

Minimal Indices for Successor Search

[Full Version]

Sarel Cohen, Amos Fiat, Moshik Hershcovitch, and Haim Kaplan

School of Computer Science, Tel Aviv University
{sarelcoh, fiat, haimk}@post.tau.ac.il
moshik1@gmail.com

Abstract. We give a new successor data structure which improves upon the index size of the Pătraşcu-Thorup data structures, reducing the index size from $O(nw^{4/5})$ bits to $O(n \log w)$ bits, with optimal probe complexity. Alternatively, our new data structure can be viewed as matching the space complexity of the (probe-suboptimal) z -fast trie of Belazzougui et al. Thus, we get the best of both approaches with respect to both probe count and index size. The penalty we pay is an extra $O(\log w)$ inter-register operations. Our data structure can also be used to solve the weak prefix search problem, the index size of $O(n \log w)$ bits is known to be optimal for any such data structure.

The technical contributions include highly efficient single word indices, with out-degree $w/\log w$ (compared to the $w^{1/5}$ out-degree of fusion tree based indices). To construct such high efficiency single word indices we devise highly efficient bit selectors which, we believe, are of independent interest.

Keywords: Predecessor Search, Succinct Data Structures, Cell Probe Model, Fusion Trees, Tries, Word RAM model

1 Introduction

A fundamental problem in data structures is the successor problem: given a RAM with w bit word operations, and n keys (each w bits long), give a data structure that answers successor queries efficiently. We distinguish between the space occupied by the n input keys themselves, which is $O(nw)$ bits, and the additional space requires by the data structure which we call the *index*. The two other performance measures of the data structure which are of main interest are how many accesses to memory (called *probes*) it performs per query, and the query time or the total number of machine operations performed per query, which could be larger than the number of probes. We can further distinguish between probes to the index and probes to the input keys themselves. The motivation is that if the index is small and fits in cache probes to the index would be cheaper. We focus on constructing a data structure for the successor problem that requires sublinear $o(nw)$ extra bits.

The simplest successor data structure is a sorted list, this requires no index, and performs $O(\log n)$ probes and $O(\log n)$ operations per binary search.

This high number of probes that are widely dispersed can makes this solution inefficient for large data sets.

Fusion trees of Fredman and Willard [9] (see also [10]) reduce the number of probes and time to $O(\log_w n)$. A fusion tree node has outdegree $B = w^{1/5}$ and therefore fusion trees require only $O(nw/B) = O(nw^{4/5})$ extra bits.

Another famous data structure is the y -fast trie of Willard [17]. It requires linear space ($O(nw)$ extra bits) and $O(\log w)$ probes and time per query.

Pătraşcu and Thorup [13] solve the successor problem optimally (to within an $O(1)$ factor) for any possible point along the probe count/space tradeoff, and for any value of n and w . However, they do not distinguish between the space required to store the input and the extra space required for the index. They consider only the total space which cannot be sublinear.

Pătraşcu and Thorup’s linear space data structure for successor search is an improvement of three previous data-structures and achieves the following bounds.

1. For values of n such that $\log n \in [0, \frac{\log^2 w}{\log \log w}]$ their data structure is a fusion tree and therefore the query time is $O(\log_w n)$. This bound increases monotonically with n .
2. For n such that $\log n \in [\frac{\log^2 w}{\log \log w}, \sqrt{w}]$ their data structure is a generalization of the data structure of Beame & Fich [2] that is suitable for linear space, and has the bound $O(\frac{\log w}{\log \log w - \log \log \log n})$. This bound increases from $O(\frac{\log w}{\log \log w})$ at the beginning of this range to $O(\log w)$ at the end of the range.
3. For values of n such that $\log n \in [\sqrt{w}, w]$ their data structure is a slight improvement of the van Emde Boas (vEB) data structure [16] and has the bound of $O(\max\{1, \log(\frac{w - \log n}{\log w})\})$. This bound decreases with n from $O(\log w)$ to $O(1)$.

A recent data structure of Belazzougui et al. [4] called the probabilistic z -fast trie, reduces the extra space requirement to $O(n \log w)$ bits, but requires a (sub-optimal) expected $O(\log w)$ probes (and $O(\log n)$ probes in the worst case). See Table 1 for a detailed comparison between various data structures for the successor problem with respect to the space and probe parameters under consideration.

Consider the following multilevel scheme to reduce index size: (a) partition the keys into consecutive sets of $w^{1/5}$ keys, (b) build a Fusion tree index structure for each such set (one w bit word), and (c) index the smallest key in every such group using any linear space data structure. The number of fusion tree nodes that we need $n/w^{1/5}$ and the total space required for these nodes and the data structure that is indexing them is $O(nw^{4/5})$.

This standard bucketing trick shows that we can get indices of smaller size by constructing a “fusion tree node” of larger outdegree. That is we seek a data structure, which we refer to as a *word-index*, that by using $O(1)$ words can answer successor queries with respect to as many keys as possible.

Data Structure	Ref.	Index size (in bits)		# Non-index Probes	Total # Probes	# operations
Binary Search		–		$O(\log n)$	$O(\log n)$	$O(\#\text{probes})$
van Emde Boas	[16]	$O(2^w)$		$O(1)$	$O(\log w)$	$O(\#\text{probes})$
x -fast trie	[17]	$O(nw^2)$		$O(1)$	$O(\log w)$	$O(\#\text{probes})$
y -fast trie	[17]	$O(nw)$		$O(1)$	$O(\log w)$	$O(\#\text{probes})$
x -fast trie on “splitters” $\text{poly}(w)$ apart	Folklore	$O(n/\text{poly}(w))$		$O(\log w)$	$O(\log w)$	$O(\#\text{probes})$
Beame & Fich	[2]	$\Theta(n^{1+\epsilon}w)$		$O(1)$	$O\left(\frac{\log w}{\log \log w}\right)$	$O(\#\text{probes})$
Fusion Trees	[9]	$O(nw^{4/5})$		$O(1)$	$O\left(\frac{\log n}{\log w}\right)$	$O(\#\text{probes})$
z -fast trie	[4,3,5]	$O(n \log w)$	$\frac{\text{exp.}}{\text{w.c.}}$	$\frac{O(1)}{O(\log n)}$	$\frac{O(\log w)}{O(\log n)}$	$O(\#\text{probes})$
Pătraşcu & Thorup	[13]	$O(nw)$ or $O(nw^{4/5})$		$O(1)$	Optimal given linear space	$O(\#\text{probes})$
Pătraşcu & Thorup + γ -nodes	This Paper	$O(n \log w)$		$O(1)$	Optimal given linear space	$O(\#\text{probes}$ $+ \log w)$

Table 1. Requirements of various data structures for the successor problem. The word length is w and the number of keys is n . Indexing groups of $w/\log w$ consecutive keys with our new word indices we can reduce the space of any of the linear space data structures above to $O(n \log w)$ bits while keeping the number of probes the same and increasing the query time by $O(\log w)$.

Our main contribution is such a word index that can handle $w/\log w$ keys (rather than $w^{1/5}$ for fusion trees).¹ However, this new highly compact index requires $\Theta(\log w)$ operations per search (versus the $O(1)$ operations required by Fusion trees).

Using these word indices we obtain, as described above, a (deterministic) data structure that, for any n, w , answers a successor query with an optimal number of probes (within an $O(1)$ factor), and requires only $O(n \log w)$ extra bits. We remark that we only probe $O(1)$ non-index words (which is true of Pătraşcu-Thorup data structures as well, with minor modifications). The penalty we pay is an additional $O(\log w)$ in the time complexity.

Indices of small size are particularly motivated today by the multicore shared memory architectures abundant today [7,14]. When multiple cores access shared cache/memory, contention arises. Such contention is deviously problematic because it may cause serialization of memory accesses, making a mockery of multicore parallelism. Multiple memory banks and other hardware are attempts to deal with such problems, to various degrees. Thus, the goals of reducing the index size, so it fits better in fast caches, reducing the number of probes extraneous to the index, and the number of probes within the index, become critical.

2 High level overview of our results and their implications

Computation model: We assume a RAM model of computation with w bits per word. A key (or query) is one word (w bits long). We can operate on the

¹ The $w/\log w$ keys take more than $O(1)$ words but are not considered part of the word index.

registers using at least a basic instruction set consisting of (as defined in [6]): Direct and indirect addressing, conditional jump, and a number of *inter-register operations*, including addition, subtraction, bitwise Boolean operations and left and right shifts. All operations are unit cost. One of our construction does not require multiplication.

We give three variants of high outdegree single word indices which we call α nodes, β nodes, and γ nodes. Each of these structures index $w/\log w$ keys and answer successor queries using only $O(1)$ w -bit words, $O(\log w)$ time, and $O(1)$ extra-index probes (in expectation for α and β nodes, worst case for γ nodes) to get at most two of the $w/\log w$ keys.

The α node is simply a z -fast trie ([4]) applied to $w/\log w$ keys. Given the small number of keys, the z -fast trie can be simplified. A major component of the z -fast trie involves translating a prefix match (between a query and a trie node) to the query rank. As there are only $w/\log w$ keys involved, we can discard this part of the z -fast trie and store ranks explicitly in $O(1)$ words.

Based on a different set of ideas, β nodes are arguably simpler than the z -fast trie, and have the same performance as the α nodes. As β -nodes are not our penultimate construction, the full description of β -nodes is in Appendix B.

Our penultimate variant, γ nodes, has the advantages that it is deterministic and gives worst case $O(1)$ non-index probes, and, furthermore, requires no multiplication.

To get the γ nodes we introduce highly efficient bit-selectors (see section 2.2) that may be of independent interest. Essentially, a bit-selector selects a multiset of bits from a binary input string and outputs a rearrangement of these bits within a shorter output string.

Thorup [15] proved that it is impossible to have $O(1)$ time successor search in a “standard AC(0) model”, for any non-constant number of keys, unless one uses enormous space, $2^{\Omega(w)}$, where w is the number of bits per word. This means that it would be impossible to derive an improved γ -node (or Fusion tree node) with $O(1)$ time successor search in the “standard AC(0) model”.

2.1 Succinct successor data structure

As mentioned in the introduction we obtain using our word indices a successor data structure that requires $O(n \log w)$ bits in addition to the input keys. The idea is standard and simple: We divide the keys into consecutive chunks of size $w/\log w$ keys each. We index each chunk with one of our word indices and index the chunks (that is the first key in each chunk) using another linear space data structure. This has the following consequences depending upon the linear space data structure which we use to index the chunks. (We henceforth refer to our γ -nodes, but similar results can be obtained using either α or β nodes in expectation.)

Fusion Trees + γ -nodes: This data structure answers successor queries with $O(\log_w n)$ probes, and $O(\log_w n + \log w)$ time.

The optimal structure of Pătraşcu & Thorup + γ -nodes: Here the number of probes to answer a query is optimal, the time is $O(\#probes + \log w)$.

y -fast-trie + γ -nodes: This gives an improvement upon the recently introduced [probabilistic] z -fast-trie, [4,3] (we omit the “probabilistic” prefix hereinafter). The worst-case probes and query time improves from $O(\log n)$ to $O(1)$ probes and $O(\log w)$ query time, and the data structure is deterministic.

The weak prefix search problem: In this problem the query is as follows. Given a bit-string p , such that p is the prefix of at least one key among the n input keys, return the range of ranks of those input keys having p as a prefix.

It is easy to modify the index of our successor data structures to a new data structure for “weak prefix search”. We construct a word x containing the query p padded to the right with trailing zeros, and a word y containing the query p padded to the right with trailing ones. Searching for the rank of the successor of x in S and the rank of the predecessor of y in S gives the required range.

We note that we can carry out the search of the successor of x and the predecessor of y without accessing the keys indexed by the γ nodes. As we will see, our γ nodes implement a succinct blind tree. Searching a blind trie for the right rank of the successor typically requires accessing one of the indexed keys. But, as implicitly used in [5], this access can be avoided if the query is a padded prefix of an indexed key such as x and y above. This implies that the keys indexed by the γ nodes can in fact be discarded and not stored at all. We get a data structure of overall size $O(n \log w)$ bits for weak prefix search.

Belazzougui et al., [5], show that any data structure supporting “weak prefix search” must have size $\Omega(n \log w)$ bits. Hence, our index size is optimal for this related problem.

2.2 Introducing Bit-Selectors and Building a (k, k) -Bit Selector

To construct the γ -nodes we define and construct bit selectors as follows. A (k, L) bit-selector, $1 \leq k \leq L \leq w$, consists of a preprocessing phase and a query phase, (see Figure 2):

- The preprocessing phase: The input is a sequence of length k (with repetitions),

$$I = I[1], I[2], \dots, I[k],$$

where $0 \leq I[j] \leq w - 1$ for all $j = 1, \dots, k$. Given I , we compute the following:

- A sequence of k strictly increasing indices, $0 \leq j_1 < j_2 < \dots < j_k \leq L - 1$, and,
 - An $O(1)$ word data structure, $D(I)$.
- The query phase: given an input word x , and using $D(I)$, produces an output word y such that

$$\begin{aligned} y_{j_\ell} &= x_{I[\ell]}, & 1 \leq \ell \leq k, \\ y_m &= 0, & m \in \{0, \dots, w - 1\} - \{j_\ell\}_{\ell=1}^k. \end{aligned}$$

One main technically difficult result is a construct for (k, k) bit-selectors for all $1 \leq k \leq w/\log w$ (Section 3). The bit selector query time is $O(\log w)$, while the probe complexity and space are constant.

With respect to upper bounds, Brodник, Miltersen, and Munro [6], give several bit manipulation primitives, similar to some of the components we use for bit-selection, but the setting of [6] is somewhat different, and there is no attempt to optimize criteria such as memory probes and index size. The use of Benes networks to represent permutations also appears in Munro et. al [12].

Note that, for (k, k) -bit-selectors, it must be that $j_\ell = \ell - 1$, $1 \leq \ell \leq k$, independently of I . For a sequence of indices I , we define $x[I]$ to be the bits of x in these positions (ordered as in I), if I has multiplicities then $x[I]$ also has multiplicities. With this notation a (k, k) bit selector $D(I)$ computes $x[I]$ for a query x in $O(\log w)$ time.

A $(w^{1/5}, w^{4/5})$ bit-selector is implicit in fusion trees and lie at the core of the data structure. Figure 1 compares the fusion tree bit-selector with our construction.

	k	L	$ D(I) $ in words	# Operations	Multiplication?
Fusion tree bit-selector	$1 \leq k \leq w^{1/5}$	k^4	$O(1)$	$O(1)$	Yes
Our bit-selector	$1 \leq k \leq w/\log w$	k	$O(1)$	$O(\log w)$	No

Fig. 1. The bit-selector used for Fusion Trees in [9,10] vs. our bit-selector.

We remark that Andersson, Miltersen, and Thorup [1] give an $AC(0)$ implementation of fusion trees, *i.e.*, they use special purpose hardware to implement a (k, k) bit-selector (that produces a sketch of length k containing k bits of the key). Ignoring other difficulties, computing a [perfect] sketch in $AC(0)$ is easy: just lead wires connecting the source bits to the target bits. With this interpretation, our bit-selector is a software implementation in $O(\log w)$ time that implements the special purpose hardware implementation of [1].

Our bit-selectors are optimal with respect to query time, when considering implementation on a “practical RAM” (no multiplication is allowed) as defined by Miltersen [11]. This follows from Brodnik et. al. [6] (Theorem 17) who prove that in the “practical RAM” model, any (k, k) -bit-selector, with $k \geq \log^{10} w$, requires at least $\Omega(\log k)$ time per bit-selector query. (Observe that the bit-reversal of Theorem 17 in [6] is a special case of bit-selection).

3 Bit Selectors

In this section we describe both the preprocessing and selection operations for our bit-selectors. We sketch the selection process, which makes use of $D(I)$, the output of the preprocessing. A more extensive description and figures can be found in the appendix, Section C.

$D(I)$ consists of $O(1)$ words and includes precomputed constants used during the selection process. As $D(I)$ is $O(1)$ words, we assume that $D(I)$ is loaded into registers at the start of the selection process. Also, the total working memory required throughout the selection is $O(1)$ words, all of whom we assume to reside within the set of registers.

Partition the sequence $\sigma = 0, 1, \dots, w - 1$ into $w/\log w$ blocks (consecutive, disjoint, subsequences of σ), each of length $\log w$. Let B_j denote the j th block of a word, i.e., $B_j = j \log w, j \log w + 1, \dots, (j + 1) \log w - 1$, $0 \leq j \leq w/\log w - 1$.

Given an input word x and the precomputed $D(I)$, the selection process goes through the seven phases sketched below.

In this high level explanation we give an example input using the following parameters: The word length $w = 16$ bits, a bit index requires $\log w = 4$ bits, I consists of $w/\log w = 4$ indices (with repetitions). A “block” consists of $\log w = 4$ bits, and there are $w/\log w = 4$ blocks.

As a running example let the input word be $x = 1000\ 1101\ 1110\ 0011$ and let $I = \langle 0, 15, 12, 15 \rangle$, the required output is $x[I] = 1101$.

Phase 0: Zero irrelevant bits. We take the mask M with ones at positions in I , and set $x = x \text{ AND } M$. For our example this gives

Input : $M = 1000\ 0000\ 0000\ 1001$, $x = 1000\ 1101\ 1110\ 0011$;

Phase 0: $M = 1000\ 0000\ 0000\ 1001$, $x = 1000\ 0000\ 0000\ 0001$.

Phase 1: Packing blocks to the Left: All bits of x whose index belongs to some block are shifted to the left within the block. We modify the mask M accordingly. Let the number of such bits in block j be b_j . This phase transforms M and x as follows:

Phase 0: $M = 1000\ 0000\ 0000\ 1001$, $x = 1000\ 0000\ 0000\ 0001$;

Phase 1: $M = 1000\ 0000\ 0000\ 1100$, $x = 1000\ 0000\ 0000\ 0100$;

Note that $b_0 = 1$, $b_1 = b_2 = 0$, and $b_3 = 2$. Phase 1 requires $O(\log w)$ operations on a constant number of words (or registers).

Phase 2: Sorting Blocks in descending order of b_j (defined in Phase 1 above). This phase transforms M and x as follows:

Phase 1: $M = 1000\ 0000\ 0000\ 1100$, $x = 1000\ 0000\ 0000\ 0100$;

Phase 2: $M = 1100\ 1000\ 0000\ 0000$, $x = 0100\ 1000\ 0000\ 0000$;

Technically, phase 2 uses a Benes network to sort the blocks in descending order of b_j , in our running example this means block 3 should come first, then block 0, then blocks 2 and 3 in arbitrary order. Brodnik, Miltersen, and Munro [6] show how to simulate a Benes network on bits of a word, we extend this so as to sort entire blocks of $\log w$ bits.

The precomputed $D(I)$ includes $O(1)$ words to encode this Benes network. Phase 2 requires $O(\log w)$ bit operations on $O(1)$ words.

Phase 3: Dispersing bits: reorganize the word produced in Phase 2 so that each of the different bits whose index is in I will occupy the leftmost bit of a unique block. As there may be less distinct indices in I than blocks, some of the blocks may be empty, and these will be the rightmost blocks. This process requires $O(\log w)$ word operations to reposition the bits. This phase transforms M and x as follows:

Phase 2: $M = \mathbf{1100\ 1000\ 0000\ 0000}$, $x = \mathbf{0100\ 1000\ 0000\ 0000}$;

Phase 3: $M = \mathbf{1000\ 1000\ 1000\ 0000}$, $x = \mathbf{0000\ 1000\ 1000\ 0000}$;

Phase 4: Packing bits. The goal now is to move the bits positioned by Phase 3 at the leftmost bits of the leftmost r blocks (r being the number of indices in I without repetitions). Again, by appropriate bit manipulation, this can be done with $O(\log w)$ word operations (see appendix). This phase transforms M and x as follows:

Phase 3: $M = \mathbf{1000\ 1000\ 1000\ 0000}$, $x = \mathbf{0000\ 1000\ 1000\ 0000}$;

Phase 4: $M = \mathbf{1110\ 0000\ 0000\ 0000}$, $x = \mathbf{0110\ 0000\ 0000\ 0000}$;

We remark that if $r = k$, *i.e.*, if I contains no duplicate indices, then we can skip Phases 5 and 6 whose purpose is to duplicate those bits required several times in I .

Phase 5: Spacing the bits. Once again, we simulate a Benes network on the k leftmost bits. The purpose of this permutation is to space out and rearrange the bits so that bits who appear multiple times in I are placed so that multiple copies can be made.

In our running example, phase 5 changes neither M nor x , but this is coincidental – for other inputs ($I' \neq I$) phase 5 would not be the identity function. Phase 5 is yet another application of a Benes network and requires $O(\log w)$ word operations.

Phase 6: Duplicating bits - we duplicate the bits for which space was prepared during Phase 5. This phase transforms M and x as follows:

Phase 5: $M = \mathbf{1110\ 0000\ 0000\ 0000}$, $x = \mathbf{0110\ 0000\ 0000\ 0000}$;

Phase 6: $M = \mathbf{1111\ 0000\ 0000\ 0000}$, $x = \mathbf{0111\ 0000\ 0000\ 0000}$;

Technically, phase 6 makes use of shift and OR operations, where the shifts are decreasing powers of two.

Phase 7: Final positioning: The bits are all now in the k leftmost positions of a word, every bit appears the same number of times it's index appears in I , and we need to run one last Benes network simulation so as to permute these k bits. This permutation gives the final outcome. This phase transforms M and x as follows:

Phase 6: $M = \mathbf{1111\ 0000\ 0000\ 0000}$, $x = \mathbf{0111\ 0000\ 0000\ 0000}$;

Phase 7: $M = \mathbf{1111\ 0000\ 0000\ 0000}$, $x = \mathbf{1101\ 0000\ 0000\ 0000}$;

Note the leftmost $|I| = w/\log w = 4$ bits of x contain the required output of the bit selector.

4 γ -nodes

In this section we use the $(w/\log w, w/\log w)$ -bit-selector, described above, to build a γ -node defined as follows.

Definition 1. A γ -node answers successor queries over a static set S of at most $w/\log w$ w -bit keys. The γ -node uses a compact index of $O(1)$ w -bit words, in addition to the input S . Successor queries perform $O(1)$ word probes, and $O(\log w)$ operations.

We describe the γ -node data structure in stages, beginning with a *slow* γ -node below. A slow γ -node is defined as a γ -node but performs $O(w/\log w)$ operations rather than $O(\log w)$.

4.1 Construction of Slow γ -nodes

We build a blind trie over the set of keys $S = y_1 < y_2, \dots, < y_k$, $k \leq w/\log w$. We denote this trie by $T(S)$. The trie $T(S)$ is a full binary tree with k leaves, each corresponds to a key, and $k-1$ internal nodes. (We do not think of the keys as part of the trie.) We store $T(S)$ in $O(1)$ w -bit words. (The keys, of course require $|S|w$ bits.) $T(S)$ has the following structure:

1. Each internal node of $T(S)$ has pointers to its left and right children.
2. An internal node u includes a bit index, i_u , in the range $0, \dots, w-1$, i_u is the length of the longest common prefix of the keys associated with the leaves in the subtree rooted at u .
3. Key y_i corresponds to the i th leaf from left to right. We store i in this leaf and denote this leaf by $\ell(y_i)$.
4. Keys associated with descendants of the left-child of u have bit i_u equals to zero. Analogously, keys associated with descendants of the right-child of u have bit i_u equals to one.

In addition to $T(S)$, we assume that the keys in S are stored in memory, consecutively in sorted order.

Indices both in internal nodes and leaves are in the range $0, \dots, w-1$ and thereby require $O(\log w)$ bits. Since $T(S)$ has $O(w/\log w)$ nodes, a pointer to a node also requires $O(\log w)$ bits. Thus, in total, each node in $T(S)$ requires only $O(\log w)$ bits. It follows that $T(S)$ (internal nodes and leaves) requires only $O(w)$ bits (or, equivalently, can be packed into $O(1)$ words).

Fundamentally, a *blind-search* follows a root to leaf path in blind trie $T(S)$, ignoring intermediate bits. Searching $T(S)$ for a query x always ends at leaf of the trie (which contains the index of some key). Let $\text{bs}(x, S)$ denote the index stored at this leaf, and let $\text{bkey}(x)$ be $y_{\text{bs}(x, S)}$. I.e., blind search for query x in $T(S)$ leads to a leaf that points to $\text{bkey}(x)$. In general, $\text{bkey}(x)$ is *not* the answer to the successor query, but it does have the longest common prefix of x amongst all keys in S . (See [8].)

To arrive at the successor of x , we retrieve $\text{bkey}(x)$ and compute its longest common prefix with x . Let b be the next bit of x , after $\text{LCP}(x, \text{bkey}(x))$. We use b to pad the remaining bits, let \parallel denote concatenation, and let

$$z = \text{LCP}(x, \text{bkey}(x)) \parallel b^{w-|\text{LCP}(x, \text{bkey}(x))|}.$$

Finally, we perform a second blind-search on z . The result of this second search gives us the index of the successor to x to within ± 1 .

Overall, the number of probes required for such a search is $O(1)$. However, the computation time is equal to the length of the longest root to leaf path in $T(S)$, which is $O(w/\log w)$.

4.2 Improving the running time

Using our $(w/\log w, w/\log w)$ -bit-selector we can reduce the search time in the blind trie from $O(w/\log w)$ to $O(\log w)$ operations while still representing the trie in $O(1)$ words. For that we change the first part of the query, that is the *blind-search* for $\text{bs}(x, S)$ (the index of $\text{bkey}(x)$). Rather than walking top down along a path in the trie we use a binary search as follows.

We need the following notation. Any node $u \in T(S)$, internal node or leaf, defines a unique root to u path in $T(S)$. Denote this path by $\pi_u = v_0, v_1, \dots, v_{|\pi_u|}$ where v_0 is the root, $v_{|\pi_u|} = u$, and v_i is the parent of v_{i+1} . For any node $u \in T(S)$ let I_u be the sequence of indices i_v for all internal nodes v along π_u . Also, let ζ_u be a sequence of zeros and ones, one entry per edge in π_u , zero for an edge pointing left, one otherwise. For all $1 \leq q \leq |S|$ we define $\pi_q = \pi_{\ell(y_q)}$, $I_q = I_{\ell(y_q)}$, and $\zeta_q = \zeta_{\ell(y_q)}$. The following lemma is straightforward.

Lemma 1. *For any index $1 \leq q \leq |S|$, query x , we have that*

$$\begin{aligned} \zeta_q \text{ is lexicographically smaller than } x[I_q] &\Rightarrow y_q < \text{bkey}(x) \\ \zeta_q = x[I_q] &\Rightarrow y_q = \text{bkey}(x), \\ \zeta_q \text{ is lexicographically larger than } x[I_q] &\Rightarrow y_q > \text{bkey}(x). \end{aligned}$$

Based on Lemma 1, given query x , we can do binary search to find $\text{bs}(S, x)$:

```

L ← 1, R ← |S|, q ← ⌊(L + R)/2⌋
while  $\zeta_q \neq x[I_q]$  do
  if  $\zeta_q < x[I_q]$  then R ← q
  else L ← q
  end if
  q ← ⌊(L + R)/2⌋
end while
return q

```

Lemma 2. *The above binary search algorithm returns $\text{bs}(x, S)$ and has $O(\log w)$ iterations.*

Next we show how to implement each iteration of this binary search and compare $x[I_q]$ and ζ_q in $O(1)$ time while keeping the trie stored in $O(1)$ words.

For this we devise a sequence I of bit indices, of length $w/\log w$. Prior to running the binary search we use the bit-selector of Section 3 to compute $x[I]$ and later we use $x[I]$ to construct $x[I_q]$ in every iteration in $O(1)$ time. We extract $x[I_q]$ from $x[I]$ and retrieve ζ_q using $O(1)$ additional words. The details are as follows.

The $O(1)$ words which form the γ node: For each $1 \leq q \leq |S|$ there is a unique interval $[L_q, R_q]$ of which q may be the splitting point (i.e. $q = \lfloor (L_q + R_q)/2 \rfloor$) during the binary search. Let $\pi_q = u_1, u_2, \dots, u_t = \ell(y_q)$ be the path to q as defined above. Define $j_{L_q} \in 1, \dots, t$ to be the length of the longest common prefix of π_q and π_{L_q} . That is $u_{j_{L_q}}$ is the lowest common ancestor of the leaves $\ell(y_q)$ and $\ell(y_{L_q})$. Define j_{R_q} analogously, and let $j = \max(j_{L_q}, j_{R_q})$.

Let $\tilde{\pi}_q$ be the suffix of π_q starting at node u_{j+1} , and let \tilde{I}_q be the suffix of I_q starting at $I_q[j+1]$. (These are the indices stored in $u_{j+1}, u_{j+2}, \dots, u_{t-1}$). Similarly, let $\tilde{\zeta}_q$ be the suffix of ζ_q , starting at the j th element.

Given S , for every $1 \leq q \leq |S|$ we precompute and store the following data: $j_{L_q}, j_{R_q}, \tilde{I}_q, \tilde{\zeta}_q$. It is easy to verify that $O(1)$ words suffice to store the $4|S|$ values above. Indeed, j_{L_q} and j_{R_q} are indices in $1, \dots, |S|$, $O(\log w)$ bits each. As the number of keys $|S| \leq w/\log w$, all the j_{L_q} 's, and j_{R_q} 's fit in $O(1)$ words. Since $\tilde{\pi}_q$ paths are pairwise disjoint, the sum of their path lengths is $O(|S|) = O(w/\log w)$. Hence, storing all the sequences $\tilde{\zeta}_q$, $1 \leq q \leq w/\log w$, requires no more than $O(w/\log w)$ bits. We store the $\tilde{\zeta}_q$'s concatenated in increasing order of q in a single word Z .

The sequence I for which we construct the bit selector is the concatenation of the \tilde{I}_q sequences, in order of q . As above, it follows that I is a sequence of $O(w/\log w)$ log w -bit indices. The bit selector $D(I)$ is also stored as part of the γ node.

For each q we also compute and store the index s_q of the starting position of $\tilde{\zeta}_q$ in Z . This is the same as the index of the starting position of I_q in I . Clearly all these indices s_q can be stored in a single word.

Implementing the blind search: As we mentioned, given x as a query to the γ -node, we compute $x[I]$ (once) from x and $D(I)$, which requires $O(\log w)$ operations and no more than $O(1)$ probes.

At the start of an iteration of the binary search, we have a new value of q , and access to the following values, all of whom are in $O(1)$ registers from previous iterations:

$$x[I], \quad j_{L_q}, \quad j_{R_q}, \quad L_q, \quad R_q, \quad \zeta_{L_q}, \quad \zeta_{R_q}, \quad x[I_{L_q}], \quad x[I_{R_q}].$$

For the rest of this section let $L = L_q$ and $R = R_q$. We now compute ζ_q and $x[I_q]$. We retrieve j_L, j_R from the data-structure, and we also retrieve $x[\tilde{I}_q]$ from $x[I]$ and $\tilde{\zeta}_q$ from Z (note that $x[\tilde{I}_q]$ is stored consecutively in $x[I]$ and $\tilde{\zeta}_q$ is stored consecutively in Z , and we use s_q to know where they start).

If $j_L \geq j_R$, we compute $x[I_q] \leftarrow (x[I_L][1, \dots, j_L]) \parallel (x[\tilde{I}_q])$ and $\zeta_q \leftarrow (\zeta_L[1, \dots, j_L]) \parallel (\tilde{\zeta}_q)$. Analogously, if $j_L < j_R$, and we compute

$$\begin{aligned} x[I_q] &\leftarrow (x[I_R][1, \dots, j_R]) \parallel (x[\tilde{I}_q]) \\ \zeta_q &\leftarrow (\zeta_R[1, \dots, j_R]) \parallel (\tilde{\zeta}_q). \end{aligned}$$

All these operations are easily computed using $O(1)$ SHIFT, AND, OR operations.

5 Open Issues

1. Our (k, k) -bit selector takes $O(\log w)$ operations, which are optimal when $k \geq w^\epsilon$ for any constant $\epsilon > 0$. What can be done for smaller values of k ? (E.g., for $k = O(1)$ one can definitely do better).

2. It follows from Thorup ([15]) that, in the practical-RAM model, a search node with fan-out $\frac{w}{\log w}$ requires $\Omega(\log \log w)$ operations. Our γ nodes have fan out $w/\log w$ and require $O(\log w)$ operations. Can this gap be bridged?
3. A natural open question is if the additive $O(\log w)$ in time complexity is required or not.

Acknowledgments. We wish to thank Nir Shavit for introducing us to the problems of contention in multicore environments, for posing the question of multicore efficient data structures, and for many useful discussions. We also wish to thank Mikkel Thorup for his kindness and useful comments.

References

1. A. Andersson, P. B. Miltersen, and M. Thorup. Fusion trees can be implemented with AC(0) instructions only. *Theor. Comput. Sci.*, 215(1-2):337–344 1999.
2. P. Beame and F. E. Fich. Optimal bounds for the predecessor problem and related problems. *Journal of Computer and System Sciences*, 65(1):38–72 2002.
3. D. Belazzougui, P. Boldi, R. Pagh, and S. Vigna. Theory and practice of monotone minimal perfect hashing. *J. Exp. Algorithmics*, 16:3.2, 2011.
4. D. Belazzougui, P. Boldi, R. Pagh, and S. Vigna. Monotone minimal perfect hashing: searching a sorted table with $o(1)$ accesses. SODA, 2009, 785–794.
5. D. Belazzougui, P. Boldi, R. Pagh, and S. Vigna. Fast prefix search in little space, with applications. ESA, 2010, 427–438.
6. A. Brodnik, P. B. Miltersen, and J. I. Munro. Trans-dichotomous algorithms without multiplication - some upper and lower bounds. WADS, 1997, 426–439.
7. U. Drepper. What every programmer should know about memory, 2007, <http://lwn.net/Articles/250967/>.
8. P. Ferragina and R. Grossi. The string B-tree: a new data structure for string search in external memory and its applications. *J. ACM*, 46:236–280, 1999.
9. M. L. Fredman and D. E. Willard. Surpassing the information theoretic bound with fusion trees. *Journal of Computer and System Sciences*, 47(3):424 – 436, 1993.
10. M. L. Fredman and D. E. Willard. Trans-dichotomous Algorithms for Minimum Spanning Trees and Shortest Paths. FOCS, 1990, 719 – 725.
11. P. B. Miltersen. Lower bounds for static dictionaries on RAMs with bit operations but no multiplication. ICALP, 1996, 442–453.
12. J. I. Munro, R. Raman, V. Raman, and S. S. Rao. Succinct representations of permutations. ICALP, 2033, 345–356.
13. M. Pătraşcu and M. Thorup. Time-space trade-offs for predecessor search. STOC, 2006, 232–240.
14. N. Shavit. Data structures in the multicore age. *Commun. ACM*, 54(3):76–84, 2011.
15. M. Thorup. On AC0 implementations of fusion trees and atomic heaps. SODA, 2003, 699–707.
16. P. van Emde Boas. Preserving order in a forest in less than logarithmic time and linear space. *Inf. Process. Lett.*, 6(3):80–82, 1977.
17. D. E. Willard. Log-logarithmic worst-case range queries are possible in space $t(n)$. *Information Processing Letters*, 17(2):81–84, 1983.

keys. The size of a β -node is $O(1)$, the query time is $O(\log w)$ (in expectation and w.h.p) and the number of probes outside the index is $O(1)$ (in expectation and w.h.p).

We start by describing a non-succinct version of a β -structure, which we refer to as B -structure, and then we describe how to transform the non-succinct B -structure into a succinct β -structure which occupies only $O(n(\log w + \log n))$ bits.

B.1 The Prefix Partitioning Lemma

Let us start by defining a prefix-partition operator \sqsubseteq : Let $S \in \{0, 1\}^*$ be an arbitrary set of binary strings, and let $p \in \{0, 1\}^*$ be a binary string, we partition S into two subsets: $S_{\sqsubseteq p}$ and $S_{\not\sqsubseteq p}$.

$S_{\sqsubseteq p}$ is the set of all the elements of S which start with p , and $S_{\not\sqsubseteq p} = S - S_{\sqsubseteq p}$, is the set of all the elements of S which don't start with p .

The following lemma proves that there exists a prefix p which partitions S into two approximately equal subsets.

Theorem 1. *For every set S of binary strings, $|S| \geq 2$, there exists a binary string $p \in \{0, 1\}^*$ s.t. $\frac{1}{3}|S| \leq |S_{\sqsubseteq p}| \leq \frac{2}{3}|S|$ and $\frac{1}{3}|S| \leq |S_{\not\sqsubseteq p}| \leq \frac{2}{3}|S|$*

Proof. Let initially $p = \epsilon$ be the empty string. While $(|S_{\sqsubseteq p}| > \frac{2}{3}|S|)$, if $|S_{\sqsubseteq p0}| \geq |S_{\sqsubseteq p1}|$ add a 0-bit to the end of p , otherwise add a 1-bit to the end of p . We stop the loop at the first time that $|S_{\sqsubseteq p}| \leq \frac{2}{3}|S|$, and it is easy to verify that at that point we also have that $|S_{\not\sqsubseteq p}| \leq \frac{2}{3}|S|$.

B.2 Construction of a B -structure

Let S_0 be the initial set of n w -bit binary strings. Assume without loss of generality that $0^w \in S_0$. We can assume that, since if $0^w \notin S_0$ then $rank_{S_0}(x) = rank_{S_0 \cup \{0\}}(x) - 1$ for every $x > 0$. The B -structure is a binary tree containing a prefix of the strings in S_0 in each internal node, and at most 3 keys of S_0 at each leaf.

We define the B -node recursively for a subset $S \subseteq S_0$ (starting with $S = S_0$). Let p be a prefix of a key in S as in Lemma 1, define $p_{min}, p_{max} \in S$ to be the minimum and maximum keys respectively in S which starts with p . Store p in the root of the B -structure of S , the right child is a B -structure of $S_R = S_{\sqsubseteq p} \cup \text{StrictPredecessor}(S, p_{min})$ (we define here $\text{StrictPredecessor}(S, p_{min})$ to be the maximal key in S which is smaller than p_{min} , or p_{min} itself if it is the minimal key in S), the left child is a B -structure of $S_L = S_{\not\sqsubseteq p} \cup p_{max}$. We "associate" S with the root, S_R with its right child, and S_L with its left child. When $|S| \leq 3$, we stop the recursion and store the (at most 3) keys of S in the leaf of the B -structure.

According to Lemma 1, $|S_R|, |S_L| \leq \frac{2}{3}|S| + 1$, and since $|S_L| + |S_R| \leq |S| + 2$, we get that the height of the resulting tree is $O(\log |S_0|)$, and the number of nodes in the tree is $O(|S_0|)$.

B.3 Querying the B -node

Let $x \in \{0,1\}^w$ be the query word. Start the search in the root. The root contains a prefix p , if x starts with p continue the search in the right child, otherwise, continue the search in the left child. Continue the search similarly in every internal node that we reach, until we reach a leaf.

In the leaf at most 3 keys are stored, denote them by $s_1 < s_2 < s_3$. If $s_1 \leq x < s_2$ output $\text{rank}_{S_0}(s_1)$ (that is, the rank of s_1 in S_0). If $s_2 \leq x < s_3$ output $\text{rank}_{S_0}(s_2)$. Otherwise, $x > s_3$ output $\text{rank}_{S_0}(s_3)$.

B.4 Correctness of the B -structure

Given $x \in \{0,1\}^w$ we need to prove that the output of the search procedure is $\text{rank}_{S_0}(x)$.

Let $y = \text{Pred}(S_0, x)$ be the predecessor of x in S_0 . At the end of the search we reach a leaf which stores between 1 and 3 elements of S_0 . We need to prove that y is one of these keys. Let (v_0, v_1, \dots, v_k) be the root-to-leaf path traversed during the search. Let p_i be the prefix stored at v_i and let S_i be the subset of S_0 associated with v_i , for $i = 0, 1, \dots, k$. We prove by induction on i , for $i = 0, 1, \dots, k$, that y is a member of S_i . This is correct at the root ($i=0$), since $y \in S_0$ by our assumption that $0^w \in S_0$. For the inductive step, we assume $y \in S_i$, and prove that $y \in S_{i+1}$.

Lemma 3. *Let $y = \text{Pred}(S_0, x)$. If $y \in S_i$ then $y \in S_{i+1}$.*

Proof. Let $p_{min}, p_{max} \in S_i$ be the minimum and maximum keys in S_i respectively which start with p_i .

If v_{i+1} is a right child, then x starts with p_i and $S_{i+1} = (S_i)_{\sqsubseteq p_i} \cup \text{StrictPredecessor}(S_i, p_{min})$. If $y \in (S_i)_{\sqsubseteq p_i}$ we are done, otherwise it must be that $y = \text{StrictPredecessor}(S_i, p_{min})$ since x starts with p_i and y is its predecessor in S_i .

If v_{i+1} is a left child, then x doesn't start with p_i and $S_{i+1} = (S_i)_{\not\sqsubseteq p_i} \cup p_{max}$. If $x < p_{min}$ then y doesn't start with p_i and hence $y \in (S_i)_{\not\sqsubseteq p_i}$. If $x > p_{max}$ then either $y = p_{max}$, or $y > p_{max}$ and then $y \in (S_i)_{\not\sqsubseteq p_i}$. We get that in all the cases, $y \in (S_i)_{\not\sqsubseteq p_i} \cup p_{max}$ as required.

B.5 Making it succinct: from B -structure to β -structure

We now describe the succinct variant of the B -structure, which we call β -structure. Its index occupies $O(n(\log w + \log n))$ bits, and its search time is $O(\log n)$ (w.h.p), and the number of probes outside the index is $O(1)$ (w.h.p).

Each node of the β -structure occupies $\log w + \log n$ bits, defined as follows:

- **Inner nodes:** Replace every prefix p in the B -structure with a pair $\langle |p|, h(p) \rangle$, $|p|$ is the length of the prefix p ($\log w$ bits) and $h(p)$ being a signature of p of length $(2 \log n)$ bits, computed using a universal hash function. To test if x starts with p check if $h(x[1..|p|]) = h(p)$. If so, then x starts with p with probability at least $1 - 2^{-2 \log n} = 1 - (\frac{1}{n})^2$, otherwise x doesn't start with p .

- **Leaves:** In the leaves, replace the $O(1)$ (at most 3) keys stored at each leaf with their rank in S_0 ($O(\log n)$ bits). In the search procedure, when reaching a leaf, retrieve these keys (from the static sorted list of the keys of S_0) using their ranks in $O(1)$ word-accesses.

Finally, at the end of the search assume the algorithm suggests that $\text{rank}_{S_0}(x) = i$. We can test if it's correct using $O(1)$ word-accesses by checking that $s_i \leq x < s_{i+1}$ (recall the notation $S_0 = s_1 < s_2 < \dots < s_n$). With probability at most $2^{-2^{\log n}} \cdot \log n < \frac{1}{n}$ we will get an error at some node along the search path. When an error is detected we do a binary search to find the predecessor of x among the static set of n keys, this takes $O(\log n)$ time and probes. So the binary search takes $o(1)$ on average. When no error occurs, the search time is $O(\log w)$ (this happens with probability at least $1 - \frac{1}{n}$). Hence, the query time is $O(\log w)$ (in expectation and w.h.p), and we access only $O(1)$ words outside the index (in expectation and w.h.p).

C Detailed Description of the $(w/\log w, w/\log w)$ Bit-Selector

In this section we describe both the preprocessing and selection operations for our bit-selectors. We follow the selection process, which makes use of $D(I)$, the output of the preprocessing. While describing the selection process we specify the different components of $D(I)$.

$D(I)$ consists of $O(1)$ words and includes precomputed constants used during the selection process. As $D(I)$ is $O(1)$ words, we assume that $D(I)$ is loaded into registers at the start of the selection process. Also, the total working memory required throughout the selection is $O(1)$ words, all of whom we assume to reside within the set of registers.

Let $I = I[0], I[1], \dots, I[k-1]$, $k \leq w/\log w$, be the sequence of bit indices to be selected from some input word x . I may contain repetitions. The indices $I[j]$ range in $[0, w-1]$, where bit zero is the most significant bit (on the left in the Figures). Let r be the number of distinct values in $I[0], \dots, I[k-1]$.

Partition the sequence $\sigma = 0, 1, \dots, w-1$ into $w/\log w$ blocks (consecutive, disjoint, subsequences of σ), each of length $\log w$. Let B_j denote the j th block of a word, i.e., $B_j = j \log w, j \log w + 1, \dots, (j+1) \log w - 1$, $0 \leq j \leq w/\log w - 1$.

We define the following notation. For sequences σ and τ , $\sigma \cap \tau$ denotes a subsequence of σ consisting of those values that appear somewhere in τ . We recall that for a sequence of indices σ , we define $x[\sigma]$ to be the bits of x in these positions (ordered as in σ), if σ has multiplicities then $x[\sigma]$ also has multiplicities.

An assignment to $x[\sigma]$, such as $x[\sigma] \leftarrow b_1, b_2, \dots, b_{|\sigma|}$, $b_i \in \{0, 1\}$, is shorthand notation for $x[\sigma[1]] \leftarrow b_1, x[\sigma[2]] \leftarrow b_2, \dots, x[\sigma[|\sigma|]] \leftarrow b_{|\sigma|}$. (Assignment to $x[\sigma]$ makes sense if σ has no multiplicities).

Also, given a word z , let $z \gg i$ denote a right shift of z by i bits, and $z \ll i$ a left shift by i bits.

Given an input word x and the precomputed $D(I)$, the selection process goes through the seven phases described below.

C.1 The ever changing x and I

As we process the various phases and sub phases of the bit selection, the original bits of x are permuted, duplicated, or set to zero.

Phases 1 – 5 and 7 simply permute the bits of x . Each permutation is performed in $O(\log w)$ operations. Phase 6 duplicates some of the bits in x (those bits with multiplicity > 1 in the sequence I). Each of the phases requires $O(1)$ precomputed words throughout its execution.

Let x_0 be the original word and $I_0 = I$ be the original sequence of indices. Moreover, let x_t be the word x after phases 1 to t , and let I_t be a sequence of indices such that for all $j = 0, \dots, k - 1$, $x_t[I_t[j]] = x_0[I_0[j]]$.

For any t , $0 < t \neq 6$, imagine that I_t is obtained by changing I_{t-1} so as to reflect the bit permutation performed on x_{t-1} to get x_t . This permutation on I_{t-1} need not be actually done, these permutations are implicitly used by the bit selection algorithm.

During phase 6, where bits are duplicated so that the number of copies of each bit is equal to the multiplicity of the index of the bit in I_0 (or I_5), imagine that I_6 is produced from I_5 by removing multiplicities and substituting $i + j - 1$ for the j th appearance of index i in I_5 .

It follows that for all $0 \leq t \leq 5$, $0 \leq j \leq k - 1$, the multiplicity of $I_t[j]$ is equal to the multiplicity of $I_0[j]$. For $t = 6, 7$, $0 \leq j \leq k - 1$, the multiplicity of $I_t[j]$ is one.

Initially, all bits not appearing in I_0 are set to zero simply by setting $x_0 \leftarrow x \text{ AND } M$ where M is a mask with its i th bit equal to one iff i appears in I .

The final output of the bit selection, from left to right, is a word $x_0[I] \parallel 0^{w-|I|}$.

For brevity, we use x as a continuously changing variable throughout the description of the different phases. The sequences I_t are needed during the preprocessing phase, the query phase requires only a constant number of precomputed words. We describe how the preprocessing phase keeps track of the various I_t sequences and the permutations applied to x implicitly through the description of the phases.

C.2 Phase 1: Packing blocks to the left (Figure 3)

We now describe the procedure for rearranging the bits of x so that for all blocks B , the bits $x[B \cap I]$ are assigned to the leftmost positions of $x[B]$, preserving their order. This will be done for all blocks in parallel by the inherent parallelism of word operations.

For block B , let $\text{suff}_i(B)$ be the length i suffix of B .

Phase 1 requires $\log w$ subphases. We maintain the following invariant after subphase i , $1 \leq i \leq \log w$: The bits $x[\text{suff}_i(B) \cap I]$ are assigned to the leftmost positions of $x[\text{suff}_i(B)]$ and the other bits of $x[\text{suff}_i(B)]$ are set to zero. Bits of x whose indices are not in $\text{suff}_i(B)$ do not change. Note that this invariant initially holds for $i = 1$.

At subphase i , for $i = 2, \dots, \log w$, for each block B whose i th largest index is not in I_0 we assign $x[\text{suff}_{i-1}(B)] \parallel '0'$ to $x[\text{suff}_i(B)]$.

	Block 0						Block 1						...	Block $w/\log w-1$						Block $w/\log w$					
X:	x_0	x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8	x_9	x_{10}	x_{11}		x_{48}	x_{49}	x_{50}	x_{51}	x_{52}	x_{53}	x_{54}	x_{55}	x_{56}	x_{57}	x_{58}	x_{59}
Mask	0	0	1	0	0	1	0	1	0	1	0	0	...	0	0	0	1	0	0	0	0	0	1	0	0
	0	0	x_2	0	0	x_5	0	x_7	0	x_9	0	0		0	0	0	x_{51}	0	0	0	0	0	x_{57}	0	0
	0	0	x_2	0	x_5	0	0	x_7	0	x_9	0	0		0	0	0	x_{51}	0	0	0	0	0	x_{57}	0	0
	0	0	x_2	x_5	0	0	0	x_7	0	x_9	0	0		0	0	0	x_{51}	0	0	0	0	0	x_{57}	0	0
	0	0	x_2	x_5	0	0	0	x_7	x_9	0	0	0		0	0	x_{51}	0	0	0	0	0	x_{57}	0	0	0
	0	x_2	x_5	0	0	0	0	x_7	x_9	0	0	0		0	x_{51}	0	0	0	0	0	x_{57}	0	0	0	0
Output:	x_2	x_5	0	0	0	0	x_7	x_9	0	0	0	0		x_{51}	0	0	0	0	0	x_{57}	0	0	0	0	0

Fig. 3. An illustration of Phase 1.

Let Z_i be a word with 1 at the i th largest index of each block, and zeros elsewhere. We need Z_i during subphase i of Phase 1. Z_1 can be constructed on the fly in a register, in time $O(\log w)$, Z_i is simply a left shift of Z_1 by $i - 1$. Let $L_i = \overline{M} \text{ AND } Z_i$, and let $S_i = L_i - (L_i \gg (i - 1))$. See Figure 4.

The i th subphase is as follows: We compute $y_1 = x \text{ AND } S_i$ which gives the bits that have to be left shift by one position, and we compute $y_2 = x \text{ AND } \overline{S}_i$ which gives a word containing the bits which are to remain in their positions. Finally, we set $x = (y_1 \ll 1) \text{ OR } y_2$. See Figure 5.

C.3 Phase 2: Sorting blocks by size (Figure 6)

We permute $x[B_0], \dots, x[B_{\frac{w}{\log w}-1}]$ such that they are in non-increasing order of $|B_j \cap I_1|$. Note that we know this permutation when preprocessing I_1 . We implement this step using a simulation of a Benes-network (described in Section C.9). This simulation requires $O(\log w)$ operations, and uses $O(1)$ precomputed constants stored in $D(I)$, and $O(1)$ registers.

C.4 Phase 3: Dispersing bits (Figure 7)

Recall that r is the number of distinct values in I . Let $\sigma = B_0[0], B_1[0], \dots, B_{r-1}[0]$, i.e., σ is a sequence of the first (and smallest) index in every block. For any sequence π , define $\tau(\pi) = \tau_1, \tau_2, \dots, \tau_r$ be a subsequence of π where $\pi[j]$ is discarded if $\pi[j'] = \pi[j]$ for some $j' < j$.

In Phase 3 we disperse the bits of $x[I_2]$, so that

$$x[\sigma] \leftarrow x[\tilde{\tau}],$$

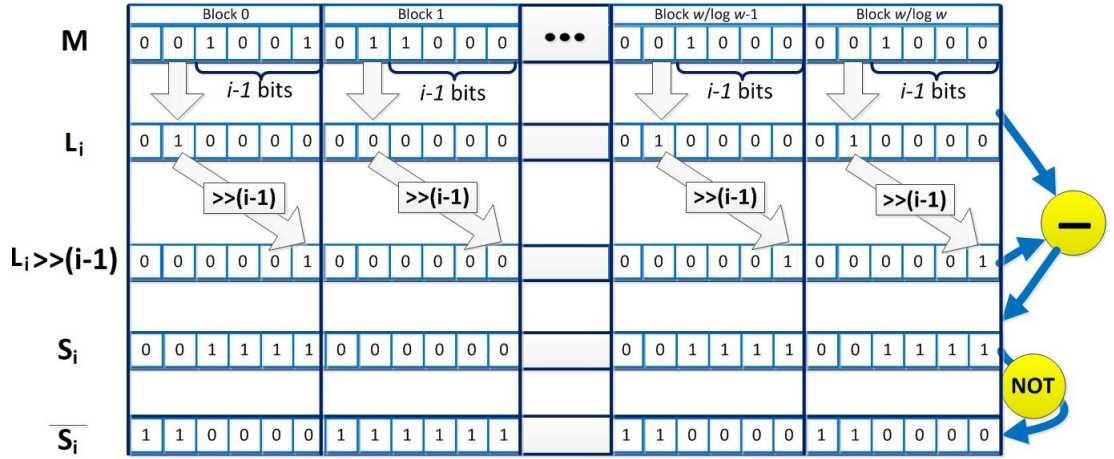


Fig. 4. The masks used during Phase 1.

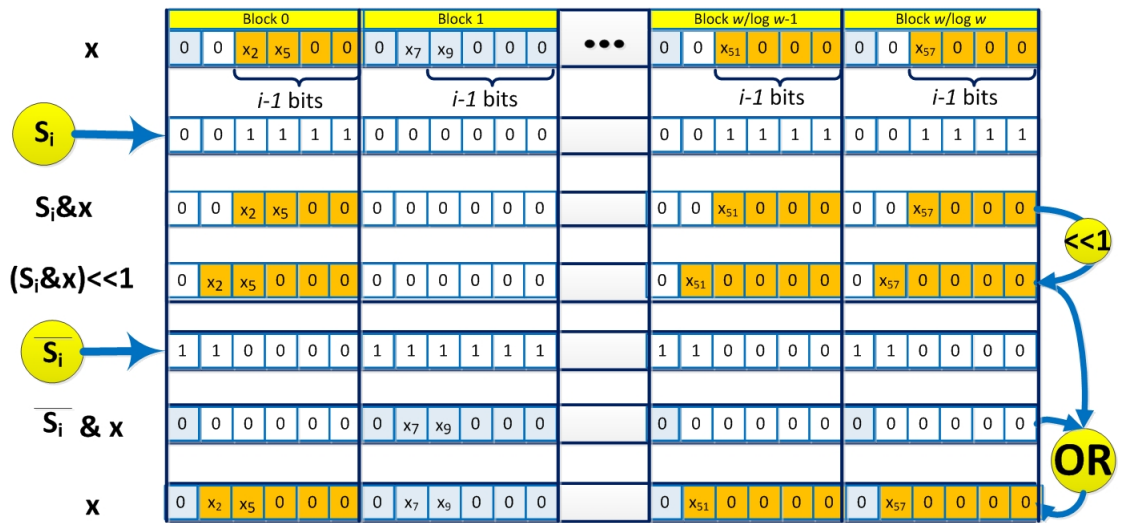


Fig. 5. A subphase of Phase 1.

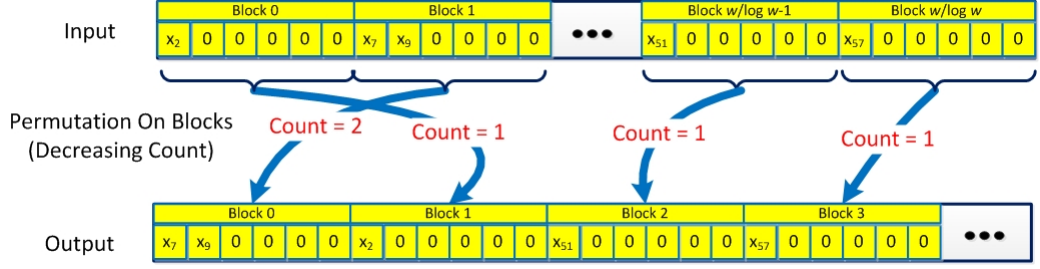


Fig. 6. An illustration of Phase 2.

for some sequence $\tilde{\tau}$ produced by some permutation on the order of $\tau(I)$. The description of $\tilde{\tau}$, is implicit in the description of Phase 3 below.

Following Phase 2, we have that $|B_0 \cap I_2| \geq |B_1 \cap I_2| \geq \dots \geq |B_{w/\log w-1} \cap I_2|$. Therefore, for $1 \leq i \leq \log w$, we can define

$$a_i = |\{j \mid |B_j \cap I_2| \geq i\}|.$$

Let $A_i = \sum_{\ell=1}^i a_\ell$ for $i = 1, \dots, \log w$, and define $A_0 = 0$.

We can now define the sequence σ_i , $1 \leq i \leq \log w$,

$$\sigma_i = \sigma[A_{i-1} \log w], \sigma[(A_{i-1} + 1) \log w], \dots, \sigma[(A_i - 1) \log w],$$

and the sequence

$$\xi_i = B_0[i-1], B_1[i-1], \dots, B_{a_{i-1}}[i-1], \quad 1 \leq i \leq \log w.$$

Phase 3 has $\log w$ subphases. Subphase i of Phase 3 performs the assignment $x[\sigma_i] \leftarrow x[\xi_i]$, this assignment can be implemented using $O(1)$ operations.

Isolate the bits to be moved (indices ξ_i), shift them to their new locations (indices σ_i , note that $\sigma_i[j] - \xi[j] = \sigma_i[j'] - \xi[j']$ for all $1 \leq j, j' \leq |\sigma_i|$), producing word y . Next, update x by setting $x[\xi_i]$ to zero and taking the OR with y .

The values $A[i] \log w$ are stored as part of $D(I)$ (in total $\log^2 w$ bits). These values suffice so as to generate all the masks and operations required in Phase 3.

C.5 Phase 4: Packing bits (Figure 8)

Let σ and $\tau(\pi)$ be as defined at the start of Phase 3. in Phase 4 we “push” the bits $x[\tau(I_3)]$ to the left, i.e.,

$$x[0, 1, \dots, r-1] \leftarrow x[\tilde{\tau}],$$

where $\tilde{\tau}$ is a sequence produced by some permutation on the order of $\tau(I_3)$. As in Phase 3, the description of $\tilde{\tau}$, is implicit in the description of Phase 4 below. Note that $\tau(I_3)$ is a permutation of σ .




Fig. 7. An illustration of Phase 3.

There are $\log w$ subphases in Phase 4.

Let $q = \lceil r/\log w \rceil$. In subphases $1, \dots, \log w - 1$ we fill $x[B_0], x[B_1], \dots, x[B_{q-1}]$, with some permutation of the first $q \log w$ bits of $x[\sigma]$. The last subphase is used to copy the $k < \log w$ leftover bits of $x[\sigma]$ into $x[B_q[0], B_q[1], \dots, B_q[k-1]]$.

For $1 \leq i < \log w$ define the sequences $v_i = i, \log w + i, \dots, (q-1)\log w + i$ and $\zeta_i = iq \log w, (i+1)\log w, \dots, ((i+1)q-1)\log w$.

In Subphase i of Phase 4 we perform the assignment

$$x[v_i] \leftarrow x[\zeta_i].$$

To do this using word operations, we first isolate the bits of $x[\zeta_i]$, shift them so as to be in their target locations, v_i , and “or” them into place.

The last subphase copies the remaining bits one by one, for a total of $O(\log w)$ operations.

C.6 Phase 5: Spacing the bits (Figure 9)

At the end of Phase 4 $x[0, 1, \dots, r-1]$ is a permutation of the bits of $x[\tau[I_4]]$, and $x[j], j \geq r$, are zero. Our goal is now to space the bits so as to make space for duplication of those bits whose indices appear multiple times in I_4 .

In this phase we space the bits $x[0, 1, \dots, r-1]$ by “inserting” $j-1$ zeros between $x[\ell]$ and $x[\ell+1]$ iff ℓ appears j times in I_4 . We do this by permuting the bits of $x[0, 1, \dots, k]$. There is one unique permutation that achieves this goal. This is done by simulating a Benes sorting network, in time $O(\log w)$, and using only $O(1)$ precomputed constants and $O(1)$ registers.

C.7 Phase 6: Duplicating bits (Figure 10)

For a sequence ϱ let $m_\ell(\varrho)$ be the number of occurrences of ℓ in ϱ . At the end of Phase 5, for every $\ell \in I_5$ such that $m_\ell(I_5) > 1$ we have that $x[\ell+1, \ell+2, \dots, \ell+$

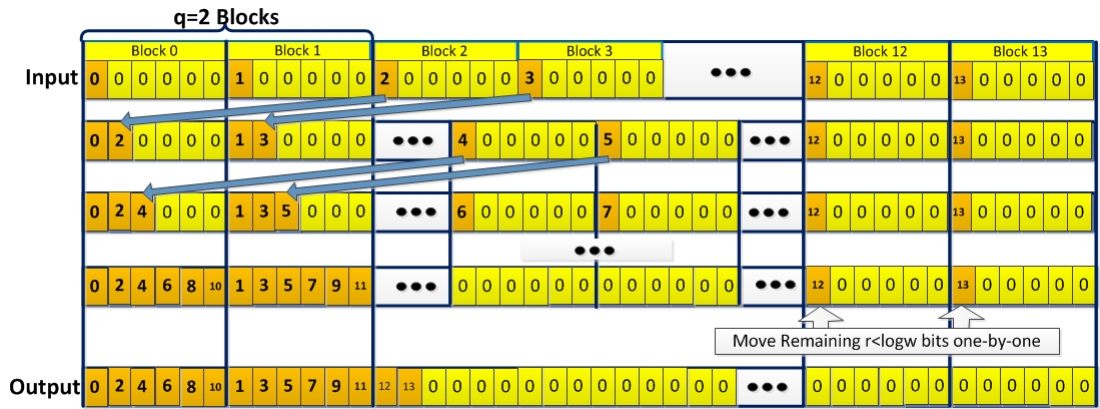


Fig. 8. An illustration of Phase 4.




Fig. 9. An illustration of Phase 5.

$m_\ell(I_5) - 1]$ contain zeros and none of the indices $\ell + 1, \ell + 2, \dots, \ell + m_\ell(I_5) - 1$ appear in I_5 .

Phase 6 consists of $\log w$ subphases, $i = 1, \dots, \log w$. Subphase i duplicates a subset of the bits of x specified by a $w/\log w$ bit mask M_i . All these masks are precomputed at preprocessing and store in a single word with $D(I)$.

We compute the masks M_i as follows. Let I_6^i be the sequence which describes the positions of the bits of I at the end of subphase i , and let $I_6^0 = I_5$. When a bit is copied we split its remaining multiplicity among the two copies.

Let $\Delta_i = 2^{\log w - i}$. Subphase $i = 1, \dots, \log w$ duplicates those bits $x[\ell]$ for which

$$m_\ell(I_6^{i-1}) > \Delta_i. \quad (1)$$

So M_i is set to one in all positions ℓ for which Equation (1) holds and it is zero in all other places.

I_6^i is computed from I_6^{i-1} as follows: For every ℓ that appears somewhere in I_6^{i-1} let $i_1 < i_2 < \dots < i_t$ be the indices of all occurrences of ℓ in I_6^{i-1} . Let $I_6^i[i_j] = \ell$ for $j < \Delta_i$ (unchanged from $I_6^{i-1}[i_j]$), and set $I_6^i[i_j] = \ell + \Delta_i$ otherwise. This effectively splits the multiplicity of ℓ between ℓ and its new copy $\ell + \Delta_i$. Bit ℓ has now multiplicity $\Delta_i - 1$ and bit $\ell + \Delta_i$ has the remaining multiplicity.

At query time in subphase i we set $x = (x \text{ OR } ((x \text{ AND } M_i) \gg \Delta_i))$.




Fig. 10. An illustration of Phase 6.

C.8 Phase 7: Final Positioning (Figure 11)

At this stage we need permute the bits $x[1, \dots, k]$, so as to get the final output. Note that I_6 is a permutation and its inverse permutation, I_6^{-1} is the permutation we need to apply to x . This too requires simulation of a Benes network, see Section C.9.




Fig. 11. An illustration of Phase 7.

C.9 Permuting elements in a word by simulating a Benes network

We show how to prepare a set C of $O(1)$ words such that given C a Benes network implementing a given permutation σ can be applied to a word x in $O(\log w)$ operations (shift, and, or).

We use such networks in two contexts:

- To permute the $b \leq w/\log w$ leftmost bits of x . We need this in Phases 5 and 7 of bit selection.
- To permute $w/\log w$ blocks of bits (each block of length $\log w$). We need this during Phase 2 of bit selection.




Fig. 12. One of many variants for the Benes network.

Overview of the Benes-Network Assume that n is a power of 2. A Benes network, $B(n)$, of size n consists of two Benes networks of size $n/2$, $B_u(n/2)$, and $B_d(n/2)$. For $1 \leq i \leq n/2$, inputs i and $i + n/2$ of $B(n)$ can be routed

to the i th input of $B_U(n/2)$ or to the i th input of $B_D(n/2)$. The outputs are connected similarly. For every $1 \leq i \leq n/2$ we define inputs i and $i + n/2$ as *mates*, analogously we define outputs i and $i + n/2$ as *mates*. Note that mates cannot both be routed to the same subnetwork. See Figure 12.

The looping algorithm: A Benes network can realize any permutation σ of its inputs as follows. We start with an arbitrary input, say 1, and route it to $B_U(n/2)$. This implies that the output $\sigma(1)$ is also routed to $B_U(n/2)$. The mate o of $\sigma(1)$ must then be routed to $B_D(n/2)$. This implies that $\sigma^{-1}(o)$ is routed to $B_D(n/2)$. If the mate of $\sigma^{-1}(o)$ is 1 we “completed a cycle” and we start again with an arbitrary input which we haven’t routed yet. Otherwise, if the mate of $\sigma^{-1}(o)$ is not 1 then we route this mate to $B_U(n/2)$ and repeat the process.

Levels of the Benes network: If we lay out the Benes network then the 1st level of the recursion above gives us 2 “stages” consisting of n 2×2 switches, stage #1 connected to the inputs to $B_U(n/2)$ and $B_D(n/2)$, and stage # $2 \log n - 1$ connecting $B_U(n/2)$ and $B_D(n/2)$ to the outputs. Opening the recursion gives us $2 \log n - 1$ stages, each consisting of n 2×2 switches.

To implement any specific permutation, one needs to set each of these switches.

Permuting the $b = w / \log w$ leftmost bits of the word We now describe an $O(1)$ word representation for any permutation σ on $b = w / \log w$ elements that allows us to apply σ to the b leftmost bits of a query word x while doing only $O(\log b)$ operations. We obtain this data structure by encoding the Benes network for σ in $O(1)$ words. To answer a query we use this encoding to apply each of the $2 \log(b) - 1 = O(\log b)$ stages of the Benes network for σ to the leftmost b bits of x . Every stage requires $O(1)$ operations giving a total of $O(\log b) = O(\log w)$ operations.

During preprocessing we prepare two $b \times (2 \log b - 1)$ binary matrices Dir and C . Both Dir and C have $2b \log b - b \leq 2w$ bits, so they can fit into 2 w -bit words. The j th column of these matrices correspond to stage j of the Benes network, the i th row of these matrices corresponds to the i th input of the stage. Pictorially, we imagine that inputs are numbered top-down.

Recall that the mate of input i in stage j is some other input i' of stage j . If $(i' < i)$ we define $\text{Dir}_{i,j} = 0$, otherwise, $(i' > i)$, and we define $\text{Dir}_{i,j} = 1$, this is defined for $i = 1, \dots, b, j = 1, \dots, 2 \log b - 1$.

The matrix C is computed as follows: $C_{i,j} = 0$ if input i of stage j routes to input i of stage $j + 1$ (i.e. goes “straight”), and $C_{i,j} = 1$ otherwise.

We pack the $b \times (2 \log b - 1)$ binary matrix C into 2 w -bit words C_1, C_2 as follows:

$$\begin{aligned} C_1[1, 2, \dots, w] &\leftarrow C_{1,1}, \dots, C_{1,b}, C_{2,1}, \\ &\quad \dots, C_{2,b}, \dots, C_{\log b,1}, \dots, C_{\log b,b}; \\ C_2[1, 2, \dots, w] &\leftarrow C_{\log b+1,1}, \dots, C_{\log b+1,b}, \\ &\quad \dots, C_{2 \log b-1,1}, \dots, C_{2 \log b-1,b}; \end{aligned}$$

We pack the matrix Dir into words Dir_1 and Dir_2 analogously.




Fig. 13. Applying stage 1 of the Benes network to x .

During query processing we apply stage i (for $i = 1 \dots \log b$) of the Benes network by computing

$$\begin{aligned}
 x &= \left(x \text{ AND } \overline{C_1[1, \dots, b]} \right) \\
 &\text{OR} \left((x \text{ AND } C_1[1, \dots, b] \text{ AND } \text{Dir}_1[1, \dots, b]) \gg (b/2^i) \right) \\
 &\text{OR} \left((x \text{ AND } C_1[1, \dots, b] \text{ AND } \overline{\text{Dir}_1[1, \dots, b]}) \ll (b/2^i) \right).
 \end{aligned}$$

This should be parsed as follows:

- $(x \text{ AND } \overline{C_1[1, \dots, b]})$ gives the bits of x that are not going to change position at stage i .
- $(x \text{ AND } C_1[1, \dots, b] \text{ AND } \text{Dir}_1[1, \dots, b]) \gg (b/2^i)$ takes the bits of x that are to move “up” at stage i and shifts them accordingly.
- $(x \text{ AND } C_1[1, \dots, b] \text{ AND } \overline{\text{Dir}_1[1, \dots, b]}) \ll (b/2^i)$ takes the bits of x that are to move “down” at stage i and shifts them accordingly.

In preparation for the next stage we also compute $C_1 = C_1 \ll b$, $\text{Dir}_1 = \text{Dir}_1 \ll b$ to prepare the control bits for the next stage of the Benes network.

Analogously, during stages $i = \log b + 1, \log b + 2, \dots, 2 \log b - 1$ we use the words C_2 and Dir_2 rather than C_1 and Dir_1 .

An example of applying stage 1 of a Benes network of size 8 is shown in figure 13.

Permuting the $w/\log w$ leftmost blocks of the word To operate the permutation on blocks of bits, we need masks that replicate the appropriate $C_{i,j}$ and $\text{Dir}_{i,j}$ values $\log w$ times so that they operate upon all bits of the block simultaneously and not only on one single bit. To precompute and store such replications in advance requires $O(\log w)$ words of storage, and we allow, in total,

only $O(1)$ words of storage for the entire bit selection. Thus, we need compute these “expansions” on the fly, and in $O(\log w)$ operations.

We now add all-zero columns on the left of matrices C and Dir so that each of them they have exactly $2\log w$ columns. This is well defined because $2\log b - 1 \leq 2\log w$. Let these new matrices be \tilde{C} and \tilde{Dir} . Let \tilde{C}^R be the rightmost $\log w$ columns of \tilde{C} , and let \tilde{C}^L be the leftmost $\log w$ columns of \tilde{C} . Also, let \tilde{Dir}^R , and \tilde{Dir}^L be defined analogously.

Previously, we packed the C and Dir matrices into words (C_1, C_2) and (Dir_1, Dir_2) , respectively, column by column. To perform Block permutations we do so row by row as follows: The matrix \tilde{C}^L is packed into the word C_1 , row by row. Likewise, \tilde{C}^R is packed, row by row, into C_2 , \tilde{Dir}^L into Dir_1 and \tilde{Dir}^R into Dir_2 .

The C and Dir bits associated with stage j of the Benes network will be spaced out, $\log w$ bits apart. See Figure 14. For $j < \log w$ these bits are in C_1 and Dir_1 .

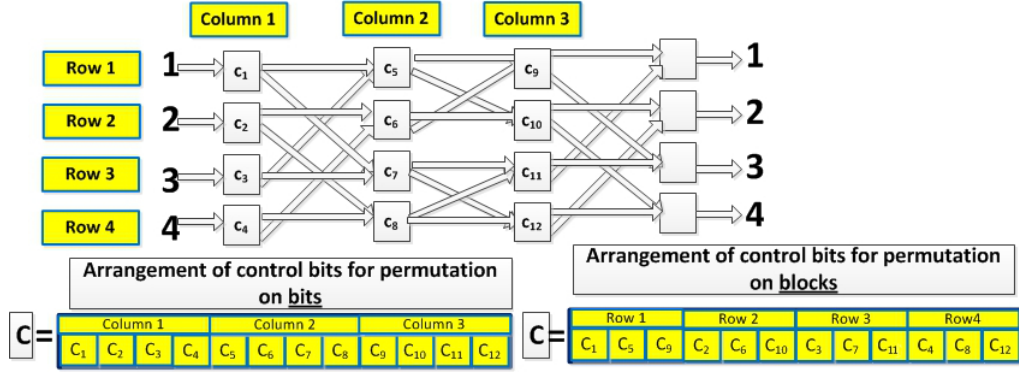


Fig. 14. Benes network control bits $C_{i,j}$ and $Dir_{i,j}$.

Given C_i or Dir_i , we seek to isolate and replicate the bits associated with stage $1 \leq j \leq \log w$. We define a transformation $g : \{0, 1\}^w \times \{1, \dots, \log w\}$ such that for any w bit word z , and any $1 \leq j \leq \log w$, $g(z, j)$ is a mask such that for any block B , all bits of $g(z, j)[B]$ are equal to $z[B[j]]$.

We compute $g(z, j)$ in $O(1)$ time as follows: Let Z_j , $1 \leq j \leq \log w$, be a bit pattern with $w/\log w$ 1's at the j th index of every block. The operation $y_0 = Z_j$ AND z isolates the j 'th bits of every block in z . Let $y_1 = y_0 \ll j - 1$ and $y_2 = y_0 \gg (\log w - j)$, let $y_3 = y_1 - y_2$. Blocks B for which the bit $z[B[j]] = 1$, now have $y_3[B] = 011 \dots 1$, blocks B for which the bit was zero now have $y_3[B]$ consisting only of zeros. Finally, set $y_4 = y_3$ OR y_1 . Now, all bits of $y_4[B]$ are equal to the bit $z[B[j]]$.

To simulate the Benes network and sort blocks rather than bits, for stages $j \leq \log w$ we use the masks $g(C_1, j)$ and $g(Dir_1, j)$, analogously to our use of the

masks $C_1[1, \dots, b]$ and $\text{Dir}_1[1, \dots, b]$ as used in Equation 2. For stages $j > \log w$ we use $g(C_2, j - \log w)$ and $g(\text{Dir}_2, j - \log w)$ analogously to the use of $C_2[1, \dots, b]$ and $\text{Dir}_2[1, \dots, b]$.

Given this transformation, we can simulate the Benes network in parallel, on entire blocks, and permute blocks at no greater cost than permuting *bits*.