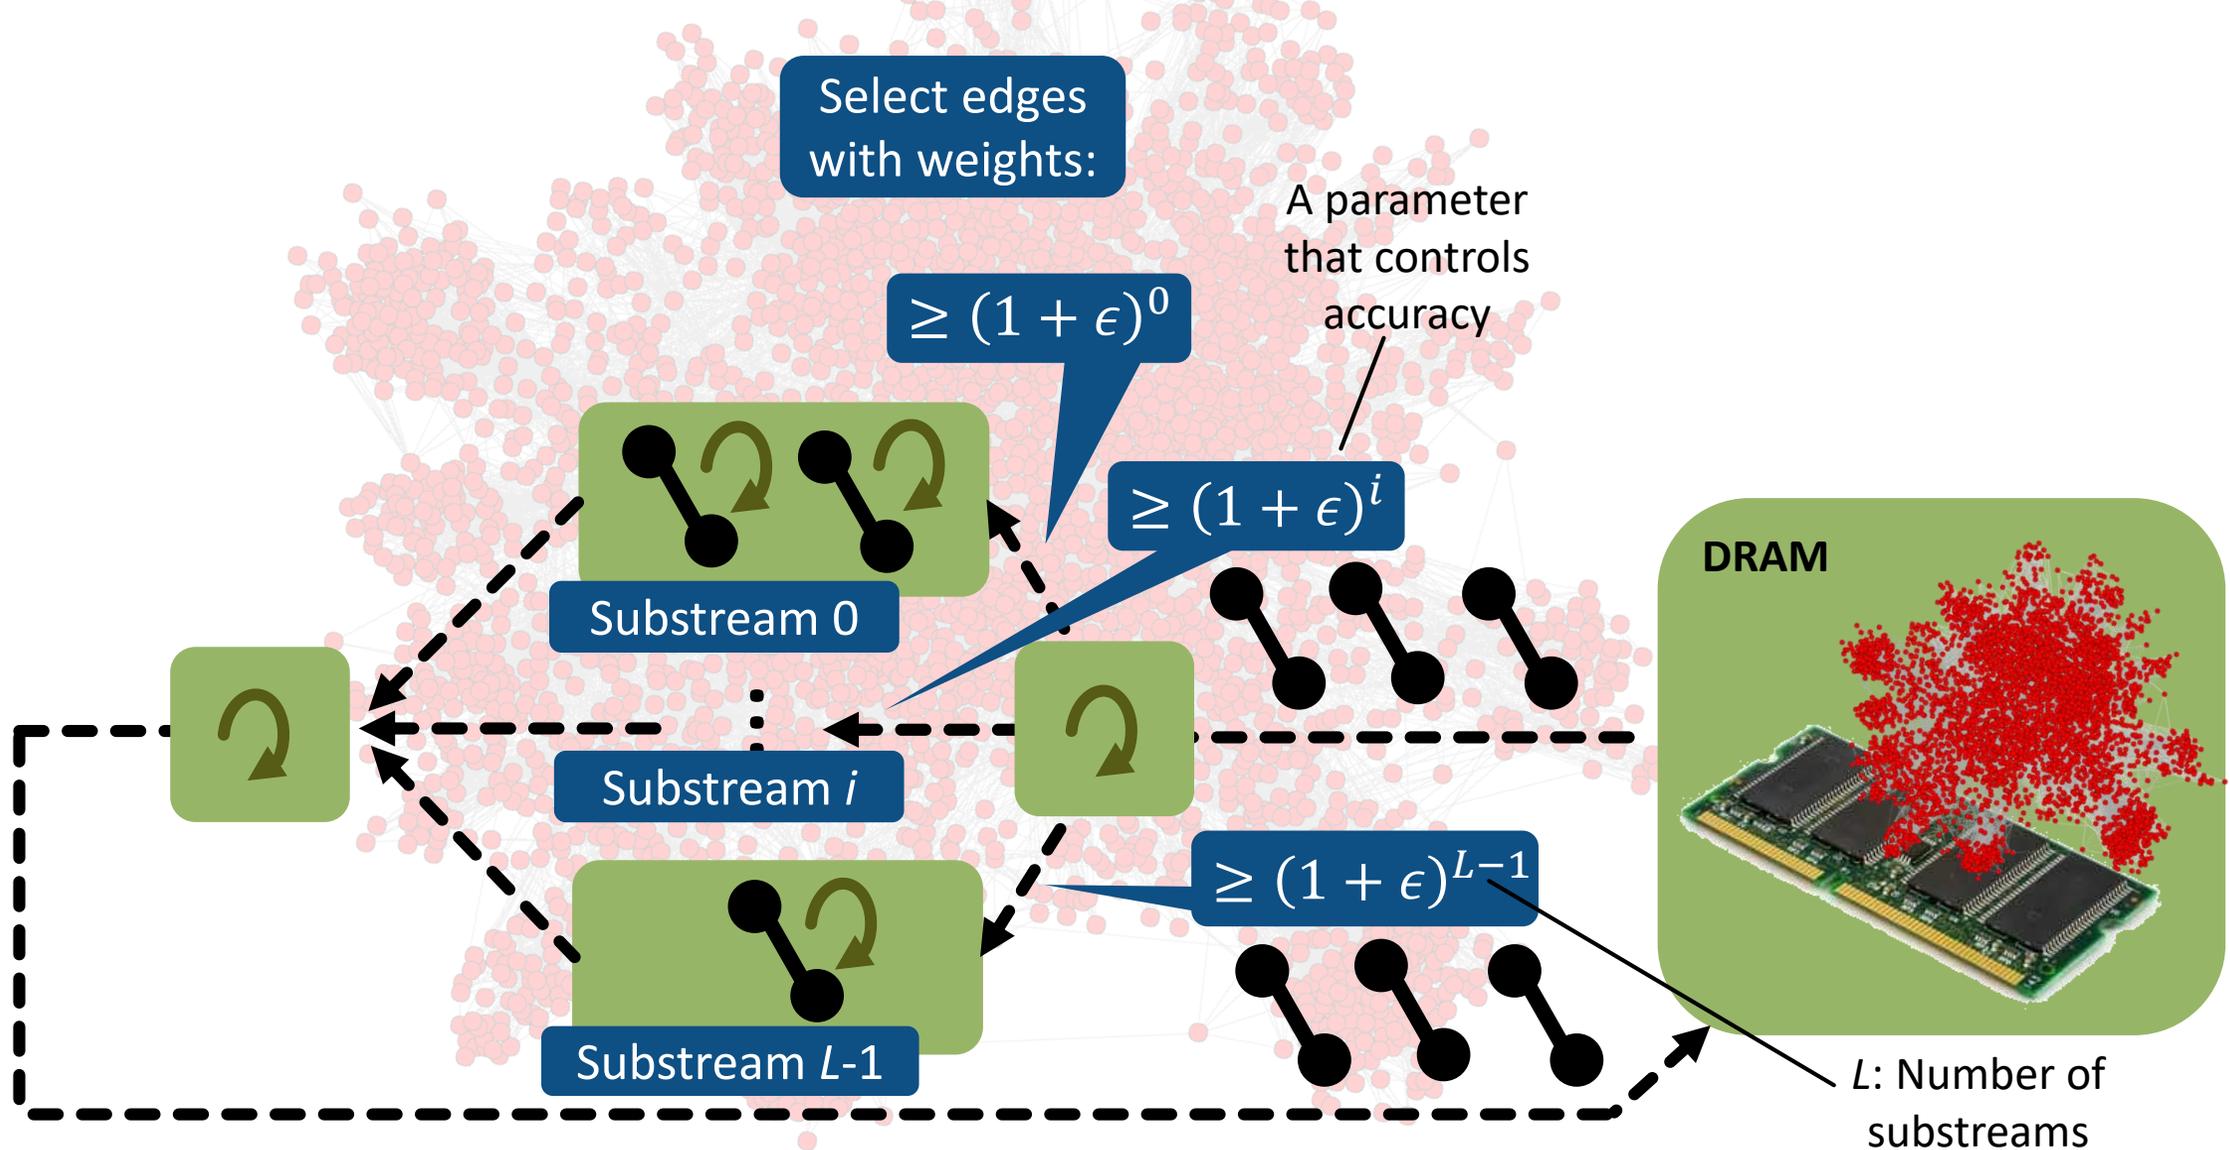


Substream-Centric Graph Processing + Crouch and Stubbs Matchings [1]



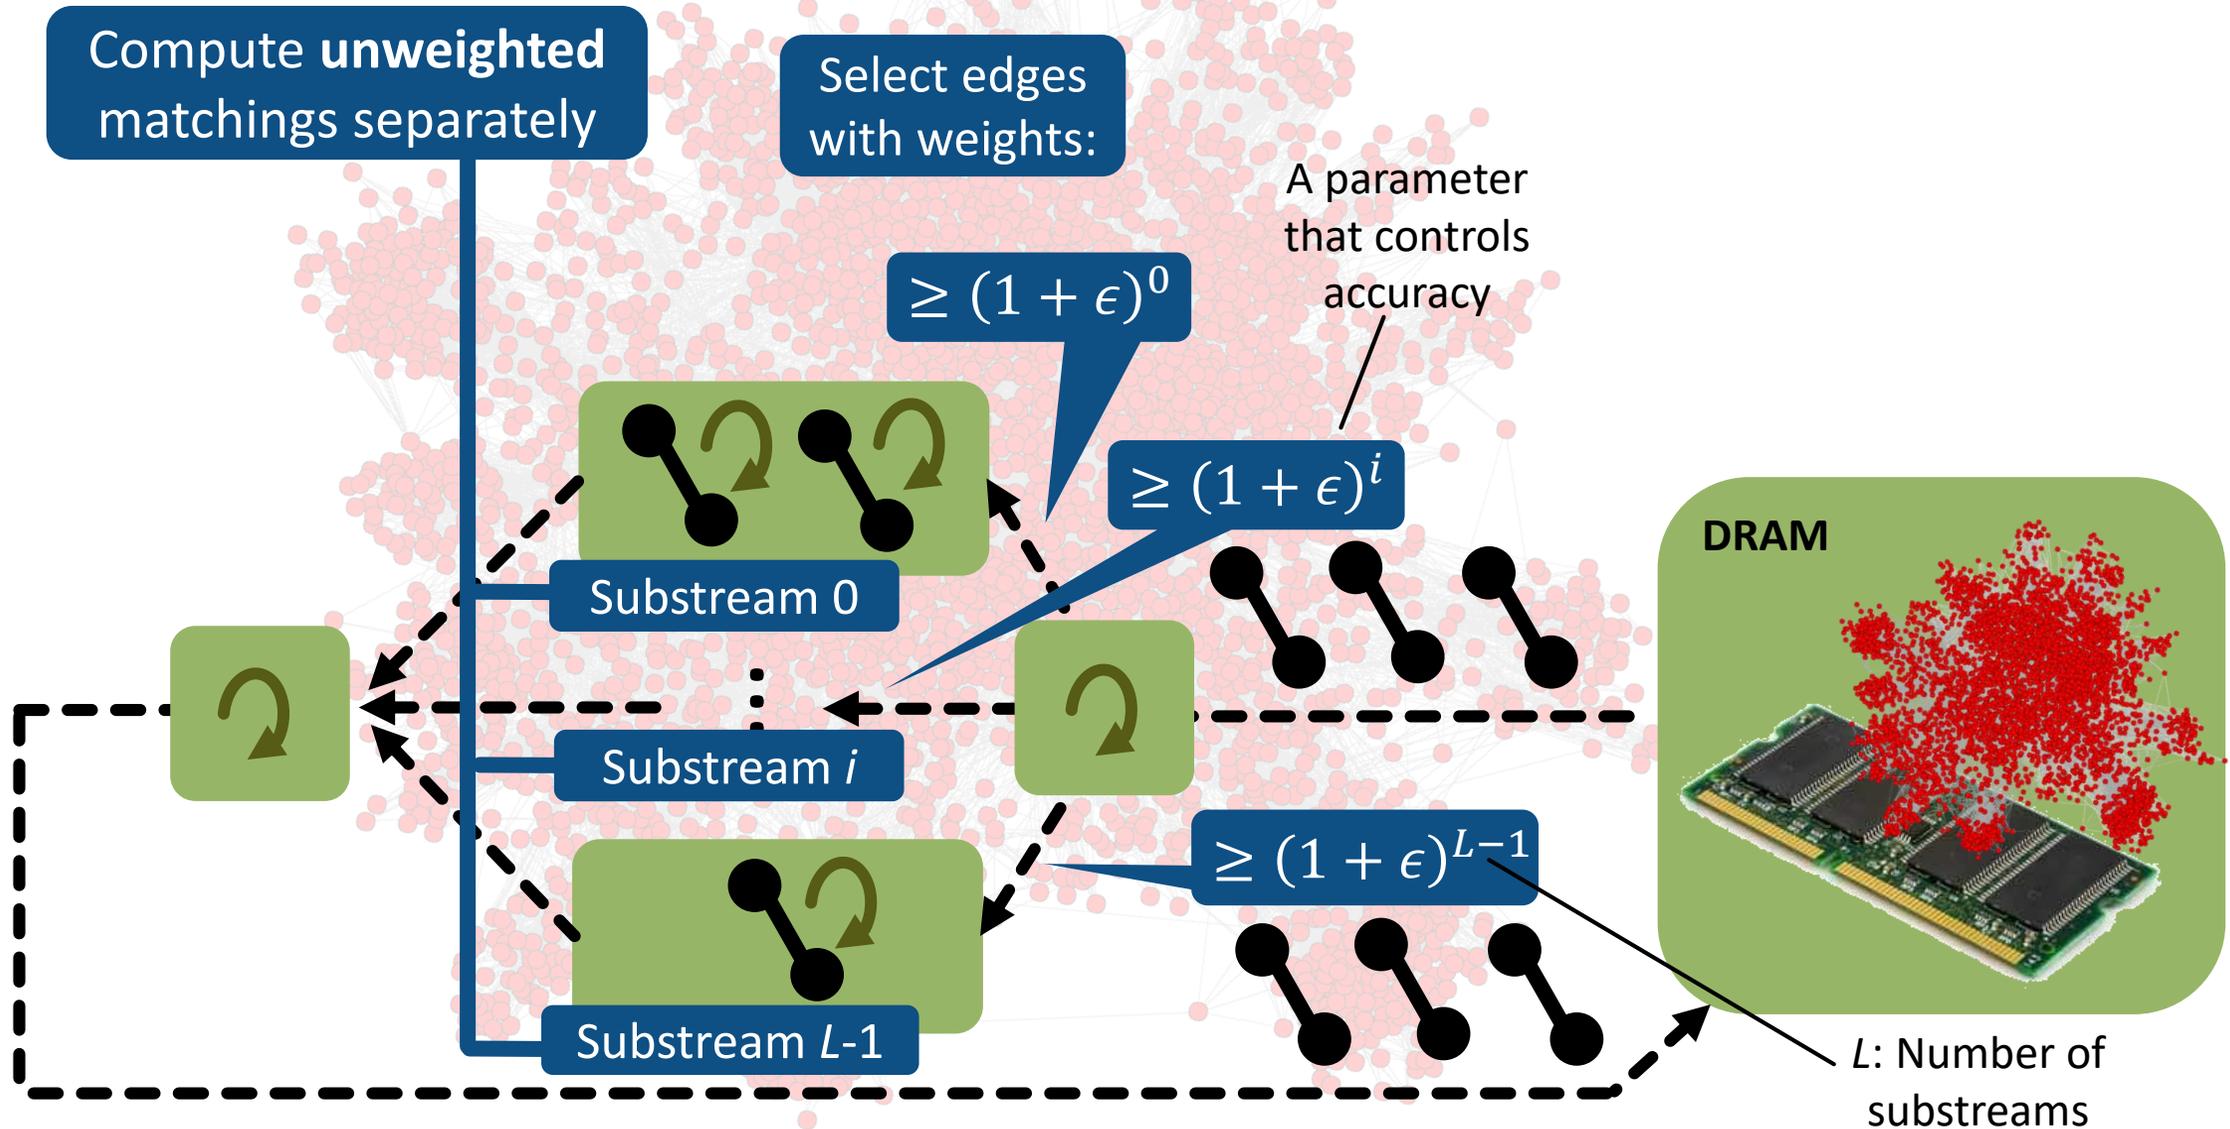
[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

Substream-Centric Graph Processing + Crouch and Stubbs Matchings [1]



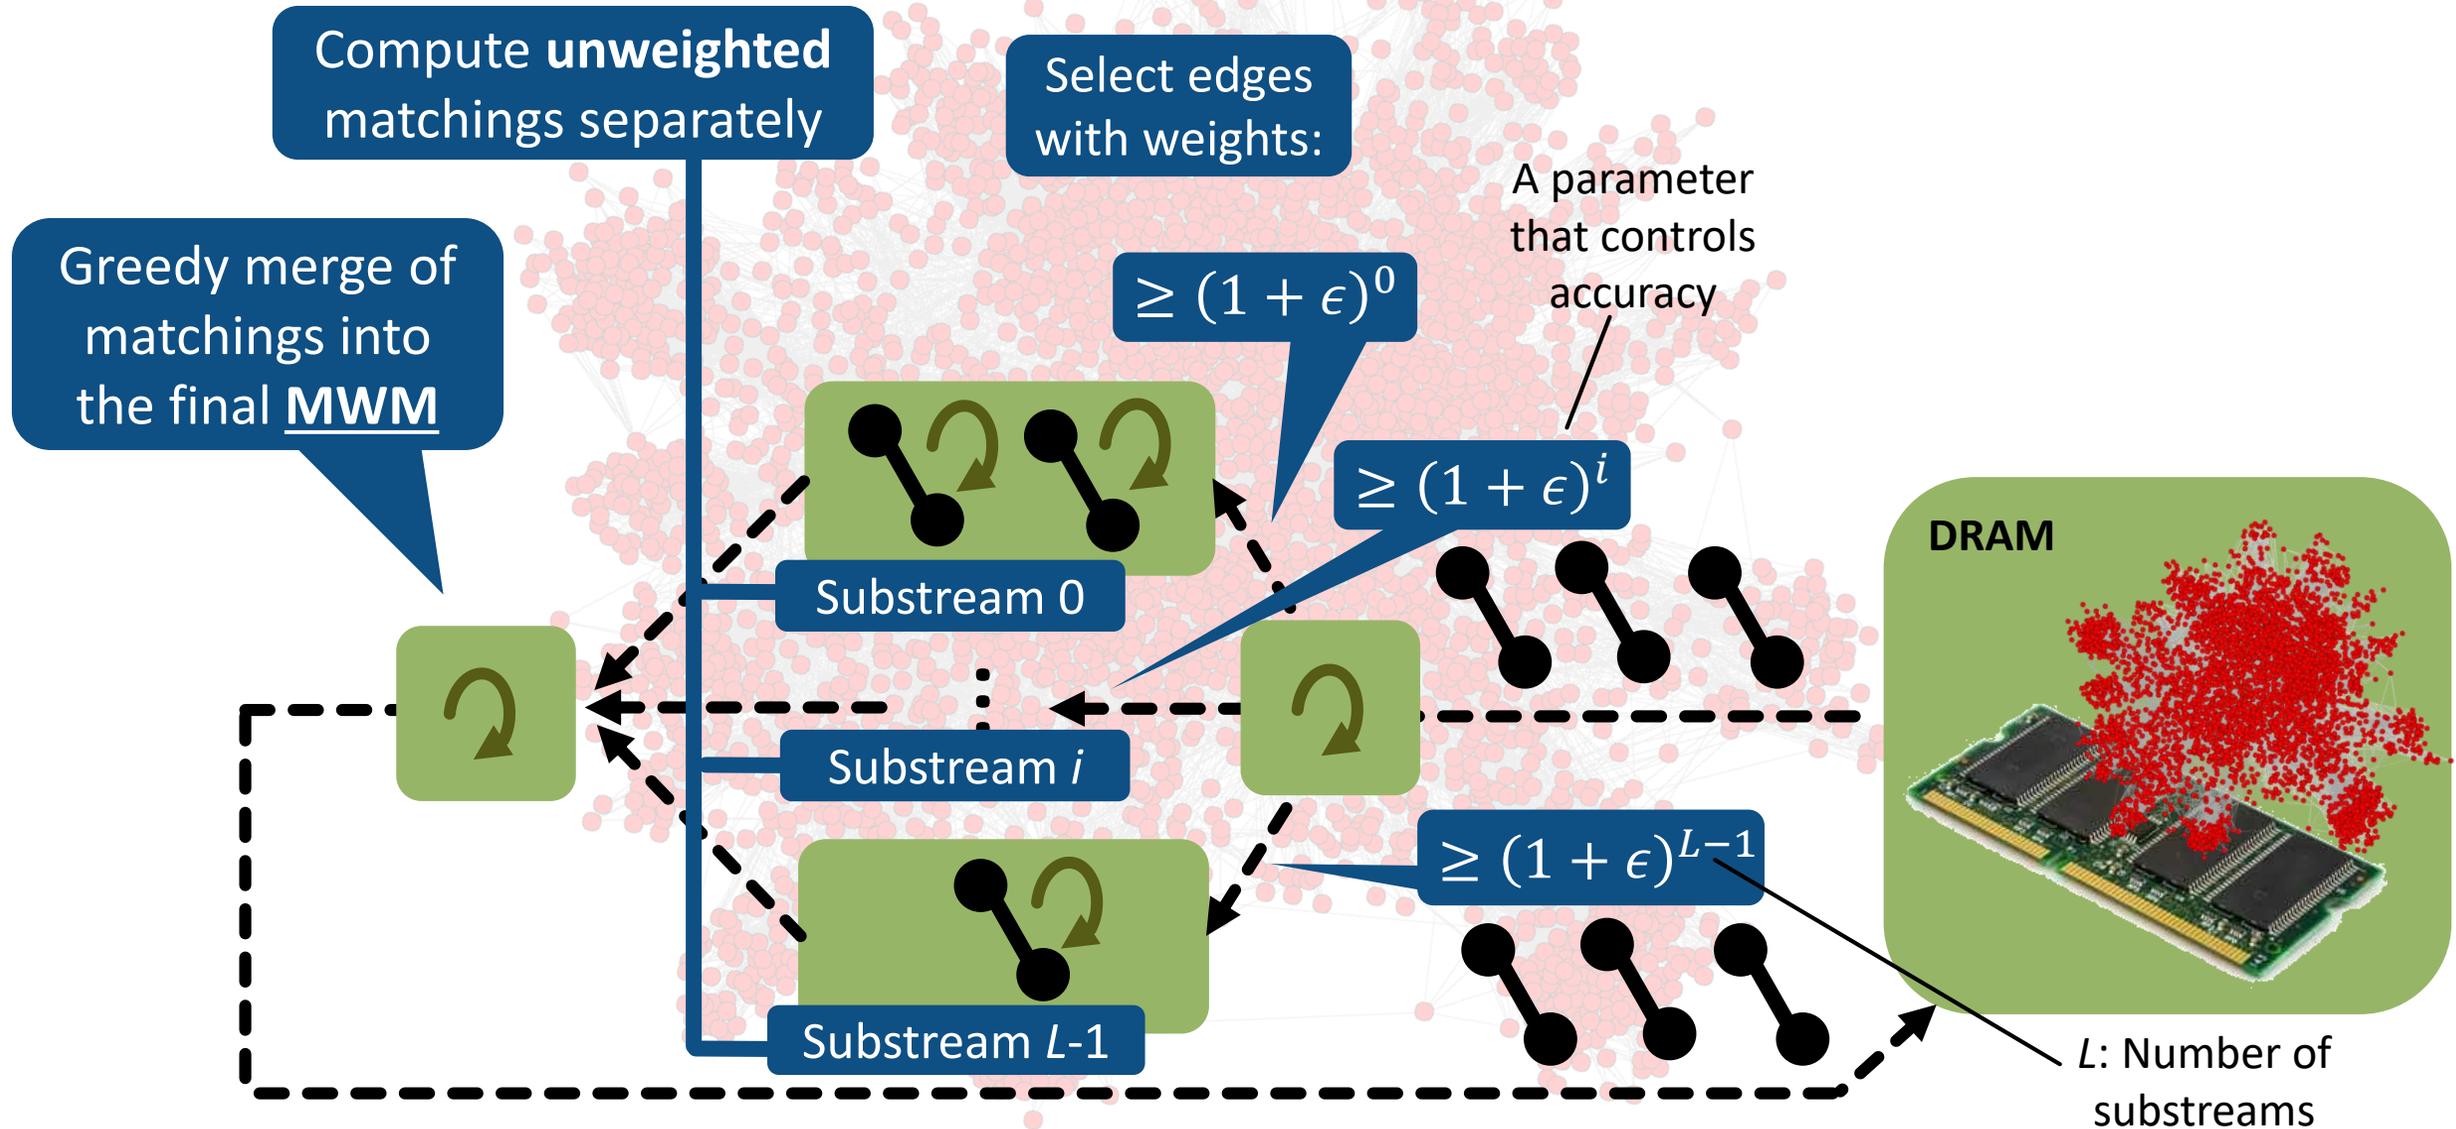
[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

Substream-Centric Graph Processing + Crouch and Stubbs Matchings [1]



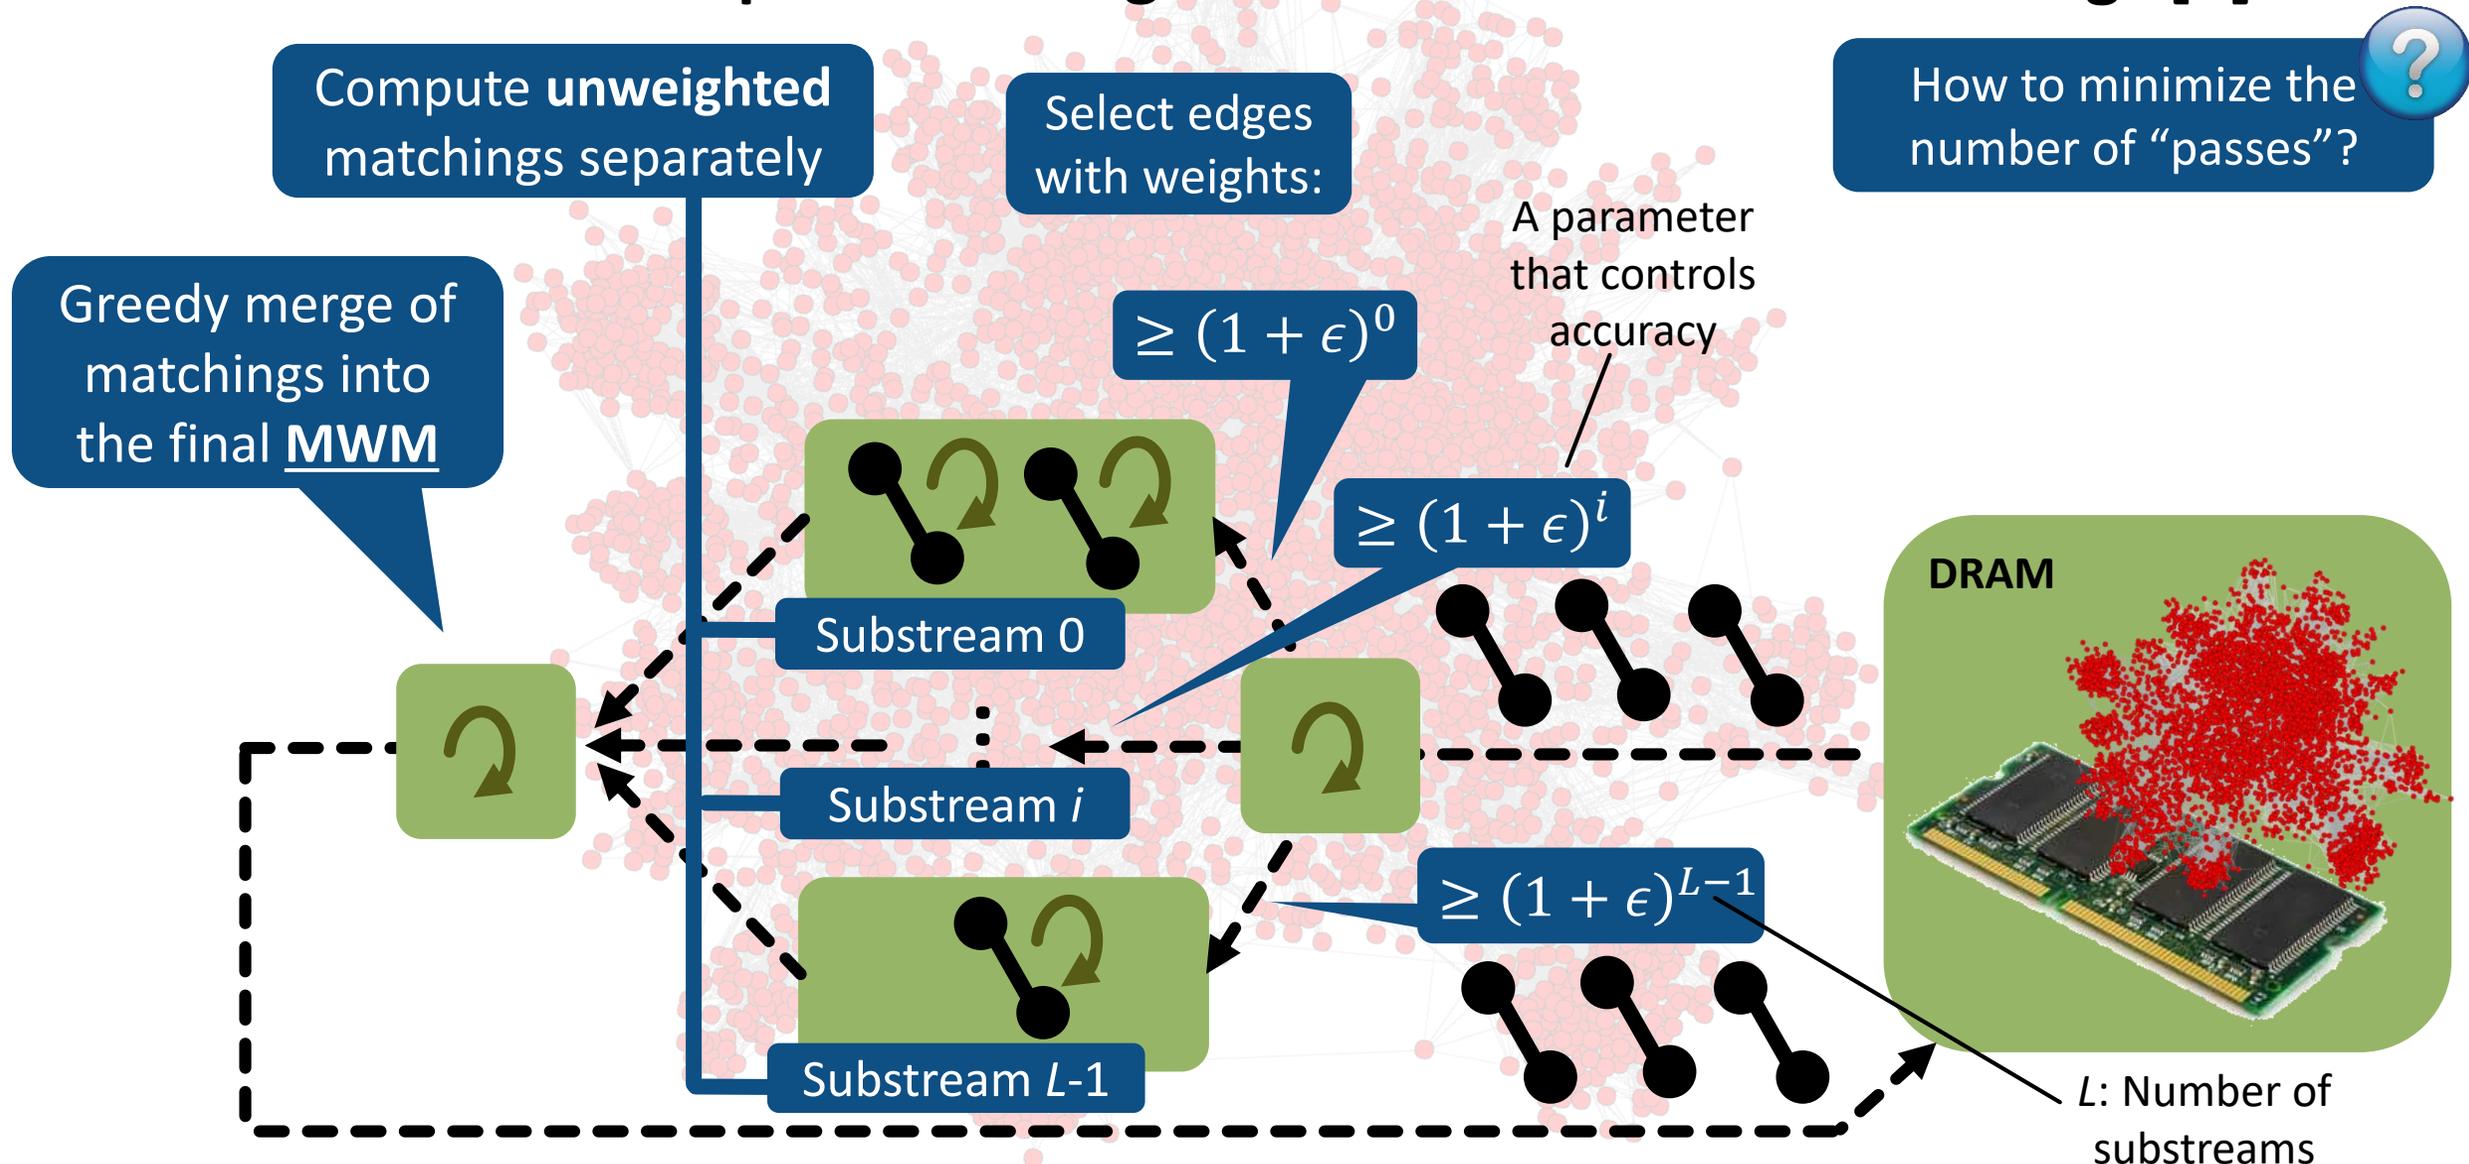
[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

Substream-Centric Graph Processing + Crouch and Stubbs Matchings [1]



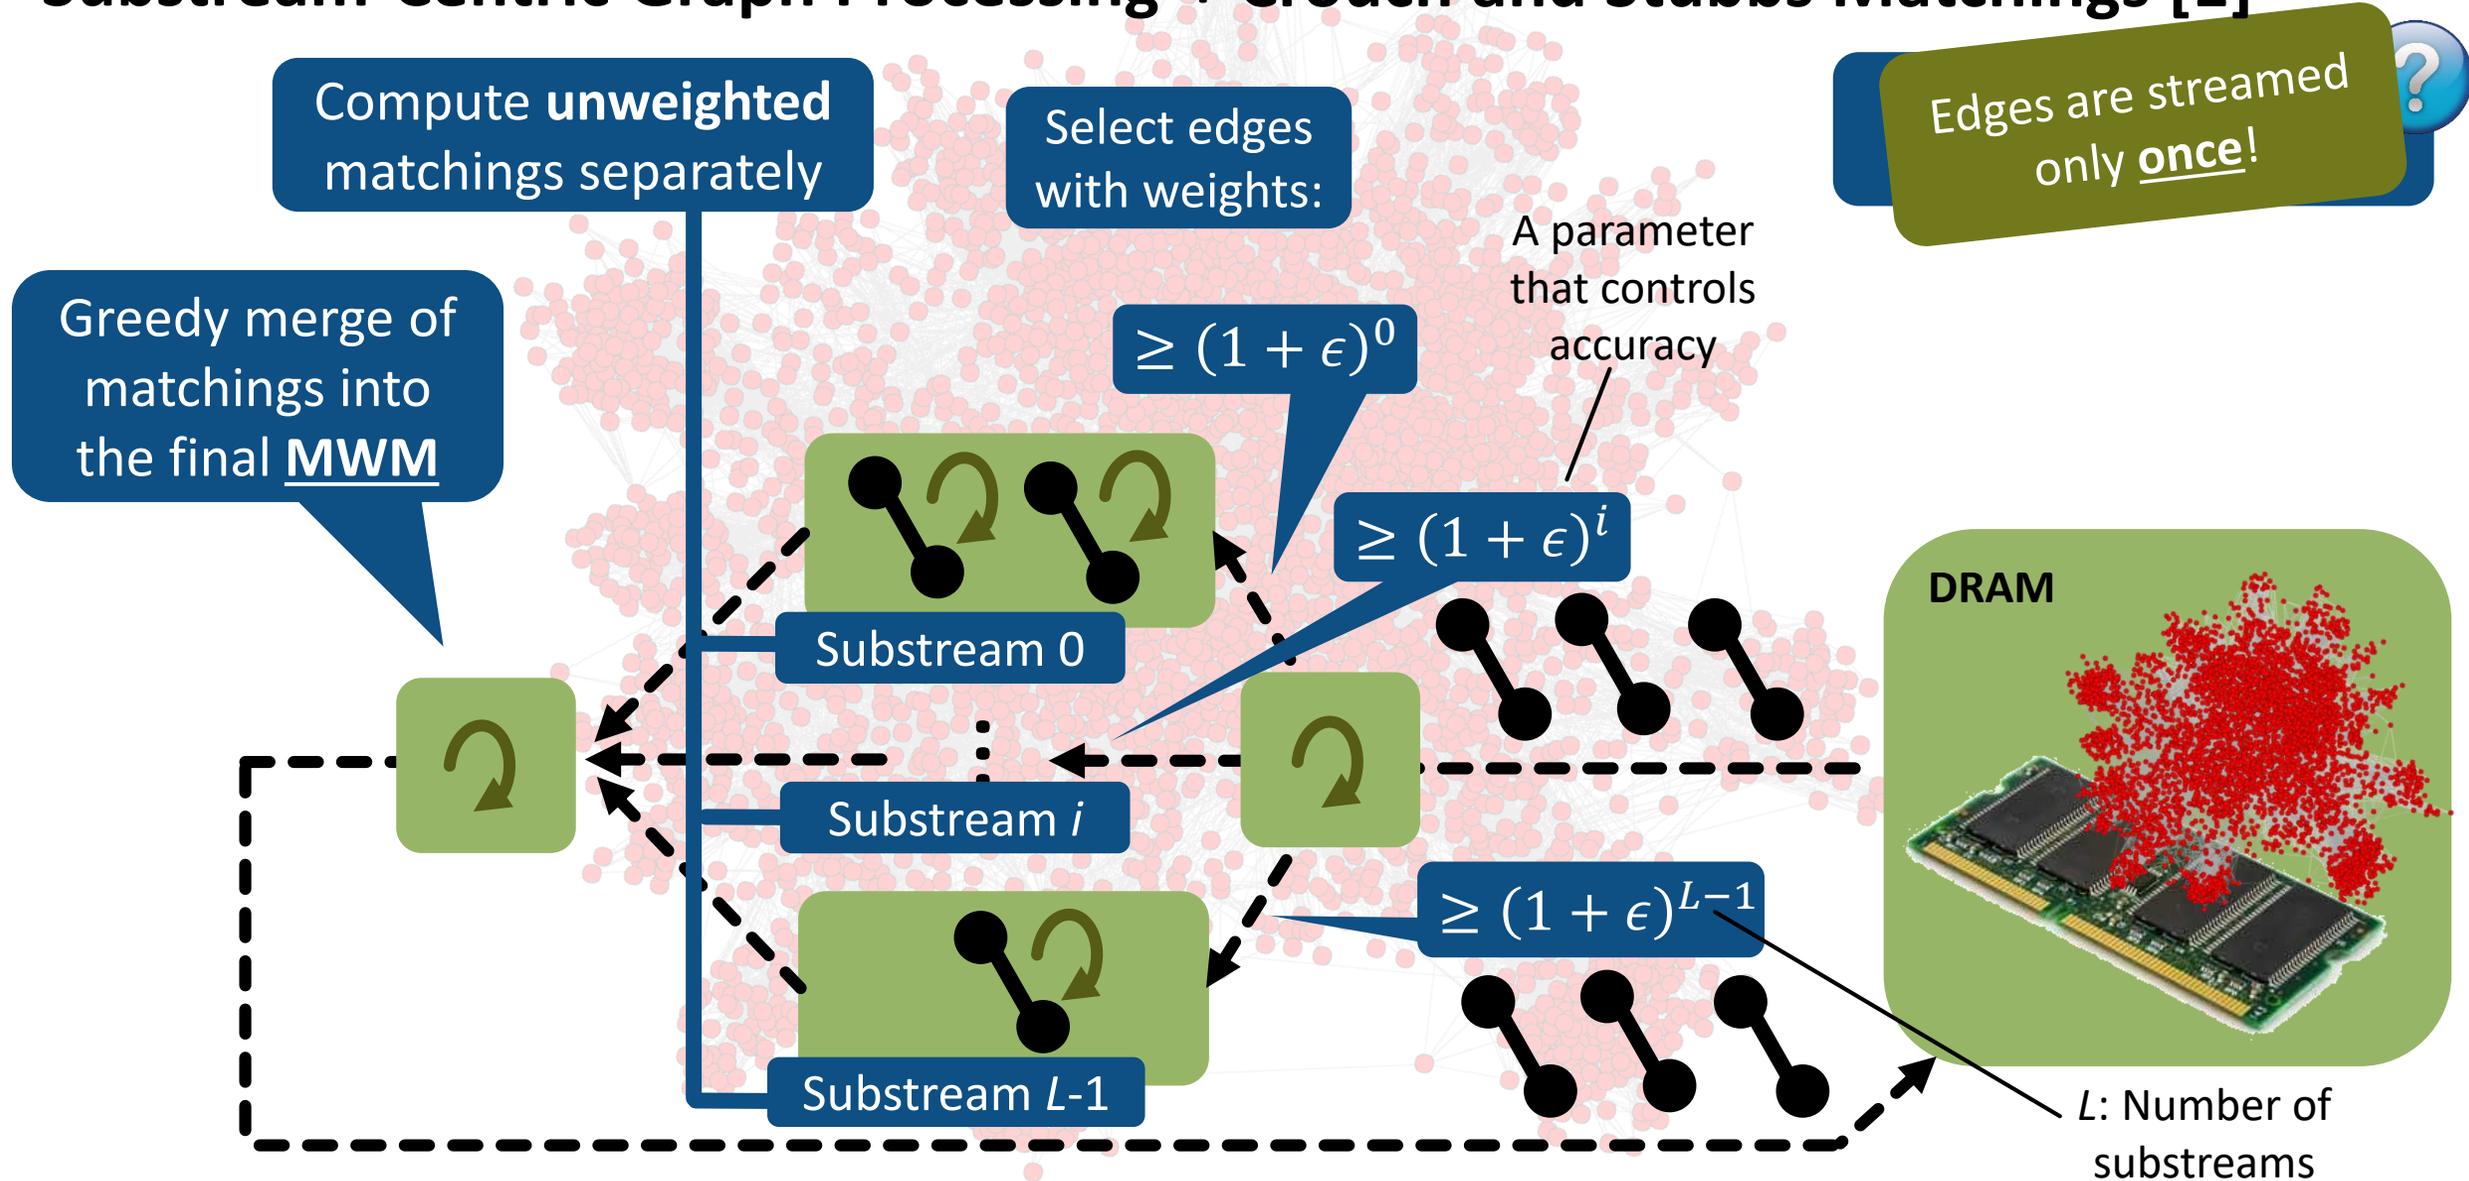
[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

Substream-Centric Graph Processing + Crouch and Stubbs Matchings [1]



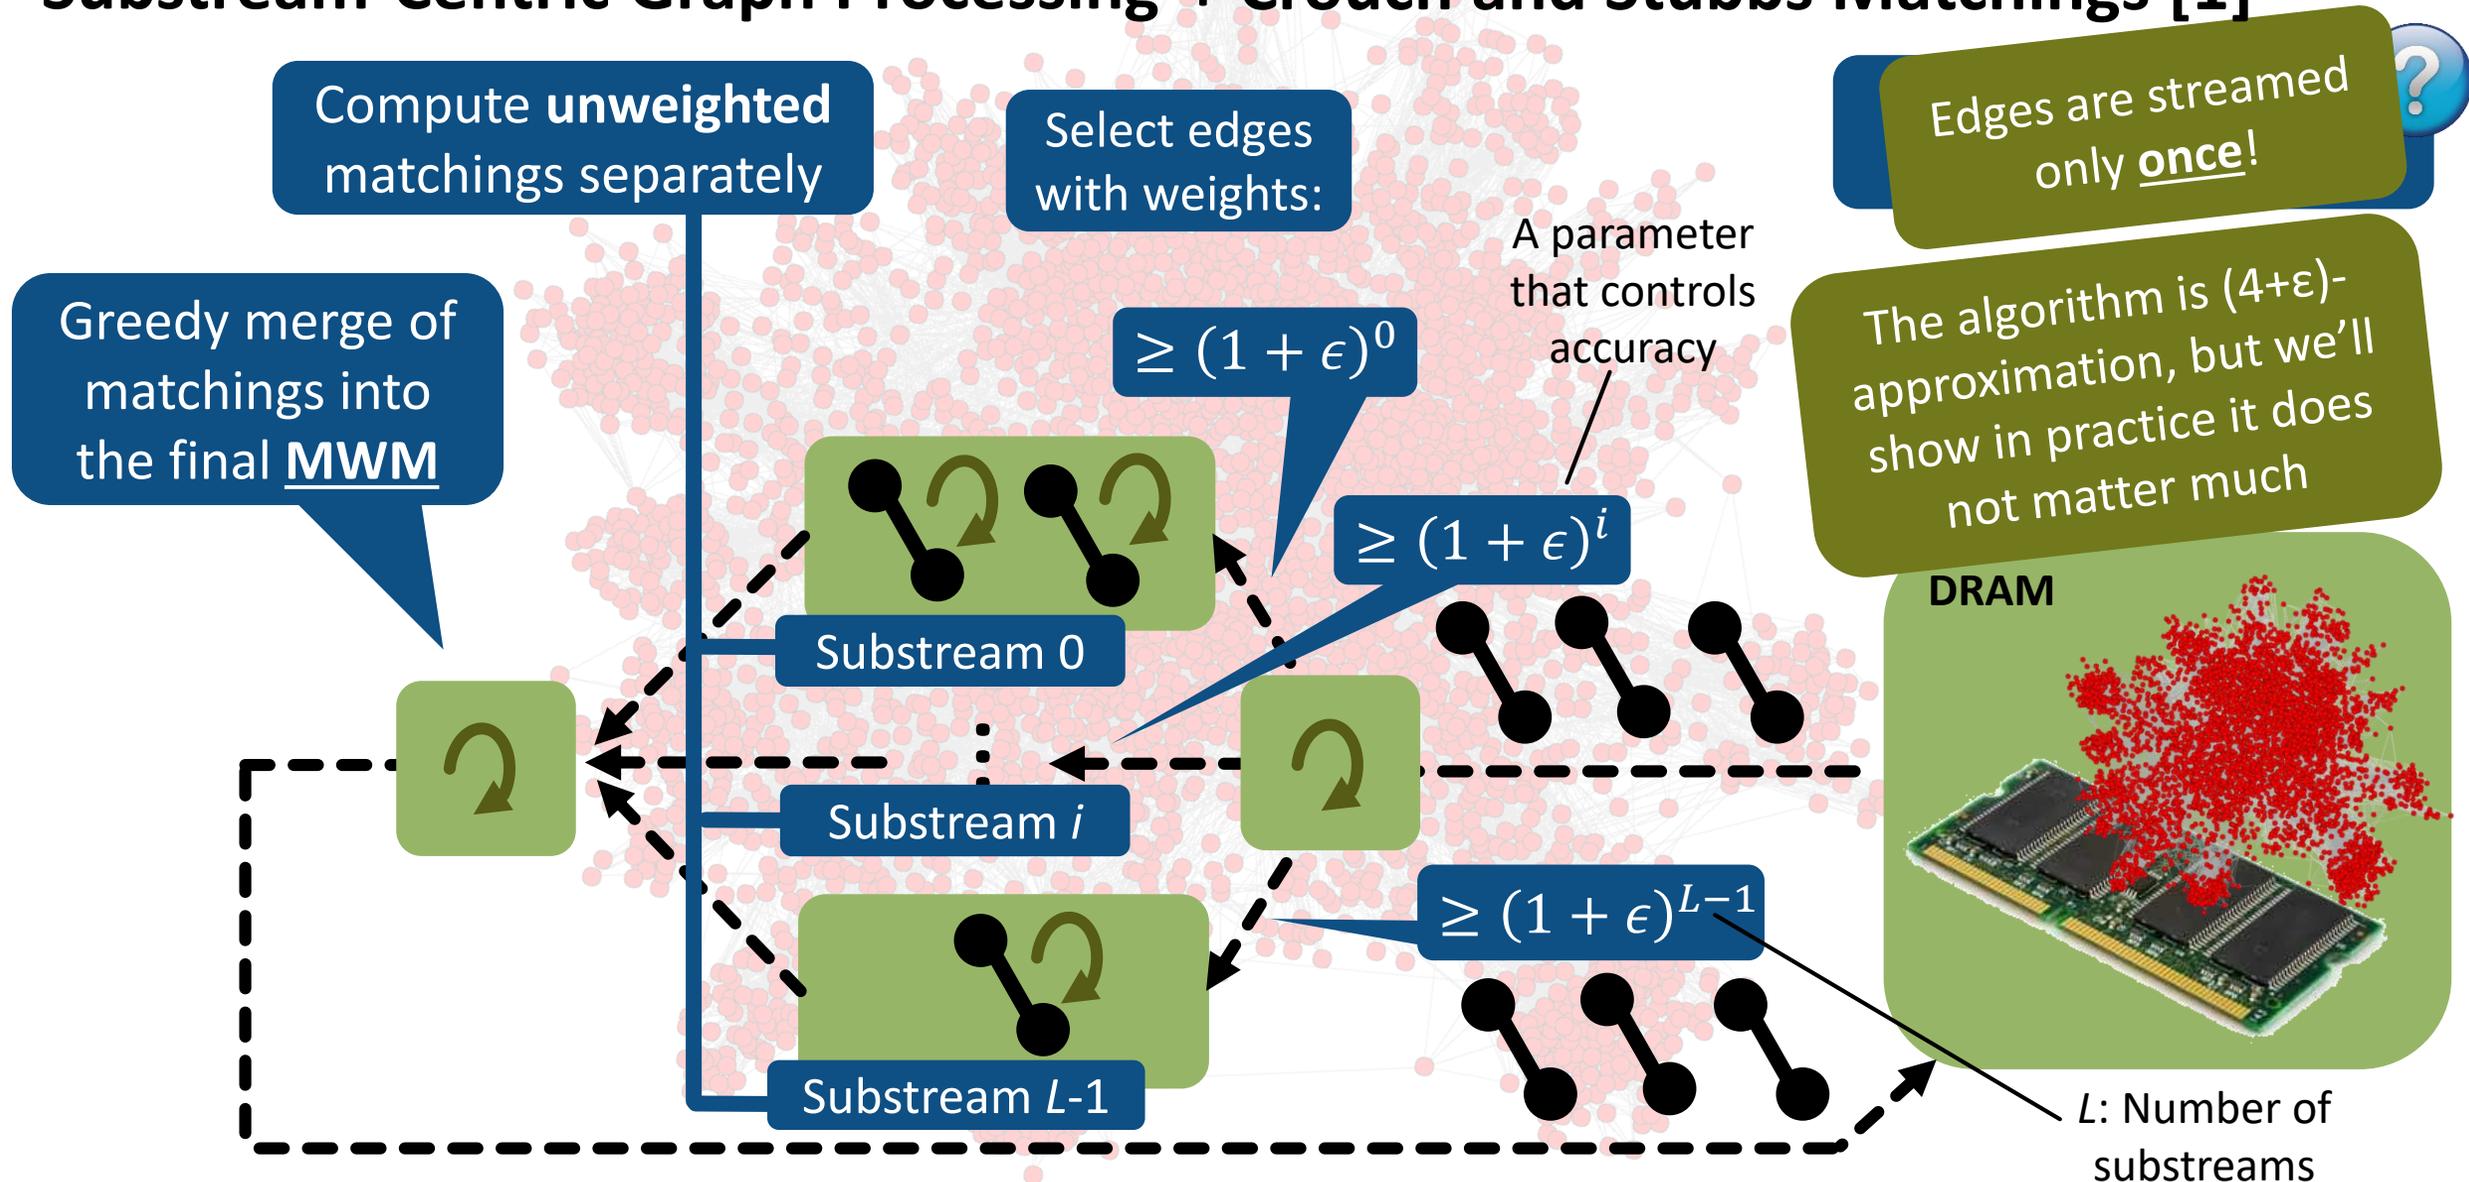
[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

Substream-Centric Graph Processing + Crouch and Stubbs Matchings [1]



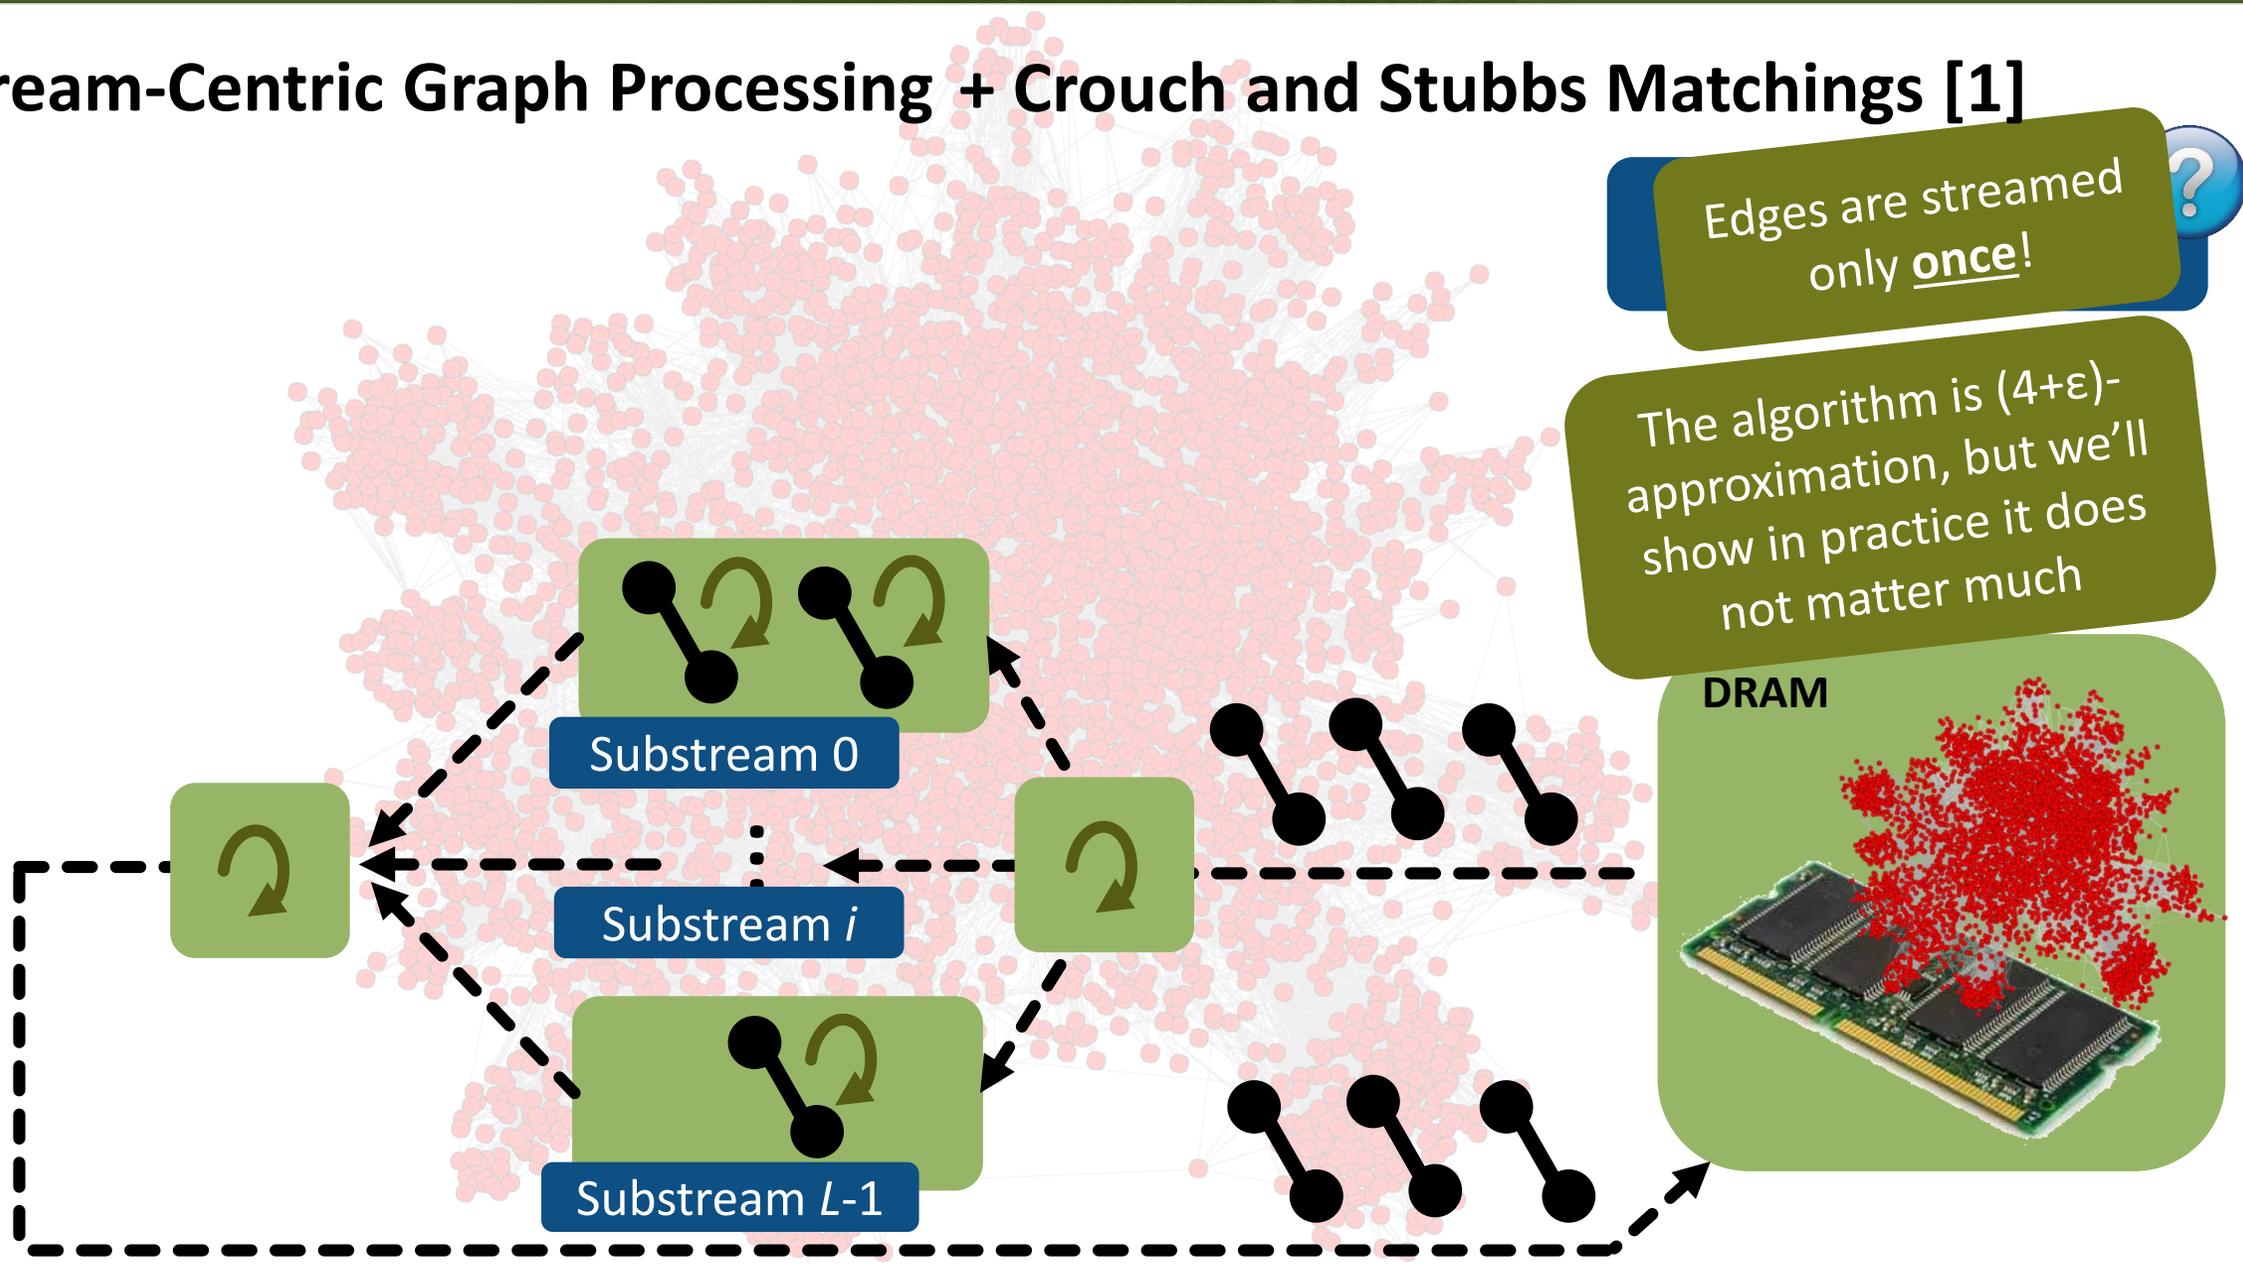
[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

Substream-Centric Graph Processing + Crouch and Stubbs Matchings [1]



[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

Substream-Centric Graph Processing + Crouch and Stubbs Matchings [1]

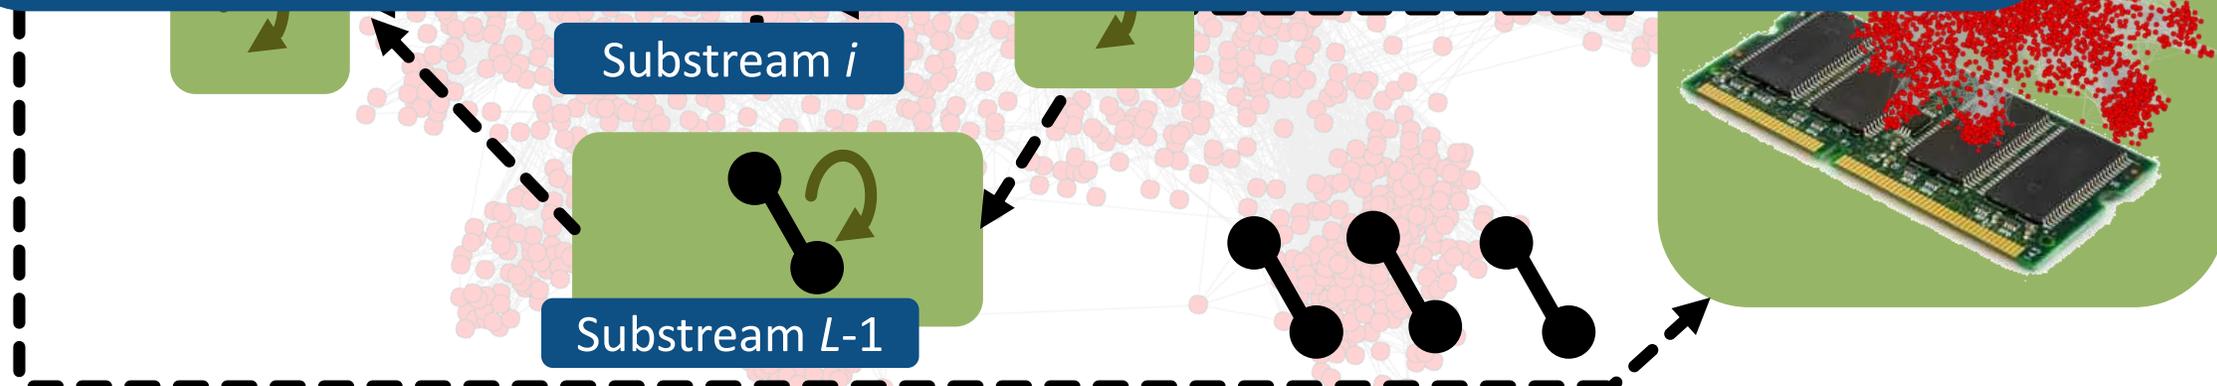


[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

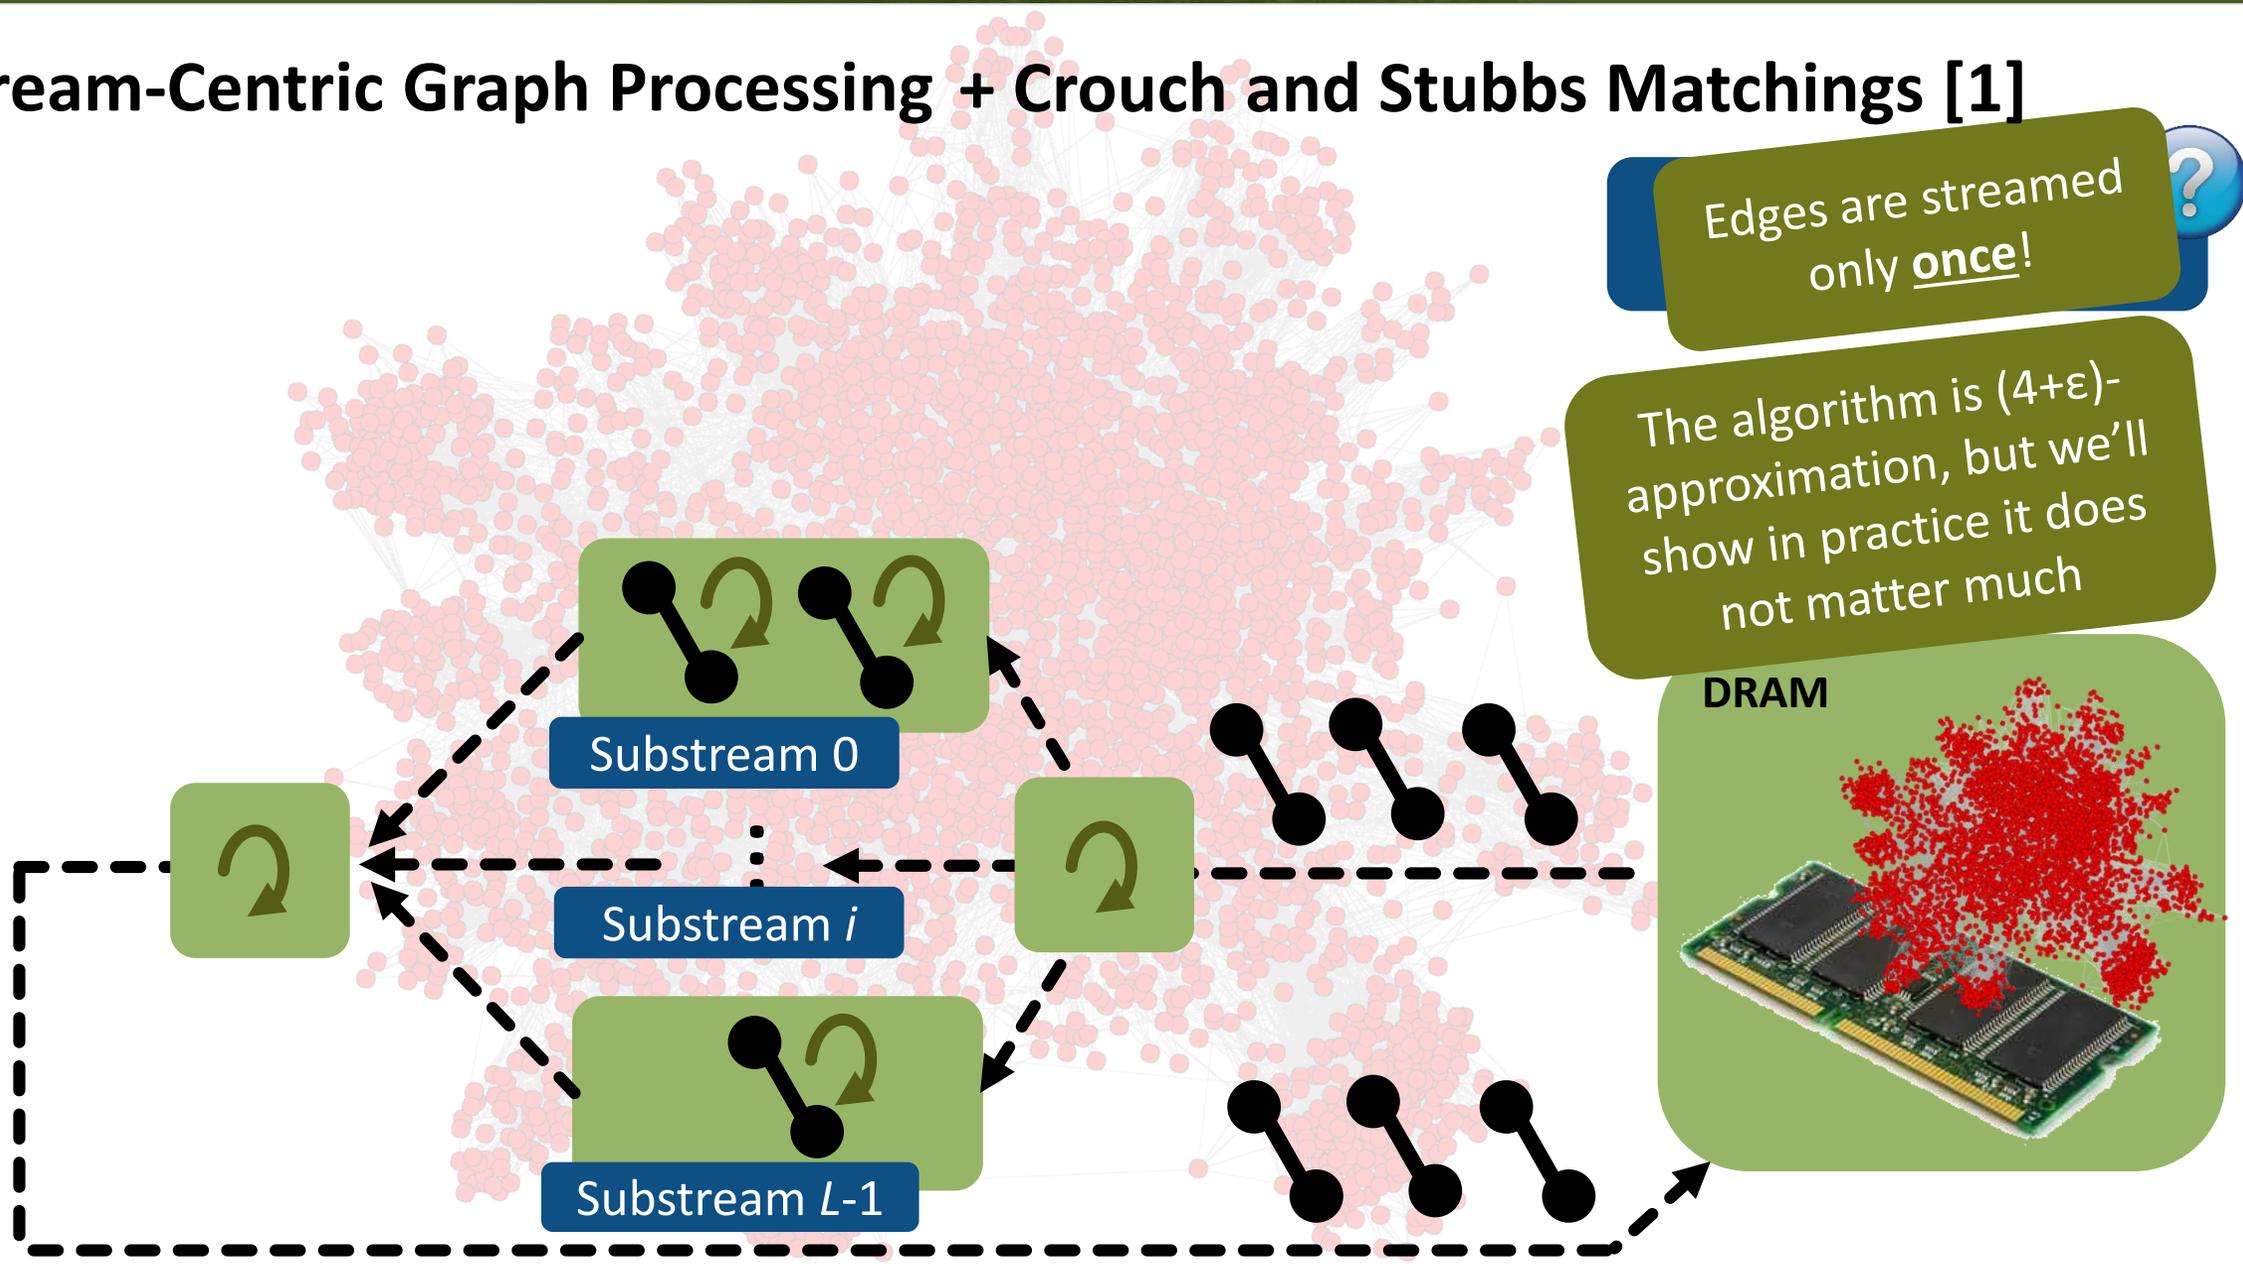
Substream-Centric Graph Processing + Crouch and Stubbs Matchings [1]

Edges are streamed only once!

Mapping the algorithm to the „right” hardware configuration

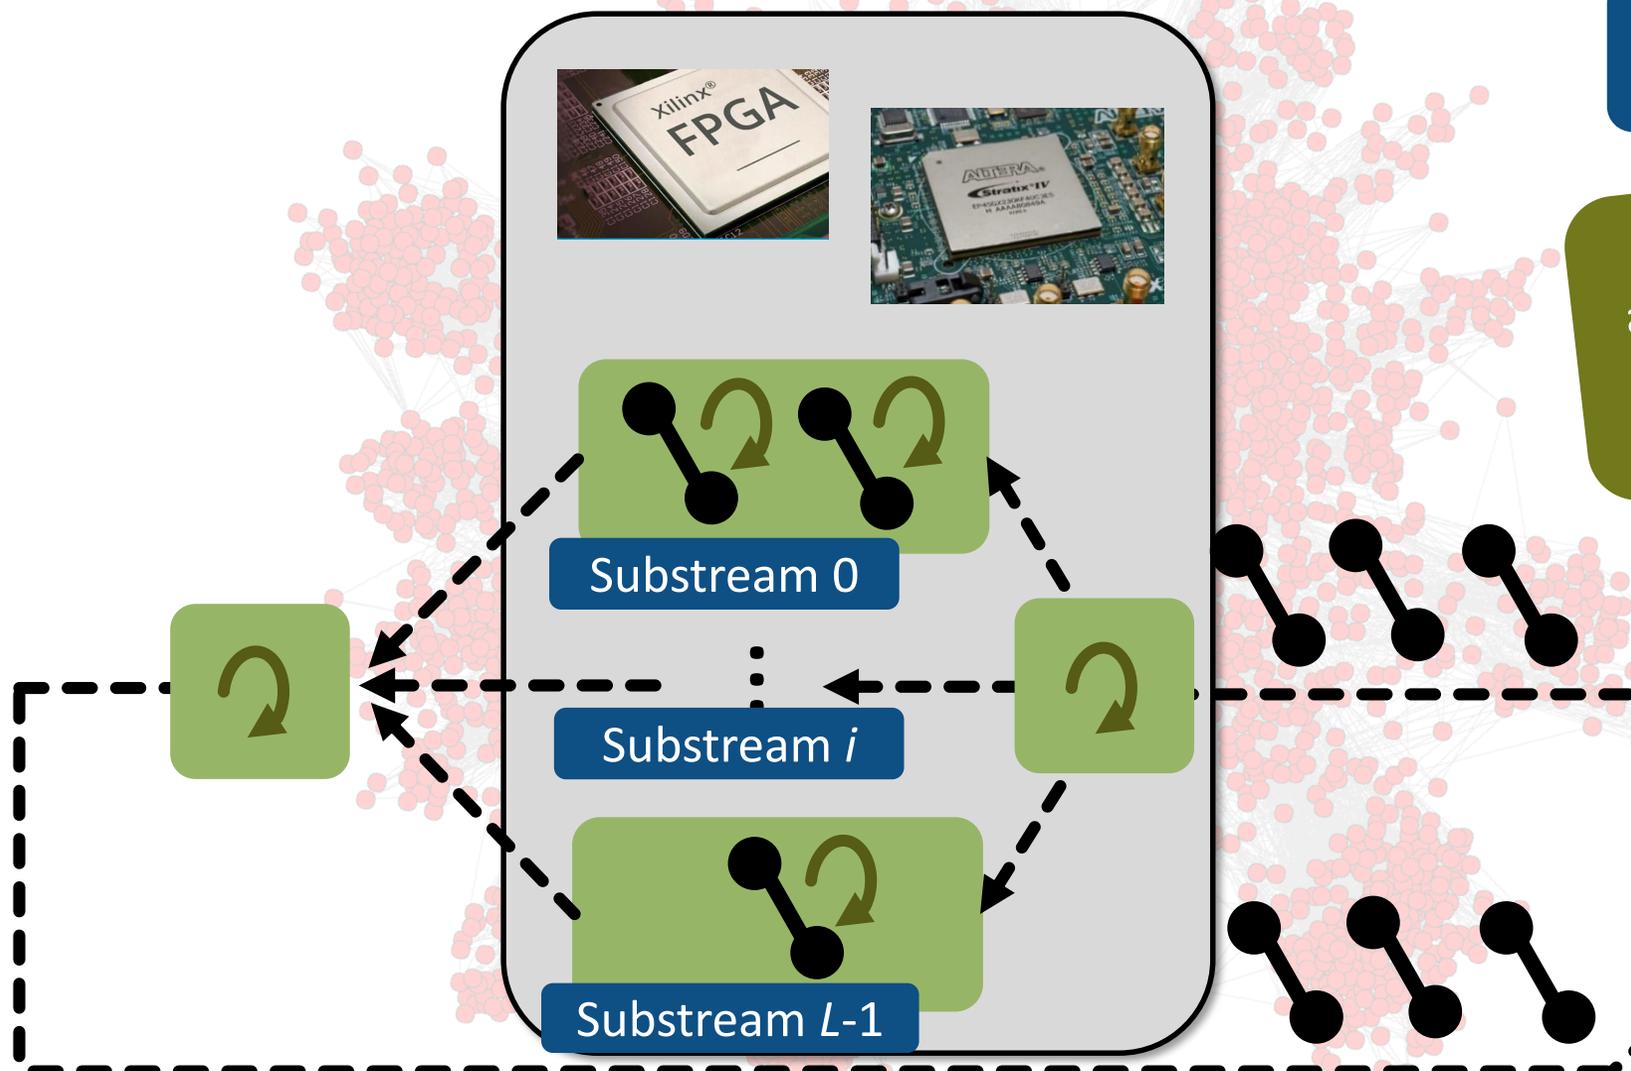


Substream-Centric Graph Processing + Crouch and Stubbs Matchings [1]



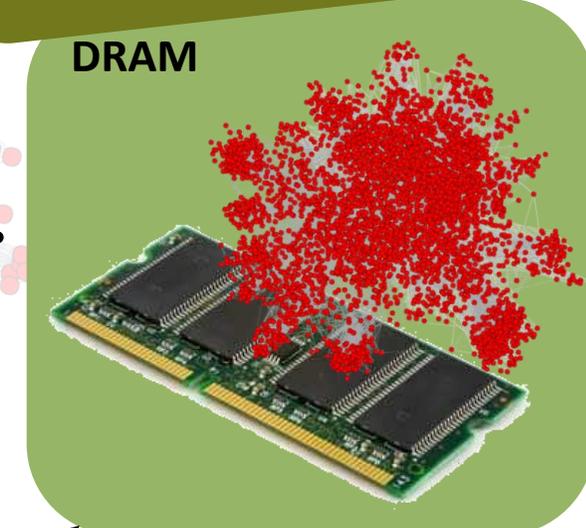
[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

Substream-Centric Graph Processing + Crouch and Stubbs Matchings [1]



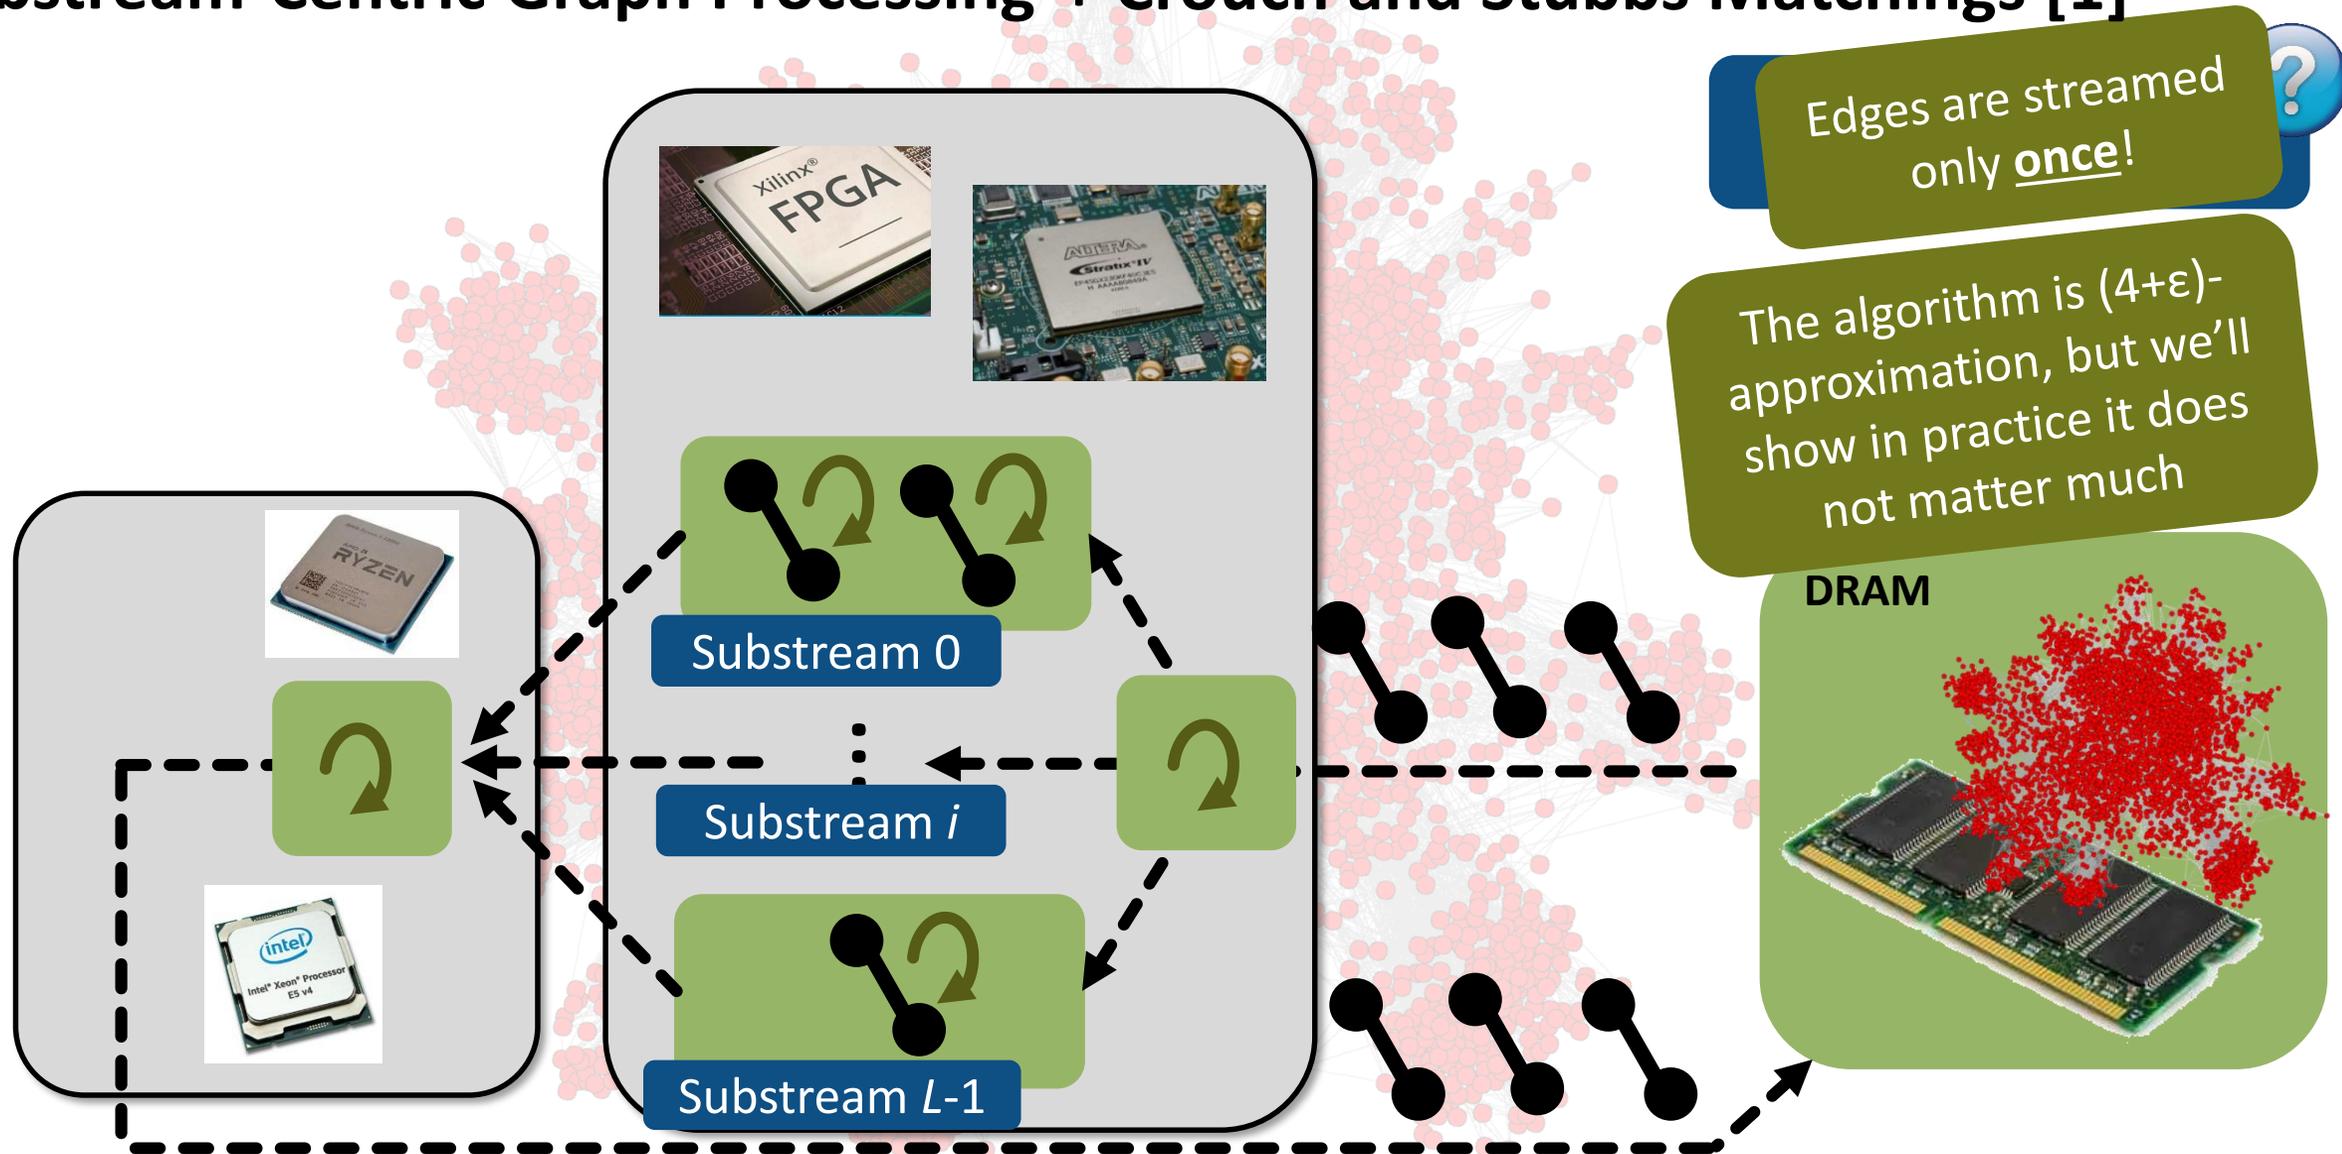
Edges are streamed only once!

The algorithm is $(4+\epsilon)$ -approximation, but we'll show in practice it does not matter much

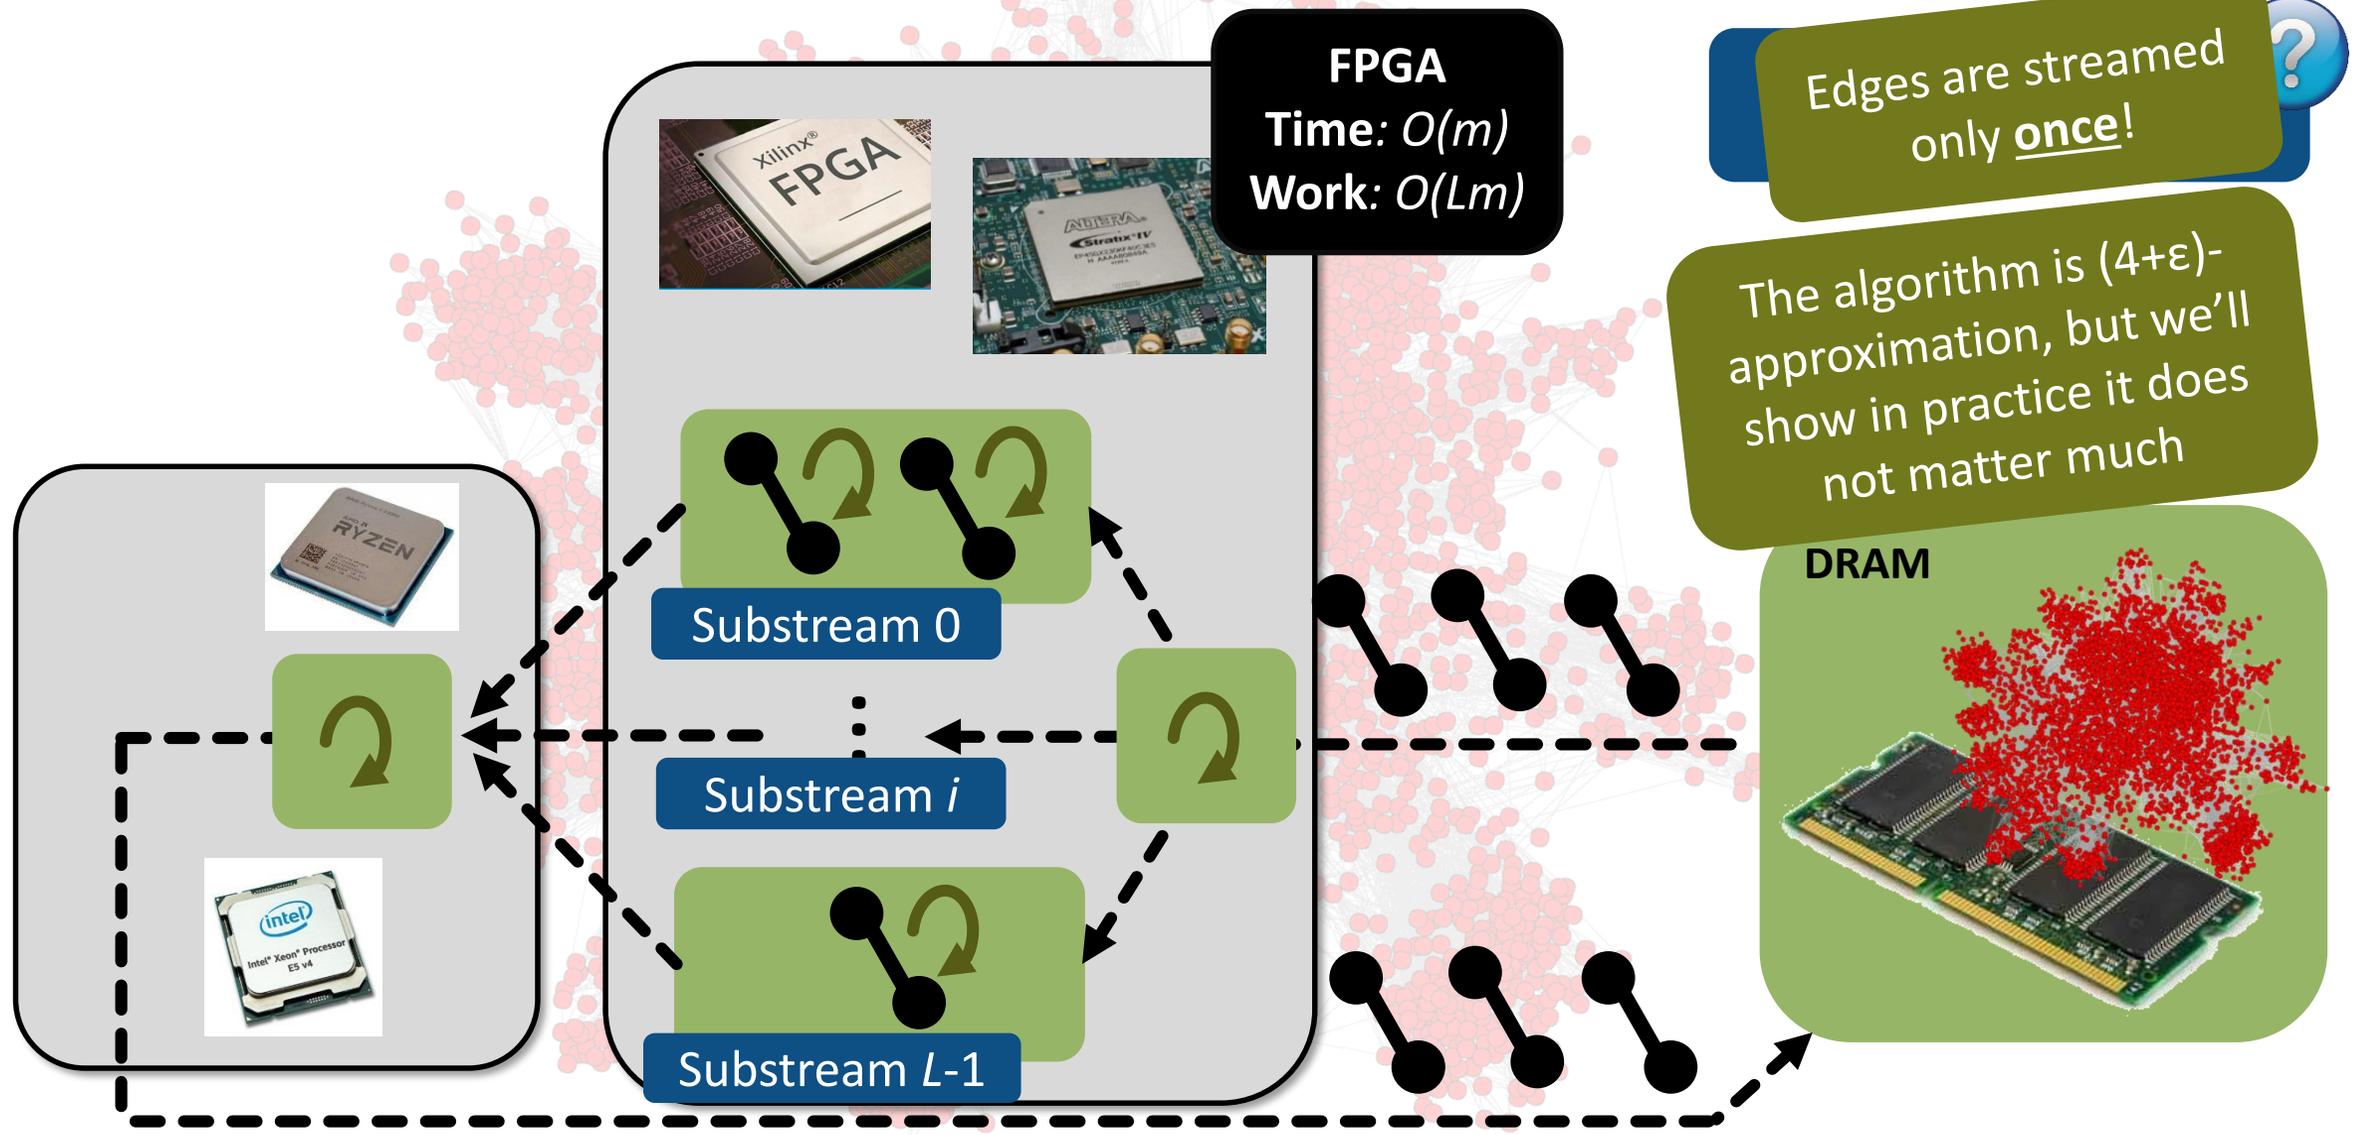


[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

Substream-Centric Graph Processing + Crouch and Stubbs Matchings [1]

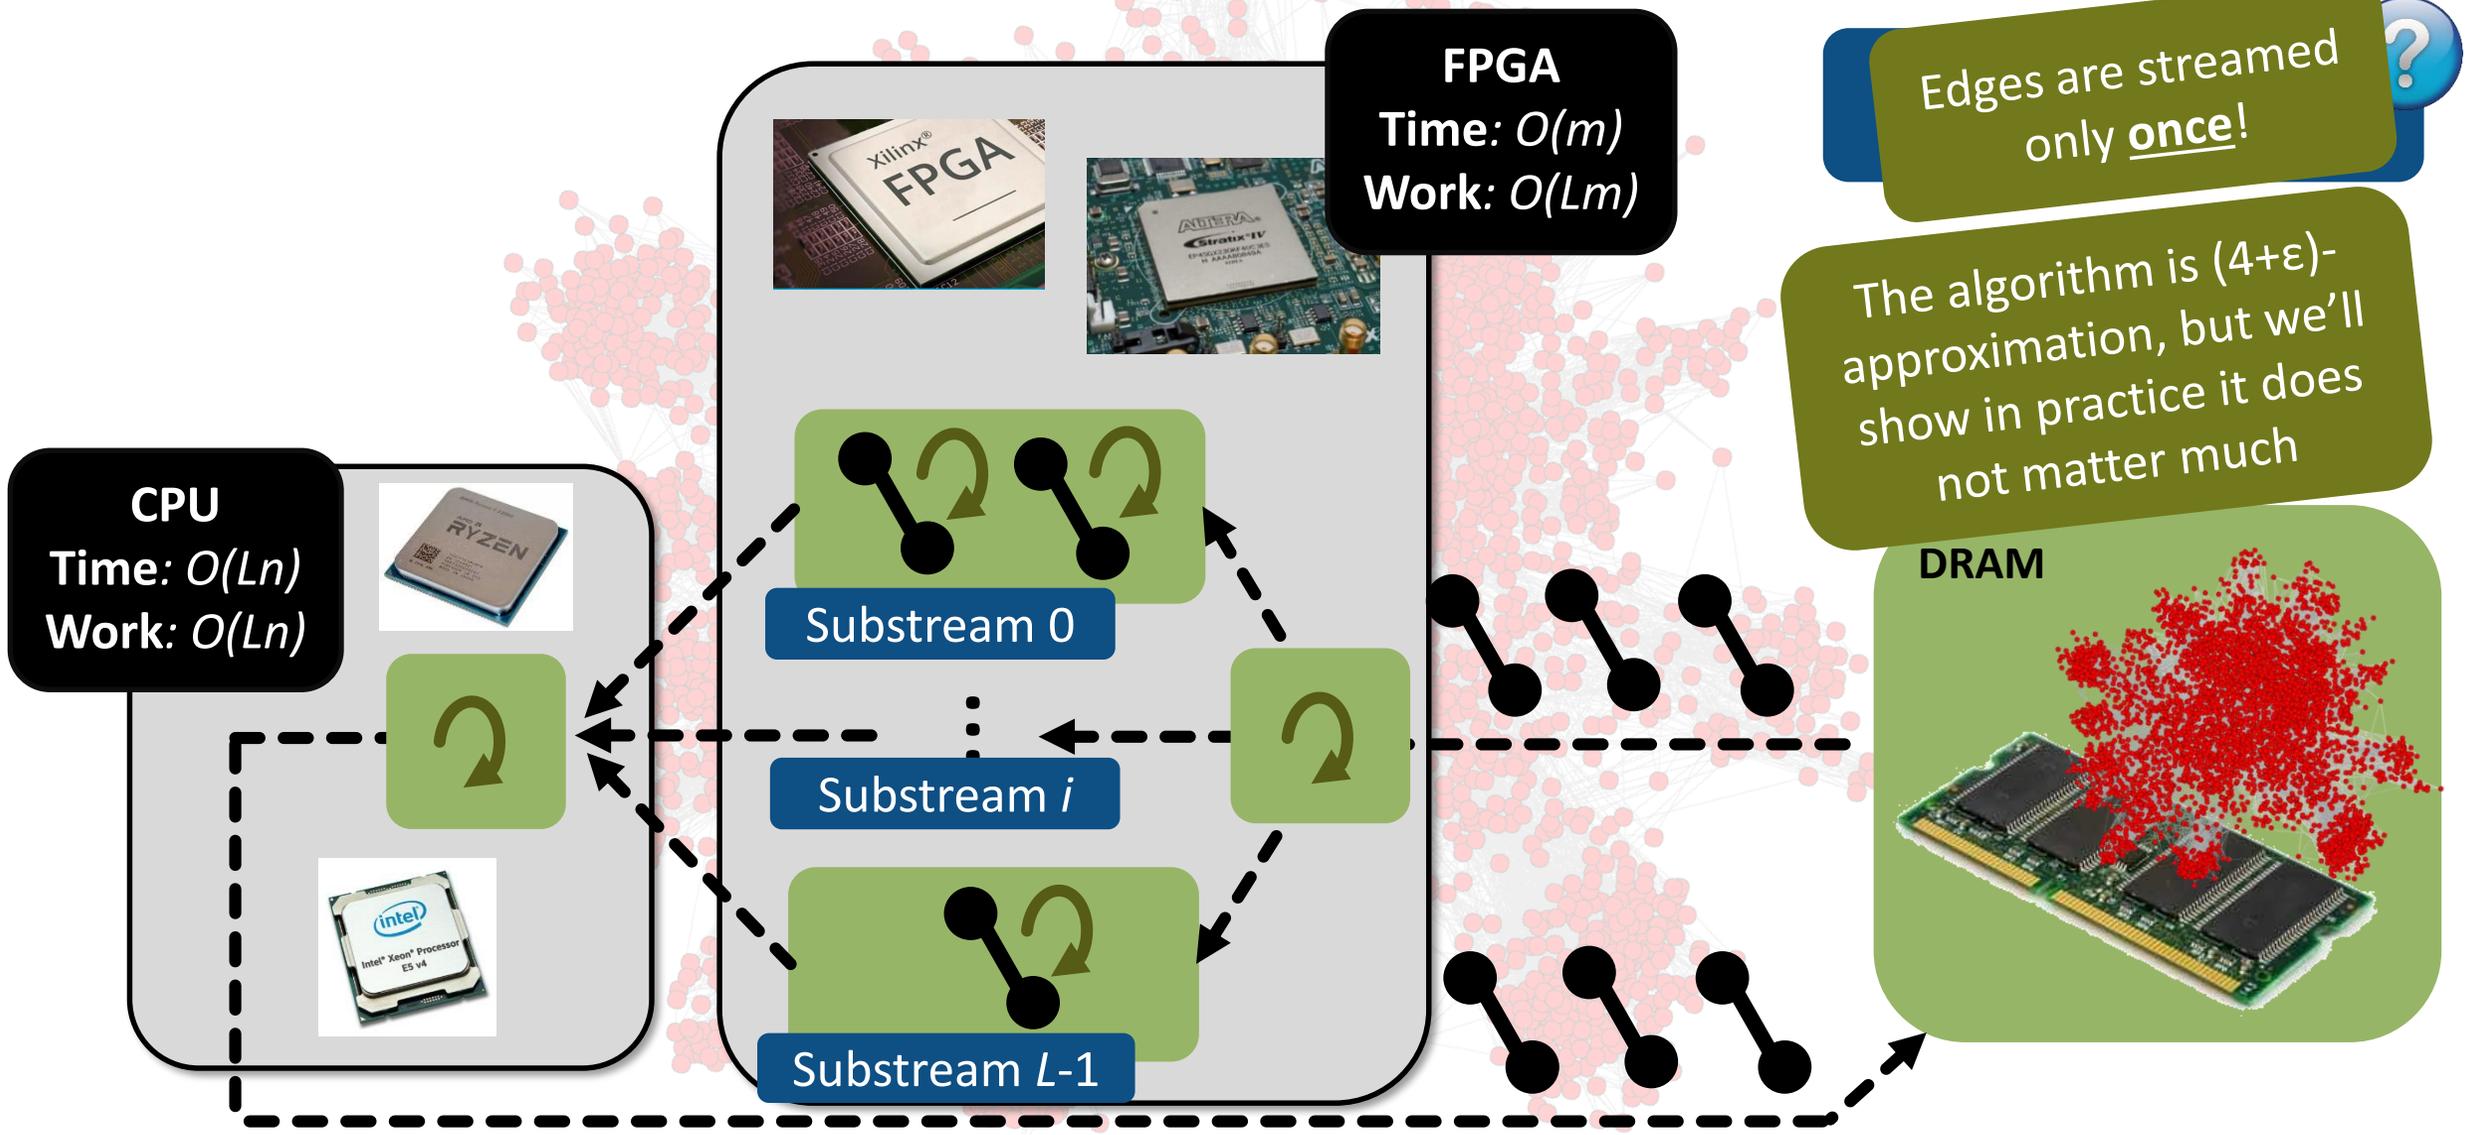


Substream-Centric Graph Processing + Crouch and Stubbs Matchings [1]



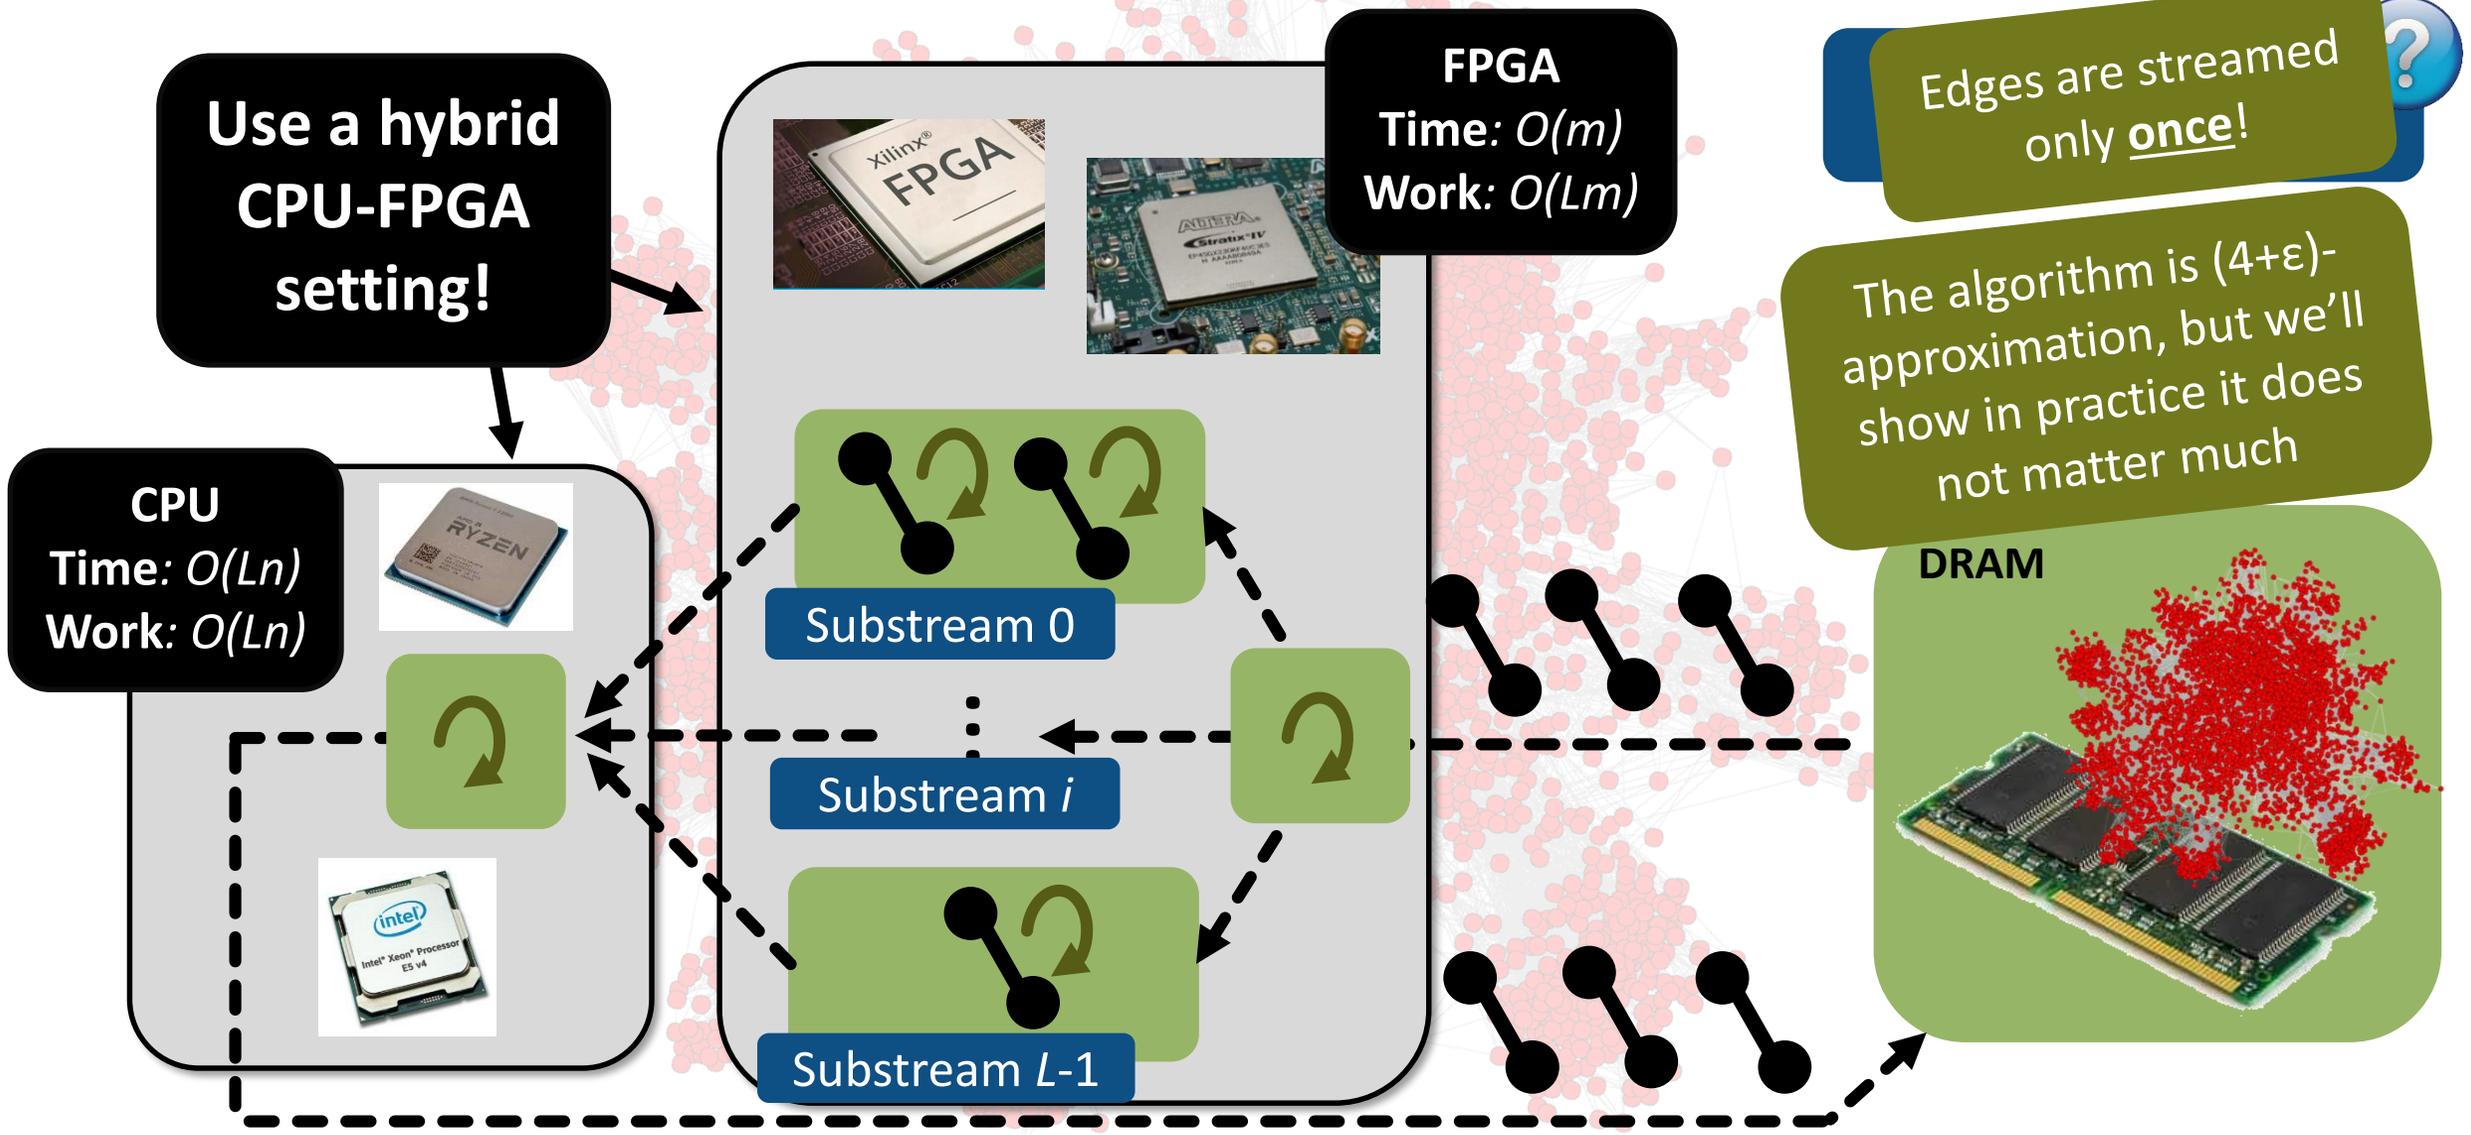
[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

Substream-Centric Graph Processing + Crouch and Stubbs Matchings [1]



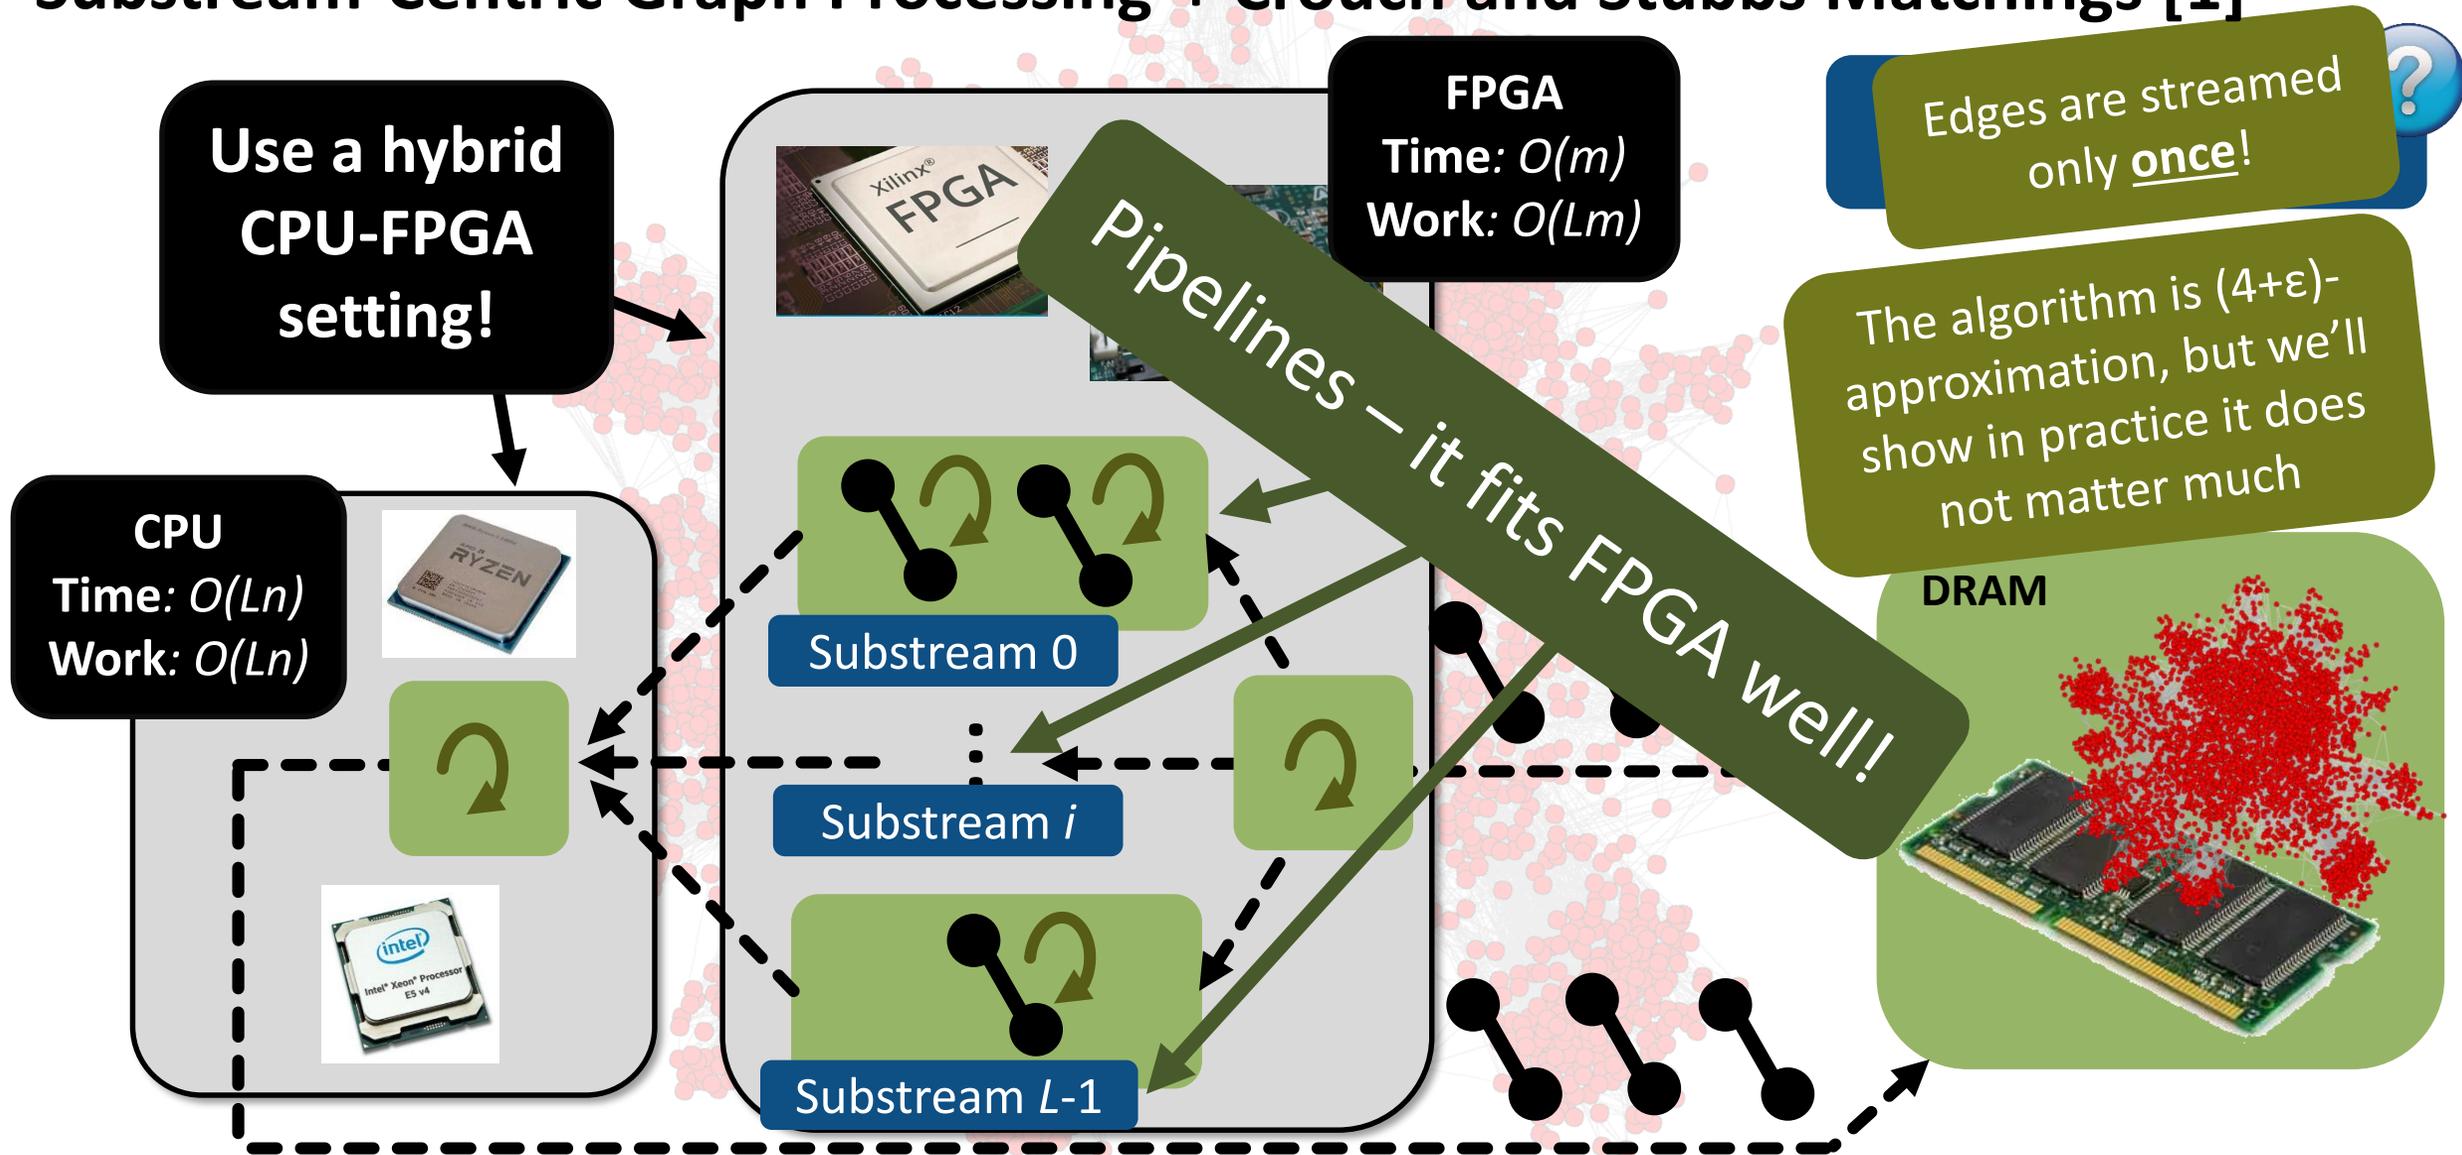
[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

Substream-Centric Graph Processing + Crouch and Stubbs Matchings [1]

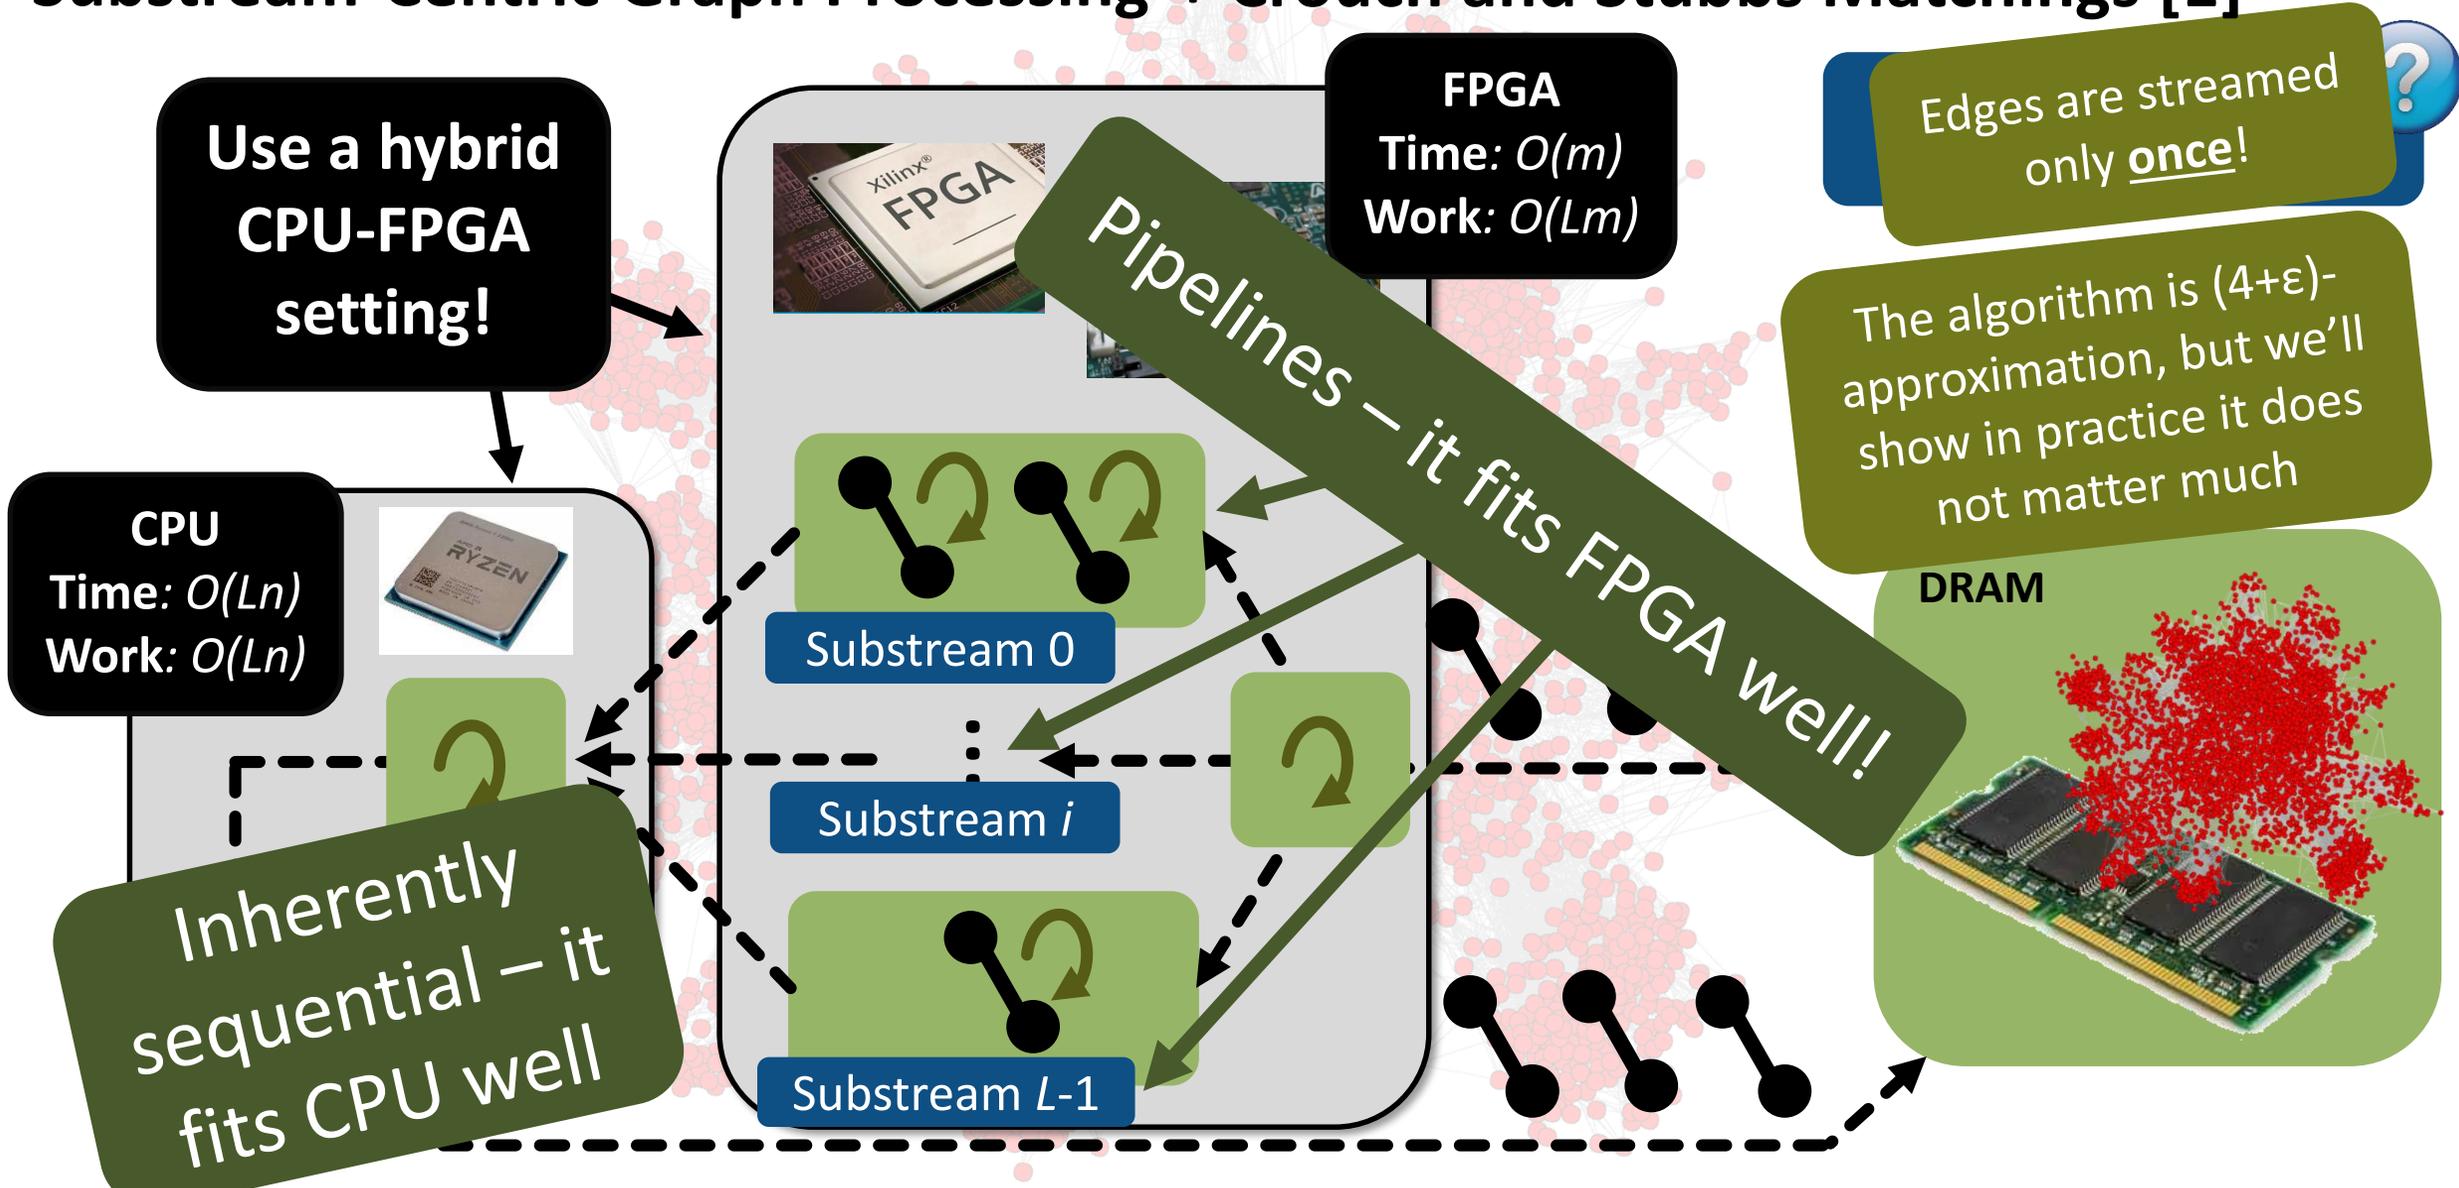


[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

Substream-Centric Graph Processing + Crouch and Stubbs Matchings [1]



Substream-Centric Graph Processing + Crouch and Stubbs Matchings [1]



[1] M. Crouch and D. M. Stubbs. Improved streaming Algorithms for weighted Matching, via unweighted Matching. LIPIcs-Leibniz Informatics. 2014.

PERFORMANCE ANALYSIS

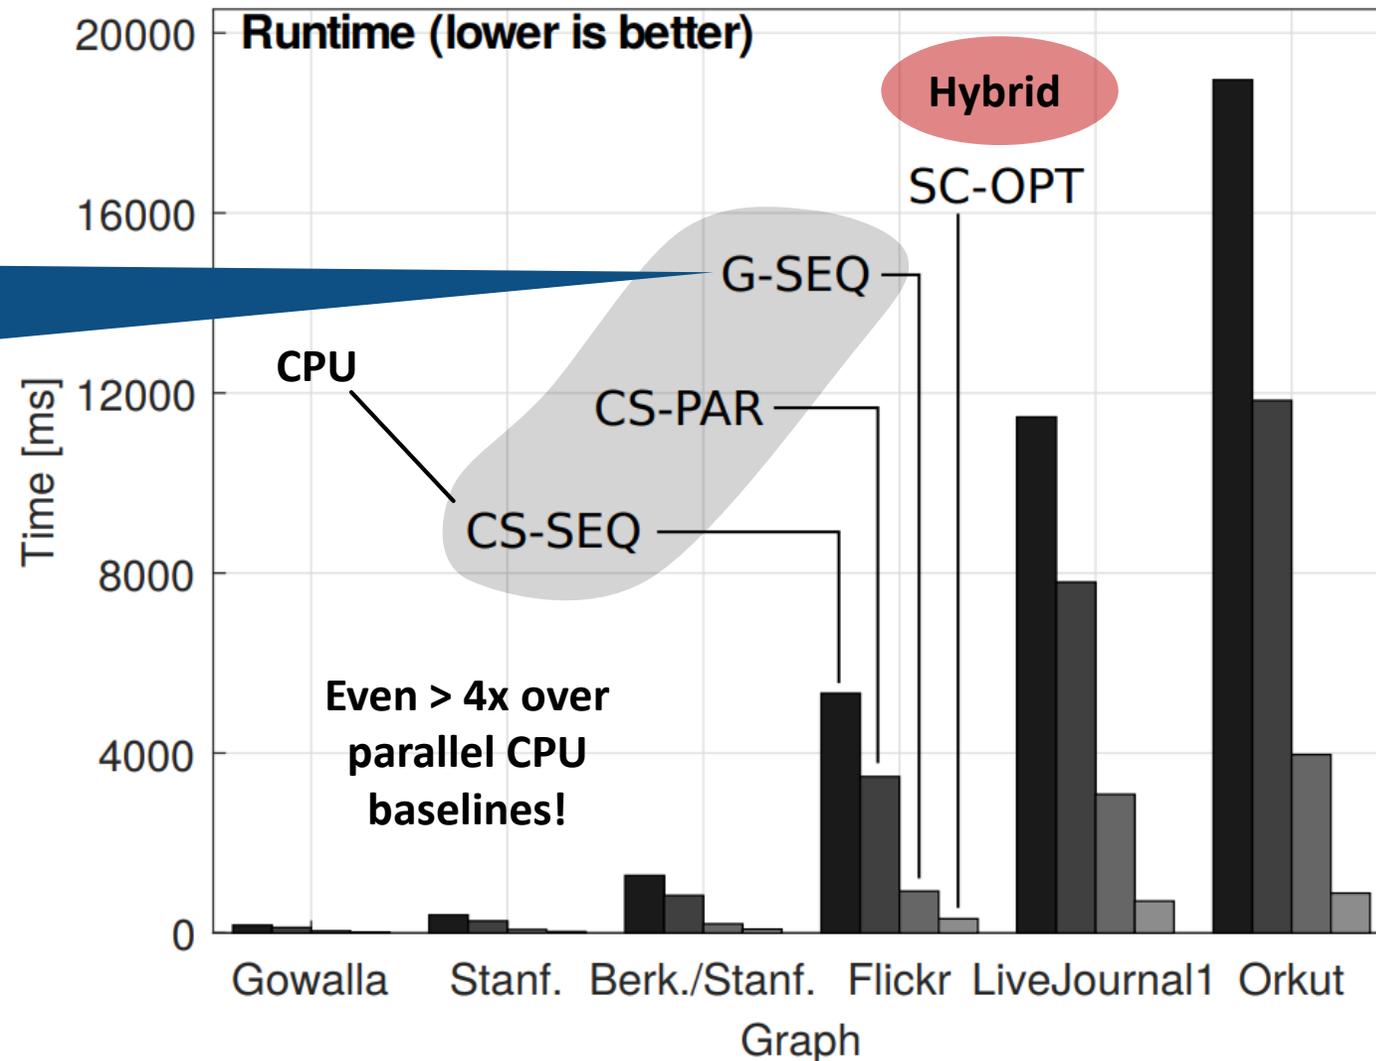
VARIOUS GRAPHS

Parameters:
 #Substreams = 64, #Threads = 4

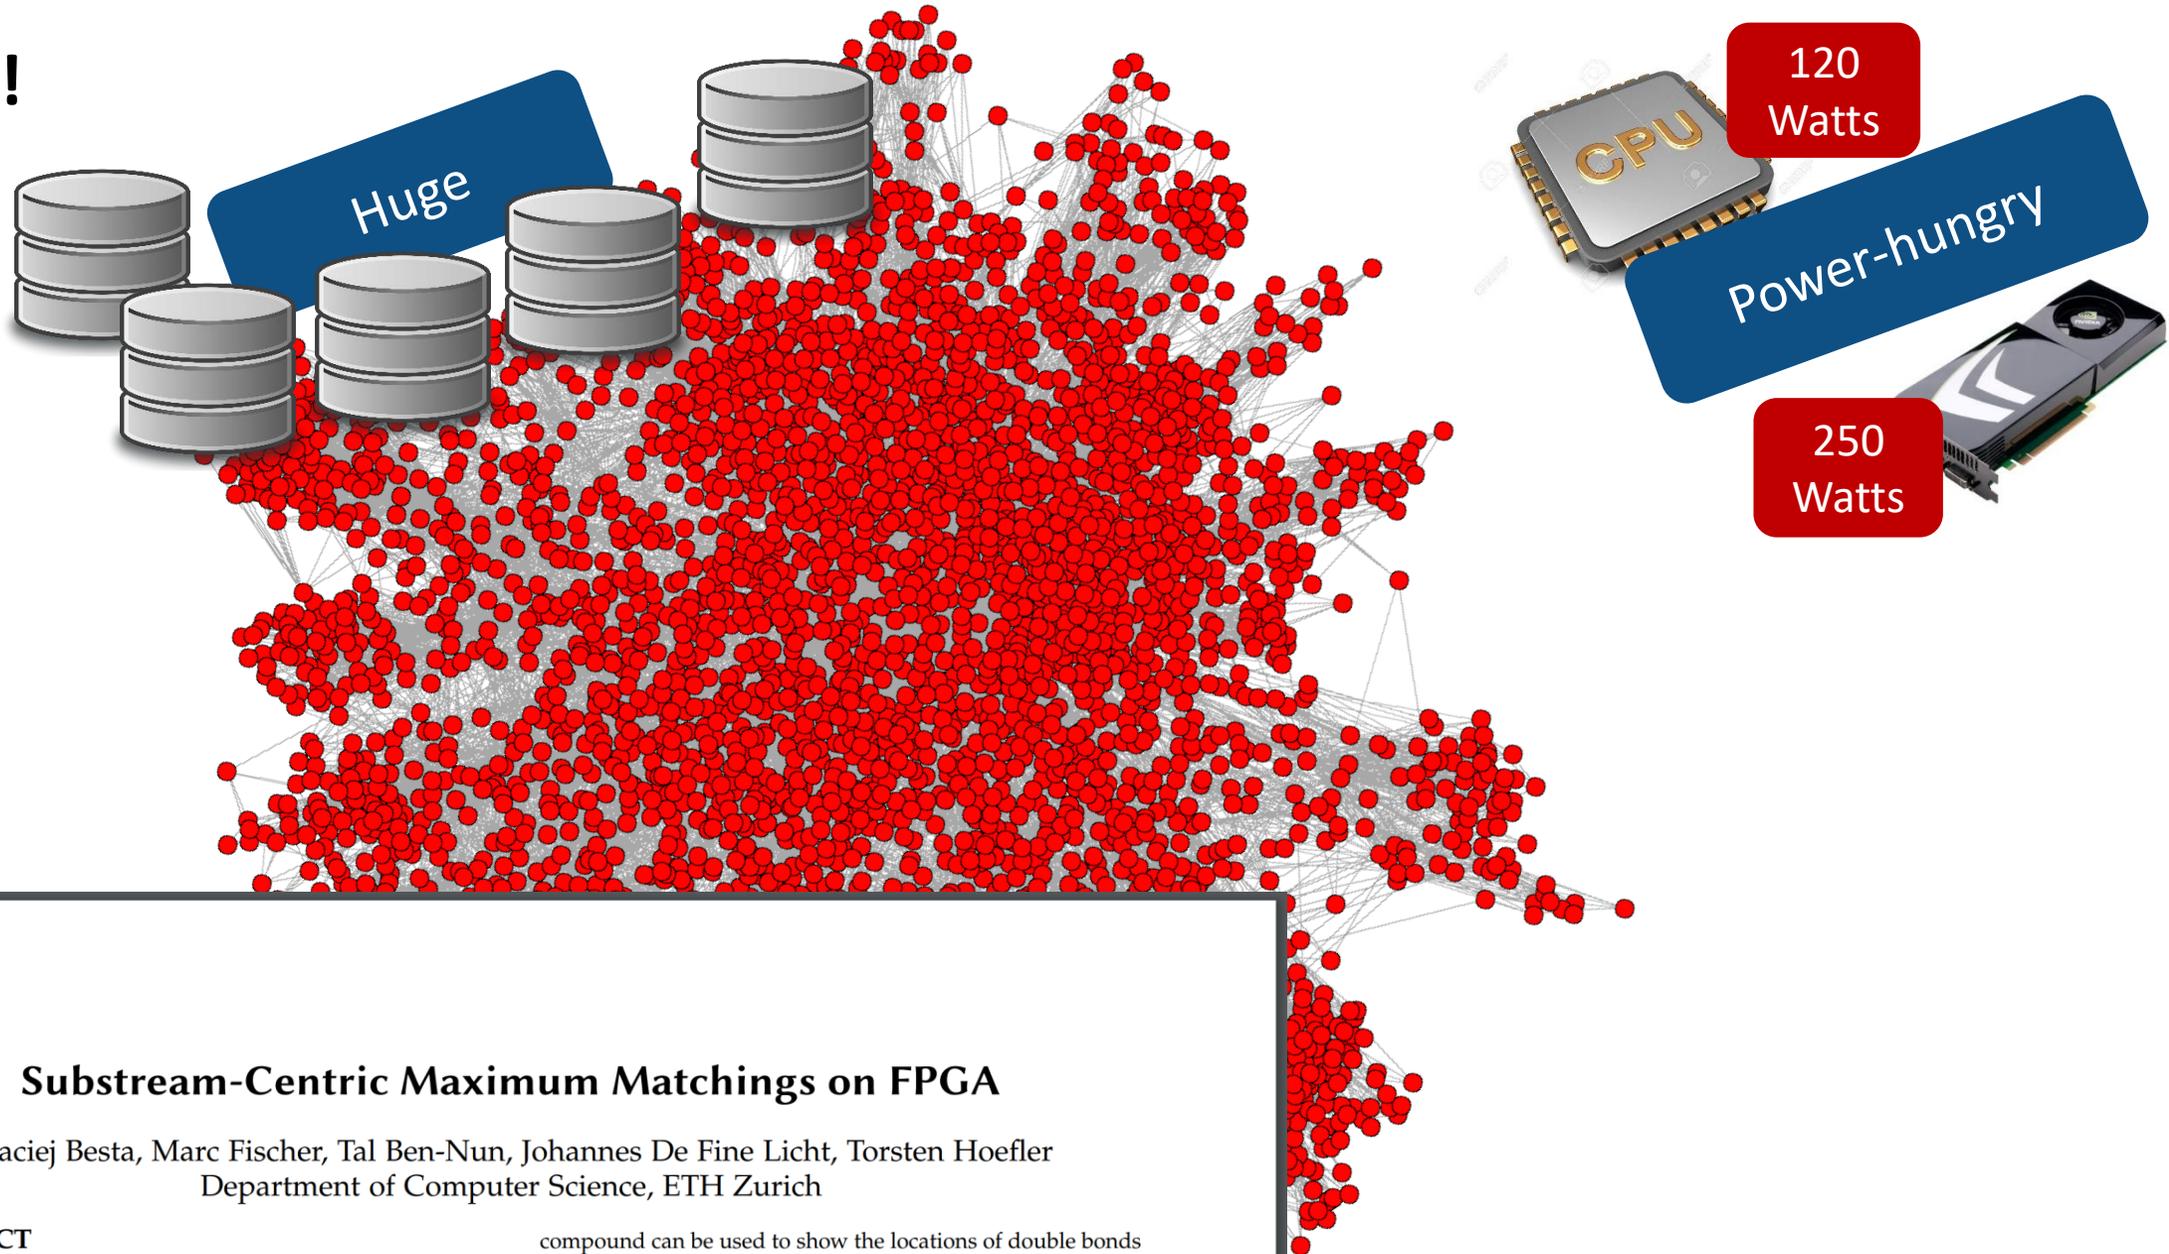
SC-OPT secures highest performance

State-of-the-art MWM algorithm, space-optimal, **time-optimal $O(m)$** , $(2+\epsilon)$ -approximation

Graph	Type	m	n
Kronecker	Synthetic power-law	$\approx 48n$	$2^k; k = 16, \dots, 21$
Gowalla	Social network	950,327	196,591
Flickr	Social network	33,140,017	2,302,925
LiveJournal1	Social network	68,993,773	4,847,571
Orkut	Social network	117,184,899	3,072,441
Stanford	Hyperlink graph	2,312,497	281,903
Berkeley	Hyperlink graph	7,600,595	685,230
arXiv hep-th	Citation graph	352,807	27,770



Problems!



Substream-Centric Maximum Matchings on FPGA

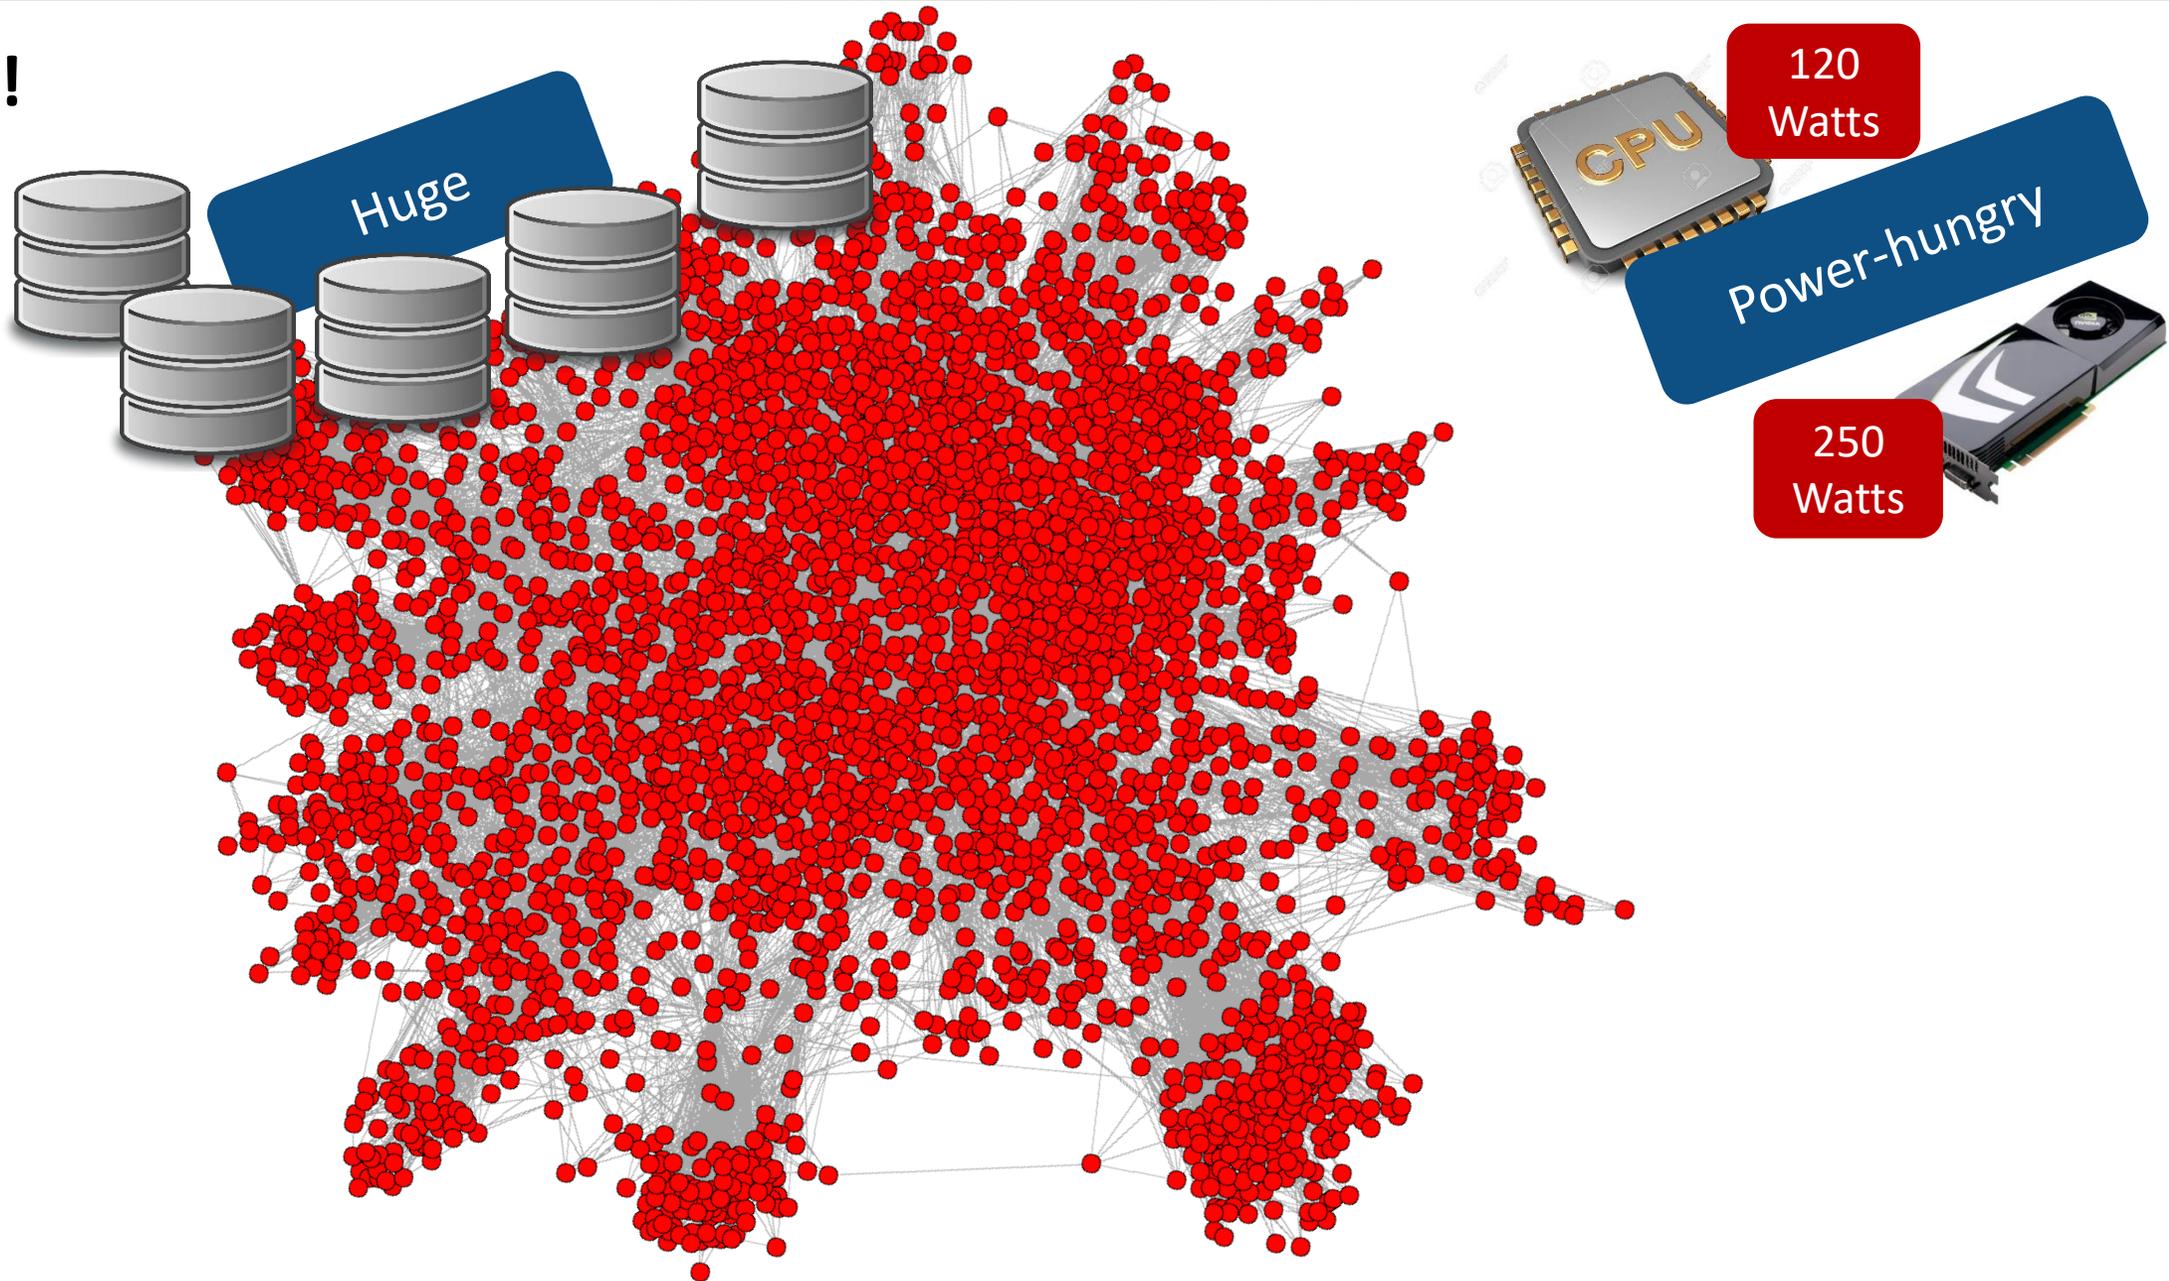
Maciej Besta, Marc Fischer, Tal Ben-Nun, Johannes De Fine Licht, Torsten Hoefler
Department of Computer Science, ETH Zurich

ABSTRACT

Developing high-performance and energy-efficient algorithms for maximum matchings is becoming increasingly important in social network analysis, computational sciences, schedul

compound can be used to show the locations of double bonds in the chemical structure [59]. As deriving the exact MM is usually computationally expensive, significant focus has been placed on developing fast approximate solutions [17].

Problems!



Problems!



Huge

Power-hungry

120 Watts

250 Watts

Communication-heavy

Synchronization-heavy

Irregular

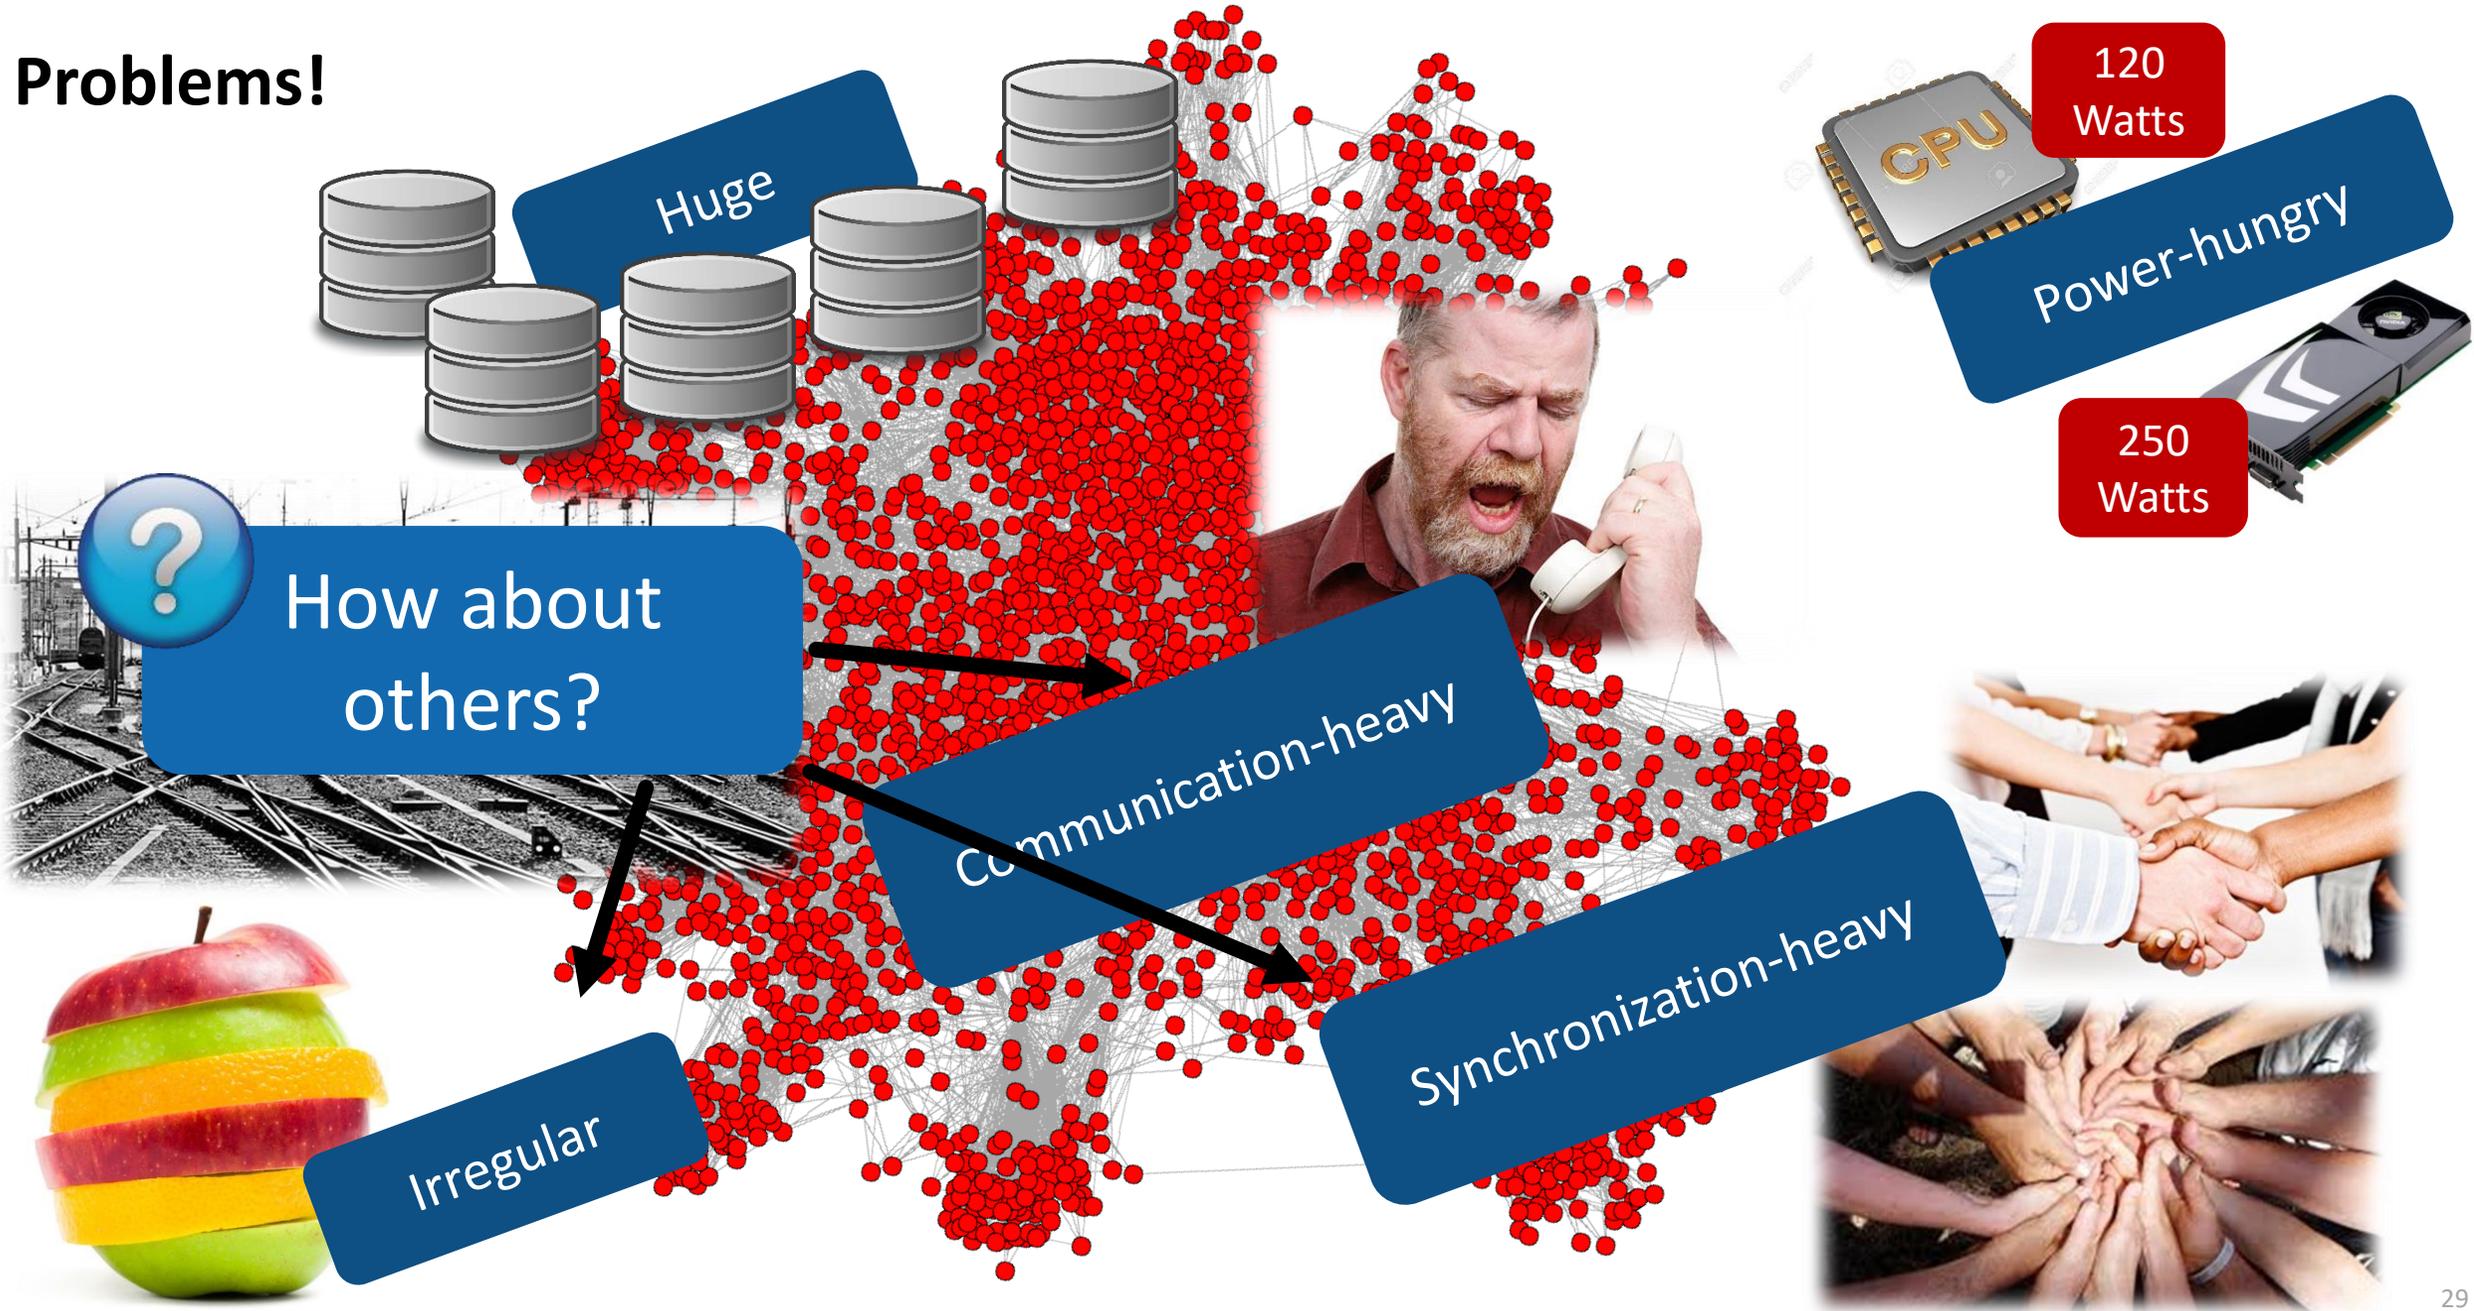
It's Complicated...

CPU

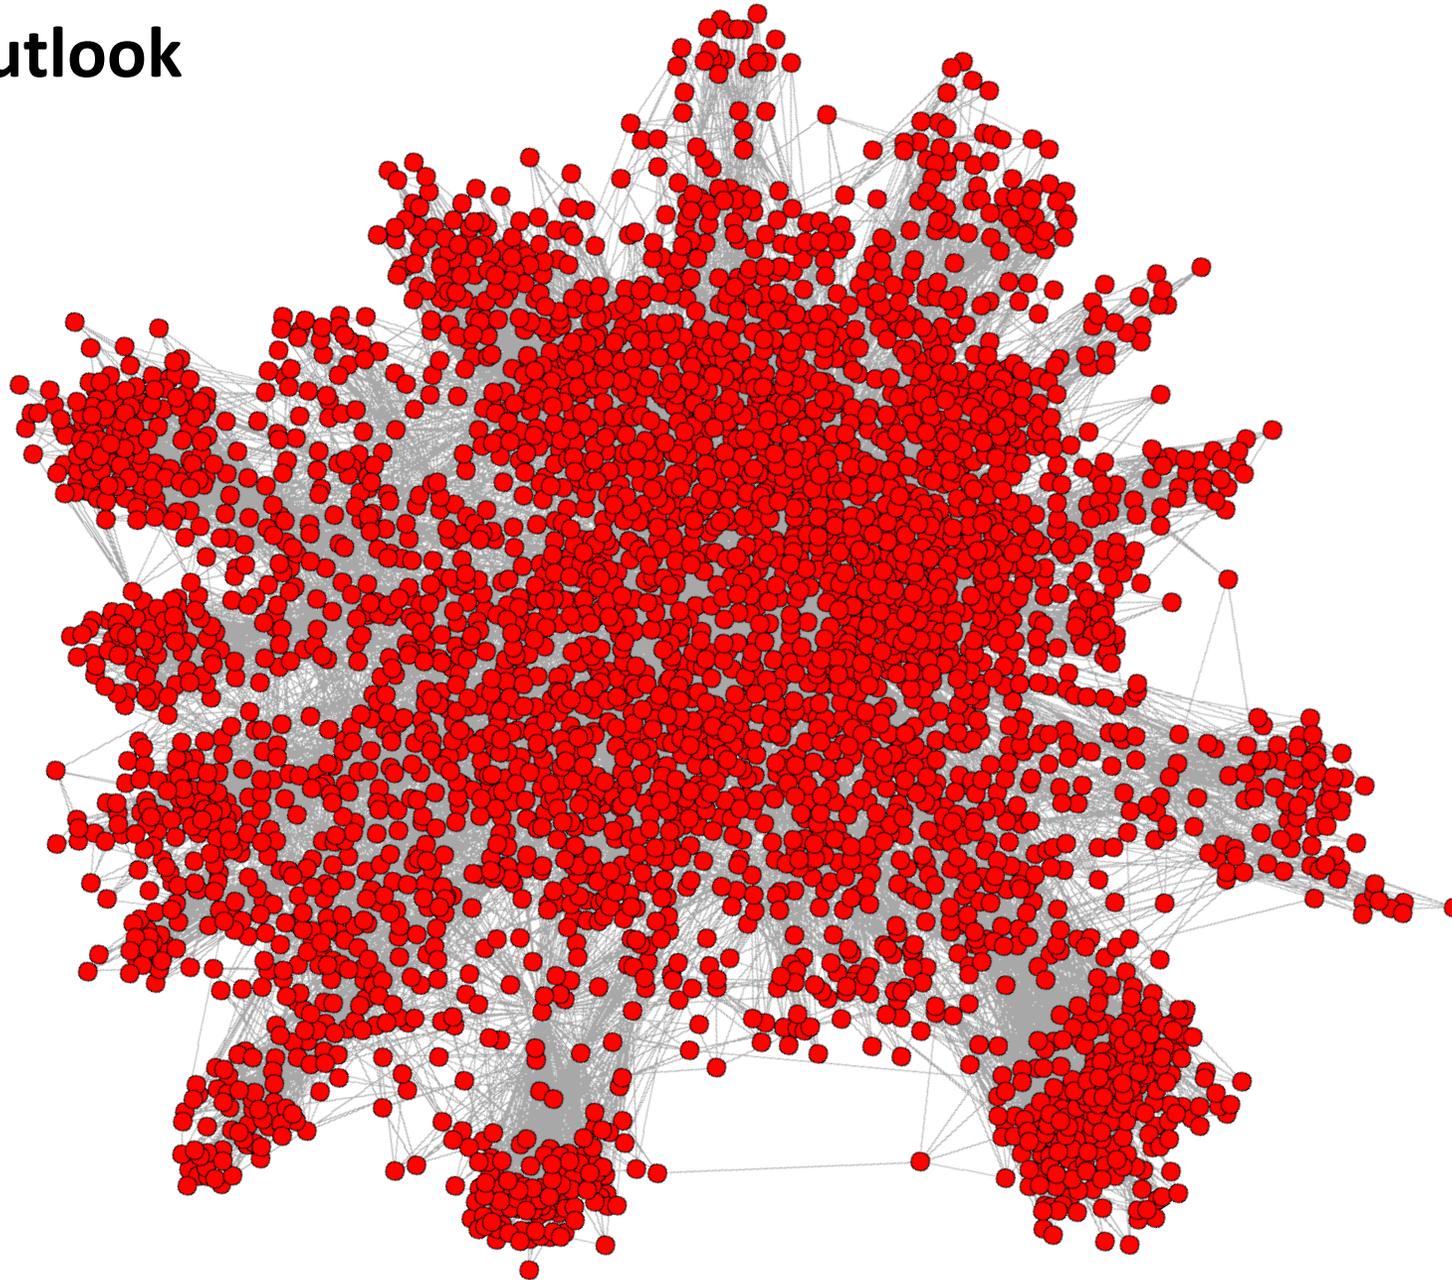
120 Watts

250 Watts

Problems!



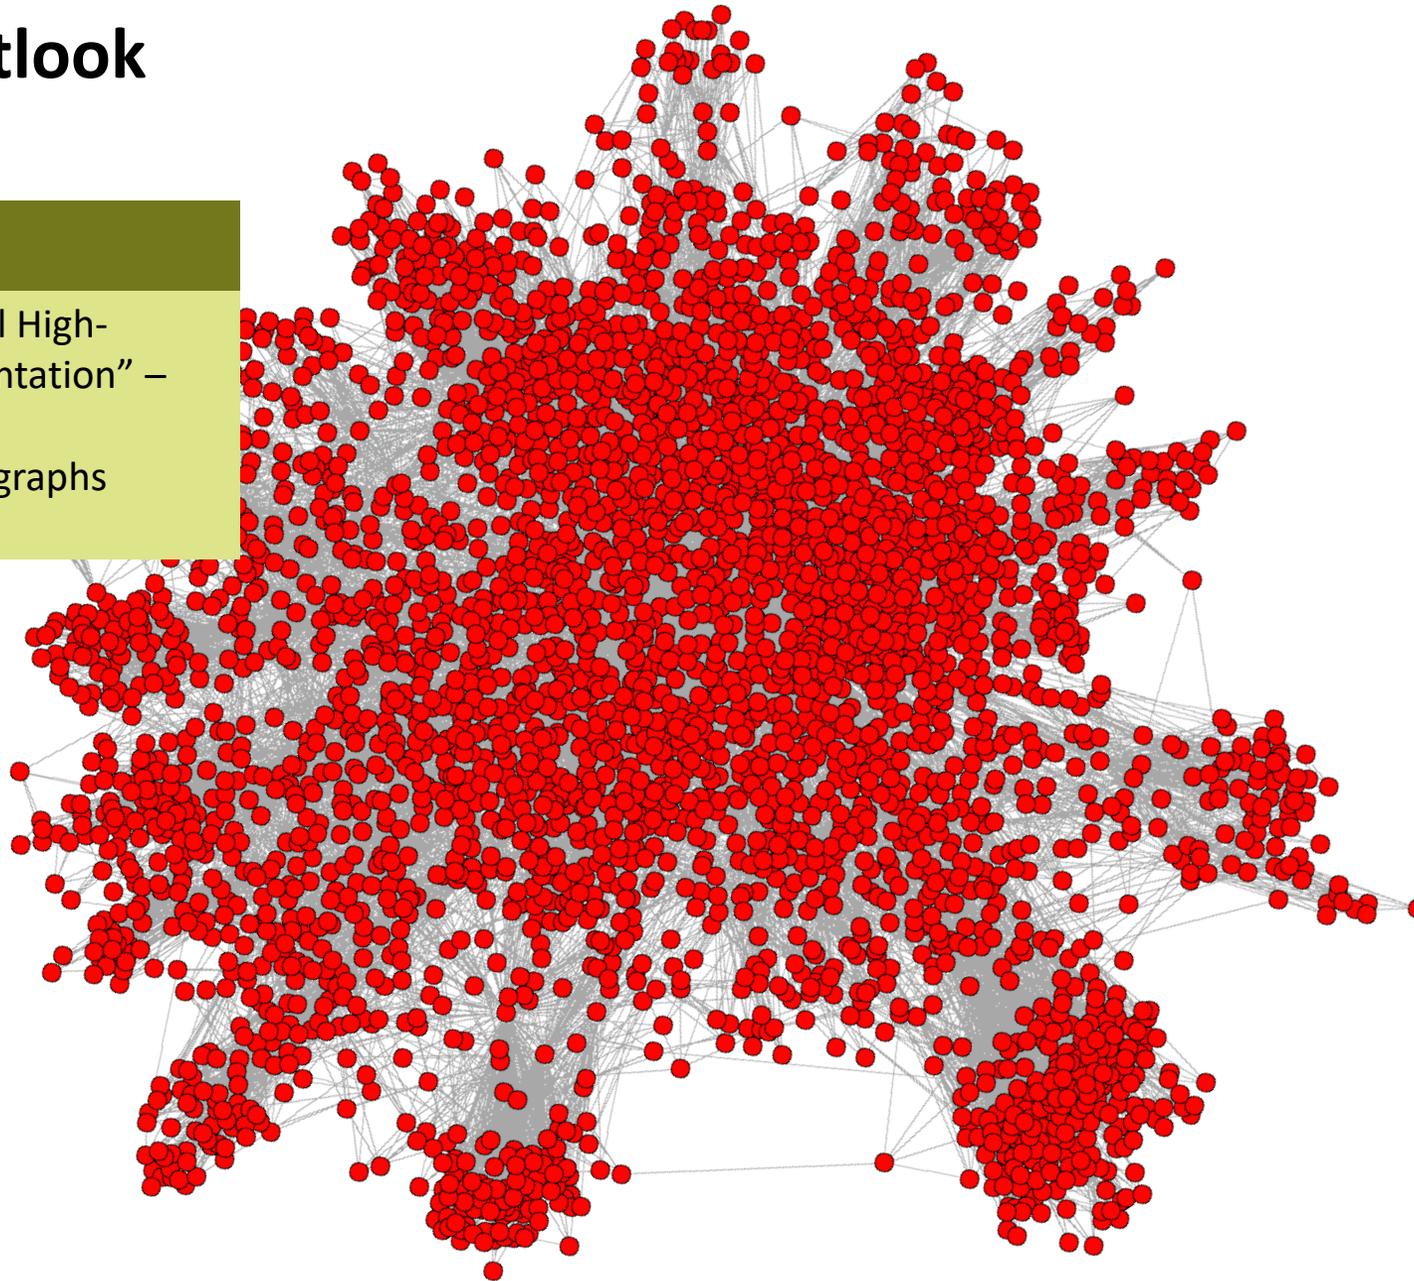
Summary and outlook



Summary and outlook

Log(Graph)

- “Log(Graph): A Near-Optimal High-Performance Graph Representation” – ACM PACT’18
- Minimal storage bounds for graphs during processing



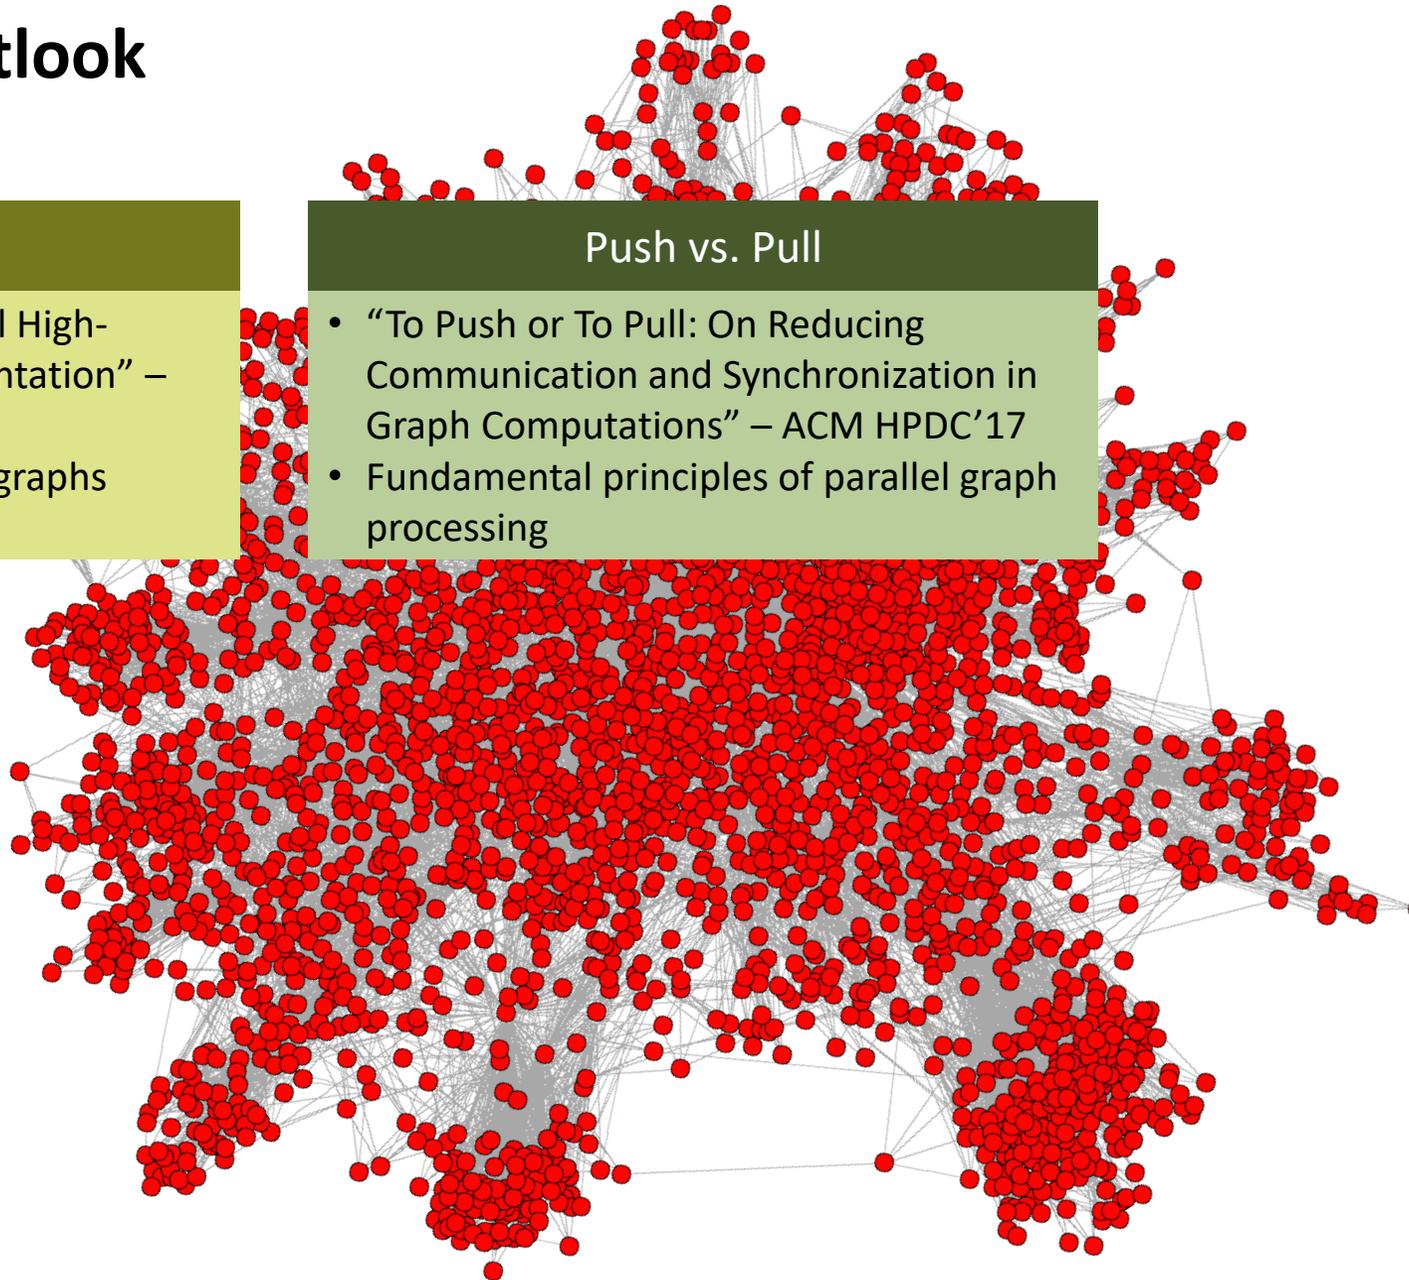
Summary and outlook

Log(Graph)

- “Log(Graph): A Near-Optimal High-Performance Graph Representation” – ACM PACT’18
- Minimal storage bounds for graphs during processing

Push vs. Pull

- “To Push or To Pull: On Reducing Communication and Synchronization in Graph Computations” – ACM HPDC’17
- Fundamental principles of parallel graph processing



Summary and outlook

Log(Graph)

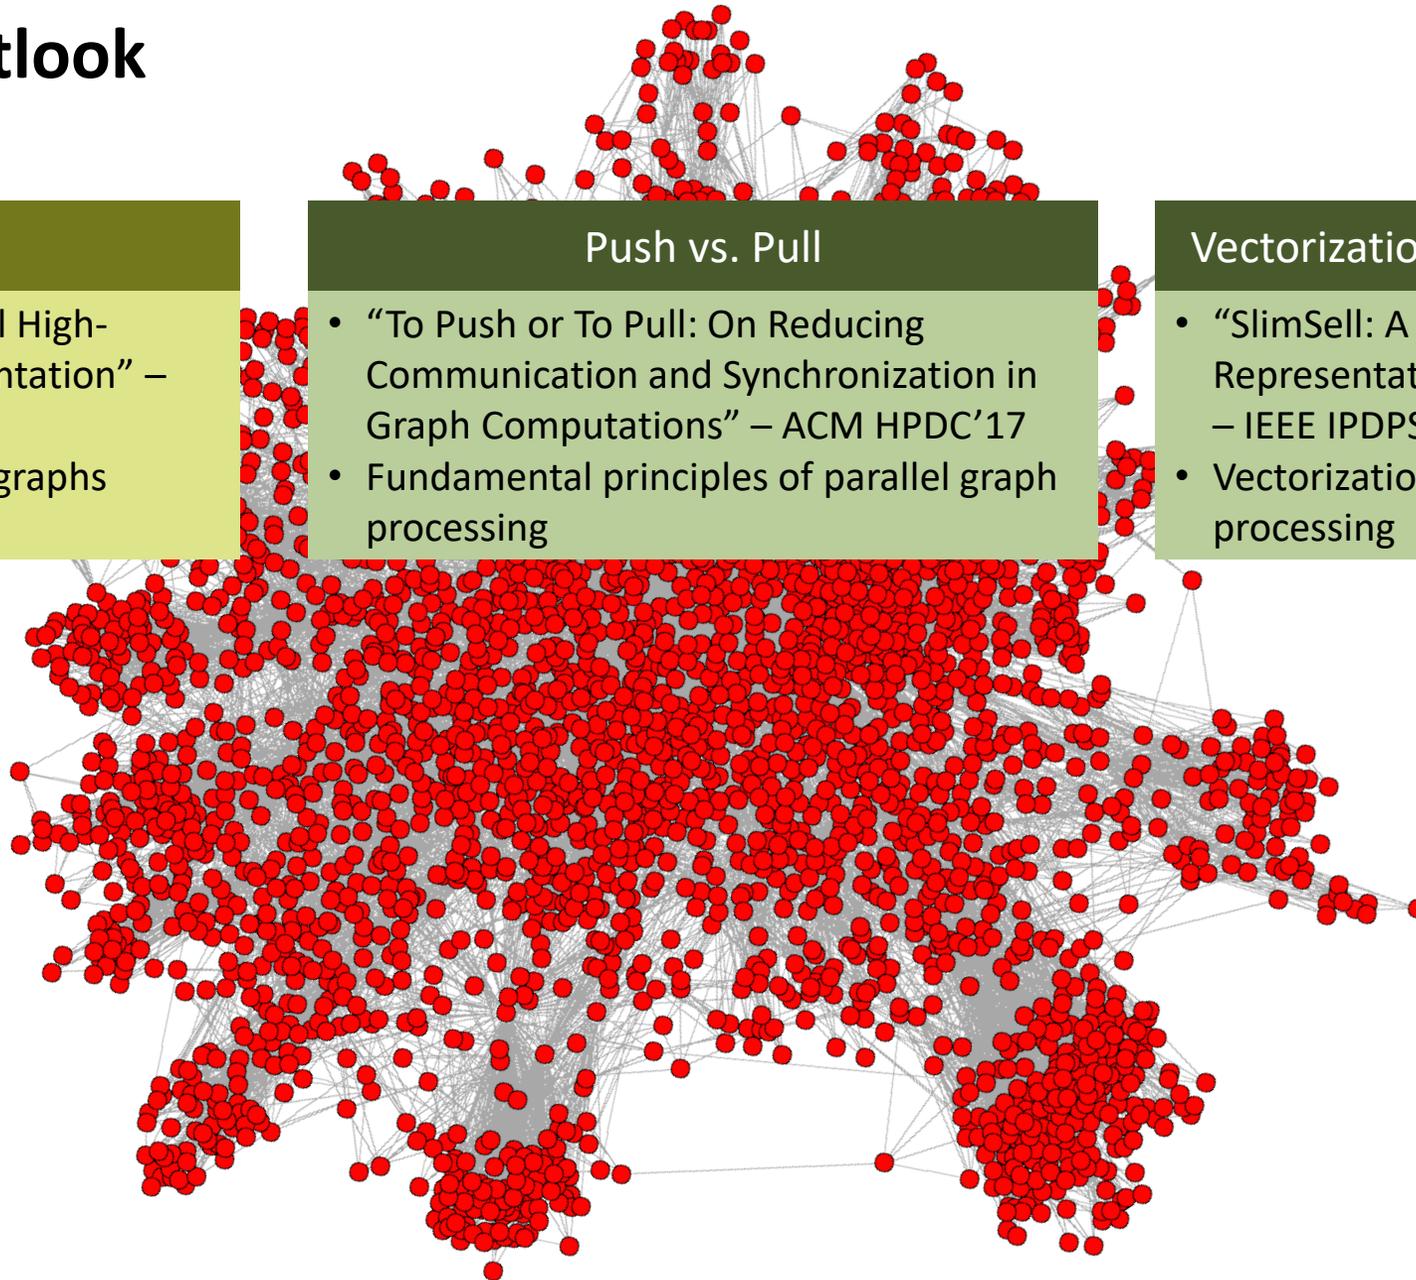
- “Log(Graph): A Near-Optimal High-Performance Graph Representation” – ACM PACT’18
- Minimal storage bounds for graphs during processing

Push vs. Pull

- “To Push or To Pull: On Reducing Communication and Synchronization in Graph Computations” – ACM HPDC’17
- Fundamental principles of parallel graph processing

Vectorization of Graph Computations

- “SlimSell: A Vectorized Graph Representation for Breadth-First Search” – IEEE IPDPS’17
- Vectorization schemes for parallel graph processing



Summary and outlook

Log(Graph)

- “Log(Graph): A Near-Optimal High-Performance Graph Representation” – ACM PACT’18
- Minimal storage bounds for graphs during processing

Push vs. Pull

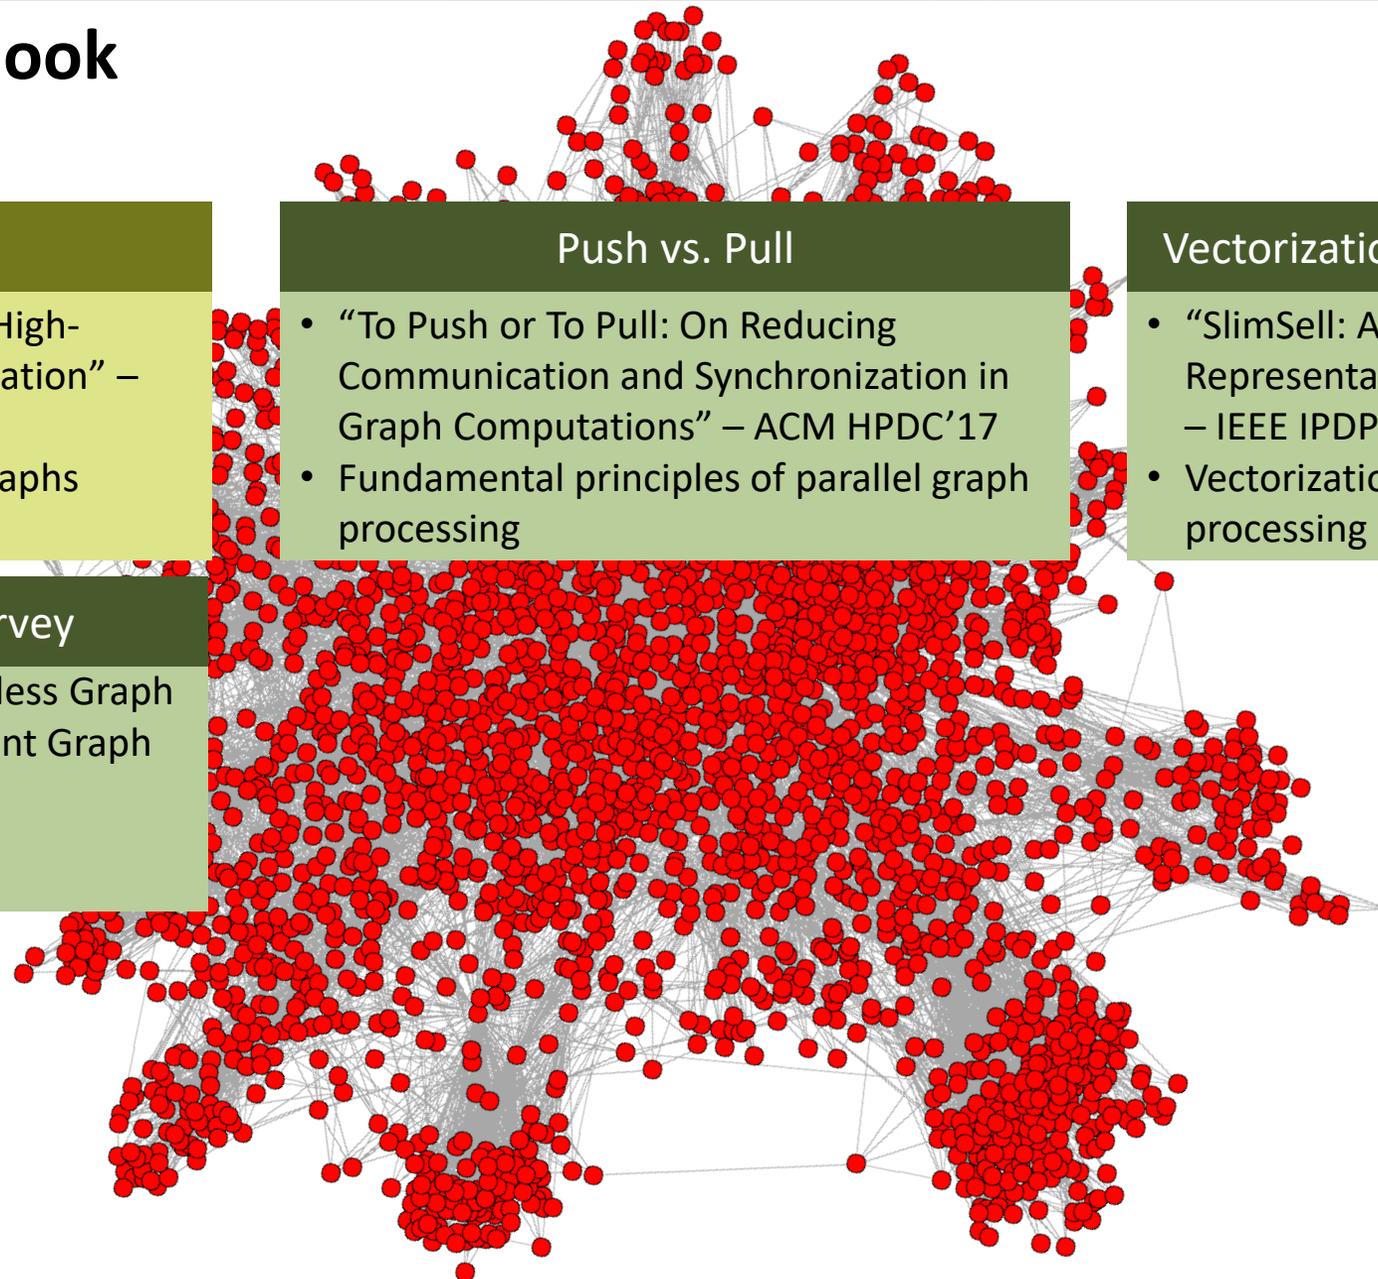
- “To Push or To Pull: On Reducing Communication and Synchronization in Graph Computations” – ACM HPDC’17
- Fundamental principles of parallel graph processing

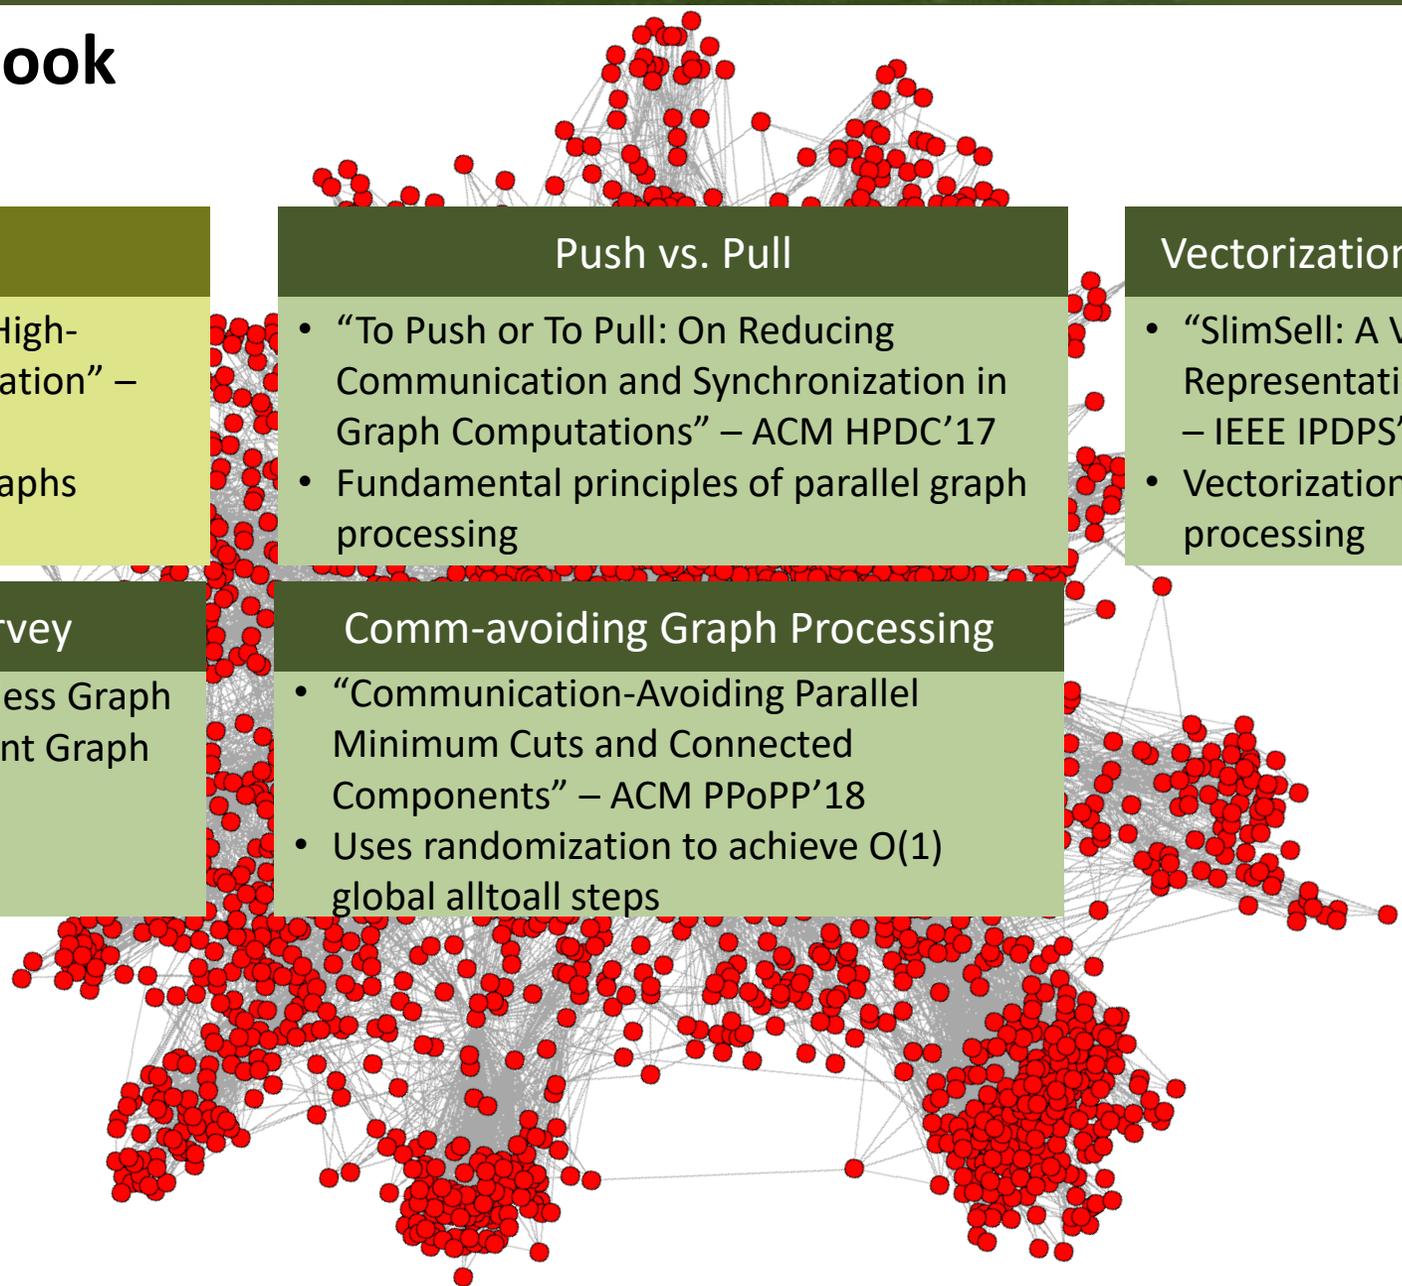
Vectorization of Graph Computations

- “SlimSell: A Vectorized Graph Representation for Breadth-First Search” – IEEE IPDPS’17
- Vectorization schemes for parallel graph processing

Graph Compression Survey

- “Survey and Taxonomy of Lossless Graph Compression and Space-Efficient Graph Representations” – arXiv
- Comprehensive overview
 - 54 pages, 465 references





Summary and outlook

Log(Graph)

- “Log(Graph): A Near-Optimal High-Performance Graph Representation” – ACM PACT’18
- Minimal storage bounds for graphs during processing

Push vs. Pull

- “To Push or To Pull: On Reducing Communication and Synchronization in Graph Computations” – ACM HPDC’17
- Fundamental principles of parallel graph processing

Vectorization of Graph Computations

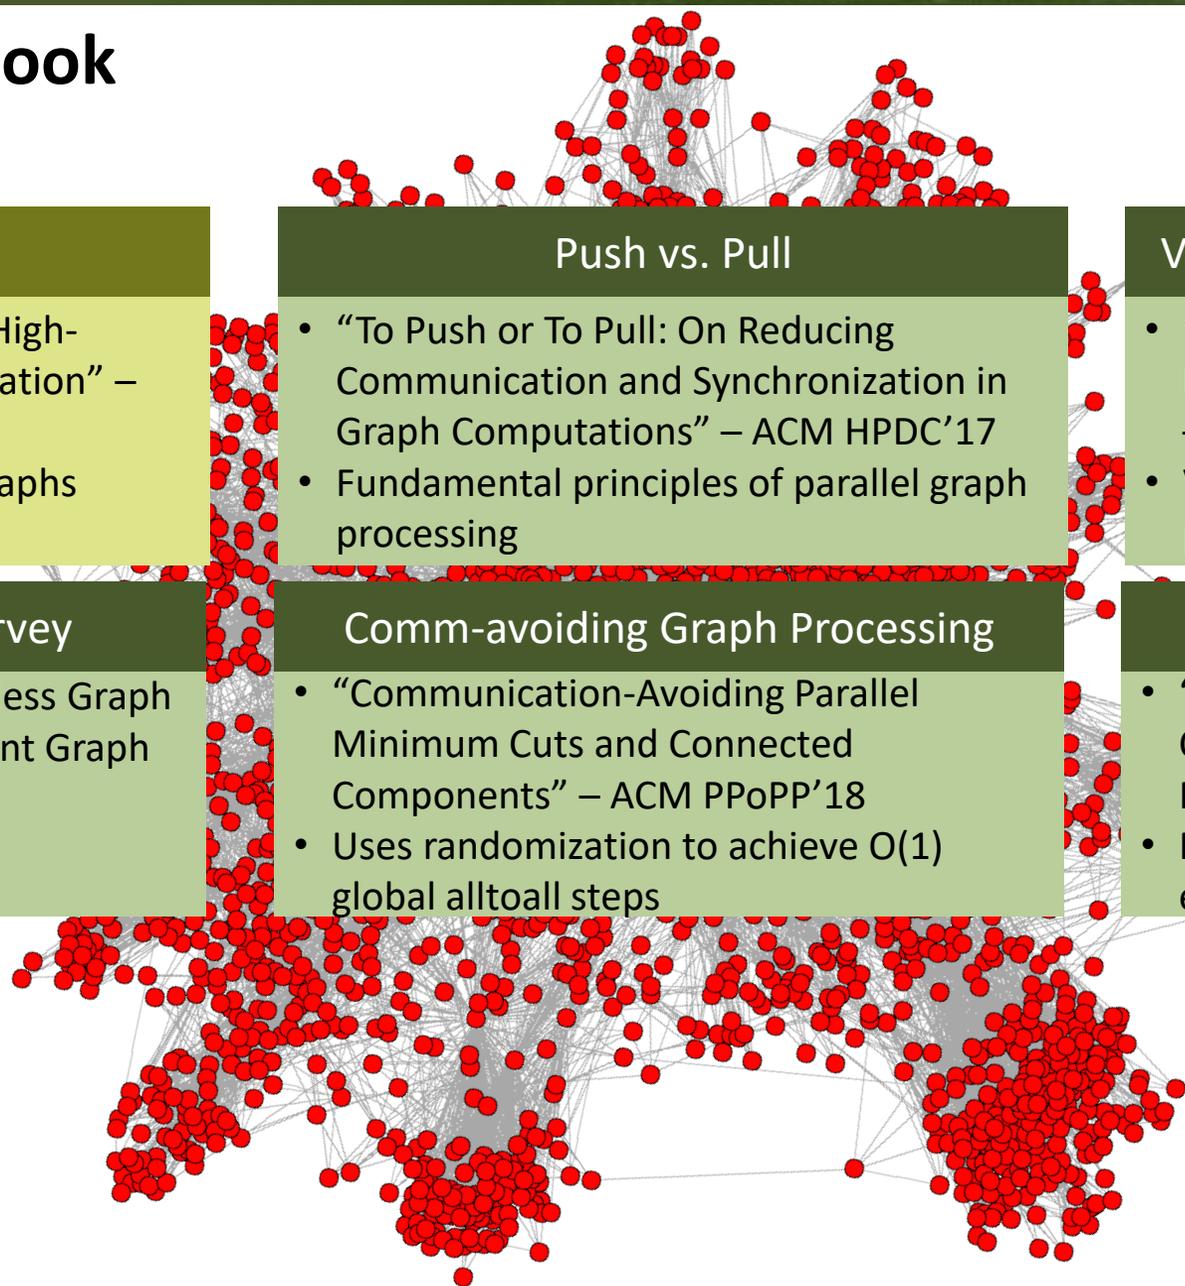
- “SlimSell: A Vectorized Graph Representation for Breadth-First Search” – IEEE IPDPS’17
- Vectorization schemes for parallel graph processing

Graph Compression Survey

- “Survey and Taxonomy of Lossless Graph Compression and Space-Efficient Graph Representations” – arXiv
- Comprehensive overview
 - 54 pages, 465 references

Comm-avoiding Graph Processing

- “Communication-Avoiding Parallel Minimum Cuts and Connected Components” – ACM PPOPP’18
- Uses randomization to achieve $O(1)$ global alltoall steps



Summary and outlook

Log(Graph)

- “Log(Graph): A Near-Optimal High-Performance Graph Representation” – ACM PACT’18
- Minimal storage bounds for graphs during processing

Push vs. Pull

- “To Push or To Pull: On Reducing Communication and Synchronization in Graph Computations” – ACM HPDC’17
- Fundamental principles of parallel graph processing

Vectorization of Graph Computations

- “SlimSell: A Vectorized Graph Representation for Breadth-First Search” – IEEE IPDPS’17
- Vectorization schemes for parallel graph processing

Graph Compression Survey

- “Survey and Taxonomy of Lossless Graph Compression and Space-Efficient Graph Representations” – arXiv
- Comprehensive overview
 - 54 pages, 465 references

Comm-avoiding Graph Processing

- “Communication-Avoiding Parallel Minimum Cuts and Connected Components” – ACM PPOPP’18
- Uses randomization to achieve $O(1)$ global alltoall steps

Algebraic Betweenness Centrality

- “Scaling Betweenness Centrality using Communication-Efficient Sparse Matrix Multiplication” – ACM SC’17
- More on the algebraic view – complex example, large-scale sparse matrices

Summary and outlook

Log(Graph)

- “Log(Graph): A Near-Optimal High-Performance Graph Representation” – ACM PACT’18
- Minimal storage bounds for graphs during processing

Graph Compression Survey

- “Survey and Taxonomy of Lossless Graph Compression and Space-Efficient Graph Representations” – arXiv
- Comprehensive overview
 - 54 pages, 465 references

Streaming Graphs on FPGA

- “Substream-Centric Maximum Matchings on FPGA” – FPGA’19
- New paradigm for parallelizing across substreams
 - Integrates with pipelining in HW/FPGA
 - Blueprint for efficient processing in HW

Push vs. Pull

- “To Push or To Pull: On Reducing Communication and Synchronization in Graph Computations” – ACM HPDC’17
- Fundamental principles of parallel graph processing

Comm-avoiding Graph Processing

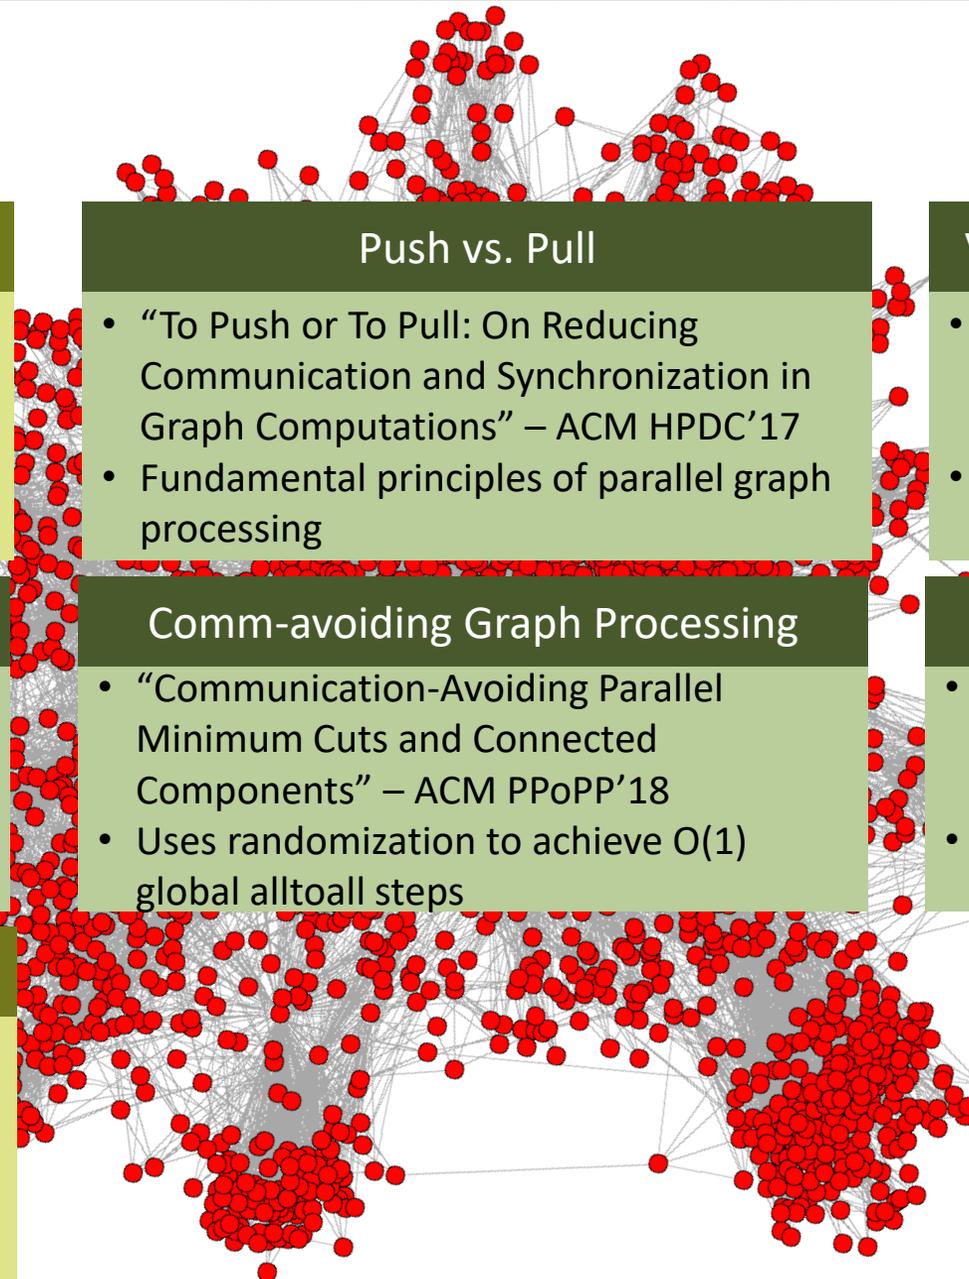
- “Communication-Avoiding Parallel Minimum Cuts and Connected Components” – ACM PPOPP’18
- Uses randomization to achieve $O(1)$ global alltoall steps

Vectorization of Graph Computations

- “SlimSell: A Vectorized Graph Representation for Breadth-First Search” – IEEE IPDPS’17
- Vectorization schemes for parallel graph processing

Algebraic Betweenness Centrality

- “Scaling Betweenness Centrality using Communication-Efficient Sparse Matrix Multiplication” – ACM SC’17
- More on the algebraic view – complex example, large-scale sparse matrices



Summary and outlook

Log(Graph)

- “Log(Graph): A Near-Optimal High-Performance Graph Representation” – ACM PACT’18
- Minimal storage bounds for graphs during processing

Graph Compression Survey

- “Survey and Taxonomy of Lossless Graph Compression and Space-Efficient Graph Representations” – arXiv
- Comprehensive overview
 - 54 pages, 465 references

Streaming Graphs on FPGA

- “Substream-Centric Maximum Matchings on FPGA” – FPGA’19
- New paradigm for parallelizing across substreams
 - Integrates with pipelining in HW/FPGA
 - Blueprint for efficient processing in HW

Push vs. Pull

- “To Push or To Pull: On Reducing Communication and Synchronization in Graph Computations” – ACM HPDC’17
- Fundamental principles of parallel graph processing

Comm-avoiding Graph Processing

- “Communication-Avoiding Parallel Minimum Cuts and Connected Components” – ACM PPOPP’18
- Uses randomization to achieve $O(1)$ global alltoall steps

Streaming Graphs Survey

- “Theory and Practice of Streaming Graph Processing” – soon on arXiv
- Overview of streaming algorithms, approximations, research gaps
 - Way forward for FPGA?

Vectorization of Graph Computations

- “SlimSell: A Vectorized Graph Representation for Breadth-First Search” – IEEE IPDPS’17
- Vectorization schemes for parallel graph processing

Algebraic Betweenness Centrality

- “Scaling Betweenness Centrality using Communication-Efficient Sparse Matrix Multiplication” – ACM SC’17
- More on the algebraic view – complex example, large-scale sparse matrices

Summary and outlook

Log(Graph)

- “Log(Graph): A Near-Optimal High-Performance Graph Representation” – ACM PACT’18
- Minimal storage bounds for graphs during processing

Graph Compression Survey

- “Survey and Taxonomy of Lossless Graph Compression and Space-Efficient Graph Representations” – arXiv
- Comprehensive overview
 - 54 pages, 465 references

Streaming Graphs on FPGA

- “Substream-Centric Maximum Matchings on FPGA” – FPGA’19
- New paradigm for parallelizing across substreams
 - Integrates with pipelining in HW/FPGA
 - Blueprint for efficient processing in HW

Push vs. Pull

- “To Push or To Pull: On Reducing Communication and Synchronization in Graph Computations” – ACM HPDC’17
- Fundamental principles of parallel graph processing

Comm-avoiding Graph Processing

- “Communication-Avoiding Parallel Minimum Cuts and Connected Components” – ACM PPOPP’18
- Uses randomization to achieve $O(1)$ global alltoall steps

Streaming Graphs Survey

- “Theory and Practice of Streaming Graph Processing” – soon on arXiv
- Overview of streaming algorithms, approximations, research gaps
 - Way forward for FPGA?

Vectorization of Graph Computations

- “SlimSell: A Vectorized Graph Representation for Breadth-First Search” – IEEE IPDPS’17
- Vectorization schemes for parallel graph processing

Algebraic Betweenness Centrality

- “Scaling Betweenness Centrality using Communication-Efficient Sparse Matrix Multiplication” – ACM SC’17
- More on the algebraic view – complex example, large-scale sparse matrices

Graphs on FPGA Survey

- “Graph Processing on FPGAs: Taxonomy, Survey, Challenges” – arXiv
- Relatively young field of graph processing on FPGAs / in hardware
 - Identify research opportunities

Summary and outlook

Log(Graph)

- “Log(Graph): A Near-Optimal High-Performance Graph Representation” – ACM PACT’18
- Minimal storage bounds for graphs during processing

Graph Compression Survey

- “Survey and Taxonomy of Lossless Graph Compression and Space-Efficient Graph Representations” – arXiv
- Comprehensive overview
 - 54 pages, 465 references

Streaming Graphs on FPGA

- “Substream-Centric Maximum Matchings on FPGA” – FPGA’19
- New paradigm for parallelizing across substreams
 - Integrates with pipelining in HW/FPGA
 - Blueprint for efficient processing in HW

Push vs. Pull

- “To Push or To Pull: On Reducing Communication and Synchronization in Graph Computations” – ACM HPDC’17
- Fundamental principles of parallel graph processing

Comm-avoiding Graph Processing

- “Communication-Avoiding Parallel Minimum Cuts and Connected Components” – ACM PPOPP’18
- Uses randomization to achieve $O(1)$ global alltoall steps

Streaming Graphs Survey

- “Theory and Practice of Streaming Graph Processing” – soon on arXiv
- Overview of streaming algorithms, approximations, research gaps
 - Way forward for FPGA?

- “Accelerating Irregular Computations with Hardware Transactional Memory and Active Messages” – ACM HPDC’15

Vectorization of Graph Computations

- “SlimSell: A Vectorized Graph Representation for Breadth-First Search” – IEEE IPDPS’17
- Vectorization schemes for parallel graph processing

Algebraic Betweenness Centrality

- “Scaling Betweenness Centrality using Communication-Efficient Sparse Matrix Multiplication” – ACM SC’17
- More on the algebraic view – complex example, large-scale sparse matrices

Graphs on FPGA Survey

- “Graph Processing on FPGAs: Taxonomy, Survey, Challenges” – arXiv
- Relatively young field of graph processing on FPGAs / in hardware
 - Identify research opportunities

Summary and outlook

Log(Graph)

- “Log(Graph): A Near-Optimal High-Performance Graph Representation” – ACM PACT’18
- Minimal storage bounds for graphs during processing

Graph Compression Survey

- “Survey and Taxonomy of Lossless Graph Compression and Space-Efficient Graph Representations” – arXiv
- Comprehensive overview
 - 54 pages, 465 references

Streaming Graphs on FPGA

- “Substream-Centric Maximum Matchings on FPGA” – FPGA’19
- New paradigm for parallelizing across substreams
 - Integrates with pipelining in HW/FPGA
 - Blueprint for efficient processing in HW

Slim Graph

- “Slim Graph: Practical Lossy Graph Compression for Approximate Graph Processing, Storage, and Analytics” – ACM/IEEE SC’19

Push vs. Pull

- “To Push or To Pull: On Reducing Communication and Synchronization in Graph Computations” – ACM HPDC’17
- Fundamental principles of parallel graph processing

Comm-avoiding Graph Processing

- “Communication-Avoiding Parallel Minimum Cuts and Connected Components” – ACM PPOPP’18
- Uses randomization to achieve $O(1)$ global alltoall steps

Streaming Graphs Survey

- “Theory and Practice of Streaming Graph Processing” – soon on arXiv
- Overview of streaming algorithms, approximations, research gaps
 - Way forward for FPGA?

Hardware Transactions for Graphs

- “Accelerating Irregular Computations with Hardware Transactional Memory and Active Messages” – ACM HPDC’15

Vectorization of Graph Computations

- “SlimSell: A Vectorized Graph Representation for Breadth-First Search” – IEEE IPDPS’17
- Vectorization schemes for parallel graph processing

Algebraic Betweenness Centrality

- “Scaling Betweenness Centrality using Communication-Efficient Sparse Matrix Multiplication” – ACM SC’17
- More on the algebraic view – complex example, large-scale sparse matrices

Graphs on FPGA Survey

- “Graph Processing on FPGAs: Taxonomy, Survey, Challenges” – arXiv
- Relatively young field of graph processing on FPGAs / in hardware
 - Identify research opportunities

Summary and outlook

Log(Graph)

- “Log(Graph): A Near-Optimal High-Performance Graph Representation” – ACM PACT’18
- Minimal storage bounds for graphs during processing

Graph Compression Survey

- “Survey and Taxonomy of Lossless Graph Compression and Space-Efficient Graph Representations” – arXiv
- Comprehensive overview
 - 54 pages, 465 references

Streaming Graphs on FPGA

- “Substream-Centric Maximum Matchings on FPGA” – FPGA’19
- New paradigm for parallelizing across substreams
 - Integrates with pipelining in HW/FPGA
 - Blueprint for efficient processing in HW

Slim Graph

- “Slim Graph: Practical Lossy Graph Compression for Approximate Graph Processing, Storage, and Analytics” – ACM/IEEE SC’19

Push vs. Pull

- “To Push or To Pull: On Reducing Communication and Synchronization in Graph Computations” – ACM HPDC’17
- Fundamental principles of parallel graph processing

Comm-avoiding Graph Processing

- “Communication-Avoiding Parallel Minimum Cuts and Connected Components” – ACM PPOPP’18
- Uses randomization to achieve $O(1)$ global alltoall steps

Streaming Graphs Survey

- “Theory and Practice of Streaming Graph Processing” – soon on arXiv
- Overview of streaming algorithms, approximations, research gaps
 - Way forward for FPGA?

Hardware Transactions for Graphs

- “Accelerating Graph Computations with Hardware” – **...and others 😊** and Active

Vectorization of Graph Computations

- “SlimSell: A Vectorized Graph Representation for Breadth-First Search” – IEEE IPDPS’17
- Vectorization schemes for parallel graph processing

Algebraic Betweenness Centrality

- “Scaling Betweenness Centrality using Communication-Efficient Sparse Matrix Multiplication” – ACM SC’17
- More on the algebraic view – complex example, large-scale sparse matrices

Graphs on FPGA Survey

- “Graph Processing on FPGAs: Taxonomy, Survey, Challenges” – arXiv
- Relatively young field of graph processing on FPGAs / in hardware
 - Identify research opportunities

Summary and outlook

Log(Graph)

- “Log(Graph): A Near-Optimal High-Performance Graph Representation” – ACM PACT’18
- Minimal storage bounds for graphs during processing

Graph Compression Survey

- “Survey and Taxonomy of Lossless Graph Compression and Space-Efficient Graph Representations” – arXiv
- Comprehensive overview
 - 54 pages, 465 references

Streaming Graphs on FPGA

- “Substream-Centric Maximum Matchings on FPGA” – FPGA’19
- New paradigm for parallelizing across substreams
 - Integrates with pipelining in HW/FPGA
 - Blueprint for efficient processing in HW

Slim Graph

- “Slim Graph: Practical Lossy Graph Compression for Approximate Graph Processing, Storage, and Analytics” – ACM/IEEE SC’19

Push vs. Pull

- “To Push or To Pull: On Reducing

Hardware Transactions for Graphs

- “Accelerating Graph Computations with Hardware Transactions and Active ...and others 😊

Vectorization of Graph Computations

- “SlimSell: A Vectorized Graph Representation for Breadth-First Search” – IPDPS’17
- Vectorization schemes for parallel graph processing

Algebraic Betweenness Centrality

- “Computing Betweenness Centrality using Communication-Efficient Sparse Matrix Multiplication” – ACM SC’17
- Focus on the algebraic view – complex, large-scale sparse matrices

Streaming Graphs Survey

- “Theory and Practice of Streaming Graph Processing” – soon on arXiv
- Overview of streaming algorithms, approximations, research gaps
 - Way forward for FPGA?

Graphs on FPGA Survey

- “Graph Processing on FPGAs: Taxonomy, Survey, Challenges” – arXiv
- Relatively young field of graph processing on FPGAs / in hardware
 - Identify research opportunities

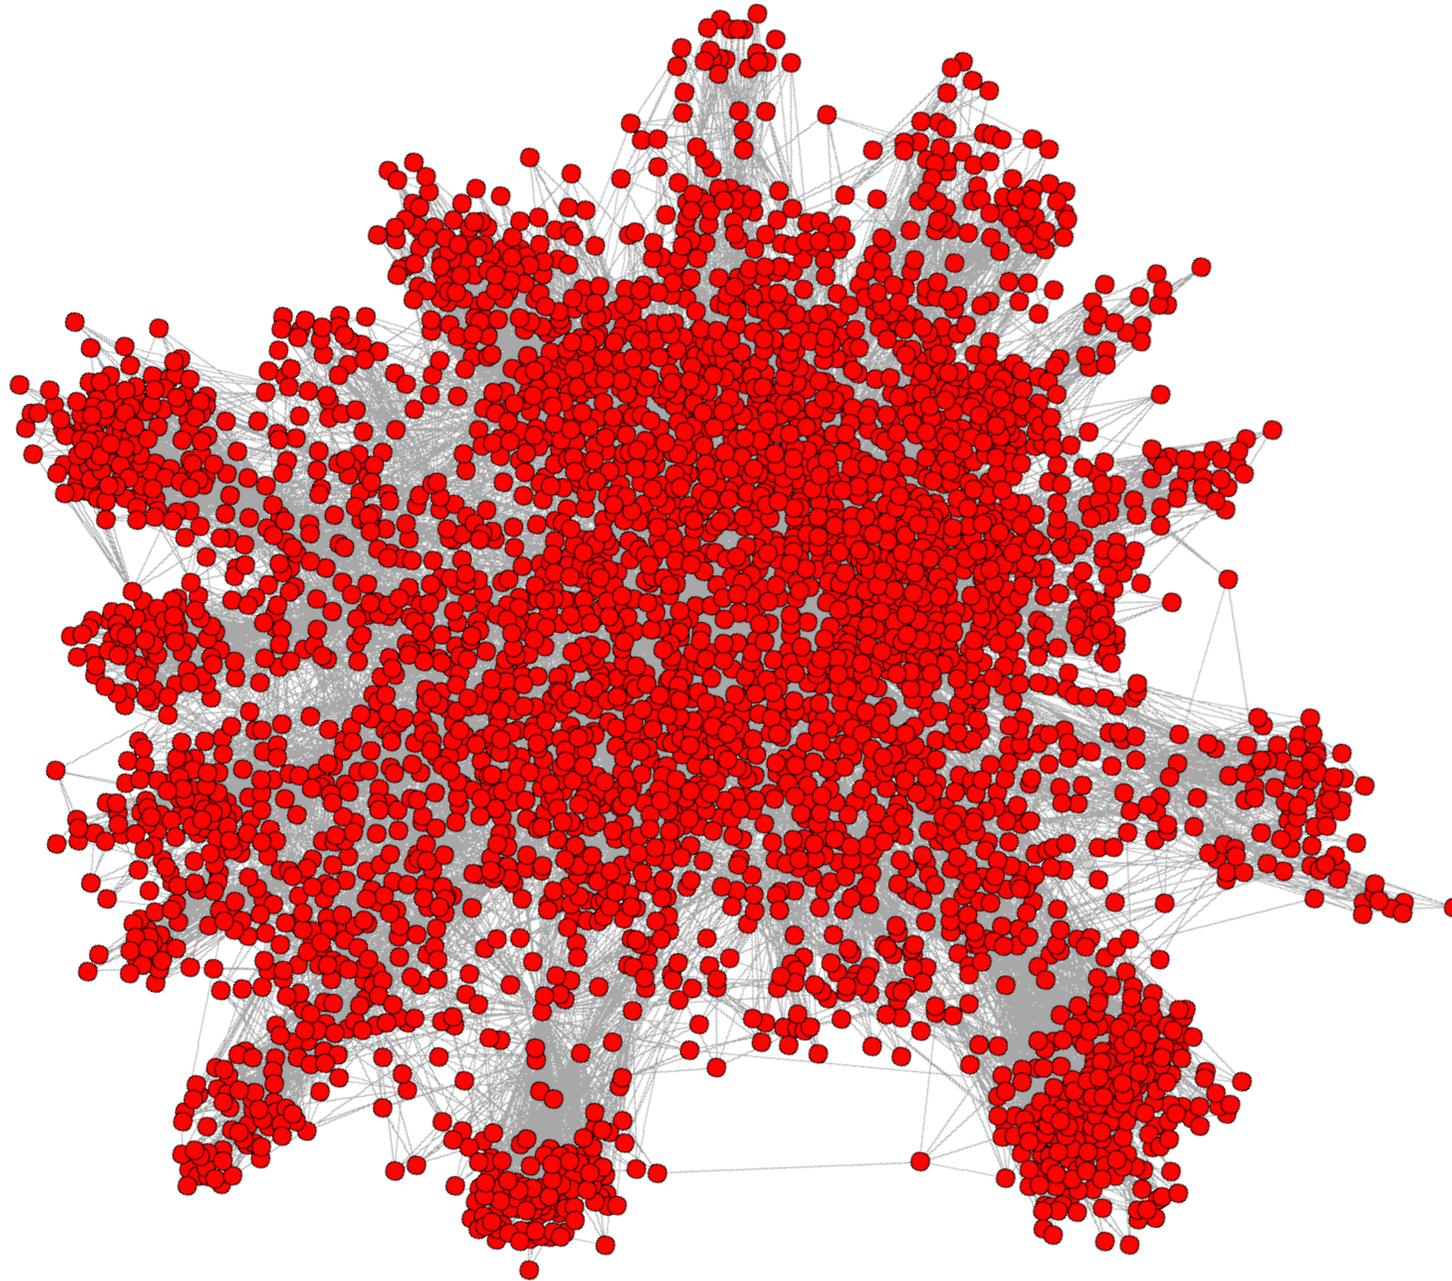
A big challenges ahead: develop a framework to integrate all techniques!

SPCL’s approach: *stateful dataflow graphs*

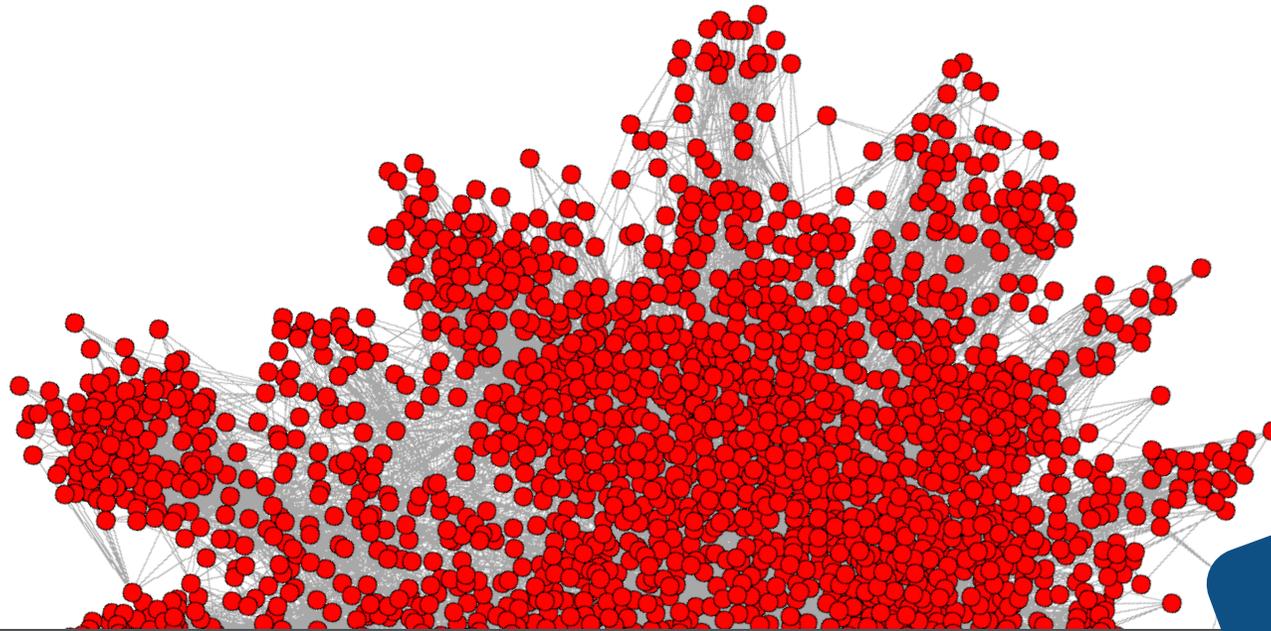
We’re always hiring excellent PhD students and postdocs at SPCL/ETH at spcl.inf.ethz.ch/Jobs

Backup

Problems!



Problems!



Synchronization-heavy

To Push or To Pull: On Reducing Communication and Synchronization in Graph Computations

Maciej Besta¹, Michał Podstawski^{2 3}, Linus Groner¹, Edgar Solomonik⁴, Torsten Hoefler¹

¹ Department of Computer Science, ETH Zurich; ² Perform Group Katowice; ³ Katowice Institute of Information Technologies;

⁴ Department of Computer Science, University of Illinois at Urbana-Champaign

maciej.best@inf.ethz.ch, michal.podstawski@performgroup.com, gronerl@student.ethz.ch, solomon2@illinois.edu, htor@inf.ethz.ch

ABSTRACT

We reduce the cost of communication and synchronization in graph processing by analyzing the fastest way to process graphs: pushing the updates to a shared state or pulling the updates to a private state. We investigate the applicability of this push-pull dichotomy to

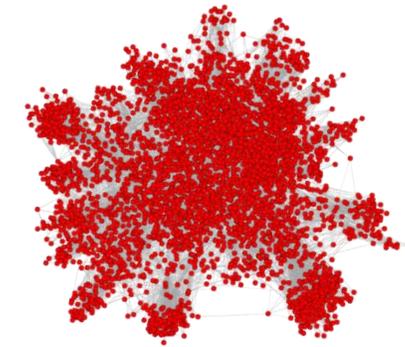
can either push v 's rank to update v 's neighbors, or it can pull the ranks of v 's neighbors to update v [52]. Despite many differences between PR and BFS (e.g., PR is not a traversal), PR can similarly be viewed in the push-pull dichotomy.

This notion sparks various questions. Can pushing and pulling



PAGERANK

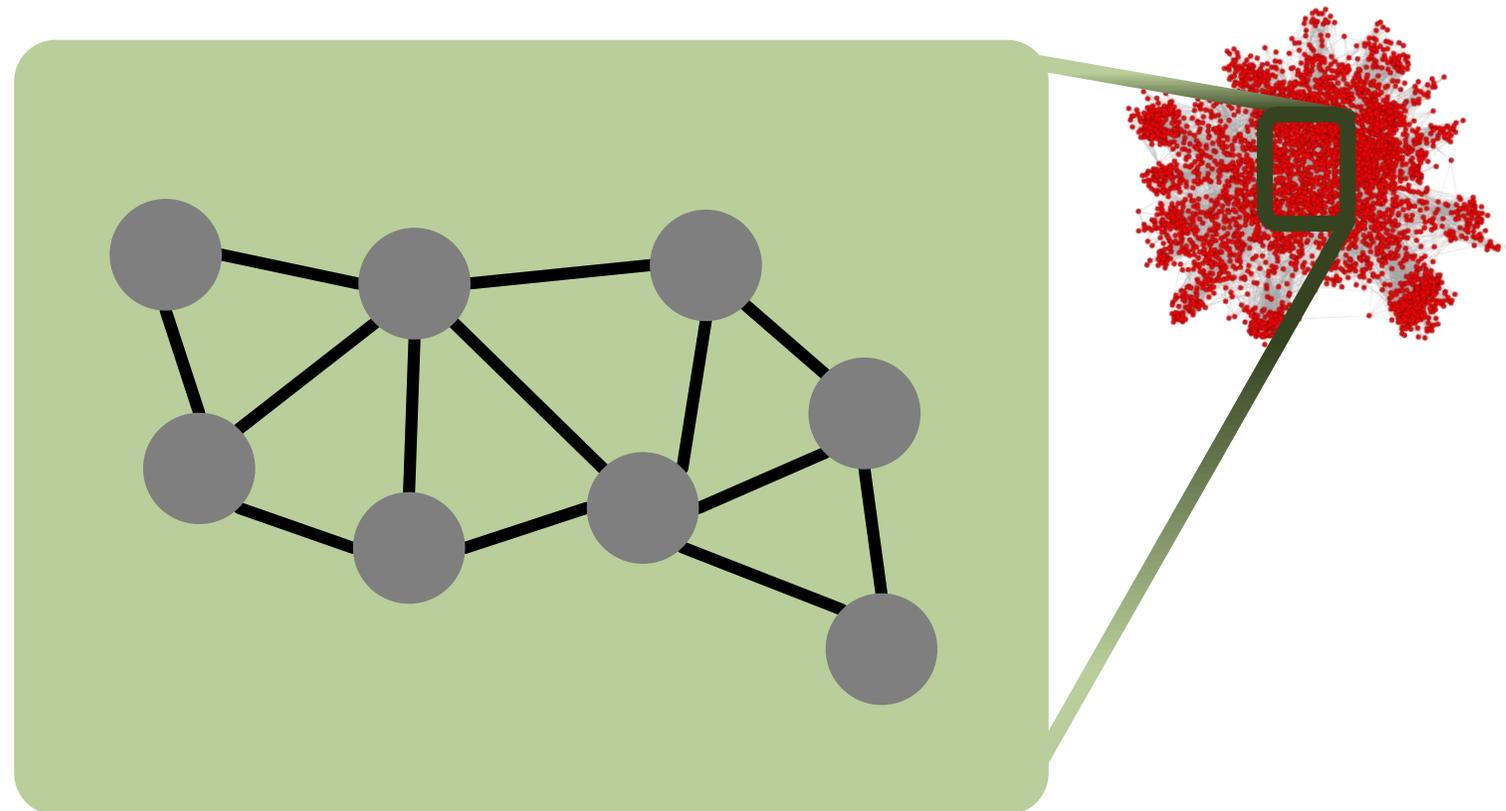
PAGERANK



PAGERANK

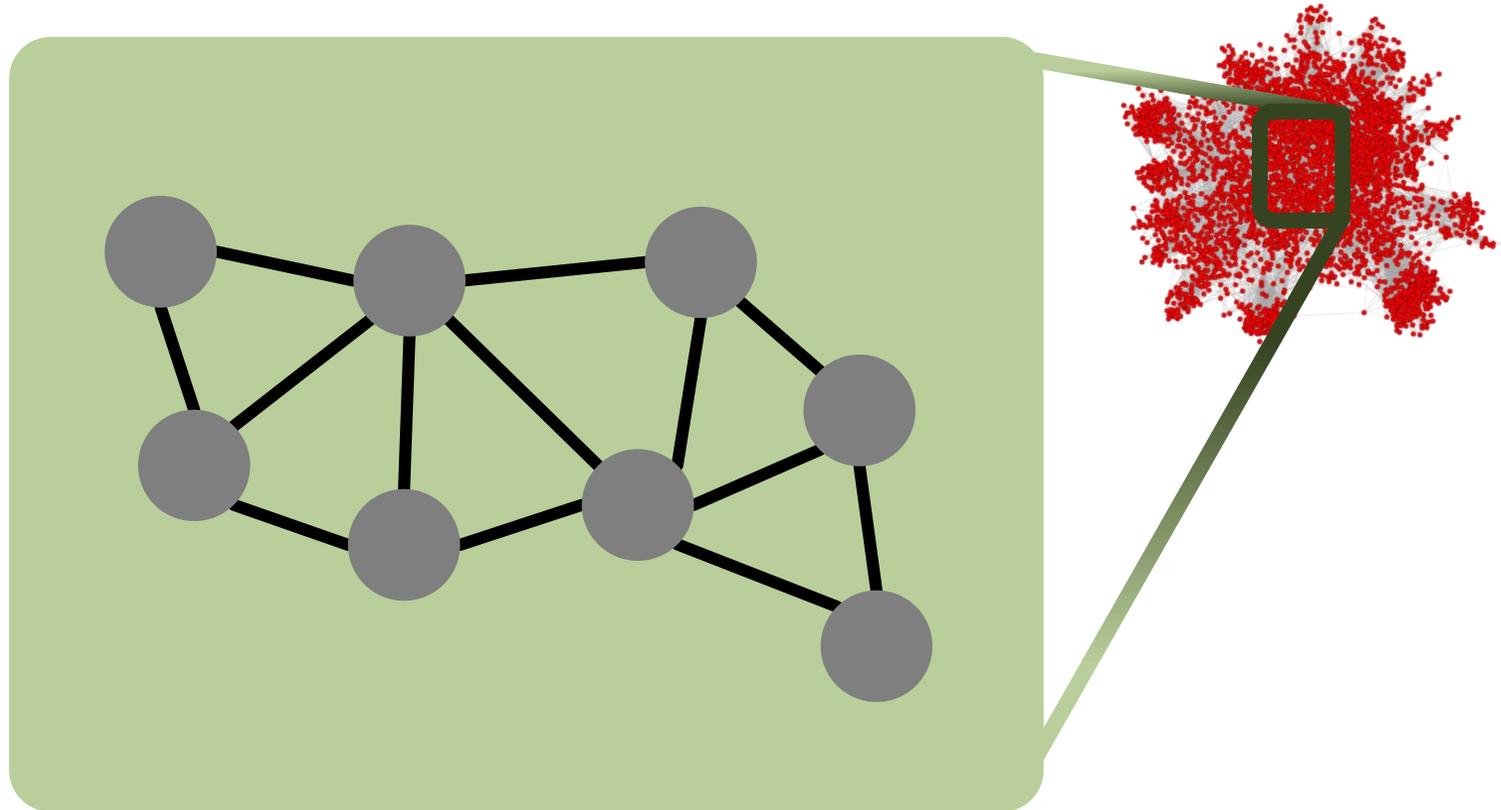


PAGERANK



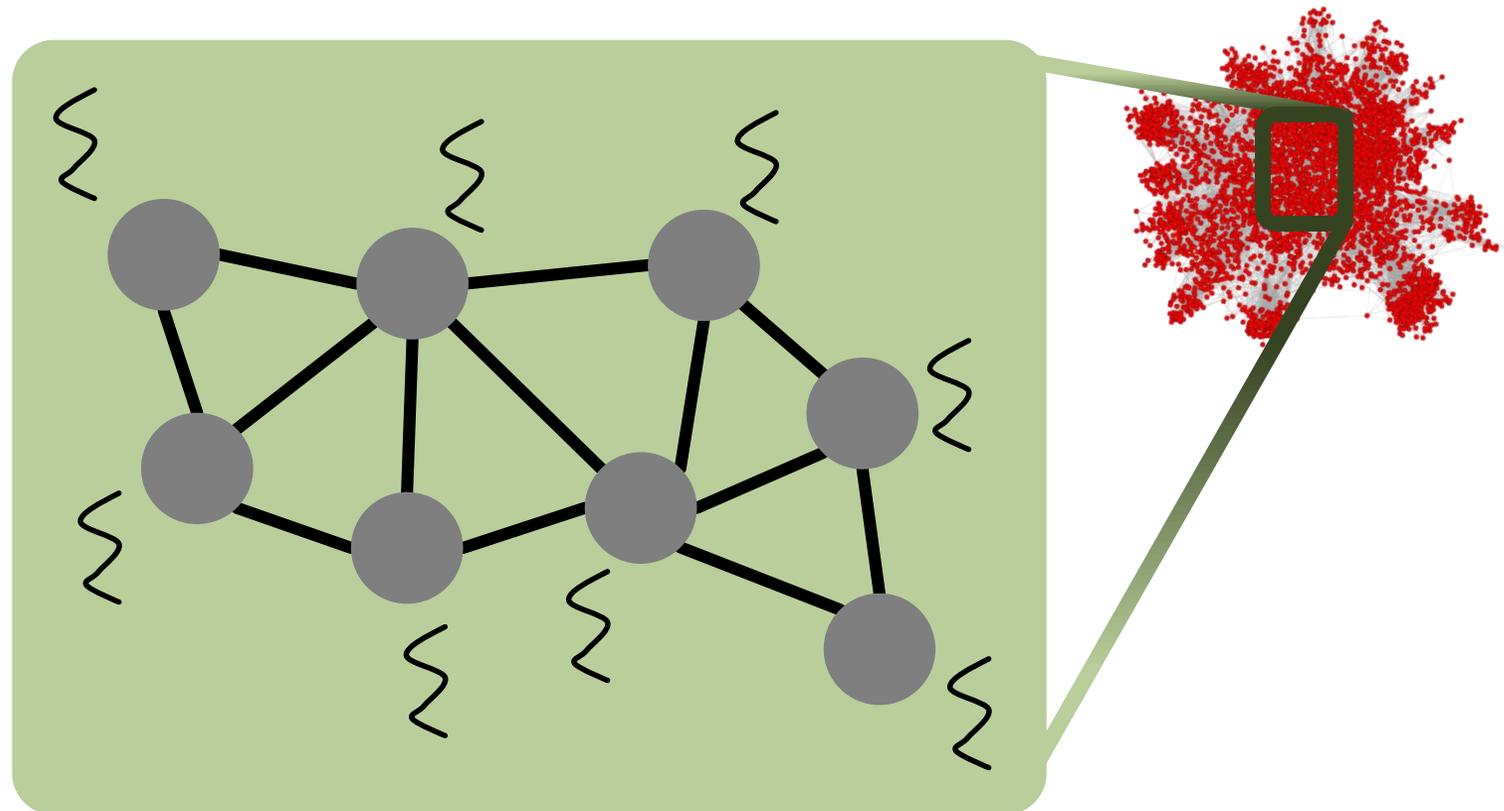
PAGERANK

P threads are used



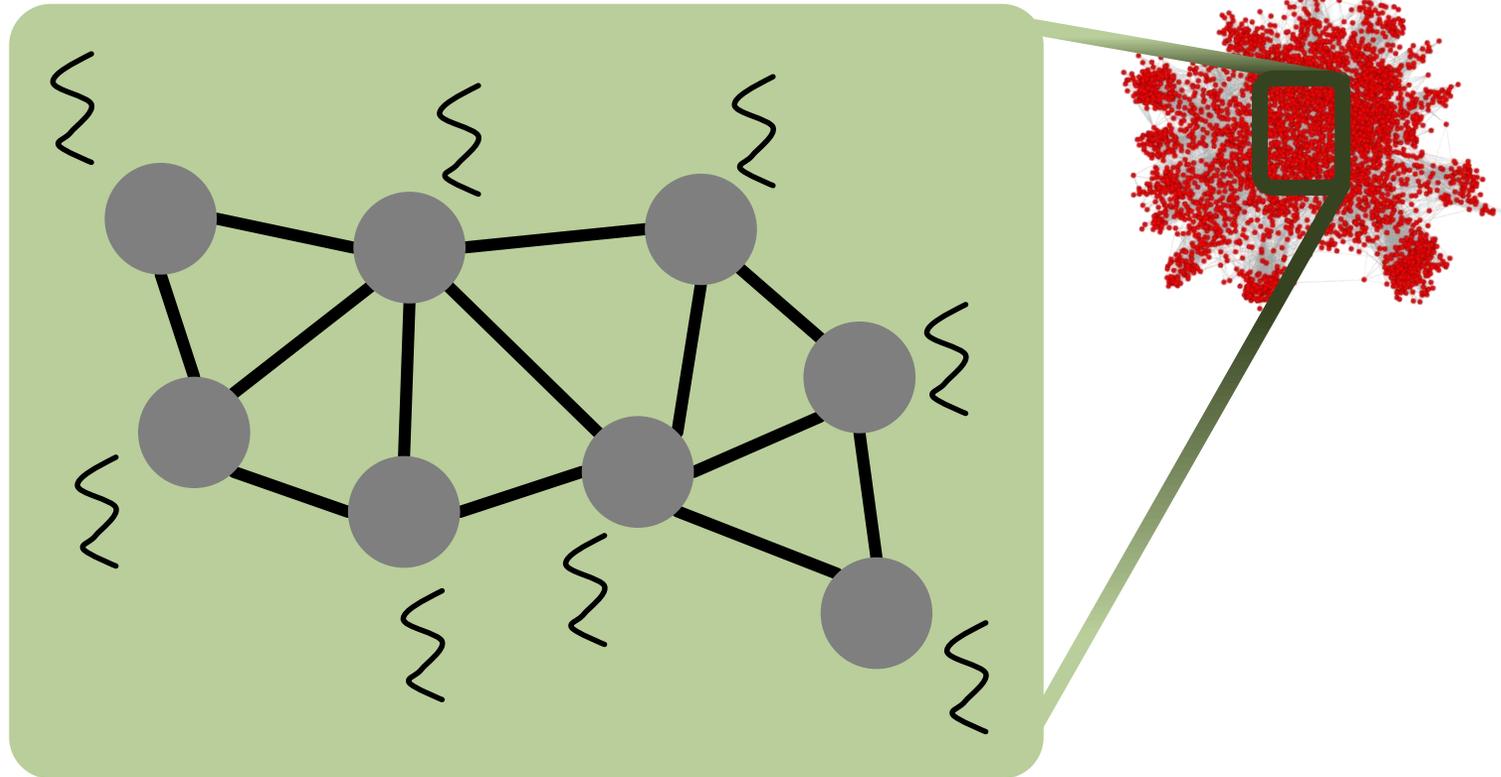
PAGERANK

P threads are used



PAGERANK

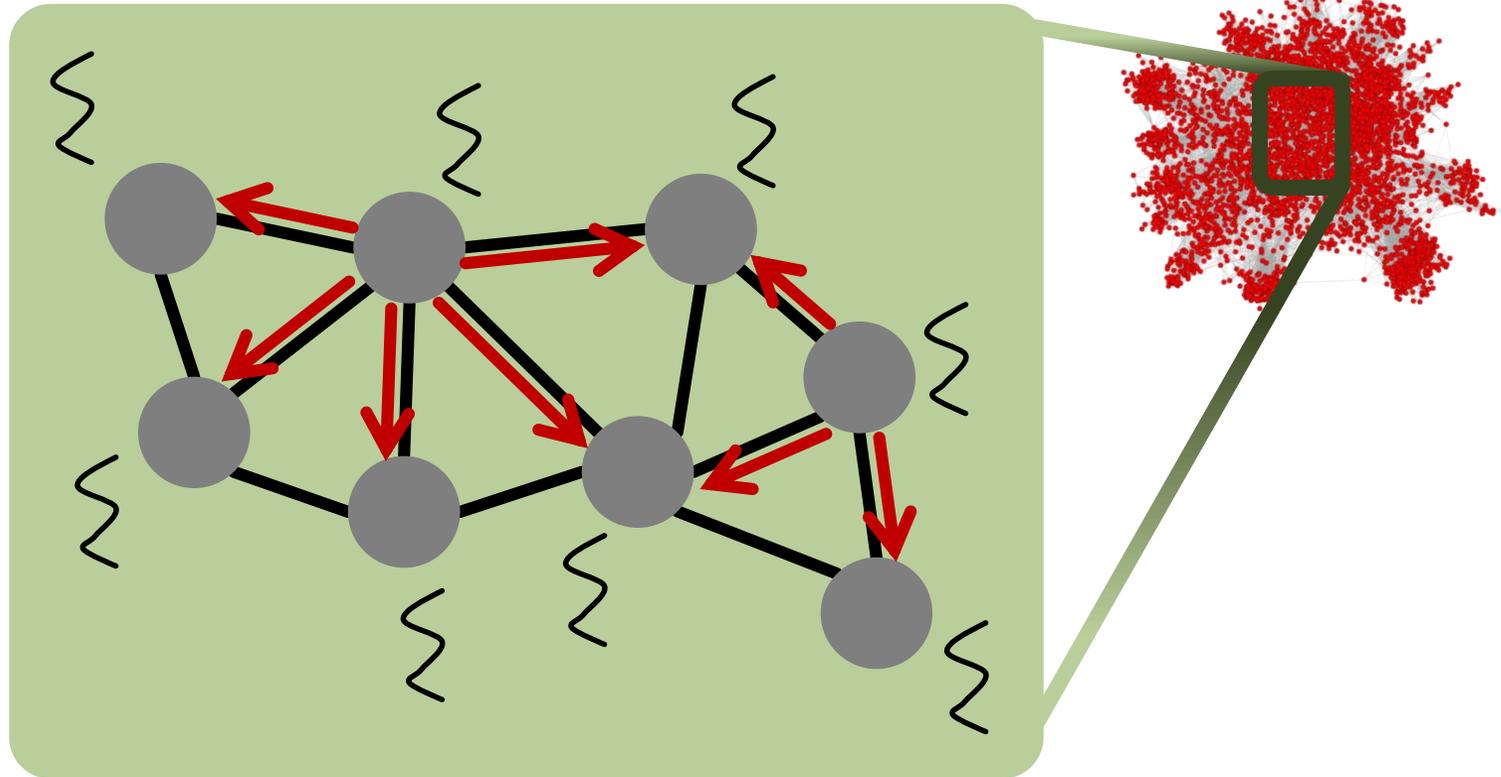
P threads are used



Pushing

PAGERANK

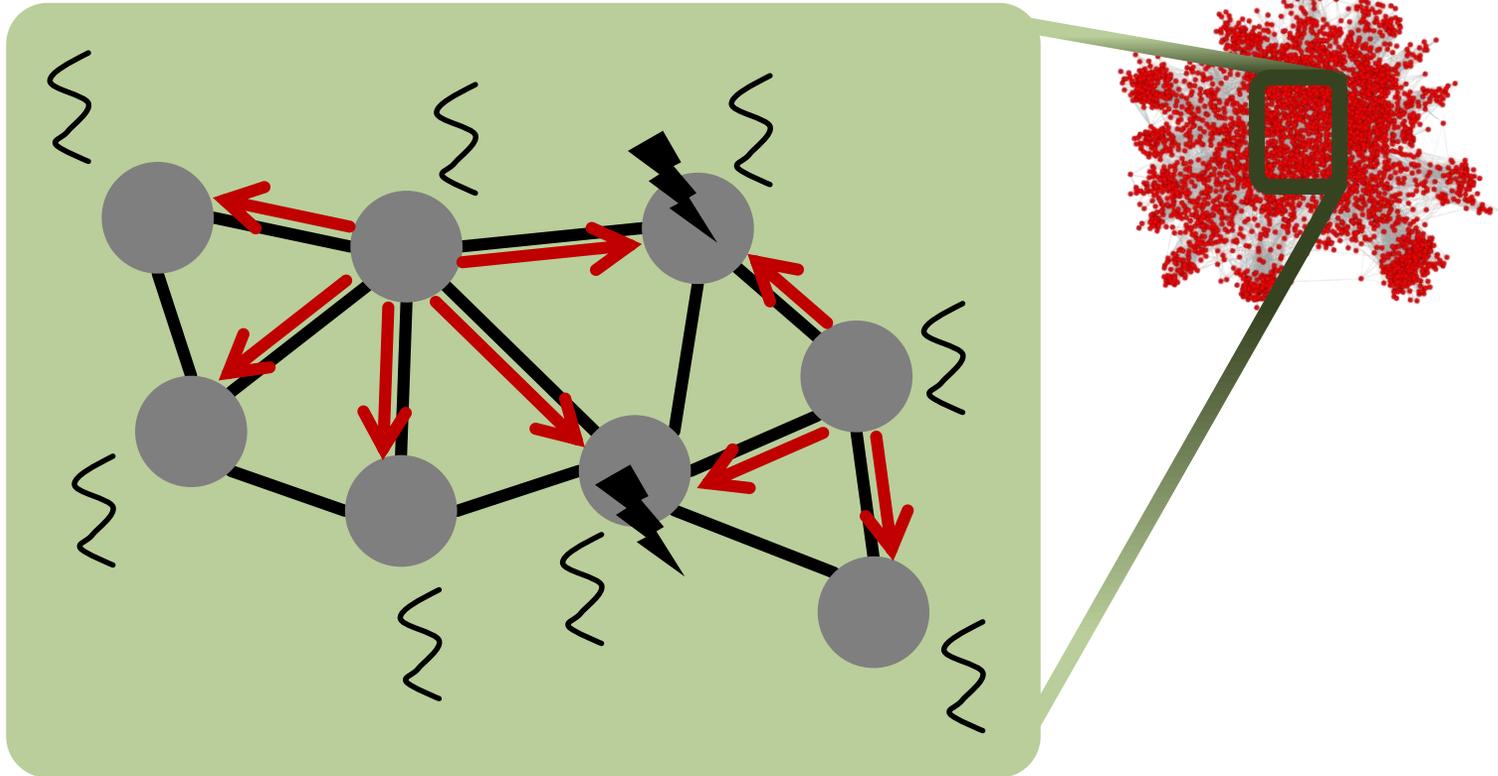
P threads are used



Pushing

PAGERANK

P threads are used

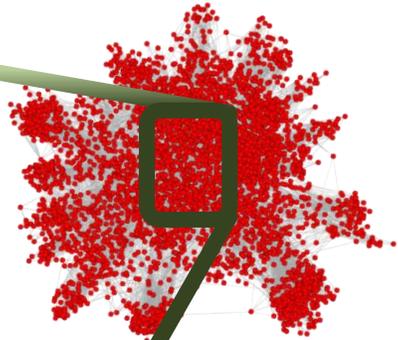
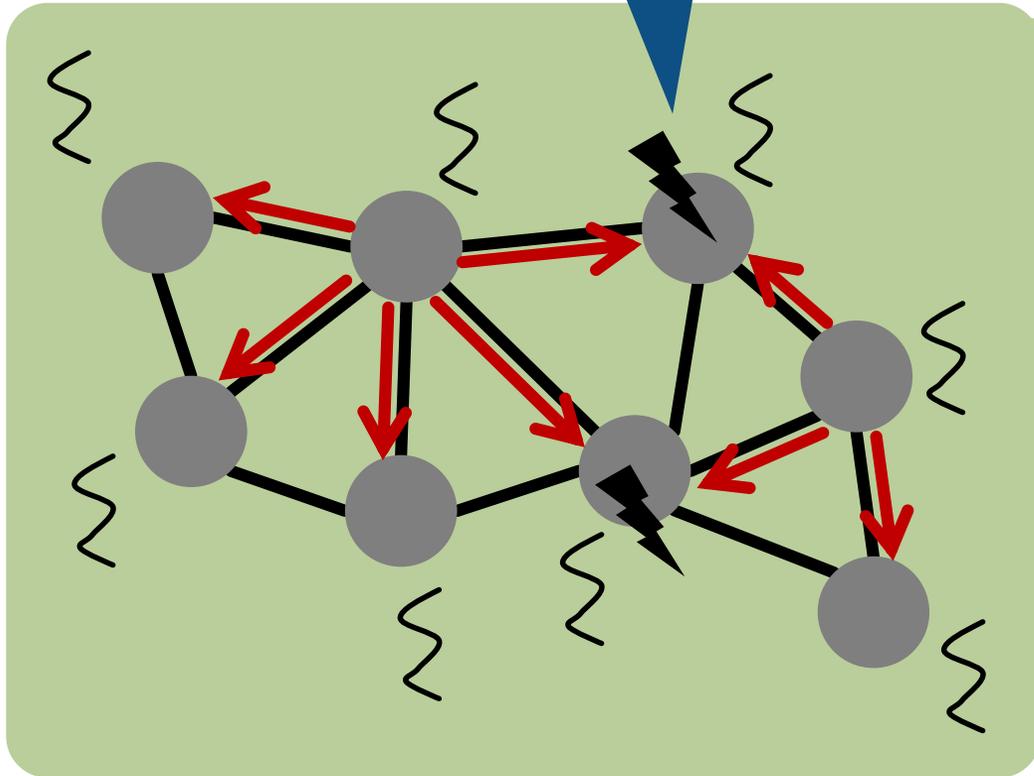


Pushing

PAGERANK

P threads are used

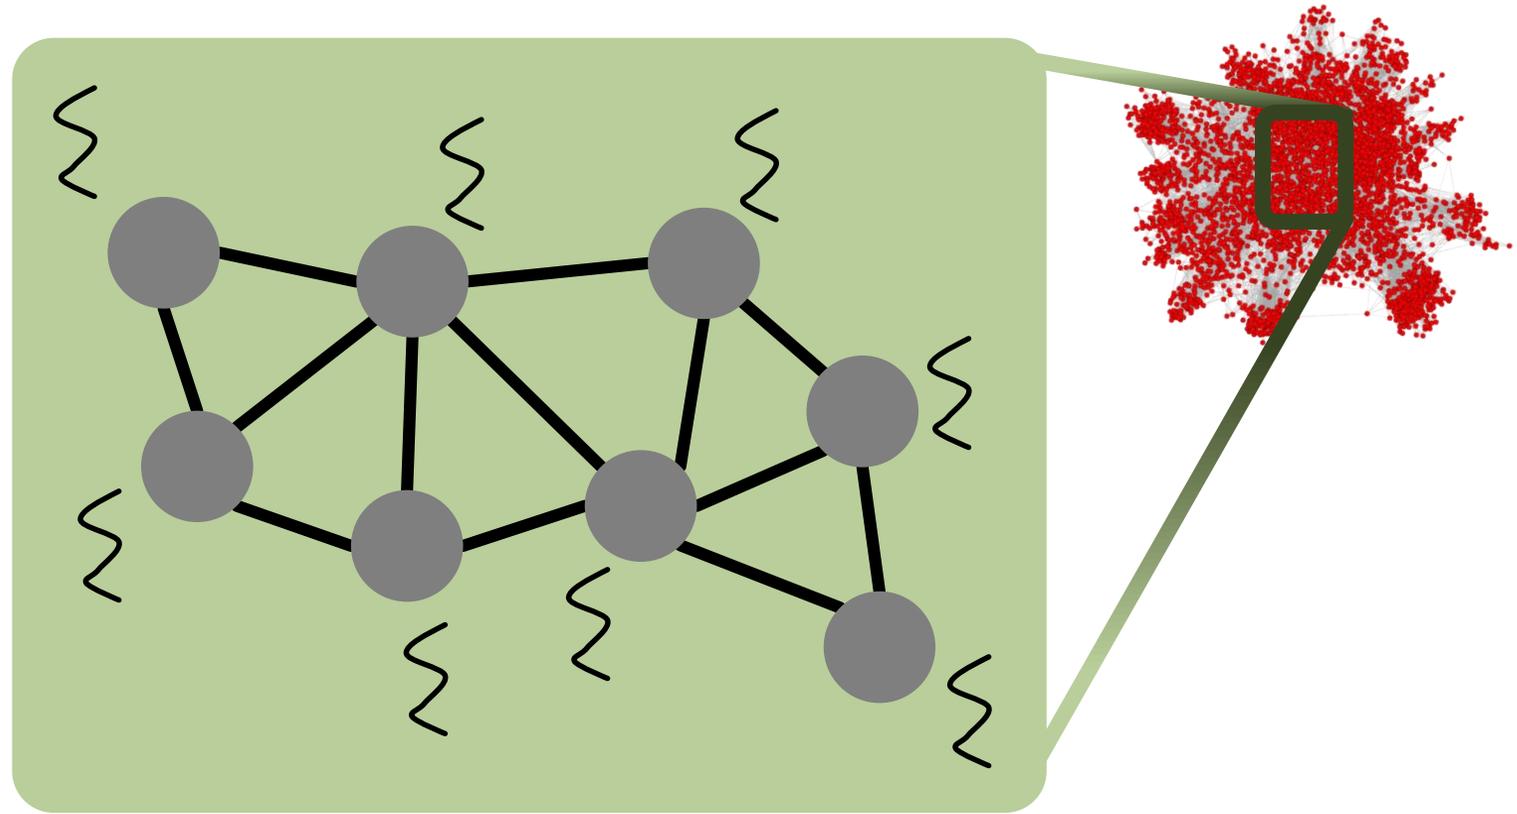
Write conflicts



Pushing

PAGERANK

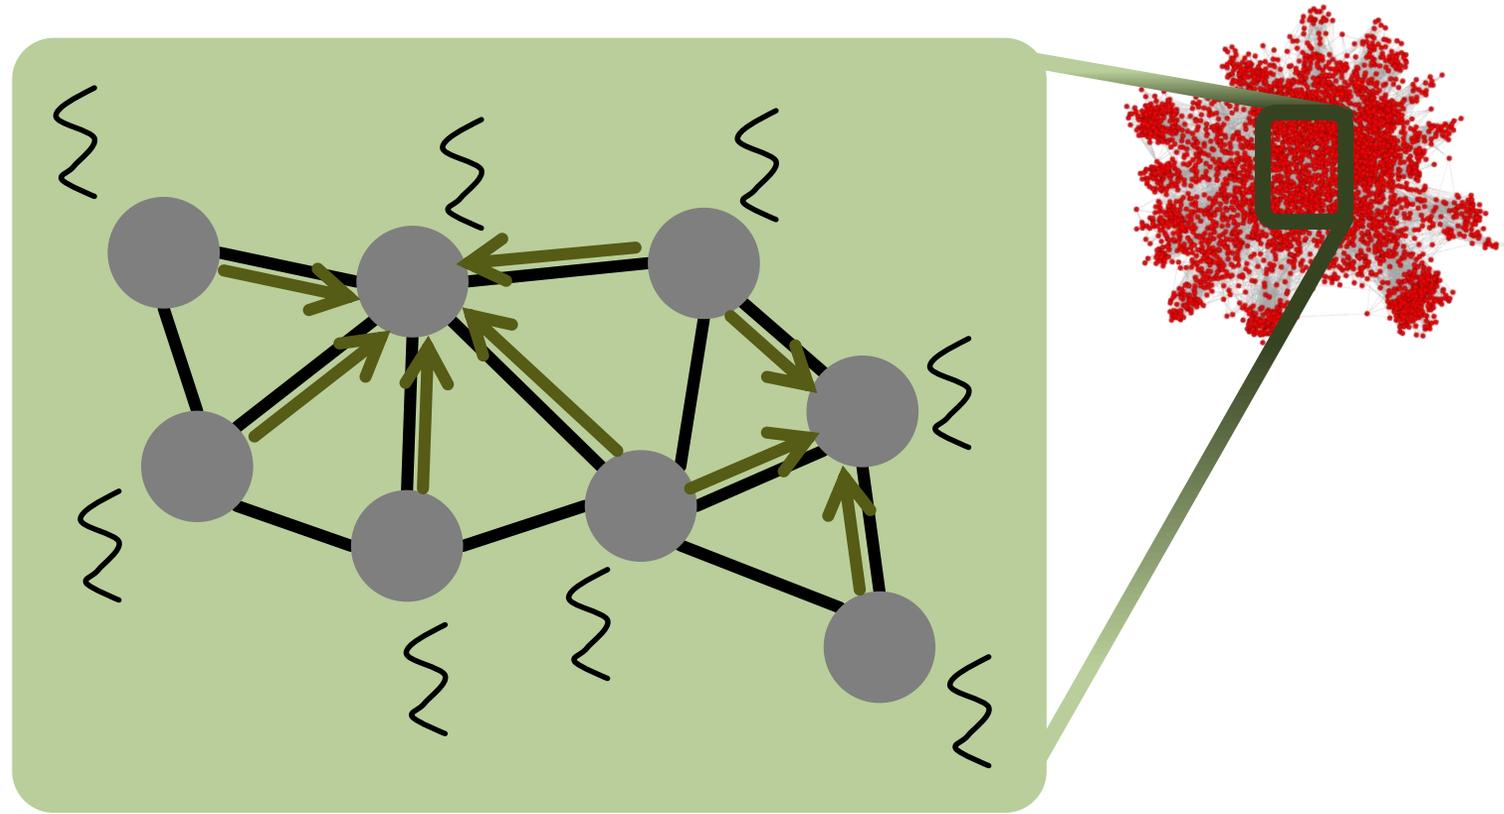
P threads are used



Pulling

PAGERANK

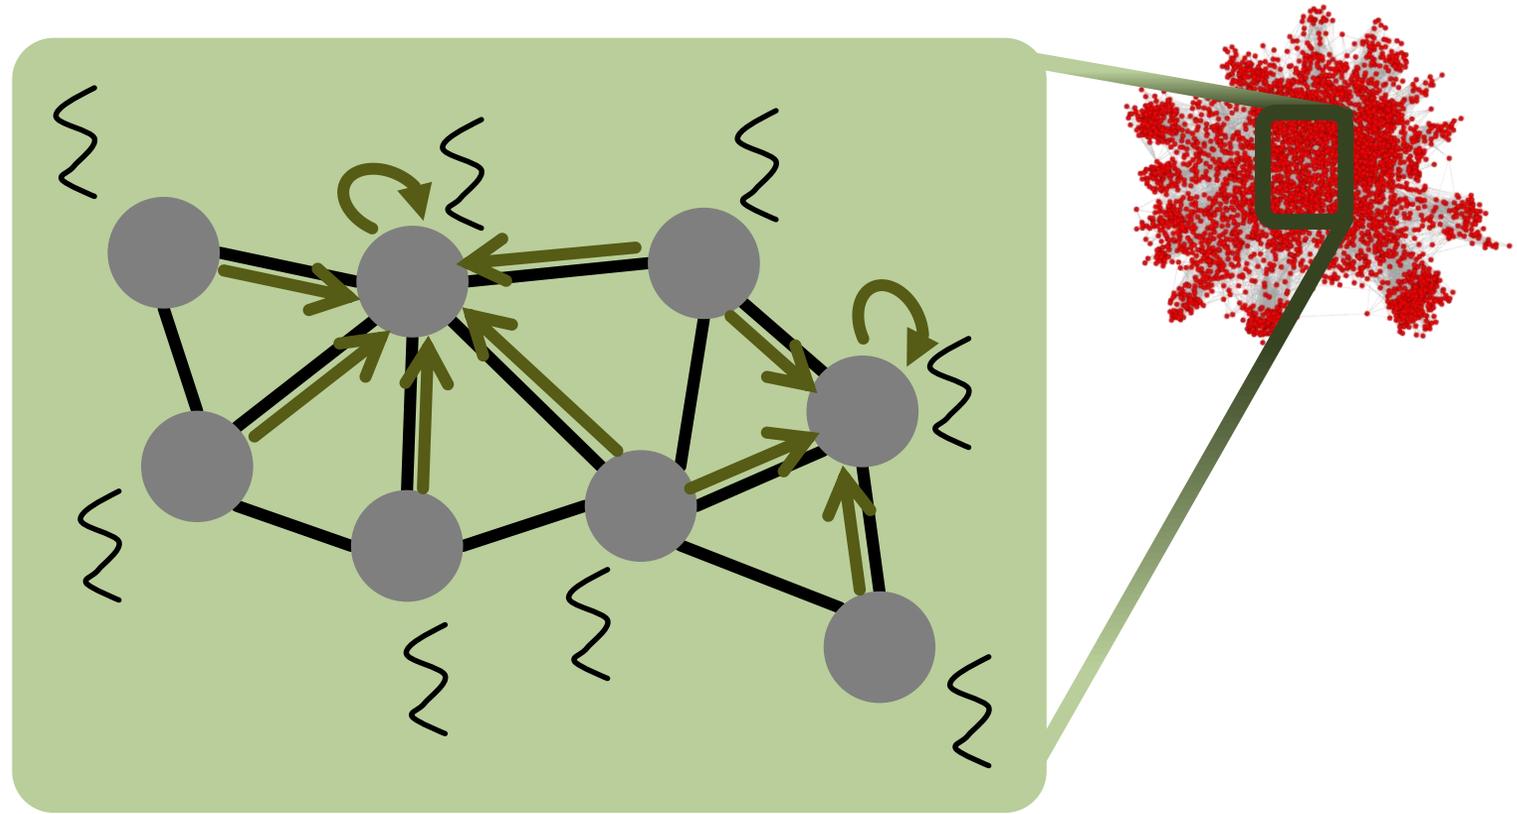
P threads are used



Pulling

PAGERANK

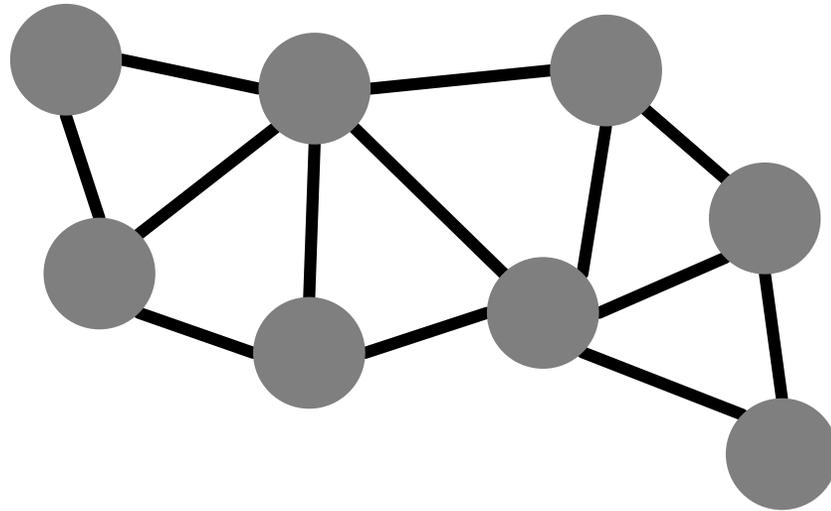
P threads are used



Pulling

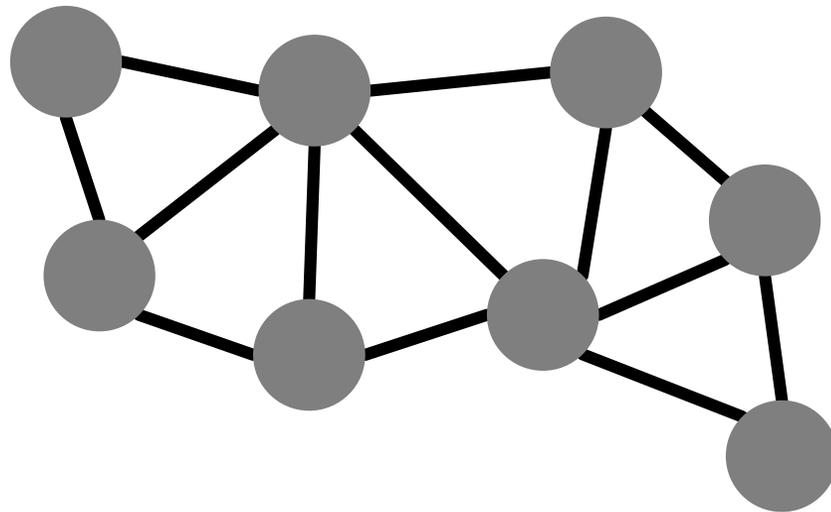
BFS

TOP-DOWN VS. BOTTOM-UP [1]



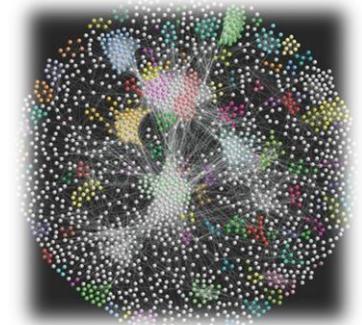
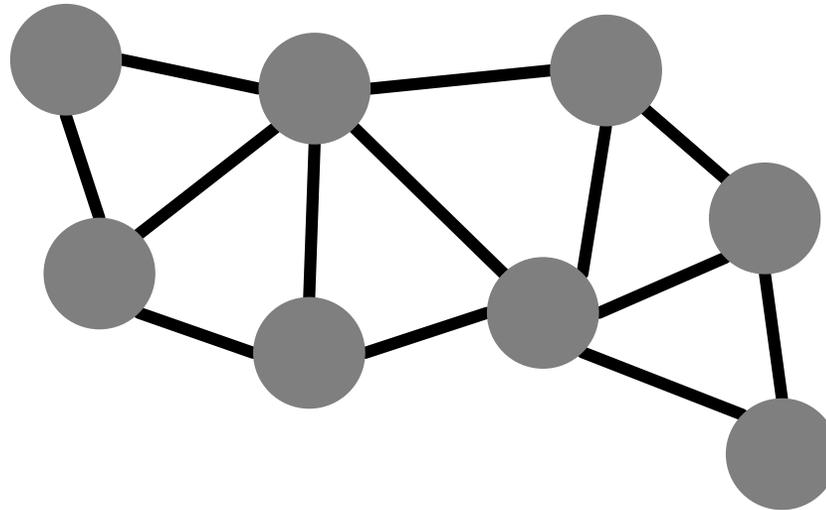
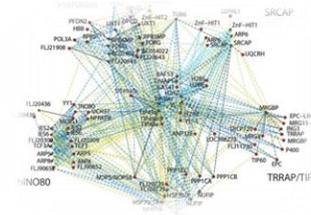
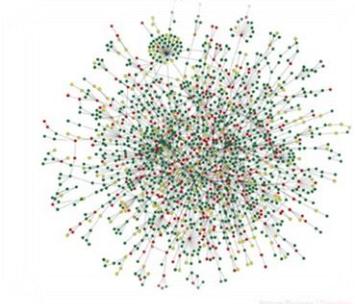
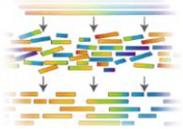
BFS

TOP-DOWN VS. BOTTOM-UP [1]



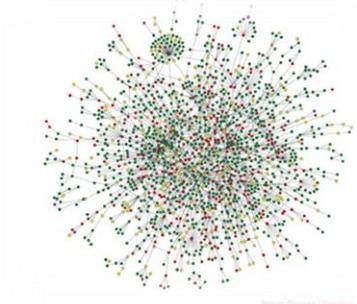
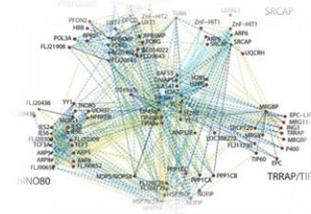
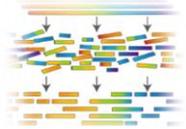
BFS

TOP-DOWN VS. BOTTOM-UP [1]

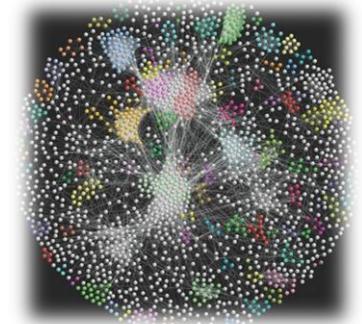
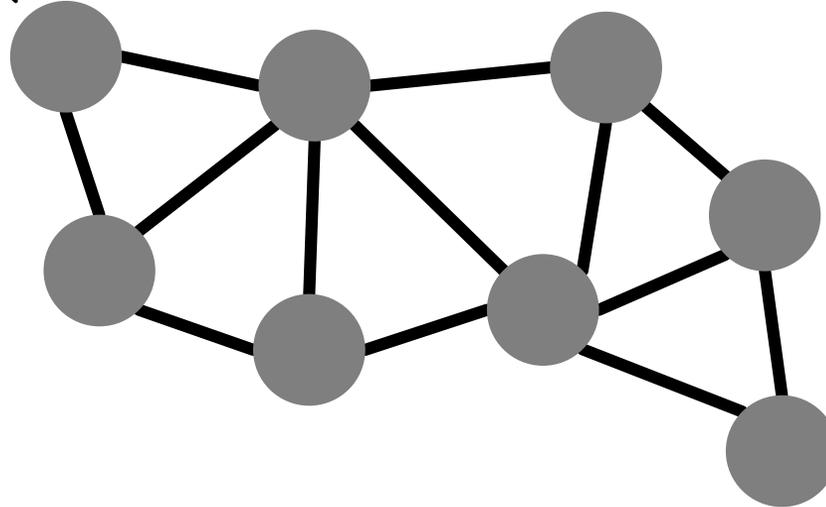


BFS

TOP-DOWN VS. BOTTOM-UP [1]

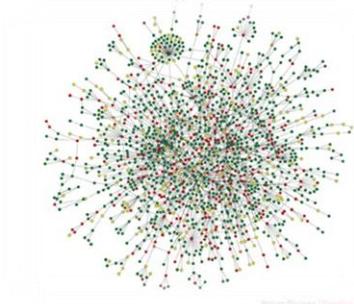
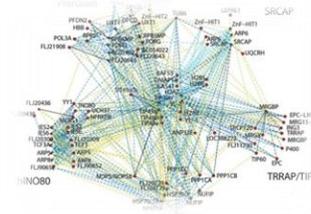
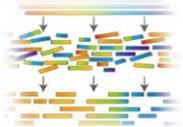


Root r

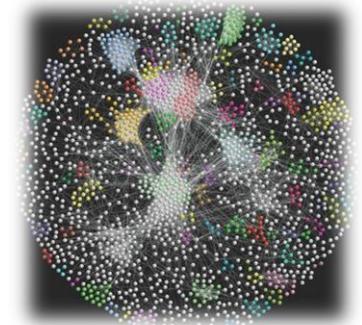
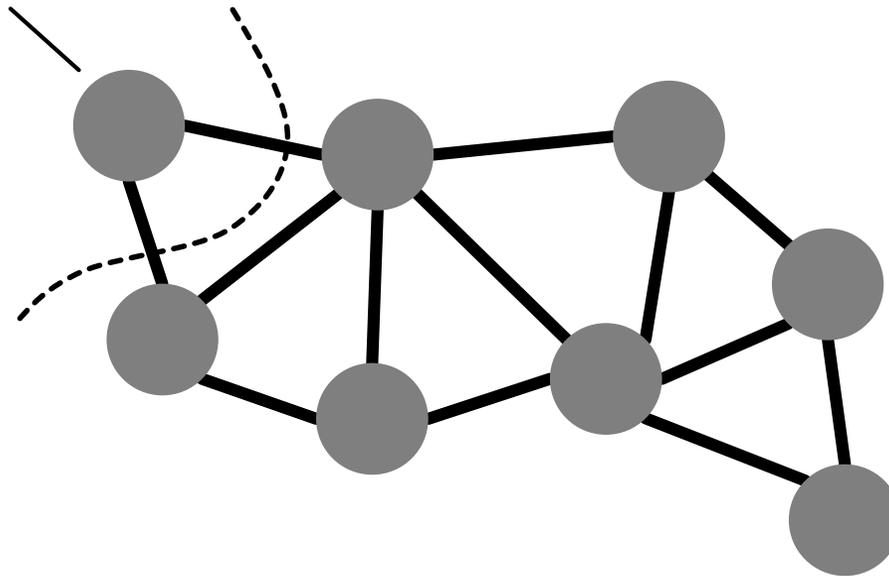


BFS

TOP-DOWN VS. BOTTOM-UP [1]

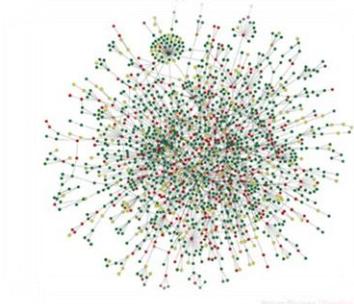
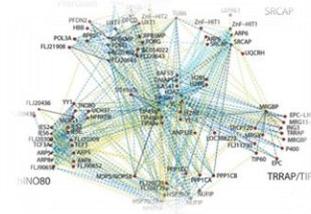
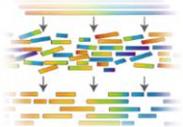


Root r

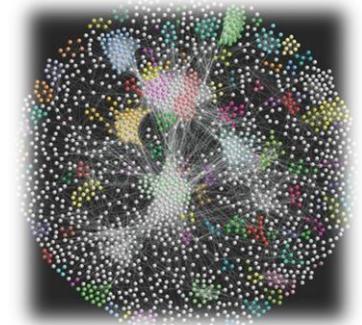
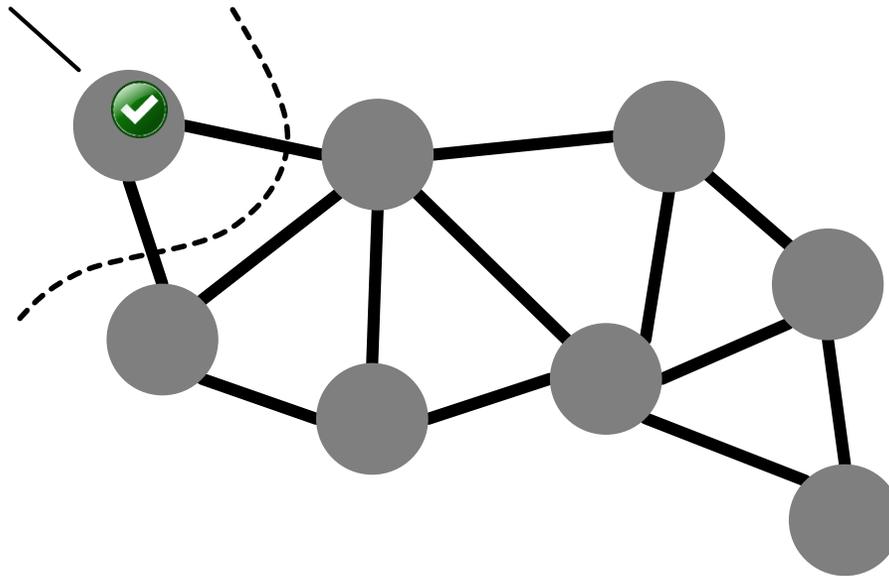


BFS

TOP-DOWN VS. BOTTOM-UP [1]

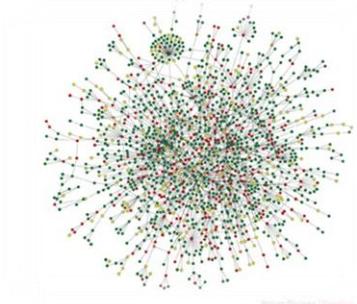
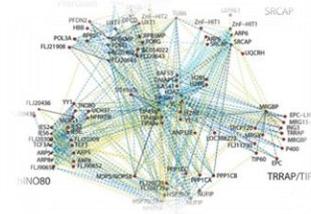
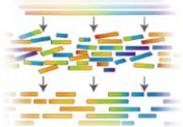


Root r

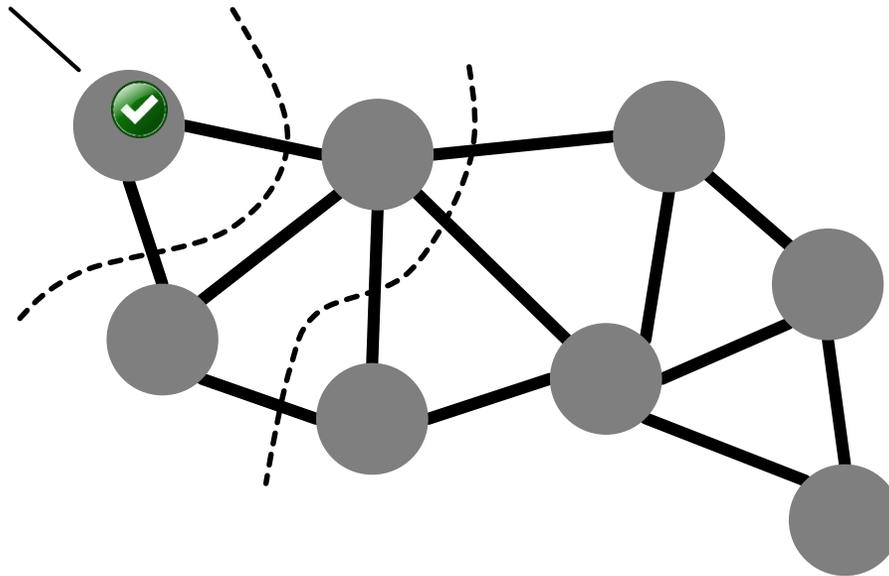


BFS

TOP-DOWN VS. BOTTOM-UP [1]

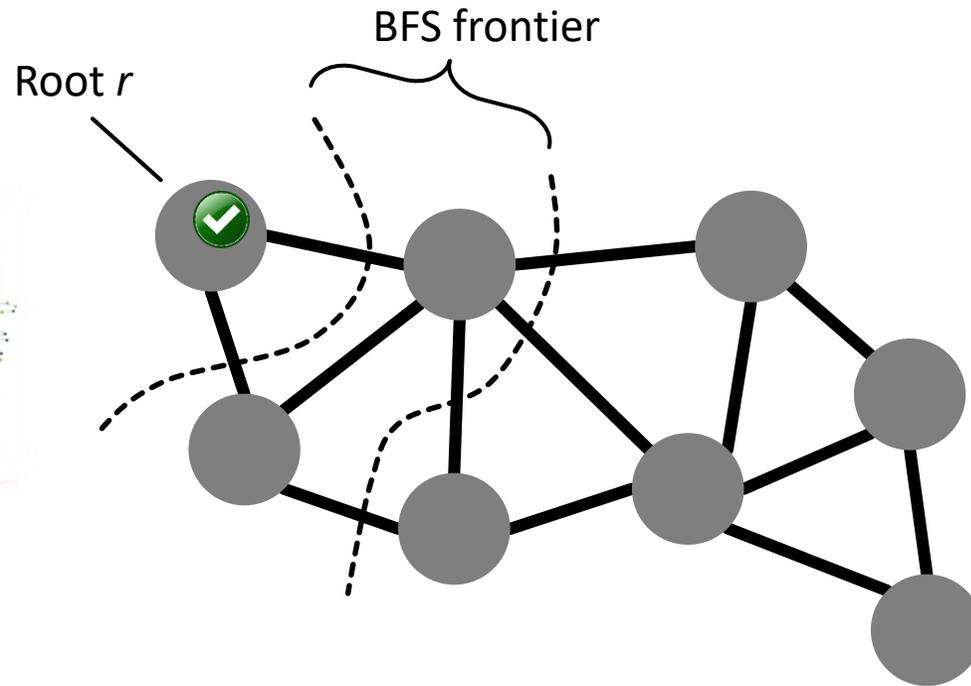
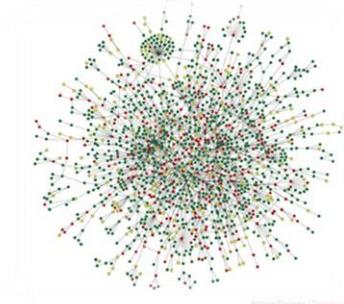
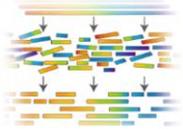
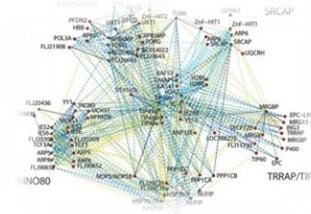


Root r



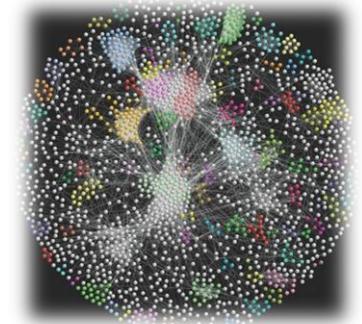
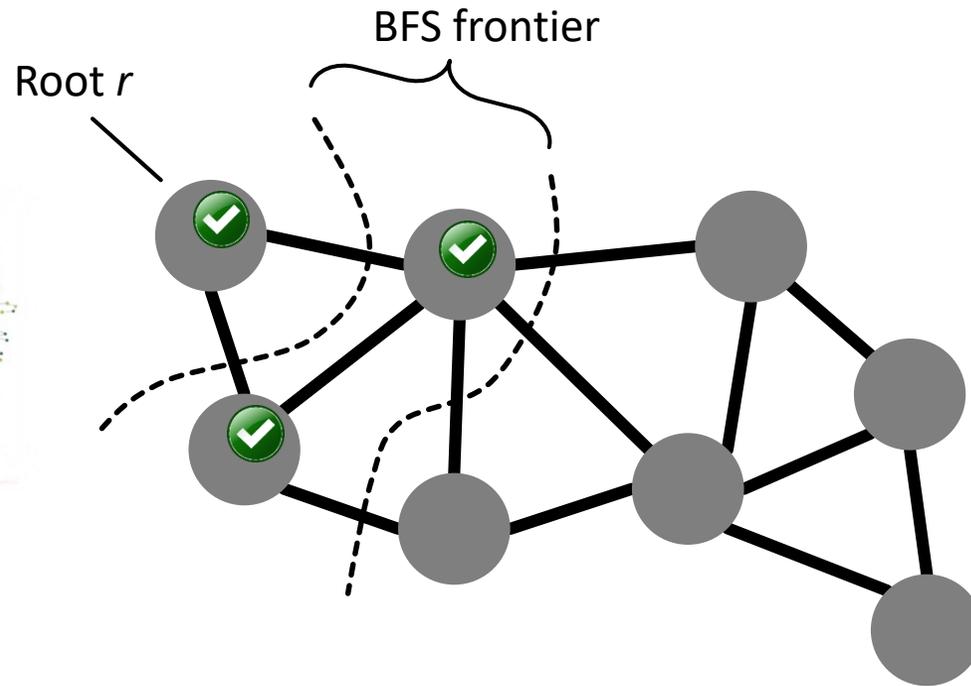
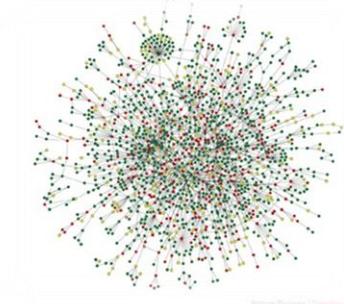
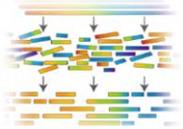
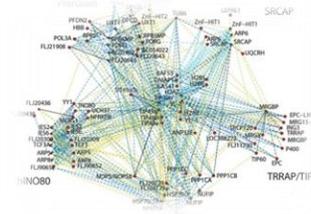
BFS

TOP-DOWN VS. BOTTOM-UP [1]



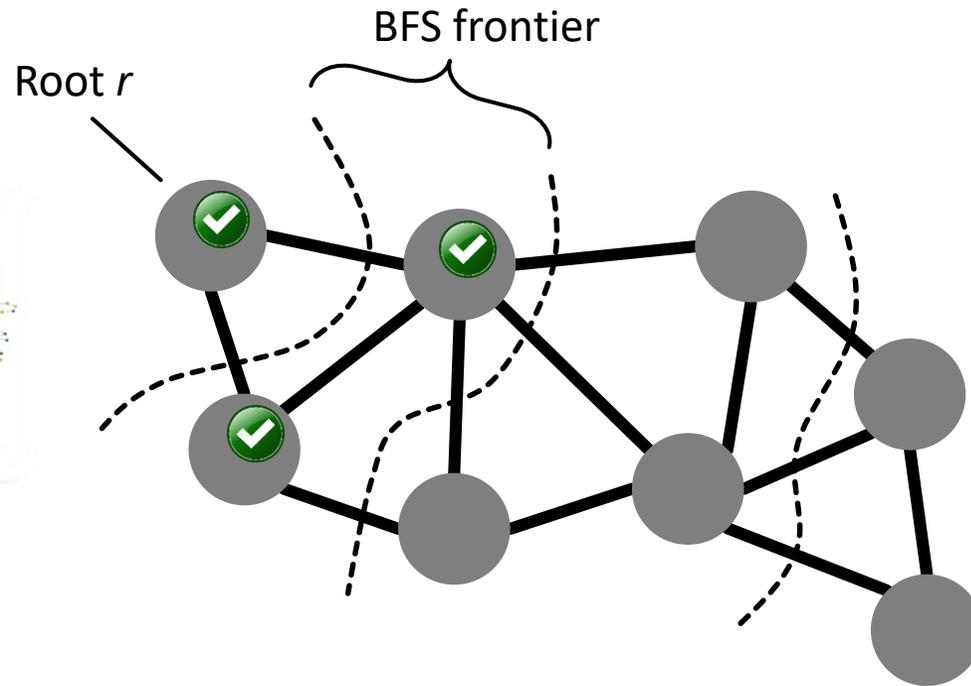
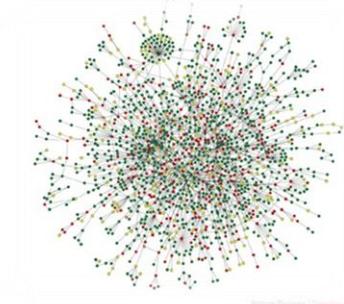
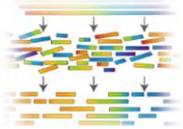
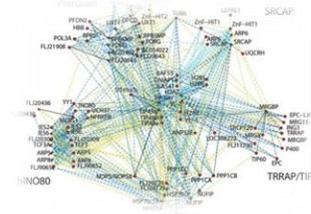
BFS

TOP-DOWN VS. BOTTOM-UP [1]



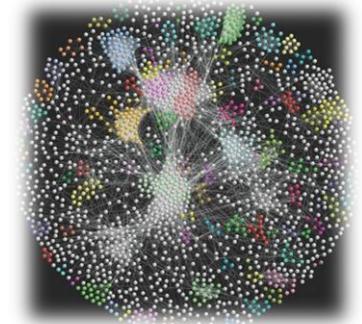
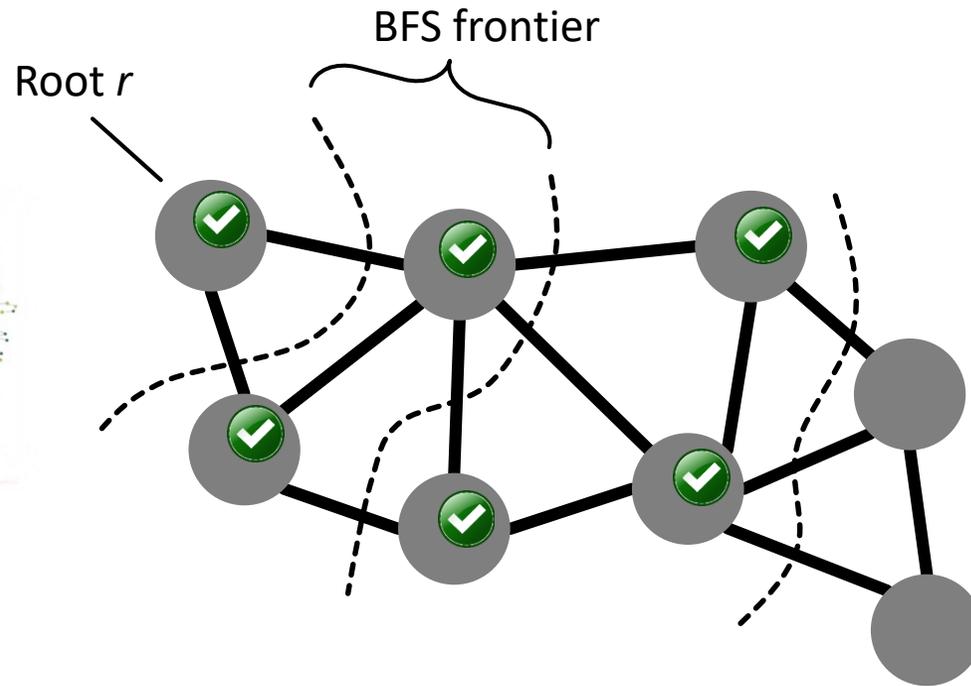
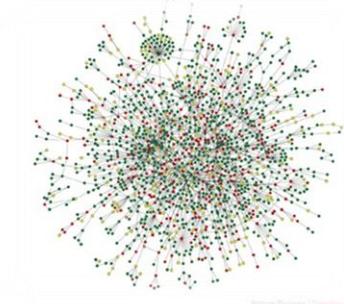
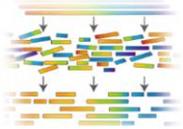
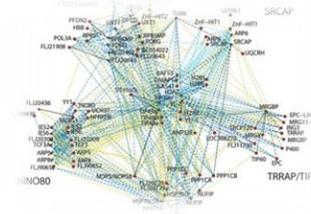
BFS

TOP-DOWN VS. BOTTOM-UP [1]



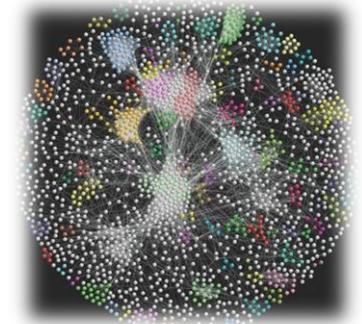
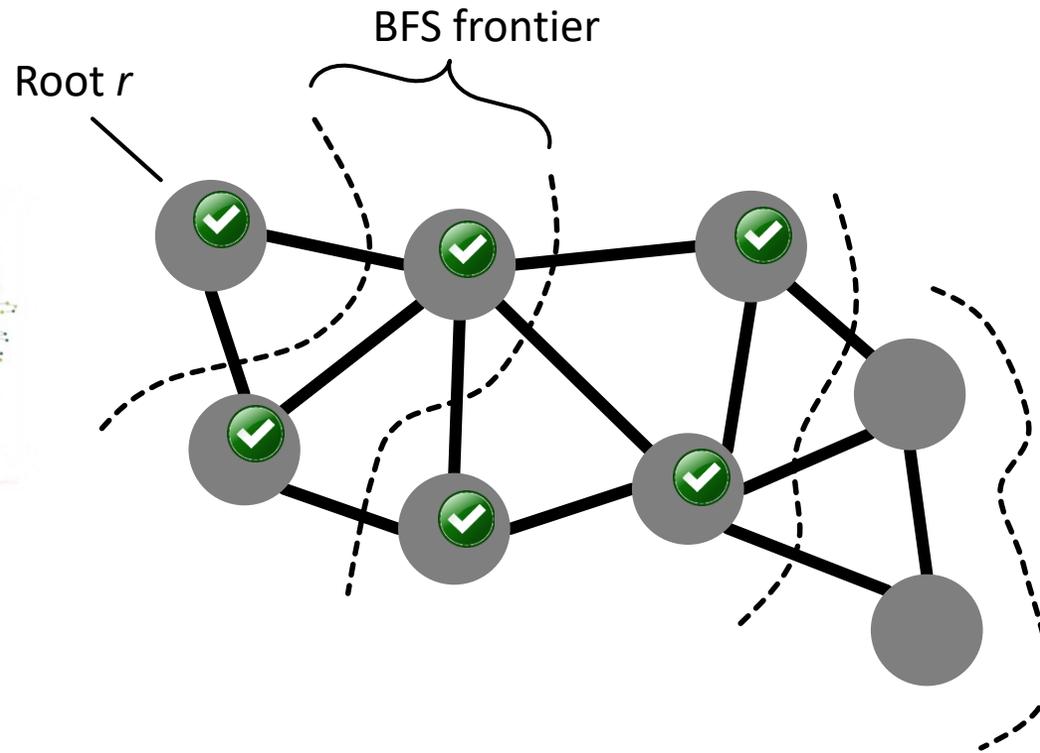
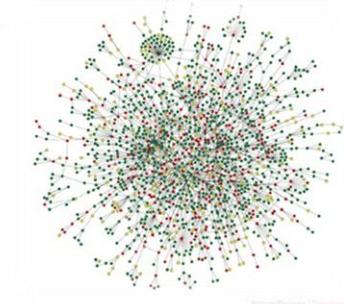
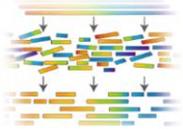
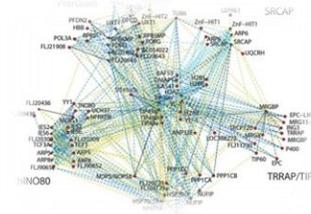
BFS

TOP-DOWN VS. BOTTOM-UP [1]



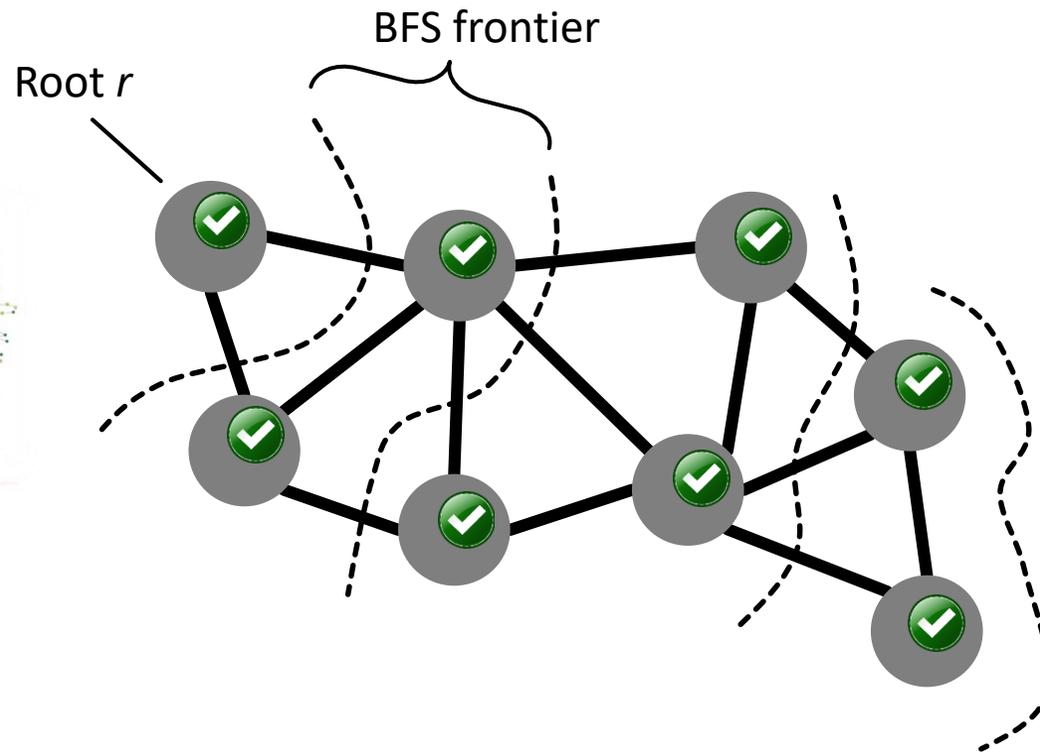
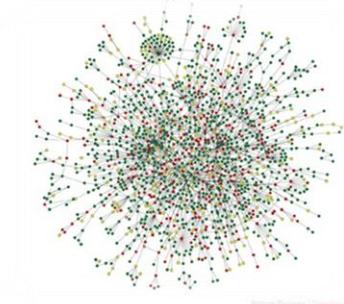
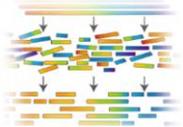
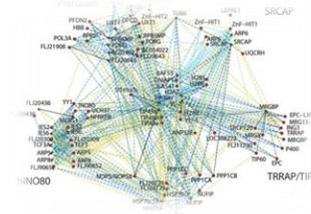
BFS

TOP-DOWN VS. BOTTOM-UP [1]



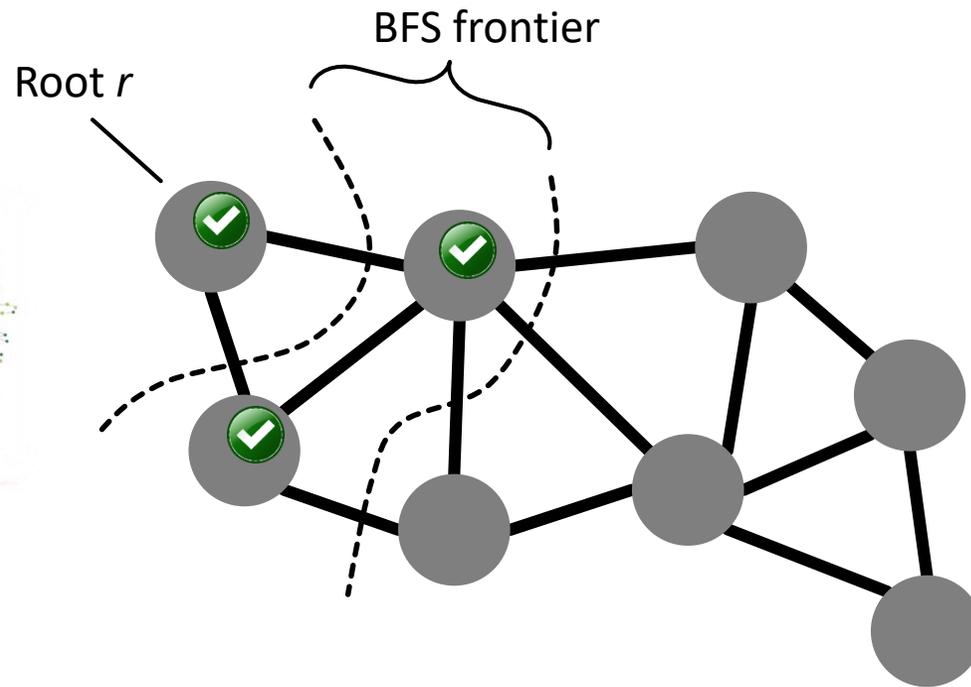
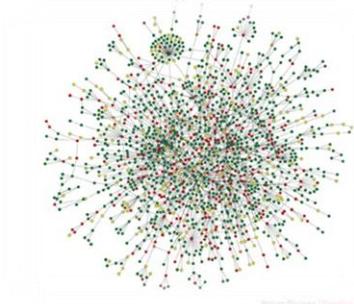
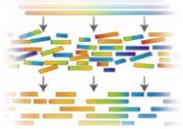
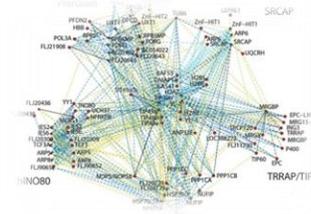
BFS

TOP-DOWN VS. BOTTOM-UP [1]



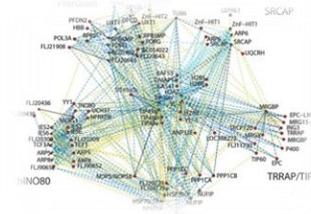
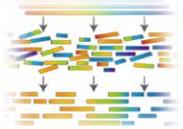
BFS

TOP-DOWN VS. BOTTOM-UP [1]

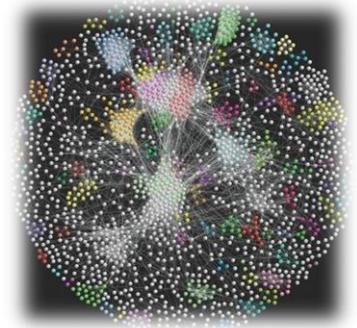
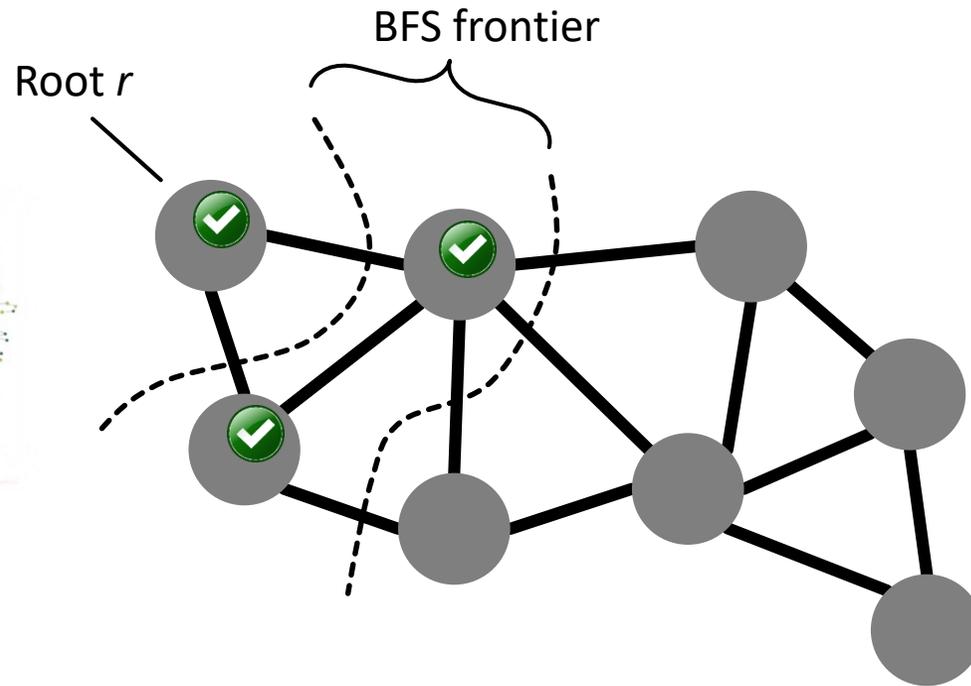
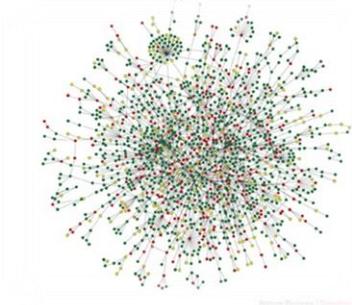


BFS

TOP-DOWN VS. BOTTOM-UP [1]

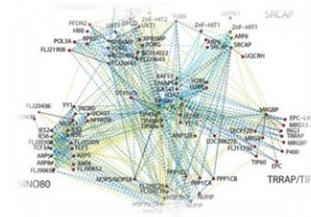
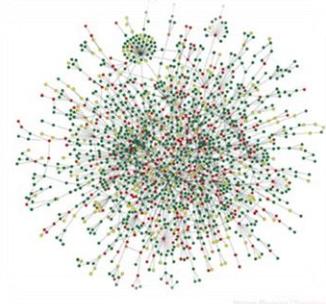
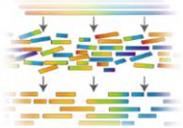


Pushing or pulling when expanding a frontier

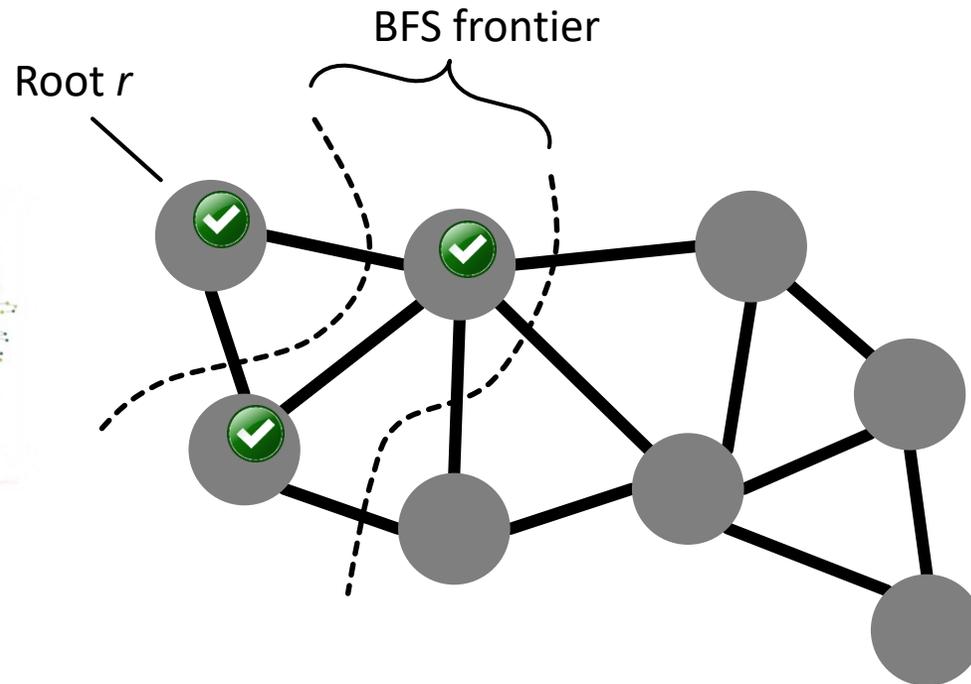


BFS

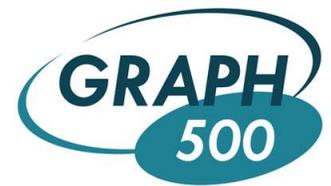
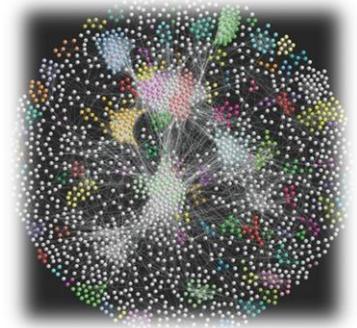
TOP-DOWN VS. BOTTOM-UP [1]



Pushing or pulling when expanding a frontier

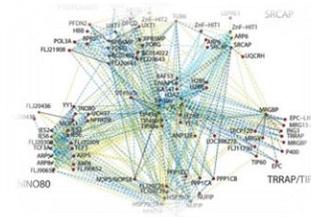
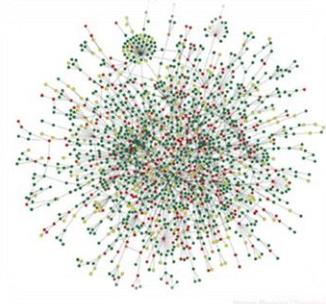
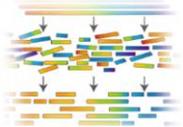


Pushing

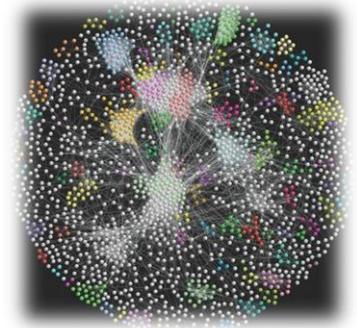
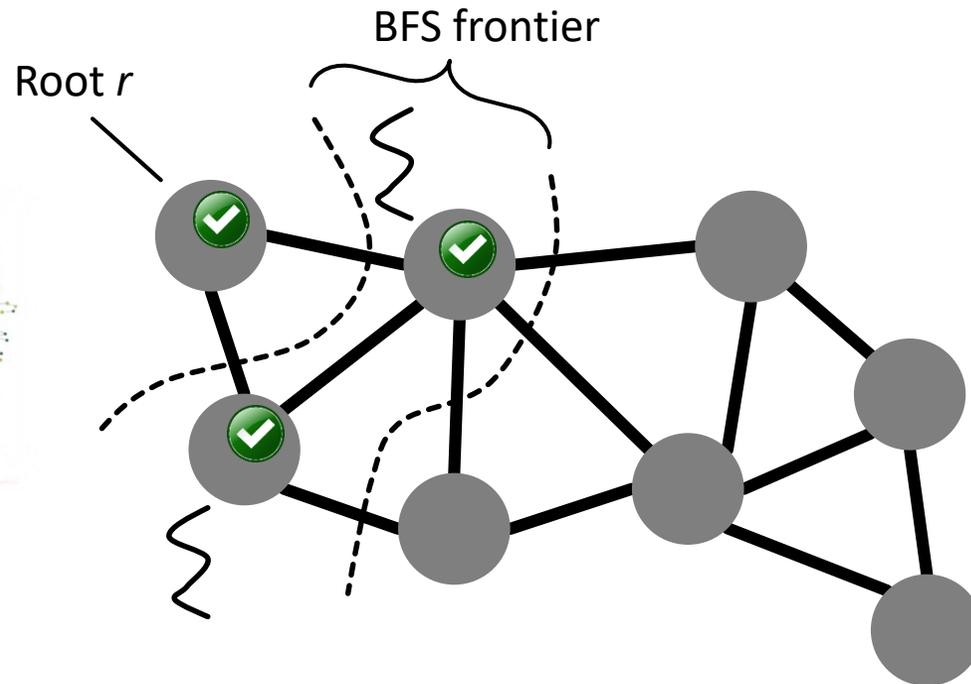


BFS

TOP-DOWN VS. BOTTOM-UP [1]



Pushing or pulling when expanding a frontier

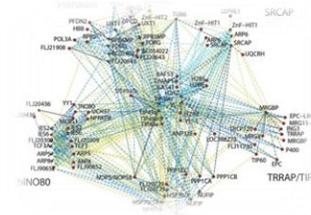
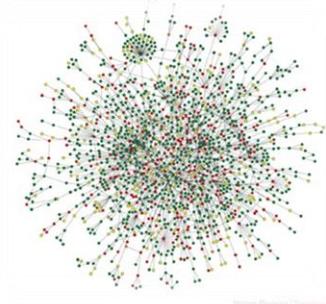
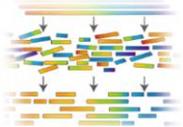


Pushing

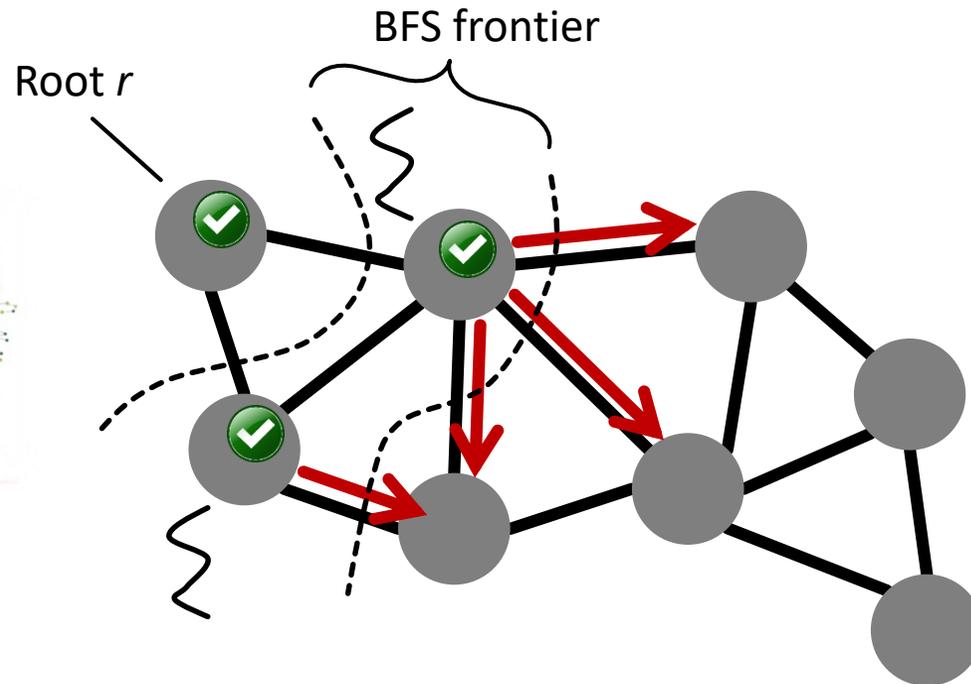


BFS

TOP-DOWN VS. BOTTOM-UP [1]



Pushing or pulling when expanding a frontier

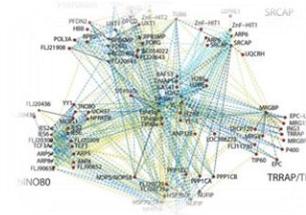
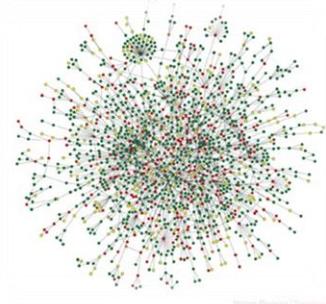
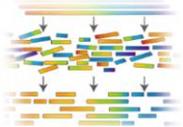


Pushing

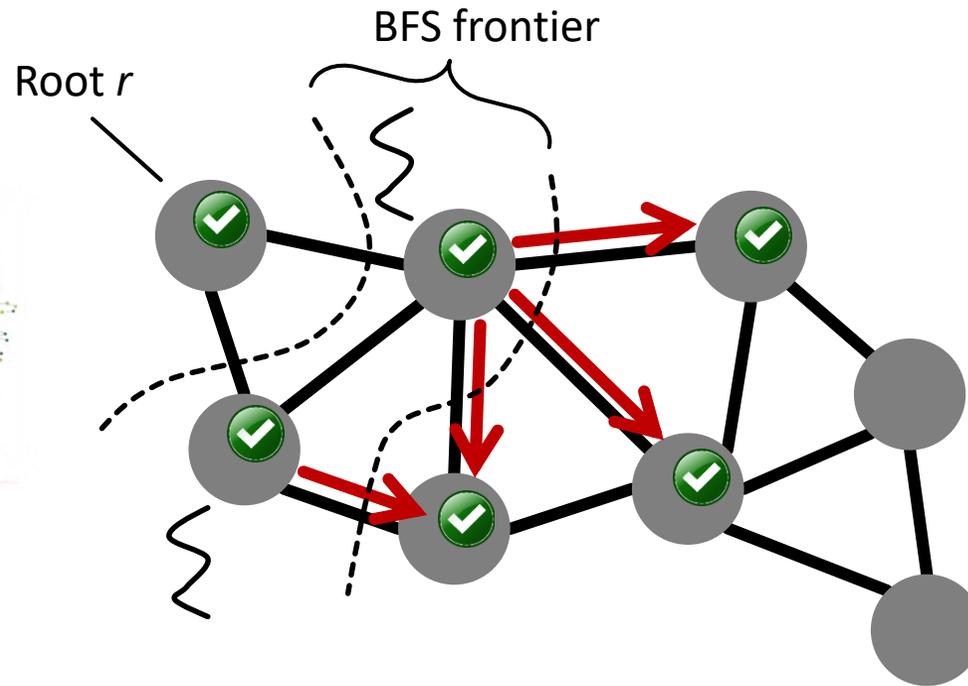


BFS

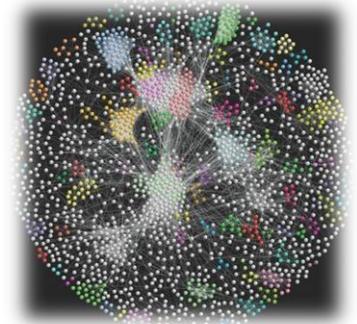
TOP-DOWN VS. BOTTOM-UP [1]



Pushing or pulling when expanding a frontier

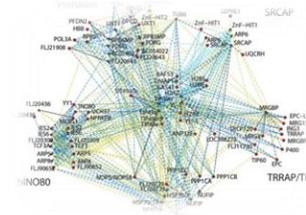
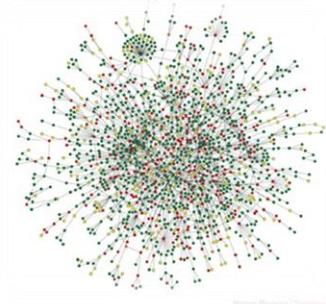
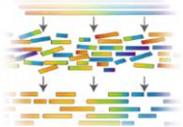


Pushing

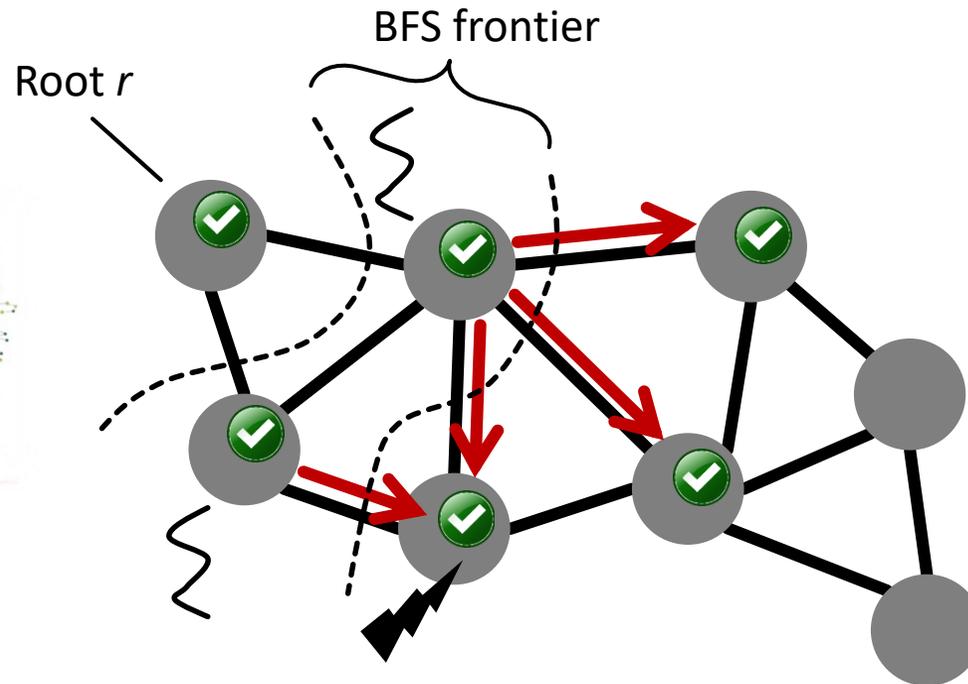


BFS

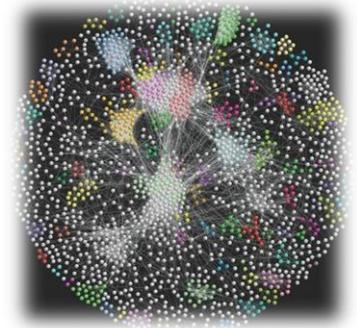
TOP-DOWN VS. BOTTOM-UP [1]



Pushing or pulling when expanding a frontier

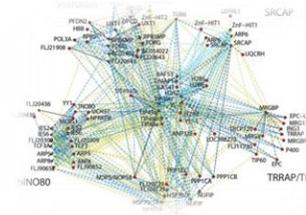
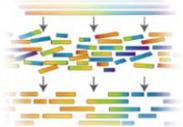


Pushing

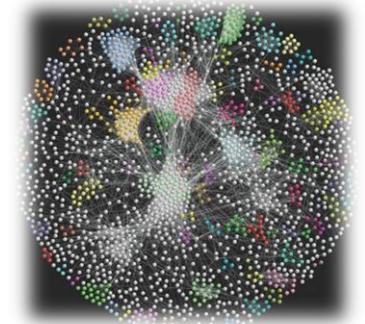
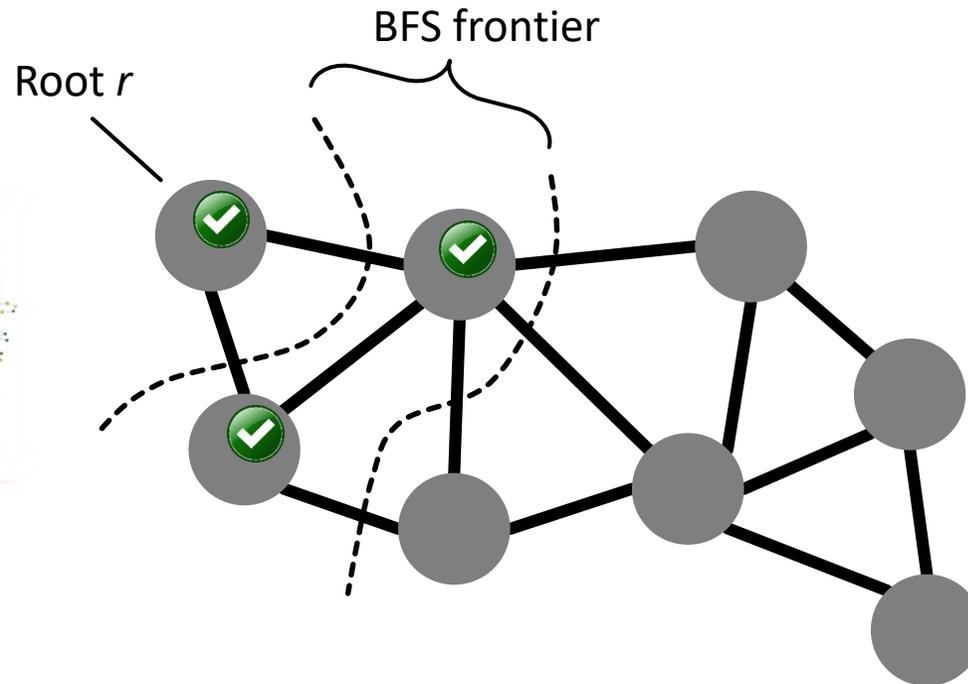
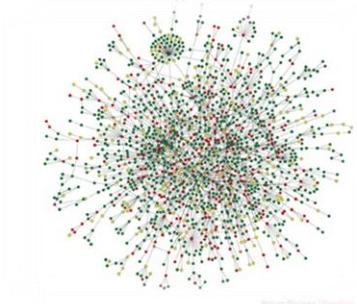


BFS

TOP-DOWN VS. BOTTOM-UP [1]

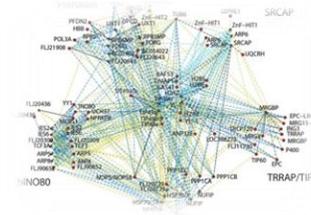
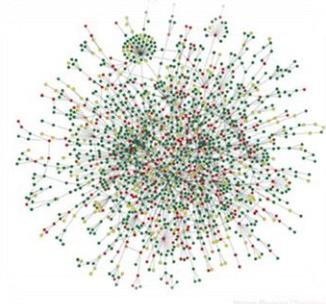
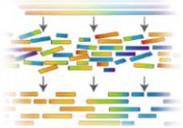


Pushing or pulling when expanding a frontier

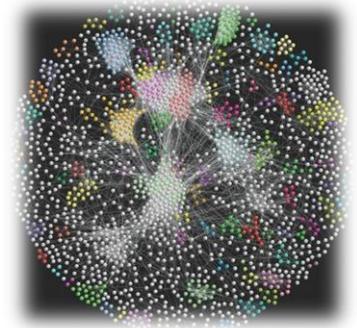
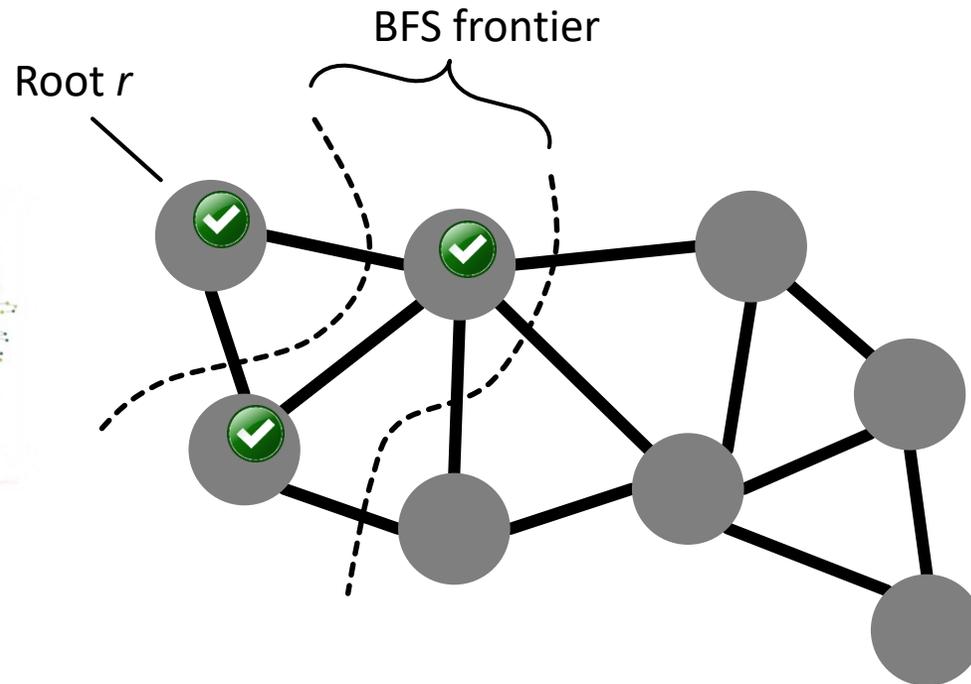


BFS

TOP-DOWN VS. BOTTOM-UP [1]



Pushing or pulling when expanding a frontier

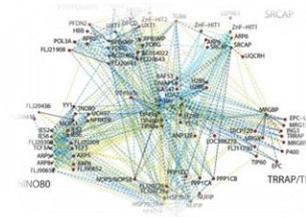
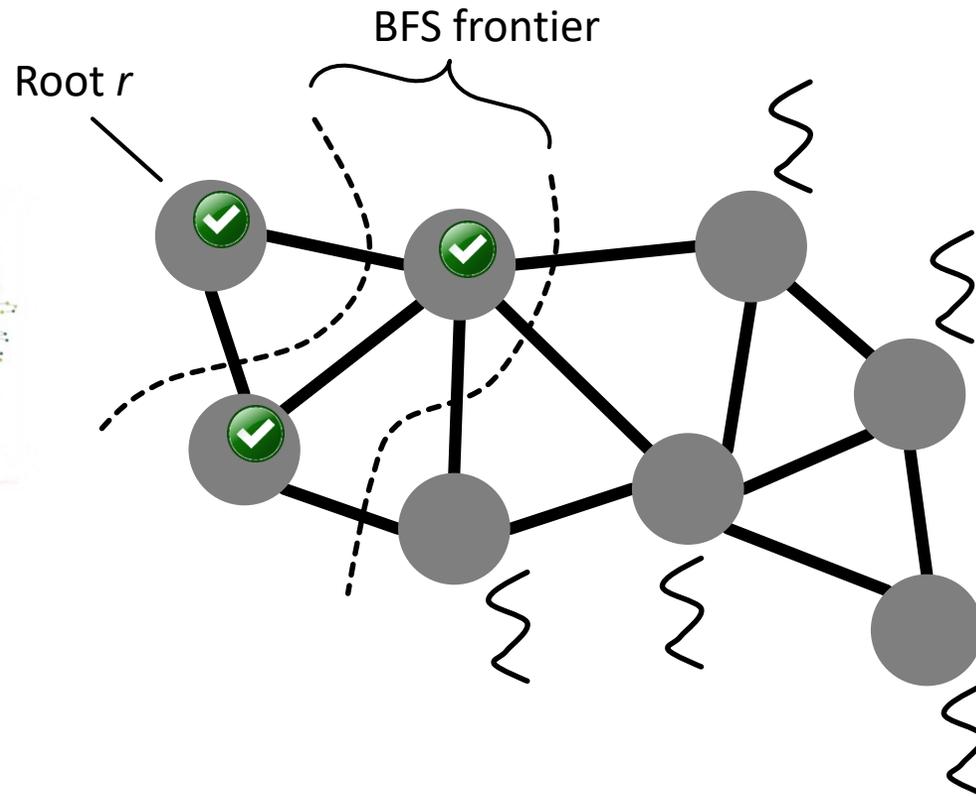
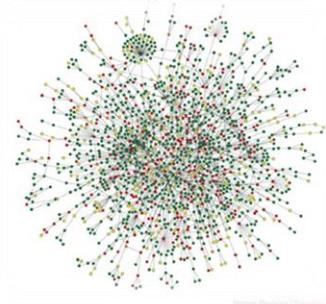
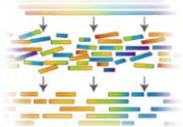


Pulling

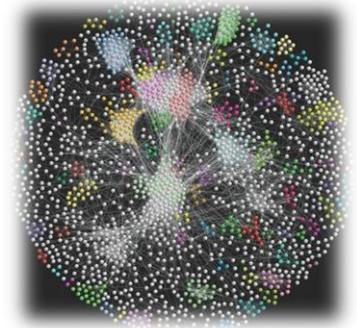


BFS

TOP-DOWN VS. BOTTOM-UP [1]



Pushing or pulling when expanding a frontier

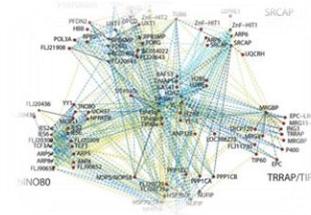
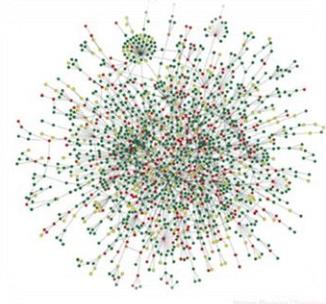
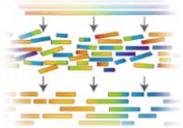


Pulling

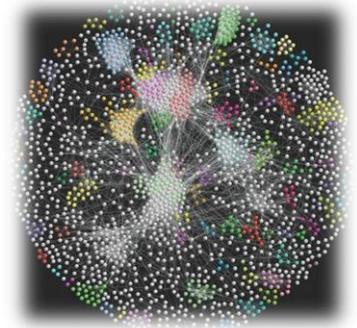
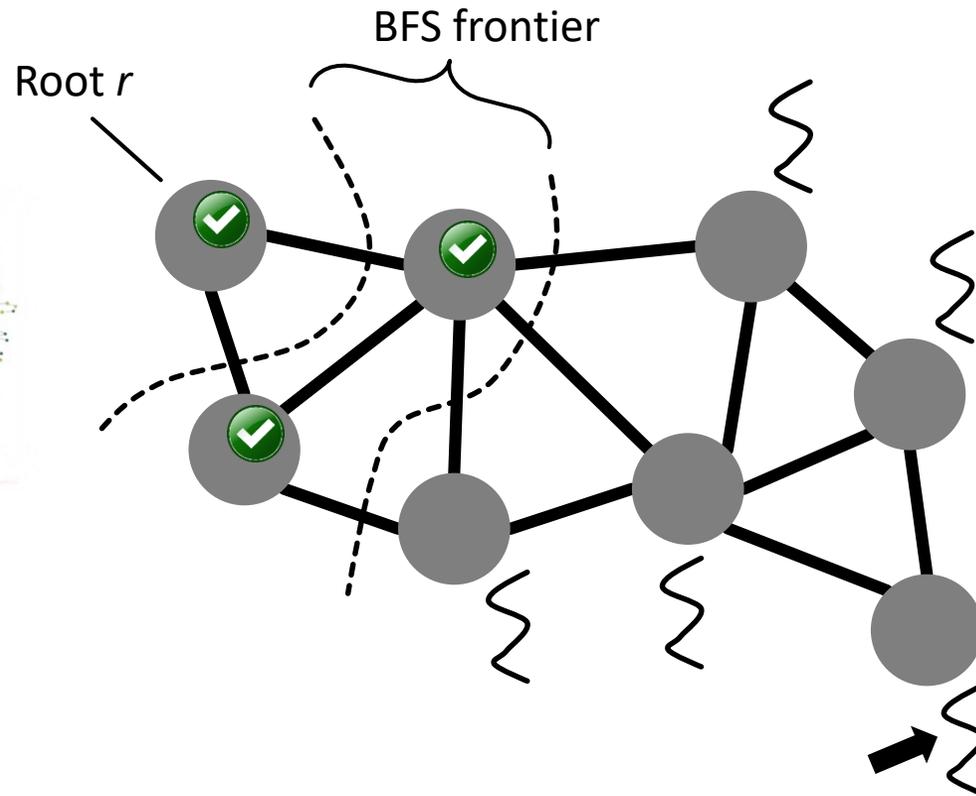


BFS

TOP-DOWN VS. BOTTOM-UP [1]



Pushing or pulling when expanding a frontier

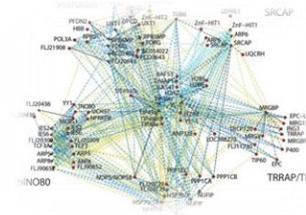
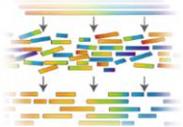


Pulling

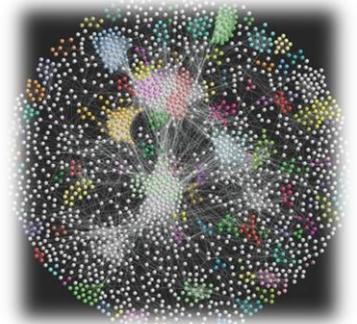
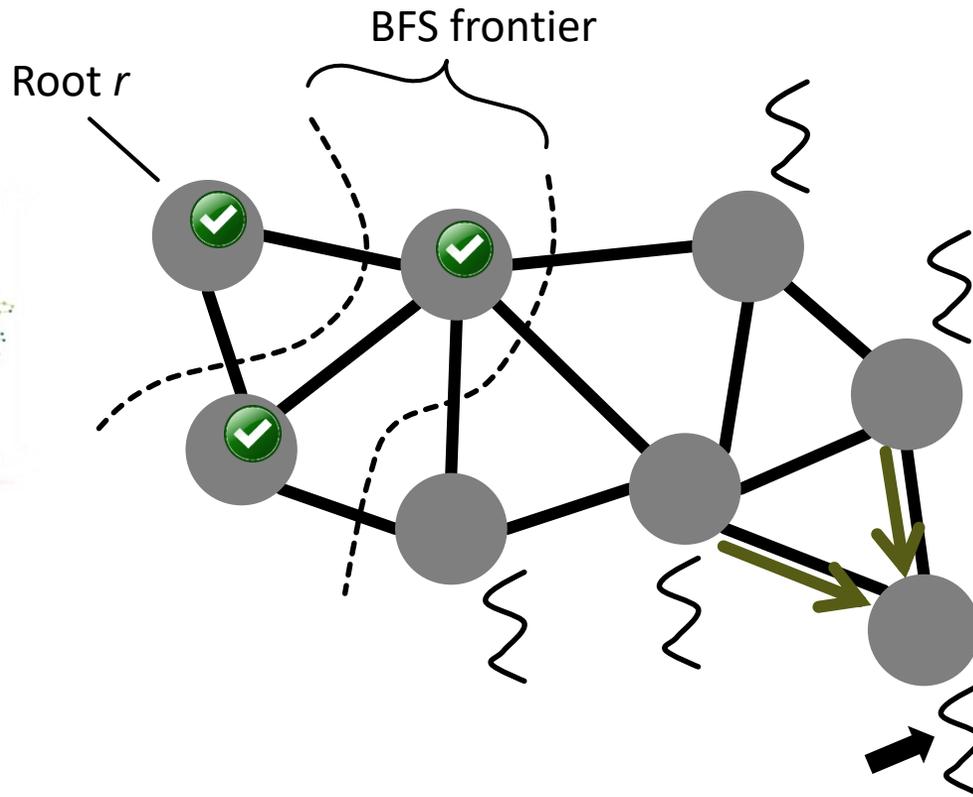
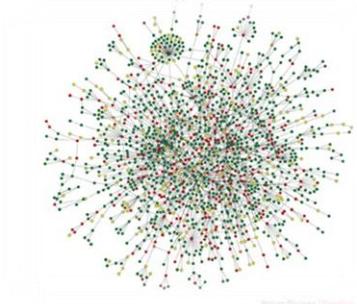


BFS

TOP-DOWN VS. BOTTOM-UP [1]



Pushing or pulling when expanding a frontier

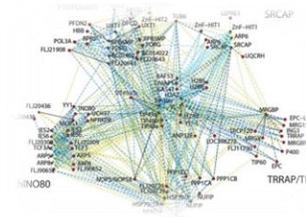
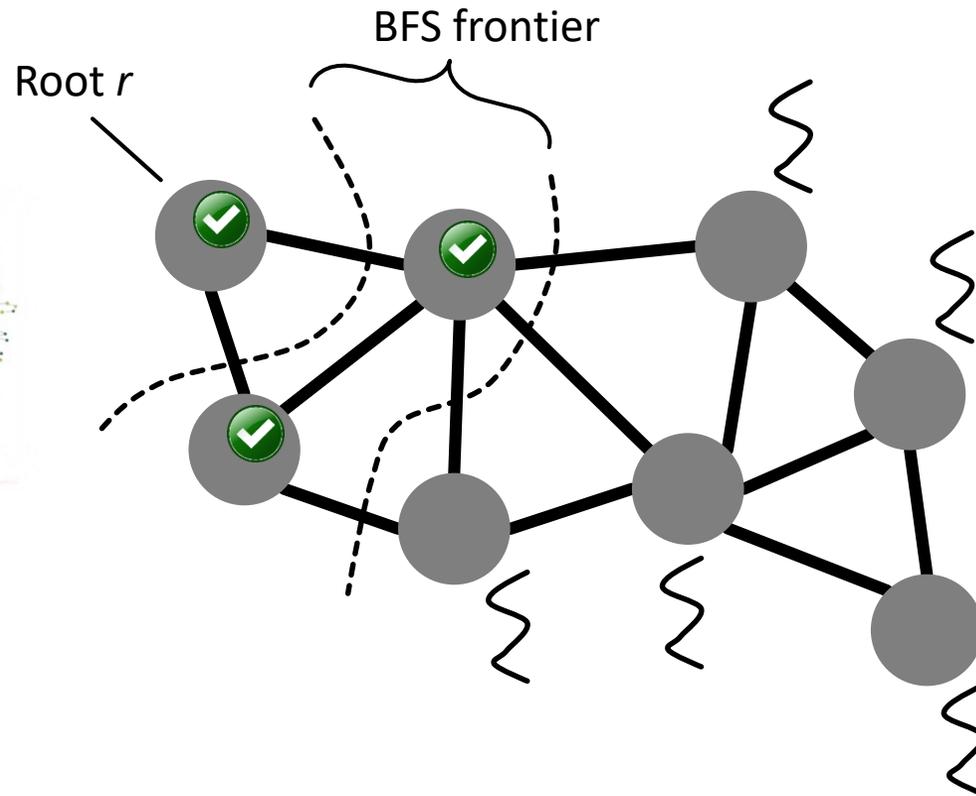
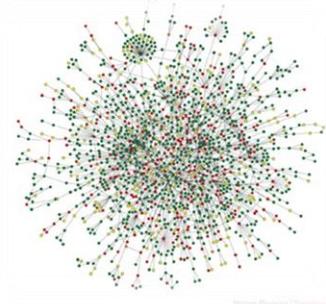
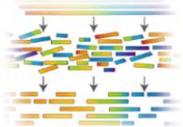


Pulling

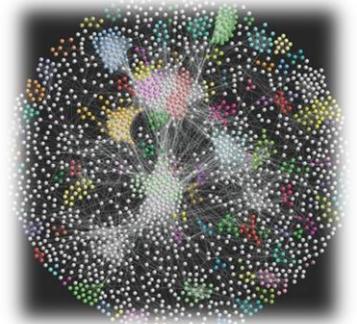


BFS

TOP-DOWN VS. BOTTOM-UP [1]



Pushing or pulling when expanding a frontier

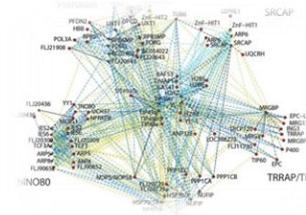
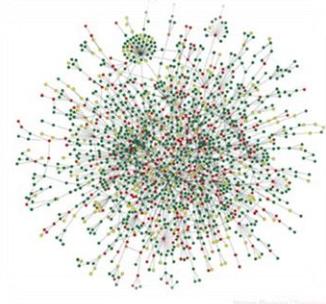
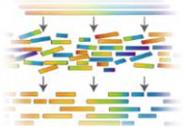


Pulling

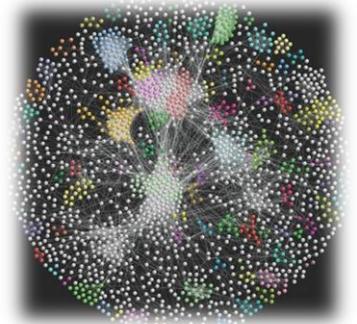
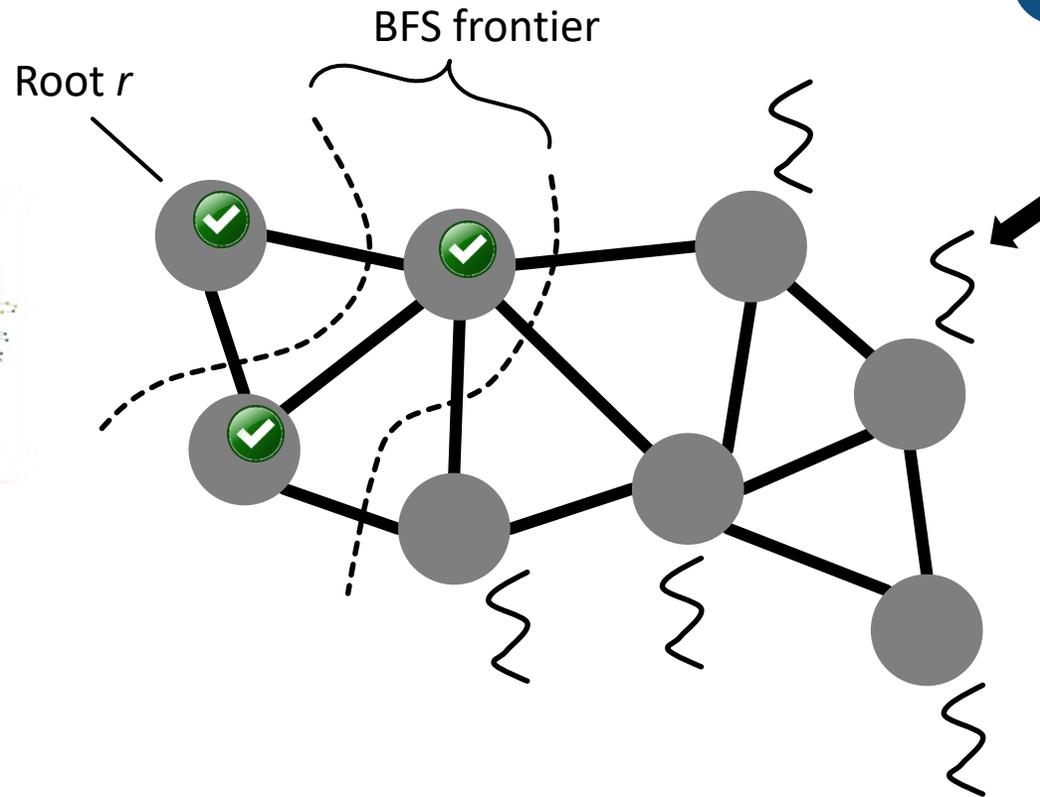


BFS

TOP-DOWN VS. BOTTOM-UP [1]



Pushing or pulling when expanding a frontier

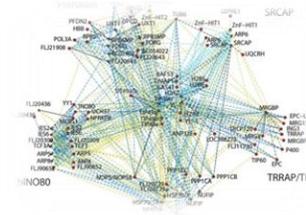
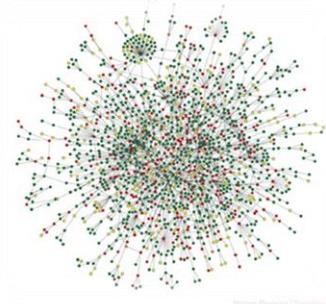
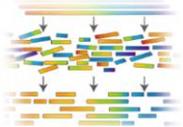


Pulling

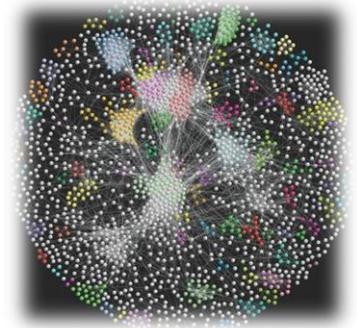
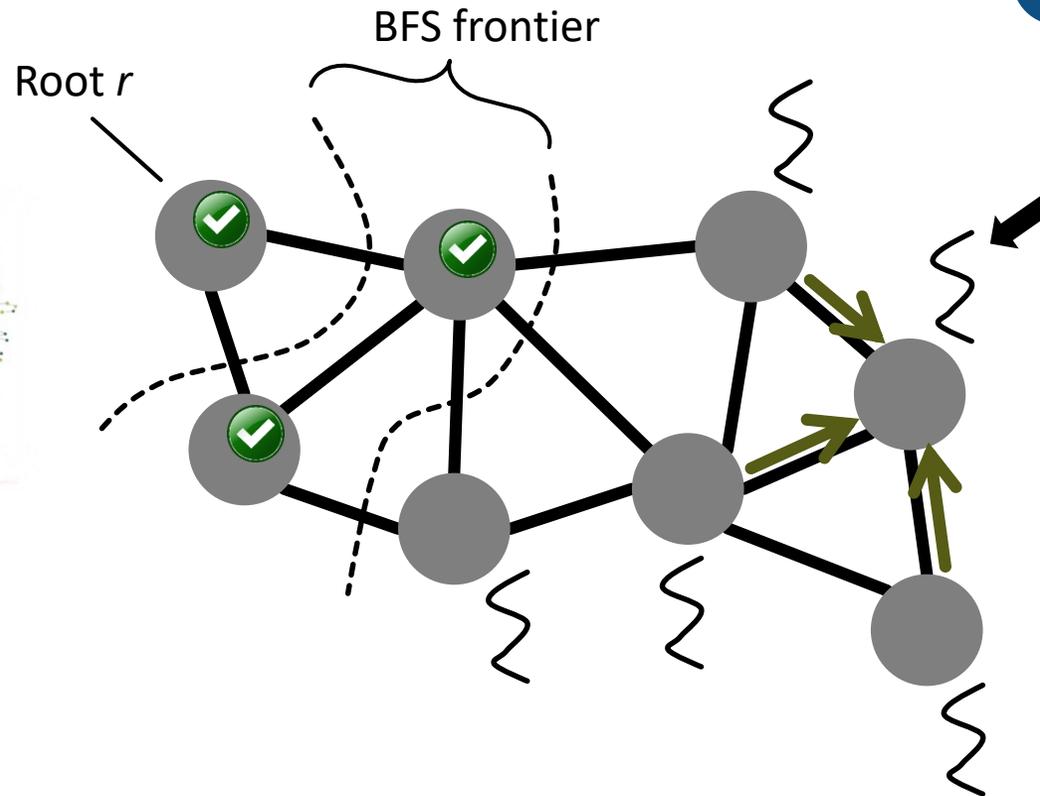


BFS

TOP-DOWN VS. BOTTOM-UP [1]



Pushing or pulling when expanding a frontier

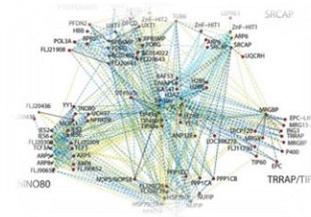
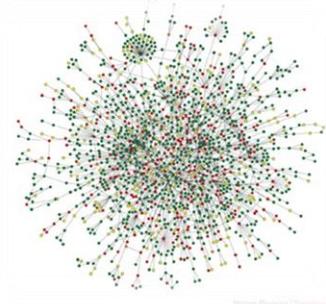
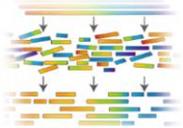


Pulling

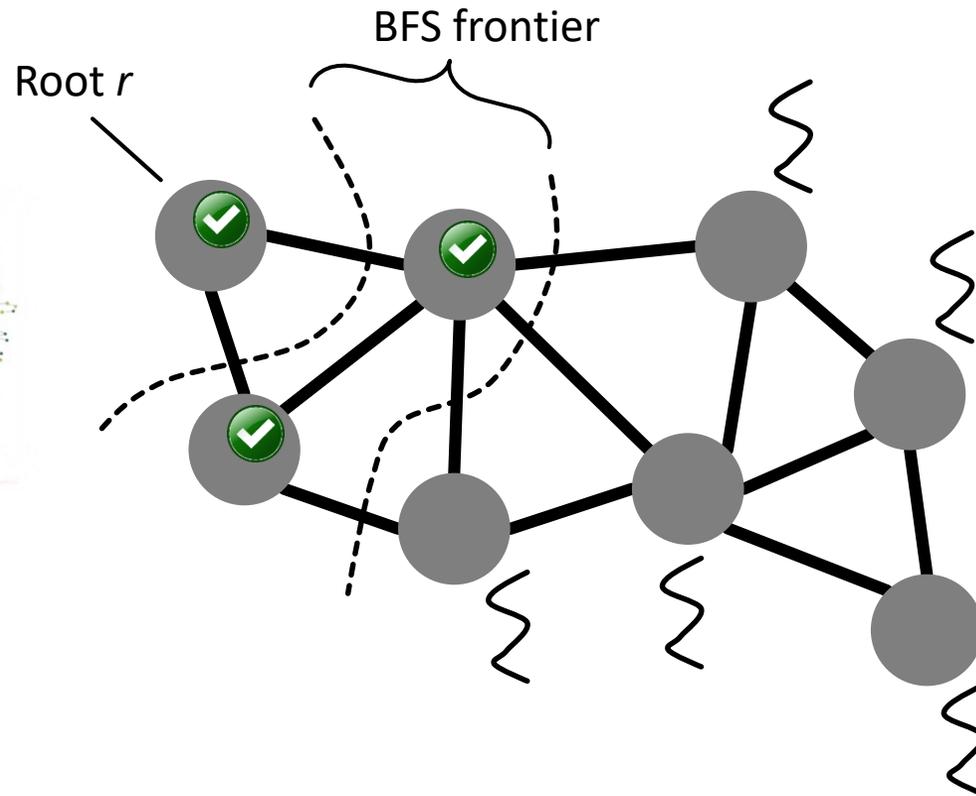


BFS

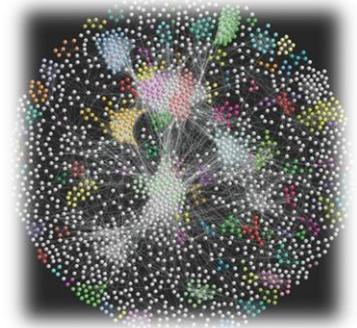
TOP-DOWN VS. BOTTOM-UP [1]



Pushing or pulling when expanding a frontier

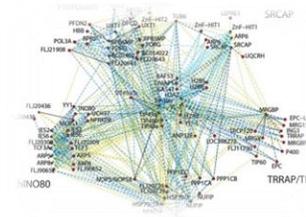
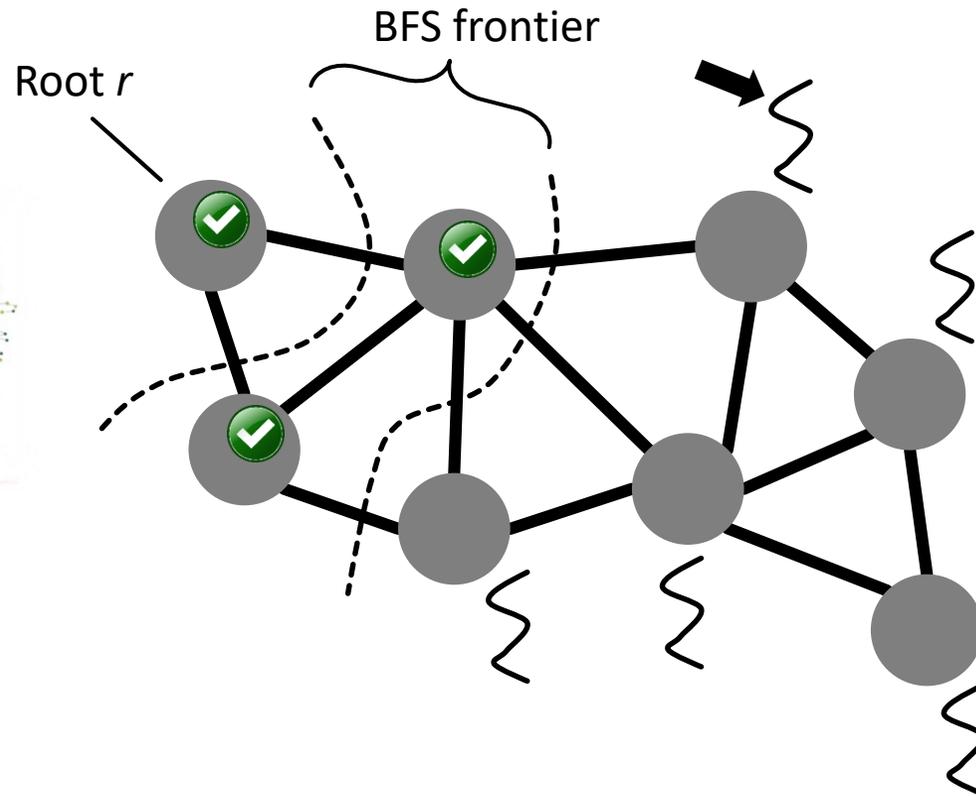
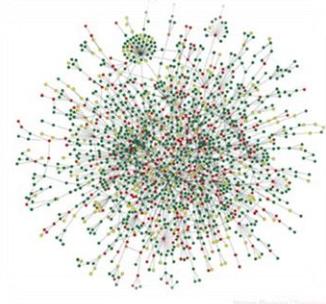
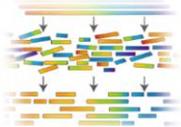


Pulling

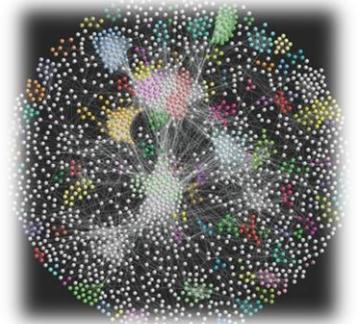


BFS

TOP-DOWN VS. BOTTOM-UP [1]



Pushing or pulling when expanding a frontier

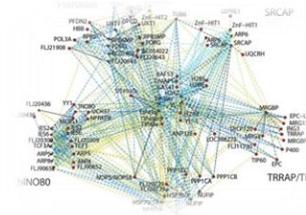
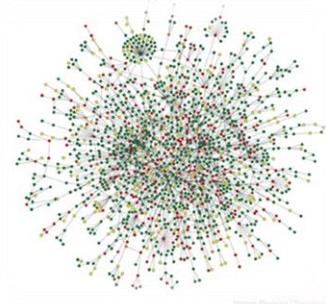
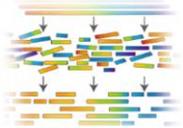


Pulling

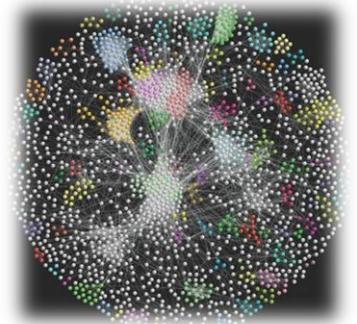
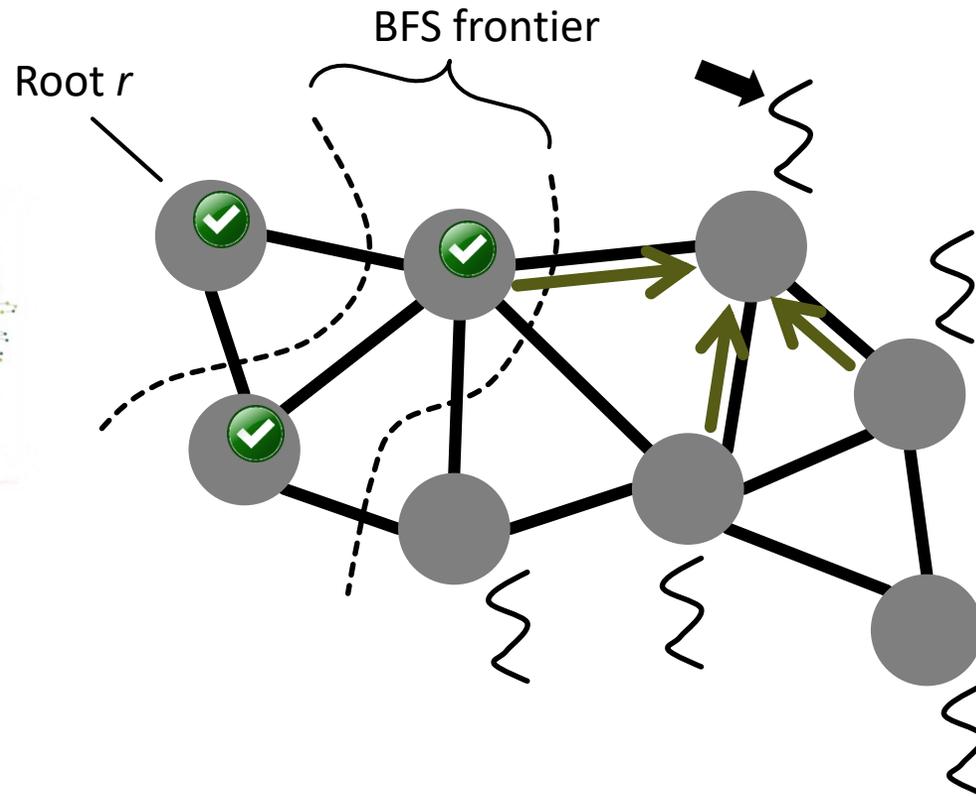


BFS

TOP-DOWN VS. BOTTOM-UP [1]



Pushing or pulling when expanding a frontier

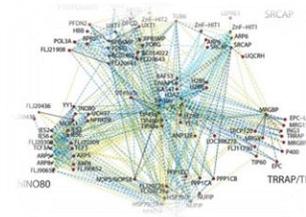
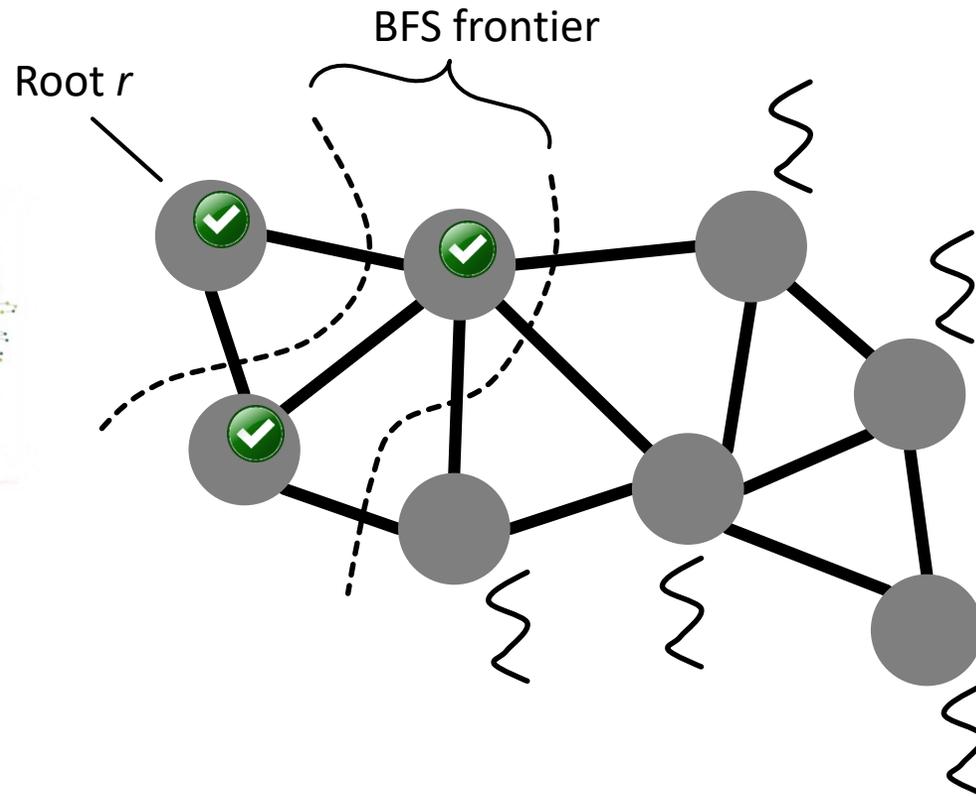
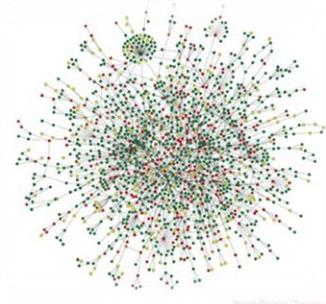
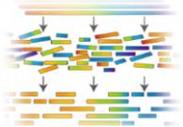


Pulling

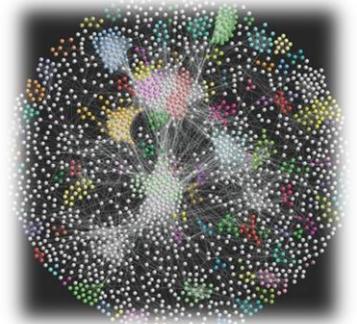


BFS

TOP-DOWN VS. BOTTOM-UP [1]



Pushing or pulling when expanding a frontier

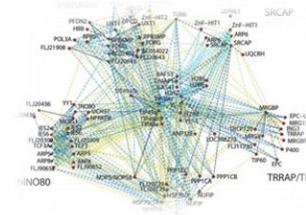
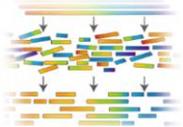


Pulling

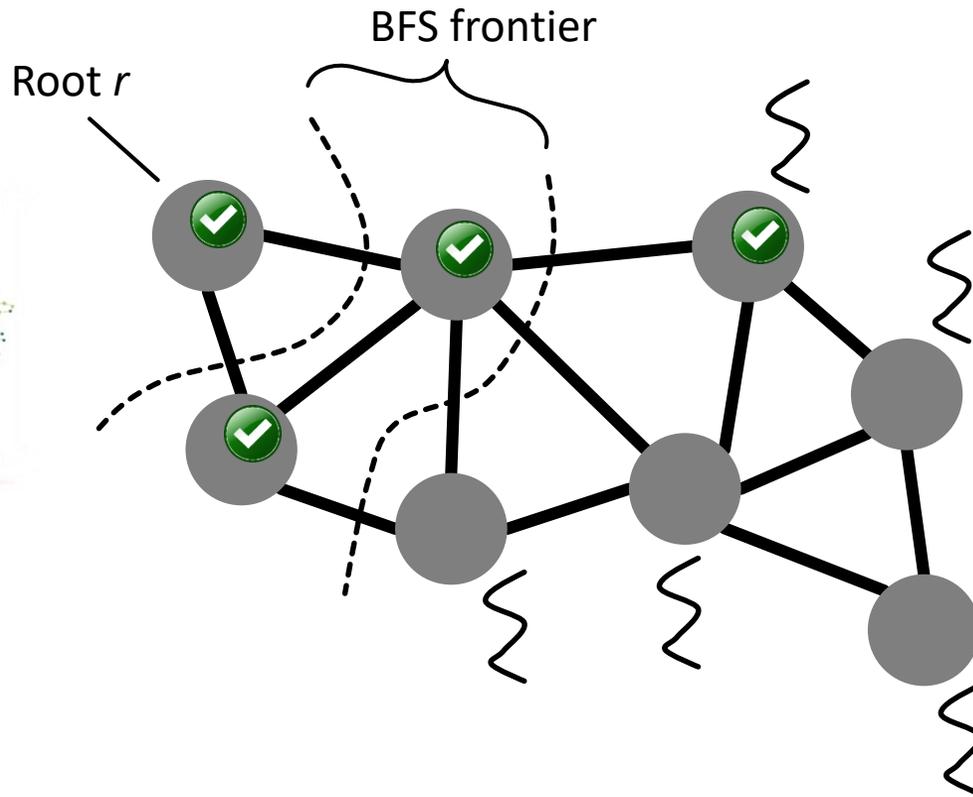
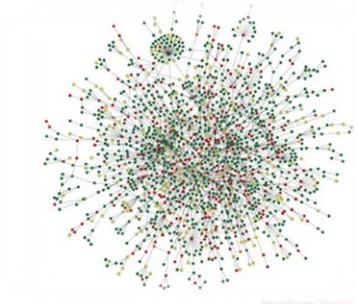


BFS

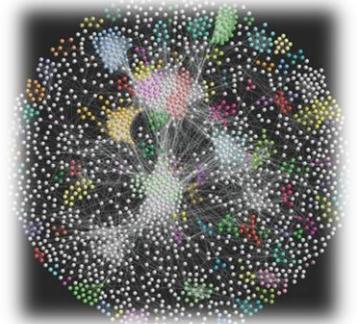
TOP-DOWN VS. BOTTOM-UP [1]



Pushing or pulling when expanding a frontier

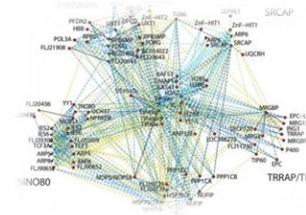
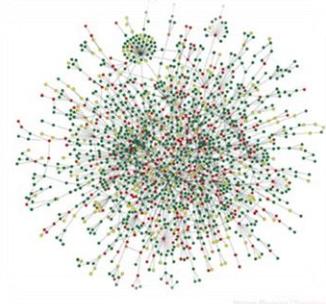
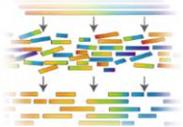


Pulling

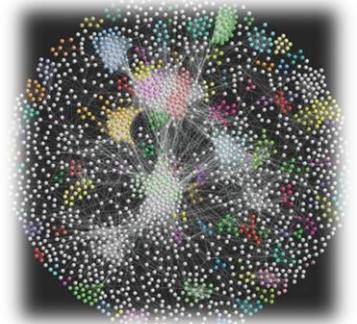
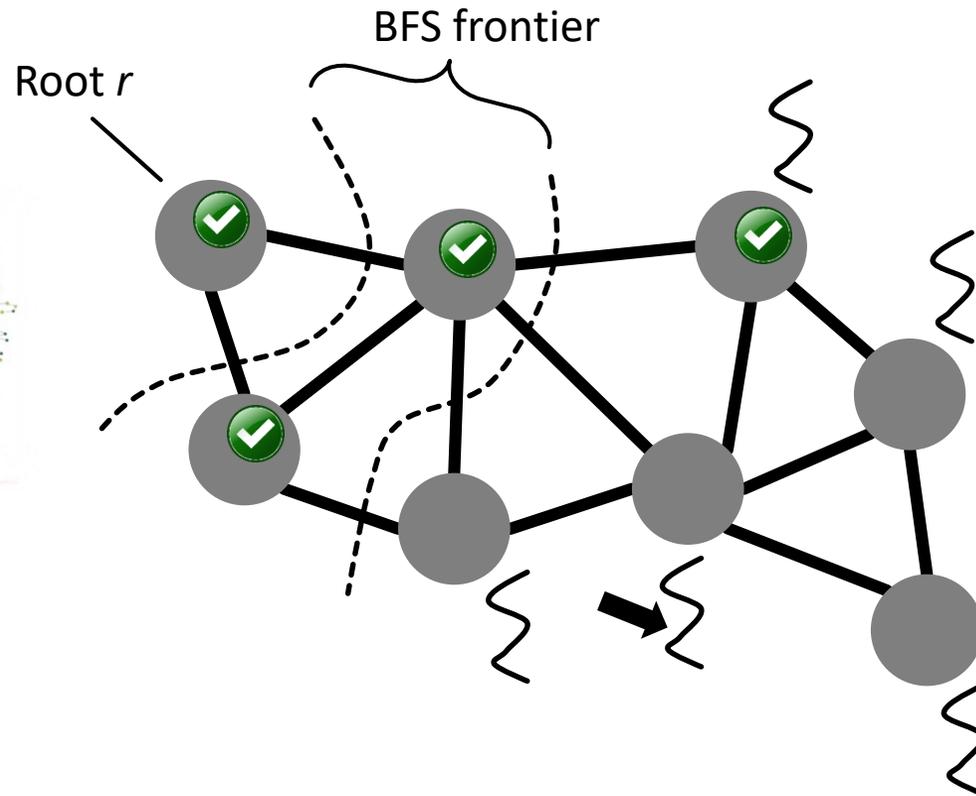


BFS

TOP-DOWN VS. BOTTOM-UP [1]



Pushing or pulling when expanding a frontier

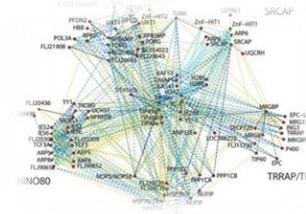
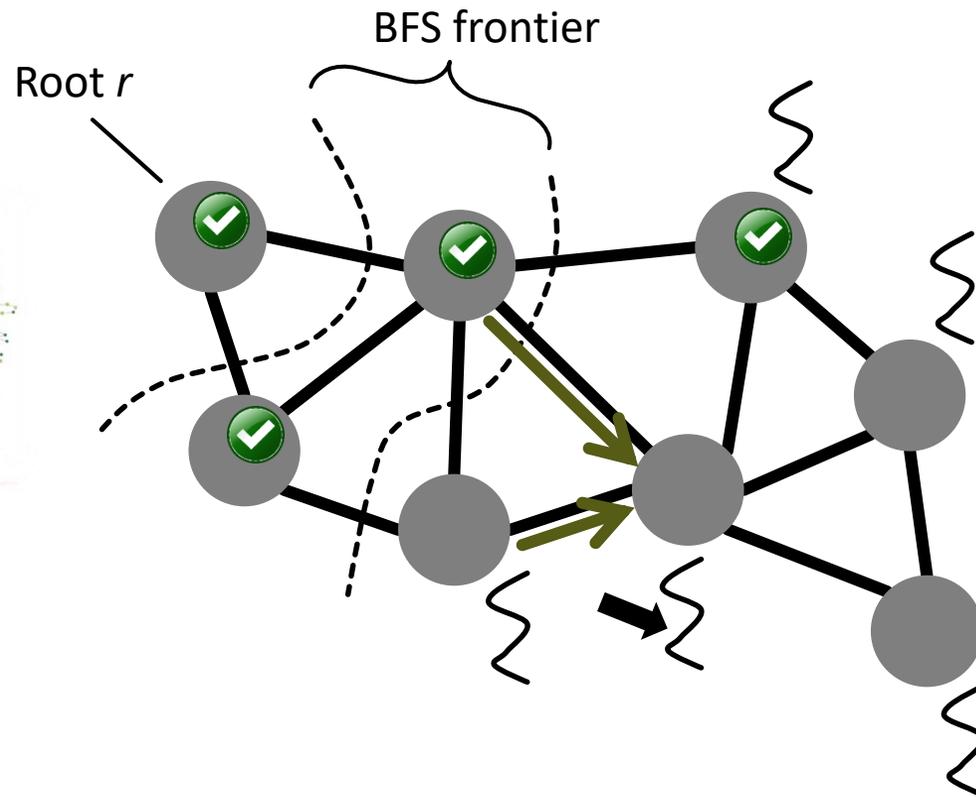
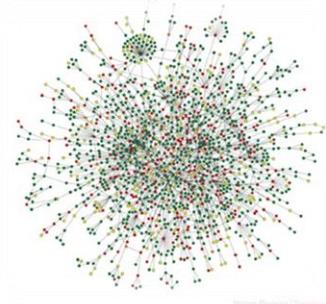
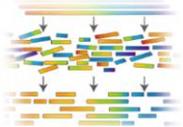


Pulling

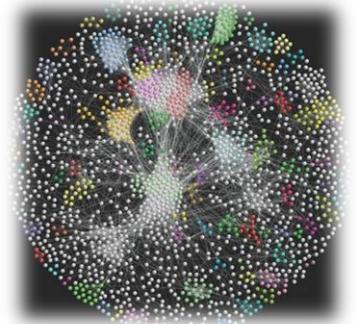


BFS

TOP-DOWN VS. BOTTOM-UP [1]



Pushing or pulling when expanding a frontier



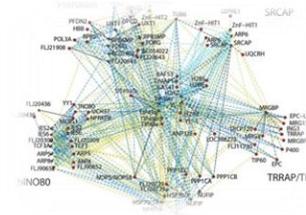
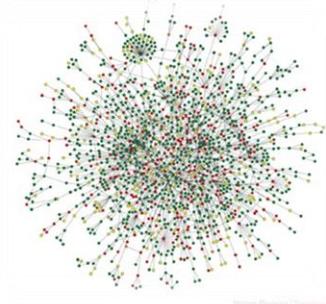
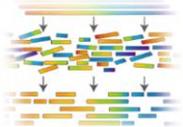
Pulling



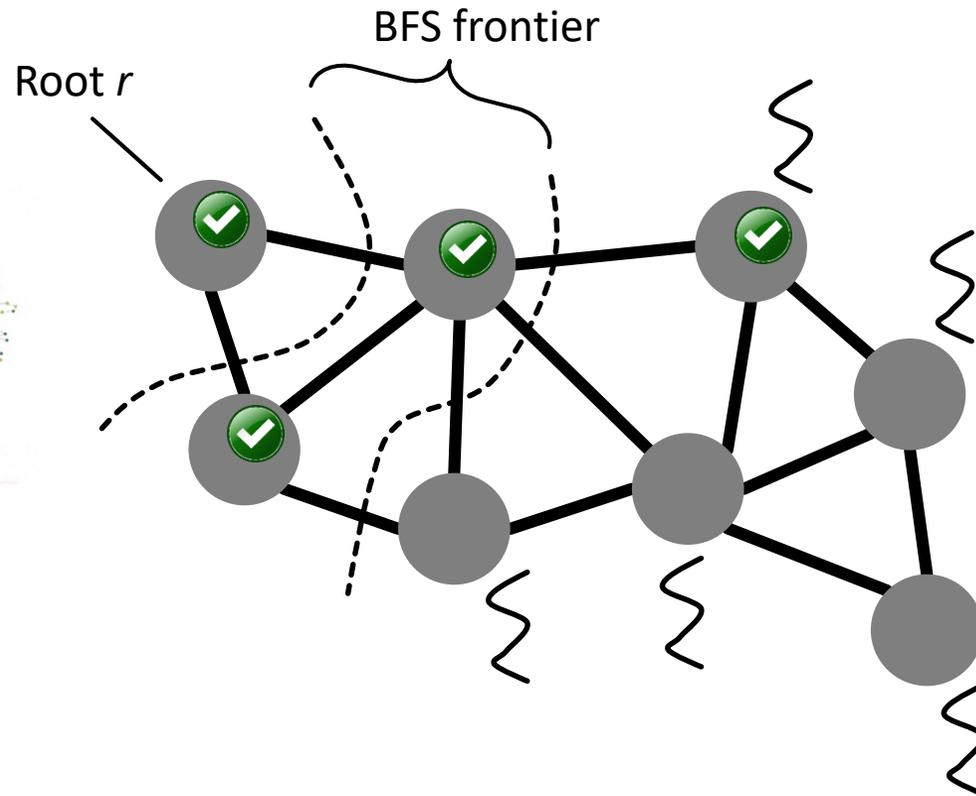
[1] S. Beamer, K. Asanović, and D. Patterson. Direction-optimizing breadth-first search. SC12.

BFS

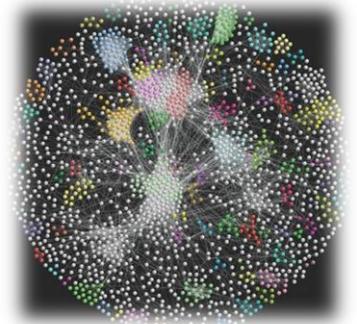
TOP-DOWN VS. BOTTOM-UP [1]



Pushing or pulling when expanding a frontier

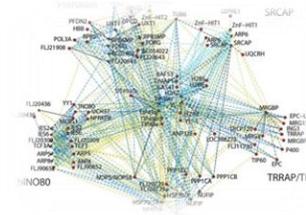
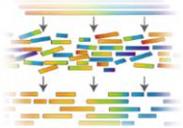


Pulling

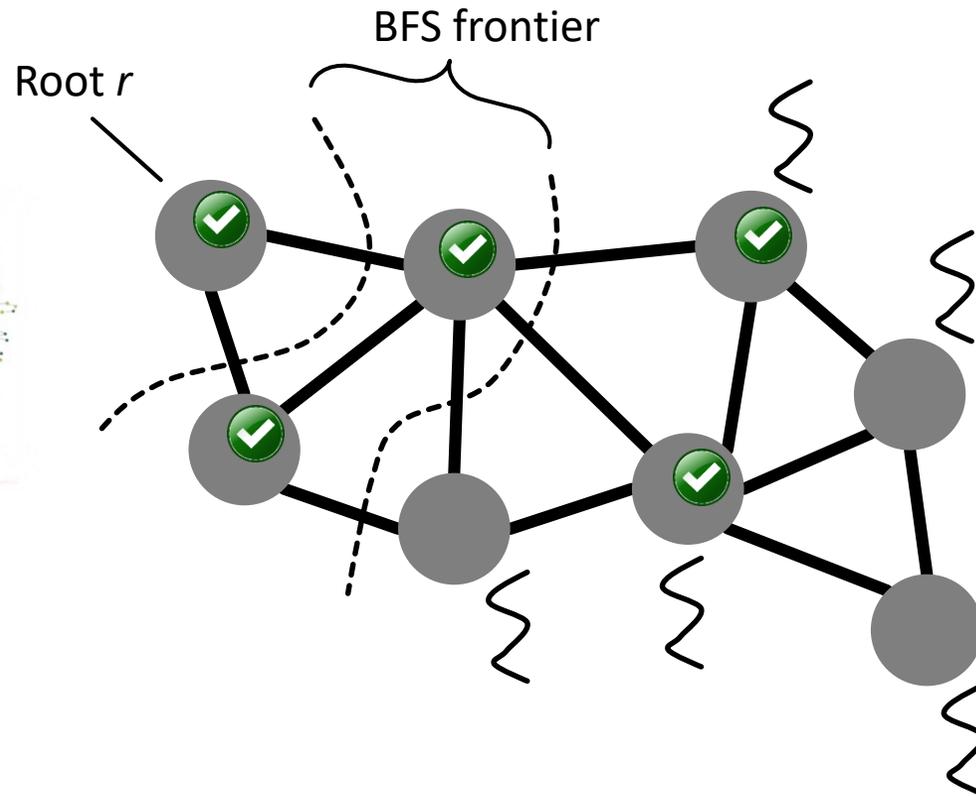
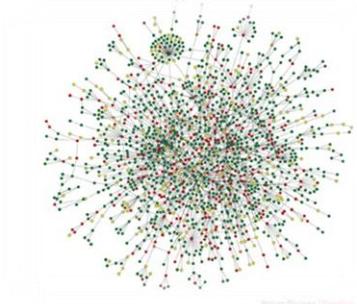


BFS

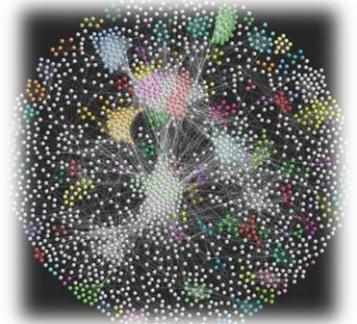
TOP-DOWN VS. BOTTOM-UP [1]



Pushing or pulling when expanding a frontier



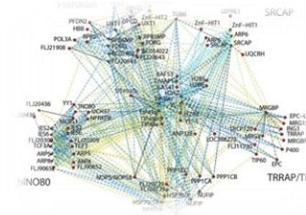
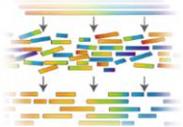
Pulling



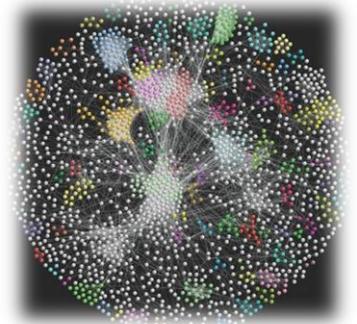
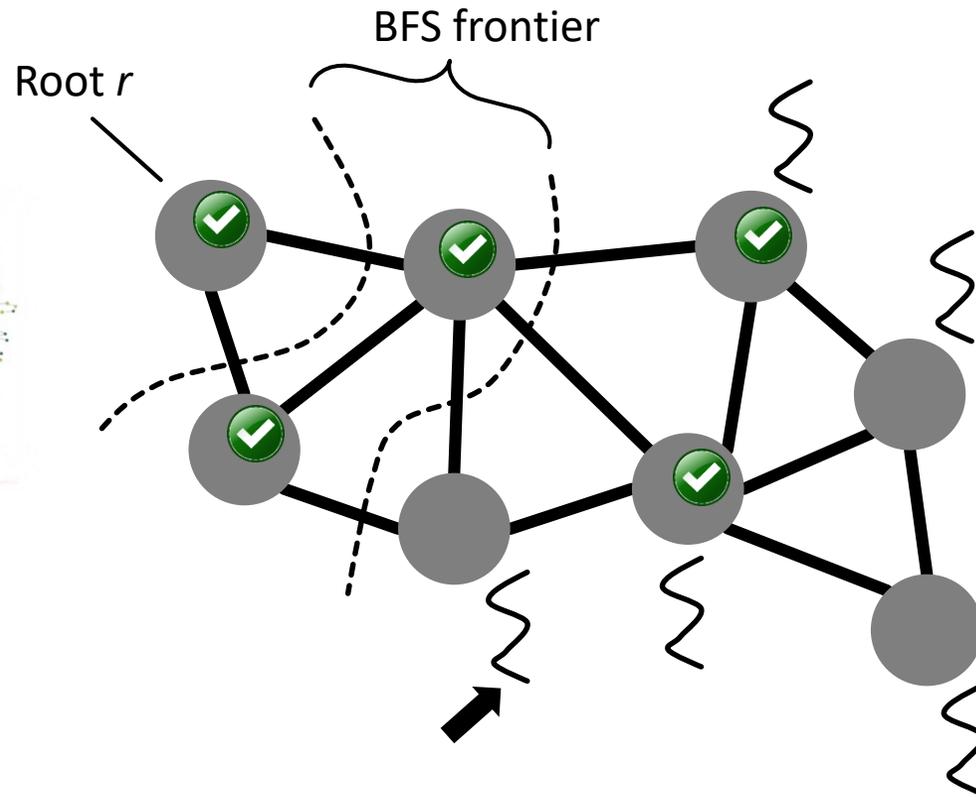
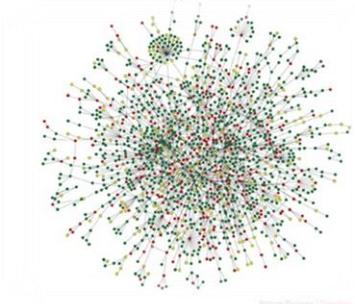
[1] S. Beamer, K. Asanović, and D. Patterson. Direction-optimizing breadth-first search. SC12.

BFS

TOP-DOWN VS. BOTTOM-UP [1]



Pushing or pulling when expanding a frontier

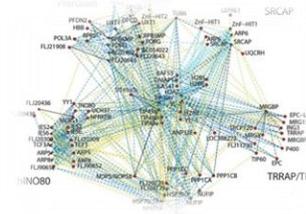
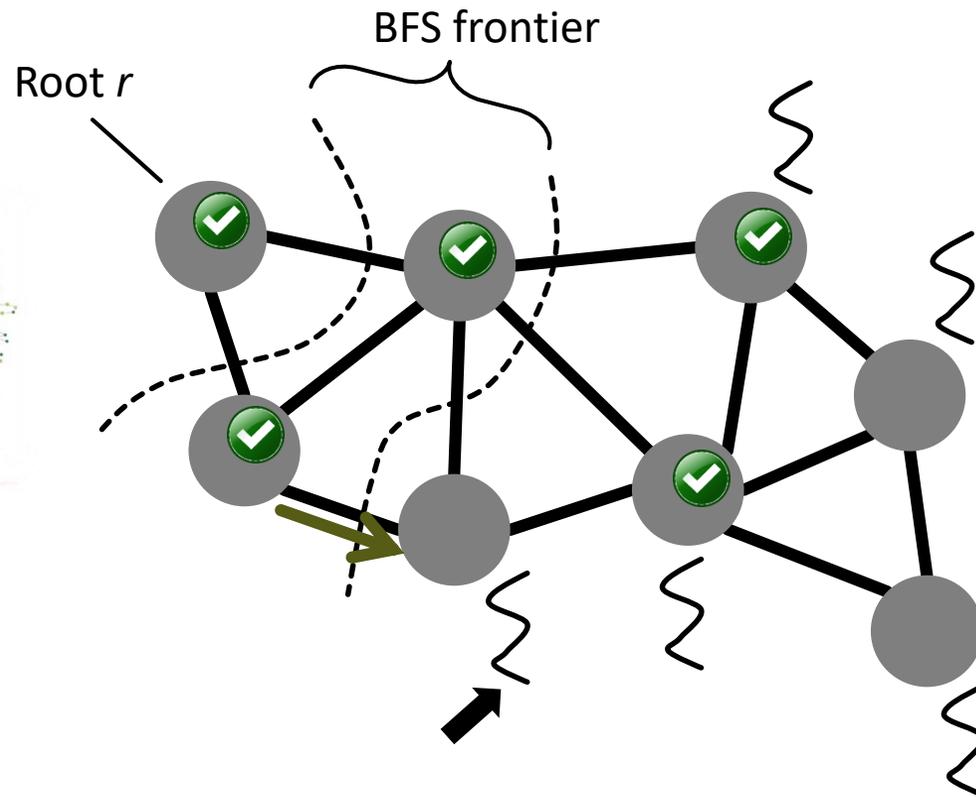
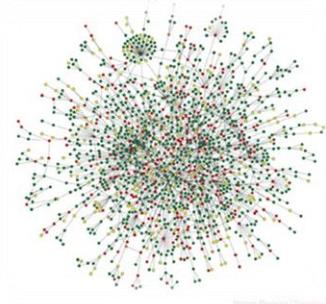
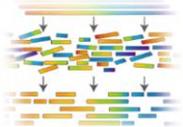


Pulling

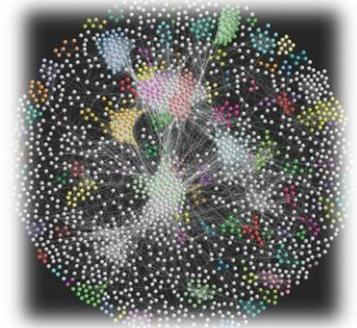


BFS

TOP-DOWN VS. BOTTOM-UP [1]



Pushing or pulling when expanding a frontier

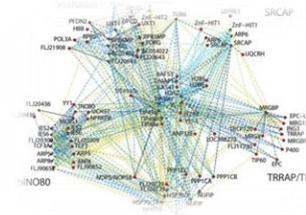
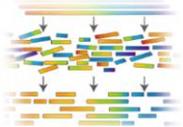


Pulling

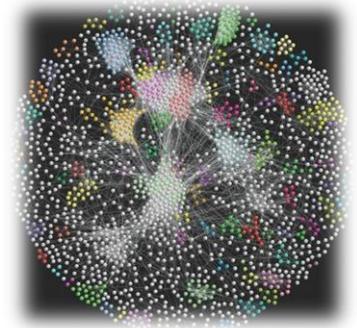
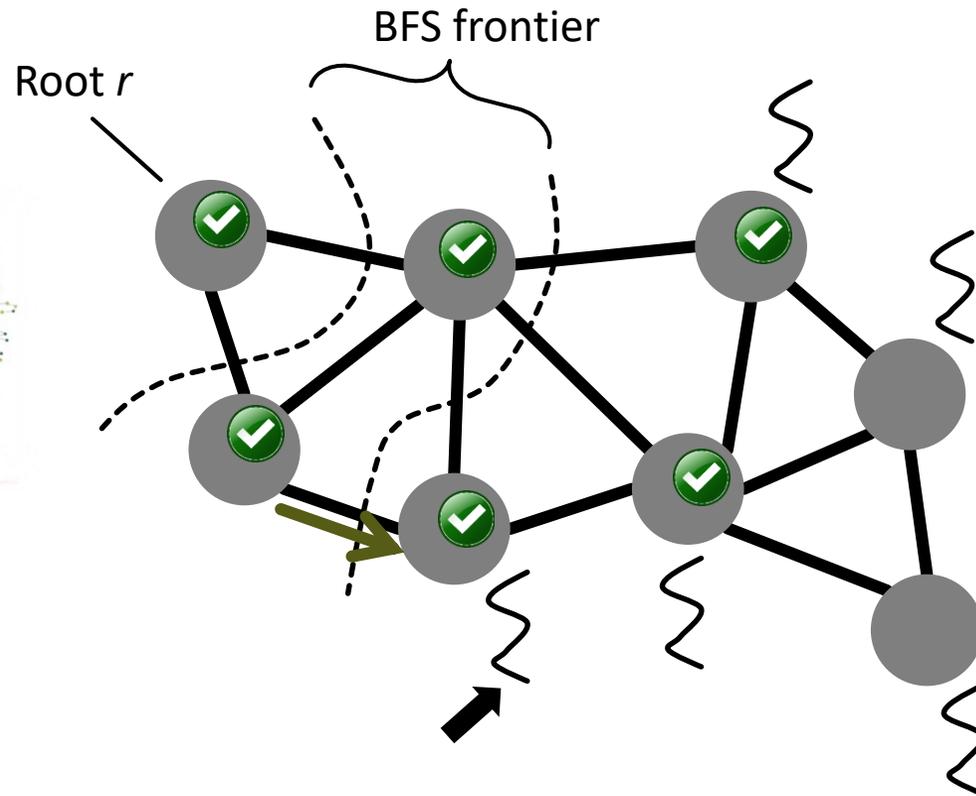
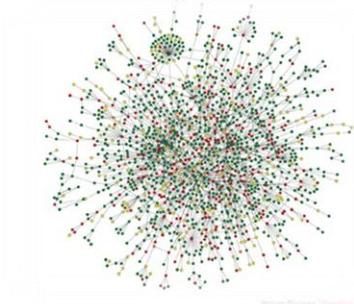


BFS

TOP-DOWN VS. BOTTOM-UP [1]



Pushing or pulling when expanding a frontier

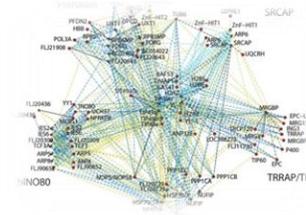
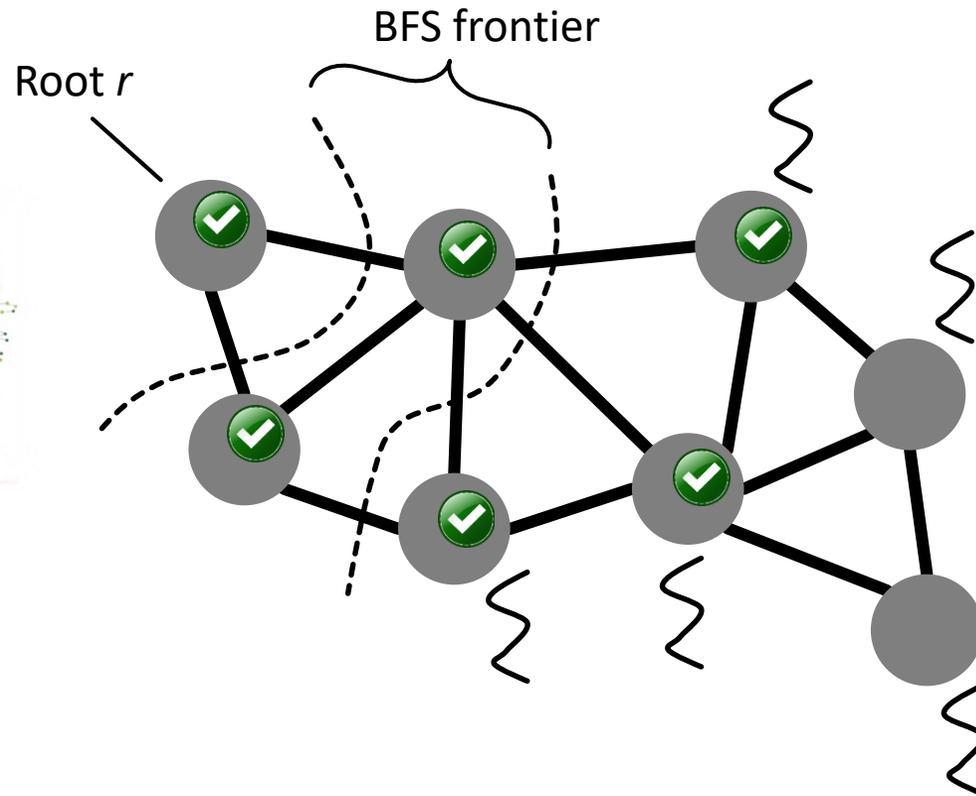
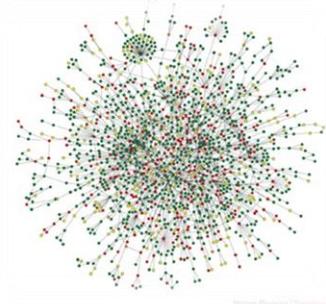
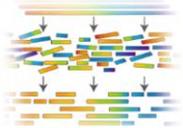


Pulling

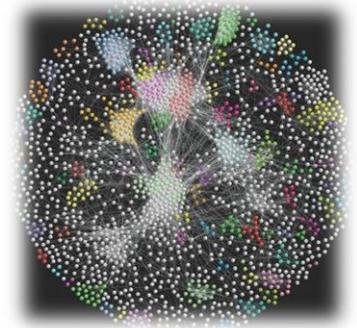


BFS

TOP-DOWN VS. BOTTOM-UP [1]



Pushing or pulling when expanding a frontier



Pulling



OTHER ALGORITHMS & FORMULATIONS

OTHER ALGORITHMS & FORMULATIONS

Triangle Counting

```
1 /* Input: a graph G. Output: An array
2 * tc[1..n] that each vertex belongs to
3
4 function TC(G) {tc[1..n] = [0..0]}
5   for v in V do in par
6     for w1 in N(v) do [in par]
7       for w2 in N(v) do [in par]
8         if adj(w1,w2) R update_tc();
9   tc[1..n] = [tc[1]/2 .. tc[n]/2]; }
10 function update_tc() {
11   {++tc[w1]; /* or ++tc[w2]. */} W I
12   {++tc[v];}
13 }
```

Δ-Stepping

```
1 /* Input: a graph G, a vertex r, the Δ parameter.
2 Output: An array of distances d */
3
4 function Δ-Stepping(G, r, Δ){
5   bckt=[∞..∞]; d=[∞..∞]; active=[false..false];
6   bckt_set={0}; bckt[r]=0; d[r]=0; active[r]=true; itr=0;
7
8   for b in bckt_set do { //For every bucket do...
9     do {bckt_empty = false; //Process b until it is empty.
10      process_buckets();} while(!bckt_empty); } }
11
12 function process_buckets() {
13   for v in bckt_set[b] do in par
14     if(bckt[v]==b && (itr == 0 or active[v])) {
15       active[v] = false; //Now, expand v's neighbors.
16       for w in N(v) {weight = d[v] + W(v,w);
17         if(weight < d[w]) { R //Proceed to relax w.
18           new_b = weight/Δ; bckt[v] = new_b;
19           bckt_set[new_b] = bckt_set[new_b] U {w};
20           d[w] = weight; W I;
21           if(bckt[w]==b) { active[w]=true; bckt_empty=true; } } R
22
23   for v in V do in par
24     if(d[v] > b) {for w in N(v) do {
25       if(bckt[w] == b && (active[w]
26         weight = d[w] + W(v,w) R;
27         if(weight < d[v]) {d[v]=weight
28         if(bckt[v] > new_b) {
```

BFS

```
1 /* Input: a graph G
2 R0 for each
3 * Output: B[1..n]
```

BC (algebraic notation)

```
1 /* Input: a graph G. Output: centrality scores bc[1..n].
2
3 function BC(G) { bc[1..n] = [0..0]}
4 Define Π so that any Π ⊃ u = (index_u, pred_u, mult_u, part_u);
5 Define u ← pred v with u, v ∈ Π so that u becomes
6   u = (index_u, pred_u ∪ index_v, mult_u + mult_v, part_u);
7 Define u ← part v with u, v ∈ Π so that u becomes
8   u = (index_u, pred_u, mult_u, part_u + (mult_u/mult_v)(1+part_v));
9
10 for s in V do [in par] {
11   ready = [1, ..., 1]; ready[s] = 0;
12   R = BFS(G, ready, [(1, 0, 0, 0)..(n, 0, 0, 0)] ← ready);
13   Define graph G' = (V, E') where (u, v) ∈ E'
14   Let ready[u] be the in-degree of u in V
15   R = BFS(G', ready, R, ← part);
16   for (index_u, pred_u, mult_u, part_u) ∈ R do [in
17     bc[u] += part_u; }
```

Betweenness Centrality (BC)

```
1 /* Input: a graph G. Output: centrality scores bc[1..n]. */
2
3 function BC(G) { bc[1..n] = [0..0]}
4 for s in V do [in par] {
5   for t in V do in par {
6     pred[t]=succ[t]=0; σ[t]=0; dist[t]=∞; PART 1: INITIALIZATION
7     σ[s]=enqueued=1; dist[s]=itr=0; δ[1..n]=[0..0]
8     Q[0]={s}; Q_1[1..p]=pred_1[1..p]=succ_1[1..p]=[0..0];
9   }
10   while enqueued > 0 do PART 2: COUNTING SHORTEST PATHS
11     count_shortest_paths();
12   --itr
13   while itr > 0 do PART 3: DEPENDENCY ACCUMULATION
14     accumulate_dependencies();
15   }
16 function count_shortest_paths() { enqueued = 0;
17 #if defined PUSHING_IN_PART_2
18   for v in Q[itr] do in par {
19     for w in N(v) do [in par] {
20       if dist[w] == ∞ R {
21         Q_1[itr + 1] = Q_1[itr + 1] U {w} I;
22         dist[w] = dist[v] + 1 W I; ++enqueued;
23         if dist[w] == dist[v] + 1 R {
24           σ[w] += σ[v]; pred_1[w] = pred_1[w] U {v};
25         }
26       }
27     }
28   }
29 #endif
30 }
```

PageRank

```
1 /* Input: a graph G, a number of steps L, the damp parameter f
2 Output: An array of ranks pr[1..n] */
3
4 function PR(G,L,f) {
5   pr[1..n] = [f..f]; //Initialize PR values.
6   for(l=1; l < L; ++l) {
7     new_pr[1..n] = [0..0];
8     for v in V do in par {
9       update_pr(); new_pr[v] += (1-f)/n; pr[v] = new_pr[v];
10    } }
11
12 function update_pr() {
13   for u in N(v) do [in par] {
14     {new_pr[u] += (f*pr[v])/d(v) W f;} PUSHING
15     {new_pr[v] += (f*pr[u])/d(u) R;} PULLING
16   } }
17 }
```

Boruvka MST

```
1 function MST_Boruvka(G) {
2   sv_flag=[1..v]; sv=[{1}..{v}]; MST=[0..0];
3   avail_svsvs={1..n}; max_e_wgt=max_{v,w in V} (W(v,w) + 1);
4
5   while avail_svsvs.size() > 0 do {avail_svsvs_new = 0;
6     for flag in avail_svsvs do in par {min_e_wgt[flag] = max_e_wgt;
7     for flag in avail_svsvs do in par {
8       for v in sv[flag] do {
9         for w in N(v) do [in par] {
10          if (sv_flag[w] ≠ flag) ∧
11             (W(v,w) < min_e_wgt[sv_flag[w]]) R {
12            min_e_wgt[sv_flag[w]] = W(v,w) W I;
13            min_e_v[sv_flag[w]] = w; min_e_w[sv_flag[w]] = w W I;
14            new_flag[sv_flag[w]] = flag W I; }
15          if (sv_flag[w] ≠ flag) ∧ (W(v,w) < min_e_wgt[flag]) R {
16            min_e_wgt[flag] = W(v,w); min_e_v[flag] = v;
17            min_e_w[flag] = w; new_flag[flag] = sv_flag[w]; } R
18        } } }
19     while flag = merge_order.pop() do {
20       neigh_flag = sv_flag[min_e_w[flag]];
21       for v in sv[flag] do sv_flag[flag] = sv_flag[neigh_flag];
22       sv[neigh_flag] = sv[flag] U sv[neigh_flag];
23       MST[neigh_flag] = MST[flag] U MST[neigh_flag]
24         U { (min_e_v[flag], min_e_w[flag]) }; }
```

Graph Coloring

```
1 // Input: a graph G. Output: An array of vertex colors c[1..n].
2 // In the code, the details of functions seq_color_partition and
3 // init are omitted due to space constrains.
4
5 function Boman-GC(G) {
6   done = false; c[1..n] = [0..0]; //No vertex is colored yet
7   //avail[i][j]=1 means that color j can be used for vertex i.
8   avail[1..n][1..C] = [1..1][1..1]; init(B, S);
9   while (!done) {
10    for P in S do in par {seq_color_partition(P);
11    fix_conflicts(); } }
12
13 function fix_conflicts() {
14   for v in B in par do {for u in N(v) do
15     if (c[u] == c[v]) {
16       {avail[u][c[v]] = 0 W I;} PUSHING
17       {avail[v][c[v]] = 0 R I;} PULLING
18     }
19   }
```

OTHER ALGORITHMS & FORMULATIONS

Triangle Counting

```

1 /* Input: a graph G. Output: An array
2 * tc[1..n] that each vertex belongs to
3
4 function TC(G) {tc[1..n] = [0..0]}
5   for v in V do in par
6     for w1 in N(v) do [in par
7       for w2 in N(v) do [in par
8         if adj(w1,w2) R update_tc();
9       tc[1..n] = [tc[1]/2 .. tc[n]/2]; }
10  function update_tc() {
11    {++tc[w1]; /* or ++tc[w2]. */} W I
12    {++tc[v];}
13  }

```

Δ -Stepping

```

1 /* Input: a graph G, a vertex r, the  $\Delta$  parameter.
2 Output: An array of distances d */
3
4 function  $\Delta$ -Stepping(G, r,  $\Delta$ ){
5   bckt=[ $\infty$ .. $\infty$ ]; d=[ $\infty$ .. $\infty$ ]; active=[false..false];
6   bckt_set={0}; bckt[r]=0; d[r]=0; active[r]=true; itr=0;
7
8   for b in bckt_set do { //For every bucket do...
9     do {bckt_empty = false; //Process b until it is empty.
10      process_buckets();} while (!bckt_empty); } }
11
12 function process_buckets() {
13   for v in bckt_set[b] do in par
14     if(bckt[v]==b && (itr == 0 or active[v])) {
15       active[v] = false; //Now, expand v's neighbors.
16       for w in N(v) {weight = d[v] + W(v,w);
17         if(weight < d[w]) { R //Proceed to relax w.
18           new_b = weight/ $\Delta$ ; bckt[v] = new_b;
19           bckt_set[new_b] = bckt_set[new_b] U {w};
20           d[w] = weight; W I;
21           if(bckt[w]==b) R {active[w]=true; bckt_empty=true;}} R
22   for v in V do in par
23     if(d[v] > b) {for w in N(v) do {
24       if(bckt[w] == b && (active[w]
25       weight = d[w] + W(v,w) R;
26       if(weight < d[v]) {d[v]=weight
27       if(bckt[v] > new_b) {

```

BFS

```

1 /* Input: a graph G. Output: centrality scores bc[1..n].
2 R0 for each
3 * Output: R[1..n]

```

BC (algebraic notation)

```

1 /* Input: a graph G. Output: centrality scores bc[1..n].
2
3 function BC(G) { bc[1..n] = [0..0]}
4 Define  $\Pi$  so that any  $\Pi \ni u = (\text{index}_u, \text{pred}_u, \text{mult}_u, \text{part}_u)$ ;
5 Define  $u \Leftarrow \text{pred } v$  with  $u, v \in \Pi$  so that  $u$  becomes
6    $u = (\text{index}_u, \text{pred}_u \cup \text{index}_v, \text{mult}_u + \text{mult}_v, \text{part}_u)$ ;
7 Define  $u \Leftarrow \text{part } v$  with  $u, v \in \Pi$  so that  $u$  becomes
8    $u = (\text{index}_u, \text{pred}_u, \text{mult}_u, \text{part}_u + (\text{mult}_u / \text{mult}_v)(1 + \text{part}_v))$ ;
9
10 for s in V do [in par] {
11   ready = [1, ..., 1]; ready[s] = 0;
12   R = BFS(G, ready, [(1, 0, 0, 0)..(n, 0, 0, 0)]);
13   Define graph  $G' = (V, E')$  where  $(u, v) \in E'$ 
14   Let ready[u] be the in-degree of  $u \in V$ 
15   R = BFS( $G'$ , ready, R,  $\Leftarrow \text{part}$ );
16   for  $(\text{index}_u, \text{pred}_u, \text{mult}_u, \text{part}_u) \in R$  do [in
17     bc[u] += part_u; }

```

Betweenness Centrality (BC)

```

1 /* Input: a graph G. Output: centrality scores bc[1..n]. */
2
3 function BC(G) { bc[1..n] = [0..0]}
4 for s in V do [in par] {
5   for t in V do in par {
6     pred[t]=succ[t]=0;  $\sigma$ [t]=0; dist[t]= $\infty$ ; PART 1: INITIALIZATION
7      $\sigma$ [s]=enqueued=1; dist[s]=itr=0;  $\delta$ [1..n]=[0..0]
8     Q[0]={s}; Q_1[1..p]=pred_1[1..p]=succ_1[1..p]=[0..0];
9   }
10  while enqueued > 0 do PART 2: COUNTING SHORTEST PATHS
11    count_shortest_paths();
12  --itr
13  while itr > 0 do PART 3: DEPENDENCY ACCUMULATION
14    accumulate_dependencies();
15  }
16 function count_shortest_paths() { enqueued = 0;
17 #if defined PUSHING_IN_PART_2
18   for v in Q[itr] do in par {
19     for w in N(v) do [in par] {
20       if dist[w] ==  $\infty$  R {
21         Q_1[itr + 1] = Q_1[itr + 1] U {w} I;
22         dist[w] = dist[v] + 1 W I; ++enqueued;
23         if dist[w] == dist[v] + 1 R {
24            $\sigma$ [w] +=  $\sigma$ [v] W I; pred_1[w] = pred_1[v] U {v};

```

Graph Coloring

```

1 // Input: a graph G. Output: An array of vertex colors c[1..n].
2 // In the code, the details of functions seq_color_partition and
3 // init are omitted due to space constrains.
4
5 function Boman-GC(G) {
6   done = false; c[1..n] = [0..0]; //No vertex is colored yet
7   //avail[i][j]=1 means that color j can be used for vertex i.
8   avail[1..n][1..C] = [1..1][1..1]; init(B, S);
9   while (!done) {
10     for P in S do in par {seq_color_partition(P);}
11     fix_conflicts(); } }
12
13 function fix_conflicts() {
14   for v in B in par do {for u in N(v) do
15     if (c[u] == c[v]) {
16       {avail[u][c[v]] = 0} W I; PUSHING
17       {avail[v][c[v]] = 0} R I; PULLING
18     }
19   }

```

PageRank

```

1 /* Input: a graph G, a number of steps L, the damp parameter f
2 Output: An array of ranks pr[1..n] */
3
4 function PR(G,L,f) {
5   pr[1..v] = [f..f]
6   for(l=1; l < L; l++)
7     new_pr[1..n] =
8     for v in V do in par
9       update_pr(v);
10  } }
11
12 function update_pr(v) {
13   for u in N(v) do [
14     {new_pr[u] += (f*pr[v])/d(u)} W I; PUSHING
15     {new_pr[v] += (f*pr[u])/d(u)} R; PULLING
16  ] }
17 }

```

Boruvka MST

```

1 function MST_Boruvka(G) {
2   sv_flag=[1..v]; sv=[{1}..{v}]; MST=[0..0];
3   avail_svsvs={1..n}; max_e_wgt=max_{v,w in V} (W(v,w) + 1);
4   while avail_svsvs.size() > 0 do {avail_svsvs_new = 0;
5     for flag in avail_svsvs do in par {min_e_wgt[flag] = max_e_wgt;
6     for v in V do in par {
7       par {
8         flag  $\wedge$  (W(v,w) < min_e_wgt[flag]) R {
9         min_e_wgt[flag] = W(v,w) W I; PUSHING
10        min_e_wgt[flag] = W(v,w); min_e_v[flag] = v; W I;
11        min_e_wgt[flag] = W(v,w); min_e_v[flag] = v; W I;
12        min_e_wgt[flag] = W(v,w); min_e_v[flag] = v; W I;
13        min_e_wgt[flag] = W(v,w); min_e_v[flag] = v; W I;
14        min_e_wgt[flag] = W(v,w); min_e_v[flag] = v; W I;
15        min_e_wgt[flag] = W(v,w); min_e_v[flag] = v; W I;
16        min_e_wgt[flag] = W(v,w); min_e_v[flag] = v; W I;
17        min_e_wgt[flag] = W(v,w); min_e_v[flag] = v; W I;
18        min_e_wgt[flag] = W(v,w); min_e_v[flag] = v; W I;
19        min_e_wgt[flag] = W(v,w); min_e_v[flag] = v; W I;
20        min_e_wgt[flag] = W(v,w); min_e_v[flag] = v; W I;
21        min_e_wgt[flag] = W(v,w); min_e_v[flag] = v; W I;
22        min_e_wgt[flag] = W(v,w); min_e_v[flag] = v; W I;
23        min_e_wgt[flag] = W(v,w); min_e_v[flag] = v; W I;
24        min_e_wgt[flag] = W(v,w); min_e_v[flag] = v; W I;

```

Check out the paper 😊

PUSHING VS. PULLING

GENERIC DIFFERENCES

PUSHING VS. PULLING

GENERIC DIFFERENCES



What pushing vs.
pulling *really* is?

PUSHING VS. PULLING

GENERIC DIFFERENCES



What pushing vs.
pulling *really* is?

- Vertices: $v \in V$
- $t \rightsquigarrow v \Leftrightarrow t$ modifies v
- $t[v]$: a thread that owns v

PUSHING VS. PULLING

GENERIC DIFFERENCES



What pushing vs.
pulling *really* is?

Algorithm uses pushing \Leftrightarrow
 $(\exists t \exists v \in V: t \rightsquigarrow v \wedge t \neq t[v])$

- Vertices: $v \in V$
- $t \rightsquigarrow v \Leftrightarrow t$ modifies v
- $t[v]$: a thread that owns v

PUSHING VS. PULLING

GENERIC DIFFERENCES



What pushing vs.
pulling *really* is?

Algorithm uses pushing \Leftrightarrow

$$(\exists t \exists v \in V: t \rightsquigarrow v \wedge t \neq t[v])$$

Algorithm uses pulling \Leftrightarrow

$$(\forall t \forall v \in V: t \rightsquigarrow v \Rightarrow t = t[v])$$

- Vertices: $v \in V$
- $t \rightsquigarrow v \Leftrightarrow t$ modifies v
- $t[v]$: a thread that owns v

PUSHING VS. PULLING

GENERIC DIFFERENCES



What pushing vs. pulling *really* is?

Algorithm uses pushing \Leftrightarrow
 $(\exists t \exists v \in V: t \rightsquigarrow v \wedge t \neq t[v])$

Algorithm uses pulling \Leftrightarrow
 $(\forall t \forall v \in V: t \rightsquigarrow v \Rightarrow t = t[v])$

- Vertices: $v \in V$
- $t \rightsquigarrow v \Leftrightarrow t$ modifies v
- $t[v]$: a thread that owns v



This *is* the actual dichotomy

PUSHING VS. PULLING

GENERIC DIFFERENCES



What pushing vs. pulling *really* is?

- Vertices: $v \in V$
- $t \rightsquigarrow v \Leftrightarrow t$ modifies v
- $t[v]$: a thread that owns v

Algorithm uses pushing \Leftrightarrow
 $\sim \left[(\exists t \exists v \in V: t \rightsquigarrow v \wedge t \neq t[v]) \right]$

Algorithm uses pulling \Leftrightarrow
 $(\forall t \forall v \in V: t \rightsquigarrow v \Rightarrow t = t[v])$



This *is* the actual dichotomy

PUSHING VS. PULLING

GENERIC DIFFERENCES



What pushing vs. pulling *really* is?

- Vertices: $v \in V$
- $t \rightsquigarrow v \Leftrightarrow t$ modifies v
- $t[v]$: a thread that owns v

$$\sim \left[\begin{array}{l} \text{Algorithm uses pushing} \Leftrightarrow \\ (\exists t \exists v \in V: t \rightsquigarrow v \wedge t \neq t[v]) \end{array} \right]$$

$$\parallel$$

$$\begin{array}{l} \text{Algorithm uses pulling} \Leftrightarrow \\ (\forall t \forall v \in V: t \rightsquigarrow v \Rightarrow t = t[v]) \end{array}$$



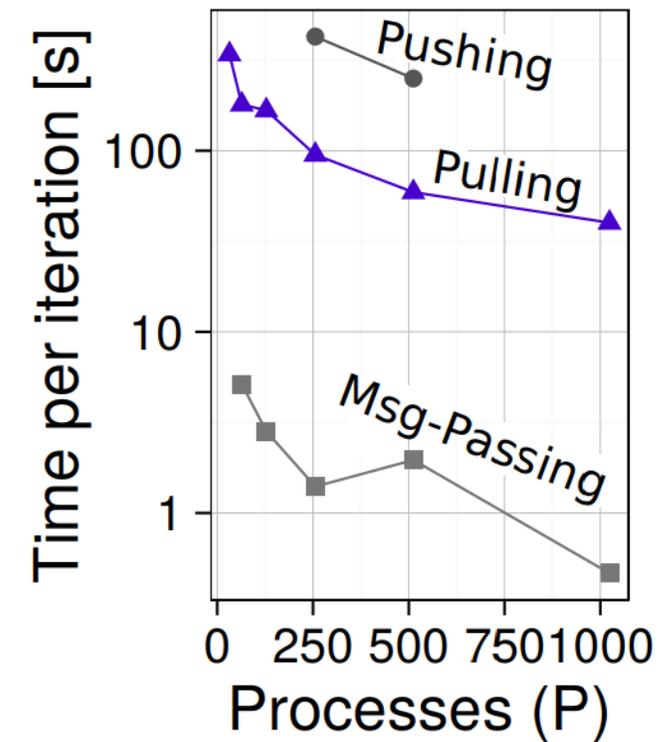
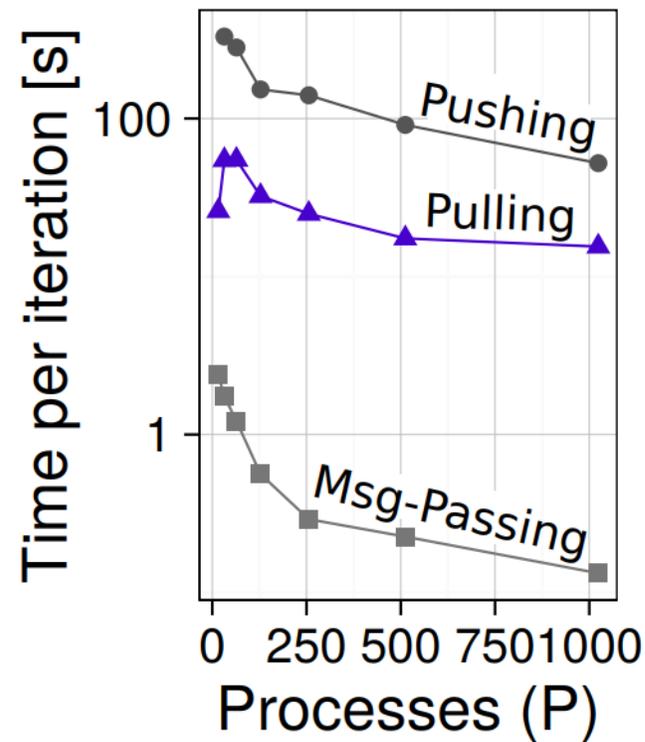
This *is* the actual dichotomy

PERFORMANCE ANALYSIS

PAGERANK

Kronecker graphs

Distributed-Memory



PERFORMANCE ANALYSIS

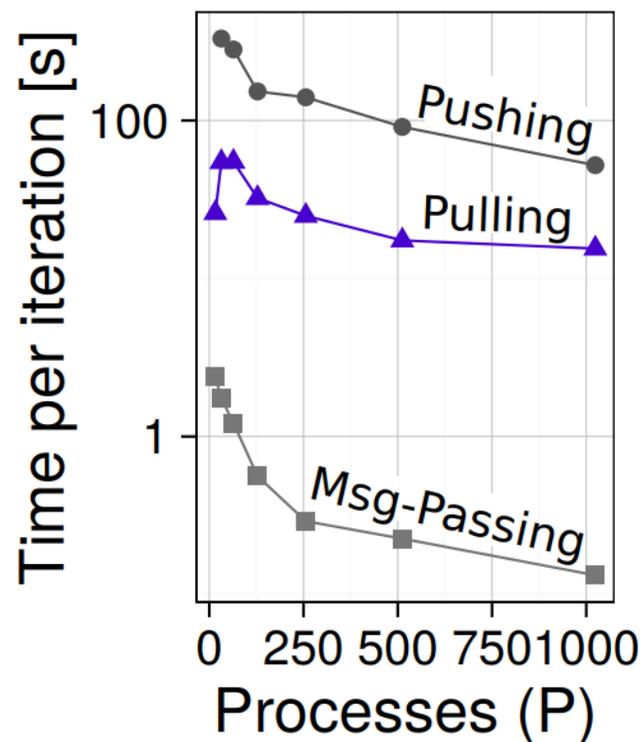
PAGERANK

Kronecker graphs

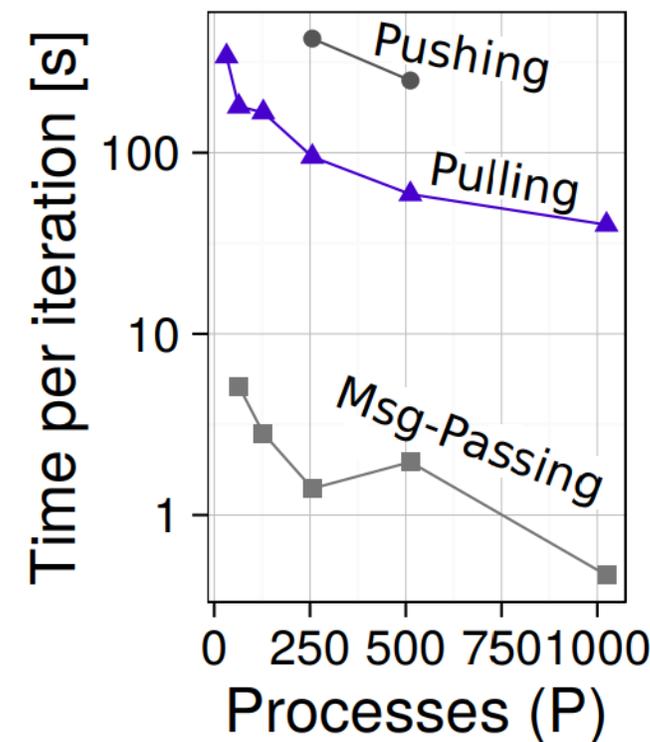
Distributed-Memory



$n = 2^{25}, m = 2^{27}$



$n = 2^{27}, m = 2^{29}$



PERFORMANCE ANALYSIS

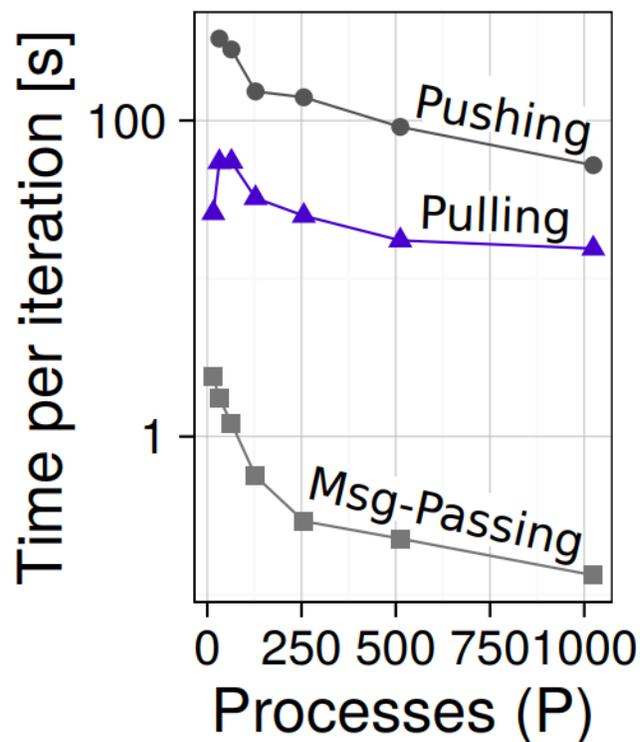
PAGERANK

Kronecker graphs

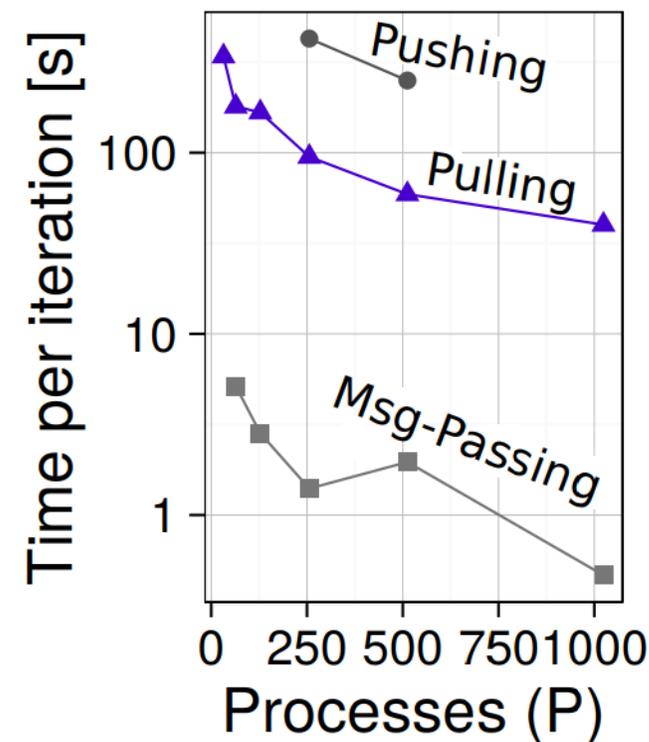
Distributed-Memory

! Msg-Passing fastest

$$n = 2^{25}, m = 2^{27}$$



$$n = 2^{27}, m = 2^{29}$$



PERFORMANCE ANALYSIS

PAGERANK

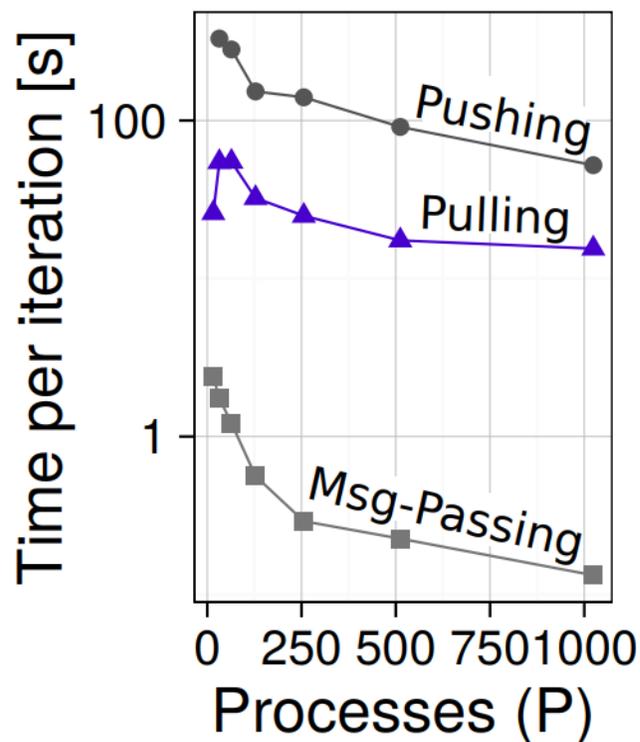
Kronecker graphs

Distributed-Memory

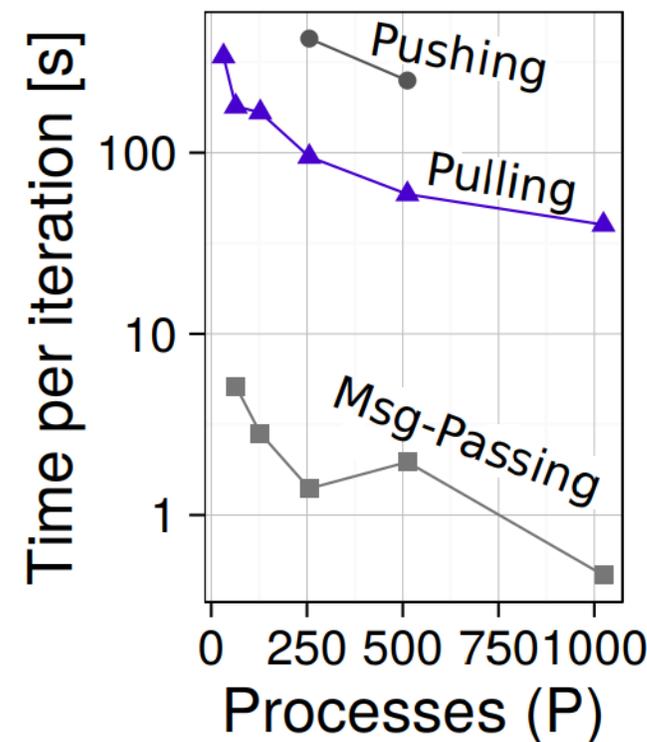
! Msg-Passing fastest

Pulling incurs more communication while pushing expensive underlying locking

$n = 2^{25}, m = 2^{27}$



$n = 2^{27}, m = 2^{29}$



PERFORMANCE ANALYSIS

PAGERANK

Kronecker graphs

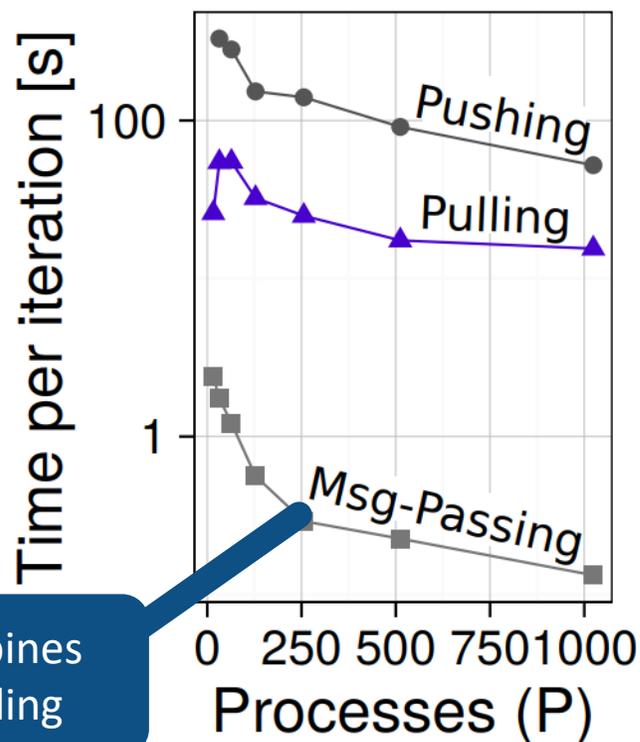
Distributed-Memory

! Msg-Passing fastest

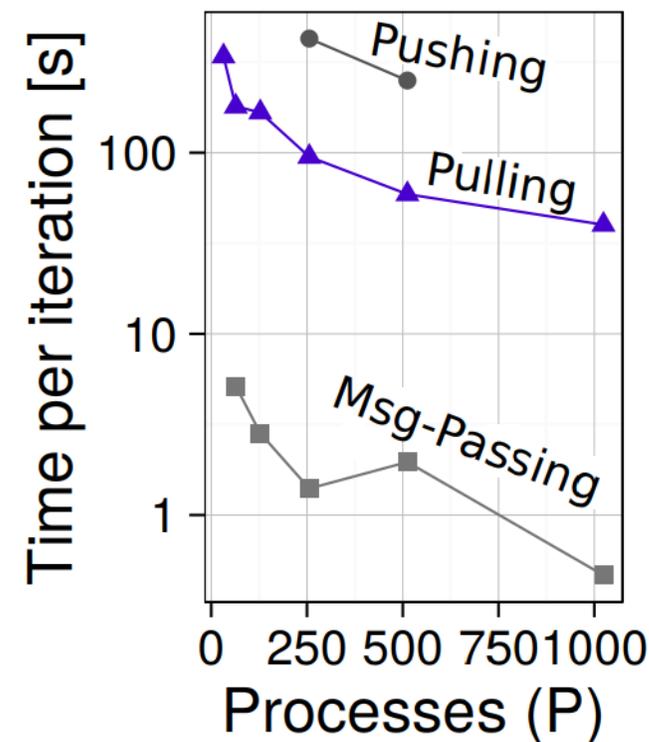
Pulling incurs more communication while pushing expensive underlying locking

! Collectives: combines pushing and pulling

$n = 2^{25}, m = 2^{27}$



$n = 2^{27}, m = 2^{29}$



To Push or To Pull?

To Push or To Pull?

If the complexities
match: pull

To Push or To Pull?

If the complexities
match: pull

Otherwise: push

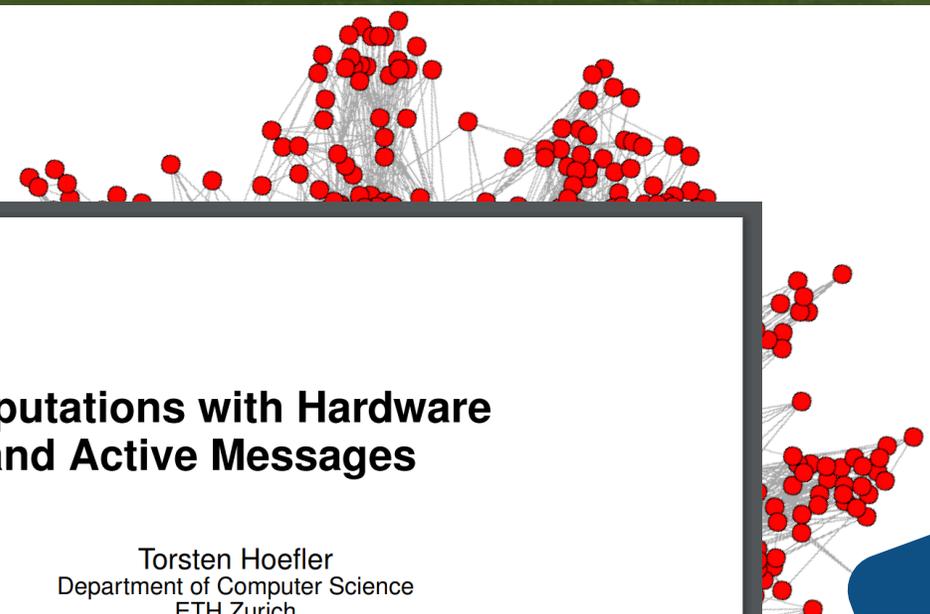
To Push or To Pull?

If the complexities
match: pull

Otherwise: push

+ check your
hardware 😊

Moving on ...



ynchronization-heavy

Accelerating Irregular Computations with Hardware Transactional Memory and Active Messages

Maciej Besta
Department of Computer Science
ETH Zurich
Universitätsstr. 6, 8092 Zurich

Torsten Hoefler
Department of Computer Science
ETH Zurich

To Push or To Synchron

Maciej Besta¹, Michał
¹ Department of Computer Science
⁴ Departm
maciej.bestam@inf.ethz.ch, michal.p

ABSTRACT

We reduce the cost of communication and processing by analyzing the fastest way to the updates to a shared state or pulling state. We investigate the applicability of

High-Performance Distributed RMA Locks

Patrick Schmid*
Department of Computer Science
ETH Zurich
patrick.schmid@ieffects.com

Maciej Besta*
Department of Computer Science
ETH Zurich
bestam@inf.ethz.ch

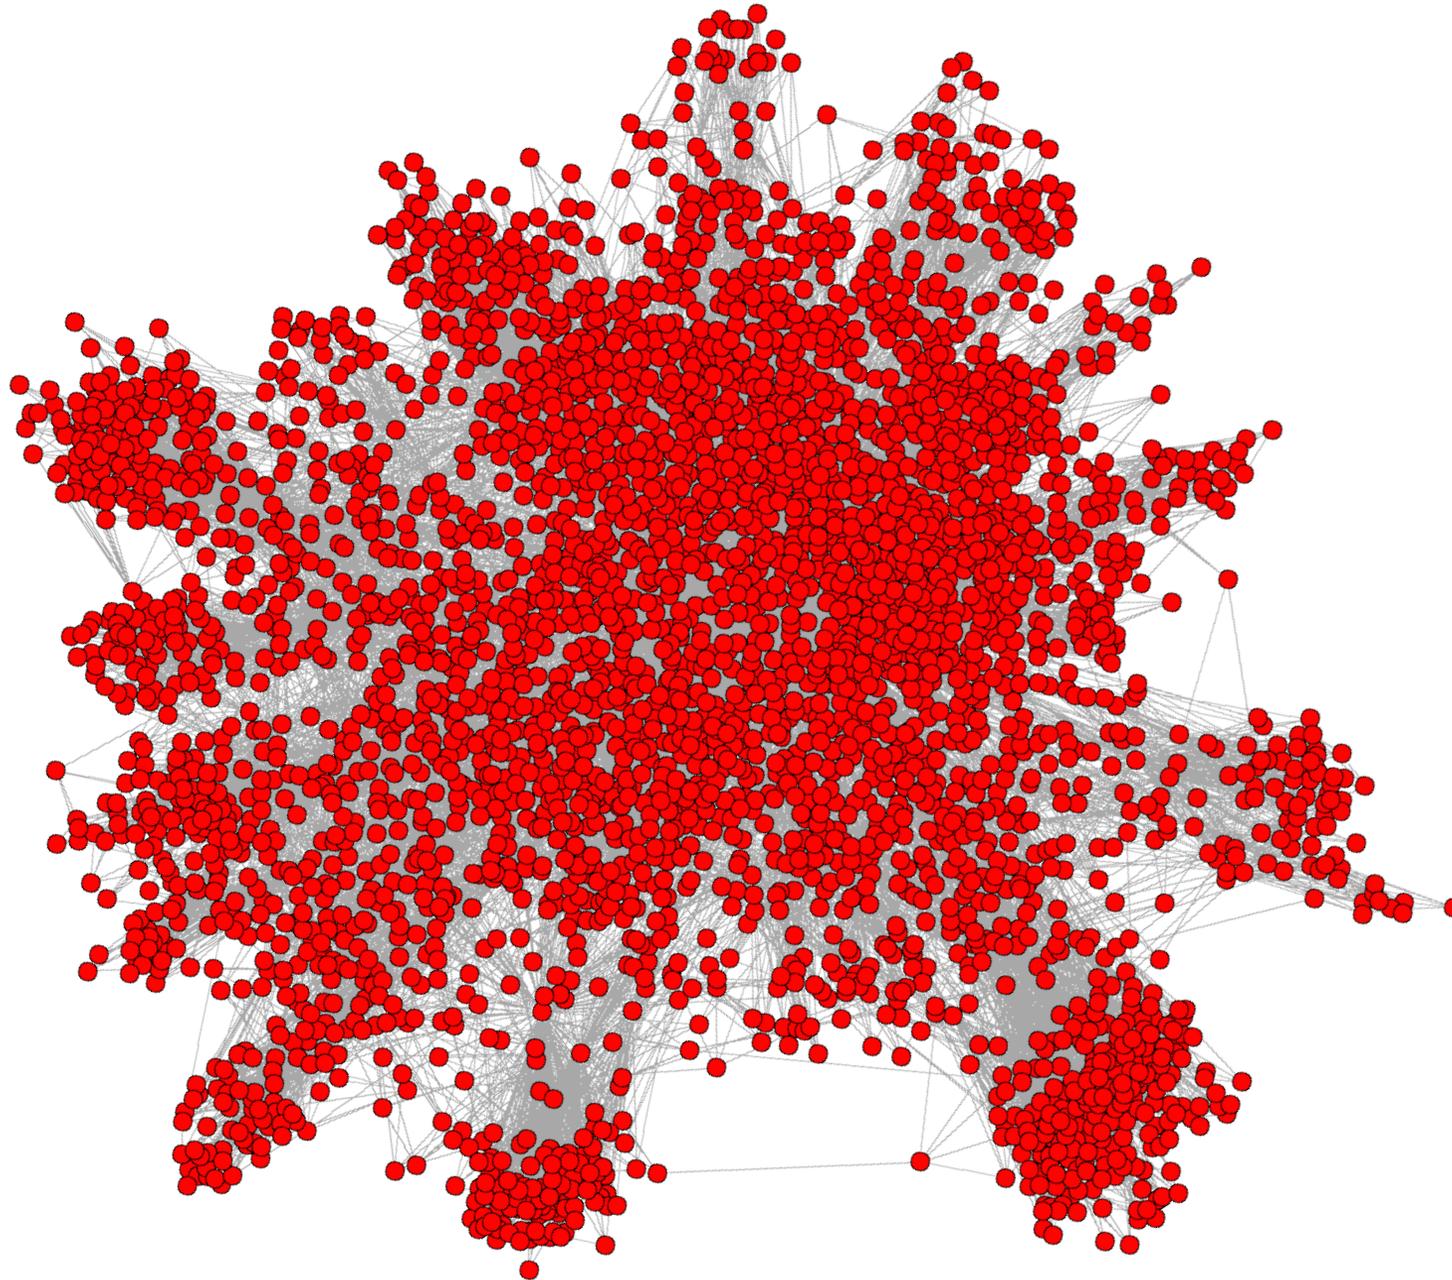
Torsten Hoefler
Department of Computer Science
ETH Zurich
htor@inf.ethz.ch

ABSTRACT

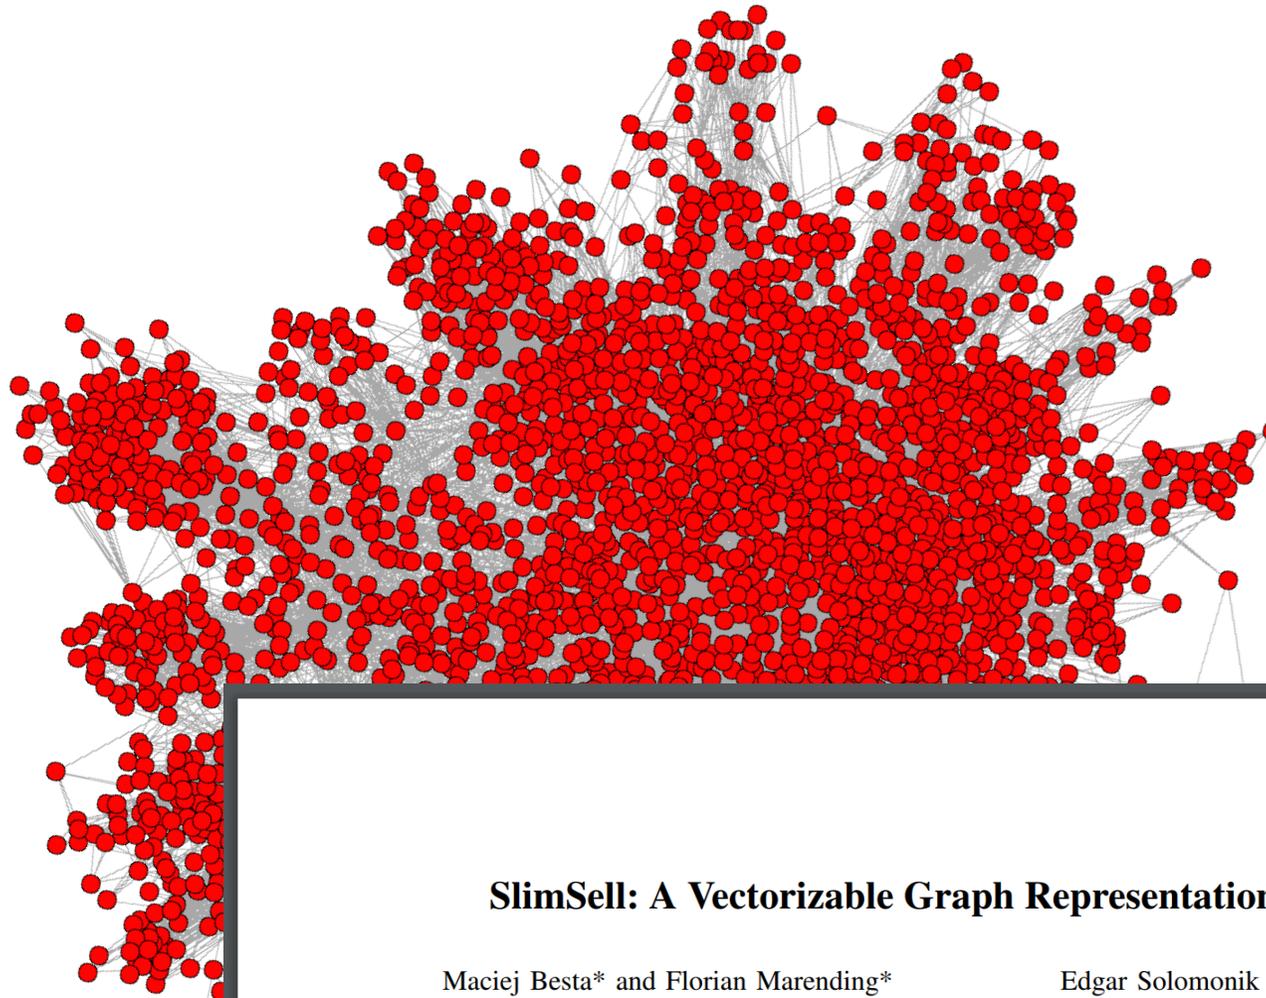
We propose a topology-aware distributed Reader-Writer lock

cesses competing for the same lock. Assume that two of them (A and B) run on one socket and the remaining two (C

Moving on ...



Moving on ...



Irregular

SlimSell: A Vectorizable Graph Representation for Breadth-First Search

Maciej Besta* and Florian Marending*
Department of Computer Science
 ETH Zurich
 {maciej.best@inf, floriama@student}.ethz.ch

Edgar Solomonik
Department of Computer Science
 University of Illinois Urbana-Champaign
 solomon2@illinois.edu

Torsten Hoefler
Department of Computer Science
 ETH Zurich
 htor@inf.ethz.ch

Abstract—Vectorization and GPUs will profoundly change graph processing. Traditional graph algorithms tuned for 32- or 64-bit based memory accesses will be inefficient on architectures with 512-bit wide (or larger) instruction units that are already present in the Intel Knights Landing (KNL) manycore GPU. Anticipating this shift, we propose SlimSell, a

a dense vector (SpMV) or a sparse matrix and a sparse vector (SpMSpV). BFS based on SpMV (BFS-SpMV) uses no explicit locking or atomics and has a succinct description as well as good locality [13]. Yet, it needs more work than traditional BFS and BFS based on SpMSpV [29]

VECTORIZATION

VECTORIZATION

- Deployed in various hardware

VECTORIZATION

- Deployed in various hardware
- Becoming more popular

VECTORIZATION

- Deployed in various hardware
- Becoming more popular



$C = 8$ (SIMD width)

VECTORIZATION

- Deployed in various hardware
- Becoming more popular



AVX

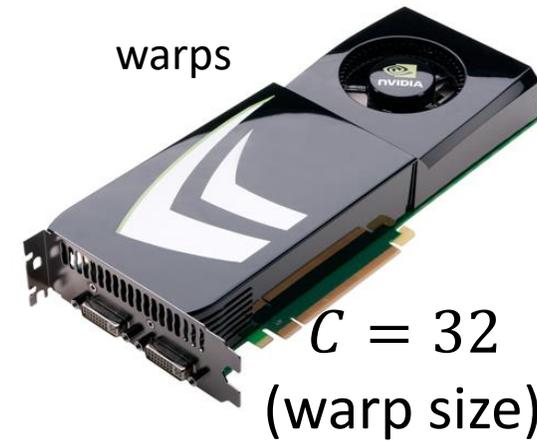
$C = 16$ (SIMD width)



$C = 8$ (SIMD width)

VECTORIZATION

- Deployed in various hardware
- Becoming more popular

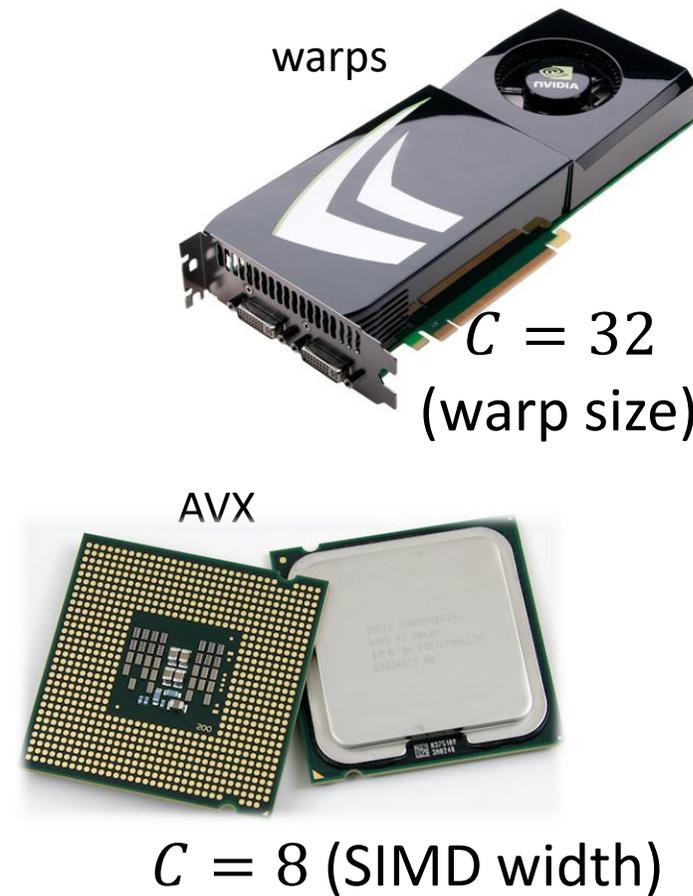


C : „Chunk” size: SIMD width (CPUs, KNLs), warp size (GPUs)

$C = 8$ (SIMD width)

VECTORIZATION

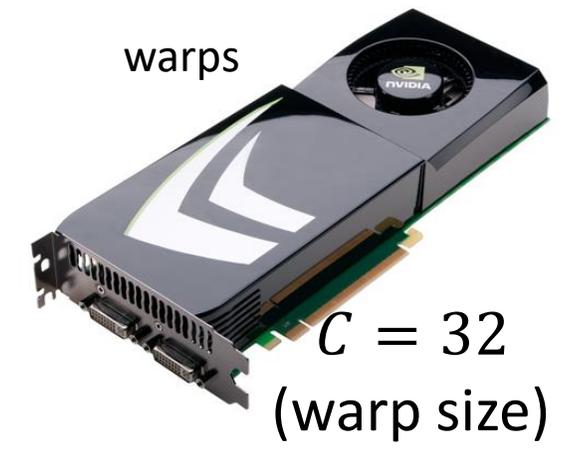
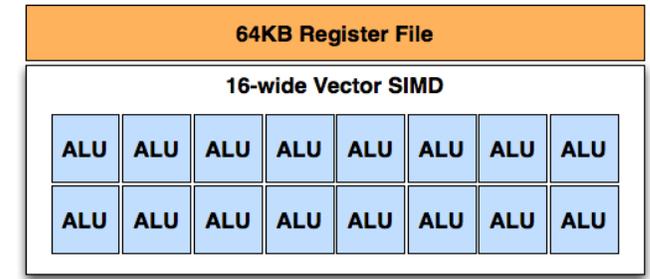
- Deployed in various hardware
- Becoming more popular
- Offers a lot of „regular” compute power



C : „Chunk” size: SIMD width (CPUs, KNLs), warp size (GPUs)

VECTORIZATION

- Deployed in various hardware
- Becoming more popular
- Offers a lot of „regular” compute power

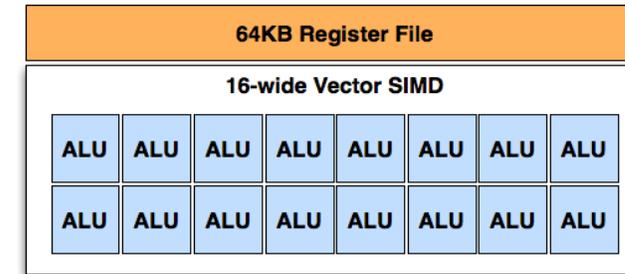


C : „Chunk” size: SIMD width (CPUs, KNLs), warp size (GPUs)

$C = 8$ (SIMD width)

VECTORIZATION

- Deployed in various hardware
- Becoming more popular
- Offers a lot of „regular” compute power

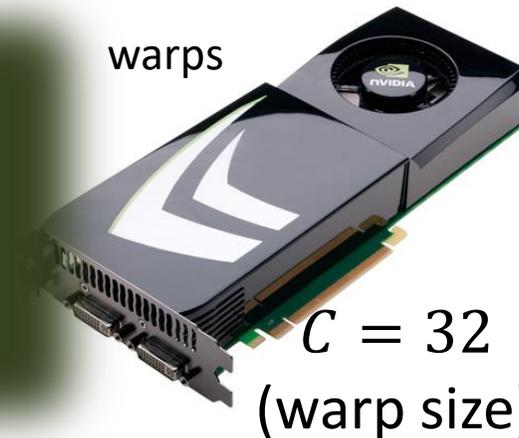


Regular



AVX

$C = 16$ (SIMD width)



warps

$C = 32$
(warp size)



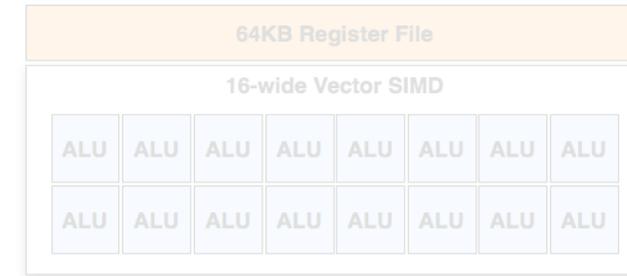
AVX

$C = 8$ (SIMD width)

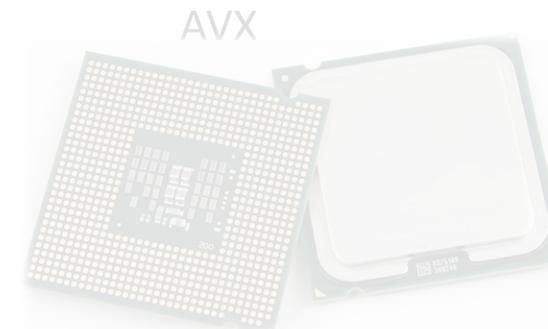
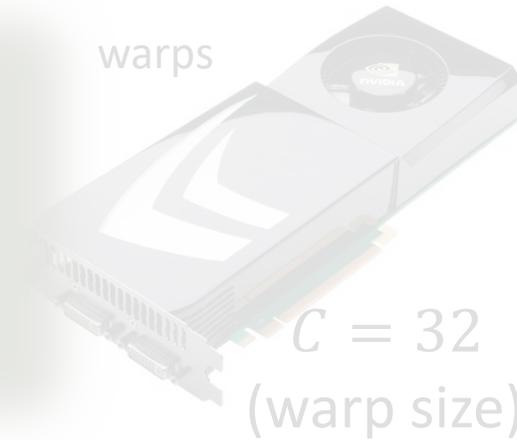
C : „Chunk” size: SIMD width (CPUs, KNLs), warp size (GPUs)

VECTORIZATION

- Deployed in various hardware
- Becoming more popular
- Offers a lot of „regular” compute power



Regular



C : „Chunk” size: SIMD width (CPUs, KNLs), warp size (GPUs)

$C = 8$ (SIMD width)

VECTORIZATION

- Deployed in various hardware
- Becoming more popular
- Offers a lot of „regular” compute power



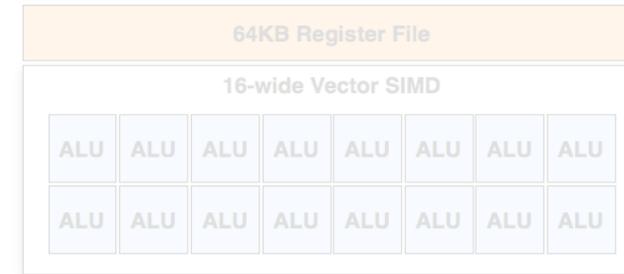
Regu +



AVX

$C = 16$ (SIMD width)

C : „Chunk” size: SIMD width (CPUs, KNLs), warp size (GPUs)



VECTORIZATION

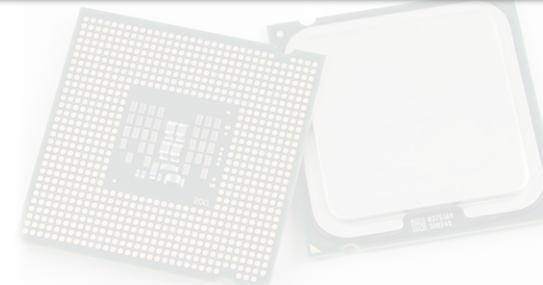
- Deployed in various hardware
- Becoming more popular
- Offers a lot of „regular” compute power

Regular

$C = 16$

$C = 32$

$C = 8$



$C = 8$ (SIMD width)

BREADTH-FIRST SEARCH

TRADITIONAL FORMULATION

BREADTH-FIRST SEARCH

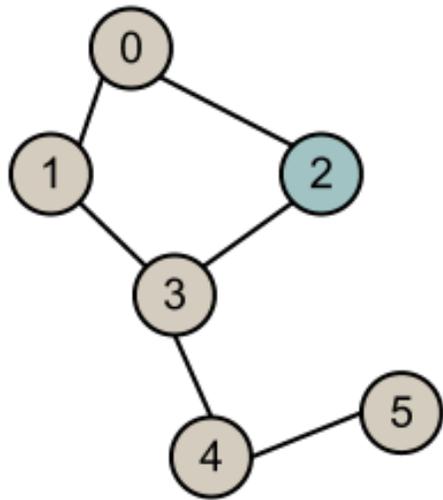
TRADITIONAL FORMULATION

- BFS is based on primitives such as queues

BREADTH-FIRST SEARCH

TRADITIONAL FORMULATION

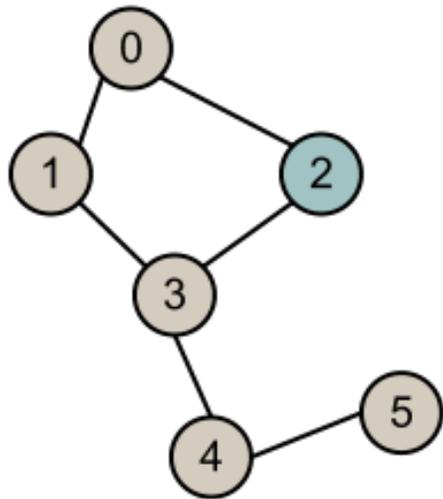
- BFS is based on primitives such as queues



BREADTH-FIRST SEARCH

TRADITIONAL FORMULATION

- BFS is based on primitives such as queues

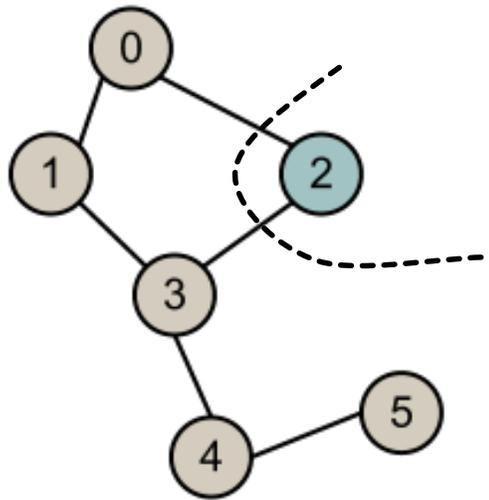


1) $F = \{\}$

BREADTH-FIRST SEARCH

TRADITIONAL FORMULATION

- BFS is based on primitives such as queues

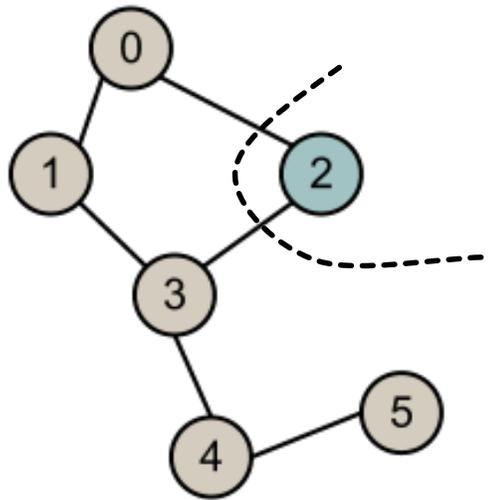


1) $F = \{\}$

BREADTH-FIRST SEARCH

TRADITIONAL FORMULATION

- BFS is based on primitives such as queues

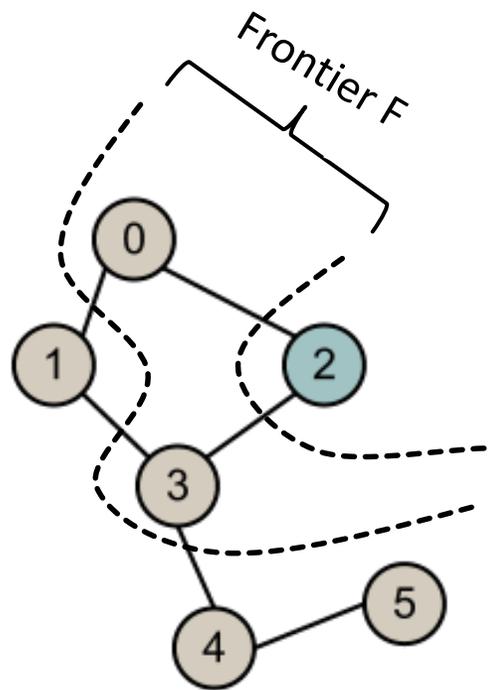


- 1) $F = \{\}$
- 2) $F = \{2\}$

BREADTH-FIRST SEARCH

TRADITIONAL FORMULATION

- BFS is based on primitives such as queues

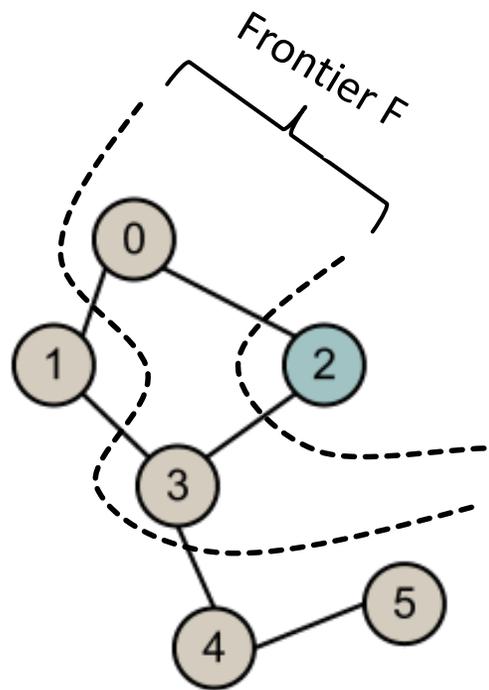


- 1) $F = \{\}$
- 2) $F = \{2\}$

BREADTH-FIRST SEARCH

TRADITIONAL FORMULATION

- BFS is based on primitives such as queues

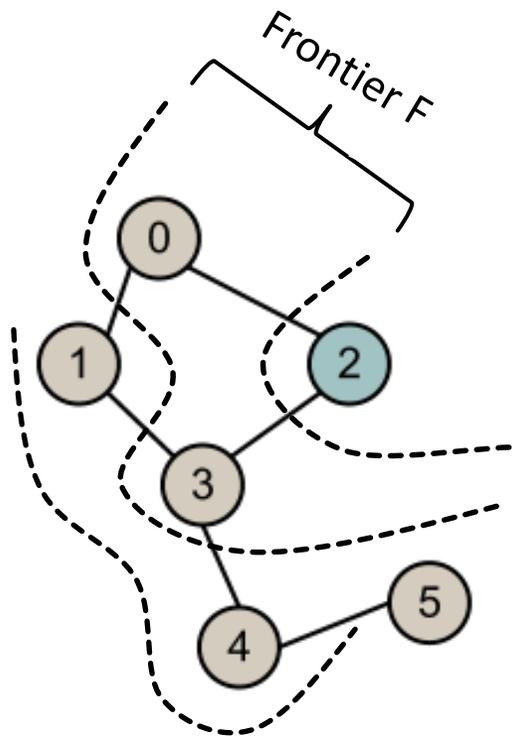


- 1) $F = \{\}$
- 2) $F = \{2\}$
- 3) $F = \{0,3\}$

BREADTH-FIRST SEARCH

TRADITIONAL FORMULATION

- BFS is based on primitives such as queues

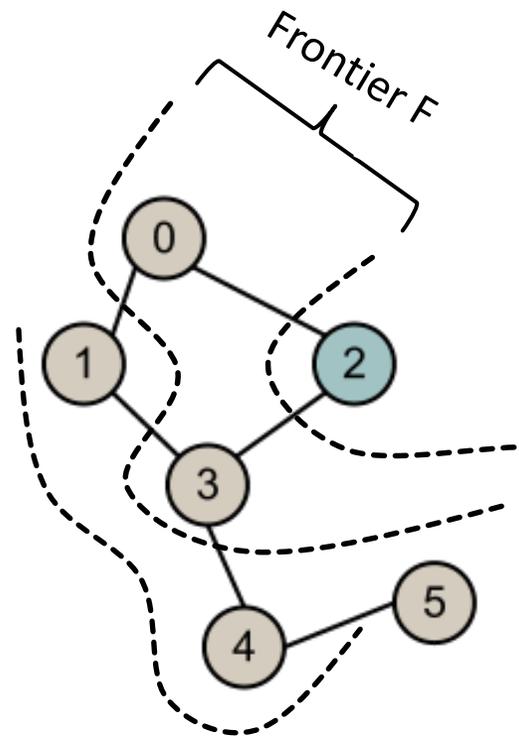


- 1) $F = \{\}$
- 2) $F = \{2\}$
- 3) $F = \{0,3\}$

BREADTH-FIRST SEARCH

TRADITIONAL FORMULATION

- BFS is based on primitives such as queues

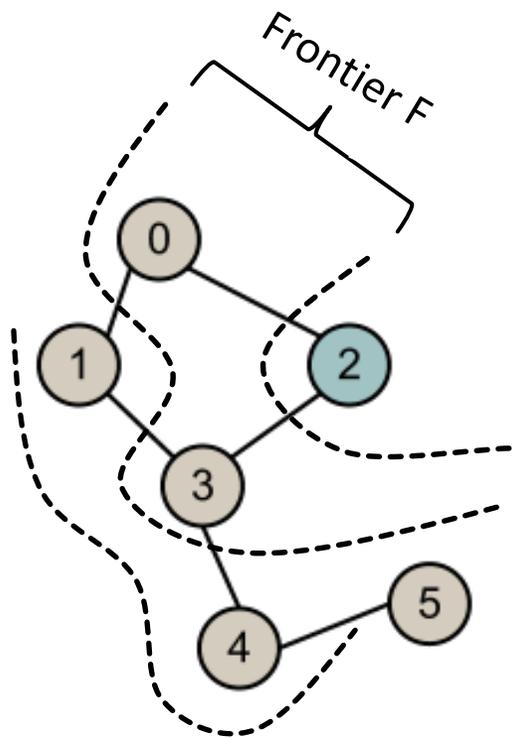


- 1) $F = \{\}$
- 2) $F = \{2\}$
- 3) $F = \{0,3\}$
- 4) $F = \{1,4\}$

BREADTH-FIRST SEARCH

TRADITIONAL FORMULATION

- BFS is based on primitives such as queues



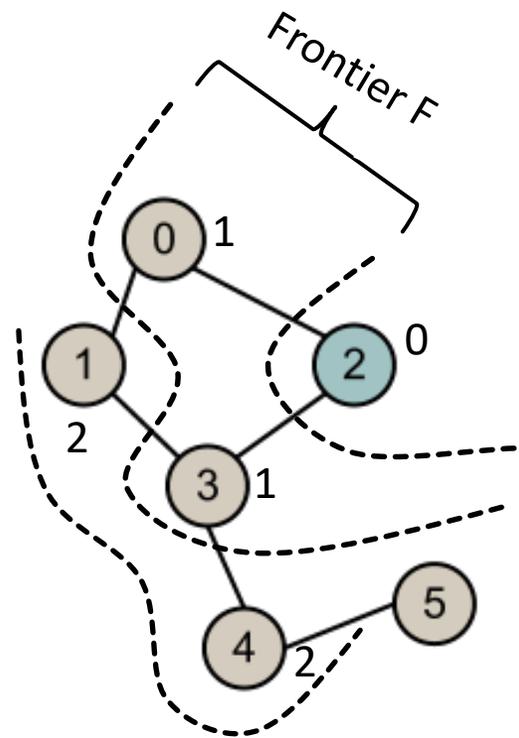
- 1) $F = \{\}$
- 2) $F = \{2\}$
- 3) $F = \{0,3\}$
- 4) $F = \{1,4\}$

! Distances from the root

BREADTH-FIRST SEARCH

TRADITIONAL FORMULATION

- BFS is based on primitives such as queues



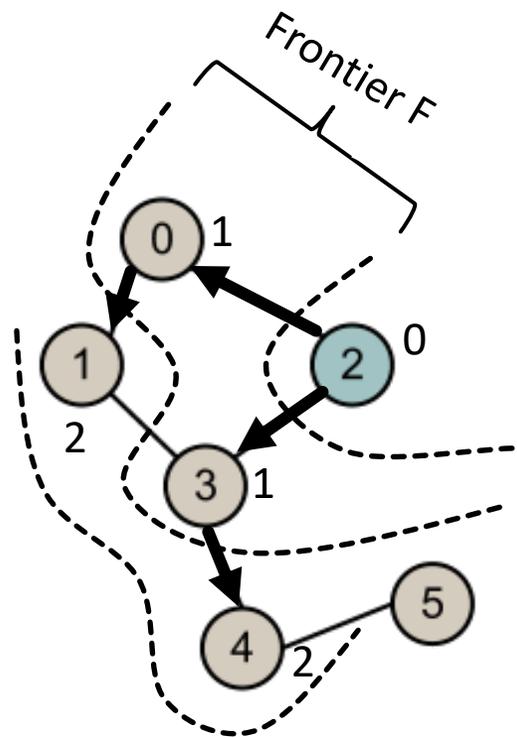
- 1) $F = \{\}$
- 2) $F = \{2\}$
- 3) $F = \{0, 3\}$
- 4) $F = \{1, 4\}$

! Distances from the root

BREADTH-FIRST SEARCH

TRADITIONAL FORMULATION

- BFS is based on primitives such as queues



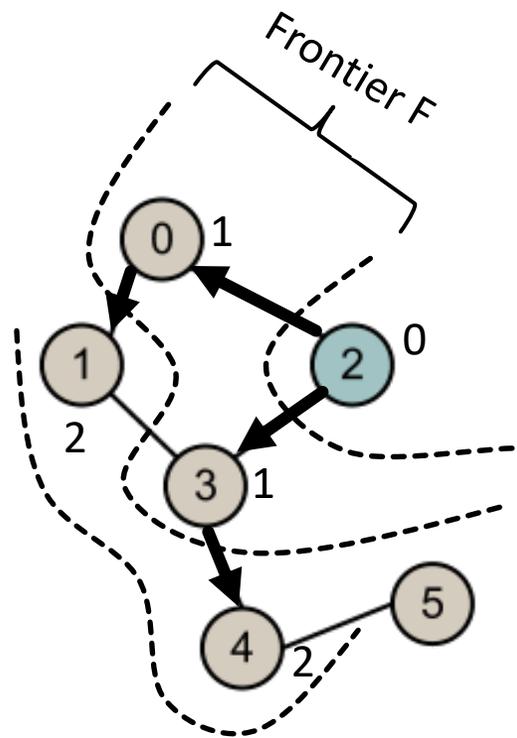
- 1) $F = \{\}$
- 2) $F = \{2\}$
- 3) $F = \{0, 3\}$
- 4) $F = \{1, 4\}$

! Distances from the root

BREADTH-FIRST SEARCH

TRADITIONAL FORMULATION

- BFS is based on primitives such as queues



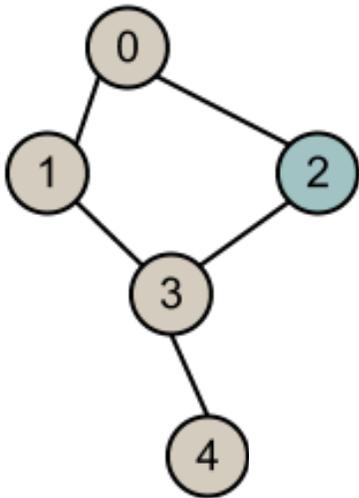
- 1) $F = \{0\}$
- 2) $F = \{2\}$
- 3) $F = \{0, 3\}$
- 4) $F = \{1, 4\}$

! Distances from the root

! Parents (predecessors) in the traversal tree

BREADTH-FIRST SEARCH

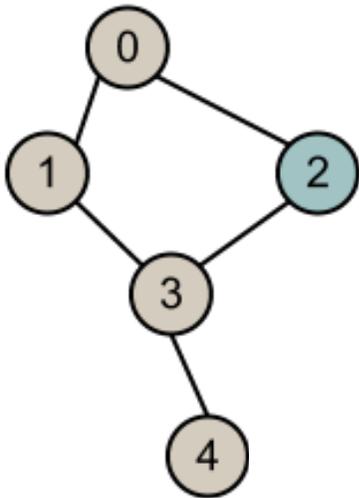
ALGEBRAIC FORMULATION



BREADTH-FIRST SEARCH

ALGEBRAIC FORMULATION

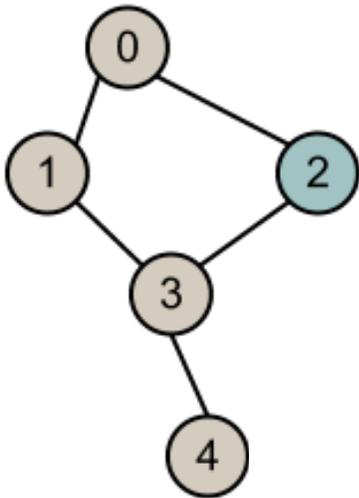
- BFS is a series of matrix-vector products



BREADTH-FIRST SEARCH

ALGEBRAIC FORMULATION

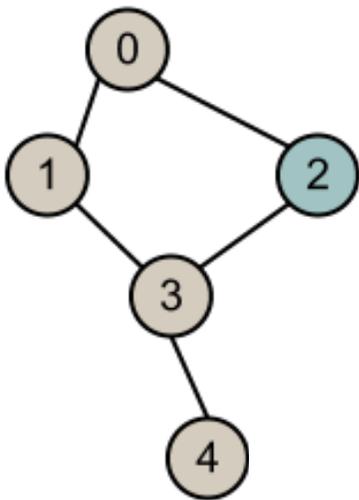
- BFS is a series of matrix-vector products
- Graph is modeled by an adjacency matrix



BREADTH-FIRST SEARCH

ALGEBRAIC FORMULATION

- BFS is a series of matrix-vector products
- Graph is modeled by an adjacency matrix



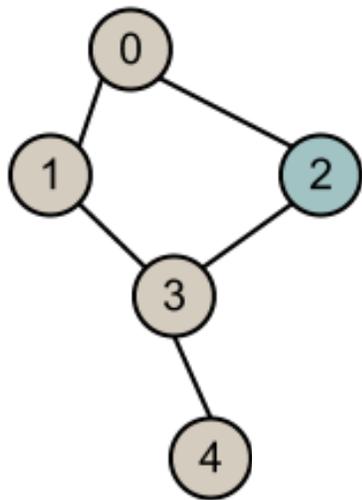
Adjacency Matrix:

$$A = \begin{bmatrix} 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$

BREADTH-FIRST SEARCH

ALGEBRAIC FORMULATION

- BFS is a series of matrix-vector products
- Graph is modeled by an adjacency matrix
- Multiplication is done over a semiring



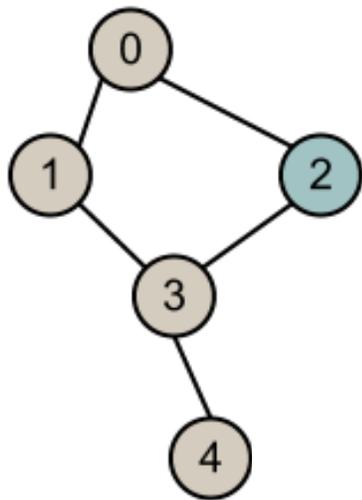
Adjacency Matrix:

$$A = \begin{bmatrix} 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$

BREADTH-FIRST SEARCH

ALGEBRAIC FORMULATION

- BFS is a series of matrix-vector products
- Graph is modeled by an adjacency matrix
- Multiplication is done over a semiring



Adjacency Matrix:

$$A = \begin{bmatrix} 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$

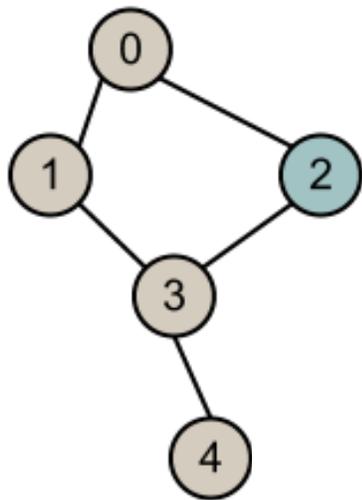
Semiring:

$$(\mathbb{R}, op_1, op_2, el_1, el_2)$$

BREADTH-FIRST SEARCH

ALGEBRAIC FORMULATION

- BFS is a series of matrix-vector products
- Graph is modeled by an adjacency matrix
- Multiplication is done over a semiring



Adjacency Matrix:

$$A = \begin{bmatrix} 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$

Semiring:

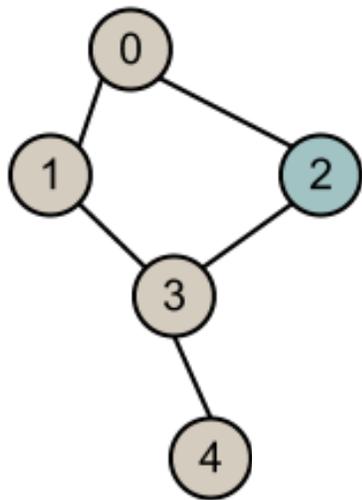
$$(\mathbb{R}, op_1, op_2, el_1, el_2)$$

$$\begin{bmatrix} 1 & 2 \\ 3 & 1 \end{bmatrix} \begin{bmatrix} 0 \\ 2 \end{bmatrix} = \begin{bmatrix} 4 \\ 2 \end{bmatrix}$$

BREADTH-FIRST SEARCH

ALGEBRAIC FORMULATION

- BFS is a series of matrix-vector products
- Graph is modeled by an adjacency matrix
- Multiplication is done over a semiring



Adjacency Matrix:

$$A = \begin{bmatrix} 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$

Semiring:

$$(\mathbb{R}, op_1, op_2, el_1, el_2)$$

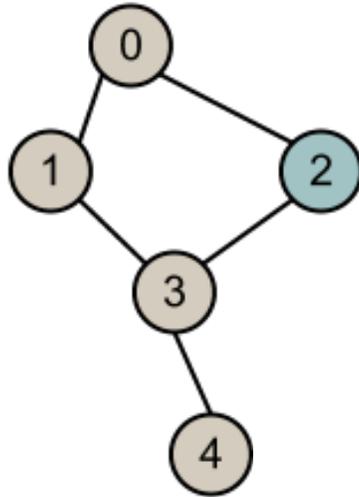
$$(\mathbb{R}, +, :, 0, 1)$$

$$\begin{bmatrix} 1 & 2 \\ 3 & 1 \end{bmatrix} \begin{bmatrix} 0 \\ 2 \end{bmatrix} = \begin{bmatrix} 4 \\ 2 \end{bmatrix}$$

BREADTH-FIRST SEARCH

ALGEBRAIC FORMULATION

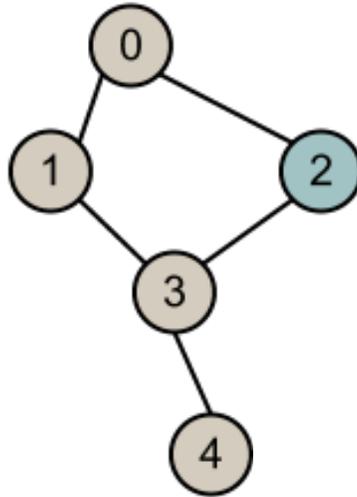
Tropical Semiring
 $(\mathbb{R} \cup \{\infty\}, \min, +, \infty, 0)$



BREADTH-FIRST SEARCH

ALGEBRAIC FORMULATION

Tropical Semiring
 $(\mathbb{R} \cup \{\infty\}, \min, +, \infty, 0)$



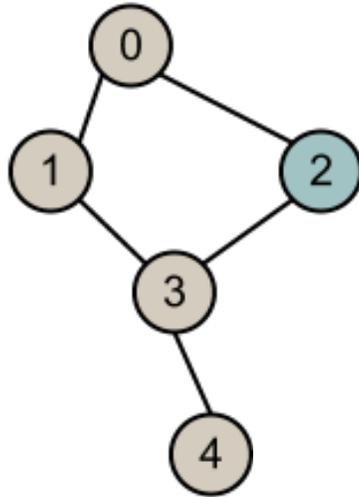
$$A' = \begin{bmatrix} 0 & 1 & 1 & \infty & \infty \\ 1 & 0 & \infty & 1 & \infty \\ 1 & \infty & 0 & 1 & \infty \\ \infty & 1 & 1 & 0 & 1 \\ \infty & \infty & \infty & 1 & 0 \end{bmatrix}$$

BREADTH-FIRST SEARCH

ALGEBRAIC FORMULATION

Tropical Semiring
 $(\mathbb{R} \cup \{\infty\}, \min, +, \infty, 0)$

$$A = \begin{bmatrix} 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$



$$A' = \begin{bmatrix} 0 & 1 & 1 & \infty & \infty \\ 1 & 0 & \infty & 1 & \infty \\ 1 & \infty & 0 & 1 & \infty \\ \infty & 1 & 1 & 0 & 1 \\ \infty & \infty & \infty & 1 & 0 \end{bmatrix}$$

BREADTH-FIRST SEARCH

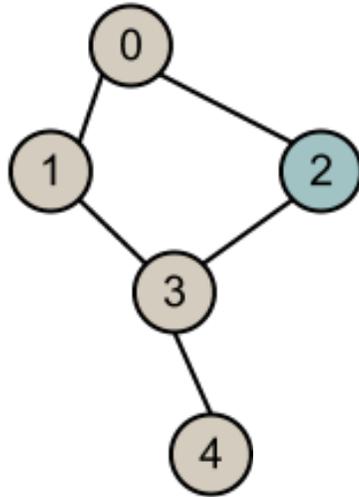
ALGEBRAIC FORMULATION

Tropical Semiring
 $(\mathbb{R} \cup \{\infty\}, \min, +, \infty, 0)$

$$A = \begin{bmatrix} 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$

Usually stored using a
sparse format

$$A' = \begin{bmatrix} 0 & 1 & 1 & \infty & \infty \\ 1 & 0 & \infty & 1 & \infty \\ 1 & \infty & 0 & 1 & \infty \\ \infty & 1 & 1 & 0 & 1 \\ \infty & \infty & \infty & 1 & 0 \end{bmatrix}$$

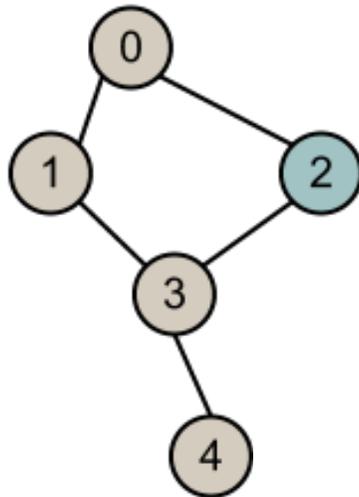


BREADTH-FIRST SEARCH

ALGEBRAIC FORMULATION

Tropical Semiring
 $(\mathbb{R} \cup \{\infty\}, \min, +, \infty, 0)$

$$A = \begin{bmatrix} 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$



Usually stored using a
sparse format

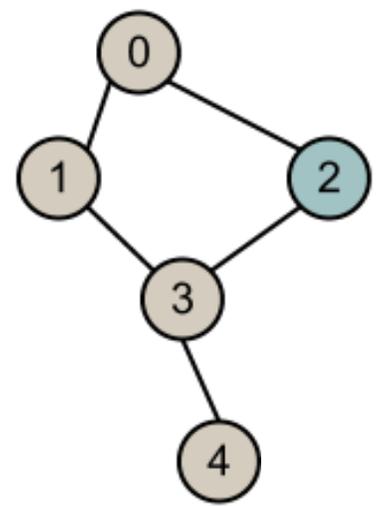
$$A' = \begin{bmatrix} 0 & 1 & 1 & \infty & \infty \\ 1 & 0 & \infty & 1 & \infty \\ 1 & \infty & 0 & 1 & \infty \\ \infty & 1 & 1 & 0 & 1 \\ \infty & \infty & \infty & 1 & 0 \end{bmatrix}$$

$$f_0 = \begin{pmatrix} \infty \\ \infty \\ 0 \\ \infty \\ \infty \end{pmatrix}$$

BREADTH-FIRST SEARCH

ALGEBRAIC FORMULATION

$$A = \begin{bmatrix} 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$



Tropical Semiring
 $(\mathbb{R} \cup \{\infty\}, \min, +, \infty, 0)$

Usually stored using a
sparse format

$$A' = \begin{bmatrix} 0 & 1 & 1 & \infty & \infty \\ 1 & 0 & \infty & 1 & \infty \\ 1 & \infty & 0 & 1 & \infty \\ \infty & 1 & 1 & 0 & 1 \\ \infty & \infty & \infty & 1 & 0 \end{bmatrix}$$

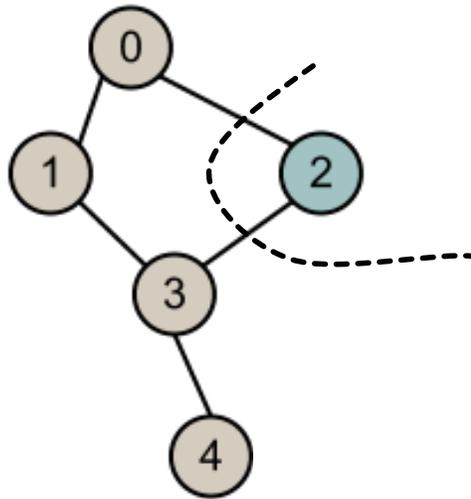
Stored with a dense or a
sparse format

$$f_0 = \begin{pmatrix} \infty \\ \infty \\ 0 \\ \infty \\ \infty \end{pmatrix}$$

BREADTH-FIRST SEARCH

ALGEBRAIC FORMULATION

$$A = \begin{bmatrix} 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$



Usually stored using a
sparse format

$$A' = \begin{bmatrix} 0 & 1 & 1 & \infty & \infty \\ 1 & 0 & \infty & 1 & \infty \\ 1 & \infty & 0 & 1 & \infty \\ \infty & 1 & 1 & 0 & 1 \\ \infty & \infty & \infty & 1 & 0 \end{bmatrix}$$

Tropical Semiring
($\mathbb{R} \cup \{\infty\}, \min, +, \infty, 0$)

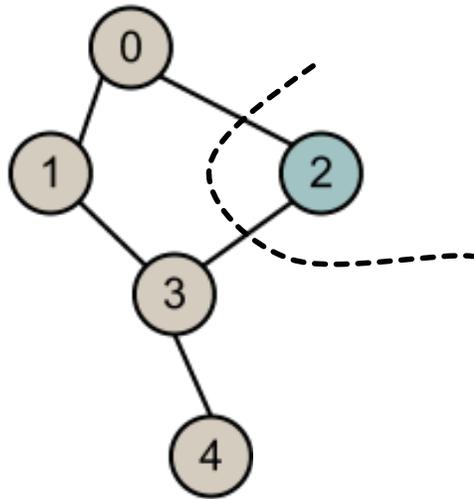
Stored with a dense or a
sparse format

$$f_0 = \begin{pmatrix} \infty \\ \infty \\ 0 \\ \infty \\ \infty \end{pmatrix}$$

BREADTH-FIRST SEARCH

ALGEBRAIC FORMULATION

$$A = \begin{bmatrix} 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$



Tropical Semiring
 $(\mathbb{R} \cup \{\infty\}, \min, +, \infty, 0)$

Usually stored using a
 sparse format

$$A' = \begin{bmatrix} 0 & 1 & 1 & \infty & \infty \\ 1 & 0 & \infty & 1 & \infty \\ 1 & \infty & 0 & 1 & \infty \\ \infty & 1 & 1 & 0 & 1 \\ \infty & \infty & \infty & 1 & 0 \end{bmatrix}$$

Stored with a dense or a
 sparse format

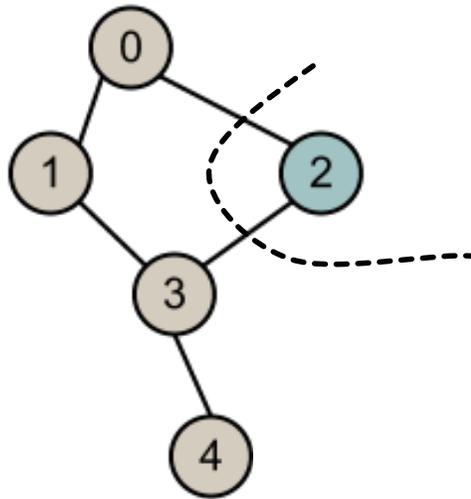
$$f_0 = \begin{pmatrix} \infty \\ \infty \\ 0 \\ \infty \\ \infty \end{pmatrix}$$

$$f_1 = A'^T \otimes_T f_0 = \begin{pmatrix} \\ \\ \\ \\ \end{pmatrix}$$

BREADTH-FIRST SEARCH

ALGEBRAIC FORMULATION

$$A = \begin{bmatrix} 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$



Usually stored using a sparse format

$$A' = \begin{bmatrix} 0 & 1 & 1 & \infty & \infty \\ 1 & 0 & \infty & 1 & \infty \\ 1 & \infty & 0 & 1 & \infty \\ \infty & 1 & 1 & 0 & 1 \\ \infty & \infty & \infty & 1 & 0 \end{bmatrix}$$

Stored with a dense or a sparse format

$$f_0 = \begin{pmatrix} \infty \\ \infty \\ 0 \\ \infty \\ \infty \end{pmatrix}$$

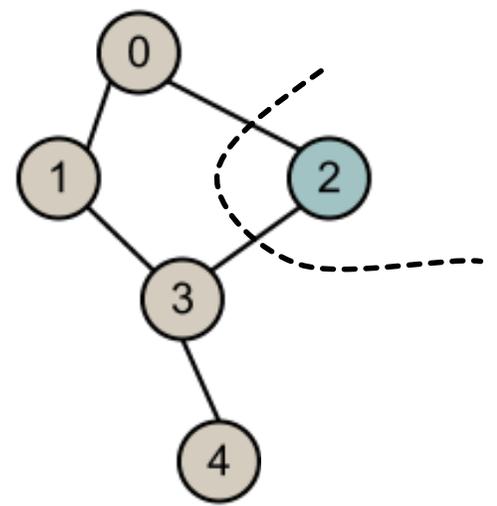
$$f_1 = A'^T \otimes_T f_0 = \begin{pmatrix} \\ \\ \\ \\ \end{pmatrix}$$

Tropical Semiring
 $(\mathbb{R} \cup \{\infty\}, \min, +, \infty, 0)$

BREADTH-FIRST SEARCH

ALGEBRAIC FORMULATION

$$A = \begin{bmatrix} 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$



Tropical Semiring
 $(\mathbb{R} \cup \{\infty\}, \min, +, \infty, 0)$

Usually stored using a sparse format

$$A' = \begin{bmatrix} 0 & 1 & 1 & \infty & \infty \\ 1 & 0 & \infty & 1 & \infty \\ 1 & \infty & 0 & 1 & \infty \\ \infty & 1 & 1 & 0 & 1 \\ \infty & \infty & \infty & 1 & 0 \end{bmatrix}$$

Stored with a dense or a sparse format

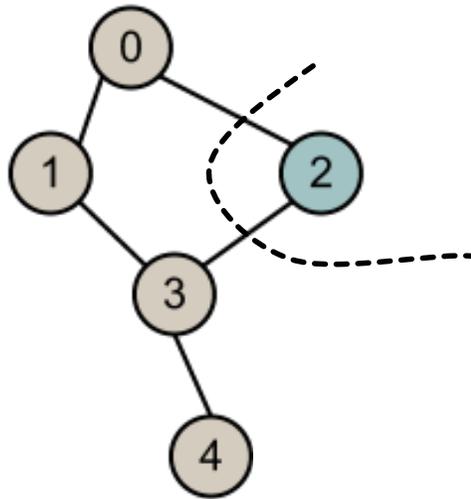
$$f_0 = \begin{pmatrix} \infty \\ \infty \\ 0 \\ \infty \\ \infty \end{pmatrix}$$

$$f_1 = A'^T \otimes_T f_0 = \begin{pmatrix} 1 \end{pmatrix}$$

BREADTH-FIRST SEARCH

ALGEBRAIC FORMULATION

$$A = \begin{bmatrix} 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$



Usually stored using a sparse format

$$A' = \begin{bmatrix} 0 & 1 & 1 & \infty & \infty \\ 1 & 0 & \infty & 1 & \infty \\ 1 & \infty & 0 & 1 & \infty \\ \infty & 1 & 1 & 0 & 1 \\ \infty & \infty & \infty & 1 & 0 \end{bmatrix}$$

Stored with a dense or a sparse format

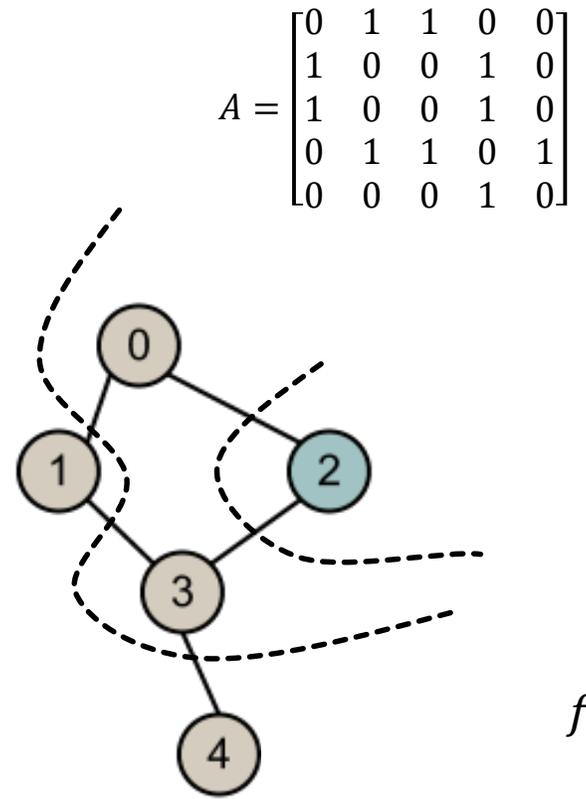
$$f_0 = \begin{pmatrix} \infty \\ \infty \\ 0 \\ \infty \\ \infty \end{pmatrix}$$

$$f_1 = A'^T \otimes_T f_0 = \begin{pmatrix} 1 \\ \infty \\ 0 \\ 1 \\ \infty \end{pmatrix}$$

Tropical Semiring
 $(\mathbb{R} \cup \{\infty\}, \min, +, \infty, 0)$

BREADTH-FIRST SEARCH

ALGEBRAIC FORMULATION



$$A = \begin{bmatrix} 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$

Tropical Semiring
 $(\mathbb{R} \cup \{\infty\}, \min, +, \infty, 0)$

Usually stored using a sparse format

$$A' = \begin{bmatrix} 0 & 1 & 1 & \infty & \infty \\ 1 & 0 & \infty & 1 & \infty \\ 1 & \infty & 0 & 1 & \infty \\ \infty & 1 & 1 & 0 & 1 \\ \infty & \infty & \infty & 1 & 0 \end{bmatrix}$$

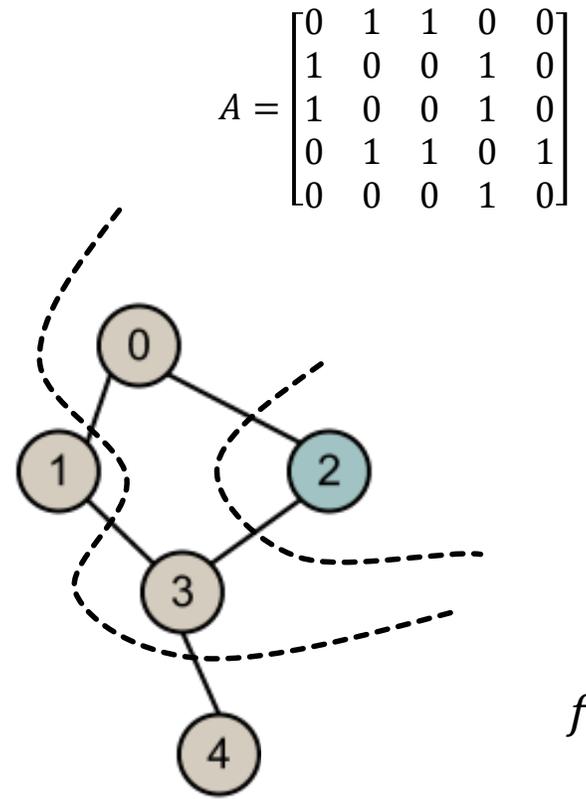
Stored with a dense or a sparse format

$$f_0 = \begin{pmatrix} \infty \\ \infty \\ 0 \\ \infty \\ \infty \end{pmatrix}$$

$$f_1 = A'^T \otimes_T f_0 = \begin{pmatrix} 1 \\ \infty \\ 0 \\ 1 \\ \infty \end{pmatrix}$$

BREADTH-FIRST SEARCH

ALGEBRAIC FORMULATION



$$A = \begin{bmatrix} 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$

Tropical Semiring
 $(\mathbb{R} \cup \{\infty\}, \min, +, \infty, 0)$

Usually stored using a sparse format

$$A' = \begin{bmatrix} 0 & 1 & 1 & \infty & \infty \\ 1 & 0 & \infty & 1 & \infty \\ 1 & \infty & 0 & 1 & \infty \\ \infty & 1 & 1 & 0 & 1 \\ \infty & \infty & \infty & 1 & 0 \end{bmatrix}$$

Stored with a dense or a sparse format

$$f_0 = \begin{pmatrix} \infty \\ \infty \\ 0 \\ \infty \\ \infty \end{pmatrix}$$

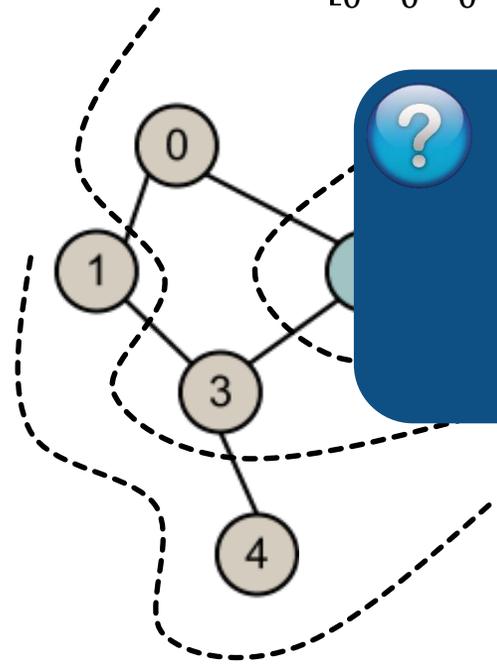
$$f_1 = A'^T \otimes_T f_0 = \begin{pmatrix} 1 \\ \infty \\ 0 \\ 1 \\ \infty \end{pmatrix}$$

$$f_2 = \begin{pmatrix} 1 \\ 2 \\ 0 \\ 1 \\ 2 \end{pmatrix}$$

BREADTH-FIRST SEARCH

ALGEBRAIC FORMULATION

$$A = \begin{bmatrix} 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$



How to do this in practice?

Usually stored using a sparse format

$$\begin{bmatrix} 0 & 1 & 1 & \infty & \infty \end{bmatrix}$$

$$f_1 = A'^T \otimes_T f_0 = \begin{pmatrix} \infty \\ 0 \\ 1 \\ \infty \end{pmatrix}$$

Tropical Semiring
 $(\mathbb{R} \cup \{\infty\}, \min, +, \infty, 0)$

Stored with a dense or a sparse format

$$f_0 = \begin{pmatrix} \infty \\ \infty \\ 0 \\ \infty \\ \infty \end{pmatrix}$$

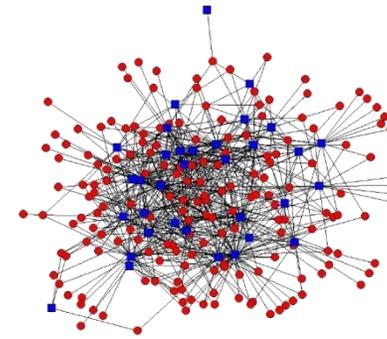
$$f_2 = \begin{pmatrix} 1 \\ 2 \\ 0 \\ 1 \\ 2 \end{pmatrix}$$

GRAPH REPRESENTATIONS

COMPRESSED SPARSE ROW (CSR)

GRAPH REPRESENTATIONS

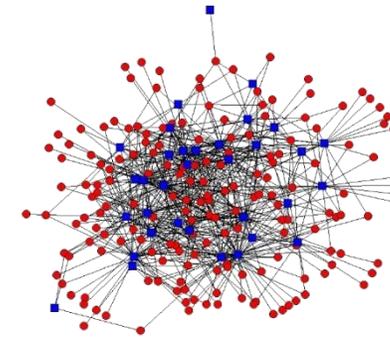
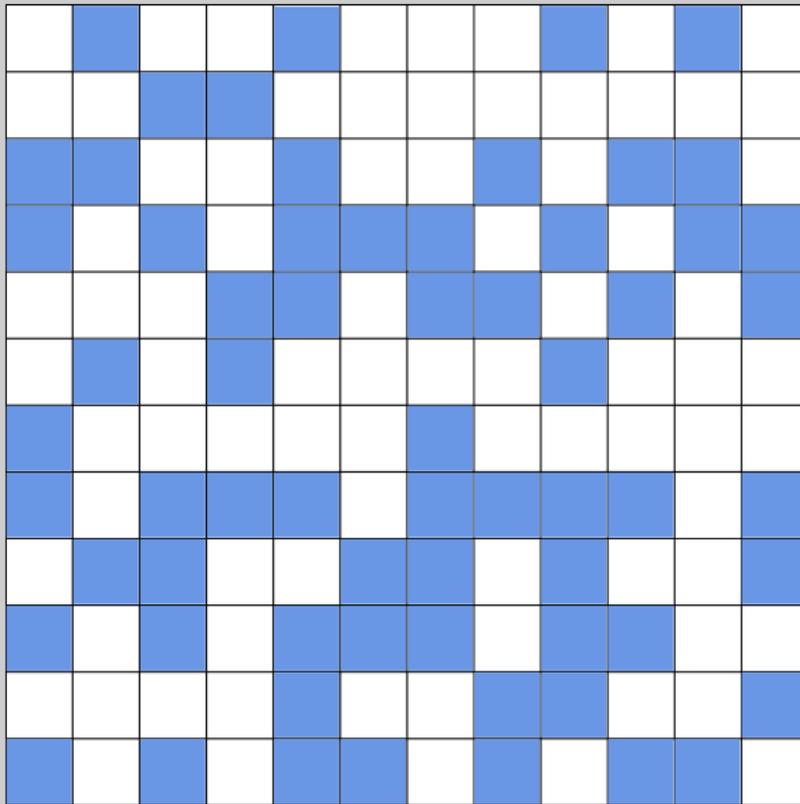
COMPRESSED SPARSE ROW (CSR)



GRAPH REPRESENTATIONS

COMPRESSED SPARSE ROW (CSR)

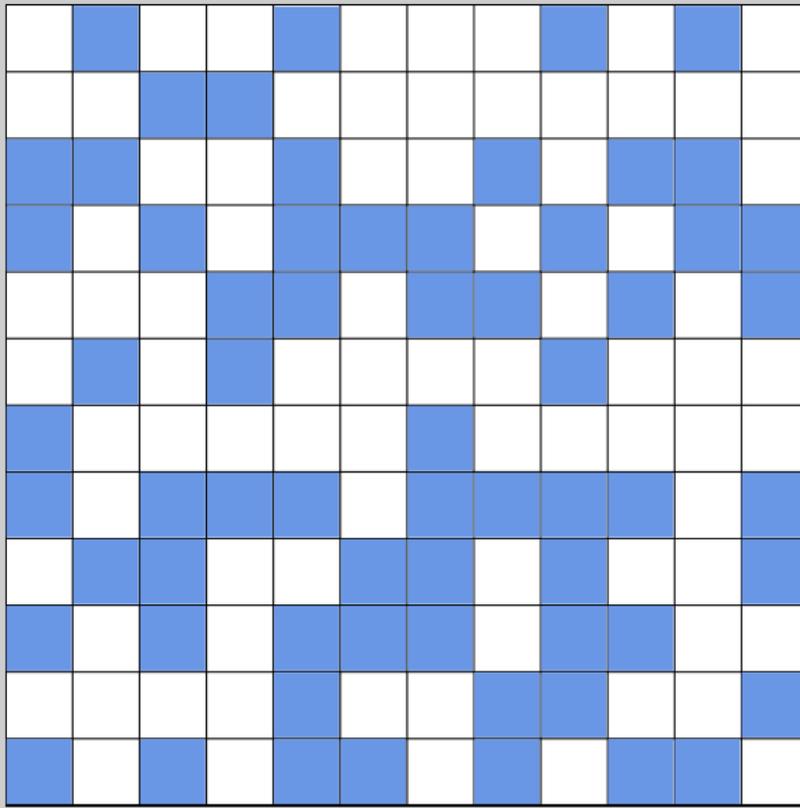
Adjacency matrix



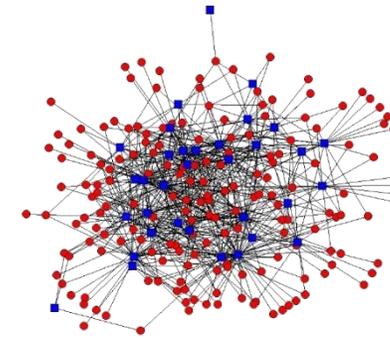
GRAPH REPRESENTATIONS

COMPRESSED SPARSE ROW (CSR)

Adjacency matrix



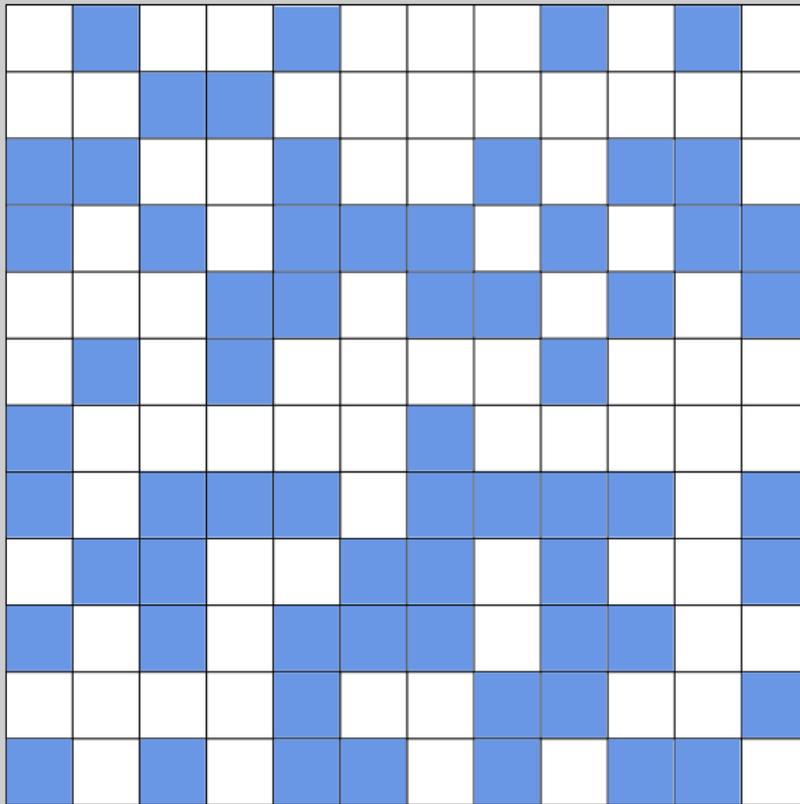
Non-zeros



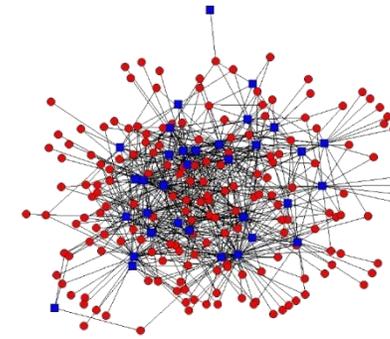
GRAPH REPRESENTATIONS

COMPRESSED SPARSE ROW (CSR)

Adjacency matrix



Non-zeros



Non-zeros are stored in
the *val* array

size: $2m$ cells

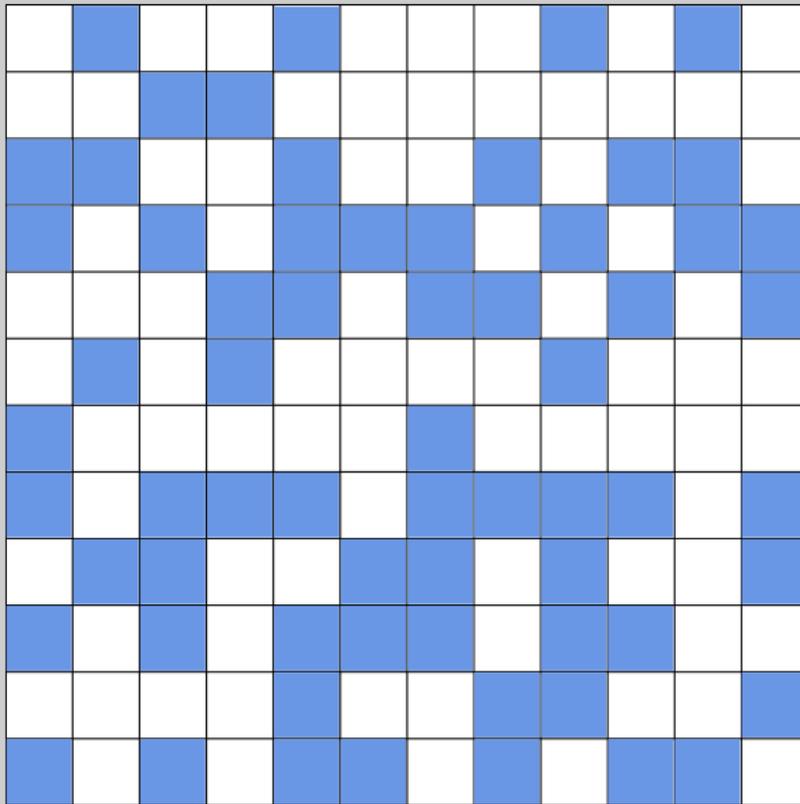


n : number of vertices
 m : number of edges

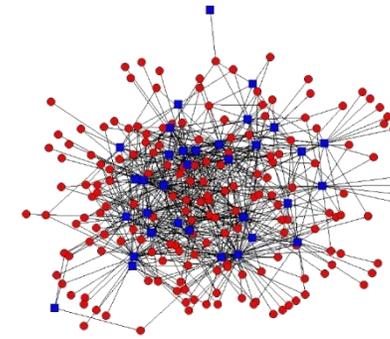
GRAPH REPRESENTATIONS

COMPRESSED SPARSE ROW (CSR)

Adjacency matrix



Non-zeros



Non-zeros are stored in
the *val* array

size: $2m$ cells



Column indices stored
in the *col* array

size: $2m$ cells

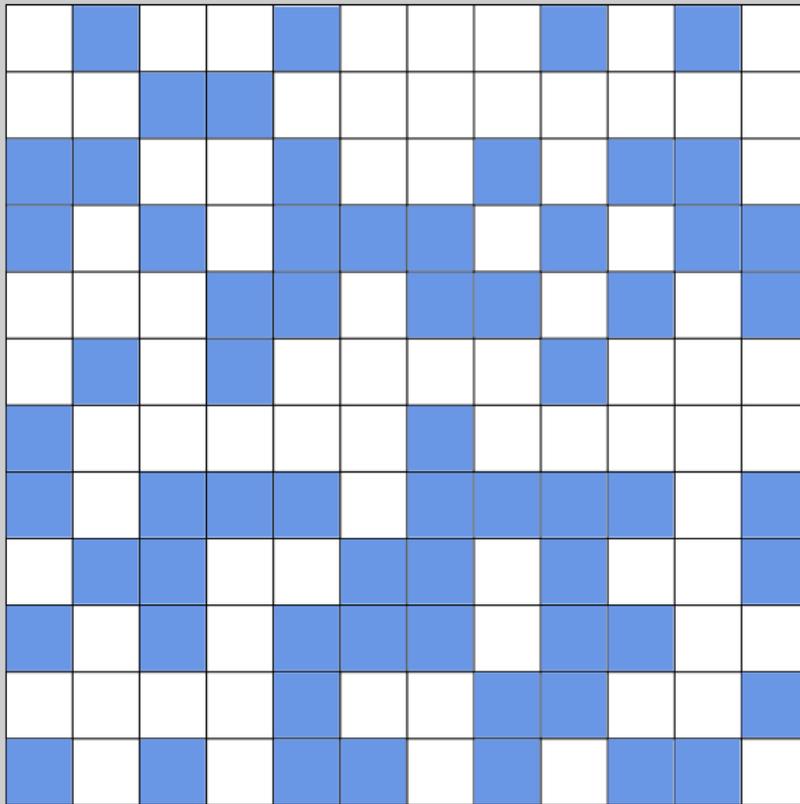


n : number of vertices
 m : number of edges

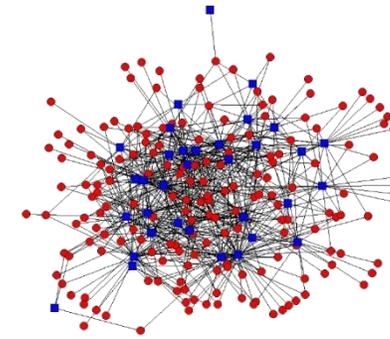
GRAPH REPRESENTATIONS

COMPRESSED SPARSE ROW (CSR)

Adjacency matrix



Non-zeros



Non-zeros are stored in
the *val* array

size: $2m$ cells



Column indices stored
in the *col* array

size: $2m$ cells



Row indices are stored
in the *row* array

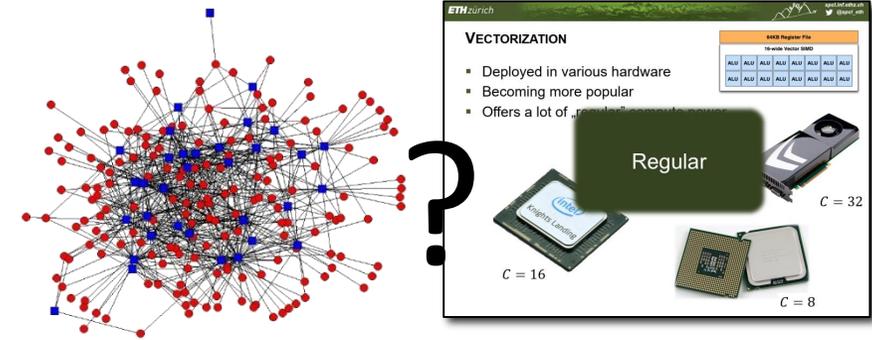
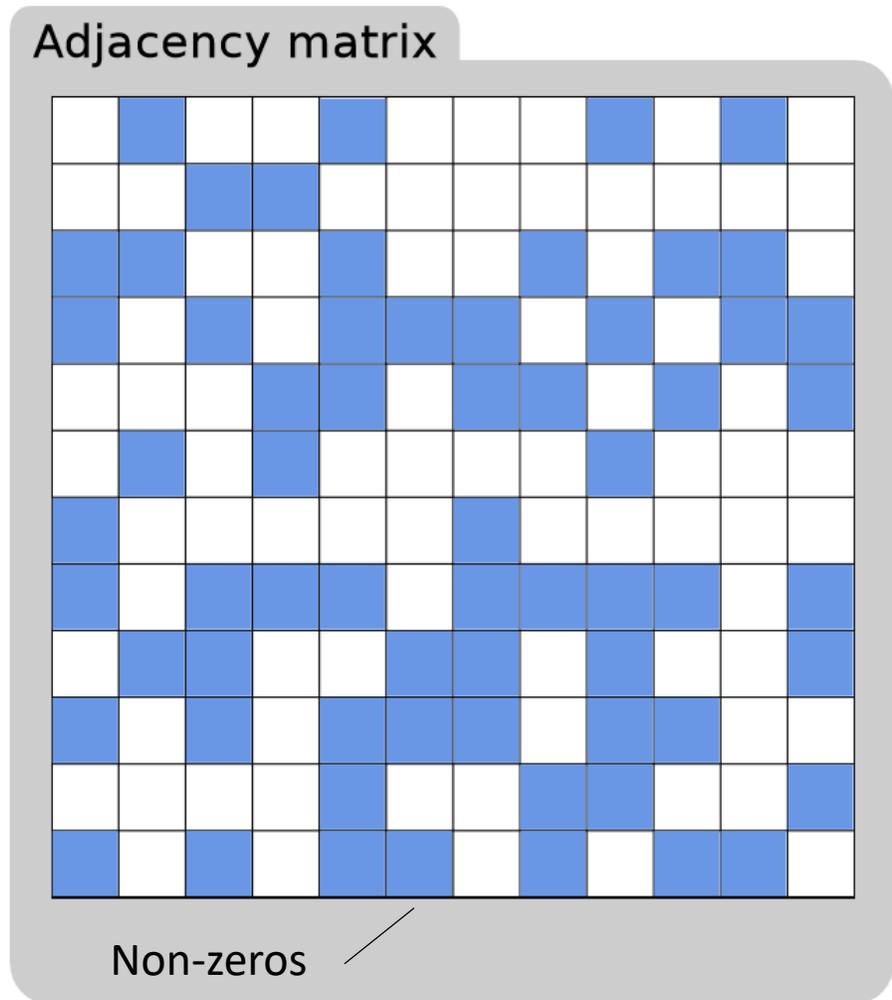
size: n cells



n : number of vertices
 m : number of edges

GRAPH REPRESENTATIONS

COMPRESSED SPARSE ROW (CSR)



Non-zeros are stored in
the *val* array size: $2m$ cells



Column indices stored
in the *col* array size: $2m$ cells



Row indices are stored
in the *row* array size: n cells



n : number of vertices
 m : number of edges

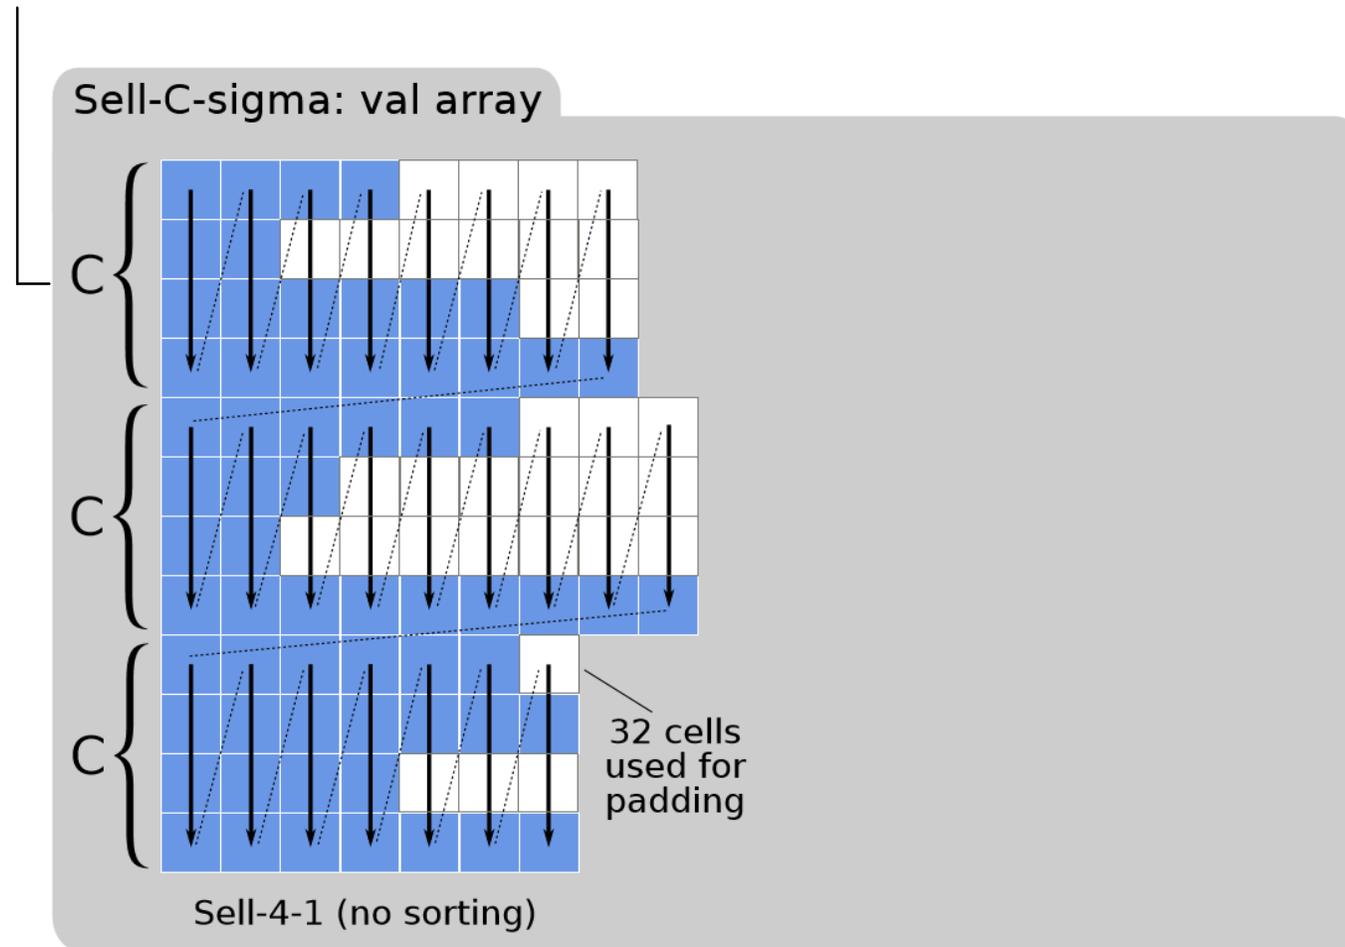
GRAPH REPRESENTATIONS

SELL-C-SIGMA

GRAPH REPRESENTATIONS

SELL-C-SIGMA

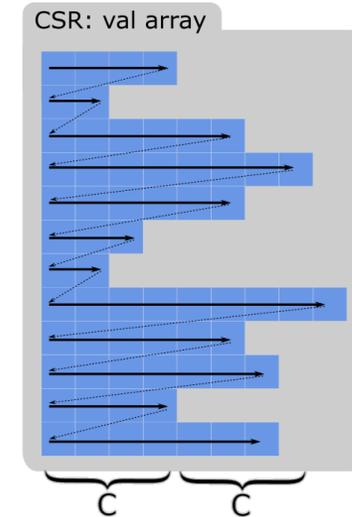
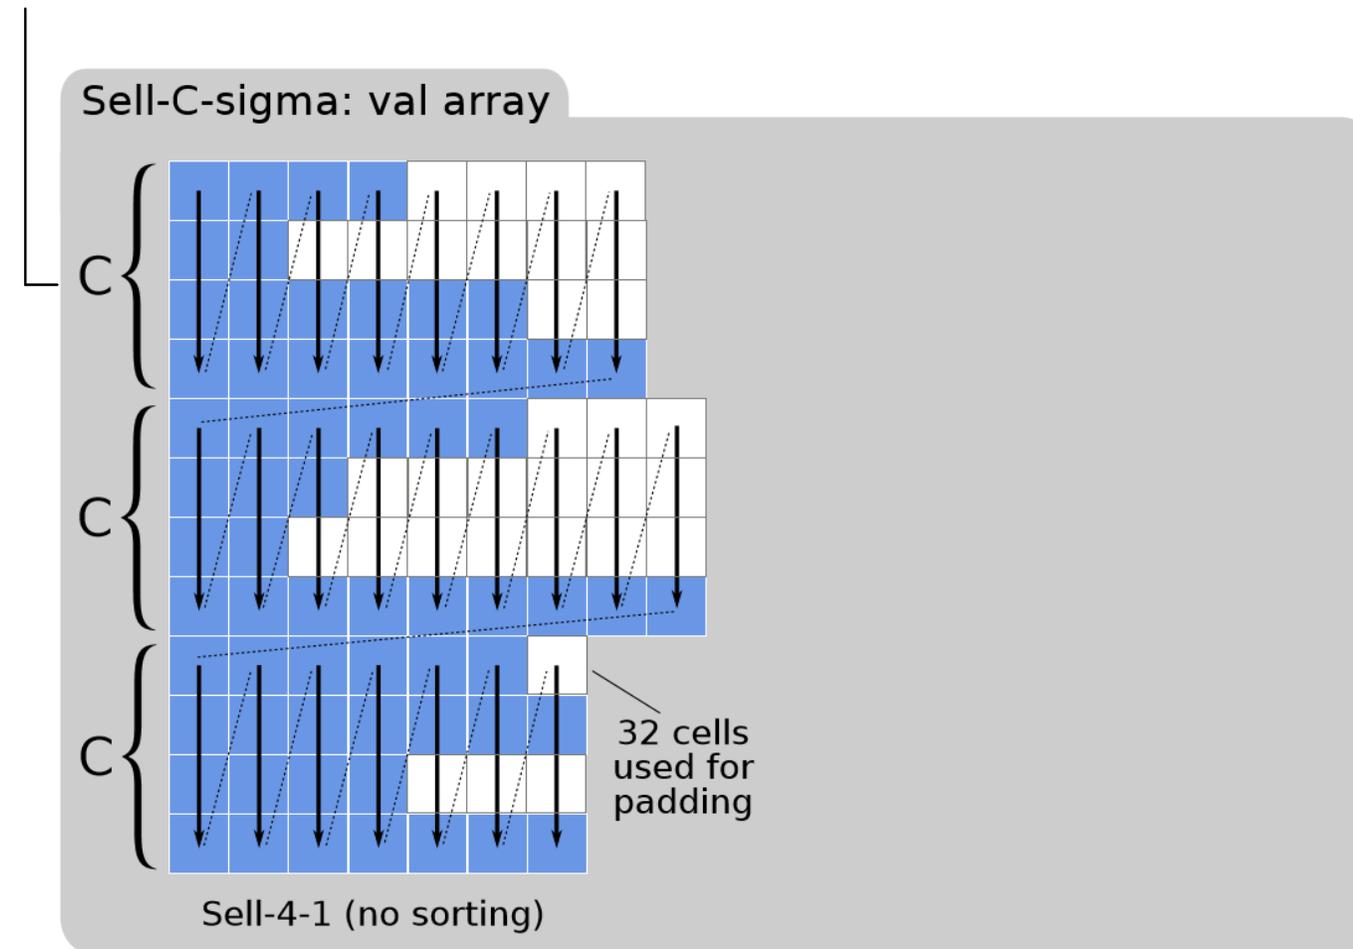
chunk size



GRAPH REPRESENTATIONS

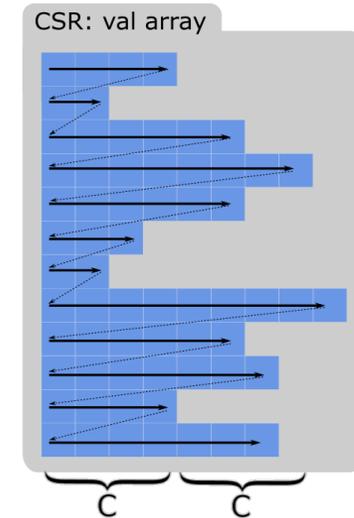
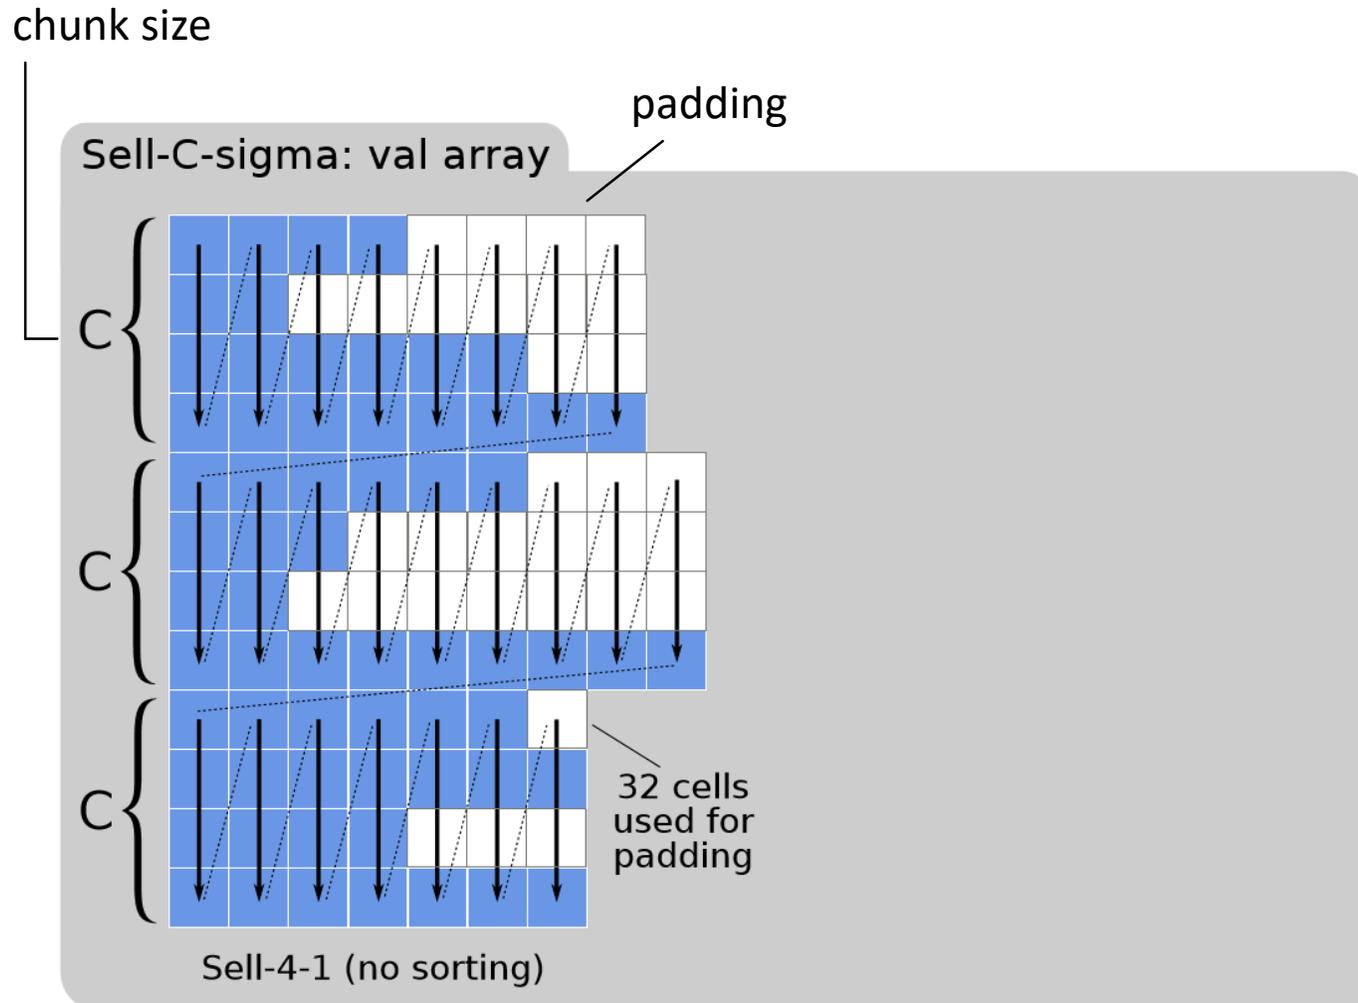
SELL-C-SIGMA

chunk size



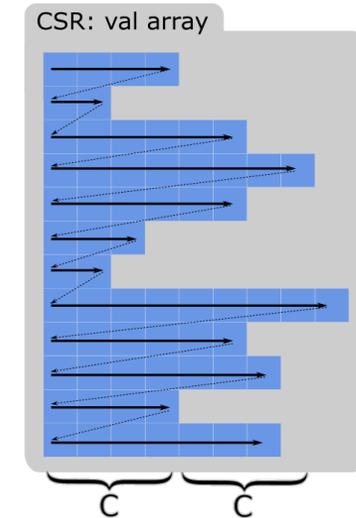
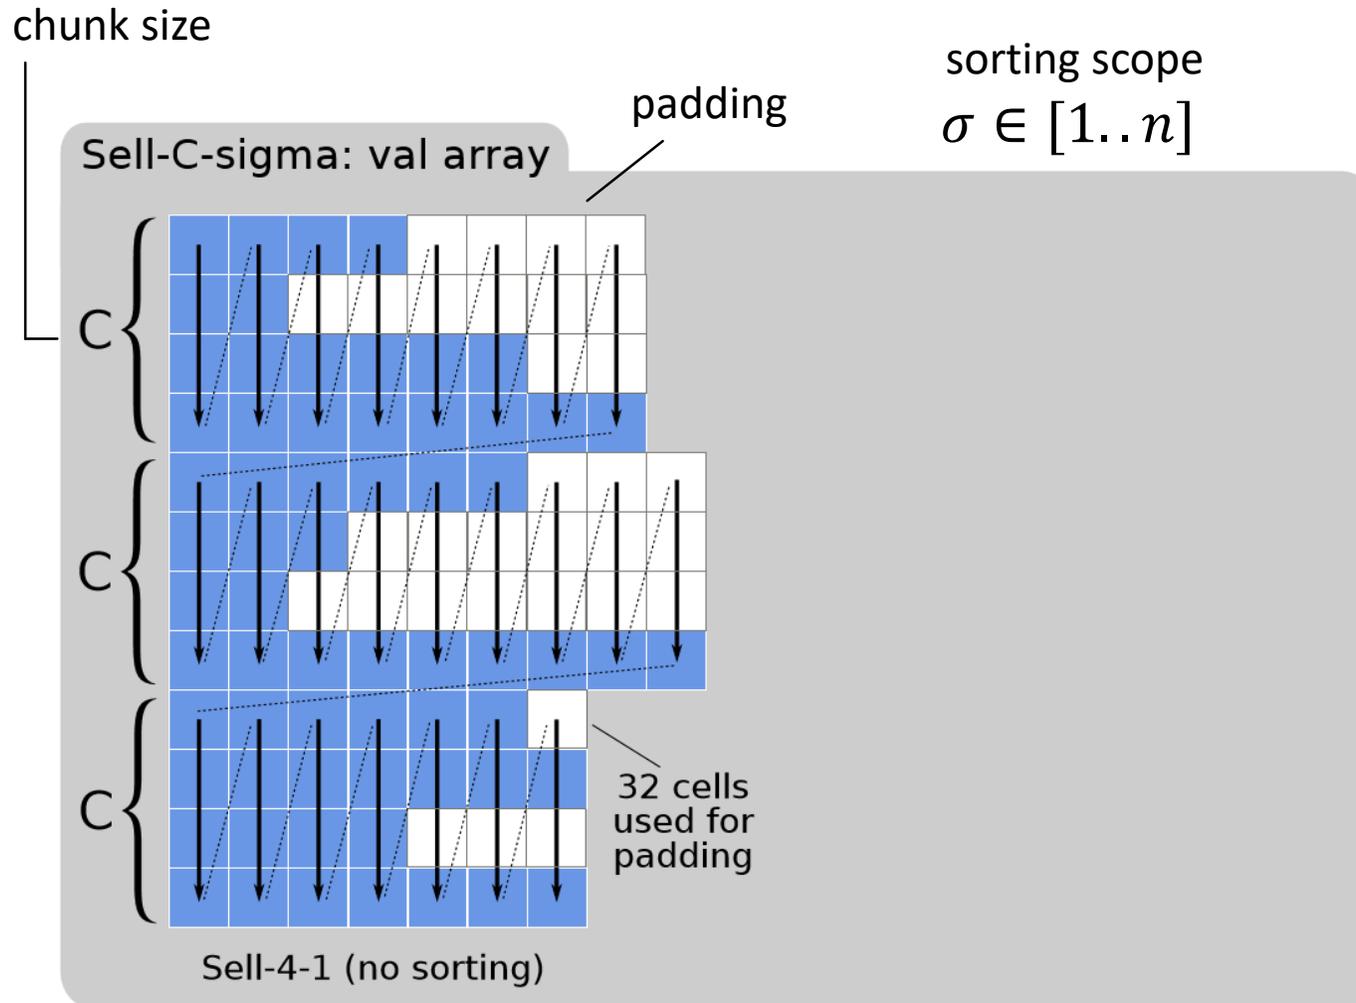
GRAPH REPRESENTATIONS

SELL-C-SIGMA



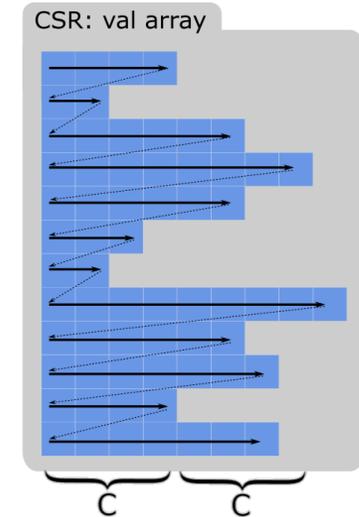
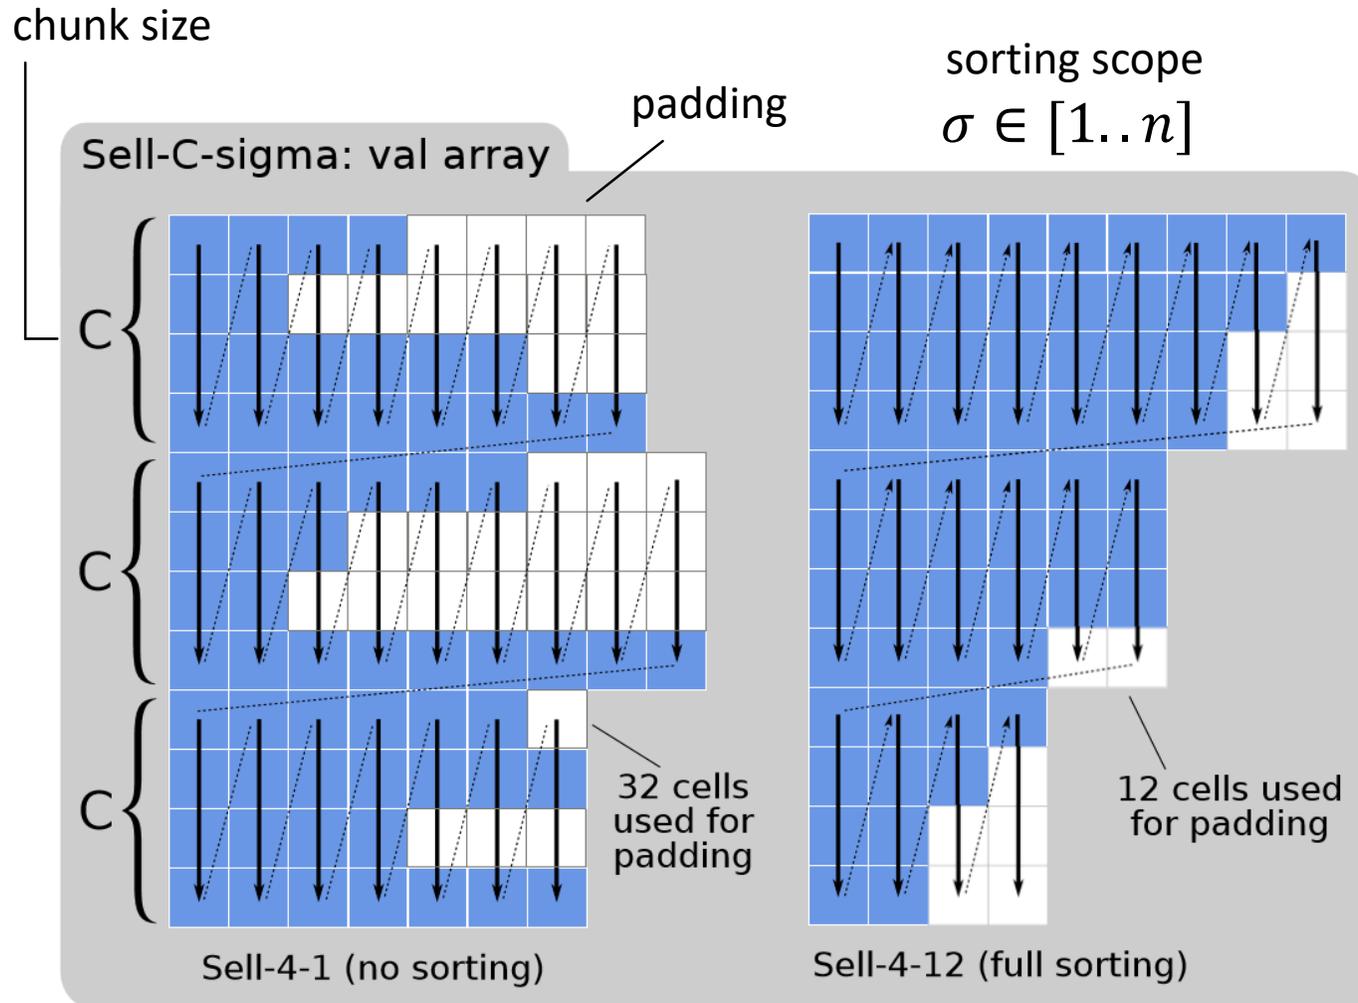
GRAPH REPRESENTATIONS

SELL-C-SIGMA



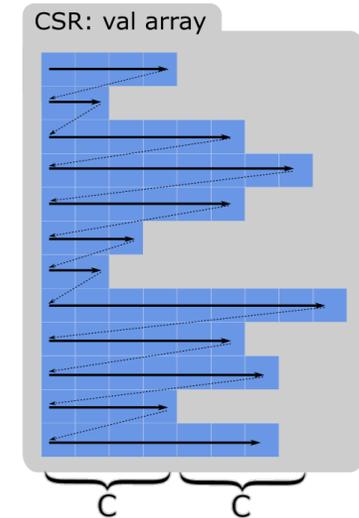
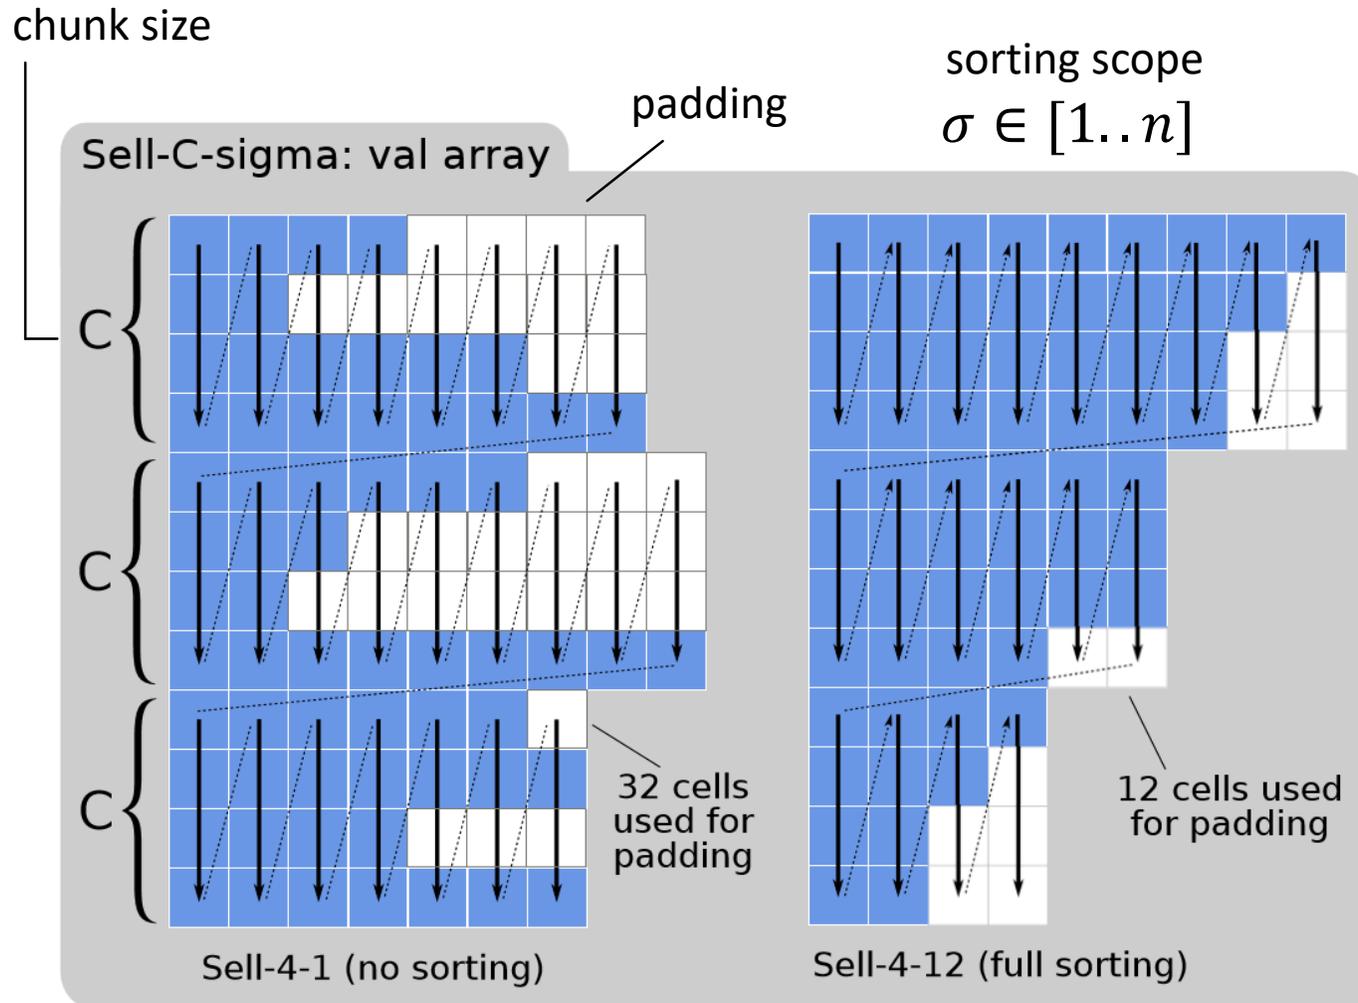
GRAPH REPRESENTATIONS

SELL-C-SIGMA



GRAPH REPRESENTATIONS

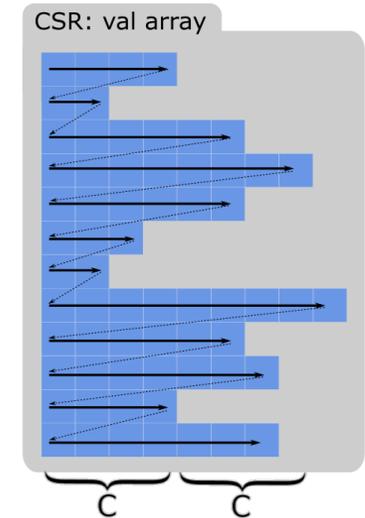
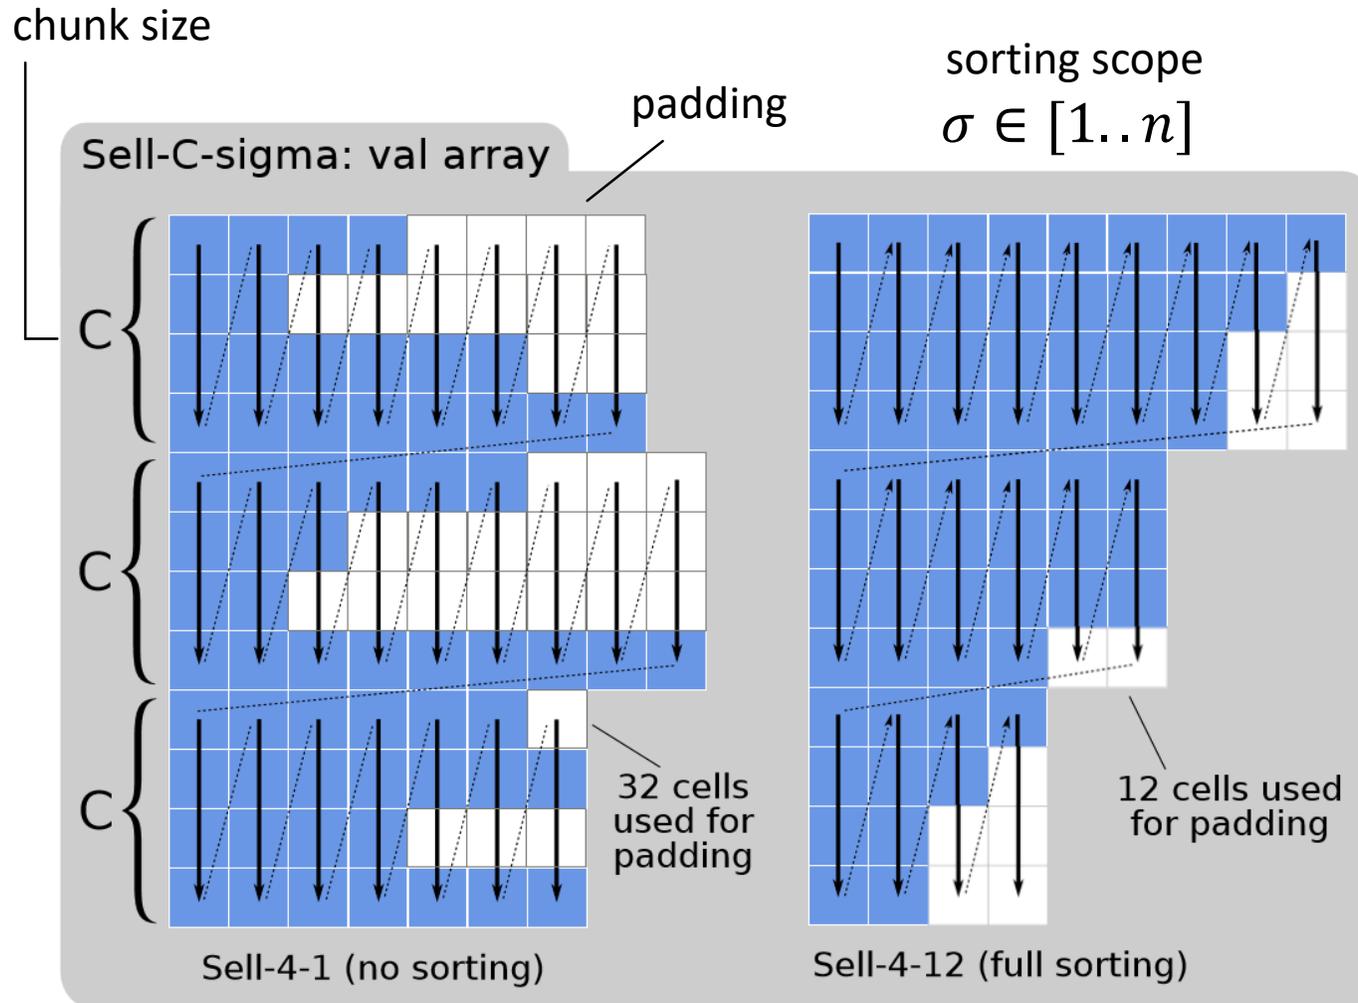
SELL-C-SIGMA



Reductions fast with SIMD operations

GRAPH REPRESENTATIONS

SELL-C-SIGMA

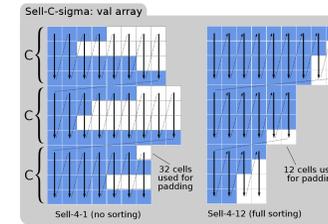


Reductions fast with SIMD operations

Portable

SELL-C-SIGMA + SEMIRINGS + (...) = SLIMSELL FORMULATIONS

SELL-C-SIGMA + SEMIRINGS + (...) = SLIMSELL FORMULATIONS



+

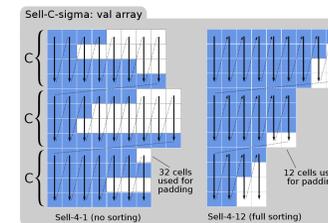
- $(X, op_1, op_2, el_1, el_2)$
- $(\mathbb{R} \cup \{\infty\}, min, +, \infty, 0)$
- $(\mathbb{R}, +, \cdot, 0, 1)$
- $(\{0,1\}, |, \&, 0, 1)$
- $(\mathbb{R}, max, \cdot, -\infty, 1)$

SELL-C-SIGMA + SEMIRINGS + (...) = SLIMSELL

FORMULATIONS

```

12 // Compute  $x_k$  (versions differ based on the used semiring):
13 #ifndef USE_TROPICAL_SEMIRING
14     x = MIN(ADD(rhs, vals), x);
15 #elif defined USE_BOOLEAN_SEMIRING
16     x = OR(AND(rhs, vals), x);
17 #elif defined USE_SELMAX_SEMIRING
18     x = MAX(MUL(rhs, vals), x);
19 #endif
20     index += C;
21 }
22 // Now, derive  $f_k$  (versions differ based on the used semiring):
23 #ifndef USE_TROPICAL_SEMIRING
24     STORE(&fk[i*C], x); // Just a store.
25 #elif defined USE_BOOLEAN_SEMIRING
26     // First, derive  $f_k$  using filtering.
27     V g = LOAD(&gk-1[i*C]); // Load the filter  $g_{k-1}$ .
28     x = CMP(AND(x, g), [0,0,...,0], NEQ); STORE(&xk[i*C], x);
29
30     // Second, update distances  $d$ ; depth is the iteration number.
31     V x_mask = x; x = MUL(x, [depth,...,depth]);
32     x = BLEND(LOAD(&d[i*C]), x, x_mask); STORE(&d[i*C], x);
33
34     // Third, update the filtering term.
35     g = AND(NOT(x_mask), g); STORE(&gk[i*C], g);
36 #elif defined USE_SELMAX_SEMIRING:
37     // Update parents.
38     V pars = LOAD(&pk-1[i*C]); // Load the required part of  $p_{k-1}$ 
39     V pnz = CMP(pars, [0,0,...,0], NEQ);
40     pars = BLEND([0,0,...,0], pars, pnz); STORE(&pk[i*C], pars);
41
42     // Set new  $x_k$  vector.
43     V tmpnz = CMP(x, [0,0,...,0], NEQ);
44     x = BLEND(x, &v[i*C], tmpnz); STORE(&xk[i*C], x);
45 #endif
    
```



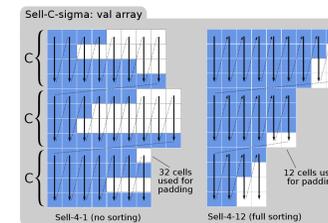
+

 $(X, op_1, op_2, el_1, el_2)$
 $(\mathbb{R} \cup \{\infty\}, min, +, \infty, 0)$
 $(\mathbb{R}, +, 0, 1)$
 $(\{0,1\}, |, \&, 0, 1)$
 $(\mathbb{R}, max, -, -\infty, 1)$

SELL-C-SIGMA + SEMIRINGS + (...) = SLIMSELL FORMULATIONS

```

12 // Compute  $x_k$  (versions differ based on the used semiring):
13 #ifndef USE_TROPICAL_SEMIRING
14     x = MIN(ADD(rhs, vals), x);
15 #elif defined USE_BOOLEAN_SEMIRING
16     x = OR(AND(rhs, vals), x);
17 #elif defined USE_SELMAX_SEMIRING
18     x = MAX(MUL(rhs, vals), x);
19 #endif
20     index += C;
21 }
22 // Now, derive  $f_k$  (versions differ based on the used semiring):
23 #ifndef USE_TROPICAL_SEMIRING
24     STORE(&fk[i*C], x); // Just a store.
25 #elif defined USE_BOOLEAN_SEMIRING
26     // First, derive  $f_k$  using filtering.
27     V g = LOAD(&gk-1[i*C]); // Load the filter  $g_{k-1}$ .
28     x = CMP(AND(x, g), [0,0,...0], NEQ); STORE(&xk[i*C], x);
29
30     // Second, update distances  $d$ ; depth is the iteration number.
31     V x_mask = x; x = MUL(x, [depth,...,depth]);
32     x = BLEND(LOAD(&d[i*C]), x, x_mask); STORE(&d[i*C], x);
33
34     // Third, update the filtering term.
35     g = AND(NOT(x_mask), g); STORE(&gk[i*C], g);
36 #elif defined USE_SELMAX_SEMIRING:
37     // Update parents.
38     V pars = LOAD(&pk-1[i*C]); // Load the required part of  $p_{k-1}$ 
39     V pnz = CMP(pars, [0,0,...,0], NEQ);
40     pars = BLEND([0,0,...,0], pars, pnz); STORE(&pk[i*C], pars);
41
42     // Set new  $x_k$  vector.
43     V tmpnz = CMP(x, [0,0,...,0], NEQ);
44     x = BLEND(x, &v[i*C], tmpnz); STORE(&xk[i*C], x);
45 #endif
    
```



+

 $(X, op_1, op_2, el_1, el_2)$
 $(\mathbb{R} \cup \{\infty\}, min, +, \infty, 0)$
 $(\mathbb{R}, +, 0, 1)$
 $(\{0,1\}, |, \&, 0, 1)$
 $(\mathbb{R}, max, -, -\infty, 1)$

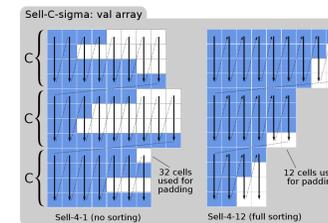
! Detailed formulations are in the paper 😊

SELL-C-SIGMA + SEMIRINGS + (...) = SLIMSELL

FORMULATIONS

```

12 // Compute  $\mathbf{x}_k$  (versions differ based on the used semiring):
13 #ifndef USE_TROPICAL_SEMIRING
14     x = MIN(ADD(rhs, vals), x);
15 #elif defined USE_BOOLEAN_SEMIRING
16     x = OR(AND(rhs, vals), x);
17 #elif defined USE_SELMAX_SEMIRING
18     x = MAX(MUL(rhs, vals), x);
19 #endif
20     index += C;
21 }
22 // Now, derive  $\mathbf{f}_k$  (versions differ based on the used semiring):
23 #ifndef USE_TROPICAL_SEMIRING
24     STORE(&fk[i*C], x); // Just a store.
25 #elif defined USE_BOOLEAN_SEMIRING
26     // First, derive  $\mathbf{f}_k$  using filtering.
27     V g = LOAD(&gk-1[i*C]); // Load the filter  $\mathbf{g}_{k-1}$ .
28     x = CMP(AND(x, g), [0,0,...0], NEQ); STORE(&xk[i*C], x);
29
30     // Second, update distances  $\mathbf{d}$ ; depth is the iteration number.
31     V x_mask = x; x = MUL(x, [depth,...,depth]);
32     x = BLEND(LOAD(&d[i*C]), x, x_mask); STORE(&d[i*C], x);
33
34     // Third, update the filtering term.
35     g = AND(NOT(x_mask), g); STORE(&gk[i*C], g);
36 #elif defined USE_SELMAX_SEMIRING:
37     // Update parents.
38     V pars = LOAD(&pk-1[i*C]); // Load the required part of  $\mathbf{p}_{k-1}$ 
39     V pnz = CMP(pars, [0,0,...,0], NEQ);
40     pars = BLEND([0,0,...,0], pars, pnz); STORE(&pk[i*C], pars);
41
42     // Set new  $\mathbf{x}_k$  vector.
43     V tmpnz = CMP(x, [0,0,...,0], NEQ);
44     x = BLEND(x, &v[i*C], tmpnz); STORE(&xk[i*C], x);
45 #endif
    
```



+

$$\begin{aligned}
 &(X, op_1, op_2, el_1, el_2) \\
 &(\mathbb{R} \cup \{\infty\}, \min, +, \infty, 0) \\
 &(\mathbb{R}, +, 0, 1) \\
 &(\{0,1\}, \wedge, \&, 0, 1) \\
 &(\mathbb{R}, \max, -, -\infty, 1)
 \end{aligned}$$

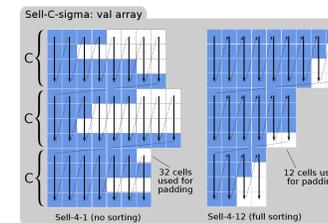
! What vector operations are required for each semiring when using Sell-C-sigma

! Detailed formulations are in the paper 😊

SELL-C-SIGMA + SEMIRINGS + (...) = SLIMSELL FORMULATIONS

```

12 // Compute  $x_k$  (versions differ based on the used semiring):
13 #ifndef USE_TROPICAL_SEMIRING
14     x = MIN(ADD(rhs, vals), x);
15 #elif defined USE_BOOLEAN_SEMIRING
16     x = OR(AND(rhs, vals), x);
17 #elif defined USE_SELMAX_SEMIRING
18     x = MAX(MUL(rhs, vals), x);
19 #endif
20     index += C;
21 }
22 // Now, derive  $f_k$  (versions differ based on the used semiring):
23 #ifndef USE_TROPICAL_SEMIRING
24     STORE(&fk[i*C], x); // Just a store.
25 #elif defined USE_BOOLEAN_SEMIRING
26     // First, derive  $f_k$  using filtering.
27     V g = LOAD(&gk-1[i*C]); // Load the filter  $g_{k-1}$ .
28     x = CMP(AND(x, g), [0,0,...0], NEQ); STORE(&xk[i*C], x);
29
30     // Second, update distances  $d$ ; depth is the iteration number.
31     V x_mask = x; x = MUL(x, [depth,...,depth]);
32     x = BLEND(LOAD(&d[i*C]), x, x_mask); STORE(&d[i*C], x);
33
34     // Third, update the filtering term.
35     g = AND(NOT(x_mask), g); STORE(&gk[i*C], g);
36 #elif defined USE_SELMAX_SEMIRING:
37     // Update parents.
38     V pars = LOAD(&pk-1[i*C]); // Load the required part of  $p_{k-1}$ 
39     V pnz = CMP(pars, [0,0,...,0], NEQ);
40     pars = BLEND([0,0,...,0], pars, pnz); STORE(&pk[i*C], pars);
41
42     // Set new  $x_k$  vector.
43     V tmpnz = CMP(x, [0,0,...,0], NEQ);
44     x = BLEND(x, &v[i*C], tmpnz); STORE(&xk[i*C], x);
45 #endif
    
```



+

$$(X, op_1, op_2, el_1, el_2)$$

$$(\mathbb{R} \cup \{\infty\}, \min, +, \infty, 0)$$

$$(\mathbb{R}, +, 0, 1)$$

$$(\{0,1\}, \&, 0, 1)$$

$$(\mathbb{R}, \max, -, -\infty, 1)$$

! What vector operations are required for each semiring when using Sell-C-sigma

! Detailed formulations are in the paper 😊

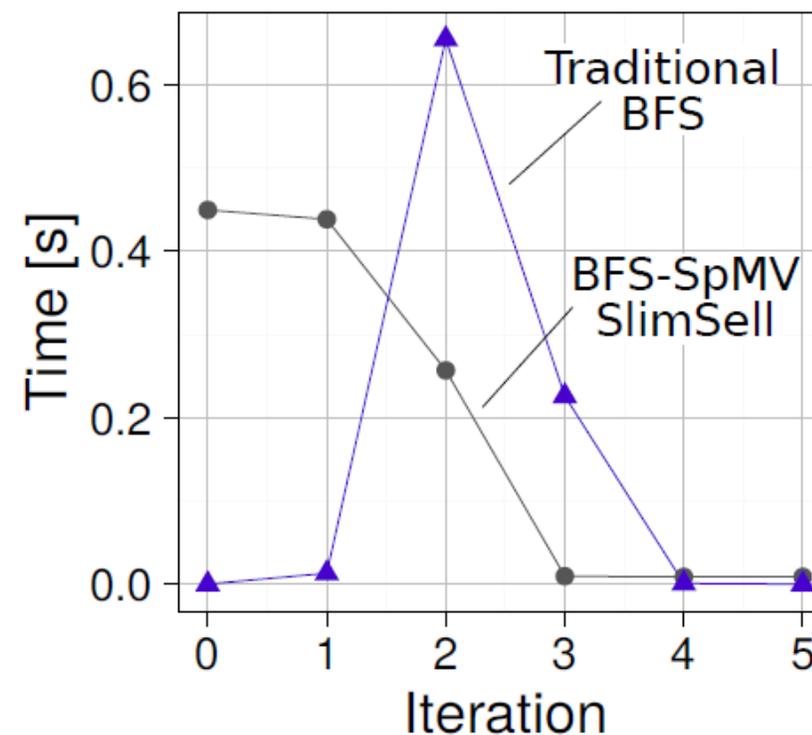
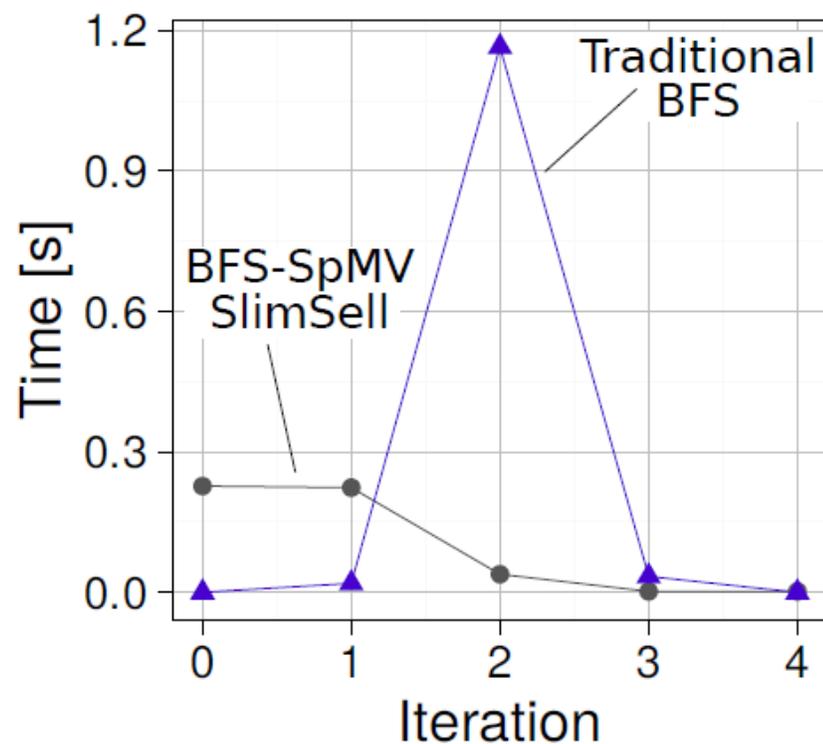
PERFORMANCE ANALYSIS

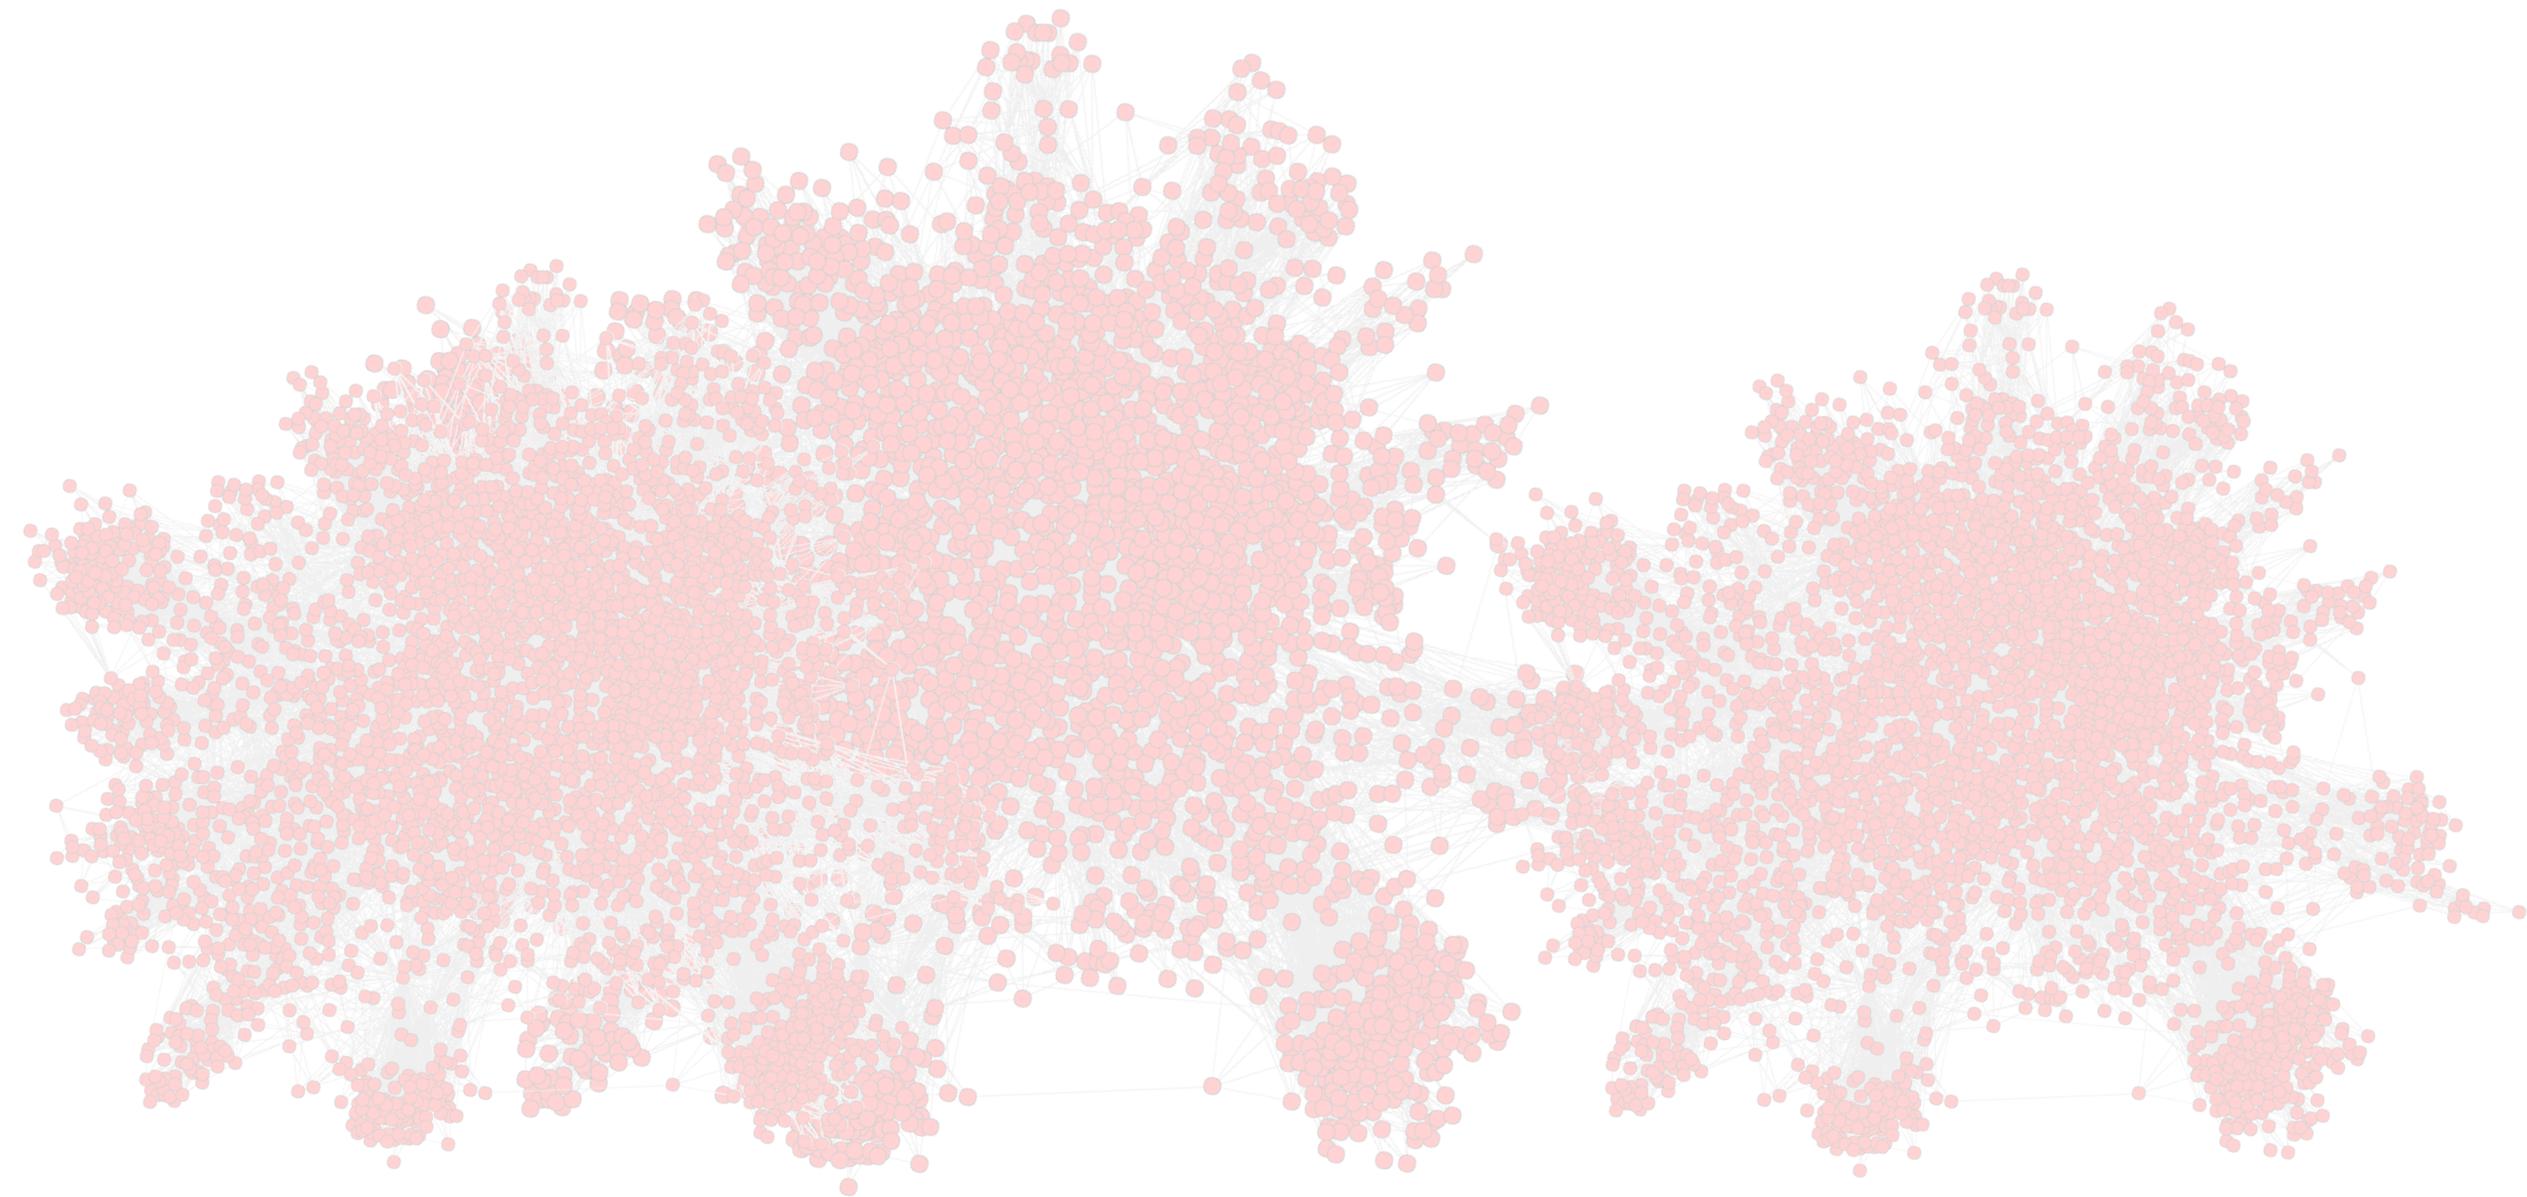
COMPARISON TO GRAPH500

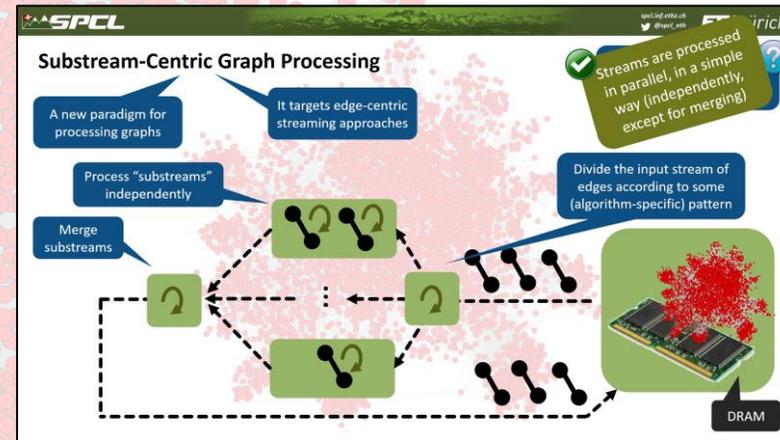
Kronecker power-law graphs



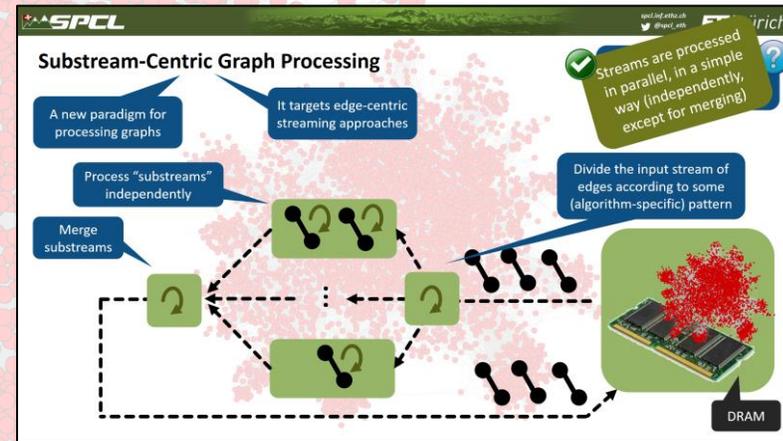
Intel KNL, $C = 16$
 $\log \sigma \in \{20,21,22\}$
 Dynamic scheduling







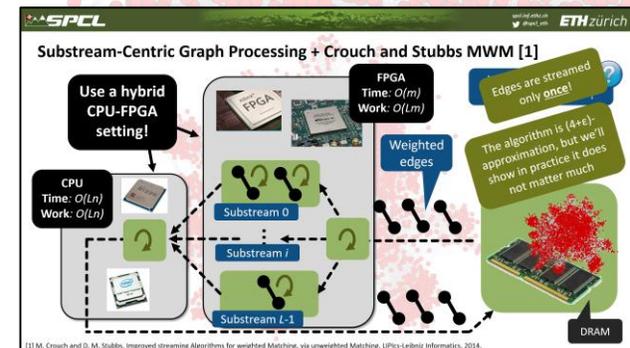
SUBSTREAM-CENTRIC GRAPH PROCESSING PARADIGM:
EXPOSES PARALLELISM, ENABLES EASY PIPELINING, SUPPORTS APPROXIMATION

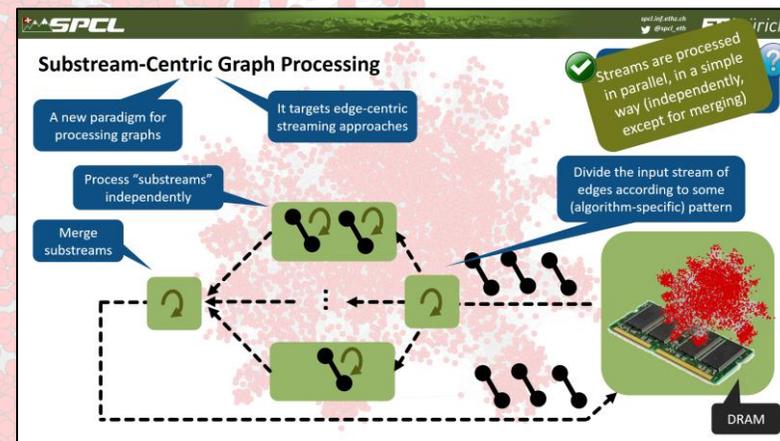


SUBSTREAM-CENTRIC GRAPH PROCESSING PARADIGM:

EXPOSES PARALLELISM, ENABLES EASY PIPELINING, SUPPORTS APPROXIMATION

THEORY-INSPIRED MWM APPROXIMATE ALGORITHM ON A HYBRID CPU-FPGA SETTING

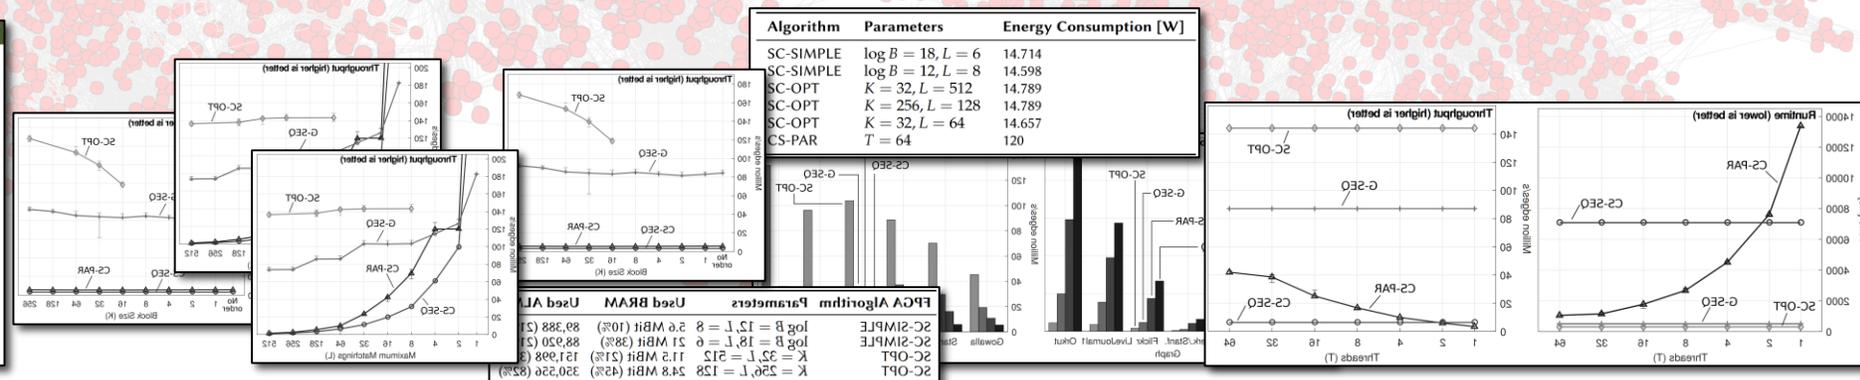
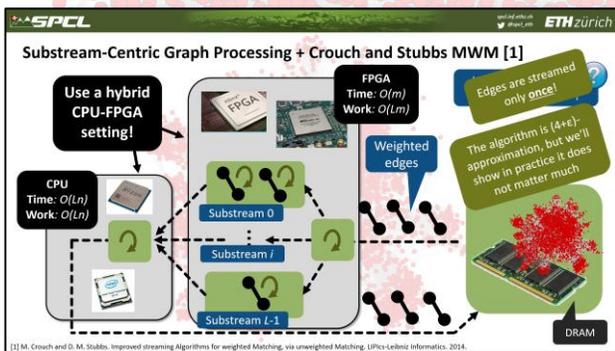


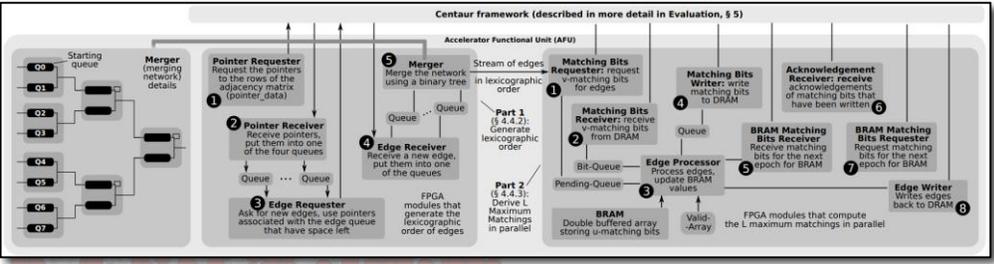


SUBSTREAM-CENTRIC GRAPH PROCESSING PARADIGM:

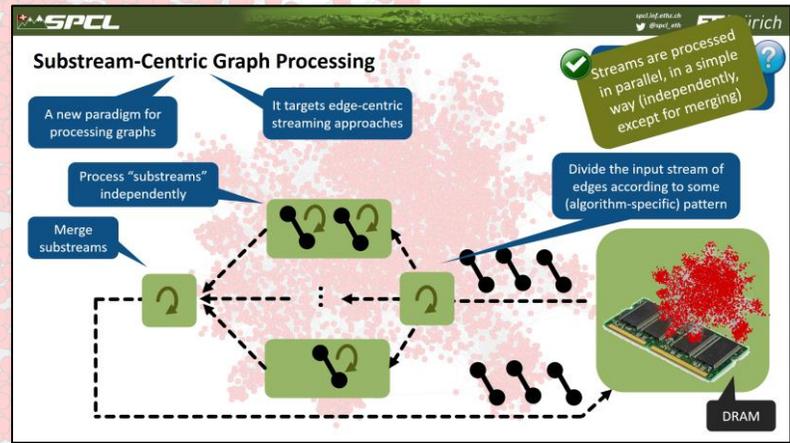
EXPOSES PARALLELISM, ENABLES EASY PIPELINING, SUPPORTS APPROXIMATION

THEORY-INSPIRED MWM APPROXIMATE ALGORITHM ON A HYBRID CPU-FPGA SETTING



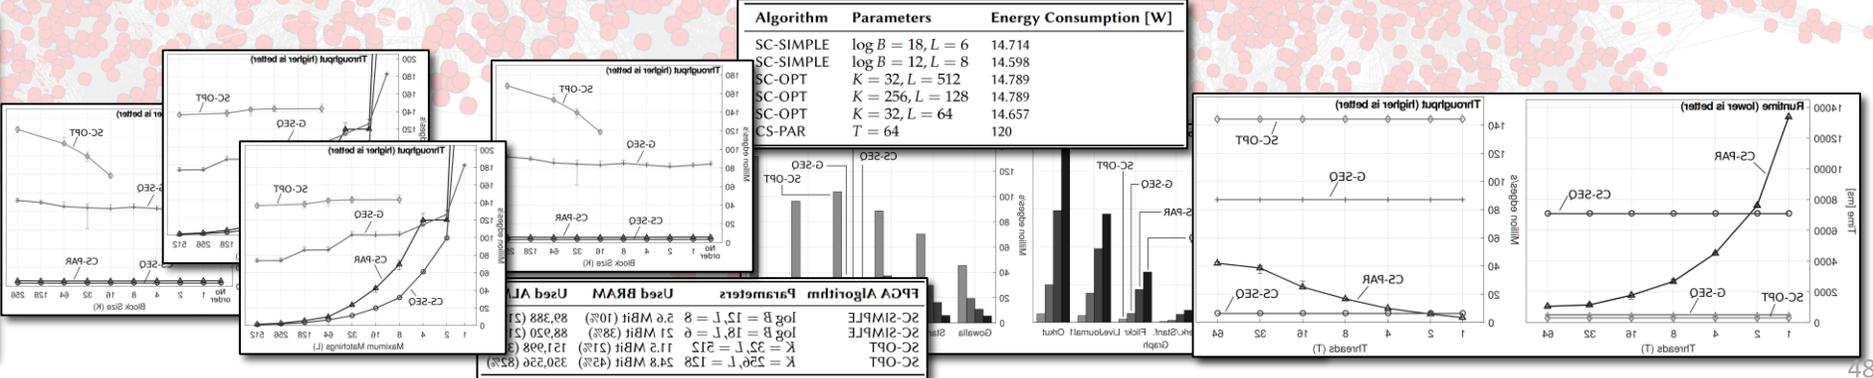
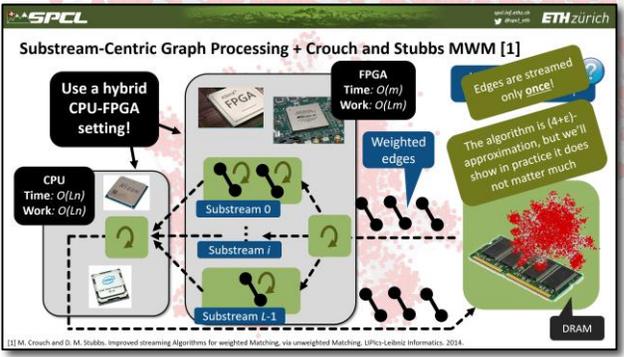


GENERIC FPGA DESIGN, CODE AVAILABLE



THEORY-INSPIRED MWM APPROXIMATE ALGORITHM ON A HYBRID CPU-FPGA SETTING

SUBSTREAM-CENTRIC GRAPH PROCESSING PARADIGM: EXPOSES PARALLELISM, ENABLES EASY PIPELINING, SUPPORTS APPROXIMATION



Substream-Centric MWM: FPGA design

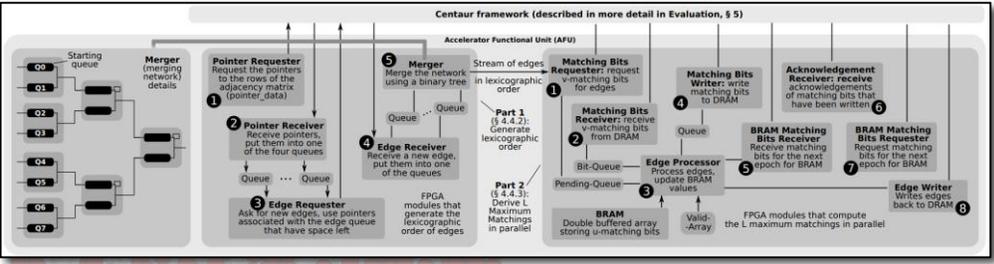
Vectorization

Blocking / Tiling

Prefetching

Pipelining

No, we will not analyze this now
All the details are in the paper.
Let's focus on the key FPGA design ideas and optimizations



OTHER ALGORITHMS, PROBLEMS, ANALYSES

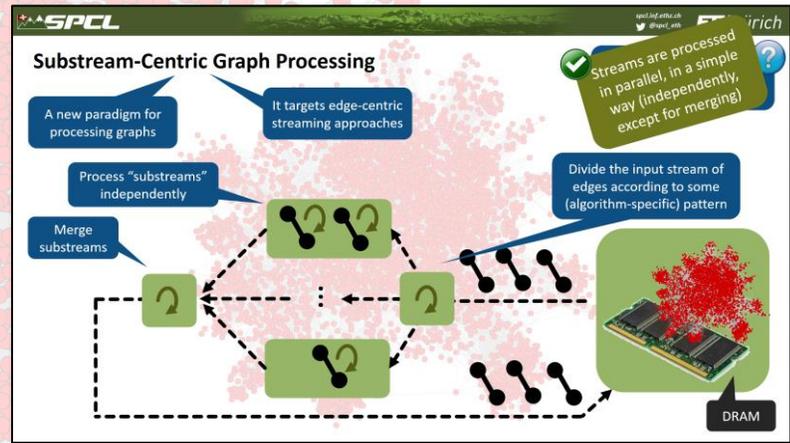
Well, not enough time to present ☹️

In addition to MWM, we also analyzed more graph problems

<https://arxiv.org/abs/...>

Graph Processing on FPGAs: Taxonomy, Survey, Challenges

GENERIC FPGA DESIGN, CODE AVAILABLE



GENERALIZABILITY TO OTHER GRAPH PROBLEMS AND SETTINGS

THEORY-INSPIRED MWM APPROXIMATE ALGORITHM ON A HYBRID CPU-FPGA SETTING

SUBSTREAM-CENTRIC GRAPH PROCESSING PARADIGM:

EXPOSES PARALLELISM, ENABLES EASY PIPELINING, SUPPORTS APPROXIMATION

Substream-Centric Graph Processing + Crouch and Stubbs MWM [1]

Use a hybrid CPU-FPGA setting!

CPU Time: $O(Ln)$
Work: $O(Ln)$

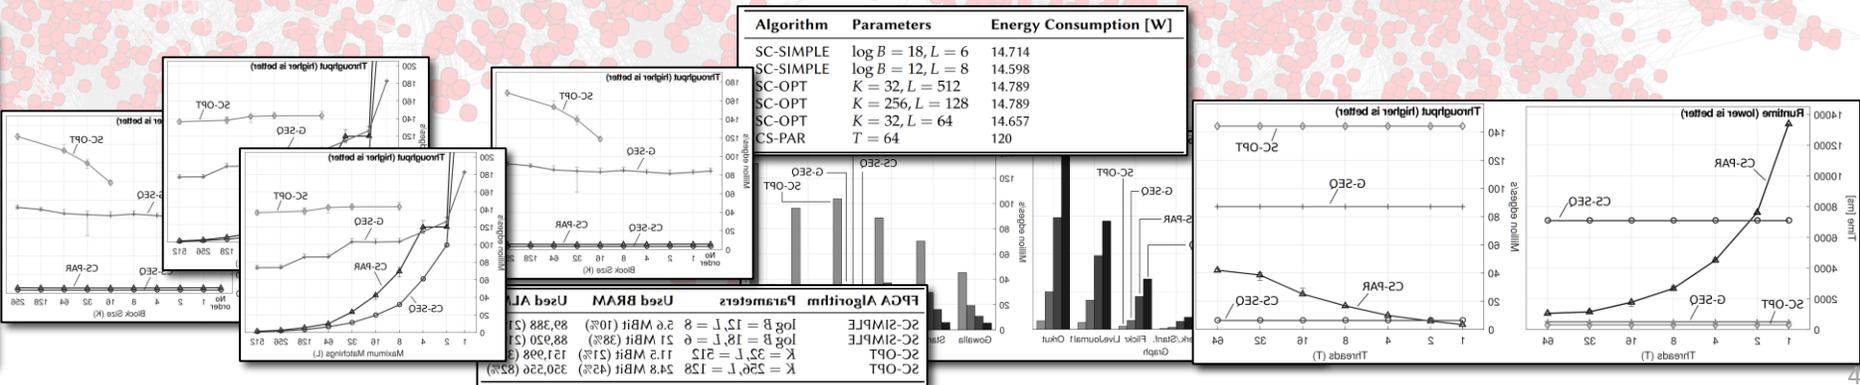
FPGA Time: $O(m)$
Work: $O(Lm)$

Edges are streamed only once!

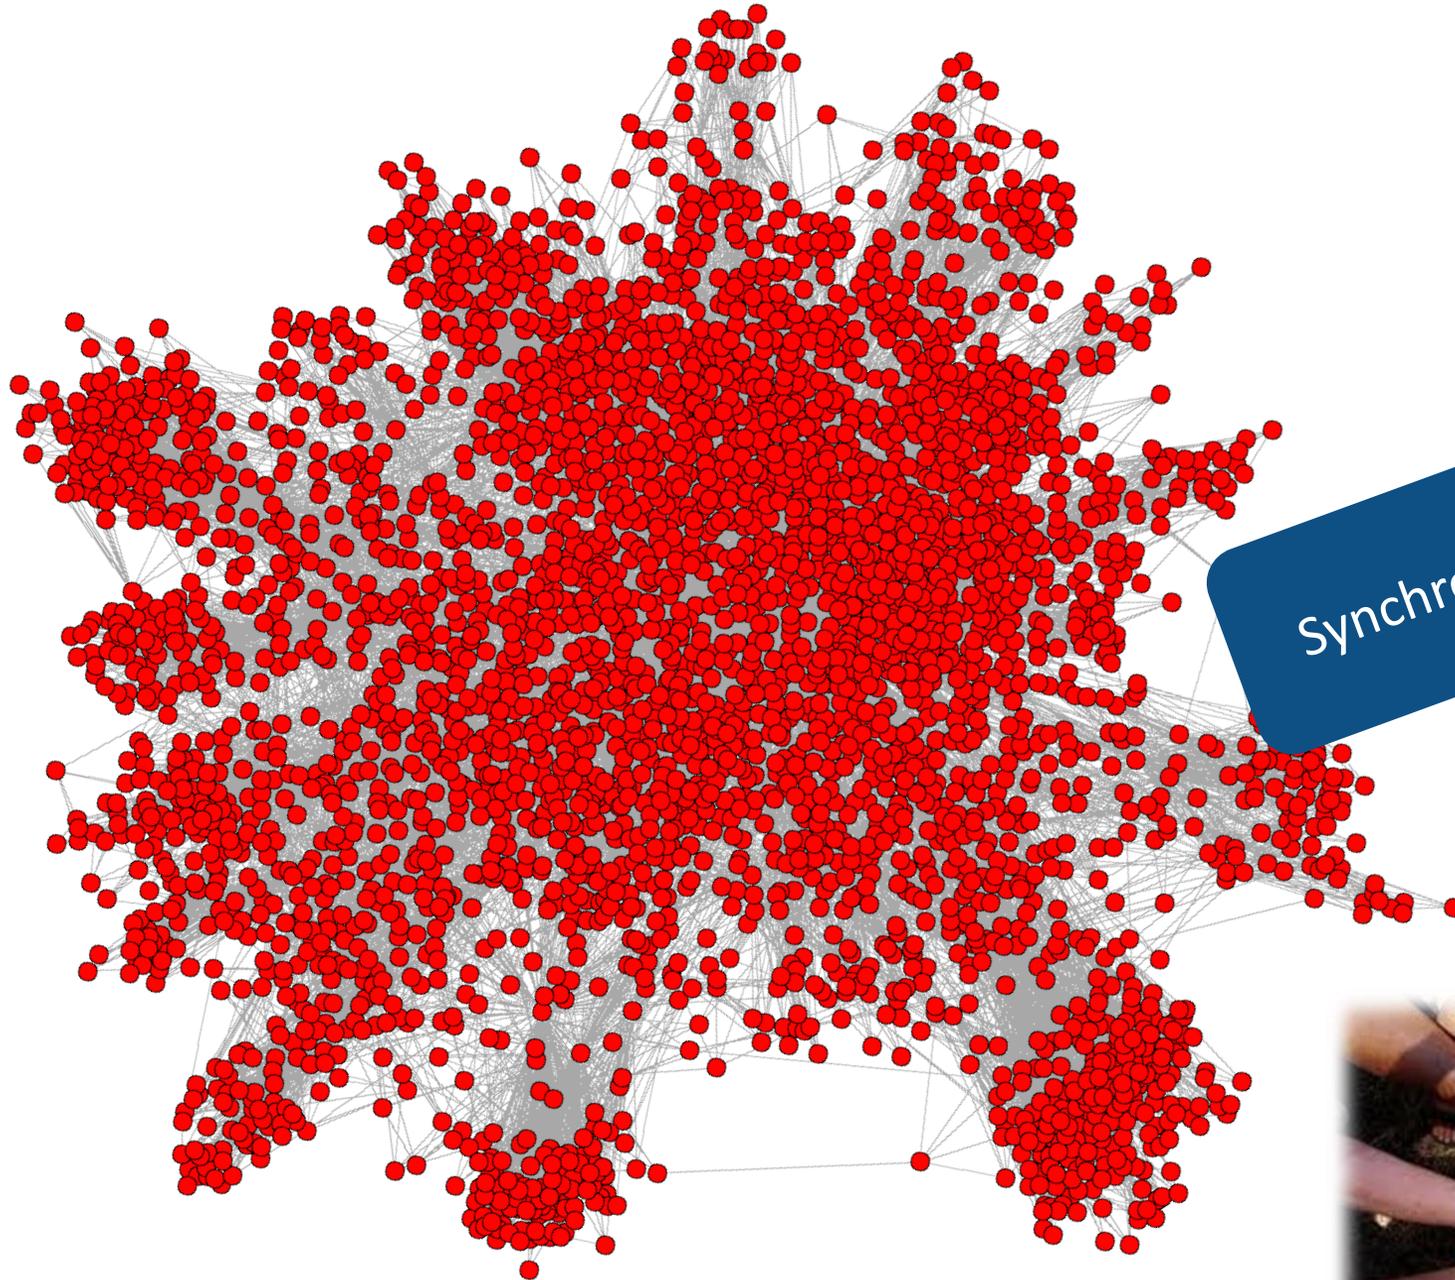
Weighted edges

The algorithm is $(4+\epsilon)$ -approximation, but we'll show in practice it does not matter much

DRAM



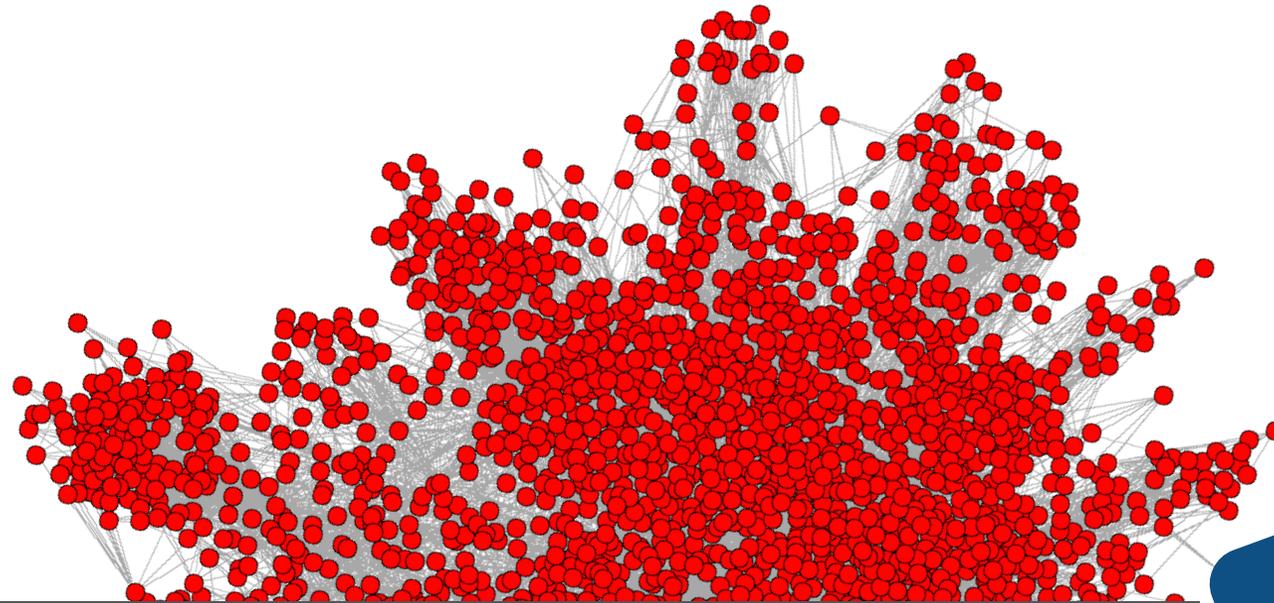
Problems!



Synchronization-heavy



Problems!



Synchronization-heavy



Accelerating Irregular Computations with Hardware Transactional Memory and Active Messages

Maciej Besta
Department of Computer Science
ETH Zurich
Universitätstr. 6, 8092 Zurich, Switzerland
maciej.best@inf.ethz.ch

Torsten Hoefler
Department of Computer Science
ETH Zurich
Universitätstr. 6, 8092 Zurich, Switzerland
htor@inf.ethz.ch

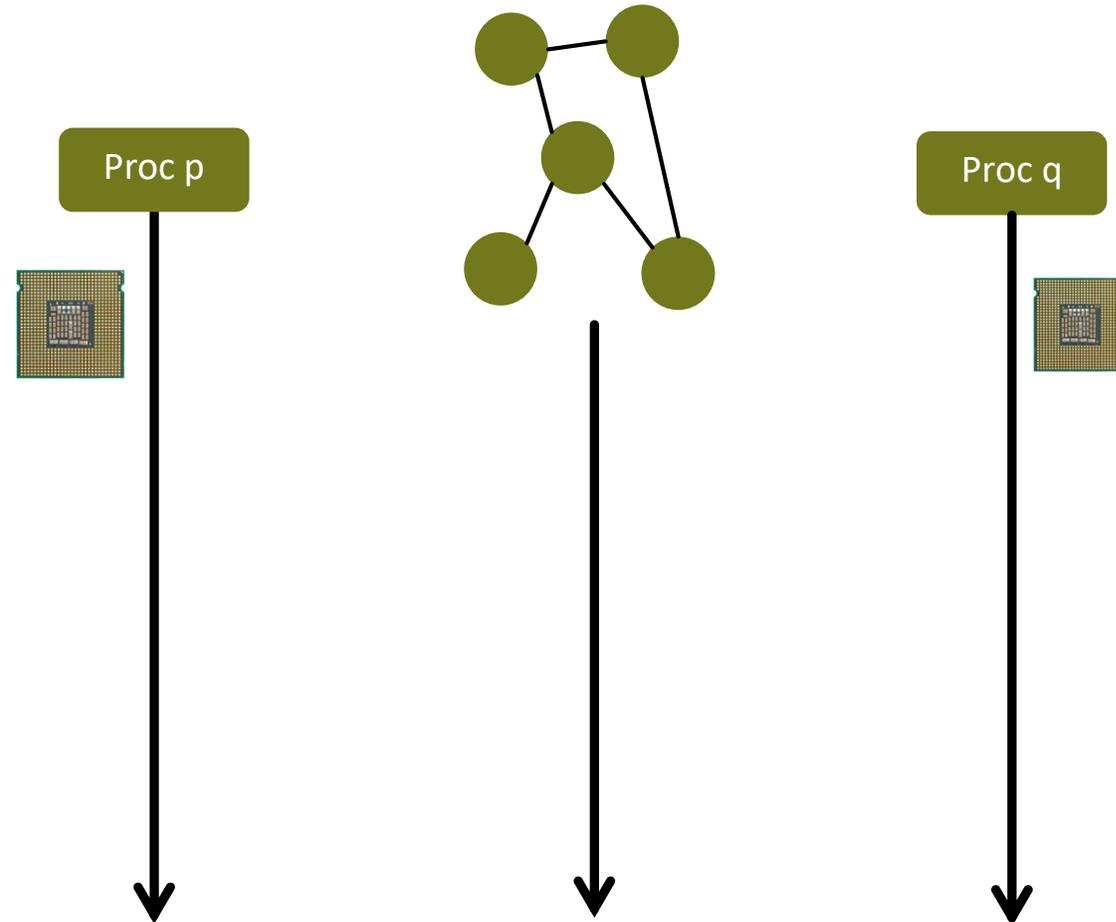
ABSTRACT

We propose Atomic Active Messages (AAM), a mechanism that accelerates irregular graph computations on both shared- and distributed-memory machines. The key idea behind AAM is that hardware transactional memory (HTM) can be used for simple and efficient processing of irregular structures in highly parallel environments. We illustrate techniques such as coarsening and coalescing that enable hard-

tion and become visible to other threads *atomically*. Available HTM implementations show promising performance in scientific codes and industrial benchmarks [40, 36]. In this work, we show that the ease of programming and performance benefits are even more promising for fine-grained, irregular, and data-driven graph computations.

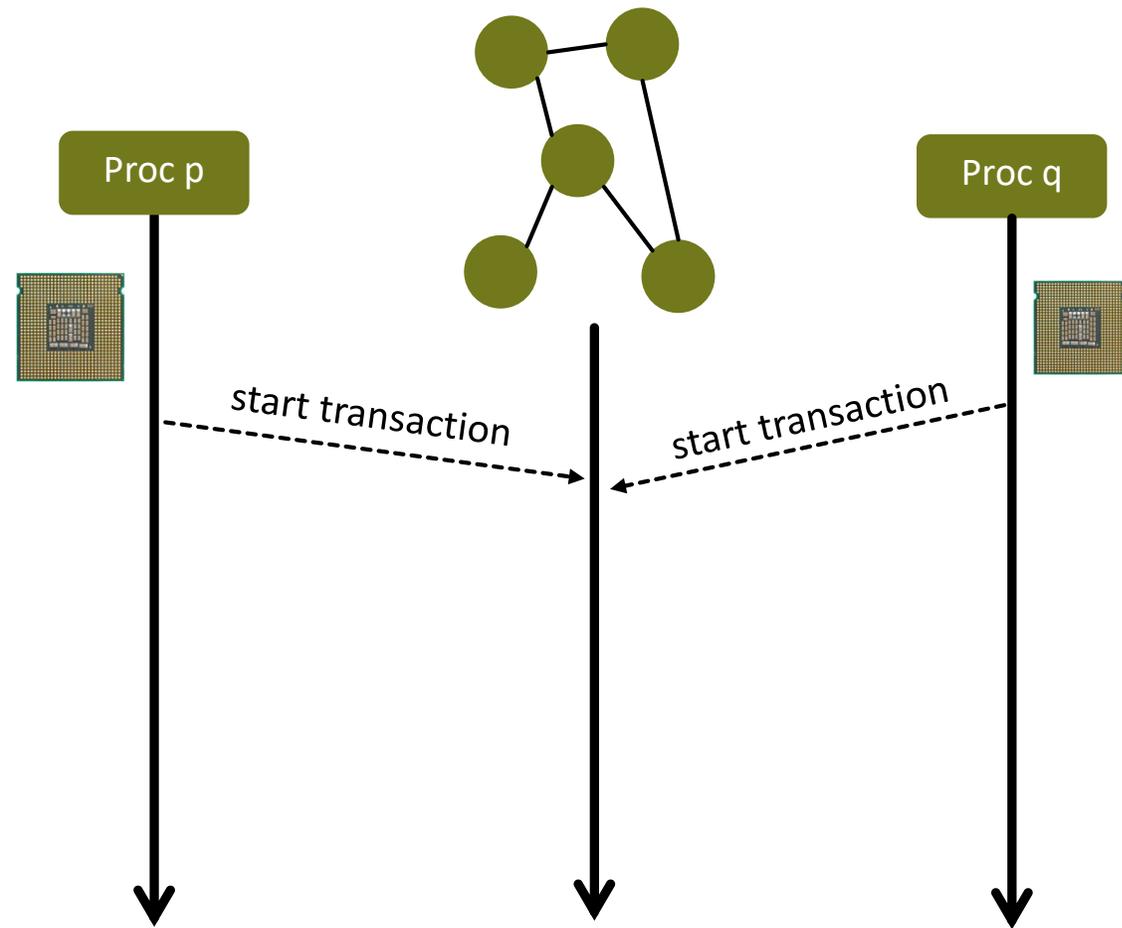
Another challenge of graph analytics is the size of the input that often requires distributed memory machines [27]. Such machines generally contain many more compute nodes

TRANSACTIONAL MEMORY [1]



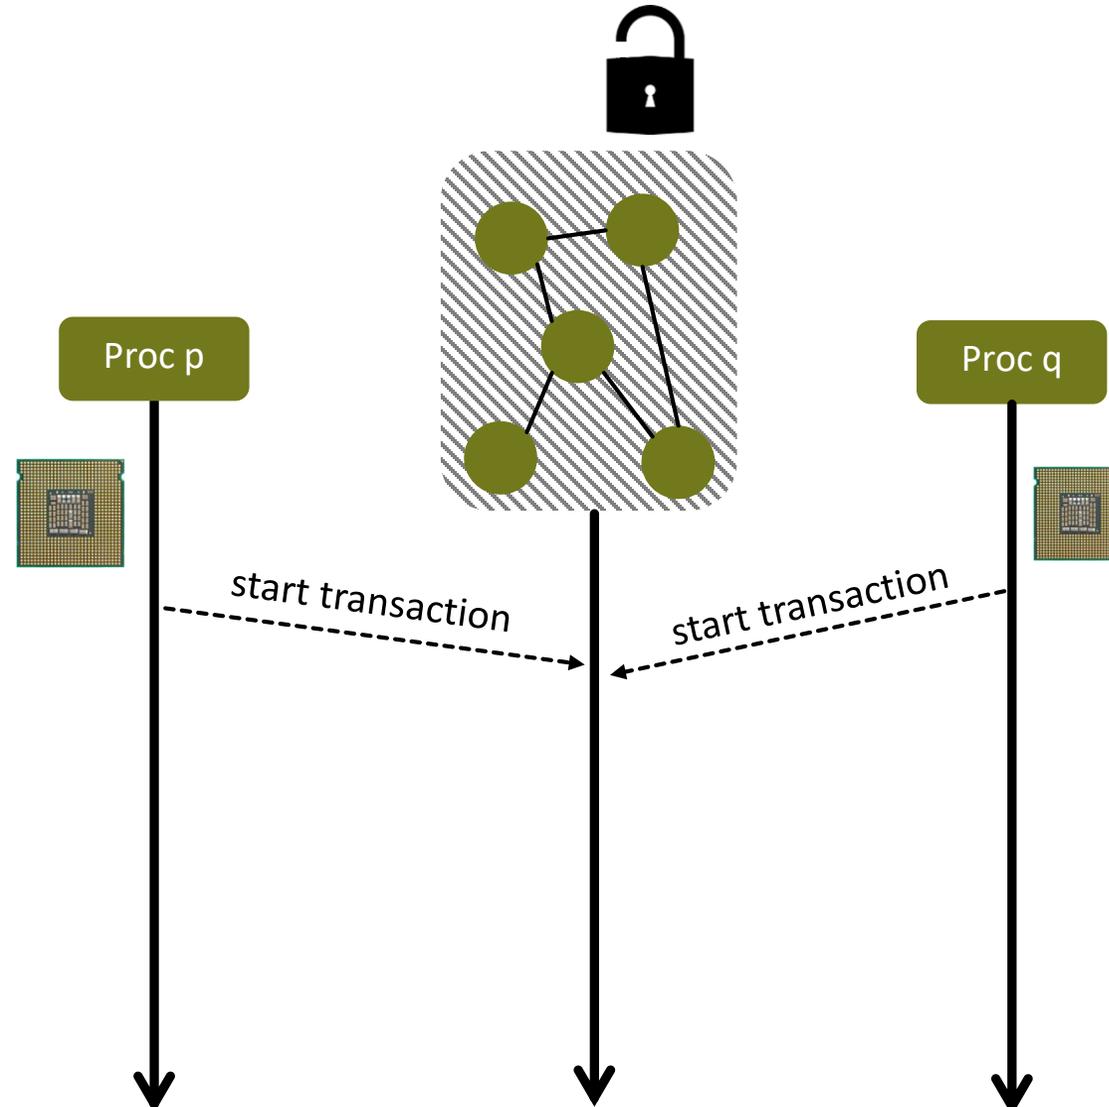
[1] N. Shavit and D. Touitou. Software transactional memory. PODC'95.

TRANSACTIONAL MEMORY [1]



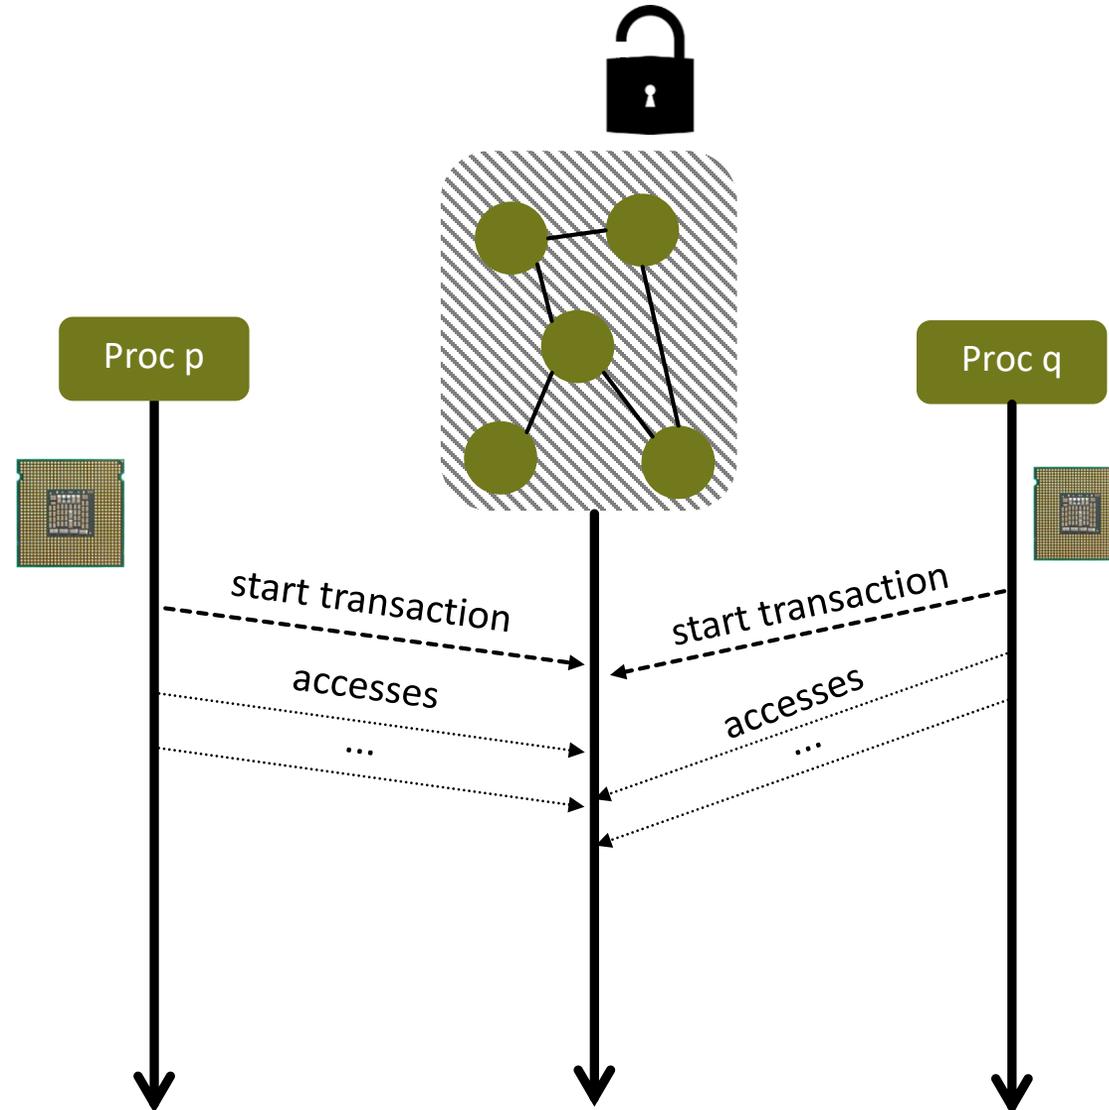
[1] N. Shavit and D. Touitou. Software transactional memory. PODC'95.

TRANSACTIONAL MEMORY [1]



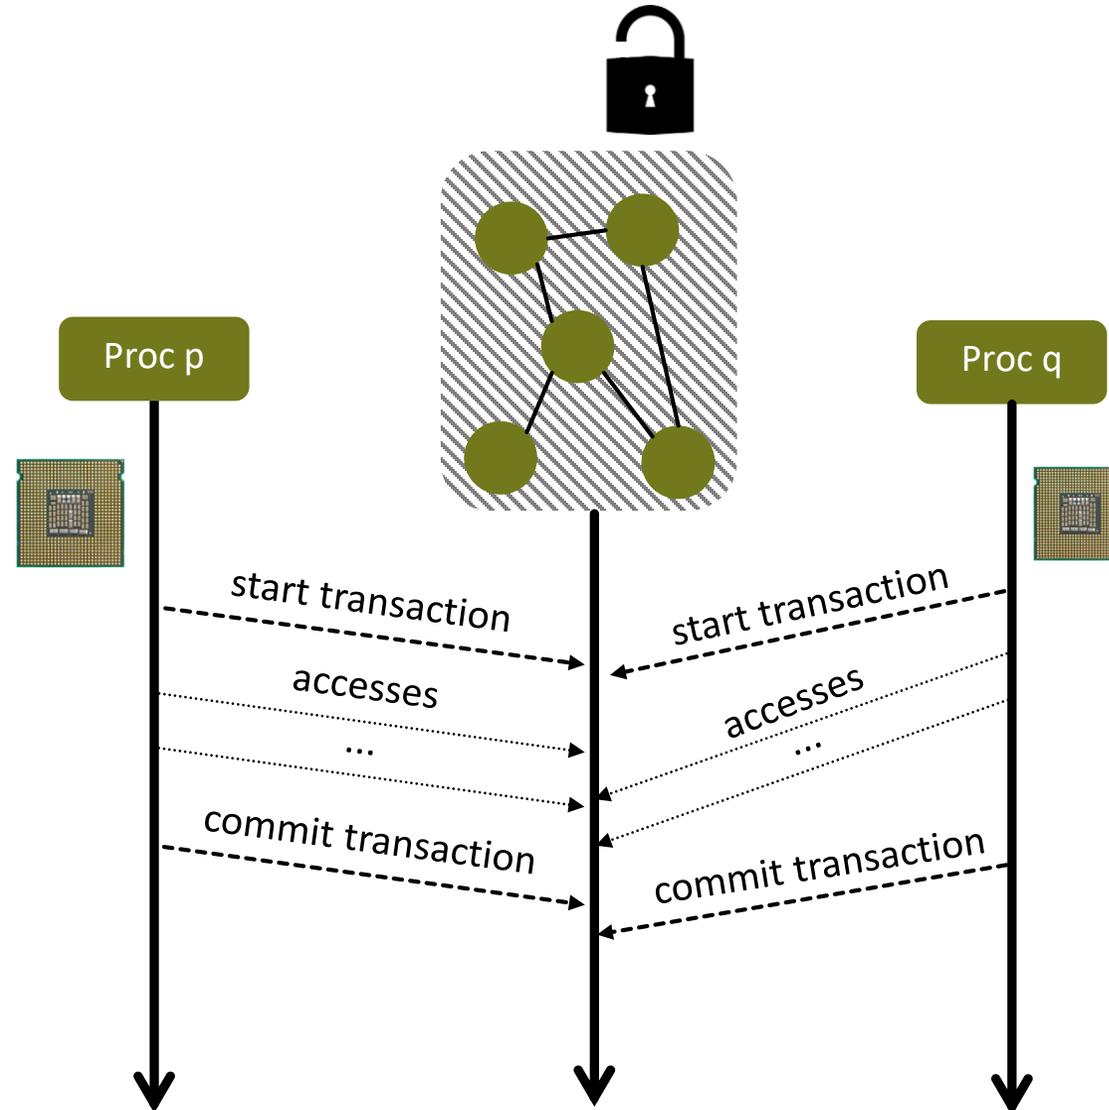
[1] N. Shavit and D. Touitou. Software transactional memory. PODC'95.

TRANSACTIONAL MEMORY [1]



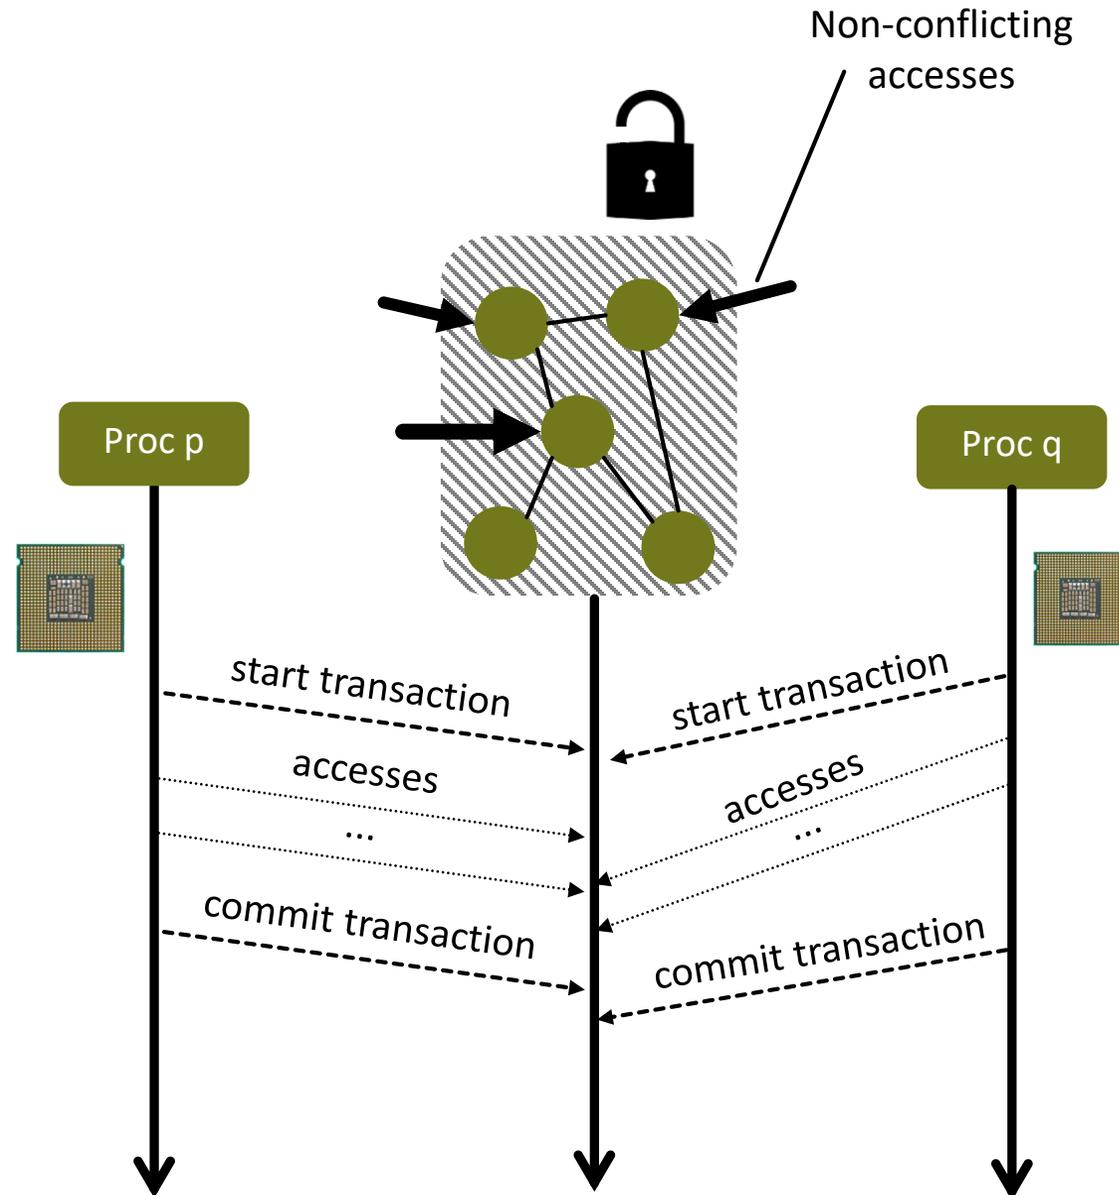
[1] N. Shavit and D. Touitou. Software transactional memory. PODC'95.

TRANSACTIONAL MEMORY [1]



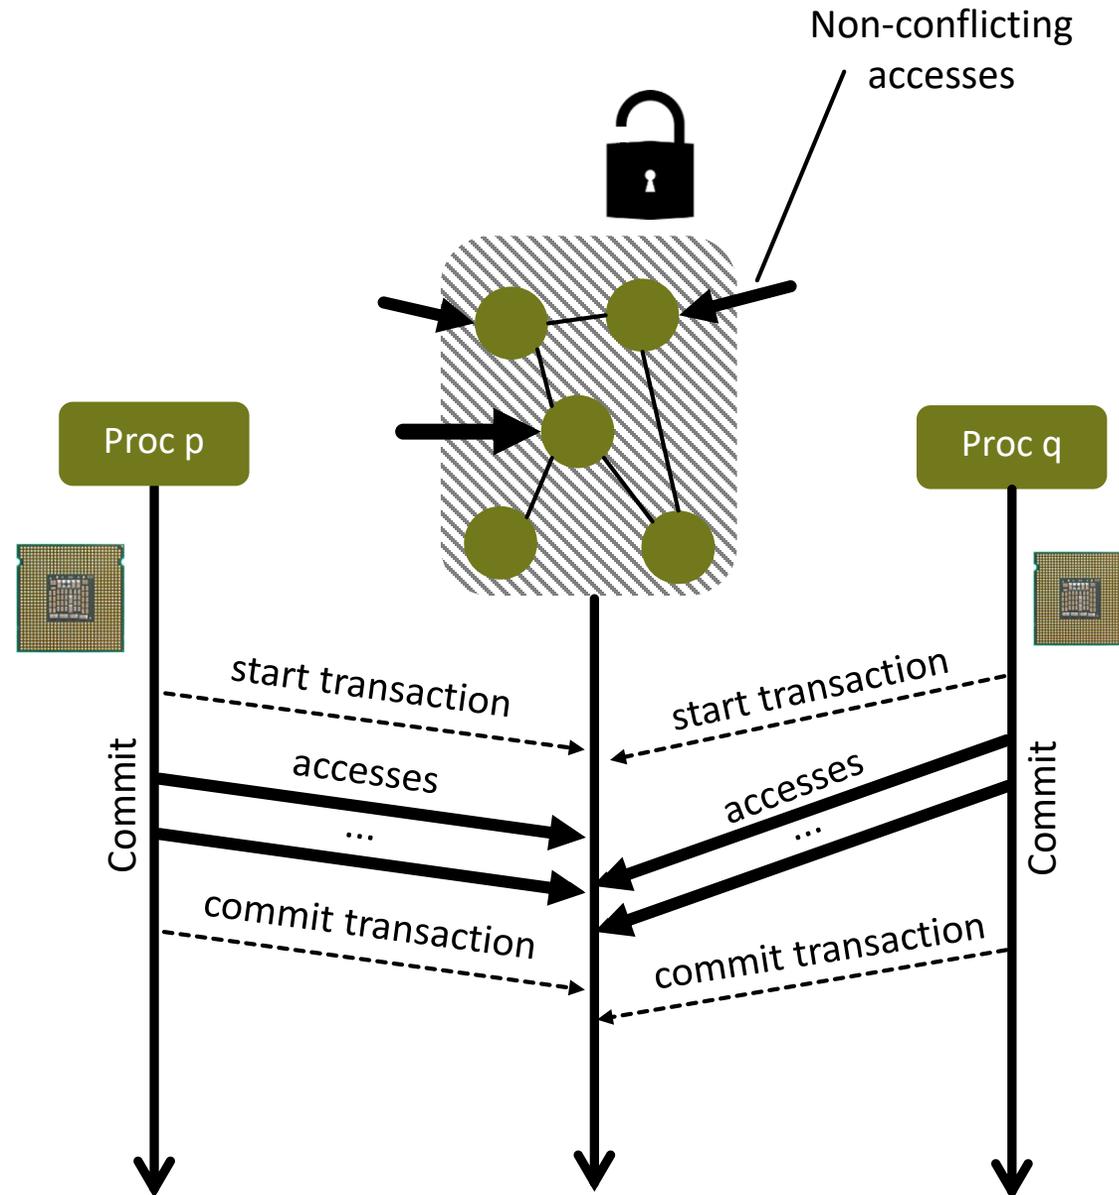
[1] N. Shavit and D. Touitou. Software transactional memory. PODC'95.

TRANSACTIONAL MEMORY [1]



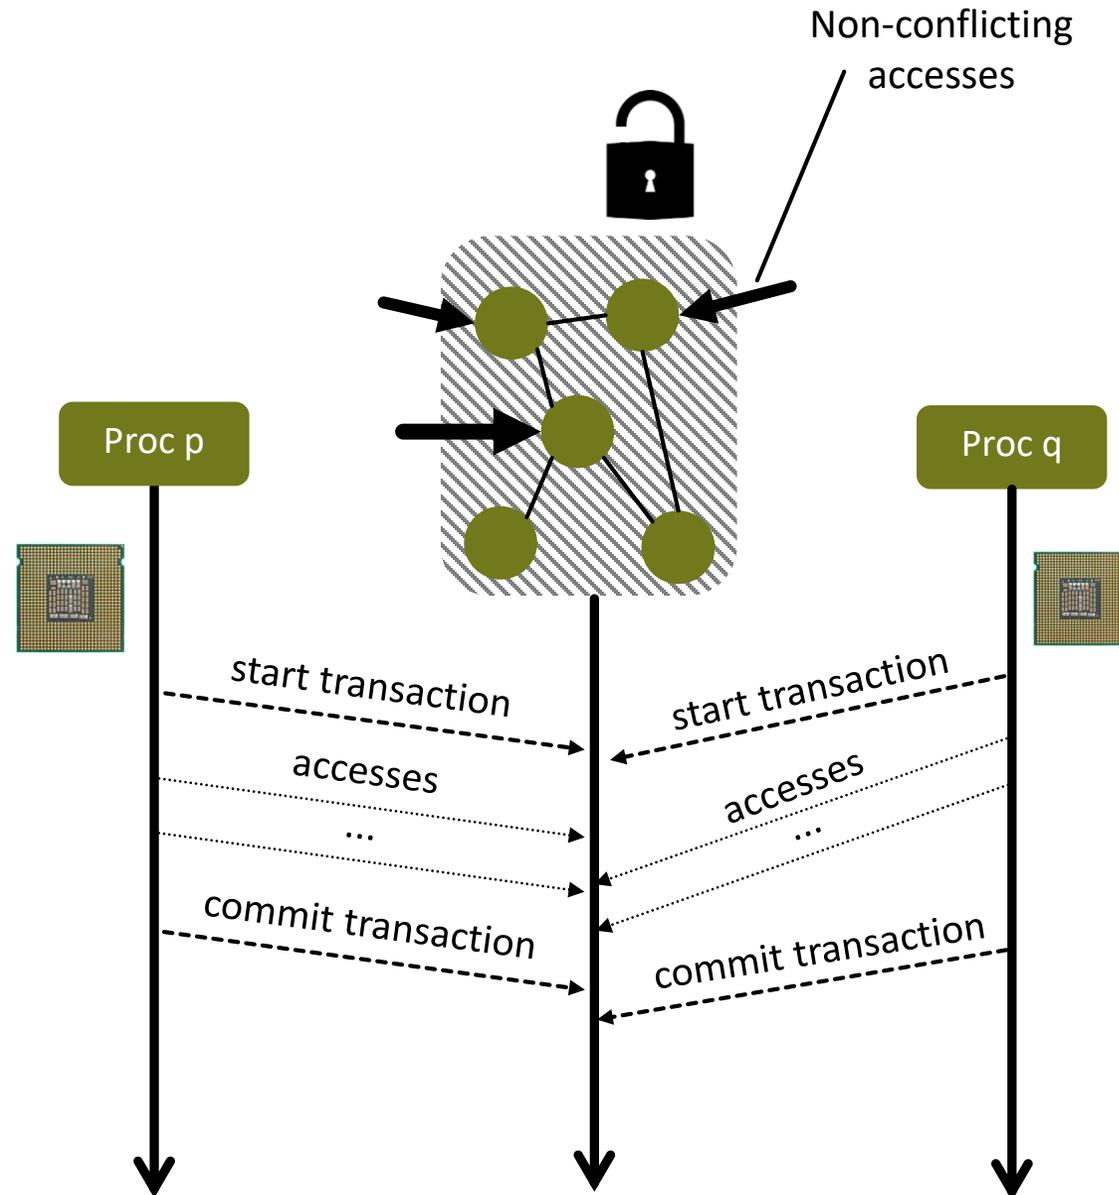
[1] N. Shavit and D. Touitou. Software transactional memory. PODC'95.

TRANSACTIONAL MEMORY [1]



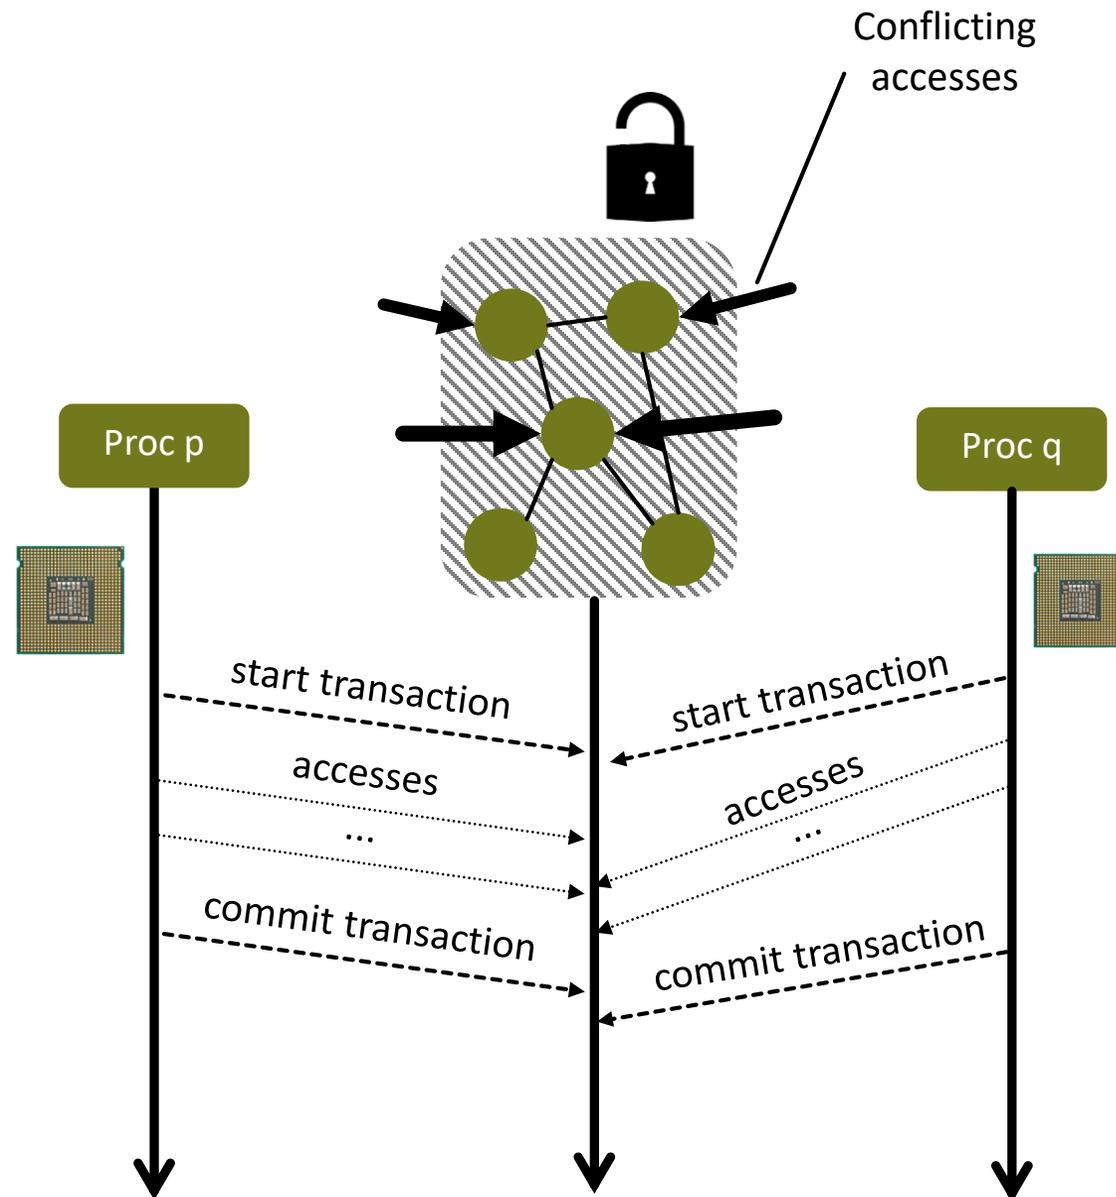
[1] N. Shavit and D. Touitou. Software transactional memory. PODC'95.

TRANSACTIONAL MEMORY [1]



[1] N. Shavit and D. Touitou. Software transactional memory. PODC'95.

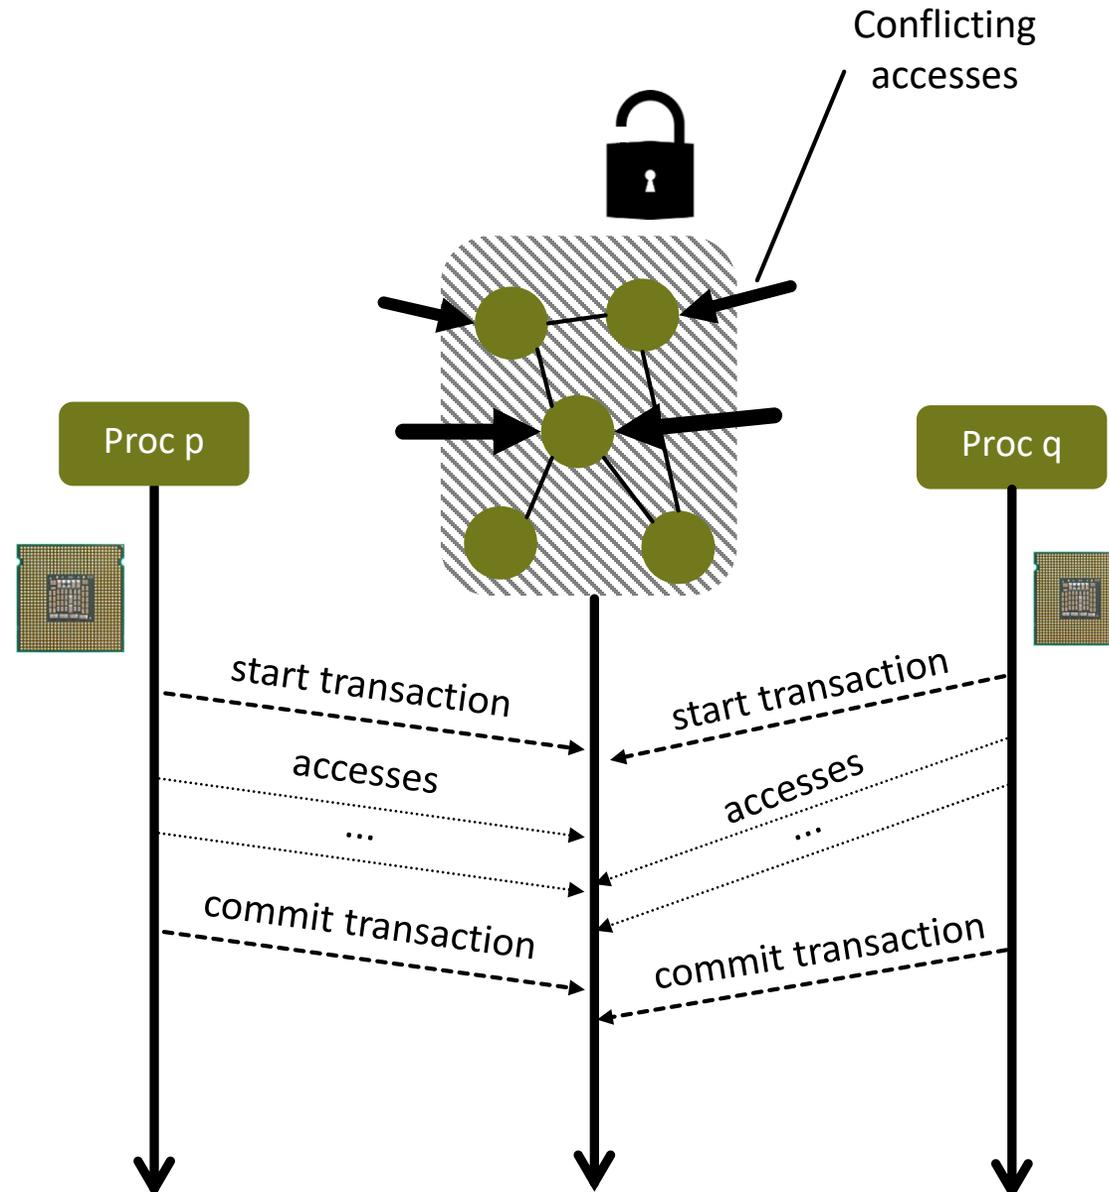
TRANSACTIONAL MEMORY [1]



[1] N. Shavit and D. Touitou. Software transactional memory. PODC'95.

TRANSACTIONAL MEMORY [1]

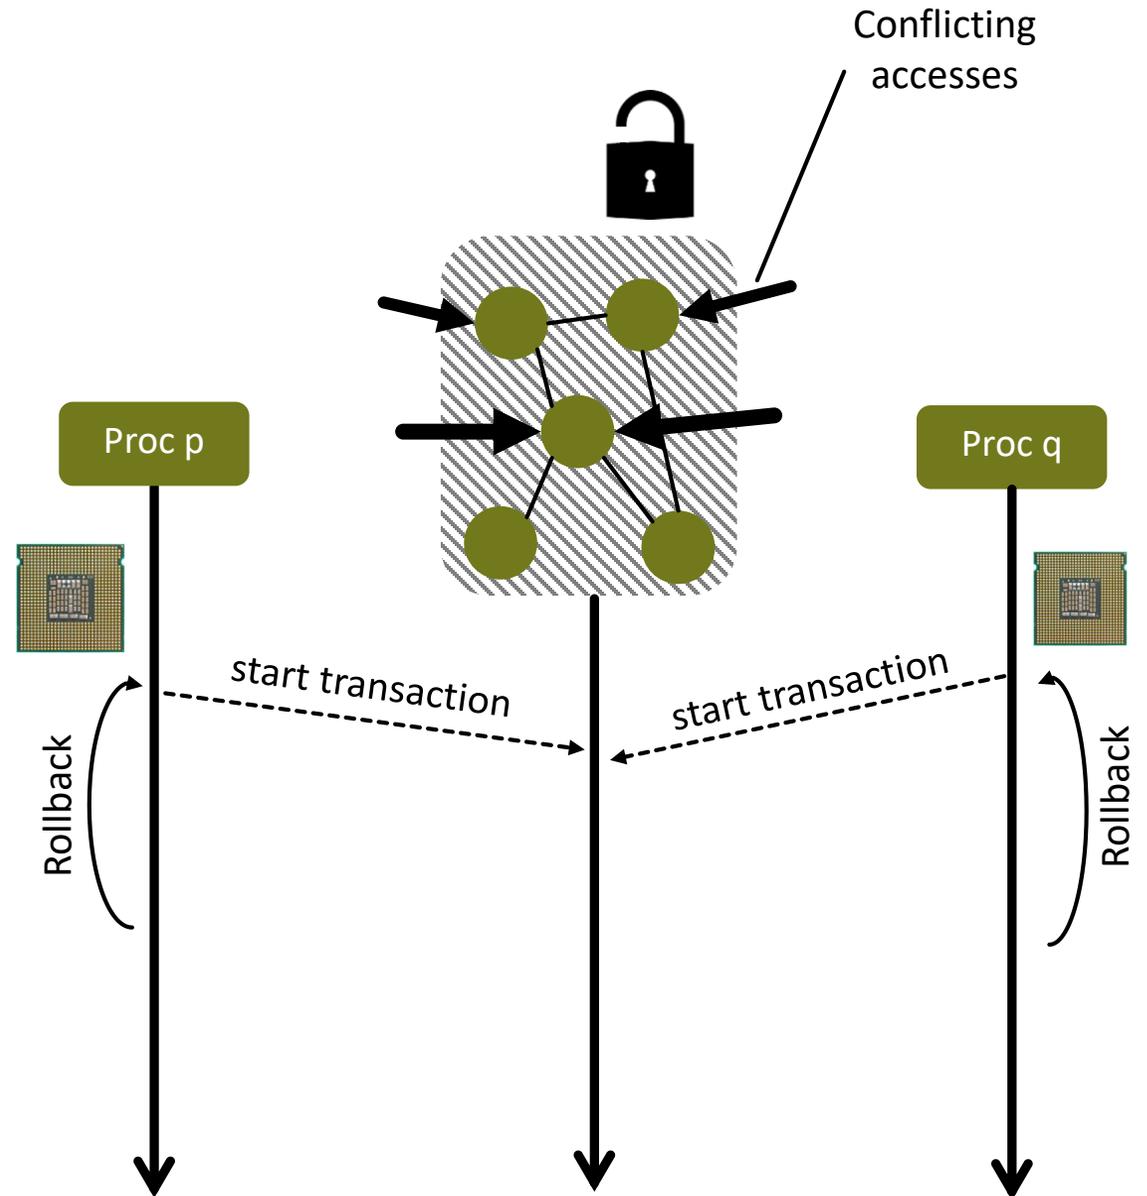
Conflicts solved with rollbacks and/or serialization.



[1] N. Shavit and D. Touitou. Software transactional memory. PODC'95.

TRANSACTIONAL MEMORY [1]

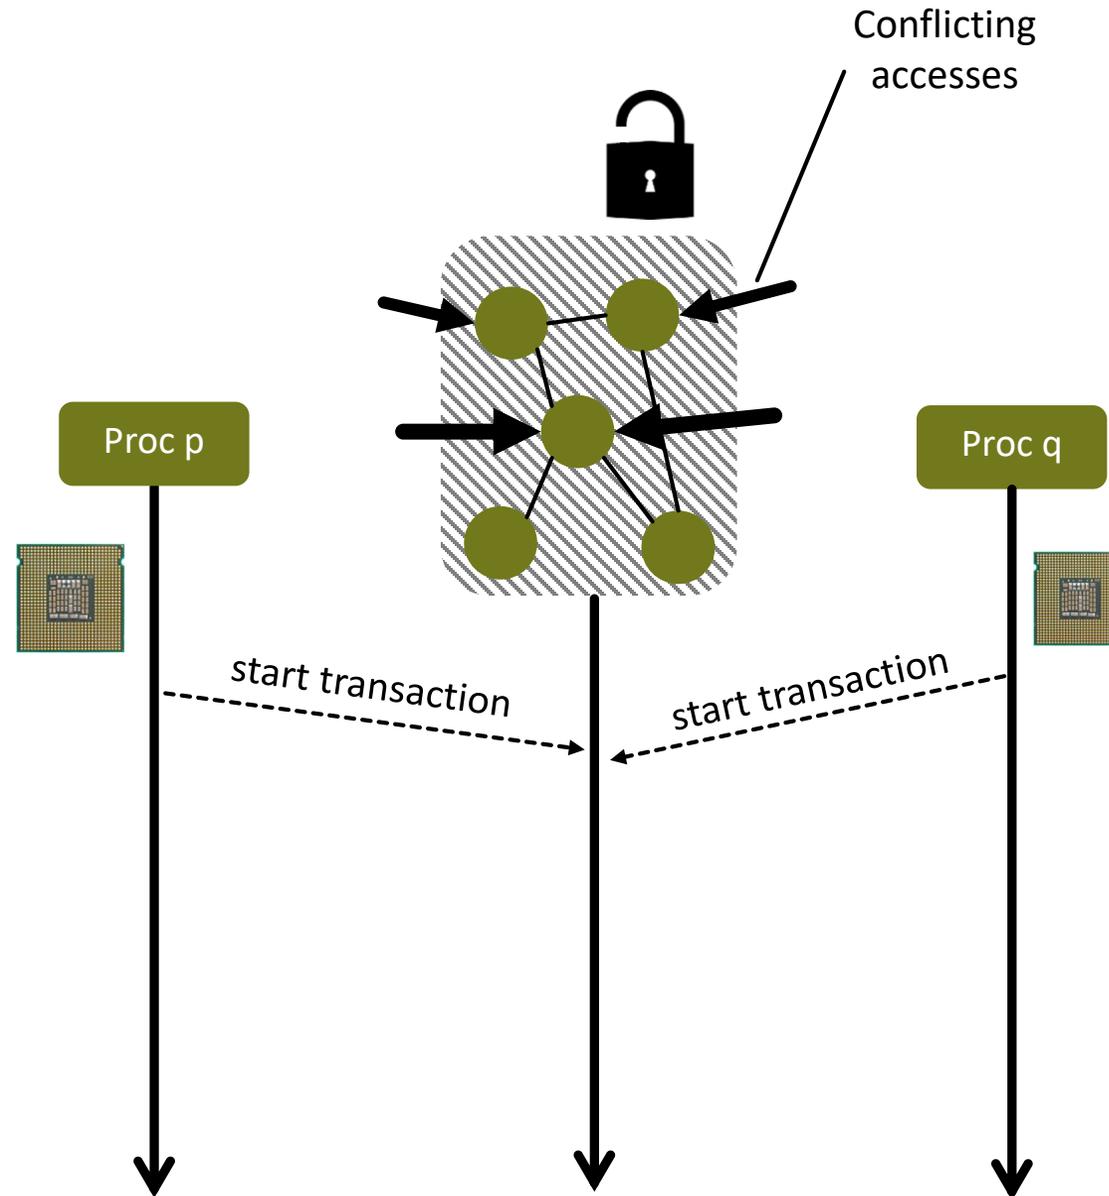
Conflicts solved with rollbacks and/or serialization.



[1] N. Shavit and D. Touitou. Software transactional memory. PODC'95.

TRANSACTIONAL MEMORY [1]

Conflicts solved with rollbacks and/or serialization.

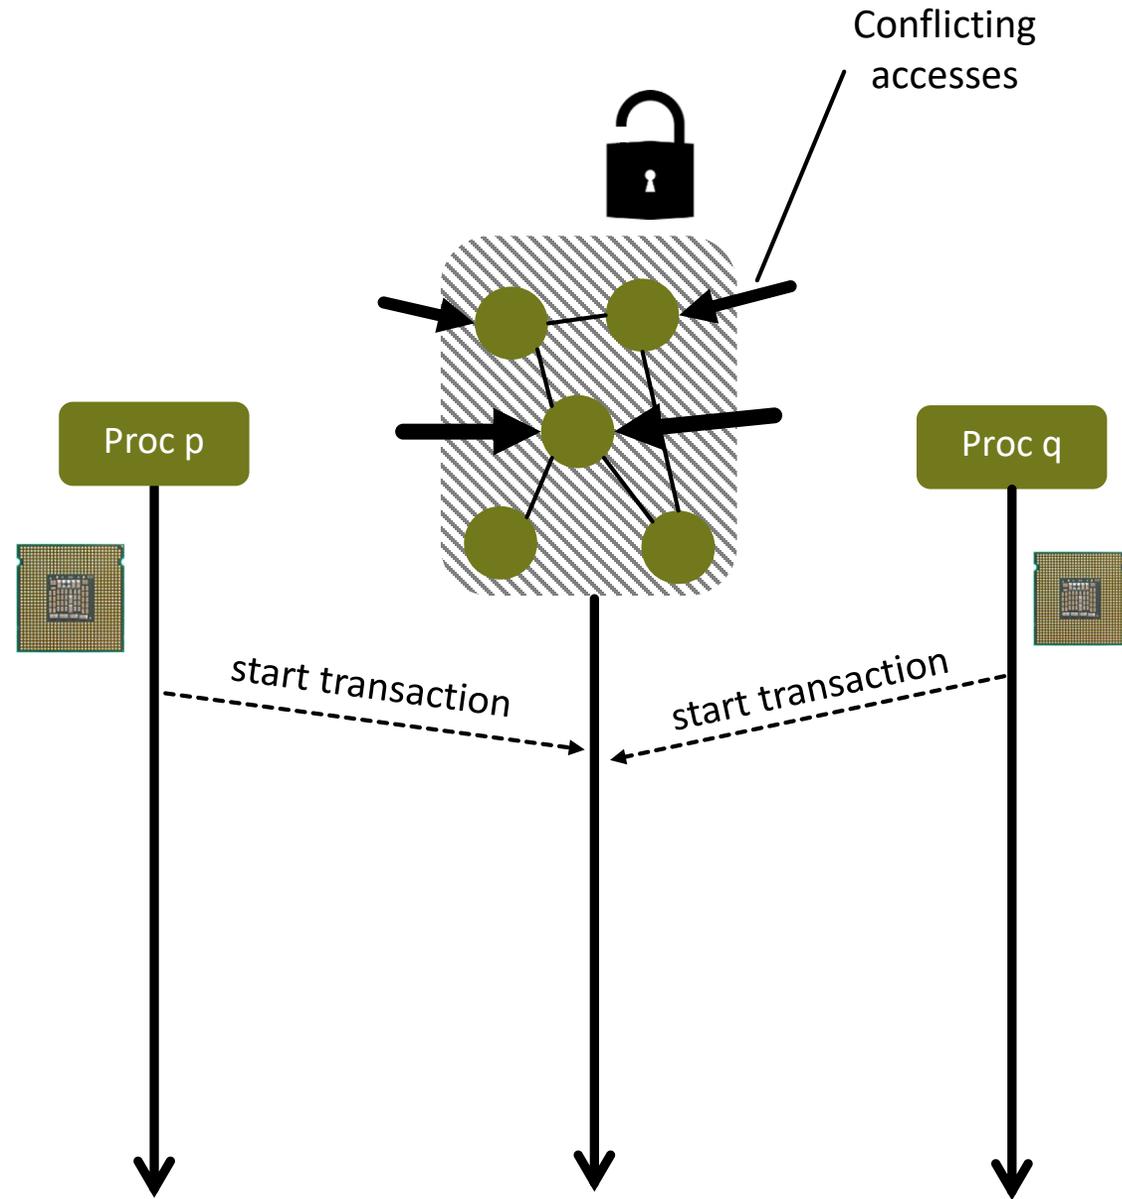


[1] N. Shavit and D. Touitou. Software transactional memory. PODC'95.

TRANSACTIONAL MEMORY [1]

Conflicts solved with rollbacks and/or serialization.

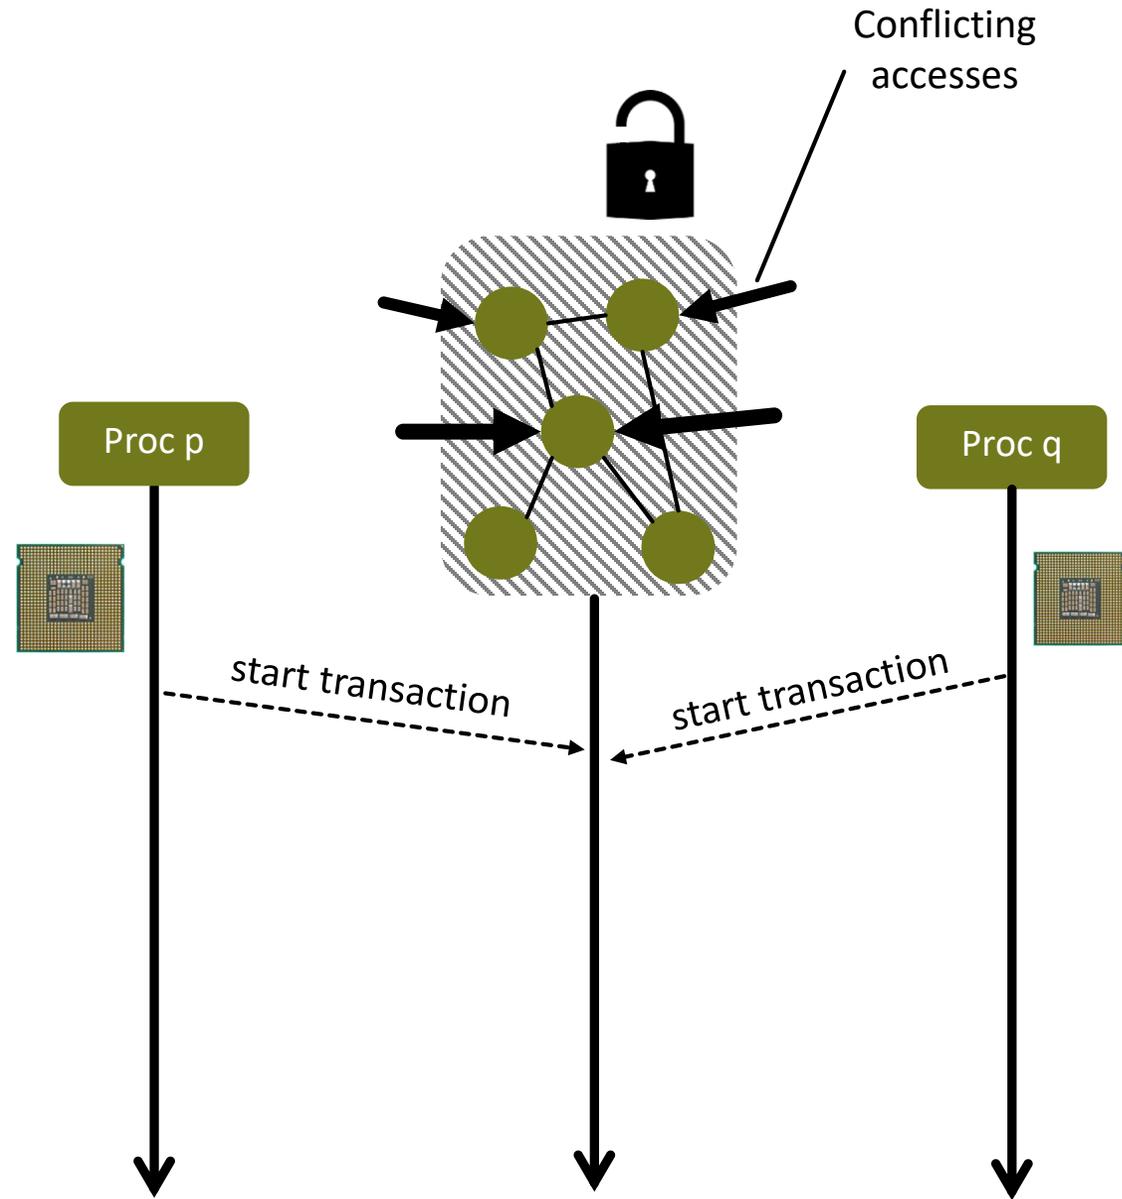
Simple protocols



[1] N. Shavit and D. Touitou. Software transactional memory. PODC'95.

TRANSACTIONAL MEMORY [1]

Conflicts solved with rollbacks and/or serialization.



Software overheads

Simple protocols

[1] N. Shavit and D. Touitou. Software transactional memory. PODC'95.

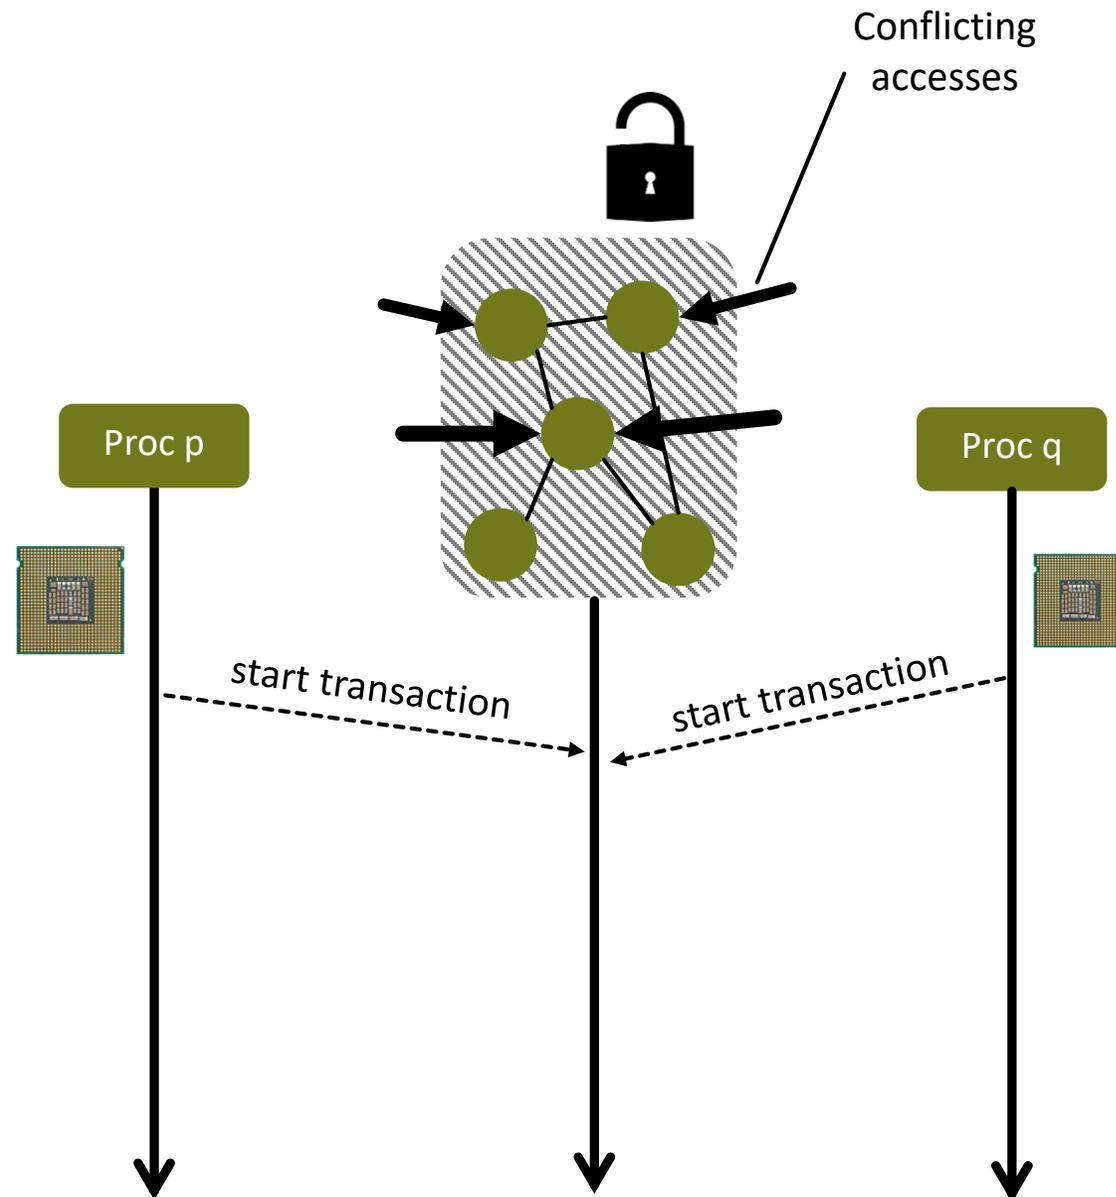
TRANSACTIONAL MEMORY [1]

Conflicts solved with rollbacks and/or serialization.

Software overheads

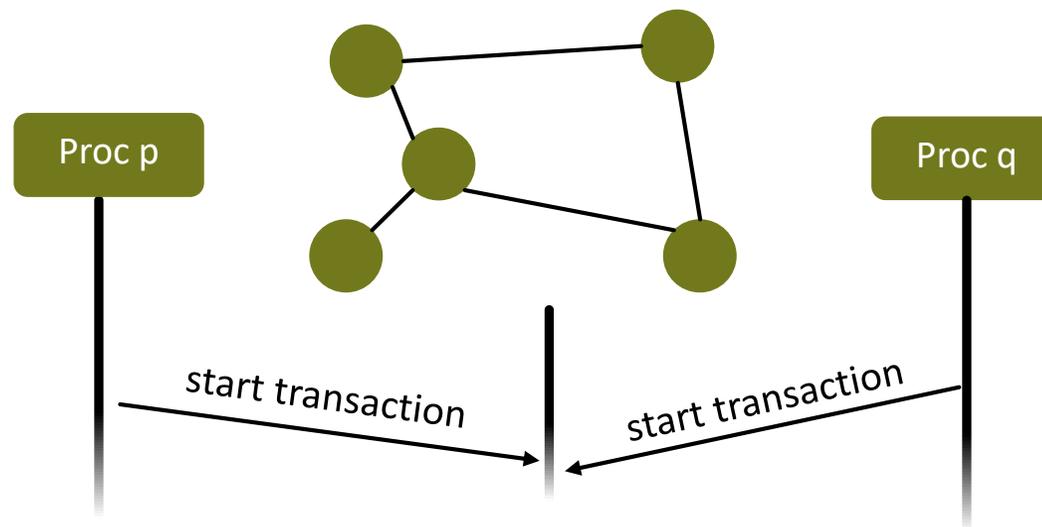
High performance?
For graphs?

Simple protocols



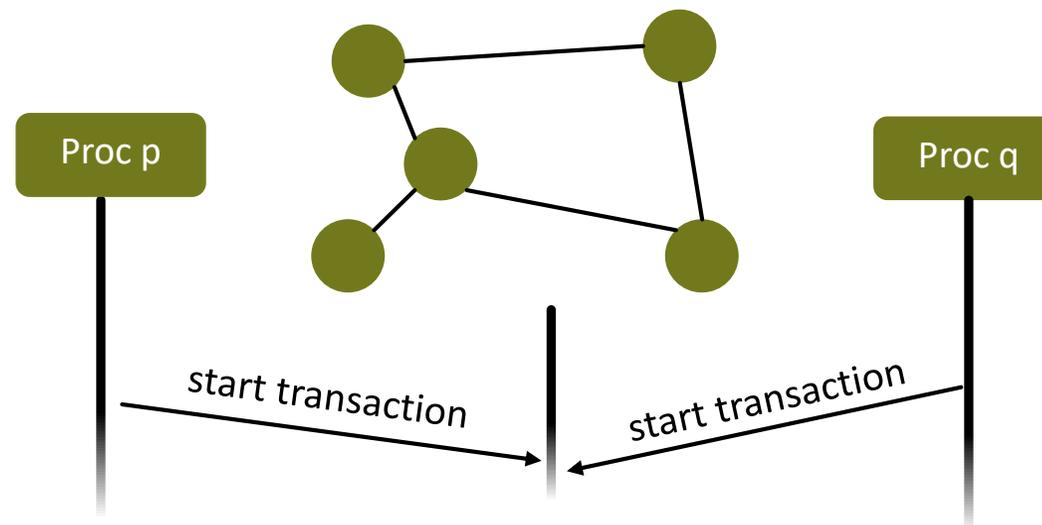
[1] N. Shavit and D. Touitou. Software transactional memory. PODC'95.

SHARED- & DISTRIBUTED-MEMORY MACHINES



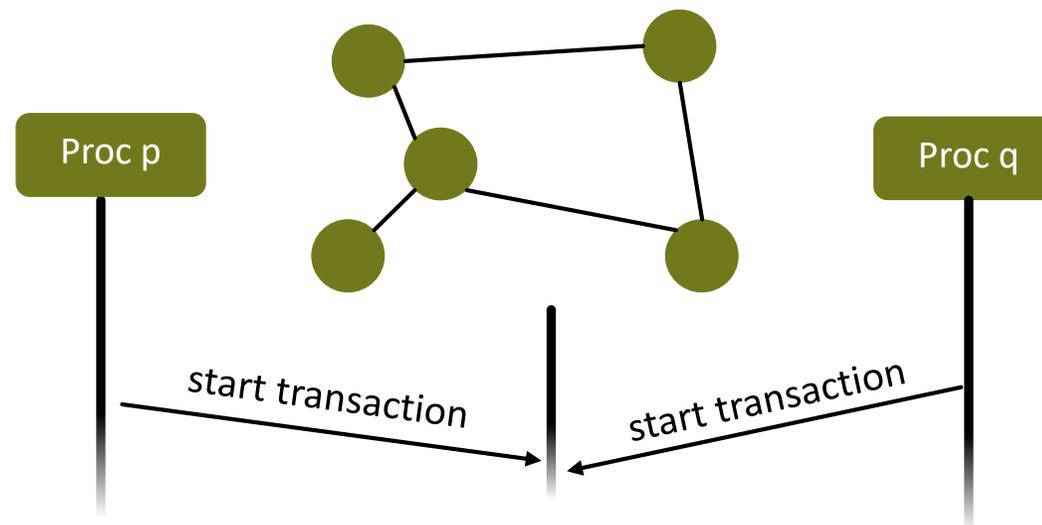
SHARED- & DISTRIBUTED-MEMORY MACHINES

- HTM works fine for single shared-memory domains



SHARED- & DISTRIBUTED-MEMORY MACHINES

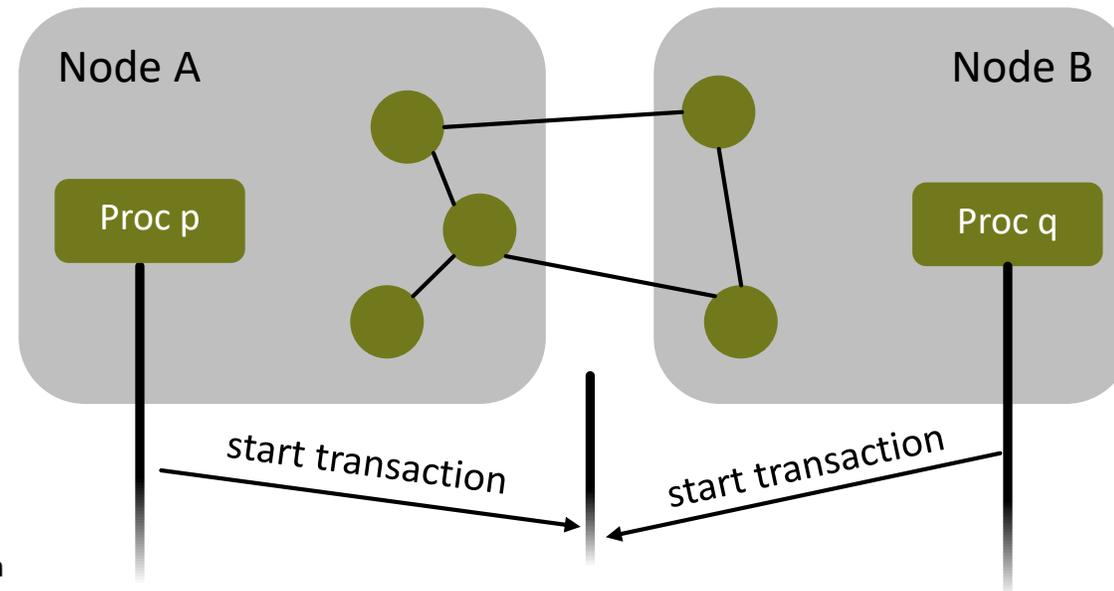
- HTM works fine for single shared-memory domains
 - Most graphs fit in such machines [1]



[1] Y. Perez et al. Ringo: Interactive Graph Analytics on Big-Memory Machines. SIGCOMM'14.

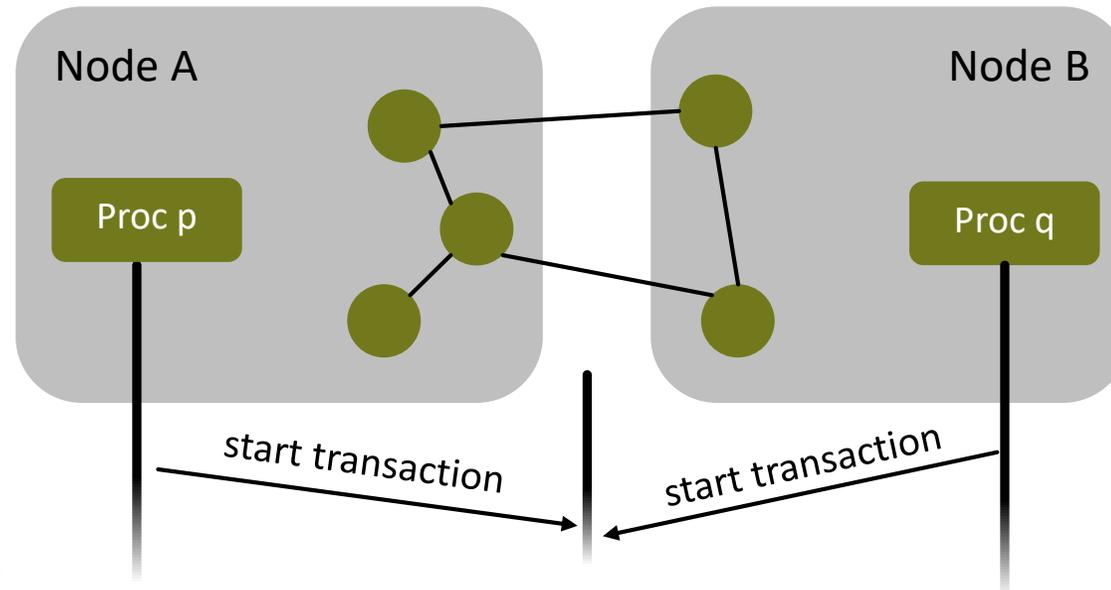
SHARED- & DISTRIBUTED-MEMORY MACHINES

- HTM works fine for single shared-memory domains
 - Most graphs fit in such machines [1]



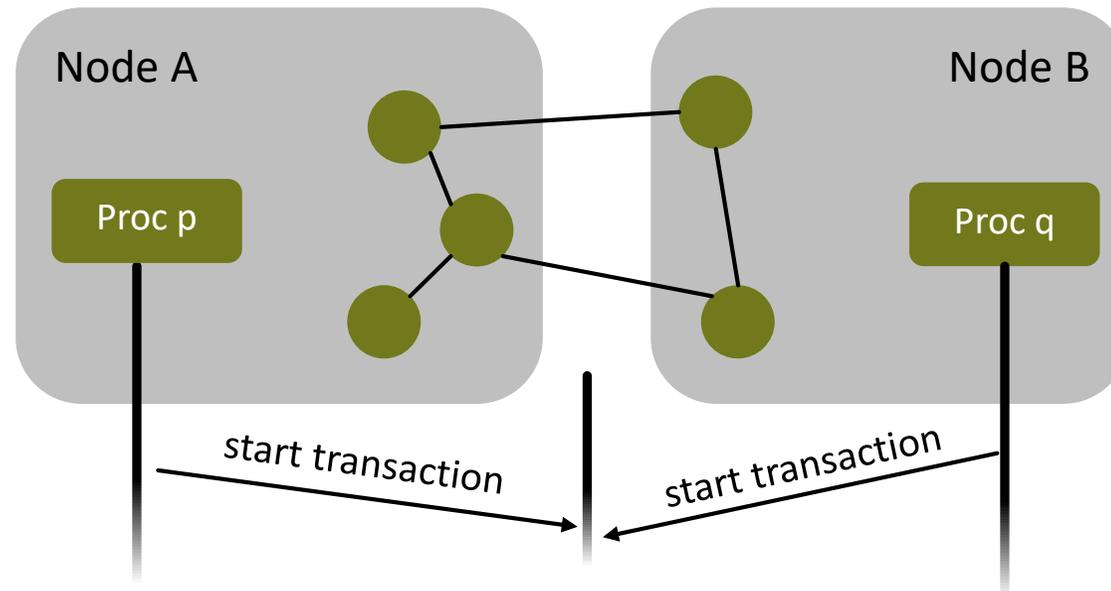
SHARED- & DISTRIBUTED-MEMORY MACHINES

- HTM works fine for single shared-memory domains
 - Most graphs fit in such machines [1]
- However, some do not:



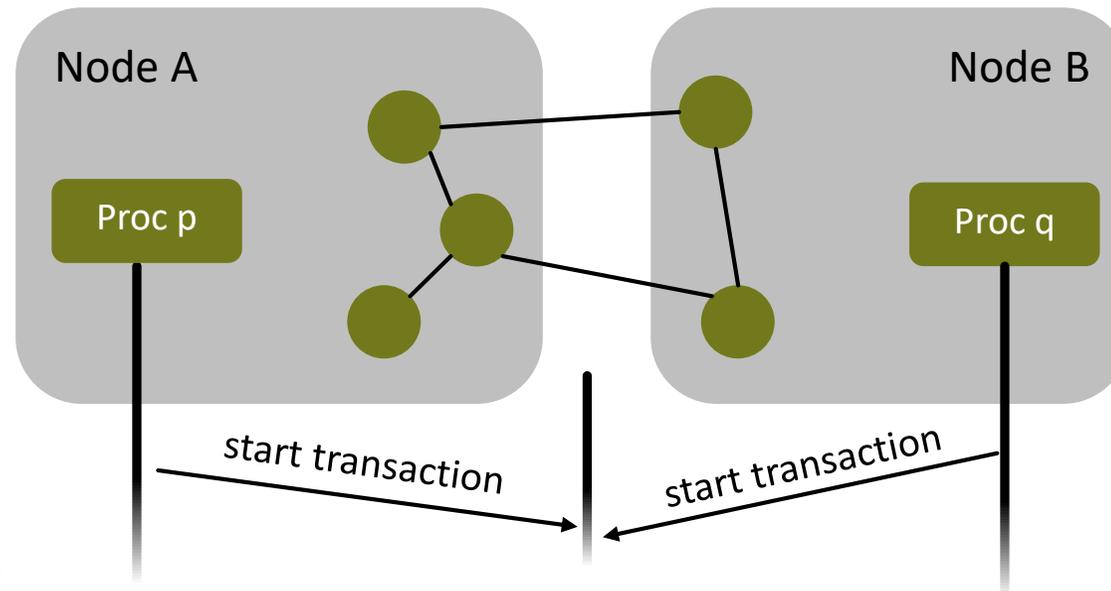
SHARED- & DISTRIBUTED-MEMORY MACHINES

- HTM works fine for single shared-memory domains
 - Most graphs fit in such machines [1]
- However, some do not:
 - Very large instances



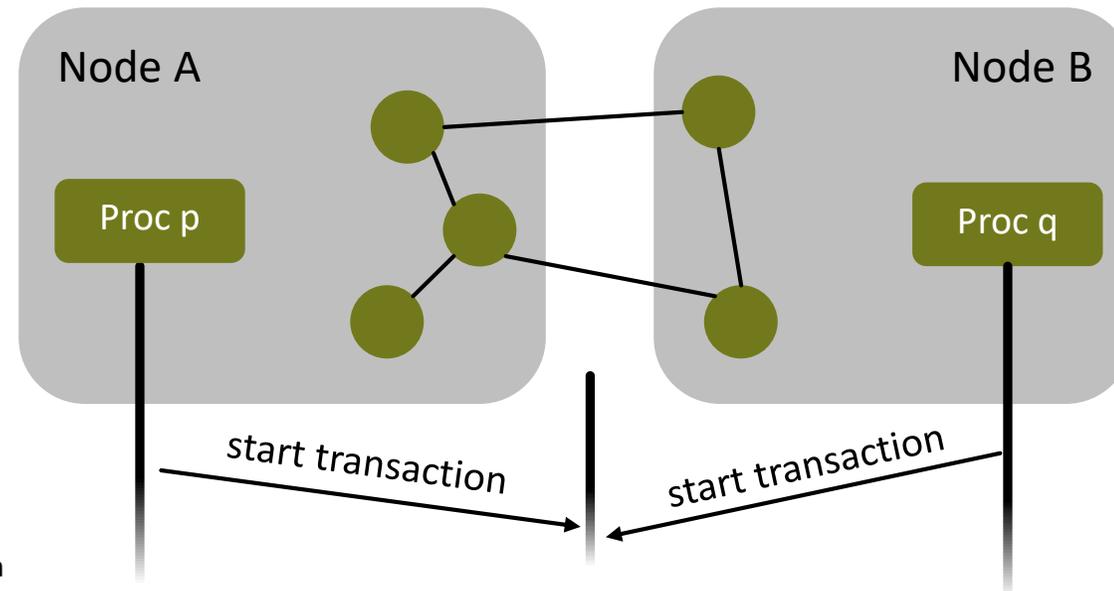
SHARED- & DISTRIBUTED-MEMORY MACHINES

- HTM works fine for single shared-memory domains
 - Most graphs fit in such machines [1]
- However, some do not:
 - Very large instances
 - Rich vertex/edge data



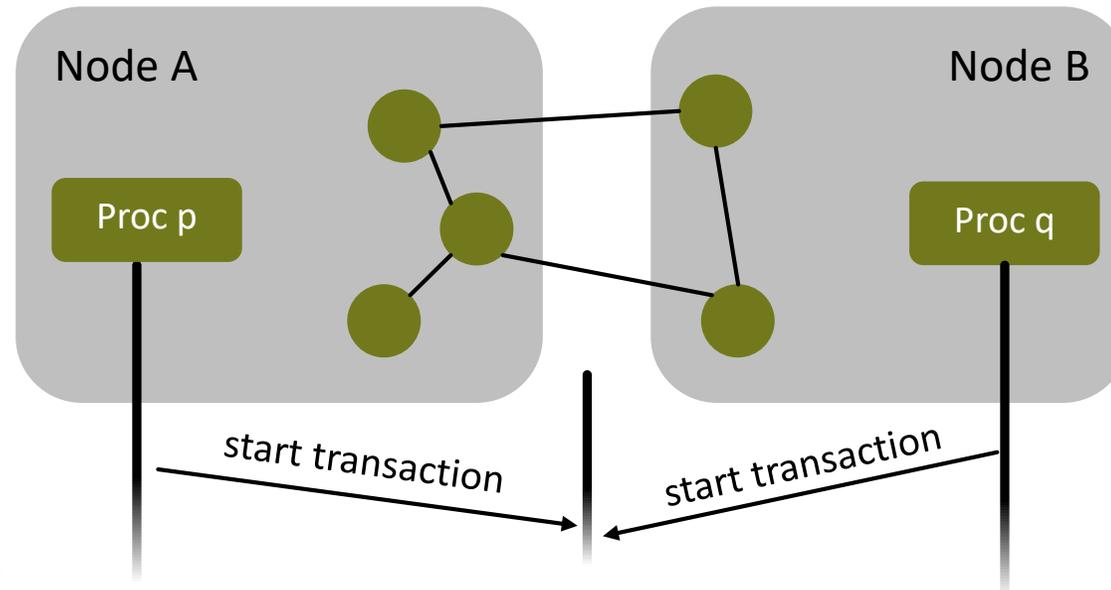
SHARED- & DISTRIBUTED-MEMORY MACHINES

- HTM works fine for single shared-memory domains
 - Most graphs fit in such machines [1]
- However, some do not:
 - Very large instances
 - Rich vertex/edge data
- Fat nodes with lots of RAM still expensive (\$35K for a machine with 1TB of RAM [1])

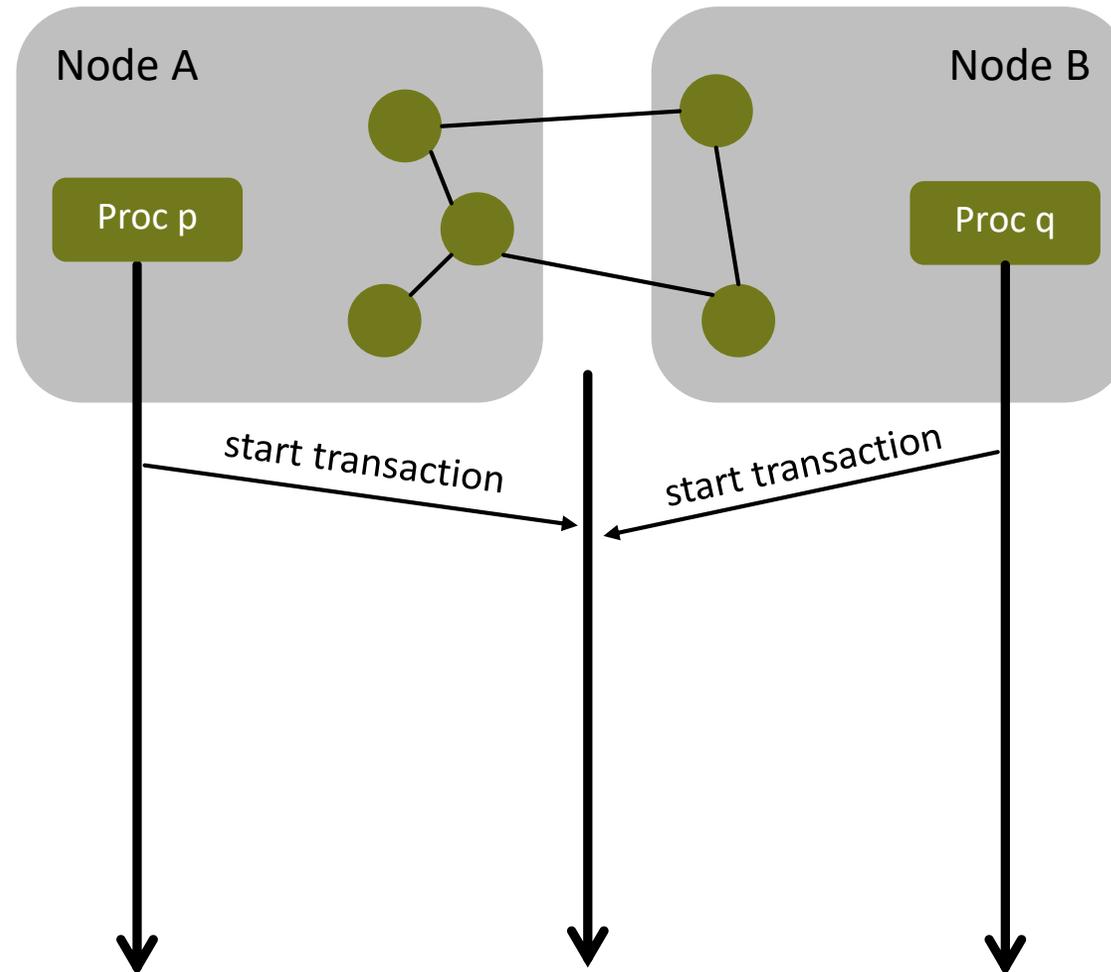


SHARED- & DISTRIBUTED-MEMORY MACHINES

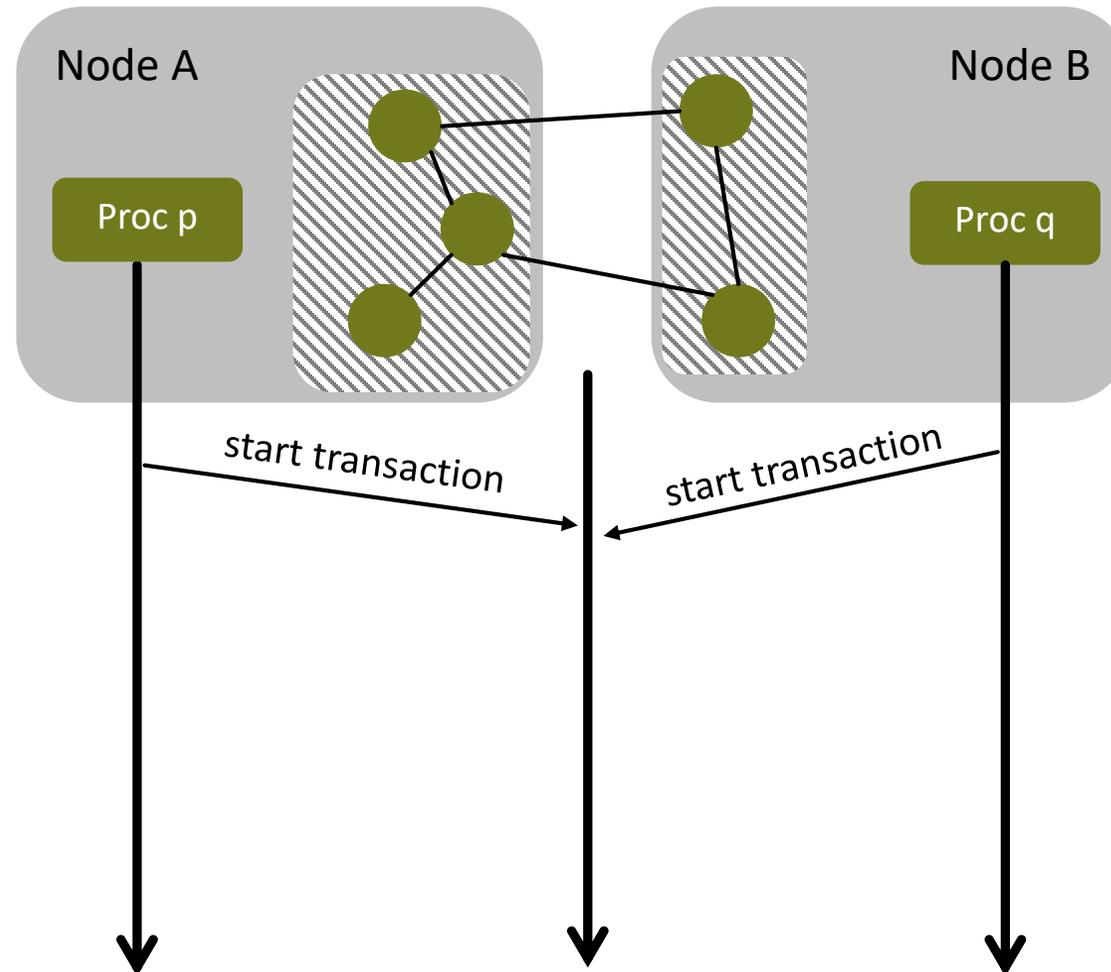
- HTM works fine for single shared-memory domains
 - Most graphs fit in such machines [1]
- However, some do not:
 - Very large instances
 - Rich vertex/edge data
- Fat nodes with lots of RAM still expensive (\$35K for a machine with 1TB of RAM [1])



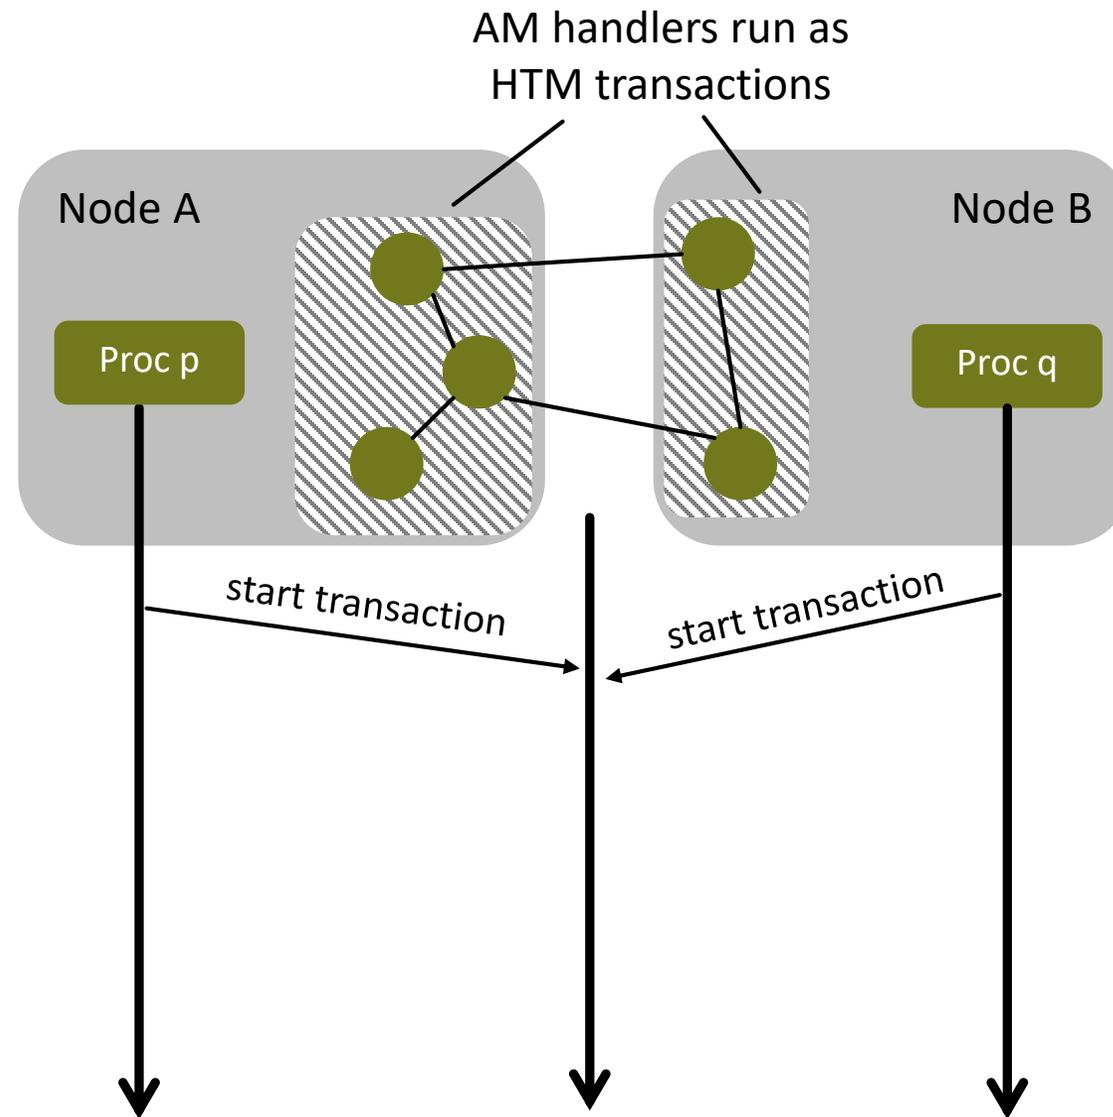
AM + HTM = ...



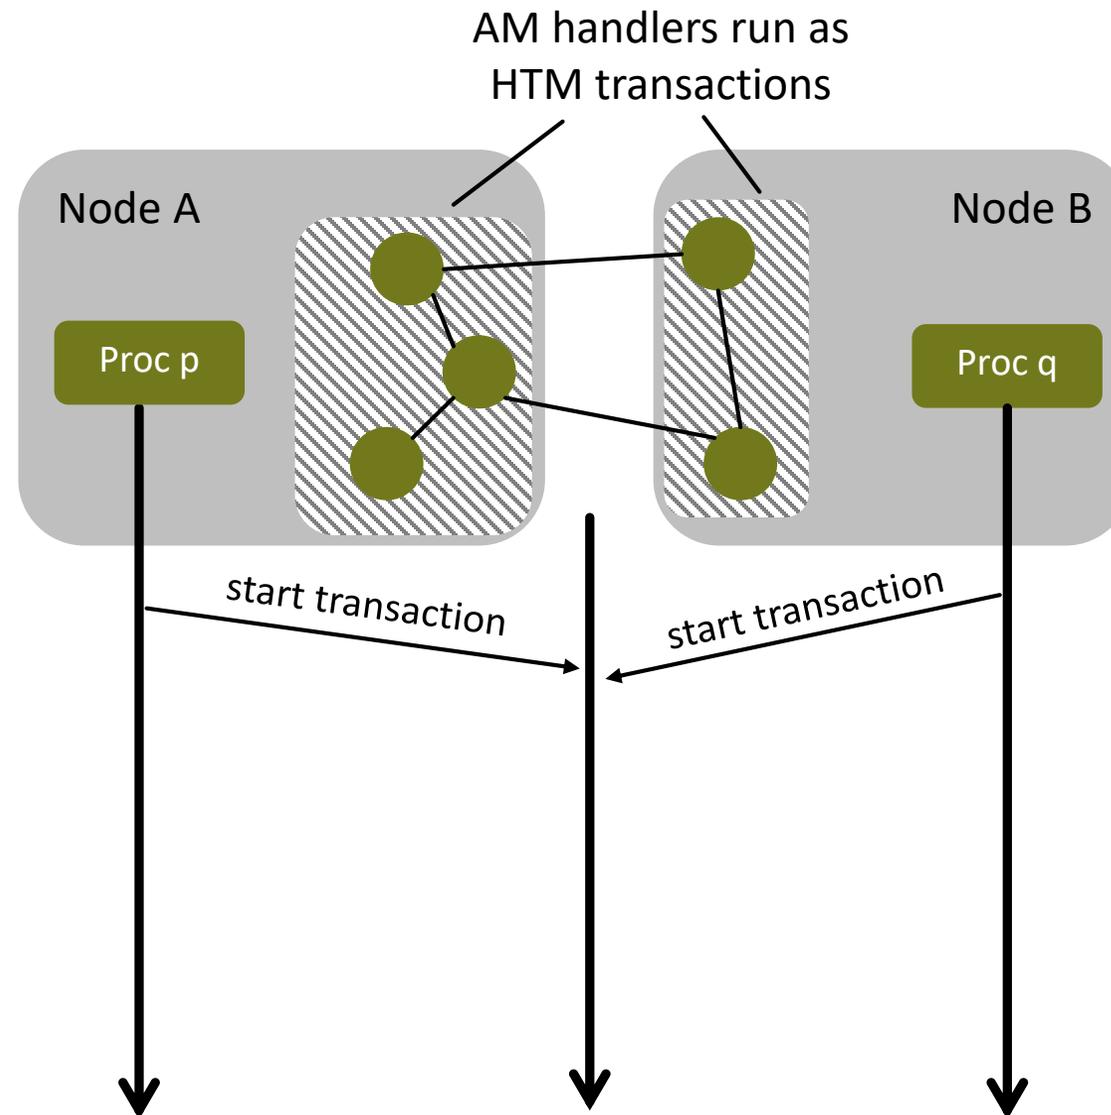
AM + HTM = ...



AM + HTM = ...

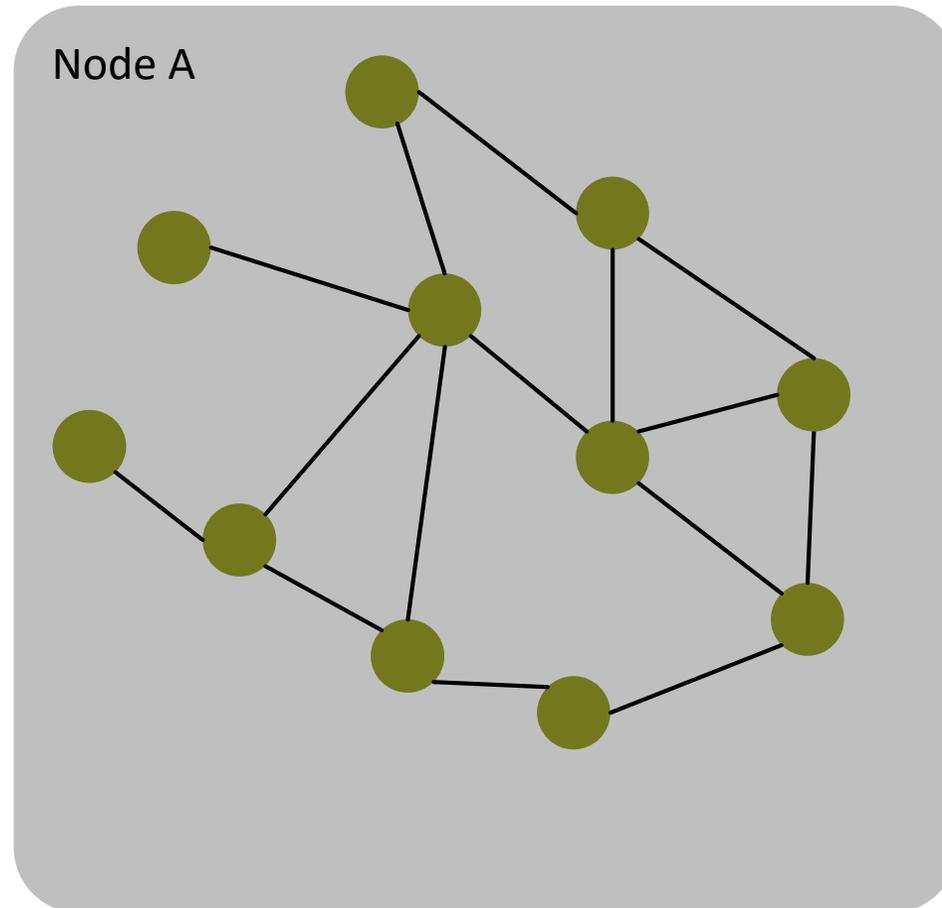


AM + HTM = ATOMIC ACTIVE MESSAGES



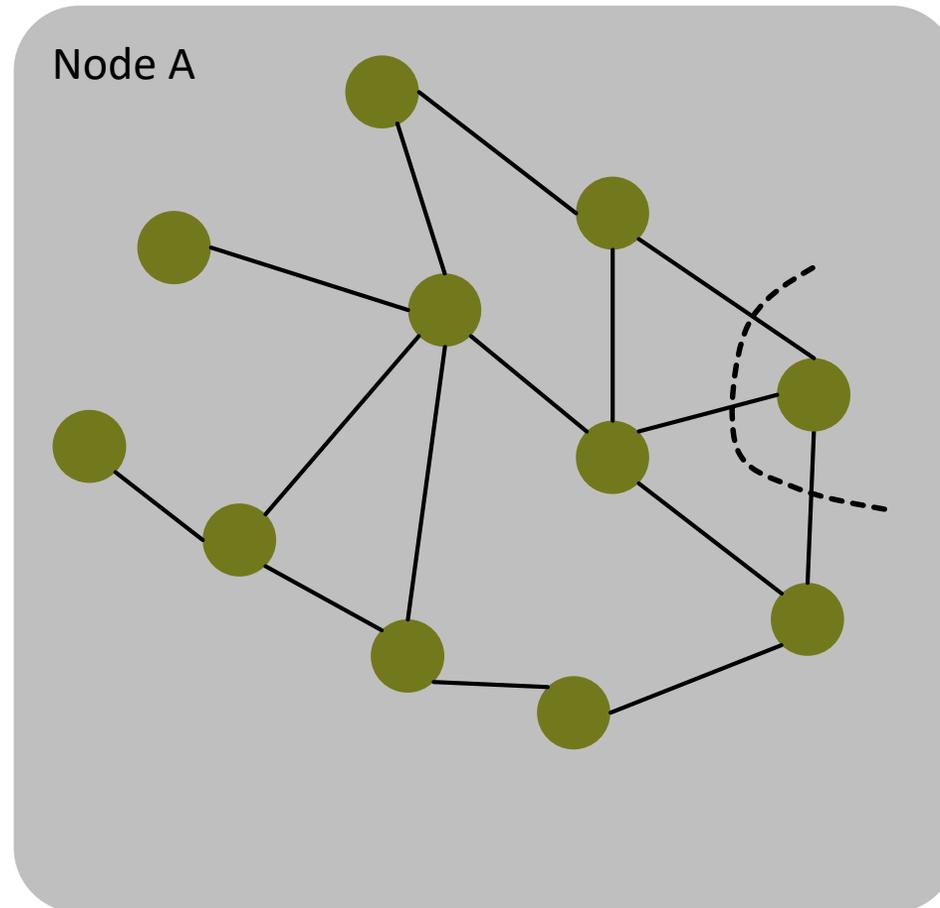
ACCESSING MULTIPLE VERTICES ATOMICALLY

Example: BFS



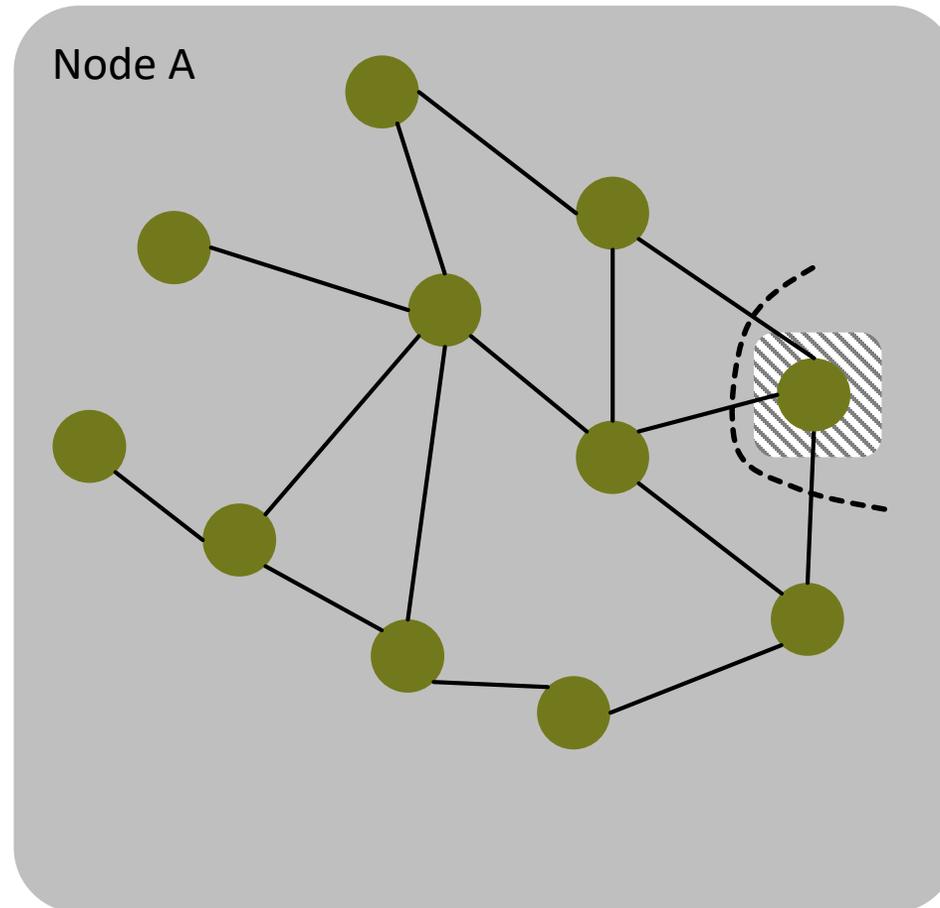
ACCESSING MULTIPLE VERTICES ATOMICALLY

Example: BFS



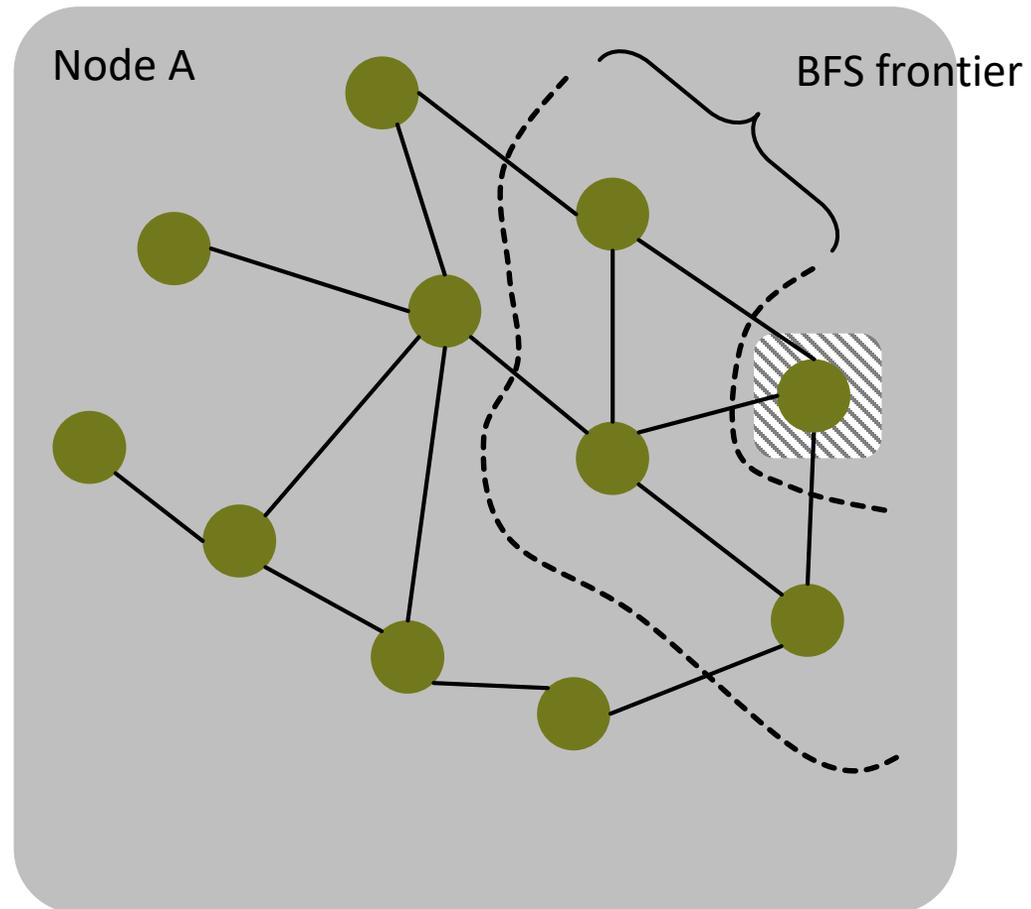
ACCESSING MULTIPLE VERTICES ATOMICALLY

Example: BFS



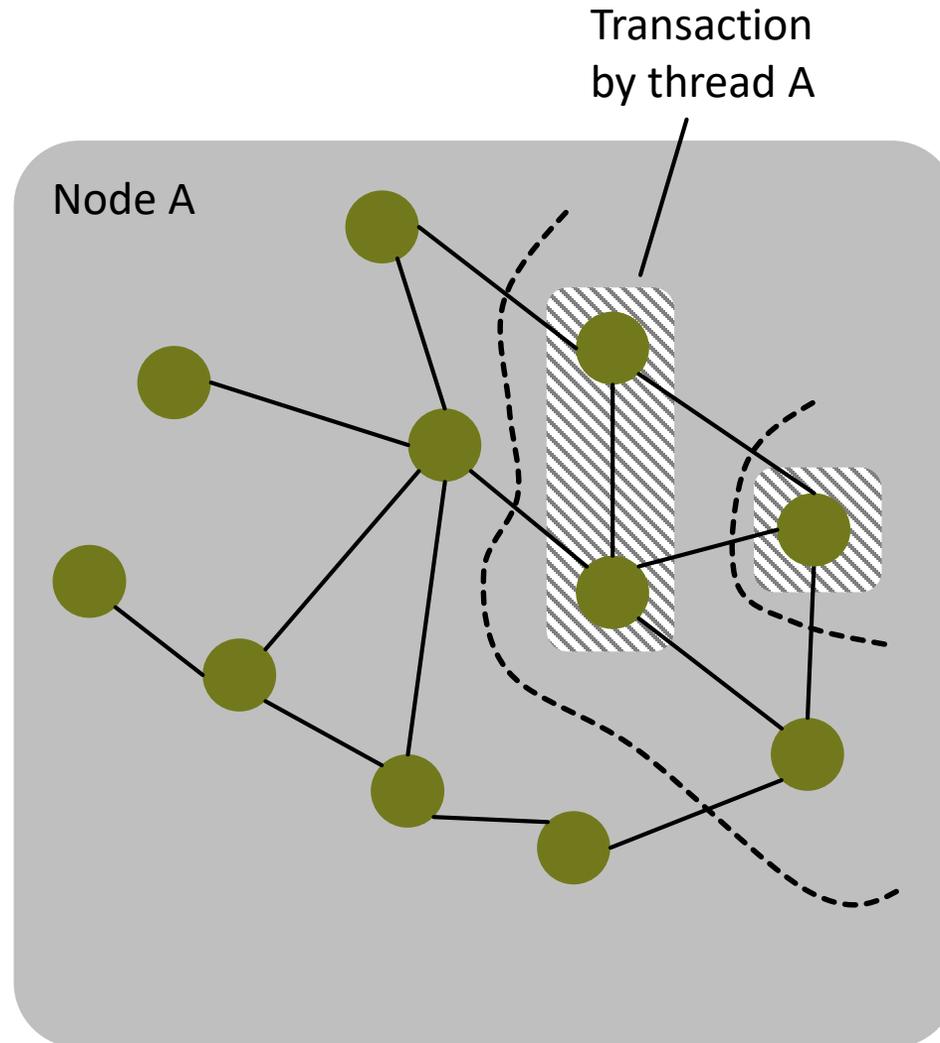
ACCESSING MULTIPLE VERTICES ATOMICALLY

Example: BFS



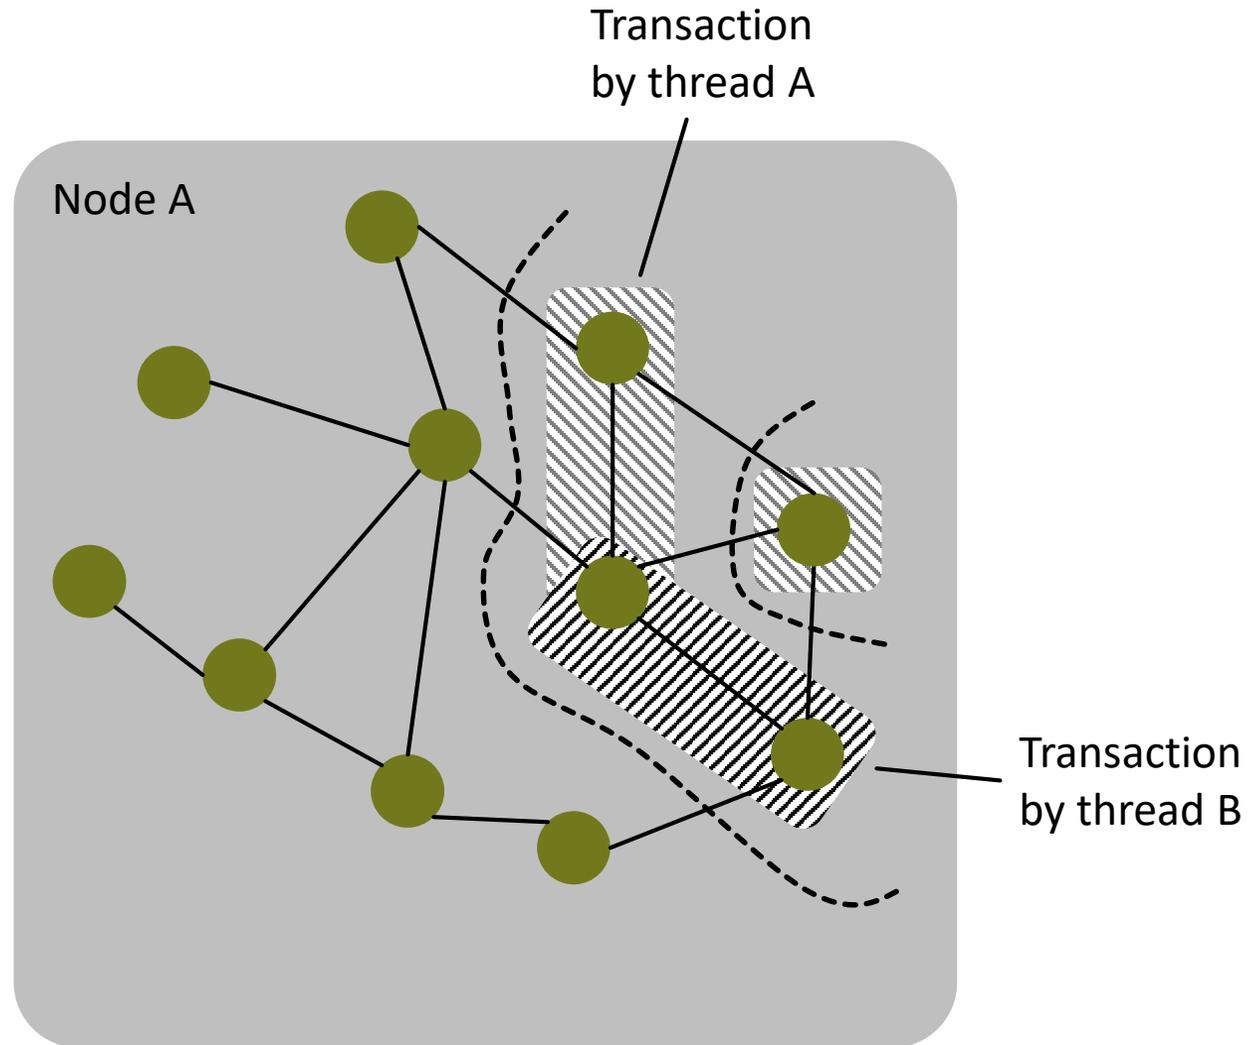
ACCESSING MULTIPLE VERTICES ATOMICALLY

Example: BFS



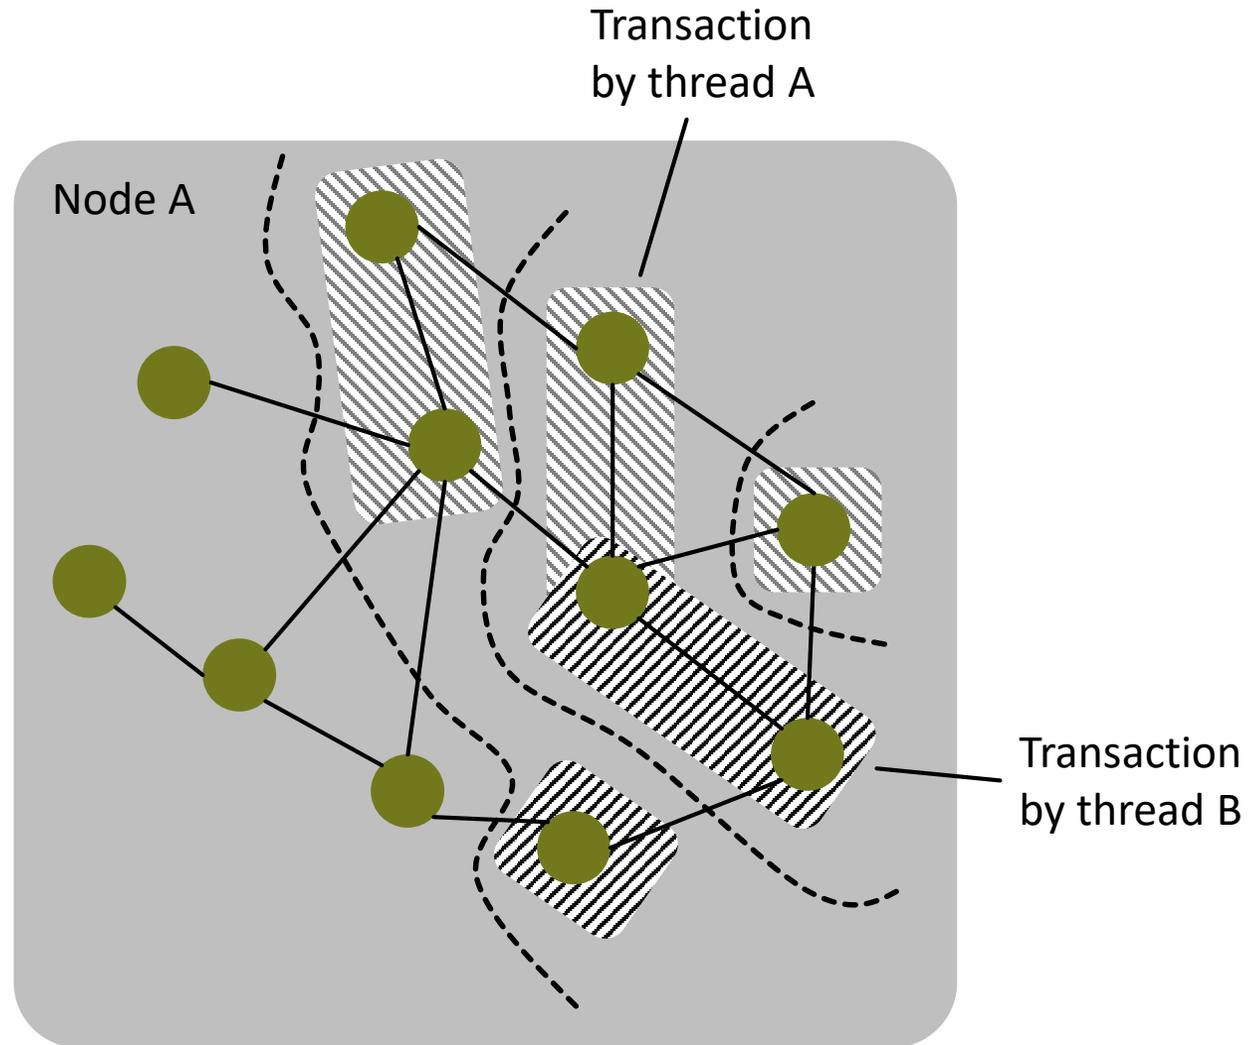
ACCESSING MULTIPLE VERTICES ATOMICALLY

Example: BFS



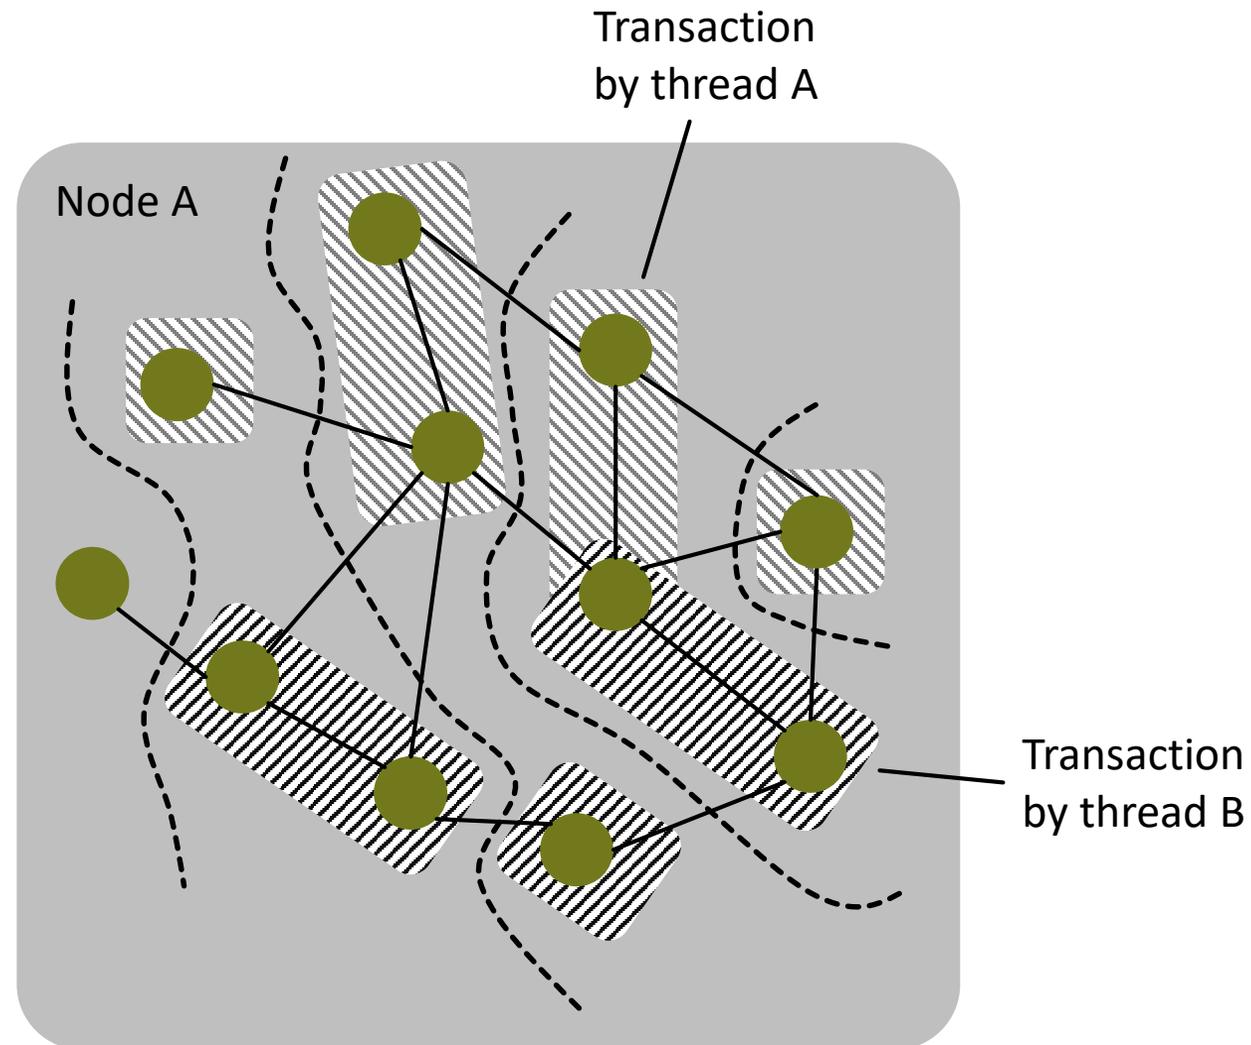
ACCESSING MULTIPLE VERTICES ATOMICALLY

Example: BFS



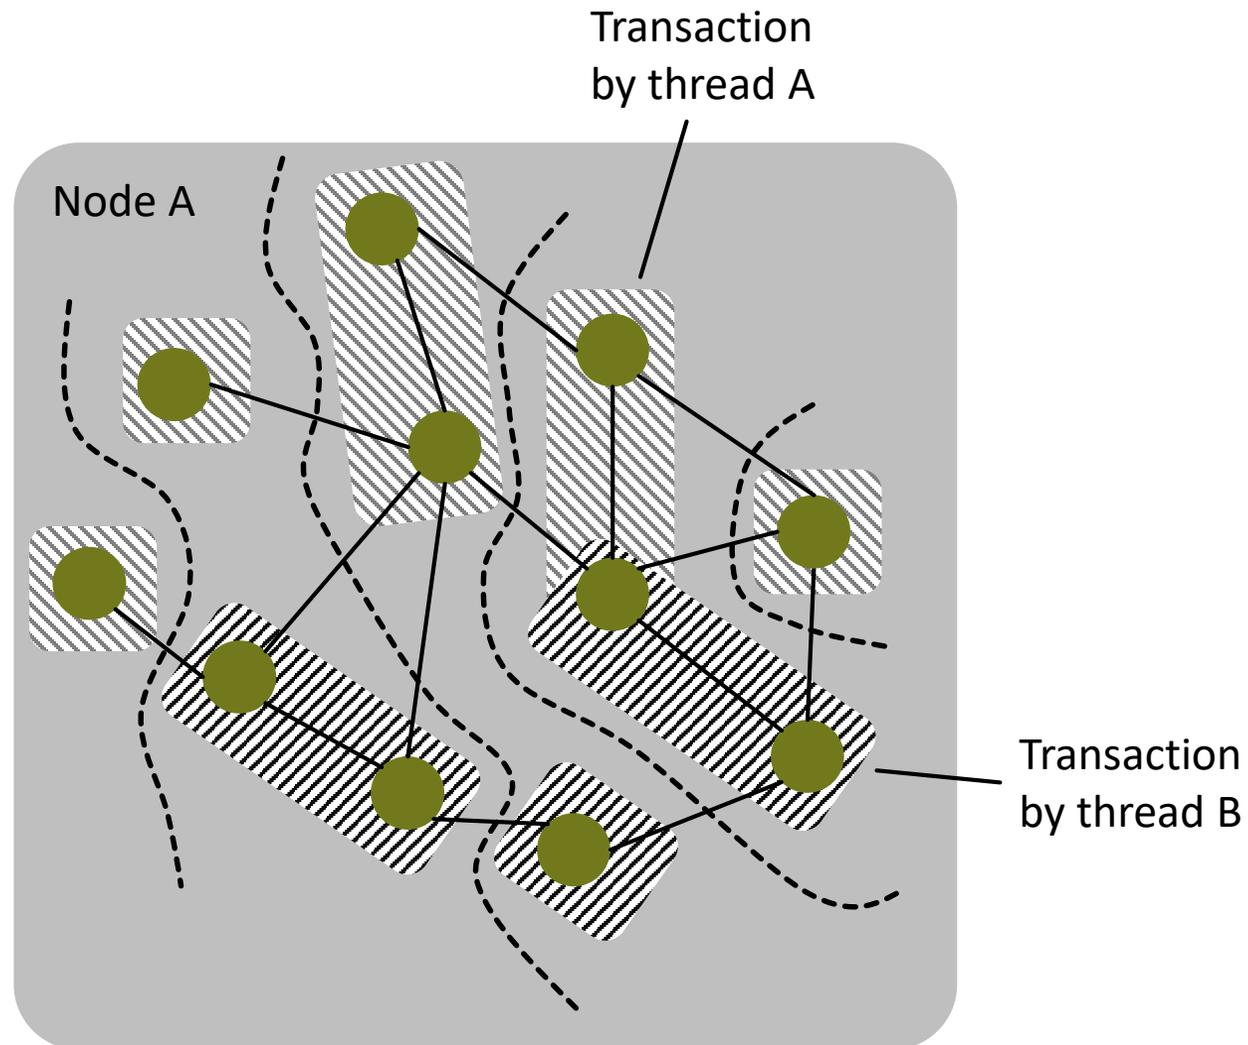
ACCESSING MULTIPLE VERTICES ATOMICALLY

Example: BFS



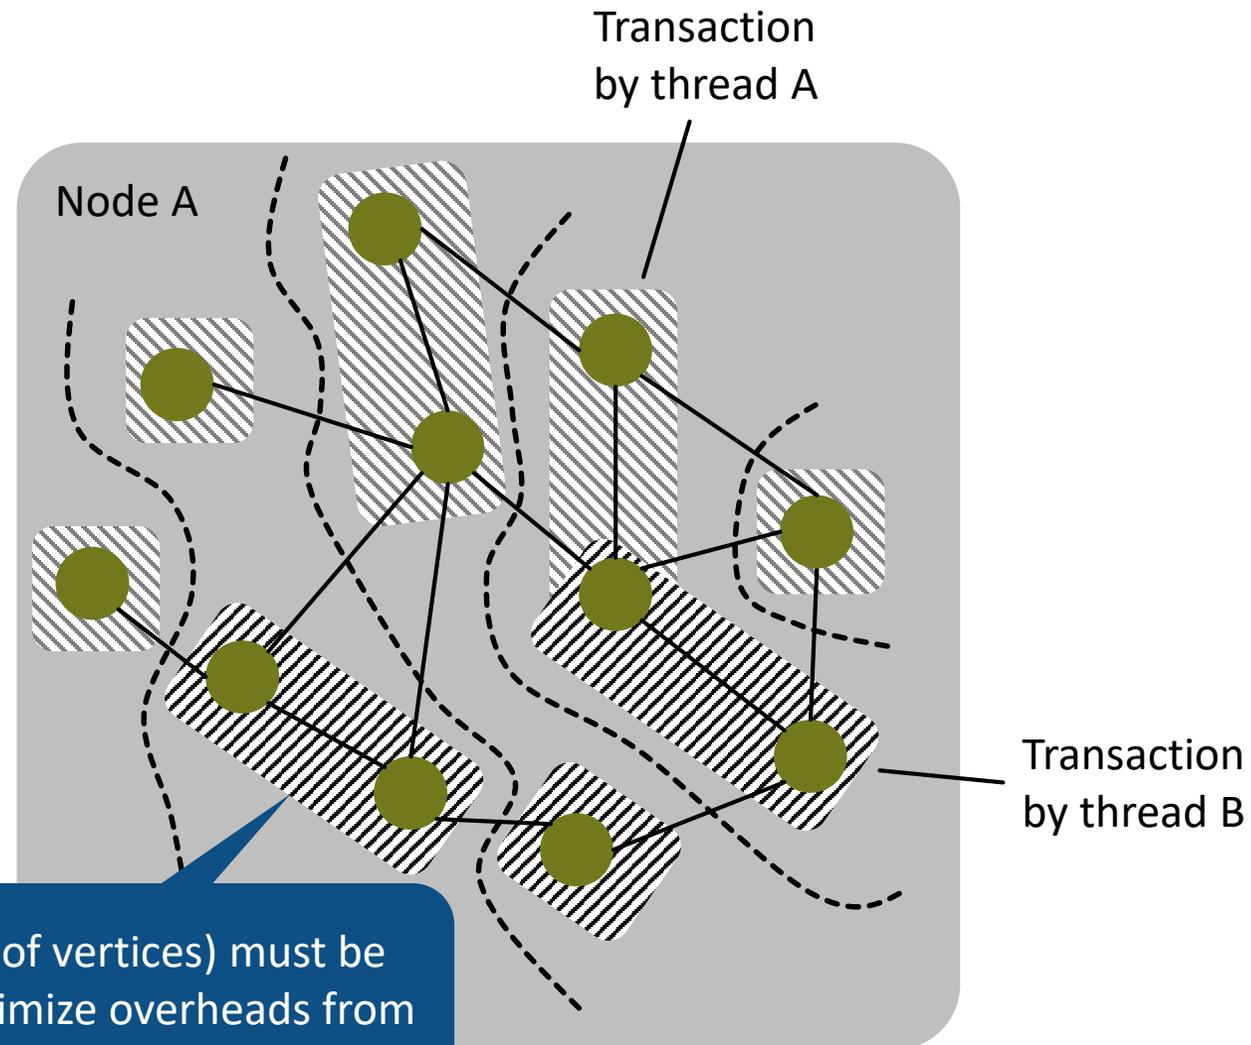
ACCESSING MULTIPLE VERTICES ATOMICALLY

Example: BFS



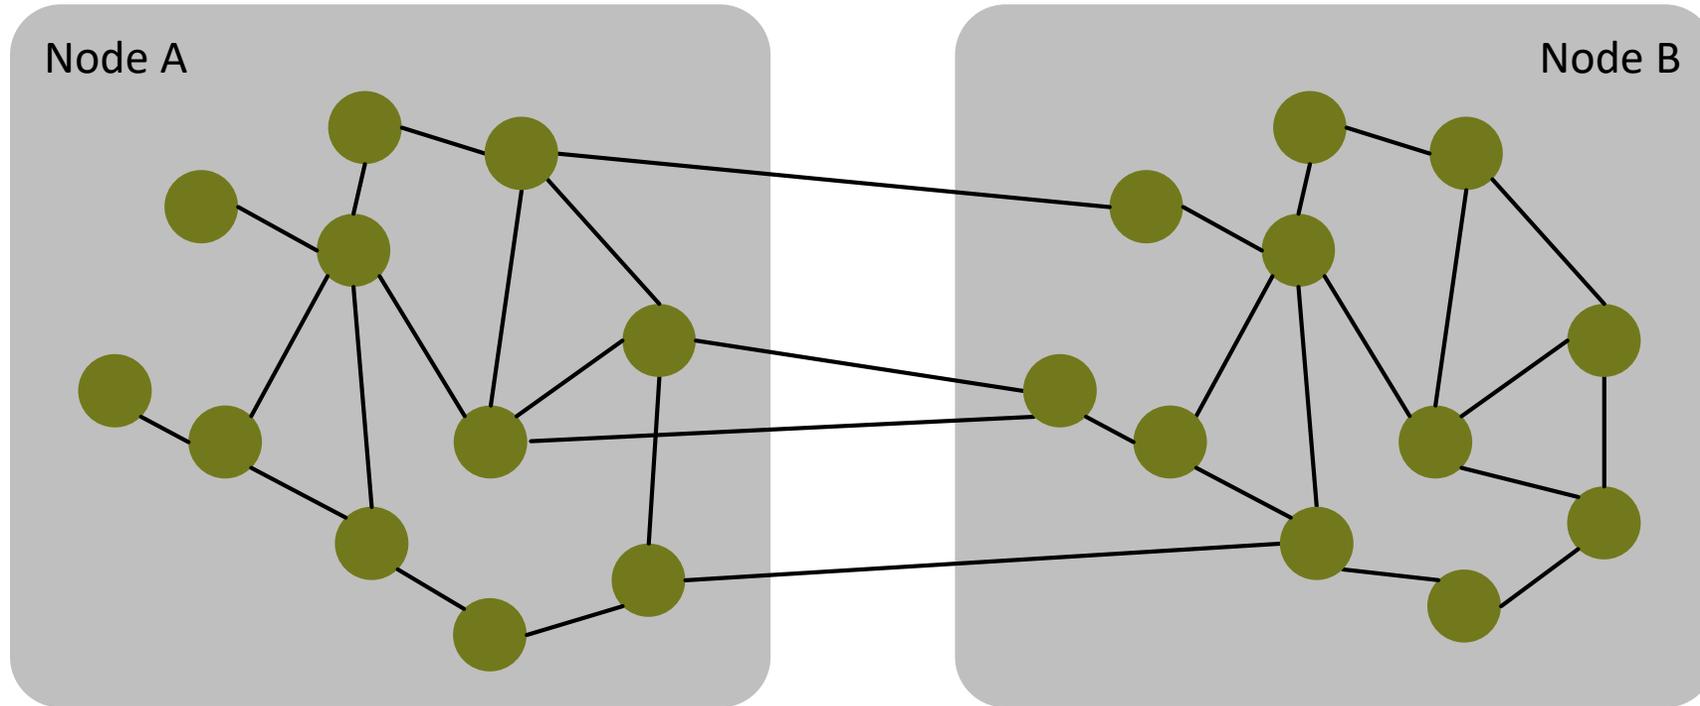
ACCESSING MULTIPLE VERTICES ATOMICALLY

Example: BFS

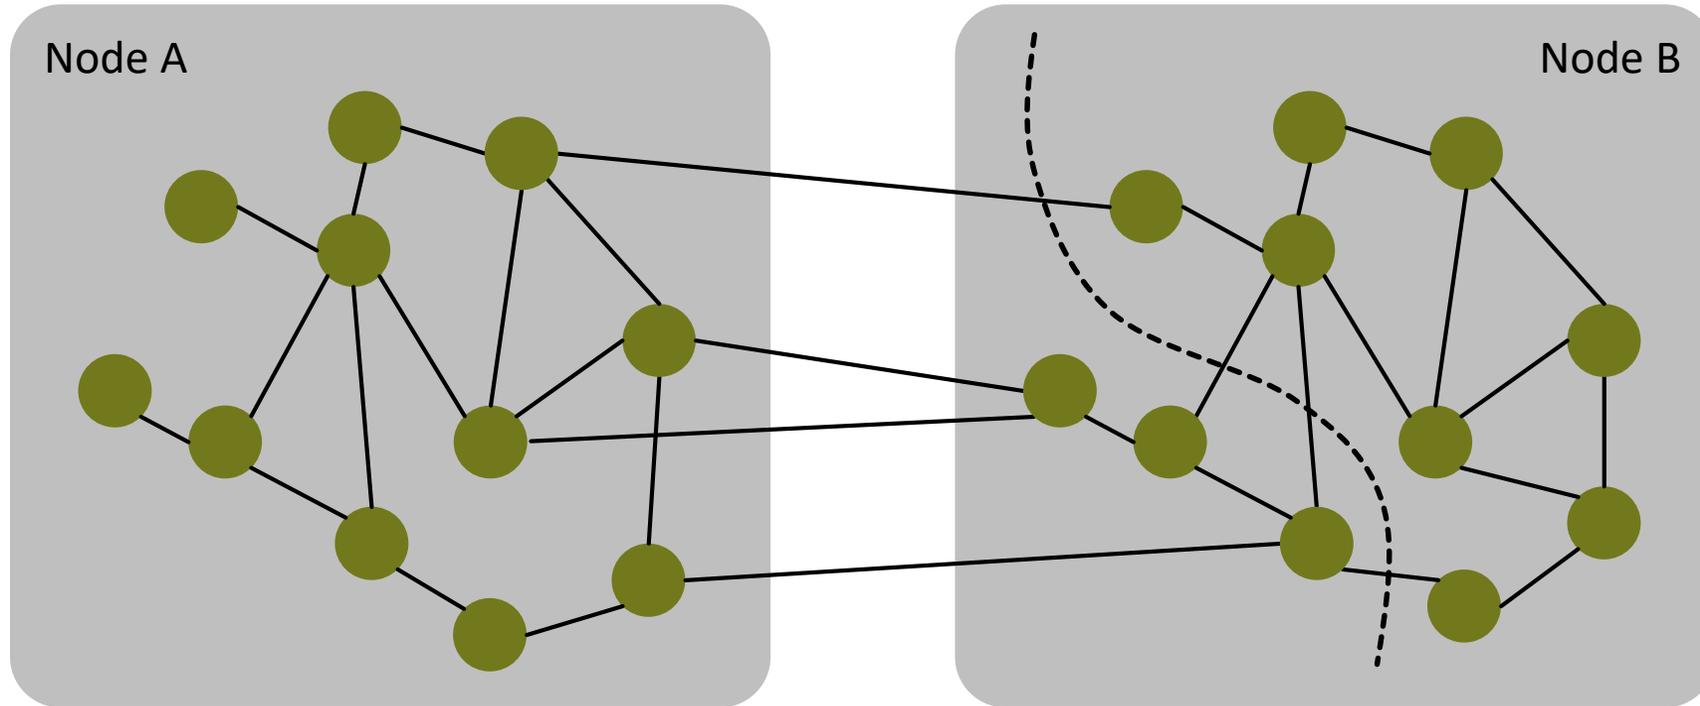


Size (the number of vertices) must be appropriate to minimize overheads from both commits and rollbacks

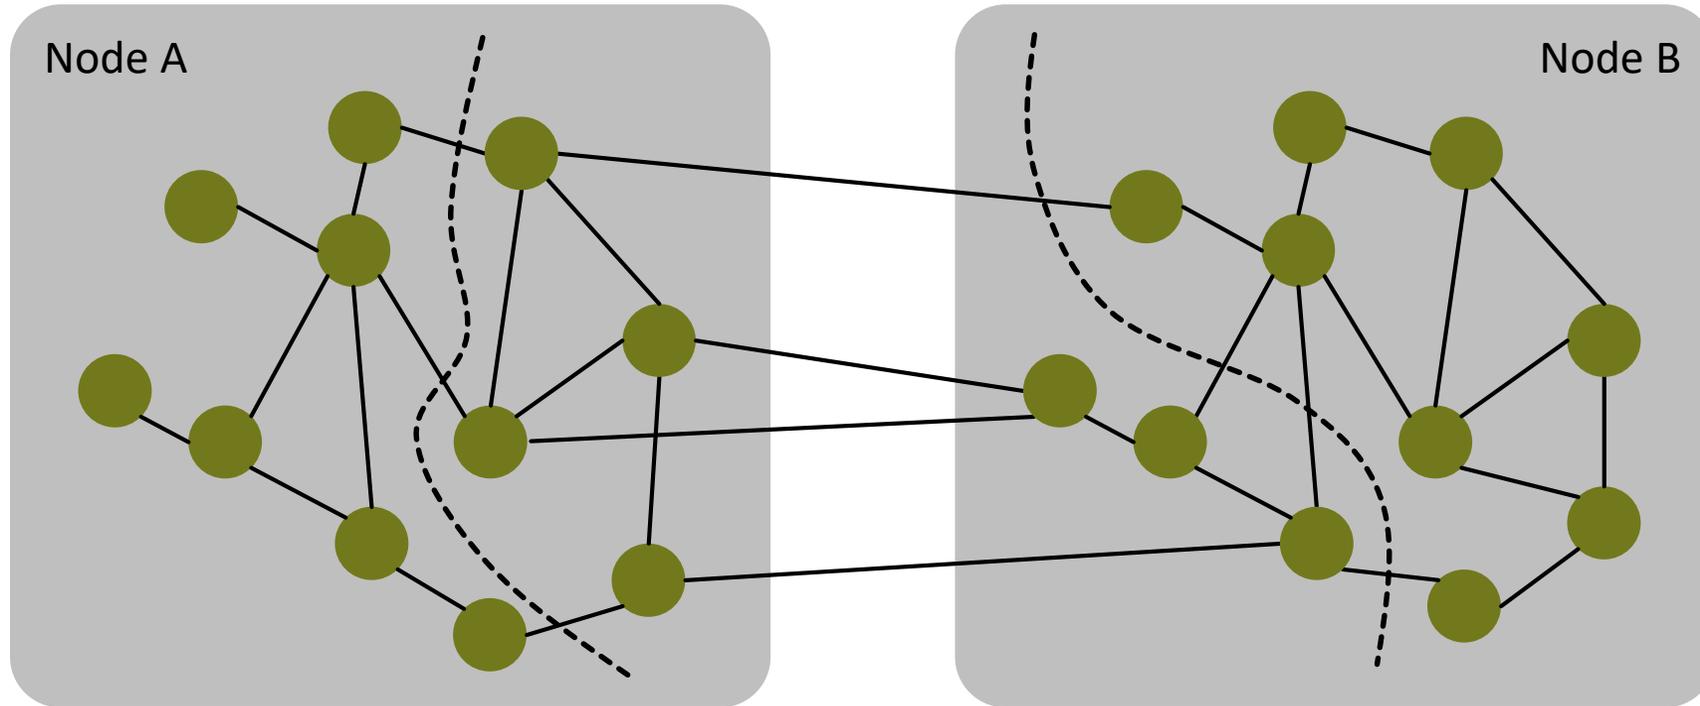
EXECUTING TRANSACTIONS ON MULTIPLE NODES



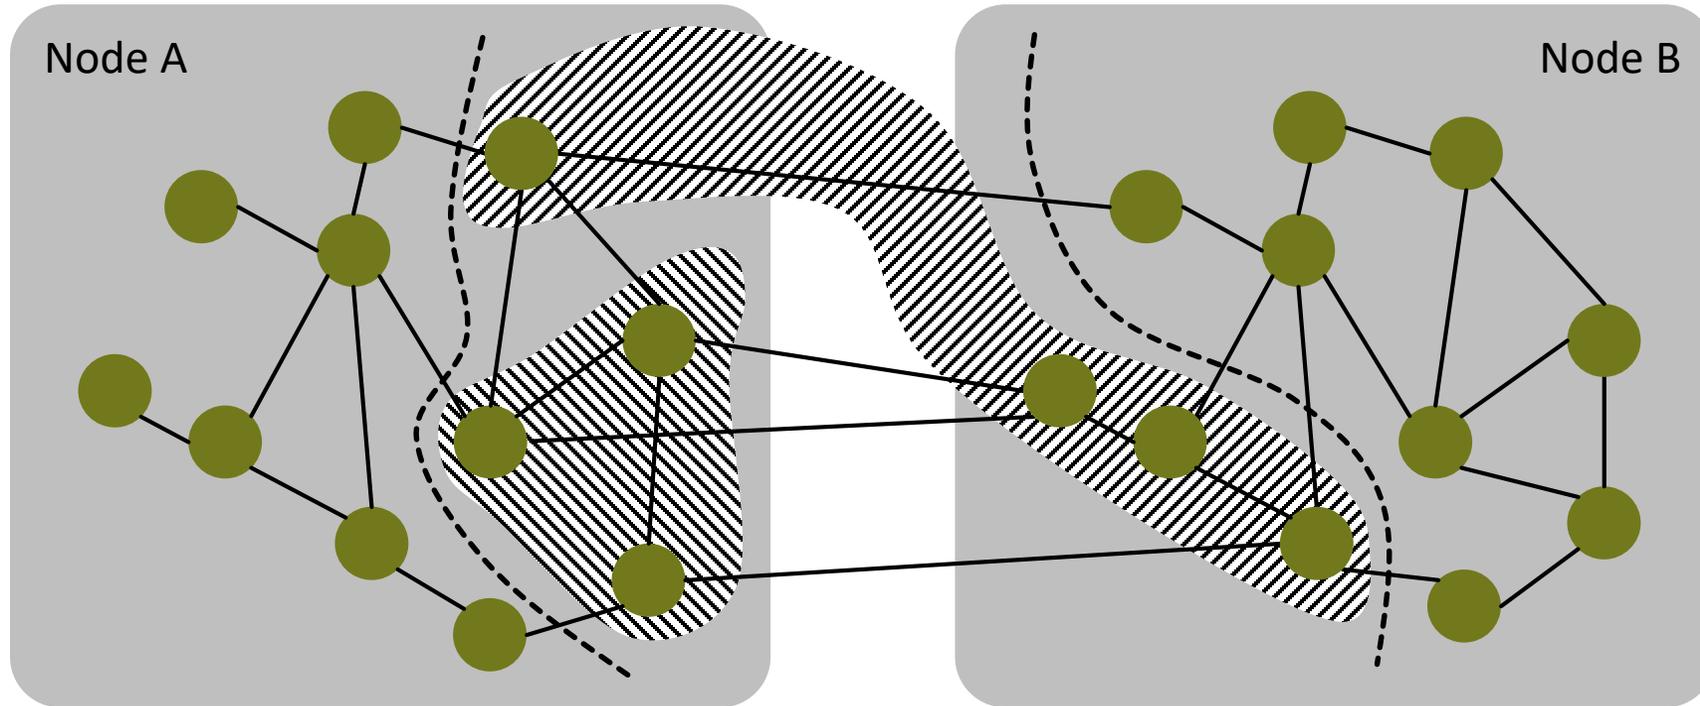
EXECUTING TRANSACTIONS ON MULTIPLE NODES



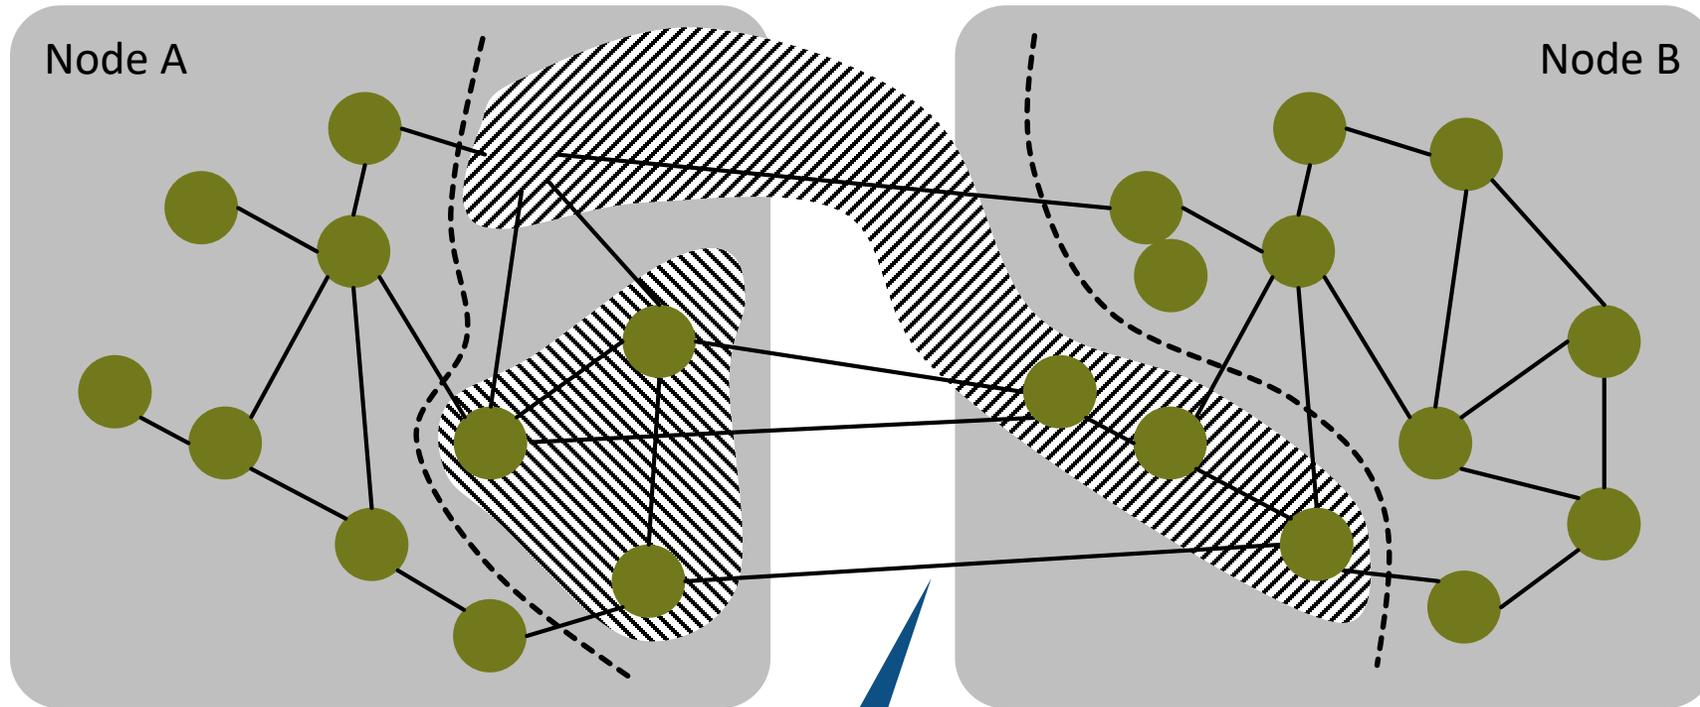
EXECUTING TRANSACTIONS ON MULTIPLE NODES



EXECUTING TRANSACTIONS ON MULTIPLE NODES

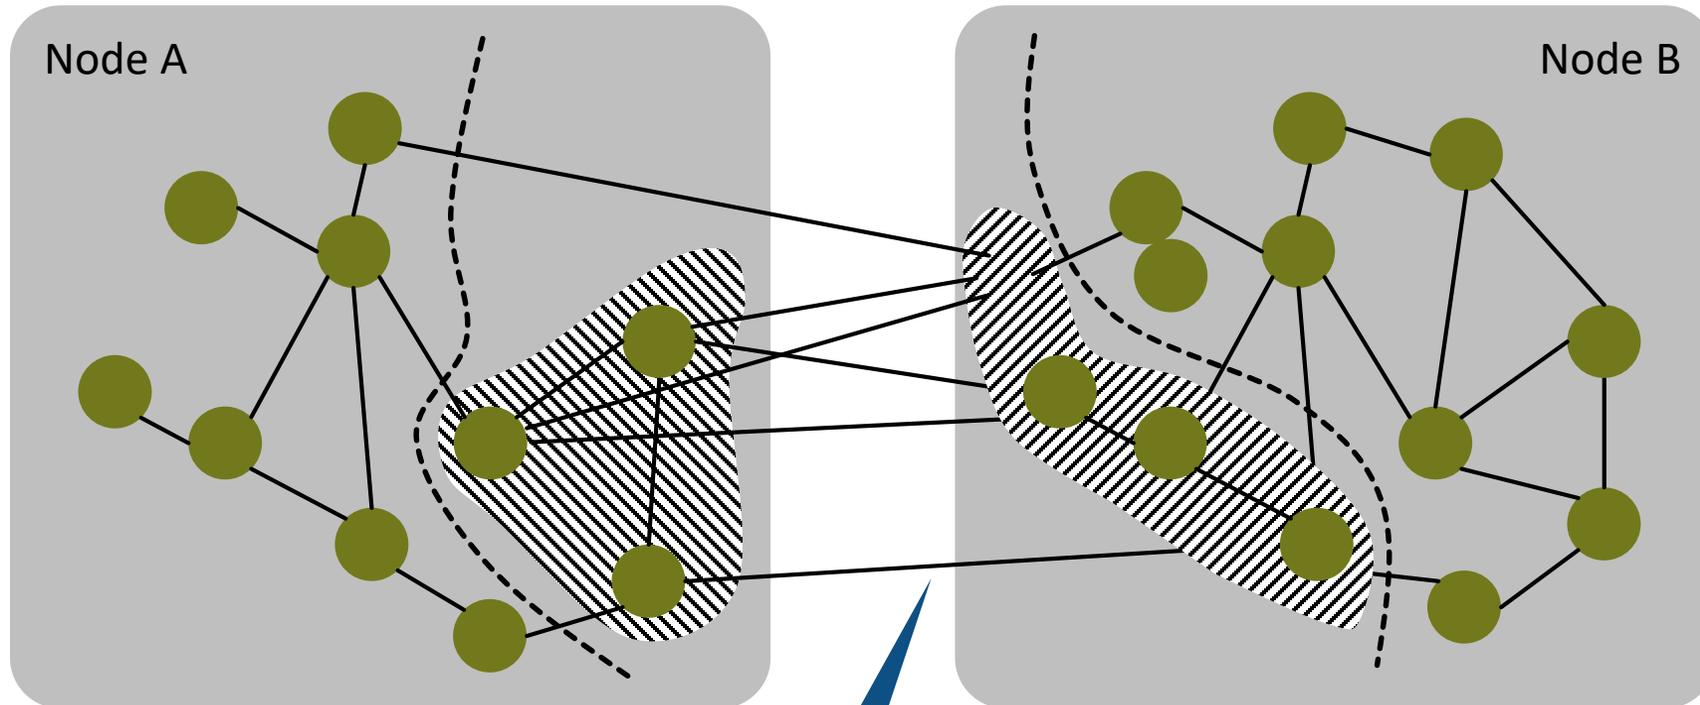


EXECUTING TRANSACTIONS ON MULTIPLE NODES



! vertices must be appropriately relocated to enable execution of a hardware transaction

EXECUTING TRANSACTIONS ON MULTIPLE NODES



! vertices must be appropriately relocated to enable execution of a hardware transaction

PERFORMANCE MODEL

ATOMICS VS TRANSACTIONS

PERFORMANCE MODEL

ATOMICS VS TRANSACTIONS

Time to modify N
vertices with atomics:

$$T_{AT}(N) = A_{AT}N + B_{AT}$$

PERFORMANCE MODEL

ATOMICS VS TRANSACTIONS

Time to modify N
vertices with atomics:

$$T_{AT}(N) = A_{AT}N + B_{AT}$$

Startup
overheads

PERFORMANCE MODEL

ATOMICS VS TRANSACTIONS

Time to modify N
vertices with atomics:

$$T_{AT}(N) = A_{AT}N + B_{AT}$$

Overhead
per vertex

Startup
overheads

PERFORMANCE MODEL

ATOMICS VS TRANSACTIONS

Time to modify N
vertices with atomics:

$$T_{AT}(N) = A_{AT}N + B_{AT}$$

Overhead
per vertex

Startup
overheads

Time to modify N vertices
with a transaction

$$T_{HTM}(N) = A_{HTM}N + B_{HTM}$$

PERFORMANCE MODEL

ATOMICS VS TRANSACTIONS

Time to modify N
vertices with atomics:

$$T_{AT}(N) = A_{AT}N + B_{AT}$$

Overhead
per vertex

Startup
overheads

Time to modify N vertices
with a transaction

$$T_{HTM}(N) = A_{HTM}N + B_{HTM}$$

Overhead
per vertex

Startup
overheads

PERFORMANCE MODEL

ATOMICS VS TRANSACTIONS

Time to modify N
vertices with atomics:

$$T_{AT}(N) = A_{AT}N + B_{AT}$$

Overhead
per vertex

Startup
overheads

Time to modify N vertices
with a transaction

$$T_{HTM}(N) = A_{HTM}N + B_{HTM}$$

Overhead
per vertex

Startup
overheads

We predict that:

$$B_{AT} < B_{HTM}$$

$$A_{AT} > A_{HTM}$$

PERFORMANCE MODEL

ATOMICS VS TRANSACTIONS

Time to modify N vertices with atomics:

$$T_{AT}(N) = A_{AT}N + B_{AT}$$

Overhead per vertex

Startup overheads

Time to modify N vertices with a transaction

$$T_{HTM}(N) = A_{HTM}N + B_{HTM}$$

Overhead per vertex

Startup overheads

We predict that:

$$B_{AT} < B_{HTM}$$

$$A_{AT} > A_{HTM}$$



Transaction startup overheads dominate

PERFORMANCE MODEL

ATOMICS VS TRANSACTIONS

Time to modify N vertices with atomics:

$$T_{AT}(N) = A_{AT}N + B_{AT}$$

Overhead per vertex

Startup overheads

Time to modify N vertices with a transaction

$$T_{HTM}(N) = A_{HTM}N + B_{HTM}$$

Overhead per vertex

Startup overheads

! Transactions' cost grows slower

We predict that:

$$B_{AT} < B_{HTM}$$

$$A_{AT} > A_{HTM}$$

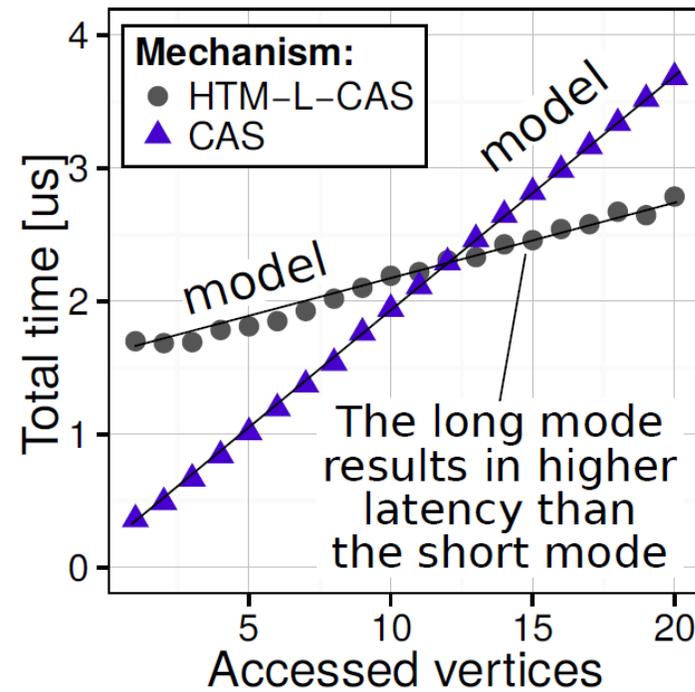
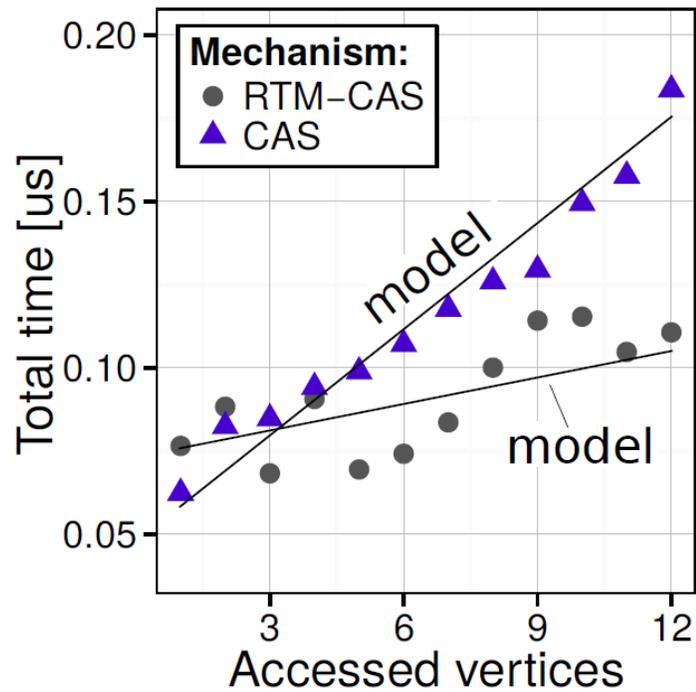
! Transaction startup overheads dominate

PERFORMANCE MODEL

ATOMICS VS TRANSACTIONS

PERFORMANCE MODEL

ATOMICS VS TRANSACTIONS

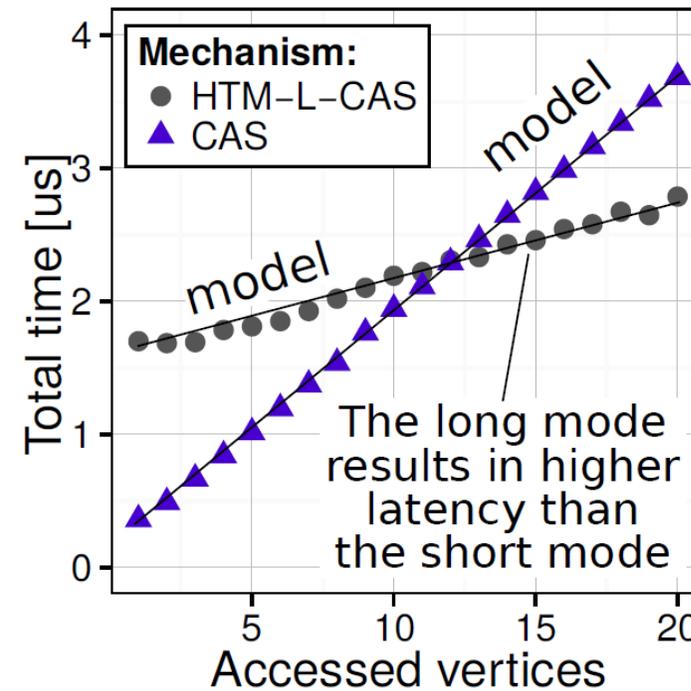
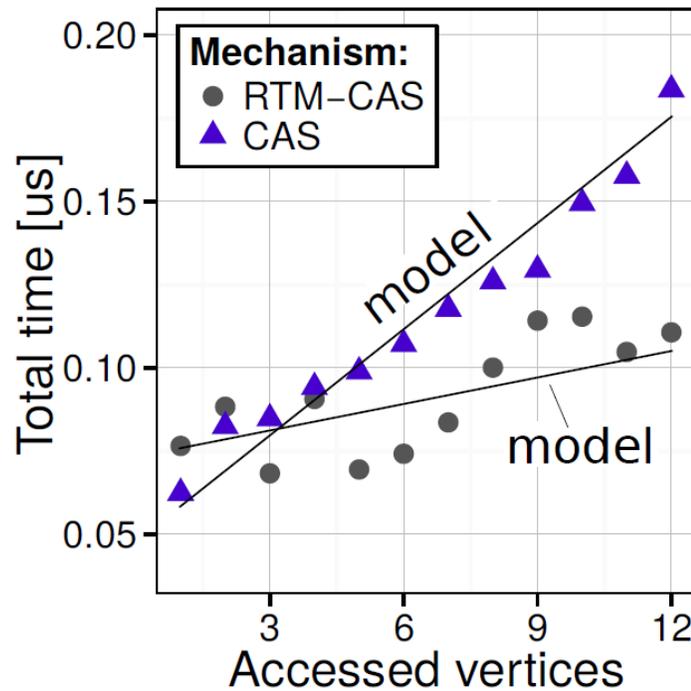


PERFORMANCE MODEL

ATOMICS VS TRANSACTIONS

Indeed:

$$B_{AT} < B_{HTM}$$

$$A_{AT} > A_{HTM}$$


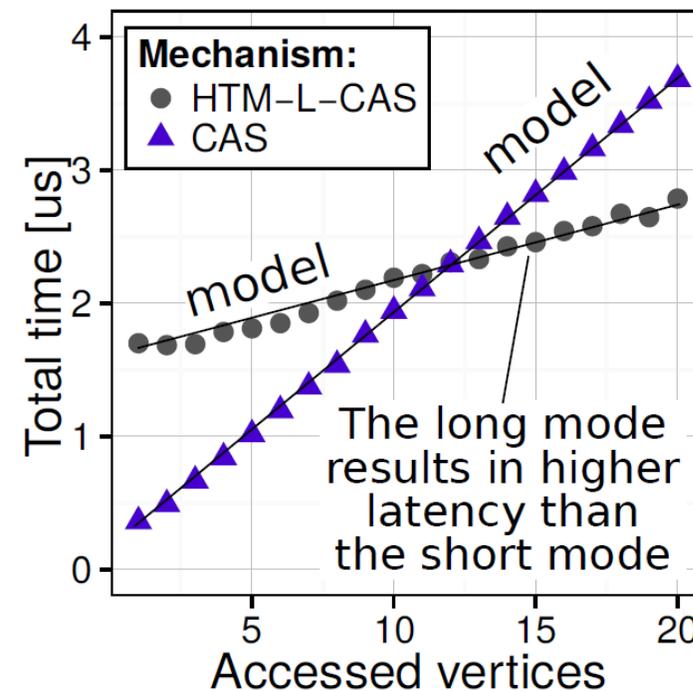
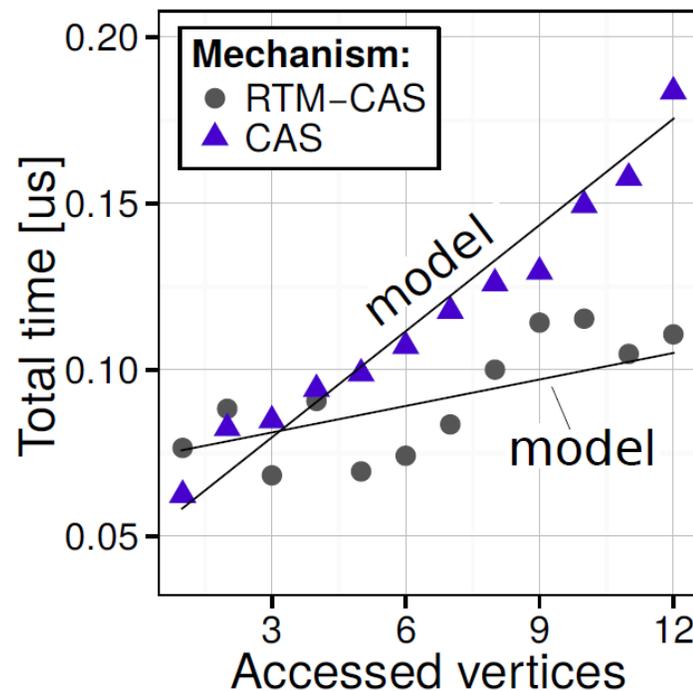
PERFORMANCE MODEL

ATOMICS VS TRANSACTIONS

- Can we amortize HTM startup/commit overheads with larger transaction sizes?

Indeed:

$$B_{AT} < B_{HTM}$$

$$A_{AT} > A_{HTM}$$


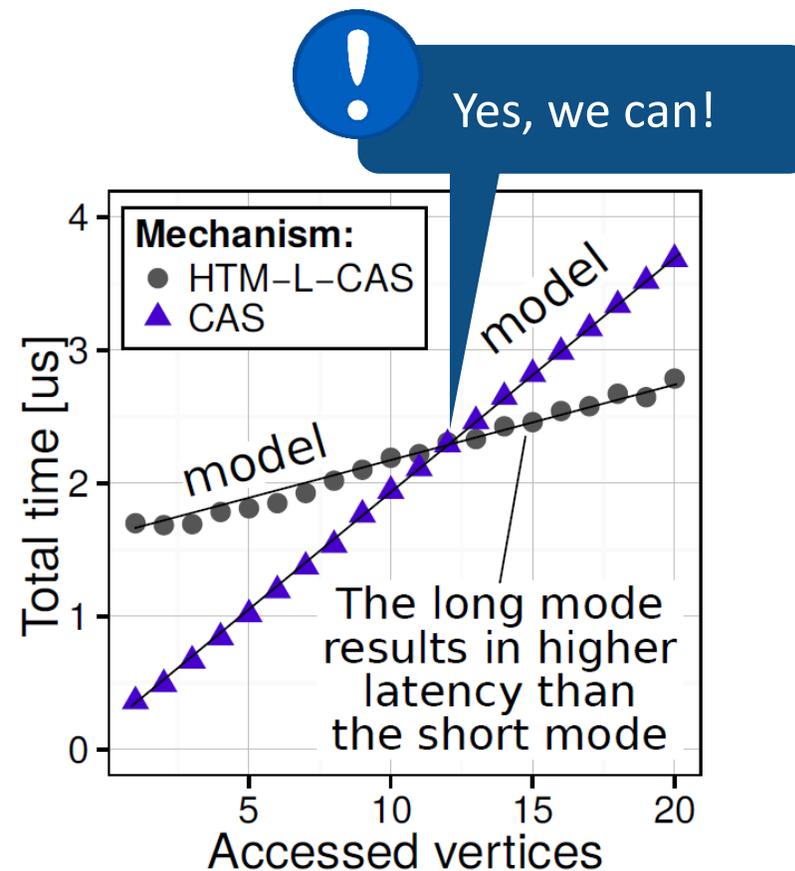
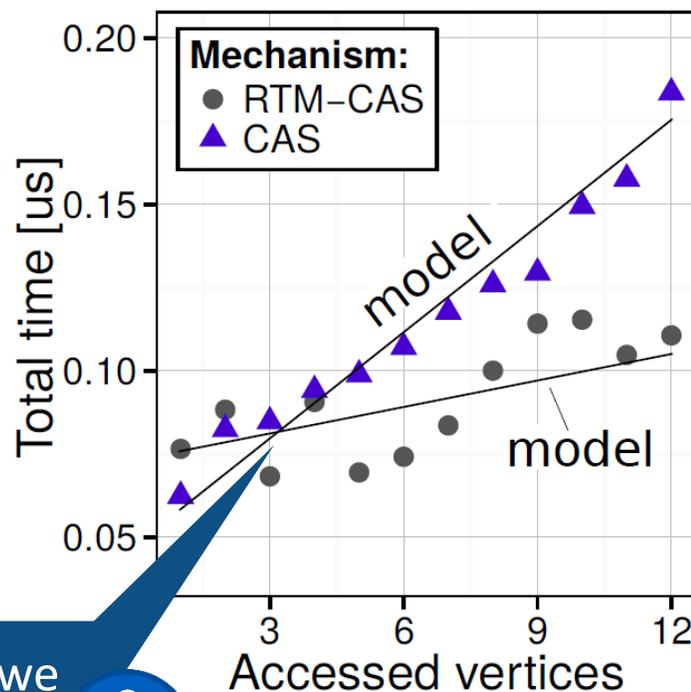
PERFORMANCE MODEL

ATOMICS VS TRANSACTIONS

- Can we amortize HTM startup/commit overheads with larger transaction sizes?

Indeed:

$$B_{AT} < B_{HTM}$$

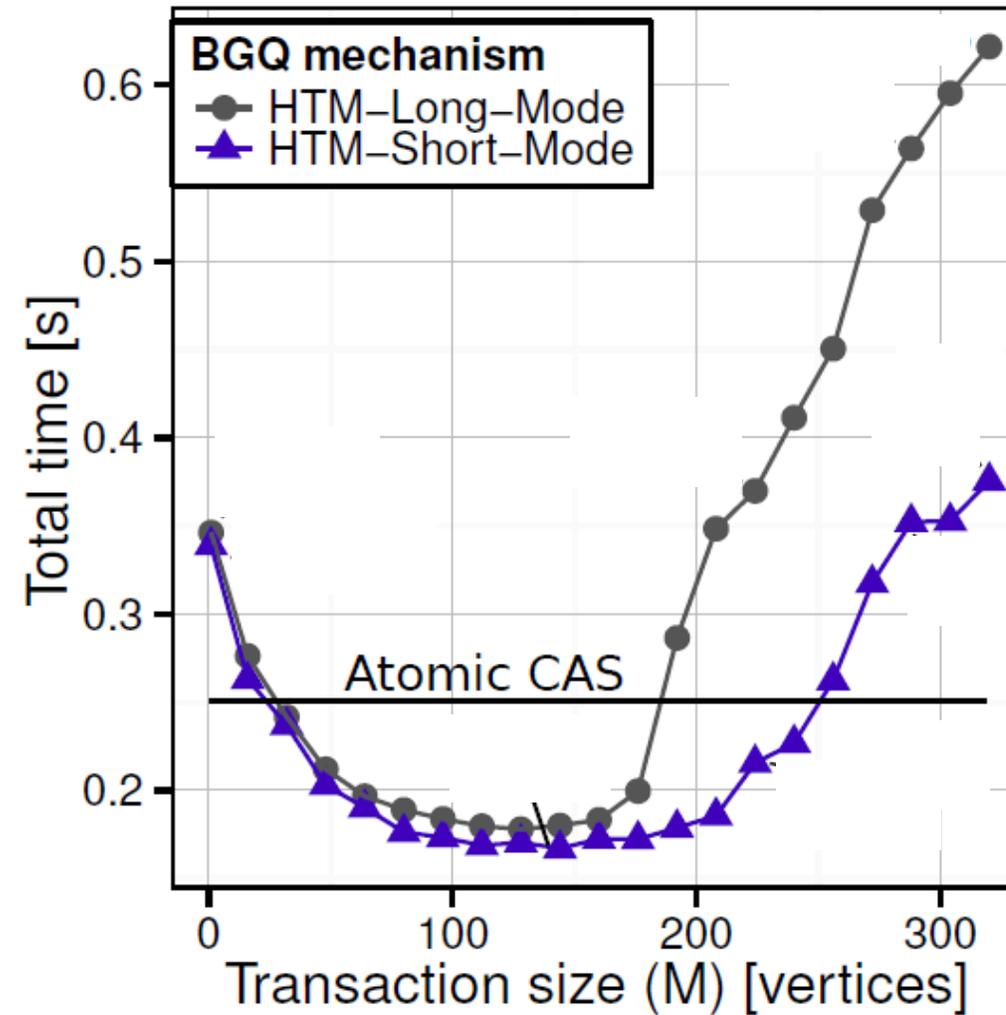
$$A_{AT} > A_{HTM}$$


MULTI-VERTEX TRANSACTIONS

MARKING VERTICES AS VISITED

MULTI-VERTEX TRANSACTIONS

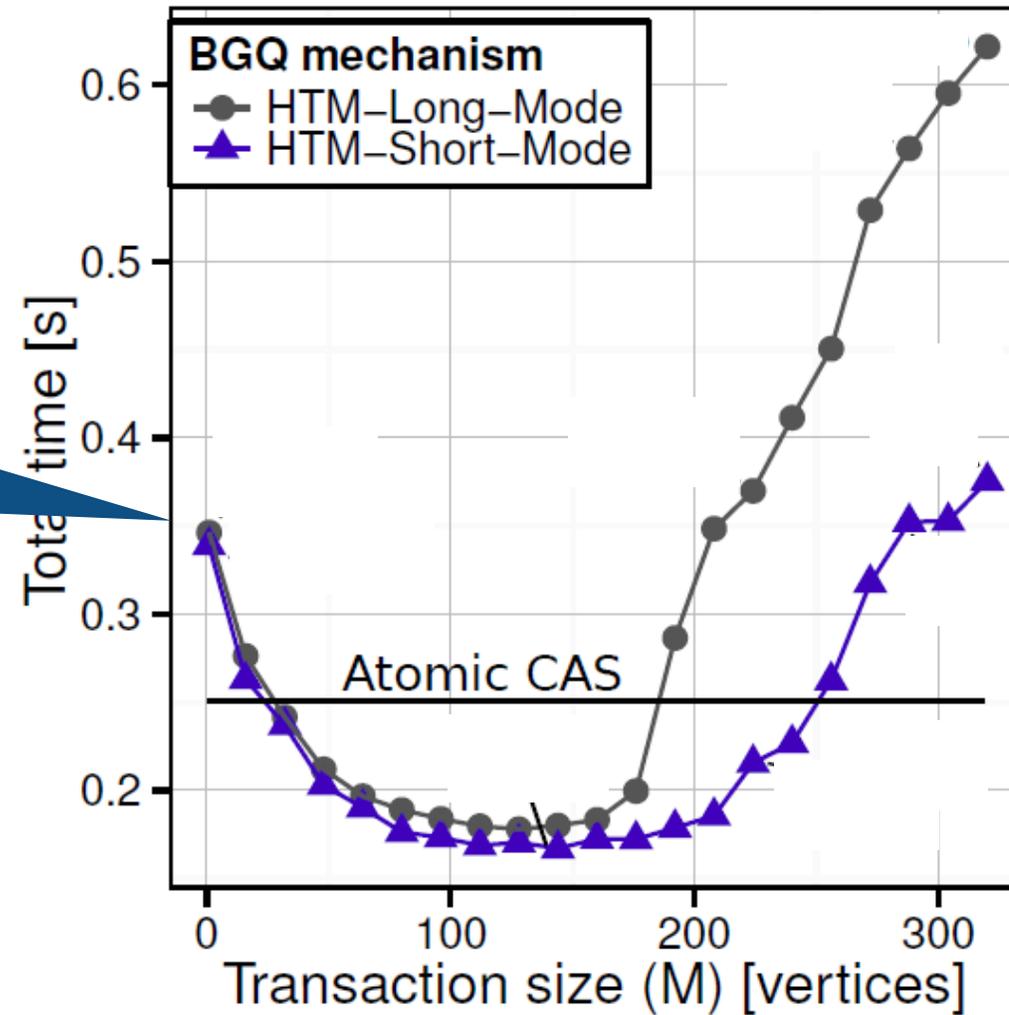
MARKING VERTICES AS VISITED



MULTI-VERTEX TRANSACTIONS

MARKING VERTICES AS VISITED

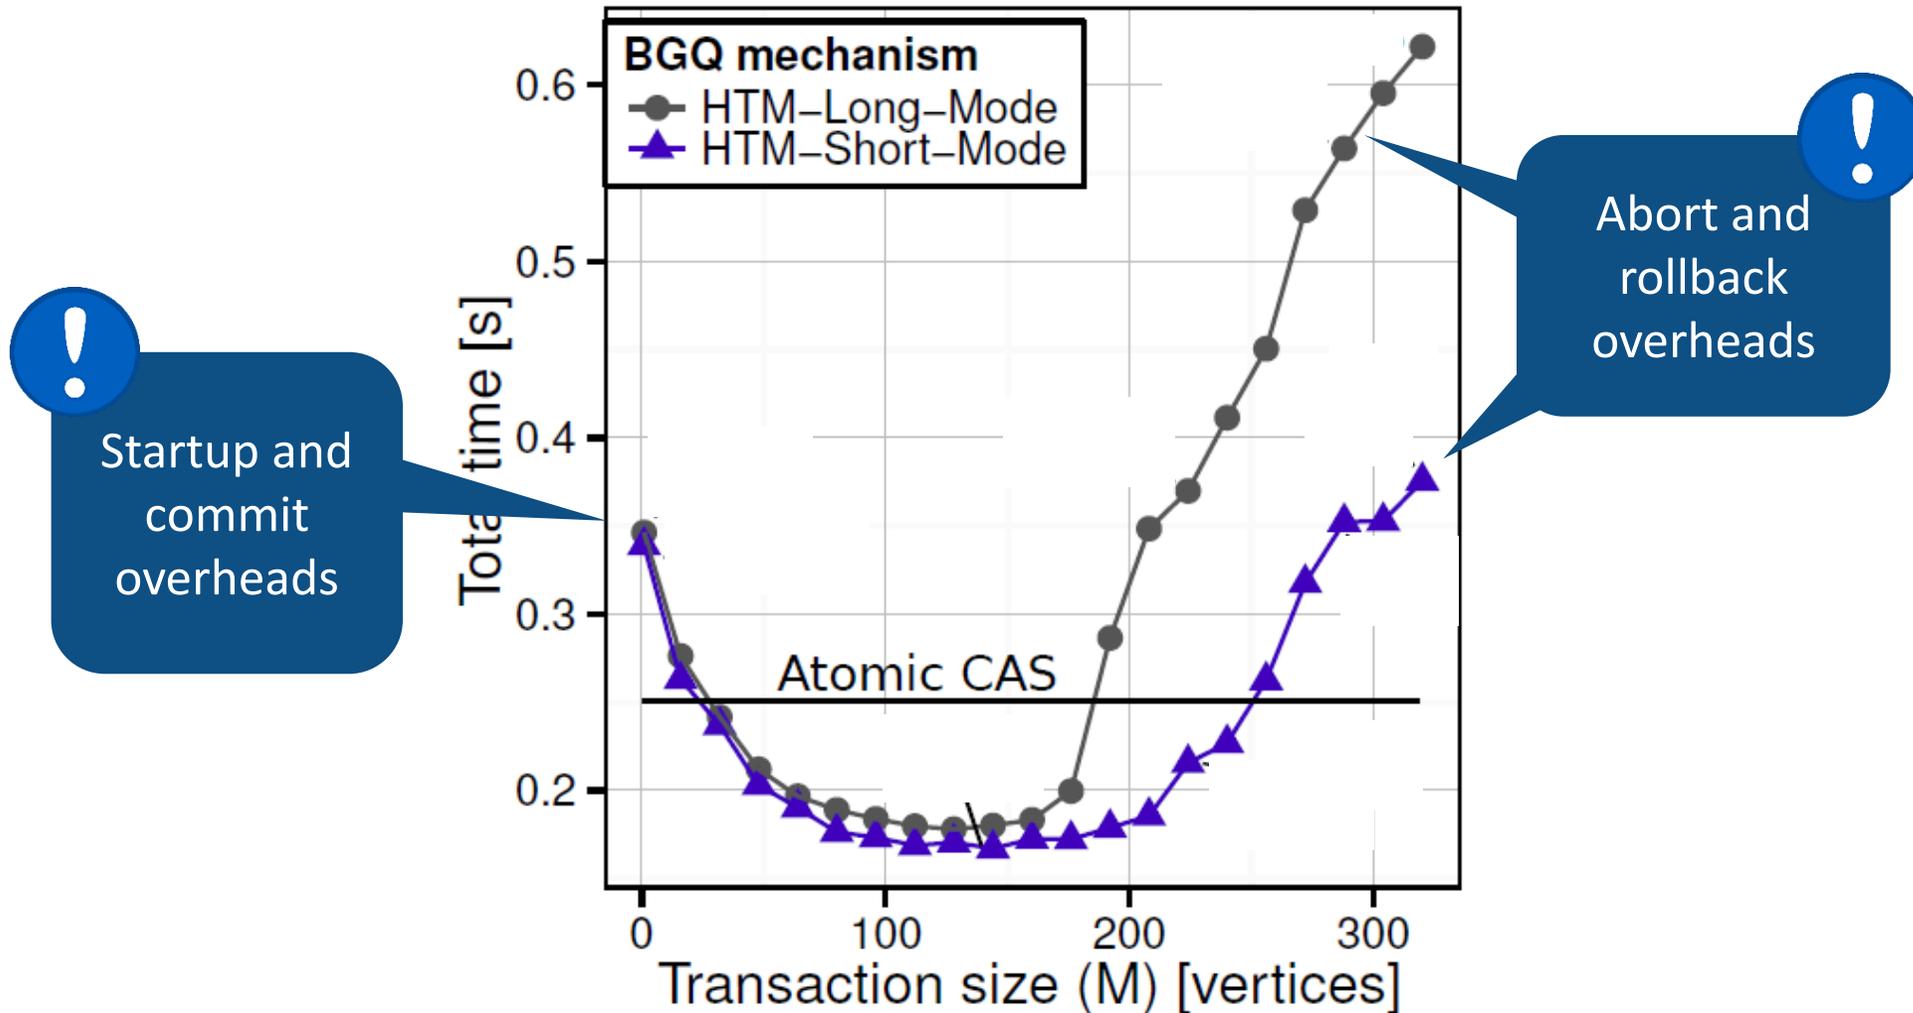
! Startup and commit overheads





MULTI-VERTEX TRANSACTIONS

MARKING VERTICES AS VISITED





MULTI-VERTEX TRANSACTIONS

MARKING VERTICES AS VISITED

