Space- and Time-Efficient Data Structures for Massive Datasets

Giulio Ermanno Pibiri

Referee Daniel Lemire Supervisor Rossano Venturini Referee Simon Gog



Department of Computer Science University of Pisa

08/03/2019



The increase of data and, hence, information does **not** scale with technology.



The increase of data and, hence, information does **not** scale with technology.



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

Niklaus Wirth, A Plea for Lean Software



The increase of data and, hence, information does **not** scale with technology.



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

Niklaus Wirth, A Plea for Lean Software

Even more relevant today!







Achieved results

Clustered Elias-Fano Indexes

Journal paper

Giulio Ermanno Pibiri and Rossano Venturini ACM Transactions on Information Systems (TOIS) Full paper, 34 pages, 2017.

Dynamic Elias-Fano Representation

Conference paper Giulio Ermanno Pibiri and Rossano Venturini Annual Symposium on Combinatorial Pattern Matching (CPM) Full paper, 14 pages, 2017.

On Optimally Partitioning Variable-Byte Codes

Giulio Ermanno Pibiri and Rossano Venturini IEEE Transactions on Knowledge and Data Engineering (TKDE). To appear. Full paper, 12 pages, 2019.

Fast Dictionary-based Compression for Inverted Indexes

Giulio Ermanno Pibiri, Matthias Petri and Alistair Moffat ACM Conference on Web Search and Data Mining (WSDM) Full paper, 9 pages, 2019.

Efficient Data Structures for Massive N-Gram Datasets

Giulio Ermanno Pibiri and Rossano Venturini ACM Conference on Research and Development in Information Retrieval (SIGIR) Full paper, 10 pages, 2017.

Handling Massive N-Gram Datasets Efficiently

Giulio Ermanno Pibiri and Rossano Venturini ACM Transactions on Information Systems (TOIS). To appear. Full paper, 41 pages, 2019.

Conference paper

Conference paper

Journal paper

Journal paper

Achieved results



Full paper, 41 pages, 2019.

Achieved results



Problem 1

Consider a sorted integer sequence.

Problem 1

Consider a sorted integer sequence.

How to represent it as a bit-vector where each original integer is uniquely-decodable, using **as few as possible** bits?

How to maintain **fast decompression speed**?

Ubiquity



Inverted indexes

The inverted index is the *de-facto* data structure at the basis of every large-scale retrieval system.

Inverted indexes

The inverted index is the *de-facto* data structure at the basis of every large-scale retrieval system.





Inverted indexes

The inverted index is the *de-facto* data structure at the basis of every large-scale retrieval system.



Many solutions

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

- Elias' Gamma and Delta
- Variable-Byte Family
- Binary Interpolative Coding
- Simple Family
- PForDelta
- QMX
- Elias-Fano
- Partitioned Elias-Fano

~1970

Many solutions

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







Is it possible to design an encoding that is **as small as BIC** and **much faster**?

1



Is it possible to design an encoding that is **as small as BIC** and **much faster**? Is it possible to design an encoding that is **as fast as VByte** and **much smaller**?

2



Is it possible to design an encoding that is **as small as BIC** and **much faster**? Is it possible to design an encoding that is **as fast as VByte** and **much smaller**?

2

What about **both** objectives at the same time?!

3



WSDM 2019

Every encoder represents each sequence individually.

Every encoder represents each sequence individually.

Encode **clusters** of (similar) inverted lists.

Every encoder represents each sequence individually.

Encode **clusters** of (similar) inverted lists.



reference list

Every encoder represents each sequence individually.

Encode **clusters** of (similar) inverted lists.



reference list

Every encoder represents each sequence individually.

Encode **clusters** of (similar) inverted lists.



The majority of values are **small** (*very* small indeed).

The majority of values are **small** (*very* small indeed).



The majority of values are **small** (*very* small indeed).



Encode **dense** regions with unary codes, **sparse** regions with VByte.

The majority of values are **small** (*very* small indeed).



Encode **dense** regions with unary codes, **sparse** regions with VByte.

Optimal partitioning in linear time and constant space.

Compression ratio improves by **2X**.

Query processing speed and sequential decoding (almost) **not affected**.

If we consider subsequences of *d*-gaps in inverted lists, these are **repetitive** across the whole inverted index.

If we consider subsequences of *d*-gaps in inverted lists, these are **repetitive** across the whole inverted index.

Put the **most frequent patterns** in a *dictionary* of size *k*. Then encode inverted lists as sequences of log₂ *k*-bit codewords.

If we consider subsequences of *d*-gaps in inverted lists, these are **repetitive** across the whole inverted index.

Put the **most frequent patterns** in a *dictionary* of size *k*. Then encode inverted lists as sequences of log₂ *k*-bit codewords.



If we consider subsequences of *d*-gaps in inverted lists, these are **repetitive** across the whole inverted index.

Put the **most frequent patterns** in a *dictionary* of size *k*. Then encode inverted lists as sequences of log₂ *k*-bit codewords.



Close to the most space-efficient representation (~**7%** away from BIC).

Almost **as fast as** the fastest SIMD-ized decoders.
The bigger picture



The bigger picture



The bigger picture



Problem 2

Consider a large text.

Problem 2

Consider a large text.

How to represent all its substrings of size $1 \le k \le N$ words for fixed N (e.g., N = 5), using **as few as possible** bits?

Fast access to individual N-grams?

How to **estimate** the probability of occurrence of the patterns under a given probability model?

Indexing





~6% of the books ever published

N	number of N-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 billions of N-grams!

The number of words following a given context is **small**.

The number of words following a given context is **small**.

k = 1



The number of words following a given context is **small**.

k = 1



The number of words following a given context is **small**.

k = 1



The number of words following a given context is **small**.

k = 1



The number of words following a given context is **small**.

k = 1



Map a word ID to the **position** it takes within its *sibling* IDs (the IDs following a context of fixed length *k*).

The (Elias-Fano) context-based remapped trie is **as fast as** the fastest competitor, but up to **65% smaller**.

The number of words following a given context is **small**.

k = 1



Map a word ID to the **position** it takes within its *sibling* IDs (the IDs following a context of fixed length *k*).

The (Elias-Fano) context-based remapped trie is **as fast as** the fastest competitor, but up to **65% smaller**. The (Elias-Fano) context-based remapped trie is **even smaller** than the most space-efficient competitors, that are lossy and with false-positives allowed, and up to **5X faster**.

To compute the modified Kneser-Ney probabilities of the N-grams, the fastest algorithm in the literature uses **3 sorting steps** in external memory.



To compute the modified Kneser-Ney probabilities of the N-grams, the fastest algorithm in the literature uses **3 sorting steps** in external memory.



Computing the distinct left extensions.



Suffix order



Context order

To compute the modified Kneser-Ney probabilities of the N-grams, the fastest algorithm in the literature uses **3 sorting steps** in external memory.



Computing the distinct left extensions.



Suffix order



Context order

To compute the modified Kneser-Ney probabilities of the N-grams, the fastest algorithm in the literature uses **3 sorting steps** in external memory.



Computing the distinct left extensions.



Suffix order



Context order

To compute the modified Kneser-Ney probabilities of the N-grams, the fastest algorithm in the literature uses **3 sorting steps** in external memory.



10

Х

С

Α

В

12

Х

X

Х

Х

С

Х

Х

С

Α

4

Α

Χ

Х

Х

Х

Computing the distinct left extensions.



To compute the modified Kneser-Ney probabilities of the N-grams, the fastest algorithm in the literature uses **3 sorting steps** in external memory.



Computing the distinct left extensions.



To compute the modified Kneser-Ney probabilities of the N-grams, the fastest algorithm in the literature uses **3 sorting steps** in external memory.



Computing the distinct left extensions.



Suffix order



Context order

Using a scan of the block and O(|V|) space.

To compute the modified Kneser-Ney probabilities of the N-grams, the fastest algorithm in the literature uses **3 sorting steps** in external memory.



Computing the distinct left extensions.



Suffix order



Context order

Using a scan of the block and O(|V|) space.









A	4
В	2
С	2
Х	4





























Rebuilding the last level of the trie.



Estimation runs 4.5X faster with billions of strings.

Take-home messages

 Efficiency to deliver better services by using less resources. Impact is far reaching and implies substantial economic gains.

• Compression is mandatory if your data are "big".

• Experiments are primary: design driven by numbers.

Thanks for your attention, time, patience!

Any questions?

High level thesis

Data Structures + Data Compression = Fast Algorithms

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

space and time-efficient



context





P("data" | "space and time-efficient") \approx

f("space and time-efficient data")

f("space and time-efficient")



 $\textbf{P}(\text{``data'' I ``space and time-efficient''}) ~\approx$

f("space and time-efficient data")

f("space and time-efficient")






Problem 3



- van Emde Boas Trees
- X/Y-Fast Tries
- Fusion Trees
- Exponential Search Trees
-
 - + time
 - space
 - + dynamic

Elias-Fano encoding

- EF(S(n,u)) = n log(u/n) + 2n bits
 to encode a sorted integer
 sequence S
- O(1) Access
- $O(1 + \log(u/n))$ **Predecessor**
 - + time
 - + space
 - static

Problem 3



Can we grab the best from both?

Dynamic inverted indexes

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

Append-only inverted indexes.









Integer dictionaries in succinct space (CPM 2017)



Integer dictionaries in succinct space (CPM 2017)

