A One-Dimensional Physically Universal Cellular Automaton*

Ville Salo** and Ilkka Törmä

Department of Matematics and Statistics,
University of Turku, Finland
{vosalo,iatorm}@utu.fi

Abstract. Physical universality of a cellular automaton was defined by Janzing in 2010 as the ability to implement an arbitrary transformation of spatial patterns. In 2014, Schaeffer gave a construction of a two-dimensional physically universal cellular automaton. We construct a one-dimensional version of the automaton and a reversibly universal automaton.

Keywords: cellular automaton, physical universality, reversibility

1 Introduction

A cellular automaton (CA) is a finite or infinite lattice of deterministic finite state machines with identical interaction rules, which, at discrete time steps, update their states simultaneously based on those of their neighbors. They are an idealized model of massively parallel computation. From another point of view, the local updates can be seen as particle interactions, and the CA is then a kind of physical law, or dynamics, governing the universe of all state configurations.

We study the notion of *physical universality* of cellular automata, introduced by Janzing in [4], which combines the two viewpoints in a novel and interesting way. Intuitively, a cellular automaton is physically universal if, given a finite subset D of the lattice and a function h on the shape-D patterns over the states of the CA, one can build a 'machine' in the universe of the CA that, under the dynamics of the CA, decomposes any given pattern P and replaces it by $h(P)$.

A crucial point in this definition is that we need to perform arbitrary computation on all patterns, not only carefully constructed ones. This has quite serious implications: The machine M that takes apart an arbitrary pattern of shape D and replaces it by its image under an arbitrary function h is not in any way special, and in particular we are not allowed to have separate 'machine states' and 'data states' with the former operating on the latter. Instead, we can also think of M as a pattern, and must be able to construct a larger machine M' that takes M apart and reassembles it in an arbitrary way.

* Research supported by the Academy of Finland Grant 131558
** The author was partially supported by CONICYT Proyecto Anillo ACT 1103

This notion differs essentially from most existing notions of universality for CA such as *intrinsic universality* [7], universality in terms of traces as discussed by Delvenne et al. in [3], and the several different versions of computational universality [2,6]. In these notions, one can usually implement the computations and simulations in a well-behaved subset of configurations. Physical universality bears more resemblance to the *universal constructor machines* of Von Neumann [12], which construct copies of themselves under the dynamics of a particular cellular automaton, and were the initial motivation for the definition of CA. Another property of CA with a similar flavor is *universal pattern generation* as discussed in [5], meaning the property of generating all finite patterns from a given simple initial configuration.

In Janzing's work [4] some results were already proved about physically universal CA, but it was left open whether such an object actually exists. A two-dimensional physically universal CA was constructed by Schaeffer in [8] (see also [9,10]), but it was left open whether this CA can be made one-dimensional. We construct such a CA, solving the question in the positive. We also outline the construction of a *reversibly physically universal* CA, solving another open problem of [8].

2 Definitions

We now define the terms and notation used in this article. Let A be a finite set, called the *state set*, and $\mathbb{Z}^d$ a d-dimensional grid; we will mostly restrict our attention to the case $d = 1$. We call $A^{\mathbb{Z}^d}$ the *d-dimensional full shift over A*, whose elements are *configurations*.

A *cellular automaton* (CA for short) is a map $f : A^{\mathbb{Z}^d} \to A^{\mathbb{Z}^d}$ defined by a finite *neighborhood* $\{\boldsymbol{n}_1, \ldots, \boldsymbol{n}_k\} \subset \mathbb{Z}^d$ and a *local function* $F : A^k \to A$, so that $f(x)_{\boldsymbol{v}} = F(x_{\boldsymbol{v}+\boldsymbol{n}_1}, \ldots, x_{\boldsymbol{v}+\boldsymbol{n}_k})$ holds for all $x \in A^{\mathbb{Z}^d}$ and $\boldsymbol{v} \in \mathbb{Z}^d$. It is *reversible* if there is another CA $g : A^{\mathbb{Z}^d} \to A^{\mathbb{Z}^d}$ such that $f \circ g = g \circ f = \mathrm{id}$. An example of a reversible CA is the *shift by $\boldsymbol{n} \in \mathbb{Z}^d$*, defined by $\sigma^{\boldsymbol{n}}(x)_{\boldsymbol{v}} = x_{\boldsymbol{v}+\boldsymbol{n}}$.

Other examples of reversible CA can be constructed as follows. Let the state set A be a Cartesian product $\prod_{i=1}^{k} A_i$ with projection maps $\pi_i : A \to A_i$, let $\boldsymbol{n}_1, \ldots, \boldsymbol{n}_k \in \mathbb{Z}^d$ be arbitrary vectors, and let $\gamma : A \to A$ be a bijection. Then the CA f defined by $f(x)_{\boldsymbol{v}} = \gamma(\pi_1(x_{\boldsymbol{v}+\boldsymbol{n}_1}), \ldots, \pi_k(x_{\boldsymbol{v}+\boldsymbol{n}_k}))$ is reversible. We call f a *partitioned CA*. The components A_i are called the *tracks* of f, and the numbers $\boldsymbol{n}_i$ are *shifts* associated to the tracks. In the CA, the tracks are first shifted individually by the vectors $\boldsymbol{n}_i$, and then the bijection γ is applied to every cell.

A CA f is *physically universal* if the following condition holds. For all finite domains $D, E \subset \mathbb{Z}^d$, and all functions $h : A^D \to A^E$, there exists a partial configuration $x \in A^{\mathbb{Z}^d \setminus D}$ and a time step $t \in \mathbb{N}$ such that for all $P \in A^D$, we have $f^t(x \cup P)_E = h(P)$. We think of the partial configuration x as a 'gadget' that implements the function h: if any pattern $P \in A^D$ is placed in the unspecified part of x and the CA f is applied exactly t times, the image $h(P)$ appears on the domain E. We say f is *efficiently physically universal* if t is polynomial in the diameter of $D \cup E$ and the computational complexity of h according to some 'reasonable' complexity measure. In this article, we use circuit complexity, or more precisely, the number of binary NAND gates needed to implement h. One could reasonably require also that the configuration x is computed efficiently from the circuit presentation of the function h, which can be seen as a uniformity condition in the sense of circuit complexity. Our proof gives a polynomial time algorithm for this. See Section 9 for a discussion.

3 The Cellular Automaton

Our physically universal automaton is a partitioned CA f defined as follows.

- The state set is $A = \{0,1\}^4$.
- The shifts are 2, 1, -1 and -2, and we denote $\mathcal{S} = \{2, 1, -1, -2\}$.
- For each $a, b \in \{0,1\}$ bijection γ maps the state $(1, a, b, 1)$ to $(1, b, a, 1)$, and $(a, 1, 1, b)$ to $(b, 1, 1, a)$. Everything else is mapped to itself.

Alternatively, f is the CA with neighborhood $\{-2, -1, 1, 2\}$ and local rule

$$(\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c}, \boldsymbol{d}) \mapsto \begin{cases} (\pi_4(\boldsymbol{d}), 1, 1, \pi_1(\boldsymbol{a})), & \text{if } \pi_2(\boldsymbol{b}) = \pi_3(\boldsymbol{c}) = 1, \\ (1, \pi_3(\boldsymbol{c}), \pi_2(\boldsymbol{b}), 1), & \text{if } \pi_1(\boldsymbol{a}) = \pi_4(\boldsymbol{d}) = 1, \\ (\pi_1(\boldsymbol{a}), \pi_2(\boldsymbol{b}), \pi_3(\boldsymbol{c}), \pi_4(\boldsymbol{d})), & \text{otherwise} \end{cases}$$

where $\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c}, \boldsymbol{d} \in A$ are length-4 Boolean vectors.

Intuitively, in the CA f there are four kinds of particles, which are represented by 1-symbols on the four tracks: fast right (track 1, speed 2), slow right (track 2, speed 1), slow left (track 3, speed -1) and fast left (track 4, speed -2). At most one particle of each kind can be present in a cell. On each step, every particle jumps from its position $n \in \mathbb{Z}$ to $n + s$, where $s \in \mathcal{S}$ is its speed. After that, if two fast or two slow particles are present in the same cell, then the direction of every particle of the other speed is reversed. Such interactions are depicted in Figure 1. This resembles the two-dimensional CA of Schaeffer, where particles move in four directions (NE, NW, SE, SW) with speed one, and the head-on collision of two particles causes other particles in the same cell to make a u-turn.

4 The Logical Cellular Automaton

For the proof of physical universality, we define another CA on an infinite state set. Denote the *ternary conditional operator* by $p(a, b, c) = (a \wedge b) \vee (\neg a \wedge c)$ for $a, b, c \in \{0,1\}$. That is, $p(a, b, c)$ is equal to b if $a = 1$, and to c otherwise. In many programming languages, $p(a, b, c)$ is denoted by $a \mathbin{?} b : c$.

Definition 1. *Let $\mathcal{V} = \{\alpha_1, \alpha_2, \ldots\}$ be an infinite set of variables, and denote by $\mathcal{F}$ the set of Boolean functions over finitely many variables of $\mathcal{V}$. The* logical extension *of f is the CA-like function $\hat{f} : \hat{A}^{\mathbb{Z}} \to \hat{A}^{\mathbb{Z}}$ on the infinite state set $\hat{A} = \mathcal{F}^4$, where the four tracks are first shifted as in f, and then the function*

$$(a, b, c, d) \mapsto (p(b \wedge c, d, a), p(a \wedge d, c, b), p(a \wedge d, b, c), p(b \wedge c, a, d))$$

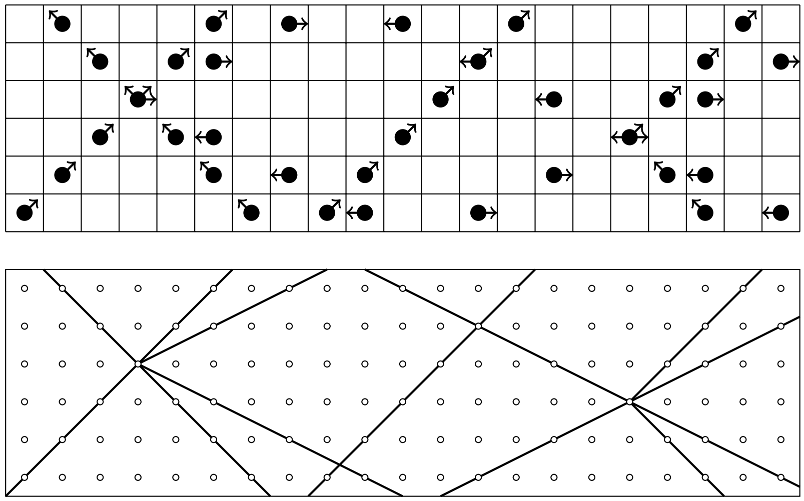

Fig. 1: A sample spacetime diagram of f showing different particle interactions, and a schematic version below it. Time increases upward. Particles are represented by arrowed bullets. For example, a bullet with arrows to the east, northeast and northwest represents three particles moving at speeds 2, 1 and -1, respectively.

is applied to each coordinate. A valuation *is a function $v : \mathcal{V} \to \{0, 1\}$. It extends to $\mathcal{F}$ and then into a function $v : \hat{A}^{\mathbb{Z}} \to A^{\mathbb{Z}}$ in the natural way.*

The logical extension simulates multiple spacetime diagrams of f: one can see that the definition of $\hat{f}$ is equal to that of f, except that each particle is replaced by a Boolean formula that corresponds to the conditional presence or absence of a particle. We think of A as a subset of $\hat{A}$ containing the constant 0 or constant 1 function in each track. Note that $\hat{f}$ is also reversible, and we denote by $\hat{f}^{-1}$ its inverse function. See Figure 2 for a spacetime diagram of $\hat{f}$.

The following result holds basically by construction.

Lemma 1. *Let $x \in \hat{A}^{\mathbb{Z}}$ be a configuration, and let $v : \mathcal{V} \to \{0, 1\}$ be a valuation, so that $v(x) \in A^{\mathbb{Z}}$. Then for all $t \in \mathbb{Z}$ we have $f^t(v(x)) = v(\hat{f}^t(x))$.*

The idea of the proof of physical universality of f using this new CA is the following. We may assume that $D = E = [0, n-1]$ in the definition of physical universality, for some $n \in \mathbb{N}$. Then, we construct a spacetime diagram of $\hat{f}$ with the following properties. First, in the initial configuration $x \in \hat{A}^{\mathbb{Z}}$, the cells of the interval $[0, n-1]$ contain $4n$ distinct variables from $\mathcal{V}$. All other cells of x contain either 0 or 1. There also exists $t > 0$ such that $\hat{f}^t(x)_{[0,n-1]}$ contains the Boolean functions computing the function h in the definition of physical universality. In the course of the construction, we define which cells of x contain a 1.

The construction proceeds in five stages, which are depicted in Figure 3:

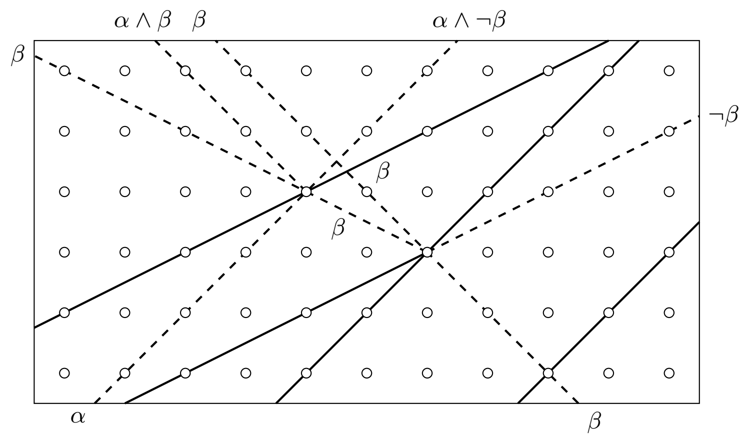

Fig. 2: A schematic spacetime diagram of $\hat{f}$, where $\alpha, \beta \in \mathcal{F}$. Particles are represented by solid lines, and Boolean particles by dotted lines. Note that Boolean particles can be created, even though f itself conserves the number of particles.

- *Diffusion Stage*, where the Boolean particles in the input area D disperse into the environment. This stage does not require any auxiliary particles, as Lemma 2 in Section 5 will show.
- *Collection Stage*, where the dispersed particles are rerouted to travel in the same direction at the same speed. This stage is implemented in Section 6.2.
- *Computation Stage*, where the NAND-gates of the circuit representing h are applied one by one to the travelling Boolean particles. This stage is implemented in Section 6.3.
- *Assembly Stage*, where the Boolean particles resulting from the computation are rerouted again toward the output area E.
- *Reverse Diffusion Stage*, where the particles converge on the output area and produce the desired pattern encoding h exactly at time t.

Definition 2. *We introduce the following terminology for the construction.*

- *The* configuration of interest*, denoted by $x \in \hat{A}^{\mathbb{Z}}$, initially contains the 'fully general' state $(\alpha_{4i}, \alpha_{4i+1}, \alpha_{4i+2}, \alpha_{4i+3})$ in every cell $i \in [0, n-1]$, and 0 everywhere else. During the construction, we change the cells of x from 0 to 1, but keep referring to it as x, so some of the definitions below depend on the stage of the proof. The* spacetime diagram of interest *is defined similarly.*
- *A* (spacetime) position *is an element of $\mathbb{Z} \times \mathcal{S}$ ($\mathbb{Z} \times \mathbb{Z} \times \mathcal{S}$, respectively), representing a (spacetime) point that may contain a particle of certain speed. Note that time is bi-infinite, since our cellular automata are reversible.*
- *There is a* Boolean particle *at spacetime position (i, t, s) if $\pi_s(\hat{f}^t(x)_i)$ is not the constant 0 function, and a* particle *if it is the constant 1 function.*
- *There is a* collision *at coordinate $(i, t) \in \mathbb{Z} \times \mathbb{Z}$ if $\hat{f}^t(x)_i$ contains at least three Boolean particles, and a* crossing *if there are at least two Boolean particles.*

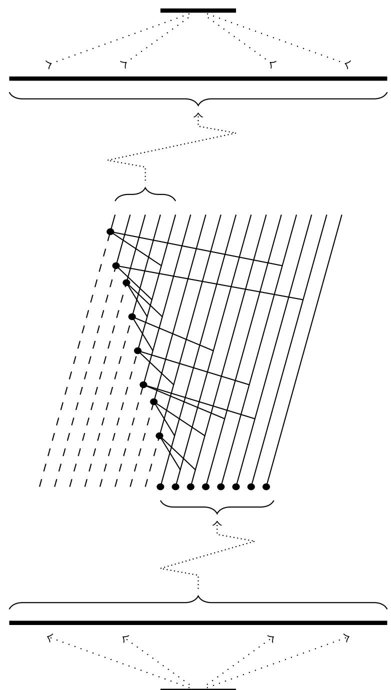

Fig. 3: A schematic diagram of the construction, not drawn to scale.

- *The* input *is the pattern* $x_{[0,n-1]} \in \hat{A}^n$*.*
- *The* gadget *is the contents of* x *outside* $[0, n-1]$*, an element of* $A^{\mathbb{Z} \setminus [0,n-1]}$*.*
- *A* line *is a subset of* $\mathbb{Z} \times \mathbb{N}$ *of the form* $L = L(i, t, s) = \{(i + st, t) \mid t \in \mathbb{N}\}$ *for some speed* $s \in \mathcal{S}$ *and* $i \in \mathbb{Z}$*. It is* occupied (in a time interval $I \subset \mathbb{Z}$) *if one of its coordinates (in the region* $\mathbb{Z} \times I$*) contains a crossing or a Boolean particle of speed* s*. We denote* $L^{(t)} = i + st$ *and* $L^t = (L^{(t)}, t, s)$*. The set of occupied lines in the spacetime diagram of interest is denoted* $\mathcal{L}_{\text{occ}}$*.*

For example, there are three crossings in Figure 2, two of which are collisions. The highest intersection of two dashed lines is not a crossing, as it does not take place at an actual coordinate (one of the white circles). Every line segment in the figure defines an infinite occupied line.

5 The Diffusion Lemma

As stated above, we initialize the gadget to the all-0 partial configuration, in which situation we have the following lemma. It states that any finite set of particles in the CA f eventually stop interacting and scatter to infinity. The corresponding result is proved for the physically universal CA of [8] by considering an *abstract CA* over the state set $\{0, \frac{1}{2}, 1\}$, with the interpretation that $\frac{1}{2}$ can be either 0 or 1, and this lack of information is suitably propagated in collisions. In our version, the role the new state $\frac{1}{2}$ is played by Boolean particles.

Lemma 2 (Diffusion Lemma). *Let $x \in \hat{A}^{\mathbb{Z}}$ be such that $x_i = 0$ for all $i \notin [0, n-1]$. Then there are $O(n^2)$ crossings in the two-directional spacetime diagram of x under $\hat{f}$ that all happen in a time window of length $O(n)$. For all other times $t \in \mathbb{Z}$, there are $O(n)$ Boolean particles in $\hat{f}^t(x)$, and they are contained in the interval $[-2|t|, n + 2|t|]$.*

Proof. We prove the claim in the positive direction of time. By induction, one sees that after any $t \geq 0$ steps, there can be no right-going Boolean particles in the cells $(-\infty, t-1]$, and no left-going particles in the cells $[n-t, \infty)$. After these sets intersect at time $\lceil n/2 \rceil$, there are no collisions, so no new Boolean particles are created. Thus the number of Boolean particles is at most $6n$, that is, twice the length of the segment of $\hat{f}^{\lceil n/2 \rceil}(x)$ that may contain Boolean particles. We may also have $O(n^2)$ crossings between Boolean particles going in the same direction with different speeds. Thus there are $O(n^2)$ crossings in total. □

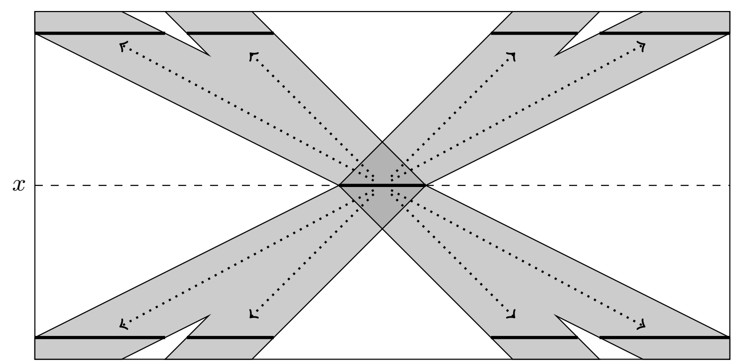

Fig. 4: An illustration of Lemma 2. The dashed line is the configuration x, and the thick segment is the interval $[0, n-1]$. All collisions take place in the dark gray area, and Boolean particles may occur in the light gray area. The eight horizontal line segments are the scattering Boolean particles, grouped by speed.

At this point, in the configuration of interest, we have an empty gadget, and the spacetime diagram contains $O(n)$ Boolean particles at any given time. Since the CA $\hat{f}$ is reversible, the values of the corresponding Boolean expressions determine the values of the original variables.

6 Manipulating the Spacetime Diagram

6.1 Controlled Modifications

In this section, we introduce new particles in the gadget that will collide with the existing Boolean particles and create new ones. This is called *scheduling* *collisions*. We never schedule a collision on an occupied line, and never add a Boolean particle on an existing crossing. This is formalized in the following definition.

Definition 3. *We say a modification of the gadget is* $(m, n, \mathcal{L}, \mathcal{T}, t, I)$-controlled *for numbers* $m, n \in \mathbb{N}$*, sets of lines* $\mathcal{L}$ *and* $\mathcal{T}$*, time* $t \in \mathbb{Z}$*, and interval* $I \subset \mathbb{Z}$*, if the following conditions hold:*

1. *the modification consists of adding at most* m *particles to the gadget,*
2. *at most* m *new occupied lines and* n *new crossings are introduced,*
3. *no existing crossings become collisions,*
4. *no line in* $\mathcal{L} \cup \mathcal{L}_{\mathrm{occ}}$ *is occupied by a new Boolean particle or a new crossing,*
5. *no spacetime position in the set*

$$F(\mathcal{T}, t, I) = \{(i, t, s') \mid i \in I, s' \in D\} \setminus \mathcal{T}^t \tag{1}$$

   gets a new Boolean particle, and

6. *no line in* $\mathcal{T}$ *gets a collision after time* t*.*

If the conditions hold in the time interval $(-\infty, t]$ *(in particular, condition 6 need not hold at all), then the modification is* weakly $(m, n, \mathcal{L}, \mathcal{T}, t, I)$-controlled.

In practice, a controlled modification is one where we add to the gadget a finite number of particles that affect the spacetime diagram of interest only where we want it to be affected: the spacetime positions in $\mathcal{T}^t$. The numbers m and n control the amount of new objects added to the diagram (recall that our goal is the keep its size polynomial). An existing crossings should not become a collision, since that could affect the labels of the crossing Boolean particles, which may store information about the input pattern. The lines in $\mathcal{L} \cup \mathcal{L}_{\mathrm{occ}}$ and the positions near $\mathcal{T}$, that is, those in $F(\mathcal{T}, t, I)$, are 'protected' from accidentally obtaining any auxiliary Boolean particles created in the modification. The role of $\mathcal{L}$ and I is to guarantee an empty area in the spacetime diagram where we may reroute more Boolean particles at a later stage of the construction.

The notion of controlled modifications allows us to prove simple lemmas that affect one Boolean particle at a time, and use them repeatedly as 'black boxes' to perform more complex modifications without explicitly specifying the path of every particle in the spacetime diagram. This makes our proof more flexible, but also somewhat technical.

Definition 4. *Let* $j, t \in \mathbb{Z}$*. The* positive cone rooted at (j, t) *is the set of spacetime coordinates* $\mathcal{C}(j, t) = \{(j + j', t + t') \mid t' \in \mathbb{N}, j' \in [-t', t']\}$*.*

The following lemmas are parametrized by the numbers m_i for $i = 1, 2, 3, 4$. They denote, intuitively, the number of existing crossings, the number of occupied lines, the size of the additional 'forbidden area,' and the number of particles involved, respectively. Also, the expression 'if $t' = \Omega(N(m_1, m_2, m_3, m_4))$ is large enough, then $P(t')$ holds' for $t' \in \mathbb{N}$, a function $N : \mathbb{N}^4 \to \mathbb{N}$ and a property P means that there exists a number $T \leq K \cdot N(m_1, m_2, m_3, m_4) + K'$, where $K, K' > 0$ are constants independent of the m_i, such that $P(t')$ holds for all $t' \geq T$.

6.2 Moving the Boolean Particles

We now prove that we can add a finite number of particles to the gadget, so that after some number of steps, a collection of Boolean particles is 'moved' onto any desired lines, at the same time. We do this one particle at a time, and without interfering with the trajectories of any other existing particles. The statements of the following lemmas refer to labels of spacetime positions instead of Boolean particles, since we want to also handle 'Boolean particles' whose label happens to be 0. However, in the intuitive explanations we still refer to Boolean particles.

Lemma 3. *Suppose we have a spacetime diagram of interest with m_1 crossings and m_2 occupied lines, let L_p be a line that contains no collisions after time t, and let $\beta \in \mathcal{F}$ be the label of the spacetime position L_p^t. Let $\mathcal{L}$ be a collection of m_3 lines not containing L_p. Let $L \notin \mathcal{L}$ be an unoccupied line that passes through some spacetime coordinate $(j', t') \in \mathcal{C}(j, t)$ with $t' > t$. Let $\mathcal{T}$ be a set of $O(m_3)$ lines containing L, and let $I \subset \mathbb{Z}$ be an interval of length $O(m_3)$. If $t' = t + \Omega(m_1 + m_2 + m_3)$ is large enough, then there is an $(O(1), O(m_2), \mathcal{L}, \mathcal{T}, t', I)$-controlled modification after which the spacetime position $L^{t'}$ has label β. The same holds if the line L is unoccupied only in the time interval $(-\infty, t']$, but the modification is weakly controlled.*

The lemma intuitively states that we can almost completely freely change the position and direction of a Boolean particle by introducing a constant number of auxiliary particles, while preserving the already constructed gadget and introducing relatively few new crossings. The point of the set of forbidden lines $\mathcal{L}$ and positions I is that we can in fact move an arbitrary number of Boolean particles at the same time. The idea is that we move them by applying Lemma 3 repeatedly to one Boolean particle at a time, always adding the target lines of the remaining ones into the protected set $\mathcal{L}$. This guarantees that the lines are not accidentally occupied too early.

Corollary 1. *Suppose we have a spacetime diagram of interest with m_1 crossings and m_2 occupied lines, and for each $k = 1, \ldots, m_4$, a line L_{p_k} that contains no collisions after time t such that the spacetime position $L_{p_k}^t$ has label $\beta_k \in \mathcal{F}$. Let $\mathcal{L}$ be a collection of m_3 lines not containing any of the L_{p_k}. Let $L_k \notin \mathcal{L}$ be unoccupied and mutually disjoint lines that pass through some spacetime coordinates $(j'_k, t') \in \mathcal{C}(L_{p_k}^{(t)}, t)$ with $t' > t$. Denote $\mathcal{T} = \{L_1, \ldots, L_{m_4}\}$, and let $I \subset \mathbb{Z}$ be an interval of length $O(m_3)$. If $t' = t + \Omega(m_1 + m_3 + m_4(m_2 + m_4))$ is large enough, then there is an $(O(m_4), O(m_4(m_2 + m_4)), \mathcal{L}, \mathcal{T}, t', I)$-controlled modification after which each spacetime position $L_k^{t'}$ has label β_k. The same holds if the lines L_k are unoccupied only in the time interval $(-\infty, t']$, but the modification is weakly controlled.*

6.3 Computing with the Boolean Particles

Next, we will do some computation with the Boolean particles. Namely, we show that the NAND of two Boolean particles can be computed nondestructively in any spacetime position, as long as we have enough time, and the target line is in the intersection of the cones rooted at the input particles.

Lemma 4. *Suppose we have a spacetime diagram of interest with m_1 crossings and m_2 occupied lines, and two distinct lines L_1 and L_2 of direction 1 containing no collisions after time t such that the spacetime coordinates L_1^t, L_2^t have labels $\beta_1, \beta_2 \in \mathcal{F}$. Let $\mathcal{L}$ be a set of m_3 lines not containing L_1 and L_2, and let $L \notin \mathcal{L}$ be an unoccupied line of slope 1 to the left of L_1 and L_2 that passes through some spacetime coordinate $(j', t') \in \mathcal{C}(L_1^{(t)}, t) \cap \mathcal{C}(L_2^{(t)}, t)$ with $t' > t$. Let $\mathcal{T} = \{L_1, L_2, L\}$. If $t' = t + \Omega(m_1 + m_2 + m_3)$ is large enough, then there is an $(O(1), O(m_2), \mathcal{L}, \mathcal{T}, t', \emptyset)$-controlled modification after which the spacetime positions $L_1^{t'}$, $L_2^{t'}$ and $L^{t'}$ have labels β_1, β_2 and $\neg(\beta_1 \wedge \beta_2)$, respectively.*

Similarly to the fact that any number of particles can be moved, we can now compute an arbitrary Boolean function, given enough time and space.

Corollary 2. *Suppose we have a spacetime diagram of interest with m_1 collisions and m_2 occupied lines. Let $j \in \mathbb{Z}$, and for every $k = 1, \dots, m_4$, let $\beta_k \in \mathcal{F}$ be the label of the spacetime position $p_k = (j + k - 1, t, 1)$. Suppose further that the lines L_{p_k} passing through the p_k do not contain collisions after time t. Let $H : \{0,1\}^{m_4} \to \{0,1\}^m$ with $m = O(m_4)$ be a Boolean function realizable with C NAND-gates. Let $\mathcal{L}$ be a set of m_3 lines not containing the L_{p_k}, let $L_1, \dots, L_{C+m} \notin \mathcal{L}$ be unoccupied lines of slope 1 that pass through the spacetime coordinates $(j-1, t), \dots, (j - C - m, t)$, and let $t' > t$. Let also $\mathcal{T} = \{L_C, \dots, L_{C+m}\}$. If*

$$t' = t + \Omega((C + m_4)(C^2 + m_1 + m_3 + (C + m_4)(m_2 + m_4)))$$

is large enough, there is an $(O(C + m_4), O((C + m_2 + m_4)(C + m_4)), \mathcal{L}, \mathcal{T}, t', \emptyset)$-controlled modification after which the spacetime positions $L_{C+\ell}^{t'}$ have labels $H(\beta_1, \dots, \beta_m)_\ell$.

7 Physical Universality

Combining the lemmas of the previous section, we obtain our main result, the efficient physical universality of f.

Theorem 1. *The cellular automaton f is efficiently physically universal.*

8 A reversibly physically universal CA

In [8], other open questions were also posed, and here we answer one of them.

Call a CA $f : A^{\mathbb{Z}^d} \to A^{\mathbb{Z}^d}$ *reversibly physically universal*, if for any finite domain $D \subset \mathbb{Z}^d$ and any bijection $\theta : A^D \to A^D$, there exist partial configurations $x, y \in A^{\mathbb{Z}^d \setminus D}$ and a time t such that for all patterns $P \in A^D$ we have

$f^t(x \cup P) = y \cup \theta(P)$. In other words, the input cannot permanently affect the gadget during the computation.

The intuition behind this alternative notion is that after the CA has produced the pattern $\theta(P)$, the state of the 'computational machinery' should not depend on P: the entropy present in P does not permanently leak out of the domain D. It was conjectured by Schaeffer in [8] that a two-dimensional reversibly universal CA exists; we construct a one-dimensional one.

Theorem 2. *There exists a one-dimensional reversibly physically universal CA.*

9 Final remarks

Our proof of the physical universality of f can be turned into a polynomial time algorithm that, given a circuit computing the function $h : A^D \to A^E$ in the definition of physical universality, computes the corresponding gadget and the polynomially bounded number t. The main issue with this is that occupied lines and crossings cannot be enumerated in polynomial time unless $P = NP$, since it requires checking whether an arbitrary NAND-expression is satisfiable. This issue can be avoided by keeping track of a set of lines $\mathcal{L}_{occ+}$ that contains at least all occupied lines, as well as a set of coordinates $X \subset \mathbb{Z}^2$ that contains at least all crossings. When invoking Lemma 3 and Lemma 4 to modify the configuration of interest, we add to the protected set $\mathcal{L}$ all lines in $\mathcal{L}_{occ+}$ and all lines that pass through a coordinate in X. The set $\mathcal{L}_{occ+}$ is then amended with those lines that 'locally' appear to be occupied based on the added particles, and similarly for X. It can be verified that this approach does not change the asymptotic bounds in the construction, so the resulting configuration is polynomial in size.

The above algorithm will need polynomial space, as it compares the new positions of auxiliary particles to all existing ones. To construct the gadget in logarithmic space, it might be necessary to fix particular choices of where the auxiliary particles are put. We have chosen the more abstract route in the hope that our methods generalize more directly to a larger class of CA.

The existence of a physically universal CA was asked in [4] without fixing the number of states. Our CA has 16 states and radius 2. It would be interesting to find the minimal number of states and the minimal radii for physically universal CA. Of course, one can make any physically universal CA have radius 1 by passing to a blocking presentation, but this increases the number of states. From our CA, one obtains a physically universal radius-1 CA with 256 states.

Question 1. Are there physically universal CA on two states? Which combinations of state set and radius allow physical universality?

A long list of open questions about physical universality is also given in [8]. Finally, it would be interesting to explore the connections between physical universality and the other types of universalities mentioned in Section 1. For example, a reversible cellular automaton is not intrinsically universal, since it cannot simulate a non-reversible automaton, but could a physically universal CA be able to simulate all reversible automata?

Acknowledgments

We are thankful to Charalampos Zinoviadis for introducing this problem to us, and for many fruitful discussions on the proof, and Luke Schaeffer for his Golly implementation of our physically universal CA. We would also like to thank Scott Aaronson for popularizing the concept in his blog [1].

References

1. Scott Aaronson. Shtetl-Optimized – the blog of Scott Aaronson. `http://www.scottaaronson.com/blog/?p=1896`. Accessed: 2014-09-17.
2. Matthew Cook. Universality in elementary cellular automata. *Complex Systems*, 15(1):1–40, 2004.
3. Jean-Charles Delvenne, Petr Kurka, and Vincent D. Blondel. Decidability and universality in symbolic dynamical systems. *Fundam. Inform.*, 74(4):463–490, 2006.
4. D. Janzing. Is there a physically universal cellular automaton or Hamiltonian? *ArXiv e-prints*, September 2010.
5. Jarkko Kari. Universal pattern generation by cellular automata. *Theoretical Computer Science*, 429(0):180 – 184, 2012. Magic in Science.
6. Turlough Neary and Damien Woods. P-completeness of cellular automaton Rule 110. In *Automata, languages and programming. Part I*, volume 4051 of *Lecture Notes in Comput. Sci.*, pages 132–143. Springer, Berlin, 2006.
7. Nicolas Ollinger. The intrinsic universality problem of one-dimensional cellular automata. In Helmut Alt and Michel Habib, editors, *STACS 2003*, volume 2607 of *Lecture Notes in Computer Science*, pages 632–641. Springer Berlin Heidelberg, 2003.
8. Luke Schaeffer. A physically universal cellular automaton. *Electronic Colloquium on Computational Complexity (ECCC)*, 21:84, 2014.
9. Luke Schaeffer. A physically universal cellular automaton. In *ITCS'15—Proceedings of the 6th Innovations in Theoretical Computer Science*, pages 237–246. ACM, New York, 2015.
10. Luke Schaeffer. A physically universal quantum cellular automaton. In *Cellular automata and discrete complex systems*, volume 9099 of *Lecture Notes in Comput. Sci.*, pages 46–58. Springer, Heidelberg, 2015.
11. Tommaso Toffoli. Reversible computing. In Jaco de Bakker and Jan van Leeuwen, editors, *Automata, Languages and Programming*, volume 85 of *Lecture Notes in Computer Science*, pages 632–644. Springer Berlin Heidelberg, 1980.
12. John von Neumann. *Theory of Self-Reproducing Automata.* University of Illinois Press, Champaign, IL, USA, 1966.