Space and Time-Efficient Data Structures for Massive Datasets

Giulio Ermanno Pibiri giulio.pibiri@di.unipi.it

> Supervisor Rossano Venturini

Computer Science Department University of Pisa

17/10/2016





The increase of information does **not** scale with technology.

Evidence

The increase of information does **not** scale with technology.



"Software is getting slower more rapidly than hardware becomes faster." Niklaus Wirth, A Plea for Lean Software

Evidence

The increase of information does **not** scale with technology.



"Software is getting slower more rapidly than hardware becomes faster."

Niklaus Wirth, A Plea for Lean Software

Data Structures PERFORMANCE

"how quickly a program does its work" - faster work



Algorithms EFFICIENCY

"how much work is required by a program" - **less** work

Data Structures PERFORMANCE

"how quickly a program does its work" - faster work



$$\quad \longleftrightarrow \quad$$

Algorithms EFFICIENCY

"how much work is required by a program" - less work

Data Structures PERFORMANCE

"how quickly a program does its work" - faster work



$$\quad \longleftrightarrow \quad$$

Algorithms EFFICIENCY

"how much work is required by a program" - **less** work

Data Compression

Data Structures PERFORMANCE

"how quickly a program does its work" - faster work



$$\quad \longleftrightarrow \quad$$

Algorithms EFFICIENCY

"how much work is required by a program" - **less** work

Data Compression

+ space - time

Data Structures PERFORMANCE

"how quickly a program does its work" - faster work





Algorithms EFFICIENCY

"how much work is required by a program" - less work

Data Compression

spacetime

Dichotomy Problem

Small VS Fast?

Dichotomy Problem

Small VS Fast? Choose one.

Dichotomy Problem

Small VS Fast? Choose one.

NO

High Level Thesis

Data Structures + Data Compression - Faster Algorithms

High Level Thesis

Data Structures + Data Compression - Faster Algorithms



"Space optimization is closely related to time optimization in a disk memory."

Donald E. Knuth, The Art of Computer Programming







Numbers are taken from: <u>https://gist.github.com/jboner/2841832</u>





Design **space-efficient** *ad-hoc* data structures, both from a theoretical *and* practical perspective, that support **fast data extraction**.



Design **space-efficient** *ad-hoc* data structures, both from a theoretical *and* practical perspective, that support **fast data extraction**.

Data compression & Fast Retrieval together.

Mature algorithmic solutions *now* ready for technology transfer.

Proposal

Must exploit properties of the addressed problem.

Design **space-efficient** *ad-hoc* data structures, both from a theoretical *and* practical perspective, that support **fast data extraction**.

Data compression & Fast Retrieval together.

Mature algorithmic solutions *now* ready for technology transfer.

Because engineered data structures are the ones the *boost* the *availability* and *wealth* of information around us.

Because engineered data structures are the ones the *boost* the *availability* and *wealth* of information around us.

Inverted Indexes
N-grams
B-trees
....

Because engineered data structures are the ones the *boost* the *availability* and *wealth* of information around us.

- Inverted Indexes
- N-grams
- B-trees
- . . .





Because engineered data structures are the ones the *boost* the *availability* and *wealth* of information around us.

- Inverted Indexes
- N-grams
- B-trees

....





Because engineered data structures are the ones the *boost* the *availability* and *wealth* of information around us.

- Inverted Indexes
- N-grams
- B-trees

....



Because engineered data structures are the ones the *boost* the *availability* and *wealth* of information around us.

- Inverted Indexes
- N-grams
- B-trees
- . . .



facebook Google Stopbox



Because engineered data structures are the ones the *boost* the *availability* and *wealth* of information around us.

- Inverted Indexes
- N-grams
- B-trees

...

facebook Google Stopbox

Because engineered data structures are the ones the *boost* the *availability* and *wealth* of information around us.





- Inverted Indexes
- N-grams
- B-trees

•





Can't we use existing libraries?

Can't we use existing libraries?

C++ Standard Template Library (STL)?

Can't we use existing libraries?

C++ Standard Template Library (STL)?

std::list std::stack std::queue std::map std::unordered map

Can't we use existing libraries?

C++ Standard Template Library (STL)?

std::list std::stack std::queue std::map



std::unordered map

Can't we use existing libraries?

C++ Standard Template Library (STL)?

std::list std::stack std::queue std::map



std::unordered map

Prefer cache-friendly (non discontiguous) data structures. Always.

Use std::vector.

Example: vector of strings VS string pool

Prefer cache-friendly data structures.


Prefer cache-friendly data structures.



vector of offsets

memory

0	1	2	3
4	5	6	7
8	9	10	11
12	13	14	15
16	17	18	19
20	21	22	23
24	25	26	27
28	29	30	31
32	33	34	35
36	37	38	39
40	41	42	

Prefer cache-friendly data structures.



vector of offsets

0
5
8
12
18
20
25
26
28
34
35
38
40

memory

0	1	2	3
4	5	6	7
8	9	10	11
12	13	14	15
16	17	18	19
20	21	22	23
24	25	26	27
28	29	30	31
32	33	34	35
36	37	38	39
40	41	42	

Prefer cache-friendly data structures.





Offsets instead of pointers. Contiguous memory layout.





Prefer cache-friendly data structures.



giulio@xor:~/sxlm/build\$./string_vector_benchmark 50000000 ~/random_strings.50M.128
2016-10-12 09:43:31: Loading strings
2016-10-12 09:43:39: Scanning strings
3224822962
2016-10-12 09:43:46: elapsed 6.705554 [sec]

Prefer cache-friendly data structures.



3224822962

2016-10-12 09:43:46: elapsed 6.705554 [sec]













Given a *textual collection* D, each document can be seen as a (multi-)set of terms. The set of terms occurring in D is the *lexicon* T.

For each term t in T we store in a list L_t the identifiers of the documents in which t appears.

Given a *textual collection* D, each document can be seen as a (multi-)set of terms. The set of terms occurring in D is the *lexicon* T.

For each term t in T we store in a list L_t the identifiers of the documents in which t appears.



Given a *textual collection* D, each document can be seen as a (multi-)set of terms. The set of terms occurring in D is the *lexicon* T.

For each term t in T we store in a list L_t the identifiers of the documents in which t appears.



Given a *textual collection* D, each document can be seen as a (multi-)set of terms. The set of terms occurring in D is the *lexicon* T.

For each term t in T we store in a list L_t the identifiers of the documents in which t appears.



Given a *textual collection* D, each document can be seen as a (multi-)set of terms. The set of terms occurring in D is the *lexicon* T.

For each term t in T we store in a list L_t the identifiers of the documents in which t appears.



2

2







Inverted Indexes owe their popularity to the *efficient resolution of queries*, such as: "return me all documents in which terms $\{t_1, ..., t_k\}$ occur".



posting lists intersection

General Problem

Consider a sequence S[0,n) of n *positive* and *monotonically increasing integers*, i.e., S[i] \leq S[i+1] for 0 \leq i < n-1, possibly repeated.

How to represent it as a *bit vector* in which each original integer is *self-delimited*, using as few as possible bits?

General Problem

Consider a sequence S[0,n) of n *positive* and *monotonically increasing integers*, i.e., S[i] \leq S[i+1] for 0 \leq i < n-1, possibly repeated.

How to represent it as a *bit vector* in which each original integer is *self-delimited*, using as few as possible bits?

Huge research corpora describing different space/time trade-offs.

- Elias gamma/delta [Elias, TIT 1975]
- Variable Byte [Salomon, Springer 2007]
- Binary Interpolative Coding [Moffat and Stuiver, IRJ 2000]
- Simple-9 [Anh and Moffat, IRJ 2005]
- PForDelta [Zukowski et al., ICDE 2006]
- OptPFD [Yan *et al.*, WWW 2009]
- Simple-16 [Anh and Moffat, SPE 2010]
- Varint-G8IU [Stepanov et al., CIKM 2011]
- Elias-Fano [Vigna, WSDM 2013]
- Partitioned Elias-Fano [Ottaviano and Venturini, SIGIR 2014]

General Problem

Consider a sequence S[0,n) of n *positive* and *monotonically increasing integers*, i.e., S[i] \leq S[i+1] for 0 \leq i < n-1, possibly repeated.

How to represent it as a *bit vector* in which each original integer is *self-delimited*, using as few as possible bits?

Huge research corpora describing different space/time trade-offs.

- Elias gamma/delta [Elias, TIT 1975]
- Variable Byte [Salomon, Springer 2007]
- Binary Interpolative Coding [Moffat and Stuiver, IRJ 2000]
- Simple-9 [Anh and Moffat, IRJ 2005]
- PForDelta [Zukowski et al., ICDE 2006]
- OptPFD [Yan *et al.*, WWW 2009]
- Simple-16 [Anh and Moffat, SPE 2010]
- Varint-G8IU [Stepanov et al., CIKM 2011]
- Elias-Fano [Vigna, WSDM 2013]
- Partitioned Elias-Fano [Ottaviano and Venturini, SIGIR 2014]



https://blogs.dropbox.com/tech/2016/09/improving-the-performance-of-full-text-search/

Improving the performance of full-text search

Adam Faulkner | September 7, 2016



For Firefly, Dropbox full text-search engine, speed has always been a priority.

They *were unable to scale* because of the dimension of their (distributed) inverted index. Consequence? Query time latencies deteriorate **from 250ms to 1s**.



They *were unable to scale* because of the dimension of their (distributed) inverted index. Consequence? Query time latencies deteriorate **from 250ms to 1s**.



They *were unable to scale* because of the dimension of their (distributed) inverted index. Consequence? Query time latencies deteriorate **from 250ms to 1s**.

Solution?



Consequence? Query time latencies deteriorate from 250ms to 1s.

Solution?

Compress the index to reduce I/O pressure.



Compress the index to reduce I/O pressure.

Data Structures + Data Compression - Faster Algorithms

Elias-Fano - Genesis



Peter Elias [1923 - 2001] Robert Fano [1917 - 2016]

Robert Fano. *On the number of bits required to implement an associative memory*. Memorandum 61, Computer Structures Group, MIT (1971).

Peter Elias. *Efficient Storage and Retrieval by Content and Address of Static Files*. Journal of the ACM (JACM) 21, 2, 246–260 (1974).

Elias-Fano - Genesis



Peter Elias [1923 - 2001] Robert Fano [1917 - 2016]

Robert Fano. *On the number of bits required to implement an associative memory*. Memorandum 61, Computer Structures Group, MIT (1971).

Peter Elias. *Efficient Storage and Retrieval by Content and Address of Static Files*. Journal of the ACM (JACM) 21, 2, 246–260 (1974).



Sebastiano Vigna. Quasi-Succinct Indices.

In Proceedings of the 6-th ACM International Conference on Web Search and Data Mining (WSDM), 83-92 (2013).

40 years later!

Elias-Fano - Encoding example

2

Elias-Fano - Encoding example

7

2

43 u = (8



Elias-Fano - Encoding example



Elias-Fano - Encoding example



21 7
high

lg n



L = 011100111101110111101011

high

lg n

3



L = 011100111101110111101011



L = 011100111101110111101011



lg n

3

1



L = 011100111101110111101011

lg n



L = 011100111101110111101011

missing

high bits



L = 011100111101110111101011







$$EF(S[0,n)) = n \left[\lg \frac{u}{n} \right] + 2n \text{ bits}$$

3

$$EF(S[0,n)) = n \left[\lg \frac{u}{n} \right] + 2n \text{ bits}$$

3

$$EF(S[0,n)) = n \left[\lg \frac{u}{n} \right] + 2n \text{ bits}$$

$$|\chi| = \begin{pmatrix} u+n \\ n \end{pmatrix}$$

3

$$EF(S[0,n)) = n \left[lg \frac{u}{n} \right] + 2n \text{ bits}$$

$$\begin{vmatrix} \chi \\ = \begin{pmatrix} u+n \\ n \end{pmatrix}$$
$$\left[\lg \begin{pmatrix} u+n \\ n \end{pmatrix} \right] \approx n \lg \frac{u+n}{n}$$

3

$$EF(S[0,n)) = n \left[lg \frac{u}{n} \right] + 2n$$
 bits

$$\begin{vmatrix} x \\ x \end{vmatrix} = \begin{pmatrix} u+n \\ n \end{pmatrix}$$
$$\left[\lg \begin{pmatrix} u+n \\ n \end{pmatrix} \right] \approx \ln \lg \frac{u+n}{n}$$

$$EF(S[0,n)) = n \left[lg \frac{u}{n} \right] + 2n$$
 bits

(less than half a bit away [Elias, JACM 1974])

3

$$\begin{vmatrix} \chi \\ = \begin{pmatrix} u+n \\ n \end{pmatrix}$$
$$\left[\lg \begin{pmatrix} u+n \\ n \end{pmatrix} \right] \approx \ln \lg \frac{u+n}{n}$$

$$EF(S[0,n)) = n \left[\lg \frac{u}{n} \right] + 2n \text{ bits}$$

(less than half a bit away [Elias, JACM 1974])

3

$$\begin{vmatrix} \chi \\ = \begin{pmatrix} u+n \\ n \end{pmatrix}$$
$$\left[\lg \begin{pmatrix} u+n \\ n \end{pmatrix} \right] \approx n \lg \frac{u+n}{n}$$

3

$$EF(S[0,n)) = n \left[lg \frac{u}{n} \right] + 2n \text{ bits}$$

$$\begin{vmatrix} \chi \\ = \begin{pmatrix} u+n \\ n \end{pmatrix}$$
$$\left[\lg \begin{pmatrix} u+n \\ n \end{pmatrix} \right] \approx n \lg \frac{u+n}{n}$$

$$EF(S[0,n)) = n \left[\lg \frac{u}{n} \right] + 2n \text{ bits}$$

3

$$EF(S[0,n)) = n \left[\lg \frac{u}{n} \right] + 2n \text{ bits}$$

access to each S[i] in O(1) worst-case

3

$$EF(S[0,n)) = n \left[\lg \frac{u}{n} \right] + 2n \text{ bits}$$

access to each S[i] in O(1) worst-case

 $predecessor(x) = max\{S[i] | S[i] < x\}$ successor(x) = min{S[i] | S[i] ≥ x} queries in O(Ig $\frac{u}{n}$) worst-case

3

$$EF(S[0,n)) = n \left[\lg \frac{u}{n} \right] + 2n \text{ bits}$$

access to each S[i] in O(1) worst-case

predecessor(x) = max{S[i] | S[i] < x} successor(x) = min{S[i] | S[i] \ge x} queries in O(Ig $\frac{u}{n}$) worst-case

3

$$EF(S[0,n)) = n \left[\lg \frac{u}{n} \right] + 2n \text{ bits}$$

access to each S[i] in O(1) worst-case

$$\begin{aligned} \mathsf{bredecessor}(\mathsf{x}) &= \max\{\mathsf{S}[\mathsf{i}] \mid \mathsf{S}[\mathsf{i}] < \mathsf{x} \}\\ \mathsf{successor}(\mathsf{x}) &= \min\{\mathsf{S}[\mathsf{i}] \mid \mathsf{S}[\mathsf{i}] \ge \mathsf{x} \}\\ \mathsf{queries in } O\left(\mathsf{Ig}\,\frac{\mathsf{u}}{\mathsf{n}}\right) \text{ worst-case} \end{aligned}$$

but...

3

$$EF(S[0,n)) = n \left[\lg \frac{u}{n} \right] + 2n \text{ bits}$$

access to each S[i] in O(1) worst-case

$$\begin{aligned} \mathsf{bredecessor}(\mathsf{x}) &= \max\{\mathsf{S}[i] \mid \mathsf{S}[i] < \mathsf{x} \}\\ \mathsf{successor}(\mathsf{x}) &= \min\{\mathsf{S}[i] \mid \mathsf{S}[i] \ge \mathsf{x} \}\\ \mathsf{queries in } O\left(\mathsf{Ig}\,\frac{\mathsf{u}}{\mathsf{n}}\right) \text{ worst-case} \end{aligned}$$

but...

need o(n) bits more to support *fast* rank/select primitives on bitvector H

$$EF(S[0,n)) = n \left[\lg \frac{u}{n} \right] + 2n \text{ bits}$$

access to each S[i] in O(1) worst-case

predecessor(x) = max{S[i] | S[i] < x} successor(x) = min{S[i] | S[i] \ge x} queries in O(Ig $\frac{u}{n}$) worst-case

but...

need o(n) bits more to support *fast* rank/select primitives on bitvector H

Every encoder represents each sequence individually.

No exploitation of redundancy.

Every encoder represents each sequence individually.

No exploitation of redundancy.



Every encoder represents each sequence individually.

No exploitation of redundancy.



Idea: encode **clusters** of posting lists.

2

cluster of posting lists

2

cluster of posting lists



unbounded universe



cluster of posting lists

reference list

unbounded universe



cluster of posting lists

lg u bits

reference list

R





reference list



unbounded universe





2



Problems

- 1. how to build clusters
- 2. how to synthesise the reference list



Time VS Space tradeoffs by varying reference size





	MIN	MID MAX			MIN	MID	MAX	
PEF	2.94 (+5.60%)	2.94 (+7.91%) 2.94 (+10.95%)	PEF	4.80 (+2.13%)	4.80 (+3.98%	b) 4.80 (+6.25%)	
CPEF	2.78	2.72	2.65	CPEF	4.70	4.62	4.52	
BIC	2.80 (+0.53%)	2.80 (+2.74%) 2.80 (+5.63%)	BIC	4.27 (-9.22%)	4.27 (-7.58%	b) 4.27 (-5.56%)	
	(a) Gov2		(b) ClueWeb09				

Table 2: Bits per posting in selected trade-off points.



Figure 3: Timings for AND queries by varying the reference size on Gov2 and ClueWeb09, using the query set TREC 06.

		MIN		MID		MAX			MIN		MID	MAX
05	PEF	14.6	(17.5%)	14.6	(29.0%)	14.6 (-49.7%)	05	PEF	3.7	(30.4%)	3.7 (-37.5%)	3.7 (-52.1%)
REC (CPEF	17.7		20.6		29.1	SEC (CPEF	5.3		5.9	7.8
Ē	BIC	41.1	(+131.9%)	41.1	(+99.5%)	41.1 (+41.3%)	F	BIC	10.5	(+96.2%)	10.5 (+76.2%)	10.5 (+35.0%)
90	PEF	17.7	(16.6%)	17.7	(29.1%)	17.7 (-50.3%)	90	PEF	6.1	(27.4%)	6.1 (-35.2%)	6.1 (-49.1%)
REC (CPEF	21.2		25.0		35.6	REC (CPEF	8.3		9.3	11.9
F	BIC	55.1	(+159.7%)	55.1	(+120.8%)	55.1 (+54.7%)	F	BIC	18.5	(+122.6%)	$18.5 \ (+98.6\%)$	$18.5 \ (+56.0\%)$
	(a) ClueWeb09						(b) Gov2					

Table 3: Timings in milliseconds for AND queries on ClueWeb09 and Gov2, using query sets TREC 05 and TREC 05. In parentheses we show the relative percentage against CPEF.
Time VS Space tradeoffs by varying reference size



Figure 2: Bits per posting of Gov2 and ClueWeb09 by varying the reference size.

	MIN MID		MAX			MIN	MID	MAX
PEF	2.94 (+5.60%) 2.	94 (+7.91%	b) 2.94 (+10.95%)		PEF	4.80 (+2.13%)	4.80 (+3.989	6) 4.80 (+6.25%)
CPEF	2.78 2.	72	2.65		CPEF	4.70	4.62	4.52
BIC	2.80 (+0.53%) 2.	80 (+2.74%	b) 2.80 (+5.63%)		BIC	4.27 (-9.22%)	4.27 (-7.589	6) 4.27 (-5.56%)
	(a) Go	ov2				(b) C	lueWeb09	

Table 2: Bits per posting in selected trade-off points.



3

Figure 3: Timings for AND queries by varying the reference size on Gov2 and ClueWeb09, using the query set TREC 06.

		I	MIN	Ν	MID	MAX			1	MIN	MID	MAX
35	PEF	14.6	(17.5%)	14.6	(29.0%)	14.6 (-49.7%)	22	PEF	3.7	(30.4%)	3.7 (-37.5%)	3.7 (-52.1%)
SEC (CPEF	17.7		20.6		29.1	SEC (CPEF	5.3		5.9	7.8
Ē	BIC	41.1	(+131.9%)	41.1	(+99.5%)	41.1 (+41.3%)	F	BIC	10.5	(+96.2%)	10.5 (+76.2%)	10.5 (+35.0%)
90	PEF	17.7	(16.6%)	17.7	(29.1%)	17.7 (-50.3%)	90	PEF	6.1	(27.4%)	6.1 (-35.2%)	6.1 (-49.1%)
SEC (CPEF	21.2		25.0		35.6	SEC (CPEF	8.3		9.3	11.9
F	BIC	55.1	(+159.7%)	55.1	(+120.8%)	55.1 (+54.7%)	F	BIC	18.5	(+122.6%)	$18.5 \ (+98.6\%)$	18.5 (+56.0%)
			(a) Cli	leWel	o09					(b)	Gov2	

Table 3: Timings in milliseconds for AND queries on ClueWeb09 and Gov2, using query sets TREC 05 and TREC 05. In parentheses we show the relative percentage against CPEF.

Time VS Space tradeoffs by varying reference size



Figure 2: Bits per posting of Gov2 and ClueWeb09 by varying the reference size.

	MIN MID		MAX	MAX		MIN	MID	MAX
PEF	2.94 (+5.60%) 2.9	4 (+7.91%) 2.94 (+10.9	5%)	PEF	4.80 (+2.13%) 4.80 (+3.989	%) 4.80 (+6.25%)
CPEF	2.78 2.7	2	2.65		CPEF	4.70	4.62	4.52
BIC	2.80 (+0.53%) 2.8	0 (+2.74%	a) 2.80 (+5.6	3%)	BIC	4.27 (-9.22%)) 4.27 (-7.589	6) 4.27 (-5.56%)
	(a) Gov	v2				(b) C	lueWeb09	

Table 2: Bits per posting in selected trade-off points.

Always better than PEF (by up to 11%) and better than BIC (by up to 6.25%)



3

Figure 3: Timings for AND queries by varying the reference size on Gov2 and ClueWeb09, using the query set TREC 06.

		1	MIN	Ν	٨ID	MAX				MIN	MID	MAX
35	PEF	14.6	(-17.5%)	14.6	(29.0%)	14.6 (-49.7%)	22	PEF	3.7	(30.4%)	3.7 (-37.5%)	3.7 (-52.1%)
SEC (CPEF	17.7		20.6		29.1	REC (CPEF	5.3		5.9	7.8
F	BIC	41.1	(+131.9%)	41.1	(+99.5%)	41.1 (+41.3%)	F	BIC	10.5	(+96.2%)	10.5 (+76.2%)	10.5 (+35.0%)
90	PEF	17.7	(16.6%)	17.7	(29.1%)	17.7 (-50.3%)	9	PEF	6.1	(27.4%)	6.1 (-35.2%)	6.1 (-49.1%)
SEC (CPEF	21.2		25.0		35.6	SEC (CPEF	8.3		9.3	11.9
F	BIC	55.1	(+159.7%)	55.1	(+120.8%)	55.1 (+54.7%)		BIC	18.5	(+122.6%)	$18.5 \ (+98.6\%)$	$18.5 \ (+56.0\%)$
			(a) Clu	JeWel	09					(b)	Gov2	

Table 3: Timings in milliseconds for AND queries on ClueWeb09 and Gov2, using query sets TREC 05 and TREC 05. In parentheses we show the relative percentage against CPEF.

Time VS Space tradeoffs by varying reference size



Figure 2: Bits per posting of Gov2 and ClueWeb09 by varying the reference size.

	MIN	MID	MAX		MIN	MID	MAX
PEF	2.94 (+5.60%)	2.94 (+7.91%	b) 2.94 (+10.95%)	PEF	4.80 (+2.13%)	4.80 (+3.989	6) 4.80 (+6.25%)
CPEF	2.78	2.72	2.65	CPEF	4.70	4.62	4.52
BIC	2.80 (+0.53%)	2.80 (+2.74%	b) 2.80 (+5.63%)	BIC	4.27 (-9.22%)	4.27 (-7.589	6) 4.27 (-5.56%)
	(a)) Gov2			(b) CI	ueWeb09	

Table 2: Bits per posting in selected trade-off points.

Always better than PEF (by up to 11%) and better than BIC (by up to 6.25%)



3

Figure 3: Timings for AND queries by varying the reference size on Gov2 and ClueWeb09, using the query set TREC 06.

	MIN	MID	MAX			N	ЛIN	MID	MAX
හ PEF	14.6 (-17.5%)	14.6 (-29.0%)	14.6 (-49.7%)	05	PEF	3.7	(—30.4%)	3.7 (-37.5%)	3.7 (-52.1%)
U CPEF	17.7	20.6	29.1	REC	CPEF	5.3		5.9	7.8
БВС	41.1 (+131.9%)	41.1 (+99.5%)	41.1 (+41.3%)	F	BIC	10.5	(+96.2%)	10.5 (+76.2%)	10.5 (+35.0%)
8 PEF	17.7 (-16.6%)	17.7 (-29.1%)	17.7 (-50.3%)	90	PEF	6.1	(27.4%)	6.1 (-35.2%)	6.1 (-49.1%)
U CPEF	21.2	25.0	35.6	REC (CPEF	8.3		9.3	11.9
F BIC	55.1 (+159.7%)	55.1 (+120.8%)	55.1 (+54.7%)	F	BIC	18.5	(+122.6%)	18.5 (+98.6%)	18.5 (+56.0%)
	(a) Clu	eWeb09		_			(b)	Gov2	

Table 3: Timings in milliseconds for AND queries on ClueWeb09 and Gov2, using query sets TREC 05 and TREC 05. In parentheses we show the relative percentage against CPEF.

Time VS Space tradeoffs by varying reference size



Figure 2: Bits per posting of Gov2 and ClueWeb09 by varying the reference size.

	MIN MID		MAX		MIN	MID	MAX
PEF	2.94 (+5.60%) 2.	94 (+7.91%	b) 2.94 (+10.95%)	PEF	4.80 (+2.13%)	4.80 (+3.989	(+6.25%) 4.80
CPEF	2.78 2.	72	2.65	CPEF	4.70	4.62	4.52
BIC	2.80 (+0.53%) 2.	80 (+2.74%	b) 2.80 (+5.63%)	BIC	4.27 (-9.22%)	4.27 (-7.589	6) 4.27 (-5.56%)
	(a) G	ov2			(b) CI	ueWeb09	

Table 2: Bits per posting in selected trade-off points.

Always better than PEF (by up to 11%) and better than BIC (by up to 6.25%)



З

Figure 3: Timings for AND queries by varying the reference size on Gov2 and ClueWeb09, using the query set TREC 06.

	MIN	MID	MAX			Ν	ЛIN	MID	MAX
ഴ PEF	14.6 (-17.5%)	14.6 (-29.0%)	14.6 (-49.7%)	ا ي	PEF	3.7	(30.4%)	3.7 (-37.5%)	3.7 (-52.1
CPEF	17.7	20.6	29.1	REC (CPEF	5.3		5.9	7.8
БІС	41.1 (+131.9%)	41.1 (+99.5%)	41.1 (+41.3%)	F	віс	10.5	(+96.2%)	10.5 (+76.2%)	10.5 (+35.0
e PEF	17.7 (-16.6%)	17.7 (-29.1%)	17.7 (-50.3%)	ا ي	PEF	6.1	(27.4%)	6.1 (-35.2%)	6.1 (-49.1
Ö CPEF	21.2	25.0	35.6	SEC (CPEF	8.3		9.3	11.9
BIC	55.1 (+159.7%)	55.1 (+120.8%)	55.1 (+54.7%)	Ē	віс	18.5	(+122.6%)	18.5 (+98.6%)	18.5 (+56.0
(a) ClueWeb09				(b) Gov2					

Much faster than BIC (103% on average) Slightly slower than PEF (20% on average)

21

Elias-Fano matches the *information theoretic minimum*.

n lg(u/n) + 2n + o(n) bits

Elias-Fano matches the *information theoretic minimum*.

n lg(u/n) + 2n + o(n) bits

- O(1) random access
- O(lg(u/n)) predecessor/successor

Elias-Fano matches the *information theoretic minimum*.

n lg(u/n) + 2n + o(n) bits

- O(1) random access
- O(lg(u/n)) predecessor/successor

Static succinct data structure. NO dynamic updates.

Elias-Fano matches the *information theoretic minimum*.

n lg(u/n) + 2n + o(n) bits

- O(1) random access
- O(lg(u/n)) predecessor/successor

Static succinct data structure. NO dynamic updates.

- **vEB Trees** [van Emde Boas, FOCS 1975]
- x/y-Fast Tries [Willard, IPL 1983]
- Fusion Trees [Fredman and Willard, JCSS 1993]
- **Exponential Search Trees** [Andersson and Thorup, JACM 2007]

- Dynamic
- Most of them take optimal time

Elias-Fano matches the *information theoretic minimum*.

n lg(u/n) + 2n + o(n) bits

- O(1) random access
- O(lg(u/n)) predecessor/successor

Static succinct data structure. NO dynamic updates.

- **vEB Trees** [van Emde Boas, FOCS 1975]
- x/y-Fast Tries [Willard, IPL 1983]
- Fusion Trees [Fredman and Willard, JCSS 1993]
- **Exponential Search Trees** [Andersson and Thorup, JACM 2007]

- Dynamic
- Most of them take optimal time

O(n lg u) bits

(or even worse)

The dynamic dictionary problem consists in representing a set S of n objects so that the following operations are supported.

- insert(x) inserts x in S
- delete(x) deletes x from S
- search(x) checks whether x belongs to S
- minimum() returns the minimum element of S
- maximum() returns the maximum element of S
- predecessor(x) returns $max{y \in S : y < x}$
- successor(x) returns $min\{y \in S : y \ge x\}$

The dynamic dictionary problem consists in representing a set S of n objects so that the following operations are supported.

- insert(x) inserts x in S
- delete(x) deletes x from S
- search(x) checks whether x belongs to S
- minimum() returns the minimum element of S
- maximum() returns the maximum element of S
- predecessor(x) returns $max{y \in S : y < x}$
- successor(x) returns $min\{y \in S : y \ge x\}$

The dynamic dictionary problem consists in representing a set S of n objects so that the following operations are supported.

- insert(x) inserts x in S
- delete(x) deletes x from S
- search(x) checks whether x belongs to S
- minimum() returns the minimum element of S
- maximum() returns the maximum element of S
- predecessor(x) returns max{ $y \in S : y < x$ }
- successor(x) returns $min\{y \in S : y \ge x\}$

The Dynamic List Representation problem

2

[Fredman and Saks, STC 1989]

Given a list S of n sorted integer, support the following operations

- access(i) return the i-th smallest element of S
- insert(x) inserts x in S
- delete(x) deletes x from S

under the assumption that $w \leq \lg^{\gamma} n$ for some γ .

The dynamic dictionary problem consists in representing a set S of n objects so that the following operations are supported.

- insert(x) inserts x in S
- delete(x) deletes x from S
- search(x) checks whether x belongs to S
- minimum() returns the minimum element of S
- maximum() returns the maximum element of S
- predecessor(x) returns max{ $y \in S : y < x$ }
- successor(x) returns $min\{y \in S : y \ge x\}$

The Dynamic List Representation problem

2

[Fredman and Saks, STC 1989]

Given a list S of n sorted integer, support the following operations

- access(i) return the i-th smallest element of S
- insert(x) inserts x in S
- delete(x) deletes x from S

under the assumption that $w \leq \lg^{\gamma} n$ for some γ .

The dynamic dictionary problem consists in representing a set S of n objects so that the following operations are supported.

- insert(x) inserts x in S
- delete(x) deletes x from S
- search(x) checks whether x belongs to S
- minimum() returns the minimum element of S
- maximum() returns the maximum element of S
- predecessor(x) returns max{ $y \in S : y < x$ }
- successor(x) returns $min\{y \in S : y \ge x\}$

The Dynamic List Representation problem

2

[Fredman and Saks, STC 1989]

Given a list S of n sorted integer, support the following operations

- access(i) return the i-th smallest element of S
- insert(x) inserts x in S
- delete(x) deletes x from S

under the assumption that $w \leq \lg^{\gamma} n$ for some γ .

The dynamic dictionary problem consists in representing a set S of n objects so that the following operations are supported.

- insert(x) inserts x in S
- delete(x) deletes x from S
- search(x) checks whether x belongs to S
- minimum() returns the minimum element of S
- maximum() returns the maximum element of S
- predecessor(x) returns max{ $y \in S : y < x$ }
- successor(x) returns $min\{y \in S : y \ge x\}$

[Patrascu and Thorup, STC 2007]

Optimal space/time trade-off for a *static* data structure taking $m = n2^{a}w$ bits, where *a* is the number of bits necessary to represent the mean number of bits per integer, i.e., a = lg(m/n) - lg w

$$\Theta\left(\min\left\{\log_{w} n, \lg\frac{w-\lg n}{a}, \frac{\lg\frac{w}{a}}{\lg\left(\frac{a}{\lg n}\lg\frac{w}{a}\right)}, \frac{\lg\frac{w}{a}}{\lg\left(\lg\frac{w}{a}/\lg\frac{\lg n}{a}\right)}\right\}\right)$$

The Dynamic List Representation problem

2

[Fredman and Saks, STC 1989]

Given a list S of n sorted integer, support the following operations

- access(i) return the i-th smallest element of S
- insert(x) inserts x in S
- delete(x) deletes x from S

under the assumption that $w \leq \lg^{\gamma} n$ for some γ .

The dynamic dictionary problem consists in representing a set S of n objects so that the following operations are supported.

- insert(x) inserts x in S
- delete(x) deletes x from S
- search(x) checks whether x belongs to S
- minimum() returns the minimum element of S
- maximum() returns the maximum element of S
- predecessor(x) returns max{ $y \in S : y < x$ }
- successor(x) returns $min\{y \in S : y \ge x\}$

[Patrascu and Thorup, STC 2007]

Optimal space/time trade-off for a *static* data structure taking $m = n2^{a}w$ bits, where *a* is the number of bits necessary to represent the mean number of bits per integer, i.e., a = lg(m/n) - lg w

$$\Theta\left(\min\left\{\log_{w} n, \lg\frac{w-\lg n}{a}, \frac{\lg\frac{w}{a}}{\lg\left(\frac{a}{\lg n}\lg\frac{w}{a}\right)}, \frac{\lg\frac{w}{a}}{\lg\left(\lg\frac{w}{a}\right)}, \frac{\lg\frac{w}{a}}{\lg\left(\lg\frac{w}{a}/\lg\frac{\lg n}{a}\right)}\right\}\right)$$

The Dynamic List Representation problem

2

[Fredman and Saks, STC 1989]

Given a list S of n sorted integer, support the following operations

- access(i) return the i-th smallest element of S
- insert(x) inserts x in S
- delete(x) deletes x from S

under the assumption that $w \leq \lg^{\gamma} n$ for some γ .

Goals

Goals

n lg(u/n) + 2n + o(n) bits

Goals







- 1. Extend the *static* Elias-Fano representation to support predecessor and successor queries in optimal worst-case O(lg lg n) time.
- 2. Maintain S in a *fully dynamic fashion*, supporting in optimal worst-case time all the operations defined in the Dynamic Dictionary and Dynamic List Representation problems.

• optimal time/space trade-off for successor search [Patrascu and Thorup, STC 2007]

• y-fast tries [Willard, IPL 1983]

▶ **Theorem 1.** There exists a data structure representing an ordered set S(n, u) of n integers drawn from a polynomial universe of size $u = n^{\gamma}$, for any $\gamma = \Theta(1)$, that takes $\mathsf{EF}(S(n, u)) + o(n)$ bits of space and supports Access in $\mathcal{O}(1)$ worst-case and Predecessor/Successor queries in optimal $\mathcal{O}(\min\{1 + \log \frac{u}{n}, \log \log n\})$ worst-case time.

• optimal time/space trade-off for successor search [Patrascu and Thorup, STC 2007]

• y-fast tries [Willard, IPL 1983]

▶ Theorem 1. There exists a data structure representing an ordered set S(n, u) of n integers drawn from a polynomial universe of size $u = n^{\gamma}$, for any $\gamma = \Theta(1)$, that takes $\mathsf{EF}(S(n, u)) + o(n)$ bits of space and supports Access in $\mathcal{O}(1)$ worst-case and Predecessor/Successor queries in optimal $\mathcal{O}(\min\{1 + \log \frac{u}{n}, \log \log n\})$ worst-case time.

Idea: divide the sequence into **blocks** and use a **y-fast trie** to index the blocks.

Results - Dynamic Elias-Fano

- optimal time/space trade-off for successor search [Patrascu and Thorup, STC 2007]
- **y-fast tries** [Willard, IPL 1983]
- **dynamic prefix-sum data structure** [Bille *et al.*, arXiv preprint 2015]

▶ Lemma 4. The total order of the blocks of C can be maintained by using a data structure that takes $O(\operatorname{polylog} n \cdot \log \log n)$ bits of space and supports the following operations in $O(\log \log n)$ worst-case time: Search(x) which returns a pointer to the block containing the integer x; Access(i) which returns the i-th integer of the total order; Insert/Delete of a block.

▶ Theorem 3. There exists a data structure representing an ordered set S(n, u) of n integers drawn from a polynomial universe of size $u = n^{\gamma}$, for any $\gamma = \Theta(1)$, that takes EF(S(n, u)) + o(n) bits of space and supports: Access in $O(\log n / \log \log n)$ worst-case; Insert/Delete in $O(\log n / \log \log n)$ amortized; Minimum/Maximum in O(1) and Predecessor/Successor queries in $O(\min\{1 + \log \frac{u}{n}, \log \log n\})$ worst-case time. These time bounds are optimal.

Results - Dynamic Elias-Fano

- optimal time/space trade-off for successor search [Patrascu and Thorup, STC 2007]
- **y-fast tries** [Willard, IPL 1983]
- **dynamic prefix-sum data structure** [Bille *et al.*, arXiv preprint 2015]

▶ Lemma 4. The total order of the blocks of C can be maintained by using a data structure that takes $O(\operatorname{polylog} n \cdot \log \log n)$ bits of space and supports the following operations in $O(\log \log n)$ worst-case time: Search(x) which returns a pointer to the block containing the integer x; Access(i) which returns the i-th integer of the total order; Insert/Delete of a block.

▶ Theorem 3. There exists a data structure representing an ordered set S(n, u) of n integers drawn from a polynomial universe of size $u = n^{\gamma}$, for any $\gamma = \Theta(1)$, that takes EF(S(n, u)) + o(n) bits of space and supports: Access in $O(\log n / \log \log n)$ worst-case; Insert/Delete in $O(\log n / \log \log n)$ amortized; Minimum/Maximum in O(1) and Predecessor/Successor queries in $O(\min\{1 + \log \frac{u}{n}, \log \log n\})$ worst-case time. These time bounds are optimal.

Idea: use a 2-level indexing data structure.

- First level indexes **blocks** using a **y-fast trie** and the **dynamic prefix-sum** data structure by Bille *et al*.
- Second level indexes **mini blocks** using the data structure of the Lemma.

Strings of at most N words. N typically ranges from 1 to 5.

Strings of at most N words. N typically ranges from 1 to 5.

Different algorithms devised to solve the same problem often differ dramatically in their efficiency. These differences can be much more significant than differences due to hardware and software.

As an example, in Chapter 2, we will see two algorithms for sorting. The first, known as *insertion sort*, takes time roughly equal to c_1n^2 to sort n items, where c_1 is a constant that does not depend on n. That is, it takes time roughly proportional to n^2 . The second, *merge sort*, takes time roughly equal to $c_2 n \lg n$, where $\lg n$ stands for $\log_2 n$ and c_2 is another constant that also does not depend on n. Insertion sort typically has a smaller constant factor than merge sort, so that $c_1 < c_2$. We shall see that the constant factors can have far less of an impact on the running time than the dependence on the input size n. Let's write insertion sort's running time as $c_1 n \cdot n$ and merge sort's running time as $c_2 n \cdot \lg n$. Then we see that where insertion sort has a factor of n in its running time, merge sort has a factor of $\lg n$, which is much smaller. (For example, when n = 1000, $\lg n$ is approximately 10, and when n equals one million, $\lg n$ is approximately only 20.) Although insertion sort usually runs faster than merge sort for small input sizes, once the input size n becomes large enough, merge sort's advantage of $\lg n$ vs. n will more than compensate for the difference in constant factors. No matter how much smaller c_1 is than c_2 , there will always be a crossover point beyond which merge sort is faster.

Strings of at most N words.

N typically ranges from 1 to 5.

Different algorithms devised to solve the same problem often differ dramatically in their efficiency. These differences can be much more significant than differences due to hardware and software.

As an example, in Chapter 2, we will see two algorithms for sorting. The first, known as *insertion sort*, takes time roughly equal to c_1n^2 to sort n items, where c_1 is a constant that does not depend on n. That is, it takes time roughly proportional to n^2 . The second, *merge sort*, takes time roughly equal to $c_2 n \lg n$, where $\lg n$ stands for $\log_2 n$ and c_2 is another constant that also does not depend on n. Insertion sort typically has a smaller constant factor than merge sort, so that $c_1 < c_2$. We shall see that the constant factors can have far less of an impact on the running time than the dependence on the input size n. Let's write insertion sort's running time as $c_1 n \cdot n$ and merge sort's running time as $c_2 n \cdot \lg n$. Then we see that where insertion sort has a factor of n in its running time, merge sort has a factor of $\lg n$, which is much smaller. (For example, when n = 1000, $\lg n$ is approximately 10, and when n equals one million, $\lg n$ is approximately only 20.) Although insertion sort usually runs faster than merge sort for small input sizes, once the input size nbecomes large enough, merge sort's advantage of $\lg n$ vs. n will more than compensate for the difference in constant factors. No matter how much smaller c_1 is than c_2 , there will always be a crossover point beyond which merge sort is faster.

N = 1

Strings of at most N words.

N typically ranges from 1 to 5.

Different algorithms devised to solve the same problem often differ dramatically in their efficiency. These differences can be much more significant than differences due to hardware and software.

As an example, in Chapter 2, we will see two algorithms for sorting. The first, known as *insertion sort*, takes time roughly equal to c_1n^2 to sort n items, where c_1 is a constant that does not depend on n. That is, it takes time roughly proportional to n^2 . The second, *merge sort*, takes time roughly equal to $c_2 n \lg n$, where $\lg n$ stands for $\log_2 n$ and c_2 is another constant that also does not depend on n. Insertion sort typically has a smaller constant factor than merge sort, so that $c_1 < c_2$. We shall see that the constant factors can have far less of an impact on the running time than the dependence on the input size n. Let's write insertion sort's running time as $c_1 n \cdot n$ and merge sort's running time as $c_2 n \cdot \lg n$. Then we see that where insertion sort has a factor of n in its running time, merge sort has a factor of $\lg n$, which is much smaller. (For example, when n = 1000, $\lg n$ is approximately 10, and when n equals one million, $\lg n$ is approximately only 20.) Although insertion sort usually runs faster than merge sort for small input sizes, once the input size n becomes large enough, merge sort's advantage of $\lg n$ vs. n will more than compensate for the difference in constant factors. No matter how much smaller c_1 is than c_2 , there will always be a crossover point beyond which merge sort is faster.

N = 1

Strings of at most N words. N typically ranges from 1 to 5.

Different algorithms devised to solve the same problem often differ dramatically in their efficiency. These differences can be much more significant than differences due to hardware and software.

As an example, in Chapter 2, we will see two algorithms for sorting. The first, known as *insertion sort*, takes time roughly equal to c_1n^2 to sort n items, where c_1 is a constant that does not depend on n. That is, it takes time roughly proportional to n^2 . The second, *merge sort*, takes time roughly equal to $c_2 n \lg n$, where $\lg n$ stands for $\log_2 n$ and c_2 is another constant that also does not depend on n. Insertion sort typically has a smaller constant factor than merge sort, so that $c_1 < c_2$. We shall see that the constant factors can have far less of an impact on the running time than the dependence on the input size n. Let's write insertion sort's running time as $c_1 n \cdot n$ and merge sort's running time as $c_2 n \cdot \lg n$. Then we see that where insertion sort has a factor of n in its running time, merge sort has a factor of $\lg n$, which is much smaller. (For example, when n = 1000, $\lg n$ is approximately 10, and when n equals one million, lg n is approximately only 20.) Although insertion sort usually runs faster than merge sort for small input sizes, once the input size n becomes large enough, merge sort's advantage of $\lg n$ vs. n will more than compensate for the difference in constant factors. No matter how much smaller c_1 is than c_2 , there will always be a crossover point beyond which merge sort is faster.

N = 1

different

Strings of at most N words. N typically ranges from 1 to 5.

Different algorithms devised to solve the same problem often differ dramatically in their efficiency. These differences can be much more significant than differences due to hardware and software.

As an example, in Chapter 2, we will see two algorithms for sorting. The first, known as *insertion sort*, takes time roughly equal to c_1n^2 to sort n items, where c_1 is a constant that does not depend on n. That is, it takes time roughly proportional to n^2 . The second, *merge sort*, takes time roughly equal to $c_2 n \lg n$, where $\lg n$ stands for $\log_2 n$ and c_2 is another constant that also does not depend on n. Insertion sort typically has a smaller constant factor than merge sort, so that $c_1 < c_2$. We shall see that the constant factors can have far less of an impact on the running time than the dependence on the input size n. Let's write insertion sort's running time as $c_1 n \cdot n$ and merge sort's running time as $c_2 n \cdot \lg n$. Then we see that where insertion sort has a factor of n in its running time, merge sort has a factor of $\lg n$, which is much smaller. (For example, when n = 1000, $\lg n$ is approximately 10, and when n equals one million, lg n is approximately only 20.) Although insertion sort usually runs faster than merge sort for small input sizes, once the input size n becomes large enough, merge sort's advantage of $\lg n$ vs. n will more than compensate for the difference in constant factors. No matter how much smaller c_1 is than c_2 , there will always be a crossover point beyond which merge sort is faster.

N = 1

different

Strings of at most N words. N typically ranges from 1 to 5.

Different algorithms devised to solve the same problem often differ dramatically in their efficiency. These differences can be much more significant than differences due to hardware and software.

As an example, in Chapter 2, we will see two algorithms for sorting. The first, known as *insertion sort*, takes time roughly equal to c_1n^2 to sort n items, where c_1 is a constant that does not depend on n. That is, it takes time roughly proportional to n^2 . The second, *merge sort*, takes time roughly equal to $c_2 n \lg n$, where $\lg n$ stands for $\log_2 n$ and c_2 is another constant that also does not depend on n. Insertion sort typically has a smaller constant factor than merge sort, so that $c_1 < c_2$. We shall see that the constant factors can have far less of an impact on the running time than the dependence on the input size n. Let's write insertion sort's running time as $c_1 n \cdot n$ and merge sort's running time as $c_2 n \cdot \lg n$. Then we see that where insertion sort has a factor of n in its running time, merge sort has a factor of $\lg n$, which is much smaller. (For example, when n = 1000, $\lg n$ is approximately 10, and when n equals one million, $\lg n$ is approximately only 20.) Although insertion sort usually runs faster than merge sort for small input sizes, once the input size n becomes large enough, merge sort's advantage of $\lg n$ vs. n will more than compensate for the difference in constant factors. No matter how much smaller c_1 is than c_2 , there will always be a crossover point beyond which merge sort is faster.

N = 1

different algorithms

Strings of at most N words. N typically ranges from 1 to 5.

Different algorithms devised to solve the same problem often differ dramatically in their efficiency. These differences can be much more significant than differences due to hardware and software.

As an example, in Chapter 2, we will see two algorithms for sorting. The first, known as *insertion sort*, takes time roughly equal to c_1n^2 to sort n items, where c_1 is a constant that does not depend on n. That is, it takes time roughly proportional to n^2 . The second, *merge sort*, takes time roughly equal to $c_2 n \lg n$, where $\lg n$ stands for $\log_2 n$ and c_2 is another constant that also does not depend on n. Insertion sort typically has a smaller constant factor than merge sort, so that $c_1 < c_2$. We shall see that the constant factors can have far less of an impact on the running time than the dependence on the input size n. Let's write insertion sort's running time as $c_1 n \cdot n$ and merge sort's running time as $c_2 n \cdot \lg n$. Then we see that where insertion sort has a factor of n in its running time, merge sort has a factor of $\lg n$, which is much smaller. (For example, when n = 1000, $\lg n$ is approximately 10, and when n equals one million, $\lg n$ is approximately only 20.) Although insertion sort usually runs faster than merge sort for small input sizes, once the input size n becomes large enough, merge sort's advantage of $\lg n$ vs. n will more than compensate for the difference in constant factors. No matter how much smaller c_1 is than c_2 , there will always be a crossover point beyond which merge sort is faster.

N = 1

different algorithms

Strings of at most N words. N typically ranges from 1 to 5.

Different algorithms devised to solve the same problem often differ dramatically in their efficiency. These differences can be much more significant than differences due to hardware and software.

As an example, in Chapter 2, we will see two algorithms for sorting. The first, known as *insertion sort*, takes time roughly equal to c_1n^2 to sort n items, where c_1 is a constant that does not depend on n. That is, it takes time roughly proportional to n^2 . The second, *merge sort*, takes time roughly equal to $c_2 n \lg n$, where $\lg n$ stands for $\log_2 n$ and c_2 is another constant that also does not depend on n. Insertion sort typically has a smaller constant factor than merge sort, so that $c_1 < c_2$. We shall see that the constant factors can have far less of an impact on the running time than the dependence on the input size n. Let's write insertion sort's running time as $c_1 n \cdot n$ and merge sort's running time as $c_2 n \cdot \lg n$. Then we see that where insertion sort has a factor of n in its running time, merge sort has a factor of $\lg n$, which is much smaller. (For example, when n = 1000, $\lg n$ is approximately 10, and when n equals one million, lg n is approximately only 20.) Although insertion sort usually runs faster than merge sort for small input sizes, once the input size n becomes large enough, merge sort's advantage of $\lg n$ vs. n will more than compensate for the difference in constant factors. No matter how much smaller c_1 is than c_2 , there will always be a crossover point beyond which merge sort is faster.

N = 1

different algorithms devised

Strings of at most N words. N typically ranges from 1 to 5.

Different algorithms devised to solve the same problem often differ dramatically in their efficiency. These differences can be much more significant than differences due to hardware and software.

As an example, in Chapter 2, we will see two algorithms for sorting. The first, known as *insertion sort*, takes time roughly equal to c_1n^2 to sort n items, where c_1 is a constant that does not depend on n. That is, it takes time roughly proportional to n^2 . The second, *merge sort*, takes time roughly equal to $c_2 n \lg n$, where $\lg n$ stands for $\log_2 n$ and c_2 is another constant that also does not depend on n. Insertion sort typically has a smaller constant factor than merge sort, so that $c_1 < c_2$. We shall see that the constant factors can have far less of an impact on the running time than the dependence on the input size n. Let's write insertion sort's running time as $c_1 n \cdot n$ and merge sort's running time as $c_2 n \cdot \lg n$. Then we see that where insertion sort has a factor of n in its running time, merge sort has a factor of $\lg n$, which is much smaller. (For example, when n = 1000, $\lg n$ is approximately 10, and when n equals one million, lg n is approximately only 20.) Although insertion sort usually runs faster than merge sort for small input sizes, once the input size nbecomes large enough, merge sort's advantage of $\lg n$ vs. n will more than compensate for the difference in constant factors. No matter how much smaller c_1 is than c_2 , there will always be a crossover point beyond which merge sort is faster.

N = 1

different algorithms devised

Strings of at most N words. N typically ranges from 1 to 5.

Different algorithms devised to solve the same problem often differ dramatically in their efficiency. These differences can be much more significant than differences due to hardware and software.

As an example, in Chapter 2, we will see two algorithms for sorting. The first, known as *insertion sort*, takes time roughly equal to c_1n^2 to sort n items, where c_1 is a constant that does not depend on n. That is, it takes time roughly proportional to n^2 . The second, *merge sort*, takes time roughly equal to $c_2 n \lg n$, where $\lg n$ stands for $\log_2 n$ and c_2 is another constant that also does not depend on n. Insertion sort typically has a smaller constant factor than merge sort, so that $c_1 < c_2$. We shall see that the constant factors can have far less of an impact on the running time than the dependence on the input size n. Let's write insertion sort's running time as $c_1 n \cdot n$ and merge sort's running time as $c_2 n \cdot \lg n$. Then we see that where insertion sort has a factor of n in its running time, merge sort has a factor of $\lg n$, which is much smaller. (For example, when n = 1000, $\lg n$ is approximately 10, and when n equals one million, $\lg n$ is approximately only 20.) Although insertion sort usually runs faster than merge sort for small input sizes, once the input size nbecomes large enough, merge sort's advantage of $\lg n$ vs. n will more than compensate for the difference in constant factors. No matter how much smaller c_1 is than c_2 , there will always be a crossover point beyond which merge sort is faster.

N = 1

different algorithms devised
Strings of at most N words. N typically ranges from 1 to 5.

Different algorithms devised to solve the same problem often differ dramatically in their efficiency. These differences can be much more significant than differences due to hardware and software.

As an example, in Chapter 2, we will see two algorithms for sorting. The first, known as *insertion sort*, takes time roughly equal to c_1n^2 to sort n items, where c_1 is a constant that does not depend on n. That is, it takes time roughly proportional to n^2 . The second, *merge sort*, takes time roughly equal to $c_2 n \lg n$, where $\lg n$ stands for $\log_2 n$ and c_2 is another constant that also does not depend on n. Insertion sort typically has a smaller constant factor than merge sort, so that $c_1 < c_2$. We shall see that the constant factors can have far less of an impact on the running time than the dependence on the input size n. Let's write insertion sort's running time as $c_1 n \cdot n$ and merge sort's running time as $c_2 n \cdot \lg n$. Then we see that where insertion sort has a factor of n in its running time, merge sort has a factor of $\lg n$, which is much smaller. (For example, when n = 1000, $\lg n$ is approximately 10, and when n equals one million, $\lg n$ is approximately only 20.) Although insertion sort usually runs faster than merge sort for small input sizes, once the input size nbecomes large enough, merge sort's advantage of $\lg n$ vs. n will more than compensate for the difference in constant factors. No matter how much smaller c_1 is than c_2 , there will always be a crossover point beyond which merge sort is faster.

N = 1

different algorithms devised

Strings of at most N words. N typically ranges from 1 to 5.

Different algorithms devised to solve the same problem often differ dramatically in their efficiency. These differences can be much more significant than differences due to hardware and software.

As an example, in Chapter 2, we will see two algorithms for sorting. The first, known as *insertion sort*, takes time roughly equal to c_1n^2 to sort n items, where c_1 is a constant that does not depend on n. That is, it takes time roughly proportional to n^2 . The second, *merge sort*, takes time roughly equal to $c_2 n \lg n$, where $\lg n$ stands for $\log_2 n$ and c_2 is another constant that also does not depend on n. Insertion sort typically has a smaller constant factor than merge sort, so that $c_1 < c_2$. We shall see that the constant factors can have far less of an impact on the running time than the dependence on the input size n. Let's write insertion sort's running time as $c_1 n \cdot n$ and merge sort's running time as $c_2 n \cdot \lg n$. Then we see that where insertion sort has a factor of n in its running time, merge sort has a factor of $\lg n$, which is much smaller. (For example, when n = 1000, $\lg n$ is approximately 10, and when n equals one million, $\lg n$ is approximately only 20.) Although insertion sort usually runs faster than merge sort for small input sizes, once the input size nbecomes large enough, merge sort's advantage of $\lg n$ vs. n will more than compensate for the difference in constant factors. No matter how much smaller c_1 is than c_2 , there will always be a crossover point beyond which merge sort is faster.

N = 1

different algorithms devised

Strings of at most N words. N typically ranges from 1 to 5.

Different algorithms devised to solve the same problem often differ dramatically in their efficiency. These differences can be much more significant than differences due to hardware and software.

As an example, in Chapter 2, we will see two algorithms for sorting. The first, known as *insertion sort*, takes time roughly equal to c_1n^2 to sort n items, where c_1 is a constant that does not depend on n. That is, it takes time roughly proportional to n^2 . The second, *merge sort*, takes time roughly equal to $c_2 n \lg n$, where $\lg n$ stands for $\log_2 n$ and c_2 is another constant that also does not depend on n. Insertion sort typically has a smaller constant factor than merge sort, so that $c_1 < c_2$. We shall see that the constant factors can have far less of an impact on the running time than the dependence on the input size n. Let's write insertion sort's running time as $c_1 n \cdot n$ and merge sort's running time as $c_2 n \cdot \lg n$. Then we see that where insertion sort has a factor of n in its running time, merge sort has a factor of $\lg n$, which is much smaller. (For example, when n = 1000, $\lg n$ is approximately 10, and when n equals one million, $\lg n$ is approximately only 20.) Although insertion sort usually runs faster than merge sort for small input sizes, once the input size nbecomes large enough, merge sort's advantage of $\lg n$ vs. n will more than compensate for the difference in constant factors. No matter how much smaller c_1 is than c_2 , there will always be a crossover point beyond which merge sort is faster.

N = 1 N = 2 different

algorithms devised

Strings of at most N words. N typically ranges from 1 to 5.

Different algorithms devised to solve the same problem often differ dramatically in their efficiency. These differences can be much more significant than differences due to hardware and software.

As an example, in Chapter 2, we will see two algorithms for sorting. The first, known as *insertion sort*, takes time roughly equal to c_1n^2 to sort n items, where c_1 is a constant that does not depend on n. That is, it takes time roughly proportional to n^2 . The second, *merge sort*, takes time roughly equal to $c_2 n \lg n$, where $\lg n$ stands for $\log_2 n$ and c_2 is another constant that also does not depend on n. Insertion sort typically has a smaller constant factor than merge sort, so that $c_1 < c_2$. We shall see that the constant factors can have far less of an impact on the running time than the dependence on the input size n. Let's write insertion sort's running time as $c_1 n \cdot n$ and merge sort's running time as $c_2 n \cdot \lg n$. Then we see that where insertion sort has a factor of n in its running time, merge sort has a factor of $\lg n$, which is much smaller. (For example, when n = 1000, $\lg n$ is approximately 10, and when n equals one million, $\lg n$ is approximately only 20.) Although insertion sort usually runs faster than merge sort for small input sizes, once the input size nbecomes large enough, merge sort's advantage of $\lg n$ vs. n will more than compensate for the difference in constant factors. No matter how much smaller c_1 is than c_2 , there will always be a crossover point beyond which merge sort is faster.

N = 1 N = 2 different algorithms

devised

Strings of at most N words. N typically ranges from 1 to 5.

Different algorithms devised to solve the same problem often differ dramatically in their efficiency. These differences can be much more significant than differences due to hardware and software.

As an example, in Chapter 2, we will see two algorithms for sorting. The first, known as *insertion sort*, takes time roughly equal to c_1n^2 to sort n items, where c_1 is a constant that does not depend on n. That is, it takes time roughly proportional to n^2 . The second, *merge sort*, takes time roughly equal to $c_2 n \lg n$, where $\lg n$ stands for $\log_2 n$ and c_2 is another constant that also does not depend on n. Insertion sort typically has a smaller constant factor than merge sort, so that $c_1 < c_2$. We shall see that the constant factors can have far less of an impact on the running time than the dependence on the input size n. Let's write insertion sort's running time as $c_1 n \cdot n$ and merge sort's running time as $c_2 n \cdot \lg n$. Then we see that where insertion sort has a factor of n in its running time, merge sort has a factor of $\lg n$, which is much smaller. (For example, when n = 1000, $\lg n$ is approximately 10, and when n equals one million, $\lg n$ is approximately only 20.) Although insertion sort usually runs faster than merge sort for small input sizes, once the input size nbecomes large enough, merge sort's advantage of $\lg n$ vs. n will more than compensate for the difference in constant factors. No matter how much smaller c_1 is than c_2 , there will always be a crossover point beyond which merge sort is faster.

N = 1 N = 2

different algorithms devised different algorithms

Strings of at most N words. N typically ranges from 1 to 5.

Different algorithms devised to solve the same problem often differ dramatically in their efficiency. These differences can be much more significant than differences due to hardware and software.

As an example, in Chapter 2, we will see two algorithms for sorting. The first, known as *insertion sort*, takes time roughly equal to c_1n^2 to sort n items, where c_1 is a constant that does not depend on n. That is, it takes time roughly proportional to n^2 . The second, *merge sort*, takes time roughly equal to $c_2 n \lg n$, where $\lg n$ stands for $\log_2 n$ and c_2 is another constant that also does not depend on n. Insertion sort typically has a smaller constant factor than merge sort, so that $c_1 < c_2$. We shall see that the constant factors can have far less of an impact on the running time than the dependence on the input size n. Let's write insertion sort's running time as $c_1 n \cdot n$ and merge sort's running time as $c_2 n \cdot \lg n$. Then we see that where insertion sort has a factor of n in its running time, merge sort has a factor of $\lg n$, which is much smaller. (For example, when n = 1000, $\lg n$ is approximately 10, and when n equals one million, lg n is approximately only 20.) Although insertion sort usually runs faster than merge sort for small input sizes, once the input size n becomes large enough, merge sort's advantage of $\lg n$ vs. n will more than compensate for the difference in constant factors. No matter how much smaller c_1 is than c_2 , there will always be a crossover point beyond which merge sort is faster.

N = 2

different algorithms

different algorithms devised

N = 1

Strings of at most N words. N typically ranges from 1 to 5.

Different algorithms devised to solve the same problem often differ dramatically in their efficiency. These differences can be much more significant than differences due to hardware and software.

As an example, in Chapter 2, we will see two algorithms for sorting. The first, known as *insertion sort*, takes time roughly equal to c_1n^2 to sort n items, where c_1 is a constant that does not depend on n. That is, it takes time roughly proportional to n^2 . The second, *merge sort*, takes time roughly equal to $c_2 n \lg n$, where $\lg n$ stands for $\log_2 n$ and c_2 is another constant that also does not depend on n. Insertion sort typically has a smaller constant factor than merge sort, so that $c_1 < c_2$. We shall see that the constant factors can have far less of an impact on the running time than the dependence on the input size n. Let's write insertion sort's running time as $c_1 n \cdot n$ and merge sort's running time as $c_2 n \cdot \lg n$. Then we see that where insertion sort has a factor of n in its running time, merge sort has a factor of $\lg n$, which is much smaller. (For example, when n = 1000, $\lg n$ is approximately 10, and when n equals one million, lg n is approximately only 20.) Although insertion sort usually runs faster than merge sort for small input sizes, once the input size nbecomes large enough, merge sort's advantage of $\lg n$ vs. n will more than compensate for the difference in constant factors. No matter how much smaller c_1 is than c_2 , there will always be a crossover point beyond which merge sort is faster.

N = 1

different algorithms devised N = 2

different algorithms algorithms devised

Strings of at most N words. N typically ranges from 1 to 5.

Different algorithms devised to solve the same problem often differ dramatically in their efficiency. These differences can be much more significant than differences due to hardware and software.

As an example, in Chapter 2, we will see two algorithms for sorting. The first, known as *insertion sort*, takes time roughly equal to c_1n^2 to sort n items, where c_1 is a constant that does not depend on n. That is, it takes time roughly proportional to n^2 . The second, *merge sort*, takes time roughly equal to $c_2 n \lg n$, where $\lg n$ stands for $\log_2 n$ and c_2 is another constant that also does not depend on n. Insertion sort typically has a smaller constant factor than merge sort, so that $c_1 < c_2$. We shall see that the constant factors can have far less of an impact on the running time than the dependence on the input size n. Let's write insertion sort's running time as $c_1 n \cdot n$ and merge sort's running time as $c_2 n \cdot \lg n$. Then we see that where insertion sort has a factor of n in its running time, merge sort has a factor of $\lg n$, which is much smaller. (For example, when n = 1000, $\lg n$ is approximately 10, and when n equals one million, $\lg n$ is approximately only 20.) Although insertion sort usually runs faster than merge sort for small input sizes, once the input size nbecomes large enough, merge sort's advantage of $\lg n$ vs. n will more than compensate for the difference in constant factors. No matter how much smaller c_1 is than c_2 , there will always be a crossover point beyond which merge sort is faster.

N = 1

different algorithms devised N = 2

different algorithms algorithms devised

Strings of at most N words. N typically ranges from 1 to 5.

Different algorithms devised to solve the same problem often differ dramatically in their efficiency. These differences can be much more significant than differences due to hardware and software.

As an example, in Chapter 2, we will see two algorithms for sorting. The first, known as *insertion sort*, takes time roughly equal to c_1n^2 to sort n items, where c_1 is a constant that does not depend on n. That is, it takes time roughly proportional to n^2 . The second, *merge sort*, takes time roughly equal to $c_2 n \lg n$, where $\lg n$ stands for $\log_2 n$ and c_2 is another constant that also does not depend on n. Insertion sort typically has a smaller constant factor than merge sort, so that $c_1 < c_2$. We shall see that the constant factors can have far less of an impact on the running time than the dependence on the input size n. Let's write insertion sort's running time as $c_1 n \cdot n$ and merge sort's running time as $c_2 n \cdot \lg n$. Then we see that where insertion sort has a factor of n in its running time, merge sort has a factor of $\lg n$, which is much smaller. (For example, when n = 1000, $\lg n$ is approximately 10, and when n equals one million, $\lg n$ is approximately only 20.) Although insertion sort usually runs faster than merge sort for small input sizes, once the input size nbecomes large enough, merge sort's advantage of $\lg n$ vs. n will more than compensate for the difference in constant factors. No matter how much smaller c_1 is than c_2 , there will always be a crossover point beyond which merge sort is faster.

N = 1

different algorithms devised different algorithms algorithms devised

devised to

N = 2

Strings of at most N words. N typically ranges from 1 to 5.

Different algorithms devised to solve the same problem often differ dramatically in their efficiency. These differences can be much more significant than differences due to hardware and software.

As an example, in Chapter 2, we will see two algorithms for sorting. The first, known as *insertion sort*, takes time roughly equal to c_1n^2 to sort n items, where c_1 is a constant that does not depend on n. That is, it takes time roughly proportional to n^2 . The second, *merge sort*, takes time roughly equal to $c_2 n \lg n$, where $\lg n$ stands for $\log_2 n$ and c_2 is another constant that also does not depend on n. Insertion sort typically has a smaller constant factor than merge sort, so that $c_1 < c_2$. We shall see that the constant factors can have far less of an impact on the running time than the dependence on the input size n. Let's write insertion sort's running time as $c_1 n \cdot n$ and merge sort's running time as $c_2 n \cdot \lg n$. Then we see that where insertion sort has a factor of n in its running time, merge sort has a factor of $\lg n$, which is much smaller. (For example, when n = 1000, $\lg n$ is approximately 10, and when n equals one million, $\lg n$ is approximately only 20.) Although insertion sort usually runs faster than merge sort for small input sizes, once the input size nbecomes large enough, merge sort's advantage of $\lg n$ vs. n will more than compensate for the difference in constant factors. No matter how much smaller c_1 is than c_2 , there will always be a crossover point beyond which merge sort is faster.

N = 1

different algorithms devised different algorithms algorithms devised

devised to

N = 2

Strings of at most N words. N typically ranges from 1 to 5.

Different algorithms devised to solve the same problem often differ dramatically in their efficiency. These differences can be much more significant than differences due to hardware and software.

As an example, in Chapter 2, we will see two algorithms for sorting. The first, known as *insertion sort*, takes time roughly equal to c_1n^2 to sort n items, where c_1 is a constant that does not depend on n. That is, it takes time roughly proportional to n^2 . The second, *merge sort*, takes time roughly equal to $c_2 n \lg n$, where $\lg n$ stands for $\log_2 n$ and c_2 is another constant that also does not depend on n. Insertion sort typically has a smaller constant factor than merge sort, so that $c_1 < c_2$. We shall see that the constant factors can have far less of an impact on the running time than the dependence on the input size n. Let's write insertion sort's running time as $c_1 n \cdot n$ and merge sort's running time as $c_2 n \cdot \lg n$. Then we see that where insertion sort has a factor of n in its running time, merge sort has a factor of $\lg n$, which is much smaller. (For example, when n = 1000, $\lg n$ is approximately 10, and when n equals one million, $\lg n$ is approximately only 20.) Although insertion sort usually runs faster than merge sort for small input sizes, once the input size nbecomes large enough, merge sort's advantage of $\lg n$ vs. n will more than compensate for the difference in constant factors. No matter how much smaller c_1 is than c_2 , there will always be a crossover point beyond which merge sort is faster.

N = 1

N = 2

different algorithms algorithms devised devised to

Ν	number of grams
1	8761
2	38900
3	61516
4	70186
5	73187

Strings of at most N words. N typically ranges from 1 to 5.



Strings of at most N words. N typically ranges from 1 to 5.



Google Books

~6% of the books ever published

Strings of at most N words. N typically ranges from 1 to 5.



Google Books

~6% of the books ever published

Ν	number of grams
1	24,359,473
2	667,284,771
3	7,397,041,901
4	1,644,807,896
5	1,415,355,596

More than 11 billion grams.

2

Word prediction.

2

Word prediction.

space and time-efficient



2

Word prediction.

space and time-efficient

?

context









The latest news from Research at Google

All Our N-gram are Belong to You

Thursday, August 03, 2006

Posted by Alex Franz and Thorsten Brants, Google Machine Translation Team

Here at Google Research we have been using word n-gram models for a variety of R&D projects, such as statistical machine translation, speech recognition, spelling correction, entity detection, information extraction, and others. While such models have usually been estimated from training corpora containing at most a few billion words, we have been harnessing the vast power of Google's datacenters and distributed processing infrastructure to process larger and larger training corpora. We found that there's no data like more data, and scaled up the size of our data by one order of magnitude, and then another, and then one more - resulting in a training corpus of *one trillion words* from public Web pages.



The latest news from Research at Google

All Our N-gram are Belong to You

Thursday, August 03, 2006

Posted by Alex Franz and Thorsten Brants, Google Machine Translation Team

Here at Google Research we have been using word n-gram models for a variety of R&D projects, such as statistical machine translation, speech recognition, spelling correction, entity detection, information extraction, and others. While such models have usually been estimated from training corpora containing at most a few billion words, we have been harnessing the vast power of Google's datacenters and distributed processing infrastructure to process larger and larger training corpora. We found that there's no data like more data, and scaled up the size of our data by one order of magnitude, and then another, and then one more - resulting in a training corpus of *one trillion words* from public Web pages.



The latest news from Research at Google



All Our N-gram are Belong to You

Thursday, August 03, 2006

Posted by Alex Franz and Thorsten Brants, Google Machine Translation Team

Here at Google Research we have been using word n-gram models for a variety of R&D projects, such as statistical machine translation, speech recognition, spelling correction, entity detection, information extraction, and others. While such models have usually been estimated from training corpora containing at most a few billion words, we have been harnessing the vast power of Google's datacenters and distributed processing infrastructure to process larger and larger training corpora. We found that there's no data like more data, and scaled up the size of our data by one order of magnitude, and then another, and then one more - resulting in a training corpus of *one trillion words* from public Web pages.

What can I help you with?



What can I help you with?



Store massive N-grams datasets such that given a pattern, we can return its frequency count at light speed. Store massive N-grams datasets such that given a pattern, we can return its frequency count at light speed.

Efficient map.

Store massive N-grams datasets such that given a pattern, we can return its frequency count at light speed.

Efficient map.

Data Structures + Data Compression - Faster Algorithms

5

Open-addressing VS **Tries**

5

Open-addressing VS **Tries**

10100101	24
10001011	24
00001010	582
11011110	24
00010101	582
01010011	36352

5

Open-addressing VS **Tries**






















Hash-based VS Trie-based

10100101	24
10001011	24
00001010	582
11011110	24
00010101	582
01010011	36352



VS

Hash-based

Open addressing?

24	10100101
24	10001011
582	00001010
24	11011110
582	00010101
36352	01010011



VS

Hash-based

Open addressing? Data structure is *static*. **Minimal Perfect Hashing.**

24	10100101
24	10001011
582	00001010
24	11011110
582	00010101
36352	01010011



VS

Hash-based

Open addressing? Data structure is *static*. **Minimal Perfect Hashing.**

10100101	24
10001011	24
00001010	582
11011110	24
00010101	582
01010011	36352



VS

Hash-based

Open addressing? Data structure is *static*. **Minimal Perfect Hashing.**

10100101	24
10001011	24
00001010	582
11011110	24

00010101	582
01010011	36352



VS

Hash-based

Open addressing? Data structure is *static*. **Minimal Perfect Hashing.**

10100101	24
10001011	24
00001010	582
11011110	24
00010101	582
01010011	36352



VS

Hash-based

Open addressing? Data structure is *static*. **Minimal Perfect Hashing.**







VS

Hash-based

Open addressing? Data structure is *static*. **Minimal Perfect Hashing.**

















Random access.

tongrams - Preliminary results

2

n	Number of <i>n</i> -grams	Maximum frequency count	Unique frequency counts	[lg] of unique frequency counts
1	24,359,472	468, 491, 999, 592	246,588	18
2	5,089,239	155, 178, <mark>1</mark> 63	44,822	16
3	52,635,338	102, 329, 901	71,690	17
4	11, 149, 161	6,401,274	21,127	15
5	8,261,975	958,556	12, 171	14
Total	101, 495, 185	468, 491, 999, 592	266,760	19

Table 4: Basic statistics for the GoogleWeb1T subset.

		Total space in GBs	Bytes per gram		Lookup time $[\mu s]$	
SH	KenLM	2.570	27.19		0.248	
НA	sxlm	1.012	10.43	(61.64%)	0.242	(2.42%)
ЯE	KenLM	1.829	21.5		1.272	
ТR	sxlm	0.541	5.7	(73.34%)	1.229	(3.38%)

Table 5: Bytes per grams and average lookup time in μ s for the GoogleWeb1T subset.

tongrams - Preliminary results

2

n	Number of <i>n</i> -grams	Maximum frequency count	Unique frequency counts	[lg] of unique
	" granns	nequency count	nequency counts	nequency counts
1	24, 359, 472	468, 491, 999, 592	246,588	18
2	5,089,239	155, 178, 163	44,822	16
3	52,635,338	102, 329, 901	71,690	17
4	11, 149, 161	6,401,274	21,127	15
5	8,261,975	958,556	12,171	14
Total	101, 495, 185	468, 491, 999, 592	266,760	19

Table 4: Basic statistics for the GoogleWeb1T subset.

		Total space in GBs	Bytes per gram		Lookup time $[\mu s]$	
SH	KenLM	2.570	27.19		0.248	
HA	sxlm	1.012	10.43	(61.64%)	0.242	(2.42%)
R	KenLM	1.829	21.5		1.272	
ТR	sxlm	0.541	5.7	(73.34%)	1.229	(3.38%)

Table 5: Bytes per grams and average lookup time in μ s for the GoogleWeb1T subset.

tongrams - Preliminary results

2

n	Number of <i>n</i> -grams	Maximum frequency count	Unique frequency counts	[lg] of unique frequency counts
	0	1	1 7	1
1	24,359,472	468, 491, 999, 592	246,588	18
2	5,089,239	155, 178, 163	44,822	16
3	52,635,338	102, 329, 901	71,690	17
4	11, 149, 161	6,401,274	21,127	15
5	8,261,975	958,556	12,171	14
Total	101, 495, 185	468, 491, 999, 592	266,760	19

Table 4: Basic statistics for the GoogleWeb1T subset.

		Total space in GBs	Bytes per gram		Lookup time $[\mu s]$	
SH	KenLM	2.570	27.19	X2.6	0.248	
ΗA	sxlm	1.012	10.43	(61.64%)	0.242	(2.42%)
TRIE	KenLM	1.829	21.5	X3.8	1.272	
	sxlm	0.541	5.7	(73.34%)	1.229	(3.38%)

Table 5: Bytes per grams and average lookup time in μ s for the GoogleWeb1T subset.

(Some) Future Research Problems

1

Dynamic Inverted Indexes.

Dynamic Inverted Indexes.

Classic solution: use two indexes. One is big and **cold**; the other is small and **hot**. **Merge** them periodically.

Dynamic Inverted Indexes.

Classic solution: use two indexes. One is big and **cold**; the other is small and **hot**. **Merge** them periodically.







(Some) Future Research Problems

2

Compressed B-trees.

2

Compressed B-trees.

Problem: maintain a dictionary on disk. Motivations: databases and file-systems.

2

Compressed B-trees.

Problem: maintain a dictionary on disk. Motivations: databases and file-systems.

"Fancy indexing structures may be a luxury now, but they will be essential by the decade's end."

(Some) Future Research Problems

Compressed B-trees.

Problem: maintain a dictionary on disk. Motivations: databases and file-systems.

"Fancy indexing structures may be a luxury now, but they will be essential by the decade's end."



Michael Bender Stony Brook University



Martin Farach-Colton

Rutgers University





Bradley Kuszmaul

MIT Laboratory for Computer Science

(Some) Future Research Problems

3

Fast Successor for IP-lookup.

3

Fast Successor for IP-lookup.

Successor search is what routers do for every incoming packet. Hence, the most run algorithm in the world.

Time and space efficiency is crucial.

Successor search is what routers do for every incoming packet. Hence, the most run algorithm in the world.

Time and space efficiency is crucial.

We can directly jump to the position of the first address having the same *Ig n* bits as the searched pattern in O(1) using the powerful search capabilities of Elias-Fano.

Successor search is what routers do for every incoming packet. Hence, the most run algorithm in the world.

Time and space efficiency is crucial.

We can directly jump to the position of the first address having the same *Ig n* bits as the searched pattern in O(1) using the powerful search capabilities of Elias-Fano.

Successor search is what routers do for every incoming packet. Hence, the most run algorithm in the world.

Time and space efficiency is crucial.

We can directly jump to the position of the first address having the same *Ig n* bits as the searched pattern in O(1) using the powerful search capabilities of Elias-Fano.

Successor search is what routers do for every incoming packet. Hence, the most run algorithm in the world.

Time and space efficiency is crucial.

We can directly jump to the position of the first address having the same *Ig n* bits as the searched pattern in O(1) using the powerful search capabilities of Elias-Fano.

> 11101110001000 Build an index on zeros.

Successor search is what routers do for every incoming packet. *Hence*, the most run algorithm in the world.

Time and space efficiency is crucial.

We can directly jump to the position of the first address having the same *Ig n* bits as the searched pattern in O(1) using the powerful search capabilities of Elias-Fano.

> 11101110001000 Build an index on zeros.

 $p = select_0(h_x) - h_x$

Successor search is what routers do for every incoming packet. Hence, the most run algorithm in the world.

Time and space efficiency is crucial.

We can directly jump to the position of the first address having the same *Ig n* bits as the searched pattern in O(1) using the powerful search capabilities of Elias-Fano.

1110111010001000Build an index on zeros.x = 001100 (12)

 $p = select_0(h_x) - h_x$

0	0	0	0	1	1
0	0	0	1	0	0
0	0	0	1	1	1
0	0	1	1	0	1
0	0	1	1	1	0
0	0	1	1	1	1
0	1	0	1	0	1
1	0	1	0	1	1
Fast Successor for IP-lookup.

Successor search is what routers do for every incoming packet. *Hence*, the most run algorithm in the world.

Time and space efficiency is crucial.

We can directly jump to the position of the first address having the same *Ig n* bits as the searched pattern in O(1) using the powerful search capabilities of Elias-Fano.

3

111011100001000 Build an index on zeros. x = 001100 (12)

 $p = select_0(h_x) - h_x$

Fast Successor for IP-lookup.

Successor search is what routers do for every incoming packet. *Hence*, the most run algorithm in the world.

Time and space efficiency is crucial.

We can directly jump to the position of the first address having the same *Ig n* bits as the searched pattern in O(1) using the powerful search capabilities of Elias-Fano.

> 1110111010001000 Build an index on zeros.

 $p = select_0(h_x) - h_x$

3

x = 001100(12)

Thanks for your attention, time, patience!

Any questions?