



MACIEJ BESTA, DIMITRI STANOJEVIC, TIJANA ZIVIC, JAGPREET SINGH, MAURICE HOEROLD, TORSTEN HOEFLER

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




Large graphs...

Large graphs...




Large graphs...

Running on...



Large graphs...




Running on...




Large graphs...

Running on...



Used in...

Large graphs...





Running on...



Used in...



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



Large graphs...

Huge





Running on...



Used in...



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



Large graphs...





Running on...



Used in...



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



Large graphs...



en	wikipedia_euets_(en)	<input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	50,757,442	572,591,212
TW	Twitter (WWW)	<input checked="" type="checkbox"/> <input type="checkbox"/> <input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	41,652,230	1,468,365,182
TF	Twitter (MPI)	<input checked="" type="checkbox"/> <input type="checkbox"/> <input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	52,579,682	1,963,263,821
FR	Friendster	<input checked="" type="checkbox"/> <input type="checkbox"/> <input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	68,349,466	2,586,147,869
UL	UK domain (2007)	<input checked="" type="checkbox"/> <input type="checkbox"/> <input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	105,153,952	3,301,876,564



KONECT graph datasets

en	wikipedia edits (en)	<input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	50,757,442	512,591,212
TW	Twitter (WWW)	<input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	41,652,230	1,468,365,182
TF	Twitter (MPI)	<input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	52,579,682	1,963,263,821
FR	Friendster	<input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	68,349,466	2,586,147,869
UL	UK domain (2007)	<input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	105,153,952	3,301,876,564



KONECT graph datasets

Graph500 Benchmark



Top Ten from June 2018 BFS

RANK	MACHINE	VENDOR	INSTALLATION SITE	LOCATION	COUNTRY	YEAR	NUMBER OF NODES	NUMBER OF CORES	SCALE	GTEPS
1	K computer	Fujitsu	RIKEN Advanced Institute for Computational Science (AICS)	Kobe Hyogo	Japan	2011	82944	663552	40	38621.4
2	Sunway TaihuLight	NRCPC	National Supercomputing Center in Wuxi	Wuxi	China	2015	40768	10599680	40	23755.7
3	DOE/NNSA/LLNL Sequoia	IBM	Lawrence Livermore National Laboratory	Livermore CA	USA	2012	98304	1572864	41	23751
4	DOE/SC/Argonne National Laboratory Mira	IBM	Argonne National Laboratory	Chicago IL	USA	2012	49152	786432	40	14982

en	wikipedia edits (en)	<input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	50,757,442	572,591,212
TW	Twitter (WWW)	<input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	41,652,230	1,468,365,182
TF	Twitter (MPI)	<input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	52,579,682	1,963,263,821
FR	Friendster	<input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	68,349,466	2,586,147,869
UL	UK domain (2007)	<input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	105,153,952	3,301,876,564



KONECT graph datasets

Graph500 Benchmark

Webgraph datasets

Graph	Crawl date	Nodes	Arcs
uk-2014	2014	787801471	47614527250
eu-2015	2015	1070557254	91792261600
gsh-2015	2015	988490691	33877399152
uk-2014-host	2014	4769354	50829923
eu-2015-host	2015	11264052	386915963
gsh-2015-host	2015	68660142	1802747600
uk-2014-tpd	2014	1766010	18244650
eu-2015-tpd	2015	6650532	170145510
gsh-2015-tpd	2015	30809122	602119716
clueweb12	2012	978408098	42574107469
uk-2002	2002	18520486	298113762



LOCATION	COUNTRY	YEAR	NUMBER OF NODES	NUMBER OF CORES	SCALE	GTEPS
Hyogo	Japan	2011	82944	663552	40	38621.4
	China	2015	40768	10599680	40	23755.7
rmore CA	USA	2012	98304	1572864	41	23751
ago IL	USA	2012	49152	786432	40	14982

en	wikipedia_euets_(en)	<input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	50,757,442	572,591,212
TW	Twitter (WWW)	<input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	41,652,230	1,468,365,182
TF	Twitter (MPI)	<input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	52,579,682	1,963,263,821
FR	Friendster	<input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	68,349,466	2,586,147,869
UL	UK domain (2007)	<input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	105,153,952	3,301,876,564

Web data commons datasets

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



Graph500 Benchmark

Webgraph datasets



Graph	Crawl date	Nodes	Arcs
uk-2014	2014	787801471	47614527250
eu-2015	2015	1070557254	91792261600
gsh-2015	2015	988490691	33877399152
uk-2014-host	2014	4769354	50829923
eu-2015-host	2015	11264052	386915963
gsh-2015-host	2015	68660142	1802747600
uk-2014-tpd	2014	1766010	18244650
eu-2015-tpd	2015	6650532	170145510
gsh-2015-tpd	2015	30809122	602119716
clueweb12	2012	978408098	42574107469
uk-2002	2002	18520486	298113762

LOCATION	COUNTRY	YEAR	NUMBER OF NODES	NUMBER OF CORES	SCALE	GTEPS
Hyogo	Japan	2011	82944	663552	40	38621.4
	China	2015	40768	10599680	40	23755.7
rmore CA	USA	2012	98304	1572864	41	23751
ago IL	USA	2012	49152	786432	40	14982

Large graphs...





Running on...



Used in...



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



Large graphs...



Running on...





Compression incurs
expensive decompression



Used in...



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



Large graphs...



Running on...



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



Compression incurs expensive decompression



Used in...



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





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? 😊

Counting bounds.
They are **logarithmic**
(one needs at least $\log |S|$
bits to store an object
from an arbitrary set S)

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



The storage **lower bound** Ω



Which one? 😊



Counting bounds.
They are **logarithmic**
(one needs at least $\log |S|$
bits to store an object
from an arbitrary set S)



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

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

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



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

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



The storage **lower bound** Ω



Key idea

Which one? 😊



Counting bounds.
They are **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

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

Which one? 😊



Counting bounds. They are **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? 😊




Counting bounds. They are **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? 😊

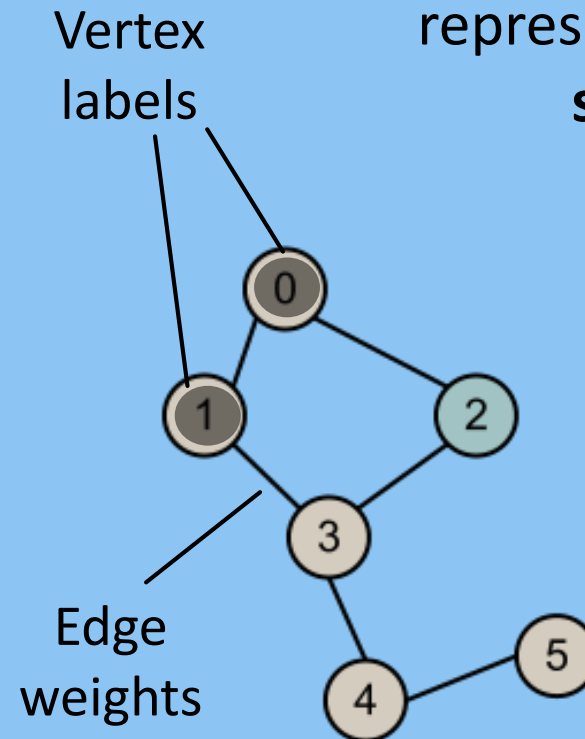


Counting bounds. They are **logarithmic** (one needs at least $\log |S|$ bits to store an object from an arbitrary set S)



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? 😊

Counting bounds. They are **logarithmic** (one needs at least $\log |S|$ bits to store an object from an arbitrary set S)

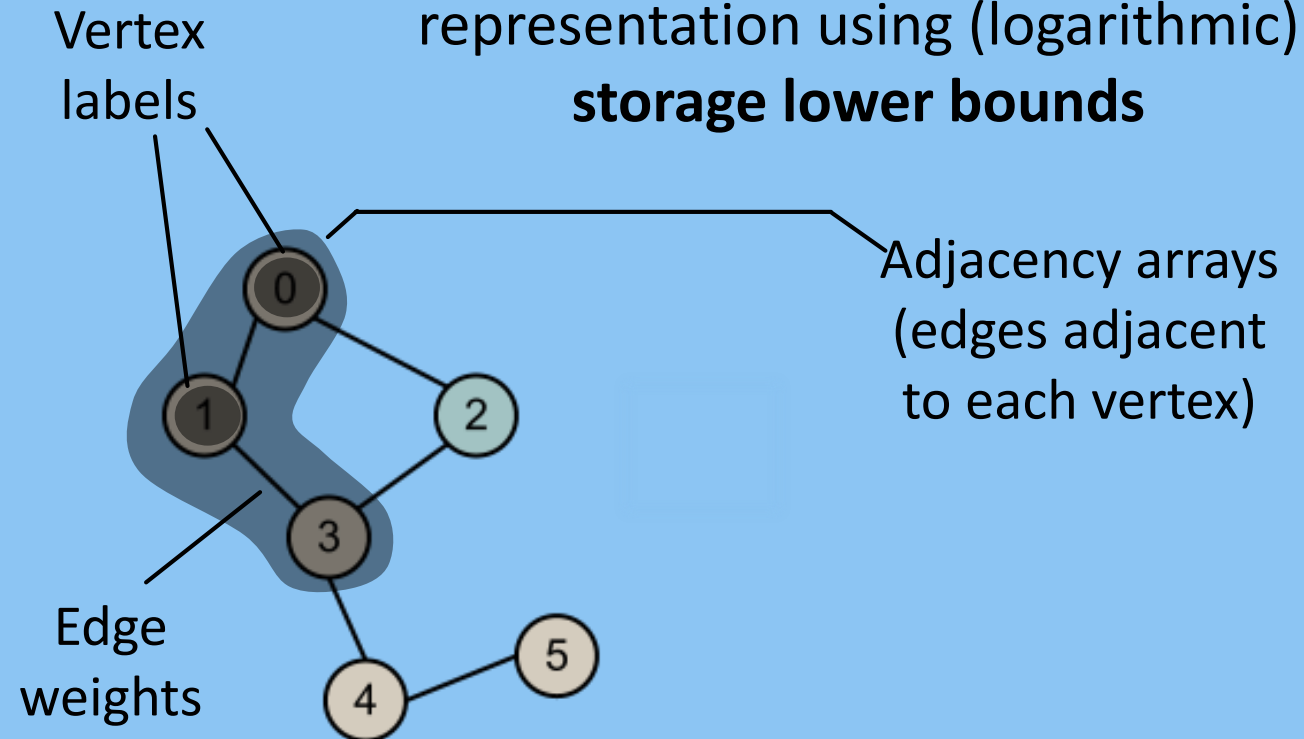
$$S = \{x_1, x_2, x_3, \dots\}$$

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



Key idea

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



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

The storage **lower bound** Ω

Which one? 😊

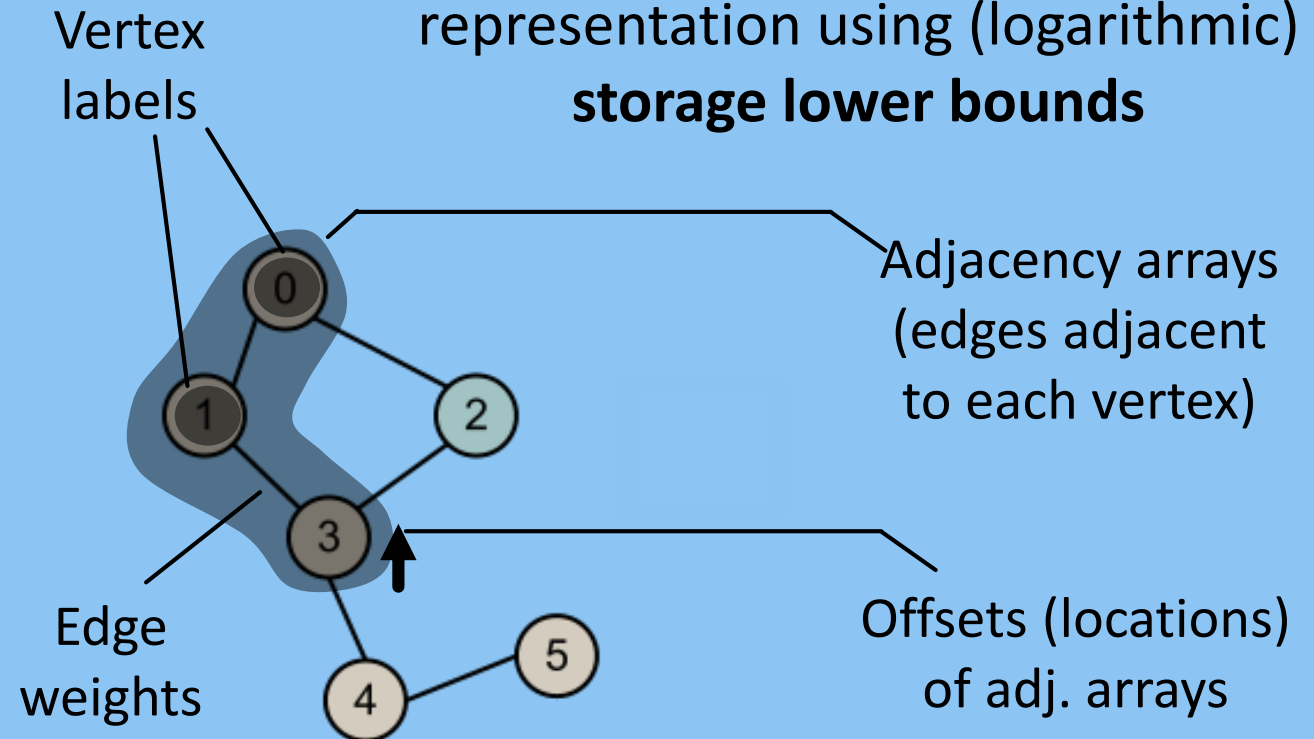
Counting bounds. They are **logarithmic** (one needs at least $\log |S|$ bits to store an object from an arbitrary set S)

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

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

Key idea

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



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

The storage **lower bound** Ω

Which one? 😊

Counting bounds. They are **logarithmic** (one needs at least $\log |S|$ bits to store an object from an arbitrary set S)


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

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

💡 Key idea

Log (Vertex labels)

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



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

The storage **lower bound** Ω

Which one? 😊

Counting bounds. They are **logarithmic** (one needs at least $\log |S|$ bits to store an object from an arbitrary set S)

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

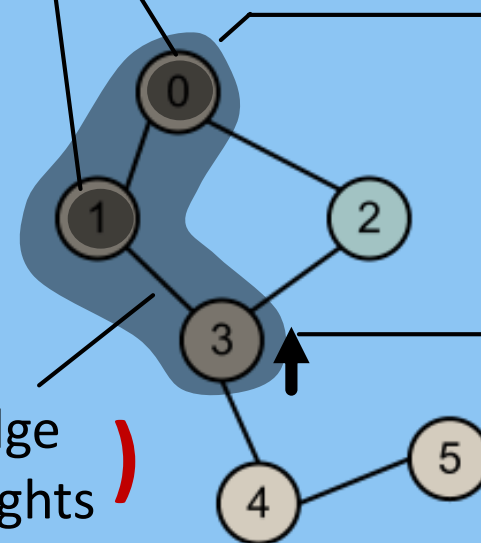
$x_1 \rightarrow$	0 ... 01
$x_2 \rightarrow$	0 ... 10
$x_3 \rightarrow$	0 ... 11
	...

Key idea

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

Log (Vertex labels)

Log (Edge weights)



Adjacency arrays (edges adjacent to each vertex)

Offsets (locations) of adj. arrays

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

The storage **lower bound** Ω

Which one? 😊

Counting bounds. They are **logarithmic** (one needs at least $\log |S|$ bits to store an object from an arbitrary set S)

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

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

Key idea

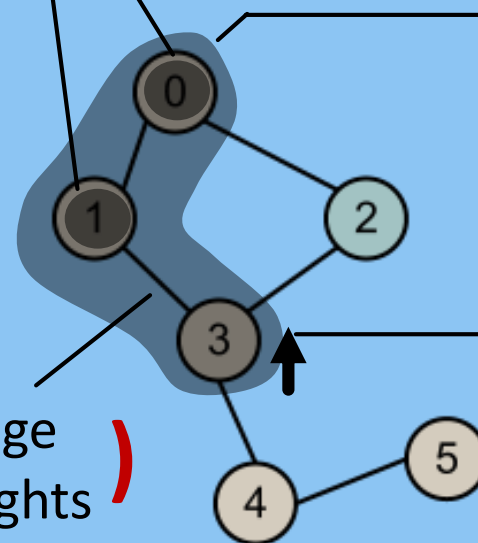
Encode different parts of a graph representation using (logarithmic) **storage lower bounds**

Log (Vertex labels)

Log ((edges adjacent to each vertex))

Log (Edge weights)

Offsets (locations) of adj. arrays



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

The storage **lower bound** Ω

Which one? 😊

Counting bounds. They are **logarithmic** (one needs at least $\log |S|$ bits to store an object from an arbitrary set S)

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

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

Key idea


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

Log (Vertex labels)


Log (Edge weights)

Log (Adjacency arrays (edges adjacent to each vertex))

Log (Offsets (locations) of adj. arrays)




ADJACENCY ARRAY GRAPH REPRESENTATION



ADJACENCY ARRAY GRAPH REPRESENTATION

Representation



ADJACENCY ARRAY GRAPH REPRESENTATION

Representation

0


1

2

3

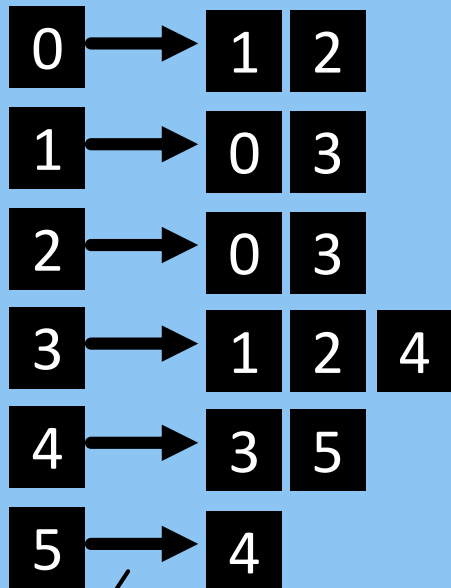
4

5




ADJACENCY ARRAY GRAPH REPRESENTATION

Representation




Offsets

Adjacency arrays
(edges adjacent
to each vertex)




ADJACENCY ARRAY GRAPH REPRESENTATION

Representation




Offsets

Adjacency arrays
(edges adjacent
to each vertex)



ADJACENCY ARRAY GRAPH REPRESENTATION


Representation



Offsets


Adjacency arrays
(edges adjacent
to each vertex)

Physical realization



ADJACENCY ARRAY GRAPH REPRESENTATION



Representation



Offsets


Adjacency arrays
(edges adjacent
to each vertex)

Physical realization



ADJACENCY ARRAY GRAPH REPRESENTATION



Representation



Offsets


Adjacency arrays
(edges adjacent
to each vertex)

Physical realization



ADJACENCY ARRAY GRAPH REPRESENTATION



Representation



Offsets


Adjacency arrays
(edges adjacent
to each vertex)

Physical realization



ADJACENCY ARRAY GRAPH REPRESENTATION


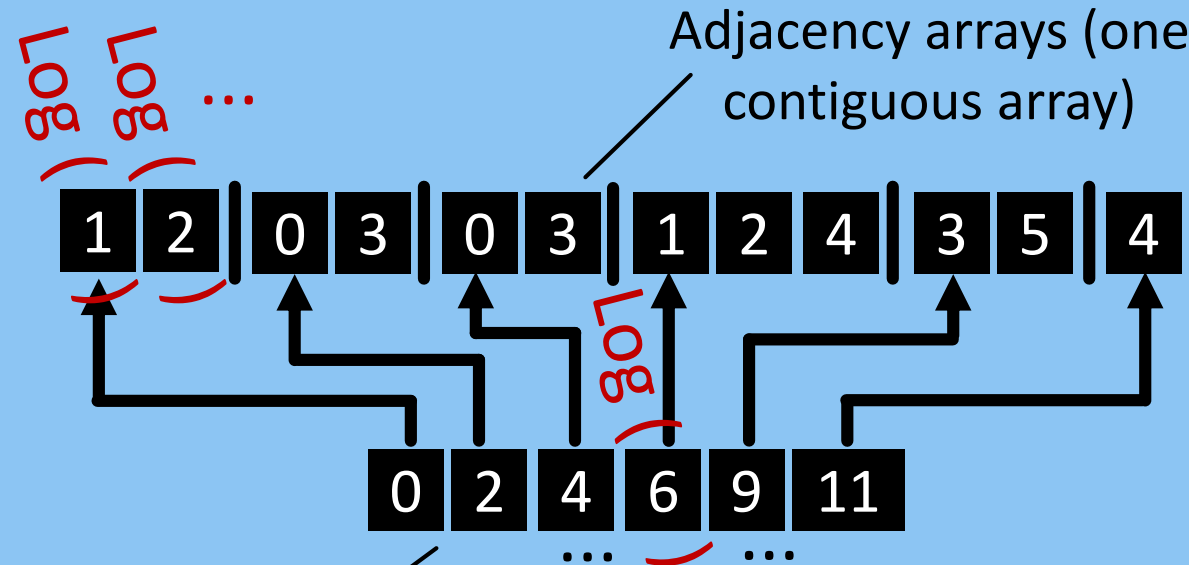
Representation



Offsets


Adjacency arrays
(edges adjacent
to each vertex)

Physical realization



ADJACENCY ARRAY GRAPH REPRESENTATION


Representation




Offsets

Adjacency arrays
(edges adjacent
to each vertex)

Physical realization




Offsets (another contiguous array)



ADJACENCY ARRAY GRAPH REPRESENTATION

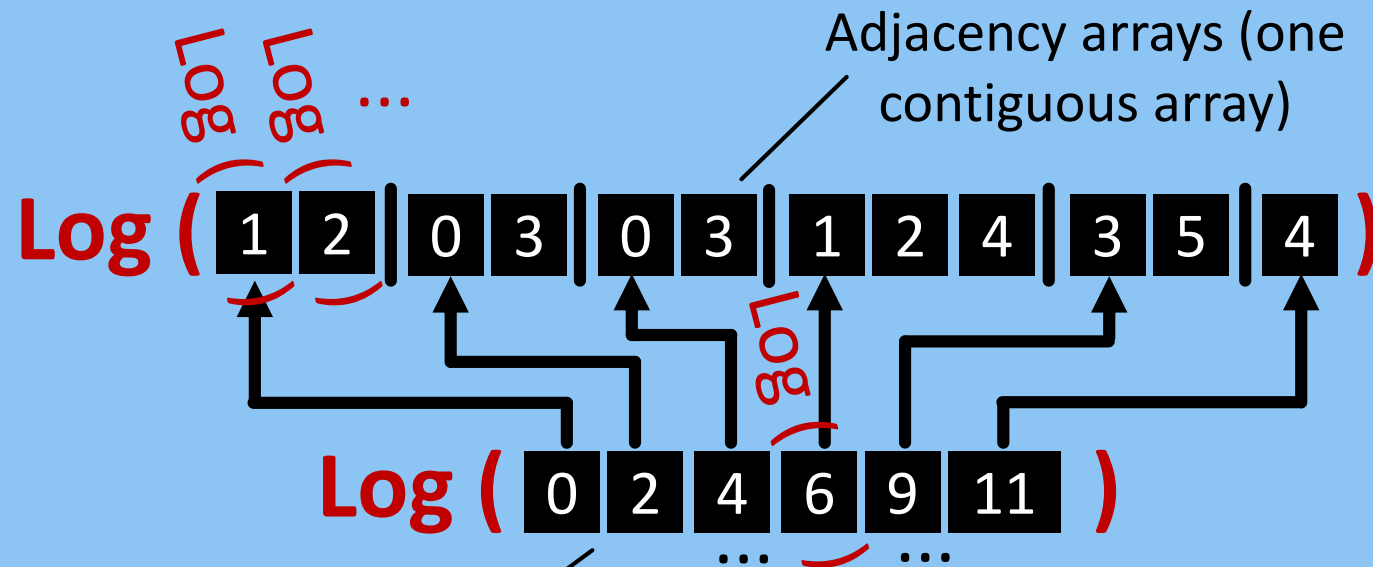
Representation




Offsets

Adjacency arrays
(edges adjacent
to each vertex)


Physical realization




Offsets (another contiguous array)




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)
[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)
[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: $\widehat{N}_v \ll n$

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




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: $\widehat{N}_v \ll n$

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




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



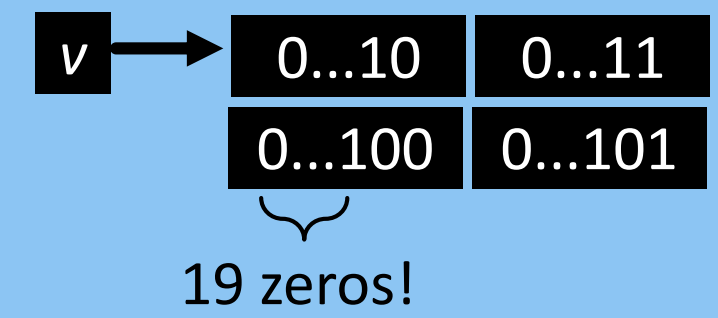
Lower bounds (local)

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



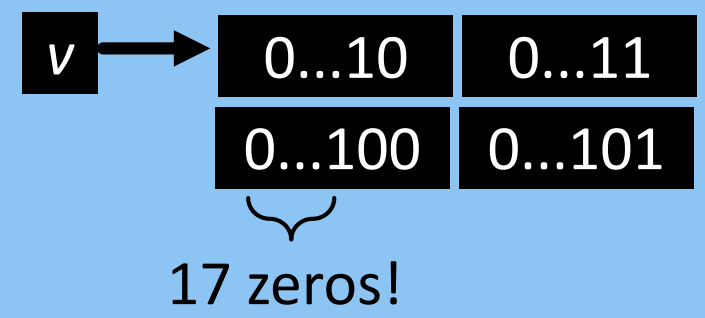
Lower bounds (local): problem

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:



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



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)

...Use Integer Linear Programming (ILP)!

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)

! ...Use Integer Linear Programming (ILP)!

Lower bounds (local) enhanced with ILP

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)

...Use Integer Linear Programming (ILP)!

Lower bounds (local) enhanced with ILP

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

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)

...Use Integer Linear Programming (ILP)! 

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) enhanced with ILP

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




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

...Use Integer Linear Programming (ILP)! 

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) enhanced with ILP

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

2 3 4 5 1M


v → 2 3 4 5 1M

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

...Use Integer Linear Programming (ILP)! 

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) enhanced with ILP

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

Permute(**2 3 4 5 1M**) =
 (simultaneously for all other neighborhoods)




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

...Use Integer Linear Programming (ILP)!

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) enhanced with ILP

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

$$\text{Permute}(\boxed{2} \boxed{3} \boxed{4} \boxed{5} \boxed{1M}) = \boxed{?} \boxed{?} \boxed{?} \boxed{?} \boxed{?}$$

(simultaneously for all other neighborhoods) $\leq 100?$




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

...Use Integer Linear Programming (ILP)!

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) enhanced with ILP

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



Permute(**2 3 4 5 1M**) = **? ? ? ? ?**

(simultaneously for all other neighborhoods)

$\leq 100?$

Heuristics:


$$\min \sum_{v \in V} \widehat{N}_v \frac{1}{d_v}$$

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

...Use Integer Linear Programming (ILP)!

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) enhanced with ILP

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



Permute(**2 3 4 5 1M**) = **? ? ? ? ?**

(simultaneously for all other neighborhoods)

$\leq 100?$

Heuristics:

$$\min \sum_{v \in V} \widehat{N}_v \frac{1}{d_v}$$


Inverse of the neighborhood size

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

...Use Integer Linear Programming (ILP)!

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) enhanced with ILP

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



Permute(**2 3 4 5 1M**) = **? ? ? ? ?**

(simultaneously for all other neighborhoods)

$\leq 100?$

Intuition:
maximum labels in new neighborhoods will be **smaller**

Heuristics:

$$\min \sum_{v \in V} \widehat{N}_v \frac{1}{d_v}$$

Inverse of the neighborhood size

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


Formal analyses

Symbols

\widehat{W} : max edge weight,

n : #vertices,

p, α, β : constants



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

\mathcal{F} Formal analyses


Power-law graphs

Symbols

\widehat{W} : max edge weight,

n : #vertices,

p, α, β : constants



Random uniform graphs

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

Formal analyses

Power-law graphs

The probability that a vertex has degree d is:


$$\alpha d^\beta$$

Symbols

\widehat{W} : max edge weight,

n : #vertices,

p, α, β : constants



Random uniform graphs


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

\mathcal{F} Formal analyses

Power-law graphs

The probability that a vertex has degree d is:

$$\alpha d^\beta$$




Symbols

\widehat{W} : max edge weight,

n : #vertices,

p, α, β : constants



Random uniform graphs


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

\mathcal{F} Formal analyses

Power-law graphs

The probability that a vertex has degree d is:

$$\alpha d^{-\beta}$$




Symbols

\widehat{W} : max edge weight,

n : #vertices,

p, α, β : constants



Random uniform graphs

The probability that a vertex has degree d is:

$$pd$$


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

\mathcal{F} Formal analyses

Power-law graphs

The probability that a vertex has degree d is:

$$\alpha d^\beta$$




Symbols

\widehat{W} : max edge weight,

n : #vertices,


p, α, β : constants



Random uniform graphs

The probability that a vertex has degree d is:

$$pd$$




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

fff Formal analyses

Power-law graphs

The probability that a vertex has degree d is:

$$\alpha d^{-\beta}$$



Expected size of the adjacency array


$$E[|\mathcal{A}|] \approx \frac{\alpha}{2-\beta} \left(\left(\frac{\alpha n \log n}{\beta-1} \right)^{\frac{2-\beta}{\beta-1}} - 1 \right) \left(\lceil \log n \rceil + \lceil \log \widehat{W} \rceil \right)$$

Symbols

\widehat{W} : max edge weight,

n : #vertices,

p, α, β : constants




Random uniform graphs

The probability that a vertex has degree d is:

$$pd$$

Expected size of the adjacency array



$$E[|\mathcal{A}|] = \left(\lceil \log n \rceil + \lceil \log \widehat{W} \rceil \right) pn^2$$

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

fff Formal analyses: more
(check the paper 😊)

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

fff Formal analyses: more (check the paper 😊)

$$|\mathcal{A}| = \sum_{v \in V} \left(d_v \lceil \log \widehat{N}_v \rceil + \lceil \log \log \widehat{N}_v \rceil \right)$$

$$|\mathcal{A}| = n \left\lceil \log \frac{n}{\mathcal{H}} \right\rceil + \mathcal{H} \lceil \log \mathcal{H} \rceil$$

$$|\mathcal{A}| = \sum_{v \in V} \left(d_v \left(\lceil \log \widehat{N}_v \rceil + \lceil \log \widehat{\mathcal{W}} \rceil \right) + \lceil \log \log \widehat{N}_v \rceil + \lceil \log \log \widehat{\mathcal{W}} \rceil \right)$$

$$E[|\mathcal{A}|] \approx \frac{\alpha}{2-\beta} \left(\left(\frac{\alpha n \log n}{\beta-1} \right)^{\frac{2-\beta}{\beta-1}} - 1 \right) \left(\lceil \log n \rceil + \lceil \log \widehat{\mathcal{W}} \rceil \right)$$

$$E[|\mathcal{O}|] = n \left\lceil \log \left(2pn^2 \right) \right\rceil = n \lceil \log 2p + 2 \log n \rceil$$

$$\forall v, u \in V (u \in N_v) \Rightarrow \lceil \mathcal{N}(u) \rceil \leq \widehat{N}_v$$

$$|\mathcal{A}| = \sum_{v \in V} \left(d_v \lceil \log \widehat{N}_v \rceil + \lceil \log \log \widehat{N}_v \rceil \right)$$

$$|\mathcal{A}| = 2m \left(\lceil \log n \rceil + \lceil \log \widehat{\mathcal{W}} \rceil \right)$$

$$E[|\mathcal{A}|] = \left(\lceil \log n \rceil + \lceil \log \widehat{\mathcal{W}} \rceil \right) pn^2$$

1 **Log** (Vertex labels), **Log** (Edge weights)
 Formal analyses: more (check the paper 😊)

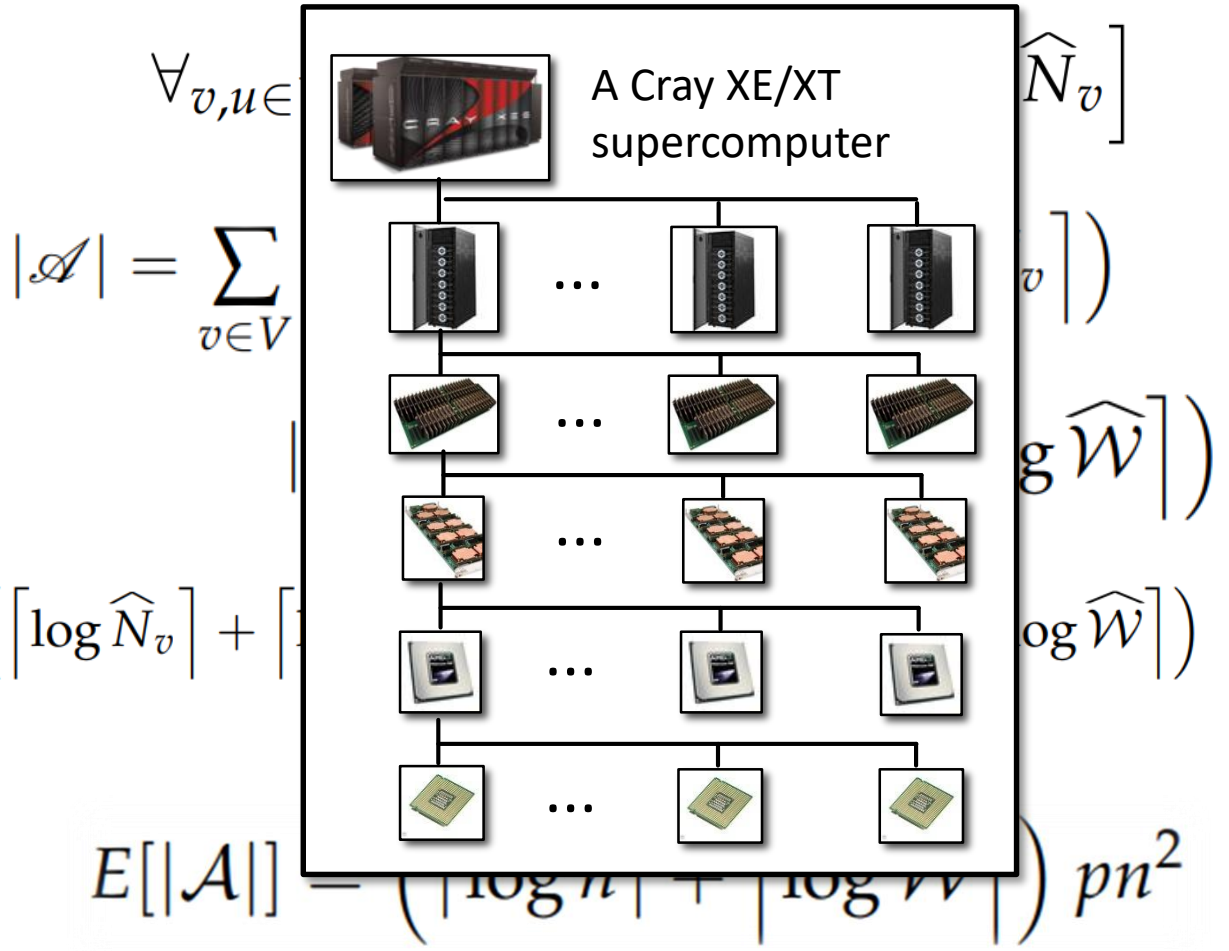
$$E[|\mathcal{O}|] = n \lceil \log(2pn^2) \rceil = n \lceil \log 2p + 2 \log n \rceil$$

$$|\mathcal{A}| = \sum_{v \in V} \left(d_v \lceil \log \widehat{N}_v \rceil + \lceil \log \log \widehat{N}_v \rceil \right)$$

$$|\mathcal{A}| = n \lceil \log \frac{n}{\mathcal{H}} \rceil + \mathcal{H} \lceil \log \mathcal{H} \rceil$$


$$|\mathcal{A}| = \sum_{v \in V} \left(d_v \left(\lceil \log \widehat{N}_v \rceil + \lceil \log \log \widehat{N}_v \rceil \right) \right)$$

$$E[|\mathcal{A}|] \approx \frac{\alpha}{2-\beta} \left(\left(\frac{\alpha n \log n}{\beta-1} \right)^{\frac{2-\beta}{\beta-1}} - 1 \right) \left(\lceil \log n \rceil + \lceil \log \widehat{W} \rceil \right)$$



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


 Key methods



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

 Key methods


Use the **BEXTR** bitwise operation to help extract an arbitrary sequence of bits



1 $\text{Log} \left(\begin{matrix} \text{Vertex} \\ \text{labels} \end{matrix} \right), \text{Log} \left(\begin{matrix} \text{Edge} \\ \text{weights} \end{matrix} \right)$

🔧 Key methods

Use the **BEXTR** bitwise operation to help extract an arbitrary sequence of bits




```


1 /* v_ID is an opaque type for IDs of vertices. */
2 v_ID Ni,v(v_ID v, int32_t i, int64_t* O, int64_t* A, int8_t s){
3   int64_t exactBitOffset = s * (O[v] + i);
4   int8_t* address = (int8_t*) A + (exactBitOffset >> 3);
5   int64_t distance = exactBitOffset & 7;
6   int64_t value = ((int64_t*) (address))[0];
7   return _bextr_u64(value, distance, s); }

```

1 $\text{Log} \left(\begin{matrix} \text{Vertex} \\ \text{labels} \end{matrix} \right), \text{Log} \left(\begin{matrix} \text{Edge} \\ \text{weights} \end{matrix} \right)$

 Key methods

Use the **BEXTR** bitwise operation to help extract an arbitrary sequence of bits 



Return i -th neighbor of vertex v


```


1 /* v_ID is an opaque type for IDs of vertices. */
2 v_ID Ni,v(v_ID v, int32_t i, int64_t* O, int64_t* A, int8_t s){
3   int64_t exactBitOffset = s * (O[v] + i);
4   int8_t* address = (int8_t*) A + (exactBitOffset >> 3);
5   int64_t distance = exactBitOffset & 7;
6   int64_t value = ((int64_t*) (address))[0];
7   return _bextr_u64(value, distance, s); }

```

1 $\text{Log} \left(\begin{matrix} \text{Vertex} \\ \text{labels} \end{matrix} \right), \text{Log} \left(\begin{matrix} \text{Edge} \\ \text{weights} \end{matrix} \right)$

 Key methods

Use the **BEXTR** bitwise operation to help extract an arbitrary sequence of bits 



Return i -th neighbor of vertex v

Pointer to the offset array


```


1 /* v_ID is an opaque type for IDs of vertices. */
2 v_ID Ni,v(v_ID v, int32_t i, int64_t* O, int64_t* A, int8_t s){
3   int64_t exactBitOffset = s * (O[v] + i);
4   int8_t* address = (int8_t*) A + (exactBitOffset >> 3);
5   int64_t distance = exactBitOffset & 7;
6   int64_t value = ((int64_t*) (address))[0];
7   return _bextr_u64(value, distance, s); }

```

1 $\text{Log} \left(\begin{matrix} \text{Vertex} \\ \text{labels} \end{matrix} \right), \text{Log} \left(\begin{matrix} \text{Edge} \\ \text{weights} \end{matrix} \right)$

 Key methods

Use the **BEXTR** bitwise operation to help extract an arbitrary sequence of bits 



Return i -th neighbor of vertex v

Pointer to the offset array

Pointer to the adjacency array


```


1 /* v_ID is an opaque type for IDs of vertices. */
2 v_ID Ni,v(v_ID v, int32_t i, int64_t* O, int64_t* A, int8_t s){
3   int64_t exactBitOffset = s * (O[v] + i);
4   int8_t* address = (int8_t*) A + (exactBitOffset >> 3);
5   int64_t distance = exactBitOffset & 7;
6   int64_t value = ((int64_t*) (address))[0];
7   return _bextr_u64(value, distance, s); }

```


1 $\text{Log} \left(\begin{matrix} \text{Vertex} \\ \text{labels} \end{matrix} \right), \text{Log} \left(\begin{matrix} \text{Edge} \\ \text{weights} \end{matrix} \right)$

 Key methods

Use the **BEXTR** bitwise operation to help extract an arbitrary sequence of bits 



Return i -th neighbor of vertex v

Pointer to the offset array

Pointer to the adjacency array

$s = \lceil \log n \rceil$


```


1 /* v_ID is an opaque type for IDs of vertices. */
2 v_ID Ni,v(v_ID v, int32_t i, int64_t* O, int64_t* A, int8_t s){
3   int64_t exactBitOffset = s * (O[v] + i);
4   int8_t* address = (int8_t*) A + (exactBitOffset >> 3);
5   int64_t distance = exactBitOffset & 7;
6   int64_t value = ((int64_t*) (address))[0];
7   return _bextr_u64(value, distance, s); }

```

1 $\text{Log} \left(\begin{matrix} \text{Vertex} \\ \text{labels} \end{matrix} \right), \text{Log} \left(\begin{matrix} \text{Edge} \\ \text{weights} \end{matrix} \right)$

 Key methods

Use the **BEXTR** bitwise operation to help extract an arbitrary sequence of bits 



Return i -th neighbor of vertex v

Derive exact offset (in bits) to the neighbor label

Pointer to the offset array

Pointer to the adjacency array

$s = \lceil \log n \rceil$


```

1 /* v_ID is an opaque type for IDs of vertices. */
2 v_ID Ni,v(v_ID v, int32_t i, int64_t* O, int64_t* A, int8_t s){
3   int64_t exactBitOffset = s * (O[v] + i);
4   int8_t* address = (int8_t*) A + (exactBitOffset >> 3);
5   int64_t distance = exactBitOffset & 7;
6   int64_t value = ((int64_t*) (address))[0];
7   return _bextr_u64(value, distance, s); }

```

1 $\text{Log} \left(\begin{matrix} \text{Vertex} \\ \text{labels} \end{matrix} \right), \text{Log} \left(\begin{matrix} \text{Edge} \\ \text{weights} \end{matrix} \right)$
 Key methods

Use the **BEXTR** bitwise  operation to help extract an arbitrary sequence of bits



Return i -th neighbor of vertex v

Derive exact offset (in bits) to the neighbor label

Pointer to the offset array

Pointer to the adjacency array

$s = \lceil \log n \rceil$

```


1 /* v_ID is an opaque type for IDs of vertices. */
2 v_ID Ni,v(v_ID v, int32_t i, int64_t* O, int64_t* A, int8_t s){
3   int64_t exactBitOffset = s * (O[v] + i);
4   int8_t* address = (int8_t*) A + (exactBitOffset >> 3);
5   int64_t distance = exactBitOffset & 7;
6   int64_t value = ((int64_t*) (address))[0];
7   return _bextr_u64(value, distance, s); }

```

Get the closest byte alignment

1 $\text{Log} \left(\begin{matrix} \text{Vertex} \\ \text{labels} \end{matrix} \right), \text{Log} \left(\begin{matrix} \text{Edge} \\ \text{weights} \end{matrix} \right)$
 Key methods

Use the **BEXTR** bitwise operation to help extract an arbitrary sequence of bits 



Return i -th neighbor of vertex v

Derive exact offset (in bits) to the neighbor label

Pointer to the offset array

Pointer to the adjacency array

$s = \lceil \log n \rceil$

```

1 /* v_ID is an opaque type for IDs of vertices. */
2 v_ID Ni,v(v_ID v, int32_t i, int64_t* O, int64_t* A, int8_t s){
3   int64_t exactBitOffset = s * (O[v] + i);
4   int8_t* address = (int8_t*) A + (exactBitOffset >> 3);
5   int64_t distance = exactBitOffset & 7;
6   int64_t value = ((int64_t*) (address))[0];
7   return _bextr_u64(value, distance, s); }

```


Get the closest byte alignment

Get the distance from the byte alignment

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

Key methods

Use the **BEXTR** bitwise operation to help extract an arbitrary sequence of bits



Return i -th neighbor of vertex v

Derive exact offset (in bits) to the neighbor label

Pointer to the offset array

Pointer to the adjacency array

$s = \lceil \log n \rceil$

```

1 /* v_ID is an opaque type for IDs of vertices. */
2 v_ID Ni,v(v_ID v, int32_t i, int64_t* O, int64_t* A, int8_t s){
3   int64_t exactBitOffset = s * (O[v] + i);
4   int8_t* address = (int8_t*) A + (exactBitOffset >> 3);
5   int64_t distance = exactBitOffset & 7;
6   int64_t value = ((int64_t*) (address))[0];
7   return _bextr_u64(value, distance, s); }

```

Get the closest byte alignment


Get the distance from the byte alignment

Access the derived 64-bit value

1 $\text{Log} \left(\begin{matrix} \text{Vertex} \\ \text{labels} \end{matrix} \right), \text{Log} \left(\begin{matrix} \text{Edge} \\ \text{weights} \end{matrix} \right)$

Key methods

Use the **BEXTR** bitwise operation to help extract an arbitrary sequence of bits



Return i -th neighbor of vertex v

Derive exact offset (in bits) to the neighbor label

Pointer to the offset array

Pointer to the adjacency array

$s = \lceil \log n \rceil$

```

1 /* v_ID is an opaque type for IDs of vertices. */
2 v_ID Ni,v(v_ID v, int32_t i, int64_t* O, int64_t* A, int8_t s){
3   int64_t exactBitOffset = s * (O[v] + i);
4   int8_t* address = (int8_t*) A + (exactBitOffset >> 3);
5   int64_t distance = exactBitOffset & 7;
6   int64_t value = ((int64_t*) (address))[0];
7   return _bextr_u64(value, distance, s); }

```

Get the closest byte alignment


Get the distance from the byte alignment

Shift the derived 64-bit value by d bits and mask it with BEXTR

Access the derived 64-bit value


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* →
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 $x \rightarrow$ 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

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.

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)



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)

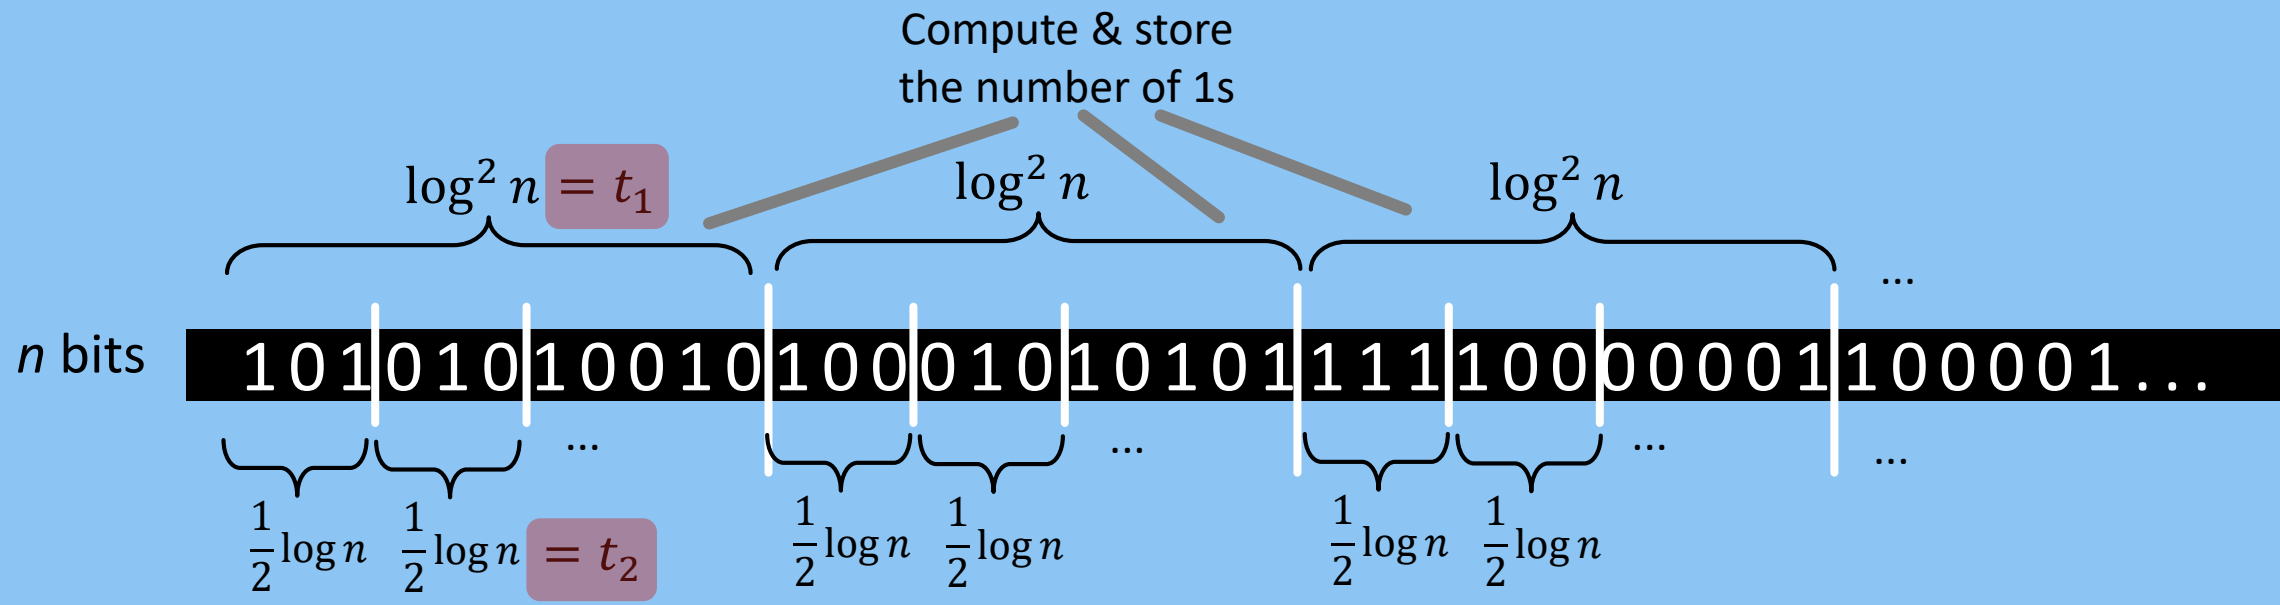
...Encode the resulting bit vectors as succinct bit vectors [1]



Succinct bit vectors

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

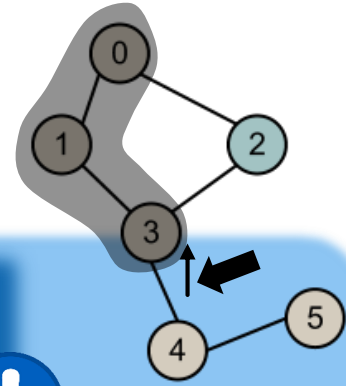
= small + fast (hopefully)



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

2 Log (Offset structure)

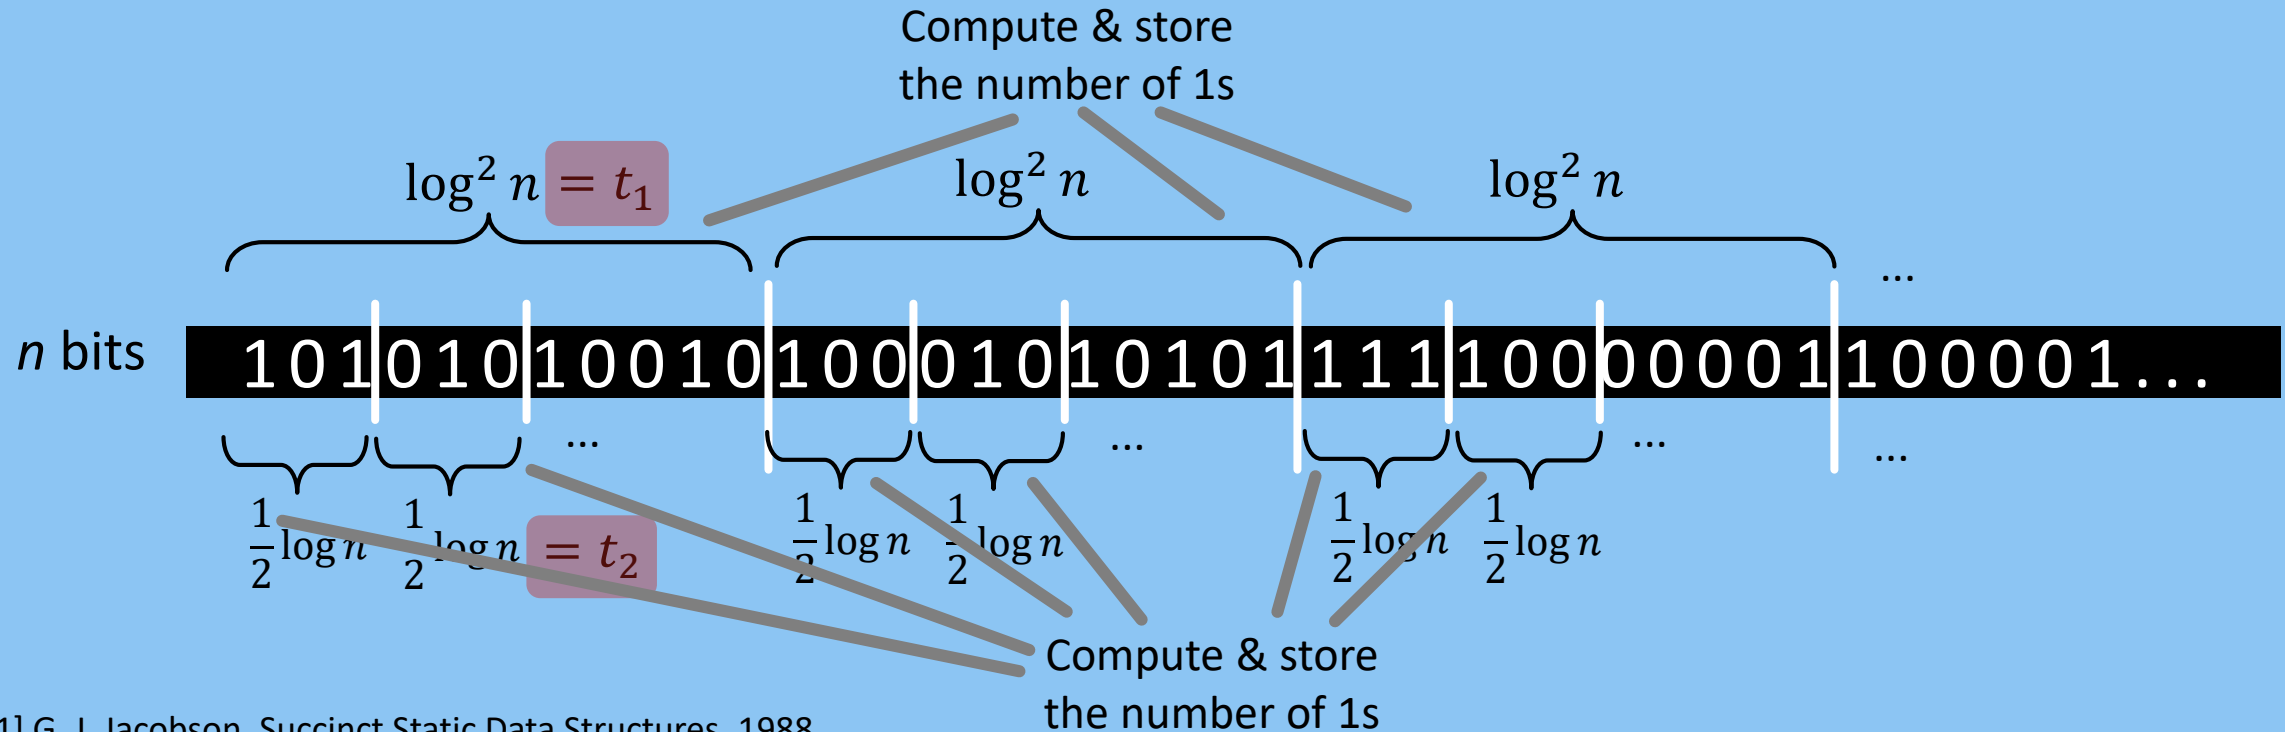
...Encode the resulting bit vectors as succinct bit vectors [1]



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] **!**

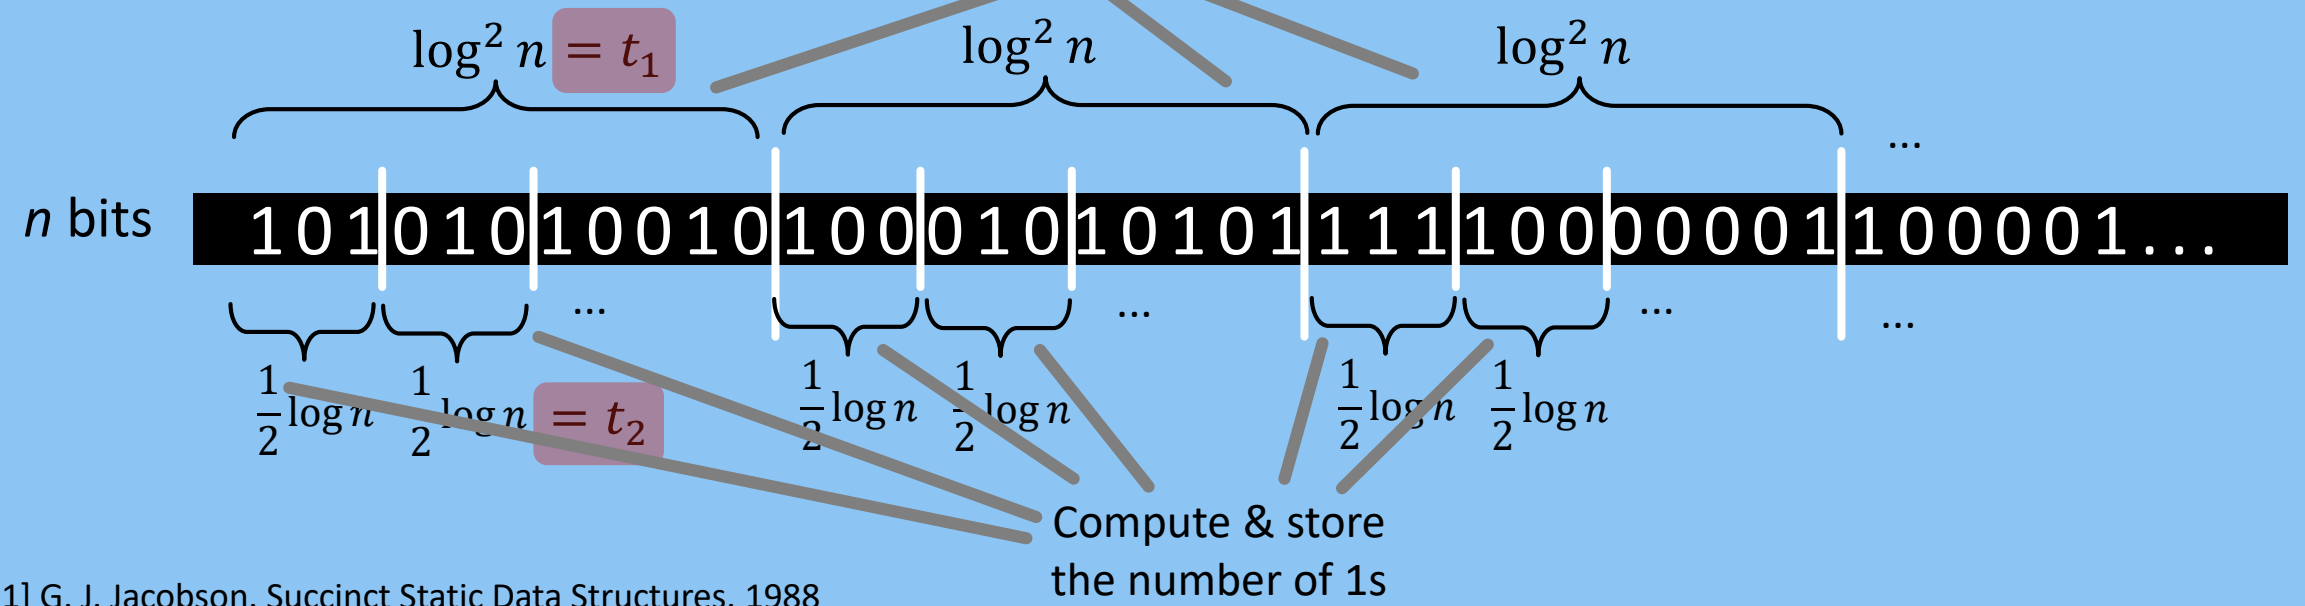


Succinct bit vectors

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 $= O\left(\frac{n}{t_1} \log n\right) = O\left(\frac{n}{\log n}\right) = o(n)$

n bits




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




2 Log (Offset structure)


...Encode the resulting bit vectors as
succinct bit vectors




Formal analyses



2 Log (Offset structure)

...Encode the resulting bit vectors as **succinct bit vectors** 




Formal analyses

\mathcal{O}	ID	Asymptotic size [bits]	Exact size [bits]	select or $\mathcal{O}[v]$
Pointer array	ptrW	$O(Wn)$	$W(n + 1)$	$O(1)$
Plain [44]	bvPL	$O\left(\frac{Wm}{B}\right)$	$\frac{2Wm}{B}$	$O(1)$
Interleaved [44]	bvIL	$O\left(\frac{Wm}{B} + \frac{Wm}{L}\right)$	$2Wm\left(\frac{1}{B} + \frac{64}{L}\right)$	$O\left(\log \frac{Wm}{B}\right)$
Entropy based [31, 78]	bvEN	$O\left(\frac{Wm}{B} \log \frac{Wm}{B}\right)$	$\approx \log\left(\frac{2Wm}{Bn}\right)$	$O\left(\log \frac{Wm}{B}\right)$
Sparse [76]	bvSD	$O\left(n + n \log \frac{Wm}{Bn}\right)$	$\approx n\left(2 + \log \frac{2Wm}{Bn}\right)$	$O(1)$
B-tree based [1]	bvBT	$O\left(\frac{Wm}{B}\right)$	$\approx 1.1 \cdot \frac{2Wm}{B}$	$O(\log n)$
Gap-compressed [1]	bvGC	$O\left(\frac{Wm}{B} \log \frac{Wm}{Bn}\right)$	$\approx 1.3 \cdot \frac{2Wm}{B} \log \frac{2Wm}{Bn}$	$O(\log n)$

2 Log (Offset structure)

...Encode the resulting bit vectors as **succinct bit vectors**




Formal analyses

Check the paper for details 😊

\mathcal{O}	ID	Asymptotic size [bits]	Exact size [bits]	select or $\mathcal{O}[v]$
Pointer array	ptrW	$O(Wn)$	$W(n + 1)$	$O(1)$
Plain [44]	bvPL	$O\left(\frac{Wm}{B}\right)$	$\frac{2Wm}{B}$	$O(1)$
Interleaved [44]	bvIL	$O\left(\frac{Wm}{B} + \frac{Wm}{L}\right)$	$2Wm\left(\frac{1}{B} + \frac{64}{L}\right)$	$O\left(\log \frac{Wm}{B}\right)$
Entropy based [31, 78]	bvEN	$O\left(\frac{Wm}{B} \log \frac{Wm}{B}\right)$	$\approx \log\left(\frac{2Wm}{Bn}\right)$	$O\left(\log \frac{Wm}{B}\right)$
Sparse [76]	bvSD	$O\left(n + n \log \frac{Wm}{Bn}\right)$	$\approx n\left(2 + \log \frac{2Wm}{Bn}\right)$	$O(1)$
B-tree based [1]	bvBT	$O\left(\frac{Wm}{B}\right)$	$\approx 1.1 \cdot \frac{2Wm}{B}$	$O(\log n)$
Gap-compressed [1]	bvGC	$O\left(\frac{Wm}{B} \log \frac{Wm}{Bn}\right)$	$\approx 1.3 \cdot \frac{2Wm}{B} \log \frac{2Wm}{Bn}$	$O(\log n)$

2 Log (Offset structure)

...Encode the resulting bit vectors as **succinct bit vectors**




Formal analyses

Check the paper for details 😊

\mathcal{O}	ID	Asymptotic size [bits]	Exact size [bits]	select or $\mathcal{O}[v]$
Pointer array	ptrW	$O(Wn)$	$W(n + 1)$	$O(1)$
Plain [44]		$O(Wm)$	$2Wm$	$O(1)$
Interleaved [44]				$O\left(\log \frac{Wm}{B}\right)$
Entropy based [31]				$O\left(\log \frac{Wm}{B}\right)$
Sparse [76]				
B-tree based [1]	bvBT	$O\left(\frac{Wm}{B}\right)$	$\approx 1.1 \cdot \frac{Wm}{B}$	$O(\log n)$
Gap-compressed [1]	bvGC	$O\left(\frac{Wm}{B} \log \frac{Wm}{Bn}\right)$	$\approx 1.3 \cdot \frac{2Wm}{B} \log \frac{2Wm}{Bn}$	$O(\log n)$


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

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)

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

(simultaneously for all other neighborhoods)



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)

(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

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

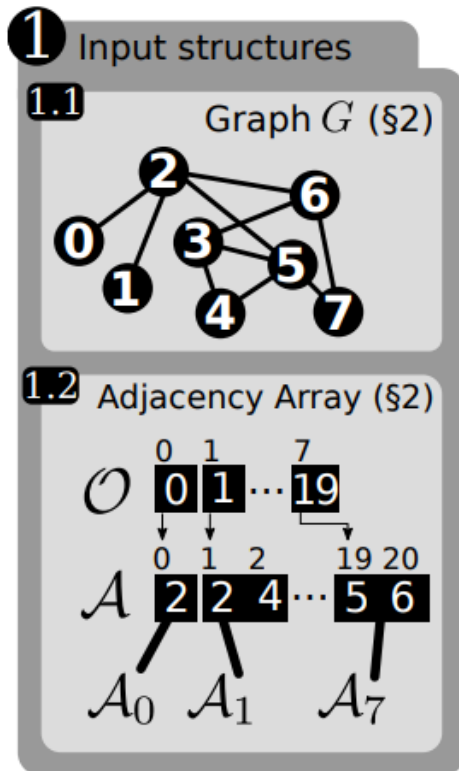
(2) Encode new labels with gap encoding (differences between consecutive labels instead of full labels)



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

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

Remove leading bits (simple bit packing)

2.1 (§3.1) U

2.3 ...globally (§3.2.1)

2.7 (§3.5) Analyze

2.4 ...locally (§3.2.2)

2.10 (§3.8) Ensure

2.5 ...on DM systems



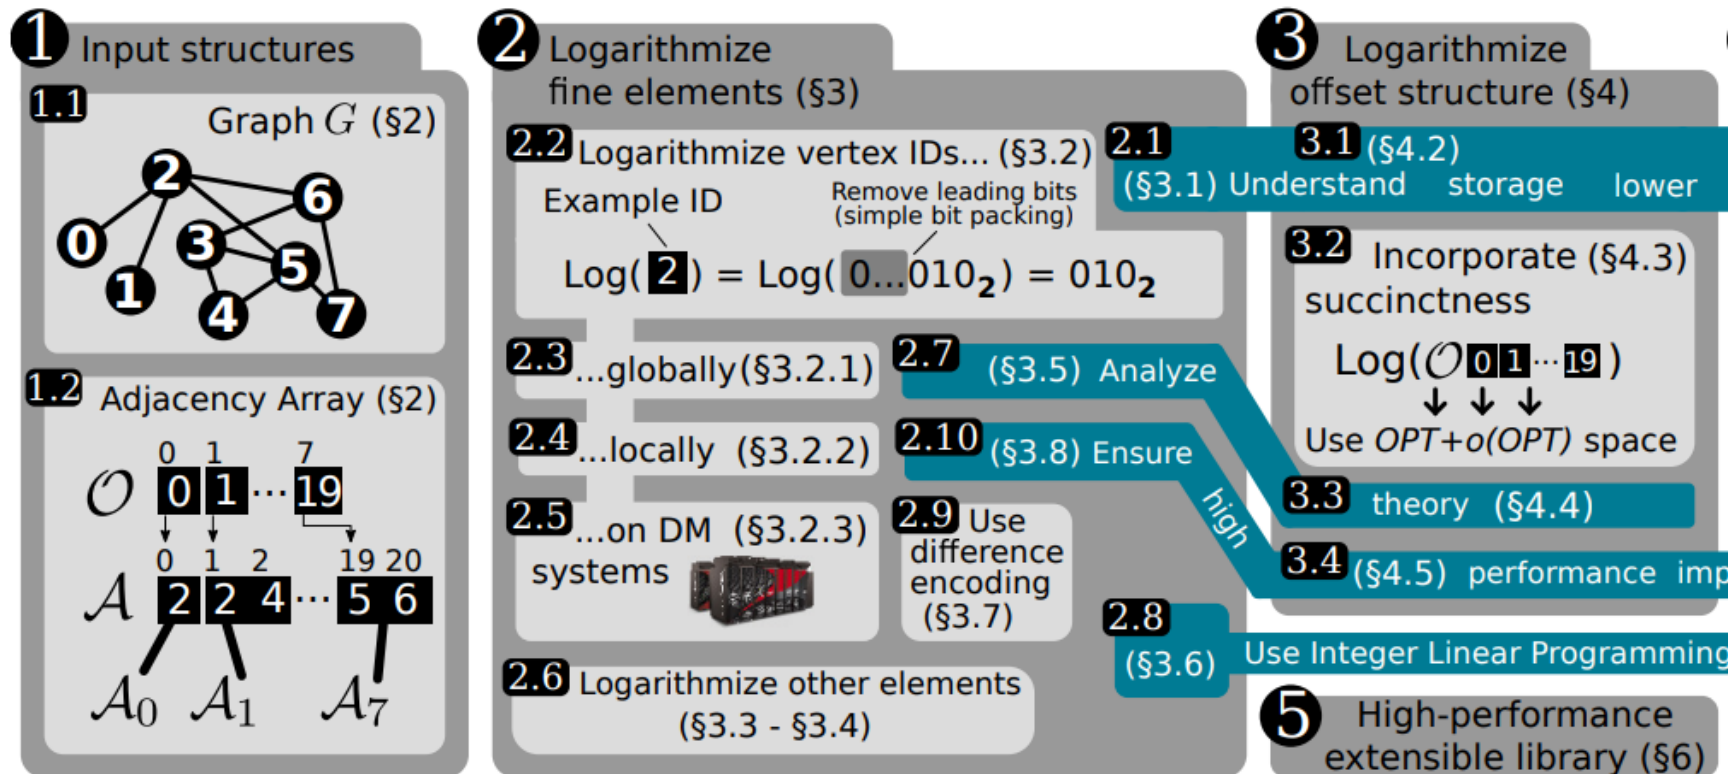
2.9 Use difference encoding (§3.7)

2.8 (§3.6) U

2.6 Logarithmize other elements (§3.3 - §3.4)

high


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)

2.2 Logarithmize vertex IDs... (§3.2)

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

Remove leading bits (simple bit packing)

2.3 ...globally (§3.2.1)

2.4 ...locally (§3.2.2)

2.5 ...on DM systems



2.6 Logarithmize other elements (§3.3 - §3.4)

2.7 (§3.5) Analyze

2.10 (§3.8) Ensure

2.9 Use difference encoding (§3.7)

2.8

(§3.6) Use Integer Linear Programming (ILP)

3 Logarithmize offset structure (§4)

2.1 (2.1) Understand storage lower bounds

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 implementation

4 Logarithmize adjacency structure (§5)

4.1 (§5.1) Unify 4.3-4.4 with $P+T$ (§5.2)

4.3 Incorporate (§5.3) recursive bisectioning

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

4.6 Use DM (§5.4)

4.8 (§6)

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

5 High-performance extensible library (§6)


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

Remove leading bits (simple bit packing)

2.3 ...globally (§3.2.1)

2.4 ...locally (§3.2.2)

2.5 ...on DM (§3.2.3) systems



2.6 Logarithmize other elements (§3.3 - §3.4)

2.7 (§3.5) Analyze

2.10 (§3.8) Ensure

2.9 Use difference encoding (§3.7)

2.8 (§3.6)

3 Logarithmize offset structure (§4)

3.1 (§4.2) Understand storage lower

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

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

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

4.6 Use DM (§5.4)

4.8 (§6)

4.4 (§5.3.1) ...use RB


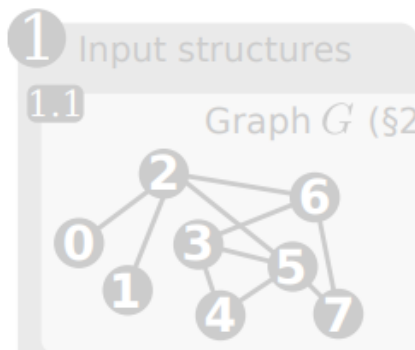
4.4 (§5.3.1) ...use RB

4.5 (§5.3.2) ...use BRB

4.7

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

OVERVIEW OF FULL LOG(GRAPH) DESIGN



2 Logarithmize fine elements (§3)

2.1 Logarithmize vertex IDs... (§3.2)

Example ID $\text{Log}(2) = \text{Log}(010) = 010$

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)

2.5 Use difference encoding (§3.7)

2.6 Logarithmize other elements (§3.3 - §3.4)

! Looks complex 😊

3 Logarithmize offset structure (§4)

3.1 Understand storage lower bounds (§4.2)

3.2 Incorporate (§4.3)

3.3 theory (§4.4)

3.4 performance implementation (§4.5)

3.5 High-performance extensible library (§6)

4 Logarithmize adjacency structure (§5)

4.1 Unify bounds (§5.1)

4.2 Unify with $P+T$ (§5.2)

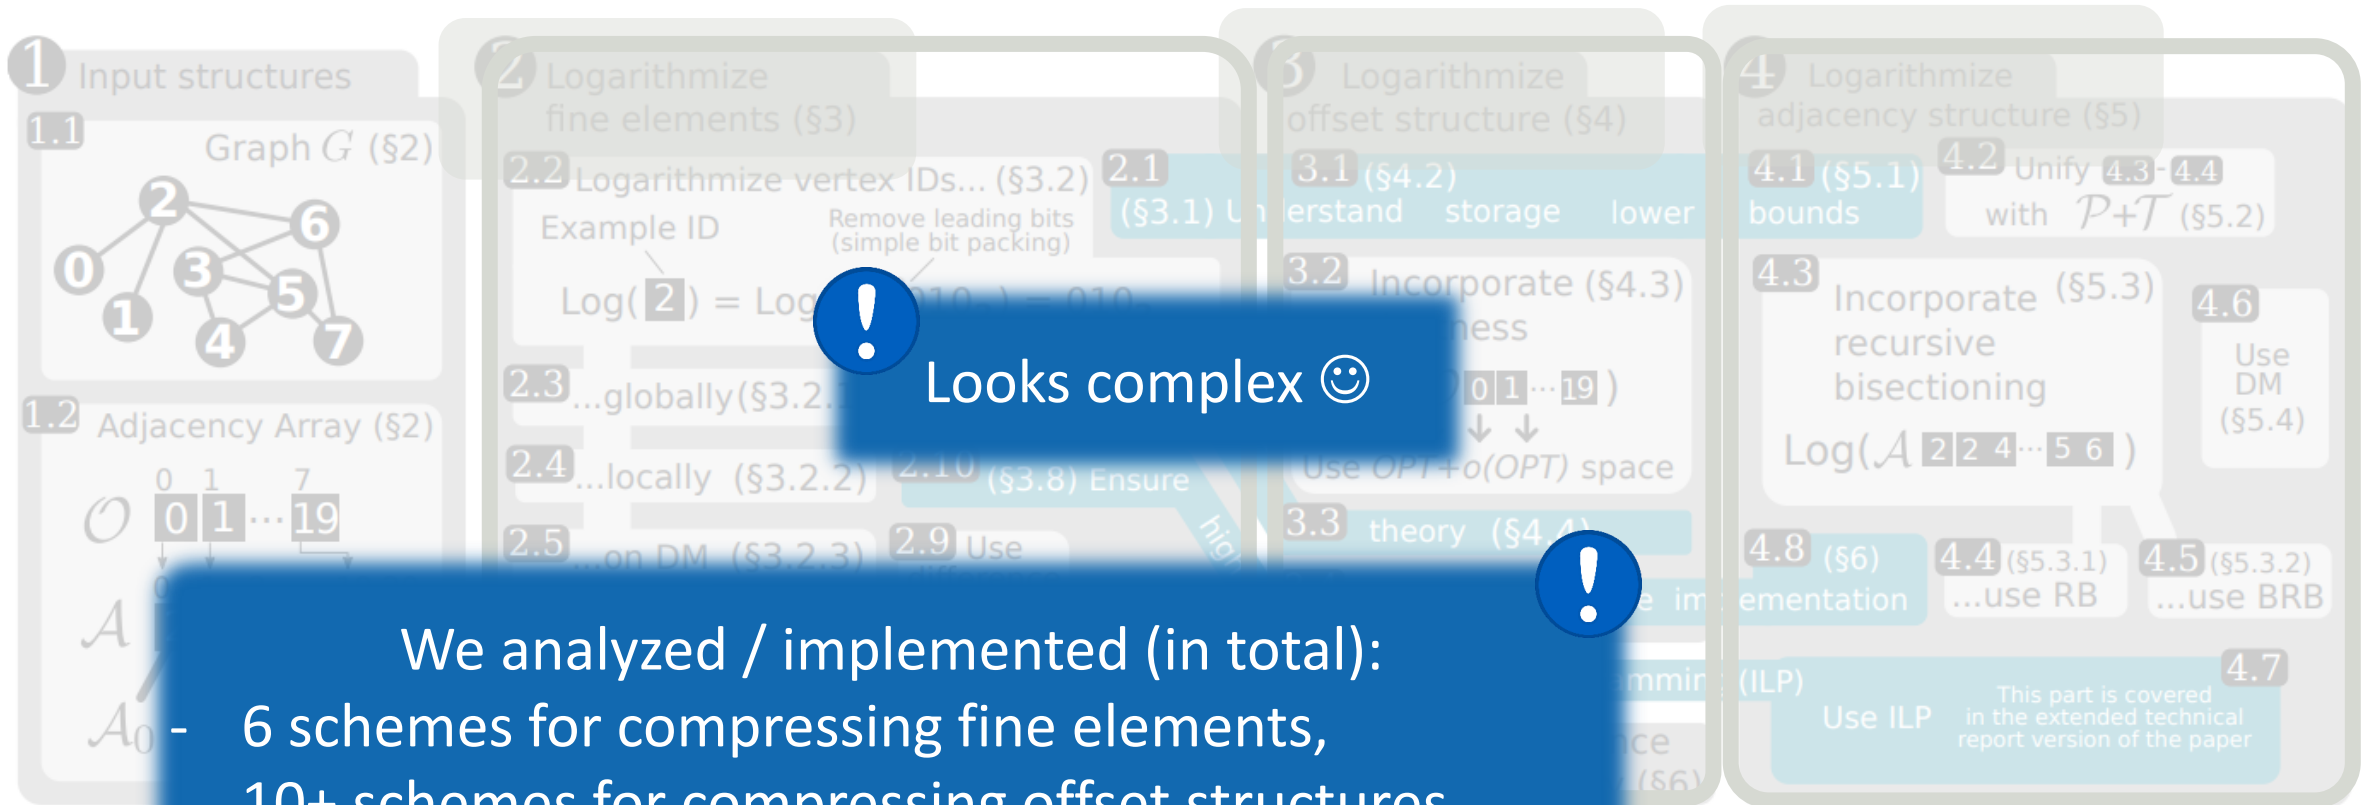
4.3 Incorporate recursive bisectioning (§5.3)

4.4 Use DM (§5.4)

4.5 Use ILP (§5.3.2)

4.6 Use ILP 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)

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 ...on DM (§3.2.3)

2.8

2.9 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)

4.2 Unify 4.3-4.4 with $P+T$ (§5.2)

4.3 Incorporate (§5.3)

4.4 recursive bisectioning

4.5 Use DM (§5.4)

4.6

4.7

4.8 (§6)

4.9 ...use RB

4.10 (§5.3.1)

4.11 (§5.3.2)

4.12 ...use BRB

4.13 Use ILP

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

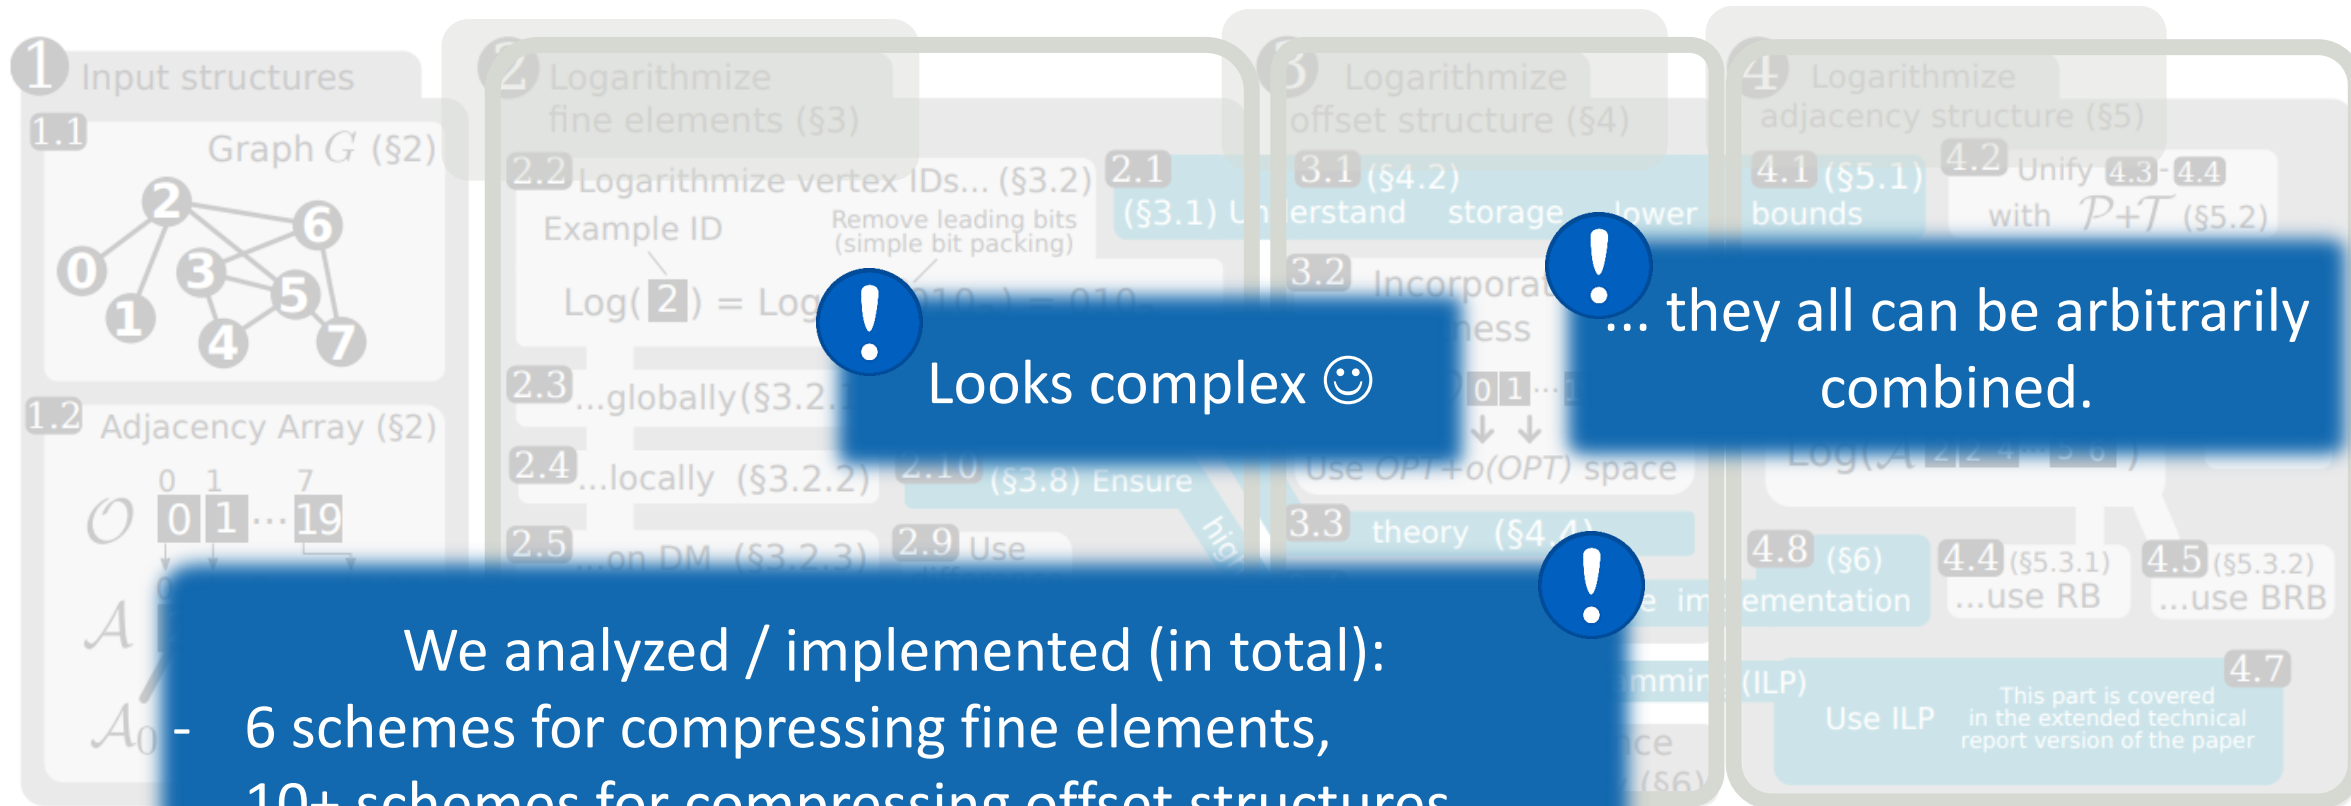
! Looks complex 😊

!

We analyzed / implemented (in total):

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

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 Remove leading bits (simple bit packing)

2.3 $\text{Log}(2) = \text{Log}(10) = 010$

2.4 ...globally (§3.2.1)

2.5 ...locally (§3.2.2)

2.9 Use

3 Logarithmize offset structure (§4)

3.1 (§4.2)

3.2 Incorporate

3.3 theory (§4.4)

4 Logarithmize adjacency structure (§5)

4.1 (§5.1)

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

4.8 (§6)

! Looks complex 😊

! ... they all can be arbitrarily combined.

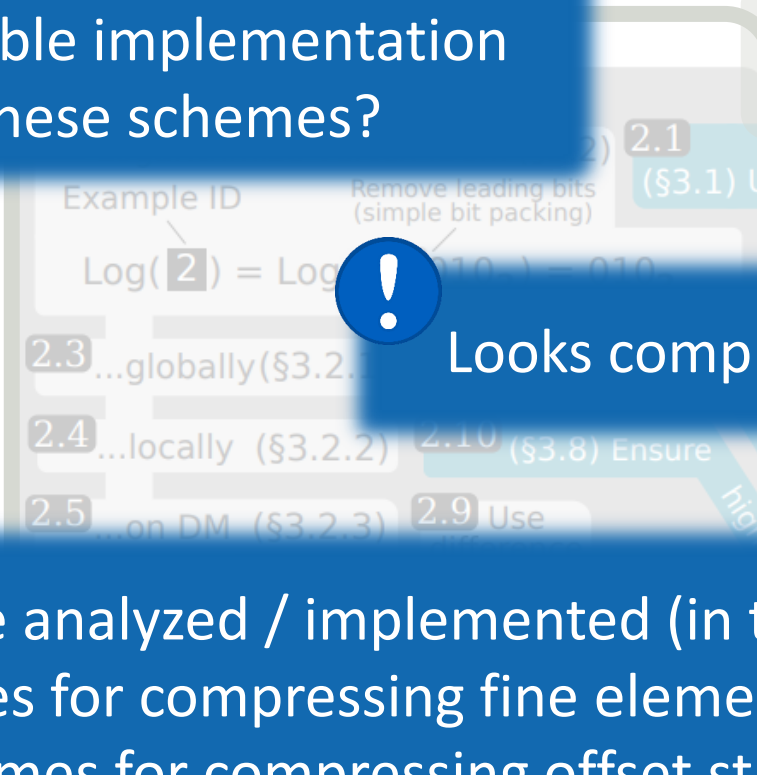
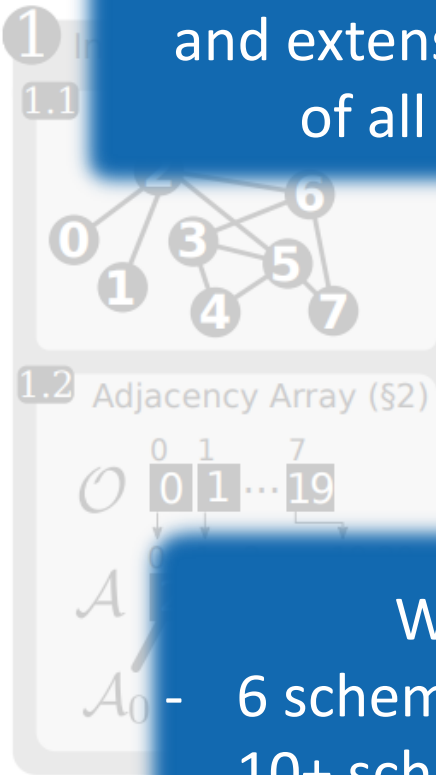
We analyzed / implemented (in total):

- 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 This part is covered 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 😊

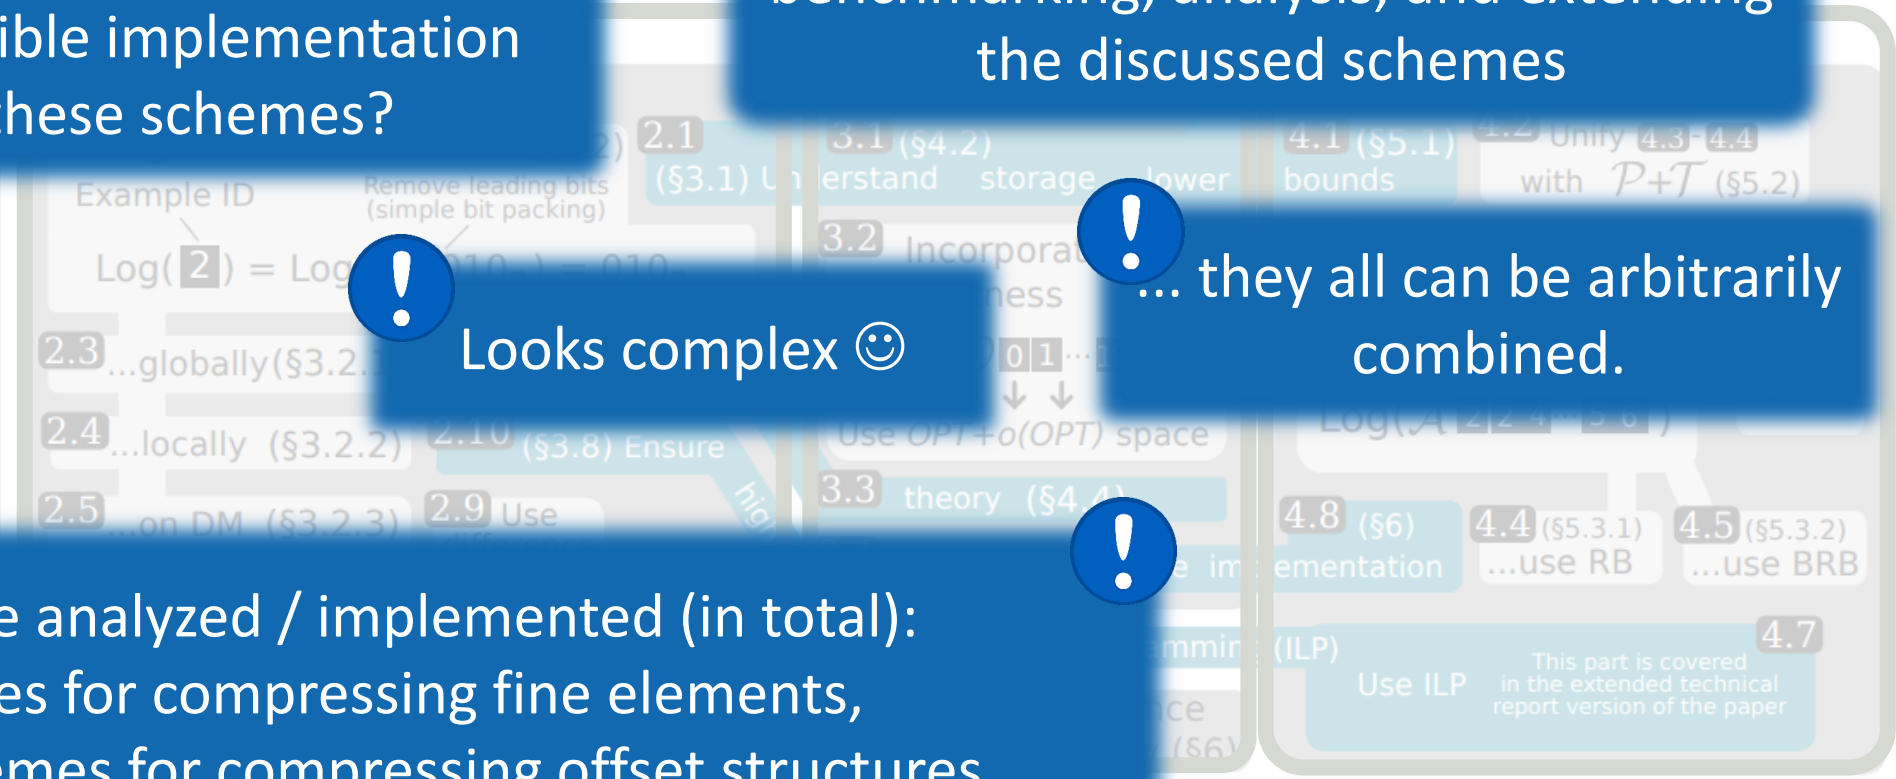


... they all can be arbitrarily combined.



We analyzed / implemented (in total):

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



PERFORMANCE ANALYSIS

TYPES OF MACHINES



CSCS Cray Piz Daint

PERFORMANCE ANALYSIS

TYPES OF MACHINES



CSCS Cray Piz Daint

PERFORMANCE ANALYSIS

TYPES OF MACHINES



CSCS Cray Piz Daint

HP "fat server" (DL360)

PERFORMANCE ANALYSIS

TYPES OF GRAPHS

PERFORMANCE ANALYSIS


TYPES OF GRAPHS

Synthetic graphs

PERFORMANCE ANALYSIS

TYPES OF GRAPHS

Synthetic graphs




Kronecker [1]


PERFORMANCE ANALYSIS

TYPES OF GRAPHS

Synthetic graphs



Kronecker [1]



Erdős-Rényi [2]


[1] J. Leskovec et al. Kronecker Graphs: An Approach to Modeling Networks. J. Mach. Learn. Research. 2010.

[2] P. Erdos and A. Renyi. On the evolution of random graphs. Pub. Math. Inst. Hun. A. Science. 1960.


PERFORMANCE ANALYSIS

TYPES OF GRAPHS

Synthetic graphs



Kronecker [1]



Erdős-Rényi [2]

Real-world graphs (SNAP [3], KONECT [4], Webgraph [5], DIMACS [6])

[3] SNAP. <https://snap.stanford.edu>


[1] J. Leskovec et al. Kronecker Graphs: An Approach to Modeling Networks. J. Mach. Learn. Research. 2010.

[2] P. Erdos and A. Renyi. On the evolution of random graphs. Pub. Math. Inst. Hun. A. Science. 1960.


PERFORMANCE ANALYSIS

TYPES OF GRAPHS

Synthetic graphs




Kronecker [1]



Erdős-Rényi [2]

Real-world graphs (SNAP [3], KONECT [4], Webgraph [5], DIMACS [6])




Web graphs



Road networks



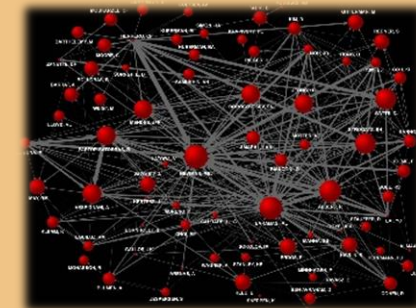
Social networks



Purchase networks



Communication graphs



Citation graphs

[1] J. Leskovec et al. Kronecker Graphs: An Approach to Modeling Networks. J. Mach. Learn. Research. 2010.

[2] P. Erdos and A. Renyi. On the evolution of random graphs. Pub. Math. Inst. Hun. A. Science. 1960.

[3] SNAP. <https://snap.stanford.edu>

[4] KONECT. <https://konect.cc>

[5] DIMACS Challenge

[6] Webgraphs. <https://law.di.unimi.it/datasets.php>

PERFORMANCE ANALYSIS

ALGORITHMS

PERFORMANCE ANALYSIS

ALGORITHMS


**Connected
Components**
(Shiloach-Vishkin [1])

PERFORMANCE ANALYSIS

ALGORITHMS

Connected Components

(Shiloach-Vishkin [1])




[1] Y. Shiloach, U. Vishkin. An $O(\log n)$ parallel connectivity algorithm. 1980.

PERFORMANCE ANALYSIS

ALGORITHMS

Connected Components

(Shiloach-Vishkin [1])




[1] Y. Shiloach, U. Vishkin. An $O(\log n)$ parallel connectivity algorithm. 1980.

PERFORMANCE ANALYSIS

ALGORITHMS

Connected Components

(Shiloach-Vishkin [1])




BFS (direction optimization [2])

[1] Y. Shiloach, U. Vishkin. An $O(\log n)$ parallel connectivity algorithm. 1980.

[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.


PERFORMANCE ANALYSIS

ALGORITHMS



Connected Components
 (Shiloach-Vishkin [1])

BFS (direction optimization [2])




[1] Y. Shiloach, U. Vishkin. An $O(\log n)$ parallel connectivity algorithm. 1980.

[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.


PERFORMANCE ANALYSIS

ALGORITHMS



Connected Components
(Shiloach-Vishkin [1])

BFS (direction optimization [2])




[1] Y. Shiloach, U. Vishkin. An $O(\log n)$ parallel connectivity algorithm. 1980.


[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.

PERFORMANCE ANALYSIS

ALGORITHMS



Connected Components
(Shiloach-Vishkin [1])




BFS (direction optimization [2])

[1] Y. Shiloach, U. Vishkin. An $O(\log n)$ parallel connectivity algorithm. 1980.


[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.

PERFORMANCE ANALYSIS

ALGORITHMS



Connected Components
(Shiloach-Vishkin [1])

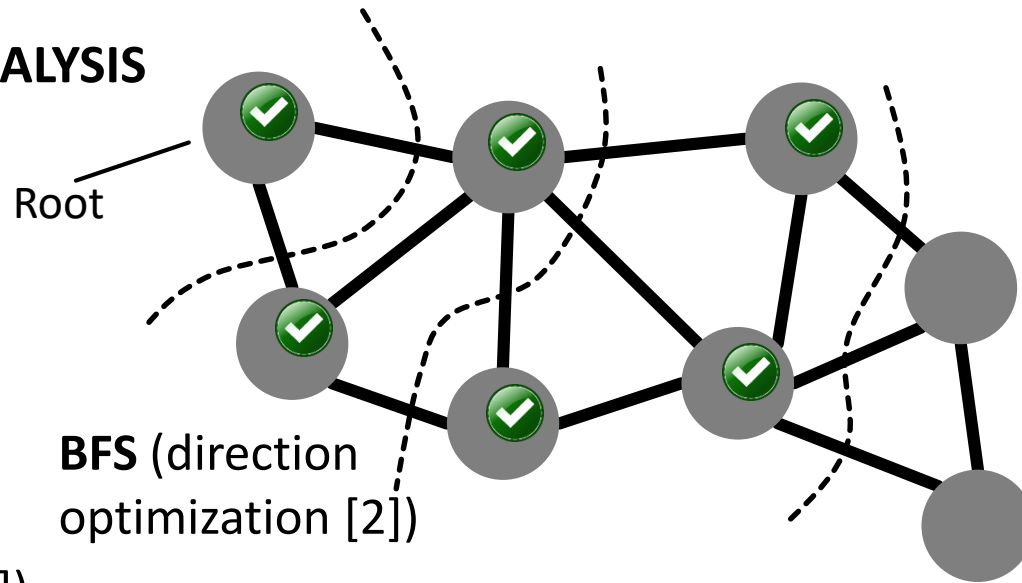


[1] Y. Shiloach, U. Vishkin. An $O(\log n)$ parallel connectivity algorithm. 1980.


[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.

PERFORMANCE ANALYSIS

ALGORITHMS



Connected Components
(Shiloach-Vishkin [1])




BFS (direction optimization [2])

[1] Y. Shiloach, U. Vishkin. An $O(\log n)$ parallel connectivity algorithm. 1980.


[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.

PERFORMANCE ANALYSIS

ALGORITHMS



Connected Components
(Shiloach-Vishkin [1])




BFS (direction optimization [2])

[1] Y. Shiloach, U. Vishkin. An $O(\log n)$ parallel connectivity algorithm. 1980.


[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.

PERFORMANCE ANALYSIS

ALGORITHMS



Connected Components
(Shiloach-Vishkin [1])



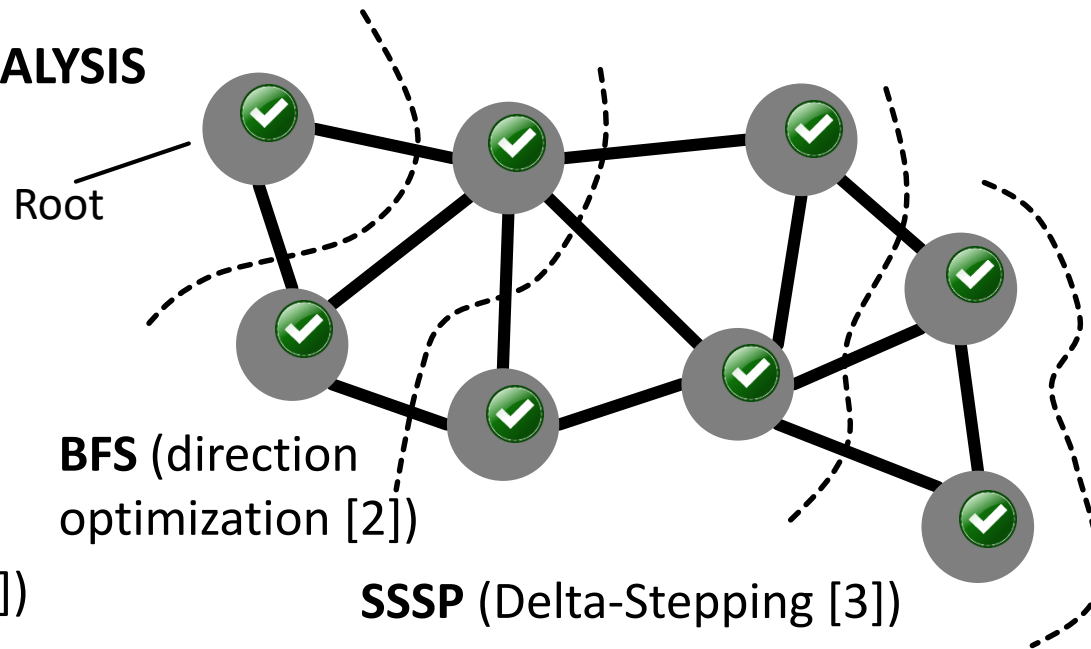
[1] Y. Shiloach, U. Vishkin. An $O(\log n)$ parallel connectivity algorithm. 1980.

[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.

[3] U. Meyer, P. Sanders. Delta-Stepping: A Parallelizable Shortest Path Algorithm. 2003.

PERFORMANCE ANALYSIS

ALGORITHMS




Connected Components
(Shiloach-Vishkin [1])

BFS (direction optimization [2])

SSSP (Delta-Stepping [3])

Triangle Counting




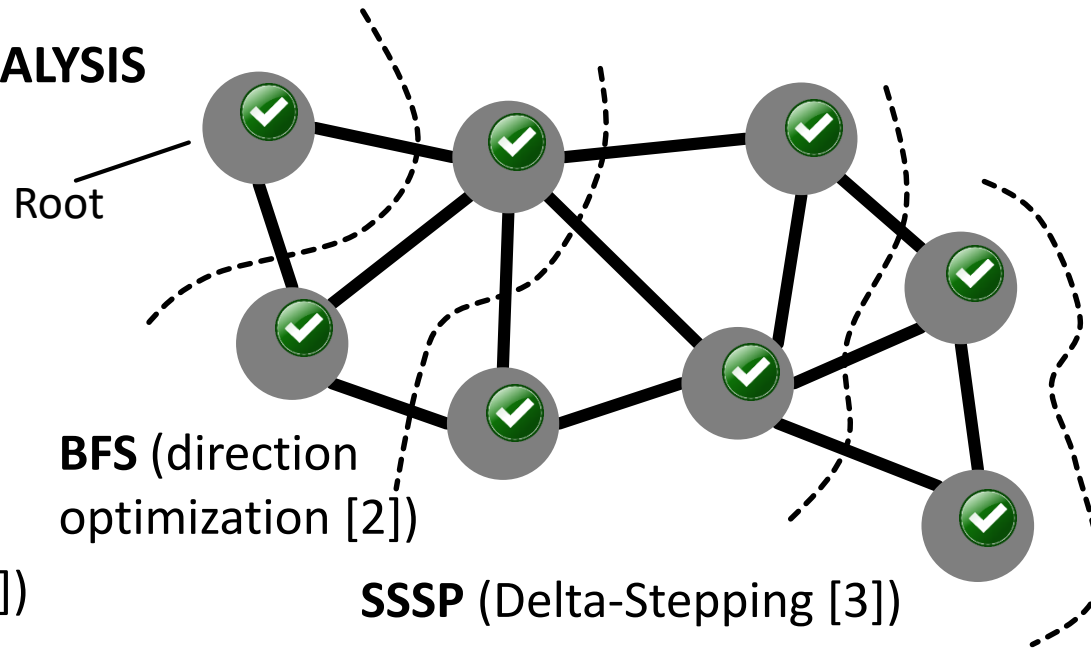
[1] Y. Shiloach, U. Vishkin. An $O(\log n)$ parallel connectivity algorithm. 1980.

[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.

[3] U. Meyer, P. Sanders. Delta-Stepping: A Parallelizable Shortest Path Algorithm. 2003.


PERFORMANCE ANALYSIS

ALGORITHMS



Connected Components

(Shiloach-Vishkin [1])



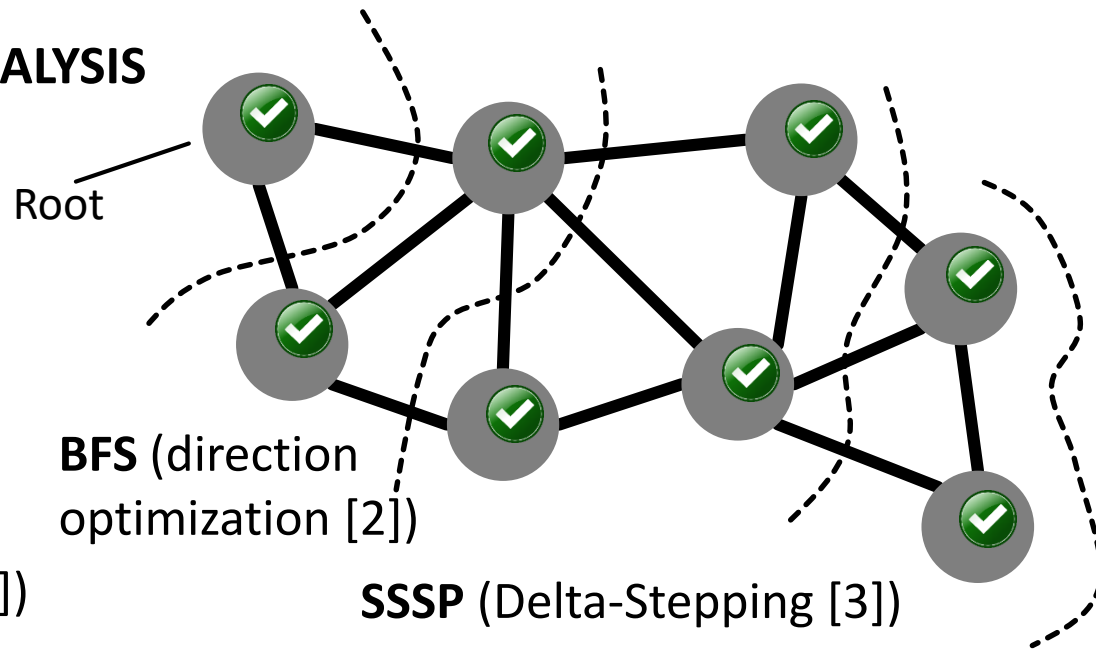
[1] Y. Shiloach, U. Vishkin. An $O(\log n)$ parallel connectivity algorithm. 1980.

[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.

[3] U. Meyer, P. Sanders. Delta-Stepping: A Parallelizable Shortest Path Algorithm. 2003.



PERFORMANCE ANALYSIS

ALGORITHMS



Connected Components

(Shiloach-Vishkin [1])



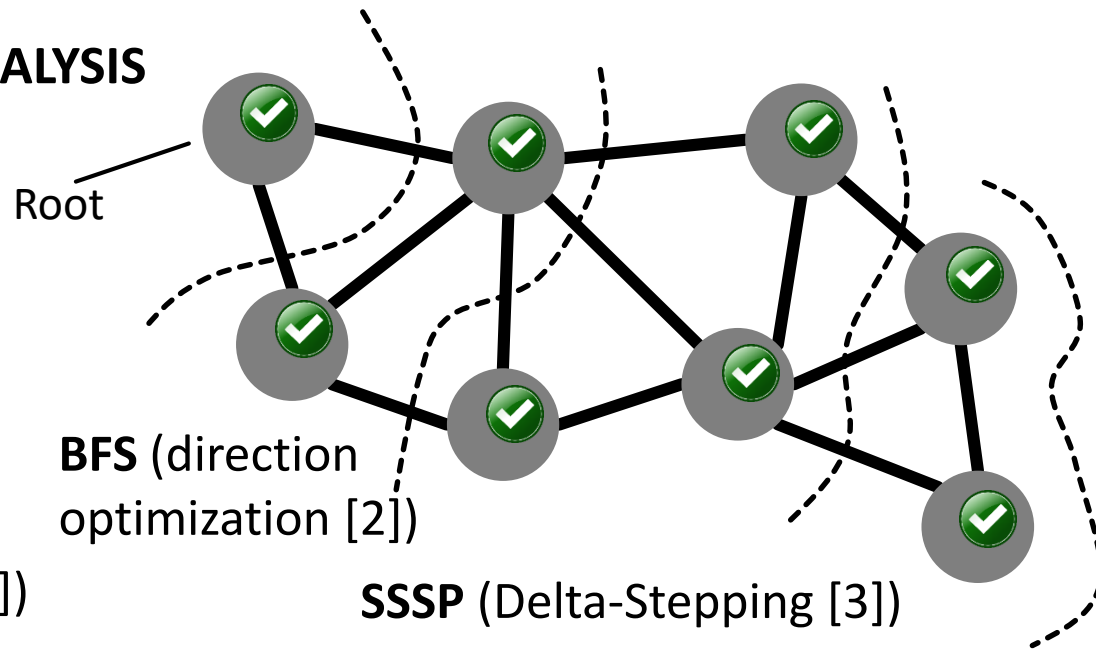
[1] Y. Shiloach, U. Vishkin. An $O(\log n)$ parallel connectivity algorithm. 1980.

[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.

[3] U. Meyer, P. Sanders. Delta-Stepping: A Parallelizable Shortest Path Algorithm. 2003.

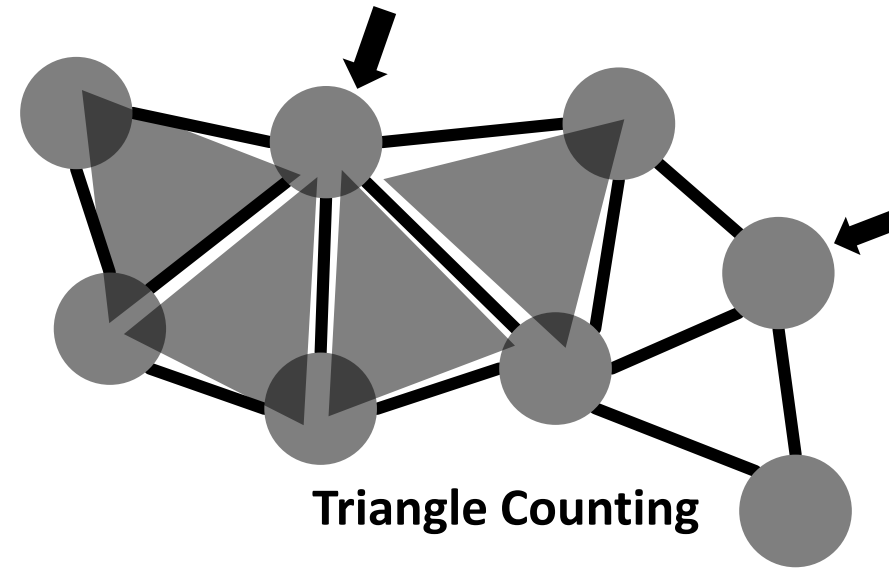

PERFORMANCE ANALYSIS

ALGORITHMS



Connected Components

(Shiloach-Vishkin [1])




[1] Y. Shiloach, U. Vishkin. An $O(\log n)$ parallel connectivity algorithm. 1980.

[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.

[3] U. Meyer, P. Sanders. Delta-Stepping: A Parallelizable Shortest Path Algorithm. 2003.



PERFORMANCE ANALYSIS

ALGORITHMS



Connected Components

(Shiloach-Vishkin [1])




[1] Y. Shiloach, U. Vishkin. An $O(\log n)$ parallel connectivity algorithm. 1980.

[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.

[3] U. Meyer, P. Sanders. Delta-Stepping: A Parallelizable Shortest Path Algorithm. 2003.



PERFORMANCE ANALYSIS

ALGORITHMS



Connected Components

(Shiloach-Vishkin [1])




[1] Y. Shiloach, U. Vishkin. An $O(\log n)$ parallel connectivity algorithm. 1980.

[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.

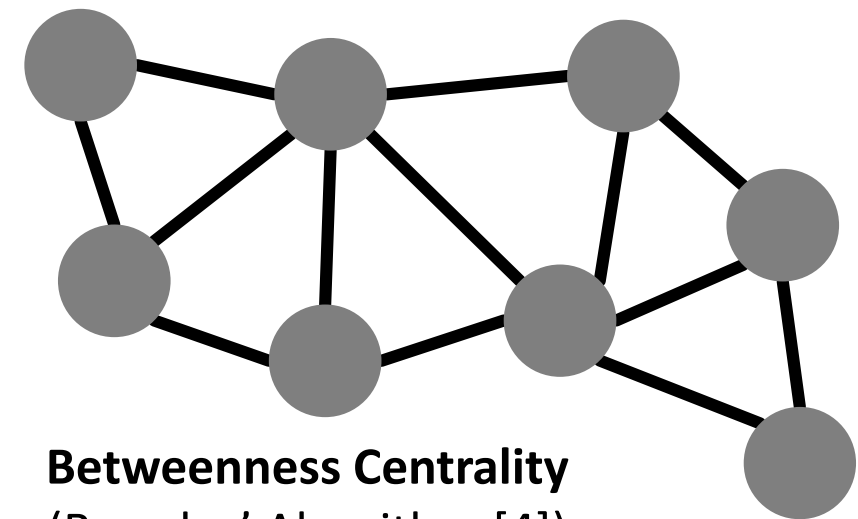
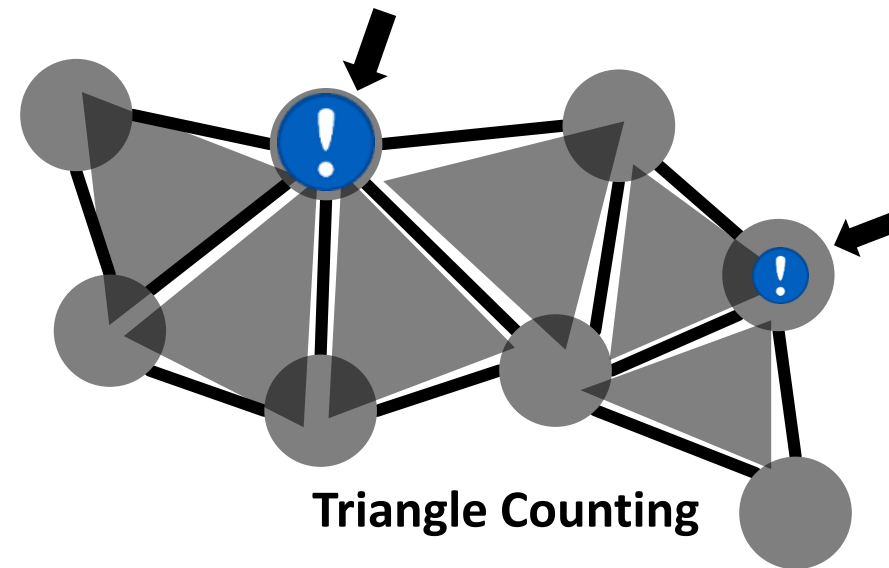

[3] U. Meyer, P. Sanders. Delta-Stepping: A Parallelizable Shortest Path Algorithm. 2003.

PERFORMANCE ANALYSIS

ALGORITHMS



Connected Components
(Shiloach-Vishkin [1])



[1] Y. Shiloach, U. Vishkin. An $O(\log n)$ parallel connectivity algorithm. 1980.


[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.

[3] U. Meyer, P. Sanders. Delta-Stepping: A Parallelizable Shortest Path Algorithm. 2003.




[4] U. Brandes. A Faster Algorithm for Betweenness Centrality. 2001.

PERFORMANCE ANALYSIS

ALGORITHMS



Connected Components
(Shiloach-Vishkin [1])



[1] Y. Shiloach, U. Vishkin. An $O(\log n)$ parallel connectivity algorithm. 1980.


[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.

[3] U. Meyer, P. Sanders. Delta-Stepping: A Parallelizable Shortest Path Algorithm. 2003.




[4] U. Brandes. A Faster Algorithm for Betweenness Centrality. 2001.

PERFORMANCE ANALYSIS

ALGORITHMS



Connected Components
(Shiloach-Vishkin [1])



[1] Y. Shiloach, U. Vishkin. An $O(\log n)$ parallel connectivity algorithm. 1980.


[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.

[3] U. Meyer, P. Sanders. Delta-Stepping: A Parallelizable Shortest Path Algorithm. 2003.

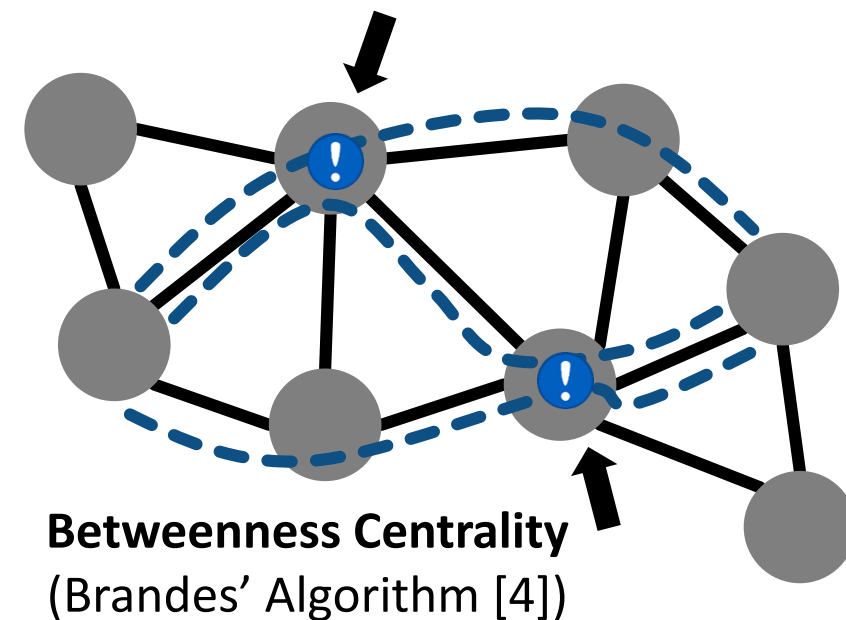


[4] U. Brandes. A Faster Algorithm for Betweenness Centrality. 2001.

PERFORMANCE ANALYSIS

ALGORITHMS



Connected Components
(Shiloach-Vishkin [1])



[1] Y. Shiloach, U. Vishkin. An $O(\log n)$ parallel connectivity algorithm. 1980.

[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.



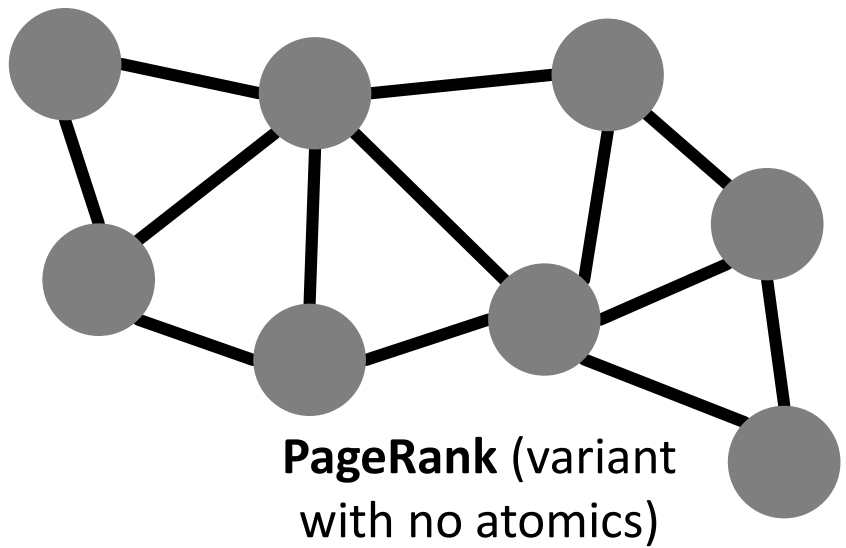
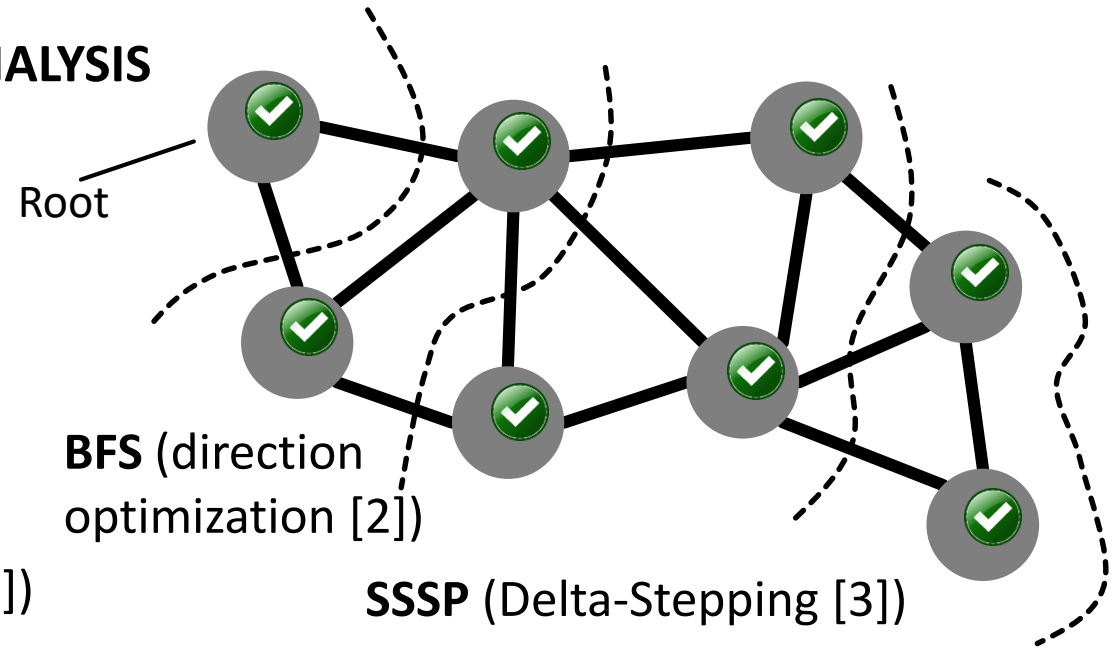

[3] U. Meyer, P. Sanders. Delta-Stepping: A Parallelizable Shortest Path Algorithm. 2003.

[4] U. Brandes. A Faster Algorithm for Betweenness Centrality. 2001.

PERFORMANCE ANALYSIS

ALGORITHMS

Connected Components
(Shiloach-Vishkin [1])



[1] Y. Shiloach, U. Vishkin. An $O(\log n)$ parallel connectivity algorithm. 1980.

[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.

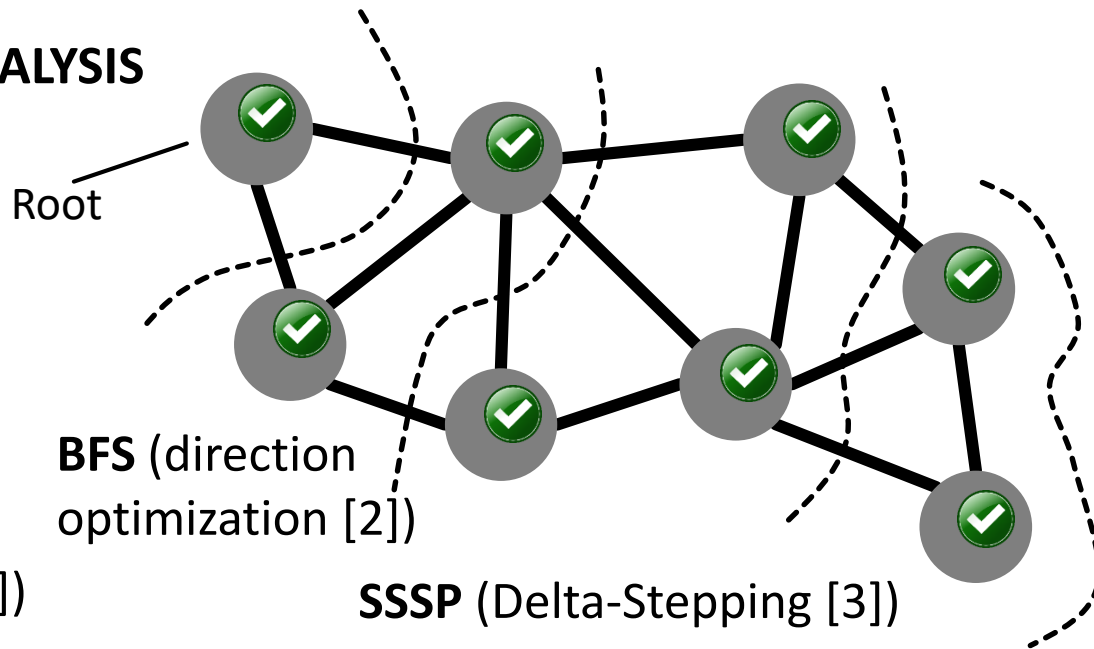

[3] U. Meyer, P. Sanders. Delta-Stepping: A Parallelizable Shortest Path Algorithm. 2003.

[4] U. Brandes. A Faster Algorithm for Betweenness Centrality. 2001.

PERFORMANCE ANALYSIS

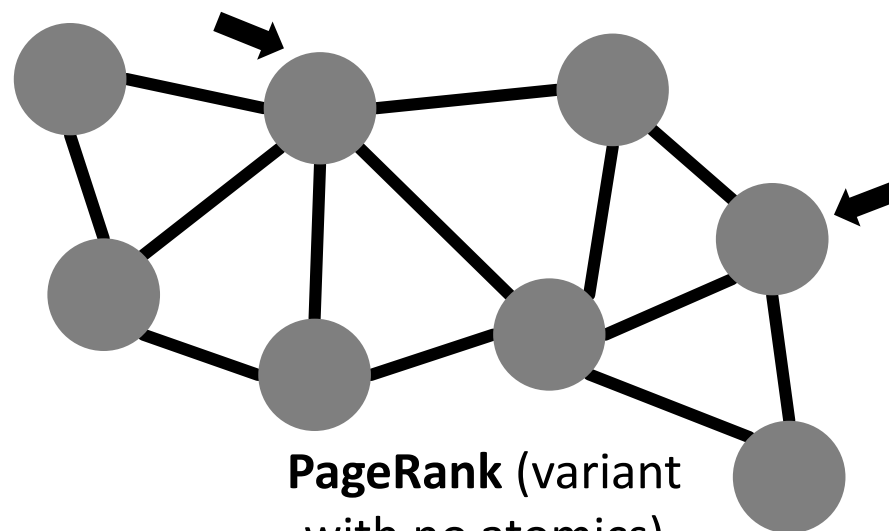
ALGORITHMS

Connected Components
(Shiloach-Vishkin [1])




BFS (direction optimization [2])

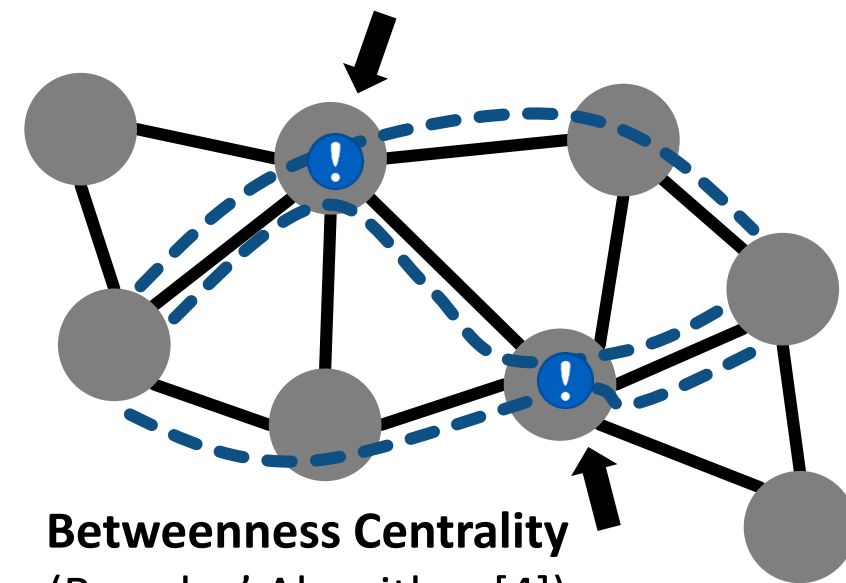
SSSP (Delta-Stepping [3])



PageRank (variant with no atomics)



Triangle Counting



Betweenness Centrality
(Brandes' Algorithm [4])

[1] Y. Shiloach, U. Vishkin. An $O(\log n)$ parallel connectivity algorithm. 1980.

[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.



[3] U. Meyer, P. Sanders. Delta-Stepping: A Parallelizable Shortest Path Algorithm. 2003.

[4] U. Brandes. A Faster Algorithm for Betweenness Centrality. 2001.

PERFORMANCE ANALYSIS

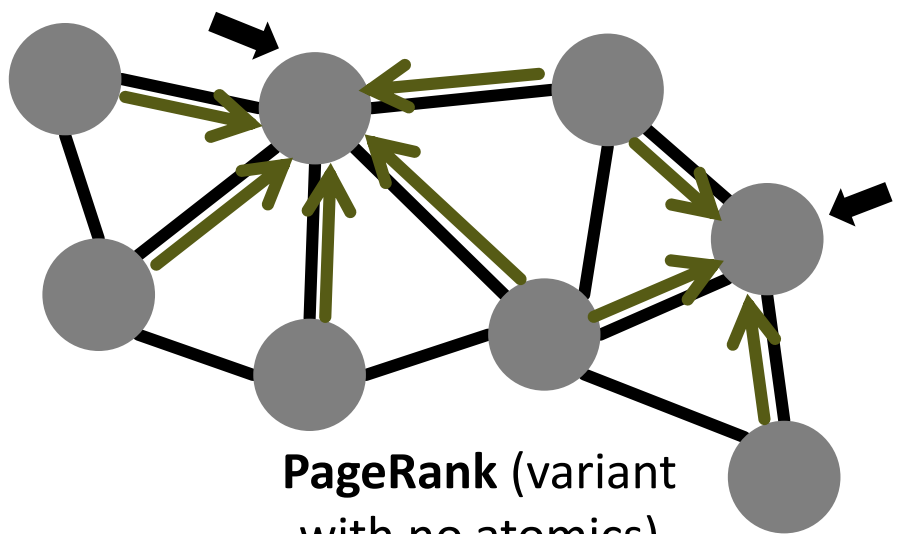
ALGORITHMS

Connected Components
(Shiloach-Vishkin [1])




BFS (direction optimization [2])


SSSP (Delta-Stepping [3])



PageRank (variant with no atomics)



Triangle Counting



Betweenness Centrality
(Brandes' Algorithm [4])

[1] Y. Shiloach, U. Vishkin. An $O(\log n)$ parallel connectivity algorithm. 1980.

[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.


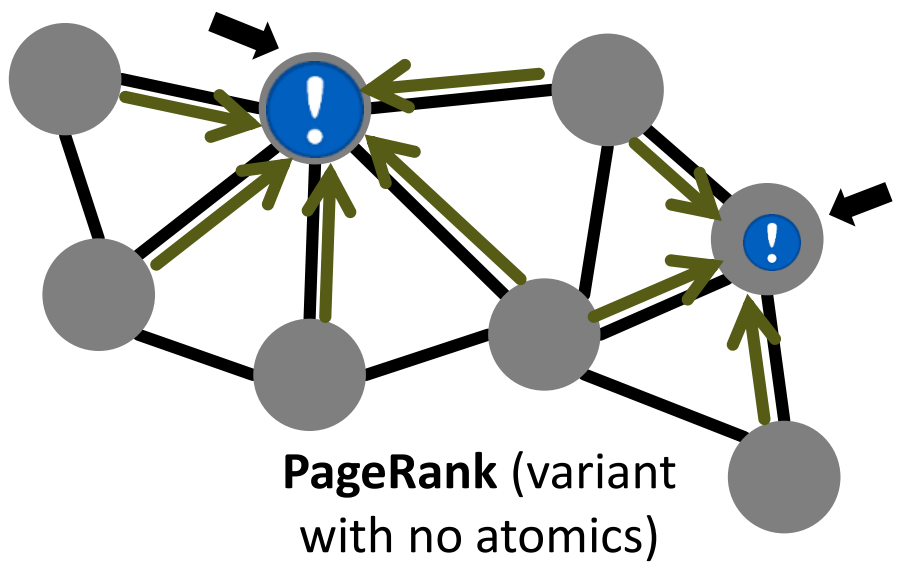



[3] U. Meyer, P. Sanders. Delta-Stepping: A Parallelizable Shortest Path Algorithm. 2003.

[4] U. Brandes. A Faster Algorithm for Betweenness Centrality. 2001.

PERFORMANCE ANALYSIS

ALGORITHMS

Connected Components
(Shiloach-Vishkin [1])



[1] Y. Shiloach, U. Vishkin. An $O(\log n)$ parallel connectivity algorithm. 1980.

[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.

[3] U. Meyer, P. Sanders. Delta-Stepping: A Parallelizable Shortest Path Algorithm. 2003.

[4] U. Brandes. A Faster Algorithm for Betweenness Centrality. 2001.

PERFORMANCE ANALYSIS

COMPARISON TARGETS

PERFORMANCE ANALYSIS

COMPARISON TARGETS



GAPBS: Graph Algorithm Platform Benchmark Suite [1].
Comparison to a traditional adjacency array implementation

[1] S. Beamer, K. Asanovic, and D. Patterson. The GAP benchmark suite. arXiv preprint arXiv:1508.03619, 2015.

PERFORMANCE ANALYSIS

COMPARISON TARGETS



Zlib [2].

Comparison to a traditional
compression scheme



GAPBS: Graph Algorithm Platform Benchmark Suite [1].
Comparison to a traditional adjacency array implementation

[1] S. Beamer, K. Asanovic, and D. Patterson. The GAP benchmark suite. arXiv preprint arXiv:1508.03619, 2015.

[2] P. Deutsch and J.-L. Gailly. ZLIB Compressed Data Format Specification, 1996.

PERFORMANCE ANALYSIS

COMPARISON TARGETS



Zlib [2].

Comparison to a traditional compression scheme



GAPBS: Graph Algorithm Platform Benchmark Suite [1].
Comparison to a traditional adjacency array implementation



WebGraph Library [3]

Comparison to a state-of-the-art graph compression scheme

[1] S. Beamer, K. Asanovic, and D. Patterson. The GAP benchmark suite. arXiv preprint arXiv:1508.03619, 2015.

[2] P. Deutsch and J.-L. Gailly. ZLIB Compressed Data Format Specification, 1996.

[3] P. Boldi and S. Vigna. The WebGraph Framework I: compression techniques. WWW, 2004.

PERFORMANCE ANALYSIS

COMPARISON TARGETS



Zlib [2].

Comparison to a traditional compression scheme



GAPBS: Graph Algorithm Platform Benchmark Suite [1].
Comparison to a traditional adjacency array implementation



WebGraph Library [3]

Comparison to a state-of-the-art graph compression scheme



Recursive Partitioning [4].

Comparison to a tuned scheme for compressing adjacency data


[1] S. Beamer, K. Asanovic, and D. Patterson. The GAP benchmark suite. arXiv preprint arXiv:1508.03619, 2015.

[2] P. Deutsch and J.-L. Gailly. ZLIB Compressed Data Format Specification, 1996.

[3] P. Boldi and S. Vigna. The WebGraph Framework I: compression techniques. WWW, 2004.

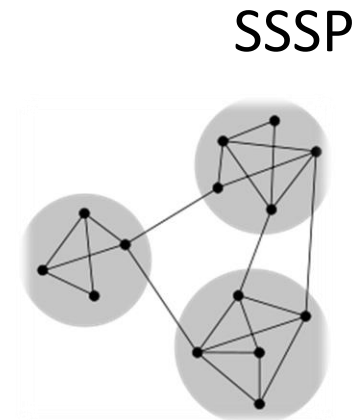

[4] D. K. Blandford, G. E. Blelloch, and I. A. Kash. Compact Representations of Separable Graphs. SODA, 2003.

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**




Log(Graph) accelerates GAPBS

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



SSSP
Kronecker graphs
Number of vertices: 4M

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




Kronecker graphs
Number of vertices: 4M

**Log(Graph)
accelerates GAPBS**


**Both storage and performance
are improved simultaneously**

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




2 Log (Offset structure) Storage

2 Log (Offset structure) Storage



Various real-world graphs


2 Log (Offset structure) Storage



Lots of data 😊

Conclusions:


2 Log (Offset structure) Storage



Lots of data 😊

Conclusions:


2 Log (Offset structure) Storage




ptr64, ptr32: traditional array of offsets
ptrLogn: separate compression of each offset
bvPL: plain bit vectors
bvIL: compact bit vectors
bvEN, bvSD: succinct bit vectors

Lots of data 😊
 Conclusions:

2 Log (Offset structure) Storage




ptr64, ptr32: traditional array of offsets
 ptrLogn: separate compression of each offset
 bvPL: plain bit vectors
 bvIL: compact bit vectors
 bvEN, bvSD: succinct bit vectors




Lots of data 😊
 Conclusions:

2 Log (Offset structure) Storage




ptr64, ptr32: traditional array of offsets
ptrLogn: separate compression of each offset
bvPL: plain bit vectors
bvIL: compact bit vectors
bvEN, bvSD: succinct bit vectors




Lots of data 😊
 Conclusions:

Succinct bit vectors consistently ensure best storage reductions

2 Log (Offset structure) Storage



ptr64, ptr32: traditional array of offsets
 ptrLogn: separate compression of each offset
 bvPL: plain bit vectors
 bvIL: compact bit vectors
 bvEN, bvSD: succinct bit vectors




Lots of data 😊
 Conclusions:

Succinct bit vectors consistently ensure best storage reductions

The main reason: succinct designs work well for sparse bit vectors, and graphs „that matter“ are sparse

2 Log (Offset structure) Performance



Accessing randomly selected neighbors



Kronecker graphs
Number of vertices: 4M

2 Log (Offset structure) Performance

Accessing randomly selected neighbors





Kronecker graphs
 Number of vertices: 4M

2 Log (Offset structure) Performance

Accessing randomly selected neighbors

- ptr64**: traditional array of offsets
- bvPL**: plain bit vectors
- bvIL**: compact bit vectors
- bvEN, bvSD**: succinct bit vectors
- zlib(.)**: zlib-compressed variants




Kronecker graphs
 Number of vertices: 4M


2 Log (Offset structure) Performance

Accessing randomly selected neighbors

Lots of data again 😊 Conclusions:



- ptr64**: traditional array of offsets
- bvPL**: plain bit vectors
- bvIL**: compact bit vectors
- bvEN, bvSD**: succinct bit vectors
- zlib(.)**: zlib-compressed variants



Kronecker graphs
 Number of vertices: 4M


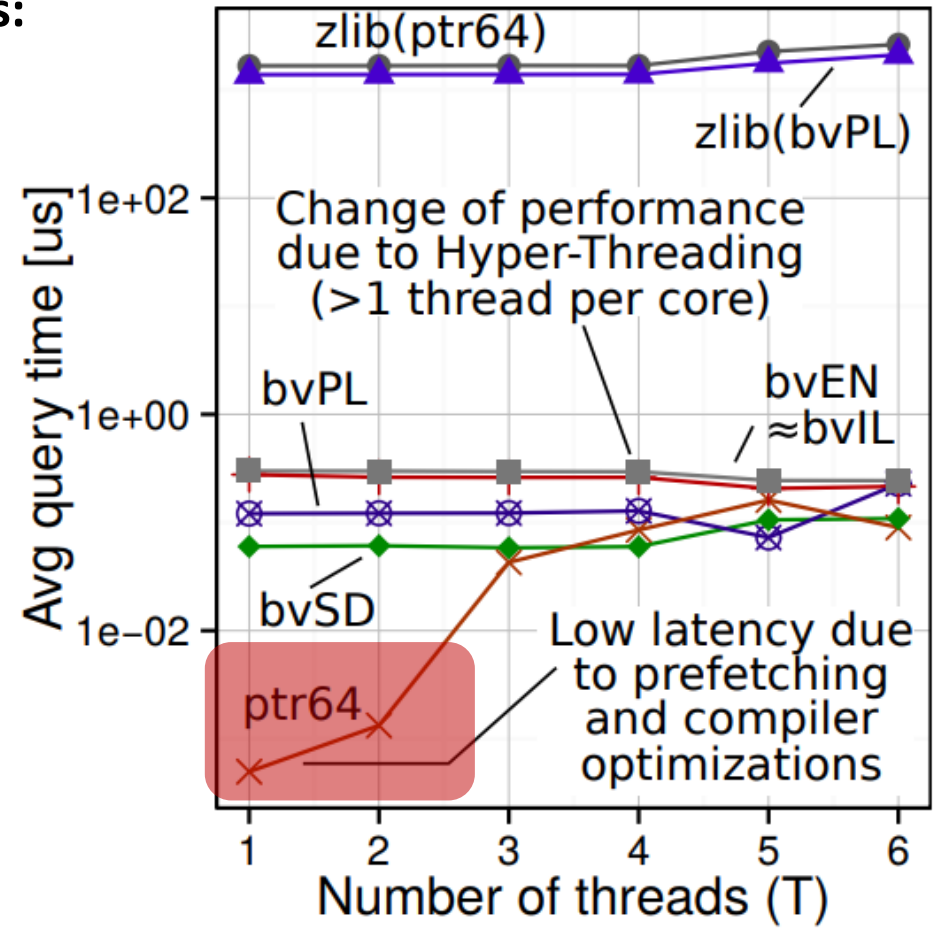
2 Log (Offset structure) Performance

Accessing randomly selected neighbors

Lots of data again 😊 Conclusions:

In sequential settings (or settings with low parallelism), simple offset arrays are the fastest

- ptr64**: traditional array of offsets
- bvPL**: plain bit vectors
- bvIL**: compact bit vectors
- bvEN, bvSD**: succinct bit vectors
- zlib(.)**: zlib-compressed variants



Kronecker graphs
Number of vertices: 4M


2 Log (Offset structure) Performance

Accessing randomly selected neighbors


Lots of data again 😊 Conclusions:

In sequential settings (or settings with low parallelism), **simple offset arrays are the fastest**

Once parallelism overheads kick in, performance of accessing succinct bit vectors and offset arrays **becomes comparable**



- ptr64**: traditional array of offsets
- bvPL**: plain bit vectors
- bvIL**: compact bit vectors
- bvEN, bvSD**: succinct bit vectors
- zlib(.)**: zlib-compressed variants



Kronecker graphs
Number of vertices: 4M


2 Log (Offset structure) Performance

Accessing randomly selected neighbors

Lots of data again 😊 Conclusions:


In sequential settings (or settings with low parallelism), simple offset arrays are the fastest

Once parallelism overheads kick in, performance of accessing succinct bit vectors and offset arrays becomes comparable



- ptr64**: traditional array of offsets
- bvPL**: plain bit vectors
- bvIL**: compact bit vectors
- bvEN, bvSD**: succinct bit vectors
- zlib(.)**: zlib-compressed variants

bvSD: the fastest and (usually) the smallest




Kronecker graphs
Number of vertices: 4M

3 **Log** (Adjacency structure) **Storage, Performane**

3 **Log** (Adjacency structure) **Storage, Performane**

Trad: Traditional adjacency array
DMd / DMf: Degree Minimizing (without / with gap encoding)
WG: WebGraph compression
BRB, RB: Schemes targeting certain specific classes of graphs




3

Log (Adjacency structure)


Storage, Performane

Trad: Traditional adjacency array
DMd / DMf: Degree Minimizing (without / with gap encoding)
WG: WebGraph compression
BRB, RB: Schemes targeting certain specific classes of graphs



3 **Log** (Adjacency structure) **Storage, Performane**

Trad: Traditional adjacency array
DMd / DMf: Degree Minimizing (without / with gap encoding)
WG: WebGraph compression
BRB, RB: Schemes targeting certain specific classes of graphs




Lots of data 😊

Conclusions:

3 **Log** (Adjacency structure) **Storage, Performane**

Trad: Traditional adjacency array
DMd / DMf: Degree Minimizing (without / with gap encoding)
WG: WebGraph compression
BRB, RB: Schemes targeting certain specific classes of graphs




Lots of data 😊

Conclusions:

WebGraph
best for web
graphs 😊

3 **Log** (Adjacency structure) **Storage, Performane**

Trad: Traditional adjacency array
DMd / DMf: Degree Minimizing (without / with gap encoding)
WG: WebGraph compression
BRB, RB: Schemes targeting certain specific classes of graphs



Lots of data 😊


Conclusions:

WebGraph best for web graphs 😊

BRB, RB: various tradeoffs but very expensive preprocessing (details in the paper)

3 **Log** (Adjacency structure) **Storage, Performane**

Trad: Traditional adjacency array
DMd / DMf: Degree Minimizing (without / with gap encoding)
WG: WebGraph compression
BRB, RB: Schemes targeting certain specific classes of graphs



Lots of data 😊

Conclusions:


WebGraph best for web graphs 😊

DMd: much better than DMf, often comparable to WG

BRB, RB: various tradeoffs but very expensive preprocessing (details in the paper)

3 **Log** (Adjacency structure) **Storage, Performane**

Trad: Traditional adjacency array
DMd / DMf: Degree Minimizing (without / with gap encoding)
WG: WebGraph compression
BRB, RB: Schemes targeting certain specific classes of graphs




Lots of data 😊

Conclusions:

WebGraph best for web graphs 😊


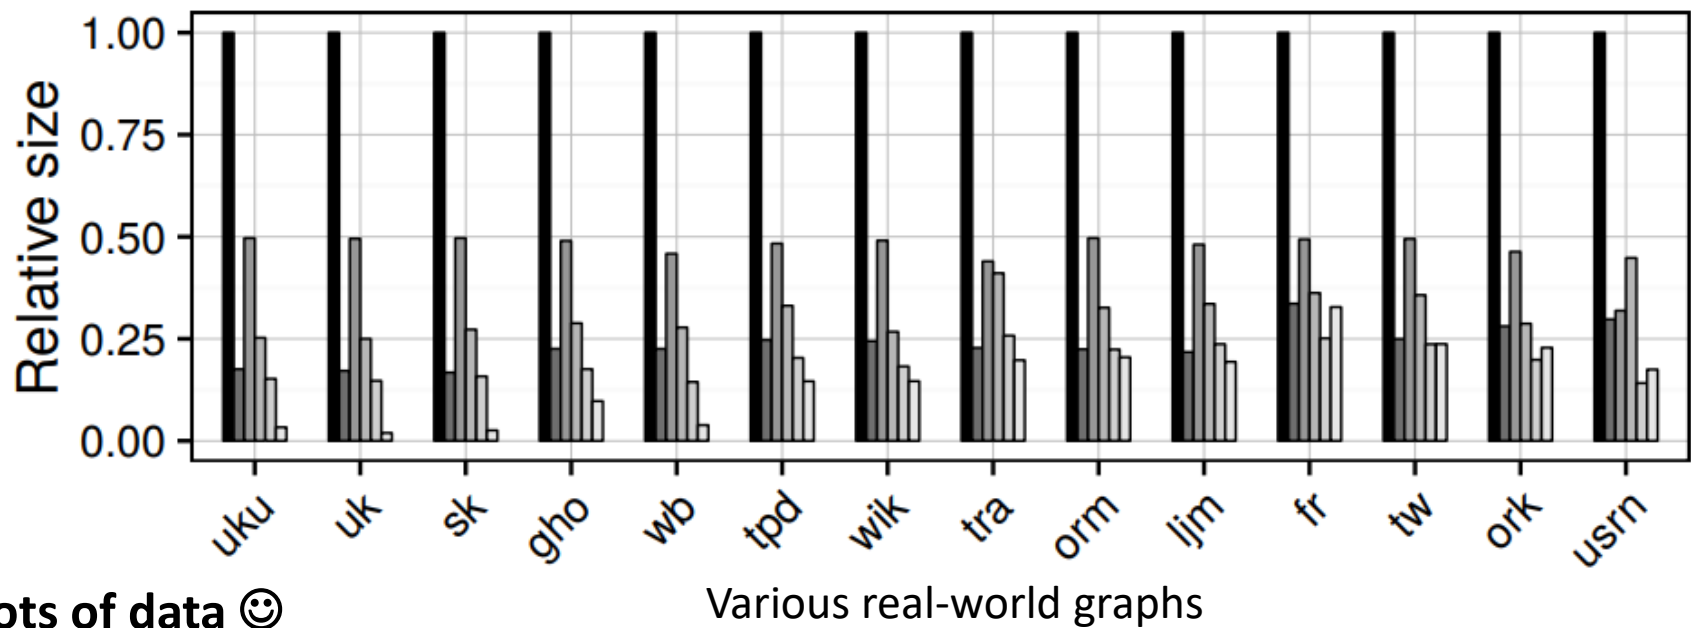
DMd: much better than DMf, often comparable to WG

BRB, RB: various tradeoffs but very expensive preprocessing (details in the paper)



3 **Log** (Adjacency structure) **Storage, Performane**

Trad: Traditional adjacency array
DMd / DMf: Degree Minimizing (without / with gap encoding)
WG: WebGraph compression
BRB, RB: Schemes targeting certain specific classes of graphs



WebGraph is the slowest, DM somewhat slower than Trad


Lots of data 😊

Conclusions:

WebGraph best for web graphs 😊

DMd: much better than DMf, often comparable to WG

BRB, RB: various tradeoffs but very expensive preprocessing (details in the paper)



**Takeaway (Results): Log(Graph) ensures
Space-Performance sweetspot (tunable!)**



Key insight (vertex labels)

20-35% storage reductions
(compared to uncompressed
data) and **negligible**
decompression overheads

**Takeaway (Results): Log(Graph) ensures
Space-Performance sweetspot (tunable!)**

! Key insight (vertex labels)

20-35% storage reductions (compared to uncompressed data) and **negligible** decompression overheads

Takeaway (Results): Log(Graph) ensures Space-Performance sweetspot (tunable!)

! Key insight (offsets)

Up to >90% storage reductions (compared to uncompressed data) and comparable performance to that of uncompressed data accesses (in parallel environments)

! Key insight (vertex labels)

20-35% storage reductions (compared to uncompressed data) and **negligible** decompression overheads

! Key insight (adjacency data)

80% storage reductions (compared to uncompressed data) and up to **>2x** speedup over modern graph compression schemes (Webgraph)

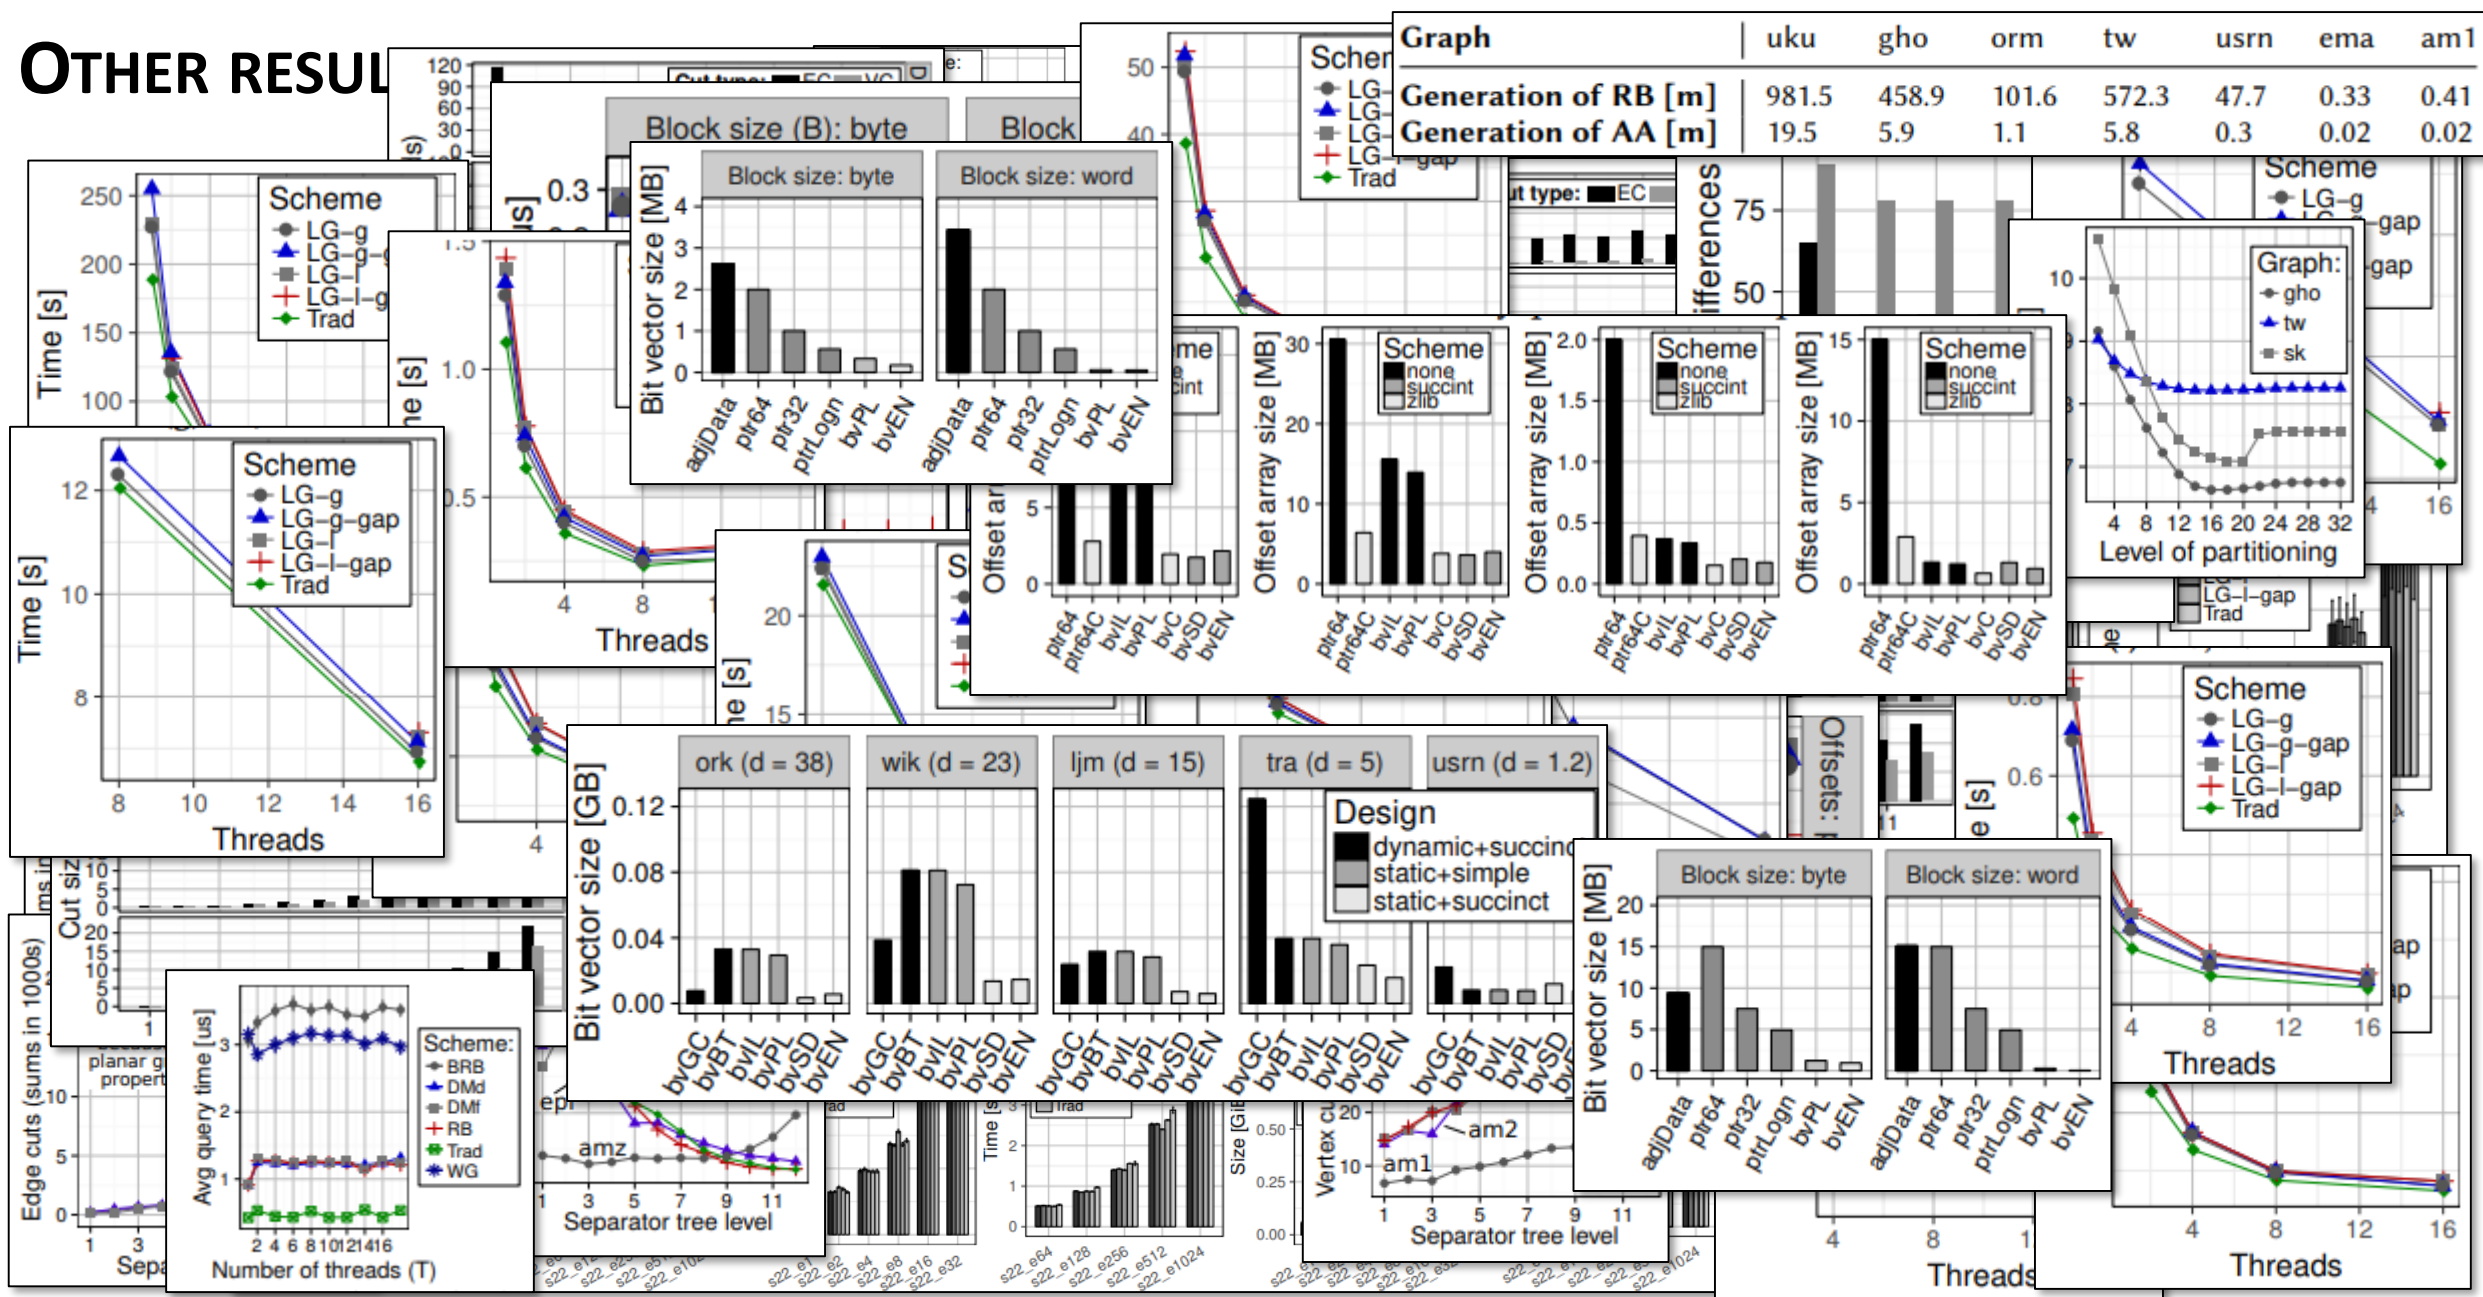
Takeaway (Results): Log(Graph) ensures Space-Performance sweetspot (tunable!)

! Key insight (offsets)

Up to >90% storage reductions (compared to uncompressed data) and comparable performance to that of uncompressed data accesses (in parallel environments)

OTHER RESULTS

OTHER RESULT





WHAT IS LOG(GRAPH)?

A NEAR-OPTIMAL GRAPH REPRESENTATION

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

The storage lower bound Ω

Which one? ?

Counting bounds. They are logarithmic (one needs at least $\log |S|$ bits to store an object from an arbitrary set S)

Key idea

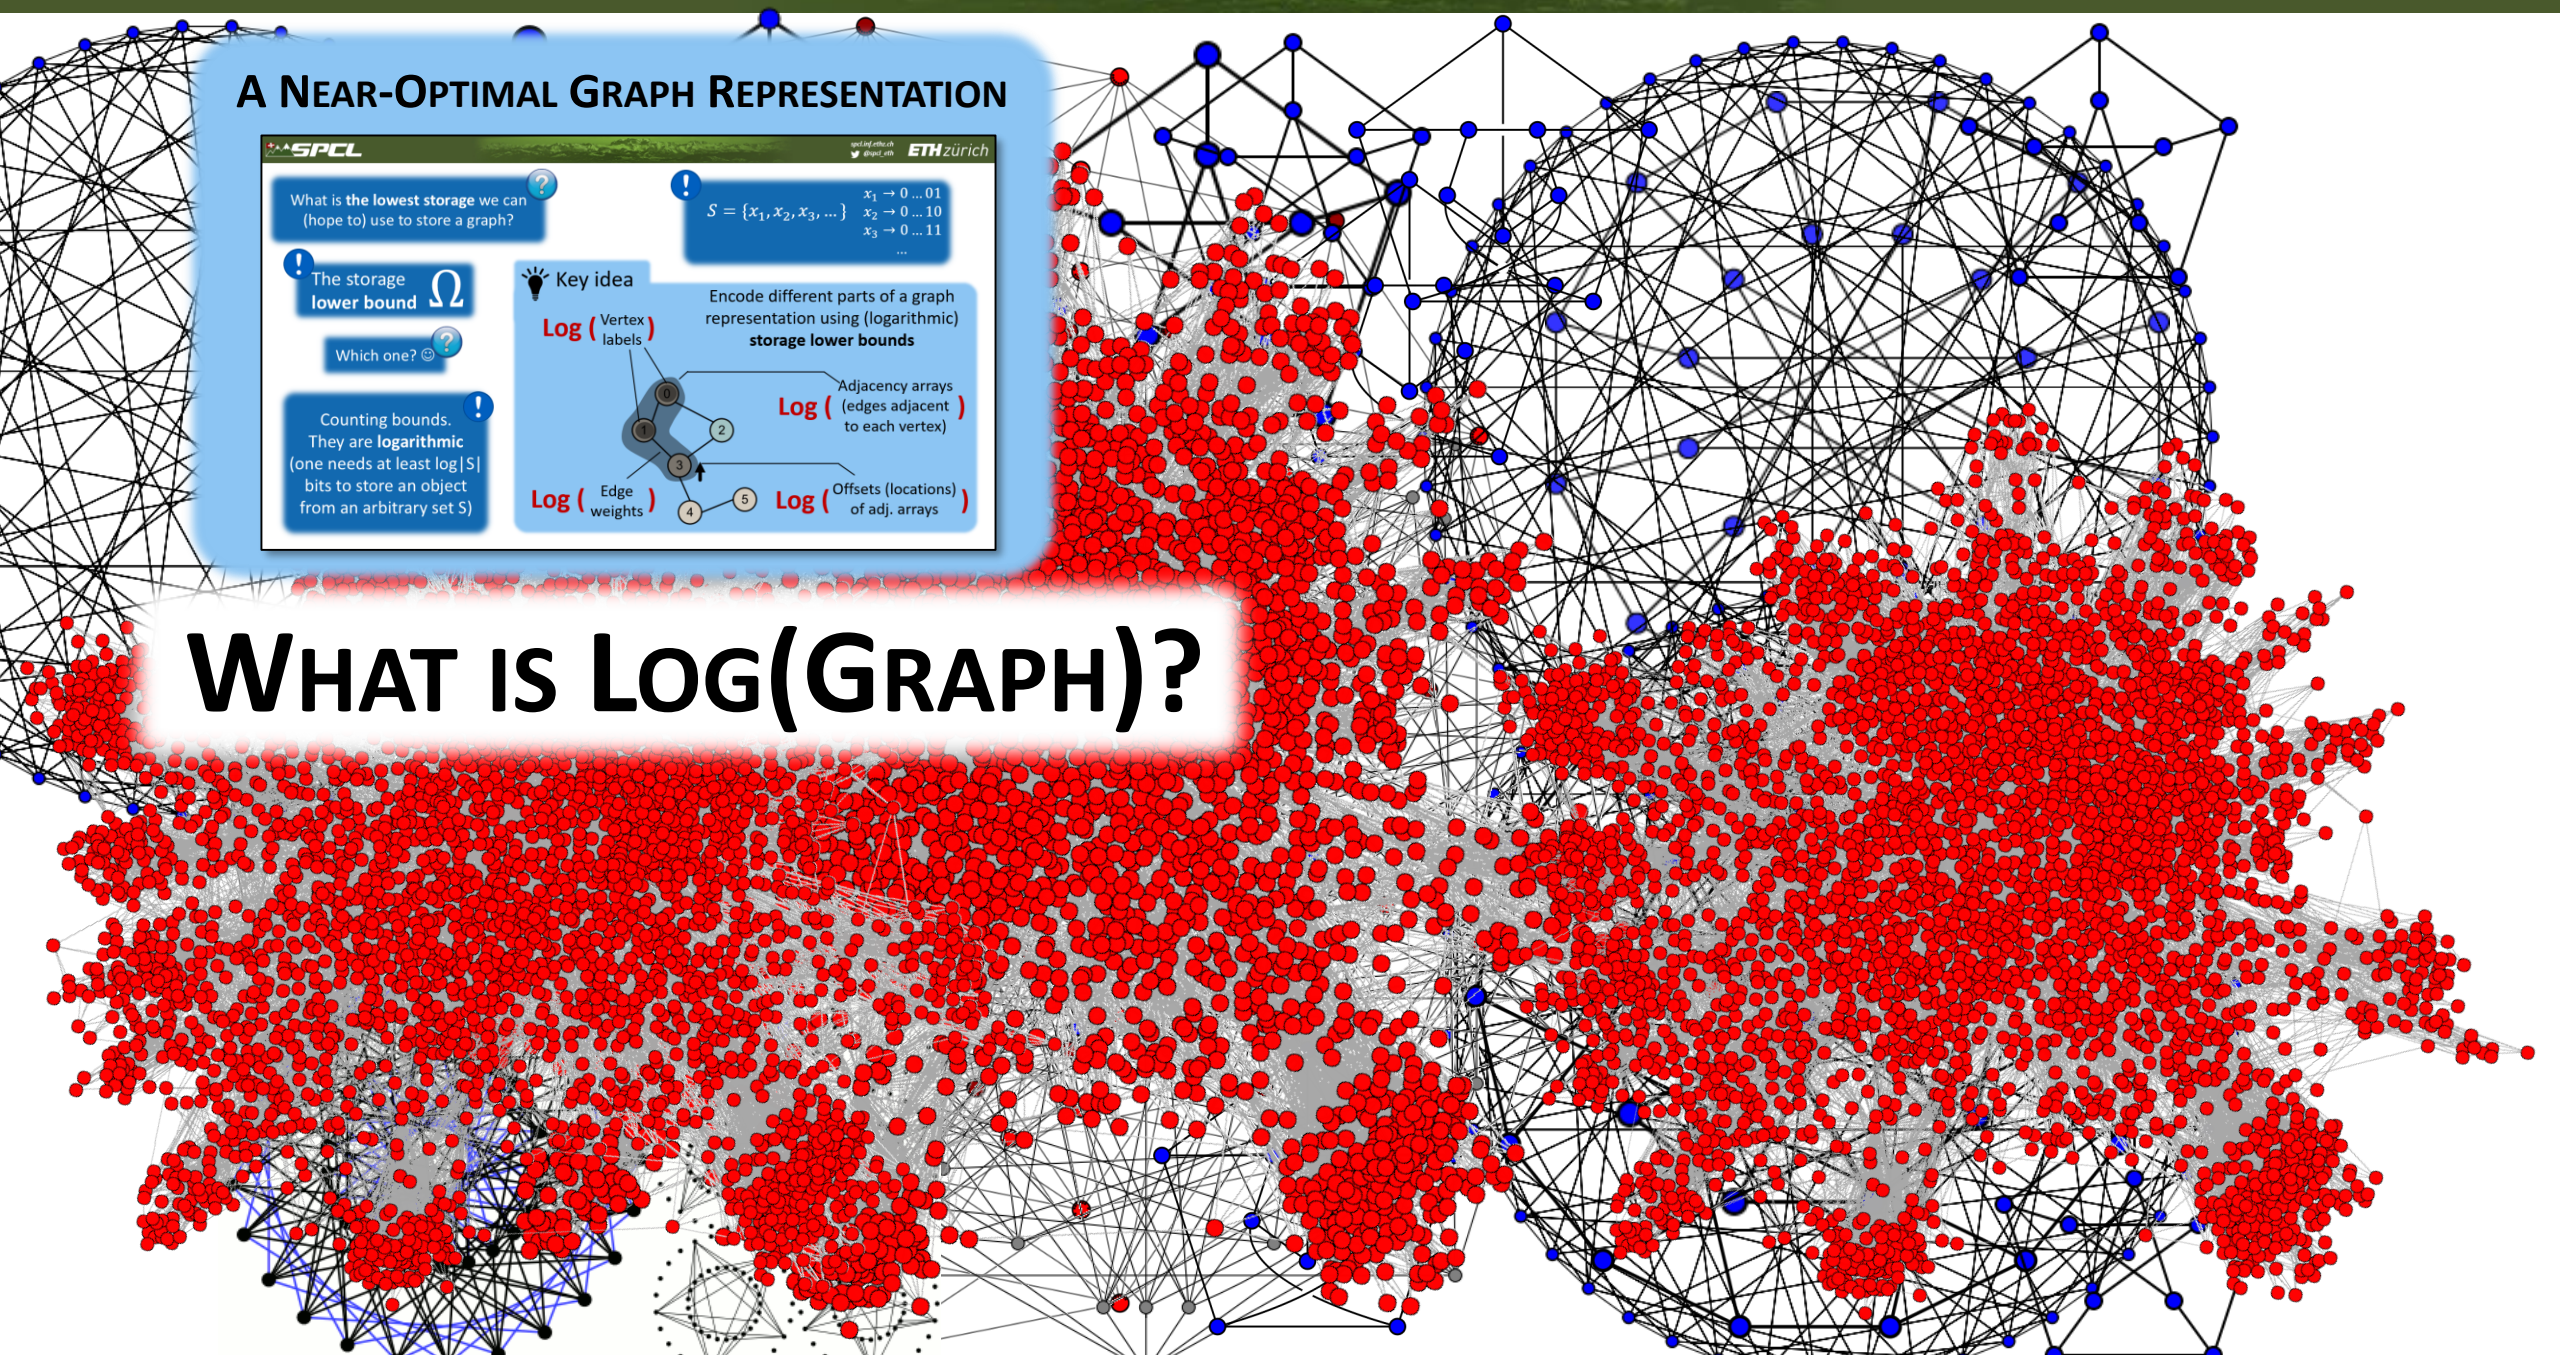
Encode different parts of a graph representation using (logarithmic) storage lower bounds

- $\text{Log}(\text{Vertex labels})$
- $\text{Log}(\text{Adjacency arrays (edges adjacent to each vertex)})$
- $\text{Log}(\text{Edge weights})$
- $\text{Log}(\text{Offsets (locations) of adj. arrays})$

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

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

WHAT IS LOG(GRAPH)?



A NEAR-OPTIMAL GRAPH REPRESENTATION

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

The storage lower bound Ω

Which one? ?

Counting bounds. They are 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

$\text{Log}(\text{Vertex labels})$
 $\text{Log}(\text{Edge weights})$
 $\text{Log}(\text{Adjacency arrays (edges adjacent to each vertex)})$
 $\text{Log}(\text{Offsets (locations) of adj. arrays})$

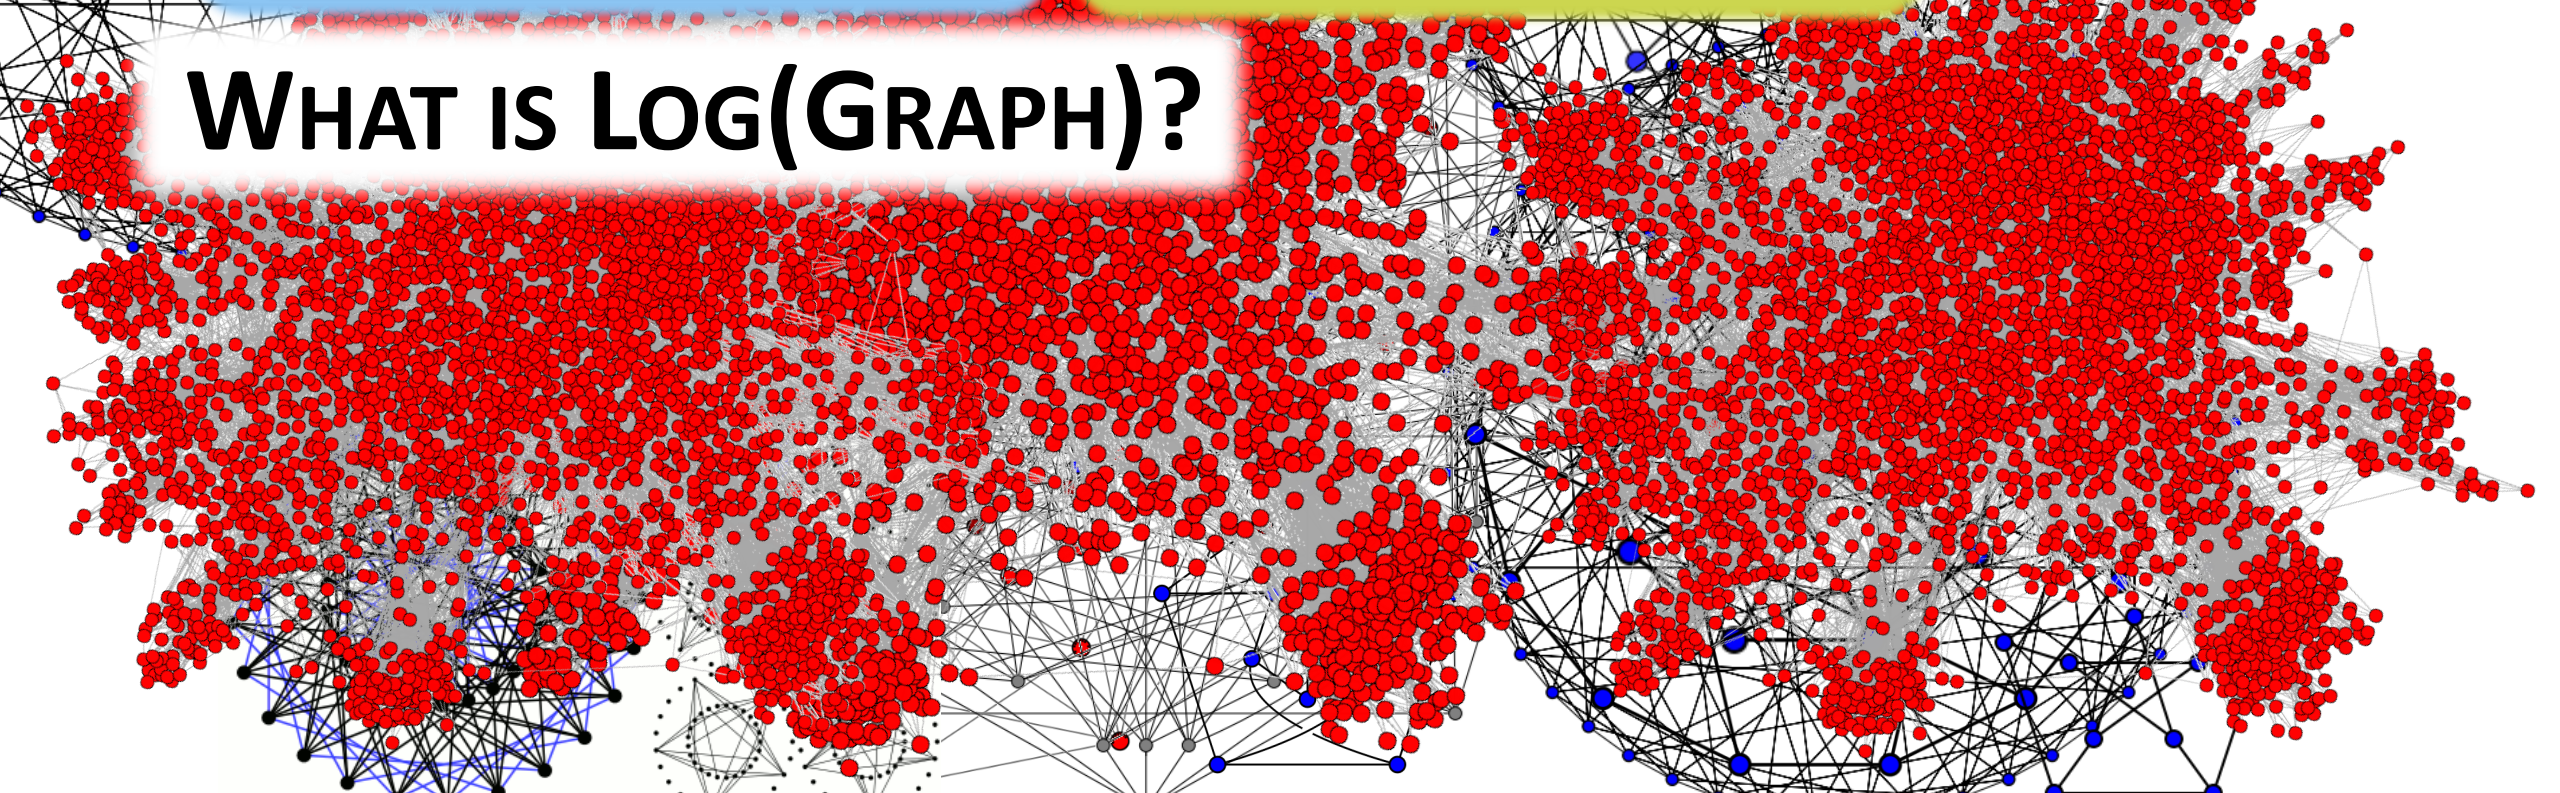
$S = \{x_1, x_2, x_3, \dots\}$
 $x_1 \rightarrow 0 \dots 01$
 $x_2 \rightarrow 0 \dots 10$
 $x_3 \rightarrow 0 \dots 11$
 \dots

AN EXTENSIBLE GRAPH REPRESENTATION

OVERVIEW OF FULL LOG(GRAPH) DESIGN

- Input structures
 - Graph G (S2)
 - Adjacency Array (S2)
- Logarithmize fine elements (S3)
 - Logarithmize vertex IDs... (S3.2)
 - Example ID: $\text{Log}(2) = \text{Log}(0\dots010_2) = 010_2$
 - globally (S3.2.1)
 - locally (S3.2.2)
 - on DM (S3.2.3) systems
 - Logarithmize other elements (S3.3 - S3.4)
- Logarithmize offset structure (S4)
 - Understand storage lower bounds (S4.1)
 - incorporate (S4.3) succinctness
 - theory (S4.4)
 - performance implementation (S4.5)
- Logarithmize adjacency structure (S5)
 - Unify (S5.1) with $P+T$ (S5.2)
 - incorporate recursive bisectioning (S5.3)
 - Use DIM (S5.4)
 - Use RB (S5.3.1)
 - Use BRB (S5.3.2)
- High-performance extensible library (S6)
 - Use ILP

WHAT IS LOG(GRAPH)?



A NEAR-OPTIMAL GRAPH REPRESENTATION

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

The storage lower bound Ω

Which one? ?

Counting bounds. They are 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

$S = \{x_1, x_2, x_3, \dots\}$
 $x_1 \rightarrow 0 \dots 01$
 $x_2 \rightarrow 0 \dots 10$
 $x_3 \rightarrow 0 \dots 11$
 \dots

$\log(\text{Vertex labels})$
 $\log(\text{Adjacency arrays (edges adjacent to each vertex)})$
 $\log(\text{Edge weights})$
 $\log(\text{Offsets (locations) of adj. arrays})$

AN EXTENSIBLE GRAPH REPRESENTATION

OVERVIEW OF FULL LOG(GRAPH) DESIGN

- Input structures
- Logarithmize fine elements (\$3)
- Logarithmize offset structure (\$4)
- Logarithmize adjacency structure (\$5)

Log(2) = $\log(0 \dots 010_2) = 010_2$

Log(Adjacency Array (\$2))

Logarithmize vertex IDs... (\$3.2)

Logarithmize other elements (\$3.3 - \$3.4)

High-performance extensible library (\$6)

WHAT IS LOG(GRAPH)?

A HIGH-PERFORMANCE GRAPH REPRESENTATION

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

Log(Graph) accelerates GAPBS

Both storage and performance are improved simultaneously

“LG”: Log(Graph)
 Trad: Traditional (non compressed, GAPBS)
 “g”: global scheme
 “l”: local scheme

Time [s]

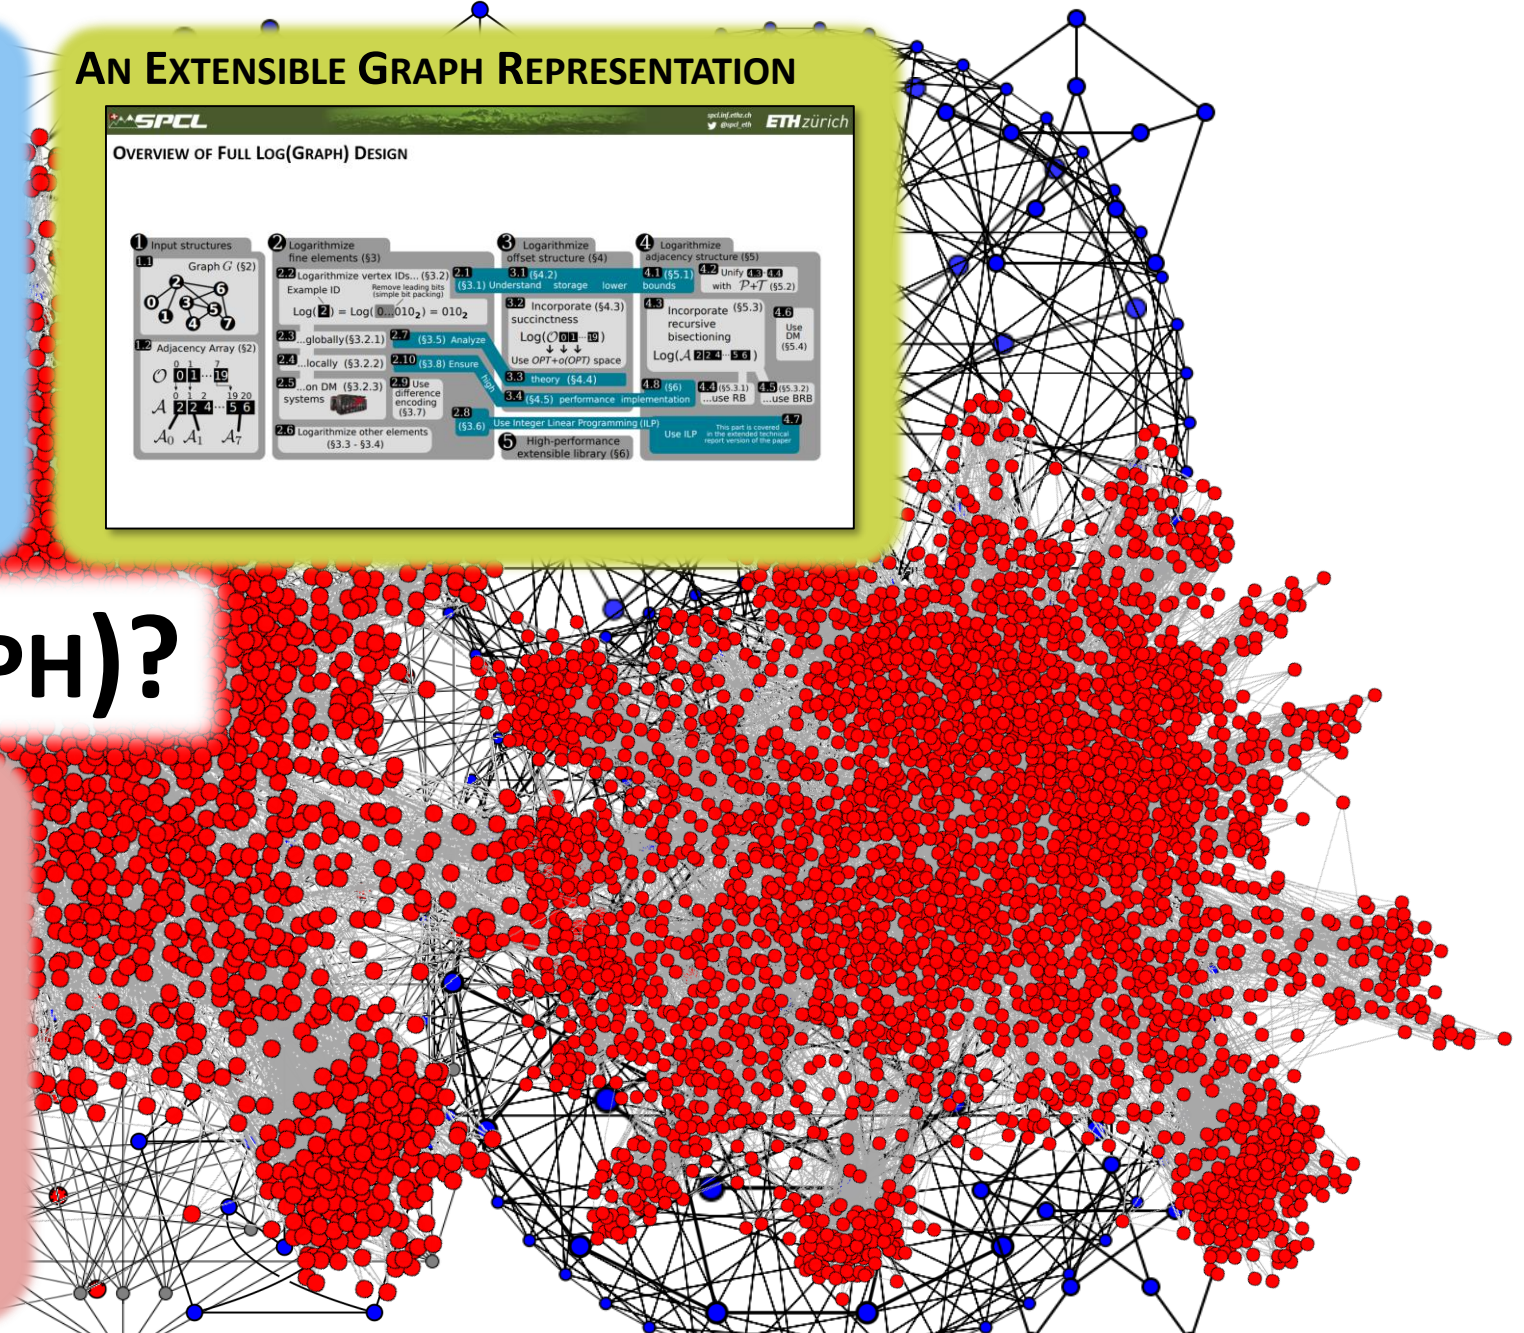
Sparse graphs

Dense graphs

Number of edges per vertex: 1, 2, 4, 8, 16, 32

Number of edges per vertex: 64, 128, 256, 512, 1024

Kronecker graphs
Number of vertices: 4M



A NEAR-OPTIMAL GRAPH REPRESENTATION

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

The storage lower bound Ω

Which one? ?

Counting bounds. They are 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

$S = \{x_1, x_2, x_3, \dots\}$
 $x_1 \rightarrow 0 \dots 01$
 $x_2 \rightarrow 0 \dots 10$
 $x_3 \rightarrow 0 \dots 11$
 \dots

$\log(\text{Vertex labels})$
 $\log(\text{Adjacency arrays (edges adjacent to each vertex)})$
 $\log(\text{Edge weights})$
 $\log(\text{Offsets (locations) of adj. arrays})$

AN EXTENSIBLE GRAPH REPRESENTATION

OVERVIEW OF FULL LOG(GRAPH) DESIGN

- Input structures
- Logarithmize fine elements (\$3)
- Logarithmize offset structure (\$4)
- Logarithmize adjacency structure (\$5)

2.1 Logarithmize vertex IDs... (\$3.2) Example ID $\log(2) = \log(0\dots010_2) = 010_2$

2.2 globally (\$3.2.1) 2.7 (\$3.5) Analyze 3.2 Incorporate (\$4.3) succinctness $\log(0\dots011\dots)$

2.3 locally (\$3.2.2) 2.10 (\$3.8) Ensure Use OPT+ α (OPT) space 3.3 theory (\$4.4)

2.4 on DM (\$3.2.3) 2.9 Use difference encoding (\$3.7) 3.4 (\$4.5) performance implementation 4.8 (\$6) 4.9 (\$5.3.1) use RB 4.5 (\$5.3.2) use BRB

2.5 Logarithmize other elements (\$3.3 - \$3.4) 2.8 (\$3.6) 5 Use Integer Linear Programming (ILP) Use ILP

4.1 (\$5.1) Unify with $P+T$ (\$5.2) 4.2 (\$5.3) Incorporate recursive bisectioning $\log(A)$ 4.6 (\$5.4) Use DM (\$5.4)

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

WHAT IS LOG(GRAPH)?

A HIGH-PERFORMANCE GRAPH REPRESENTATION

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

Log(Graph) accelerates GAPBS

“LG”: Log(Graph)
 Trad: Traditional (non compressed, GAPBS)
 “g”: global scheme
 “l”: local scheme

Both storage and performance are improved simultaneously

Time [s]

Sparse graphs

Dense graphs

Number of edges per vertex: 1 2 4 8 16 32

Number of vertices: 64 128 256 512 1024

Kronecker graphs
Number of vertices: 4M

Scheme: LG-g, LG-g-gap, LG-l, LG-l-gap, Trad

A CONDENSED GRAPH REPRESENTATION

3 $\log(\text{Adjacency structure})$ Storage

Trad: Traditional adjacency array
 DMd / DMf: Degree Minimizing (without / with gap encoding)
 WG: WebGraph compression
 BRB, RB: Schemes targeting certain specific classes of graphs

Relative size

Lots of data ?

Conclusions: WebGraph best for web graphs ?

DMf: much better than DMd, often comparable to others

BRB, RB: various tradeoffs but very expensive preprocessing (details in the paper)

Various real-world graphs: uk, uk, sk, ghno, Wb, lpd, wik, tra, orm, ljm, fr, tw, ork, Usrn

Scheme: Trad, DMd, DMf, BRB, RB, WG

A NEAR-OPTIMAL GRAPH REPRESENTATION

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

The storage lower bound Ω

Which one? ?

Counting bounds. They are 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

$S = \{x_1, x_2, x_3, \dots\}$
 $x_1 \rightarrow 0 \dots 01$
 $x_2 \rightarrow 0 \dots 10$
 $x_3 \rightarrow 0 \dots 11$
 \dots

$\log(\text{Vertex labels})$
 $\log(\text{Edge weights})$
 $\log(\text{Adjacency arrays (edges adjacent to each vertex)})$
 $\log(\text{Offsets (locations) of adj. arrays})$

AN EXTENSIBLE GRAPH REPRESENTATION

OVERVIEW OF FULL LOG(GRAPH) DESIGN

- Input structures
- Logarithmize fine elements (\$3)
- Logarithmize offset structure (\$4)
- Logarithmize adjacency structure (\$5)

2.1 Logarithmize vertex IDs... (\$3.2)
 Example ID: $\log(2) = \log(0 \dots 010_2) = 010_2$

2.2 globally (\$3.2.1) 2.7 (\$3.5) Analyze
 2.3 locally (\$3.2.2) 2.10 (\$3.8) Ensure succinctness
 2.4 on DM (\$3.2.3) 2.9 Use difference encoding (\$3.7)

3.1 (\$4.2) 3.2 Incorporate (\$4.3) Use $OPT+o(OPT)$ space
 3.3 theory (\$4.4) 3.4 (\$4.5) performance implementation

4.1 (\$5.1) 4.2 Unify (\$5.2) with $P+T$ (\$5.2)
 4.3 Incorporate recursive bisectioning
 4.4 (\$5.3) Use DM (\$5.4)
 4.5 (\$5.3.2) 4.6 (\$5.3.2) 4.7

5 High-performance extensible library (\$6)

WHAT IS LOG(GRAPH)?

A HIGH-PERFORMANCE GRAPH REPRESENTATION

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

Log(Graph) accelerates GAPBS

Both storage and performance are improved simultaneously

“LG”: Log(Graph)
 Trad: Traditional (non compressed, GAPBS)
 “g”: global scheme
 “l”: local scheme

Time [s]

Sparse graphs

Dense graphs

Number of edges per vertex: 1 2 4 8 16 32

Number of vertices: 64 128 256 512 1024

Kronecker graphs
 Number of vertices: 4M

Scheme: LG-g, LG-g-gap, LG-l, LG-l-gap, Trad

A CONDENSED GRAPH REPRESENTATION

3 $\log(\text{Adjacency structure})$ Storage

Trad: Traditional adjacency array
 DMd / DMf: Degree Minimizing (without / with gap encoding)
 WG: WebGraph compression
 BRB, RB: Schemes targeting certain specific classes of graphs

Relative size

Lots of data !

Conclusions: WebGraph best for web graphs !

DMf: much better than DMd, often comparable to others

BRB, RB: various tradeoffs but very expensive preprocessing (details in the paper)

Various real-world graphs: uklu, uk, sk, ghno, Wb, lpd, wik, tra, orm, ljm, fr, tw, ork, Usrn

Scheme: Trad, DMd, DMf, BRB, RB, WG

Website:
<http://spcl.inf.ethz.ch/Research/Performance/LogGraph>

A NEAR-OPTIMAL GRAPH REPRESENTATION

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

The storage lower bound Ω

Which one? ?

Counting bounds. They are 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

$S = \{x_1, x_2, x_3, \dots\}$
 $x_1 \rightarrow 0 \dots 01$
 $x_2 \rightarrow 0 \dots 10$
 $x_3 \rightarrow 0 \dots 11$
 \dots

$\log(\text{Vertex labels})$
 $\log(\text{Adjacency arrays (edges adjacent to each vertex)})$
 $\log(\text{Edge weights})$
 $\log(\text{Offsets (locations) of adj. arrays})$

AN EXTENSIBLE GRAPH REPRESENTATION

OVERVIEW OF FULL LOG(GRAPH) DESIGN

- Input structures (Graph G (\$2))
- Logarithmize fine elements (\$3)
- Logarithmize offset structure (\$4)
- Logarithmize adjacency structure (\$5)

Additional steps include: Logarithmize vertex IDs, globally/locality, on DM systems, Logarithmize other elements, Incorporate recursive bisectioning, and Use Integer Linear Programming (ILP).

WHAT IS LOG(GRAPH)?

A HIGH-PERFORMANCE GRAPH REPRESENTATION

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

Log(Graph) accelerates GAPBS

Both storage and performance are improved simultaneously

“LG”: Log(Graph)
 Trad: Traditional (non compressed, GAPBS)
 “g”: global scheme
 “l”: local scheme

Number of vertices: 4M

A CONDENSED GRAPH REPRESENTATION

3 $\log(\text{Adjacency structure})$ Storage

Trad: Traditional adjacency array
 DMd / DMf: Degree Minimizing (without / with gap encoding)
 WG: WebGraph compression
 BRB, RB: Schemes targeting certain specific classes of graphs

Relative size

Various real-world graphs

Conclusions:
 WebGraph best for web graphs
 DMf: much better than DMd, often comparable to others
 BRB, RB: various tradeoffs but very expensive preprocessing (details in the paper)

COMING SOON

Website:
<http://spcl.inf.ethz.ch/Research/Performance/LogGraph>

A NEAR-OPTIMAL GRAPH REPRESENTATION

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

The storage lower bound Ω

Which one? ?

Counting bounds. They are 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

- $\text{Log}(\text{Vertex labels})$
- $\text{Log}(\text{Adjacency arrays (edges adjacent to each vertex)})$
- $\text{Log}(\text{Edge weights})$
- $\text{Log}(\text{Offsets (locations) of adj. arrays})$

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

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

AN EXTENSIBLE GRAPH REPRESENTATION

OVERVIEW OF FULL LOG(GRAPH) DESIGN

- Input structures
- Logarithmize fine elements (\$3)
- Logarithmize offset structure (\$4)
- Logarithmize adjacency structure (\$5)

Log(2) = $\text{Log}(0 \dots 010_2) = 010_2$

Log(1) = $\text{Log}(0 \dots 001_2) = 001_2$

Log(4) = $\text{Log}(0 \dots 100_2) = 100_2$

Log(8) = $\text{Log}(0 \dots 1000_2) = 1000_2$

Log(16) = $\text{Log}(0 \dots 10000_2) = 10000_2$

Log(32) = $\text{Log}(0 \dots 100000_2) = 100000_2$

Log(64) = $\text{Log}(0 \dots 1000000_2) = 1000000_2$

Log(128) = $\text{Log}(0 \dots 10000000_2) = 10000000_2$

Log(256) = $\text{Log}(0 \dots 100000000_2) = 100000000_2$

Log(512) = $\text{Log}(0 \dots 1000000000_2) = 1000000000_2$

Log(1024) = $\text{Log}(0 \dots 10000000000_2) = 10000000000_2$



COMING SOON

Website:

<http://spcl.inf.ethz.ch/Research/Performance/LogGraph>

WHAT IS LOG(GRAPH)?

Thank you for your attention

A HIGH-PERFORMANCE GRAPH REPRESENTATION

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

Log(Graph) accelerates GAPBS

Both storage and performance are improved simultaneously

“LG”: Log(Graph)
“Trad”: Traditional (non compressed, GAPBS)
“g”: global scheme
“l”: local scheme

Time [s]

Sparse graphs

Dense graphs

Number of edges per vertex: 1, 2, 4, 8, 16, 32

Number of vertices: 64, 128, 256, 512, 1024

Kronecker graphs

Number of vertices: 4M

BFS

A CONDENSED GRAPH REPRESENTATION

3 $\text{Log}(\text{Adjacency structure})$ Storage

Trad: Traditional adjacency array
DMd / DMf: Degree Minimizing (without / with gap encoding)
WG: WebGraph compression
BRB, RB: Schemes targeting certain specific classes of graphs

Relative size

Lots of data ?


Conclusions:

- WebGraph best for web graphs ?
- DMf: much better than DMd, often comparable to others
- BRB, RB: various tradeoffs but very expensive preprocessing (details in the paper)


Various real-world graphs

Uku, Uk, sk, ghno, Wb, lpd, Wik, Ita, orm, ljm, fr, tw, ork, Usrn


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

\widehat{W} : max edge weight,


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

\widehat{W} : max edge weight,


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)

$\lceil \log n \rceil$

Symbols

\widehat{W} : max edge weight,


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)

$\lceil \log n \rceil$ $\lceil \log \hat{W} \rceil$

Symbols

\hat{W} : max edge weight,


n : #vertices,

m : #edges,

d_v : degree of vertex v ,

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

\hat{N}_v : maximum among N_v



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

Lower bounds (global)

$\lceil \log n \rceil$ $\lceil \log \hat{W} \rceil$

This is it?
Not really 😊



Symbols

\hat{W} : max edge weight,


n : #vertices,

m : #edges,

d_v : degree of vertex v ,

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

\hat{N}_v : maximum among N_v



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

Lower bounds (global)

$\lceil \log n \rceil$ $\lceil \log \hat{W} \rceil$

This is it?
Not really 😊



Symbols

\hat{W} : max edge weight,


n : #vertices,

m : #edges,

d_v : degree of vertex v ,

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

\hat{N}_v : maximum among N_v



Lower bounds (local)

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

Lower bounds (global)

$\lceil \log n \rceil$ $\lceil \log \hat{W} \rceil$

This is it?
Not really 😊



Symbols

\hat{W} : max edge weight,


n : #vertices,

m : #edges,

d_v : degree of vertex v ,

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

\hat{N}_v : maximum among N_v



Lower bounds (local)

Assume:

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

Lower bounds (global)

$\lceil \log n \rceil$ $\lceil \log \hat{W} \rceil$

This is it?
Not really 😊



Symbols

\hat{W} : max edge weight,


n : #vertices,

m : #edges,

d_v : degree of vertex v ,

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

\hat{N}_v : maximum among N_v



Lower bounds (local)

Assume:

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

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

Lower bounds (global)

$\lceil \log n \rceil$ $\lceil \log \hat{W} \rceil$

This is it?
Not really 😊



Symbols

\hat{W} : max edge weight,


n : #vertices,

m : #edges,

d_v : degree of vertex v ,

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

\hat{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)

$\lceil \log n \rceil$ $\lceil \log \hat{W} \rceil$

This is it?
Not really 😊



Symbols

\hat{W} : max edge weight,


n : #vertices,

m : #edges,

d_v : degree of vertex v ,

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

\hat{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: $\hat{N}_v \ll n$

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

Lower bounds (global)

$\lceil \log n \rceil$ $\lceil \log \widehat{W} \rceil$

This is it?
Not really 😊



Symbols

\widehat{W} : max edge weight,


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)


$$\lceil \log n \rceil \quad \lceil \log \hat{W} \rceil$$

This is it?
Not really 😊



Symbols

- \hat{W} : max edge weight,
- n : #vertices,
- m : #edges,
- d_v : degree of vertex v ,
- N_v : neighbors (adj. array) of vertex v ,
- \hat{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: $\hat{N}_v \ll n$

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



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

Lower bounds (global)


$$\lceil \log n \rceil \quad \lceil \log \hat{W} \rceil$$

This is it?
Not really 😊



Symbols

- \hat{W} : max edge weight,
- n : #vertices,
- m : #edges,
- d_v : degree of vertex v ,
- N_v : neighbors (adj. array) of vertex v ,
- \hat{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: $\hat{N}_v \ll n$

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




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

Lower bounds (global)
 $\lceil \log n \rceil$ $\lceil \log \hat{W} \rceil$

This is it?
 Not really 😊

Symbols

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



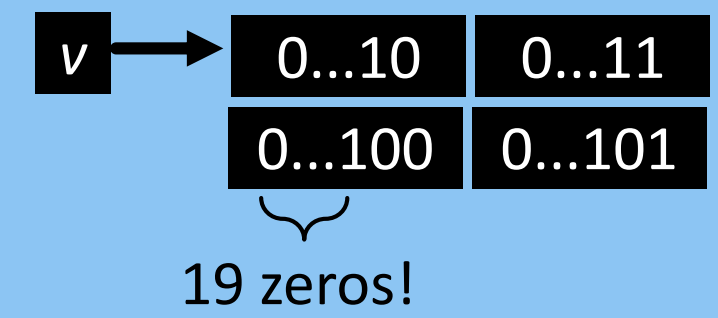
Lower bounds (local)

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: $\hat{N}_v \ll n$



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




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

Lower bounds (global)
 $\lceil \log n \rceil$ $\lceil \log \hat{W} \rceil$

This is it?
 Not really 😊

Symbols

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



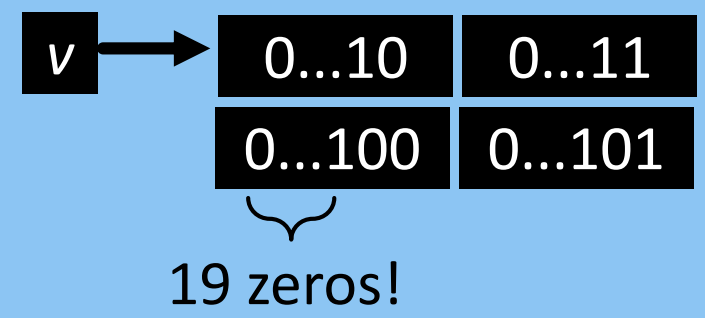
Lower bounds (local)

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: $\hat{N}_v \ll n$



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




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

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

Symbols

n : #vertices,
 m : #edges,
 H : number of compute nodes,
 H_i : number of machine elements at level i ,
 N : number of machine levels




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

This is it? Still not really 😊



Symbols

n : #vertices,
 m : #edges,
 H : number of compute nodes,
 H_i : number of machine elements at level i ,
 N : number of machine levels



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


This is it? Still not really 😊



Lower bounds (local):
distributed memories

Symbols

n : #vertices,
 m : #edges,
 H : number of compute nodes,
 H_i : number of machine elements at level i ,
 N : number of machine levels



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


This is it? Still not really 😊 

**Lower bounds (local):
distributed memories**



A Cray XE/XT supercomputer


Symbols
 n : #vertices,
 m : #edges,
 H : number of compute nodes,
 H_i : number of machine elements at level i ,
 N : number of machine levels




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

This is it? Still not really 😊

Symbols
 n : #vertices,
 m : #edges,
 H : number of compute nodes,
 H_i : number of machine elements at level i ,
 N : number of machine levels




**Lower bounds (local):
distributed memories**




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

This is it? Still not really 😊 ?

Symbols
 n : #vertices,
 m : #edges,
 H : number of compute nodes,
 H_i : number of machine elements at level i ,
 N : number of machine levels




**Lower bounds (local):
distributed memories**




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

This is it? Still not really 😊 ?

Symbols
 n : #vertices,
 m : #edges,
 H : number of compute nodes,
 H_i : number of machine elements at level i ,
 N : number of machine levels




Lower bounds (local): distributed memories



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

This is it? Still not really 😊 ?

Symbols
 n : #vertices,
 m : #edges,
 H : number of compute nodes,
 H_i : number of machine elements at level i ,
 N : number of machine levels




Lower bounds (local): distributed memories

 A Cray XE/XT supercomputer

4 cabinets:



3 chassis:



8 blades:



4 nodes:




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

This is it? Still not really 😊 ?

Symbols

- n : #vertices,
- m : #edges,
- H : number of compute nodes,
- H_i : number of machine elements at level i ,
- N : number of machine levels


Lower bounds (local): distributed memories




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

This is it? Still not really 😊 ?

Symbols
 n : #vertices,
 m : #edges,
 H : number of compute nodes,
 H_i : number of machine elements at level i ,
 N : number of machine levels



Lower bounds (local): distributed memories




4 cabinets:
 3 chassis:
 8 blades:
 4 nodes:
 32 cores:


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

This is it? Still not really 😊 ?

Symbols
 n : #vertices,
 m : #edges,
 H : number of compute nodes,
 H_i : number of machine elements at level i ,
 N : number of machine levels



Lower bounds (local): distributed memories



4 cabinets:
 3 chassis:
 8 blades:
 4 nodes:
 $H = 4$
 32 cores:


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

This is it? Still not really 😊

Symbols
 n : #vertices,
 m : #edges,
 H : number of compute nodes,
 H_i : number of machine elements at level i ,
 N : number of machine levels

Lower bounds (local): distributed memories

The number of vertices that can be stored in the memory of one node:



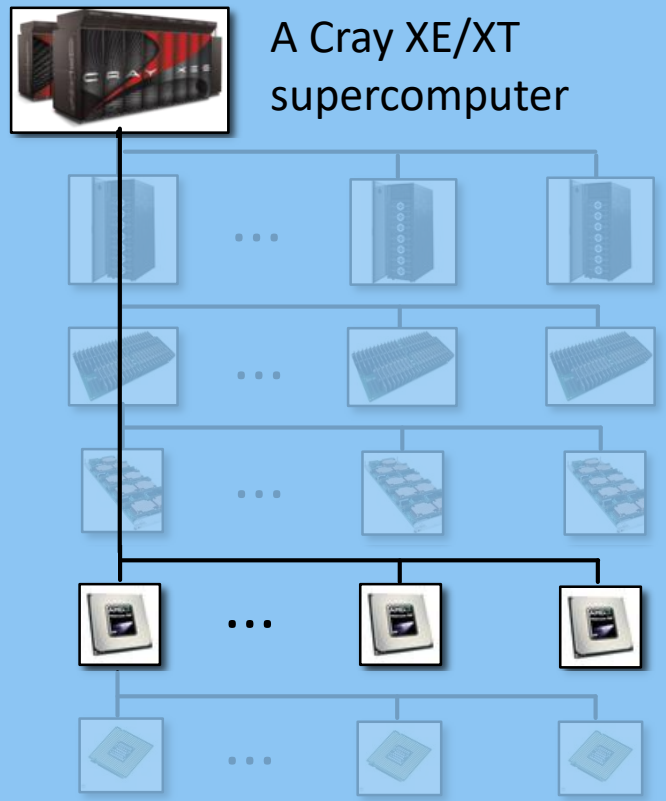
1 **Log** (Vertex labels), **Log** (Edge weights)

This is it? Still not really 😊

Symbols
 n : #vertices,
 m : #edges,
 H : number of compute nodes,
 H_i : number of machine elements at level i ,
 N : number of machine levels

Lower bounds (local): distributed memories


The number of vertices that can be stored in the memory of one node: $\frac{n}{H}$



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

This is it? Still not really 😊 ?


Symbols
 n : #vertices,
 m : #edges,
 H : number of compute nodes,
 H_i : number of machine elements at level i ,
 N : number of machine levels



Lower bounds (local): distributed memories

The number of vertices that can be stored in the memory of one node: $\frac{n}{H}$

The „intra-node” vertex label thus takes [bits]: $\left\lceil \log \frac{n}{H} \right\rceil$




4 cabinets:
 3 chassis:
 8 blades:
 4 nodes:
 $H = 4$
 32 cores:

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

This is it? Still not really 😊 ?

Symbols
 n : #vertices,
 m : #edges,
 H : number of compute nodes,
 H_i : number of machine elements at level i ,
 N : number of machine levels




Lower bounds (local): distributed memories

The number of vertices that can be stored in the memory of one node: $\frac{n}{H}$

The „intra-node” vertex label thus takes [bits]: $\lceil \log \frac{n}{H} \rceil$

The „inter-node” vertex label is unique for a whole node and it takes [bits]: $\lceil \log H \rceil$



A Cray XE/XT supercomputer

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

This is it? Still not really 😊 ?

Symbols


- n : #vertices,
- m : #edges,
- H : number of compute nodes,
- H_i : number of machine elements at level i ,
- N : number of machine levels

Lower bounds (local): distributed memories

The number of vertices that can be stored in the memory of one node: $\frac{n}{H}$

The „intra-node” vertex label thus takes [bits]: $\left\lceil \log \frac{n}{H} \right\rceil$

The „inter-node” vertex label is unique for a whole node and it takes [bits]: $\lceil \log H \rceil$



The **total size of the adjacency arrays** is thus [bits]:

$$n \left\lceil \log \frac{n}{H} \right\rceil + H \lceil \log H \rceil$$

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

This is it? Still not really 😊 ?

Symbols


- n : #vertices,
- m : #edges,
- H : number of compute nodes,
- H_i : number of machine elements at level i ,
- N : number of machine levels

Lower bounds (local): distributed memories

The number of vertices that can be stored in the memory of one node: $\frac{n}{H}$

The „intra-node” vertex label thus takes [bits]: $\left\lceil \log \frac{n}{H} \right\rceil$

The „inter-node” vertex label is unique for a whole node and it takes [bits]: $\lceil \log H \rceil$



The **total size of the adjacency arrays** is thus [bits]:

$$n \left\lceil \log \frac{n}{H} \right\rceil + H \lceil \log H \rceil$$

We also generalize this to arbitrarily many levels (details in the paper 😊) and derive the total size:

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

This is it? Still not really 😊 ?

Symbols


- n : #vertices,
- m : #edges,
- H : number of compute nodes,
- H_i : number of machine elements at level i ,
- N : number of machine levels

Lower bounds (local): distributed memories

The number of vertices that can be stored in the memory of one node: $\frac{n}{H}$

The „intra-node” vertex label thus takes [bits]: $\left\lceil \log \frac{n}{H} \right\rceil$

The „inter-node” vertex label is unique for a whole node and it takes [bits]: $\lceil \log H \rceil$



The **total size of the adjacency arrays** is thus [bits]:

$$n \left\lceil \log \frac{n}{H} \right\rceil + H \lceil \log H \rceil$$

We also generalize this to arbitrarily many levels (details in the paper 😊) and derive the total size:

$$n \left\lceil \log \frac{n}{H_N} \right\rceil + \sum_{j=2}^{N-1} H_j \lceil \log H_j \rceil$$

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

fff Formal analyses: more
(check the paper 😊)

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

fff Formal analyses: more (check the paper 😊)

$$|\mathcal{A}| = \sum_{v \in V} \left(d_v \left\lceil \log \widehat{N}_v \right\rceil + \left\lceil \log \log \widehat{N}_v \right\rceil \right)$$

$$|\mathcal{A}| = n \left\lceil \log \frac{n}{\mathcal{H}} \right\rceil + \mathcal{H} \left\lceil \log \mathcal{H} \right\rceil$$

$$E[|\mathcal{A}|] \approx \frac{\alpha}{2-\beta} \left(\left(\frac{\alpha n \log n}{\beta-1} \right)^{\frac{2-\beta}{\beta-1}} - 1 \right) \left(\left\lceil \log n \right\rceil + \left\lceil \log \widehat{\mathcal{W}} \right\rceil \right)$$

$$E[|\mathcal{O}|] = n \left\lceil \log (2pn^2) \right\rceil = n \left\lceil \log 2p + 2 \log n \right\rceil$$

$$\forall v, u \in V (u \in N_v) \Rightarrow \left\lceil \mathcal{N}(u) \right\rceil \leq \widehat{N}_v$$

$$|\mathcal{A}| = \sum_{v \in V} \left(d_v \left\lceil \log \widehat{N}_v \right\rceil + \left\lceil \log \log \widehat{N}_v \right\rceil \right)$$

$$|\mathcal{A}| = 2m \left(\left\lceil \log n \right\rceil + \left\lceil \log \widehat{\mathcal{W}} \right\rceil \right)$$

$$|\mathcal{A}| = \sum_{v \in V} \left(d_v \left(\left\lceil \log \widehat{N}_v \right\rceil + \left\lceil \log \widehat{\mathcal{W}} \right\rceil \right) + \left\lceil \log \log \widehat{N}_v \right\rceil + \left\lceil \log \log \widehat{\mathcal{W}} \right\rceil \right)$$

$$E[|\mathcal{A}|] = \left(\left\lceil \log n \right\rceil + \left\lceil \log \widehat{\mathcal{W}} \right\rceil \right) pn^2$$

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

Formal analyses: more (check the paper 😊)

```

1 /* Input: graph G, Output: a new relabeling  $\mathcal{N}(v), \forall v \in V$ . */
2 void relabel(G) {
3     ID[0..n-1] = [0..n-1]; //An array with vertex IDs.
4     D[0..n-1] = [d_0..d_{n-1}]; //An array with degrees of vertices.
5     //An auxiliary array for determining if a vertex was relabeled:
6     visit[0..n-1] = [false..false];
7     nl = 1; //An auxiliary variable ``new label``.
8     sort(ID); sort(D);
9     for(int i = 1; i < n; ++i) //For each vertex...
10    for(int j = 0; j < D[i]; ++j) { //For each neighbor...
11        int id = N_{j,ID[i]}; //N_{j,ID[i]} is jth neighbor of vertex with ID ID[i]
12        if(visit[id] == false) {
13             $\mathcal{N}(id) = nl++$ ;
14            visit[id] = true;
15        }
16    for(int i = 1; i < n; ++i)
17        if(visit[i] == false)
18             $\mathcal{N}(id) = nl++$ ;
19 }
    
```

$$E[|\mathcal{O}|] = n \lceil \log(2pn^2) \rceil = n \lceil \log 2p + 2 \log n \rceil$$

$$\forall v, u \in V (u \in N_v) \Rightarrow \lceil \mathcal{N}(u) \rceil \leq \widehat{N}_v$$

$$|\mathcal{O}| = \sum_{v \in V} \left(d_v \lceil \log \widehat{N}_v \rceil + \lceil \log \log \widehat{N}_v \rceil \right)$$


$$|\mathcal{A}| = 2m \left(\lceil \log n \rceil + \lceil \log \widehat{\mathcal{W}} \rceil \right)$$

$$|\mathcal{A}| = \sum_{v \in V} \left(d_v \lceil \log \widehat{N}_v \rceil + \lceil \log \widehat{\mathcal{W}} \rceil \right) + \lceil \log \log \widehat{N}_v \rceil + \lceil \log \log \widehat{\mathcal{W}} \rceil$$

$$E[|\mathcal{A}|] = \left(\lceil \log n \rceil + \lceil \log \widehat{\mathcal{W}} \rceil \right) pn^2$$


2 Log (Offset structure)

...Encode the resulting bit vectors as
succinct bit vectors



2 Log (Offset structure)


...Encode the resulting bit vectors as
succinct bit vectors



Formal analyses

2 Log (Offset structure)

...Encode the resulting bit vectors as succinct bit vectors




Formal analyses

\mathcal{O}	ID	Asymptotic size [bits]	Exact size [bits]	select or $\mathcal{O}[v]$
Pointer array	ptrW	$O(Wn)$	$W(n + 1)$	$O(1)$
Plain [44]	bvPL	$O\left(\frac{Wm}{B}\right)$	$\frac{2Wm}{B}$	$O(1)$
Interleaved [44]	bvIL	$O\left(\frac{Wm}{B} + \frac{Wm}{L}\right)$	$2Wm\left(\frac{1}{B} + \frac{64}{L}\right)$	$O\left(\log \frac{Wm}{B}\right)$
Entropy based [31, 78]	bvEN	$O\left(\frac{Wm}{B} \log \frac{Wm}{B}\right)$	$\approx \log\left(\frac{2Wm}{Bn}\right)$	$O\left(\log \frac{Wm}{B}\right)$
Sparse [76]	bvSD	$O\left(n + n \log \frac{Wm}{Bn}\right)$	$\approx n\left(2 + \log \frac{2Wm}{Bn}\right)$	$O(1)$
B-tree based [1]	bvBT	$O\left(\frac{Wm}{B}\right)$	$\approx 1.1 \cdot \frac{2Wm}{B}$	$O(\log n)$
Gap-compressed [1]	bvGC	$O\left(\frac{Wm}{B} \log \frac{Wm}{Bn}\right)$	$\approx 1.3 \cdot \frac{2Wm}{B} \log \frac{2Wm}{Bn}$	$O(\log n)$

2 Log (Offset structure)

...Encode the resulting bit vectors as **succinct bit vectors** 




Formal analyses

Check the paper for details 😊

\mathcal{O}	ID	Asymptotic size [bits]	Exact size [bits]	select or $\mathcal{O}[v]$
Pointer array	ptrW	$O(Wn)$	$W(n + 1)$	$O(1)$
Plain [44]	bvPL	$O\left(\frac{Wm}{B}\right)$	$\frac{2Wm}{B}$	$O(1)$
Interleaved [44]	bvIL	$O\left(\frac{Wm}{B} + \frac{Wm}{L}\right)$	$2Wm\left(\frac{1}{B} + \frac{64}{L}\right)$	$O\left(\log \frac{Wm}{B}\right)$
Entropy based [31, 78]	bvEN	$O\left(\frac{Wm}{B} \log \frac{Wm}{B}\right)$	$\approx \log\left(\frac{2Wm}{Bn}\right)$	$O\left(\log \frac{Wm}{B}\right)$
Sparse [76]	bvSD	$O\left(n + n \log \frac{Wm}{Bn}\right)$	$\approx n\left(2 + \log \frac{2Wm}{Bn}\right)$	$O(1)$
B-tree based [1]	bvBT	$O\left(\frac{Wm}{B}\right)$	$\approx 1.1 \cdot \frac{2Wm}{B}$	$O(\log n)$
Gap-compressed [1]	bvGC	$O\left(\frac{Wm}{B} \log \frac{Wm}{Bn}\right)$	$\approx 1.3 \cdot \frac{2Wm}{B} \log \frac{2Wm}{Bn}$	$O(\log n)$

2 Log (Offset structure)

...Encode the resulting bit vectors as **succinct bit vectors**



Formal analyses

Check the paper for details ☺



\mathcal{O}	ID	Asymptotic size [bits]	Exact size [bits]	<i>select</i> or $\mathcal{O}[v]$
Pointer array	ptrW	$O(Wn)$	$W(n + 1)$	$O(1)$
Plain [44]	bvPL	$O\left(\frac{Wm}{B}\right)$	$\frac{2Wm}{B}$	$O(1)$
Interleaved [44]	bvIL	$O\left(\frac{Wm}{B} + \frac{Wm}{L}\right)$	$2Wm\left(\frac{1}{B} + \frac{64}{L}\right)$	$O\left(\log \frac{Wm}{B}\right)$
Entropy based [31, 78]	bvEN	$O\left(\frac{Wm}{B} \log \frac{Wm}{B}\right)$	$\approx \log\left(\frac{2Wm}{Bn}\right)$	$O\left(\log \frac{Wm}{B}\right)$
Sparse [76]	bvSD	$O\left(n + n \log \frac{Wm}{Bn}\right)$	$\approx n\left(2 + \log \frac{2Wm}{Bn}\right)$	$O(1)$
B-tree based [1]	bvBT	$O\left(\frac{Wm}{B}\right)$	$\approx 1.1 \cdot \frac{2Wm}{B}$	$O(\log n)$
Gap-compressed [1]	bvGC	$O\left(\frac{Wm}{B} \log \frac{Wm}{Bn}\right)$	$\approx 1.3 \cdot \frac{2Wm}{B} \log \frac{2Wm}{Bn}$	$O(\log n)$

Key methods

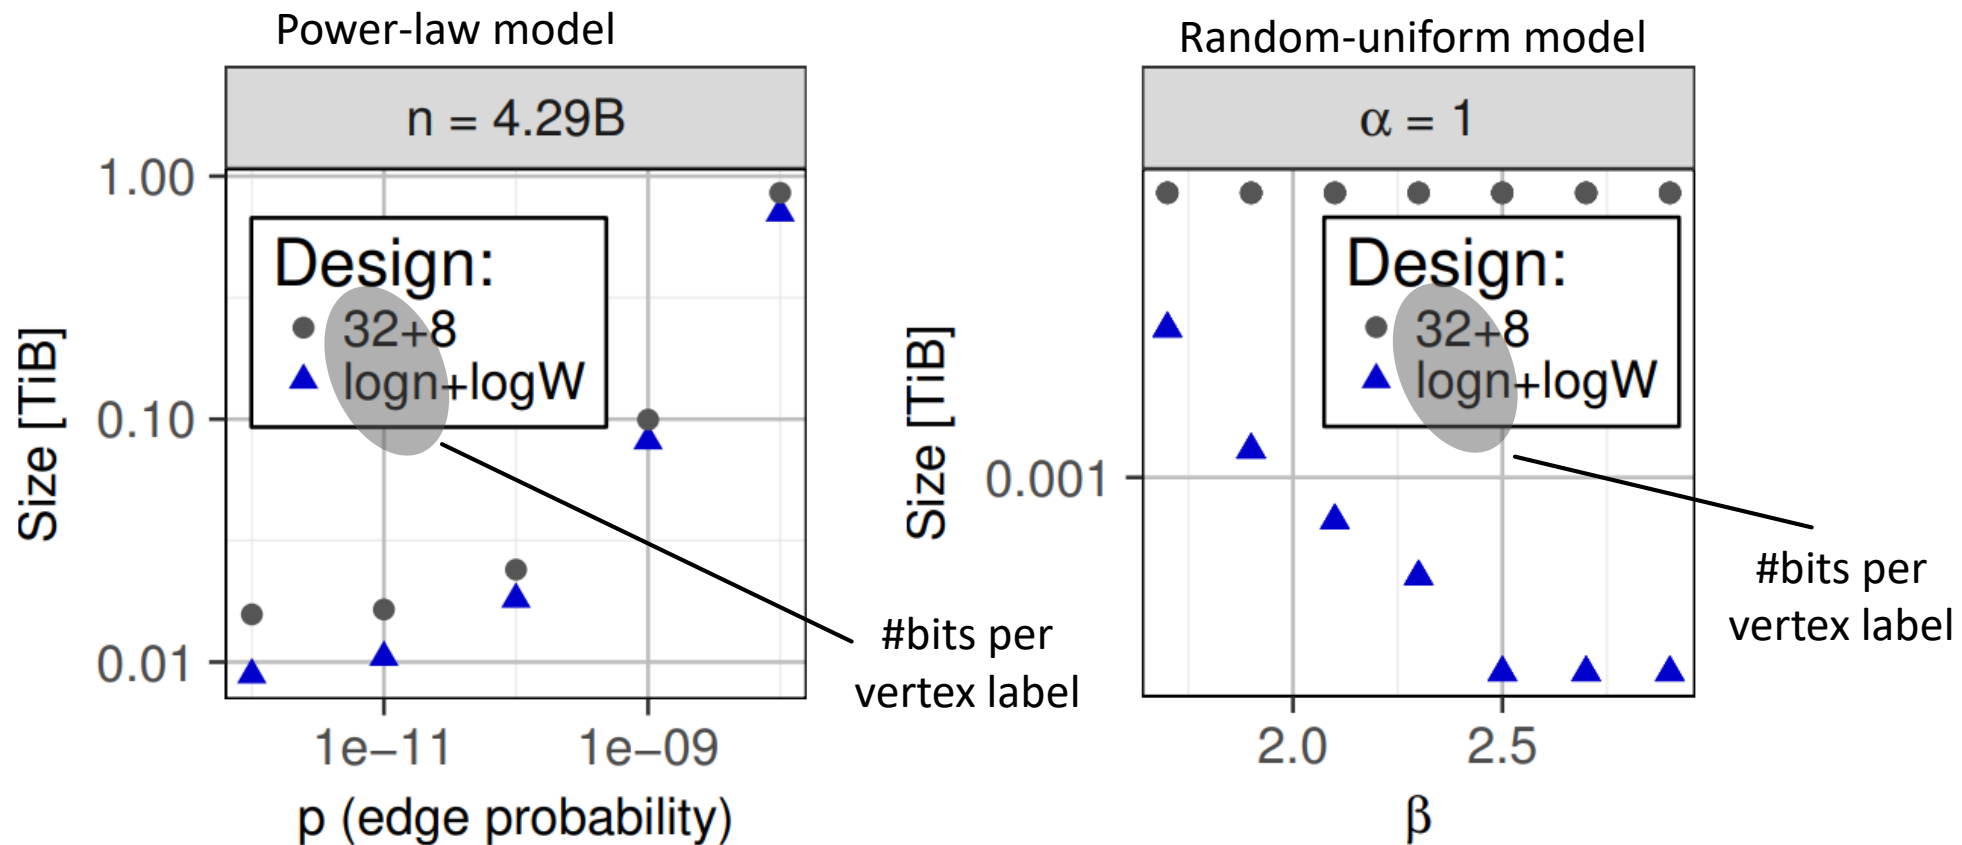
Use the *sdsl-lite* sequential library of succinct bit vectors [1] and investigate if it fares well when being accessed by multiple threads

[1] S. Gog. SDSL-Lite Succinct Library. 2015.

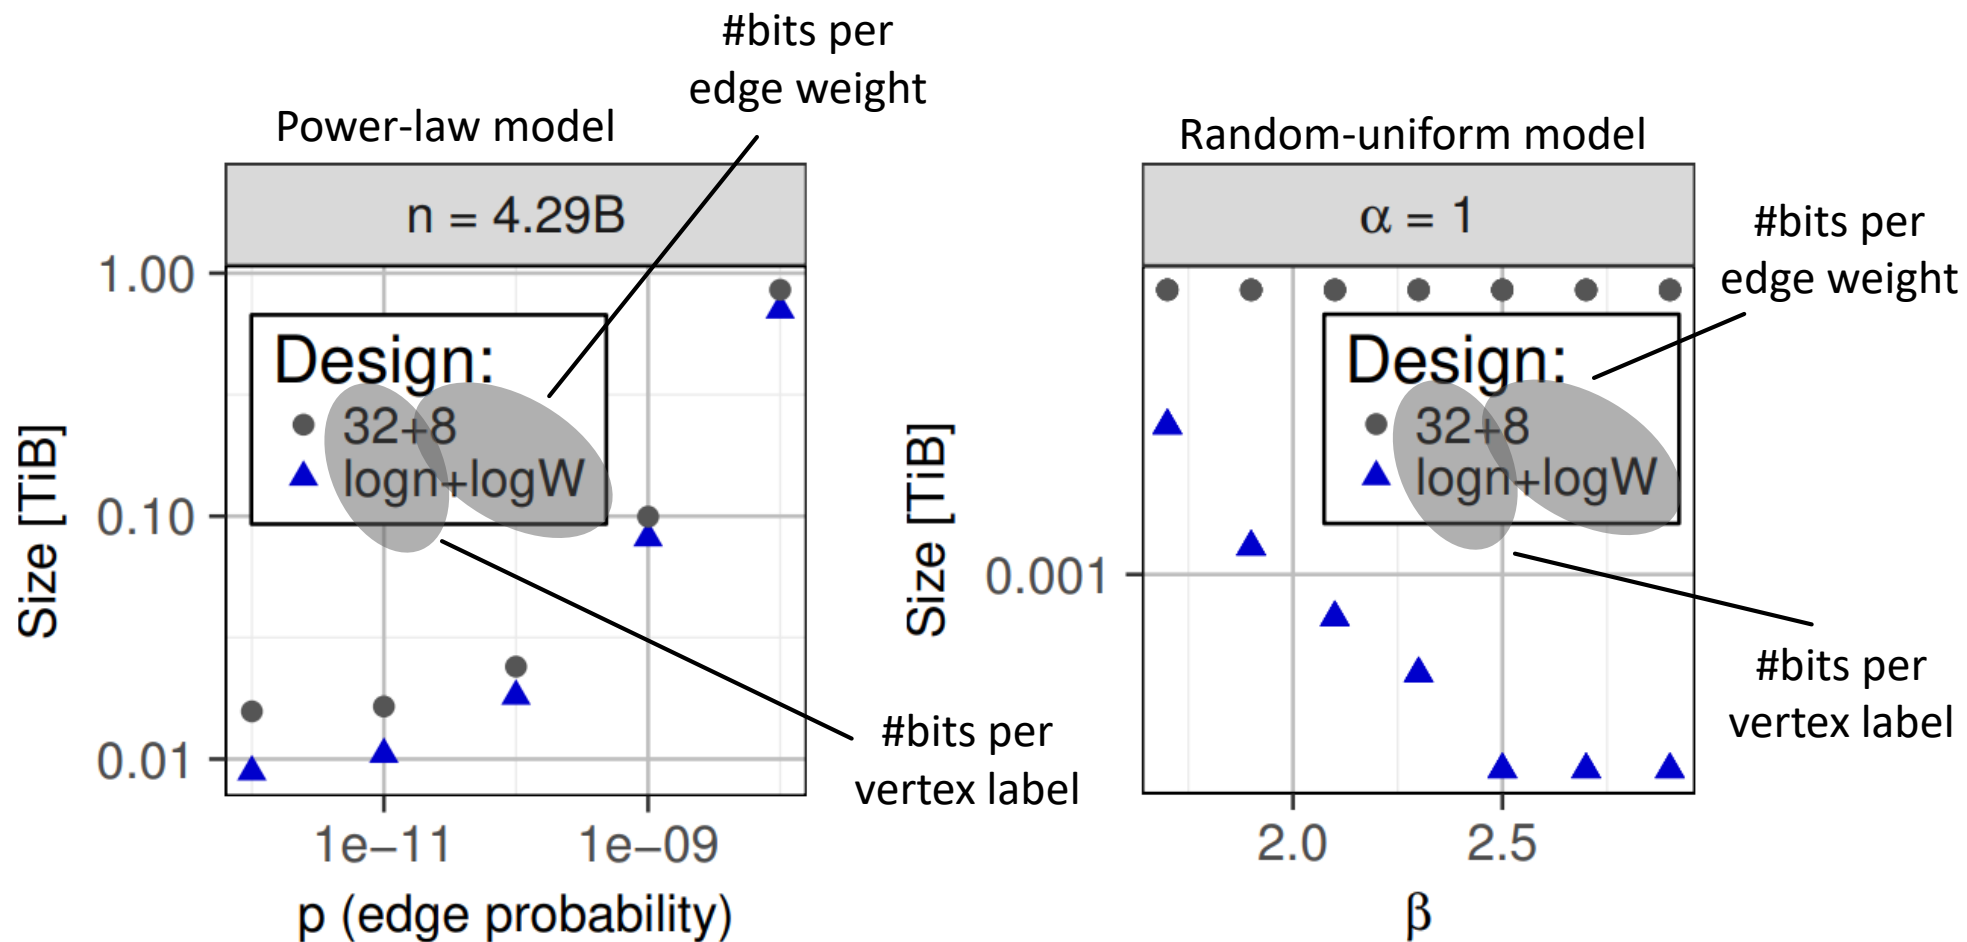
1
Log (Vertex labels), Log (Edge weights) Storage



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

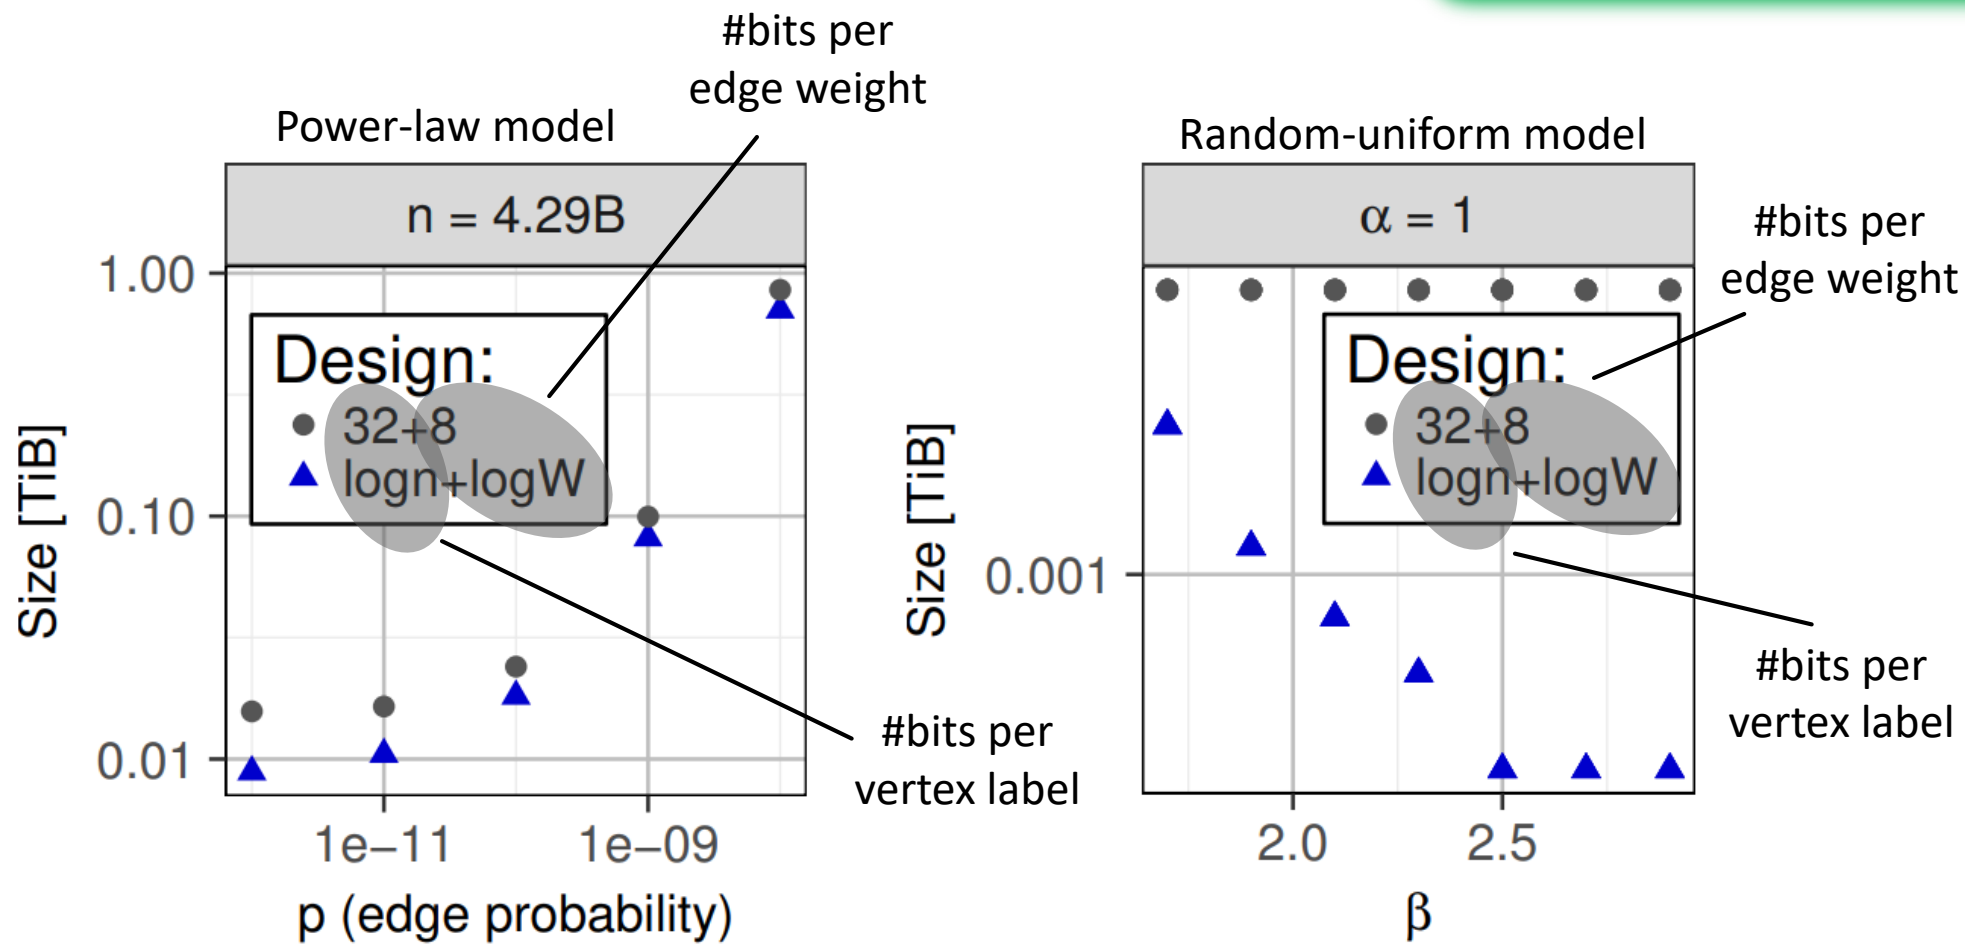


1 $\text{Log}(\text{Vertex labels}), \text{Log}(\text{Edge weights})$ Storage





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

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




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




Kronecker graphs
Number of vertices: 4M

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




Log(Graph) accelerates GAPBS




Kronecker graphs
Number of vertices: 4M

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



Log(Graph)
accelerates GAPBS

Both storage and
performance
are improved
simultaneously



Kronecker graphs
Number of vertices: 4M

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

Betweenness Centrality


“**LG**”: Log(Graph)

Trad: Traditional
(non compressed,
GAPBS)

“**g**”: global scheme

“**l**”: local scheme

“**gap**”: additional
gap encoding



Kronecker graphs


Number of vertices: 4M

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


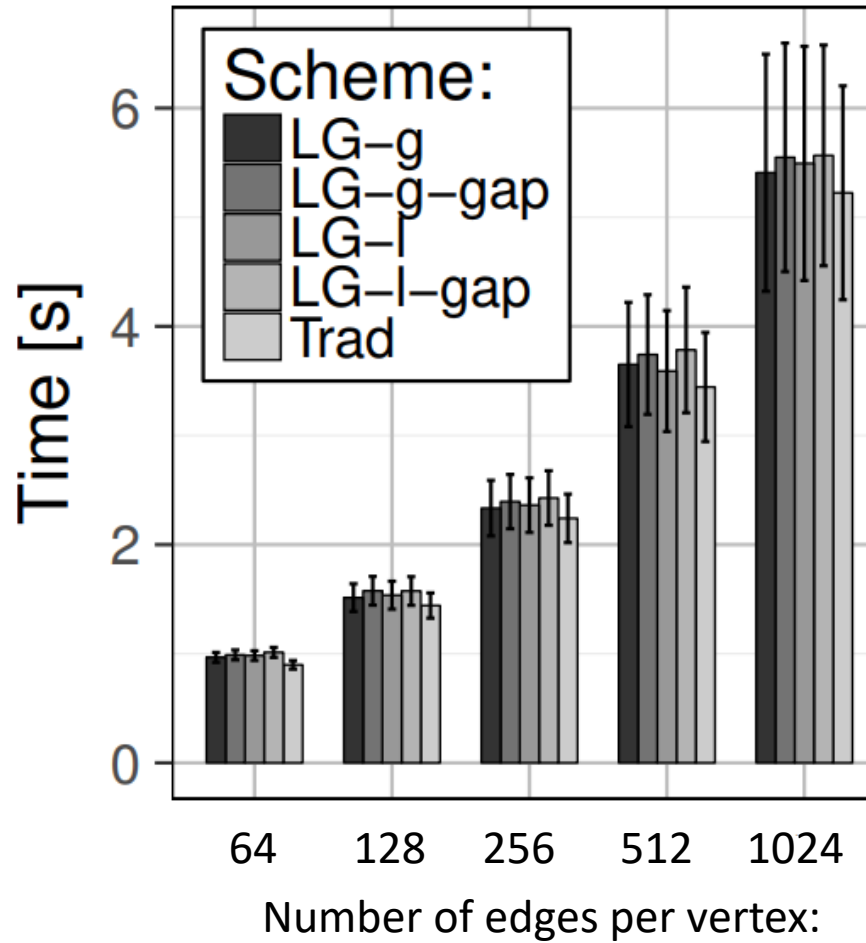
Betweenness Centrality

“**LG**”: Log(Graph)
Trad: Traditional (non compressed, GAPBS)
“**g**”: global scheme
“**l**”: local scheme
“**gap**”: additional gap encoding

Sparse graphs



Dense graphs



Kronecker graphs
Number of vertices: 4M

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

Betweenness Centrality

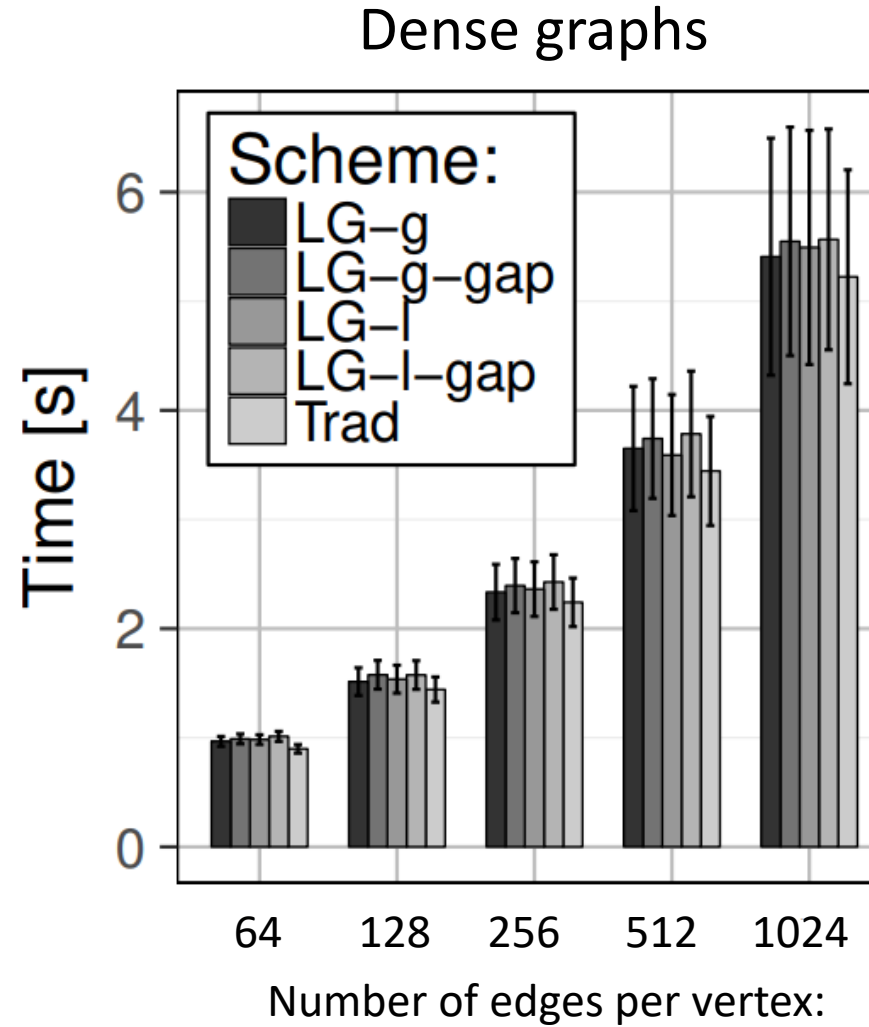

“**LG**”: Log(Graph)

Trad: Traditional (non compressed, GAPBS)


“**g**”: global scheme

“**l**”: local scheme

“**gap**”: additional gap encoding




Log(Graph) incurs negligible overheads



Kronecker graphs
 Number of vertices: 4M

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

BFS

“**LG**”: Log(Graph)Trad: Traditional
(non compressed,
GAPBS)“**g**”: global scheme“**l**”: local scheme“**gap**”: additional
gap encoding

Kronecker graphs



Number of vertices: 4M

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


BFS

Sparse graphs

Dense graphs



“LG”: Log(Graph)
Trad: Traditional (non compressed, GAPBS)
“g”: global scheme
“l”: local scheme
“gap”: additional gap encoding





Kronecker graphs
 Number of vertices: 4M

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

BFS


Sparse graphs

Dense graphs



Both storage and performance are improved simultaneously

“LG”: Log(Graph)
Trad: Traditional (non compressed, GAPBS)
“g”: global scheme
“l”: local scheme



Kronecker graphs
Number of vertices: 4M

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


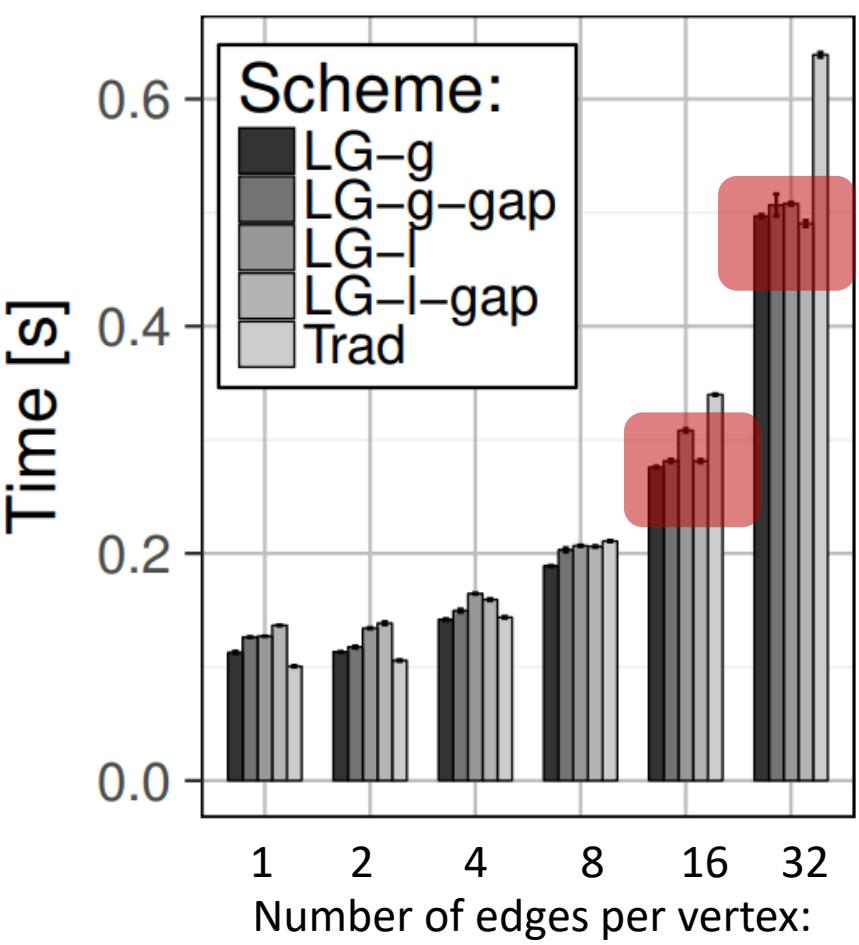
Log(Graph) accelerates GAPBS

BFS


“**LG**”: Log(Graph)
Trad: Traditional (non compressed, GAPBS)
 “**g**”: global scheme
 “**l**”: local scheme

Sparse graphs

Dense graphs




Both storage and performance are improved simultaneously




Kronecker graphs
 Number of vertices: 4M

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



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




The amount of communicated data is consistently reduced by ~37%

3 Log (Adjacency structure) Storage


3 **Log** (Adjacency structure) **Storage**

Trad: Traditional adjacency array
DMd / DMf: Degree Minimizing (without / with gap encoding)
WG: WebGraph compression
BRB, RB: Schemes targeting certain specific classes of graphs




3 **Log** (Adjacency structure) **Storage**

Trad: Traditional adjacency array
DMd / DMf: Degree Minimizing (without / with gap encoding)
WG: WebGraph compression
BRB, RB: Schemes targeting certain specific classes of graphs



3 **Log** (Adjacency structure) **Storage**

Trad: Traditional adjacency array
DMd / DMf: Degree Minimizing (without / with gap encoding)
WG: WebGraph compression
BRB, RB: Schemes targeting certain specific classes of graphs




Lots of data 😊

Conclusions:

3 **Log** (Adjacency structure) **Storage**

Trad: Traditional adjacency array
DMd / DMf: Degree Minimizing (without / with gap encoding)
WG: WebGraph compression
BRB, RB: Schemes targeting certain specific classes of graphs




Lots of data 😊

Conclusions: **WebGraph best for web graphs 😊**

3 **Log** (Adjacency structure) **Storage**

Trad: Traditional adjacency array
DMd / DMf: Degree Minimizing (without / with gap encoding)
WG: WebGraph compression
BRB, RB: Schemes targeting certain specific classes of graphs




Lots of data 😊

Conclusions: **WebGraph best for web graphs 😊**

BRB, RB: various tradeoffs but very expensive preprocessing (details in the paper)

3 **Log** (Adjacency structure) **Storage**

Trad: Traditional adjacency array
DMd / DMf: Degree Minimizing (without / with gap encoding)
WG: WebGraph compression
BRB, RB: Schemes targeting certain specific classes of graphs



Lots of data 😊

Conclusions:

WebGraph best for web graphs 😊

DMd: much better than DMf, often comparable to others

BRB, RB: various tradeoffs but very expensive preprocessing (details in the paper)

3 Log (Adjacency structure) Performance


3 Log (Adjacency structure) Performance

Trad: Traditional adjacency array

DMd / DMf: Degree Minimizing (without / with gap encoding)


WG: WebGraph compression

RB: Scheme targeting certain specific classes of graphs



3 Log (Adjacency structure) Performance


Trad: Traditional adjacency array
DMd / DMf: Degree Minimizing (without / with gap encoding)
WG: WebGraph compression
RB: Scheme targeting certain specific classes of graphs



WebGraph is the slowest

3 Log (Adjacency structure) Performance

Trad: Traditional adjacency array
DMd / DMf: Degree Minimizing (without / with gap encoding)
WG: WebGraph compression
RB: Scheme targeting certain specific classes of graphs




WebGraph is the slowest

DM, RB: comparable

Log(Graph) full design...


Log(Graph) full design...




Log(Graph) full 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 $\text{Log}(2) = \text{Log}(0\dots010_2) = 010_2$

Remove leading bits (simple bit packing)

2.1 (§3.1) U

2.3 ...globally (§3.2.1)

2.7 (§3.5) Analyze

2.4 ...locally (§3.2.2)

2.10 (§3.8) Ensure

2.5 ...on DM systems



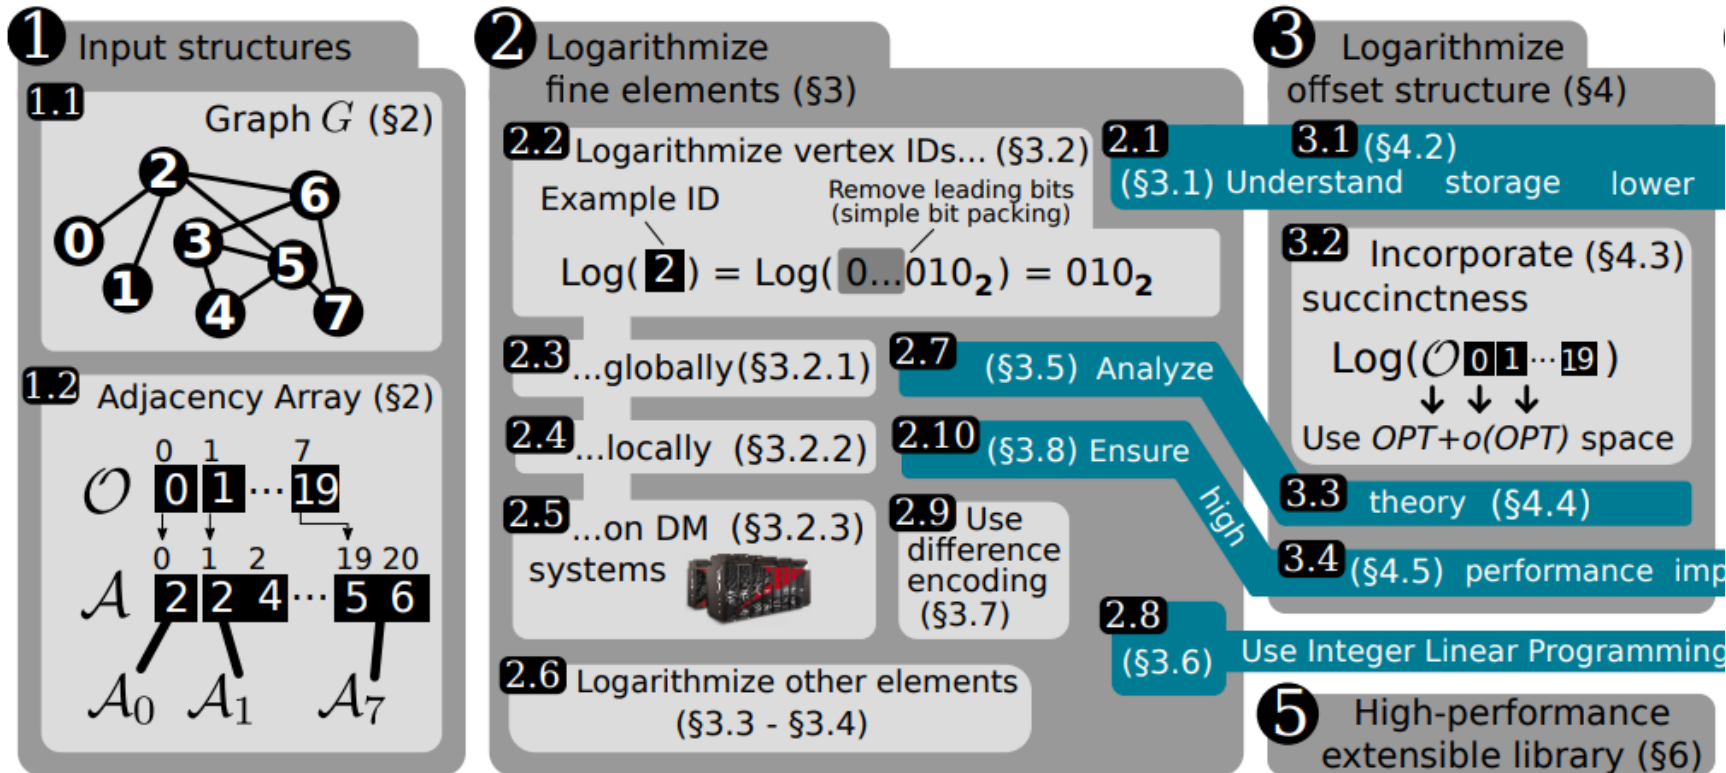
2.9 Use difference encoding (§3.7)

2.8 (§3.6) U

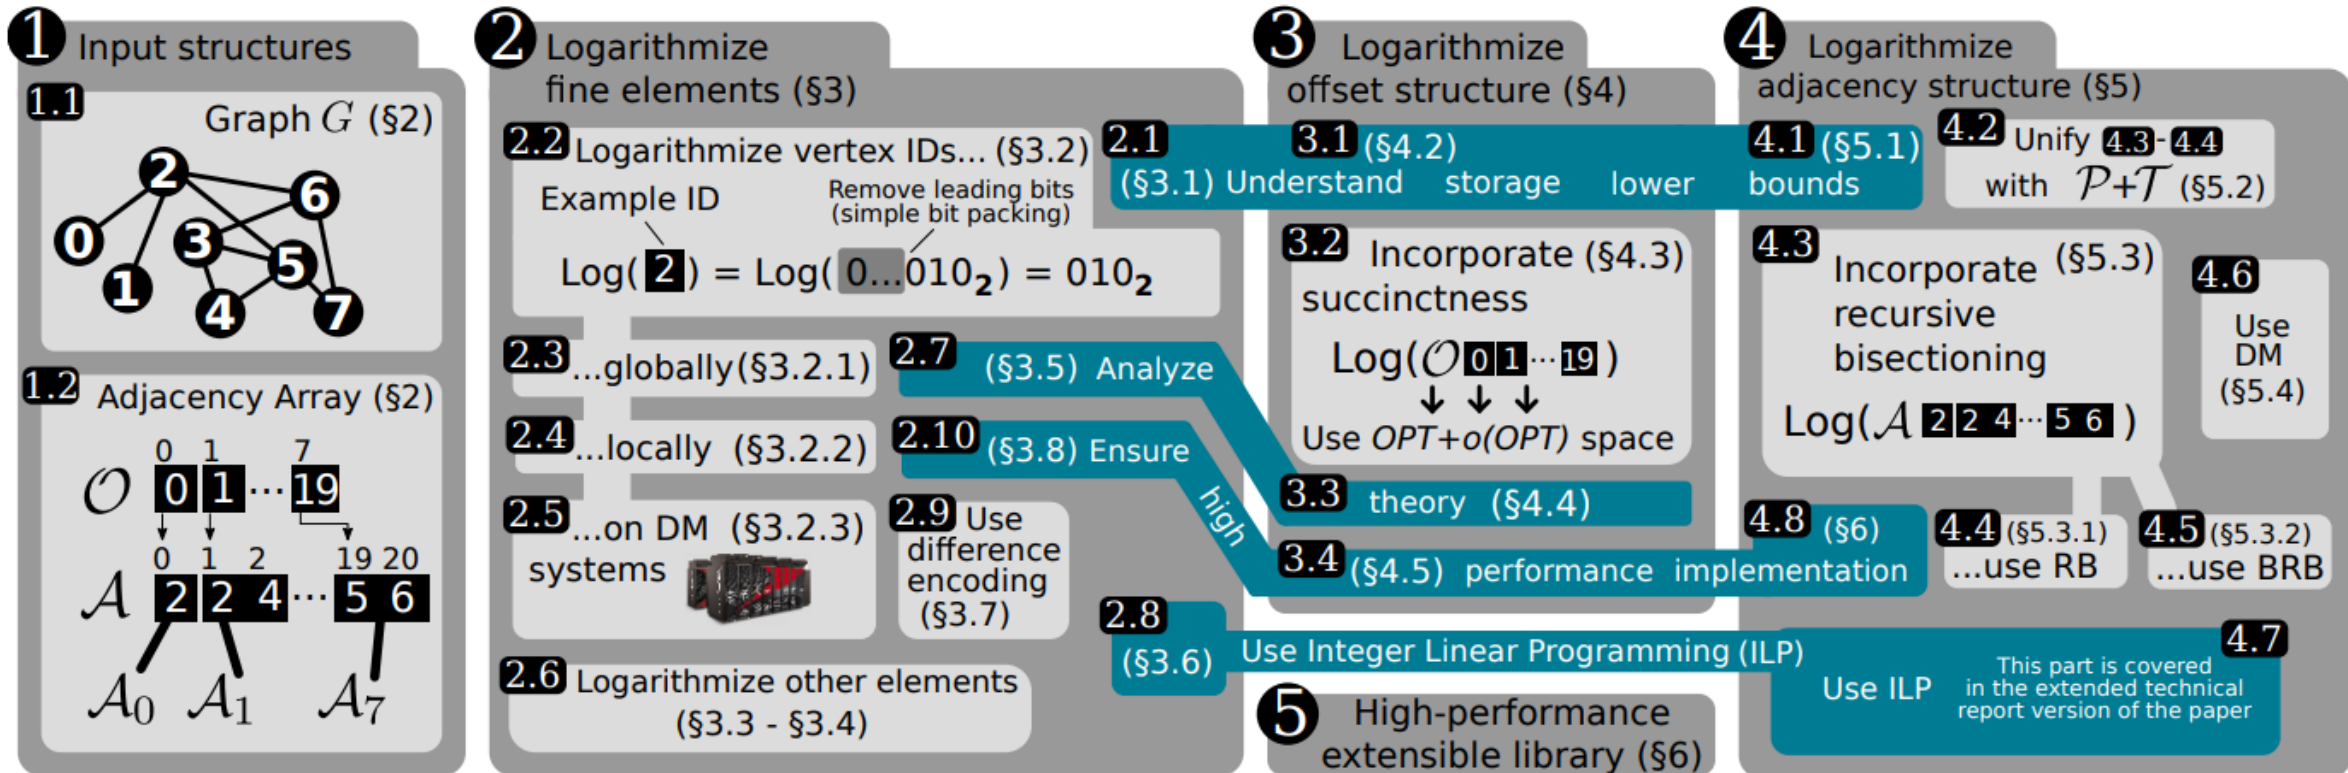
2.6 Logarithmize other elements (§3.3 - §3.4)

high

Log(Graph) full design...




Log(Graph) full design...




Log(Graph) full 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 $\text{Log}(2) = \text{Log}(0\dots010_2) = 010_2$

Remove leading bits (simple bit packing)

2.3 ...globally (§3.2.1)

2.4 ...locally (§3.2.2)

2.5 ...on DM systems



2.6 Logarithmize other elements (§3.3 - §3.4)

2.7 (§3.5) Analyze

2.10 (§3.8) Ensure

2.9 Use difference encoding (§3.7)

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 implementation

high

2.8 (§3.6) Use Integer Linear Programming (ILP)

5 High-performance extensible library (§6)

4 Logarithmize adjacency structure (§5)

4.1 (§5.1)

4.3 Incorporate (§5.3) recursive bisectioning

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

4.8 (§6)

4.4 (§5.3.1) ...use RB

4.5 (§5.3.2) ...use BRB

Use ILP

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

4.2 Unify 4.3-4.4 with $P+T$ (§5.2)

4.6 Use DM (§5.4)


4.5 (§5.3.2) ...use BRB

4.7


Log(Graph) full 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 $\text{Log}(2) = \text{Log}(0\dots010_2) = 010_2$

Remove leading bits (simple bit packing)

2.3 ...globally (§3.2.1)

2.4 ...locally (§3.2.2)

2.5 ...on DM (§3.2.3) systems



2.6 Logarithmize other elements (§3.3 - §3.4)

2.7 (§3.5) Analyze

2.10 (§3.8) Ensure

2.9 Use difference encoding (§3.7)

2.8 (§3.6)

3 Logarithmize offset structure (§4)

3.1 (§4.2) Understand storage lower

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

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

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

4.6 Use DM (§5.4)

4.8 (§6)

4.4 (§5.3.1) ...use RB

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

Log(Graph) full design...

Understand storage lower bounds and the theory

1 Input structures

1.1 Graph G (§2)

1.2 Adjacency Array (§2)

O 0 1 ... 7

0 1 2 ... 19 20

A 2 2 4 ... 5 6

A_0 A_1 ... A_7

2 Logarithmize fine elements (§3)

2.1 Logarithmize vertex IDs... (§3.2)

Example ID $\log(2) = \log(0\dots010_2) = 010_2$

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 Logarithmize other elements (§3.3 - §3.4)

2.7 (§3.5) Analyze

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

$\log(O 0 1 \dots 19)$

Use $OPT + o(OPT)$ space

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

$\log(A 2 2 4 \dots 5 6)$

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

Log(Graph) full design...

Understand storage lower bounds and the theory

1 Input structures

1.1 Graph G (§2)

1.2 Adjacency Array (§2)

2 Logarithmize fine elements (§3)

2.1 Understand storage lower bounds (§3.1)

2.2 Logarithmize vertex IDs... (§3.2)

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

2.3 ...globally (§3.2.1)

2.4 ...locally (§3.2.2)

2.5 ...on DM (§3.2.3) systems

2.6 Logarithmize other elements (§3.3 - §3.4)

2.7 (§3.5) Analyze

2.8 (§3.6) Use Integer Linear Programming (ILP)

2.9 Use difference encoding (§3.7)

2.10 (§3.8) Ensure high-performance implementation

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 implementation

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

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

4.4 (§5.3.1) ...use RB

4.5 (§5.3.2) ...use BRB

4.6 Use DM (§5.4)

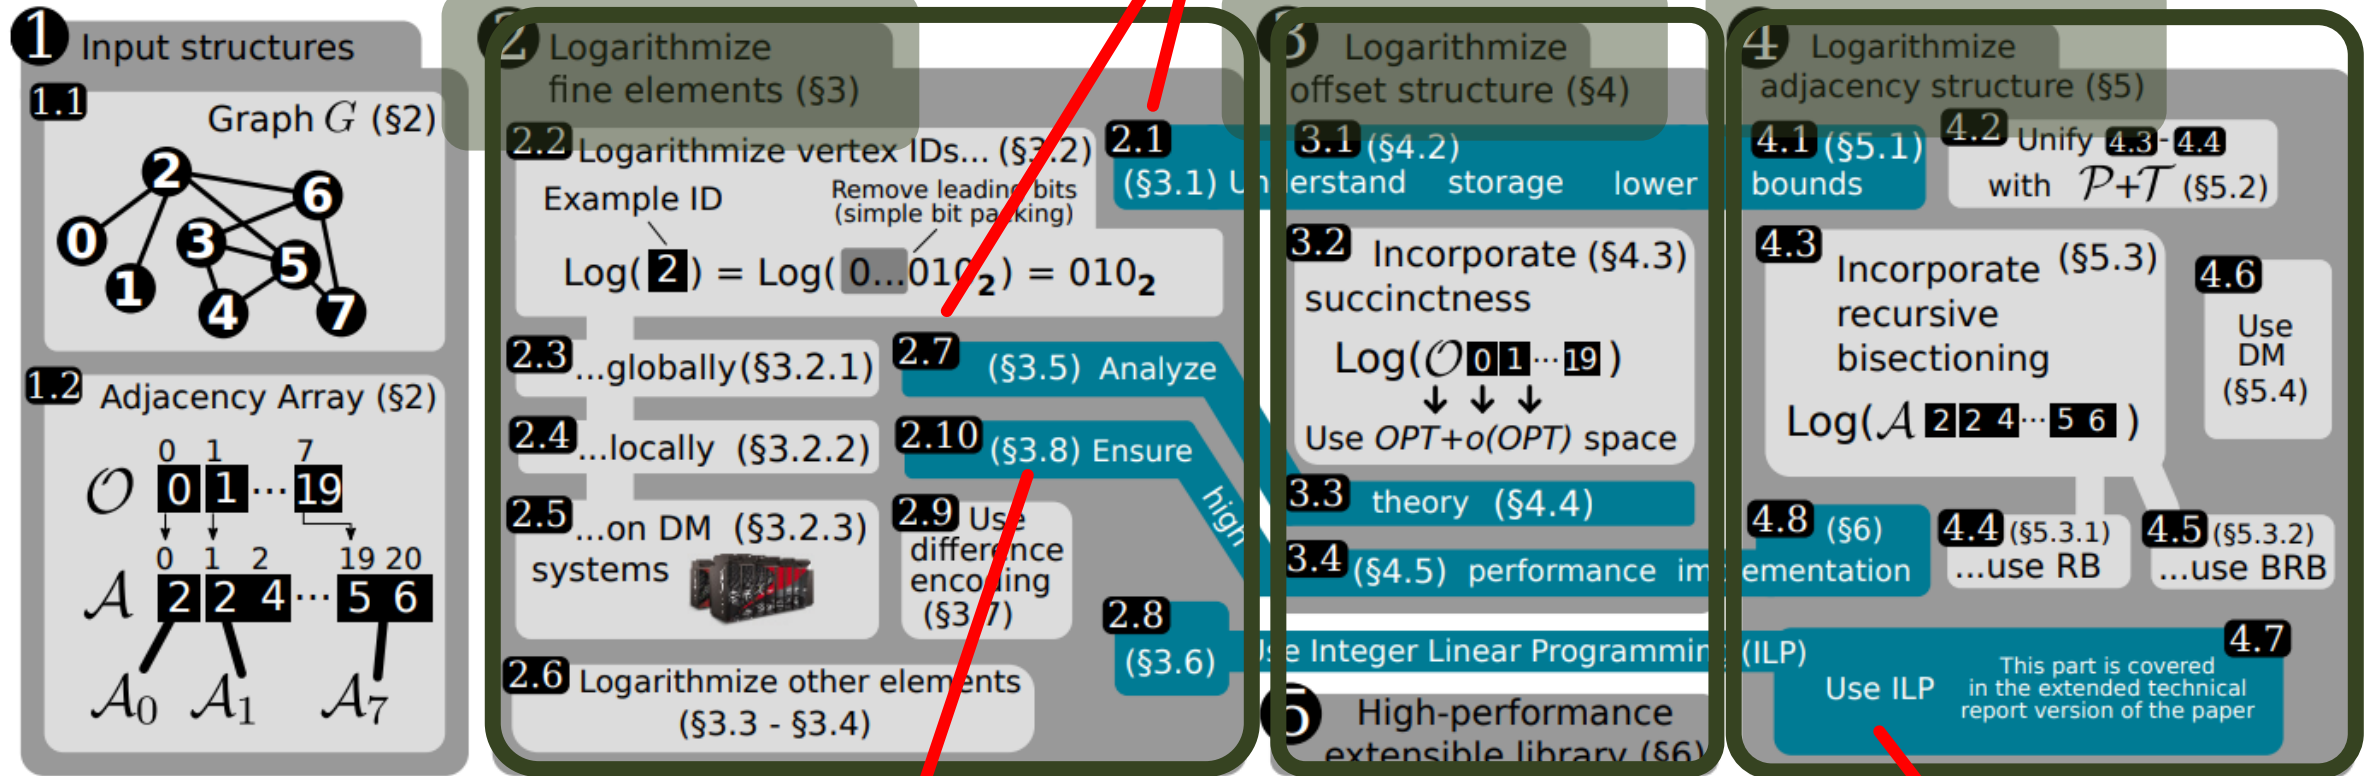
4.7 Use ILP

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

Ensure high-performance implementation

Log(Graph) full design...

Understand storage lower bounds and the theory



Ensure high-performance implementation

Use Integer Linear Programming (ILP) for more storage reductions

 Key method (vertex labels)

Key method (vertex labels)

Bit packing: use $\lceil \log n \rceil$ bits
for one vertex label

🔧 Key method (vertex labels)

Bit packing: use $\lceil \log n \rceil$ bits
for one vertex label

**Modern bitwise
operations**



⚡ Key method (vertex labels)

Bit packing: use $\lceil \log n \rceil$ bits for one vertex label

Modern bitwise operations



⚡ Key method (offsets)

🔧 Key method (vertex labels)

Bit packing: use $\lceil \log n \rceil$ bits for one vertex label

Modern bitwise operations



🔧 Key method (offsets)

Succinct bit vectors:

understand state-of-the-art designs and use the best ones in a given context

\mathcal{O}	ID	Asymptotic size [bits]	Exact size [bits]	select or $\mathcal{O}[v]$
Pointer array	ptrW	$O(Wn)$	$W(n+1)$	$O(1)$
Plain [44]	bvPL	$O\left(\frac{Wm}{B}\right)$	$\frac{2Wm}{B}$	$O(1)$
Interleaved [44]	bvIL	$O\left(\frac{Wm}{B} + \frac{Wm}{L}\right)$	$2Wm\left(\frac{1}{B} + \frac{64}{L}\right)$	$O\left(\log \frac{Wm}{B}\right)$
Entropy based [31, 78]	bvEN	$O\left(\frac{Wm}{B} \log \frac{Wm}{B}\right)$	$\approx \log\left(\frac{2Wm}{B}\right)$	$O\left(\log \frac{Wm}{B}\right)$
Sparse [76]	bvSD	$O\left(n + n \log \frac{Wm}{Bn}\right)$	$\approx n\left(2 + \log \frac{2Wm}{Bn}\right)$	$O(1)$
B-tree based [1]	bvBT	$O\left(\frac{Wm}{B}\right)$	$\approx 1.1 \cdot \frac{2Wm}{B}$	$O(\log n)$
Gap-compressed [1]	bvGC	$O\left(\frac{Wm}{B} \log \frac{Wm}{Bn}\right)$	$\approx 1.3 \cdot \frac{2Wm}{B} \log \frac{2Wm}{Bn}$	$O(\log n)$

🔧 Key method (vertex labels)

Bit packing: use $\lceil \log n \rceil$ bits for one vertex label

Modern bitwise operations



🔧 Key method (neighborhoods)

🔧 Key method (offsets)

Succinct bit vectors: understand state-of-the-art designs and use the best ones in a given context

\mathcal{O}	ID	Asymptotic size [bits]	Exact size [bits]	select or $\mathcal{O}[v]$
Pointer array	ptrW	$O(Wn)$	$W(n+1)$	$O(1)$
Plain [44]	bvPL	$O\left(\frac{Wm}{B}\right)$	$\frac{2Wm}{B}$	$O(1)$
Interleaved [44]	bvIL	$O\left(\frac{Wm}{B} + \frac{Wm}{L}\right)$	$2Wm\left(\frac{1}{B} + \frac{64}{L}\right)$	$O\left(\log \frac{Wm}{B}\right)$
Entropy based [31, 78]	bvEN	$O\left(\frac{Wm}{B} \log \frac{Wm}{B}\right)$	$\approx \log\left(\frac{2Wm}{B}\right)$	$O\left(\log \frac{Wm}{B}\right)$
Sparse [76]	bvSD	$O\left(n + n \log \frac{Wm}{Bn}\right)$	$\approx n\left(2 + \log \frac{2Wm}{Bn}\right)$	$O(1)$
B-tree based [1]	bvBT	$O\left(\frac{Wm}{B}\right)$	$\approx 1.1 \cdot \frac{2Wm}{B}$	$O(\log n)$
Gap-compressed [1]	bvGC	$O\left(\frac{Wm}{B} \log \frac{Wm}{Bn}\right)$	$\approx 1.3 \cdot \frac{2Wm}{B} \log \frac{2Wm}{Bn}$	$O(\log n)$

⚡ Key method (vertex labels)

Bit packing: use $\lceil \log n \rceil$ bits for one vertex label

Modern bitwise operations



⚡ Key method (neighborhoods)

Recursive partitioning: use representations that assume more about graph structure to enable better bounds

⚡ Key method (offsets)

Succinct bit vectors: understand state-of-the-art designs and use the best ones in a given context

\mathcal{O}	ID	Asymptotic size [bits]	Exact size [bits]	select or $\mathcal{O}[v]$
Pointer array	ptrW	$O(Wn)$	$W(n+1)$	$O(1)$
Plain [44]	bvPL	$O\left(\frac{Wm}{B}\right)$	$\frac{2Wm}{B}$	$O(1)$
Interleaved [44]	bvIL	$O\left(\frac{Wm}{B} + \frac{Wm}{L}\right)$	$2Wm\left(\frac{1}{B} + \frac{64}{L}\right)$	$O\left(\log \frac{Wm}{B}\right)$
Entropy based [31, 78]	bvEN	$O\left(\frac{Wm}{B} \log \frac{Wm}{B}\right)$	$\approx \log\left(\frac{2Wm}{Bn}\right)$	$O\left(\log \frac{Wm}{B}\right)$
Sparse [76]	bvSD	$O\left(n + n \log \frac{Wm}{Bn}\right)$	$\approx n\left(2 + \log \frac{2Wm}{Bn}\right)$	$O(1)$
B-tree based [1]	bvBT	$O\left(\frac{Wm}{B}\right)$	$\approx 1.1 \cdot \frac{2Wm}{B}$	$O(\log n)$
Gap-compressed [1]	bvGC	$O\left(\frac{Wm}{B} \log \frac{Wm}{Bn}\right)$	$\approx 1.3 \cdot \frac{2Wm}{B} \log \frac{2Wm}{Bn}$	$O(\log n)$

⚡ Key method (vertex labels)

Bit packing: use $\lceil \log n \rceil$ bits for one vertex label

Modern bitwise operations



⚡ Key method (neighborhoods)

Recursive partitioning: use representations that assume more about graph structure to enable better bounds

C++ templates to reduce overheads in performance-critical kernels

⚡ Key method (offsets)

Succinct bit vectors:

understand state-of-the-art designs and use the best ones in a given context

\mathcal{O}	ID	Asymptotic size [bits]	Exact size [bits]	select or $\mathcal{O}[v]$
Pointer array	ptrW	$O(Wn)$	$W(n+1)$	$O(1)$
Plain [44]	bvPL	$O\left(\frac{Wm}{B}\right)$	$\frac{2Wm}{B}$	$O(1)$
Interleaved [44]	bvIL	$O\left(\frac{Wm}{B} + \frac{Wm}{L}\right)$	$2Wm\left(\frac{1}{B} + \frac{64}{L}\right)$	$O\left(\log \frac{Wm}{B}\right)$
Entropy based [31, 78]	bvEN	$O\left(\frac{Wm}{B} \log \frac{Wm}{B}\right)$	$\approx \log\left(\frac{2Wm}{Bn}\right)$	$O\left(\log \frac{Wm}{B}\right)$
Sparse [76]	bvSD	$O\left(n + n \log \frac{Wm}{Bn}\right)$	$\approx n\left(2 + \log \frac{2Wm}{Bn}\right)$	$O(1)$
B-tree based [1]	bvBT	$O\left(\frac{Wm}{B}\right)$	$\approx 1.1 \cdot \frac{2Wm}{B}$	$O(\log n)$
Gap-compressed [1]	bvGC	$O\left(\frac{Wm}{B} \log \frac{Wm}{Bn}\right)$	$\approx 1.3 \cdot \frac{2Wm}{B} \log \frac{2Wm}{Bn}$	$O(\log n)$