

REAL-TIME SORTING OF BINARY NUMBERS ON ONE-DIMENSIONAL CA

THOMAS WORSCH ¹ AND HIDENOSUKE NISHIO ²

¹ KIT (Karlsruhe Institute of Technology)
E-mail address: `worsch@kit.edu`

² Kyoto University
E-mail address: `yra05762@nifty.com`

ABSTRACT. A new fast (real time) sorter of binary numbers by one-dimensional cellular automata is proposed. It sorts a list of n numbers represented by k -bits each in exactly nk steps. This is only one step more than a lower bound.

1. Introduction

Sorting is one of the most fundamental subjects of computer science and many sorting algorithms including sorting arrays and networks can for example be found in volume 3 of Knuth's TAOCP [2]. However, for cellular automata there are only a few papers on this important topic. It should be pointed out that the algorithm described by one of the authors in an earlier paper [3] has running $3nk$ (despite the title starting with the words "real time").

We are not aware of any speedup techniques which would allow to turn this CA or any other solving the problem into one running in real-time, i. e. exactly the number of steps which is the length of the input. In the present paper we propose a sorting algorithm of binary numbers and its implementation on one-dimensional CA with nearest neighbors, which sorts n numbers of k bits each in exactly nk steps.

Sequential comparison based sorting algorithms need time $\Omega(n \log n)$ where n is the number of elements to be sorted and it is assumed that each comparison can be done in constant time independent of the size of the elements. The latter assumption is also usually made for parallel sorting algorithms. On linear arrays odd-even transposition sort needs exactly n steps. But there the additional assumption is made that from the beginning each processor knows the parity of its own address (assuming those are e. g. 1 to n). Of course on parallel models from the second machine class [1] like PRAM there are algorithms running in poly-logarithmic (or even logarithmic) time. But for these models one has to assume non-local communication in constant time when embedded into Euclidean space [4].

The rest of the paper is organized as follows. In Section 2 we precisely state the problem and the results obtained in this paper. In Section 3 we give a short proof for the lower bound of the sorting problem. The main aspects of our algorithm are presented in Section 4. The first version does not achieve a running time which

matches the lower bound. For that two modifications are needed which are described in Section 5.

2. Statement of problem and results

We are considering one-dimensional CA with von Neumann neighborhood of radius 1 and assume that the reader is familiar with these concepts. Since we will not define our CA on such a low level there is no need to introduce any related formalism.

We also assume that the reader is familiar with the firing squad synchronization problem [5]. If a block of k cells needs to be synchronized and there are generals at both ends, then synchronization can be achieved in exactly k steps.

The inputs for our CA are provided as finite words with all surrounding cells in a quiescent state. Those cells will never be used during computations.

The inputs which have to be processed by our CA are n numbers of equal length k , with the most and least significant bits marked as such.

Problem 2.1. The input alphabet is $A = \{0, 1, \langle 0, \langle 1, 0 |, 1 | \rangle\}$.

Each input w that has to be processed properly is of the form

$$w = w_1 \cdots w_n = \langle x_1 y_1 z_1 | \cdots \langle x_n y_n z_n |$$

for some $n \geq 1$ and $k \geq 2$ where all $x_i, z_i \in \{0, 1\}$ and all $y_i \in \{0, 1\}^{k-2}$. Each w_i is the binary representation of a non-negative integer (also denoted w_i) with most significant bit x_i and least significant bit z_i .

For every such input after a finite number of steps a stable configuration of the form $w_{\sigma(1)} \cdots w_{\sigma(n)}$ has to be reached where σ is a permutation of the numbers $1, \dots, n$ such that $w_{\sigma(i)} \leq w_{\sigma(i+1)}$ holds for all $1 \leq i < n$.

We note that it would have been sufficient to mark either the most or the least significant bits, because the other end of each number can then always be identified by looking at neighbor cells.

The above problem statement also excludes the case of 1-bit inputs. For those the “traffic rule” 184 can be easily extended to do sorting, taking into account quiescent neighbors. The resulting CA works as follows: A cell in state 1 (0) becomes 0 (1) if its right (left) neighbor is 0, otherwise it keeps its current state.

In the following we always call a sequence of cells which initially stores one input number w_i the *block* i (or simply a block).

It is clear that it can be necessary to move a number from block 1 to block n . This immediately gives a lower bound of $(n - 1)k$ steps for the sorting time. We shall see, that one can do slightly better:

Theorem 2.2. *Every CA solving Problem 2.1 needs at least time $nk - 1$.*

Until now the fastest sorting algorithms known needed time cnk for some constant $c > 1$ with no obvious possibility to speed up the computation to run in nk steps. The main contribution of the present paper therefore is the following:

Theorem 2.3. *There is a CA (which does not depend on n or k) with von Neumann neighborhood of radius 1 solving the Sorting Problem 2.1 in exactly nk steps.*

3. Lower bound on sorting time

First consider the input $w = w_1 w_2 \cdots w_n$ where $w_1 = \langle 01^{k-2}1 |$ and $w_2 = \cdots = w_n = \langle 10^{k-2}0 |$. Clearly this input sequence is already sorted and the rightmost bit of the output is a 0.

If on the other hand we flip the leftmost bit of the first block only and consider the input

$$w'_1 \cdot w_2 \cdots w_n = \langle 11^{k-2}1 | \cdot \langle 10^{k-2}0 | \cdots \langle 10^{k-2}0 |$$

then the correct sorted output is

$$w_2 \cdots w_n \cdot w'_1 = \langle 10^{k-2}0 | \cdots \langle 10^{k-2}0 | \cdot \langle 11^{k-2}1 |$$

That is, by changing only the leftmost bit of the input the rightmost bit of the output must change. Hence no CA correctly solving the sorting problem can be faster than the distance between leftmost and rightmost bit which is $nk - 1$.

This proves Theorem 2.2.

4. The base sorting algorithm

The goal of this section is to prove a weakened version of Theorem 2.3.

Lemma 4.1. *There is a CA (which does not depend on n or k) with the von Neumann neighborhood of radius 1 solving the sorting Problem 2.1 in exactly $k + nk$ steps.*

Before going into details and explaining some aspects of the CA on the cell level, we describe the main idea on the level of numbers. In particular we will employ a well-known simple algorithm for parallel sorting.

4.1. Odd-even transposition sort

Throughout this section, one can assume that N is an even number; this is the case needed for the CA below.

Assume that N numbers a_1, a_2, \dots, a_N are given, arranged in an array of processors. In addition each processor (or each number) has a direction (indicated by arrows below). In each step each pair of adjacent processors whose arrows point to each other exchange their numbers. Both processors compare the two numbers; the left one keeps the smaller number, the right one the larger number, and both change their direction to the other neighbor. Thus one step can for example look like this:

$$\begin{array}{l}
 t \quad \boxed{\overleftarrow{a_1}} \quad \boxed{\overrightarrow{a_2}} \quad \boxed{\overleftarrow{a_3}} \quad \boxed{\overrightarrow{a_4}} \quad \boxed{\overleftarrow{a_5}} \quad \boxed{\overrightarrow{a_6}} \quad \dots \dots \quad \boxed{\overleftarrow{a_{N-1}}} \quad \boxed{\overrightarrow{a_N}} \\
 t+1 \quad \boxed{\overrightarrow{b_1}} \quad \boxed{\overleftarrow{b_2}} \quad \boxed{\overrightarrow{b_3}} \quad \boxed{\overleftarrow{b_4}} \quad \boxed{\overrightarrow{b_5}} \quad \boxed{\overleftarrow{b_6}} \quad \dots \dots \quad \boxed{\overrightarrow{b_{N-1}}} \quad \boxed{\overleftarrow{b_N}}
 \end{array}$$

$$\text{where } b_i = \begin{cases} \min(a_i, a_{i+1}) & \text{if } a_i \text{ points to the right} \\ \max(a_{i-1}, a_i) & \text{if } a_i \text{ points to the left} \end{cases}$$

A missing neighboring number to the left is treated as if it were $-\infty$ and a missing neighboring number to the right is treated as if it were ∞ .

It is known that odd-even transposition sort always produces the correct result after exactly N steps (see e.g.[2]).

4.2. Outline of the base algorithm

The algorithm which will be the basis for the improved construction in Section 5 will simply work as follows. Given an input of the form

$$\boxed{w_1} \quad \boxed{w_2} \quad \boxed{w_3} \quad \boxed{w_4} \quad \dots \quad \boxed{w_{n-1}} \quad \boxed{w_n}$$

first each number is copied and the two copies get opposite directions assigned. This will take k steps in the CA. Instead of storing two copies side by side they are stored in parallel:

$$\begin{array}{|c|c|c|c|} \hline \overleftarrow{w_1} & \overleftarrow{w_2} & \overleftarrow{w_3} & \overleftarrow{w_4} \\ \hline \overrightarrow{w_1} & \overrightarrow{w_2} & \overrightarrow{w_3} & \overrightarrow{w_4} \\ \hline \end{array} \quad \dots \quad \begin{array}{|c|c|} \hline \overleftarrow{w_{n-1}} & \overleftarrow{w_n} \\ \hline \overrightarrow{w_{n-1}} & \overrightarrow{w_n} \\ \hline \end{array}$$

These are now treated as $N = 2n$ numbers that are sorted using odd-even transposition sort. In the end one would get

$$\begin{array}{|c|c|c|c|} \hline \overleftarrow{w_{\sigma(1)}} & \overleftarrow{w_{\sigma(2)}} & \overleftarrow{w_{\sigma(3)}} & \overleftarrow{w_{\sigma(4)}} \\ \hline \overrightarrow{w_{\sigma(1)}} & \overrightarrow{w_{\sigma(2)}} & \overrightarrow{w_{\sigma(3)}} & \overrightarrow{w_{\sigma(4)}} \\ \hline \end{array} \quad \dots \quad \begin{array}{|c|c|} \hline \overleftarrow{w_{\sigma(n-1)}} & \overleftarrow{w_{\sigma(n)}} \\ \hline \overrightarrow{w_{\sigma(n-1)}} & \overrightarrow{w_{\sigma(n)}} \\ \hline \end{array}$$

where σ again denotes the “sorting permutation” as in Problem 2.1.

During the last sorting step the lower parts and the arrows are deleted, and the required output is obtained.

It will become clear in the next subsection why it is actually useful to first copy each number and then seemingly spend twice as much time for the $N = 2n$ sorting steps. It will be shown that each such step can be implemented in the CA in $k/2$ steps. Hence the total running time of the CA for this base version of the algorithm will be $k + nk$ steps.

4.3. Outline of the CA for the base algorithm

In this section we will describe how the base sorting algorithm can be implemented on a CA. It will need $n + 1$ phases each of which needs exactly k steps. First comes a setup phase followed by phases $1, \dots, n$. In order to avoid more complicated descriptions, throughout the rest of the paper we assume that k is even.

If k is odd the middle cell of a block plays the role of the two middle cells one would have for the (even) case $k + 1$. In that case synchronization of a block using generals at both ends needs at least time $2\frac{k+1}{2} - 2 = k - 1$ and hence is possible in time k (as is the case for even k).

Algorithm 4.2 (Setup phase).

- (1) During the setup phase the following tasks are carried out in each block:
 - By sending a signal from the left and the right end of the block the two middle cells are found. In each resulting *sub-block* the leftmost and the rightmost cell are marked as such. They are called L and R respectively.
 - Using an additional register the mirror image of the input number is computed. Below we call the register holding the original value *left* and the registers with the mirrored value *right*.
The numbers in the *left* registers will play the role of the $\overleftarrow{w_i}$ and the numbers in the *right* registers the role of the $\overrightarrow{w_i}$ used in the previous subsection.
 - Using synchronization, the preliminary phase is stopped after k steps.

- (2) The leftmost and the rightmost cell of the whole input are set up as generals. Starting with the first step of phase 1 an algorithm is started to synchronize all nk cells after nk steps.

A concrete example is shown in Figure 1 for two 6-bit numbers. The borders between sub-blocks are shown as double vertical lines. It can be noted that we also use markers for the most and least significant bits of the mirrored numbers.

	L	R	L	R	L	R	L	R				
<i>left</i>	$\langle 1$	0	0	0	0	0	$\langle 0$	1	1	1	1	1
<i>right</i>	$ 0$	0	0	0	0	1	$ 1$	1	1	1	1	0

Figure 1: An example configuration for two 6-bit numbers after the setup phase.

In addition to the registers *left* and *right* the L - and R -cells of each sub-block will make use of a register *comp* which will hold a (preliminary) comparison result (see also Figure 2). Each *comp* register can hold one of the values $=$, $<$ or $>$. Their use is described in Algorithm 4.4 below.

The core idea is the following:

- C1. Bits are shifted in the *left* and *right* registers in the corresponding directions.
 C2. Whenever the most significant bits of two numbers arrive in a pair of adjacent $R||L$ -cells, the numbers are compared sequentially bit by bit. The smaller number will be directed to move to the left and the larger one will be directed to move to the right.

Part C1 basically means that numbers are unconditionally shifted everywhere except at $R||L$ pairs. Since those are located at distance $k/2$ and numbers have length k this might look suspicious at first sight, because in general a number simultaneously gets compared at two such pairs. It will become clear later why this does not pose any problems. Ignoring it for the moment, C1 is easy to implement:

Algorithm 4.3 (Implementation of C1).

- (1) The cells which are *not* an R - or L -cell have a very simple behavior.
 - The *left* register gets its content from the *left* register of the right neighbor.
 - The *right* register gets its content from the *right* register of the left neighbor.
- (2) Analogously the *left* register of an L -cell gets its content from the *left* register of the right neighbor and the *right* register of an R -cell gets its content from the *right* register of the left neighbor.
- (3) The same holds for the *left* register of a R -cell and the *right* register of a L -cell if the *comp* register of that cell has value $=$.

First of all, this part is needed during the first sub-phase of phase 1 when the $R||L$ pairs in the middle of a block still have no meaning. As will be seen this requirement is also consistent with the rules for later sub-phases.

Each of the phases 1, \dots , n is subdivided into two sub-phases of $k/2$ steps each.

We will now describe how the comparison of numbers is done. For this we use $R.left$ to denote the *left* register of the R -cell and similarly for the other cases. It will be seen that all information needed to update $R.comp$ are also available in the neighboring L -cell, so that the invariant $R.comp = L.comp$ can be maintained. Hence it suffices to describe the case of $R.comp$:

Algorithm 4.4 (Implementation of C2). If the two bits in $R.right$ and $L.left$ are most significant ones, then the new value of $R.comp$ is determined as follows:

$$R.comp \text{ becomes } \begin{cases} < & \text{if } R.right < L.left \\ = & \text{if } R.right = L.left \\ > & \text{if } R.right > L.left \end{cases} \quad (4.1)$$

If the two bits to be compared are not most significant ones, the new value of $R.comp$ is determined as follows:

- If $R.comp$ already has value $<$ or $>$ it is not changed.
- If $R.comp$ has old value $=$ its new value is determined according to rule 4.1 above.

It remains to define how the new value for $R.left$ is computed. That is most easily described as depending on the just defined *new* value of $R.comp$:

$$R.left \text{ becomes } \begin{cases} R.right & \text{if } R.comp \text{ is now } < \\ R.right & \text{if } R.comp \text{ is now } = \\ L.left & \text{if } R.comp \text{ is now } > \end{cases}$$

Dually the new value for $L.right$ depends on the *new* value of $L.comp$ (remember that always $R.comp = L.comp$):

$$L.right \text{ becomes } \begin{cases} L.left & \text{if } L.comp \text{ is now } < \\ L.left & \text{if } L.comp \text{ is now } = \\ R.right & \text{if } L.comp \text{ is now } > \end{cases}$$

Figure 2 shows the relevant parts of computations for the comparison of two numbers. In the left part of the figure initially the larger number is on the left, in the right part the smaller number is on the left. After the comparison in both cases the smaller number is on the left. We remind the reader that at the left resp. right end of the complete input a missing number is treated as $-\infty$ resp. ∞ . Hence a number arriving at a border is simply reflected.

The following picture may be helpful: When the smaller number comes from the left and the larger from the right, then from the first (most significant) bit both numbers are reflected at the border between the R - and the L -cell. When the larger number comes from the left and the smaller from the right, then from the first (most significant) bit both numbers pass through the border. This is a correct picture even if the numbers have identical higher order bits, and hence in the beginning no cell knows which is the present case. Readers are encouraged to check Figure 2 again.

At least from a formal point of view it is now straightforward to put the pieces together. An even better intuition of how the algorithm works may arise in the subsequent Subsection 4.4 when the correctness of the algorithm will be shown.

	<i>R</i>				<i>L</i>					<i>R</i>				<i>L</i>											
<i>left</i>					$\langle 0$	0	1	1						$\langle 0$	1	0	1								
<i>right</i>	1	0	1	0																					
<i>comp</i>																									
<i>left</i>					$\langle 0$	0	1	1						$\langle 0$	1	0	1								
<i>right</i>		1	0	1	0									1	1	0	0								
<i>comp</i>					=	=											=	=							
<i>left</i>					$\langle 0$	0			1	1								$\langle 0$	0			0	1		
<i>right</i>			1	0	1	0			1	0						1	0	1	0						
<i>comp</i>					>	>			>	>								<	<			<	<		
<i>left</i>					$\langle 0$	0	1		1									$\langle 0$	0	1		1			
<i>right</i>									1	0	1	0						1	0	1	0				
<i>comp</i>					>	>			>	>								<	<			<	<		
<i>left</i>					$\langle 0$	0	1	1										$\langle 0$	0	1	1				
<i>right</i>									1	0	1	0						1	0	1	0				
<i>comp</i>					>	>			>	>								<	<			<	<		

Figure 2: Two comparisons of 4 bits. On the left hand side the number coming from the left (in the *right* registers) is larger, on the right hand side the number coming from the right (in the *left* registers). If the complete numbers were longer, the comparisons would continue analogously.

Algorithm 4.5.

- First the setup phase is done as described in Algorithm 4.2. This phase takes k steps. It is stopped at the correct time using a synchronization algorithm in each block separately. Both ends of each block have to act as generals in order to achieve the required synchronization time.
- Once all cells are synchronized they will work as described in Algorithms 4.3 and 4.4: All numbers are shifted to the left or right, and whenever two most significant bits meet, the sequential comparison of the two numbers is started. The smaller number is sent to the left and the larger to the right.
- This is repeated until the synchronization started immediately after the setup phase fires all nk cells after nk steps.

It will be shown that at that point in time the *left* registers contain the sorted numbers.

4.4. Correctness of the base sorting algorithm

The correctness of algorithm 4.5 is essentially due to the correctness of odd-even transposition sort. This is basically a proof by induction. The main parts are stated in the following Lemma.

Lemma 4.6. *Let $\mathcal{X} = |xX\rangle$ and $\mathcal{Y} = \langle Yy|$ be two numbers with the $k/2$ higher order bits denoted by capital letters and the $k/2$ lower order bits denoted by small letters. Similarly let $\mathcal{A} = \langle Aa|$ be the minimum of \mathcal{X} and \mathcal{Y} and $\mathcal{B} = |bB\rangle$ be the maximum. In other words one basic compare-and-exchange step of odd-even transposition sort transforms the pair \mathcal{X}, \mathcal{Y} into the pair \mathcal{A}, \mathcal{B} .*

If the most significant bits of $X\rangle$ and $\langle Y$ meet in an $R||L$ -pair (with the other higher order bits following) and if after $k/2$ steps the lower order bits $|x$ and $y|$ arrive in the correct order, then during the first $k/2$ steps the higher order bits of $\langle A$ and $B\rangle$ will be produced moving to the right directions, followed by the lower order bits $a|$ and $|b$ afterwards.

This is basically a restatement of the construction from Algorithm 4.4.

Since the configuration produced by the setup phase corresponds to the initial configuration for the odd-even transposition sort, and since the preconditions of the if-statement in Lemma 4.6 are met, an induction teaches that in particular the higher order bits of each number after t phases are in the sub-block corresponding to its position in odd-even transposition sort after t sorting steps.

Corollary 4.7. *For all input sequences w_1, \dots, w_n Algorithm 4.5 does sort the numbers as required in Problem 2.1.*

Proof. Since odd-even transposition sort does sort N numbers in N sorting steps, it immediately follows from Lemma 4.6 that at the end of phase n of Algorithm 4.5 the higher order bits of each of the n input numbers and of its n copies are in the correct blocks. That is, there are the same higher order bits of a number w_i in each block twice, once in the *left* registers and once in the *right* registers.

Furthermore it is clear that the most significant bit of the number stored in the *left* (resp. *right*) registers is in the L -cell (resp. R -cell) of the *block* (not sub-block). This implies that the lower order bits are in the same block:

	L	R
<i>left</i>	$\langle k/2$ higher order bits of w_i	$k/2$ lower order bits of $w_i $
<i>right</i>	$ k/2$ lower order bits of w_i	$k/2$ higher order bits of $w_i\rangle$

Therefore the *left* registers of the full block hold the correct value. ■

5. A sorting algorithm matching the lower bound

We will save k steps of the running time of Algorithm 4.5 by starting with some comparisons not after k but already after $k/2$ steps and stopping the odd-even transposition sort $k/2$ steps earlier. A detailed description will be given in Section 5.1. The resulting algorithm still computes the correct output *except* for the rightmost block. This will be fixed in Section 5.2.

5.1. Speeding up the algorithm

We describe the fast algorithm as three changes to Algorithm 4.5.

The first change is simple: Since we want to have the result after nk steps instead of $k + nk$, the synchronization of all nk cells is not started after the setup phase, but in the very first step.

The second change concerns the computation of the mirror of a bit string as required by Algorithm 4.2. It can be implemented by shifting the original to the left, and letting the L -cell of the block act as reflector sending bits back to the right in the *right* cells. This means that after $k/2$ steps the lower order bits of an input number have arrived in the *left* registers of the left sub-block and the higher order bits are its *right* registers. It is useful if the shift to the left is not done using a temporary register but *left*. In Figure 3 the resulting process is shown for two adjacent 6-bit numbers. It can be seen that due to the simultaneous shift to the left, already after $k/2$ steps for the first time most significant bits meet. (This also determines the border of the sub-blocks.) Thus comparisons can be started $k/2$ steps earlier. The rightmost block now needs some special attention. Analogously to the other blocks we assume that symbols representing the “number ∞ ” are shifted to the left from the rightmost cell. This is also depicted in Figure 3.

This completes the second change to Algorithm 4.5.

<i>left</i>	$\langle a_6$	a_5	a_4	a_3	a_2	a_1	$\langle b_6$	b_5	b_4	b_3	b_2	b_1
<i>right</i>												

<i>left</i>	a_5	a_4	a_3	a_2	a_1	$\langle b_6$	b_5	b_4	b_3	b_2	b_1	∞
<i>right</i>	a_6						b_6					

<i>left</i>	a_4	a_3	a_2	a_1	$\langle b_6$	b_5	b_4	b_3	b_2	b_1	∞	∞
<i>right</i>	a_5	a_6					b_5	b_6				

<i>left</i>	a_3	a_2	a_1	$\langle b_6$	b_5	b_4	b_3	b_2	b_1	∞	∞	∞
<i>right</i>	a_4	a_5	a_6				b_4	b_5	b_6			

Figure 3: Computing the mirror of two 6-bit numbers. Already after 3 steps most significant bits meet at the border between sub-blocks.

The third change is the most complicated. For the shifts to the right *two* registers are used, *right* and *right2*. The additional register *right2* is empty almost everywhere. After each sub-phase there are only two adjacent sub-blocks in which *right2* stores a number:

- (1) The initialization takes place in the leftmost block: During the first $k/2$ steps the reflected bits are shifted to the right in *right* and *right2*. And this is done *only* during the first $k/2$ steps, but not afterwards.
- (2) Then it happens for the first time that three most significant bits meet, two coming from the left and one coming from the right.

In such a case the comparisons are done as follows:

- The largest of the three numbers is shifted to the right in *right2*.
- The other two numbers are compared and shifted as described in Algorithm 4.4.

A consequence of these rules is, that after k steps the *right2* registers are used in the sub-blocks of block 1, after $2k$ steps in the sub-blocks of block 2, etc. and after nk steps in the sub-blocks of block n , and nowhere else.

Why do these three changes lead to a result after nk steps where in all blocks except the rightmost one the *left* registers contain the correct numbers? That the rightmost block can still be wrong can be seen in examples.

First of all, reflecting w_1 twice make sure that there are really $2n$ numbers which are sorted, each w_i twice. Therefore in the end each block will again contain twice the same number. This would in general not be the case if w_1 would be reflected only once, and the argument below would fail.

Since we are still using odd-even transposition sort, it is clear that after $k/2 + nk$ steps the correct results are obtained everywhere, but with the most significant bits in the middle of the blocks. Thus the result would look like

	L	R
<i>left</i>	$k/2$ lower order bits of w_i	$\langle k/2$ higher order bits of w_i
<i>right</i>	$k/2$ higher order bits of w_i	$ k/2$ lower order bits of w_i

How can such a left sub-block arise? Without loss of generality assume that the input numbers are pairwise different. Then in the left neighboring full block a smaller number (or $-\infty$) is present. Hence the lower left half must have been reflected and going back $k/2$ steps, the bits must have been in the upper part. Analogously the upper right half must have been in the lower part. Except in the rightmost block (where the *right2* registers are in use) there is no other possibility than that the lower order bits are in the remaining registers:

	L	R
<i>left</i>	$\langle k/2$ higher order bits of w_i	$k/2$ lower order bits of w_i
<i>right</i>	$ k/2$ lower order bits of w_i	$k/2$ higher order bits of w_i

Thus the *left* registers hold the desired result.

In the rightmost block it can happen that lower order bits are not stored in the *left* registers of the right sub-block but in the *right2* registers of the left sub-block.

5.2. Determining the rightmost output block

In order to produce the largest number in the rightmost output block we use a separate algorithm which has to be run in parallel to the one described above. It will have finished after nk steps. Remember that the input is a word

$$w = w_1 \cdots w_n = \langle x_1 y_1 z_1 | \cdots \langle x_n y_n z_n |$$

and the task is to have the maximum of the w_i be stored in the rightmost block in the end. This can be achieved as follows.

Algorithm 5.1.

- (1) During the k steps of the setup phase a signal is sent from the right end of the input until it reaches the most significant bit of w_n , marking all cells as belonging to the last block.

When the last block is synchronized after k steps, all cells have received the information and know that they have an additional task.

- (2) From the very first step *all* cells shift their input to the right using an additional register.

Hence after $1 \cdot k$ steps the number w_{n-1} reaches the last block, after $2 \cdot k$ steps number w_{n-2} reaches the last block, etc. and after $(n-1)k$ steps number w_1 reaches the last block.

- (3) The cells in the rightmost block use two additional registers for storing numbers; call them max and $next$. Register max is initialized in the very first step with w_n , register $next$ is marked as not holding a value. Whenever the rightmost block ends a phase, in register $next$ the number is stored that has arrived from the left in *right* because of the shifting.
- (4) If at the beginning of a phase register $next$ has a valid number the rightmost block computes $max \leftarrow \max(max, next)$. For this a signal is sent from left to right, that is from the most significant bit to the least significant bit, comparing $next$ and max . (While this comparison takes place the next number is already arriving in *right*.)

As long as the same bit value is found in both registers nothing is changed and the signal moves one cell to the right.

As soon as at some position for the first time different bit values are found, the following happens:

- If $next$ has a 1 bit, but max has a 0 bit, max is smaller than $next$ and this and all remaining bits are copied from $next$ to max .
- If $next$ has a 0 bit, but max has a 1 bit, max is larger than $next$ and the signal is simply killed leaving max unchanged.

It is straightforward to verify by induction that for $1 \leq i < n$ after phase i one has

$$max = \max\{w_{n-j} \mid 0 \leq j < i\}$$

and hence in the end $max = \max\{w_i \mid 1 \leq i \leq n\}$ as required. Since w_1 is copied to $next$ after $(n-1)k$ steps the final correct value is stored in max k steps later, i.e. after nk steps as required.

Taking together the changes to the base Algorithm 4.5 described in Section 5.1 and the additional algorithm just described one gets a proof of Theorem 2.3.

6. Conclusion

We have shown the sorting of n numbers with k bits can be achieved in (almost) real-time. Thus the situation is very similar to the firing squad synchronization problem: There is an algorithm which has — in our case except for one step — a running time matching a lower bound.

Clearly, the number of states per cell required by our algorithm is finite but large, at least when compared to algorithms e.g. for the synchronization problem. We do not know how much the set of states can be reduced.

The authors gratefully acknowledge a number of suggestions by the referees for improving the presentation of the algorithms.

References

- [1] Peter van Emde Boas. Machine Models and Simulations. *Handbook of Theoretical Computer Science, Volume A*, 1–66, Elsevier, 1990.
- [2] Donald Knuth. *The Art of Computer Programming; Volume 3, 2nd ed.*, Addison-Wesley, 1998.

- [3] Hidenosuke Nishio. Real time sorting of binary numbers by 1-dimensional cellular automaton. In *Proceedings of the International Symposium on Uniformly Structured Automata and Logic, Tokyo, Japan, August 21-23, 1975, IEEE Catalog Number 75 CH1052-OC*, pages 153–162, 1975.
- [4] Amir R. Schorr. Physical parallel devices are not much faster than sequential ones. *Information Processing Letters*, 17(2):103–106, 1983.
- [5] Hiroshi Umeo, Masaya Hisaoka, and Takashi Sogabe. A Survey on Optimum-Time Firing Squad Synchronization Algorithms for One-Dimensional Cellular Automata. *Int. Journal of Unconventional Computing*, 1(4):403–426, 2005.