## 7IHzürich

Maciej Besta, Dimitri Stanojevic, Tijana Zivic, Jagpreet Singh, Maurice Hoerold, Torsten Hoefler

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



ЄНzürich

## Large graphs...

## Large graphs...



## Large graphs...



## Large graphs...



## Large graphs...

## Running on...



## Large graphs...



Used in.


## $\left.\left.\frac{1}{1}\left(\frac{1}{2}\right)\right\rangle \left.+\frac{1}{2} \right\rvert\,=\frac{1}{2}\right)$



## Large graphs...

## Running on..



Used in...


## $\left.\frac{1}{1}\left|\frac{1}{2}\right| \lambda+\frac{1}{2} \right\rvert\,$



Large graphs...

Running on...


## Used in..



## $\left.\frac{1}{1}\left(\frac{1}{2}\right)+\frac{1}{2} \right\rvert\,$



ЄН_ürich

## Large graphs...



## $\tan \rightarrow \rho[L$

| en | vvinipeula euls（ell） | －ーローし |  |
| :---: | :---: | :---: | :---: |
| Tw | Twitter（WWW） | ，ロ－D | 41，652，230 1，468，365，182 |
| TF | Twitter（MPI） |  | 52，579，682 1，963，263，821 |
| FR | Friendster | ，ロ－D－L | 68，349，466 2，586，147，869 |
| uL | UK domain（2007） |  | 105，153，952 3，301，876，564 |



## tancral

| en | vviniveula Euls（ell） | －レローし |  |
| :---: | :---: | :---: | :---: |
| ${ }^{\text {TW }}$ | Twitter（WWW） | ，ロ－近哭 | 41，652，230 1，468，365，182 |
| TF | Twitter（MPI） |  | 52，579，682 1，963，263，821 |
| FR | Friendster | －D D－L | 68，349，466 2，586，147，869 |
| UL | UK domain（2007） |  | 105，153，952 3，301，876，564 |


| Graph500 Benchmark |  |  |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Top Ten from June 2018 BFS |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |  |
| 1 | кcomputer | Fujitu |  | Kobe Hyogo | Japan | 2011 | 82944 | 66352 | 40 | ${ }^{36621.4}$ |
| 2 | Sunway TaihuLight | NлCPC | $\begin{aligned} & \text { National } \\ & \text { Supercomputing } \\ & \text { Center in Wuxi } \end{aligned}$ | wuxi | China | 2015 | ${ }^{40768}$ | 10599880 | 40 | 23755.7 |
| 3 | DOE／NNSA／LLNL Sequoia | вм | Lawrence Livermore National Laboratory | Livemore CA | usa | 2012 | 93364 | 157284 | 41 | 23751 |
| 4 | DOE／SC／Argonne National Laboratory Mir | вмм | Argonne Nationa <br> Laboratory | chicagoll | usa | 2012 | 49152 | 786332 | 40 | 19882 |


| en | Vvikipeula euls（ely） | ツ－10ーレ |  |
| :---: | :---: | :---: | :---: |
| TW | Twitter（WWW） | ，$\square$ D－［i | 41，652，230 1，468，365，182 |
| TF | Twitter（MPI） | ，$\square$ D－［ | 52，579，682 1，963，263，821 |
| FR | Friendster | v $\square$ D－［ | 68，349，466 2，586，147，869 |
| UL | UK domain（2007） |  | 105，153，952 3，301，876，564 |


|  |  |  |  | $\square K O N E$ | CT gra | h da | ets |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Graph500 Benchmark |  |  |  |  |  |  |  |  |  |  |
|  | Webgraph datasets |  |  |  | $\left(\begin{array}{c} \text { GRAPH } \\ 500 \end{array}\right.$ |  |  |  |  |  |  |
| Graph | － | $\begin{array}{r} \text { Crawl date } \\ \hline 2014 \end{array}$ | $\begin{array}{\|c\|} \hline \text { Nodes } \\ \hline 787801471 \end{array}$ | Arcs |  |  |  |  |  |  |  |
| uk－2014 |  |  |  | 47614527250 | COUNTRY $\leqslant$ YEAR |  |  |  |  | scale＊gteps＊ |  |
| eu－2015 |  | 2015 | 1070557254 | 91792261600 |  |  |  |  |  |  |  |
| gsh－2015 |  | 2015 | 988490691 | 33877399152 |  |  |  |  |  |  |
| uk－2014－host |  | 2014 | 4769354 | 50829923 |  | Japan | 2011 |  |  | 66355 | 40 | 38621.4 |
| eu－2015－host |  | 2015 | 11264052 | 386915963 |  |  |  |  |  |  |  |
| gsh－2015－hos |  | 2015 | 68660142 | 1802747600 |  | China | 2015 | 40768 | 10599680 | 40 | 23755.7 |
| uk－2014－tpd |  | 2014 | 1766010 | 18244650 |  |  |  |  |  |  |  |
| eu－2015－tpd |  | 2015 | 6650532 | 170145510 |  |  |  |  |  |  |  |
| gsh－2015－tpd |  | 2015 | 30809122 | 602119716 | CA | USA | 2012 | 98304 | 1572864 | 41 | 23751 |
| cluewebl2 |  | 2012 | 978408098 | 42574107469 |  | USA | 2012 | 49152 | 786432 | 40 | 14982 |
| $\underline{\text { uk－2002 }}$ |  | 2002 | 18520486 | 298113762 |  |  |  |  |  |  |  |

## GHzürich

| $\begin{aligned} & \text { en } \\ & \text { TW } \\ & \text { TF } \\ & \text { FR } \\ & \text { UL } \end{aligned}$ | Twitter (WWW) <br> Twitter (MPI) <br> Friendster <br> UK domain (2007) |  |  | Web data commons datasets |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  |  | Granularity | \#Nodes | \#Arcs |
|  |  |  |  |  |  |  |
|  |  |  |  | 3,563 million | 128,736 million |  |
|  |  |  | NECT graph dataset |  | Host | 101 million | 2,043 million |
|  | Graph500 Ben |  |  | Pay-Level-Domain | 43 million | 623 million |


\＃｜Hzürich

| en | vvinkeula euls（ell） | vープーし |  |
| :---: | :---: | :---: | :---: |
| TW | Twitter（WWW） |  | 41，652，230 1，468，365，182 |
| TF | Twitter（MPI） | －D D－L | 52，579，682 1，963，263，821 |
| FR | Friendster | ，ロ－D－L | 68，349，466 2，586，147，869 |
| uL | UK domain（2007） |  | 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 |

Large graphs...

Running on...


## Used in..



## $\left.\frac{1}{1}\left(\frac{1}{2}\right)+\frac{1}{2} \right\rvert\,$



Running on...


## Used in..



## $\left.\frac{1}{1}\left(\frac{1}{2}\right)\right\rangle \left.+\frac{1}{2} \right\rvert\,$




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 $\boldsymbol{\Delta}$

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

## The storage lower bound

Which one?

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

The storage lower bound $\sqrt{ }$

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?

$$
S=\left\{x_{1}, x_{2}, x_{3}, \ldots\right\} \begin{aligned}
& x_{1} \rightarrow 0 \ldots 01 \\
& x_{2} \rightarrow 0 \ldots 10 \\
& x_{3} \rightarrow 0 \ldots 11
\end{aligned}
$$

## The storage lower bound 1

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?

$$
S=\left\{x_{1}, x_{2}, x_{3}, \ldots\right\} \begin{aligned}
& x_{1} \rightarrow 0 \ldots 01 \\
& x_{2} \rightarrow 0 \ldots 10 \\
& x_{3} \rightarrow 0 \ldots 11
\end{aligned}
$$

## ${ }_{\bullet}^{*}$ Key Idea

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?

$$
S=\left\{x_{1}, x_{2}, x_{3}, \ldots\right\} \begin{aligned}
& x_{1} \rightarrow 0 \ldots 01 \\
& x_{2} \rightarrow 0 \ldots 10 \\
& x_{3} \rightarrow 0 \ldots 11
\end{aligned}
$$

The storage lower bound Which one? P

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


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

$$
S=\left\{x_{1}, x_{2}, x_{3}, \ldots\right\} \begin{aligned}
& x_{1} \rightarrow 0 \ldots 01 \\
& x_{2} \rightarrow 0 \ldots 10 \\
& x_{3} \rightarrow 0 \ldots 11
\end{aligned}
$$



Which one?


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

## "'K Key idea

Vertex labels

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

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

$$
S=\left\{x_{1}, x_{2}, x_{3}, \ldots\right\} \begin{aligned}
& x_{1} \rightarrow 0 \ldots 01 \\
& x_{2} \rightarrow 0 \ldots 10 \\
& x_{3} \rightarrow 0 \ldots 11
\end{aligned}
$$



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

Vertex labels

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


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

$$
S=\left\{x_{1}, x_{2}, x_{3}, \ldots\right\} \begin{aligned}
& x_{1} \rightarrow 0 \ldots 01 \\
& x_{2} \rightarrow 0 \ldots 10 \\
& x_{3} \rightarrow 0 \ldots 11
\end{aligned}
$$

The storage lower bound Which one? (3)

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

## *" Key idea

Vertex labels

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


Adjacency arrays (edges adjacent to each vertex)

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

$$
S=\left\{x_{1}, x_{2}, x_{3}, \ldots\right\} \begin{aligned}
& x_{1} \rightarrow 0 \ldots 01 \\
& x_{2} \rightarrow 0 \ldots 10 \\
& x_{3} \rightarrow 0 \ldots 11
\end{aligned}
$$

The storage lower bound Which one? (3)

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

## "'K Key idea

Vertex labels

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


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

$$
S=\left\{x_{1}, x_{2}, x_{3}, \ldots\right\} \begin{aligned}
& x_{1} \rightarrow 0 \ldots 01 \\
& x_{2} \rightarrow 0 \ldots 10 \\
& x_{3} \rightarrow 0 \ldots 11
\end{aligned}
$$

The storage lower bound Which one? (?)

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

## 当 Key idea

$\log \binom{$ Vertex }{ labels }

Encode different parts of a graph representation using (logarithmic) storage lower bounds
 (edges adjacent to each vertex)

Offsets (locations) of adj. arrays

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

$$
S=\left\{x_{1}, x_{2}, x_{3}, \ldots\right\} \begin{aligned}
& x_{1} \rightarrow 0 \ldots 01 \\
& x_{2} \rightarrow 0 \ldots 10 \\
& x_{3} \rightarrow 0 \ldots 11
\end{aligned}
$$

The storage lower bound Which one? P

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

## "'K Key idea

$\log \binom{$ Vertex }{ labels }

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

(4)

5
Offsets (locations) of adj. arrays

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

$$
S=\left\{x_{1}, x_{2}, x_{3}, \ldots\right\} \begin{aligned}
& x_{1} \rightarrow 0 \ldots 01 \\
& x_{2} \rightarrow 0 \ldots 10 \\
& x_{3} \rightarrow 0 \ldots 11
\end{aligned}
$$

The storage lower bound Which one? P

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

## "'K Key idea

$\log \binom{$ Vertex }{ labels }


Encode different parts of a graph representation using (logarithmic) storage lower bounds Adjacency arrays
$\log ($ (edges adjacent $)$ to each vertex)

Offsets (locations) of adj. arrays

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

$$
S=\left\{x_{1}, x_{2}, x_{3}, \ldots\right\} \begin{aligned}
& x_{1} \rightarrow 0 \ldots 01 \\
& x_{2} \rightarrow 0 \ldots 10 \\
& x_{3} \rightarrow 0 \ldots 11
\end{aligned}
$$

The storage lower bound Which one? P

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

## 

$\log \binom{$ Vertex }{ labels }

## ,



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


ज1Hzürich
Adjacency Array Graph Representation

Representation


## Adjacency Array Graph Representation

Representation

## 0

1
2
3
4
5


## Adjacency Array Graph Representation



GHzürich

## Adjacency Array Graph Representation



ヨ円zürich

## Adjacency Array Graph Representation

Representation
 （edges adjacent to each vertex）

Physical realization


Adjacency Array Graph Representation

Representation
 (edges adjacent to each vertex)

Physical realization
Adjacency arrays (one contiguous array)

| 1 | 2 | 0 | 3 | 0 | 3 | 1 | 2 | 4 | 3 | 5 | 4 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |

Adjacency Array Graph Representation


Physical realization


Offsets (another contiguous array)


Adjacency Array Graph Representation

| Representation |  |
| :---: | :---: |
| 1 | 2 |
| $\rightarrow 0$ | 3 |
| $\rightarrow 0$ | 3 |
| $\rightarrow 1$ | 2 |
| $\rightarrow 3$ | 5 |
|  |  |
| Adjacency arrays (edges adjacent to each vertex) |  |

Physical realization



Adjacency arrays (one contiguous array)


Offsets (another contiguous array)

Adjacency Array Graph Representation

| Representation |  |
| :---: | :---: |
| 1 | 2 |
| $\rightarrow 0$ | 3 |
| $\rightarrow 0$ | 3 |
| $\rightarrow 1$ | 2 |
| $\rightarrow 3$ | 5 |
|  |  |
| Adjacency arrays (edges adjacent to each vertex) |  |

Physical realization



Adjacency arrays (one contiguous array)


Offsets (another contiguous array)

Adjacency Array Graph Representation

| Representation |  |
| :---: | :---: |
| $0 \longrightarrow 1$ | 2 |
| $\rightarrow 0$ | 3 |
| $\rightarrow 0$ | 3 |
| $\rightarrow 1$ | 24 |
| $\rightarrow 3$ | 5 |
| $5$ 4 <br> Offsets |  |
| Adjacency arrays (edges adjacent to each vertex) |  |

Physical realization



Adjacency arrays (one contiguous array)


Offsets (another contiguous array)

Adjacency Array Graph Representation

| Representation |  |
| :---: | :---: |
| $0 \longrightarrow 1$ | 2 |
| $\rightarrow 0$ | 3 |
| $\rightarrow 0$ | 3 |
| $\rightarrow 1$ | 24 |
| $\rightarrow 3$ | 5 |
| $5$ 4 <br> Offsets |  |
| Adjacency arrays (edges adjacent to each vertex) |  |

Physical realization



Offsets (another contiguous array)
$1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }
(1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

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

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

## Symbols <br> $n$ : \#vertices, <br> $m$ : \#edges, <br> $d_{v}$ : degree of vertex $v$, <br> $N_{v}$ : neighbors (adj. array) of vertex $v$, <br> $\widehat{N_{v}}$ : maximum among $N_{v}$



캐zürich

## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

## Lower bounds (global)

Symbols
$n$ : \#vertices,
$m$ : \#edges,

$$
d_{v}: \text { degree of vertex } v,
$$

$$
N_{v}: \text { neighbors (adj. array) of }
$$

$\widehat{N_{v}}$ : maximum among $N_{v}$


דㅐzürich

## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

## Lower bounds (global)

$\lceil\log n\rceil$

## Symbols <br> $n$ : \#vertices, <br> $m$ : \#edges, <br> $d_{v}$ : degree of vertex $v$, <br> $N_{v}$ : neighbors (adj. array) of vertex $v$, <br> $\widehat{N_{v}}$ : maximum among $N_{v}$



캐zürich

## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

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

This is it?
Not really ©

## Symbols <br> $n$ : \#vertices, <br> $m$ : \#edges, <br> $d_{v}$ : degree of vertex $v$, <br> $N_{v}$ : neighbors (adj. array) of vertex $v$, <br> $\widehat{N_{v}}$ : maximum among $N_{v}$



## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ 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)

## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ 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:

## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ 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., $V=\left\{1, \ldots, 2^{22}\right\}$


## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ 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., $V=\left\{1, \ldots, 2^{22}\right\}$
- A vertex $v$ with few neighbors: $d_{v} \ll n$


## $1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ 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
$\widehat{N_{v}}$ : maximum among $N_{v}$

캐zürich


## Lower bounds (local)

Assume:

- a graph, e.g., $V=\left\{1, \ldots, 2^{22}\right\}$
- A vertex $v$ with few neighbors: $d_{v} \ll n$
- ... all these neighbors have small labels: $\widehat{N_{v}} \ll n$


## 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., $V=\left\{1, \ldots, 2^{22}\right\}$
- A vertex $v$ with few neighbors: $d_{v} \ll n$
- ...all these neighbors have small labels: $\widehat{N_{v}} \ll n$



## 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 $N_{v}$

эПHzürich


## Lower bounds (local)

Assume:

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

- a graph, e.g., $V=\left\{1, \ldots, 2^{22}\right\}$
- A vertex $v$ with few neighbors: $d_{v} \ll n$
$-\ldots$ all these neighbors have small labels: $\widehat{N_{v}} \ll n$


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 $N_{v}$

## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ 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 $N_{v}$


캐zürich

## Lower bounds (local)

Assume:
$\left\lceil\log 2^{22}\right\rceil=22$

- a graph, e.g., $V=\left\{1, \ldots, 2^{22}\right\}$
- A vertex $v$ with few neighbors: $d_{v} \ll n$
- ...all these neighbors have small labels: $\widehat{N_{v}} \ll n$



19 zeros!

## 1) $\log \binom{$ Vertex }{ labels }, $\log \left(\begin{array}{c}\left.\begin{array}{c}\text { Edge } \\ \text { weights }\end{array}\right)\end{array}\right.$

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 $N_{v}$


эПHzürich

Lower bounds (local)
Assume:



19 zeros!
Thus, use the local bound $\left\lceil\log \widehat{N_{v}}\right\rceil$

캐zürich

## 

## This is it? <br> 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., $V=\left\{1, \ldots, 2^{22}\right\}$
- A vertex $v$ with few neighbors: $d_{v} \ll n$
- ...all these neighbors have small labels: $\widehat{N_{v}} \ll n$



## 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): problem

What if:

- a graph, e.g., $V=\left\{1, \ldots, 2^{22}\right\}$
- A vertex $v$ with few neighbors: $d_{v} \ll n$
- ...all these neighbors have small labels: $\widehat{N_{v}} \ll n$



## 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): problem

What if:

- a graph, e.g., $V=\left\{1, \ldots, 2^{22}\right\}$
- 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:

This is it? (?)

## Symbols

$n$ : \#vertices,
$m$ : \#edges,
$d_{v}$ : degree of vertex $v$,
$N_{v}$ : neighbors (adj. array) of
 Not really (:)
$\widehat{N_{v}}$ : maximum among $N_{v}$

## Lower bounds (local): problem

What if:

- a graph, e.g., $V=\left\{1, \ldots, 2^{22}\right\}$
- 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:



## Symbols

$n$ : \#vertices,
$m$ : \#edges,
$d_{v}$ : degree of vertex $v$,
$N_{v}$ : neighbors (adj. array) of
 Not really ©

$$
\text { vertex } v \text {, }
$$

$\widehat{N_{v}}$ : maximum among $N_{v}$

Lower bounds (local): problem
What if:

- a graph, e.g., $V=\left\{1, \ldots, 2^{22}\right\}$
- 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:


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



## Symbols

## This is it?



## Not really ()

vertex $v$,
$\widehat{N_{v}}$ : maximum among $N_{v}$

$n$ : \#vertices,
$m$ : \#edges,
$d_{v}$ : degree of vertex $v$,
$N_{v}$ : neighbors (adj. array) of

Lower bounds (local): problem
What if:

- a graph, e.g., $V=\left\{1, \ldots, 2^{22}\right\}$
- 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:
$\left\lceil\log 2^{20}\right\rceil=20$


17 zeros!
(1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

## Symbols <br> $n$ : \#vertices, <br> $m$ : \#edges, <br> $d_{v}$ : degree of vertex $v$, <br> $N_{v}$ : neighbors (adj. array) of vertex $v$, <br> $\widehat{N_{v}}$ : maximum among $N_{v}$



캐zürich

## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }



[^0]

## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

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

Symbols
$n$ : \#vertices,
$m$ : \#edges,
$d_{v}$ : degree of vertex $v$,
$N_{v}$ : neighbors (adj. array) of
$\widehat{N_{v}}$ : maximum among $N_{v}$


## Lower bounds (local) enhanced with ILP

$1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

## Symbols

$n$ : \#vertices,
$m$ : \#edges,
$d_{v}$ : degree of vertex $v$,
$N_{v}$ : neighbors (adj. array) of
 Programming (ILP)! 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

## $1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ 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

## $1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ 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

## $1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

## 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

(simultaneously for all other neighborhoods)

## $1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

## 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

$m$ : \#edges,
$d_{v}$ : degree of vertex $v$,
$N_{v}$ : neighbors (adj. array) of vertex $v$,
$\widehat{N_{v}}$ : maximum among $N_{v}$


דㅐzürich

## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

## Programming (ILP)! <br> ...Use Integer Linear !

## Symbols

$n$ : \#vertices,

Lower bounds (local) enhanced with ILP

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


Heuristics:

$$
\min \sum_{v \in V} \widehat{N_{v}} \frac{1}{d_{v}}
$$

## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ 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 $N_{v}$

דㅐzürich

## Programming (ILP)! <br> ...Use Integer Linear!



Lower bounds (local) enhanced with ILP
Permute vertex labels to reduce such maximum labels in as many neighborhoods as possible


## $1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

## Symbols

$n$ : \#vertices,
$m$ : \#edges,
$d_{v}$ : degree of vertex $v$,
$N_{v}$ : neighbors (adj. array) of
 Programming (ILP)!
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


Intuition:
maximum
labels in new neighborhoods will be smaller

$1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }
Symbols
$\widehat{W}$ : max edge weight,
$n$ : \#vertices,
$p, \alpha, \beta$ : constants

$1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }
$\oiint$ Formal analyses

## Power-law graphs

## Symbols

$\widehat{W}$ : max edge weight,
$n$ : \#vertices,
$p, \alpha, \beta$ : constants


## Random uniform graphs

(1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }
$\oiiint$ Formal analyses

## Symbols

$\widehat{W}$ : max edge weight,
$n$ : \#vertices,
$p, \alpha, \beta$ : constants


## Power-law graphs

The probability that a vertex has degree $d$ is:
$\alpha d^{\beta}$

## Random uniform graphs

(1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }
$\oiiint$ Formal analyses

## Symbols

$\widehat{W}$ : max edge weight,
$n$ : \#vertices,
$p, \alpha, \beta$ : constants


## Power-law graphs

The probability that a vertex has degree $d$ is:
$\alpha d^{\beta}$


## Random uniform graphs

(1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }
$\oiiint$ Formal analyses

## Symbols

$\widehat{W}$ : max edge weight,
$n$ : \#vertices,
$p, \alpha, \beta$ : constants


## Power-law graphs

The probability that a vertex has degree $d$ is:
$\alpha d^{\beta}$


## Random uniform graphs

The probability that a vertex has degree $d$ is:

$$
p d
$$

## דㅐzürich

## (1) Log $\binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights } $\oiiint$ 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,
$p, \alpha, \beta$ : constants


## Random uniform graphs

The probability that a vertex has degree $d$ is:

$$
p d
$$



## $1 \log ($ Vertex $), \log ($ Edge $\left.{ }_{\text {labels }}\right), \log \left(\begin{array}{l}\text { weights }\end{array}\right)$ $\oiiint$ Formal analyses

## Power-law graphs

The probability that a vertex has degree $d$ is:


## Symbols

$\widehat{W}$ : max edge weight,
$n$ : \#vertices,
$p, \alpha, \beta$ : constants


## Random uniform graphs

The probability that a vertex has degree $d$ is:

$$
p d
$$

Expected size of the adjacency array

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

## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

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

$\oiint$ Formal analyses: more (check the paper ©)

$$
\forall_{v, u \in V}\left(u \in N_{v}\right) \Rightarrow\left[\mathcal{N}(u) \leq \widehat{N}_{v}\right]
$$

$$
\begin{array}{rr}
|\mathscr{A}|=\sum_{v \in V}\left(d_{v}\left\lceil\log \widehat{N}_{v}\right\rceil+\left\lceil\log \log \widehat{N}_{v}\right\rceil\right) & |\mathscr{A}|=\sum_{v \in V}\left(d_{v}\left\lceil\log \widehat{N}_{v}\right\rceil+\left\lceil\log \log \widehat{N}_{v}\right\rceil\right) \\
|\mathcal{A}|=n\left\lceil\log \frac{n}{\mathcal{H}}\right\rceil+\mathcal{H}\lceil\log \mathcal{H}\rceil & |\mathcal{A}|=2 m(\lceil\log n\rceil+\lceil\log \widehat{\mathcal{W}}\rceil) \\
|\mathcal{A}|=\sum_{v \in V}\left(d_{v}\left(\left\lceil\log \widehat{N}_{v}\right\rceil+\lceil\log \widehat{\mathcal{W}}\rceil\right)+\left\lceil\log \log \widehat{N}_{v}\right\rceil+\lceil\log \log \widehat{\mathcal{W}}\rceil\right)
\end{array}
$$

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

## (1) $\log \binom{$ Vertex $)}{$ labels }, $\log \binom{$ Edge }{ weights }

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

$\oiint$ Formal analyses: more (check the paper ©)

$$
\begin{aligned}
& |\mathscr{A}|=\sum_{v \in V}\left(d_{v}\left\lceil\log \widehat{N}_{v}\right\rceil+\left\lceil\log \log \widehat{N}_{v}\right\rceil\right) \quad|\mathscr{A}|=\sum_{v \in V} \\
& |\mathcal{A}|=n\left\lceil\log \frac{n}{\mathcal{H}}\right\rceil+\mathcal{H}\lceil\log \mathcal{H}\rceil \\
& |\mathcal{A}|=\sum_{v \in V}\left(d_{v}\left(\left\lceil\log \widehat{N}_{v}\right\rceil+\right\rceil\right.
\end{aligned}
$$

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


$1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }
Key methods


ヨ|Hzürich
(1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

* Key methods

(1) $\log \left(\begin{array}{l}\text { labelts }\end{array}\right), \log \binom{$ (edge }{ weights }
\% 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 N Ni,v
3 int64_t exactBitOffset = s * (\mathcal{O}[v] + i);
4 int8_t* address = (int8_t*) \mathcal{A + (exactBitOffset >> 3);}
5 int64_t distance = exactBitOffset & 7;
6 int64_t value = ((int64_t*) (address))[0];
7 return _bextr_u64(value, distance, s); }
```

$1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }
\% 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 tove for IDs of vertices. */
```



```
3 int64_t exactBitOffset = s * (\mathcal{O}[v] + i);
4 int8_t* address = (int8_t*) \mathcal{A + (exactBitOffset >> 3);}
5 int64_t distance = exactBitOffset & 7;
6 int64_t value = ((int64_t*) (address))[0];
7 return _bextr_u64(value, distance, s); }
```

$1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }
K Key methods

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



Pointer to the offset array

Return $i$-th neighbor of vertex $v$

$1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }
K 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

$1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }
K 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 tone for IDs of ertices. */
2 v_ID $N_{i, v}\left(v_{-} I D v, i n t 32 \_t i, \operatorname{int64\_ t*} \mathcal{O}\right.$, int64_t* $\mathcal{A}$, int8_t s)\{
3 int64_t exactBitOffset $=s *(\mathcal{O}[v]+i)$;
4 int8_t* address $=\left(i n t 8 \_t *\right) \mathcal{A}+(e x a c t B i t 0 f f s e t \gg 3)$;
5 int64_t distance = exactBitOffset \& 7;
6 int64_t value = ((int64_t*) (address)) [0];
7 return _bextr_u64 (value, distance, $s$ ) ; \}
$1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }
K 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 $s=\lceil\log n\rceil$

v_ID $N_{i, v}\left(v_{-} I D v, i n t 32 \_t i, i n i 64 \_t * \mathcal{O}\right.$, int 64_t* $\mathcal{A}$, int 8_t $\left.s\right)\{$
3 int64_t exactBitOffset $=s *(\mathcal{O}[v]+i)$;
4 int 8_t* address $=$ (int8_t*) $\mathcal{A}+$ (exactBitOffset >> 3 );
5 int64_t distance = exactBitOffset \& 7;
6 int 64_t value = ((int64_t*) (address)) [0];
7 return _bextr_u64 (value, distance, $s$ ) ; \}
(1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

K 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

Get the closest byte alignment int 64_t exactBitOffset $=s *(\mathcal{O}[v]+i)$; int 8_t* address $=\left(i n t 8 \_t *\right) \mathcal{A}+(e x a c t B i t O f f s e t \gg 3)$; int 64_t distance = exactBitOffset \& 7; int64_t value = ((int64_t*) (address)) [0]; return _bextr_u64 (value, distance, s) ; \}
(1) Log $\binom{$ Vertex }{ labels }, $\log \binom{$ 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 $s=\lceil\log n\rceil$

Get the closest byte alignment int64_t exactBitOffset $=s *(\mathcal{O}[v]+i)$; int8_t* address $=\left(i n t 8 \_t *\right)(\mathcal{A}+(e x a c t B i t O f f s e t \gg 3)$; int64_t distance $=$ exactBitoffset \& 7; int64_t value = ((int64_t*) (address)) [0]; return _bextr_u64 (value, distance, s) ; \} Get the distance from the byte alignment
$1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ 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

Get the closest byte alignment int64_t exactBitOffset $=s *(\mathcal{O}[v]+i)$; int8_t* address $=\left(i n t 8 \_t *\right)(\mathcal{A}+(e x a c t B i t O f f s e t \gg 3)$; int64_t distance $=$ exactBitoffset \& 7;

6 int64_t value $=(($ int64_t*) (address)) [0];
7 return _bextr_u64 (value, distance, $s$ ) ; \}

Get the distance from the byte alignment

Access the derived 64 -bit value

## $1 \log \binom{$ Vertex }{ labels }, $\log ($ Edge labels $)$ LOg (weights <br> K 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

Get the closest byte alignment
4 int8_t* address $=\left(i n t 8 \_t *\right)(\mathcal{A}+$ (exactBitOffset $\gg 3)$;
5 int64_t distance $=$ exactBitoffset \& 7;
6 int64_t value $=(($ int64_t*) $($ address) $)[0]$;
return _bextr_u64(value, distance, s); \}

Get the distance from the byte alignment

Access the derived 64-bit value
(2) $\log$ ( offset structure)

EHzürich
(2) $\log$ ( offset structure)


## (2) LO\% (Offset structure)

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


## ज्ञzürich

## (2) Log ( offset structure)

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



Bit vectors instead of offset arrays
(2) Log ( offset structure)

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


Bit vectors instead of offset arrays


## (2) Log (offset structure)

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


Bit vectors instead of offset arrays


## (2) Log ( offset structure)

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


Bit vectors instead of offset arrays


## (2) Log ( offset structure)

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


Bit vectors instead of offset arrays

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

## (2) Log ( offset structure)

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


Bit vectors instead of offset arrays


How many 1s are set before a given i-th bit?
$i$-th set bit has a position $x \rightarrow$ 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]


## 7Hzürich

## (2) $\log$ ( offset structure)

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

(2) Log ( Offset structure)

## Succinct bit vectors

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

(2) Log ( Offset structure)

## Succinct bit vectors

They use $[\mathrm{Q}]+o(\mathrm{Q})$ bits ( $[\mathrm{Q}]$ - lower bound), they answer various queries in $o(Q)$ time. = small + fast (hopefully) !
(2) Log ( Offset structure)

## Succinct bit vectors

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

## = small + fast

 (hopefully)!(2) Log ( Offset structure)

## Succinct bit vectors

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

## Succinct bit vectors

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

$n$ bits $101010100101000101010111110000001100001 \ldots$
(2) Log ( Offset structure)

## Succinct bit vectors

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

(2) Log ( offset structure)

## Succinct bit vectors

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


## (2) Log ( offset structure)

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

## Succinct bit vectors

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


## (2) Log ( offset structure)

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

## Succinct bit vectors <br> They use $[\mathrm{Q}]+o(\mathrm{Q})$ bits ( $[\mathrm{Q}]$ - lower bound), <br> = small + fast 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 <br> They use [Q] +o(Q) bits ([Q] - lower bound), (hopefully)

# 1 


(2) Log ( offset structure)
...Encode the resulting bit vectors as succinct bit vectors [1]

## Compute \& store <br> the number of 1 s <br> Succinct bit vectors

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


## (2) Log ( offset structure)

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

## Compute \& store <br> the number of 1 s <br> Succinct bit vectors

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


## 2 Log ( Offset structure)

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



## $2 \log$ ( offset structure)

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

$\begin{gathered}\text { Compute \& store } \\ \text { the number of 1s }\end{gathered}=O\left(\frac{n}{t_{1}} \log n\right)=O\left(\frac{n}{\log n}\right)=o(n)$


## $2 \log$ ( offset structure)

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


## 2 Log ( Offset structure)

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

## Succinct bit vectors

## Total storage:

$n+o(n)+o(n)+\cdots$
They use [Q] +o(Q) bits ([Q] - lower bound), = small + fast they answer various queries in $o(Q)$ time.
$\begin{aligned} & \text { Compute \& store } \\ & \text { the number of } 1 \mathrm{~s}\end{aligned}=O\left(\frac{n}{t_{1}} \log n\right)=O\left(\frac{n}{\log n}\right)=O(n)$


## (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
$\oiint$ Formal analyses


## (2) Log ( offset structure)

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

## 丹 Formal analyses

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

## (2) Log ( offset structure)

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

## $\oiint$ Formal analyses

Check the paper for details ©


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

## (2) Log ( offset structure)

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

## Æ Formal analyses

Check the paper for details ;)


| $\mathcal{O}$ | ID | Asymptotic size [bits] | Exact size [bits] |  | select or $\mathcal{O}[v]$ |
| :---: | :---: | :---: | :---: | :---: | :---: |
| Pointer array | ptr | $O(W n)$ | $W(n+1)$ |  | $O(1)$ |
| Plain [44] <br> Interleaved [44] |  |  |  |  |  |
|  | Me wi show g $\frac{W m}{B}$ ) |  |  |  |  |
| Entropy based [3 | $\text { in pactice ooth smat and fast } \mathrm{g} \frac{W m}{B} \text { ) }$ |  |  |  |  |
| Sparse [76] |  |  |  |  |  |
| B-tree based [1] | bvbr $\left(\frac{1}{B}\right) \sim 1.1 \cdot \frac{B}{B}(\log n)$ |  |  |  |  |
| Gap-compressed [1] |  | $O\left(\frac{W m}{B} \log \frac{W m}{B n}\right)$ | $\approx 1.3 \cdot \frac{2 W m}{B} \log$ | $\frac{2 W m}{B n}$ | $O(\log n)$ |

АПНürich
(3) $\log \binom{$ Adjacency }{ structure }


F|Hzürich

## (3) $\log \binom{$ Adjacency }{ structure }

Use different relabelings


## (3) $\log \binom{$ Adjacency }{ structure }

Use different relabelings

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


## (3) $\log \binom{$ Adjacency }{ structure }

Use different relabelings

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

ㅋHzürich


More schemes that assume specific classes of graphs

## (3) $\log \binom{$ Adjacency }{ structure }

Use different relabelings

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

(simultaneously for all
other neighborhoods)

More schemes that assume specific classes of graphs

## 3 log Adjacency structure <br> Degree-Minimizing: Targeting general graphs (no assumptions on graph structure)

Use different relabelings

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

Use different relabelings


## More schemes

 that assume specific classes of graphs
## 3 Log Adjacency structure <br> (no assumptions on graph structure)


(simultaneously for all other neighborhoods)


More schemes that assume specific classes of graphs

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

## An $\rightarrow$ FPLL

## Overview of Full Log(Graph) Design

## Overview of Full Log(Graph) Design



캐zürich

## Overview of Full Log(Graph) Design



## Overview of Full Log(Graph) Design



## Overview of Full Log(Graph) Design



## Overview of Full Log(Graph) Design



## Overview of Full Log(Graph) Design



## Overview of Full Log(Graph) Design



## Overview of Full Log(Graph) Design



## Overview of Full Log(Graph) Design

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

-.. they all can be arbitrarily combined.

We analyzed / implemented (in total):
6 schemes for compressing fine elements, 10+ schemes for compressing offset structures, 4+ schemes for compressing adjacency structures

## Overview of Full Log(Graph) Design

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

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

## 2.1

- 

 ..... 10 ally -3. I

We analyzed / implemented (in total):
6 schemes for compressing fine elements, 10+ schemes for compressing offset structures, 4+ schemes for compressing adjacency structures
spcl.inf.ethz.ch

## Performance Analysis

## TYpes of machines



CSCS Cray Piz Daint



CSCS Cray Piz Daint

## 2 Performance Analysis

Types of machines

## IIIFIT


spcl.inf.ethz.ch @spcl_eth

द्मzürich


CSCS Cray Piz Daint

## Performance Analysis

Types of graphs

## Performance Analysis

Types of graphs
Synthetic graphs
ЄНzürich

## Performance Analysis

TYpes of Graphs
Synthetic graphs


## Performance Analysis

Types of graphs


## Performance Analysis

Types of graphs


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

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

[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

## TYpes of graphs



## 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 BFS (direction<br>Components optimization [2])

(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 <br> 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
BFS (direction
optimization [2])
(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
BFS (direction
optimization [2])
(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.


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

## Performance Analysis

## Algorithms

## Connected

Components
BFS (direction
optimization [2]) (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.
[3] U. Meyer, P. Sanders. Delta-Stepping: A Parallelizable Shortest Path Algorithm. 2003.

## Performance Analysis

## Connected

Components (Shiloach-Vishkin [1])


## Triangle Counting



[1] Y. Shiloach, U. Vishkin. An O (log n) parallel connectivity algorithm. 1980.
[3] U. Meyer, P. Sanders. Delta-Stepping: A Parallelizable Shortest Path Algorithm. 2003.
[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.


[1] Y. Shiloach, U. Vishkin. An O (log n) parallel connectivity algorithm. 1980.
[3] U. Meyer, P. Sanders. Delta-Stepping: A Parallelizable Shortest Path Algorithm. 2003.


[1] Y. Shiloach, U. Vishkin. An O (log n) parallel connectivity algorithm. 1980.
[3] U. Meyer, P. Sanders. Delta-Stepping: A Parallelizable Shortest Path Algorithm. 2003.
[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.


[1] Y. Shiloach, U. Vishkin. An O (log n) parallel connectivity algorithm. 1980.
[3] U. Meyer, P. Sanders. Delta-Stepping: A Parallelizable Shortest Path Algorithm. 2003.
[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.


[1] Y. Shiloach, U. Vishkin. An O (log n) parallel connectivity algorithm. 1980.
[3] U. Meyer, P. Sanders. Delta-Stepping: A Parallelizable Shortest Path Algorithm. 2003.
[2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.

[3] U. Meyer, P. Sanders. Delta-Stepping: A Parallelizable Shortest Path Algorithm. 2003. [2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.

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.


[3] U. Meyer, P. Sanders. Delta-Stepping: A Parallelizable Shortest Path Algorithm. 2003. [2] S Beamer et al. Direction-Optimizing Breadth-First Search. 2013.





## Performance Analysis

## COMPARISON TARGETS

## Performance Analysis

## Comparison Targets

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

## Performance Analysis

## Comparison Targets



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

## Performance Analysis

## Comparison Targets

## Zlib [2].

Comparison to a traditional

> GAPBS: Graph Algorithm Platform Benchmark Suite [1].

Comparison to a traditional adjacency array implementation compression scheme

## 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

## Recursive Partitioning [4].

Comparison to a tuned scheme for compressing adjacency data
(1) Log $\binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights } Storage, Performance
SSSP


Kronecker graphs Number of vertices: 4M

캐zürich

## $1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights } Storage, Performance

SSSP


Kronecker graphs Number of vertices: 4M

Number of edges per vertex

## $1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights } Storage, Performance

SSSP


Number of edges per vertex

Kronecker graphs Number of vertices: 4M

Log(Graph) consistently reduces storage overhead

## $1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights } Storage, Performance



Number of edges per vertex

## Log(Graph) accelerates GAPBS

SSSP


Kronecker graphs Number of vertices: 4M

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

## $1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights } Storage, Performance



Number of edges per vertex

## Log(Graph) accelerates GAPBS

SSSP

Kronecker graphs Number of vertices: 4M

Both storage and performance are improved simultaneously

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

## (2) Log (offset structure) Storage



Offsets:


## (2) Log (offset structure) Storage



Offsets:

## Lots of data : $^{\text {) }}$

## Conclusions:

## (2) -of Stfset 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) <br> 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) <br> 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 ${ }^{\text {© }}$ Conclusions:

Succinct bit vectors consistently ensure best storage reductions

## 2 Log (Offset structure) Storage



Lots of data $)^{-}$ Conclusions:

Succinct bit vectors consistently ensure best storage reductions
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


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

Accessing randomly selected neighbors


## (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

## Lots of data again ;) Conclusions:



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

## Lots of data again © Conclusions:

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

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

## 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

## 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

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
bvSD: the fastest and (usually) the smallest


Kronecker graphs Number of vertices: 4M

## NAREL

## (3) Log ( $\left.\begin{array}{c}\text { Adjacency } \\ \text { structure }\end{array}\right) \begin{gathered}\text { Storage, }\end{gathered}$ 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 ( $\left.\begin{array}{c}\text { Adjacency } \\ \text { structure }\end{array}\right)$ 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 $\binom{$ Adjacency }{ structure } $\begin{gathered}\text { Storage, }\end{gathered}$ 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


## Conclusions:

## (3) $\log \binom{$ Adjacency }{ structure } $\begin{gathered}\text { Storage, }\end{gathered}$ 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


## Conclusions:

WebGraph
best for web
graphs ()

## (3) Log $\binom{$ Adjacency }{ structure } $\begin{gathered}\text { Storage, } \\ \text { Performane }\end{gathered}$

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


## Conclusions:

WebGraph
best for web
graphs ()

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

## (3) $\log \binom{$ Adjacency }{ structure } $\begin{gathered}\text { Storage, }\end{gathered}$ 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


## 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 ( $\left.\begin{array}{c}\text { Adjacency } \\ \text { structure }\end{array}\right) \begin{gathered}\text { Storage, }\end{gathered}$ 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 ${ }^{\text {: }}$

## 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) $\log \binom{$ Adjacency }{ structure } $\begin{gathered}\text { Storage, } \\ \text { Performan }\end{gathered}$ 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)

WebGraph is the slowest, DM somewhat slower than Trad


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

Key insight (vertex labels)
$\mathbf{2 0 - 3 5 \%}$ 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 $\mathbf{> 9 0 \%}$ storage reductions (compared to uncompressed data) and comparable performance to that of uncompressed data accesses (in parallel environments)
$\mathbf{2 0 - 3 5 \%}$ storage reductions (compared to uncompressed data) and negligible decompression overheads
$80 \%$ storage reductions (compared to uncompressed data) and up to $>2 x$ speedup over modern graph compression schemes (Webgraph)

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

Key insight (offsets)
Up to $\mathbf{> 9 0 \%}$ storage reductions (compared to uncompressed data) and comparable performance to that of uncompressed data accesses (in parallel environments)

## Other results



A Near-Optimal Graph Representation


## What is Log(Graph)?



## A Near-Optimal Graph Representation

## An Extensible Graph Representation

| WNGPEL | Wamm Erizürich |
| :---: | :---: |
| What is the lowest storage we (hope to) use to store a grap | ! $s=\left\{x_{1}, x_{2}, x_{3}, \ldots\right\} \begin{aligned} & x_{1} \rightarrow 0 . \ldots 1 \\ & x_{2} \rightarrow 0 . \ldots 10 \\ & x_{3} \rightarrow 0 \ldots 11\end{aligned}$ |
|  |  |

## What is Log(Graph)?

\section*{8

-58
-8
-88
.8
.8 <br> A High-Performance Graph Representation <br> 

## A Near-Optimal Graph Representation



## An Extensible Graph Representation



## What is Log(Graph)?

## A High-Performance Graph Representation



A Condensed Graph Representation




## An Extensible Graph Representation



## What is Log(Graph)?

A High-Performance Graph Representation


## A Condensed Graph Representation


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


## An Extensible Graph Representation



## What is Log(Graph)?

## Thank you for your attention

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

A Condensed Graph Representation



A High-Performance Graph Representation



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

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

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

## Symbols <br> $\widehat{W}$ : max edge weight, <br> $n$ : \#vertices, <br> $m$ : \#edges, <br> $d_{v}$ : degree of vertex $v$, <br> $N_{v}$ : neighbors (adj. array) of vertex $v$, <br> $\widehat{N_{v}}$ : maximum among $N_{v}$



캐zürich
(1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

Lower bounds (global)

## Symbols <br> $\widehat{W}$ : max edge weight, <br> $n$ : \#vertices, <br> $m$ : \#edges, <br> $d_{v}$ : degree of vertex $v$, <br> $N_{v}$ : neighbors (adj. array) of vertex $v$, <br> $\widehat{N_{v}}$ : maximum among $N_{v}$



캐zürich

## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

## Lower bounds (global)

$\lceil\log n\rceil$

## Symbols <br> $\widehat{W}$ : max edge weight, <br> $n$ : \#vertices, <br> $m$ : \#edges, <br> $d_{v}$ : degree of vertex $v$, <br> $N_{v}$ : neighbors (adj. array) of vertex $v$, <br> $\widehat{N_{v}}$ : maximum among $N_{v}$



דㅐzürich
(1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

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

## Symbols <br> $\widehat{W}$ : max edge weight, <br> $n$ : \#vertices, <br> $m$ : \#edges, <br> $d_{v}$ : degree of vertex $v$, <br> $N_{v}$ : neighbors (adj. array) of vertex $v$, <br> $\widehat{N_{v}}$ : maximum among $N_{v}$



캐zürich

## (1) Log $\binom{$ Vertex }{ labels }, $\log \binom{$ 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 $N_{v}$

## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ 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)

## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ 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:

## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ 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., $V=\left\{1, \ldots, 2^{22}\right\}$


## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ 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 $N_{v}$

캐zürich


## Lower bounds (local)

Assume:

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


## $1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

## Lower bounds (global)

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

## 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 $N_{v}$


## Lower bounds (local)

Assume:

- a graph, e.g., $V=\left\{1, \ldots, 2^{22}\right\}$
- A vertex $v$ with few neighbors: $d_{v} \ll n$
- ...all these neighbors have small labels: $\widehat{N_{v}} \ll n$


## Lower bounds (global)

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

## 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 $N_{v}$

캐zürich


## Lower bounds (local)

Assume:

- a graph, e.g., $V=\left\{1, \ldots, 2^{22}\right\}$
- A vertex $v$ with few neighbors: $d_{v} \ll n$
$-\ldots$ all these neighbors have small labels: $\widehat{N_{v}} \ll n$



## Lower bounds (global)

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

## 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 $N_{v}$


## Not really (:)

ЕІІzürich

## Lower bounds (local)

Assume:

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

- a graph, e.g., $V=\left\{1, \ldots, 2^{22}\right\}$
- A vertex $v$ with few neighbors: $d_{v} \ll n$
- ...all these neighbors have small labels: $\widehat{N_{v}} \ll n$



## 1) $\log \left({ }^{\text {Vertex }}\right), \log ($ Edge labels $), \log ($ 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:
$\left\lceil\log 2^{22}\right\rceil=22$

- a graph, e.g., $V=\left\{1, \ldots, 2^{22}\right\}$
- A vertex $v$ with few neighbors: $d_{v} \ll n$
- ...all these neighbors have small labels: $\widehat{N_{v}} \ll n$



## 1) $\log \left({ }^{\text {Vertex }}\right), \log ($ Edge labels $), \log ($ 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:
$\left\lceil\log 2^{22}\right\rceil=22$

- a graph, e.g., $V=\left\{1, \ldots, 2^{22}\right\}$
- A vertex $v$ with few neighbors: $d_{v} \ll n$
- ...all these neighbors have small labels: $\widehat{N_{v}} \ll n$


19 zeros!

## $1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

## Lower bounds (global)

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

## 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 $N_{v}$

## Lower bounds (local)

Assume:


19 zeros!
Thus, use the local bound $\left\lceil\log \widehat{N_{v}}\right\rceil$
(1) Log $\binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

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

This is it? Still not really


## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

This is it? Still not really


## Lower bounds (local):

 distributed memories
## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

This is it? Still not really

| $n$ | $:$ \#ver |
| :---: | :---: |
| $m$ | $: \#$ ed |
| $H$ | $:$ num |
| $H_{i}$ | $:$ num |
|  | elem |
| $N$ | $:$ num |
|  |  |
| XE/XT |  |
| computer |  |

## Lower bounds (local): distributed memories



A Cray XE/XT
supercomputer

## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

This is it? Still not really

| $n$ | : \#vertices, Symbols |
| ---: | :--- |
| $m$ | : \#edges, |
| $H$ | : number of compute nodes |
| $H_{i}$ | : number of machine |
|  | elements at level $i$, |



## Lower bounds (local): distributed memories



## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

This is it? Still not really

| $n$ | $:$ \#vertices, Symbols |
| ---: | :--- |
| $m$ | $:$ \#edges, |
| $H$ | : number of compute nodes |
| $H_{i}$ | : number of machine |
|  | elements at level $i$, |



## Lower bounds (local): distributed memories



## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

This is it? Still not really

| $n$ | : \#vertices, Symbols |
| :--- | :--- |
| $m$ | : \#edges, |
| $H$ | : number of compute nodes |
| $H_{i}$ | : number of machine |
|  | elements at level $i$, |



## Lower bounds (local): distributed memories



## $1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

This is it? Still not really

| $n$ | : \#vertices, Symbols |
| ---: | :--- |
| $m$ | : \#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



## $1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

This is it? Still not really

| $n$ | : \#vertices, Symbols |
| ---: | :--- |
| $m$ | : \#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) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

This is it? Still not really

| $n$ | $:$ \#vertices, Symbols |
| ---: | :--- |
| $m$ | $:$ \#edges, |
| $H$ | : number of compute nodes |
| $H_{i}$ | : number of machine |
|  | elements at level $i$, |



## Lower bounds (local): distributed memories



## $1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

This is it? Still not really

| $n$ | $:$ \#vertices, Symbols |
| ---: | :--- |
| $m$ | $:$ \#edges, |
| $H$ | : number of compute nodes |
| $H_{i}$ | : number of machine |
|  | elements at level $i$, |



## Lower bounds (local): distributed memories



## ज्İzürich

## $1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

This is it? Still not really

| $n$ | $:$ \#vertices, Symbols |
| ---: | :--- |
| $m$ | $:$ \#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:


## $1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

This is it? Still not really

| $n$ | $:$ \#vertices, Symbols |
| ---: | :--- |
| $m$ | $:$ \#edges, |
| $H$ | : number of compute nodes |
| $H_{i}$ | : number of machine |
|  | elements at level $i$, |



## Lower bounds (local): distributed memories

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


## $1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

This is it? Still not really

| $n$ | $:$ \#vertices, Symbols |
| ---: | :--- |
| $m$ | $:$ \#edges, |
| $H$ | : number of compute nodes |
| $H_{i}$ | : number of machine |
|  | elements at level $i$, |



## Lower bounds (local): distributed memories

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

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


## $1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

## This is it? Still not really

| $n$ | $:$ \#vertices, Symbols |
| ---: | :--- |
| $m$ | $:$ \#edges, |
| $H$ | : number of compute nodes |
| $H_{i}$ | : number of machine |
|  | elements at level $i$, |



## Lower bounds (local): distributed memories

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

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]: $\lceil\log H\rceil$


A Cray XE/XT
supercomputer

4 cabinets:

3 chassis:

8 blades:

4 nodes:
$H=4$
32 cores:


## €Нzürich

## $1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

## This is it? Still not really

| $n$ | $:$ \#vertices, Symbols |
| ---: | :--- |
| $m$ | $:$ \#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: $\bar{H}$

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]: $\lceil\log H\rceil$


A Cray XE/XT supercomputer

4 cabinets

3 chassis:

8 blades:

4 nodes:
$H=4$
32 cores:


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

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

## $1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

## This is it? Still not really

$n$ : \#vertices,
Symbols
$m$ : \#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: $\bar{H}$

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]: $\lceil\log H\rceil$


A Cray XE/XT
supercomputer

4 cabinets:

3 chassis:

8 blades

4 nodes:
$H=4$
32 cores:


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:

## $1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

## This is it? Still not really

$n$ : \#vertices,
Symbols
$m$ : \#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: $\bar{H}$

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]: $\lceil\log H\rceil$


A Cray XE/XT
supercomputer

4 cabinets:

3 chassis:

8 blades

4 nodes:
$H=4$
32 cores:


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:

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

## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }

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

$\oiint$ Formal analyses: more (check the paper ©)

$$
\forall_{v, u \in V}\left(u \in N_{v}\right) \Rightarrow\left[\mathcal{N}(u) \leq \widehat{N}_{v}\right]
$$

$$
\begin{array}{rr}
|\mathscr{A}|=\sum_{v \in V}\left(d_{v}\left\lceil\log \widehat{N}_{v}\right\rceil+\left\lceil\log \log \widehat{N}_{v}\right\rceil\right) & |\mathscr{A}|=\sum_{v \in V}\left(d_{v}\left\lceil\log \widehat{N}_{v}\right\rceil+\left\lceil\log \log \widehat{N}_{v}\right\rceil\right) \\
|\mathcal{A}|=n\left\lceil\log \frac{n}{\mathcal{H}}\right\rceil+\mathcal{H}\lceil\log \mathcal{H}\rceil & |\mathcal{A}|=2 m(\lceil\log n\rceil+\lceil\log \widehat{\mathcal{W}}\rceil) \\
|\mathcal{A}|=\sum_{v \in V}\left(d_{v}\left(\left\lceil\log \widehat{N}_{v}\right\rceil+\lceil\log \widehat{\mathcal{W}}\rceil\right)+\left\lceil\log \log \widehat{N}_{v}\right\rceil+\lceil\log \log \widehat{\mathcal{W}}\rceil\right)
\end{array}
$$

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

## (1) $\log \binom{$ Vertex $)}{$ labels }, $\log \binom{($ Edge }{ weights }

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

## Æ Formal analyses: more (check the paper ©)

$$
\begin{aligned}
& \forall_{v, u \in V}\left(u \in N_{v}\right) \Rightarrow\left[\mathcal{N}(u) \leq \widehat{N}_{v}\right] \\
& I=\sum_{v \in V}\left(d_{v}\left\lceil\log \widehat{N}_{v}\right\rceil+\left\lceil\log \log \widehat{N}_{v}\right\rceil\right) \\
& \quad|\mathcal{A}|=2 m(\lceil\log n\rceil+\lceil\log \widehat{\mathcal{N}}\rceil) \\
& \left.\left.\widehat{N}_{v}\right\rceil+\lceil\log \widehat{\mathcal{W}}\rceil\right)+\left\lceil\log \log \widehat{N}_{v}\right\rceil+\lceil\log \log \widehat{\mathcal{W} \mid}) \\
& E[|\mathcal{A}|]=(\lceil\log n\rceil+\lceil\log \widehat{\mathcal{W}}\rceil) p n^{2}
\end{aligned}
$$

/* Input: graph $G$, Output: a new relabeling $\mathcal{N}(v), \forall v \in V$. */
void relabel (G) \{
$I D[0 . . n-1]=[0 . . n-1] ; / / A n$ array with vertex IDs.
$D[0 . . n-1]=\left[d_{0} . . d_{n-1}\right] ; / / A n$ array with degrees of vertices
//An auxiliary array for determining if a vertex was relabeled:
visit $[0 . . n-1]=[$ false..false $]$;
$n l=1 ; / / A n$ auxiliary variable ${ }^{\prime}$ new label'
sort(ID); sort (D);
for (int $i=1 ; i<n ;++i) / / F o r ~ e a c h ~ v e r t e x ~$
for (int $j=0 ; j<D[i] ;++j$ ) \{ //For each neighbor
int $i d=N_{j, I D[i]} ; / / N_{j, I D[i]}$ is $j$ th neighbor of vertex with ID $I D[i]$
if(visit $[$ id $]==$ false) \{
$\mathcal{N}(i d)=n l++$;
visit $[$ id $]=$ true;
\}\}
for (int $i=1 ; i<n ;++i$ )
if(visit $[i]==$ false)
$\mathcal{N}(i d)=n l++$;
\}

## (2) Log (Offset structure)

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


## जНzürich

## (2) Log ( offset structure)

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


## $\oiint$ Formal analyses

## (2) Log ( offset structure)

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


## $\oiint$ Formal analyses

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

## (2) Log ( offset structure)

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

$\oiint$ Formal analyses
Check the paper for details $\odot$

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

## 2 Log ( offset structure)

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

$\oiiint$ Formal analyses

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

## x. 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) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights } Storage



## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights } Storage



## (1) $\log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights } Storage



## (1) Log $\binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights } Storage

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


ヨНzürich

## (1) Log $\binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }$\quad$ Performance

SSSP


Kronecker graphs Number of vertices: 4M

Number of edges per vertex

## (1) Log $\binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights } Performance



SSSP

Number of edges per vertex

## Log(Graph) accelerates GAPBS <br> Log(Graph)

 -

Kronecker graphs Number of vertices: 4M

## $1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights } Performance



Number of edges per vertex

## Log(Graph) accelerates GAPBS

SSSP


Kronecker graphs Number of vertices: 4M

Both storage and performance are improved simultaneously

## (1) Log $\binom{$ Vertex }{ labels } , $\log \binom{$ Edge }{ weights } Performance

# Betweenness Centrality 

"LG": Log(Graph) Trad: Traditional (non compressed, GAPBS)
" g ": global scheme
" $I$ ": local scheme "gap": additional gap encoding


Kronecker graphs Number of vertices: 4M

## $1 \log \binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }$\quad$ Performance

Sparse graphs


Dense graphs

"LG": Log(Graph) Trad: Traditional (non compressed, GAPBS)
" g ": global scheme " I ": local scheme "gap": additional gap encoding


Kronecker graphs Number of vertices: 4M

## $1 \log \binom{$ Vertex }{ labels } , $\log \binom{$ Edge }{ weights }$\quad$ Performance

Sparse graphs


Dense graphs


Betweenness Centrality
"LG": $\log ($ Graph $)$ Trad: Traditional (non compressed, GAPBS)
" g ": global scheme "I": local scheme "gap": additional gap encoding

Log(Graph) incurs negligible overheads

Kronecker graphs Number of vertices: 4M

## (1) Log $\binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights } Performance

BFS
"LG": Log(Graph) Trad: Traditional (non compressed, GAPBS)
" g ": global scheme
"I": local scheme "gap": additional gap encoding


Kronecker graphs Number of vertices: 4M

## (1) Log $\binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights }$\quad$ Performance



BFS
"LG": Log(Graph) Trad: Traditional (non compressed, GAPBS)
" g ": global scheme "I": local scheme "gap": additional gap encoding


Kronecker graphs Number of vertices: 4M

## (1) Log $\binom{$ Vertex }{ labels }, $\log \binom{$ Edge }{ weights } Performance

BFS
"LG": Log(Graph) Trad: Traditional (non compressed, GAPBS)
"g": global scheme
"l". Iocal scheme
Both storage and performance are improved simultaneously


Kronecker graphs Number of vertices: 4 M

## (1) $\log \binom{$ Vertex }{ thets }, $\log \binom{$ Edge }{ wdights } labels), LOg (weights

Log(Graph) accelerates GAPBS


Performance

Dense graphs

Sparse graphs
"LG": Log(Graph) Trad: Traditional (non compressed, GAPBS)
"g": global scheme
"l". Iocal scheme
Both storage and performance are improved simultaneously


Kronecker graphs Number of vertices: 4M

## (1) $\log ($ Vertex $), \log ($ Edge $)$ Communicated labels $), \log ($ weights $)$ data



Various real-world and synthetic graphs

## (1) $\log ($ Vertex $) \log ($ Edge $)$ Communicated labels $) \operatorname{LOg}($ weights $)$ data



Various real-world and synthetic graphs

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

## NAREL

ЄHzürich


## (3) Log $\binom{$ 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 $\binom{$ 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) $\log \binom{$ 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


Scheme:

| Trad |
| :---: |
| DMd |
| DMf |
| $\square$ |
| BRB |
| $\square$ RB |
| $\square$ |

Lots of data ©
Various real-world graphs

## Conclusions:

## (3) $\log \binom{$ Adjacency }{ structure }$\quad$ 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


Scheme:

| Trad |
| :---: |
| DMd |
| DMf |
| -BRB |
| RB |
| $W G$ |

Lots of data $;$
Conclusions:

Various real-world graphs

## (3) Log $\binom{$ 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


Scheme:

| $\substack{\text { Trad } \\ \text { DMd } \\ \text { DMI } \\ \text { DMI } \\ \text { BRB } \\ \text { BB } \\ \hline W G \\ \hline \\ \hline}$ |
| :---: |

Lots of data © Conclusions:

Various real-world graphs

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

## 3 log Adjacency structure <br> 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


Scheme:

| $\substack{\text { Trad } \\ \text { DMd } \\ \text { DMI } \\ \text { DMI } \\ \text { BRB } \\ \text { BB } \\ \hline W G \\ \hline \\ \hline}$ |
| :---: |

Lots of data © Conclusions:

WebGraph best for web graphs ©

Various real-world graphs

DMd: much better than DMf, often comparable to others

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

## NAREL

## (3) Log $\binom{$ 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




Scheme:
Trad $\square$ RB $\square$ DMd
$\square \mathrm{DMf} \square \mathrm{WG}$

## (3) Log $\binom{$ 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




Scheme:
Trad $\square$ RB
$\square$ DMd
$\square$ DM
$\square W G$

## (3) Log $\binom{$ 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




Scheme:
Trad $\square$ RB
$\square$ DMd
$\square$ DM
$\square W G$

Log(Graph) full design...

Log(Graph) full design...


## Log(Graph) full design...



## Log(Graph) full design...



## Log(Graph) full design...



## Log(Graph) full design...



## Log(Graph) full design...



## Log(Graph) full design...

## Understand storage lower bounds and the theory



## Log(Graph) full design...

## Understand storage lower bounds and the theory



## Log(Graph) full design...

## Understand storage lower bounds and the theory




Ensure high-performance implementation

Use Integer Linear Programming (ILP) for more storage reductions

## $\mathcal{X}$ 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

## ※X 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 | ptr $W$ | $O(W n)$ | $W(n+1)$ | $O(1)$ |
| Plain [44] | bvPL | $O\left(\frac{W m}{B}\right)$ | $\frac{2 W m}{B}$ | $O(1)$ |
| Interleaved [44] | bvIL | $O\left(\frac{W m}{B}+\frac{W m}{L}\right)$ | $2 W m\left(\frac{1}{B}+\frac{64}{L}\right)$ | $O\left(\log \frac{W m}{B}\right)$ |
| Entropy based [31, 78] | bvEN | $O\left(\frac{W m}{B} \log \frac{W m}{B}\right)$ | $\approx \log \left(\frac{2 W m}{B}\right)$ | $O\left(\log \frac{W m}{B}\right)$ |
| Sparse [76] | bvSD | $O\left(n+n \log \frac{W m}{B n}\right)$ | $\approx n\left(2+\log \frac{2 W m}{B n}\right)$ | $O(1)$ |
| B-tree based [1] | bvBT | $O\left(\frac{W m}{B}\right)$ | $\approx 1.1 \cdot \frac{2 W m}{B}$ | $O(\log n)$ |
| Gap-compressed [1] | bvGC | $O\left(\frac{W m}{B} \log \frac{W m}{B n}\right)$ | $\approx 1.3 \cdot \frac{2 W m}{B} \log \frac{2 W m}{B n}$ | $O(\log n)$ |

天 Key method (vertex labels)

## K Key method (neighborhoods)

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

Modern bitwise operations

## K× 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 | ptr $W$ | $O(W n)$ | $W(n+1)$ | $O(1)$ |
| Plain [44] | bvPL | $O\left(\frac{W m}{B}\right)$ | $\frac{2 W m}{B}$ | $O(1)$ |
| Interleaved [44] | bvIL | $O\left(\frac{W m}{B}+\frac{W m}{L}\right)$ | $2 W m\left(\frac{1}{B}+\frac{64}{L}\right)$ | $O\left(\log \frac{W m}{B}\right)$ |
| Entropy based [31, 78] | bvEN | $O\left(\frac{W m}{B} \log \frac{W m}{B}\right)$ | $\approx \log \left(\frac{2 W m}{B}\right)$ | $O\left(\log \frac{W m}{B}\right)$ |
| Sparse [76] | bvSD | $O\left(n+n \log \frac{W m}{B n}\right)$ | $\approx n\left(2+\log \frac{2 W m}{B n}\right)$ | $O(1)$ |
| B-tree based [1] | bvBT | $O\left(\frac{W m}{B}\right)$ | $\approx 1.1 \cdot \frac{2 W m}{B}$ | $O(\log n)$ |
| Gap-compressed [1] | bvGC | $O\left(\frac{W m}{B} \log \frac{W m}{B n}\right)$ | $\approx 1.3 \cdot \frac{2 W m}{B} \log \frac{2 W m}{B n}$ | $O(\log n)$ |

* 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

## K× 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 | $O(W n)$ | $W(n+1)$ | $O(1)$ |
| Plain [44] | bvPL | $O\left(\frac{W m}{B}\right)$ | $\frac{2 W m}{B}$ | $O(1)$ |
| Interleaved [44] | bvIL | $O\left(\frac{W m}{B}+\frac{W m}{L}\right)$ | $2 W m\left(\frac{1}{B}+\frac{64}{L}\right)$ | $O\left(\log \frac{W m}{B}\right)$ |
| Entropy based [31, 78] | bvEN | $O\left(\frac{W m}{B} \log \frac{W m}{B}\right)$ | $\approx \log \left(\frac{2 W m}{B}\right)$ | $O\left(\log \frac{W m}{B}\right)$ |
| Sparse [76] | bvSD | $O\left(n+n \log \frac{W m}{B n}\right)$ | $\approx n\left(2+\log \frac{2 W m}{B n}\right)$ | $O(1)$ |
| B-tree based [1] | bvBT | $O\left(\frac{W m}{B}\right)$ | $\approx 1.1 \cdot \frac{2 W m}{B}$ | $O(\log n)$ |
| Gap-compressed [1] | bvGC | $O\left(\frac{W m}{B} \log \frac{W m}{B n}\right)$ | $\approx 1.3 \cdot \frac{2 W m}{B} \log \frac{2 W m}{B n}$ | $O(\log n)$ |

* Key method (vertex labels)

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

Modern bitwise operations

## K 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

## K× 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 | $O(W n)$ | $W(n+1)$ | $O(1)$ |
| Plain [44] | bvPL | $O\left(\frac{W m}{B}\right)$ | $\frac{2 W m}{B}$ | $O(1)$ |
| Interleaved [44] | bvIL | $O\left(\frac{W m}{B}+\frac{W m}{L}\right)$ | $2 W m\left(\frac{1}{B}+\frac{64}{L}\right)$ | $O\left(\log \frac{W m}{B}\right)$ |
| Entropy based [31, 78] | bvEN | $O\left(\frac{W m}{B} \log \frac{W m}{B}\right)$ | $\approx \log \left(\frac{2 W m}{B}\right)$ | $O\left(\log \frac{W m}{B}\right)$ |
| Sparse [76] | bvSD | $O\left(n+n \log \frac{W m}{B n}\right)$ | $\approx n\left(2+\log \frac{2 W m}{B n}\right)$ | $O(1)$ |
| B-tree based [1] | bvBT | $O\left(\frac{W m}{B}\right)$ | $\approx 1.1 \cdot \frac{2 W m}{B}$ | $O(\log n)$ |
| Gap-compressed [1] | bvGC | $O\left(\frac{W m}{B} \log \frac{W m}{B n}\right)$ | $\approx 1.3 \cdot \frac{2 W m}{B} \log \frac{2 W m}{B n}$ | $O(\log n)$ |


[^0]:    Symbols
    $n$ : \#vertices,
    $m$ : \#edges,
    $d_{v}$ : degree of vertex $v$,
    $N_{v}$ : neighbors (adj. array) of
    vertex $v$,
    $\widehat{N_{v}}$ : maximum among $N_{v}$

