The Most General Conservation Law for a Cellular Automaton

Enrico Formenti[1⋆], Jarkko Kari[2⋆⋆], and Siamak Taati[3,2]

[1] Université de Nice-Sophia Antipolis, Départment d'Informatique,
Parc Valrose, 06108 Nice Cedex 2, France
[2] Department of Mathematics, University of Turku,
FI-20014 Turku, Finland
[3] Turku Centre for Computer Science, FI-20520 Turku, Finland

Abstract. We study the group-valued and semigroup-valued conservation laws in cellular automata (CA). We provide examples to distinguish between semigroup-valued, group-valued and real-valued conservation laws. We prove that, even in one-dimensional case, it is undecidable if a CA has any non-trivial conservation law of each type. For a fixed range, each CA has a *most general* (group-valued or semigroup-valued) conservation law, encapsulating all conservation laws of that range. For one-dimensional CA the semigroup corresponding to such a most general conservation law has an effectively constructible finite presentation, while for higher-dimensional ones no such effective construction exists.

1 Introduction

Conservation laws in physics are numerical invariants of the dynamics of a system. In cellular automata, a similar concept has already been defined and studied (see e.g. [8, 3, 14, 5, 6]). We first choose a finite window, through which we can recognize the pattern made by the states of a finite number of cells on the lattice. We associate a real value to each possible local pattern that may be seen through this window, resembling the "energy" (or "mass", or ...) of that pattern. The total "energy" of a configuration is obtained by sliding this window all over the lattice and summing up the energy values of the local patterns we see. We have a conservation law for that energy provided the evolution of the CA preserves the total energy of each configuration.

In physics, conservation laws are used to write equations about the dynamics of the system. Each conserved quantity extracts certain information about the dynamics. In many cases, different conservation laws extract enough information to allow the reconstruction of the whole dynamics. In other cases, conservation laws concretize the physicist's insight into the behavior of the system by refuting those sequences of events that do not respect their preservation.

We study what happens if instead of mere real numbers, we allow energy valuations from a commutative group or semigroup. A remarkable (but trivial)

⋆ This work has been supported by the French ANR Blanc Projet "Sycomore".
⋆⋆ Research supported by the Academy of Finland grant 211967.

fact is that for each CA and a fixed window, there is a *most general* conservation law that extracts whatever information can be expressed in terms of conservation laws using that window. Any other more specific conservation law using that window can be derived from that by applying an algebraic homomorphism.

We provide examples that the group-valued conservation laws give strictly more information than the real-valued ones, and examples in which the semigroup-valued conservation laws are strictly more general than the group-valued ones. Needless to say, the semigroup-valued conservation laws can be quite expressive. Nevertheless, we prove that for one-dimensional CA, the most general conservation law of each range, and a finite presentation of the corresponding semigroup can be effectively constructed. This is a good news, because the Word Problem for commutative semigroups is also decidable (see e.g. [1]). Therefore the whole theory, in one-dimensional case, turns out to be algorithmically effective! For example we can effectively determine whether two (finite) configurations have the same total energy, or if a given CA conserves a given energy valuation. In higher dimensions, however, no such construction for the most general semigroup-valued conservation laws is possible.

Increasing the size of our window, we obtain more and more general (more and more discriminating) conservation laws. We may ask if there is a large enough window that provides the *absolutely most general* conservation law amongst all. This we still cannot answer adequately. However, we show that there are CA whose non-trivial conservation laws require very large windows. In fact, we prove that it is undecidable (even in one-dimensional case) whether a CA has any non-trivial conservation law at all, answering a formerly open question (see [10]). This is valid, no matter we are looking for real-valued, group-valued, or semigroup-valued conservation laws.

We note that the concept of the most general conservation law is closely related to the Conway's tiling group [4, 16], when we look at the space-time diagram of CA as tilings of the plane. Unlike Conway's group, in this paper we restrict our study to *commutative* groups or semigroups. That is because non-commutative conservation laws in higher-than-one dimensional CA do not make much sense. Furthermore, in the non-commutative case, the Word Problem is undecidable even for groups, and we much prefer to stay in the algorithmic realm.

2 Preliminaries

A cellular automaton (CA) is a collection of identical *cells* arranged regularly on a *lattice* where a natural notion of *neighborhood* is present. Each cell is assigned a *state* from a finite number of possible states. The state of the cells are updated synchronously, in discrete time steps, according to a *local update rule*, which takes into account the current state of each cell and its neighbors.

The cells are often indexed by $\mathbb{Z}^d$ ($d \geq 1$), where we obtain a *d-dimensional* CA. The state set is a finite set S. An assignment $c : \mathbb{Z}^d \to S$ of states to the cells of the lattice is referred to as a *configuration* (of the lattice). For each state

$q \in S$, a *q-uniform* configuration is a configuration with all cells in state q, and a *q-finite* configuration is one in which all cells but finitely many of them are in the state q. The set of all q-finite configurations is denoted by $C_q[S]$. A *pattern* over a finite set $A \subseteq \mathbb{Z}^d$ is an assignment $p : A \to S$. We use the notation $g[X]$ for the restriction of a mapping g to a subset X of its domain. Therefore, for example, $c[A]$ denotes the pattern seen over $A \subseteq \mathbb{Z}^d$ in the configuration $c \in S^{\mathbb{Z}^d}$.

The neighborhood is specified by a finite set $N \subseteq \mathbb{Z}^d$. The neighborhood of a cell $i \in \mathbb{Z}^d$ is the set $i + N = \{i + a : a \in N\}$. The local update rule is a function $f : S^N \to S$. The local rule f naturally induces a mapping $F : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$, called the *global mapping*, that maps each configuration c, to its follower configuration $F(c)$, which when starting from c, appears on the lattice after one time-step. Namely, $F(c)[i] \triangleq f(c[i + N])$; i.e., the state of the cell i in $F(c)$ is the result of the application of the local rule on the pattern of the neighborhood of i in c. We often identify a CA with its global mapping. A *quiescent* state is a state $q \in S$ such that F maps the q-uniform configuration to itself; i.e., $f(q^N) = q$. If q is a quiescent state, the image of every q-finite configuration is also q-finite.

For any $a \in \mathbb{Z}^d$, the *translation by a* is the operator $\sigma^a : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ defined by $(\sigma^a c)[i] \triangleq c[a + i]$. Notice that σ^a is a CA with neighborhood $\{a\}$. When $d = 1$, we may write σ for σ^1.

Let $A \subseteq \mathbb{Z}^d$ be a finite set. The *A-block-presentation* of a configuration $c \in S^{\mathbb{Z}^d}$ is a configuration $e \in \left(S^A\right)^{\mathbb{Z}^d}$ where $e[i] = c[i + A]$. That is, the state of the cell i in e is the overall state of the cells $i + A$ in c.

One-dimensional CA have a natural representation (upto translations) using edge-labeled De Bruijn graphs. The *De Bruijn graph* of order k ($k > 0$) over an alphabet S, is a graph $B_k[S]$ with vertex set $V = S^k$ and edge set $E = S^{k+1}$, where for any $a, b \in S$ and $u \in S^{k-1}$, there is an edge aub from au to ub.

Let $F : S^{\mathbb{Z}} \to S^{\mathbb{Z}}$ be a one-dimensional CA with neighborhood $[-l, r] = \{-l, -l + 1, \ldots, r\}$ and local rule $f : S^{[-l,r]} \to S$. For any $k \geq l + r$, the CA can be represented on the De Bruijn graph $B_k[S]$ with labeling $\lambda : E \to S^{k-(l+r)}$ which is defined as follows. For every edge $u_0 u_1 \cdots u_k \in S^{k+1}$, let $\lambda(u_0 u_1 \cdots u_k) = v_l v_{l+1} \cdots v_{k-r}$ where $v_i = f(u[i - l, i + r])$. The edge sequence $p = \{p[i]\}_{i \in \mathbb{Z}}$ of each bi-infinite path on $B_k[S]$ is the $[0, k]$-block-presentation of a unique configuration $c \in S^{\mathbb{Z}}$, while its label sequence $\lambda(p) = \{\lambda(p[i])\}_{i \in \mathbb{Z}}$ is the $[l, k - r]$-block-presentation of $F(c)$. Conversely, for every configuration $c \in S^{\mathbb{Z}}$, there is a unique infinite path on $B_k[S]$ whose edge sequence is the $[0, k]$-block-presentation of c.

A *counter machine* is a finite automaton equipped with two unbounded *counters*. The machine can increase or decrease the value of each counter, and can test if either has value zero. Counter machines are known to be equivalent in power with Turing machines — any algorithm can be implemented on a counter machine (see e.g. [11]).

3 Universal Conservation Law

Let $F : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ be a CA, and to avoid cumbersome technicalities, let us assume that the CA has a quiescent state q. Let Φ be a commutative (additive) semigroup, and $W \subseteq \mathbb{Z}^d$ a finite set, and assume that to each pattern $p \in S^W$ we have associated a value $\mu(p) \in \Phi$ as its *energy*. We would like to define the total energy $M(c)$ of a configuration $c \in S^{\mathbb{Z}^d}$ as the sum

$$\sum_{i \in \mathbb{Z}^d} \mu(c[i + W]) \tag{1}$$

by sliding the *window* W over c and adding up all the local energy values we see. However, there is no uniform way to interpret an infinite sum over a semigroup! To overcome this, we choose the following approach: we assume that Φ contains an identity element 0, and that $\mu(q^W) = 0$. Then the above sum will have a natural meaning for all q-finite configurations — only finite number of the terms have values other than 0. For all other configurations we leave the total energy undefined.[4]

We say that the CA *conserves* the energy valuation μ, if

$$M(F(c)) = M(c) \tag{2}$$

for every q-finite configuration c. Then M is a Φ-valued conserved energy defined using window W, and (2) is a *conservation law* for the CA. More concisely, the conservation law can be identified by the pair (Φ, μ).

Notice that the values of M do not depend on the displacement of the window. That is, for every translation σ^a, $M \circ \sigma^a = M$. The conservation of M by a CA F means further that its value is constant over the orbits of F; i.e., $M \circ F = M$. Therefore, the kernel of M is a partition of the orbits of F. The finer this partition, the more information the conservation law extracts about the dynamics of the CA.

An energy function which assigns the same value to every finite configuration is trivially conserved by every CA. We call such a conservation law *trivial*. Note that the total energy mapping $M : C_q[S] \to \Phi$ of a conservation law is not necessarily onto (even if the local energy function μ is onto). The uncovered part of Φ bears no information about the dynamics of F. Hence it is convenient to name the *realizable* sub-monoid $\breve{\Phi} \triangleq M(C_q[S])$ of Φ. So (Φ, μ) is trivial, if and only if, the realizable sub-monoid $\breve{\Phi}$ is trivial.

Let us now fix a CA $F : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ with a quiescent state q, and a finite window $W \subseteq \mathbb{Z}^d$. Every conserved energy valuation μ for F satisfies (2) for every q-finite configuration c. The key observation in this paper is that the largest semigroup generated by the (formal) values of μ for which (2) holds for every q-finite configuration c provides the *most general* conservation law for F defined using window W.

[4] This is not the only possible approach, but it sounds most natural to us.

Put it precisely, let $\Sigma \triangleq S^W - \{q^W\}$ be the set of non-quiescent patterns on W, and let us denote by $\mathbb{N}^\Sigma$ the free commutative (additive) monoid generated by Σ. Define the energy assignment $\mu_\circ : S^W \to \mathbb{N}^\Sigma$ with $\mu_\circ(q^W) = 0$, and $\mu_\circ(p) = p$ for any other pattern p over W, and let $M_\circ : C_q[S] \to \mathbb{N}^\Sigma$ be the corresponding total energy. Let $\cong \subseteq \mathbb{N}^\Sigma \times \mathbb{N}^\Sigma$ be the coarsest monoid congruence where $M_\circ(F(c)) \cong M_\circ(c)$ for every q-finite c. Define $\Phi_F \triangleq \mathbb{N}^\Sigma/\cong$ and let $h_\cong : \mathbb{N}^\Sigma \to \Phi_F$ be the natural homomorphism. Define $\mu_F : W \to \Phi_F$ with $\mu_F = h_\cong \circ \mu_\circ$. Clearly the pair (Φ_F, μ_F) identifies a conservation law, because

$$M_F(F(c)) = h_\cong(M_\circ(F(c))) = h_\cong(M_\circ(c)) = M_F(c) \tag{3}$$

Furthermore, for every conservation law (Φ, μ) with window W, there is a monoid homomorphism $h : \Phi_F \to \Phi$ so that $\mu = h \circ \mu_F$; schematically,

$$\begin{array}{ccc} (S^{\mathbb{Z}}, F) & \xrightarrow{M_F} & \Phi_F \\ & \searrow^{M} & \downarrow h \\ & & \Phi \end{array} \tag{4}$$

One can verify that using $\breve{\Phi}_F$ instead of Φ_F, the choice of h would be unique.

It is easy to see that when $W' \subseteq W \subseteq \mathbb{Z}^d$, the most general conservation law based on window W' is a factor of the most general conservation law based on window W; i.e., $(\Phi_F^{(W)}, \mu_F^{(W)})$ is *more general* than $(\Phi_F^{(W')}, \mu_F^{(W')})$. That is because every energy valuation on W' can be seen also as a energy valuation on W, by adding some *dummy* elements to W'.

Following a similar trail of reasoning, if instead of semigroup-valued energies we consider group-valued energies, we can define the *most general* group-valued conservation law (G_F, ρ_F) based on W for F. Likewise, we define the realizable subgroup $\breve{G}_F \triangleq P_F(C_q[S])$ of G_F, where P_F is the total energy mapping corresponding to ρ. The conservation law $(\breve{G}_F, \rho_F)$ satisfies a similar universal property (4) among group-valued conservation laws.

Example 1 (Spreading 1's). Consider the one-dimensional CA F with state set $\{0, 1\}$, neighborhood $\{-1, 0, 1\}$, and local rule $f(a, b, c) = a \vee b \vee c$; see Fig. 1 for a typical snapshot. The time axis in the figure goes downward. It is easy to see that every group-valued conservation law for F is trivial. Notice that every non-quiescent finite configuration eventually turns into a single ever-growing block of ones. In contrast, F has a non-trivial semigroup-valued conservation law. Let $\Phi = \{0, 1\}$ be the commutative semigroup with binary operation $a + b \triangleq a \vee b$. Let $W = \{0\}$ be the singleton window, and $\mu : \{0, 1\} \to \Phi$ be the identity. Under this energy valuation, the 0-uniform has total energy 0, while every other 0-finite configuration has total energy 1.

Example 2 (XOR). The second example is again a one-dimensional CA F' with binary state set, and neighborhood $\{-1, 0, 1\}$. The local rule is the XOR rule

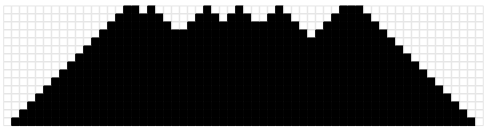

Fig. 1. A space-time snapshot from the CA in Ex. 1.

$f'(a,b,c) = a+b+c \pmod 2$. Figure 2 shows a typical snapshot. A well-known property of the XOR CA (and in general every linear CA) is its replicating behavior. Specifically, every finite pattern, after a finite number of steps, is replicated into three copies with large 0 blocks in between (Fig. 2 depicts an example. This is easy to verify using generating functions; see e.g. [15]). This implies that F' cannot have any non-trivial real-valued conservation law. On the other hand, G preserves the parity of the configurations. Let $G = \mathbb{Z}_2$ be the binary cyclic group, and consider the identity energy function $\rho : \{0,1\} \to \mathbb{Z}_2$ on window $\{0\}$. The total energy $P(c)$ is simply the parity of the number of 1's in c, and is preserved by F'.

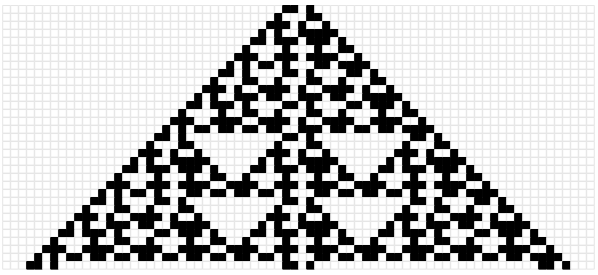

Fig. 2. A space-time snapshot from the CA in Ex. 2.

4 Semigroup-valued Conservation Laws

The definition of the most general conservation law of certain range, given above, is based on an infinite presentation of the corresponding semigroup. A standard theorem from the theory of semigroups states that any finitely generated commutative semigroup has a finite presentation (see e.g. [7]). A question arises that how one can find such a finite presentation. A finite presentation is needed if, for example, we want to algorithmically verify whether two configurations have the same total energy. It turns out there is no algorithm to construct such a finite presentation for the semigroup of the most general conservation law in

2- or higher-dimensional CA. In one-dimensional case, we can construct these semigroups effectively.

Let $F : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ be a CA with a quiescent state $q \in S$. Clearly, (q, q) is a quiescent state for the product $F \times F$. Let $\Phi = \{0, 1\}$ be the Boolean semigroup with $a + b \triangleq a \vee b$, for any $a, b \in \Phi$. Define the range-$\{0\}$ energy valuation $\mu : S \times S \to \Phi$ with

$$\mu(a, b) = \begin{cases} 0 \text{ if } a = b, \\ 1 \text{ otherwise.} \end{cases} \tag{5}$$

The energy μ is conserved by $F \times F$, if and only if, F is injective on finite configurations. According to the Garden-of-Eden theorem [12, 13], F is injective on finite configurations if and only if it is surjective. However, when $d \geq 2$, it is undecidable whether a given d-dimensional CA is surjective [9]. Therefore no algorithm could verify, for a given F, whether $F \times F$ conserves μ.

Theorem 1. *There is no algorithm, that given a 2- or higher-dimensional CA F with state set S, and an energy valuation $\mu : S \to \Phi$ from a finitely presented commutative semigroup Φ, determines if F conserves μ.*

Corollary 1. *There is no algorithm, that given a 2- or higher-dimensional CA F, computes a finite presentation of the semigroup Φ_F and the energy valuation μ_F of the most general conservation law for F with window $\{0\}$.*

Let us now focus on one-dimensional CA. Let $F : S^{\mathbb{Z}} \to S^{\mathbb{Z}}$ be a 1d CA. Without loss of generality we assume that F has a neighborhood $[-l, r] = \{-l, -l + 1, \ldots, r\}$ with $l + r \geq 0$. Let q be the designated quiescent state of F. Let $W \subseteq \mathbb{Z}$ be a finite set, and $\mu : S^W \to \Phi$ be an energy valuation on window W with values from a commutative monoid Φ. Again without loss of generality we assume that $W = [0, m) = \{0, 1, \ldots, m - 1\}$.

For $k = l+r+m$, consider the k'th order De Bruijn representation $(B_k[S], \lambda)$ of F. This has a vertex q^k, with a loop edge q^{k+1} which is labeled by q^m. Any path corresponding to a q-finite configuration starts by circulating in this loop, and after possibly passing through a finite number of other edges, eventually returns back to this loop.

To each edge $u_0u_1 \cdots u_k \in S^{k+1}$ let us assign two elements

$$\alpha(u_0u_1 \cdots u_k) \triangleq \mu(u_0u_1 \cdots u_{m-1}) \tag{6}$$

and

$$\beta(u_0u_1 \cdots u_k) \triangleq \mu(v_lv_{l+1} \cdots v_{l+m-1}) \tag{7}$$

from Φ_F, where $v_lv_{l+1} \cdots v_{k-r} = \lambda(u_0u_1 \cdots u_k)$ is the label of $u_0u_1 \cdots u_k$. The total energy of a q-finite configuration x can be calculated by adding up the values of α over the edges of the corresponding bi-infinite path on $B_k[S]$. Likewise, the sum of β values on this path gives the total energy of $F(x)$. Note that the initial and final parts of such a path, where it is circulating in the loop q^{k+1} do not contribute to the total energy, because $\mu(q^m) = 0$. For any path $p = p_1p_2 \cdots p_n$

(p_i is the i'th edge of the path), let us use the notation $\alpha(p)$ for the sum of the values of α over the edges of p; i.e.,

$$\alpha(p) \triangleq \sum_{i=1}^{n} \alpha(p_i) \tag{8}$$

and similarly for β.

The requirements imposed by the conservation of μ can now be translated in terms of the values of α and β over finite paths on the graph $B_k[S]$: The pair (Φ, μ) specifies a conservation law, if and only if, for any finite path p starting and ending at vertex q^k, $\alpha(p) = \beta(p)$.

We claim that (the proof can be found in the Appendix):

Proposition 1. *Let G be a (finite, directed) graph with vertex set V and edge set E, and Δ a finite symbol set. Let $\alpha, \beta : E \to \Delta$ and $A, B \subseteq V$. Let Φ be the largest commutative monoid generated by Δ, satisfying the equations*

$$\alpha(p) = \beta(p) \tag{9}$$

for any finite path p starting from A and ending at B. Then, there is an algorithmically constructible finite subset of the above equations, such that any commutative monoid generated by Δ satisfying those equations is a factor of Φ.

From Prop. 1, we immediately obtain what we were after in this section:

Theorem 2. *For any one-dimensional CA F and any finite window $W \subseteq \mathbb{Z}$, the semigroup Φ_F of the most general conservation law for F based on W is effectively finitely presentable.*

This does not say much about the realizable sub-monoid $\breve{\Phi}_F \subseteq \Phi_F$. For example, it is not even clear if $\breve{\Phi}_F$ is finitely generated or not. Following a similar construction, however, one can decide whether $\breve{\Phi}_F$ is trivial or not.

Proposition 2. *Let F be a one-dimensional CA, and (Φ, μ) a semigroup-valued conservation law for F. It is decidable whether $\breve{\Phi}$ is trivial.*

5 Group-valued Conservation Laws

Group-valued conservation laws are more easily tractable. There is a simple algorithm which tests whether a given CA F (of any dimension) conserves a given group-valued energy function ρ. (This is similar to the real-valued case; see e.g. [10]). In particular, for any fixed window W one can effectively construct the most general group-valued conservation law $(G_F^{(W)}, \rho_F^{(W)})$ based on W. The next challenge would be to classify all conservation laws, based on all windows. For example, given a CA, is it possible to decide if it has any conservation law at all? In this section we prove that the answer to the latter question is negative.

Lemma 1. *Let $F : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ be a 1d CA with a designated quiescent state q. Suppose that F is nilpotent over q-finite configurations. Then F does not have any non-trivial (real-valued/group-valued/semigroup-valued) conservation law.*

Theorem 3 (Blondel, Cassaigne and Nichitiu [2]). *Given a counter machine A with 2 counters and no halting state, it is undecidable whether A has a periodic orbit.*

Theorem 4. *There is no algorithm that given a one-dimensional cellular automaton F determines if F has a non-trivial (real-valued/group-valued/semigroup-valued) conservation law.*

Sketch of the proof. We show how to reduce the problem of whether a given counter machine has a periodic orbit to finding out if a 1d CA has any non-trivial conservation law. Since the former is undecidable, we conclude that so is the latter.

Let A be a counter machine with state set Q, two registers x_1 and x_2, and transition function $\delta : Q \times \{0,1\}^2 \to Q \times \{1,2\} \times \{-,0,+\}$. We construct a CA F with a designated quiescent state q such that

a) if A has a periodic configuration, F has a non-trivial (real-valued) conservation law, while
b) if A has no periodic orbit, F is nilpotent over q-finite configurations (hence, has no non-trivial conservation law).

The CA F has two states L and R which are *end-markers*. In the interval between a left end-marker L and a right end-marker R, the CA simulates the machine A. The CA also constantly verifies the syntax of the block between two end-markers, to make sure it corresponds with an actual simulation. If a syntax error is found, or if the simulation overflows the end-markers, the CA erases the whole block, by replacing the cell contents with q. Blocks not having end-markers are also erased.

If the machine A has no periodic orbit, every syntactically correct simulation of A on a finite block eventually overflows the boundaries. Therefore, every q-finite configuration eventually goes quiescent; i.e., F is nilpotent over q-finite configurations.

On the other hand, if A has a periodic configuration, one can choose a sufficiently large simulation block in F which evolves periodically, and never overflows. Let us fix a snapshot of such a periodic simulation block (including its end-markers), and denote it by B. Since no new end-marker is ever created by F, and since the end-markers block the flow of information inwards and outwards the simulation blocks, we can argue that if a simulation block in a configuration eventually turns into B, it does so independent of the rest of the configuration, and after a bounded number of steps. Let us denote by $\mathcal{B}$ the set of all simulation blocks that eventually turn into B. This is a stable set. Once we know that a block from $\mathcal{B}$ occurs in a certain position on a configuration c, we also know that a block from $\mathcal{B}$ occurs in the same position on any pre-image of c. This immediately gives a non-trivial (real-valued) conservation law that states that the number of occurrences of blocks from $\mathcal{B}$ is conserved by F. □

6 Conclusion and Open Problems

In this paper we examined a number of algorithmic problems that arise from studying the algebraic conservation laws for cellular automata. The semigroup-valued conservation laws are highly expressive, still not so tractable. The group-valued conservation laws are more expressive than the real-valued ones, yet as accessible as they are.

Since reversible CA are particularly attractive for modeling physical processes, it would be useful to examine the same problems in the restricted case of reversible CA. In particular, is it decidable whether a given *reversible* CA has any conservation law? Notice that our proof of Theorem 4 takes advantage of the existence of very long transients to construct CA whose conservation laws need very large windows. We conjecture that every reversible CA has a non-trivial conservation law with a relatively small window.

References

[1] A. P. Biryukov, *Some algorithmic problems for finitely defined commutative semigroups*, Siberian Mathematical Journal, 8 (1967), pp. 384–391.

[2] V. D. Blondel, J. Cassaigne, and C. Nichitiu, *On the presence of periodic configurations in Turing machines and in counter machines*, Theoretical Computer Science, 289 (2002), pp. 573–590.

[3] N. Boccara and H. Fukś, *Number-conserving cellular automaton rules*, Fundamenta Informaticae, 52 (2002), pp. 1–13.

[4] J. H. Conway and J. C. Lagarias, *Tiling with polyominoes and combinatorial group theory*, Journal of Combinatorial Theory A, 53 (1990), pp. 183–208.

[5] B. Durand, E. Formenti, and Z. Róka, *Number conserving cellular automata I: decidability*, Theoretical Computer Science, 299 (2003), pp. 523–535.

[6] E. Formenti and A. Grange, *Number conserving cellular automata II: dynamics*, Theoretical Computer Science, 304 (2003), pp. 269–290.

[7] P. A. Grillet, *Semigroups: An Introduction to the Structure Theory*, Dekker, 1995.

[8] T. Hattori and S. Takesue, *Additive conserved quantities in discrete-time lattice dynamical systems*, Physica D, 49 (1991), pp. 295–322.

[9] J. Kari, *Reversibility and surjectivity problems of cellular automata*, Journal of Computer and System Sciences, 48 (1994), pp. 149–182.

[10] ———, *Theory of cellular automata: A survey*, Theoretical Computer Science, 334 (2005), pp. 3–33.

[11] M. Minsky, *Computation: Finite and Infinite Machines*, Prentice-Hall, 1967.

[12] E. F. Moore, *Machine models of self-reproduction*, in Proceedings of Symposia in Applied Mathematics, AMS, 1962, pp. 17–33.

[13] J. Myhill, *The converse of Moore's Garden-of-Eden theorem*, Proceedings of the American Mathematical Society, 14 (1963), pp. 685–686.

[14] M. Pivato, *Conservation laws in cellular automata*, Nonlinearity, 15 (2002), pp. 1781–1793.

[15] A. D. Robison, *Fast computation of additive cellular automata*, Complex Systems, 1 (1987), pp. 211–216.

[16] W. P. Thurston, *Conway's tiling groups*, American Mathematical Monthly, 97 (1990), pp. 757–773.

A Appendix

A.1 Proposition 1

Proof. We start by introducing the finite subset in question. For any vertex $v \in V$, define the following three sets:

P_v: The set of all *simple* paths starting from A and ending at v.
Q_v: The set of all *simple* paths starting from v and ending at B.
C_v: The set of all *simple* cycles (including the empty one) passing through v.

For any v, the set $P_v C_v Q_v$ is finite, because each of P_v, C_v and Q_v is finite. Its elements are paths starting from A, passing through v (and possibly a cycle around v), and continuing further to end up at B. Define

$$R \triangleq \bigcup_{v \in V} P_v C_v Q_v \,.$$

We claim that if for some semigroup Φ' generated by Δ, the equation (9) holds for all paths $r \in R$, it also holds for any other path p from A to B.

The proof is by induction on the length of the path p. Note that any sufficiently short path p from A to B is simple and passes through a vertex like v. Therefore it is of the form xy where $x \in P_v$ and $y \in Q_v$. That is $p \in R$ and the equation (9) holds by the assumption.

Suppose that the equation (9) holds for all paths of length at most n, and let p be a path of length $n+1$ from A to B. If p is not simple, it contains at least one non-empty simple cycle, and hence can be written in the form xcy where x is a (not necessarily simple) path from A to a vertex v, c is a non-empty simple cycle starting and ending at v, and y is a (not necessarily simple) path from v to B. On the other hand x contains a subsequence $\tilde{x}$ which is a simple path from A to v, and we can write

$$\alpha(x) = \sum_i \alpha(x_i) = \gamma_x + \sum_i \alpha(\tilde{x}_i) = \gamma_x + \alpha(\tilde{x})$$

where $\gamma_x \in \Phi$ is the sum of α over those edges from x that are not in $\tilde{x}$. Similarly y contains a subsequence $\tilde{y}$ which is a simple path from v to B and we can write

$$\alpha(y) = \gamma_y + \alpha(\tilde{y})$$

for some $\gamma_y \in \Phi$. Writing the equation (9) for the paths $\tilde{x}\tilde{y}$ and $\tilde{x}c\tilde{y}$ (note that $\tilde{x}\tilde{y}, \tilde{x}c\tilde{y} \in R$) and xy (by induction hypothesis), we have

$$\alpha(\tilde{x}) + \alpha(\tilde{y}) = \beta(\tilde{x}) + \beta(\tilde{y}) \tag{10}$$
$$\alpha(\tilde{x}) + \alpha(c) + \alpha(\tilde{y}) = \beta(\tilde{x}) + \beta(c) + \beta(\tilde{y}) \tag{11}$$
$$\alpha(x) + \alpha(y) = \beta(x) + \beta(y) \tag{12}$$

from which we obtain that

$$\begin{aligned}\alpha(p) &= \alpha(x) + \alpha(c) + \alpha(y) \\ &= \gamma_x + \gamma_y + \alpha(\tilde{x}) + \alpha(c) + \alpha(\tilde{y}) \\ &= \gamma_x + \gamma_y + \beta(\tilde{x}) + \beta(c) + \beta(\tilde{y}) \\ &= \gamma_x + \gamma_y + \alpha(\tilde{x}) + \alpha(\tilde{y}) + \beta(c) \\ &= \alpha(x) + \alpha(y) + \beta(c) \\ &= \beta(x) + \beta(y) + \beta(c) \\ &= \beta(p)\end{aligned}$$

which is what we wanted to proof. Therefore, any equation satisfied by the semigroup Φ is also satisfied by Φ', and Φ' is a factor of Φ. □

A.2 Theorem 4

Construction. Let $E = \{L, R\}$, $X = \{0, 1\}$, $K = \{1, 2\}$ and $C = \{0, +, -, 1\}$. The state set of F is $S = \{q\} \cup E \cup Q \cup (X \times X \times K \times C)$, and its neighborhood is $N = \{-1, 0, 1\}$. Each simulation block starts with a left end-marker L, followed by an element from Q, representing the state of the counter machine, and ends with a right end-marker R. The space between the Q state and the right end-marker stores the two counters, and manages the required signaling. The first and the second components of $X \times X \times K \times C$ keep the unary value of the counters x_1 and x_2 in the form of stacks extending to the right. The K component corresponds the second component of the values of δ, indexing the counter to be increased or decreased. The C component carries a signal indicating whether the indexed counter should be increased (+), decreased (−), or left unchanged (0), and an acknowledgment signal (1) that comes back to the left to initiate the simulation of the next step. The local rule $f : S^3 \to S$ of the CA is presented in Table 1. A sample syntactically correct simulation block is depicted in Figure 3. □

$$\cdots\ q\ \ q\ \ q\ \ L\ \ x\ \begin{array}{|c|c|c|c|c|c|c|c|c|c|}\hline 0 & 0 & 0 & + & 1 & 1 & 1 & 1 & - & 1 \\ \hline 2 & 2 & 2 & 2 & 1 & 1 & 1 & 1 & 1 & 1 \\ \hline 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \hline 1 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ \hline\end{array}\ R\ \ q\ \ q\ \ q\ \cdots$$

Fig. 3. A syntactically correct simulation block in the CA described in the proof of Theorem 4. The simulated machine is in state $x \in Q$. The first counter contains 2, while the second counter has 6. There is a signal, moving to the right, commanding the second counter to increase.

a_{-1}	a_0	a_1	$f(a_{-1}a_0a_1)$	Condition
	q		q	
	L	x	L	$x \in Q$
	L	x	q	$x \notin Q$
$(0,0,k,c)$	R		R	$k \in K$ and $c \neq 0$
x	R		q	$x \notin \{0\} \times \{0\} \times K \times \{+,-,1\}$
L	x	$(b_1,b_2,k,1)$	x'	$x \in Q$ and $\delta(x,\neg b_1,\neg b_2)$ $= (x',k',c')$
L	x	(b_1,b_2,k,c)	x	$x \in Q$ and $c \neq 1$
y	x		q	$x \in Q$ but $y \neq L$
	x	R	q	$x \in Q$
x	$(b_1,b_2,k,1)$	y	(b_1,b_2,k',c')	$x \in Q$ and $y \notin \{q,L\}$ and $\delta(x,\neg b_1,\neg b_2)$ $= (x',k',c')$
(b_1',b_2',k',c')	$(b_1,b_2,k,1)$	y	(b_1,b_2,k',c')	$c' \neq 0$ and $y \notin \{q,L\}$
$(b_1',b_2',k',0)$	$(b_1,b_2,k,1)$	y	$(b_1,b_2,k,1)$	$y \notin \{q,L\}$
x	$(b_1,b_2,k,0)$	$(b_1',b_2',k',1)$	$(b_1,b_2,k,1)$	$x \notin \{q,L,R\}$
x	$(b_1,b_2,k,0)$	(b_1',b_2',k',c)	$(b_1,b_2,k,0)$	$c \neq 1$ and $x \notin \{q,L,R\}$
x	$(0,b_2,1,+)$	y	$(1,b_2,1,1)$	$x \in Q$ and $y \notin \{q,L\}$
(b_1',b_2',k',c')	$(0,b_2,1,+)$	y	$(b_1',b_2,1,1)$	$y \notin \{q,L\}$
x	$(1,b_2,1,+)$	y	$(1,b_2,1,0)$	$x \notin \{q,L,R\}$ and $y \notin \{q,L\}$
x	$(b_1,0,2,+)$	y	$(b_1,1,2,1)$	$x \in Q$ and $y \notin \{q,L\}$
(b_1',b_2',k',c')	$(b_1,0,2,+)$	y	$(b_1,b_2',2,1)$	$y \notin \{q,L\}$
x	$(b_1,1,2,+)$	y	$(b_1,1,2,0)$	$x \notin \{q,L,R\}$ and $y \notin \{q,L\}$
x	$(0,b_2,1,-)$	y	$(0,b_2,1,1)$	$x \notin \{q,L,R\}$ and $y \notin \{q,L\}$
x	$(1,b_2,1,-)$	$(0,b_2',k',c')$	$(0,b_2,1,1)$	$x \notin \{q,L,R\}$
x	$(1,b_2,1,-)$	$(1,b_2',k',c')$	$(1,b_2,1,0)$	$x \notin \{q,L,R\}$
x	$(b_1,0,1,-)$	y	$(b_1,0,1,1)$	$x \notin \{q,L,R\}$ and $y \notin \{q,L\}$
x	$(b_1,1,1,-)$	$(b_1',0,k',c')$	$(b_1,0,1,1)$	$x \notin \{q,L,R\}$
x	$(b_1,1,1,-)$	$(b_1',1,k',c')$	$(b_1,1,1,0)$	$x \notin \{q,L,R\}$
L	x		q	$x \neq Q$
	x	R	q	$x \notin \{0\} \times \{0\} \times K \times \{+,-,1\}$
y	x		q	$x \neq L$ and $y \in \{q,R\}$
	x	y	q	$x \neq R$ and $y \in \{q,L\}$

Table 1. The local rule of the CA described in the proof of Theorem 4.