


T. HOEFLER

Extreme-Scale Graphs

Invited talk at Supercomputing Frontiers and Innovation 2019, Warsaw, Poland


Special thanks to my student Maciej Besta

With contributions from 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 at SPCL and Tsinghua University




Extreme-Scale Graphs

Extreme-Scale Graphs



Extreme-Scale Graphs

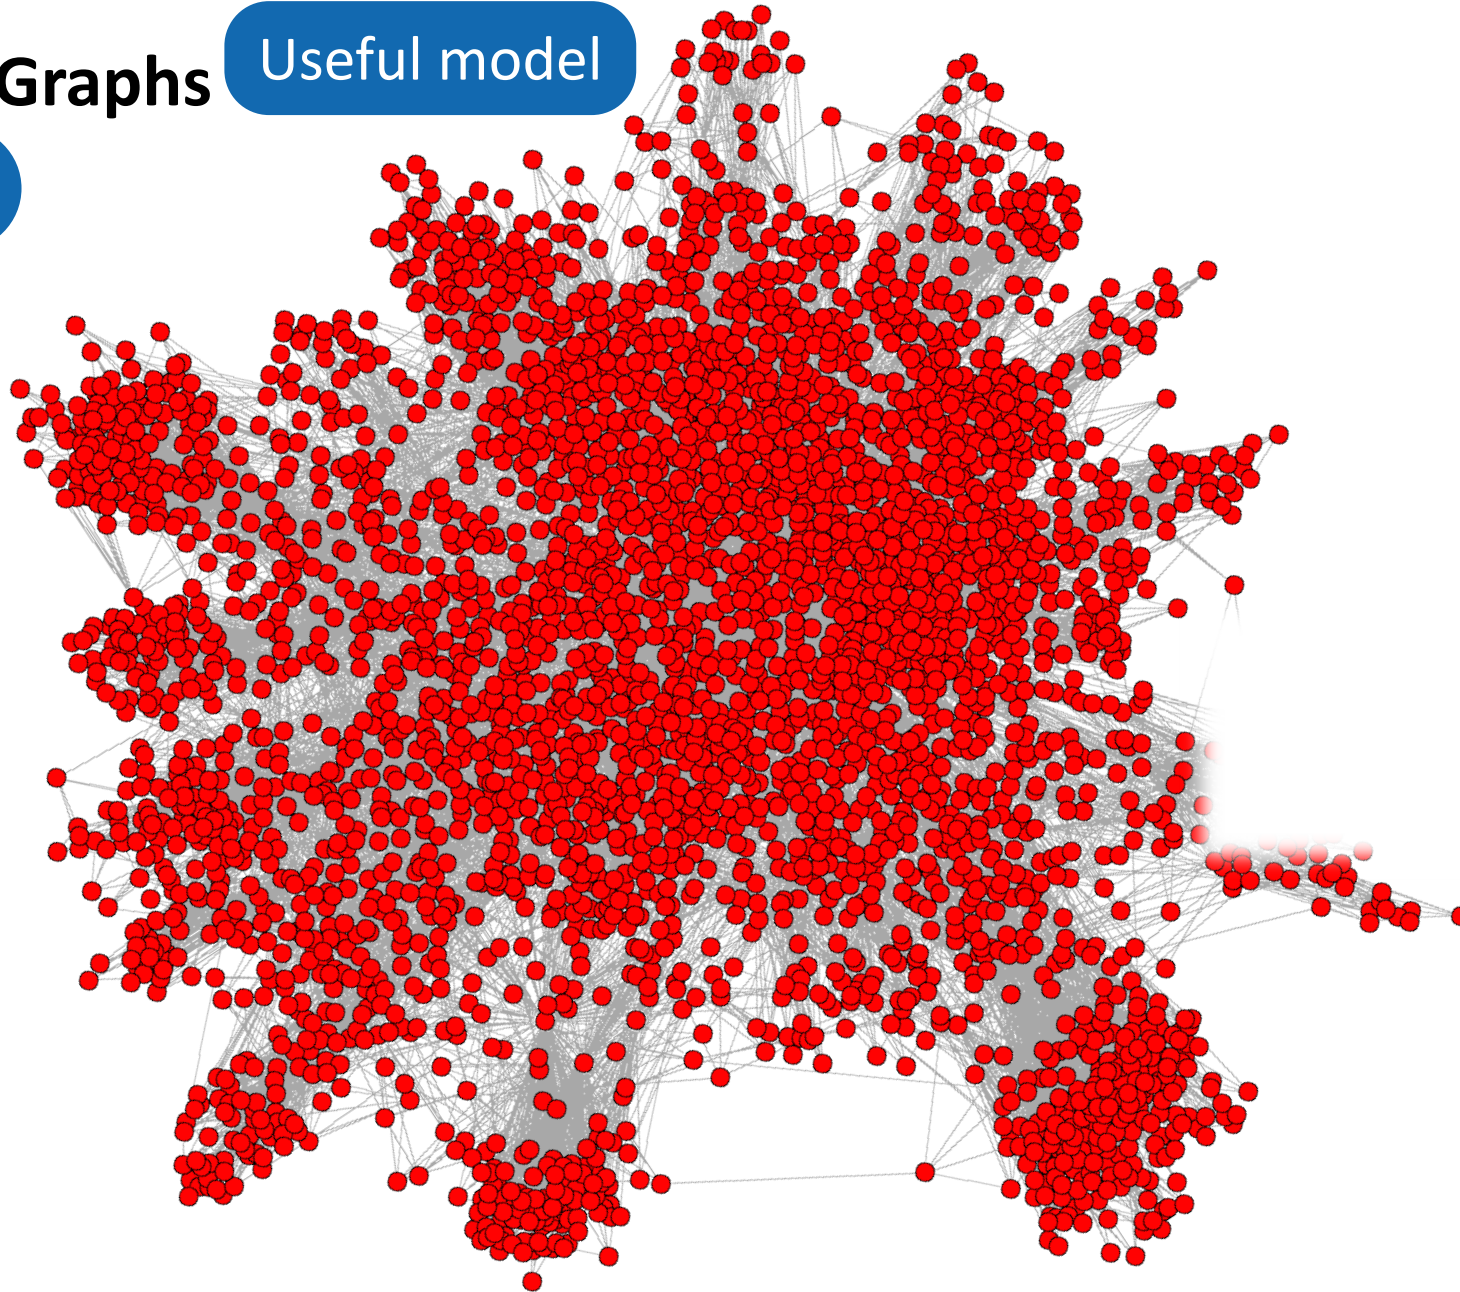
Why do we care?



Extreme-Scale Graphs

Useful model

Why do we care?

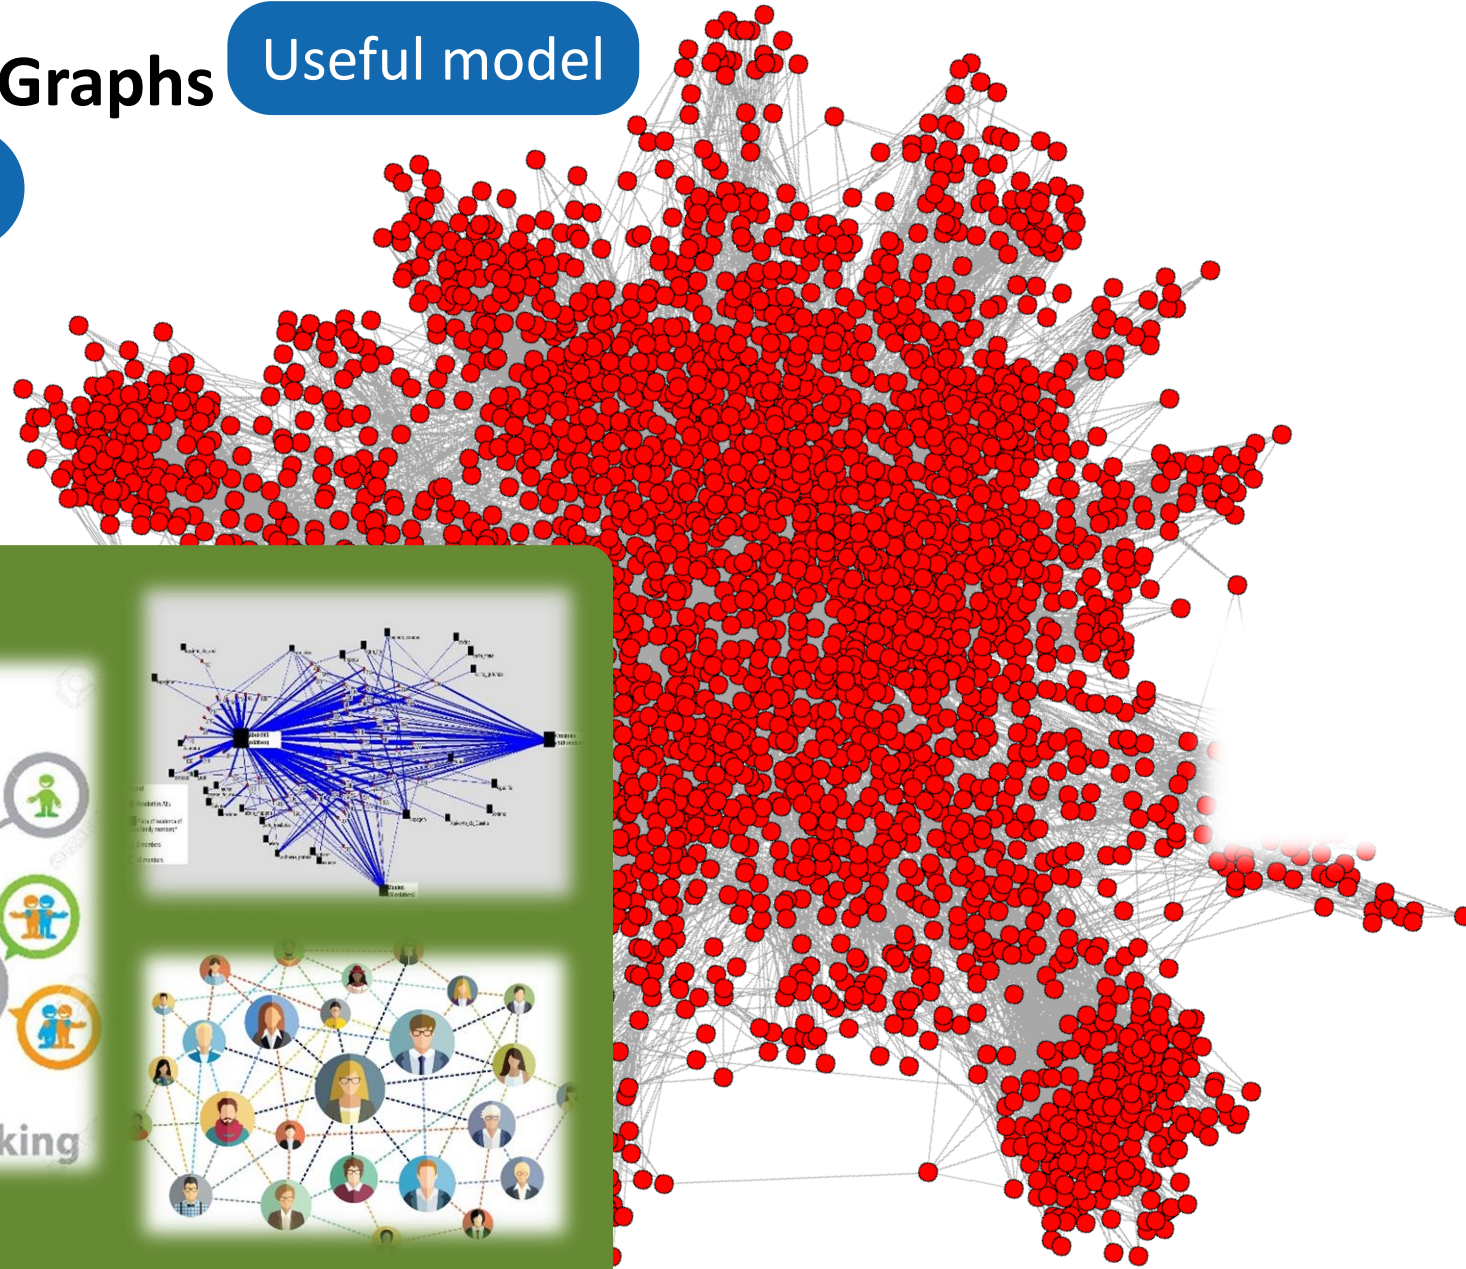
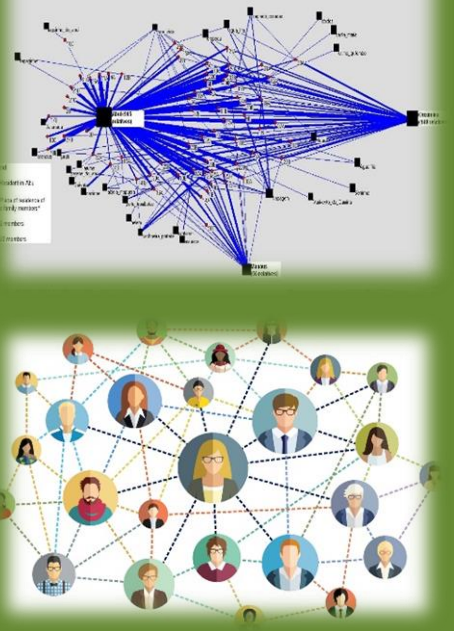


Extreme-Scale Graphs

Useful model

Why do we care?

Social networks

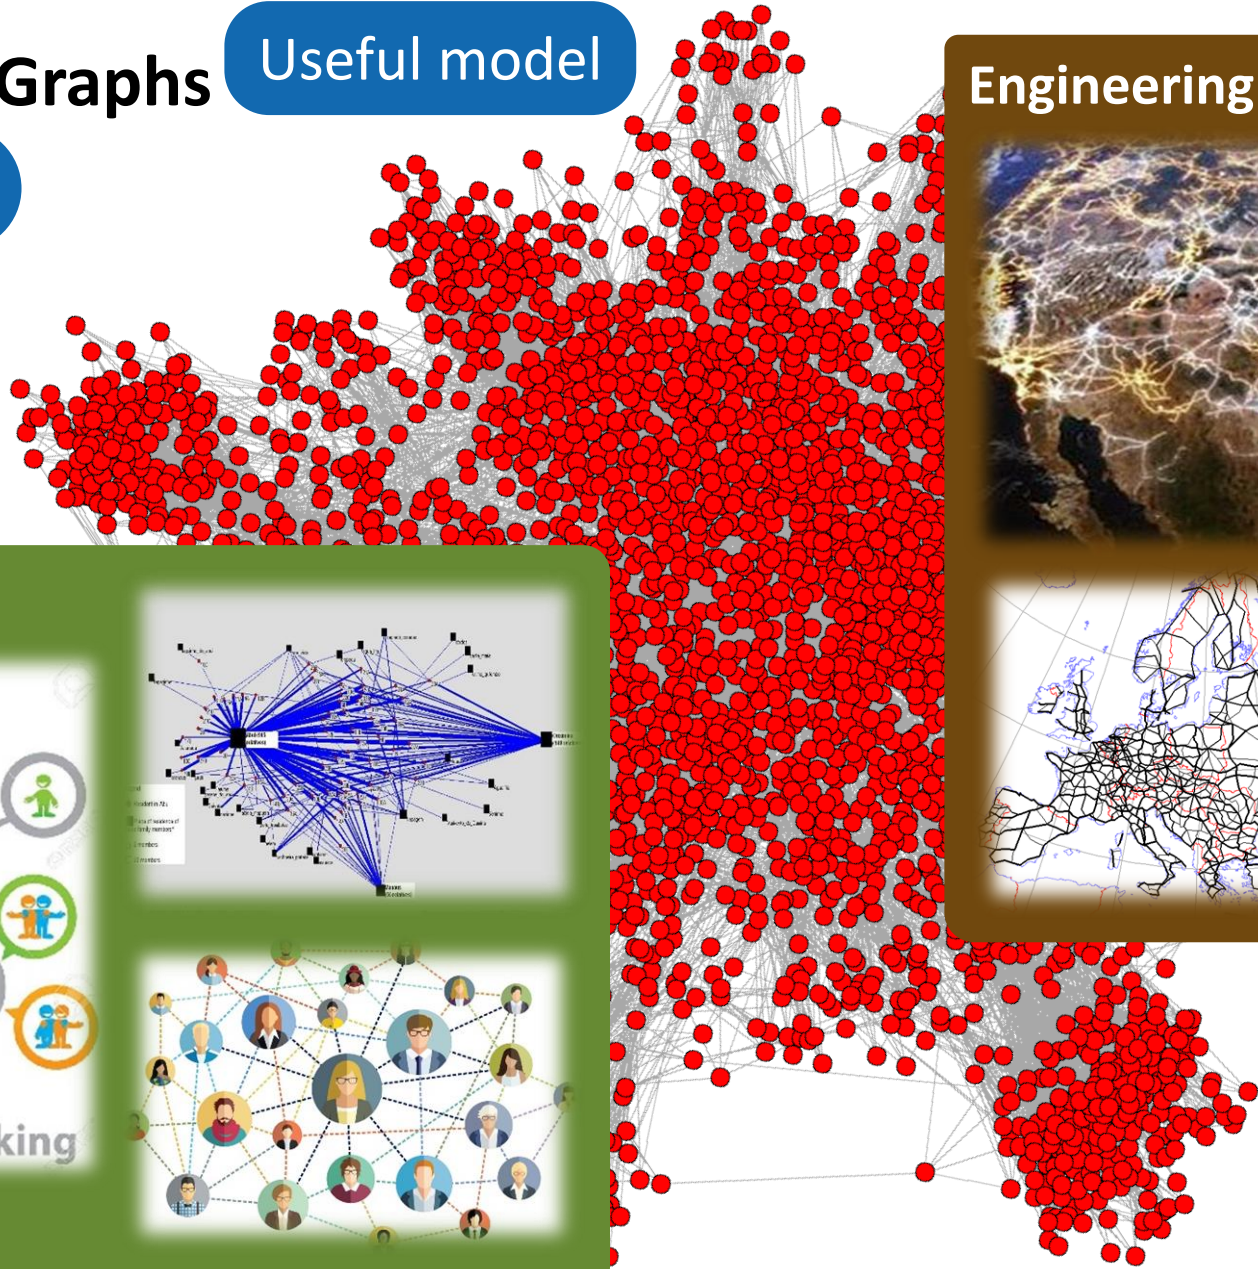



Extreme-Scale Graphs

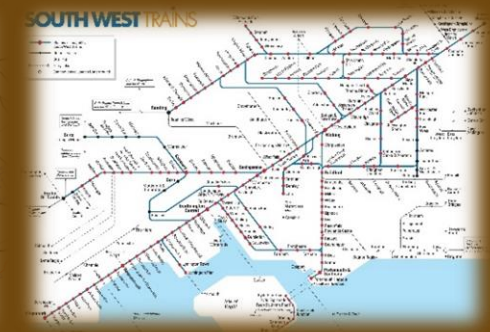
Useful model

Why do we care?

Social networks



Engineering networks



Extreme-Scale Graphs




Useful model

Why do we care?

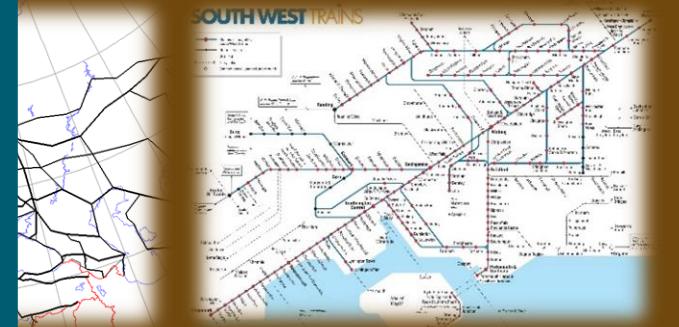
Social networks



Biological networks



Engineering networks



Extreme-Scale Graphs



Useful model

Why do we care?

Social networks



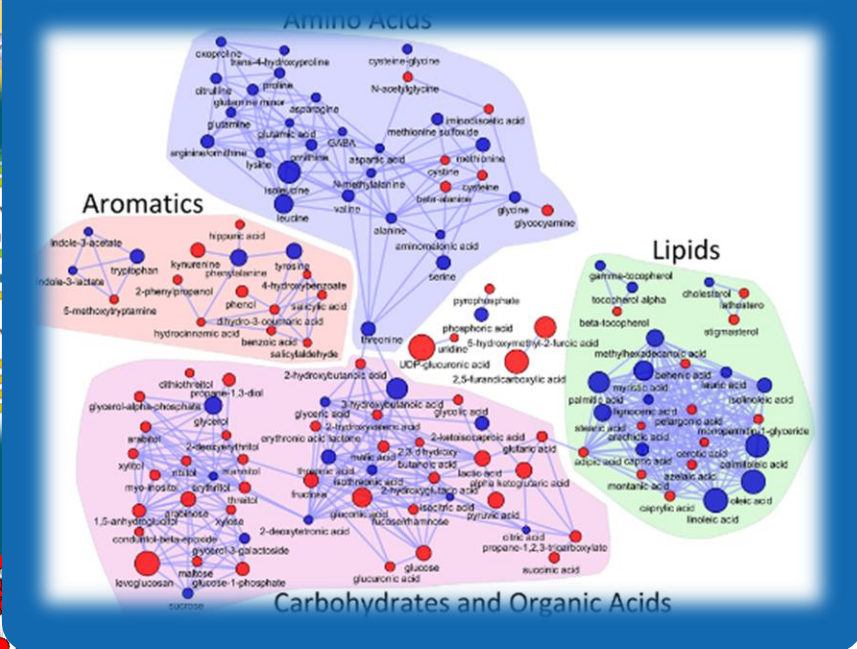
Biological networks



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


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



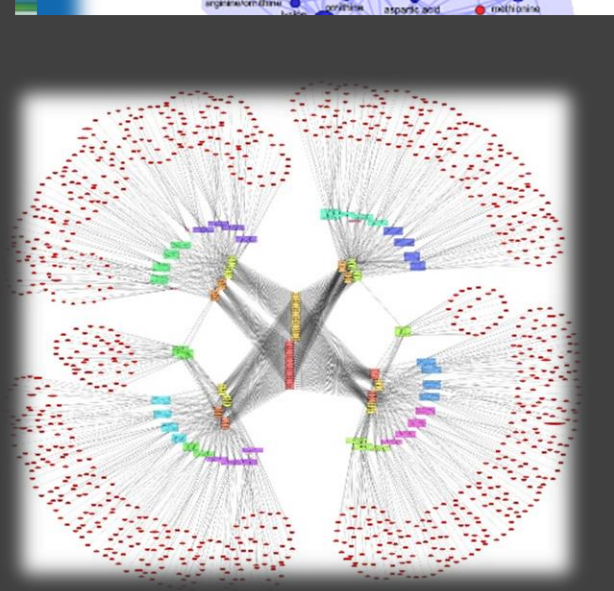
Communication networks



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

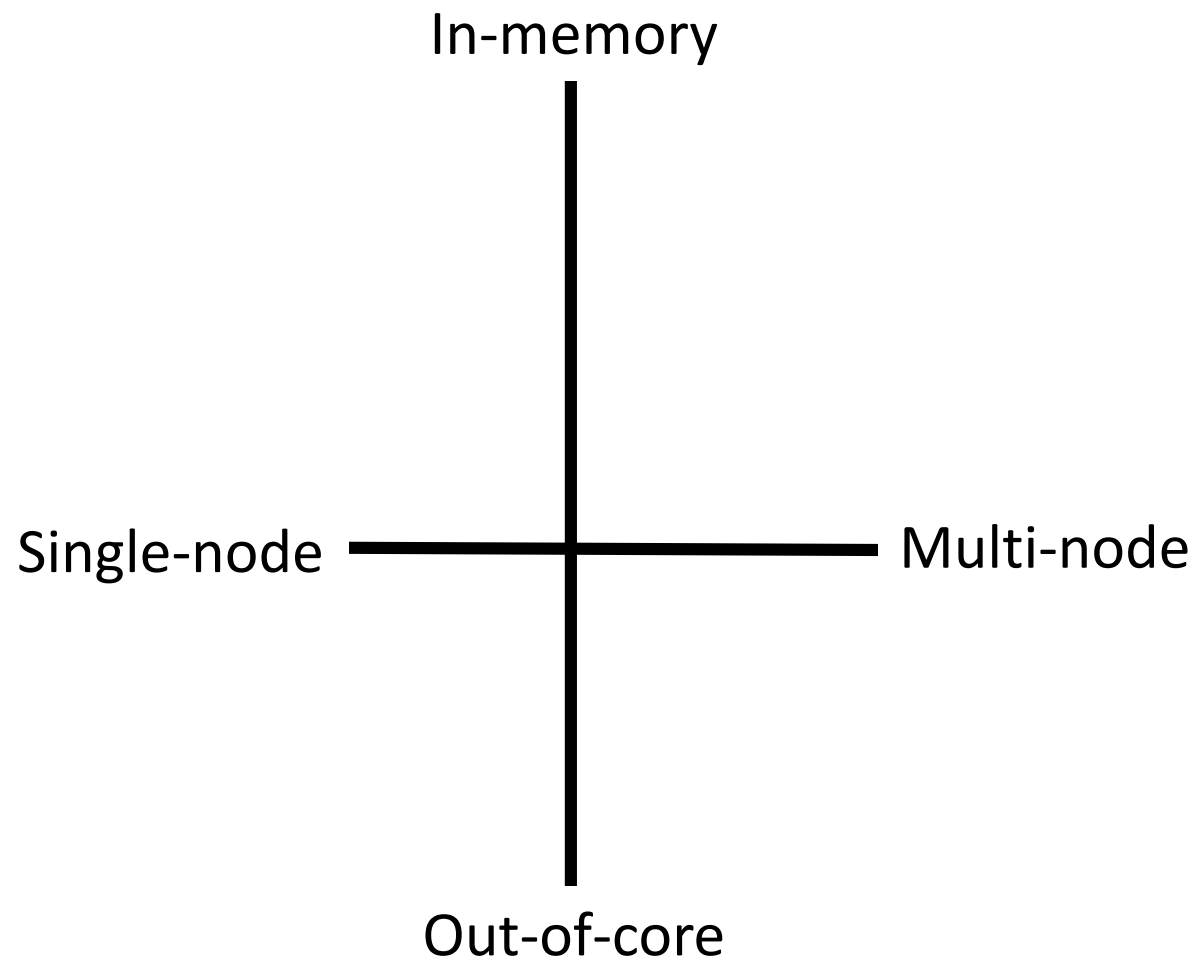
- 📍 Track: Graph Processing devroom
- 📍 Room: AW1.126
- 📅 Day: Saturday
- 🕒 Start: 12:45
- 🕒 End: 13:35



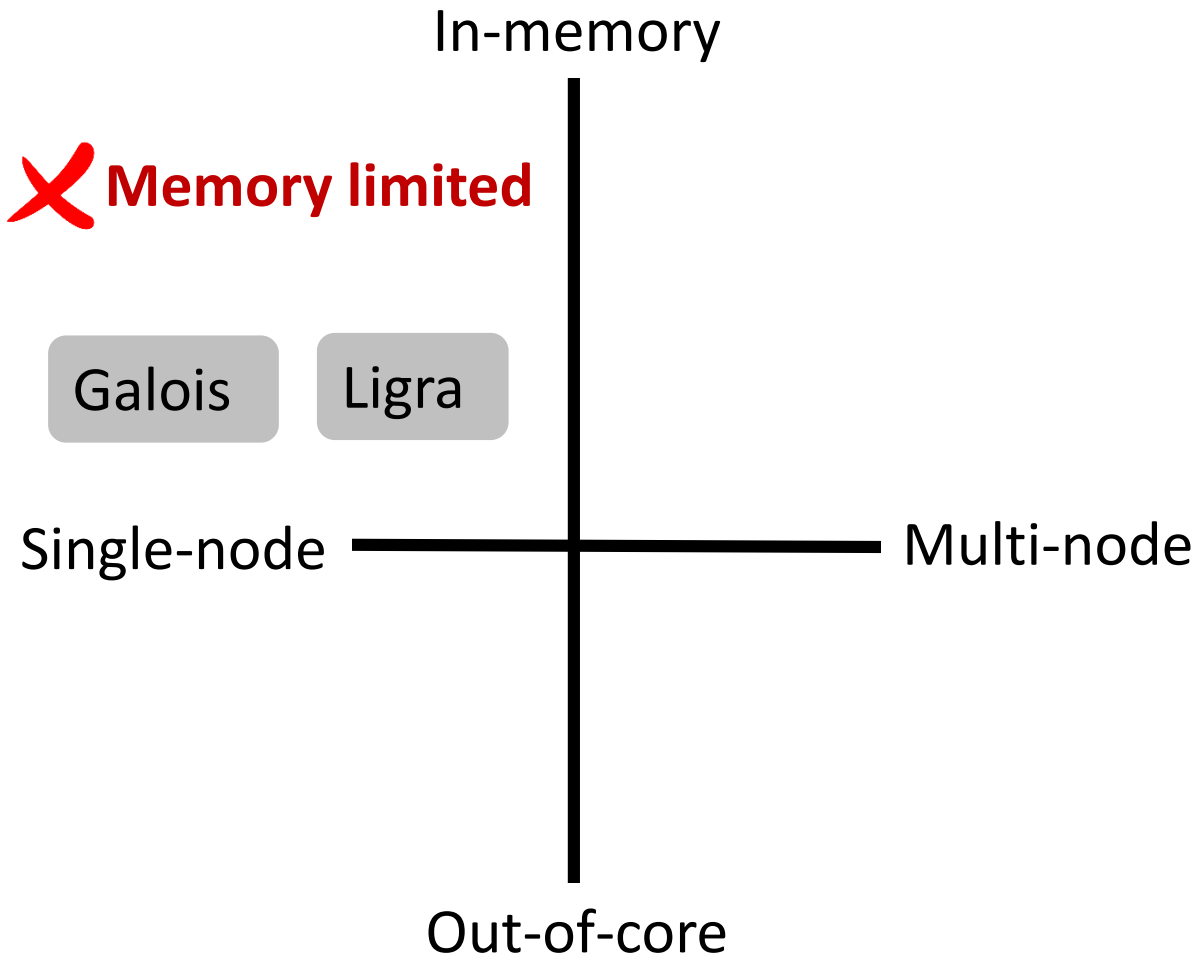
Bruno Latour wrote a book about philosophy (an inquiry into modes of existence). He decided that the paper book was no place for the numerous footnotes, documentation or glossary, instead giving access to all this information surrounding the book through a web application which would present itself as 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

Practice of Extreme-Scale Graph Processing

Practice of Extreme-Scale Graph Processing



Practice of Extreme-Scale Graph Processing



Practice of Extreme-Scale Graph Processing

In-memory

X Memory limited

Galois Ligra

Single-node

Multi-node

G-Store Mosaic Chaos

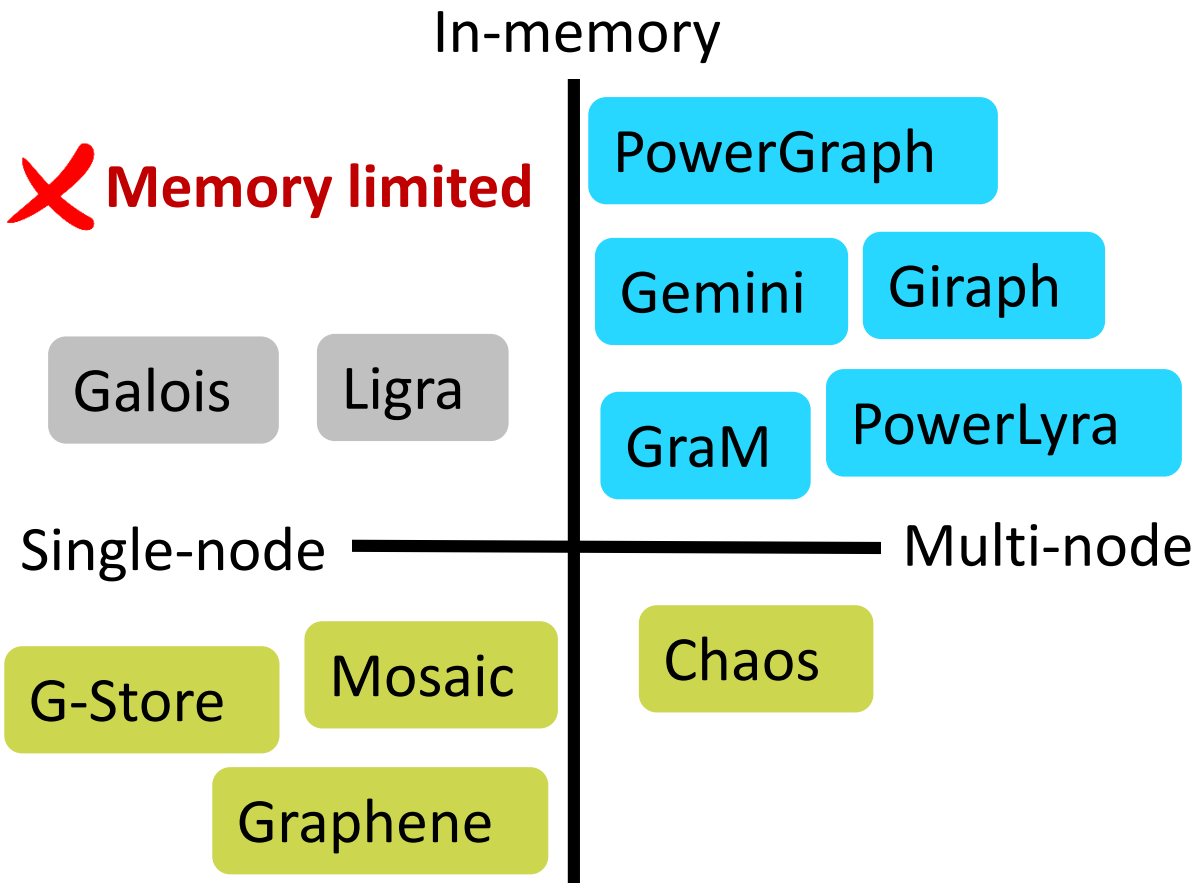
Graphene

Out-of-core

X Too slow

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

Practice of Extreme-Scale Graph Processing

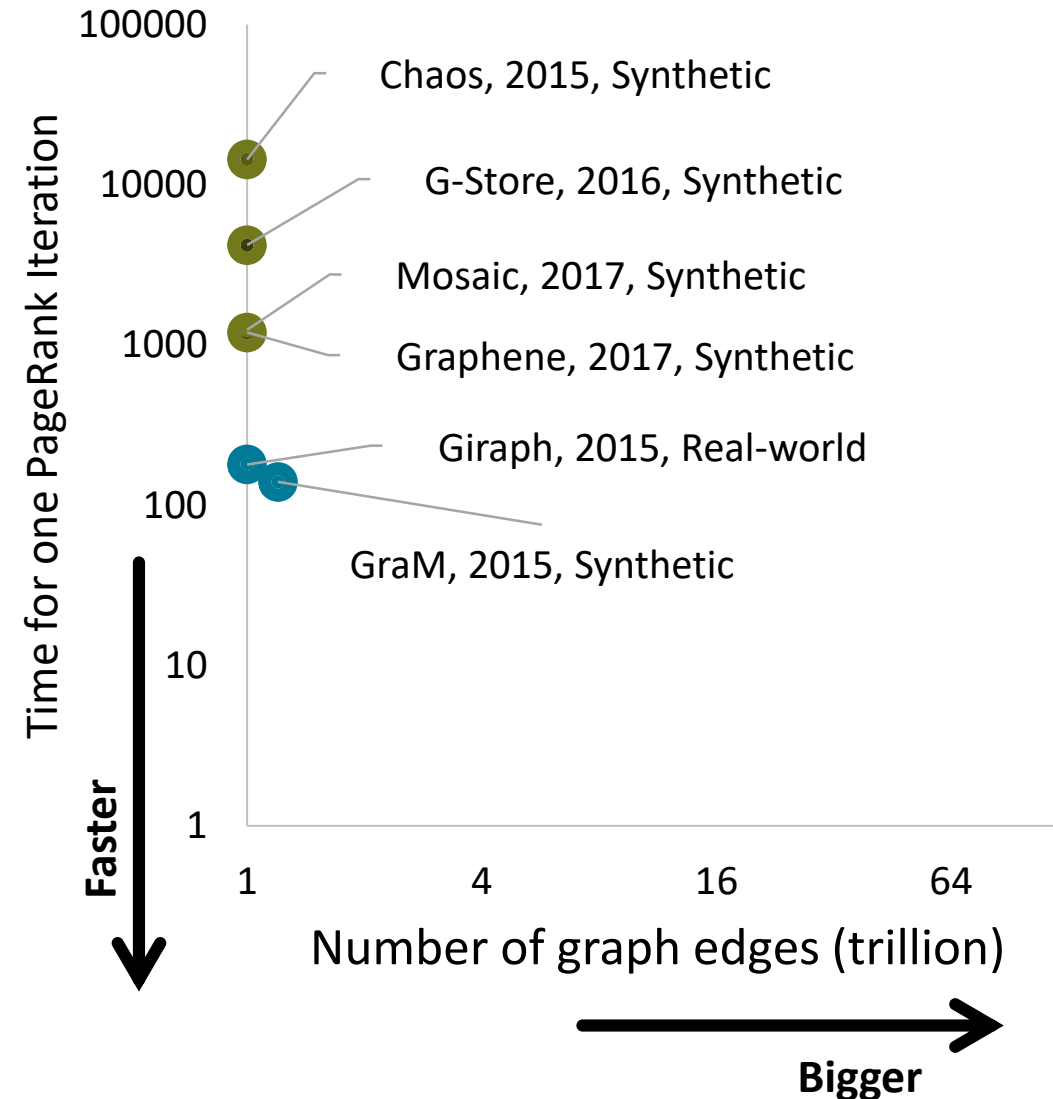
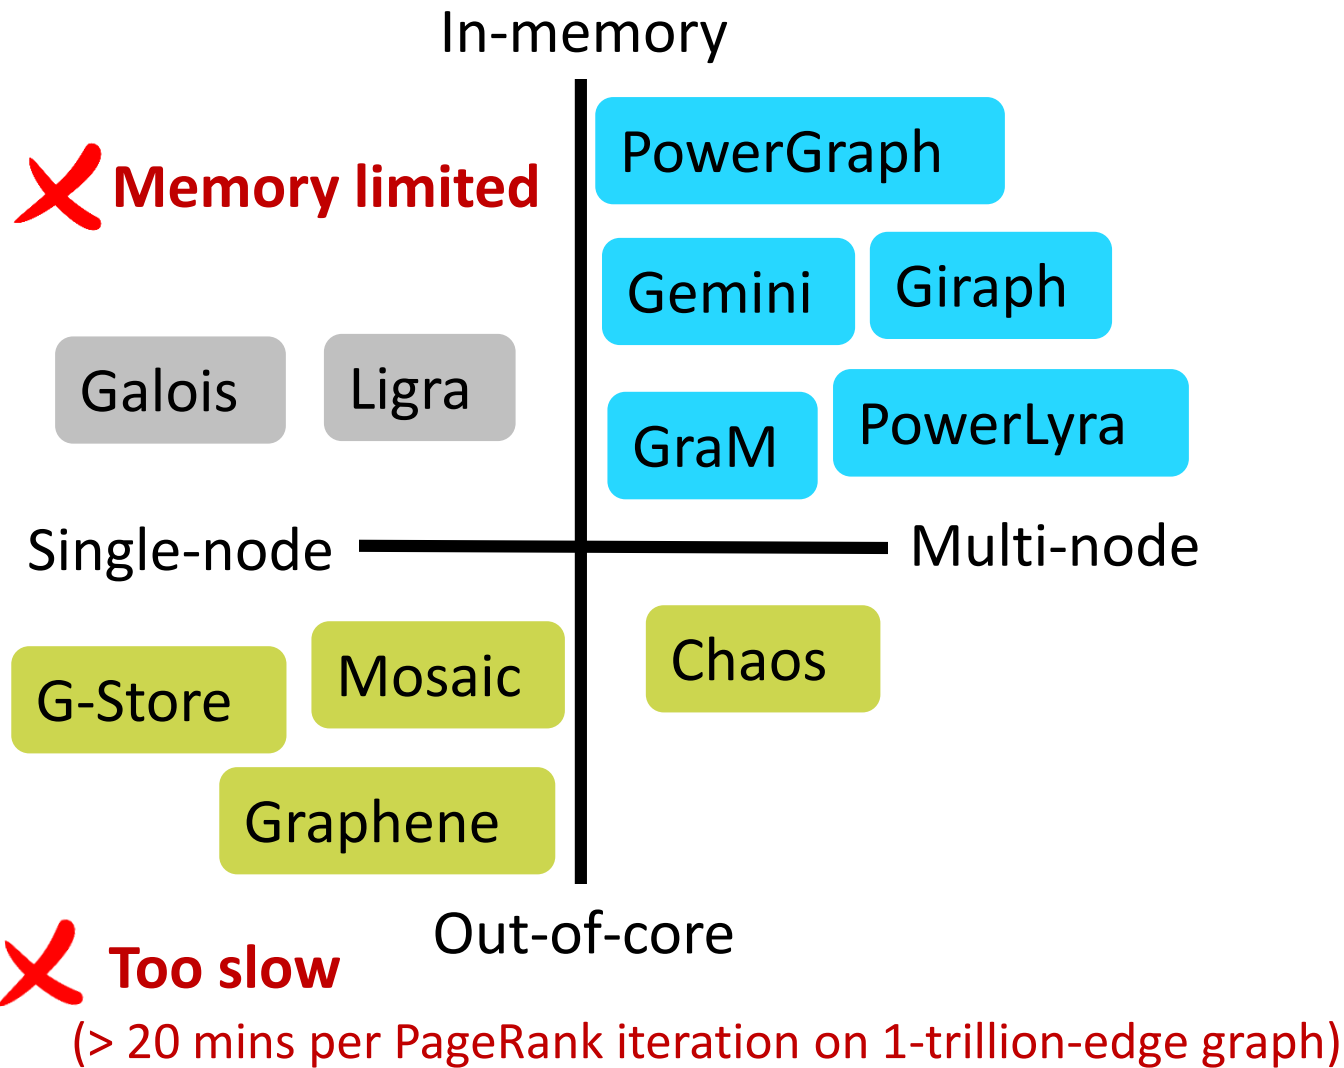


X Memory limited


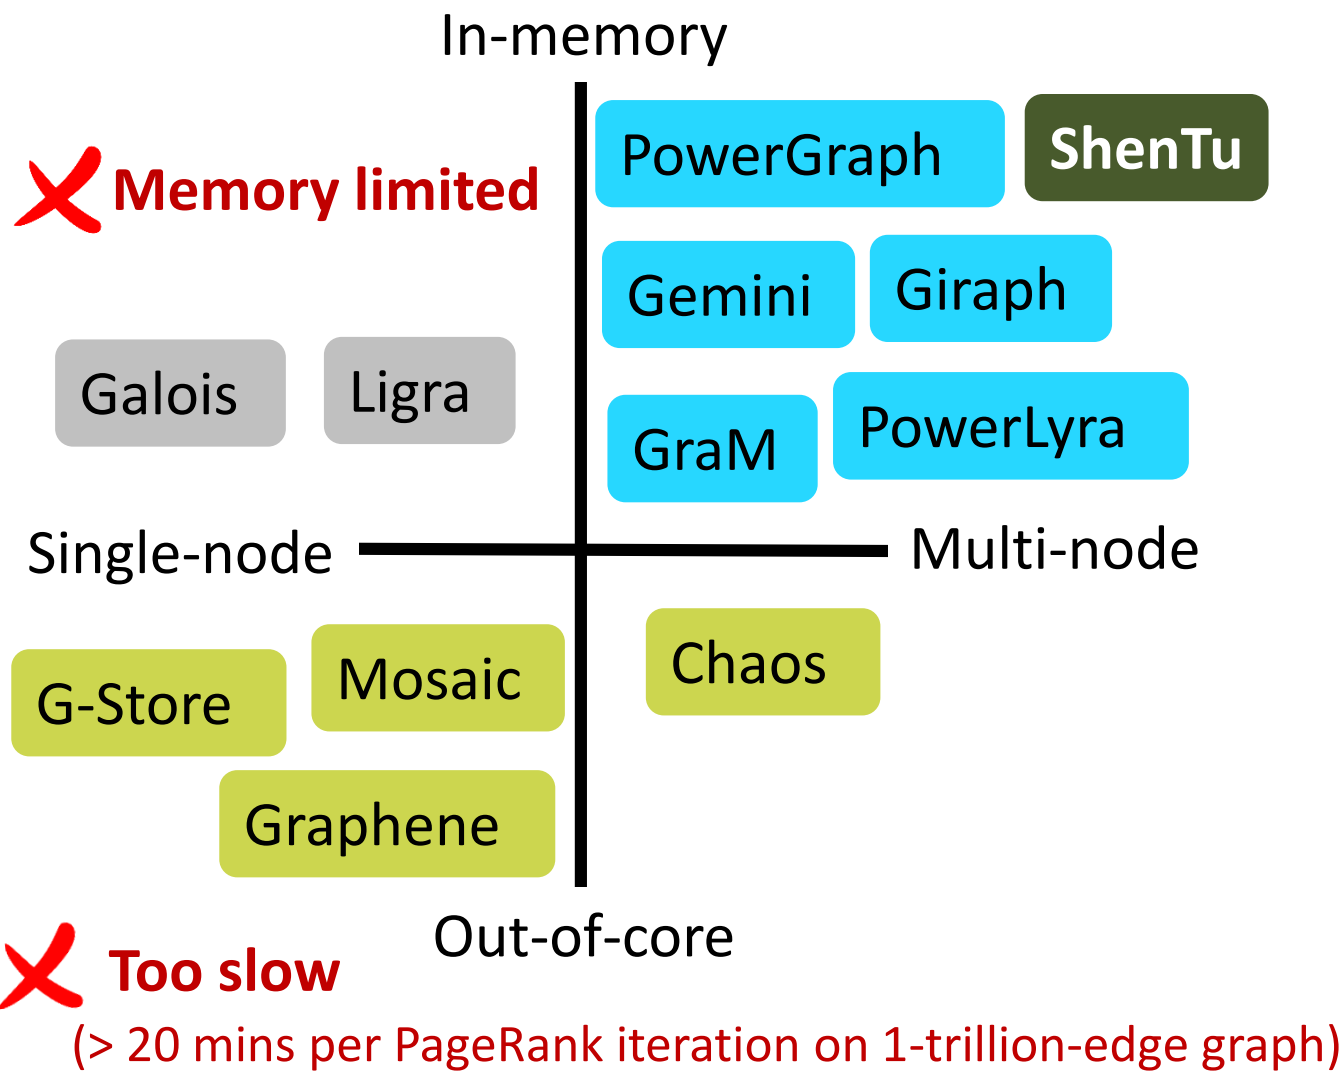
X Too slow

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


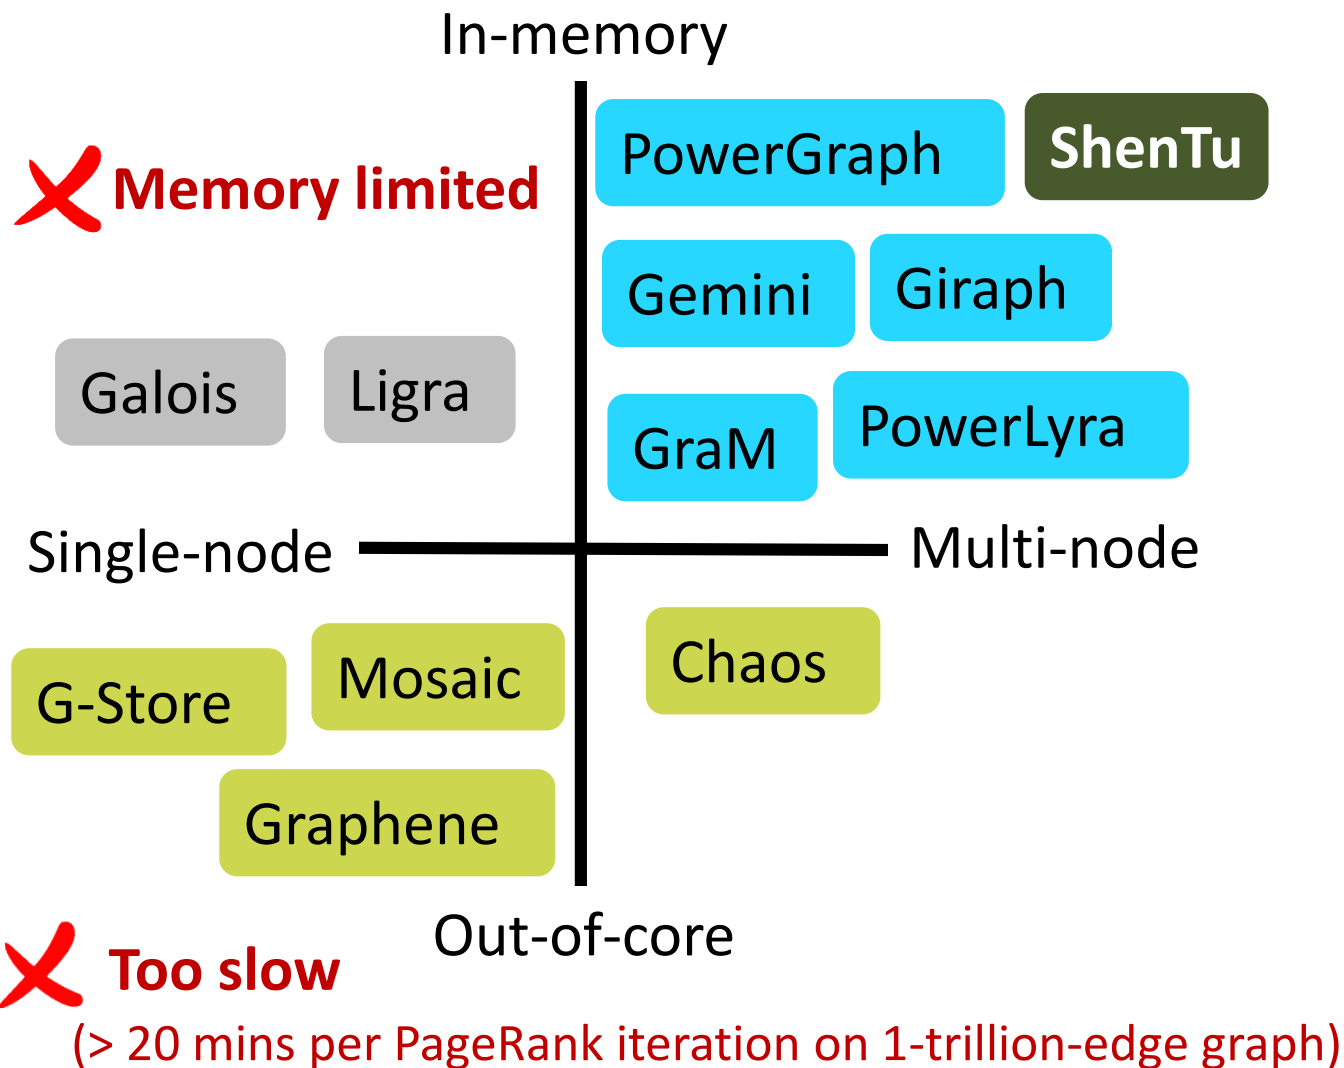
Practice of Extreme-Scale Graph Processing



Practice of Extreme-Scale Graph Processing



Practice of Extreme-Scale Graph Processing



How large are extreme-scale graphs today?

How large are extreme-scale graphs today?

$1 = 10^0$

How large are extreme-scale graphs today?

$$\begin{aligned} 1 &= 10^0 \\ 10 &= 10^1 \end{aligned}$$

How large are extreme-scale graphs today?

$$\begin{aligned} 1 &= 10^0 \\ 10 &= 10^1 \\ 100 &= 10^2 \end{aligned}$$

How large are extreme-scale graphs today?


1 = 10^0
10 = 10^1
100 = 10^2
1 000 = 10^3

How large are extreme-scale graphs today?


1 = 10^0
10 = 10^1
100 = 10^2
1 000 = 10^3
10 000 = 10^4

How large are extreme-scale graphs today?


1 = 10^0
10 = 10^1
100 = 10^2
1 000 = 10^3
10 000 = 10^4
100 000 = 10^5




How large are extreme-scale graphs today?




How large are extreme-scale graphs today?




How large are extreme-scale graphs today?



How large are extreme-scale graphs today?




How large are extreme-scale graphs today?



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

How large are extreme-scale graphs today?




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



271 billion vertices
 12 trillion edges


How large are extreme-scale graphs today?



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



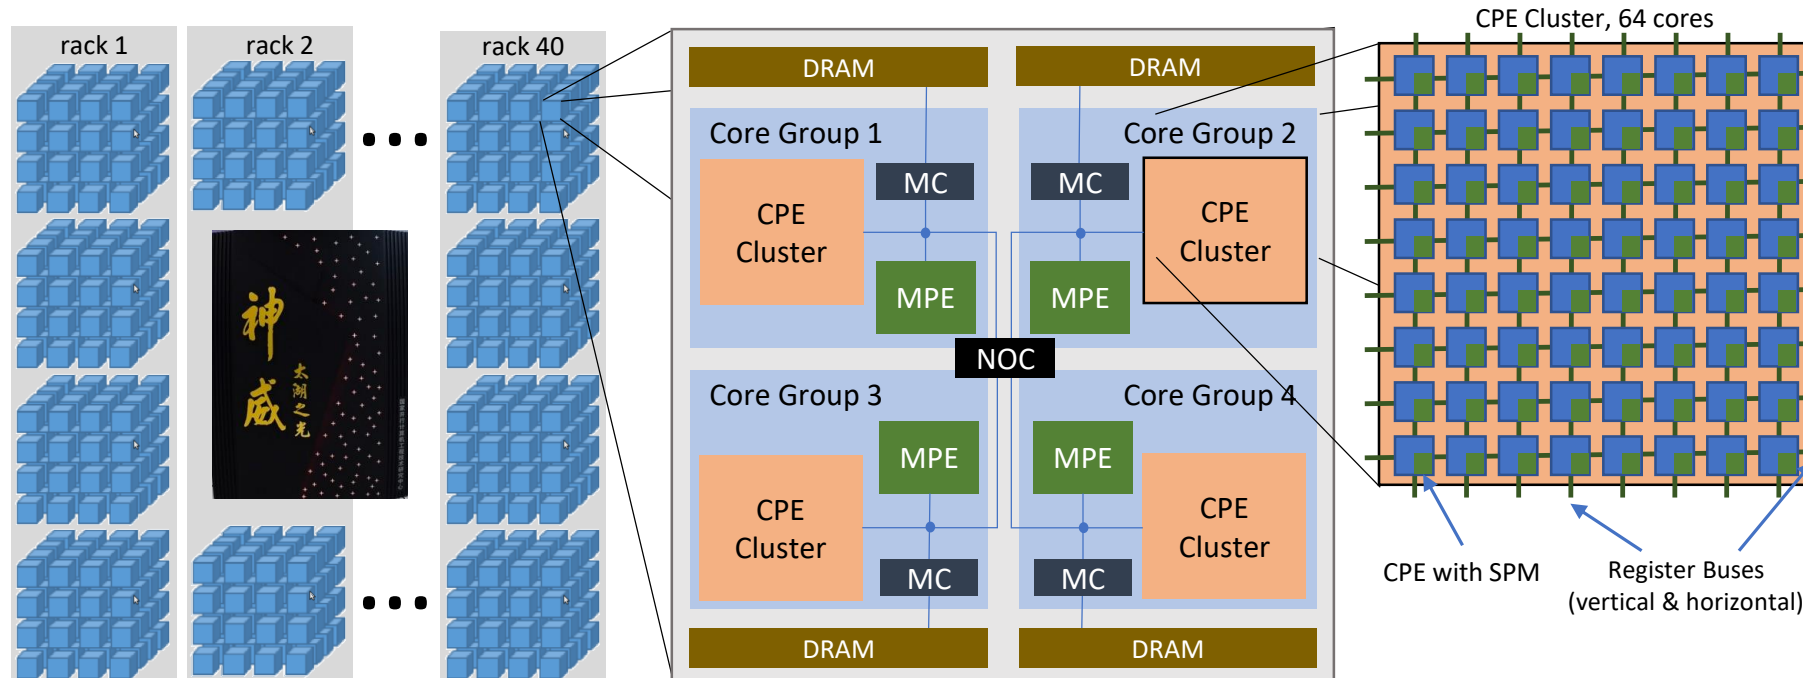
271 billion vertices
 12 trillion edges



4.4 trillion vertices
 70 trillion edges

Sunway TaihuLight

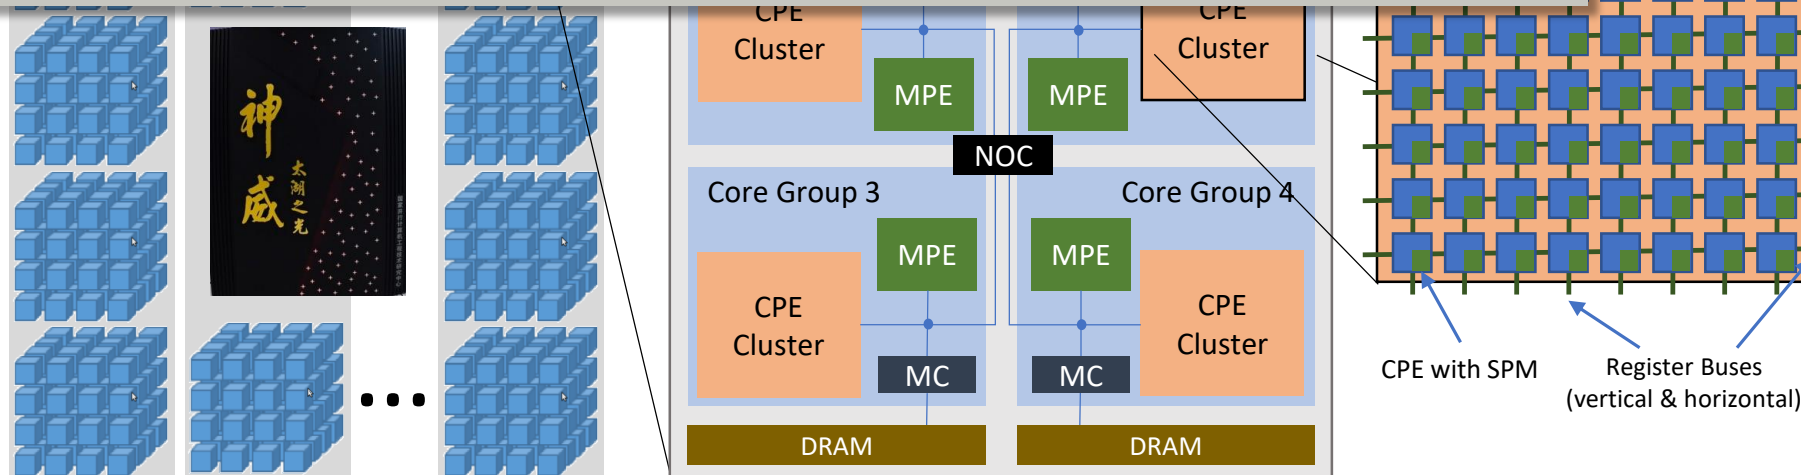
Sunway TaihuLight



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

Sunway TaihuLight

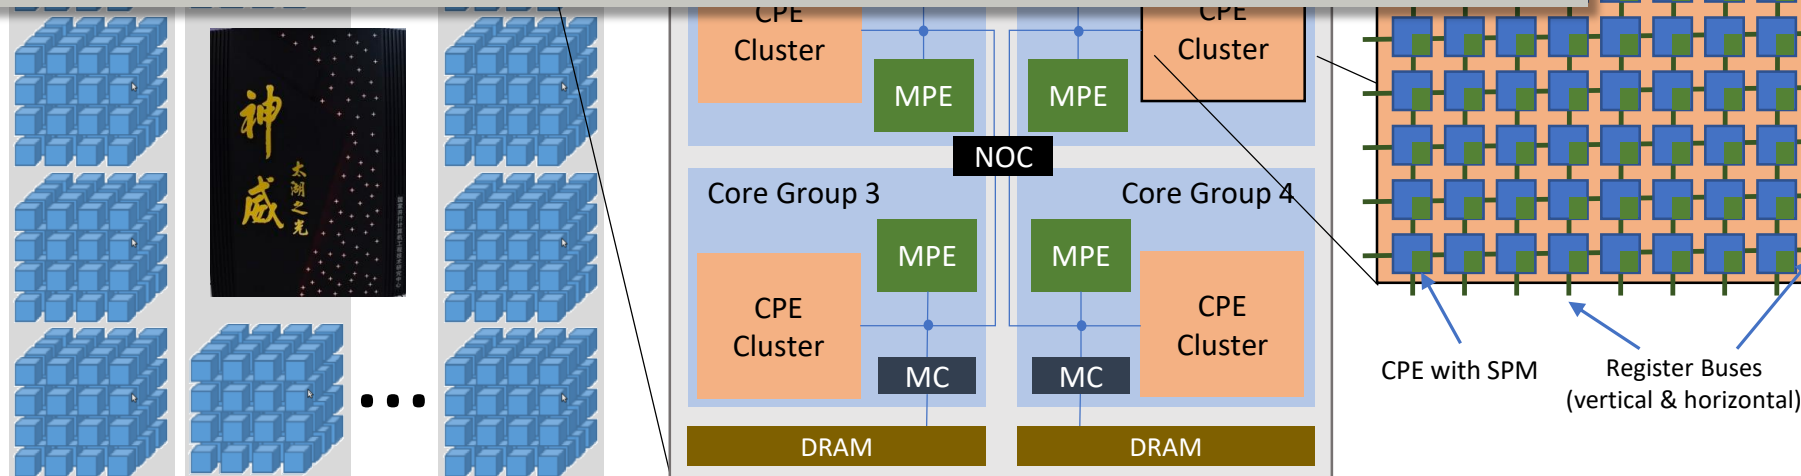
- 1/8 EFLOPS peak performance
- 1.3 PB main memory, with 5,591 TB/s bandwidth
- 70 TB/s network bisection bandwidth
- Reliable external data access, 288 GB/s IO bandwidth



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

Sunway TaihuLight

- 1/8 EFLOPS peak performance
- 1.3 PB main memory, with 5,591 TB/s bandwidth
- 70 TB/s network bisection bandwidth
- Reliable external data access, 288 GB/s IO bandwidth

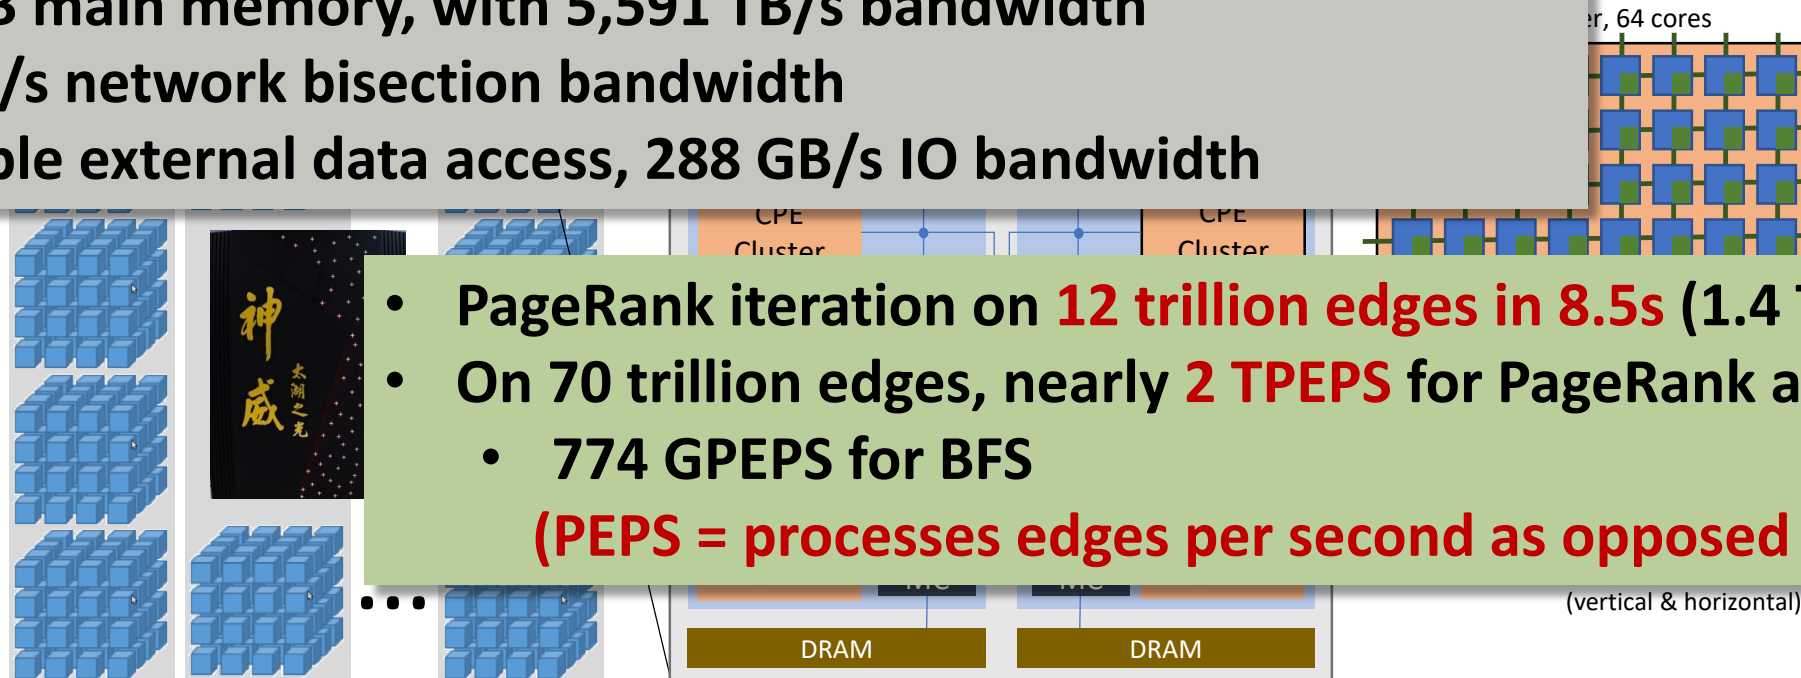


TaihuLight Top500 ranking: #3 (2018 Nov), #1 (2016, 2017)

- Handling of huge number of messages among 40,960 nodes
- Complex workload to map to its heterogeneous processing units
- Irregular data flow to be scheduled in regular accelerator cores

Sunway TaihuLight

- 1/8 EFLOPS peak performance
- 1.3 PB main memory, with 5,591 TB/s bandwidth
- 70 TB/s network bisection bandwidth
- Reliable external data access, 288 GB/s IO bandwidth




- 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)**


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

- Handling of huge number of messages among 40,960 nodes
- Complex workload to map to its heterogeneous processing units
- Irregular data flow to be scheduled in regular accelerator cores


Problems!



Problems!



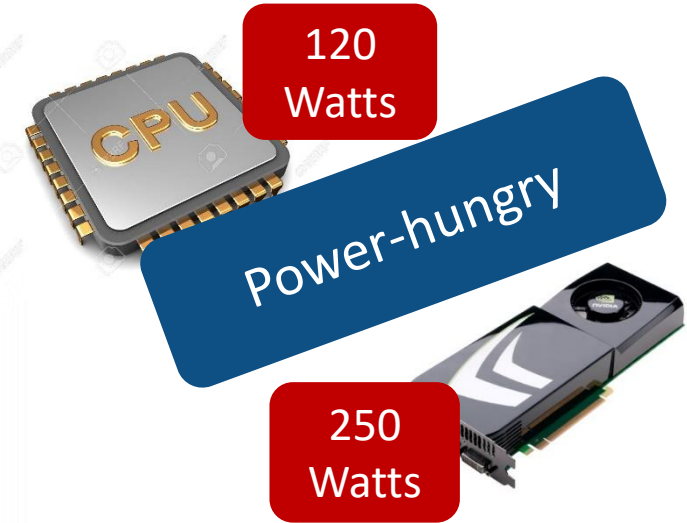
Problems!



Problems!



Problems!



120 Watts

Power-hungry

250 Watts



Problems!



Problems!



Huge

120 Watts

Power-hungry

250 Watts

Communication-heavy

Synchronization-heavy

Irregular


It's Complicated...

CPU


120 Watts

250 Watts

Problems!



Problems!



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.



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?



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

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



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?



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

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



The storage
lower bound Ω



Key idea

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?

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

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


! 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

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?



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

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

! 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

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?



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

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



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

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?



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

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



The storage **lower bound** Ω



Which one? 😊



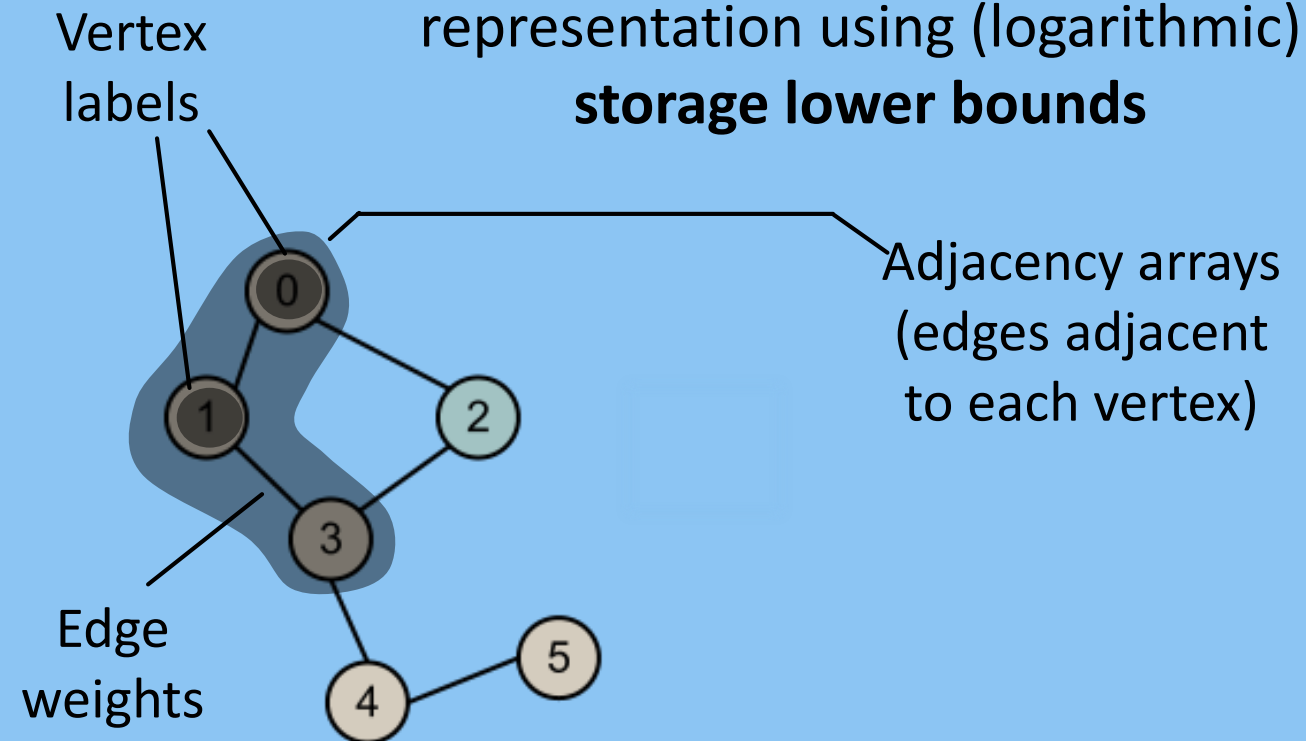
Shannon's approach **logarithmic**

(one needs at least $\log |S|$ bits to store an object from an arbitrary set S)



Key idea

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? 😊

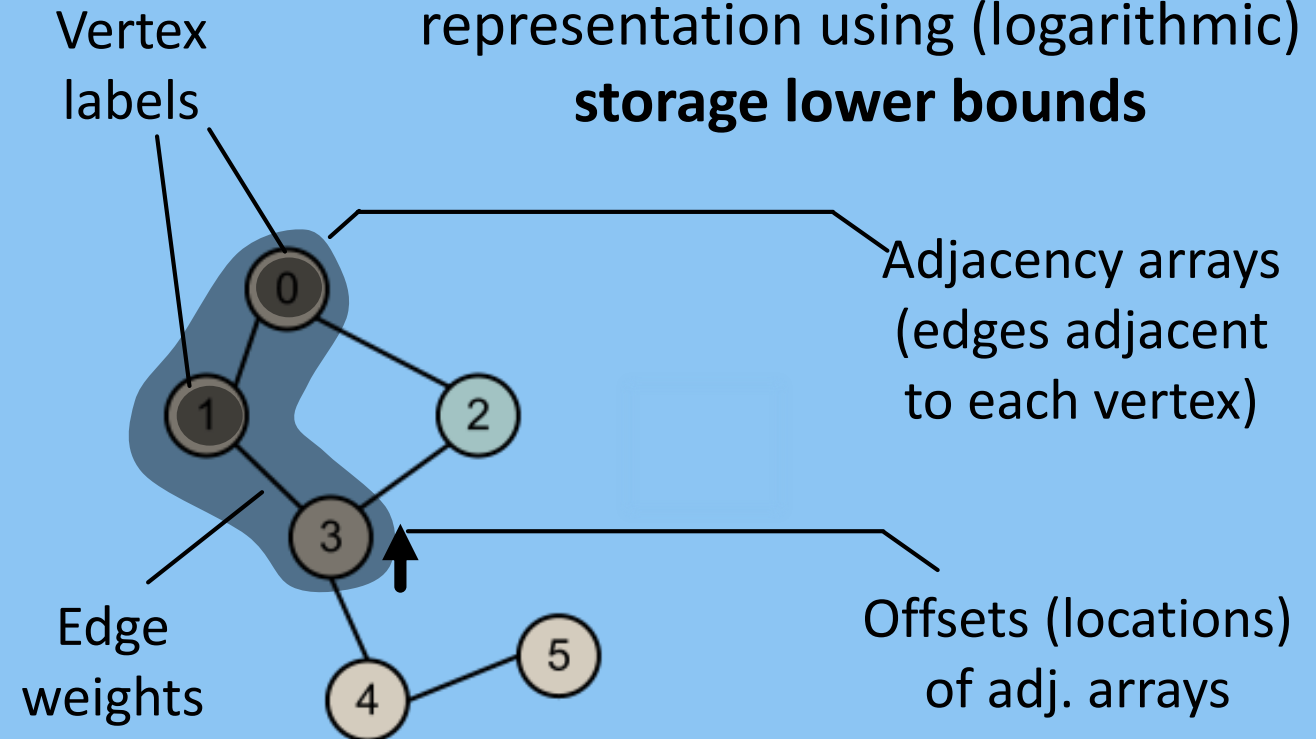
Shannon's approach **logarithmic**
(one needs at least $\log |S|$ bits to store an object from an arbitrary set S)

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

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

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

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?

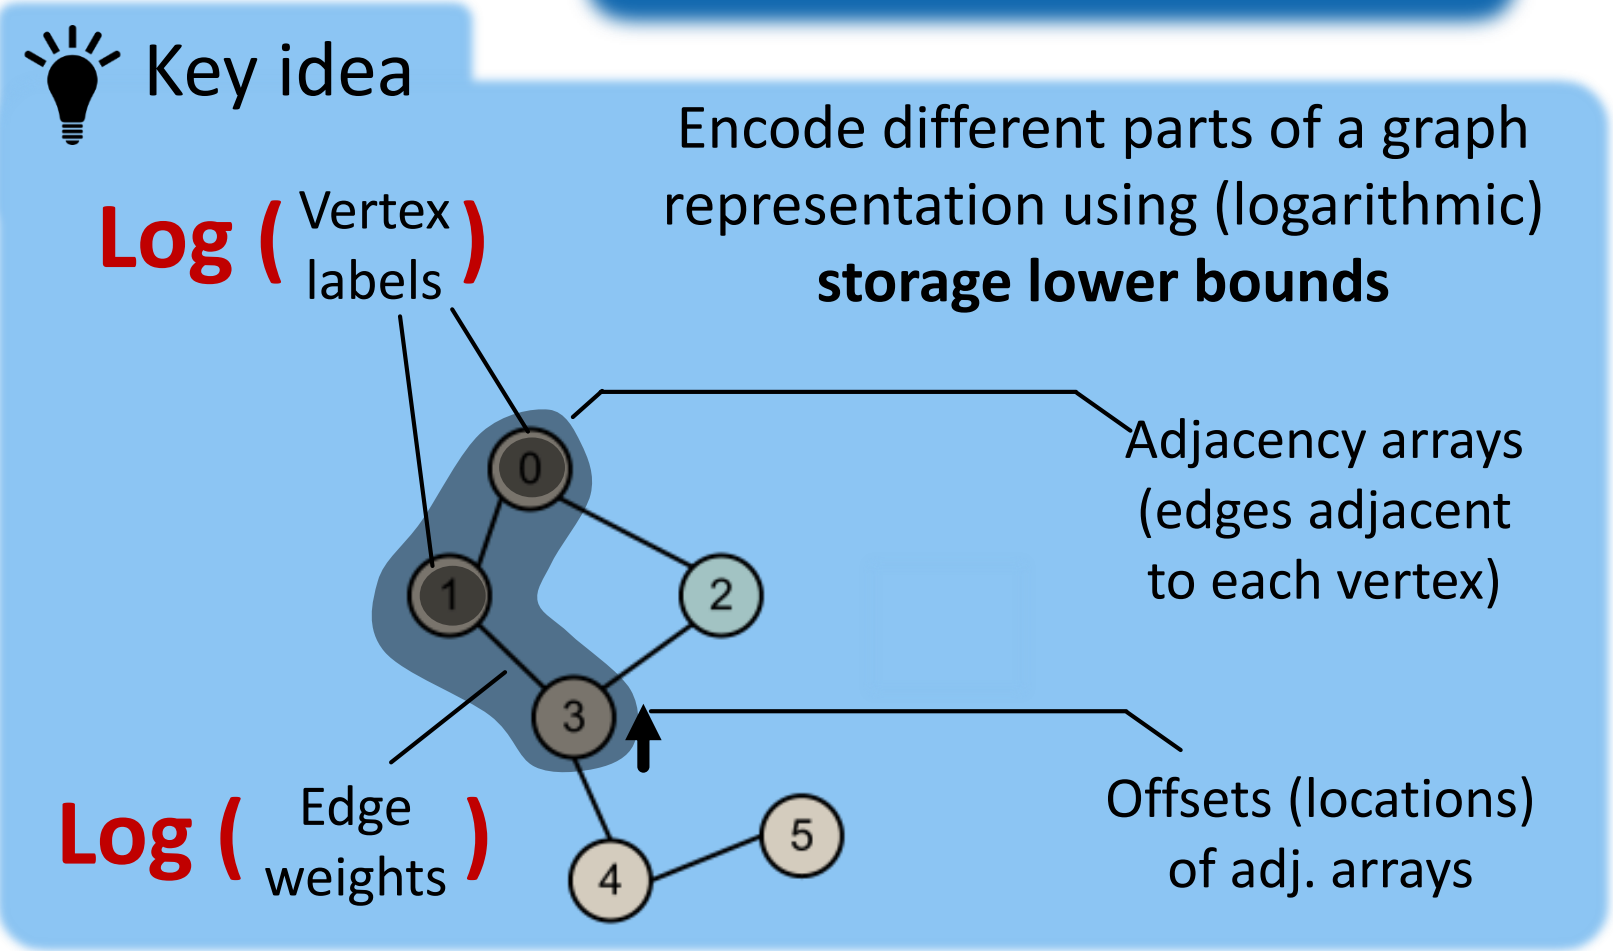
! $S = \{x_1, x_2, x_3, \dots\}$

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

! 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?



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

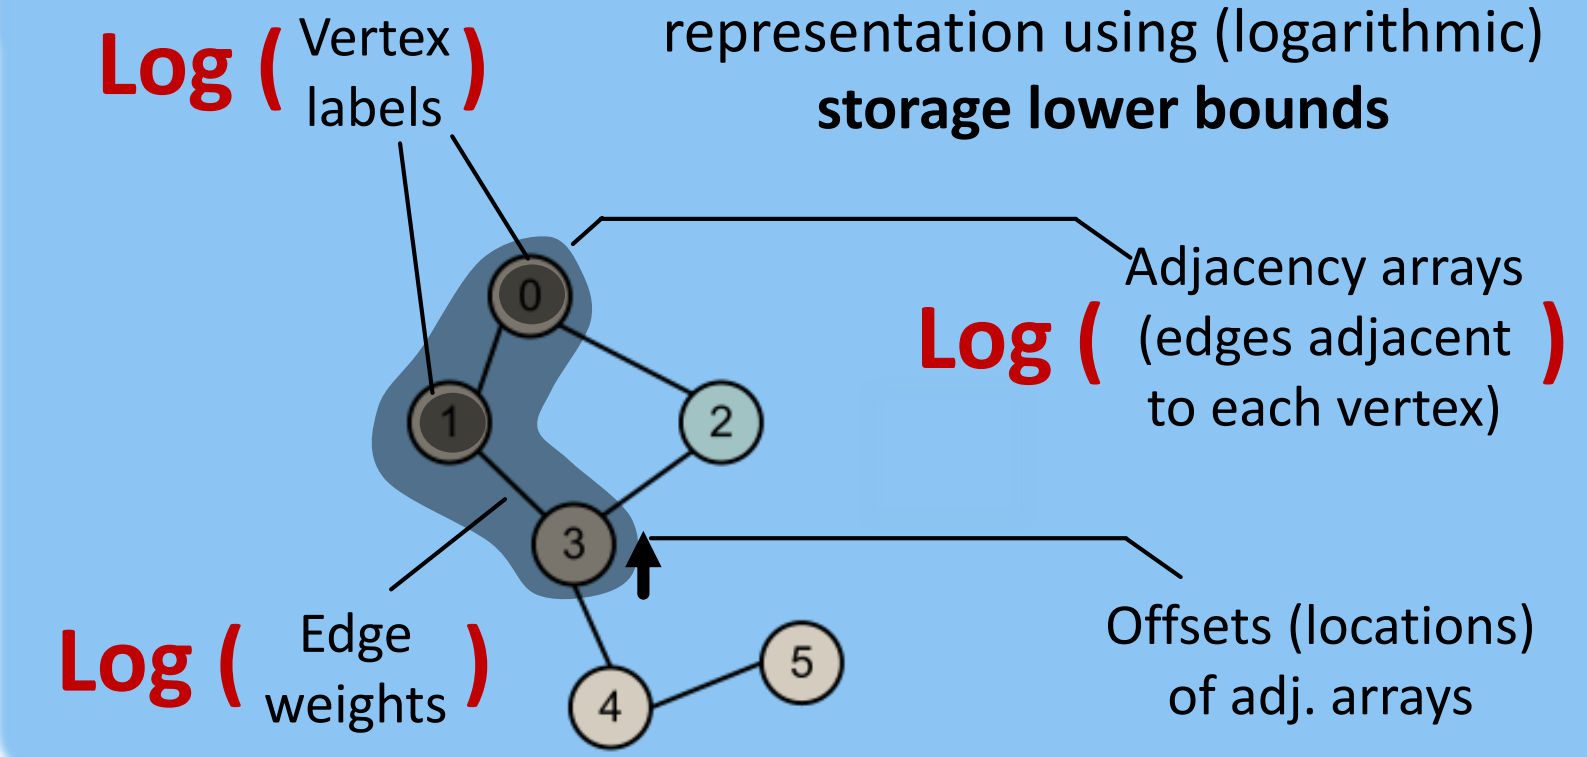
$x_1 \rightarrow 0 \dots 01$
$x_2 \rightarrow 0 \dots 10$
$x_3 \rightarrow 0 \dots 11$
...

! 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



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)

Log (Offsets (locations) of adj. arrays)




ADJACENCY ARRAY GRAPH REPRESENTATION



ADJACENCY ARRAY GRAPH REPRESENTATION

Representation



ADJACENCY ARRAY GRAPH REPRESENTATION

Representation

0


1

2

3


4

5




ADJACENCY ARRAY GRAPH REPRESENTATION

Representation




Offsets

Adjacency arrays
(vertices adjacent
to each vertex)




ADJACENCY ARRAY GRAPH REPRESENTATION

Representation




Offsets

Adjacency arrays
(vertices adjacent
to each vertex)



ADJACENCY ARRAY GRAPH REPRESENTATION


Representation



Offsets


Adjacency arrays
(vertices adjacent
to each vertex)

Physical realization



ADJACENCY ARRAY GRAPH REPRESENTATION



Representation



Offsets


Adjacency arrays
(vertices adjacent
to each vertex)

Physical realization



ADJACENCY ARRAY GRAPH REPRESENTATION


Representation




Offsets

Adjacency arrays
(vertices adjacent
to each vertex)

Physical realization




Offsets (another contiguous array)



ADJACENCY ARRAY GRAPH REPRESENTATION


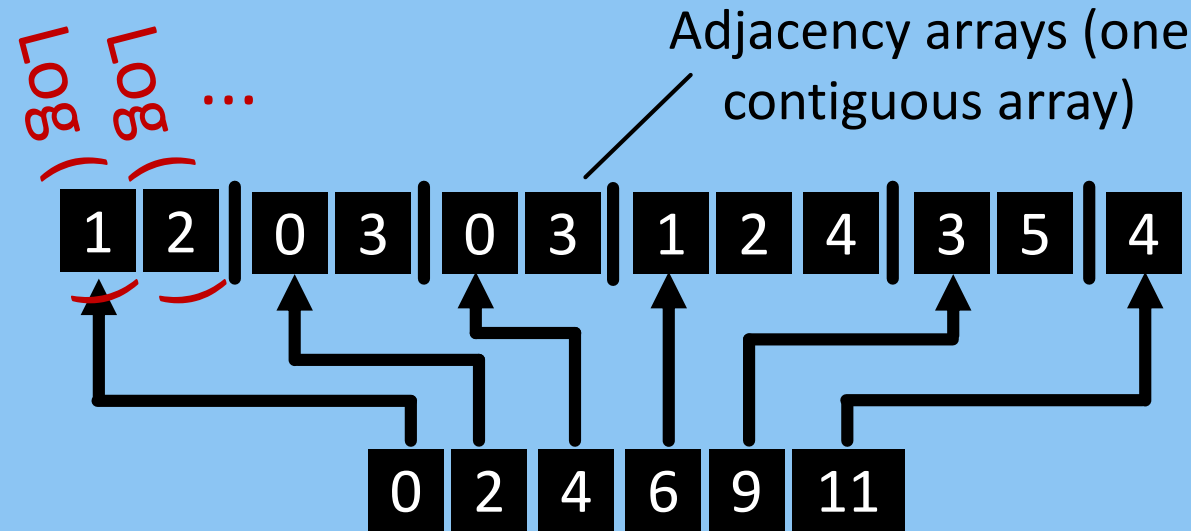
Representation



Offsets


Adjacency arrays
(vertices adjacent
to each vertex)

Physical realization



ADJACENCY ARRAY GRAPH REPRESENTATION


Representation




Offsets

Adjacency arrays
(vertices adjacent
to each vertex)

Physical realization




Offsets (another contiguous array)



ADJACENCY ARRAY GRAPH REPRESENTATION



Representation



Offsets


Adjacency arrays
(vertices adjacent
to each vertex)

Physical realization



ADJACENCY ARRAY GRAPH REPRESENTATION


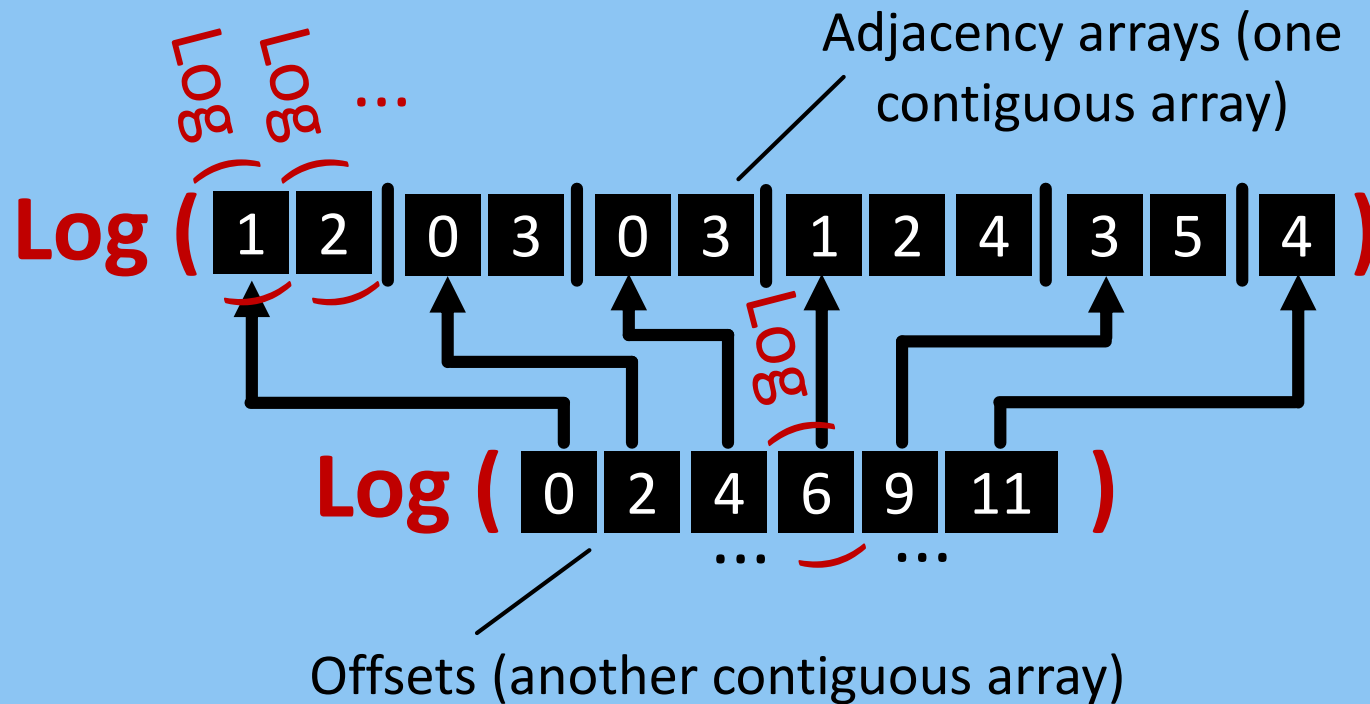
Representation




Offsets

Adjacency arrays
(vertices adjacent
to each vertex)


Physical realization




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



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



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



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

Symbols


n : #vertices,

m : #edges,

d_v : degree of vertex v ,

N_v : neighbors (adj. array) of vertex v ,

\widehat{N}_v : maximum among N_v



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

Lower bounds (global)

Symbols


n : #vertices,

m : #edges,

d_v : degree of vertex v ,

N_v : neighbors (adj. array) of vertex v ,

\widehat{N}_v : maximum among N_v



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

Lower bounds (global)
[log n]

Symbols


n : #vertices,

m : #edges,

d_v : degree of vertex v ,

N_v : neighbors (adj. array) of vertex v ,

\widehat{N}_v : maximum among N_v



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


Lower bounds (global)
[log n]

This is it?
Not really 😊



Symbols

n : #vertices,
 m : #edges,
 d_v : degree of vertex v ,
 N_v : neighbors (adj. array) of
vertex v ,
 \widehat{N}_v : maximum among N_v



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


Lower bounds (global)
[log n]

This is it?
Not really 😊

Lower bounds (local)

Symbols

n : #vertices,
 m : #edges,
 d_v : degree of vertex v ,
 N_v : neighbors (adj. array) of vertex v ,
 \widehat{N}_v : maximum among N_v



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

Lower bounds (global)
[log n]


This is it?
Not really 😊

Lower bounds (local)

Assume:

Symbols

n : #vertices,
 m : #edges,
 d_v : degree of vertex v ,
 N_v : neighbors (adj. array) of vertex v ,
 \widehat{N}_v : maximum among N_v



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

Lower bounds (global)
[log n]

This is it?
Not really 😊


Lower bounds (local)

Assume:

- a graph, e.g., $V = \{1, \dots, 2^{22}\}$

Symbols

n : #vertices,
 m : #edges,
 d_v : degree of vertex v ,
 N_v : neighbors (adj. array) of vertex v ,
 \widehat{N}_v : maximum among N_v



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


Lower bounds (global)
[log n]

This is it?
Not really 😊



Symbols

n : #vertices,
 m : #edges,
 d_v : degree of vertex v ,
 N_v : neighbors (adj. array) of vertex v ,
 \widehat{N}_v : maximum among N_v



Lower bounds (local)

Assume:

- a graph, e.g., $V = \{1, \dots, 2^{22}\}$
- A vertex v with few neighbors: $d_v \ll n$

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


Lower bounds (global)
[log n]

This is it?
Not really 😊



Symbols

n : #vertices,
 m : #edges,
 d_v : degree of vertex v ,
 N_v : neighbors (adj. array) of vertex v ,
 \widehat{N}_v : maximum among N_v



Lower bounds (local)

Assume:

- a graph, e.g., $V = \{1, \dots, 2^{22}\}$
- A vertex v with few neighbors: $d_v \ll n$
- ...all these neighbors have small labels: $\widehat{N}_v \ll n$

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


Lower bounds (global)
 $\lceil \log n \rceil$

This is it?
 Not really 😊



Symbols

n : #vertices,
 m : #edges,
 d_v : degree of vertex v ,
 N_v : neighbors (adj. array) of vertex v ,
 \widehat{N}_v : maximum among N_v



Lower bounds (local)

Assume:

- a graph, e.g., $V = \{1, \dots, 2^{22}\}$
- A vertex v with few neighbors: $d_v \ll n$
- ...all these neighbors have small labels: $\widehat{N}_v \ll n$




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

Lower bounds (global)
 $\lceil \log n \rceil$

This is it?
 Not really 😊

Symbols

- n : #vertices,
- m : #edges,
- d_v : degree of vertex v ,
- N_v : neighbors (adj. array) of vertex v ,
- \widehat{N}_v : maximum among N_v



Lower bounds (local)

Assume:

- a graph, e.g., $V = \{1, \dots, 2^{22}\}$
- A vertex v with few neighbors: $d_v \ll n$
- ...all these neighbors have small labels: $\widehat{N}_v \ll n$

$$\lceil \log 2^{22} \rceil = 22$$



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

Lower bounds (global)
 $\lceil \log n \rceil$

This is it?
 Not really 😊

Symbols

- n : #vertices,
- m : #edges,
- d_v : degree of vertex v ,
- N_v : neighbors (adj. array) of vertex v ,
- \widehat{N}_v : maximum among N_v


Lower bounds (local)

Assume:

- a graph, e.g., $V = \{1, \dots, 2^{22}\}$
- A vertex v with few neighbors: $d_v \ll n$
- ...all these neighbors have small labels: $\widehat{N}_v \ll n$



$$\lceil \log 2^{22} \rceil = 22$$




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

Lower bounds (global)
 $\lceil \log n \rceil$

This is it?
 Not really 😊

Symbols

- n : #vertices,
- m : #edges,
- d_v : degree of vertex v ,
- N_v : neighbors (adj. array) of vertex v ,
- \widehat{N}_v : maximum among N_v




Lower bounds (local)

Assume:

- a graph, e.g., $V = \{1, \dots, 2^{22}\}$
- A vertex v with few neighbors: $d_v \ll n$
- ...all these neighbors have small labels: $\widehat{N}_v \ll n$



$$\lceil \log 2^{22} \rceil = 22$$




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

Lower bounds (global)
 $\lceil \log n \rceil$

This is it?
 Not really 😊

Symbols

- n : #vertices,
- m : #edges,
- d_v : degree of vertex v ,
- N_v : neighbors (adj. array) of vertex v ,
- \widehat{N}_v : maximum among N_v




Lower bounds (local)

Assume:

- a graph, e.g., $V = \{1, \dots, 2^{22}\}$
- A vertex v with few neighbors: $d_v \ll n$
- ...all these neighbors have small labels: $\widehat{N}_v \ll n$



$$\lceil \log 2^{22} \rceil = 22$$




Thus, use the local bound $\lceil \log \widehat{N}_v \rceil$

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

This is it?
Not really 😊


Symbols

n : #vertices,
 m : #edges,
 d_v : degree of vertex v ,
 N_v : neighbors (adj. array) of vertex v ,
 \widehat{N}_v : maximum among N_v



Lower bounds (local): problem

- a graph, e.g., $V = \{1, \dots, 2^{22}\}$
- A vertex v with few neighbors: $d_v \ll n$
- ...all these neighbors have small labels: $\widehat{N}_v \ll n$




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

This is it?
Not really 😊

Symbols


n : #vertices,
 m : #edges,
 d_v : degree of vertex v ,
 N_v : neighbors (adj. array) of vertex v ,
 \widehat{N}_v : maximum among N_v



Lower bounds (local): problem

What if:

- a graph, e.g., $V = \{1, \dots, 2^{22}\}$
- A vertex v with few neighbors: $d_v \ll n$
- ...all these neighbors have small labels: $\widehat{N}_v \ll n$




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

This is it?
Not really 😊

Symbols


- n : #vertices,
- m : #edges,
- d_v : degree of vertex v ,
- N_v : neighbors (adj. array) of vertex v ,
- \widehat{N}_v : maximum among N_v



Lower bounds (local): problem

What if:

- a graph, e.g., $V = \{1, \dots, 2^{22}\}$
- A vertex v with few neighbors: $d_v \ll n$
- ...all these neighbors have small labels: $\widehat{N}_v \ll n$
- ...one neighbor has a large ID:




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

This is it?
 Not really 😊

Symbols


- n : #vertices,
- m : #edges,
- d_v : degree of vertex v ,
- N_v : neighbors (adj. array) of vertex v ,
- \widehat{N}_v : maximum among N_v



Lower bounds (local): problem

What if:

- a graph, e.g., $V = \{1, \dots, 2^{22}\}$
- A vertex v with few neighbors: $d_v \ll n$
- ...all these neighbors have small labels: $\widehat{N}_v \ll n$
- ...one neighbor has a large ID:



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

This is it?
 Not really 😊

Symbols

- n : #vertices,
- m : #edges,
- d_v : degree of vertex v ,
- N_v : neighbors (adj. array) of vertex v ,
- \widehat{N}_v : maximum among N_v


Lower bounds (local): problem

What if:

- a graph, e.g., $V = \{1, \dots, 2^{22}\}$
- A vertex v with few neighbors: $d_v \ll n$
- ...all these neighbors have small labels: $\widehat{N}_v \ll n$
- ...one neighbor has a large ID:



$$\lceil \log 2^{20} \rceil = 20$$




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

This is it?
Not really 😊

Symbols

- n : #vertices,
- m : #edges,
- d_v : degree of vertex v ,
- N_v : neighbors (adj. array) of vertex v ,
- \widehat{N}_v : maximum among N_v



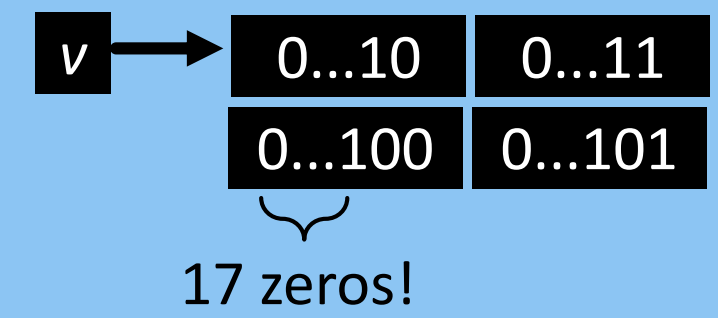
Lower bounds (local): problem

What if:

- a graph, e.g., $V = \{1, \dots, 2^{22}\}$
- A vertex v with few neighbors: $d_v \ll n$
- ...all these neighbors have small labels: $\widehat{N}_v \ll n$
- ...one neighbor has a large ID:




$$\lceil \log 2^{20} \rceil = 20$$




2 Log (Offset structure)

2 **Log** (Offset structure)




2 Log (Offset structure)

Use a **bit vector** instead of an array of offsets...



2 Log (Offset structure)


Use a **bit vector** instead of an array of offsets...




Bit vectors instead of offset arrays

2 Log (Offset structure)

Use a **bit vector** instead of an array of offsets...




Bit vectors instead of offset arrays




2 Log (Offset structure)

Use a **bit vector** instead of an array of offsets...




Bit vectors instead of offset arrays




2 Log (Offset structure)

Use a **bit vector** instead of an array of offsets...




Bit vectors instead of offset arrays




2 Log (Offset structure)

Use a **bit vector** instead of an array of offsets... 




Bit vectors instead of offset arrays




i -th set bit has a position $x \rightarrow$
 the adjacency array of a vertex i
 starts at a word x

2 Log (Offset structure)

Use a **bit vector** instead of an array of offsets...



Bit vectors instead of offset arrays




How many 1s are set before a given *i*-th bit?

i-th set bit has a position *x* → the adjacency array of a vertex *i* starts at a word *x*

2 Log (Offset structure)


...Encode the resulting bit vectors as succinct bit vectors [1]



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

2 Log (Offset structure)

...Encode the resulting bit vectors as
succinct bit vectors [1]




Succinct bit vectors

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

2 Log (Offset structure)

...Encode the resulting bit vectors as
succinct bit vectors [1]



Succinct bit vectors

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

2 Log (Offset structure)


...Encode the resulting bit vectors as
succinct bit vectors [1]



Succinct bit vectors

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


= small + fast
(hopefully)



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

2 Log (Offset structure)

...Encode the resulting bit vectors as **succinct bit vectors [1]** !



Succinct bit vectors

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


= small + fast (hopefully) !

101010100101000101010111110000001100001...

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

2 Log (Offset structure)

...Encode the resulting bit vectors as **succinct bit vectors [1]** !



Succinct bit vectors

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


= small + fast (hopefully) !

n bits **101010100101000101010111110000001100001...**

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


2 Log (Offset structure)

...Encode the resulting bit vectors as succinct bit vectors [1] !



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


= small + fast (hopefully) !



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

2 Log (Offset structure)


...Encode the resulting bit vectors as succinct bit vectors [1] !



Succinct bit vectors

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


= small + fast (hopefully) !



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

2 Log (Offset structure)


...Encode the resulting bit vectors as succinct bit vectors [1] !



Succinct bit vectors

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


= small + fast (hopefully) !



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

2 Log (Offset structure)


...Encode the resulting bit vectors as succinct bit vectors [1] !



Succinct bit vectors

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


= small + fast (hopefully) !



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

2 Log (Offset structure)


...Encode the resulting bit vectors as succinct bit vectors [1] !



Succinct bit vectors

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


= small + fast (hopefully) !



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

2 Log (Offset structure)


...Encode the resulting bit vectors as succinct bit vectors [1] !



Succinct bit vectors

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


= small + fast (hopefully) !



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

2 Log (Offset structure)


...Encode the resulting bit vectors as succinct bit vectors [1]



Succinct bit vectors

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


= small + fast (hopefully)



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

2 Log (Offset structure)


...Encode the resulting bit vectors as succinct bit vectors [1]



Succinct bit vectors

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


= small + fast (hopefully)



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

2 Log (Offset structure)

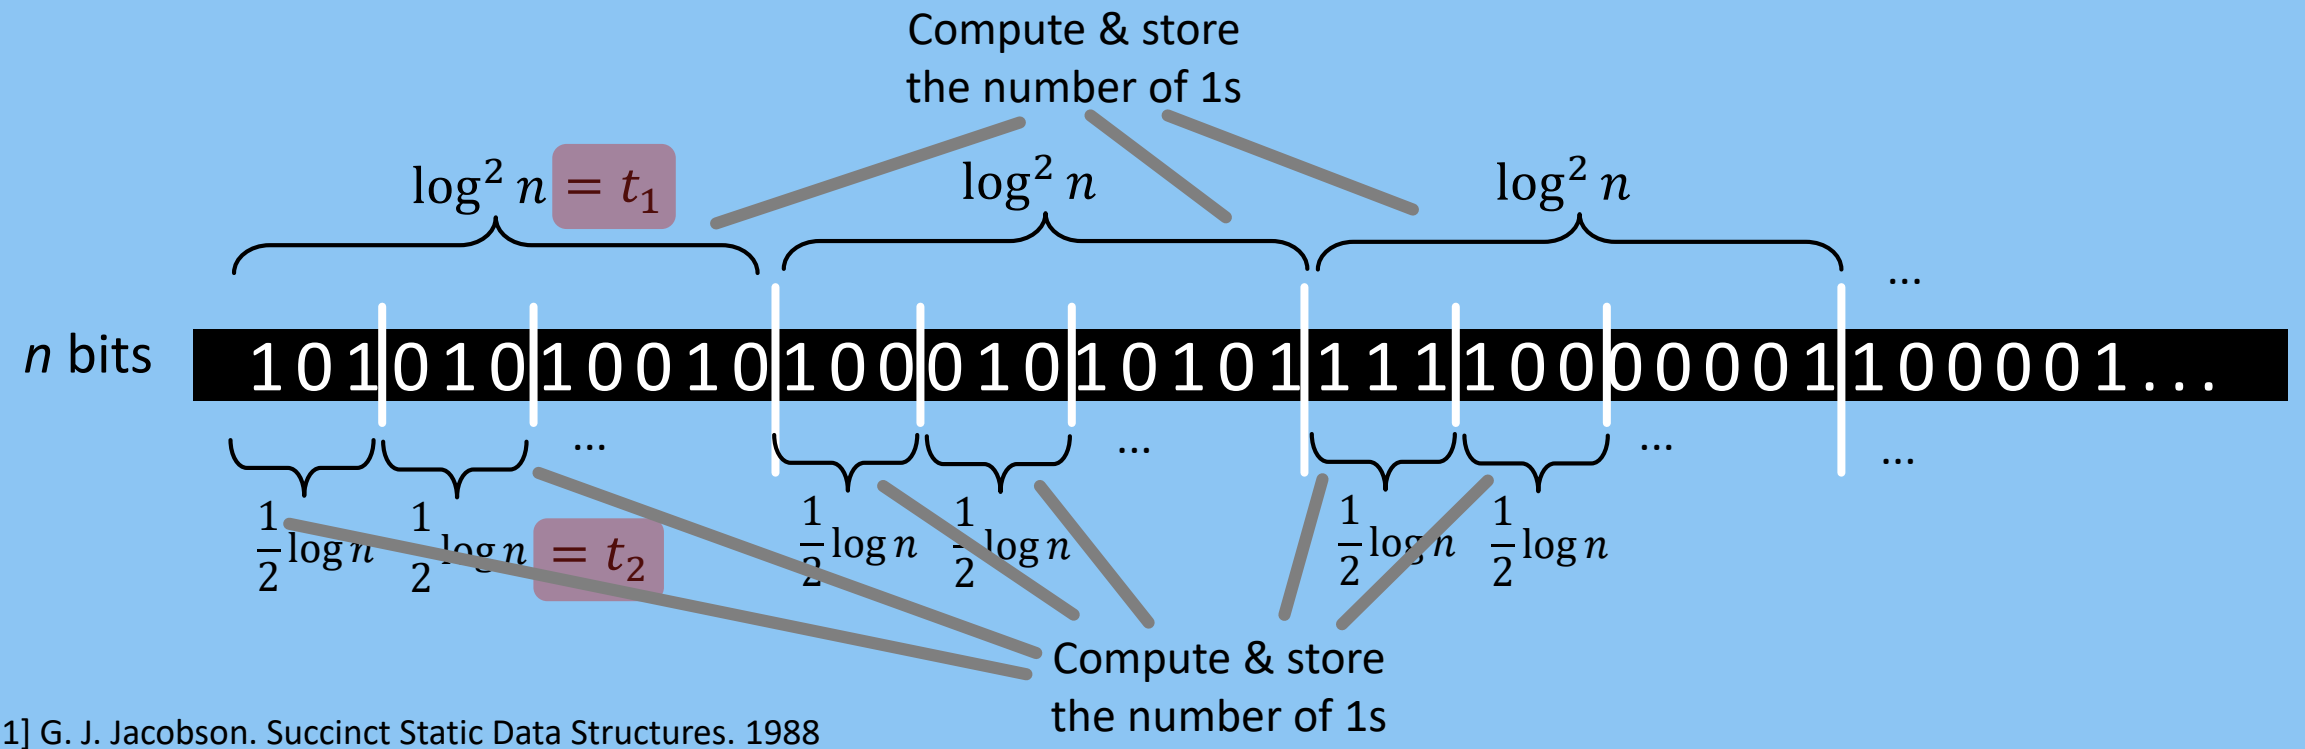
...Encode the resulting bit vectors as succinct bit vectors [1] !



Succinct bit vectors

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


= small + fast (hopefully) !



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

2 Log (Offset structure)

...Encode the resulting bit vectors as succinct bit vectors [1]




Succinct bit vectors

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

= small + fast (hopefully)


Compute & store the number of 1s $= O\left(\frac{n}{t_1} \log n\right) = O\left(\frac{n}{\log n}\right) = o(n)$



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

2 Log (Offset structure)

...Encode the resulting bit vectors as succinct bit vectors [1]




Succinct bit vectors

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

= small + fast (hopefully)

Compute & store the number of 1s $= O\left(\frac{n}{t_1} \log n\right) = O\left(\frac{n}{\log n}\right) = o(n)$




Compute & store the number of 1s $= O\left(\frac{n}{t_2} \log t_1\right) = O\left(\frac{n \log \log n}{\log n}\right) = o(n)$

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

2 Log (Offset structure)

...Encode the resulting bit vectors as succinct bit vectors [1]



Succinct bit vectors

They use $[Q] + o(Q)$ bits ($[Q]$ - lower bound), they answer various queries in $o(Q)$ time. = small + fast (hopefully)

Compute & store the number of 1s $= O\left(\frac{n}{t_1} \log n\right) = O\left(\frac{n}{\log n}\right) = o(n)$

$\log^2 n = t_1$

n bits

1 0 1 0 1 0 1 0 0 1 0 1 0 0 0 1 0 1 1 1 1 1 0 0 0 0 0 0 1 1 0 0 0 0 1 ...


$\frac{1}{2} \log n = t_2$

Compute & store the number of 1s $= O\left(\frac{n}{t_2} \log t_1\right) = O\left(\frac{n \log \log n}{\log n}\right) = o(n)$

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

2 Log (Offset structure)

...Encode the resulting bit vectors as succinct bit vectors [1]




Succinct bit vectors

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

= small + fast (hopefully)

Total storage:
 $n + o(n) + o(n) + \dots$
 $= n + o(n)$

Compute & store the number of 1s $= O\left(\frac{n}{t_1} \log n\right) = O\left(\frac{n}{\log n}\right) = o(n)$




Compute & store the number of 1s $= O\left(\frac{n}{t_2} \log t_1\right) = O\left(\frac{n \log \log n}{\log n}\right) = o(n)$

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

2 Log (Offset structure)

...Encode the resulting bit vectors as succinct bit vectors [1] !



Succinct bit vectors

They use $[Q] + o(Q)$ bits ($[Q]$ - lower bound), they answer various queries in $o(Q)$ time. = small + fast (hopefully) !

Total storage:
 $n + o(n) + o(n) + \dots$
 $= n + o(n)$!

Compute & store the number of 1s $= O\left(\frac{n}{t_1} \log n\right) = O\left(\frac{n}{\log n}\right) = o(n)$


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



Compute & store the number of 1s $= O\left(\frac{n}{t_2} \log t_1\right) = O\left(\frac{n \log \log n}{\log n}\right) = o(n)$


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

3 **Log** (Adjacency structure)



3 **Log** (Adjacency structure)

Use different relabelings




3 **Log** (Adjacency structure)

Use different relabelings



**Degree-Minimizing: Targeting general graphs
(no assumptions on graph structure)**




3 **Log** (Adjacency structure)

Use different relabelings



**Degree-Minimizing: Targeting general graphs
(no assumptions on graph structure)**



**More schemes
that assume specific
classes of graphs**

...

3 Log (Adjacency structure)


Use different relabelings



**Degree-Minimizing: Targeting general graphs
(no assumptions on graph structure)**

Permute(**2 3 4 5 1M**) = **v w x y z**

(simultaneously for all other neighborhoods)



**More schemes
that assume specific
classes of graphs**

...

3 **Log** (Adjacency structure)

Use different relabelings




Degree-Minimizing: Targeting general graphs (no assumptions on graph structure)

$$\text{Permute}(\boxed{2} \boxed{3} \boxed{4} \boxed{5} \boxed{1M}) = \boxed{v} \boxed{w} \boxed{x} \boxed{y} \boxed{z}$$

(simultaneously for all other neighborhoods)

(1) The more often a label occurs (i.e., the higher vertex degree), the smaller permuted value it receives



More schemes that assume specific classes of graphs

...

3 Log (Adjacency structure)

Use different relabelings




Degree-Minimizing: Targeting general graphs (no assumptions on graph structure)

$$\text{Permute}(\boxed{2} \boxed{3} \boxed{4} \boxed{5} \boxed{1M}) = \boxed{v} \boxed{w} \boxed{x} \boxed{y} \boxed{z}$$

(simultaneously for all other neighborhoods)

(1) The more often a label occurs (i.e., the higher vertex degree), the smaller permuted value it receives

$$\text{Gap-encode}(\boxed{v} \boxed{w} \boxed{x} \boxed{y} \boxed{z}) = \boxed{v} \boxed{w-v} \boxed{x-w} \boxed{y-x} \boxed{z-y}$$



More schemes that assume specific classes of graphs

...

3 Log (Adjacency structure)

Use different relabelings



Degree-Minimizing: Targeting general graphs (no assumptions on graph structure)


$$\text{Permute}(\boxed{2} \boxed{3} \boxed{4} \boxed{5} \boxed{1M}) = \boxed{v} \boxed{w} \boxed{x} \boxed{y} \boxed{z}$$

(simultaneously for all other neighborhoods)

(1) The more often a label occurs (i.e., the higher vertex degree), the smaller permuted value it receives

$$\text{Gap-encode}(\boxed{v} \boxed{w} \boxed{x} \boxed{y} \boxed{z}) = \boxed{v} \boxed{w-v} \boxed{x-w} \boxed{y-x} \boxed{z-y}$$


(2) Encode new labels with gap encoding (differences between consecutive labels instead of full labels)



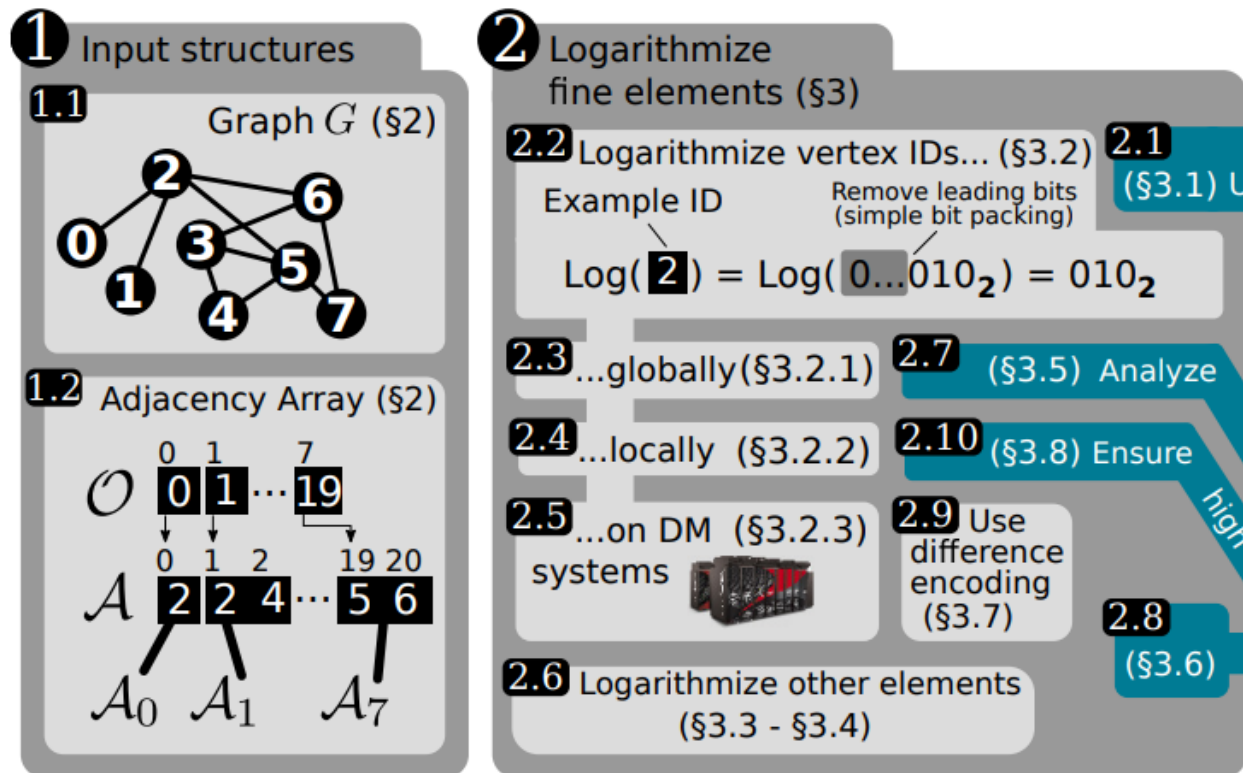
More schemes that assume specific classes of graphs
...

OVERVIEW OF FULL LOG(GRAPH) DESIGN

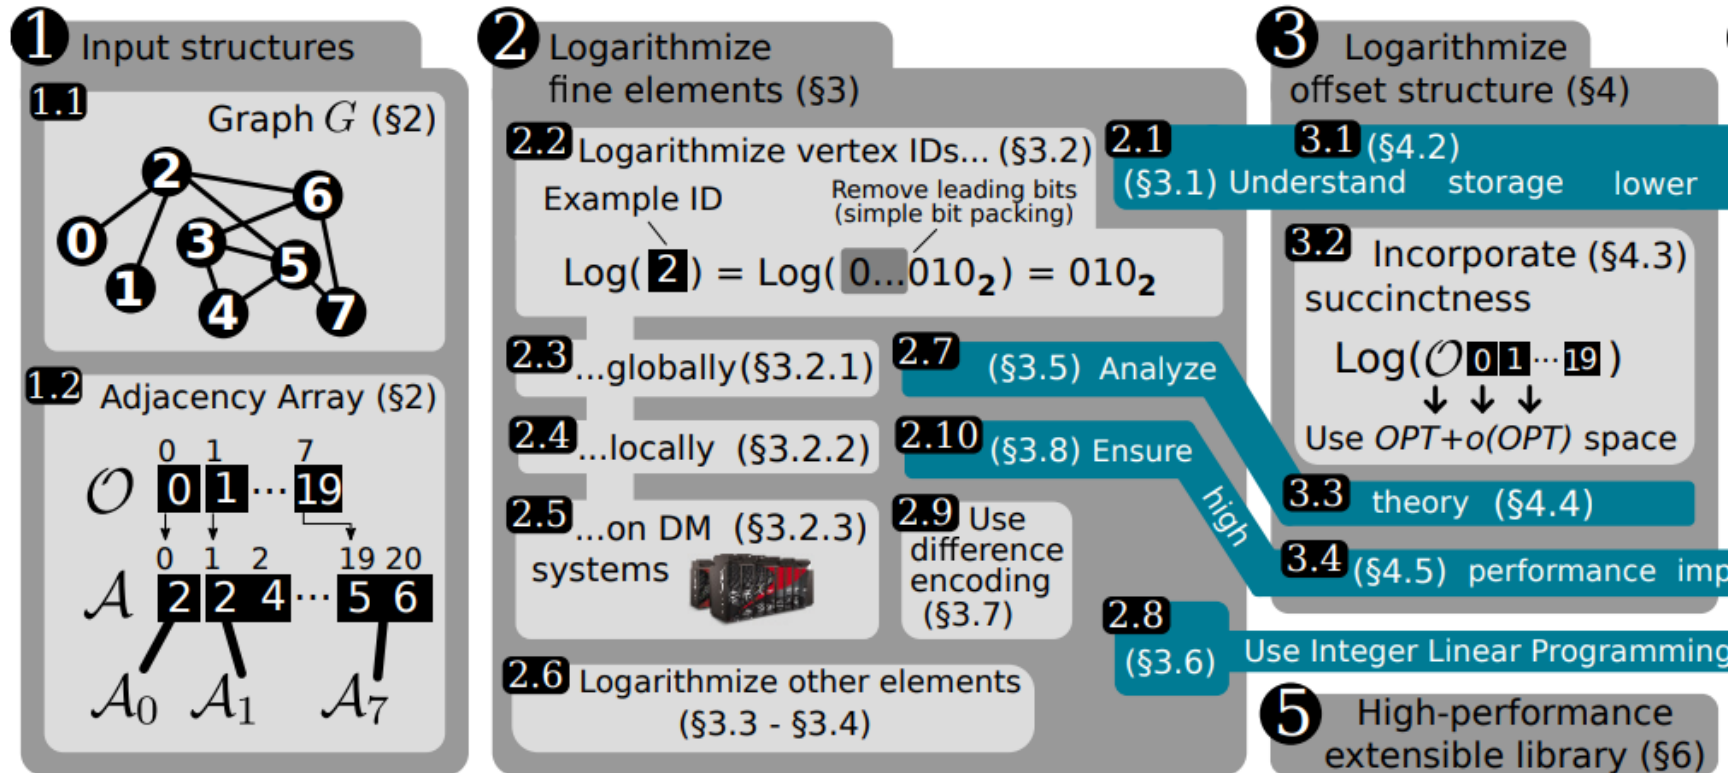
OVERVIEW OF FULL LOG(GRAPH) DESIGN



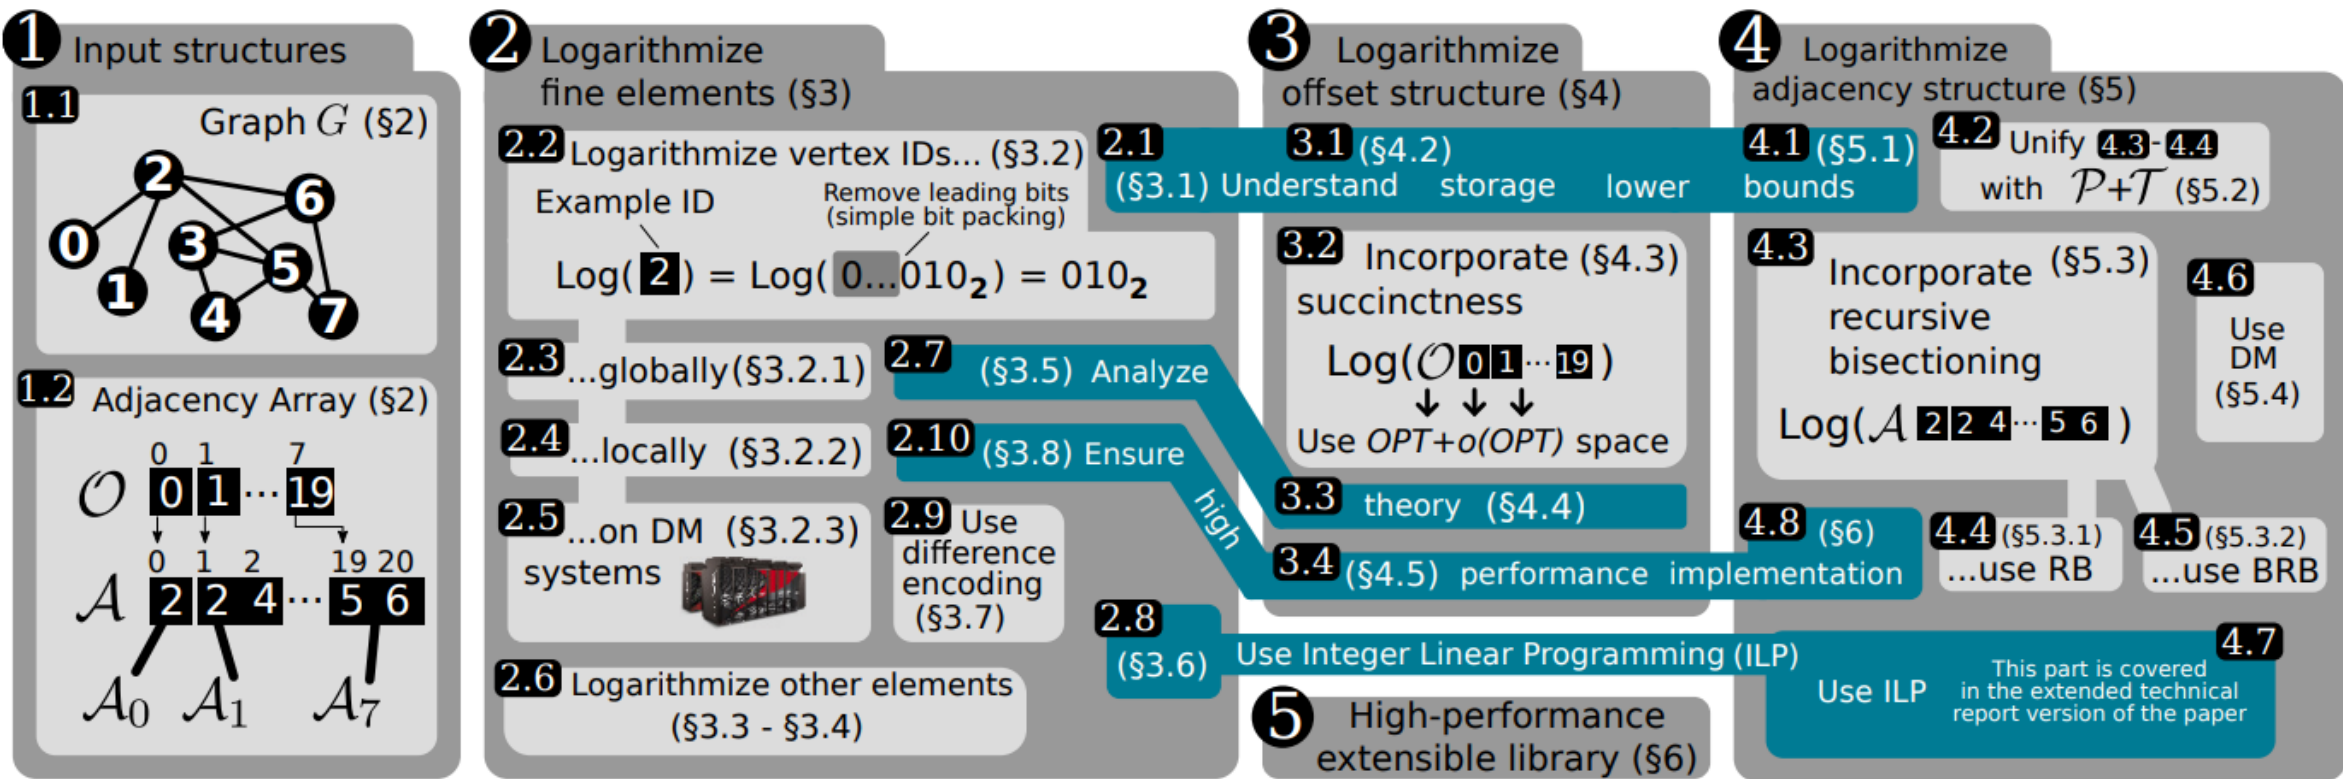
OVERVIEW OF FULL LOG(GRAPH) DESIGN



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 (§3.6)

2.9 Use difference encoding (§3.7)

2.10 (§3.8) Ensure

3 Logarithmize offset structure (§4)

3.1 (§4.2)

3.2 Incorporate (§4.3) succinctness

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

Use $OPT + o(OPT)$ space

3.3 theory (§4.4)

3.4 (§4.5) performance im

3.5 Use Integer Linear Programming (ILP)

3.6 High-performance extensible library (§6)

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

$\text{Log}(A 2 2 4 \dots 5 6)$

4.4 (§5.3.1) ...use RB

4.5 (§5.3.2) ...use BRB

4.6 Use DM (§5.4)

4.7 Use ILP

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

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: $\text{Log}(2) = \text{Log}(10) = 010$

Remove leading bits (simple bit packing)

2.2 ...globally (§3.2)

2.3 ...locally (§3.2.2)

2.4 ...on DM (§3.2.3)

2.5 systems

2.6 Logarithmize other elements (§3.3 - §3.4)

2.7 Use difference encoding (§3.7)

2.8 Use Integer Linear Programming (§3.6)

3 Logarithmize offset structure (§4)

3.1 (§4.2)

3.2 Incorporate (§4.3)

3.3 theory (§4.4)

3.4 (§4.5) performance im

3.5 High-performance extensible library (§6)

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

4.4 (§5.3.1) ...use RB

4.5 (§5.3.2) ...use BRB

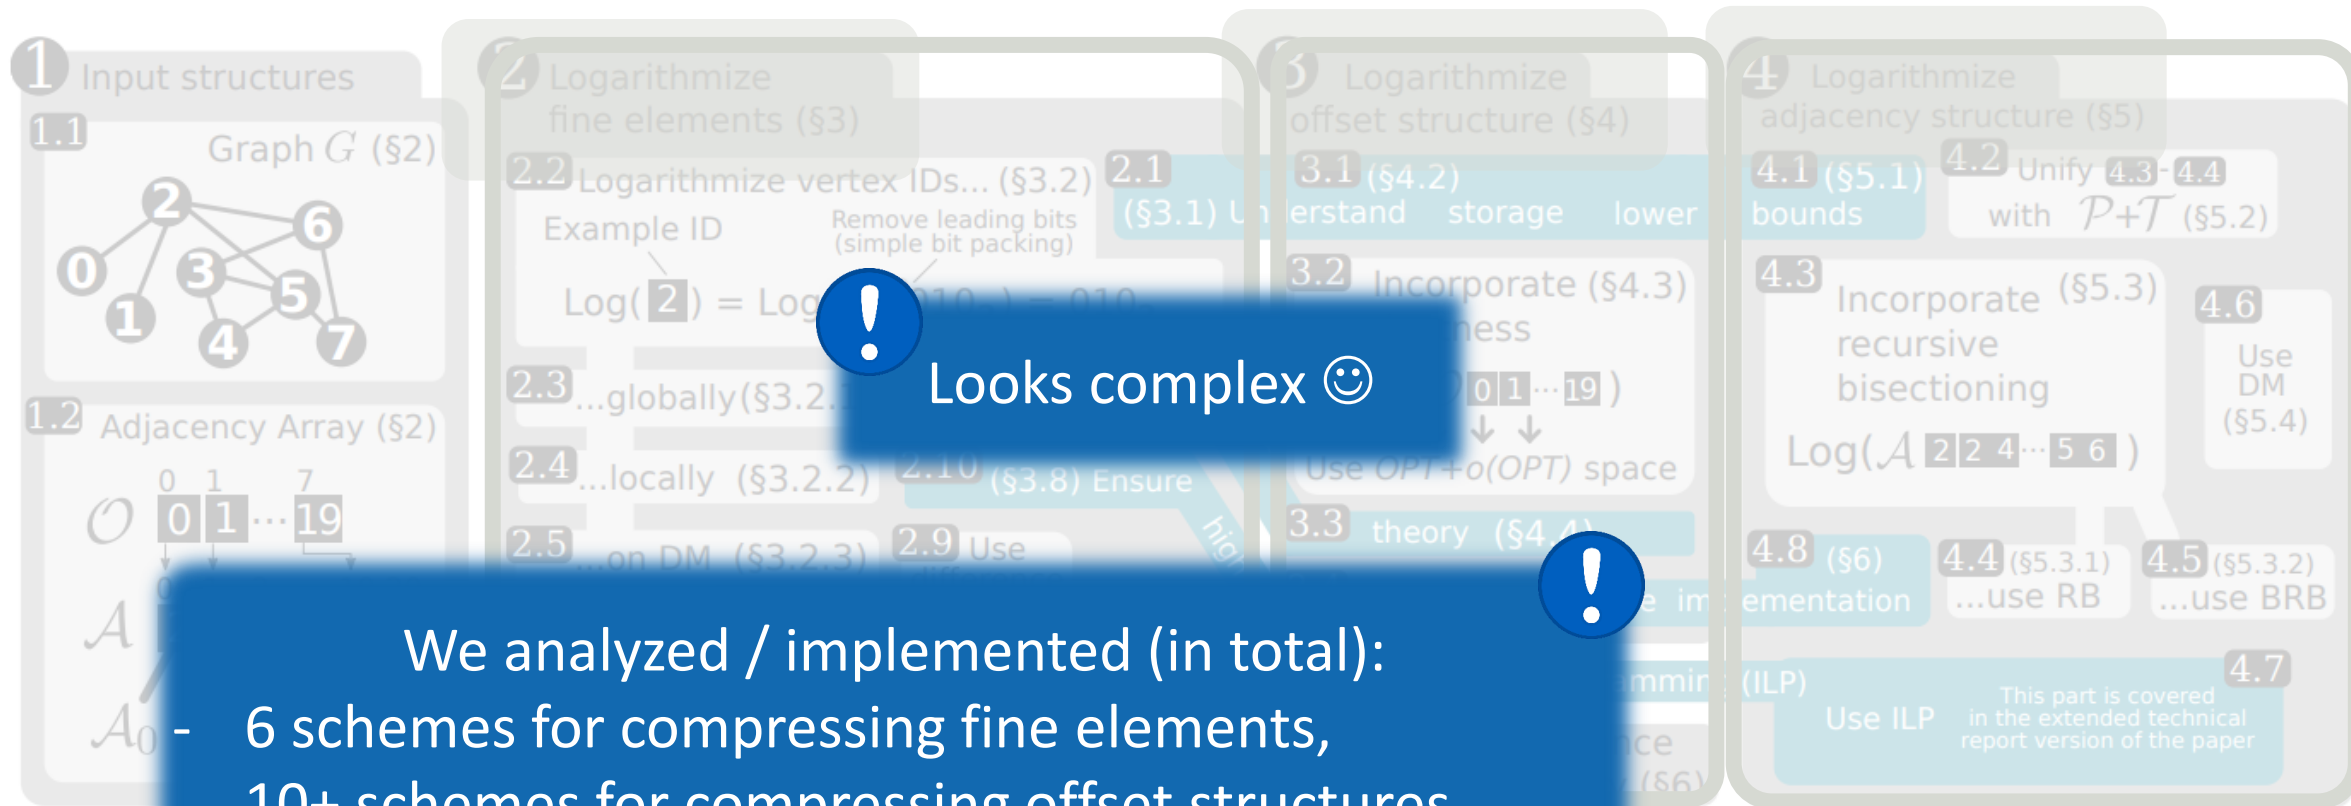
4.6 Use DM (§5.4)

4.7 Use ILP

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

Looks complex 😊

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)

2.2 Example ID

2.3 Remove leading bits (simple bit packing)

2.4 $\text{Log}(2) = \text{Log}(10) = 010$

2.5 ...globally (§3.2.1)

2.6 ...locally (§3.2.2)

2.7 Use DM (§3.2.3)

2.8 Use

3 Logarithmize offset structure (§4)

3.1 (§4.2)

3.2 Incorporate (§4.3)

3.3 theory (§4.4)

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

4.4 Use DM (§5.4)

4.5 Use BRB (§5.3.2)

4.6 Use ILP (§6)

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

4.8 ...use RB

4.9 ...use BRB

4.10 Use OPT+o(OPT) space

4.11 Ensure

4.12 Use

4.13 Use

4.14 Use

4.15 Use

4.16 Use

4.17 Use

4.18 Use

4.19 Use

4.20 Use

4.21 Use

4.22 Use

4.23 Use

4.24 Use

4.25 Use

4.26 Use

4.27 Use

4.28 Use

4.29 Use

4.30 Use

4.31 Use

4.32 Use

4.33 Use

4.34 Use

4.35 Use

4.36 Use

4.37 Use

4.38 Use

4.39 Use

4.40 Use

4.41 Use

4.42 Use

4.43 Use

4.44 Use

4.45 Use

4.46 Use

4.47 Use

4.48 Use

4.49 Use

4.50 Use

4.51 Use

4.52 Use

4.53 Use

4.54 Use

4.55 Use

4.56 Use

4.57 Use

4.58 Use

4.59 Use

4.60 Use

4.61 Use

4.62 Use

4.63 Use

4.64 Use

4.65 Use

4.66 Use

4.67 Use

4.68 Use

4.69 Use

4.70 Use

4.71 Use

4.72 Use

4.73 Use

4.74 Use

4.75 Use

4.76 Use

4.77 Use

4.78 Use

4.79 Use

4.80 Use

4.81 Use

4.82 Use

4.83 Use

4.84 Use

4.85 Use

4.86 Use

4.87 Use

4.88 Use

4.89 Use

4.90 Use

4.91 Use

4.92 Use

4.93 Use

4.94 Use

4.95 Use

4.96 Use

4.97 Use

4.98 Use

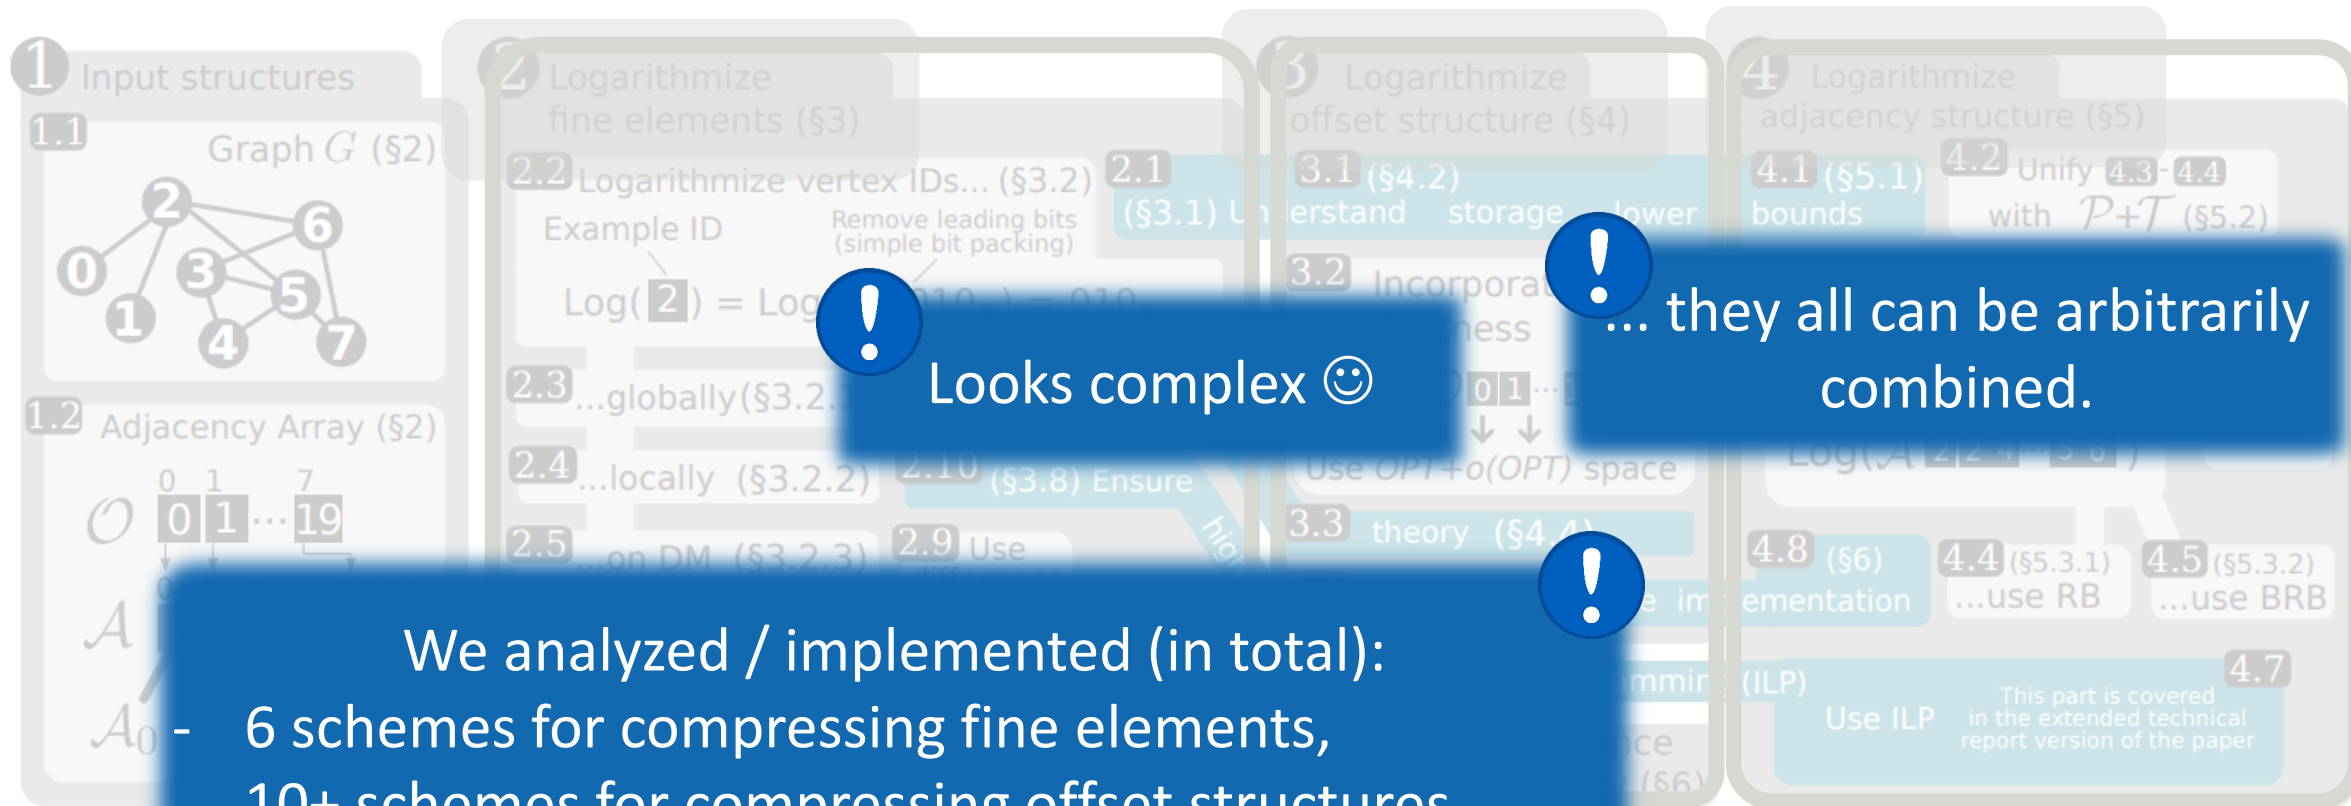
4.99 Use

4.100 Use

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



1 Input structures

1.1 Graph G (§2)

1.2 Adjacency Array (§2)

2 Logarithmize fine elements (§3)

2.2 Logarithmize vertex IDs... (§3.2)

Example ID Remove leading bits (simple bit packing)

$\text{Log}(2) = \text{Log}(10) = 010$

3 Logarithmize offset structure (§4)

3.1 (§4.2) Understand storage lower bounds

3.2 Incorporate theory

3.3 theory (§4.4)

4 Logarithmize adjacency structure (§5)

4.1 (§5.1) bounds

4.2 Unify 4.3-4.4 with $P+T$ (§5.2)

4.4 (§5.3.1) ...use RB

4.5 (§5.3.2) ...use BRB

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

! Looks complex 😊

! ... they all can be arbitrarily combined.

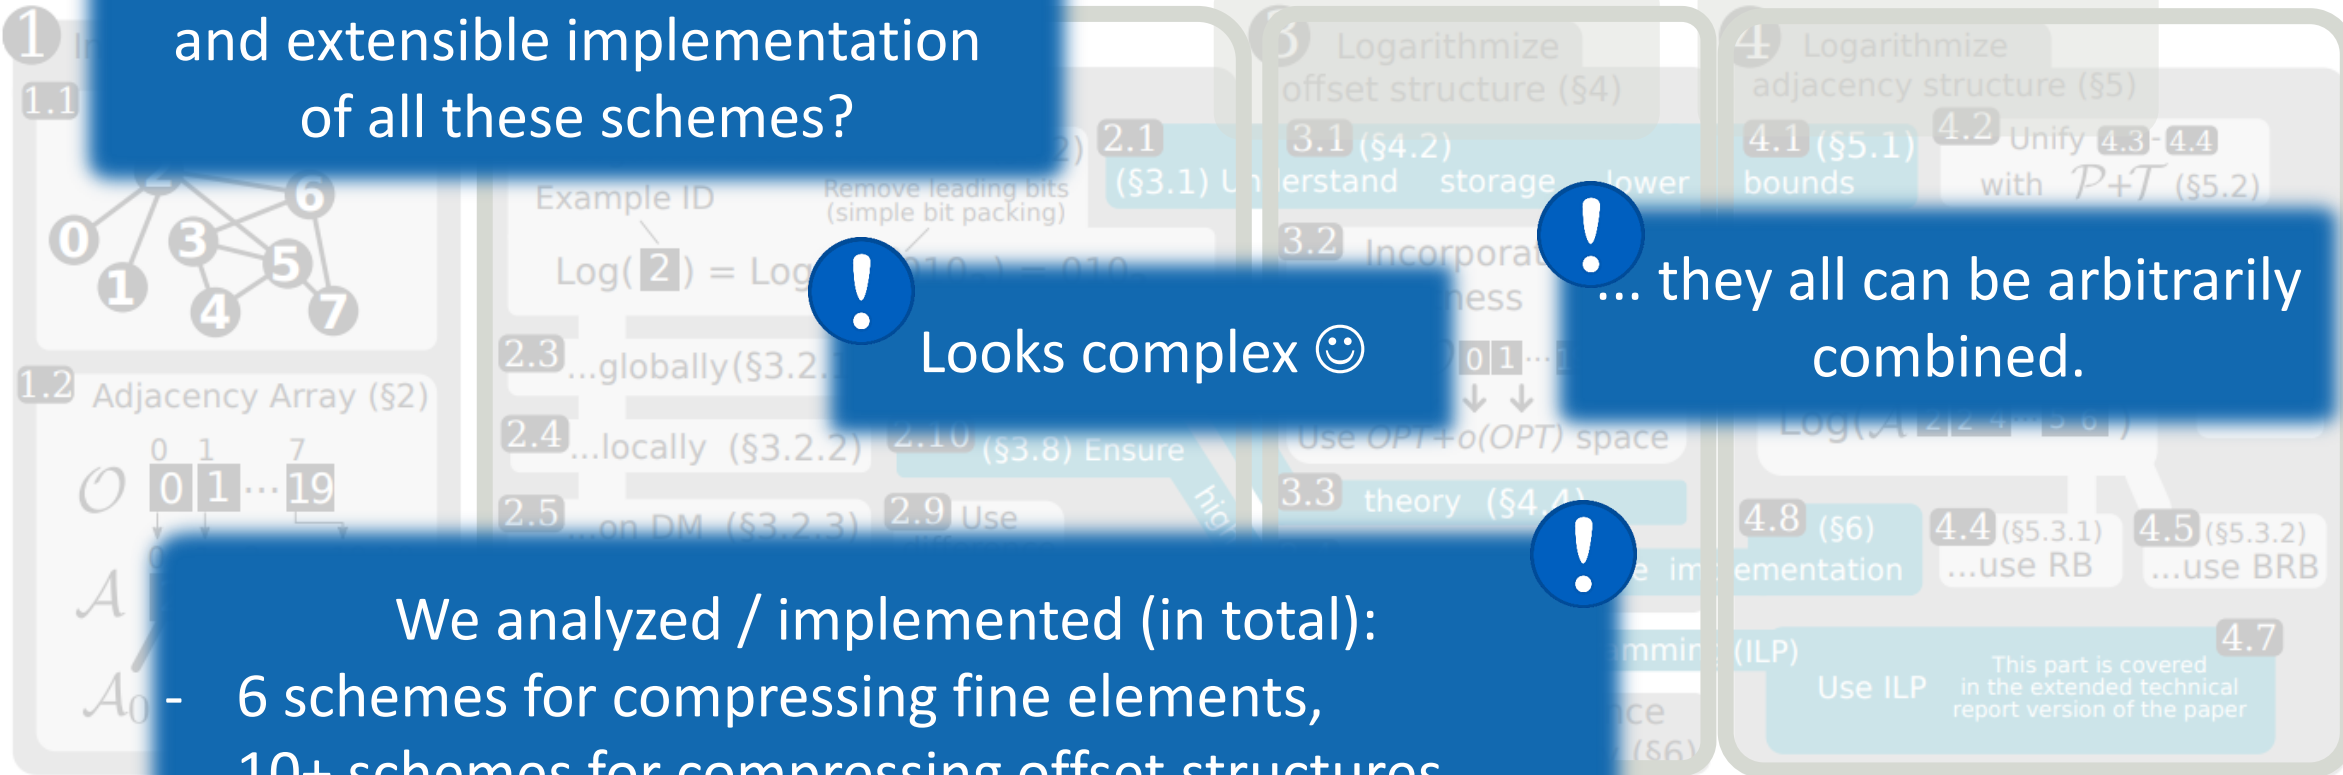
We analyzed / implemented (in total):

- 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.

! Use ILP in the extended technical report version of the paper

We analyzed / implemented (in total):

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

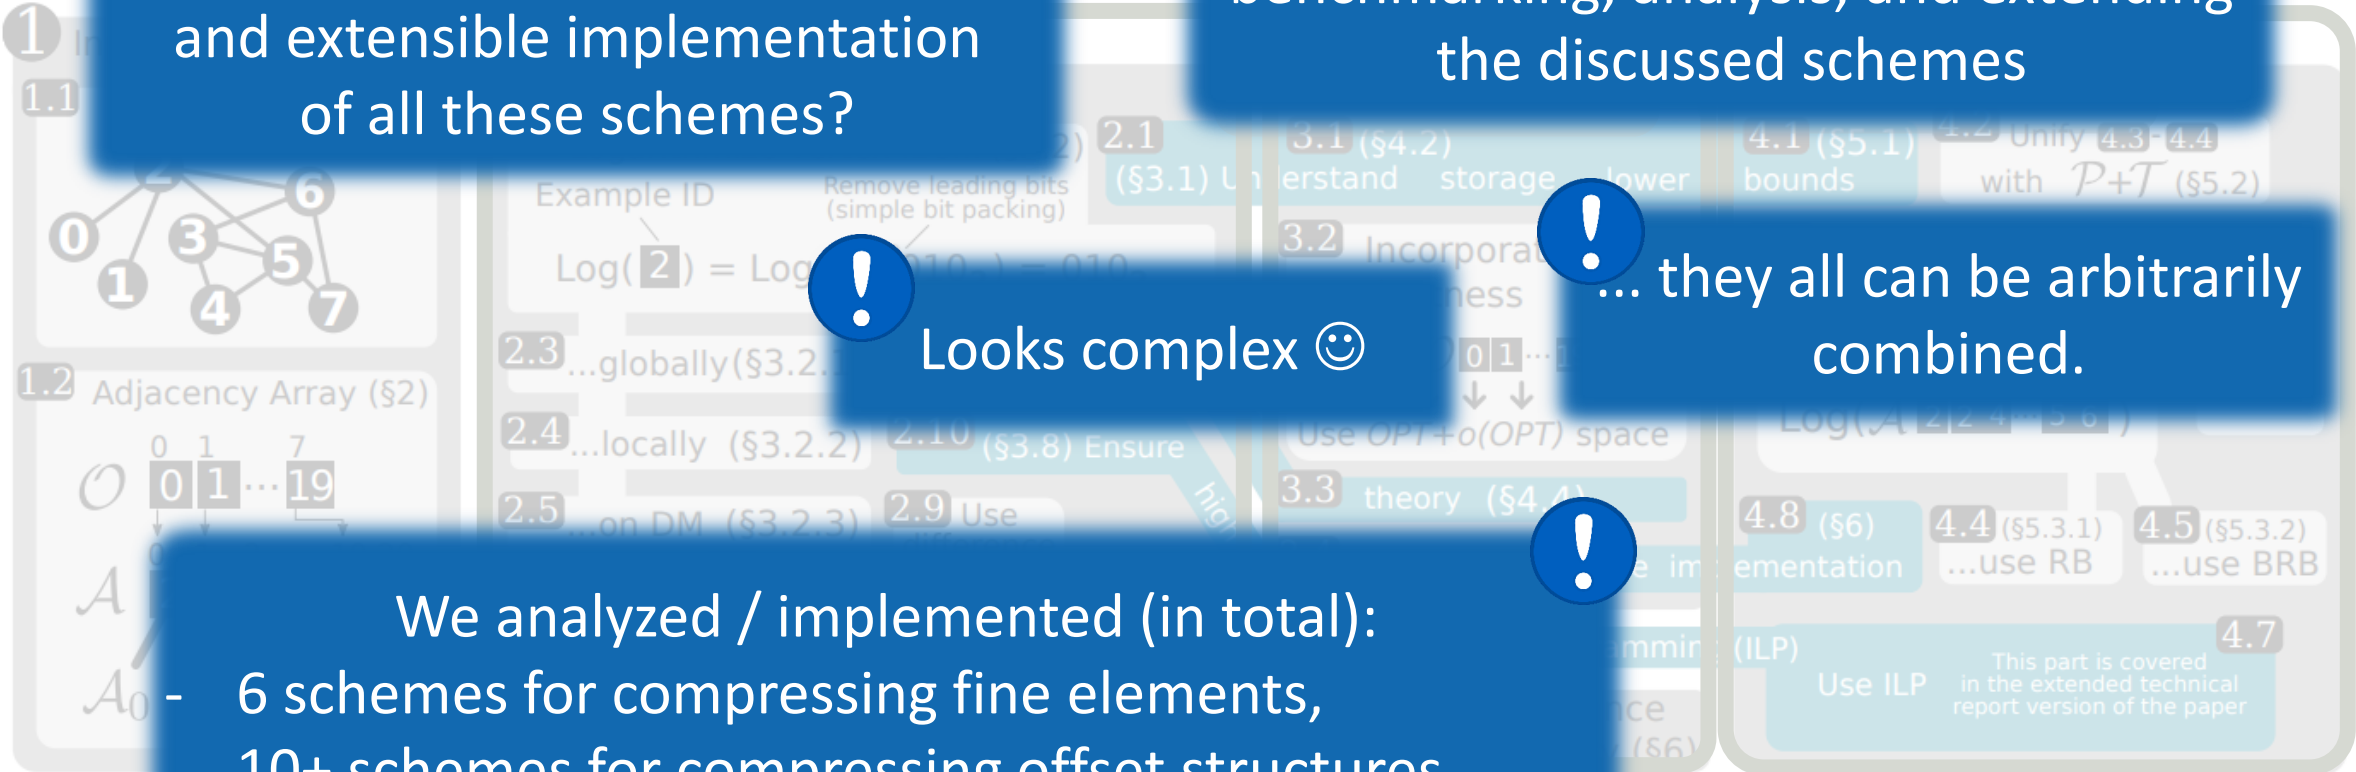
OVERVIEW OF FULL LOG(GRAPH) DESIGN



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



We use C++ templates to develop a library that facilitates implementation, benchmarking, analysis, and extending the discussed 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

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





Kronecker graphs
Number of vertices: 4M

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



Kronecker graphs
Number of vertices: 4M


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



Kronecker graphs
Number of vertices: 4M


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

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




Kronecker graphs
Number of vertices: 4M

**Log(Graph)
accelerates GAPBS**



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

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




Kronecker graphs
Number of vertices: 4M

**Log(Graph)
accelerates GAPBS**

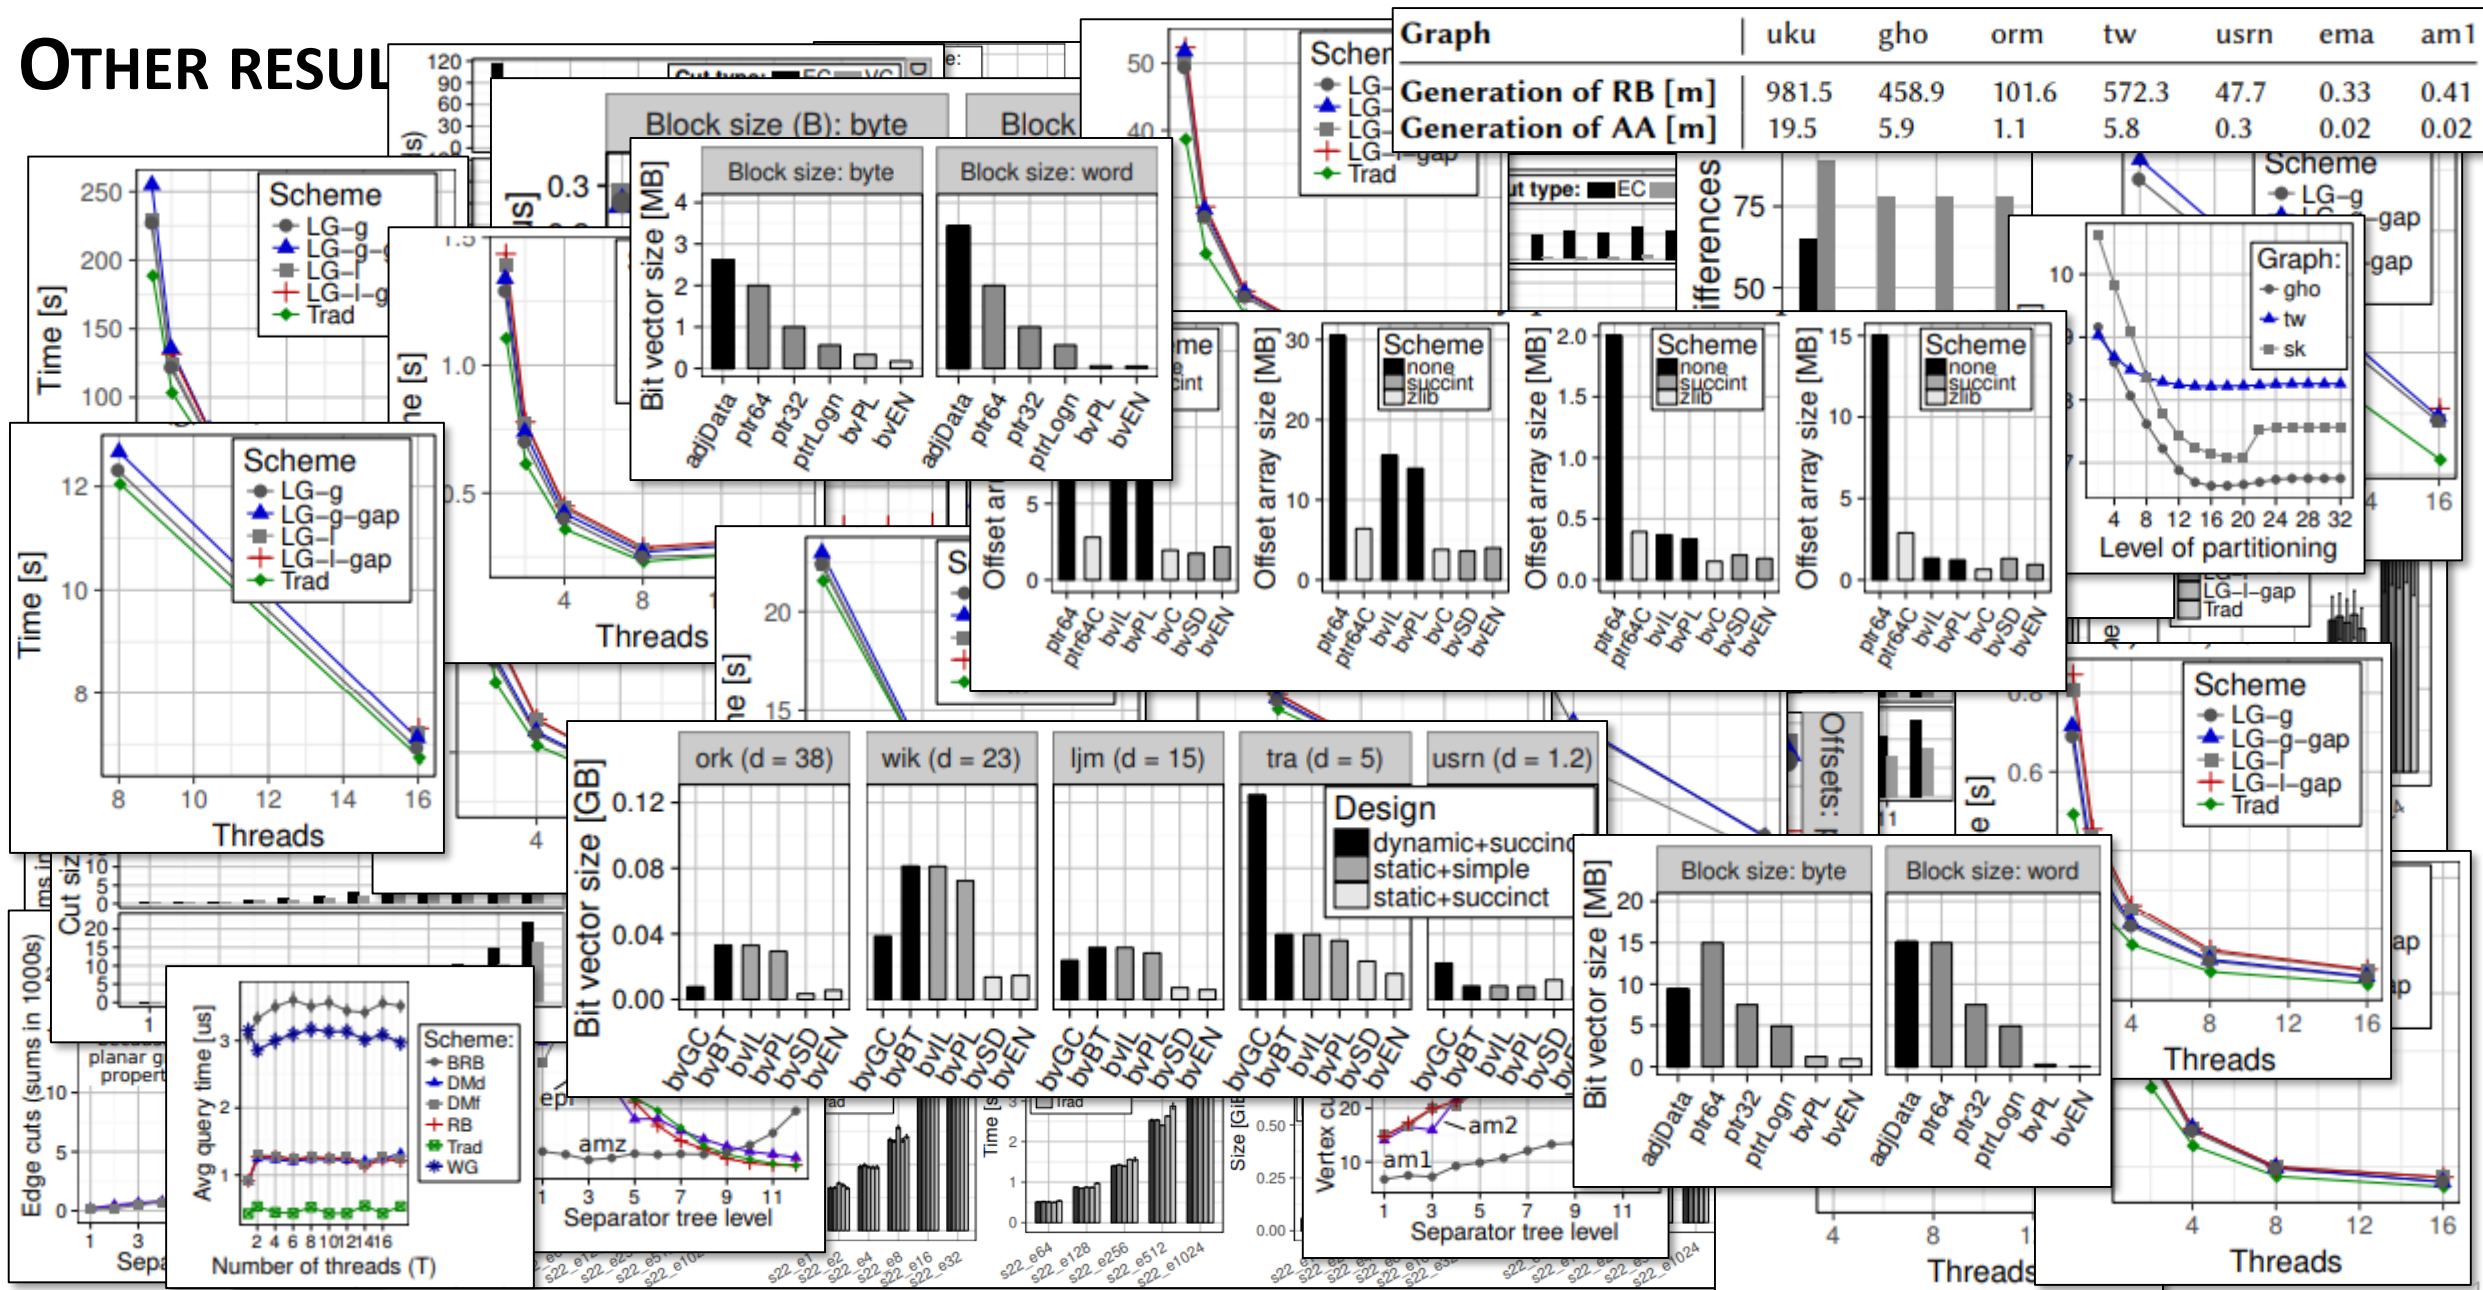
**Both storage and performance
are improved simultaneously**

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

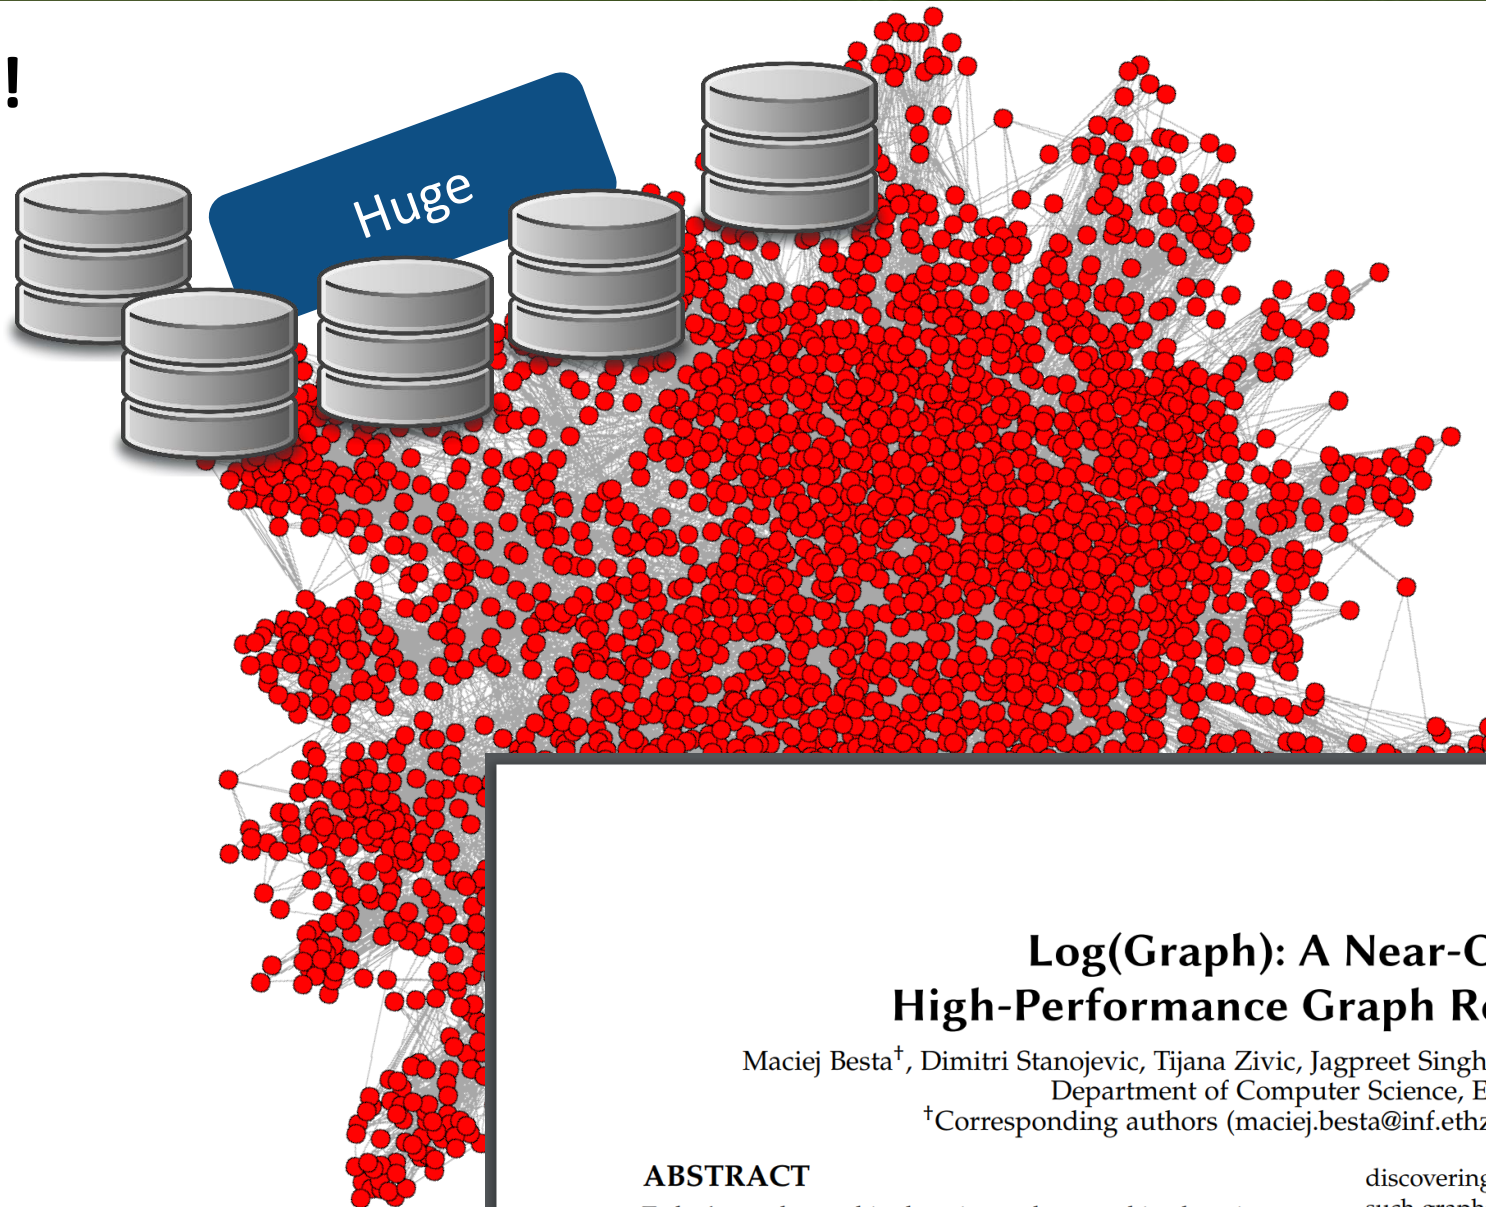


OTHER RESULTS

OTHER RESULT



Problems!



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!



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

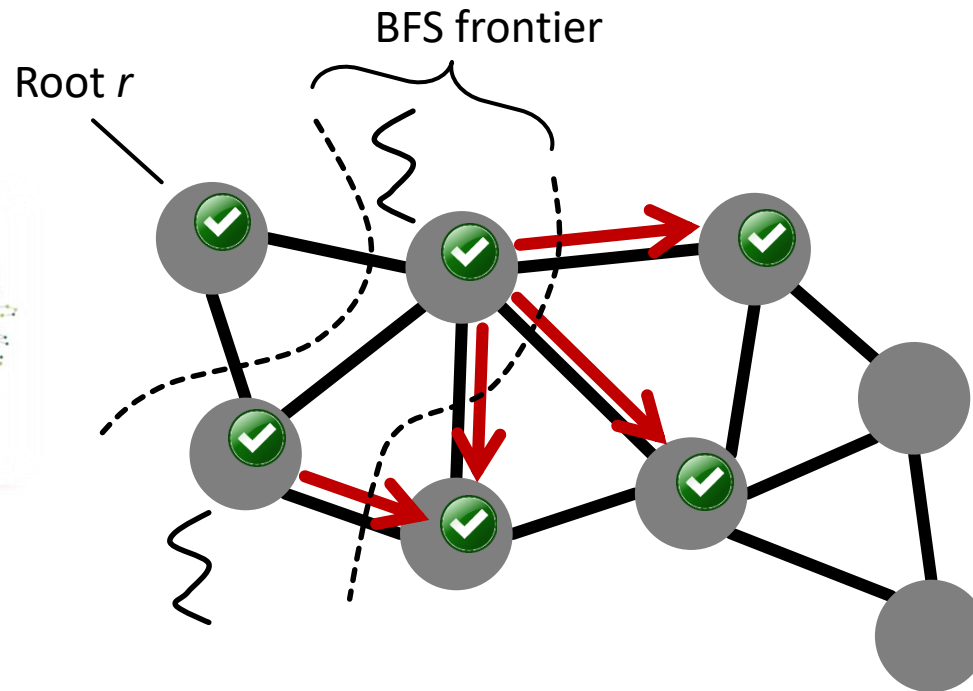


BFS


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


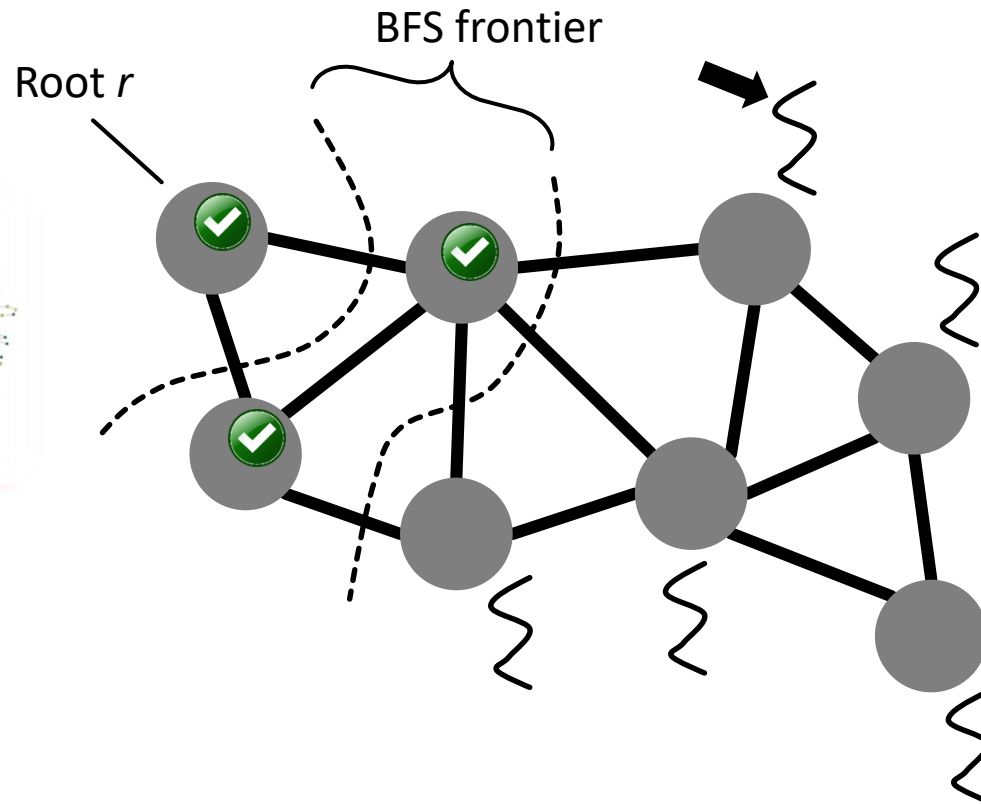


BFS

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




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



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  $\Delta$  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] +  $\mathcal{W}_{(v,w)}$ ;
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
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] +  $\mathcal{W}_{(w,v)}$  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  $\Pi$  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  $G'$ 
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]; pred_1[w] = pred_1[v] 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..v] = [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 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} ( $\mathcal{W}_{(v,w)}$  + 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]  $\neq$  flag)  $\wedge$ 
11            ( $\mathcal{W}_{(v,w)}$  < min_e_wgt[sv_flag[w]]) R {
12            min_e_wgt[sv_flag[w]] =  $\mathcal{W}_{(v,w)}$  W I; PUSHING
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]  $\neq$  flag)  $\wedge$  ( $\mathcal{W}_{(v,w)}$  < min_e_wgt[flag]) R {
16            min_e_wgt[flag] =  $\mathcal{W}_{(v,w)}$ ; min_e_v[flag] = v; PULLING
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,  $\mathcal{P}$ );
9   while (!done) {
10     for  $\mathcal{P} \in \mathcal{P}$  do in par {seq_color_partition( $\mathcal{P}$ );
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 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) R {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. 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 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)..(s, 0, 1, 0)..(n, 0, 0, 0)]);
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};

```

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;} PART 1: PUSHING
15     {new_pr[v] += (f*pr[u])/d(v) R;} PART 2: 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 ∧ (W(v,w) < min_e_wgt[flag]) R {
9         min_e_wgt[flag] = W(v,w) W I;
10        min_e_v[flag] = v; min_e_w[flag] = w;
11        min_e_v[flag] = v; min_e_w[flag] = w;
12        new_flag[flag] = sv_flag[w]; } R
13      } }
14    } }
15    while flag = merge_order.pop() do {
16      neigh_flag = sv_flag[min_e_w[flag]];
17      for v in sv[flag] do sv_flag[flag] = sv_flag[neigh_flag];
18      sv[neigh_flag] = sv[flag] U sv[neigh_flag];
19      MST[neigh_flag] = MST[flag] U MST[neigh_flag]
20      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;} PART 1: PUSHING
17       {avail[v][c[v]] = 0 R I;} PART 2: PULLING
18     } }
19 } }

```

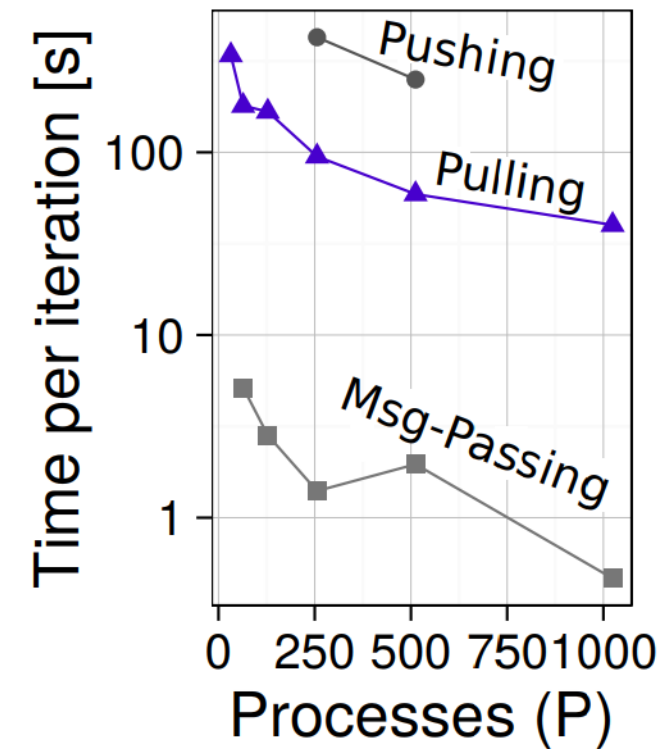

Check out the paper 😊

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

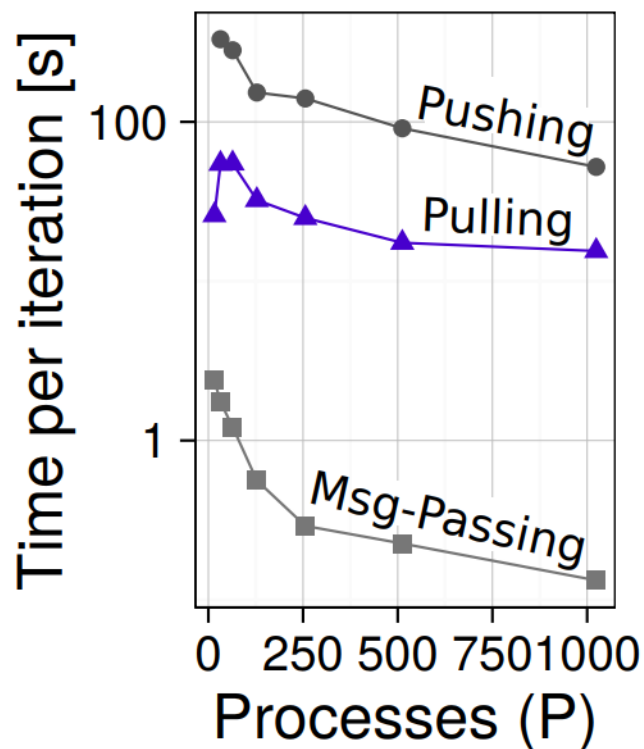
Kronecker graphs

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 ...



Accelerating Irregular Computations with Hardware Transactional Memory and Active Messages

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

Torsten Hoefler
Department of Computer Science
ETH Zurich



ynchronization-heavy

To Push or To Synchron

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

ABSTRACT

We reduce the cost of communication and processing by analyzing the fastest way to push the updates to a shared state or pulling the 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 and D) run on another socket. We investigate the applicability of

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



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


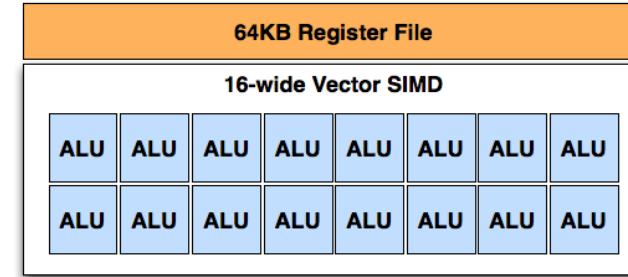
- Deployed in various hardware
- Becoming more popular
- Offers a lot of „regular” compute power



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

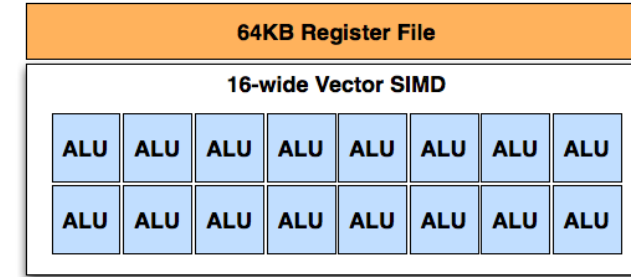


C : „Chunk” size: SIMD width (CPUs, KNLs), warp size (GPUs)

$C = 8$ (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)



ETH zürich

VECTORIZATION


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

$C = 16$

Regular

$C = 32$

$C = 8$



$C = 8$ (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

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



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


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

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

BREADTH-FIRST SEARCH

ALGEBRAIC FORMULATION

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



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


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

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

BREADTH-FIRST SEARCH

ALGEBRAIC FORMULATION

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



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

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



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

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



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



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

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

Stored with a dense or a
sparse format



?

How to do this in practice?

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

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

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


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

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)



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

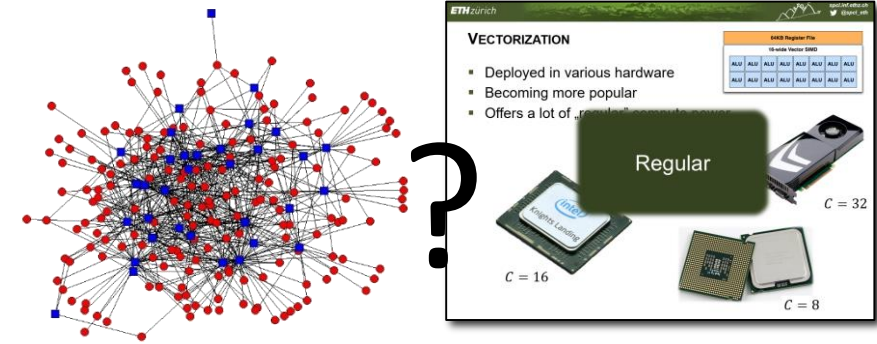

Column indices stored
in the *col* array size: $2m$ cells

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


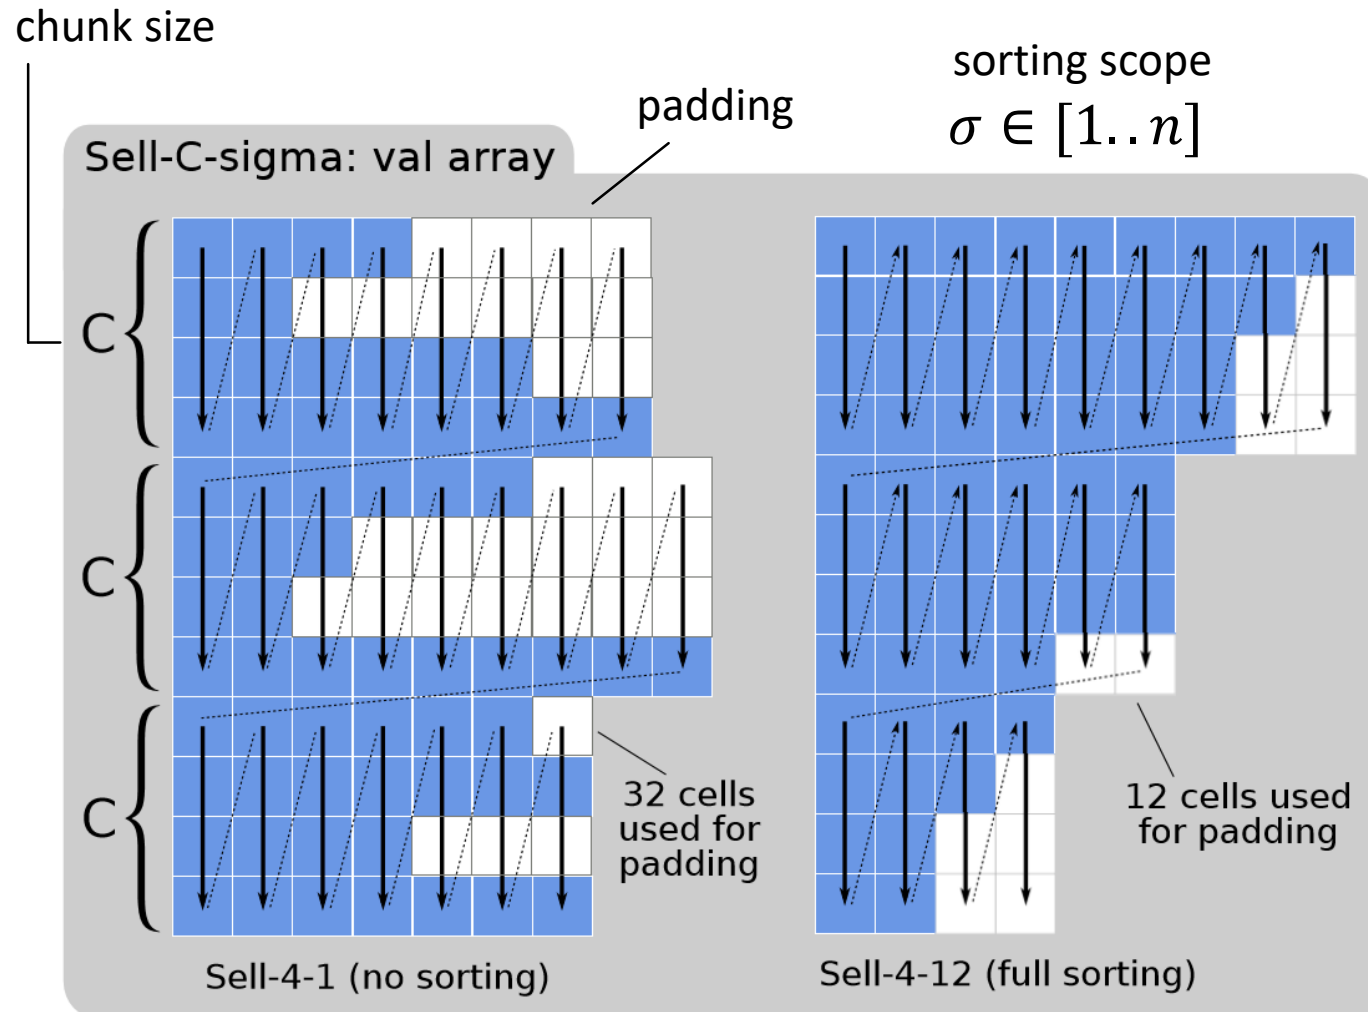
SELL-C-SIGMA



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
    
```



+

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


! 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 😊

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
    
```



+

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

! 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



Summary and outlook



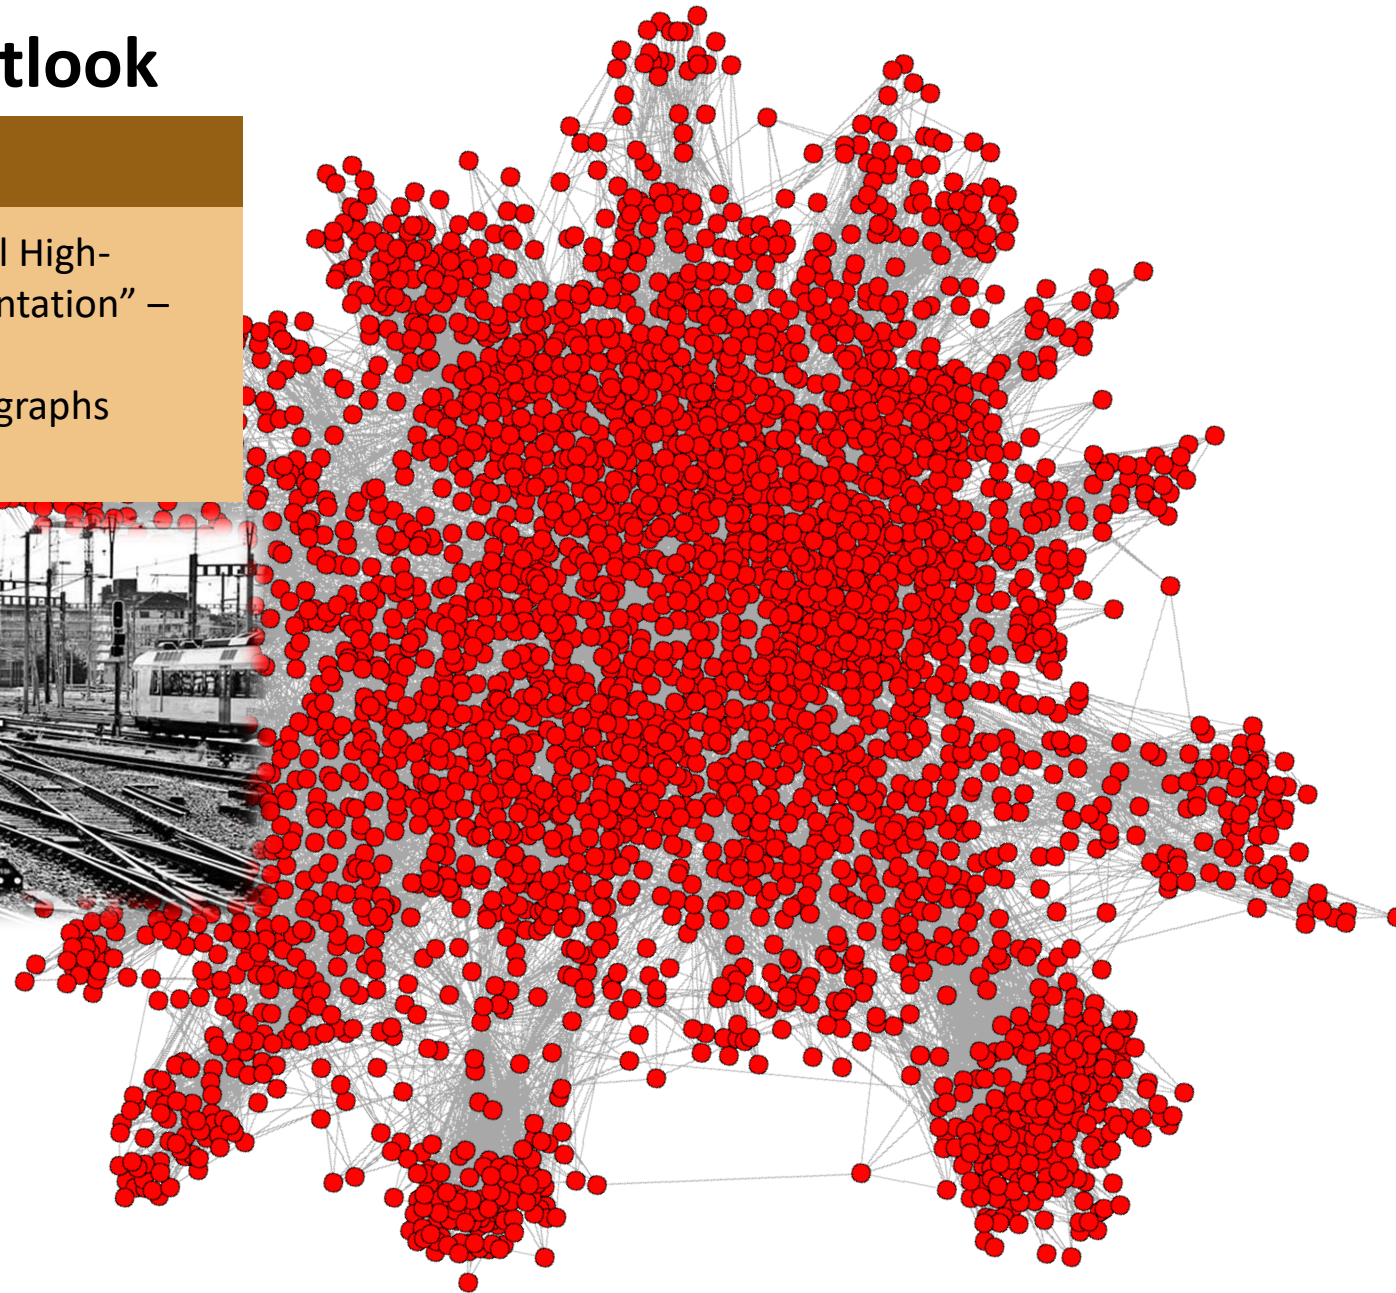
Summary and outlook



Summary and outlook

Log(Graph)

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



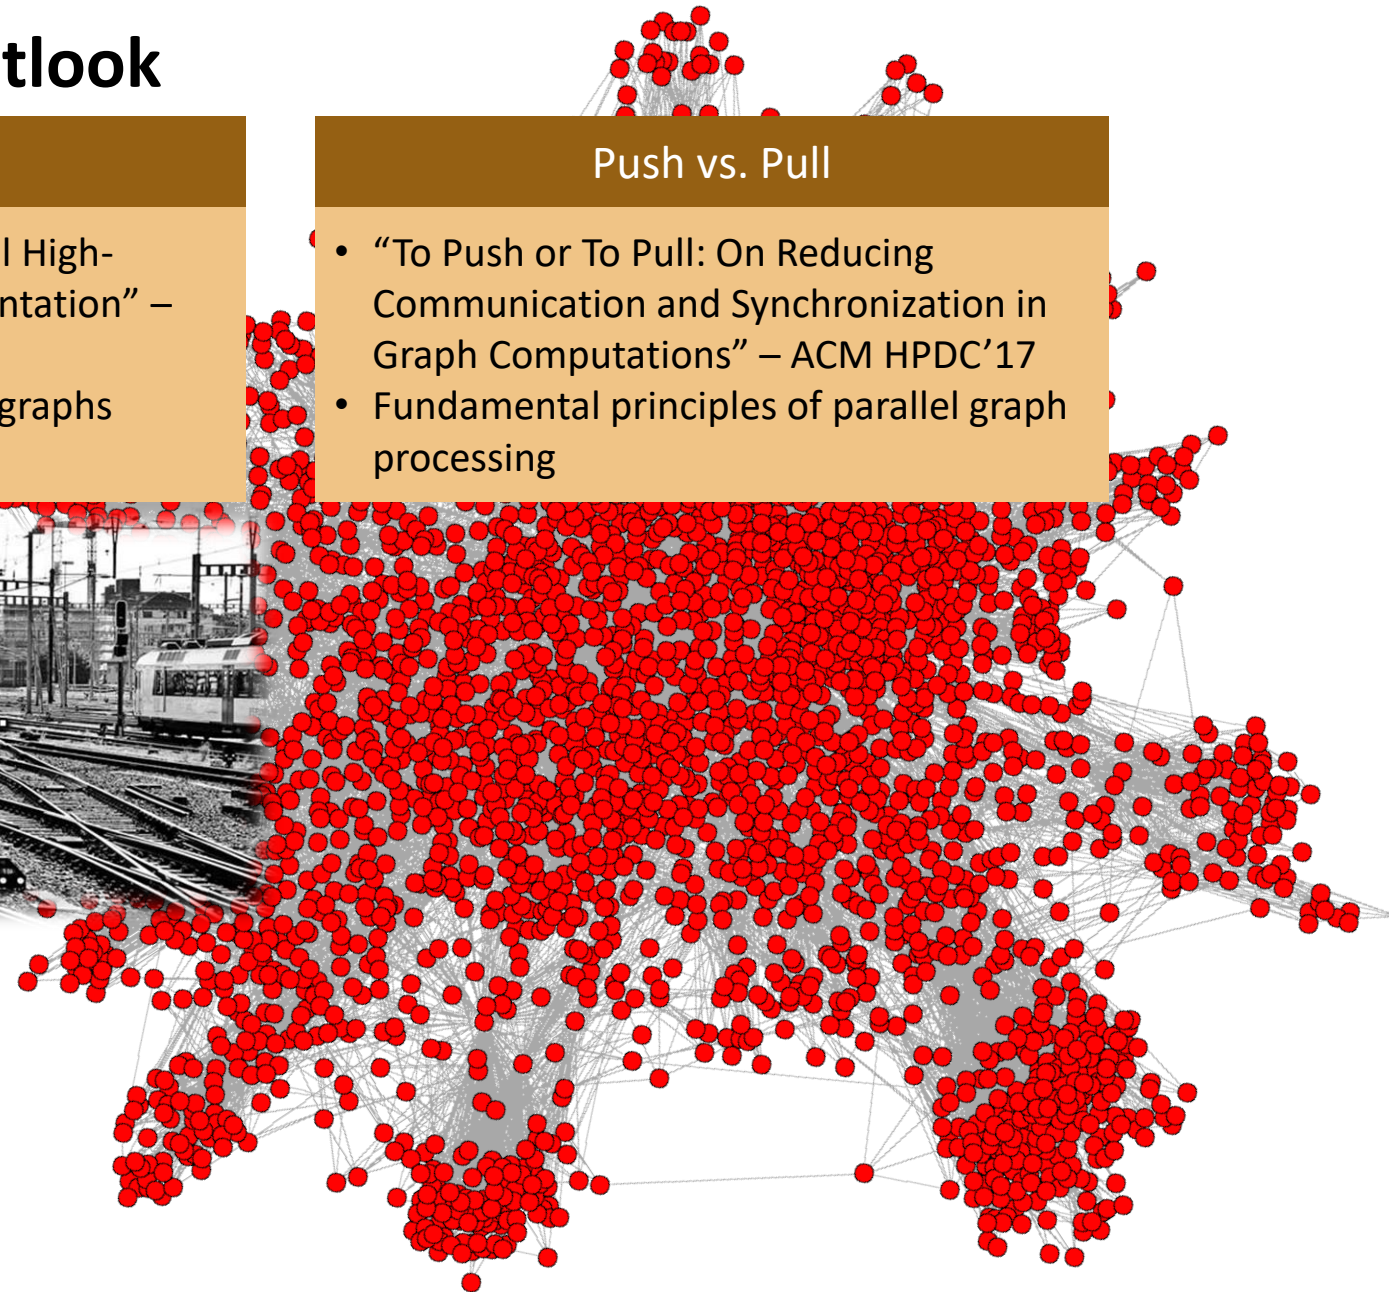
Summary and outlook

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)

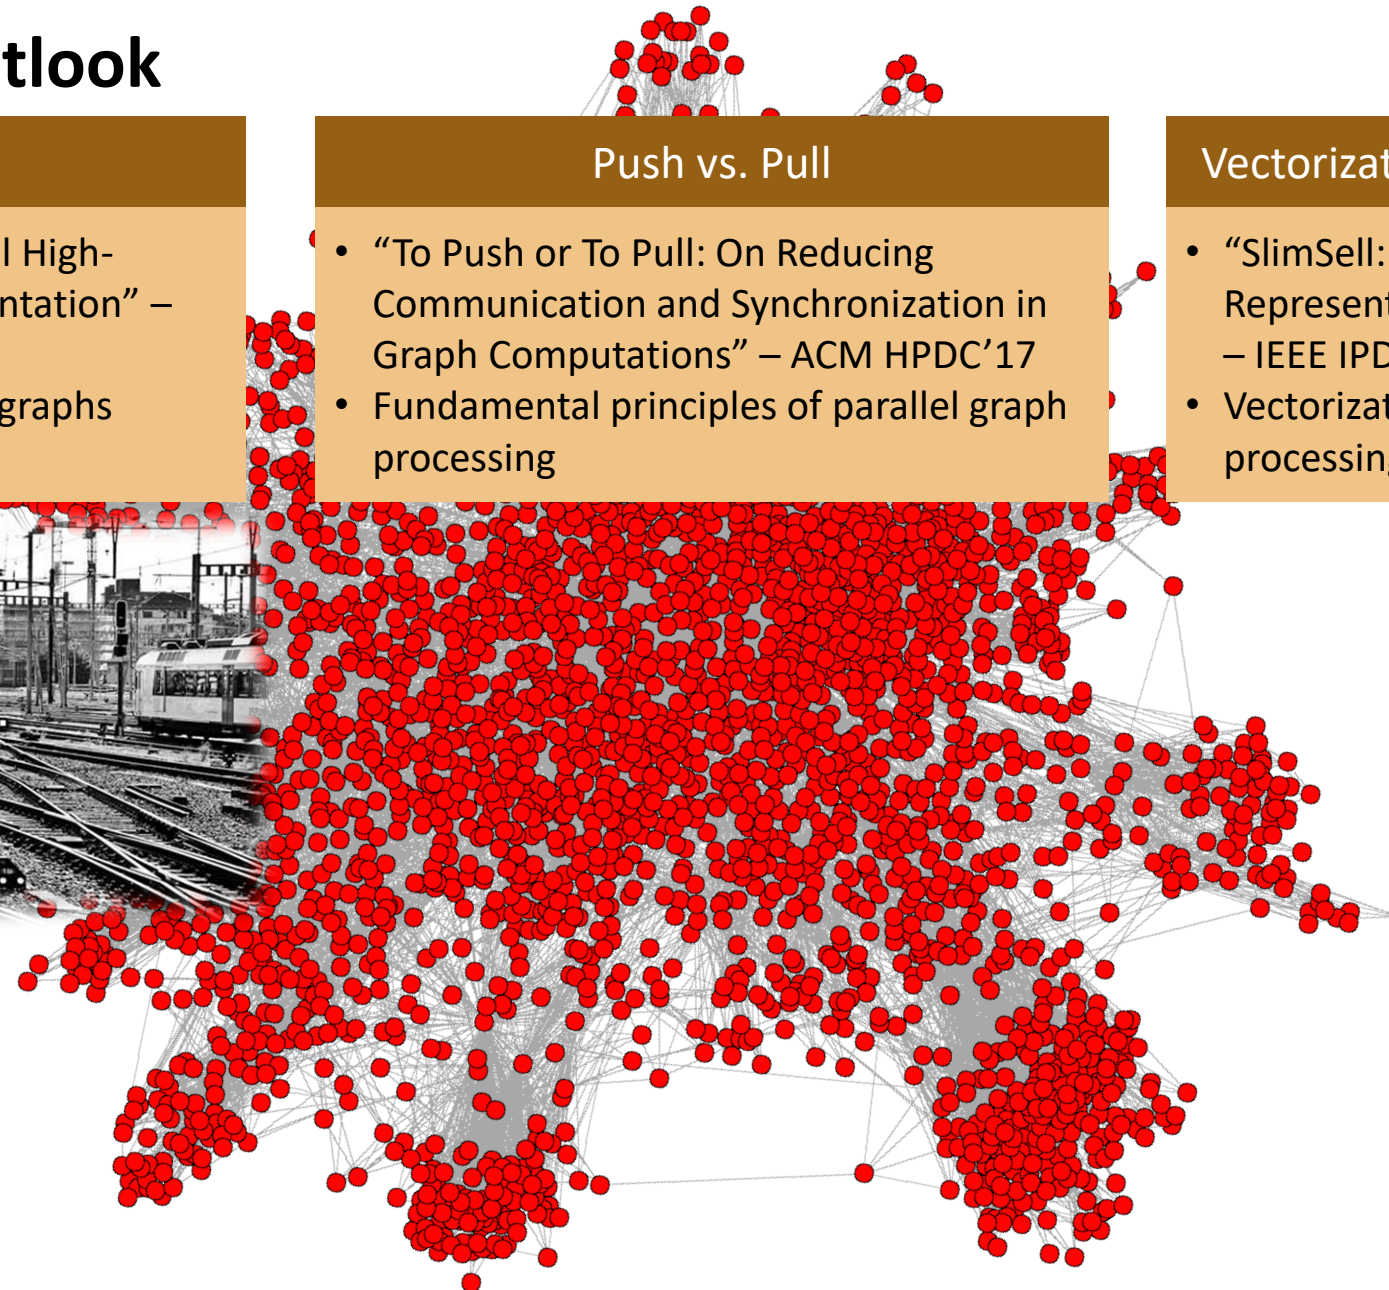
- “Log(Graph): A Near-Optimal High-Performance Graph Representation” – ACM PACT’18
- Minimal storage bounds for graphs during processing

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

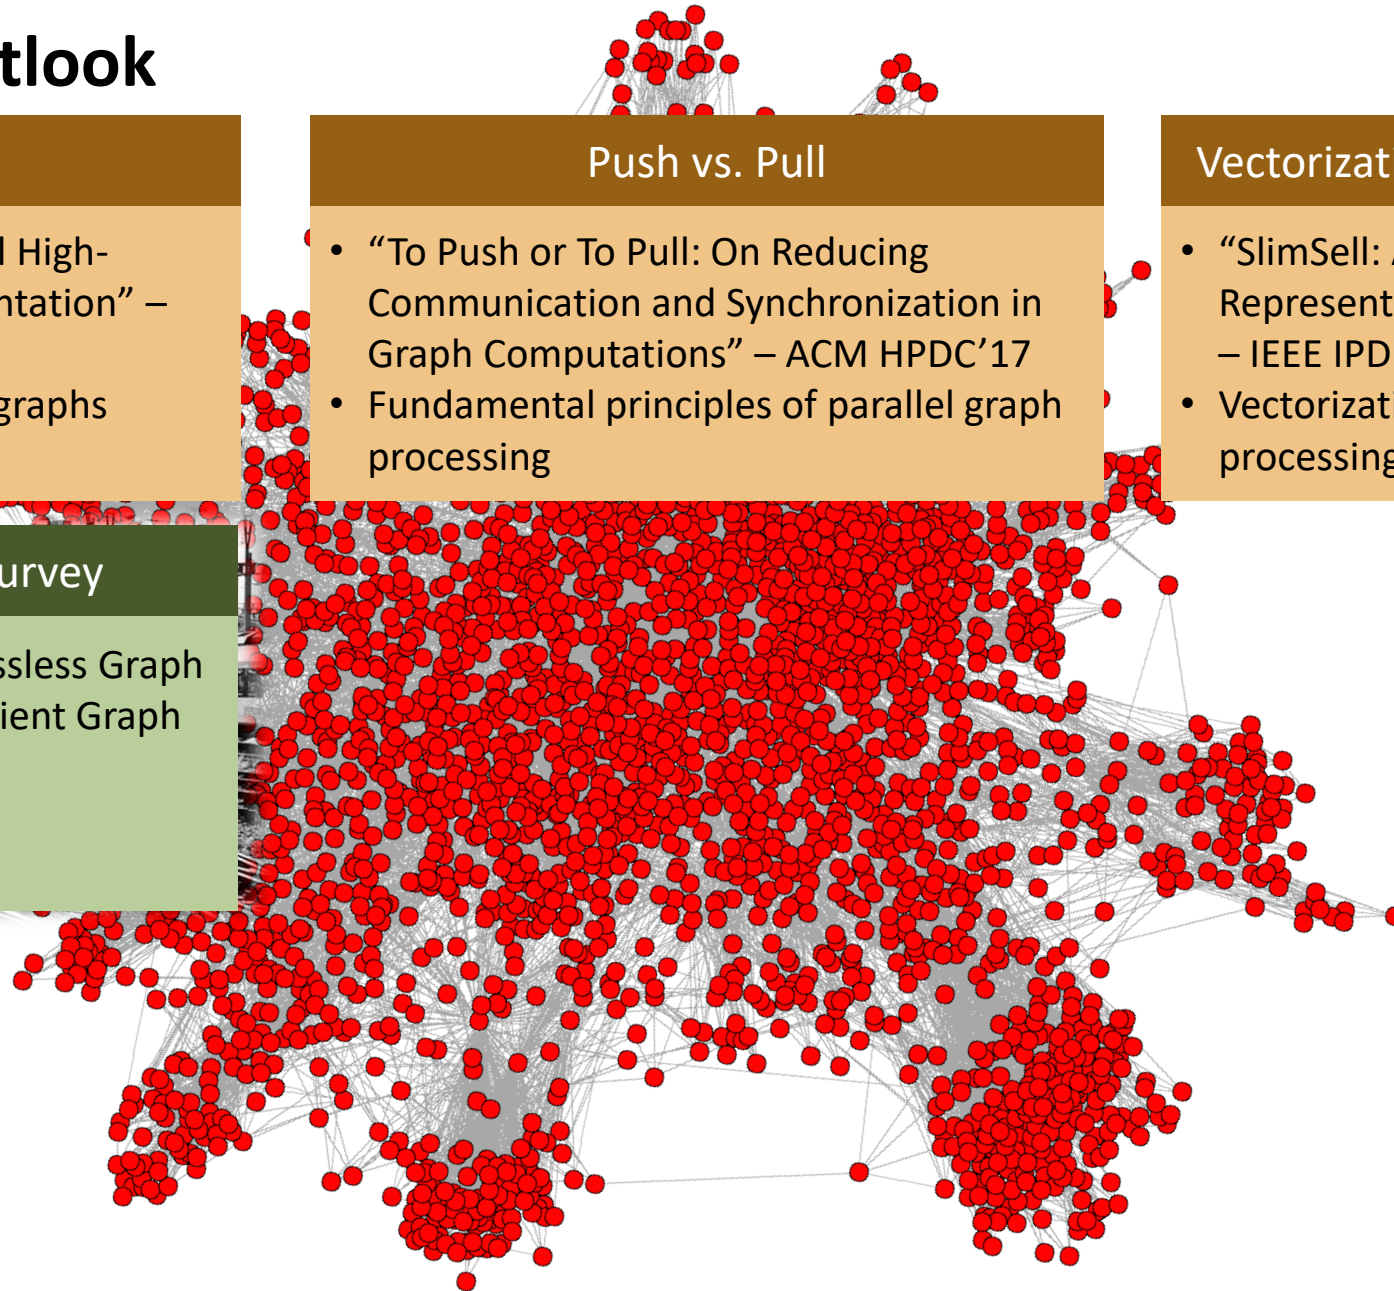
- “To Push or To Pull: On Reducing Communication and Synchronization in Graph Computations” – ACM HPDC’17
- Fundamental principles of parallel graph processing

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

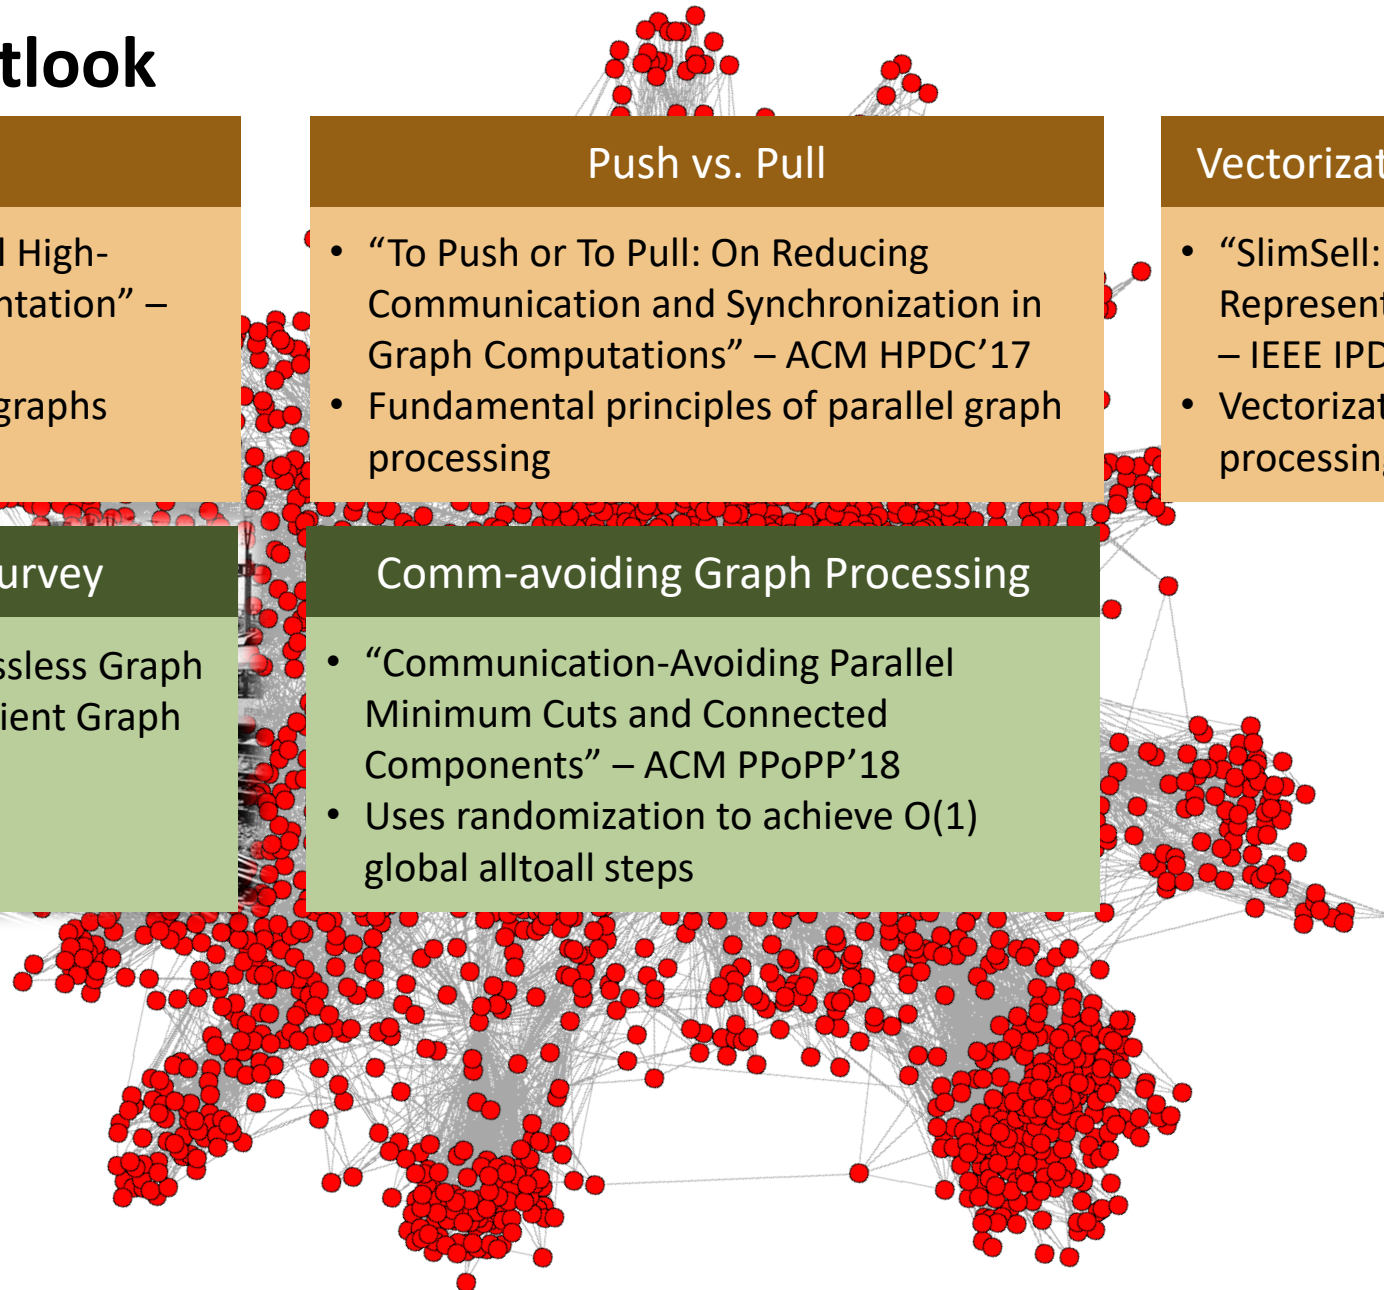
- “SlimSell: A Vectorized Graph Representation for Breadth-First Search” – IEEE IPDPS’17
- Vectorization schemes for parallel graph processing

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

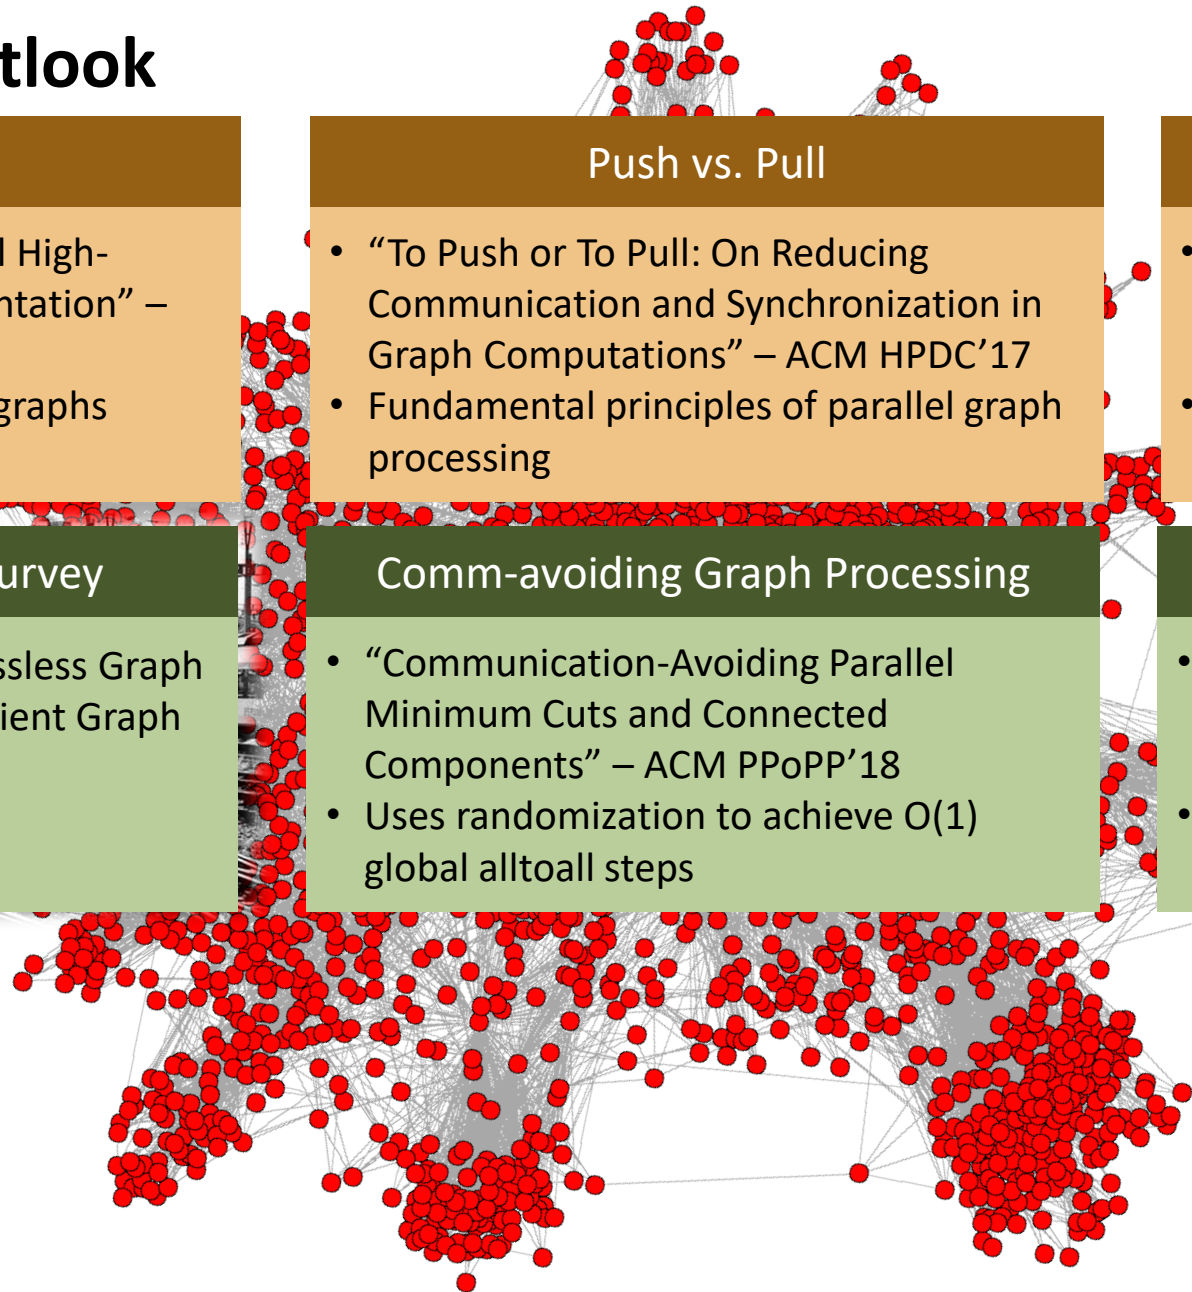
- “Survey and Taxonomy of Lossless Graph Compression and Space-Efficient Graph Representations” – arXiv
- Comprehensive overview
 - 54 pages, 465 references

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

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

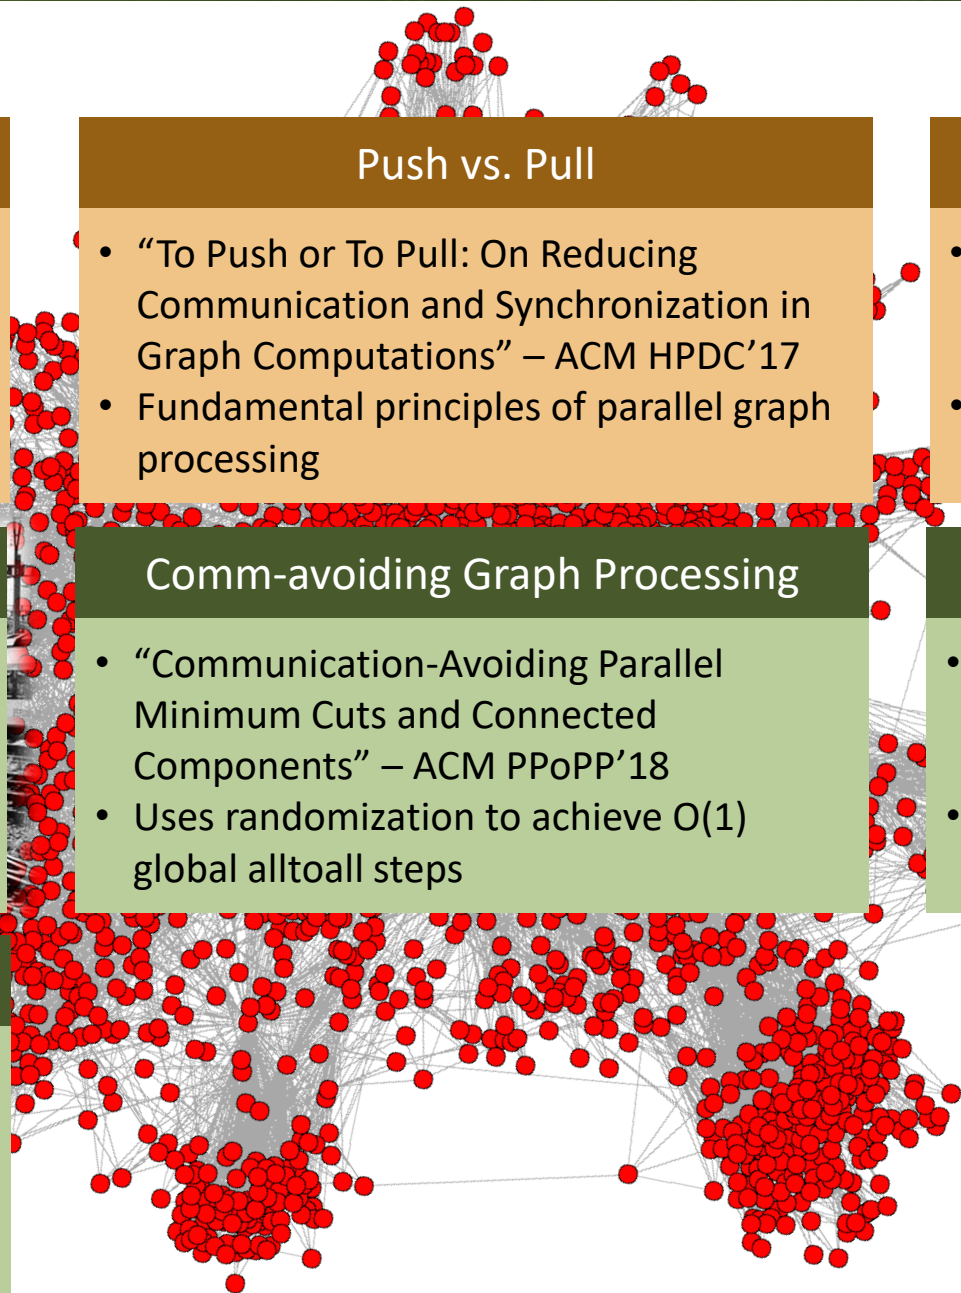
- “Communication-Avoiding Parallel Minimum Cuts and Connected Components” – ACM PPOPP’18
- Uses randomization to achieve $O(1)$ global alltoall steps

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

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



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

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

Streaming Graphs Survey

- “Survey and Taxonomy of Models and Algorithms for Streaming Graph Processing” – arXiv
- Overview of streaming algorithms, approximations, research gaps
 - Way forward for FPGA?

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

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

Streaming Graphs Survey

- “Survey and Taxonomy of Models and Algorithms for Streaming Graph Processing” – 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

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 Surveys

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

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

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

- “Survey and taxonomy of models and Algorithms for Streaming Graph Processing” – arXiv
- Overview of streaming algorithms, approximations, research gaps
 - Way forward for FPGA?

Algebraic Betweenness Centrality

...ing Betweenness Centrality using communication-Efficient Sparse Matrix multiplication” – ACM SC’17
...e on the algebraic view – complex, 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