Succinct Partial Sums and Fenwick Trees

Philip Bille, Anders Roy Christiansen, Nicola Prezza, and Frederik Rye
Skjoldjensen

Technical University of Denmark, DTU Compute, Kgs. Lyngby, Denmark
{phbi,aroy,npre,fskj}@dtu.dk

Abstract. We consider the well-studied *partial sums* problem in succint
space where one is to maintain an array of n k-bit integers subject to
updates such that partial sums queries can be efficiently answered. We
present two succint versions of the Fenwick Tree – which is known for
its simplicity and practicality. Our results hold in the encoding model
where one is allowed to reuse the space from the input data. Our main
result is the first that only requires $nk + o(n)$ bits of space while still
supporting sum/update in $\mathcal{O}(\log_b n)$ / $\mathcal{O}(b \log_b n)$ time where $2 \leq b \leq$
$\log^{\mathcal{O}(1)} n$. The second result shows how optimal time for sum/update
can be achieved while only slightly increasing the space usage to $nk +$
$o(nk)$ bits. Beyond Fenwick Trees, the results are primarily based on
bit-packing and sampling – making them very practical – and they also
allow for simple optimal parallelization.

Keywords: Partial sums, Fenwick tree, succinct, parallel

1 Introduction

Let A be an array of k-bits integers, with $|A| = n$. The *partial sums* problem is
to build a data structure maintaining A under the following operations.

- $\mathtt{sum}(i)$: return the value $\sum_{t=1}^{i} A[t]$.
- $\mathtt{search}(j)$: return the smallest i such that $\mathtt{sum}(i) \geq j$.
- $\mathtt{update}(i, \Delta)$: set $A[i] \leftarrow A[i] + \Delta$, for some Δ such that $0 \leq A[i] + \Delta < 2^k$.
- $\mathtt{access}(i)$: return $A[i]$.

Note that $\mathtt{access}(i)$ can implemented as $\mathtt{sum}(i) - \mathtt{sum}(i-1)$ and we therefore
often do not mention it explicitly.

The partial sums problem is one the most well-studied data structure prob-
lems [1, 2, 3, 4, 6, 7, 8, 9]. In this paper, we consider solutions to the partial sums
problem that are *succinct*, that is, we are interested in data structures that
use space close to the information-theoretic lower bound of nk bits. We distin-
guish between *encoding data structures* and *indexing data structures*. Indexing
data structures are required to store the input array A verbatim along with ad-
ditional information to support the queries, whereas encoding data structures
have to support operations without consulting the input array.

In the indexing model Raman et al. [8] gave a data structure that supports `sum`, `update`, and `search` in $\mathcal{O}(\log n/\log\log n)$ time while using $nk+o(nk)$ bits of space. This was improved and generalized by Hon et al. [6]. Both of these papers have the constrain $\Delta \leq \log^{\mathcal{O}(1)} n$. The above time complexity is nearly optimal by a lower bound of Patrascu and Demaine [7] who showed that `sum`, `search`, and `update` operations takes $\Theta(\log_{w/\delta} n)$ time per operation, where $w \geq \log n$ is the word size and δ is the number of bits needed to represent Δ. In particular, whenever $\Delta = \log^{\mathcal{O}(1)} n$ this bound matches the $\mathcal{O}(\log n/\log\log n)$ bound of Raman et al. [8].

Fenwick [2] presented a simple, elegant, and very practical encoding data structure. The idea is to replace entries in the input array A with partial sums that cover A in an implicit complete binary tree structure. The operations are then implemented by accessing at most $\log n$ entries in the array. The Fenwick tree uses $nk + n\log n$ bits and supports all operations in $O(\log n)$ time. In this paper we show two succinct b-ary versions of the Fenwick tree. In the first version we reduce the size of the Fenwick tree while improving the `sum` and `update` time. In the second version we obtain optimal times for `sum` and `update` without using more space than the previous best succinct solutions [6, 8]. All results in this paper are in the RAM model.

Our results We show two encoding data structures that gives the following results.

Theorem 1. *We can replace A with a succinct Fenwick tree of $nk + o(n)$ bits supporting `sum`, `update`, and `search` queries in $\mathcal{O}(\log_b n)$, $\mathcal{O}(b\log_b n)$, and $\mathcal{O}(\log n)$ time, respectively, for any $2 \leq b \leq \log^{\mathcal{O}(1)} n$.*

Theorem 2. *We can replace A with a succinct Fenwick tree of $nk + o(nk)$ bits supporting `sum` and `update` queries in optimal $\mathcal{O}(\log_{w/\delta} n)$ time and `search` queries in $\mathcal{O}(\log n)$ time.*

2 Data structure

For simplicity, assume that n is a power of 2. The Fenwick tree is an implicit data structure replacing a word-array $A[1,\ldots,n]$ as follows:

Definition 1. *Fenwick tree of A [2]. If $n = 1$, then leave A unchanged. Otherwise, divide A in consecutive non-overlapping blocks of two elements each and replace the second element $A[2i]$ of each block with $A[2i-1] + A[2i]$, for $i = 1, \ldots, n/2$. Then, recurse on the sub-array $A[2, 4, \ldots, 2i, \ldots, n]$.*

To answer $sum(i)$, the idea is to write i in binary as $i = 2^{j_1} + 2^{j_2} + \cdots + 2^{j_k}$ for some $j_1 > j_2 > \cdots > j_k$. Then there are $k \leq \log n$ entries in the Fenwick tree, that can be easily computed from i, whose values added together yield $sum(i)$. In Section 2.1 we describe in detail how to perform such accesses. As per the above definition, the Fenwick tree is an array of n indices. If represented compactly, this array can be stored in $nk+n\log n$ bits. In this section we present a generalization of Fenwick trees taking only succinct space.

2.1 Layered b-ary structure

We first observe that it is easy to generalize Fenwick trees to be b-ary, for $b \geq 2$: we divide A in blocks of b integers each, replace the first $b-1$ elements in each block with their partial sum, and fill the remaining n/b entries of A by recursing on the array A' of size n/b that stores the sums of each block. This generalization gives an array of n indices supporting sum, update, and search queries on the original array in $\mathcal{O}(\log_b n)$, $\mathcal{O}(b \log_b n)$, and $\mathcal{O}(\log n)$ time, respectively. We now show how to reduce the space of this array.

Let $\ell = \log_b n$. We represent our b-ary Fenwick tree $T_b(A)$ using $\ell + 1$ arrays (layers) $T_b^1(A), \ldots, T_b^{\ell+1}(A)$. For simplicity, we assume that $n = b^e$ for some $e \geq 0$ (the general case is then straightforward to derive). To improve readability, we define our layered structure for the special case $b = 2$, and then sketch how to extend it to the general case $b \geq 2$. Our layered structure is defined as follows. If $n = 1$, then $T_2^1(A) = A$. Otherwise:

- $T_2^{\ell+1}(A)[i] = A[(i-1) \cdot 2 + 1]$, for all $i = 1, \ldots, n/2$. Note that $T_2^{\ell+1}(A)$ contains $n/2$ elements.

- Divide A in blocks of 2 elements each, and build an array $A'[j]$ containing the $n/2$ sums of each block, i.e. $A'[j] = A[(j-1) \cdot 2 + 1] + A[(j-1) \cdot 2 + 2]$, for $j = 1, \ldots, n/2$. Then, the next layers are recursively defined as $T_2^\ell(A) \leftarrow T_2^\ell(A'), \ldots, T_2^1(A) \leftarrow T_2^1(A')$.

For general $b \geq 2$, $T_b^{\ell+1}(A)$ is an array of $\frac{n(b-1)}{b}$ elements that stores the $b-1$ partial sums of each block of b consecutive elements in A, while A' is an array of size n/b containing the complete sums of each block. In Figure 1 we report an example of our layered structure with $b = 3$. It follows that elements of $T_b^i(A)$, for $i > 1$, take at most $k + (\ell - i + 2) \log b$ bits each. Note that arrays $T_b^1(A), \ldots, T_b^{\ell+1}(A)$ can easily be packed contiguously in a word array while preserving constant-time access to each of them. This saves us $\mathcal{O}(\ell)$ words that would otherwise be needed to store pointers to the arrays. Let $S_b(n, k)$ be the space (in bits) taken by our layered structure. This function satisfies the recurrence

$$S_b(1, k) = k$$

$$S_b(n, k) = \frac{n(b-1)}{b} \cdot (k + \log b) + S_b(n/b, k + \log b)$$

Which unfolds to $S_b(n, k) = \sum_{i=1}^{\log_b n + 1} \frac{n(b-1)}{b^i} \cdot (k + i \log b)$. Using the identities $\sum_{i=1}^{\infty} 1/b^i = 1/(b-1)$ and $\sum_{i=1}^{\infty} i/b^i = b/(b-1)^2$, one can easily derive that $S_b(n, k) \leq nk + 2n \log b$.

We now show how to obtain the time bounds stated in Theorem 1. In the next section, we reduce the space of the structure without affecting query times.

Answering sum Let the notation $(x_1 x_2 \ldots x_t)_b$, with $0 \leq x_i < b$ for $i = 1, \ldots, t$, represent the number $\sum_{i=1}^{t} b^{t-i} x_i$ in base b. *sum(i)* queries on our structure are a generalization (in base b) of *sum(i)* queries on standard Fenwick trees.

Consider the base-b representation $x_1 x_2 \ldots x_{\ell+1}$ of i, i.e. $i = (x_1 x_2 \ldots x_{\ell+1})_b$ (note that we have at most $\ell+1$ digits since we enumerate indexes starting from 1). Consider now all the positions $1 \le i_1 < i_2 < \cdots < i_t \le \ell+1$ such that $x_j \ne 0$, for $j = i_1, \ldots, i_t$. The idea is that each of these positions $j = i_1, \ldots, i_t$ can be used to compute an offset o_j in $T_b^j(A)$. Then, $sum(i) = \sum_{j=i_1,\ldots,i_t} T_b^j(A)[o_j]$. The offset o_j relative to the j-th most significant (nonzero) digit of i is defined as follows. If $j = 1$, then $o_j = x_1$. Otherwise, $o_j = (b-1) \cdot (x_1 \ldots x_{j-1})_b + x_j$. Note that we scale by a factor of $b - 1$ (and not b) as the first term in this formula as each level $T^j(A)$ stores only $b - 1$ out of b partial sums (the remaining sums are passed to level $j - 1$). Note moreover that each o_j can be easily computed in constant time and *independently from the other offsets* with the aid of modular arithmetic. It follows that `sum` queries are answered in $\mathcal{O}(\log_b n)$ time. See Figure 1 for a concrete example of `sum`.

Answering update The idea for performing $update(i, \Delta)$ is analogous to that of $sum(i)$. We access all levels that contain a partial sum covering position i and update at most $b-1$ sums per level. Using the same notation as above, for each $j = i_1, \ldots, i_t$ such that $x_j \ne 0$, we update $T_b^j(A)[o_j + l] \leftarrow T_b^j(A)[o_j + l] + \Delta$ for $l = 0, \ldots, b - x_j - 1$. This procedure takes $\mathcal{O}(b \log_b n)$ time.

Answering search To answer $search(j)$ we start from $T_b^1(A)$ and simply perform a top-down traversal of the implicit B-tree of degree b defined by the layered structure. At each level, we perform $\mathcal{O}(\log b)$ steps of binary search to find the new offset in the next level. There are $\log_b n$ levels, so `search` takes overall $\mathcal{O}(\log n)$ time.
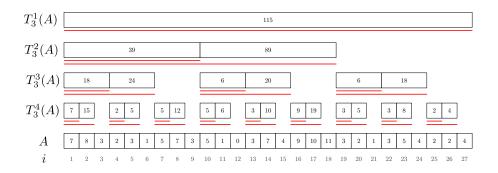

Fig. 1. Example of our layered structure with $n = 27$ and $b = 3$. Horizontal red lines show the portion of A covered by each element in $T_3^j(A)$, for $j = 1, \ldots, \log_b n + 1$. To access the i-th partial sum, we proceed as follows. Let, for example, $i = 19 = (0201)_3$. The only nonzero digits in i are the 2-nd and 4-th most significant. This gives us $o_2 = 2 \cdot (0)_3 + 2 = 2$ and $o_4 = 2 \cdot (020)_3 + 1 = 13$. Then, $sum(19) = T_3^2(A)[2] + T_3^4(A)[13] = 89 + 3 = 92$.

2.2 Sampling

Let $0 < d \leq n$ be a sample rate, where for simplicity we assume that d divides n. Given our input array A, we derive an array A' of n/d elements containing the sums of groups of d adjacent elements in A, i.e. $A'[i] = \sum_{j=1}^{d} A[(i-1) \cdot d + j]$, $i = 1, \ldots, d$. We then compact A by removing $A[j \cdot d]$ for $j = 1, \ldots, n/d$, and by packing the remaining integers in at most $nk(1 - 1/d)$ bits. We build our layered b-ary Fenwick tree $T_b(A')$ over A'. It is clear that queries on A can be solved with a query on $T_b(A')$ followed by at most d accesses on (the compacted) A. The space of the resulting data structure is $nk(1 - 1/d) + S_b(n/d, k + \log d) \leq nk + \frac{n \log d}{d} + \frac{2n \log b}{d}$ bits. In order to retain the same query times of our basic layered structure, we choose $d = (1/\epsilon) \log_b n$ for any constant $\epsilon > 0$ and obtain a space occupancy of $nk + \epsilon \left(\frac{n \log \log_b n}{\log_b n} + \frac{2n \log b}{\log_b n} \right)$ bits. For $b \leq \log^{\mathcal{O}(1)} n$, this space is $nk + o(n)$ bits. Note that—as opposed to existing succinct solutions—the low-order term does not depend on k.

3 Optimal-time `sum` and `update`

In this section we show how to obtain optimal running times for `sum` and `update` queries in the RAM model. We can directly apply the word-packing techniques described in [7] to speed-up queries; here we only sketch this strategy, see [7] for full details. Let us describe the idea on the structure of Section 2.1, and then plug in sampling to reduce space usage. We divide arrays $T_b^j(A)$ in blocks of $b - 1$ entries, and store one word (w bits) for each such block. We can pack $b - 1$ integers of at most $w/(b - 1)$ bits each (for an opportune b, read below) in the word associated with each block. Since blocks of $b - 1$ integers fit in a single word, we can easily answer `sum` and `update` queries on them in constant time. `sum` queries on our overall structure can be answered as described in Section 2.1, except that now we also need to access one of the packed integers at each level j to correct the value read from $T_b^j(A)$. To answer `update` queries, the idea is to perform `update` operations on the packed blocks of integers in constant time exploiting bit-parallelism instead of updating at most $b - 1$ values of $T_b^j(A)$. At each `update` operation, we transfer one of these integers on $T_b^j(A)$ (in a cyclic fashion) to avoid overflowing and to achieve worst-case performance. Note that each packed integer is increased by at most Δ for at most $b - 1$ times before being transferred to $T_b^j(A)$, so we get the constraint $(b - 1) \log((b - 1)\Delta) \leq w$. We choose $(b - 1) = \frac{w}{2(\log w + \delta)}$. Then, it is easy to show that the above constraint is satisfied. The number of levels becomes $\log_b n = \mathcal{O}(\log_{w/\delta} n)$. Since we spend constant time per level, this is also the worst-case time needed to answer `sum` and `update` queries on our structure. To analyze space usage we use the corrected formula

$$S_b(1, k) = k$$

$$S_b(n, k) = \frac{n(b-1)}{b} \cdot (k + \log b) + \frac{nw}{b} + S_b(n/b, k + \log b)$$

yielding $S_b(1,k) \leq nk + 2n \log b + \frac{nw}{b-1}$. Replacing $b - 1 = \frac{w}{2(\log w + \delta)}$ we achieve $nk + \mathcal{O}(n\delta + n \log w)$ bits of space.

We now apply the sampling technique of Section 2.2 with a slight variation. In order to get the claimed space/time bounds, we need to further apply bit-parallelism techniques on the packed integers stored in A: using techniques from [5], we can answer `sum`, `search`, and `update` queries in $\mathcal{O}(1)$ time on blocks of w/k integers. It follows that we can now use sample rate $d = \frac{w \log n}{k \log(w/\delta)}$ without affecting query times. After sampling A and building the Fenwick tree above described over the sums of size-d blocks of A, the overall space is $nk(1 - 1/d) + S_b(n/d, k + \log d) = nk + \frac{n \log d}{d} + \mathcal{O}(\frac{n\delta}{d} + \frac{n \log w}{d})$. Note that $d \leq \frac{w^2}{k \log(w/\delta)} \leq w^2$, so $\log d \in \mathcal{O}(\log w)$ and space simplifies to $nk + \mathcal{O}(\frac{n\delta}{d} + \frac{n \log w}{d})$. The term $\frac{n\delta}{d}$ equals $\frac{n\delta k \log(w/\delta)}{w \log n}$. Since $\delta \leq w$, then $\delta \log(w/\delta) \leq w$, and this term therefore simplifies to $\frac{nk}{\log n} \in o(nk)$. Finally, the term $\frac{n \log w}{d}$ equals $\frac{n \log w \cdot k \log(w/\delta)}{w \log n} \leq \frac{nk}{(w \log n)/(\log w)^2} \in o(nk)$. The bounds of Theorem 2 follow.

Parallelism Note that `sum` and `update` queries on our succinct Fenwick trees can be naturally parallelized as all accesses/updates on the levels can be performed independently from each other. For sum, we need $\mathcal{O}(\log \log_b n)$ further time to perform a parallel sum of the $\log_b n$ partial results. It is not hard to show that—on architectures with $\log_b n$ processors—this reduces `sum`/`update` times to $\mathcal{O}(\log \log_b n)/\mathcal{O}(b)$ and $\mathcal{O}(\log \log_{w/\delta} n)/\mathcal{O}(1)$ in Theorems 1 and 2, respectively.

References

1. Dietz, P.F.: Optimal algorithms for list indexing and subset rank. In: Proc. 1st WADS. pp. 39–46 (1989)
2. Fenwick, P.M.: A new data structure for cumulative frequency tables. Software: Practice and Experience 24(3), 327–336 (1994)
3. Fredman, M., Saks, M.: The cell probe complexity of dynamic data structures. In: Proc. 21st STOC. pp. 345–354 (1989)
4. Fredman, M.L.: The complexity of maintaining an array and computing its partial sums. Journal of the ACM (JACM) 29(1), 250–260 (1982)
5. Hagerup, T.: Sorting and searching on the word ram. In: STACS 98. pp. 366–398. Springer (1998)
6. Hon, W.K., Sadakane, K., Sung, W.K.: Succinct data structures for searchable partial sums with optimal worst-case performance. Theoretical Computer Science 412(39), 5176–5186 (2011)
7. Patrascu, M., Demaine, E.D.: Logarithmic lower bounds in the cell-probe model. SIAM Journal on Computing 35(4), 932–963 (2006)
8. Raman, R., Raman, V., Rao, S.S.: Succinct dynamic data structures. In: Workshop on Algorithms and Data Structures. pp. 426–437. Springer (2001)
9. Yao, A.C.: On the complexity of maintaining partial sums. SIAM Journal on Computing 14(2), 277–288 (1985)