

Sorting and Counting Networks of Small Depth and Arbitrary Width

Costas Busch and Maurice Herlihy
Computer Science Department,
Brown University,
Providence, RI 02912;
{cb,herlihy}@cs.brown.edu

October 28, 1999

Abstract

We present the first construction for sorting and counting networks of arbitrary width that uses both small depth and small constant factors. Let w be the product $w = p_0 \cdots p_{n-1}$, whose factors are not necessarily prime. We present a novel network construction of width w and depth $O(n^2) = O(\log^2 w)$, using comparators (or balancers) of width less than or equal to $\max(p_i)$. This construction is practical in the sense that the asymptotic notation does not hide any large constants.

An interesting aspect of this construction is that it establishes a family of sorting and counting networks of width w , one for each distinct factorization of w . A factorization in which $\max(p_i)$ is large and n is small yields a network that trades small depth for large comparators (or balancers), and a factorization where $\max(p_i)$ is small and n is large makes the opposite trade-off.

1 Introduction

A *sorting network* [2, 4, 7, 8] is a class of parallel data structures used for sorting. Sorting networks are constructed from p -input p -output synchronous switches called *p -comparators*. As illustrated on the left-hand side of Figure 1, a comparator accepts p values on its input wires, and outputs those same values in sorted order on its output wires. A *comparator network* is an acyclic network of comparators where output wires of some comparators are linked to input wires of others. The network's *input wires* are those input wires not linked to an output, and similarly for the network's *output wires*. In this paper, we restrict our attention to networks with the same number of input and output wires, called the network's *width*. Values enter the network on the input wires, one per input wire, propagate in lock-step through the comparators, and leave on the output wires, one per output wire. Each comparator reorders its input

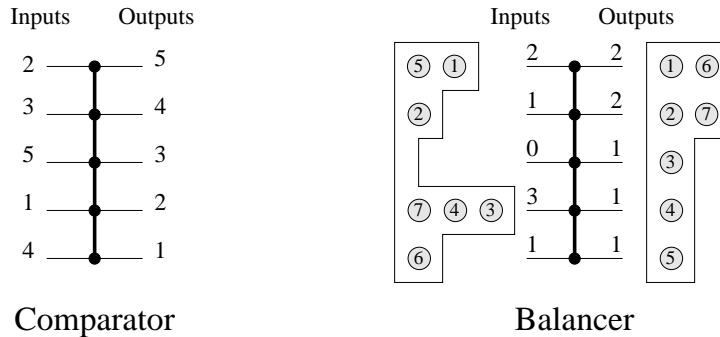



Figure 1: Comparator (left) and balancer (right)

values, sending the i -th ranked input to the i -th output wire. If the network's i -th ranked input emerges on the network's i -th output wire, then the network is a *sorting network*. The network's *depth* is the maximum number of comparators traversed by any value. A sorting network is illustrated in the top half of Figure 2.


A *counting network* [3] is a class of distributed data structures used to construct concurrent, low-contention implementations of *Fetch&Increment* counters. Counting networks are constructed from p -input p -output asynchronous switches called p -balancers. As illustrated on the right-hand side of Figure 1, a balancer accepts a stream of tokens on its p input wires, and the i -th token to enter leaves on output wire $i \bmod p$. A *balancing network* is an acyclic network of balancers where output wires of some balancers are linked to input wires of others. The network's *input wires*, *output wires*, *width*, and *depth* are defined just as for comparator networks. Tokens enter the network on the input wires, typically several per wire, propagate asynchronously through the balancers, and leave on the output wires, typically several per wire. A balancing network is a *counting network* if the overall distribution of output tokens across the output wires satisfies the *step property*: exiting tokens are divided uniformly among the output wires, and any excess tokens emerge on the upper wires (a formal definition is given below). A counting network is illustrated in the bottom half of Figure 2.

Counting and sorting networks behave differently: a sorting network of width w sorts values synchronously in batches of w , while counting networks count an arbitrary number of tokens asynchronously. Nevertheless, counting networks and sorting networks are related in a simple way: every counting network is *isomorphic* to a sorting network, that is, if we replace each balancer in a counting network with a comparator, then the result is a sorting network [3]. Figure 2 shows two isomorphic networks constructed from comparators or balancers of sizes two, three and five.

The converse, however, is false: replacing each comparator in a sorting network with a balancer does not necessarily yield a counting network. Figure 3



Sorting Network



Counting Network

Figure 2: Isomorphic Sorting and Counting Networks




Figure 3: Bubble Sorting Network

shows a simple counterexample: This network is a sorting network (based on *bubblesort*), but the figure illustrates why it is not also a counting network.

In this paper, we present new network constructions that illuminate how width, depth, and balancer size can be traded off in both sorting and counting networks. Specifically, we present the first network construction of arbitrary width w that requires both small depth and small constant factors. Let w be the product $w = p_0 \cdots p_{n-1}$, whose factors are not necessarily prime. We construct a network of width w and depth $O(n^2) = O(\log^2 w)$, using comparators or balancers of width at most $\max(p_i)$. This construction is practical in the sense that the asymptotic notation does not hide any large constants.

For brevity, the terms “counting network” and “balancer” in our constructions stand for “sorting or counting network” and “comparator or balancer”, respectively.

An interesting aspect of this construction is that it establishes a family of counting networks of width w , one for each distinct factorization of w . A factorization in which $\max(p_i)$ is large and n is small yields a network that trades small depth for large balancers, and a factorization where $\max(p_i)$ is small and n is large makes the opposite trade-off. This flexibility may be useful in practice, since experimental evidence [9] suggests that for shared-memory implementations of counting networks, optimal performance for a fixed w is achieved by balancers of intermediate size. (Each distinct ordering of a fixed set of factors also yields a different counting network, but all such networks have the same depth.)

2 Related Work

Knuth [11, Prob. 5.3.4.44] was the first to raise the question of properties of sorting networks constructed from k -comparators for $k > 2$, asking whether there are efficient ways to sort k^2 elements using k -comparators. This paper answers the natural generalization of this question to arbitrary factorizations.

There are several sorting network constructions that use comparators of size $p \geq 2$. Chvatál [6] modified the AKS sorting network to use comparators of size p instead of size 2. Tseng and Lee [17] construct a sorting network of width $w = p^k$ and depth $O(p \log^2 w)$ from comparators of size p . Parker and Parberry




Figure 4: Wires

[16] present a sorting network construction of width $w = p^k$ and depth $O(\log^2 w)$ from balancers of width p , where p must have an integer square root. Lee and Batcher [12] present a multi-way merge sorting network, a generalization of the odd-even sorting network, that could be used to construct a sorting network of arbitrary width $w = p_0, \dots, p_{n-1}$ and depth $O((\lg^2 p_m) \log^2 w)$, from balancers of size at most $\max(p_i)$, where p_m is at least as big as the median of p_0, \dots, p_{n-1} .

The first counting network constructions [3] used 2-balancers, yielding networks of width 2^n and depth $O(n^2)$. Aharonson and Attiya [1] constructed a counting network of width $w = p2^k$ and depth $O(\lg^3(w/p))$ from balancers of size 2 and p . They also construct networks of arbitrary width by taking a standard counting network and linking the excess output wires to the excess input wires, resulting in a cyclic network (our is acyclic). Busch, Hardavellas, and Mavronicolas [5] give a construction of $w = p2^k$ and depth $O(\lg^2(w/p))$ using balancers of size 2 and p . Felten, LaMarca, and Ladner [9] give a construction of width $w = 2^k$ from balancers of size 2^l , where the depth ranges from $O(1)$ to $O(\log^2 w)$ depending on the value of l , as well as a construction of width $w = p2^k$. Klugerman [10] gives a construction of arbitrary width w and depth $O((\lg w) \lg \lg w)$ from p -balancers, where p ranges over the prime factors of w . This construction is based on the AKS sorting network, and it is impractical in the sense that the constant factors are enormous.

In this abstract, we give a top-down description of the counting network construction. We focus on the modular decomposition of the network. Where alternative constructions exist, we focus on the simplest, adding a brief description of more complicated optimizations. Readers are encouraged to consult the illustrations.

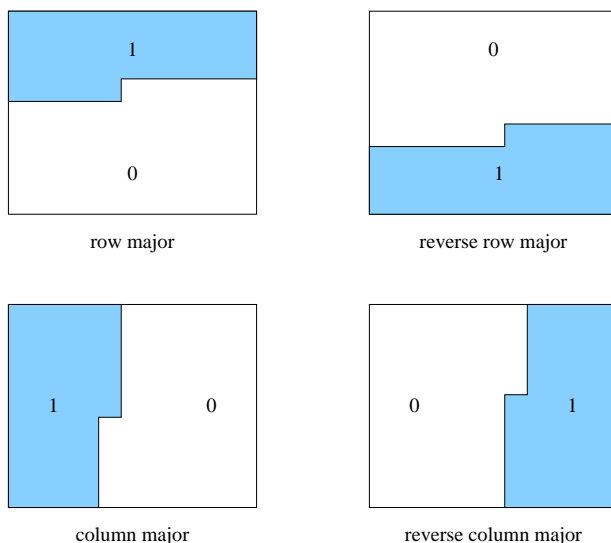


Figure 5: Matrix arrangements

3 Preliminaries

3.1 Sequences

We denote sequences of natural numbers in upper case, and elements of a sequence in lower case. For example, a sequence $X = x_0, \dots, x_{w-1}$ has length (or width) $|X| = w$. Let $\Sigma(X) = x_0 + \dots + x_{w-1}$. The subsequence $X[i, p]$ is the sequence $x_i, x_{i+p}, x_{i+2p}, \dots$

A sequence X of length w has the *step property* if $0 \leq x_i - x_j \leq 1$, for any $0 \leq i < j < w$. If X has the step property, then its *step point* is the unique index i such that $x_i < x_{i-1}$, or 0 if all x_i are equal. X is k -smooth if $|x_i - x_j| \leq k$, for any $0 \leq i, j < w$. The elements of a k -smooth sequence take values in a range $a, a+1, \dots, a+k$. Any sequence satisfying the step property is 1-smooth. In any sequence X we say that there is a *transition* between two consecutive elements x_i and x_{i+1} if their values are different. A sequence X has the *bitonic property* if it is 1-smooth and has at most two transitions. The sequences X_0, \dots, X_{p-1} have the k -staircase property if $0 \leq \Sigma(X_i) - \Sigma(X_j) \leq k$, for any $0 \leq i < j < w$.

It is often convenient to express a sequence X of length rc as an $r \times c$ matrix. There are four ways to arrange the elements of X , as shown by the following table.

x_i goes to	row	column
row major	$\lfloor i/c \rfloor$	$i \bmod c$
reverse row major	$r - \lfloor i/c \rfloor - 1$	$c - (i \bmod c) - 1$
column major	$i \bmod r$	$\lfloor i/r \rfloor$
reverse col. major	$r - (i \bmod r) - 1$	$c - \lfloor i/r \rfloor - 1$




Figure 6: A 5-balancer

These arrangements are illustrated in Figure 5, for a sequence that has the step property. In all figures, the dark region labeled “1” represents the subsequence of higher values, and the light region labeled “0” the lower values.

3.2 Balancing Networks

Henceforth, we consider balancers and balancing networks in *quiescent* states in which no tokens are traversing the network. Let x_i denote the number of tokens that have entered on wire i of p -balancer b (see Figure 6). $X = x_0, \dots, x_{p-1}$ is the *input sequence* of b (Figure 4). The output sequence is defined similarly. Input and output sequences are defined for balancing networks in the same way.

We consider the following balancing network families.

- A *counting network* $\mathcal{C}(p_0, \dots, p_{n-1})$ has input and output sequence of length $w = p_0 \cdots p_{n-1}$. The output sequence has the step property.
- A *merger network* $\mathcal{M}(p_0, \dots, p_{n-1})$ has input sequences $X_0, \dots, X_{p_{n-1}-1}$, where each $|X_i| = p_0 \cdots p_{n-2}$, and output sequence of length $p_0 \cdots p_{n-1}$. If each X_i satisfies the step property, so does the output sequence.
- A *staircase-merger network* $\mathcal{S}(r, p, q)$ has input sequences X_0, \dots, X_{q-1} , where each $|X_i| = rp$, and output sequence of length rpq . If each X_i satisfies the step property, and X_0, \dots, X_{q-1} satisfy the p -staircase property, then the output sequence satisfies the step property.
- A *two-merger network* $\mathcal{T}(p, q_0, q_1)$ has input sequences X_0 and X_1 , where $|X_0| = pq_0$ and $|X_1| = pq_1$, and output sequence of length $p(q_0 + q_1)$. If X_0 and X_1 each satisfies the step property, so does the output sequence.
- A *bitonic-converter network* $\mathcal{D}(p, q)$ has input and output sequence of length pq . If the input sequence satisfies the bitonic property then the output sequence satisfies the step property.

We use \mathcal{C} to refer to the family $\mathcal{C}(p_0, \dots, p_{n-1})$, when the exact values of the p_i are unimportant, and similarly for the other balancing network families. Denote by $\text{depth}(\mathcal{B})$ the depth of a balancing network \mathcal{B} .




Figure 7: Construction of counting network

4 A Counting Network Construction

Let $w = p_0 \cdots p_{n-1}$, and $w_i = p_0 \cdots p_i$, for $0 \leq i < n$, where $p_i \geq 2$ and $n \geq 2$. We give the construction of a counting network $\mathcal{C}(p_0, \dots, p_{n-1})$ of width w and depth $O(n^2)$. Assume for now that we are given the network $\mathcal{C}(p, q)$ with constant depth d , for any $p, q \geq 2$. As discussed below in Section 5, replacing each instance of $\mathcal{C}(p, q)$ with a single pq -balancer yields a counting network family \mathcal{K} of width w and depth $O(n^2)$ from balancers of width at most $\max(p_i p_j)$, while replacing each instance of $\mathcal{C}(p, q)$ with the novel $\mathcal{R}(p, q)$ construction described below in Section 5.3 yields a counting network family \mathcal{L} of width w and depth $O(n^2)$ from balancers of width at most $\max(p_i)$.

4.1 A Counting Network

For the construction of $\mathcal{C}(p_0, \dots, p_{n-1})$, we argue by induction on n , the length of the factorization. For the base case, where $n = 2$, the network $\mathcal{C}(p_0, p_1)$ is given by assumption.

Let $n > 2$, and $\mathcal{C}(p_0, \dots, p_{n-2})$ the counting network guaranteed by the induction hypothesis. Our construction relies on the merger network $\mathcal{M}(p_0, \dots, p_{n-1})$ constructed below in Section 4.2 (see Figure 7). Take p_{n-1} copies of $\mathcal{C}(p_0, \dots, p_{n-2})$, denoted $\mathcal{C}_0, \dots, \mathcal{C}_{p_{n-1}-1}$. Split the input sequence X of length w into subsequences $X_0, \dots, X_{p_{n-1}-1}$, each of length w_{n-2} . Direct each X_i to \mathcal{C}_i , and let Y_i be the corresponding output sequence. Each Y_i has the step property. Direct the $Y_0, \dots, Y_{p_{n-1}-1}$ to $\mathcal{M}(p_0, \dots, p_{n-1})$. The resulting output has the step property.

Next, we compute the depth of \mathcal{C} in terms of the constant depth of the staircase-merger \mathcal{S} , presented in Section 4.2, and the constant depth d of $\mathcal{C}(p_0, p_1)$.

Proposition 1 $\text{depth}(\mathcal{C}(p_0, \dots, p_{n-1})) = (n-1)d + (n^2/2 - 3n/2 + 1) \cdot \text{depth}(\mathcal{S})$.

Proof: From the inductive construction of $\mathcal{C}(p_0, \dots, p_{n-1})$ we have:

$$\text{depth}(\mathcal{C}(p_0, \dots, p_{n-1}))$$

$$\begin{aligned}
&= \text{depth}(\mathcal{C}(p_0, \dots, p_{n-2})) + \text{depth}(\mathcal{M}(p_0, \dots, p_{n-1})) \\
&= \text{depth}(\mathcal{C}(p_0, \dots, p_{n-3})) + \text{depth}(\mathcal{M}(p_0, \dots, p_{n-2})) \\
&\quad + \text{depth}(\mathcal{M}(p_0, \dots, p_{n-1})) \\
&= \dots \\
&= \text{depth}(\mathcal{C}(p_0, p_1)) + \text{depth}(\mathcal{M}(p_0, p_1, p_2)) + \dots \\
&\quad + \text{depth}(\mathcal{M}(p_0, \dots, p_{n-1})) \\
&= d + (d + (3 - 2) \cdot \text{depth}(\mathcal{S})) + \dots \\
&\quad + (d + (n - 2) \cdot \text{depth}(\mathcal{S})) \\
&\quad \text{(by Proposition 3)} \\
&= (n - 1)d + ((3 + \dots + n) - 2(n - 2)) \cdot \text{depth}(\mathcal{S}) \\
&= (n - 1)d + ((n(n + 1)/2 - 3) - 2(n - 2)) \cdot \text{depth}(\mathcal{S}) \\
&= (n - 1)d + (n^2/2 - 3n/2 + 1) \cdot \text{depth}(\mathcal{S}).
\end{aligned}$$

■

4.2 A Merger Network




Figure 8: Construction of merger network

We now show how to construct the merger network $\mathcal{M}(p_0, \dots, p_{n-1})$. This construction relies on the staircase-merger network \mathcal{S} constructed below in Section 4.3.

We argue by induction on n . For the base case, where $n = 2$, the network $\mathcal{M}(p_0, p_{n-1})$ is the network $\mathcal{C}(p_0, p_{n-1})$ (given by assumption).

Let $X_0, \dots, X_{p_{n-1}-1}$ be the input sequences and assume that each satisfies the step property. Using the induction hypothesis, construct the merger network $\mathcal{M}(p_0, \dots, p_{n-3}, p_{n-1})$. Take p_{n-2} copies of this network, denoted $\mathcal{M}_0, \dots, \mathcal{M}_{p_{n-2}-1}$. Each \mathcal{M}_i has p_{n-1} input sequences $X_0[i, p_{n-2}], \dots, X_{p_{n-1}-1}[i, p_{n-2}]$. Denote the output sequence of \mathcal{M}_i by Y_i . Now direct each Y_i to the staircase-merger $\mathcal{S}(w_{n-3}, p_{n-1}, p_{n-2})$. The final output sequence satisfies the step property. See Figure 8.

For the correctness of network \mathcal{M} we need only show that the input sequences to the staircase-merger \mathcal{S} satisfy the p_{n-1} -staircase property.

Proposition 2 *The sequences Y_i , for $0 \leq i < p_{n-2}$, satisfy the p_{n-1} -staircase property.*

Proof: Since each X_i has the step property,

$$0 \leq \Sigma(X_i[j, p_{n-2}]) - \Sigma(X_i[k, p_{n-2}]) \leq 1$$

for $0 \leq j < k < p_{n-2}$. By construction,

$$\Sigma(Y_i) = \Sigma(X_0[i, p_{n-2}]) + \cdots + \Sigma(X_{p_{n-1}-1}[i, p_{n-2}]).$$

It follows that for $0 \leq i < j < p_{n-2}$

$$\begin{aligned} \Sigma(Y_i) - \Sigma(Y_j) &= \Sigma(X_0[i, p_{n-2}]) - \Sigma(X_0[j, p_{n-2}]) + \cdots \\ &\quad + \Sigma(X_{p_{n-1}-1}[i, p_{n-2}]) - \Sigma(X_{p_{n-1}-1}[j, p_{n-2}]) \\ &\leq p_{n-1}. \end{aligned}$$

Furthermore, $\Sigma(Y_i) - \Sigma(Y_j) \geq 0$. Subsequently, the Y_i satisfy the p_{n-1} -staircase property, as needed. \blacksquare

Next, we compute the depth of \mathcal{M} in terms of the constant depth of the staircase-merger \mathcal{S} , and the constant depth d of $\mathcal{C}(p_0, p_{n-1})$.

Proposition 3 $depth(\mathcal{M}(p_0, \dots, p_{n-1})) = d + (n-2) \cdot depth(\mathcal{S})$.

Proof: From the inductive construction of $\mathcal{M}(p_0, \dots, p_{n-1})$ we have:

$$\begin{aligned} depth(\mathcal{M}(p_0, \dots, p_{n-1})) &= depth(\mathcal{M}(p_0, \dots, p_{n-3}, p_{n-1})) + depth(\mathcal{S}) \\ &= depth(\mathcal{M}(p_0, \dots, p_{n-4}, p_{n-1})) + depth(\mathcal{S}) + depth(\mathcal{S}) \\ &= \dots \\ &= depth(\mathcal{M}(p_0, \dots, p_{n-k}, p_{n-1})) + (k-2) \cdot depth(\mathcal{S}) \\ &= \dots \\ &= depth(\mathcal{M}(p_0, p_{n-1})) + (n-2) \cdot depth(\mathcal{S}) \\ &= depth(\mathcal{C}(p_0, p_{n-1})) + (n-2) \cdot depth(\mathcal{S}) \\ &= d + (n-2) \cdot depth(\mathcal{S}). \end{aligned}$$

\blacksquare




Figure 9: Construction of staircase-merger network

4.3 A Staircase-Merger

We now show how to construct a constant depth staircase-merger network $\mathcal{S}(r, p, q)$, where $r, p, q \geq 2$. This construction relies on the two-merger network \mathcal{T} constructed below in Section 4.4.

Let X_0, \dots, X_{q-1} be the input sequences and assume that they satisfy the p -staircase property and each satisfies the step property. Let A be the $rp \times q$ matrix such that column i is the sequence X_i , for all $0 \leq i < q$. Because the X_i satisfy the p -staircase property, the step points of the columns of A lie within p (modulo rp) of one another (see Figure 9 (a)). Partition A into $p \times q$ submatrices A_0, \dots, A_{r-1} (A_0 is on the top). The column step points all lie within adjacent A_i and $A_{(i+1) \bmod r}$, for some $0 \leq i < r$ (Figure 9 (b)).

Use $\mathcal{C}(p, q)$ (given by assumption) to give each A_i the step property (see Figure 9 (c), where the sequences A_i are drawn in row major order). Use a layer of two-mergers $\mathcal{T}(p, q, q)$ to merge each A_{2i} and A_{2i+1} (Figure 9 (d)), and a second layer of $\mathcal{T}(p, q, q)$ to merge each A_{2i+1} and $A_{(2i+2) \bmod r}$ (Figure 9 (e)), for $0 \leq i < \lfloor r/2 \rfloor$. If r is odd we need a third layer with one $\mathcal{T}(p, q, q)$ to merge A_0 and A_{r-1} . The resulting matrix A has the step property in row major order and this is the output sequence of the staircase-merger \mathcal{S} .

Since each two-merger \mathcal{T} has depth two, and the depth of $\mathcal{C}(p, q)$ is equal to d we have that $\text{depth}(S) \leq d + 6$. The two-mergers use balancers of width $2q$ and p . If using balancers of size at most $\max(p, q)$, substitute each $2q$ -balancer with a two-merger $\mathcal{T}(q, 1, 1)$ that uses balancers of width 2 and q , yielding $\text{depth}(S) \leq d + 9$.

4.3.1 Optimizations




Figure 10: Optimizing the construction of the staircase-merger

We can improve the depth of \mathcal{S} by replacing the two-mergers with the following construction (see Figure 10). After we have applied the network $\mathcal{C}(p, q)$ the discrepancy on the output lies between two consecutive A_i and $A_{i+1 \bmod r}$, for some i , $0 \leq i < r$. We split each resulting A_i into two equal sized upper and lower parts, respectively, of every two adjacent A_i and $A_{(i+1) \bmod r}$. This moves the discrepancy into a single A_i , for some i , $0 \leq i < r$, where now it has the form of a *bitonic sequence*. A final layer of $\mathcal{C}(p, q)$ counting networks corrects the discrepancy in the A_i . This construction gives $\text{depth}(S) = 2d + 1$. Alternatively, instead of the final layer of $\mathcal{C}(p, q)$, we can use the *bitonic-converter* $\mathcal{D}(p, q)$, described in section 4.4. This construction gives $\text{depth}(S) = d + 3$. Below we give more details for the optimized construction.

After the first layer of $C(p, q)$ networks, we split each A_i into two subsequences A_i^u and A_i^d each of size $s = \lfloor pq/2 \rfloor$, such that they contain the first s and last s elements, respectively, of A_i . Since the A_i have the step property, each of these subsequences has the step property too. A layer ℓ of 2-balancers connects the sequences A_i^d and $A_{(i+1) \bmod r}^u$, for all i , $0 \leq i < r$. Specifically, a 2-balancer connects the j th element of A_i^d with the $(s - 1 - j)$ th element of $A_{(i+1) \bmod r}^u$, for all j , $0 \leq i < s$, such that the first output of each 2-balancer is directed to north, with respect to matrix A . Next, we show that after layer ℓ the discrepancy on the output sequence spans only one A_i .

Proposition 4 *After layer ℓ the discrepancy spans only one A_i , for some i , $0 \leq i < r$, and this A_i satisfies the bitonic property.*

Proof: Before layer ℓ the discrepancy spans at most two A_i and $A_{(i+1) \bmod r}$, for some $0 \leq i < r$. First we consider the case $i \neq r - 1$ where the discrepancy spans two consecutive A_i and A_{i+1} (the other case is described below). These A_i and A_{i+1} are 1-smooth and for simplicity assume that their elements take values 0 and 1 (for higher values the analysis is similar). Denote by z_i and o_i the number of elements of A_i with value 0 and 1, respectively. Note that $z_i + o_i = pq$. By construction we have $o_i \geq o_{i+1}$. There are two possible cases: (a) $0 \leq o_i + o_{i+1} \leq pq$, and (b) $pq < o_i + o_{i+1} \leq 2pq$.

In case (a) (shown in Figure 10), we have $o_{i+1} \leq s$ and $z_i \geq o_{i+1}$. All the o_{i+1} 1s of A_{i+1} are in A_{i+1}^u and at least as many 0s are in A_i^d . Subsequently, the layer of 2-balancers ℓ , that connects the A_i^d and A_{i+1}^u , moves all the 1s from A_{i+1}^u to A_i^d . The A_i^u and A_{i+1}^d remain unaffected. The result is that A_{i+1} contains only 0s, and A_i contains o_i 1s followed by $z_i - o_{i+1}$ 0s followed by o_{i+1} 1s, and thus A_i is bitonic. Therefore, the discrepancy has moved to A_i with the form of a bitonic sequence.

In case (b), we have $z_i \leq s$ and $o_{i+1} \geq z_i$. All the z_i 0s of A_i are in A_i^d and at least as many 1s are in A_{i+1}^u . Subsequently, the layer of 2-balancers ℓ , that connects the A_i^d and A_{i+1}^u , moves all the 0s from A_i^d to A_{i+1}^u . The A_i^u and A_{i+1}^d remain unaffected. The result is that A_i contains only 1s, and A_{i+1} contains z_i 0s followed by $o_{i+1} - z_i$ 1s followed by z_{i+1} 0s, and thus A_{i+1} is bitonic. Therefore, the discrepancy has moved to A_{i+1} with the form of a bitonic sequence.

Next, consider the case $i = r - 1$, where the discrepancy before layer ℓ spans the A_0 and A_{r-1} . In this case, the combination of A_0 and A_{r-1} is 2-smooth and each A_0 and A_{r-1} is 1-smooth. For simplicity assume that the elements of these sequences take values 0, 1, and 2 (for higher values the analysis is similar). Specifically, the elements of A_0 take values 1 and 2 and the elements of A_{r-1} take values 0 and 1. Denote by o_0 and t_0 the number of elements of A_0 with value 1 and 2, respectively, and by z_{r-1} and o_{r-1} the number of elements of A_{r-1} with value 0 and 1, respectively. Note that $o_0 + t_0 = pq$ and $z_{r-1} + o_{r-1} = pq$. By construction we have $o_{r-1} \geq t_0$. Again, there are two possible cases: (a) $0 \leq t_0 + o_{r-1} \leq pq$, and (b) $pq < t_0 + o_{r-1} \leq 2pq$.

In case (a) we have $t_0 \leq s$ and $z_{r-1} \geq t_0$. All the t_0 2s of A_0 are in A_0^u and at least as many 0s are in A_{r-1}^d . Subsequently, the layer of 2-balancers ℓ , that connects the A_0^u and A_{r-1}^d , transforms the 2s of A_0^u to 1s and the same number of 0s of A_{r-1}^d to 1s. The A_0^d and A_{r-1}^u remain unaffected. The result is that A_0 contains only 1s, and A_{r-1} contains o_{r-1} 1s followed by $z_{r-1} - t_0$ 0s followed by t_0 1s, and thus A_{r-1} is bitonic. Therefore, the discrepancy has moved to A_{r-1} with the form of a bitonic sequence.

In case (b), we have $z_{r-1} \leq s$ and $t_0 \geq z_{r-1}$. All the z_{r-1} 0s of A_{r-1} are in A_{r-1}^d and at least as many 2s are in A_0^u . Subsequently, the layer of 2-balancers ℓ , that connects the A_0^u and A_{r-1}^d , transforms the 0s of A_{r-1}^d to 1s and the same number of 2s of A_0^u to 1s. The A_0^d and A_{r-1}^u remain unaffected. The result is that A_{r-1} contains only 1s, and A_0 contains z_{r-1} 1s followed by $t_0 - z_{r-1}$ 2s followed by o_0 1s, and thus A_i is bitonic. Therefore, the discrepancy has moved to A_0 with the form of a bitonic sequence. ■

Since after layer ℓ the discrepancy is in a single A_i with the form of a bitonic sequence, we can correct the discrepancy by using for each A_i either a final layer of the counting network $\mathcal{C}(p, q)$ or the bitonic-converter $\mathcal{D}(p, q)$. The resulting output sequence of the staircase-merger \mathcal{S} has the step property.

4.4 A Two-Merger and a Bitonic-Converter




Figure 11: Construction of two-merger network

First, we construct the *two-merger* network $\mathcal{T}(p, q_0, q_1)$ of depth two from $(q_0 + q_1)$ -balancers and p -balancers, where $p \geq 2$ and $q_0, q_1 \geq 1$.

Let X_0 and X_1 be the input sequences with respective lengths pq_0 and pq_1 . As illustrated in Figure 11, arrange X_0 as a $p \times q_0$ matrix in column-major form, X_1 as a $p \times q_1$ matrix in reverse column major form, and align the two matrices side by side. If we place a $(q_0 + q_1)$ -balancer across each row only one column is 1-smooth. If we then place a p -balancer across each column, the result has the step property (as a matrix in column-major form). More precisely, we have the following.

Proposition 5 *The network $\mathcal{T}(p, q_0, q_1)$ is a two-merger.*

Proof: let X_0 assume values a_0 and $a_0 + 1$, and let (r_0, c_0) be the step point's row and column in the $p \times q_0$ matrix. Define a_1 and (r_1, c_1) similarly for X_1 . Suppose $r_0 \leq r_1$ (the other case is similar). Consider the row sums for the combined $p \times (q_0 + q_1)$ matrix. Let s_r the sum of the elements of row r , for $0 \leq r \leq p - 1$. We have

row r	sum s_r
$r < r_0$	$s_r = q_0 a_0 + (c_0 + 1) + q_1 a_1 + c_1$
$r_0 \leq r < r_1$	$s_r = q_0 a_0 + c_0 + q_1 a_1 + c_1$
$r_1 \leq r$	$s_r = q_0 a_0 + c_0 + q_1 a_1 + (c_1 + 1)$.

The sequence s_0, \dots, s_{p-1} is 1-smooth. Therefore, after the first (horizontal) layer of balancers, all the step points of the balancers will appear in at most two consecutive columns (modulo $q_0 + q_1$). As a result, the matrix has a single column c such that all elements of columns to the left have some value $d + 1$, all elements to the right have value d , and all elements of column c are 1-smooth with values d or $d + 1$. After the second (vertical) layer of balancers, columns to the left and right are unaffected, but column c has the step property, and so does the resulting matrix. ■

Next, we construct the *bitonic-converter* network $\mathcal{D}(p, q)$ of depth two from q -balancers and p -balancers, where $p, q \geq 2$.

Let X be the input sequence with the bitonic property. As illustrated in Figure 12, arrange X as a $p \times q$ matrix in column-major form. If we place a q -balancer across each row only one column is 1-smooth. If we then place a p -balancer across each column, the result has the step property (as a matrix in column-major form). The detailed correctness analysis is similar to the two-merger $\mathcal{T}(p, q_0, q_1)$.

5 Specific Counting Network Constructions

5.1 The Counting Network \mathcal{K}

We construct $\mathcal{K}(p_0, \dots, p_{n-1})$, the counting network of depth $O(n^2)$ from balancers of width at most $\max(p_i p_j)$, for $0 \leq i, j < n$, where $p_i \geq 2$ and $n \geq 2$.




Figure 12: Construction of bitonic-converter network

The construction is the same with the construction of \mathcal{C} described in Section 4, where in place of each instance of $\mathcal{C}(p_i, p_j)$ we use a balancer of width $p_i p_j$ with $d = 1$. For the staircase-merger \mathcal{S} we use the optimization described in Section 4.3.1 with $\text{depth}(\mathcal{S}) = 2d + 1 = 3$, and we get for the depth of \mathcal{K} :

Proposition 6 $\text{depth}(\mathcal{K}(p_0, \dots, p_{n-1})) = 1.5n^2 - 3.5n + 2$.

Proof:

$$\begin{aligned}
 & \text{depth}(\mathcal{K}(p_0, \dots, p_{n-1})) \\
 &= \text{depth}(\mathcal{C}(p_0, \dots, p_{n-1})) \\
 &= (n-1)d + (n^2/2 - 3n/2 + 1) \cdot \text{depth}(\mathcal{S}) \\
 &\quad \text{(by Proposition 1)} \\
 &= (n-1)1 + (n^2/2 - 3n/2 + 1)3 \\
 &= 1.5n^2 - 3.5n + 2.
 \end{aligned}$$

■

5.2 The Counting Network \mathcal{L}

We construct $\mathcal{L}(p_0, \dots, p_{n-1})$, the counting network of depth $O(n^2)$ from balancers of width at most $\max(p_i)$, for $0 \leq i < n$, where $p_i \geq 2$ and $n \geq 2$.

The construction is the same with the construction of \mathcal{C} described in Section 4, where in place of each instance of network $\mathcal{C}(p_i, p_j)$ we use the counting network $\mathcal{R}(p_i, p_j)$, described below in Section 5.3 with $d = \text{depth}(\mathcal{R}(p_i, p_j)) \leq 16$. For the staircase-merger \mathcal{S} we use the optimization described in Section 4.3.1 with $\text{depth}(\mathcal{S}) = d + 3 \leq 19$, and we get for the depth of \mathcal{L} :

Theorem 7 $\text{depth}(\mathcal{L}(p_0, \dots, p_{n-1})) = 9.5n^2 - 12.5n + 3$.

Proof:

$$\begin{aligned}
 & \text{depth}(\mathcal{L}(p_0, \dots, p_{n-1})) \\
 &= \text{depth}(\mathcal{C}(p_0, \dots, p_{n-1})) \\
 &= (n-1)d + (n^2/2 - 3n/2 + 1) \cdot \text{depth}(\mathcal{S}) \\
 &\quad \text{(by Proposition 1)} \\
 &\leq (n-1)16 + (n^2/2 - 3n/2 + 1)19 \\
 &= 9.5n^2 - 12.5n + 3.
 \end{aligned}$$

■

5.3 The Counting Network $\mathcal{R}(p, q)$

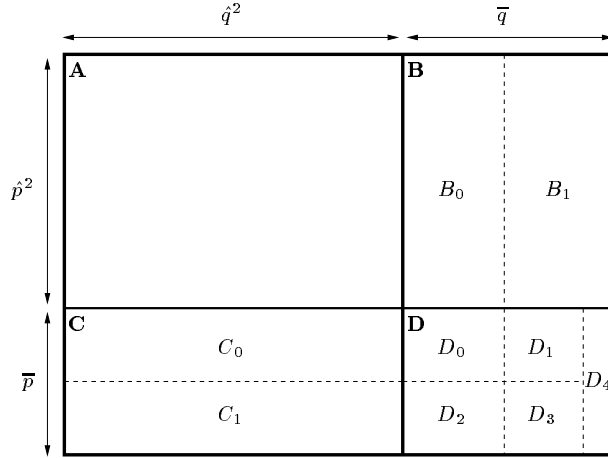


Figure 13: Construction of width- pq counting network

Let $w = pq$, where $p, q \geq 2$. We now construct a constant-depth counting network $\mathcal{R}(p, q)$ of width w from balancers of width at most $\max(p, q)$. We rely on two subsidiary networks: the two-merger network \mathcal{T} described in Section 4.4, and the counting network \mathcal{K} described in Section 5.1.

Let $\hat{p} = \lfloor \sqrt{\bar{p}} \rfloor$, and $\bar{p} = p - \hat{p}^2$. Similarly, we define \hat{q} and \bar{q} . The following inequalities hold (see the appendix):

$$\max(\hat{p}, \hat{q})^2 \leq \max(p, q) \quad (1)$$

$$\max(\hat{p}, \hat{q}) \lceil \max(\bar{p}, \bar{q})/2 \rceil \leq \max(p, q) \quad (2)$$

$$\lceil \max(\bar{p}, \bar{q})/2 \rceil \lceil \max(\bar{p}, \bar{q})/2 \rceil \leq \max(p, q) \quad (3)$$

Let X be the input sequence to $\mathcal{R}(p, q)$. Because $|X| = pq$, we can arrange X as a $p \times q$ matrix in arbitrary order. Divide X into four quadrants: A encompasses the first \hat{p}^2 rows and \hat{q}^2 columns, B the first \hat{p}^2 rows and remaining \bar{q} columns, C the remaining \bar{p} rows and first \hat{q}^2 columns, and D the remaining \bar{p} rows and \bar{q} columns. These divisions are shown as thick lines in Figure 13.

A is a sequence of length $\hat{p}\hat{p}\hat{q}\hat{q}$. We can use the constant-depth counting network $\mathcal{K}(\hat{p}, \hat{p}, \hat{q}, \hat{q})$, constructed from balancers of width at most $\max(\hat{p}^2, \hat{q}^2, \hat{p}\hat{q}) \leq \max(p, q)$ (Equation 1), to transform A into a sequence A' satisfying the step property.

Let $\bar{q}_0 = \lfloor \bar{q}/2 \rfloor$ and $\bar{q}_1 = \lceil \bar{q}/2 \rceil$. Partition B into disjoint submatrices B_0 and B_1 of respective dimensions $\hat{p}^2 \times \bar{q}_0$ and $\hat{p}^2 \times \bar{q}_1$. (These divisions are shown as dotted lines in Figure 13.) We use the constant-depth counting network $\mathcal{K}(\bar{q}_0, \hat{p}, \hat{p})$ and $\mathcal{K}(\bar{q}_1, \hat{p}, \hat{p})$, constructed from balancers of width at most $\max(\hat{p}^2, \hat{p}\bar{q}_0)$ and $\max(\hat{p}^2, \hat{p}\bar{q}_1)$, that respectively transform B_0 and B_1 into sequences B'_0 and B'_1 satisfying the step property. By Equations 1 and 2, each of these networks is constructed from balancers of width at most $\max(p, q)$. Finally, the constant-depth two-merger network $\mathcal{T}(\hat{p}^2, \bar{q}_0, \bar{q}_1)$ merges B'_0 and B'_1 to a single sequence B' satisfying the step property. This two-merger is constructed from balancers of width \hat{p}^2 and \bar{q} , each less than or equal to $\max(p, q)$ (Equation 1). In exactly the same way, C can be transformed to C' satisfying the step property.

Partition D into disjoint submatrices D_0, D_1, D_2, D_3 , and D_4 , with respective dimensions $\bar{p}_0 \times \bar{q}_0, \bar{p}_0 \times \bar{q}_0, \bar{p}_1 \times \bar{q}_0, \bar{p}_1 \times \bar{q}_0$, and $\bar{p} \times 1$. (See Figure 13.) Each of these regions can be given the step property by a single balancer of size less than or equal to $\max(p, q)$ (Equation 3). The resulting sequences can then be merged in constant depth using several copies of the two-merger network \mathcal{T} to a sequence D' satisfying the step property. These two-mergers are constructed with balancers of width less than $\max(p, q)$.

We have shown that A, B, C , and D can be transformed to A', B', C' , and D' satisfying the step property by counting networks constructed from balancers of width less than $\max(p, q)$. In the same way, two-merger networks can merge A' and B' , and (in parallel) C' and D' . Finally, a two-merger network can merge their results. These two-mergers are constructed with balancers of width less than or equal to $\max(p, q)$.

The depth of the construction of \mathcal{R} is dominated by the depth of the counting network of A plus the final two layers of two-mergers that each has depth 2. We have:

$$\text{depth}(\mathcal{R}(p, q)) = \text{depth}(\mathcal{K}(\hat{p}, \hat{p}, \hat{q}, \hat{q})) + 2\text{depth}(\mathcal{T})$$

$$\begin{aligned}
&= 1.5 \cdot 4^2 - 3.5 \cdot 4 + 2 + 2 \cdot 2 \\
&\quad \text{(by Proposition 6)} \\
&= 16.
\end{aligned}$$

If some of the variables $\hat{p}, \bar{p}, \bar{p}_0, \dots$ take the extreme values 0 or 1, then for each of the affected A, B, B_0, \dots we either do not use any network or we use a single balancer, and then we use the two-mergers accordingly. In these cases we may obtain a network construction with smaller depth. Therefore, taking into consideration all the cases, we have $\text{depth}(\mathcal{R}(p, q)) \leq 16$.

6 Discussion

We have a new construction for a family of sorting or counting networks of width $w = p_0 \cdots p_{n-1}$, and depth at most $9.5n^2 - 12.5n + 3$, from comparators or balancers of width at most $\max(p_i)$. This is the first arbitrary-width construction without enormous constant factors.

The overall network structure (Figure 7) is similar but not identical to that of the bitonic network [3, 4]. The bitonic network, however, has smaller depth by a constant factor, suggesting that further improvement in our constant terms may be possible. It remains an open problem whether the asymptotic $O(n^2)$ depth can be improved without introducing very large constants.

An interesting open question concerns the timing constraints necessary for counting networks built in this way to be linearizable (c.f., [13, 14, 15]).

References

- [1] E. Aharonson and H. Attiya, “Counting Networks with Arbitrary Fan-Out,” *Distributed Computing*, Vol. 8, pp. 163–169, 1995.
- [2] M. Ajtai, J. Komlós and E. Szemerédi, “Sorting in $c \log n$ Parallel Steps,” *Combinatorica*, Vol. 3, pp. 1–19, 1983.
- [3] J. Aspnes, M. Herlihy and N. Shavit, “Counting Networks,” *Journal of the ACM*, Vol. 41, No. 5, pp. 1020–1048, September 1994.
- [4] K.E. Batcher, “Sorting networks and their applications,” *Proceedings of the AFIPS Spring Joint Computer Conference*, Vol. 32, pp. 338–334, 1968.
- [5] C. Busch, N. Hardavellas and M. Mavronicolas, “Contention in Counting Networks,” *Proceedings of the 13th Annual ACM Symposium on Principles of Distributed Computing*, pp. 404, August 1994.
- [6] V. Chvátal, “Lecture Notes on the New AKS Sorting Network,” *Technical Report 92–29*, DIMACS Center for Discrete Mathematics and Theoretical Computer Science, June 1992.
- [7] T.H. Cormen, C.E. Leiserson, and R. L. Rivest, “Introduction to Algorithms,” MIT Press, Cambridge MA, 1990.
- [8] M. Dowd, Y. Perl, L. Rudolph, and M. Saks, “The Periodic Balanced Sorting Network,” *Journal of the ACM*, Vol. 36, No. 4, pp. 738–757, October 1989.

- [9] E. W. Felten, A. LaMarca and R. Ladner, “Building Counting Networks from Larger Balancers,” *Technical Report 93-04-09*, Department of Computer Science and Engineering, University of Washington, April 1993.
- [10] M. Klugerman, “Small-Depth Counting Networks and Related Topics,” *Ph.D. Thesis*, Department of Mathematics, Massachusetts Institute of Technology, September 1994.
- [11] D.E. Knuth, “The Art of Computer Programming Vol. 3”, Addison-Wesley, 1973.
- [12] D.-L. Lee and K.E. Batcher, “A Multiway Merge Sorting Network,” *IEEE Transactions on Parallel and Distributed Systems*, Vol. 6, No. 2, pp. 211–215, February 1995.
- [13] N. Lynch, N. Shavit, A. Shvartsman, and D. Touitou, “Counting Networks are Practically Linearizable,” *Proceedings of the 15th Annual ACM Symposium on Principles of Distributed Computing*, pp. 280–289, May 1996.
- [14] M. Mavronicolas, M. Merritt, and G. Taubenfeld, “Sequentially Consistent versus Linearizable Counting Networks,” *Proceedings of the 18th Annual ACM Symposium on Principles of Distributed Computing*, to appear, Atlanta, Georgia, May 1999.
- [15] M. Mavronicolas, M. Papatrantafileou, and Ph. Tsigas, “The Impact of Timing on Linearizability in Counting Networks,” *Proceedings of the 11th International Parallel Processing Symposium*, pp. 684–688, April 1997.
- [16] B. Parker and I. Parberry, “Constructing Counting Networks from k-Sorters,” *Information Processing Letters*, Vol. 33, pp. 157–162, 1989/90.
- [17] S.S. Tseng and R.C.T. Lee, “A Parallel Sorting Scheme whose Basic Operation Sorts n Elements,” *International Journal of Computer and Information Sciences*, Vol. 14, No. 6, pp. 455–467, 1985.

Appendix

We prove Equations 1, 2, and 3 of Section 5.3.

Take any two integers $p, q \geq 2$. Let $\hat{p} = \lfloor \sqrt{p} \rfloor$, $\bar{p} = p - \hat{p}^2$, and similarly define \hat{q} and \bar{q} for q . Let also $m = \max(p, q)$, $r = \max(\hat{p}, \hat{q})$, and $s = \max(\bar{p}, \bar{q})$.

Obviously, if $m = p$ then $r = \hat{p}$, and if $m = q$ then $r = \hat{q}$. Since $\hat{p}^2 \leq p$ and $\hat{q}^2 \leq q$, we have $r^2 \leq m$ and thus Equation 1 holds.

We continue by showing the inequality:

$$s < 2\sqrt{p} - 1 \tag{4}$$

Proof: Since $\hat{p} = \lfloor \sqrt{p} \rfloor > \sqrt{p} - 1$, we have

$$\begin{aligned} \bar{p} &= p - \hat{p}^2 \\ &< p - (\sqrt{p} - 1)^2 \\ &= 2\sqrt{p} - 1, \end{aligned}$$

and thus $\bar{p} < 2p - 1$. Similarly, $\bar{q} < 2q - 1$.

Since $\bar{p} < 2\sqrt{p} - 1$ and $p \leq m$ we have $\bar{p} < 2\sqrt{m} - 1$. Similarly, $\bar{q} < 2\sqrt{m} - 1$. Since $s = \max(\bar{p}, \bar{q})$, we have $s < 2\sqrt{m} - 1$ as needed. ■

Next, we show the correctness of Equation 2 which can be written as:

$$r\lceil s/2 \rceil \leq m \quad (5)$$

Proof: By Equation 4, we have $s/2 < \sqrt{m} - 1/2$. Subsequently, $\lceil s/2 \rceil \leq \lceil \sqrt{m} - 1/2 \rceil$. and thus $r\lceil s/2 \rceil \leq r\lceil \sqrt{m} - 1/2 \rceil$. Therefore, we only need to show that $r\lceil \sqrt{m} - 1/2 \rceil \leq m$.

First, we examine the case $\sqrt{m} - r < 1/2$. We have that $\lceil \sqrt{m} - 1/2 \rceil = r$, and thus $r\lceil \sqrt{m} - 1/2 \rceil = r^2$. Since $r \leq \sqrt{m}$, we have $r^2 \leq m$. Therefore, $r\lceil \sqrt{m} - 1/2 \rceil \leq m$, as needed.

Next, we examine the case $\sqrt{m} - r \geq 1/2$. We have $\lceil \sqrt{m} - 1/2 \rceil = r + 1$, and thus $r\lceil \sqrt{m} - 1/2 \rceil = r^2 + r$. Since $r \leq \sqrt{m} - 1/2$, we have

$$\begin{aligned} r^2 + r &\leq (\sqrt{m} - 1/2)^2 + \sqrt{m} - 1/2 \\ &= m - 1/4 \\ &\leq m. \end{aligned}$$

Therefore, $r\lceil \sqrt{m} - 1/2 \rceil \leq m$, as needed. ■

Finally, we show the correctness of Equation 3 which can be written as:

$$\lfloor s/2 \rfloor \lfloor s/2 \rfloor \leq m \quad (6)$$

Proof: By Equation 5 we only need to show that $\lfloor s/2 \rfloor \leq r$. By Equation 4, we have that $s/2 < \sqrt{m} - 1/2$. Therefore, $\lfloor s/2 \rfloor \leq \lfloor \sqrt{m} - 1/2 \rfloor$. Since $\lfloor \sqrt{m} - 1/2 \rfloor \leq r$, we have $\lfloor s/2 \rfloor \leq r$, as needed. ■