Succinct Data Structures for Text and Information Retrieval

Simon Gog[1] Matthias Petri[2]

[1]Institute of Theoretical Informatics
Karslruhe Insitute of Technology

[2]Computing and Information Systems
The University of Melbourne, Australia

December 5, 2016

Outline

1. Introduction and Motivation (15 Minutes)
2. Basic Technologies and Notation (20 Minutes)
3. Index based Pattern Matching (20 Minutes)

   Break (20 Minutes)

4. Pattern Matching using Compressed Indexes (40 Minutes)
5. Applications to NLP (30 Minutes)

Introduction and Motivation (15 Mins)

1 What

2 Why

3 Who and Where

4 Practicality

What is it?

- Data structures and algorithms for working with large data sets
- Desiderata
  - miminise space requirement
  - maintaining efficient searchability
- Compressed data structures do just this! Near-optimal compression, with minor effect on runtime
- E.g., bitvector and integer compression, wavelet trees, compressed suffix array, compressed suffix trees

Why do we need it?

- Era of 'big data': text corpora are often 100s of gigabytes or terabytes in size (e.g., CommonCrawl, Twitter)
- Even simple algorithms like counting n-grams become difficult
- One solution is to use distributed computing, however can be very inefficient
- Succinct data structures provide a compelling alternative, providing compression and efficient access
- Complex algorithms become possible in memory, rather than requiring cluster and disk access

Why do we need it?

- Era of 'big data': text corpora are often 100s of gigabytes or terabytes in size (e.g., CommonCrawl, Twitter)
- Even simple algorithms like counting n-grams become difficult
- One solution is to use distributed computing, however can be very inefficient
- Succinct data structures provide a compelling alternative, providing compression and efficient access
- Complex algorithms become possible in memory, rather than requiring cluster and disk access

E.g., Infinite order language model possible, with runtime similar to current fixed order models, and lower space requirement.

Who uses it and where is it used?

Surprisingly few applications in NLP. However has applications in...

- Bioinformatics, Genome assembly
- Information Retrieval, Graph Search (Facebook)
- Search Engine Auto-complete
- Trajectory compression and retrieval
- XML storage and retrieval (xpath queries)
- Geo-spartial databases
- ...

Practicality

The SDSL library (GitHub repo: link) contains most practical compressed structures we talk about today.

It is easy to install:

```
git clone https://github.com/simongog/sdsl-lite.git
cd sdsl-lite
./install.sh
```

Throughout this tutorial we will show how to use SDSL to create and use a variety of different compressed data structures.

SDSL Resources

Tutorial:
http://simongog.github.io/assets/data/sdsl-slides/tutorial

Cheatsheet:
http://simongog.github.io/assets/data/sdsl-cheatsheet.pdf

Examples: https://github.com/simongog/sdsl-lite/examples

Tests: https://github.com/simongog/sdsl-lite/test

Basic Technologies and Notation (20 Mins)

1. Bitvectors
2. Rank and Select
3. Succinct Tree Representations
4. Variable Size Integers

Basic Building blocks: the bitvector

Definition

A bitvector (or bit array) B of length n compactly stores n binary numbers using n bits.

Example

	0	1	2	3	4	5	6	7	8	9	10	11
B	1	1	0	0	1	1	0	1	0	1	1	0

$B[0] = 1$, $B[1] = 1$, $B[2] = 0$, $B[n-1] = B[11] = 0$ etc.

Bitvector operations

Access and Set

$B[0] = 1$, $B[0] = B[1]$

Logical Operations

A OR B, A AND B, A XOR B

Advanced Operations

POPCOUNT(B): Number of one bits set
MSB_SET(B): Most significant bit set
LSB_SET(B): Least significant bit set

Operation RANK

Definitions

$\text{RANK}_1(B, j)$: How many 1's are in $B[0, j]$

$\text{RANK}_0(B, j)$: How many 0's are in $B[0, j]$

Example

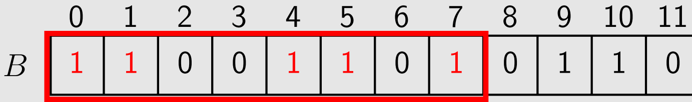

$\text{RANK}_1(B, 7) = 5$
$\text{RANK}_0(B, 7) = 8 - \text{RANK}_1(B, 7) = 3$

Operation $\textsc{Select}$

Definitions

$\textsc{Select}_1(B, j)$: Where is the j-th (start count at 1) 1 in B

$\textsc{Select}_0(B, j)$: Where is the j-th (start count at 1) 0 in B

Example

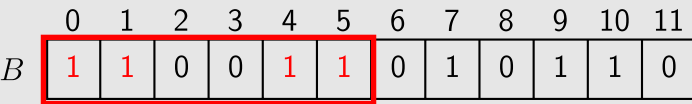

$\textsc{Select}_1(B, 4) = 5$
$\textsc{Select}_0(B, 3) = 6$

Complexity of Operations Rank and Select

Simple and Slow

Scan the whole bitvector using $O(1)$ extra space and $O(n)$ time to answer both Rank and Select

Constant time Rank

Divide bitvector into blocks. Store absolute ranks at block boundaries. Subdivide blocks into subblocks. Store ranks relative to block boundary. Subblocks are $O(\log n)$ which can be processed in constant time. Space usage: $n + o(n)$ bits. Runtime: $O(1)$. In practice: 25% extra space.

Constant time Select

Similar to Rank but more complex as blocks are based on the number of $1/0$ observed

Compressed Bitvectors

Idea

If only few 1's or clustering present in the bitvector, we can use compression techniques to substantially reduce space usage while efficiently supporting operations $\textsc{Rank}$ and $\textsc{Select}$

In Practice

Bitvector of size 1 GiB marking all uppercase letters in 8 GiB wikiepdia text:
Encodings:

- Elias-Fano ['73]: 343 MiB
- RRR ['02]: 335 MiB

Elias-Fano Coding

Elias-Fano Coding

Given a non-decreasing sequence X of length m over alphabet $[0..n]$. X can be represented using $2m + m \log \frac{n}{m} + o(m)$ bits while each element can still be accessed in constant time.

This representation can also be used to represent a bitvector (e.g. n is bitvector length, m the number of set bits, and X the position of the set bits)

How does Elias-Fano coding work?

- Divide each element into two parts: high-part and low-part.
- $\lfloor \log m \rfloor$ high-bits and $\lceil \log n \rceil - \lfloor \log m \rfloor$ low bits
- Sequence of high-parts of X is also non-decreasing.
- Gap encode the high-parts and use unary encoding to represent gaps. Call result H.
- I.e. for a gap of size g_i we use $g_i + 1$ bits (g_i zeros, 1 one).
- Sum of gaps $(= \#zeros)$ is at most $2^{\lfloor \log m \rfloor} \leq 2^{\log m} = m$
- I.e. H has size at most $2m$ ($\#zeros + \#ones$)
- Low-parts are represented explicitly.

How does Elias-Fano coding work?

$X =$ 4 13 15 24 26 27 29

00100 01101 01111 11000 11010 11011 11101

0 4 1 5 1 7 3 0 3 2 3 3 3 5

0−0 1−0 1−1 3−1 3−3 3−3 3−3

$\delta =$ 0 1 0 2 0 0 0

$H =$ 1011001111

$L =$ 4 5 7 0 2 3 5

How does Elias-Fano coding work?

Constant time access

- Add a select structure to H (Okanohara & Sadakane '07).

```
00  ACCESS(i)
01      p ← SELECT_1(H, i + 1)
02      x ← p − i
03      return x · 2^(⌈log n⌉−⌊log m⌋) + L[i]
```

Bitvectors - Practical Performance

How fast are RANK and SELECT in practice? Experiment: Cost per operation averaged over 1M executions: (code)

Uncompressed:

BV Size	Access	Rank	Select	Space
1MB	3ns	4ns	47ns	127%
10MB	10ns	14ns	85ns	126%
1GB	26ns	36ns	303ns	126%
10GB	78ns	98ns	372ns	126%

Compressed:

BV Size	Access	Rank	Select	Space
1MB	68ns	65ns	49ns	33%
10MB	99ns	88ns	58ns	30%
1GB	292ns	275ns	219ns	32%
10GB	466ns	424ns	336ns	30%

Using RANK and SELECT

- Basic building block of many compressed / succinct data structures
- Different implementations provide a variety of time and space trade-offs
- Implemented an ready to use in SDSL and many others:
  - http://github.com/simongog/sdsl-lite
  - http://github.com/facebook/folly
  - http://sux.di.unimi.it
  - http://github.com/ot/succinct
- Used in practice! For example: Facebook Graph search (Unicorn)

Succinct Tree Representations

Idea

Instead of storing pointers and objects, flatten the tree structure into a bitvector and use RANK and SELECT to navigate

From

```
typedef struct {
    void* data;      // 64 bits
    node_t* left;    // 64 bits
    node_t* right;   // 64 bits
    node_t* parent;  // 64 bits
} node_t;
```

To

Bitvector + RANK + SELECT + Data (≈ 2 bits per node)

Succinct Tree Representations

Definition: Succinct Data Structure

A succinct data structure uses space "close" to the information theoretical lower bound, but still supports operations time-efficiently.

Example: Succinct Tree Representations:

The number of unique binary trees containing n nodes is (roughly) 4^n. To differentiate between them we need at least $log_2(4^n) = 2n$ bits. Thus, a succinct tree representations should require $2n + o(n)$ bits.

LOUDS level order unary degree sequence

LOUDS

A succinct representation of a rooted, ordered tree containing nodes with arbitrary degree [Jacobson'89]

Example:

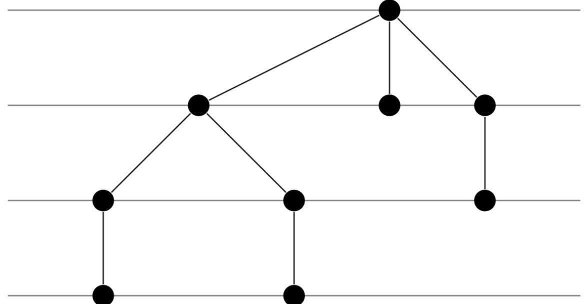

LOUDS Step 1

Add Pseudo Root:

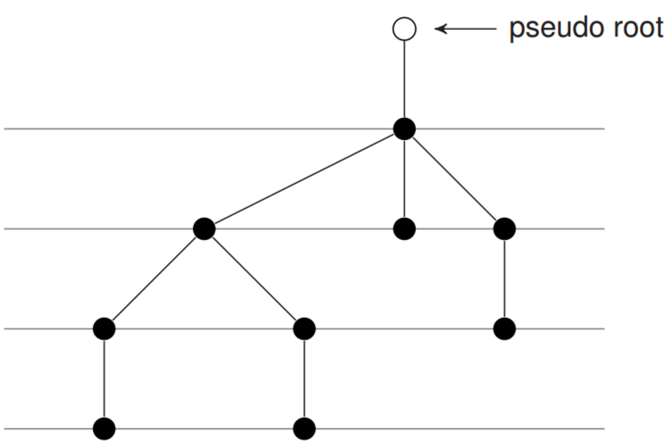

LOUDS Step 2

For each node unary encode the number of children:

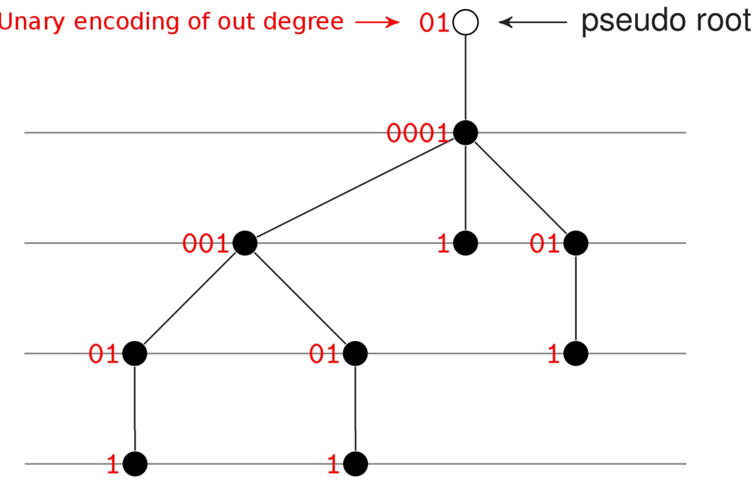

LOUDS Step 3

Write out unary encodings in level order:

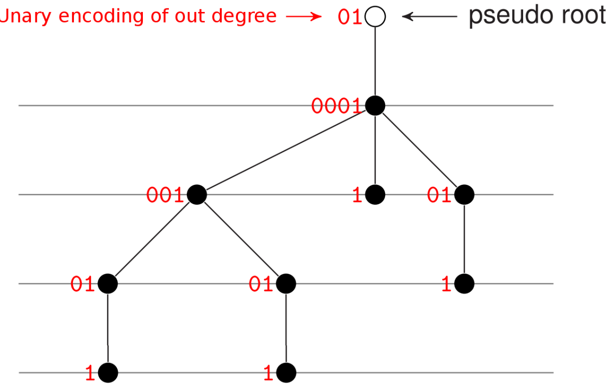

LOUDS sequence $L = 0100010011010101111$

LOUDS Nodes

- Each node (except the pseudo root) is represented twice
  - Once as “0” in the child list of its parent
  - Once as the terminal (“1”) in its child list
- Represent node v by the index of its corresponding “0”
- I.e. root corresponds to “0”
- A total of $2n$ bits are used to represent the tree shape!

LOUDS Navigation

Use RANK and SELECT to navigate the tree in constant time

Examples:

Compute node degree

```
int node_degree(int v) {
  if is_leaf(v) return 0
  id = Rank_0(L, v)
  return Select_1(L, id + 2)
     - Select_1(L, id + 1) - 1
}
```

Return the i-th child of node v

```
int child(int v, i) {
  if i > node_degree(v)
    return -1
  id = Rank_0(L, v)
  return Select_1(L, id + 1) + i
}
```

Complete construction, load, storage and navigation code of LOUDS is only 200 lines of C++ code.

Variable Size Integers

- Using 32 or 64 bit integers to store mostly small numbers is wasteful
- Many efficient encoding schemes exist to reduce space usage

Variable Byte Compression

Idea

Use variable number of bytes to represent integers. Each byte contains 7 bits "payload" and one continuation bit.

Examples

Number	Encoding	
824	00000110	10111000
5	10000101	

Storage Cost

Number Range	Number of Bytes
0 – 127	1
128 – 16383	2
16384 – 2097151	3

Variable Byte Compression - Algorithm

Encoding

1: **function** $\text{ENCODE}(x)$
2: **while** $x >= 128$ **do**
3: $\text{WRITE}(x \mod 128)$
4: $x = x \div 128$
5: **end while**
6: $\text{WRITE}(x + 128)$
7: **end function**

Decoding

1: **function** DECODE(bytes)
2: $x = 0$
3: $y =$ READ1(bytes)
4: **while** $y < 128$ **do**
5: $x = 128 \times x + y$
6: $y =$ READ1(bytes)
7: **end while**
8: $x = 128 \times x + (y - 128)$
9: **return** x
10: **end function**

Variable Sized Integer Sequences

Problem

Sequences of vbyte encoded numbers can not be accessed at arbitrary positions

Solution: Directly addressable variable-length codes (DAC)

Separate the indicator bits into a bitvector and use RANK and SELECT to access integers in $O(1)$ time. [Brisboa et al.'09]

DAC - Concept

Sample vbyte encoded sequence of integers:

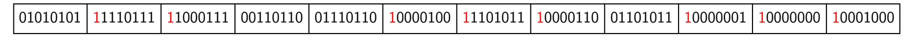

DAC restructuring of the vbyte encoded sequence of integers:

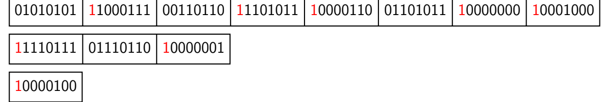

Separate the indicator bits:

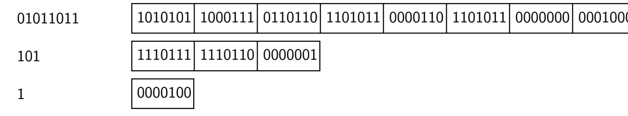

DAC - Access

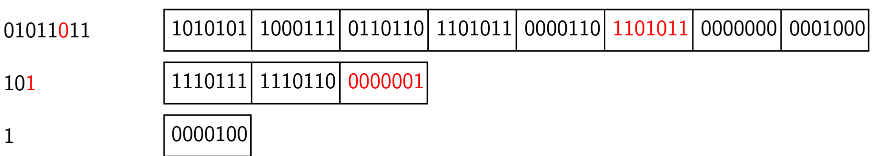

Accessing element A[5]:

- Access indicator bit of the first level at position 5: $I1[5] = 0$
- 0 in the indicator bit implies the number uses at least 2 bytes
- Perform $Rank_0(I1, 5) = 3$ to determine the number of integers in $A[0, 5]$ with at least two bytes
- Access $I2[3 - 1] = 1$ to determine that number $A[5]$ has two bytes.
- Access payloads and recover number in $O(1)$ time.

Practical Exercise

```
#include <vector>
#include "sdsl/dac_vector.hpp"

int main(int , char const *argv[])
{ using u32 = uint32_t; sdsl::int_vector<8> T;
  sdsl::load_vector_from_file(T,argv[1],1);
  std::vector<u32> counts(256*256*256,0);
  u32 cur3gram = (u32(T[0]) << 16) | (u32(T[1]) << 8);
  for(size_t i=2;i<T.size();i++) {
    cur3gram = ((cur3gram&0x0000FFFF)<<8) | u32(T[i]);
    counts[cur3gram]++;
  }
  std::cout << "u32 = " << sdsl::size_in_mega_bytes(counts);
  sdsl::dac_vector<3> dace(counts);
  std::cout << "dac = " << sdsl::size_in_mega_bytes(dace);
}
```

Code: here.

Index based Pattern Matching (20 Mins)

5 Suffix Trees

6 Suffix Arrays

7 Compressed Suffix Arrays

Pattern Matching

Definition

Given a text T of size n, find all occurrences (or just count) of pattern P of length m.

Online Pattern Matching

Preprocess P, scan T. Examples: KMP, Boyer-Moore, BMH etc. $O(n+m)$ search time.

Offline Pattern Matching

Preprocess T, Build Index. Examples: Inverted Index, Suffix Tree, Suffix Array. $O(m)$ search time.

Suffix Tree (Weiner'73)

- Data structure capable of processing T in $O(n)$ time and answering search queries in $O(n)$ space and $O(m)$ time. Optimal from a theoretical perspective.
- All suffixes of T into a trie (a tree with edge labels)
- Contains n leaf nodes corresponding to the n suffixes of T
- Search for a pattern P is performed by finding the subtree corresponding to all suffixes prefixed by P

Suffix Tree - Example

$$T = \texttt{abracadabracarab\$}$$

Suffix Tree - Example

$$T = \texttt{abracadabracarab\$}$$

Suffixes:

0	abracadabracarab$
1	bracadabracarab$
2	racadabracarab$
3	acadabracarab$
4	cadabracarab$
5	adabracarab$
6	dabracarab$
7	abracarab$
8	bracarab$
9	racarab$
10	acarab$
11	carab$
12	arab$
13	rab$
14	ab$
15	b$
16	$

Suffix Tree - Example

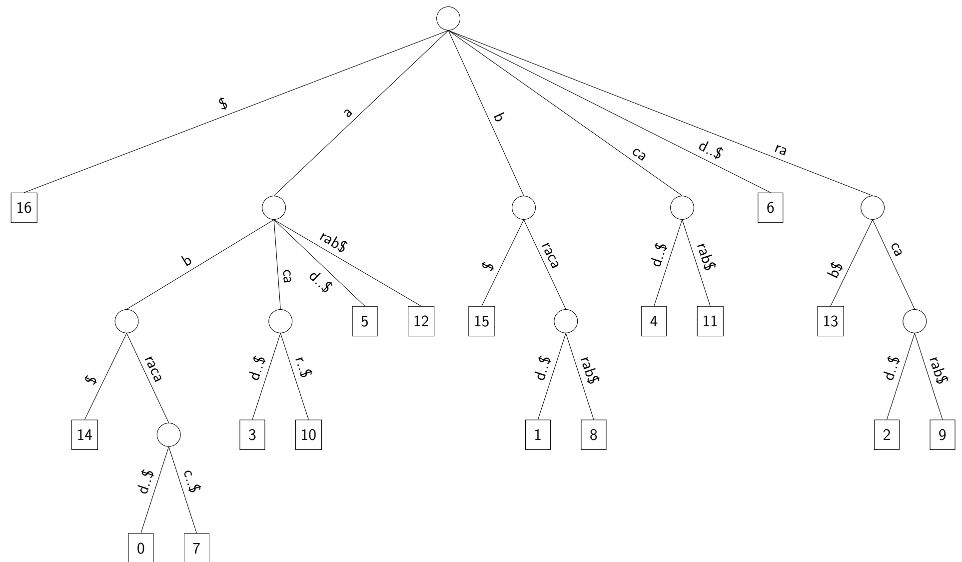

Suffix Tree - Search for "aca"

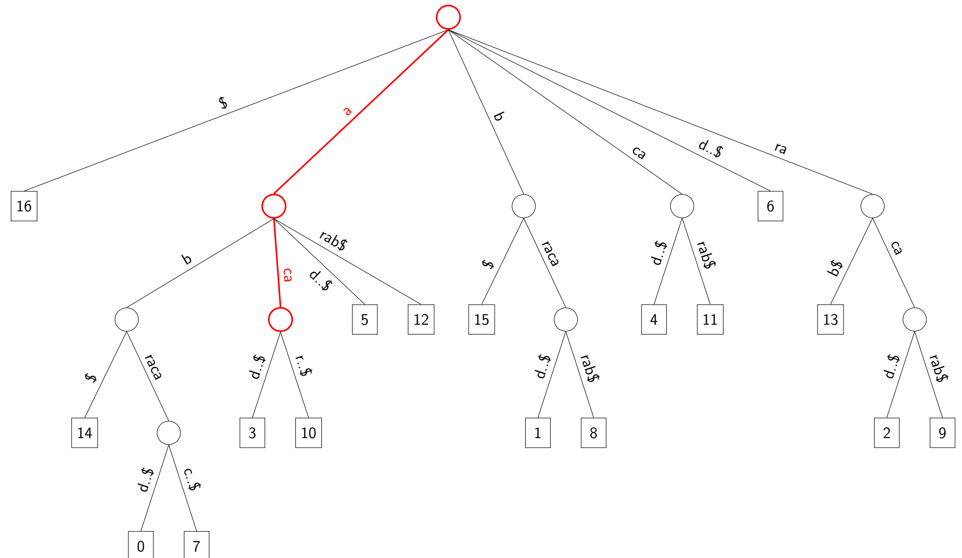

Suffix Tree - Problems

- Space usage in practice is large. $20 - 40$ times n for highly optimized implementations.
- Only useable for small datasets.

Suffix Arrays (Manber and Myers'92)

- Reduce space of Suffix Tree by only storing the n leaf pointers into the text
- Requires $n \log n$ bits for the pointers plus T to perform search
- In practice $5 - 9$n bytes for character alphabets
- Search for P using binary search

Suffix Arrays - Example

$$T = \texttt{abracadabracarab\$}$$

Suffix Arrays - Example

$T =$abracadabracarab$

Suffixes:

0	abracadabracarab$
1	bracadabracarab$
2	racadabracarab$
3	acadabracarab$
4	cadabracarab$
5	adabracarab$
6	dabracarab$
7	abracarab$
8	bracarab$
9	racarab$
10	acarab$
11	carab$
12	arab$
13	rab$
14	ab$
15	b$
16	$

Suffix Arrays - Example

$$T = \texttt{abracadabracarab\$}$$

Sorted Suffixes:

16	$
14	ab$
0	abracadabracarab$
7	abracarab$
3	acadabracarab$
10	acarab$
5	adabracarab$
12	arab$
15	b$
1	bracadabracarab$
8	bracarab$
4	cadabracarab$
11	carab$
6	dabracarab$
13	rab$
2	racadabracarab$
9	racarab$

Suffix Arrays - Example

$T =$ abracadabracarab$

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
a	b	r	a	c	a	d	a	b	r	a	c	a	r	a	b	$

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
16	14	0	7	3	10	5	12	15	1	8	4	11	6	13	2	9

Suffix Arrays - Search

$T =$abracadabracarab\$, $P =$abr

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
a	b	r	a	c	a	d	a	b	r	a	c	a	r	a	b	b

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
16	14	0	7	3	10	5	12	15	1	8	4	11	6	13	2	9

Suffix Arrays - Search

$T =$abracadabracarab\$, $P =$abr

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
a	b	r	a	c	a	d	a	b	r	a	c	a	r	a	b	$

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
16	14	0	7	3	10	5	12	15	1	8	4	11	6	13	2	9

Suffix Arrays - Search

$T =$ abracadabracarab\$, $P =$ abr

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
a	b	r	a	c	a	d	a	b	r	a	c	a	r	a	b	b

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
16	14	0	7	3	10	5	12	15	1	8	4	11	6	13	2	9

Suffix Arrays - Search

$T =$abracadabracarab\$, $P =$abr

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
a	b	r	a	c	a	d	a	b	r	a	c	a	r	a	b	b

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
16	14	0	7	3	10	5	12	15	1	8	4	11	6	13	2	9

Suffix Arrays - Search

$$T = \texttt{abracadabracarab\$},$$

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
a	b	r	a	c	a	d	a	b	r	a	c	a	r	a	b	b

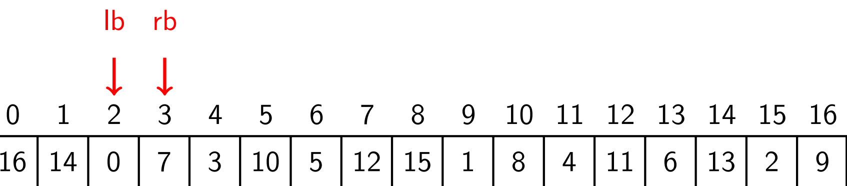

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
16	14	0	7	3	10	5	12	15	1	8	4	11	6	13	2	9

Suffix Arrays / Trees - Resource Consumption

In practice:

- Suffix Trees requires $\approx 20n$ bytes of space (for efficient implementations)
- Suffix Arrays require $5-9n$ bytes of space
- Comparable search performance

Example: 5GB English text requires 45GB for a character level suffix array index and up to 200GB for suffix trees

Suffix Arrays / Trees - Construction

In theory: Both can be constructed in optimal $O(n)$ time

In practice:

- Suffix Trees and Suffix Arrays construction can be parallelized
- Most efficient suffix array construction algorithm in practice are not $O(n)$
- Efficient semi-external memory construction algorithms exist
- Parallel suffix array construction algorithms can index 20MiB/s (24 threads) in-memory and 4MiB/s in external memory
- Suffix Arrays of terabyte scale text collection can be constructed. Practical!
- Word-level Suffix Array construction also possible.

Dilemma

- There is lots of work out there which proposes solutions for different problems based on suffix trees
- Suffix trees (and to a certain extend suffix arrays) are not really applicable for large scale problems
- However, large scale suffix arrays can be constructed efficiently without requiring large amounts of memory

Solutions:

- External or Semi-External memory representation of suffix trees / arrays

Dilemma

- There is lots of work out there which proposes solutions for different problems based on suffix trees
- Suffix trees (and to a certain extend suffix arrays) are not really applicable for large scale problems
- However, large scale suffix arrays can be constructed efficiently without requiring large amounts of memory

Solutions:

- External or Semi-External memory representation of suffix trees / arrays
- Compression?

External / Semi-External Suffix Indexes

String-B Tree [Ferragina and Grossi'99]

- Cache-Oblivious
- Uses blind-trie (succinct trie; requires verification step)
- Space requirement on disk one order of magnitude larger than text

Semi-External Suffix Array (RoSA) [Gog et al.'14]

- Compressed version of the String-B tree
- Replace blind-trie with a condensed BWT
- If pattern is frequent: Answer from in-memory structure (fast!)
- If pattern is infrequent: perform disk access

Compressed Suffix Arrays and Trees

Idea

Utilize data compression techniques to substantially reduce the space of suffix arrays/trees while retaining their functionality

Compressed Suffix Arrays (CSA):

- Use space equivalent to the compressed size of the input text. Not 4-8 times more! Example: 1GB English text compressed to roughly 300MB using gzip. CSA uses roughly 300MB (sometimes less)!
- Provide more functionality than regular suffix arrays
- Implicitly contain the original text, no need to retain it. Not needed for query processing
- Similar search efficiency than regular suffix arrays.
- Used to index terabytes of data on a reasonably powerful machine!

CSA and CST in practice using SDSL

```
1  #include "sdsl/suffix_arrays.hpp"
2  #include <iostream>
3
4  int main(int argc,char** argv) {
5      std::string input_file = argv[1];
6      std::string out_file = argv[2];
7      sdsl::csa_wt<> csa;
8      sdsl::construct(csa,input_file,1);
9      std::cout << "CSA size = "
10         << sdsl::size_in_megabytes(csa) << std::endl;
11     sdsl::store_to_file(csa,out_file);
12 }
```

Code: here.

How does it work? Find out after the break!

Break Time

See you back here in 20 minutes!

Compressed Indexes (40 Mins)

1 CSA Internals

2 BWT

3 Wavelet Trees

4 CSA Usage

5 Compressed Suffix Trees

Compressed Suffix Arrays - Overview

Two practical approaches developed independently:

- CSA-SADA: Proposed by Grossi and Vitter in 2000. Practical refinements by Sadakane also in 2000.
- CSA-WT: Also referred to as the FM-Index. Proposed by Ferragina and Manzini in 2000.

Many practical (and theoretical) improvements to compression, query and construction speed since then. Efficient implementations available in SDSL: `csa_sada<>` and `csa_wt<>`.

For now, we focus on CSA-WT.

CSA-WT or the FM-Index

- Utilizes the Burrows-Wheeler Transform (BWT) used in compression tools such as `bzip2`
- Requires RANK and SELECT on non-binary alphabets
- Heavily utilize compressed bitvector representations
- Theoretical bound on space usage related to compressibility (entropy) of the input text

The Burrows-Wheeler Transform (BWT)

- Reversible Text Permutation
- Initially proposed by Burrows and Wheeler as a compression tool. The BWT is easier to compress than the original text!
- Defined as $BWT[i] = T[SA[i] - 1 \mod n]$
- In words: $BWT[i]$ is the symbol preceding suffix $SA[i]$ in T

Why does it work? How is it related to searching?

BWT - Example

$$T = \texttt{abracadabracarab\$}$$

BWT - Example

$T =$abracadabracarab$

0	abracadabracarab$
1	bracadabracarab$
2	racadabracarab$
3	acadabracarab$
4	cadabracarab$
5	adabracarab$
6	dabracarab$
7	abracarab$
8	bracarab$
9	racarab$
10	acarab$
11	carab$
12	arab$
13	rab$
14	ab$
15	b$
16	$

BWT - Example

$$T = \texttt{abracadabracarab\$}$$

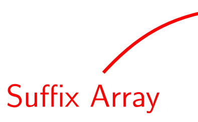

16	$
14	ab$
0	abracadabracarab$
7	abracarab$
3	acadabracarab$
10	acarab$
5	adabracarab$
12	arab$
15	b$
1	bracadabracarab$
8	bracarab$
4	cadabracarab$
11	carab$
6	dabracarab$
13	rab$
2	racadabracarab$
9	racarab$

BWT - Example

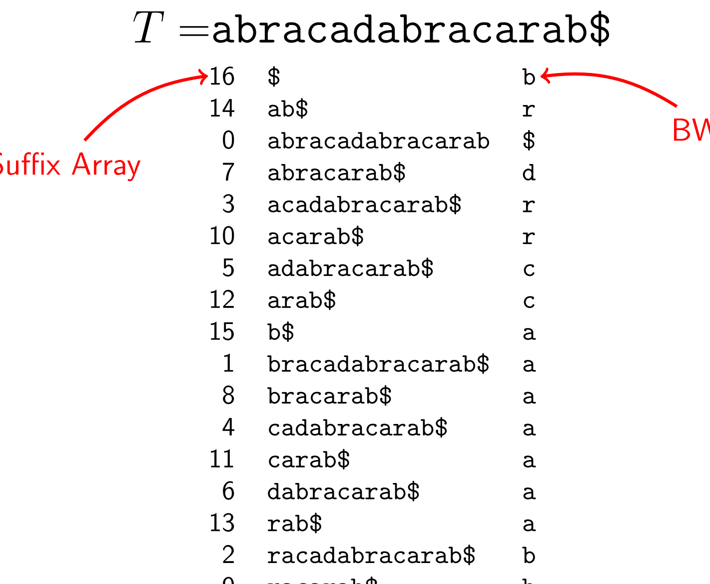

$T=$abracadabracarab$

16	$	b
14	ab$	r
0	abracadabracarab	$
7	abracarab$	d
3	acadabracarab$	r
10	acarab$	r
5	adabracarab$	c
12	arab$	c
15	b$	a
1	bracadabracarab$	a
8	bracarab$	a
4	cadabracarab$	a
11	carab$	a
6	dabracarab$	a
13	rab$	a
2	racadabracarab$	b
9	racarab$	b

BWT - Example

$$T = \texttt{abracadabracarab\$}$$

F	L
$	b
a	r
a	$
a	d
a	r
a	r
a	c
a	c
b	a
b	a
b	a
c	a
c	a
d	a
r	a
r	b
r	b

BWT - Reconstructing T from BWT

$T =$

b
r
$
d
r
r
c
c
a
a
a
a
a
a
a
b
b

BWT - Reconstructing T from BWT

$T =$

0	$	b
1	a	r
2	a	$
3	a	d
4	a	r
5	a	r
6	a	c
7	a	c
8	b	a
9	b	a
10	b	a
11	c	a
12	c	a
13	d	a
14	r	a
15	r	b
16	r	b

1. Sort BWT to retrieve first column F

BWT - Reconstructing T from BWT

$T =$ $

	F	L
0	$	b
1	a	r
2	a	$
3	a	d
4	a	r
5	a	r
6	a	c
7	a	c
8	b	a
9	b	a
10	b	a
11	c	a
12	c	a
13	d	a
14	r	a
15	r	b
16	r	b

2. Find last symbol $\$$ in F at position 0 and write to output

BWT - Reconstructing T from BWT

$T =$ b$

	F	BWT
0	$	b
1	a	r
2	a	$
3	a	d
4	a	r
5	a	r
6	a	c
7	a	c
8	b	a
9	b	a
10	b	a
11	c	a
12	c	a
13	d	a
14	r	a
15	r	b
16	r	b

2. Symbol preceding $ in T is $BWT[0] = b$. Write to output

BWT - Reconstructing T from BWT

$T =$ b$

0	$	b
1	a	r
2	a	$
3	a	d
4	a	r
5	a	r
6	a	c
7	a	c
8	b	a
9	b	a
10	b	a
11	c	a
12	c	a
13	d	a
14	r	a
15	r	b
16	r	b

3. As there is no b before $BWT[0]$, we know that this b corresponds to the first b in F at pos $F[8]$.

BWT - Reconstructing T from BWT

$T =$ ab$

	F	BWT
0	$	b
1	a	r
2	a	$
3	a	d
4	a	r
5	a	r
6	a	c
7	a	c
8	b	a
9	b	a
10	b	a
11	c	a
12	c	a
13	d	a
14	r	a
15	r	b
16	r	b

4. The symbol preceding $F[8]$ is $BWT[8] = a$. Output!

BWT - Reconstructing T from BWT

$T =$ ab$

0	$	b
1	a	r
2	a	$
3	a	d
4	a	r
5	a	r
6	a	c
7	a	c
8	b	a
9	b	a
10	b	a
11	c	a
12	c	a
13	d	a
14	r	a
15	r	b
16	r	b

5. Map that a back to F at position $F[1]$

BWT - Reconstructing T from BWT

$T =$ rab$

	F	BWT
0	$	b
1	a	r
2	a	$
3	a	d
4	a	r
5	a	r
6	a	c
7	a	c
8	b	a
9	b	a
10	b	a
11	c	a
12	c	a
13	d	a
14	r	a
15	r	b
16	r	b

6. Output $BWT[1] = r$ and map r to $F[14]$

BWT - Reconstructing T from BWT

$T =$ arab$

0	$	b
1	a	r
2	a	$
3	a	d
4	a	r
5	a	r
6	a	c
7	a	c
8	b	a
9	b	a
10	b	a
11	c	a
12	c	a
13	d	a
14	r	a
15	r	b
16	r	b

7. Output $BWT[14] = a$ and map a to $F[7]$

BWT - Reconstructing T from BWT

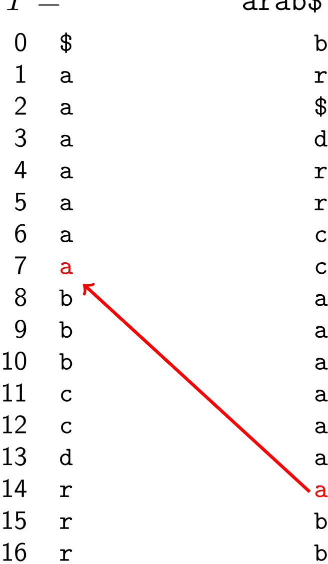

Why does $BWT[14] = a$ map to $F[7]$?

BWT - Reconstructing T from BWT

All a preceding $BWT[14] = a$ preceed suffixes smaller than $SA[14]$.

$T =$ arab$

0	$	b
1	a	r
2	a	$
3	a	d
4	a	r
5	a	r
6	a	c
7	a	c
8	b	a
9	b	a
10	b	a
11	c	a
12	c	a
13	d	a
14	r	a
15	r	b
16	r	b

BWT - Reconstructing T from BWT

Thus, among the suffixes starting with a, the one preceding $SA[14]$ must be the last one.

$T =$ arab$

0	$	b
1	a	r
2	a	$
3	a	d
4	a	r
5	a	r
6	a	c
7	a	c
8	b	a
9	b	a
10	b	a
11	c	a
12	c	a
13	d	a
14	r	a
15	r	b
16	r	b

BWT - Reconstructing T from BWT

$T =$abracadabracarab$

0	$	b
1	a	r
2	a	$
3	a	d
4	a	r
5	a	r
6	a	c
7	a	c
8	b	a
9	b	a
10	b	a
11	c	a
12	c	a
13	d	a
14	r	a
15	r	b
16	r	b

Searching using the BWT

$T =$abracadabracarab$, $P =$abr

	F	L
0	$	b
1	a	r
2	a	$
3	a	d
4	a	r
5	a	r
6	a	c
7	a	c
8	b	a
9	b	a
10	b	a
11	c	a
12	c	a
13	d	a
14	r	a
15	r	b
16	r	b

Searching using the BWT

$T=$abracadabracarab\$, $P=$abr

	F	L
0	\$	b
1	a	r
2	a	\$
3	a	d
4	a	r
5	a	r
6	a	c
7	a	c
8	b	a
9	b	a
10	b	a
11	c	a
12	c	a
13	d	a
14	r	a
15	r	b
16	r	b

Search backwards, start by finding the r interval in F

Searching using the BWT

T =abracadabracarab$, P =abr

	F	L
0	$	b
1	a	r
2	a	$
3	a	d
4	a	r
5	a	r
6	a	c
7	a	c
8	b	a
9	b	a
10	b	a
11	c	a
12	c	a
13	d	a
→ 14	r	a
15	r	b
→ 16	r	b

Search backwards, start by finding the r interval in F

Searching using the BWT

T =abracadabracarab$, P =abr

	F		L
0	$		b
1	a		r
2	a		$
3	a		d
4	a		r
5	a		r
6	a		c
7	a		c
8	b		a
9	b		a
10	b		a
11	c		a
12	c		a
13	d		a
→ 14	r		a
15	r		b
→ 16	r		b

How many bs are the r interval in $BWT[14, 16]$? 2

Searching using the BWT

T =abracadabracarab$, P =abr

0	$	b
1	a	r
2	a	$
3	a	d
4	a	r
5	a	r
6	a	c
7	a	c
8	b	a
9	b	a
10	b	a
11	c	a
12	c	a
13	d	a
→ 14	r	a
15	r	b
→ 16	r	b

How many suffixes starting with b are smaller than those 2?

1 at $BWT[0]$

Searching using the BWT

$T =$ abracadabracarab\$, $P =$ abr

Thus, all suffixes starting with br are in $SA[9, 10]$.

0	$	b
1	a	r
2	a	$
3	a	d
4	a	r
5	a	r
6	a	c
7	a	c
8	b	a
→ 9	b	a
→ 10	b	a
11	c	a
12	c	a
13	d	a
14	r	a
15	r	b
16	r	b

Searching using the BWT

T =abracadabracarab$, P =abr

	0	$	b
	1	a	r
	2	a	$
	3	a	d
	4	a	r
	5	a	r
	6	a	c
	7	a	c
	8	b	a
→	9	b	a
→	10	b	a
	11	c	a
	12	c	a
	13	d	a
	14	r	a
	15	r	b
	16	r	b

How many of the suffixes starting with br are preceded by a? 2

Searching using the BWT

$T =$abracadabracarab$, $P =$abr

How many of the suffixes smaller than br are preceded by a? 1

0	$	b
1	a	r
2	a	$
3	a	d
4	a	r
5	a	r
6	a	c
7	a	c
8	b	a
→ 9	b	a
→ 10	b	a
11	c	a
12	c	a
13	d	a
14	r	a
15	r	b
16	r	b

Searching using the BWT

T =abracadabracarab$, P =abr

0	$	b
1	a	r
→ 2	a	$
→ 3	a	d
4	a	r
5	a	r
6	a	c
7	a	c
8	b	a
9	b	a
10	b	a
11	c	a
12	c	a
13	d	a
14	r	a
15	r	b
16	r	b

There are 2 occurrences of abr in T corresponding to suffixes $SA[2, 3]$

Searching using the BWT

- We only require F and BWT to search and recover T
- We only had to count the number of times a symbol s occurs within an interval, and before that interval $BWT[i, j]$
- Equivalent to $Rank_s(BWT, i)$ and $Rank_s(BWT, j)$
- Need to perform $Rank$ on non-binary alphabets efficiently

Wavelet Trees - Overview

- Data structure to perform $Rank$ and $Select$ on non-binary alphabets of size σ in $O(\log_2 \sigma)$ time
- Decompose non-binary $Rank$ operations into binary $Rank$s via tree decomposition
- Space usage $n \log \sigma + o(n \log \sigma)$ bits. Same as original sequence + Rank + Select overhead

Wavelet Trees - Example

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
b	r	$	d	r	r	c	c	a	a	a	a	a	a	a	b	b

Symbol	Codeword
$	00
a	010
b	011
c	10
d	110
r	111

Wavelet Trees - Example

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
b	r	$	d	r	r	c	c	a	a	a	a	a	a	a	b	b
0	1	0	1	1	1	1	1	0	0	0	0	0	0	0	0	0

Wavelet Trees - Example

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
b	r	$	d	r	r	c	c	a	a	a	a	a	a	a	b	b
0	1	0	1	1	1	1	1	0	0	0	0	0	0	0	0	0

Wavelet Trees - Example

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
b	r	$	d	r	r	c	c	a	a	a	a	a	a	a	b	b
0	1	0	1	1	1	1	1	0	0	0	0	0	0	0	0	0

0	1	2	3	4	5	6	7	8	9	10
b	$	a	a	a	a	a	a	a	b	b
1	0	1	1	1	1	1	1	1	1	1

Wavelet Trees - Example

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
b	r	$	d	r	r	c	c	a	a	a	a	a	a	a	b	b
0	1	0	1	1	1	1	1	0	0	0	0	0	0	0	0	0

0	1	2	3	4	5	6	7	8	9	10
b	$	a	a	a	a	a	a	a	b	b
1	0	1	1	1	1	1	1	1	1	1

0	1	2	3	4	5
r	d	r	r	c	c
1	1	1	1	0	0

Wavelet Trees - Example

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
b	r	$	d	r	r	c	c	a	a	a	a	a	a	a	b	b
0	1	0	1	1	1	1	1	0	0	0	0	0	0	0	0	0

0	1	2	3	4	5	6	7	8	9	10
b	$	a	a	a	a	a	a	a	b	b
1	0	1	1	1	1	1	1	1	1	1

0	1	2	3	4	5
r	d	r	r	c	c
1	1	1	1	0	0

$

Wavelet Trees - Example

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
b	r	$	d	r	r	c	c	a	a	a	a	a	a	a	b	b
0	1	0	1	1	1	1	1	0	0	0	0	0	0	0	0	0

0	1	2	3	4	5	6	7	8	9	10
b	$	a	a	a	a	a	a	a	b	b
1	0	1	1	1	1	1	1	1	1	1

0	1	2	3	4	5
r	d	r	r	c	c
1	1	1	1	0	0

$

0	1	2	3	4	5	6	7	8	9
b	a	a	a	a	a	a	a	b	b
1	0	0	0	0	0	0	0	1	1

Wavelet Trees - Example

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
b	r	$	d	r	r	c	c	a	a	a	a	a	a	a	b	b
0	1	0	1	1	1	1	1	0	0	0	0	0	0	0	0	0

0	1	2	3	4	5	6	7	8	9	10
b	$	a	a	a	a	a	a	a	b	b
1	0	1	1	1	1	1	1	1	1	1

0	1	2	3	4	5
r	d	r	r	c	c
1	1	1	1	0	0

$

0	1	2	3	4	5	6	7	8	9
b	a	a	a	a	a	a	a	b	b
1	0	0	0	0	0	0	0	1	1

Wavelet Trees - Example

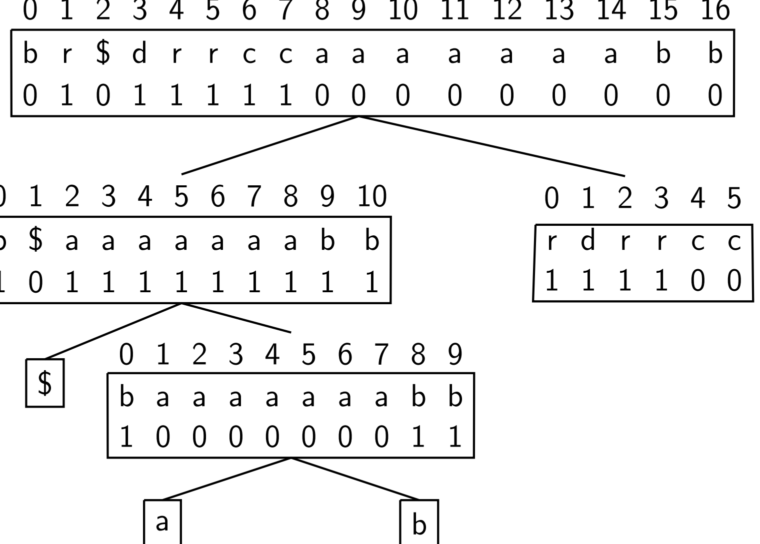

Wavelet Trees - Example

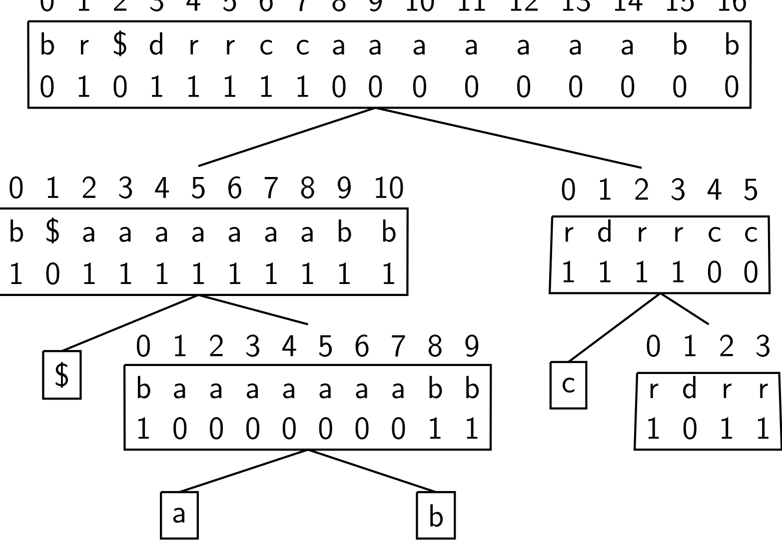

Wavelet Trees - Example

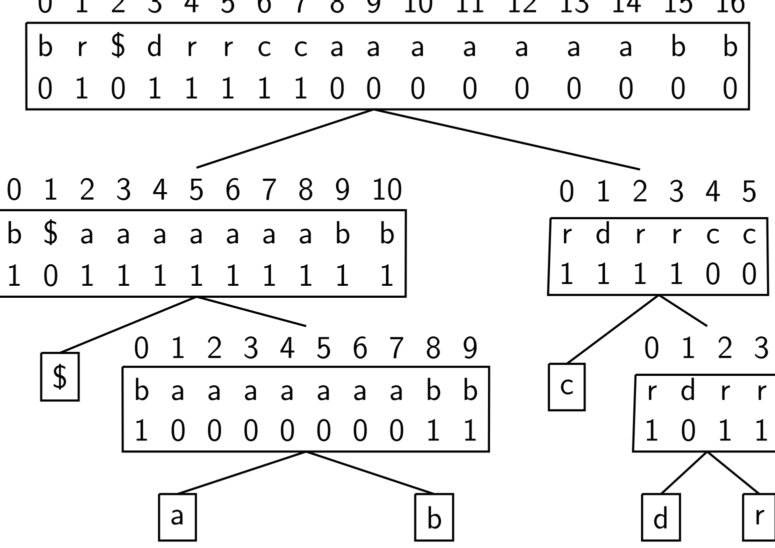

Wavelet Trees - What is actually stored

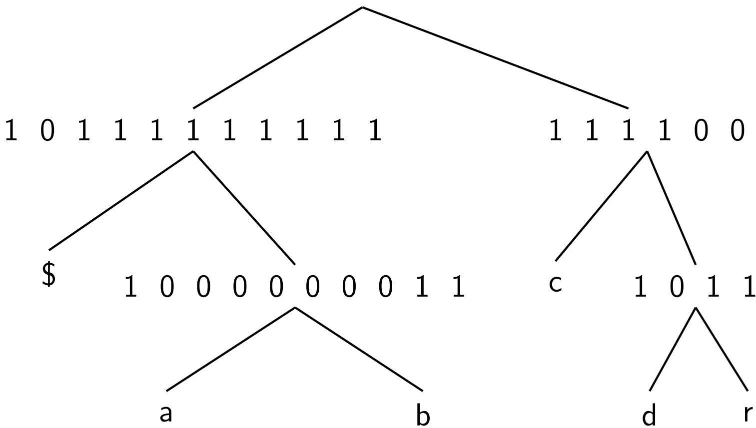

Wavelet Trees - Performing $Rank_a(BWT, 11)$

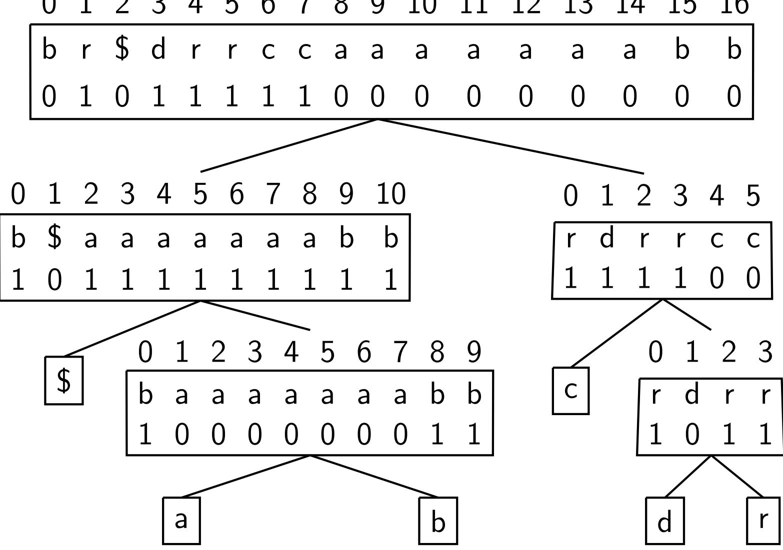

Wavelet Trees - Performing $Rank_a(BWT, 11)$

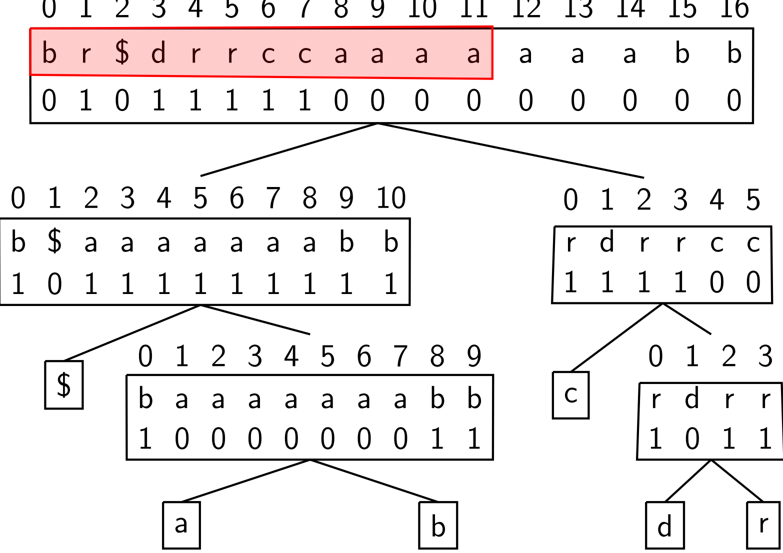

Wavelet Trees - Performing $Rank_a(BWT, 11)$

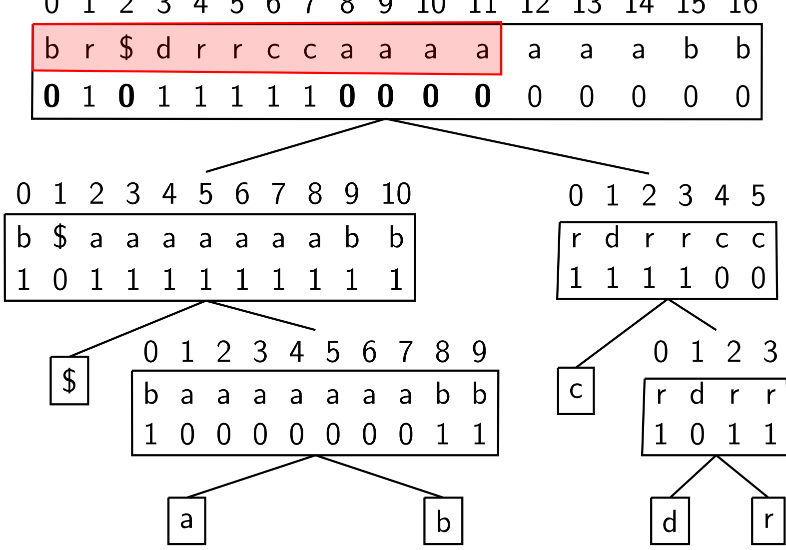

Wavelet Trees - Performing $Rank_a(BWT, 11)$

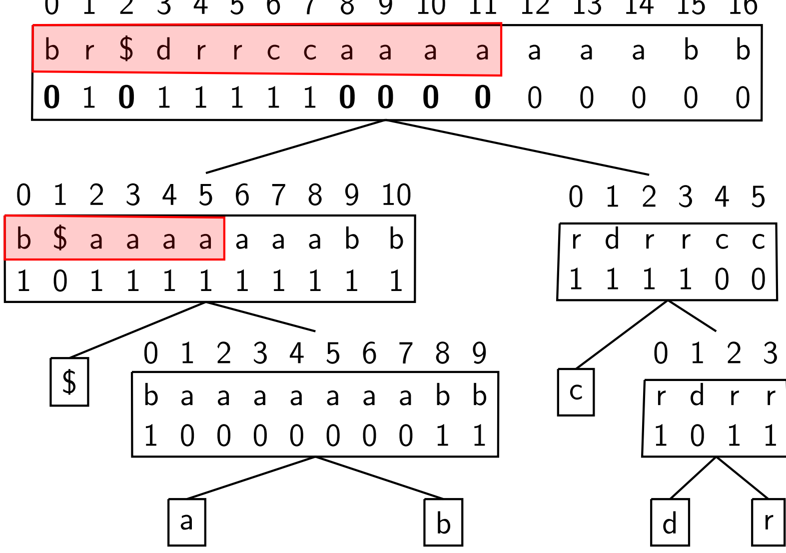

Wavelet Trees - Performing $Rank_a(BWT, 11)$

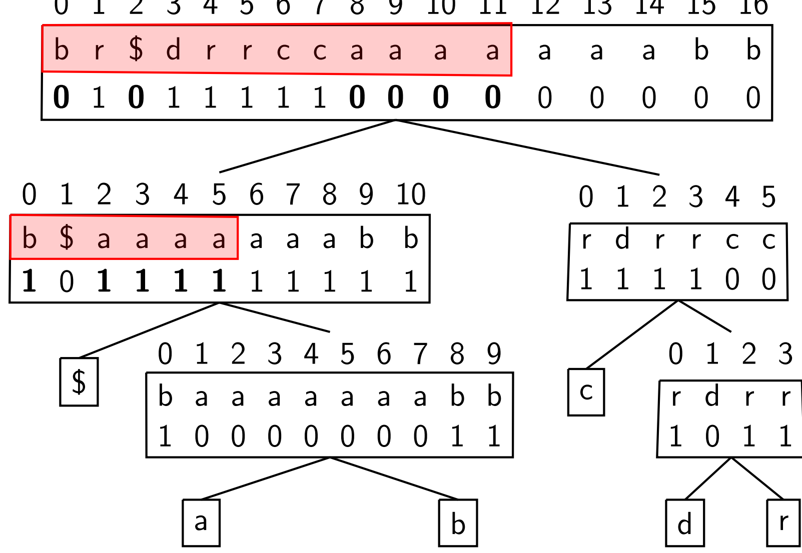

Wavelet Trees - Performing $Rank_a(BWT, 11)$

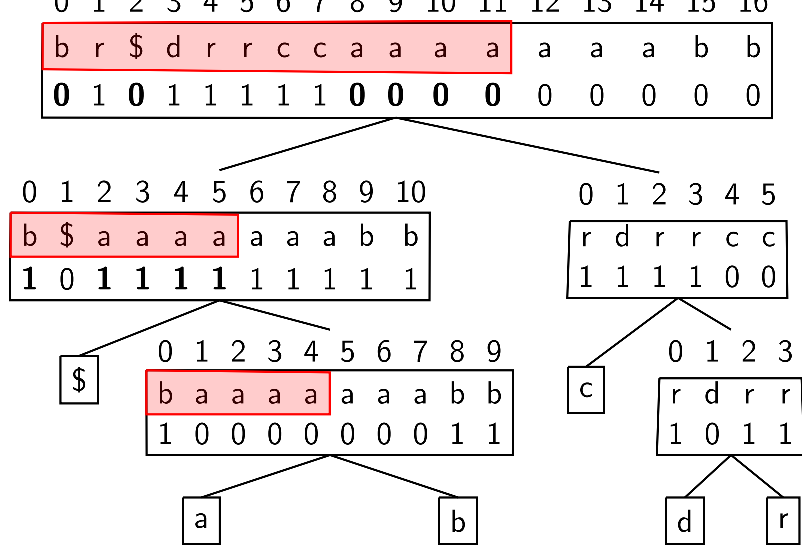

Wavelet Trees - Performing $Rank_a(BWT, 11)$

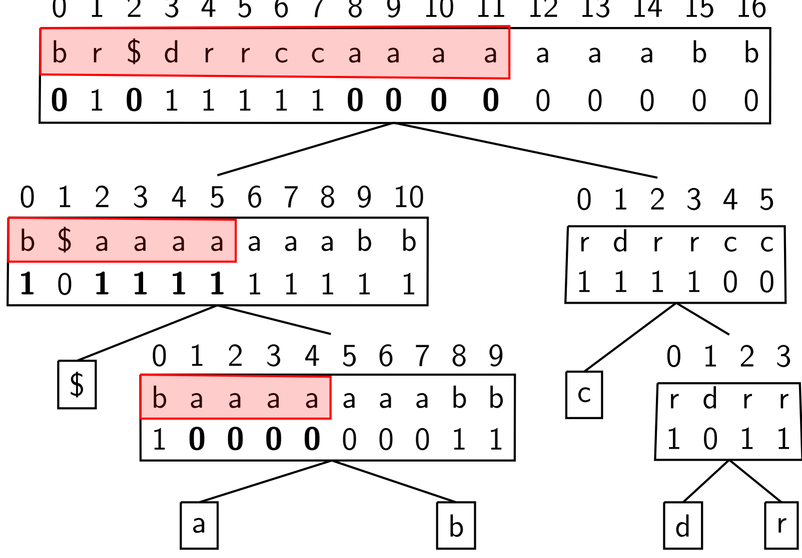

Wavelet Trees - Space Usage

Currently: $n \log \sigma + o(n \log \sigma)$ bits. Still larger than the original text!

How can we do better?

- Compressed bitvectors

Wavelet Trees - Space Usage

Currently: $n \log \sigma + o(n \log \sigma)$ bits. Still larger than the original text!

How can we do better?

- Picking the codewords for each symbol smarter!

Wavelet Trees - Space Usage

Currently

Symbol	Freq	Codeword
$	1	00
a	7	010
b	3	011
c	2	10
d	1	110
r	3	111

Bits per symbol: 2.82

Huffman Shape:

Symbol	Freq	Codeword
$	1	1100
a	7	0
b	3	101
c	2	111
d	1	1101
r	3	100

Bits per symbol: 2.29

Space usage of Huffman shaped wavelet tree:
$H_0(T)n + o(H_0(T)n)$ bits.

Even better: Huffman shape + compressed bitvectors

CSA-WT - Space Usage in practice

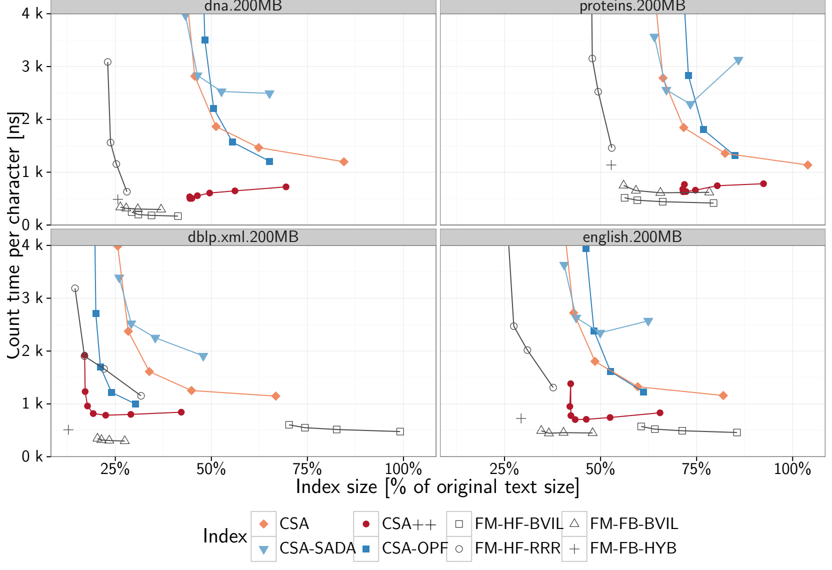

CSA-WT - Space Usage in practice (WORDS)

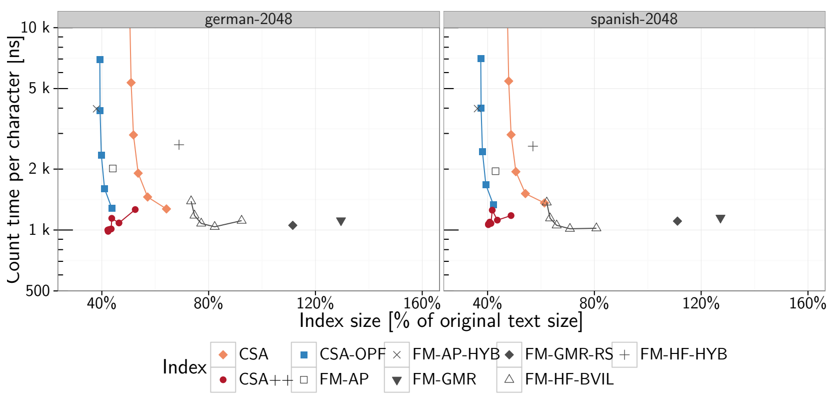

CSA-WT - Trade-offs in SDSL

```
1  #include "sdsl/suffix_arrays.hpp"
2  #include "sdsl/bit_vectors.hpp"
3  #include "sdsl/wavelet_trees.hpp"
4
5  int main(int argc,char** argv) {
6      std::string input_file = argv[1];
7      // use a compressed bitvector
8      using bv_type = sdsl::hyb_vector<>;
9      // use a huffman shaped wavelet tree
10     using wt_type = sdsl::wt_huff<bv_type>;
11     // use a wt based CSA
12     using csa_type = sdsl::csa_wt<wt_type>;
13     csa_type csa;
14     sdsl::construct(csa,input_file,1);
15     sdsl::store_to_file(csa,out_file);
16 }
```

CSA-WT - Trade-offs in SDSL

```
1  // use a regular bitvector
2  using bv_type = sdsl::bit_vector;
3  // 5% overhead rank structure
4  using rank_type = sdsl::rank_support_v5<1>;
5  // don't need select so we just use
6  // scanning which is O(n)
7  using select1_type = sdsl::select_support_scan<1>;
8  using select0_type = sdsl::select_support_scan<0>;
9  // use a huffman shaped wavelet tree
10 using wt_type  = sdsl::wt_huff<bv_type,
11                                 rank_type,
12                                 select1_type,
13                                 select0_type>;
14 using csa_type = sdsl::csa_wt<wt_type>;
15 csa_type csa;
16 sdsl::construct(csa,input_file,1);
17 sdsl::store_to_file(csa,out_file);
```

CSA-WT - Searching

```
1  int main(int argc,char** argv) {
2      std::string input_file = argv[1];
3      sdsl::csa_wt<> csa;
4      sdsl::construct(csa,input_file ,1);
5
6      std::string pattern = "abr";
7      auto nocc = sdsl::count(csa,pattern);
8      auto occs = sdsl::locate(csa,pattern);
9      for(auto& occ : occs) {
10         std::cout << "found at pos "
11                   << occ << std::endl;
12     }
13     auto snippet = sdsl::extract(csa,5,12);
14     std::cout << "snippet = '"
15               << snippet << "'" << std::endl;
16 }
```

CSA-WT - Searching - UTF-8

```
sdsl::csa_wt<> csa; //
sdsl::construct(csa, "this-file.cpp", 1);
std::cout << "count("") : "
     << sdsl::count(csa, "") << endl;
auto occs = sdsl::locate(csa, "\n");
sort(occs.begin(), occs.end());
auto max_line_length = occs[0];
for (size_t i=1; i < occs.size(); ++i)
    max_line_length = std::max(max_line_length,
                               occs[i]-occs[i-1]+1);
std::cout << "max line length : "
          << max_line_length << endl;
```

CSA-WT - Searching - Words

32 bit integer words:

```
sdsl::csa_wt_int<> csa;
// file containing uint32_t ints
sdsl::construct(csa, "words.u32", 4);
std::vector<uint32_t> pattern = {532432,43433};
std::cout << "count() : "
          << sdsl::count(csa,pattern) << endl;
```

$\log_2 \sigma$ bit words in SDSL format:

```
sdsl::csa_wt_int<> csa;
// file containing a serialized sdsl::int_vector ints
sdsl::construct(csa, "words.sdsl", 0);
std::vector<uint32_t> pattern = {532432,43433};
std::cout << "count() : "
          << sdsl::count(csa,pattern) << endl;
```

Compressed Suffix Trees

- Compressed representation of a Suffix Tree
- Internally uses a CSA
- Store extra information to represent tree shape and node depth information
- Three different CST types available in SDSL

Compressed Suffix Trees - CST

- Use a succinct tree representation to store suffix tree shape
- Compress the LCP array to store node depth information

Operations:
`root`, `parent`, `first_child`, `iterators`, `sibling`, `depth`, `node_depth`, `edge`, `children`... many more!

CST - Example

```
1  using csa_type = sdsl::csa_wt<>;
2  sdsl::cst_sct3<csa_type> cst;
3  sdsl::construct_im(cst, "ananas", 1);
4  for (auto v : cst) {
5      cout << cst.depth(v) << "-[" << cst.lb(v) << ","
6           << cst.rb(v) << "]" << endl;
7  }
8  auto v = cst.select_leaf(2);
9  for (auto it = cst.begin(v); it != cst.end(v); ++it) {
10     auto node = *it;
11     cout << cst.depth(v) << "-[" << cst.lb(v) << ","
12          << cst.rb(v) << "]" << endl;
13 }
14 v = cst.parent(cst.select_leaf(4));
15 for (auto it = cst.begin(v); it != cst.end(v); ++it) {
16     cout << cst.depth(v) << "-[" << cst.lb(v) << ","
17          << cst.rb(v) << "]" << endl;
18 }
```

CST - Space Usage Visualization

http://simongog.github.io/assets/data/space-vis.html

Applications to NLP and IR (30 Mins)

1. Applications to NLP and IR
2. LM fundamentals
3. LM complexity
4. LMs meet SA/ST
5. Experiments
6. Applications to IR

Application to NLP: language modelling

1 Applications to NLP and IR

2 LM fundamentals

3 LM complexity

4 LMs meet SA/ST

5 Experiments

6 Applications to IR

Language models

Definition

A language model defines probability $P(w_i|w_1, \ldots, w_{i-1})$, often with a Markov assumption, i.e., $P \approx P^{(k)}(w_i|w_{i-k}, \ldots, w_{i-1})$.

Example: MLE for k-gram LM

$$P^{(k)}(w_i|w_{i-k}^{i-1}) = \frac{c(w_{i-k}^{i})}{c(w_{i-k}^{i-1})}$$

- using count of context, $c(w_{i-k}^{i-1})$; and
- count of full k-gram, $c(w_{i-k}^{i})$

Notation: $w_i^j \triangleq (w_i, w_{i+1}, \ldots, w_j)$

Smoothed count-based LMs (Kneser Ney)

Kneser Ney probability compution requires the following:

$$c(w_i^j) \quad \text{basic counts}$$

$$\left.\begin{array}{c} N_{1+}(w_i^j \bullet) \\ N_{1+}(\bullet\, w_i^j) \\ N_{1+}(\bullet\, w_i^j \bullet) \\ N_1(w_i^j \bullet) \\ N_2(w_i^j \bullet) \end{array}\right\} \text{occurrence counts}$$

Other smoothing methods also require forms of occurrence counts, e.g., Good-Turing, Witten-Bell.

Construction and querying

Probabilities computed ahead of time

- Calculate a static hashtable or trie mapping k-grams to their probability and backoff values.
- Big: number of possible & observed k-grams grows with k

Querying

Lookup the longest matching span including the current token, and without the token. Probability computed from the full score and context backoff.

Query cost German Europarl, KenLM trie

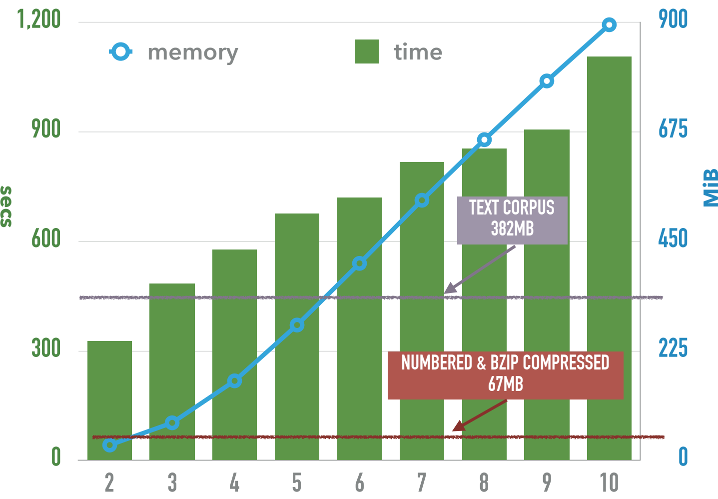

Precomputing versus on-the-fly

Precomputing approach

- Does not scale gracefully to high order m;
- Large training corpora also problematic

Can be computed directly from a CST

- CST captures unlimited order k-grams (no limit on m);
- Many (but not all) statistics cheap to retrieve
- LM probabilities computed on-the-fly

Sufficient statistics captured in suffix structures

$$T = \texttt{abracadabracarab\$}$$

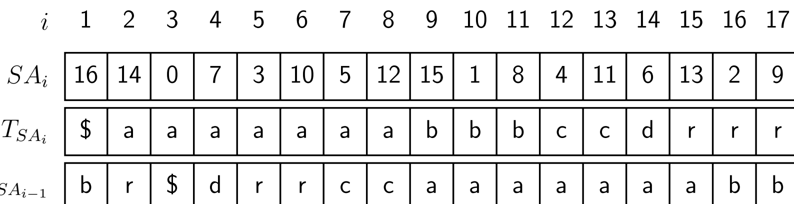

i	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
SA_i	16	14	0	7	3	10	5	12	15	1	8	4	11	6	13	2	9
T_{SA_i}	$	a	a	a	a	a	a	a	b	b	b	c	c	d	r	r	r
$T_{SA_{i-1}}$	b	r	$	d	r	r	c	c	a	a	a	a	a	a	a	b	b

Sufficient statistics captured in suffix structures

$$T = \texttt{abracadabracarab\$}$$

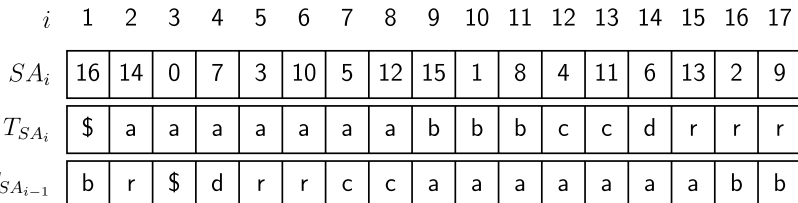

i	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
SA_i	16	14	0	7	3	10	5	12	15	1	8	4	11	6	13	2	9
T_{SA_i}	$	a	a	a	a	a	a	a	b	b	b	c	c	d	r	r	r
$T_{SA_{i-1}}$	b	r	$	d	r	r	c	c	a	a	a	a	a	a	a	b	b

Sufficient statistics captured in suffix structures

$$T = \texttt{abracadabracarab\$}$$

i	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
SA_i	16	14	0	7	3	10	5	12	15	1	8	4	11	6	13	2	9
T_{SA_i}	$	a	a	a	a	a	a	a	b	b	b	c	c	d	r	r	r
$T_{SA_{i-1}}$	b	r	$	d	r	r	c	c	a	a	a	a	a	a	a	b	b

- $c(\texttt{abra}) = 2$ from CSA
  range between $lb = 3$ and $rb = 4$, inclusive

Sufficient statistics captured in suffix structures

$T =$ abracadabracarab\$

i	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
SA_i	16	14	0	7	3	10	5	12	15	1	8	4	11	6	13	2	9
T_{SA_i}	\$	a	a	a	a	a	a	a	b	b	b	c	c	d	r	r	r
$T_{SA_{i-1}}$	b	r	\$	d	r	r	c	c	a	a	a	a	a	a	a	b	b

- $c(\texttt{abra}) = 2$ from CSA
  range between $lb = 3$ and $rb = 4$, inclusive
- $N_{1+}(\bullet\ \texttt{abra}) = 2$ from BWT (wavelet tree)
  size of set of preceeding symbols $\{\$, \text{d}\}$

Occurrence counts from the suffix tree

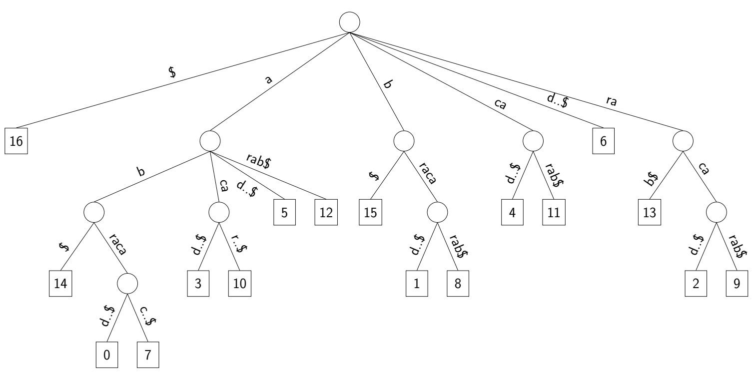

Number of proceeding symbols, $N_{1+}(\alpha\,\bullet)$, is either

Occurrence counts from the suffix tree

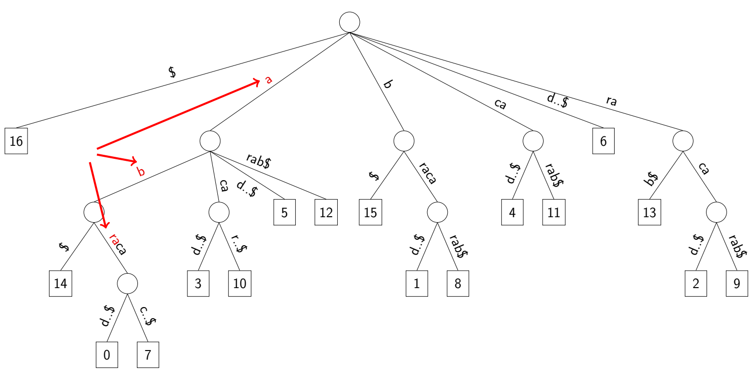

Number of proceeding symbols, $N_{1+}(\alpha\,\bullet)$, is either

- 1 if internal to an edge (e.g., $\alpha =$abra)

Occurrence counts from the suffix tree

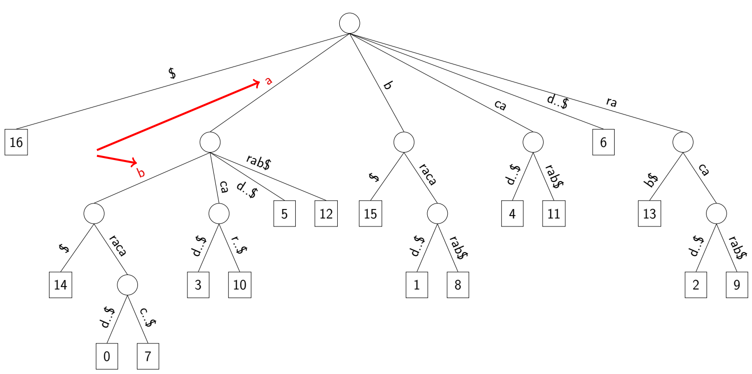

Number of proceeding symbols, $N_{1+}(\alpha\,\bullet)$, is either

- 1 if internal to an edge (e.g., $\alpha =$abra)
- $\mathrm{degree}(v)$ otherwise (e.g., $\alpha =$ab with degree 2)

Timing versus other LMs: Large DE Commoncrawl

construction load+query

100
10
1

Memory [GiB]

1 4 8 16 32 1 4 8 16 32

Input Size [GiB]

construction load+query

10k
1k
100

Time [seconds]

1 4 8 16 32 1 4 8 16 32

Input Size [GiB]

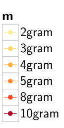

Applications to IR: Compressed Tries / Dictionaries

- Support $\textsc{lookup}(s)$ which returns unique id if string s is in dict or -1 otherwise
- Support $\textsc{retrieve}(i)$ return string with id i
- Very compact. $10\% - 20\%$ of original data
- Very fast lookup times
- Efficient construction
- MARISA trie: https://github.com/s-yata/marisa-trie
- MARISA trie stats: File: all page titles of English Wikipedia (Nov. 2012) - Size uncompressed: 191 MiB, Trie size: 48 MiB, gzip: 52 MiB

Range Minimum/Maximum Queries

- Given an array A of n items
- For any range $A[i, j]$ answer in constant time, what is the largest / smallest item in the range
- Space usage: $2n + o(n)$ bits. A not required!

Applications to IR: Top-k query completion

Conclusions / take-home message

- Basic succinct structures rely on bitvectors and operations RANK and SELECT
- More complex structures are composed of these basic building blocks
- Many trade-offs exist
- Practical, highly engineered open source implementations exist and can be used within minutes in industry and academia
- Other fields such as Information Retrieval, Bioinformatics have seen many papers using these succinct structures in recent years

Resources

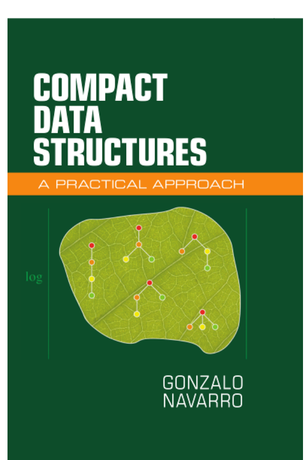

Compact Data Structures,
A practical approach
Gonzalo Navarro
ISBN 978-1-107-15238-0. 570 pages.
Cambridge University Press, 2016

Resources II

Full-day tutorial at SIGIR 2016:

Succinct Data Structures in Information Retrieval: Theory and Practice
Simon Gog and Rossano Venturini
727 slides!
More extensive coverage of different succinct structures.

Materials: http://pages.di.unipi.it/rossano/succinct-data-structures-in-information-retrieval-theory-and-practice/

Resources III

- Overview of compressed text indexes: [Ferragina et al., 2008, Navarro and Mäkinen, 2007]
- Bitvectors: [Gog and Petri, 2014]
- Document Retrieval: [Navarro, 2014a]
- Compressed Suffix Trees: [Sadakane, 2007, Ohlebusch et al., 2010]
- Wavelet Trees: [Navarro, 2014b]
- Compressed Tree Representations: [Navarro and Sadakane, 2016]

References I

Ferragina, P., González, R., Navarro, G., and Venturini, R. (2008).
Compressed text indexes: From theory to practice.
ACM J. of Exp. Algorithmics, 13.

Gog, S. and Petri, M. (2014).
Optimized succinct data structures for massive data.
Softw., Pract. Exper., 44(11):1287–1314.

Navarro, G. (2014a).
Spaces, trees and colors: The algorithmic landscape of document retrieval on sequences.
ACM Comp. Surv., 46(4.52).

Navarro, G. (2014b).
Wavelet trees for all.
Journal of Discrete Algorithms, 25:2–20.

References II

Navarro, G. and Mäkinen, V. (2007).
Compressed full-text indexes.
ACM Comp. Surv., 39(1):2.

Navarro, G. and Sadakane, K. (2016).
Compressed tree representations.
In *Encyclopedia of Algorithms*, pages 397–401.

Ohlebusch, E., Fischer, J., and Gog, S. (2010).
CST++.
In *Proceedings of the International Symposium on String Processing and Information Retrieval*.

Sadakane, K. (2007).
Compressed suffix trees with full functionality.
Theory of Computing Systems, 41(4):589–607.