A Tight Linear Bound on the Neighborhood of Inverse Cellular Automata

Eugen Czeizler[1] and Jarkko Kari[2,⋆]

[1] Department of Mathematics, FIN-20014 University of Turku, Finland, and Turku Centre for Computer Science, FIN-20520 Turku, Finland
eugenc@cs.utu.fi

[2] Department of Mathematics, FIN-20014 University of Turku, Finland, and Department of Computer Science, University of Iowa, Iowa City, IA 52242, USA
jjkari@cs.uiowa.edu

Abstract. Reversible cellular automata (RCA) are models of massively parallel computation that preserve information. They consist of an array of identical finite state machines that change their states synchronously according to a local update rule. By selecting the update rule properly the system has been made information preserving, which means that any computation process can be traced back step-by-step using an inverse automaton. We investigate the maximum range in the array that a cell may need to see in order to determine its previous state. We provide a tight upper bound on this inverse neighborhood size in the one-dimensional case: we prove that in a RCA with n states the inverse neighborhood is not wider than $n-1$, when the neighborhood in the forward direction consists of two consecutive cells. Examples are known where range $n-1$ is needed, so the bound is tight. If the forward neighborhood consists of m consecutive cells then the same technique provides the upper bound $n^{m-1}-1$ for the inverse direction.

1 Introduction

Cellular automata (CA) are discrete dynamical systems consisting of a grid of identical finite state machines whose states are updated synchronously at discrete time steps according to a local update rule. Cellular automata possess several fundamental properties of the physical world: they are massively parallel, homogeneous and all interactions are local. It is therefore not surprising that physical and biological systems have been successfully simulated using cellular automata models. The physical nature of cellular automata may have even greater importance when applied in the opposite direction, that is, when using the physics to simulate cellular automata. Many cellular automata are computationally universal — including some extremely simple ones, as reported recently

⋆ Research supported by the Academy of Finland grant 54102.

L. Caires et al. (Eds.): ICALP 2005, LNCS 3580, pp. 410–420, 2005.

by S. Wolfram [12] — so the most powerful massively parallel computers in the future may be implementations of cellular automata based on some physical phenomena of microscopic scale. Energy efficiency of such an implementation requires that the simulated universal CA obeys the rules of physics, including reversibility and conservation laws. Non-reversibility always implies energy dissipation, usually in the form of heat.

A cellular automaton is called *reversible* if there is another cellular automaton — the *inverse CA* — that computes the inverse function. The inverse CA retraces the computation steps back in time. There are simple reversible cellular automata that are computationally universal [5]. Universality is even possible in the one-dimensional space [7], that is, when the cells are organized along a line. Reversible CA have been popular topics of study since the early years of CA research, and many interesting facts have been discovered.

It is well known that injectivity and reversibility of CA are equivalent concepts: if a CA function has an inverse (i.e. it is one-to-one) then this inverse is always a CA function [2, 9]. This means that in order to backtrack the computation, each cell only needs to know the states of a finite number of its neighbors. The question this article investigates is the extent of the neighborhood that may be needed. In two- and higher dimensional cellular automata this inverse neighborhood can be extremely large: there is namely no algorithm to determine if a given CA is reversible, which means that the extent of the inverse neighborhood cannot be bounded by any computable function of the number of states [4]. In the one-dimensional case the reversibility question is decidable, and a trivial quadratic upper bound $O(n^2)$ exists [1], where n is the number of states and the neighborhood in the forward direction has been fixed to two consecutive cells.

We improve this bound to linear $n-1$ where n is the number of states. This bound is tight as examples of one-dimensional reversible CA are known whose inverse neighborhoods reach this bound [3]. If the neighborhood in the forward direction consists of m consecutive cells rather than two cells, then the same argument provides an upper bound $n^{m-1}-1$ for the inverse neighborhood. This is not known to be tight: [3] only provides examples of cellular automata with $2n$ states whose inverse neighborhoods reach this size.

2 Definitions and Basic Properties

In this section we present precise definitions and some basic properties of reversible cellular automata, and Welch sets and indices. Our proofs are based on elementary linear algebra, so we also recall some linear algebra concepts.

2.1 Cellular Automata

Formally, a one-dimensional cellular automaton, CA for short, is a 3-tuple system

$$\mathcal{A} = (S, N, f),$$

where $S = \{1, 2, \ldots, n\}$ is a finite *state set*, N is a *neighborhood vector*

$$N = (x_1, \ldots, x_m) \in \mathbb{Z}^m$$

of m distinct integers, and f is a mapping from S^m to S representing the *local update rule* of the CA. The *cells* are laid on an infinite line and are indexed by $\mathbb{Z}$, the set of integers. The neighbors of a cell situated on position $x \in \mathbb{Z}$ are all the cells on positions $x + x_i, i = 1, \ldots, m$. The local update rule f determines the future state of a cell according to the states of its neighbors.

A *configuration* c of a CA $\mathcal{A}$ is a mapping

$$c : \mathbb{Z} \to S$$

which specifies the states of all the cells. We are denoting by $\mathcal{C}$ the set of all configurations. The *global transition function*

$$G : \mathcal{C} \to \mathcal{C}$$

describes the evolution of the CA and is obtained by a simultaneous application of the local update rule f on all cells:

$$G(c)(x) = f(c(x + x_1), \ldots, c(x + x_m)),$$

for all $x \in \mathbb{Z}$. It is common to identify a cellular automaton with its global transition function G, and talk about cellular automaton function G or, when there is no risk of confusion, simply cellular automaton G.

If the neighborhood vector is $(-r, \ldots, r)$ then the CA is called *radius-r* automaton. The special case $r = 1$ is the nearest neighbor neighborhood. In this work we mainly consider CA whose neighborhood is even smaller and consists of just two consecutive integers. If $N = (0, 1)$ we say that we have a *radius-*$\frac{1}{2}$ CA. Figure 1 shows the trellis whose rows are consecutive configurations of a radius-$\frac{1}{2}$ cellular automaton, and the rows are shifted to make the neighborhood look symmetric. Note that any CA can be viewed as a radius-$\frac{1}{2}$ CA over a larger state set if we divide the configurations into sufficiently long blocks and use the blocks as "super cells". The partitioning may shift in time, but the computation is essentially the same.

Two CA are called equivalent if their global functions are identical. The following facts are easy to see: If two cellular automata are equivalent then there is a third equivalent CA whose neighborhood is the intersection of the neighborhoods of the first two CA. Hence, each CA function G has a minimal neighborhood, that is, a neighborhood that is contained in the neighborhoods of all CA that specify G. We call it the neighborhood of G. The interval from the smallest to the largest element of the minimal neighborhood is the neighborhood range for G. It is the smallest contiguous segment that can be used as the neighborhood to specify G.

A CA $\mathcal{A}$ with global function G is called *reversible*, for short RCA, if there exists another CA, called the *inverse automaton* of $\mathcal{A}$, whose global transition function is G^{-1}, the inverse of G. The minimal neighborhood of G^{-1} is called the

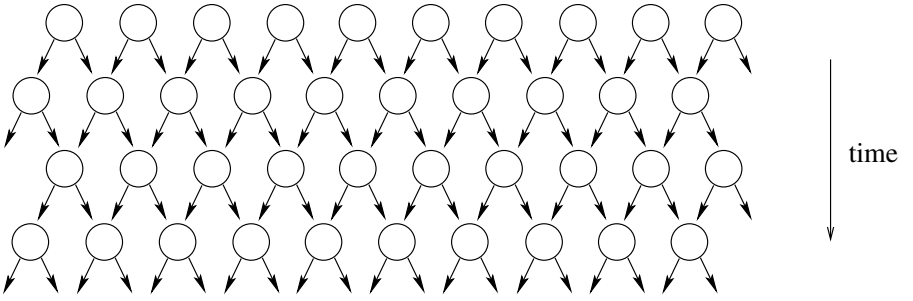

Fig. 1. Dependencies in a radius-$\frac{1}{2}$ cellular automaton

inverse neighborhood of $\mathcal{A}$ or G. Each cell can uniquely determine its previous state by looking only at the states contained in the inverse neighborhood.

A CA $\mathcal{A}$ is called *injective* (*surjective, bijective*) if its global transition rule $G : \mathcal{C} \to \mathcal{C}$ is an injective (surjective, bijective, respectively) function. It has been known since the early 60's that injective cellular automata are automatically also surjective [6, 8], while the converse is not necessarily true. It is also known that all bijective CA are reversible [2, 9]. We have

Property 1 ([2, 6, 8, 9]). *In cellular automata, reversibility, bijectivity and injectivity are equivalent. They imply surjectivity.*

2.2 Welch Sets and Indices

From now on we consider radius-$\frac{1}{2}$ RCA only. We frequently need to apply the CA on partial configuration where we only know the states on some contiguous interval. Since the exact location of the interval on the line is irrelevant, we specify such configurations as finite or infinite words. For the state set S we denote by S^* the set of all words over alphabet S, by S^k the set of words of length k, by S^ω the set of one-way infinite words that are infinite to the right, and by ${}^\omega S$ the set of words that are infinite to the left. CA $\mathcal{A} = (S, (0, 1), f)$ specifies the functions $G : S^* \longrightarrow S^*$, $G : S^\omega \longrightarrow S^\omega$ and $G : {}^\omega S \longrightarrow {}^\omega S$ (all denoted by the same symbol G) defined by

- $G(a_1 a_2 \dots a_k) = b_1 b_2 \dots b_{k-1}$ where each $b_i = f(a_i, a_{i+1})$,
- $G(a_1 a_2 \dots) = b_1 b_2 \dots$ where each $b_i = f(a_i, a_{i+1})$,
- $G(\dots a_2 a_1) = \dots b_2 b_1$ where each $b_i = f(a_{i+1}, a_i)$.

For each $w \in S^\omega$ we set

$$L(w) = \{a \in S \mid G(au) = w \text{ for some } u \in S^\omega \}$$

and call it the *left Welch set* of w. It contains all the states that were possible one time step earlier at the leftmost cell that affects w, see Figure 2. Analogously, for any $w \in {}^\omega S$ we define the *right Welch set* as

$$R(w) = \{a \in S \mid G(ua) = w \text{ for some } u \in {}^\omega S \}.$$

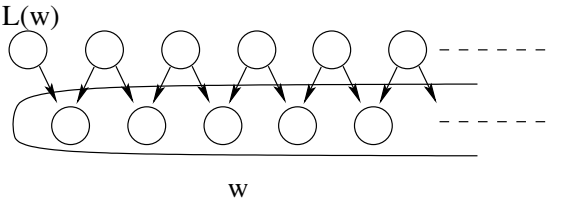

Fig. 2. The left Welch set $L(w)$ of the infinite word w consists of all possible states in the indicated cell

These sets were introduced already in [2], and have since been reinvented independently by many authors. The Welch sets have the following nice properties [2]:

Property 2. *Let $\mathcal{A} = (S, (0, 1), f)$ be reversible and let $n = |S|$ be the number of states. Then for every $w \in S^\omega$ and $v \in {}^\omega S$ we have*

$$|L(w)| \cdot |R(v)| = n.$$

Consequently, the cardinalities $|L(w)|$ and $|R(v)|$ are independent of the choice of w and v.

We denote by n_L the size of left Welch sets and by n_R the size of the right Welch sets, and call them the *left* and the *right Welch index*. Then $n_L \cdot n_R = n$.

The following result is another useful property of the Welch sets [2]:

Property 3. *Let $\mathcal{A} = (S, (0, 1), f)$ be reversible. Then for every $w \in S^\omega$ and $v \in {}^\omega S$ we have*

$$|L(w) \cap R(v)| = 1,$$

i.e. the intersection of any left Welch set with any right Welch set is a singleton.

The following proposition relates the Welch sets to the minimal inverse neighborhood of the CA:

Proposition 1. *Let $\mathcal{A} = (S, (0, 1), f)$ be reversible. Then the inverse neighborhood of G is included in the interval*

$$\{-r, \ldots, l - 1\}$$

if and only if

- *$L(uw) = L(uv)$ for all $u \in S^l$ and $w, v \in S^\omega$, and*
- *$R(wu) = R(vu)$ for all $u \in S^r$ and $w, v \in {}^\omega S$.*

Proof. First, notice that even though the given interval does not at first appear symmetric, it in fact contains l positions to the right of the cell and r positions to the left of the cell, if the cells are shifted to the right as in Figure 1.

($\Longrightarrow$) Suppose that $\{-r, \ldots, l - 1\}$ contains the inverse neighborhood of G and let $u \in S^l$ and $w, v \in S^\omega$ be arbitrary. Let us prove that $L(uw) \subseteq L(uv)$. Then by symmetry we have $L(uw) = L(uv)$.

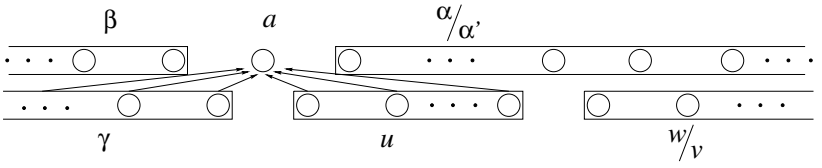

Fig. 3. Configurations γuw and γuv agree in all the positions up to $l-1$

If $a \in L(uw)$ then there exists some $\alpha \in S^\omega$ such that $G(a\alpha) = uw$. Pick an arbitrary $\beta \in {}^\omega S$ and let $\gamma = G(\beta a)$. Then $G(\beta a\alpha) = \gamma uw$, where letter a and the first letter of u are in the position 0 of the cellular array. Consider then the configuration γuv, where the first letter of u is still in position 0, see Figure 3. It agrees with γuw in positions up to $l-1$. By applying the inverse cellular automata to γuv we must therefore get state a in position 0, which means that $a \in L(uv)$.

Analogously we get the claim concerning the right Welch sets.

($\Longleftarrow$) Suppose l and r are such that for all $v \in S^l$ and $\alpha_1, \alpha_2 \in S^\omega$ we have $L(v\alpha_1) = L(v\alpha_2)$, and that for all $u \in S^r$ and $\beta_1, \beta_2 \in {}^\omega S$ we have $R(\beta_1 u) = R(\beta_2 u)$. Then the inverse function G^{-1} is computed by the cellular automaton that uses the neighborhood $(-r, \ldots, l-1)$ and has the local update rule

$$g(uv) = R(\beta u) \cap L(v\alpha)$$

for $u \in S^r$, $v \in S^l$ and all $\alpha \in S^\omega$ and $\beta \in {}^\omega S$. The above intersection always contains a unique element, due to Property 3. □

2.3 Vector Interpretation of Sets

In our proofs we take advantage of dimension arguments on vector spaces. Any subset X of the state set $S = \{1, 2, \ldots, n\}$ is interpreted as the 0-1 vector $\vec{X}$ in $\mathbb{R}^n$ whose i'th coordinate is 1 if $i \in X$ and 0 if $i \notin X$. The single element sets $\{a\}$ then correspond to the unit coordinate vectors of $\mathbb{R}^n$ and they form a basis of the vector space $\mathbb{R}^n$. Notice that for any $X, Y \subseteq S$ the inner product $\vec{X} \cdot \vec{Y}$ is the cardinality of their intersection $X \cap Y$. The vectors $\vec{L}$ and $\vec{R}$ corresponding to left and right Welch sets L and R will be called left and right Welch vectors, respectively.

Let us denote by Θ the null space $\{(0, 0, \ldots 0)\}$ and by I the one-dimensional space generated by vector $(1, 1, \ldots, 1)$. For any $U \subseteq \mathbb{R}^n$ the subspace of $\mathbb{R}^n$ generated by U is denoted as $\langle U \rangle$.

Let $\mathcal{A} = (S, (0, 1), f)$ be reversible. For any $c \in S$ we define a linear function $h_c : \mathbb{R}^n \longrightarrow \mathbb{R}^n$ as follows. For every $b \in S$ we have $h_c(\vec{b}) = \vec{H}$ where $\vec{b}$ is the basis vector corresponding to b and $H = \{a \mid f(a, b) = c\}$. This uniquely specifies the linear function h_c. Vector $\vec{X}$, corresponding to a set $X \subseteq S$ of states, is mapped according to $h_c(\vec{X}) = \sum_{b \in X} h_c(\vec{b})$. Note that $h_c(\vec{X})$ is not always a 0-1 vector, so it does not necessarily represent a set. However, the next proposition states that if L is a left Welch set then $h_c(\vec{L})$ is a 0-1 vector representing a left Welch set:

Proposition 2. *Let $\mathcal{A} = (S, (0,1), f)$ be reversible, and let $c \in S$ be arbitrary. For every $w \in S^\omega$ we have $h_c(\vec{L}(w)) = \vec{L}(cw)$.*

Proof. It is enough to show that (i) for every $a \in L(cw)$ there is a unique $b \in L(w)$ such that $f(a,b) = c$, and (ii) for any $a \notin L(cw)$ there is no $b \in L(w)$ such that $f(a,b) = c$. Parts (i) and (ii) imply then that the vector $h_c(\vec{L}(w))$ has 1 and 0 in coordinates i for all $i \in L(cw)$ and $i \notin L(cw)$, respectively.

Claim (ii) is trivial, as if there would exist $b \in L(w)$ such that $f(a,b) = c$ then $G(ab\alpha) = cw$ where $\alpha \in S^\omega$ is such that $G(b\alpha) = w$. This contradicts the assumption $a \notin L(cw)$.

Consider then claim (i). Since $a \in L(cw)$ there is some $b\alpha \in S^\omega$ such that $G(ab\alpha) = cw$. This b satisfies the condition in (i). If $b' \in L(w)$ is another state with the property $f(a,b') = c$ then $G(ab'\beta) = cw$ for some $\beta \in S^\omega$. But then $G(\gamma ab\alpha) = G(\gamma ab'\beta)$ for any $\gamma \in {}^\omega S$ which, by injectivity, implies that $b = b'$. □

Analogously, let us define linear functions $g_c(\vec{a}) = \vec{H}$ where $H = \{b \mid f(a,b) = c\}$. They naturally have the similar property concerning the right Welch sets:

Proposition 3. *Let $\mathcal{A} = (S, (0,1), f)$ be reversible, and let $c \in S$ be arbitrary. For every $w \in {}^\omega S$ we have $g_c(\vec{R}(w)) = \vec{R}(wc)$.* □

3 The Inverse Neighborhood Range

In this section we prove that the size of the inverse neighborhood range of a radius-$\frac{1}{2}$ RCA $\mathcal{A} = (S, (0,1), f)$ is less than or equal to $n-1$, where n is the number of states. We do this by creating two decreasing chains of linear subspaces of $\mathbb{R}^n$ based on the Welch sets. The first elements of the chains are the subspaces

$$\mathcal{L}_0 = \langle \vec{L}(w) - \vec{L}(v) \mid w, v \in S^\omega \rangle, \text{ and}$$
$$\mathcal{R}_0 = \langle \vec{R}(w) - \vec{R}(v) \mid w, v \in {}^\omega S \rangle,$$

that is, the spaces generated by the differences between any two left Welch vectors and any two right Welch vectors, respectively. The goal is to prove the following theorem:

Theorem 1. *Let $\mathcal{A} = (S, (0,1), f)$ be reversible, and let $\mathcal{L}_0$ and $\mathcal{R}_0$ be the subspaces defined above. Then the inverse neighborhood range of G contains at most $\dim \mathcal{L}_0 + \dim \mathcal{R}_0$ elements. More precisely, the inverse neighborhood of G is included in the interval*

$$\{-\dim \mathcal{R}_0, \ldots, \dim \mathcal{L}_0 - 1\}.$$

Proof. For every $k = 0, 1, 2, \ldots$ define the following subspaces of $\mathbb{R}^n$:

$$\mathcal{L}_k = \langle \vec{L}(uw) - \vec{L}(uv) \mid u \in S^k, w, v \in S^\omega \rangle, \text{ and}$$
$$\mathcal{R}_k = \langle \vec{R}(wu) - \vec{R}(vu) \mid u \in S^k, w, v \in {}^\omega S \rangle.$$

We make the following observations:

- There is l such that $\mathcal{L}_l = \Theta$, the null space,
- $\mathcal{L}_{k+1} \subseteq \mathcal{L}_k$ for every $k = 0, 1, 2, \ldots$, and
- if $\mathcal{L}_{k+1} = \mathcal{L}_k$ then $\mathcal{L}_j = \mathcal{L}_k$ for every $j \geq k$.

To prove the first fact, choose l and r such that the inverse neighborhood of G is included in the interval $\{-r, \ldots, l-1\}$. According to Proposition 1, $L(uw) = L(uv)$ for every $u \in S^l$ and $w, v \in S^\omega$. But then all generators of $\mathcal{L}_l$ are zero vectors, hence $\mathcal{L}_l = \Theta$.

The second fact is trivial since all the generators of $\mathcal{L}_{k+1}$ are among the generators of $\mathcal{L}_k$.

For the third fact, notice that $\vec{L}(cuw) - \vec{L}(cuv) = h_c(\vec{L}(uw) - \vec{L}(uv))$. This means that, for every $k = 0, 1, 2, \ldots$, the generators of $\mathcal{L}_{k+1}$ are obtained from the generators of $\mathcal{L}_k$ by applying the homomorphisms h_c with all $c \in S$. Consequently,

$$\begin{aligned}\mathcal{L}_{k+1} &= \langle h_c(\vec{X}) \mid c \in S, \vec{X} \text{ is a generator of } \mathcal{L}_k \rangle \\ &= \langle h_c(\vec{X}) \mid c \in S, \vec{X} \in \mathcal{L}_k \rangle.\end{aligned}$$

In other words, $\mathcal{L}_{k+1}$ is determined by $\mathcal{L}_k$. It follows that if $\mathcal{L}_{k+1} = \mathcal{L}_k$ then $\mathcal{L}_{k+2} = \mathcal{L}_{k+1}$, and therefore $\mathcal{L}_j = \mathcal{L}_k$ for all $j \geq k$.

Our three facts imply that

$$\mathcal{L}_0 \supsetneq \mathcal{L}_1 \supsetneq \mathcal{L}_2 \supsetneq \ldots \supsetneq \mathcal{L}_l = \Theta$$

for some l. Since the dimension of the subspaces decreases at every step, we must have $l \leq \dim \mathcal{L}_0$.

The analogous reasoning can be done on the right Welch sets. We conclude that there are numbers $l \leq \dim \mathcal{L}_0$ and $r \leq \dim \mathcal{R}_0$ such that $\mathcal{L}_l = \mathcal{R}_r = \Theta$. Then l has the property that $L(uw) = L(uv)$ for every $u \in S^l$ and $w, v \in S^\omega$, and r has the property that $L(wu) = L(vu)$ for every $u \in S^r$ and $w, v \in {}^\omega S$. According to Proposition 1, the inverse neighborhood of G is included in the interval $\{-r, \ldots, l-1\}$ and hence also in the interval $\{-\dim \mathcal{R}_0, \ldots, \dim \mathcal{L}_0 - 1\}$. □

Upper bounds on the dimensions of the spaces $\mathcal{L}_0$ and $\mathcal{R}_0$ provide nice limits on the inverse neighborhood:

Corollary 1. *Let $\mathcal{A} = (S, (0,1), f)$ be reversible. Then $\dim \mathcal{L}_0 + \dim \mathcal{R}_0 \leq n-1$ where n is the number of states. Hence the inverse neighborhood range of G has at most size $n-1$*

Proof. This follows from the facts that vector spaces $\mathcal{L}_0$ and $\mathcal{R}_0$ are orthogonal to each other and also to the one-dimensional space I generated by the vector $(1, 1, \ldots, 1)$. Let $\vec{L}_1 - \vec{L}_2$ and $\vec{R}_1 - \vec{R}_2$ be two arbitrary generators of $\mathcal{L}_0$ and $\mathcal{R}_0$, respectively. Their inner product is

$$(\vec{L}_1 - \vec{L}_2) \cdot (\vec{R}_1 - \vec{R}_2) = \vec{L}_1 \cdot \vec{R}_1 - \vec{L}_1 \cdot \vec{R}_2 - \vec{L}_2 \cdot \vec{R}_1 + \vec{L}_2 \cdot \vec{R}_2 = 1 - 1 - 1 + 1 = 0,$$

where we have used Property 3 of the Welch sets. So spaces $\mathcal{L}_0$ and $\mathcal{R}_0$ are orthogonal to each other. With $(1, 1, \ldots, 1)$ we get the inner product

$$(\vec{L}_1 - \vec{L}_2) \cdot (1, 1, \ldots, 1) = n_L - n_L = 0,$$

where n_L is the left Welch index. Here we used Property 2 of the Welch sets. Analogously $\mathcal{R}_0$ is seen orthogonal to I.

Now we can reason as follows: Since the three spaces are orthogonal, we have

$$\dim \mathcal{L}_0 + \dim \mathcal{R}_0 + \dim I = \dim(\mathcal{L}_0 \oplus \mathcal{R}_0 \oplus I) \leq \dim \mathbb{R}^n = n,$$

so

$$\dim \mathcal{L}_0 + \dim \mathcal{R}_0 \leq n - 1.$$

□

We can also use our theorem to bound the inverse neighborhood from either side separately:

Corollary 2. *Let $\mathcal{A} = (S, (0, 1), f)$ be reversible, and let n_L and n_R be its left and right Welch indices, respectively. Then* $\dim \mathcal{L}_0 \leq n - n_L$ *and* $\dim \mathcal{R}_0 \leq n - n_R$ *where n is the number of states. Hence the inverse neighborhood of G is contained in the interval*

$$\{n_R - n, \ldots, n - n_L - 1\}.$$

Proof. Consider the left Welch vectors $\vec{L}(u)$, $u \in S^\omega$. Each is a 0-1 vector with n_L ones. Every state belongs to some left Welch set, so each position has one in some of the vectors. Out of all this vectors, we can extract a set of linearly independent ones as follows. First, extract an arbitrarily vector. Then, for any state $a \in S$ such that the corresponding position is zero in all the vectors already selected, extract a left Welch vector having one in position a, and add it to the set of linearly independent vectors. Repeat the process until each position is covered by at least one selected vector. It is clear that the extracted vectors are linearly independent, and since each vector covers n_L positions there are at least $\frac{n}{n_L}$ vectors selected. Since $\frac{n}{n_L} = n_R$, it follows easily that there are at least n_R linearly independent left Welch vectors.

Next we use the following well known property: if $\vec{v}_1, \vec{v}_2, \ldots, \vec{v}_k$ are k linearly independent vectors, then $\vec{v}_2 - \vec{v}_1, \vec{v}_3 - \vec{v}_1, \ldots, \vec{v}_k - \vec{v}_1$ are $k - 1$ linearly independent vectors. We apply this to the n_R linearly independent left Welch vectors and obtain the result that the generators of $\mathcal{L}_0$ contain at least $n_R - 1$ linearly independent vectors, so $\dim \mathcal{L}_0 \geq n_R - 1$. On the other hand we know by Corollary 1 that $\dim \mathcal{L}_0 + \dim \mathcal{R}_0 \leq n - 1$, so

$$\dim \mathcal{R}_0 \leq (n - 1) - \dim \mathcal{L}_0 \leq (n - 1) - (n_R - 1) = n - n_R.$$

Analogously we can prove that $\dim \mathcal{L}_0 \leq n - n_L$. □

The previous corollaries were proved in [3] in the special case that one of the Welch indices is 1. This constraint simplifies the proofs very much. The techniques used in [3] were quite different. Examples were also provided in [3] of reversible CA with n states and Welch index 1 whose inverse neighborhood reached the size $n-1$. Hence the bound of Corollary 1 is tight. We do not know if there are such examples for other values of the Welch indices, and also we do not know if the bounds in Corollary 2 are tight.

4 Larger Neighborhoods

So far we have been concerned with radius-$\frac{1}{2}$ cellular automata. With larger forward neighborhoods larger inverse neighborhoods are possible. The notions of the Welch sets and indices can be generalized to such settings. Let m be the size of the neighborhood range in the forward direction, that is, m consecutive positions can be used as the forward neighborhood. Then the elements of the Welch sets are words of length $m-1$ over alphabet S, and the Welch indices n_L and n_R satisfy the relation $n_L \cdot n_R = n^{m-1}$. By a straightforward generalization of the proofs in the previous section we obtain the following results:

Theorem 2. *Let $\mathcal{A} = (S, (0, \ldots, m-1), f)$ be a reversible CA with n states and forward neighborhood range m. Then the inverse range has size at most $n^{m-1} - 1$. Moreover, the size of the left inverse neighborhood is less than or equal to $n^{m-1} - n_R$ while the size of the right inverse neighborhood is less than or equal to $n^{m-1} - n_L$, where n_L and n_R are the left and right Welch indices.*

The bound in Theorem 2 is not known to be tight. The best known examples are automata with $2n$ states whose inverse neighborhood have range n^{m-1} [3].

5 Final Remarks

We have shown that the inverse neighborhood of a one-dimensional reversible cellular automaton of size n is at most $n-1$ when the neighborhood in the forward direction consists of only two consecutive cells. We have also generalized this result for the case when the forward neighborhood is wider, i.e., if it contains m consecutive cells, then the size of the inverse neighborhood is bounded by $n^{m-1} - 1$. The proof uses several properties of the Welch sets, as well as some algebraic results concerning dimension of vector spaces.

The present paper gives rise to several open problems. E.g., we do not know if the generalized bound $n^{m-1} - 1$ for the size of the inverse neighborhood is tight. This is indeed the case if $m = 2$: see [3] for an example of a reversible cellular automaton with left and right Welch indexes equal to 1 and n respectively, and with inverse neighborhood size equal to $n - 1$. Also, for any $n_L, n_R \in \mathbb{N}$, it remains open to find examples of reversible cellular automata with left and right Welch indexes equal to n_L and n_R respectively, such that the size of the inverse neighborhood is maximal, i.e., equal to $n - 1 = n_L \cdot n_R - 1$.

There are quadratic time algorithms in the literature testing for surjectivity and injectivity of a given cellular automaton, see [1] and [10]. Although it is improbable that a linear algorithm exists, some improvements may be possible. For example, Lemma 3 from [1] can now be improved from quadratic to linear, although the time complexity of the injectivity algorithm, based on that result, does not change.

References

1. S. Amoroso and Y. Patt, Decision Procedures for Surjectivity and Injectivity of Parallel Maps for Tessellation Structures, Journal of Computer and System Sciences 6 (1972) 448–464.
2. G. Hedlund, Endomorphisms and automorphisms of shift dynamical systems, Mathematical Systems Theory 3 (1969) 320–375.
3. J. Kari, On the Inverse Neighborhood of Reversible Cellular Automata, in: Lindenmayer Systems, Impact in Theoretical Computer Science, Computer Graphics and Developmental Biology, G. Rosenberg, A. Salomaa, eds., 477–495, Springer-Verlag, Berlin-Heidenberg, 1989.
4. J. Kari, Reversibility and surjectivity problems of cellular automata, Journal of Computer and System Sciences 48 (1994) 149–182.
5. N. Margolus, Physics-like models of computation, Physica D 10 (1984) 81–95.
6. E.F. Moore, Machine Models of Self-reproduction, Proceedings of the Symposium in Applied Mathematics 14 (1962) 17–33.
7. K. Morita and M. Harao, Computation Universality of one-dimensional reversible (injective) cellular automata, IEICE Transactions E72 (1989) 758–762.
8. J. Myhill, The Converse to Moore's Garden-of-Eden Theorem, Proceedings of the American Mathematical Society 14 (1963) 685–686.
9. D. Richardson, Tessellations with Local Transformations, Journal of Computer and System Sciences 6 (1972) 373–388.
10. K. Sutner, De Bruijn graphs and linear cellular automata, Complex Systems 5 (1991) 19–31.
11. T. Toffoli and N. Margolus, Invertible cellular automata: a review, Physica D 45 (1990) 229–253.
12. S. Wolfram, A New Kind of Science, Wolfram Media, 2002.