<https://helda.helsinki.fi>

Compressed Spaced Suffix Arrays

Gagie, Travis

2017-06

Gagie , T , Manzini , G & Valenzuela , D 2017 , ' Compressed Spaced Suffix Arrays ' ,
Mathematics in Computer Science , vol. 11 , no. 2 , pp. 151-157 . <https://doi.org/10.1007/s11786-016-0283-z>

<http://hdl.handle.net/10138/312475>

<https://doi.org/10.1007/s11786-016-0283-z>

acceptedVersion

Downloaded from Helda, University of Helsinki institutional repository.

This is an electronic reprint of the original article.

This reprint may differ from the original in pagination and typographic detail.

Please cite the original version.

Compressed Spaced Suffix Arrays

Travis Gagie¹, Giovanni Manzini², and Daniel Valenzuela¹

¹ University of Helsinki, Finland

² University of Eastern Piedmont, Italy

Abstract. Spaced seeds are important tools for similarity search in bioinformatics, and using several seeds together often significantly improves their performance. With existing approaches, however, for each seed we keep a separate linear-size data structure, either a hash table or a spaced suffix array (SSA). In this paper we show how to compress SSAs relative to normal suffix arrays (SAs) and still support fast random access to them. We first prove a theoretical upper bound on the space needed to store an SSA when we already have the SA. We then present experiments indicating that our approach works even better in practice.

1 Introduction

For the problem of similarity search, we are given two texts and asked to find each sufficiently long substring of the first text that is within a certain Hamming distance of some substring of the second text. Similarity search has many applications in bioinformatics — e.g., ortholog detection, structure prediction or determining rearrangements — and has been extensively studied (see, e.g., [27]). Researchers used to first look for short substrings of the first text that occur unchanged in the second text, called seeds, then try to extend these short, exact matches in either direction to obtain longer, approximate matches. This approach is called, naturally enough, “seed and extend”. The substrings’ exact matches are found using either a hash table of the substrings with the right length, or an index structure such as a suffix array (SA).

Around the turn of the millenium, Burkhardt and Kärkkäinen [8] and Ma, Tromp and Li [22] independently proposed looking for short *subsequences* of the first text that have a certain shape and occur unchanged in the second text, and trying to extend those. A binary string encoding the shape of a subsequence, with 1s indicating positions where the characters must match and 0s indicating positions where they need not, is called a spaced seed. The total number of bits in the binary string is called the seed’s length, and the number of 1s is called its weight. The subsequences’ exact matches are found using either a hash table of the subsequences with the right shape, or a kind of modified SA called a spaced suffix array [18] (SSA).

Burkhardt and Kärkkäinen, Ma et al. and subsequent authors have shown that using spaced seeds significantly improves the performance of seeding and extending. Many papers have been written about how to design spaced seeds to minimize the number of errors (see, e.g., [7,13,17] and references therein), with the specifics depending on the model of sequence similarity and the acceptable numbers of false positives (for which the characters indicated by 1s all match but the substrings are not similar) and false negatives (for which those do not all match but the substrings are still similar) for the application in question. Regardless of the particular application, however, researchers have consistently observed that the best results are obtained using more than one seed at a time. A set of spaced seeds used in combination is called a multiple seed.

Multiple seeds are now a popular and powerful tool for similarity search, but they have a lingering flaw: we keep a hash table or SSA for each seed, and each instance of these data structures takes linear space. For example, SHRiMP2’s [12] index for the human genome takes 16 GB *for each seed*. In contrast, Bowtie 2’s [20] compressed SA for that genome takes only 2.5 GB. This is because a normal SA (which supports only substring matching) can be compressed

such that the number of bits per character is only slightly greater than the empirical entropy of the text. Unfortunately, the techniques for compressing normal SAs do not seem to apply directly to SSAs.

In this paper we show how to compress SSAs relative to normal SAs and still support fast random access to them. Whereas the normal SA for a text lists the starting points of the suffixes of that text by those suffixes' lexicographic order, the SSA for a text and a spaced seed lists the starting points of the subsequences with the right shape by those subsequences' lexicographic order. Intuitively, if the seed starts with many 1s, the SSA will be similar to the SA. In Section 2 we formalize this intuition and prove a theoretical upper bound on the space needed to store an SSA when we already have the SA, in terms of the text's length, the alphabet's size, and the seed's length and weight.

In Section 3 we present experiments showing that our approach works even better in practice. That is, even when we implement our data structures using simpler, theoretically sub-optimal components, we achieve better compression than our upper bound predicts. In fact, in practice we can even successfully apply our approach in some cases when the assumptions underlying our upper bounds are violated. However, we still want to improve our compression and random-access times for seeds with low weight-to-length ratios.

We recently learned that Peterlongo et al. [25] and Crochemore and Tischler [11] independently defined SSAs, under the names “bi-factor arrays” and “gapped suffix arrays”, for the special case in which the spaced seed has the form $1^a 0^b 1^c$. Russo and Tischler [26] showed how to represent such an SSA in asymptotically succinct space such that we can support random access to it in time logarithmic in the length of the text. We note, however, that the spaced seeds used for most applications do not have this form. We also recently learned that Battaglia et al. [4] used an idea similar to that of spaced seeds in an algorithm for finding motifs with don't-care symbols. It seems possible our results could be useful in reducing their algorithm's memory usage.

2 Theory

Suppose we want to store an SSA for a text $T[0..n-1]$ over an alphabet of size σ and a spaced seed S with length ℓ and weight w . For $i < n$, let T_i be the subsequence of $T[i..n-1]$ that contains $T[j]$ if and only if $i \leq j$ and $S[j-i] = 1$. Let T'_i be the subsequence of $T[i..n-1]$ that contains $T[j]$ if and only if $S[j-i] = 0$. Let SSA be the permutation on $\{0, \dots, n-1\}$ in which i precedes i' if either $T_i \prec T_{i'}$, or $T_i = T_{i'}$ and $T[i..n-1] \prec T[i'..n-1]$.

For example, if $T = \text{abracadabra}$ and $S = 101$ then

$T_0 = \text{ar}$	$T_6 = \text{db}$	$T'_0 = \text{b}$	$T'_6 = \text{a}$
$T_1 = \text{ba}$	$T_7 = \text{ar}$	$T'_1 = \text{r}$	$T'_7 = \text{b}$
$T_2 = \text{rc}$	$T_8 = \text{ba}$	$T'_2 = \text{a}$	$T'_8 = \text{r}$
$T_3 = \text{aa}$	$T_9 = \text{r}$	$T'_3 = \text{c}$	$T'_9 = \text{a}$
$T_4 = \text{cd}$	$T_{10} = \text{a}$	$T'_4 = \text{a}$	
$T_5 = \text{aa}$		$T'_5 = \text{d}$	

and so $\text{SSA} = [10, 3, 5, 7, 0, 8, 1, 4, 6, 9, 2]$, while $\text{SA} = [10, 7, 0, 3, 5, 8, 1, 4, 6, 9, 2]$.

If $T_i \preceq T_{i'}$ and $T'_i \preceq T'_{i'}$, then i precedes i' in both SSA and SA. In particular, if $T_i = T_{i'}$ or $T'_i = T'_{i'}$, then i and i' have the same relative order in SSA and SA. In our example, $T_3 = T_5 = \text{aa}$, so 3 precedes 5 in both SSA and SA; $T'_2 = T'_6 = T'_9 = \text{a}$, so 6 precedes 9 and 9 precedes 2 in both SSA and SA.

If we partition SSA into subsequences such that i and i' are in the same subsequence if and only if $T_i = T_{i'}$, then we can partition SA into the same subsequences. Since there are at most

$\sigma^w + w$ distinct strings T_i , our partitions each consist of at most $\sigma^w + w$ subsequences. Similarly, if we partition based on T'_i and T''_i , then our partitions each consist of at most $\sigma^{\ell-w} + \ell - w$ subsequences.

For our example, we can partition both SSA and SA into $[4, 6, 9, 2]$, for $T'_i = \mathbf{a}$; $[7, 0]$, for $T'_i = \mathbf{b}$; $[3]$, for $T'_i = \mathbf{c}$; $[5]$, for $T'_i = \mathbf{d}$; $[8, 1]$, for $T'_i = \mathbf{r}$; and $[10]$, for $T'_i = \mathbf{e}$. In this particular case, however, we could just as well partition both SSA and SA into only two common subsequences: e.g., $[10, 7, 0]$ and $[3, 5, 8, 1, 4, 6, 9, 2]$.

Consider the permutation $\text{SA}^{-1} \circ \text{SSA}$, which maps elements' positions in SSA to their positions in SA, and let ρ be the minimum number of increasing subsequences into which $\text{SA}^{-1} \circ \text{SSA}$ can be partitioned. Since any subsequence common to SSA and SA corresponds to an increasing subsequence in $\text{SA}^{-1} \circ \text{SSA}$, we have $\rho \leq \min(\sigma^w + w, \sigma^{\ell-w} + \ell - w)$. In our example, $\text{SA}^{-1} \circ \text{SSA} = [0, 3, 4, 1, 2, 5, 6, 7, 8, 9, 10]$ and $\rho = 2$.

Supowit [28] gave a simple algorithm that partitions $\text{SA}^{-1} \circ \text{SSA}$ into ρ increasing subsequences in $\mathcal{O}(n \lg \rho) \subseteq \mathcal{O}(n \min(w, \ell - w) \lg \sigma)$ time. When applied to $\text{SA}^{-1} \circ \text{SSA}$ in our example, Supowit's algorithm partitions it into $[0, 3, 4]$ and $[1, 2, 5, 6, 7, 8, 9, 10]$.

Barbay et al. [2] showed how, given a partition of $\text{SA}^{-1} \circ \text{SSA}$ into ρ increasing subsequences, we can store it in $(2 + o(1))n \lg \rho \leq (2 + o(1))n \min(w, \ell - w) \lg \sigma$ bits and support random access to it in $\mathcal{O}(\lg \lg \rho)$ time. Combining their ideas with later work by Belazzougui and Navarro [5], we can keep the same space bound and improve the time bound to $\mathcal{O}(1)$.

To do this, for $i \leq \rho$, we replace each element in the i th subsequence in $\text{SA}^{-1} \circ \text{SSA}$ by a character a_i , and store the resulting string R such that we can support random access to it and partial rank queries on it. We then permute R according to $\text{SA}^{-1} \circ \text{SSA}$ and store the resulting string R' such that we can support fast select queries on it. In our example, $R = a_1 a_2 a_2 a_1 a_1 a_2 a_2 a_2 a_2 a_2$ and $R' = a_1 a_1 a_1 a_2 a_2 a_2 a_2 a_2 a_2 a_2$.

The partial rank query $R.\text{p_rank}(i)$ returns the number of copies of $R[i]$ in $R[0..i]$, and the select query $R'.\text{select}_a(i)$ returns the position of the i th copy of a in R' . Barbay et al. noted that, for $i < n$,

$$(\text{SA}^{-1} \circ \text{SSA})[i] = R'.\text{select}_{R[i]}(R.\text{p_rank}(i)) .$$

Belazzougui and Navarro showed how we can store R in $(1 + o(1))n \lg \rho$ bits and support random access to it and partial rank queries on it in $\mathcal{O}(1)$ time, and store R' in $(1 + o(1))n \lg \rho$ bits and support select queries on it in $\mathcal{O}(1)$ time.

In summary, we can store $\text{SA}^{-1} \circ \text{SSA}$ in $(2 + o(1))n \min(w, \ell - w) \lg \sigma$ bits such that we can support random access to it in $\mathcal{O}(1)$ time. We will give a longer explanation in the full version of this paper. Since $\text{SSA} = \text{SA} \circ (\text{SA}^{-1} \circ \text{SSA})$, this gives us the following result:

Theorem 1. *Let $T[0..n - 1]$ be a text over an alphabet of size σ and let S be a spaced seed with length ℓ and weight w . If we have already stored the suffix array SA for T such that we can support random access to SA in time t_{SA} , then we can store a spaced suffix array SSA for T and S in $(2 + o(1))n \min(w, \ell - w) \lg \sigma$ bits such that we can support random access to SSA in $t_{\text{SA}} + \mathcal{O}(1)$ time.*

3 Practice

Theorem 1 says we can store SSAs for the human Y-chromosome `chrY.fa` in FASTA format (about 60 million characters over an alphabet of size 5) and SHRiMP2's three default spaced seeds — i.e., 11110111101111, 1111011100100001111 and 1111000011001101111 — in about 560 MB, in addition to the SA, whereas storing the SSAs naïvely would take about 720 MB. Storing the SSAs packed such that each entry takes $\lceil \lg 60\,000\,000 \rceil = 26$ bits would reduce this to about 580 MB.

To test our approach, we built the SSAs as described in Section 2; computed $SA^{-1} \circ SSA$, R and R' for each SSA; and stored each copy of R or R' as a wavelet tree. We chose wavelet trees because they are simple to use and often more practical than the theoretically smaller and faster data structures mentioned in Section 2. We ran all our tests described in this section on a computer with a quad-core Intel Xeon CPU with 32 GB of RAM, running Ubuntu 12.04. We used a wavelet-tree implementation from <https://github.com/fclaude/libcds> and compiled it with GNU g++ version 4.4.3 with optimization flag `-O3`.

The uncompressed SA took 226 MB, and the six wavelet trees took a total of 215 MB and performed 10 000 random accesses each in 7.67 microseconds per access. That is, we compressed the SSAs into about 60% of the space it would take to store them naïvely and, although our accesses were much slower than direct memory accesses, they were fast compared to disk accesses. Thus, our approach seems likely to be useful when a set of SSAs is slightly larger than the memory and fits only when compressed.

Using the same test setup, we then compressed SSAs for the ten spaced seeds BFAST [15, Table S3] uses for 36-base-pair Illumina reads, which all have weight 18:

1. 111111111111111111
2. 11110100110111101010101111
3. 11111111111111001111
4. 1111011101100101001111111
5. 11110111000101010000010101110111
6. 1011001101011110100110010010111
7. 1110110010100001000101100111001111
8. 11110111111111111111
9. 11011111100010110111101101
10. 111010001110001110100011011111 .

Since the first seed consists only of 1s, the SSA we would build for it is the same as the SA. The uncompressed SA again took 226 MB and the 18 wavelet trees for the other nine seeds took a total of 649 MB — so instead of 2.26 GB, we used 875 MB (about 39%) for all ten seeds — and together performed 10 000 random accesses to each of the ten SSAs in about 7 microseconds per access. The left side of the top half of Figure 1 shows how many bits per character (bpc) of the text each SSA took, and the average time per access to each SSA.

We also compressed the SSAs for the ten spaced seeds BFAST uses for 50-base-pair Illumina reads, which all have weight 22:

1. 11111111111111111111
2. 1111101110111010100101011011111
3. 1011110101101001011000011010001111111
4. 10111001101001100100111101010001011111
5. 11111011011101111011111111
6. 111111100101001000101111101110111
7. 111101011110010100010101101010111111
8. 111101101011011001100000101101001011101
9. 1111011010001000110101100101100110100111
10. 1111010010110110101110010110111011 .

Again, the first seed consists only of 1s. This time, the 18 wavelet trees for the other nine seeds took a total of 712 MB; each access took about 8 microseconds. The left side of the bottom half of Figure 1 shows how many bit per character of the text each SSA took, and the average access time per access to each SSA.

If we have a permutation π_1 on $\{0, \dots, n-1\}$ stored and π_2 is any other permutation on $\{0, \dots, n-1\}$, then we can store π_2 relative to π_1 using the ideas from Section 2. For example, we can store SSAs relative to other SSAs. Suppose we consider the size of each SSA (except the SA) when compressed relative to each other SSA (including the SA), build a minimum spanning tree rooted at the SA, and compress each SSA relative to its parent in the tree. This can reduce

seed	space (bpc)	time (μ s)	seed	reference	space (bpc)	time (μ s)
1	32.00	0	1	-	32.00	0
2	11.29	9	2	8	9.71	11
3	4.41	4	3	1	4.41	4
4	9.75	8	4	8	9.22	10
5	11.54	9	5	4	9.23	19
6	13.77	11	6	8	12.27	14
7	13.14	10	7	3	12.58	14
8	3.85	3	8	1	3.85	3
9	10.10	7	9	1	10.10	7
10	13.91	11	10	7	12.59	26

seed	space (bpc)	time (μ s)	seed	reference	space (bpc)	time (μ s)
1	32.00	0	1	-	32.00	0
2	9.03	8	2	1	9.03	6
3	12.30	10	3	1	12.30	10
4	13.86	11	4	2	12.59	18
5	8.13	7	5	1	8.13	6
6	10.80	9	6	1	10.80	8
7	11.14	8	7	1	11.14	8
8	11.09	8	8	1	11.09	9
9	11.77	9	9	8	11.34	18
10	12.54	10	10	8	8.94	17

Fig. 1. The space usage of the SSAs of the spaced seeds BFAST uses for Illumina reads, in bits per character of the text, and the average time for a random access. On top, the seeds are for 36-base-pair reads; on the bottom, the seeds are for 50-base-pair reads. On the left, all the SSAs are compressed relative to the SA; on the right, some of the SSAs are compressed relative to other SSAs.

our space usage at the cost of increasing the random-access time, as shown for the BFAST seeds on the right side of Figure 1.

There are other circumstances in which we can ignore SSAs’ semantics and consider them only as permutations. For example, spaced seeds can be generalized to subset seeds [19], such as ternary strings in which 1s indicate positions where the characters must match, 0s indicate positions where they need not, and Ts indicate positions where characters must fall within the same equivalence class (such as the pyrimidines C and T and the purines A and G). It is not difficult to generalize Theorem 1 to subset seeds — we will do so in the full version of this paper — but it is also not necessary to obtain practical results. The *Iedera* tool (available at <http://bioinfo.lifl.fr/yass/iedera.php>) generates good subset seeds.

A more challenging change is from fixed-length seeds to repetitive seeds [18]. A repetitive seed is a string in whose repetition the digits indicate which characters must match and how. For example, with respect to the repetitive spaced seed 10110, **ATCGATCGGT** matches **ACCGTTGGGA** but not **ACCGTTGAGA**. Repetitive seeds are useful when looking for approximate matches of substrings that have been extended until they become sufficiently infrequent. It is not clear how or if we can extend Theorem 1 to repetitive seeds. Nevertheless, the *LAST* tool (available at <http://last.cbrc.jp>) generates SSAs for repetitive spaced or subset seeds, which we can still try to compress in practice; see also [16,24].

Our current goal is to achieve reasonable compression and access times for a set of repetitive subset seeds that we received from Martin Frith, which have average length 19.85 and average weight about 10.44, counting “same equivalence class” digits as 0.5. Unfortunately, at the moment we use nearly 24 bits per entry in the corresponding SSAs (including the overhead for the uncompressed SA), which is only marginally better than the 26 bits we would use with simple packing. Meanwhile, random accesses take about 12 microseconds on average, which is significantly slower than access to a packed array. On the other hand, these seeds have an unusually low average weight-to-length ratio. We used *Iedera* and *LAST* to generate SSAs for a set of

eight repetitive subset seeds, with average length 17.875 and average weight 12. For these, we used only 20.15 bits per entry, with random accesses taking about 10 microseconds on average.

4 New Directions

Barbay et al.'s result holds even when we are given a partition of a permutation into ρ increasing or decreasing subsequences, and some authors [1,10] have found that using both increasing and decreasing subsequences often improves compression in practice. Computing a partition into the minimum number of such subsequences is NP-hard, however, and we see no reason why $SA^{-1} \circ SSA$ should contain long decreasing subsequences. Therefore, in this paper we considered only increasing subsequences and did not discuss other papers on compressing permutations (see [3] and references therein).

Nevertheless, while writing this paper we noticed a way to generalize slightly Barbay et al.'s result. Suppose we partition a permutation $\pi[0..n-1]$ into subsequences $\tau_0, \dots, \tau_{\rho-1}$. Let $R[0..n-1]$ be the bitvector in which $R[i] = j$ if $\pi[i]$ is in τ_j . Let $R'[0..n-1]$ be the bitvector in which $R'[i] = j$ if i is in τ_j . Let $\pi_0[0..|\tau_0|-1], \dots, \pi_{\rho-1}[0..|\tau_{\rho-1}|-1]$ be the permutations in which $\pi_i[j]$ is the number of elements strictly smaller than $\tau_i[j]$ in τ_i . Then

$$\pi[i] = R'.\text{select}_{R[i]} \left(\pi_{R[i]} \left[R.\text{p_rank}(i) - 1 \right] + 1 \right).$$

Therefore, if we can store $\pi_0, \dots, \pi_{\rho-1}$ in small space, then we can store π in small space. Notice also that, conversely,

$$\pi_j[i] = R'.\text{p_rank} \left(\pi \left[R.\text{select}_j(i+1) \right] \right) - 1;$$

this will be useful later.

For example, if

$$\begin{aligned} \pi[0..19] &= [16, 17, 12, 10, 6, 1, 9, 15, 18, 4, 14, 13, 5, 11, 19, 2, 8, 7, 0, 3] \\ \tau_0[0..8] &= [16, 12, 6, 1, 15, 13, 5, 2, 0] \\ \tau_1[0..10] &= [17, 10, 9, 18, 4, 14, 11, 19, 8, 7, 3] \end{aligned}$$

then

$$\begin{aligned} R[0..19] &= 01010010111001101101 \\ R'[0..19] &= 00011001111100100111 \\ \pi_0[0..8] &= [8, 5, 4, 1, 7, 6, 3, 2, 0] \\ \pi_1[0..10] &= [8, 5, 4, 9, 1, 7, 6, 10, 3, 2, 0] \end{aligned}$$

and thus, say,

$$\begin{aligned} \pi[6] &= R'.\text{select}_{R[6]} \left(\pi_{R[6]} \left[R.\text{p_rank}(i) - 1 \right] + 1 \right) \\ &= R'.\text{select}_1(\pi_1[2] + 1) \\ &= R'.\text{select}_1(5) \\ &= 9. \end{aligned}$$

Therefore, since the elements in π_0 and π_1 appear in lexicographic order by name, we can store π in small space.

This may seem at first like rather a pointless generalization, but consider the case in which we want to store a permutation π that is similar to a permutation $\hat{\pi}$ that we already have

stored. By similar, we mean here that we can find a long subsequences in π and $\hat{\pi}$ such that the elements in those subsequences have the same relative order. In this case, we can use $\hat{\pi}$ to compress π .

For example, if

$$\begin{aligned}\pi[0..12] &= [12, 11, 6, 3, 8, 0, 5, 7, 4, 9, 1, 10, 2] \\ \hat{\pi}[0..11] &= [11, 10, 7, 0, 3, 5, 8, 1, 4, 6, 9, 2]\end{aligned}$$

then the elements in their subsequences $[12, 11, 8, 0, 9, 1, 10, 2]$ and $[11, 10, 7, 0, 8, 1, 9, 2]$ have the same relative order, $[7, 6, 3, 0, 4, 1, 5, 2]$. Let

$$\begin{aligned}\tau_0[0..7] &= [12, 11, 8, 0, 9, 1, 10, 2] \\ \tau_1[0..4] &= [6, 3, 5, 7, 4] \\ \hat{\tau}_0[0..7] &= [11, 10, 7, 0, 8, 1, 9, 2] \\ \hat{\tau}_1[0..3] &= [3, 5, 4, 6]\end{aligned}$$

so

$$\begin{aligned}R[0..12] &= 0011001110000 \\ R'[0..12] &= 0001111100000 \\ \hat{R}[0..11] &= 000011001100 \\ \hat{R}'[0..11] &= 000111100000 \\ \pi_0[0..7] &= \hat{\pi}_0[0..7] = [7, 6, 3, 0, 4, 1, 5, 2] \\ \pi_1[0..4] &= [3, 0, 2, 4, 1] \\ \hat{\pi}_1 &= [0, 2, 1, 3],\end{aligned}$$

where we use accents to distinguish the data related to $\hat{\pi}$ from those related to π .

Because we already have $\hat{\pi}$ stored, if we store $\hat{R}$ and $\hat{R}'$, then for $0 \leq i \leq 7$ we can compute in $\mathcal{O}(1)$ time

$$\pi_0[i] = \hat{\pi}_0[i] = \hat{R}'.\text{p_rank}\left(\hat{\pi}\left[\hat{R}.\text{select}_0(i+1)\right]\right) - 1.$$

Therefore, if we also store R , R' and π_1 , then we can support access to π in $\mathcal{O}(1)$ time. That is, we can store $\pi[0..12]$ as $\tau_1[0..4]$ and four bitvectors.

We chose $\pi = [12, 11, 6, 3, 8, 0, 5, 7, 4, 9, 1, 10, 2]$ and $\hat{\pi} = [11, 10, 7, 0, 3, 5, 8, 1, 4, 6, 9, 2]$ because they are the SAs of `abrabbababra$` and `abracadabra$`. Thus, our example illustrates how we can usually store the SA of a slightly modified version of a document in small space when we already have the SA of the original: we look for a long subsequence of the new SA in which the elements have the same order as those in a subsequence of the original SA, and store the complement of that subsequence and the appropriate four bitvectors. Finding the longest subsequence to use is NP-hard, by a reduction from permutation pattern matching [6], but Léonard, Mouchard and Salson [21] showed how we can find a reasonably long one in the average case and in practice.

As a very preliminary test, we built the SA for the Wikipedia page for “Suffix array” (http://en.wikipedia.org/wiki/Suffix_array, downloaded on February 13th, 2014) and the SA for the same page with the order of the bullet points swapped in the paragraph

“Suffix arrays are closely related to suffix trees:

- Suffix arrays can be constructed by performing a depth-first traversal of a suffix tree. The suffix array corresponds to the leaf-labels given in the order in which these are visited during the traversal, if edges are visited in the lexicographical order of their first character.

- A suffix tree can be constructed in linear time by using a combination of suffix and LCP array. For a description of the algorithm, see the corresponding section in the LCP array article.”

Using the approach described above, we were able to store the SA of the modified page in about 0.4 bits per character, on top of the SA of the original page. In comparison, compressing the modified page with `p7zip` (available at <https://packages.debian.org/sid/p7zip-full>) used about 1.85 bits per character.

We hope our approach will prove useful for storing and indexing, e.g., many separate human genomes. (We note that indexing the concatenation of many human genomes has been studied previously; see, e.g., [23].) Since human genomes are similar to each other, we should be able to support fast access to a (possibly sampled) SA for each genome while using reasonable total space. Moreover, most of the genomes’ Burrows-Wheeler Transforms [9] (BWTs) should share a long common subsequence, so we should be able to store also rank and select data structures for them in reasonable total space: for each BWT, we store a bitvector indicating which characters are part of the common subsequence, and a rank and select data structure for those characters not in the common subsequence; we then need store only a single rank and select data structure for the common subsequence. Therefore, we should be able to store an FM-index [14] for each genome using reasonable total space. Investigating this possibility is beyond the scope of this paper, however, so we leave it as future work.

Acknowledgments

Many thanks to Francisco Claude, Maxime Crochemore, Matei David, Martin Frith, Costas Iliopoulos, Juha Kärkkäinen, Gregory Kucherov, Bin Ma, Ian Munro, Taku Onodera, Gonzalo Navarro, Luis Russo, German Tischler and the anonymous reviewers.

References

1. D. Arroyuelo, F. Claude, R. Dorigiv, S. Durocher, M. He, A. López-Ortiz, J. I. Munro, P. K. Nicholson, A. Salinger, and M. Skala. Untangled monotonic chains and adaptive range search. *Theoretical Computer Science*, 432:4200–4211, 2011.
2. J. Barbay, F. Claude, T. Gagie, G. Navarro, and Y. Nekrich. Efficient fully-compressed sequence representations. *Algorithmica*. To appear.
3. J. Barbay and G. Navarro. On compressing permutations and adaptive sorting. *Theoretical Computer Science*, 513:109–123, 2013.
4. G. Battaglia, D. Cangelosi, R. Grossi, and N. Pisanti. Masking patterns in sequences: A new class of motif discovery with don’t cares. *Theoretical Computer Science*, 410:4327–4340, 2009.
5. D. Belazzougui and G. Navarro. Alphabet-independent compressed text indexing. *ACM Transactions on Algorithms*. To appear.
6. P. Bose, J. F. Buss, and A. Lubiw. Pattern matching for permutations. *Information Processing Letters*, 65:277–283, 1998.
7. D. G. Brown. A survey of seeding for sequence alignment. In I. Măndoiu and A. Zelikovsky, editors, *Bioinformatics Algorithms: Techniques and Applications*, pages 126–152. Wiley-Interscience, 2008.
8. S. Burkhardt and J. Kärkkäinen. Better filtering with gapped q-grams. *Fundamenta Informaticae*, 56:51–70, 2003.
9. M. Burrows and D. J. Wheeler. A block sorting lossless data compression algorithm. Technical Report 124, Digital Equipment Corporation, 1994.
10. F. Claude and J. I. Munro. Adaptive data structures for permutations and binary relations. In *Proceedings of the 20th Symposium on String Processing and Information Retrieval (SPIRE)*, pages 64–71, 2013.
11. M. Crochemore and G. Tischler. The gapped suffix array: A new index structure for fast approximate matching. In *Proceedings of the 17th Symposium on String Processing and Information Retrieval (SPIRE)*, pages 359–364, 2010.
12. M. David, M. Dzamba, D. Lister, L. Ilie, and M. Brudno. SHRiMP2: Sensitive yet practical short read mapping. *Bioinformatics*, 27:1011–1012, 2011.

13. L. Egidi and G. Manzini. Better spaced seeds using quadratic residues. *Journal of Computer and System Sciences*, 79:1144–1155, 2013.
14. P. Ferragina and G. Manzini. Indexing compressed text. *Journal of the ACM*, 52:552–581, 2005.
15. N. Homer, B. Merriman, and S. F. Nelson. BFAST: An alignment tool for large scale genome resequencing. *PLOS One*, 4:e7767, 2009.
16. P. Horton, S. M. Kielbasa, and M. C. Frith. DisLex: A transformation for discontinuous suffix array construction. In *Proceedings of the Workshop on Knowledge, Language, and Learning in Bioinformatics (KLLBI)*, pages 1–11, 2008.
17. L. Ilie, S. Ilie, S. Khoshraftar, and A. Mansouri Bigvand. Seeds for effective oligonucleotide design. *BMC Genomics*, 12:280, 2011.
18. S. M. Kielbasa, R. Wan, K. Sato, P. Horton, and M. C. Frith. Adaptive seeds tame genomic sequence comparison. *Genome Research*, 21:487–493, 2011.
19. G. Kucherov, L. Noé, and M. A. Roytberg. A unifying framework for seed sensitivity and its application to subset seeds. *Journal of Bioinformatics and Computational Biology*, 4:553–570, 2006.
20. B. Langmead and S. L. Salzberg. Fast gapped-read alignment with Bowtie 2. *Nature Methods*, 9:357–359, 2012.
21. M. Léonard, L. Mouchard, and M. Salson. On the number of elements to reorder when updating a suffix array. *Journal of Discrete Algorithms*, 11:87–99, 2012.
22. B. Ma, J. Tromp, and M. Li. PatternHunter: faster and more sensitive homology search. *Bioinformatics*, 18:440–445, 2002.
23. V. Mäkinen, G. Navarro, J. Sirén, and N. Välimäki. Storage and retrieval of highly repetitive sequence collections. *Journal of Computational Biology*, 17:281–308, 2010.
24. T. Onodera and T. Shibuya. An index structure for spaced seed search. In *Proceedings of the 22nd International Symposium on Algorithms and Computation (ISAAC)*, pages 764–772, 2011.
25. P. Peterlongo, N. Pisanti, F. Boyer, and M.-F. Sagot. Lossless filter for finding long multiple approximate repetitions using a new data structure, the bi-factor array. In *Proceedings of the 12th Symposium on String Processing and Information Retrieval (SPIRE)*, pages 179–190, 2005.
26. L. M. S. Russo and G. Tischler. Succinct gapped suffix arrays. In *Proceedings of the 18th Symposium on String Processing and Information Retrieval (SPIRE)*, pages 290–294, 2011.
27. Y. Sun and J. Buhler. Designing multiple simultaneous seeds for DNA similarity search. *Journal of Computational Biology*, 12:847–861, 2005.
28. K. J. Supowit. Decomposing a set of points into chains, with applications to permutation and circle graphs. *Information Processing Letters*, 21:249–252, 1985.