Nat Comput
DOI 10.1007/s11047-010-9222-0

On the hierarchy of conservation laws in a cellular automaton

Enrico Formenti · Jarkko Kari · Siamak Taati

Abstract Conservation laws in cellular automata (CA) are studied as an abstraction of the conservation laws observed in nature. In addition to the usual real-valued conservation laws we also consider more general group-valued and semigroup-valued conservation laws. The (algebraic) conservation laws in a CA form a hierarchy, based on the range of the interactions they take into account. The conservation laws with smaller interaction ranges are the homomorphic images of those with larger interaction ranges, and for each specific range there is a most general law that incorporates all those with that range. For one-dimensional CA, such a most general conservation law has—even in the semigroup-valued case—an effectively constructible finite presentation, while for higher-dimensional CA such effective construction exists only in the group-valued case. It is even undecidable whether a given two-dimensional CA conserves a given semigroup-valued energy assignment. Although the local properties of this hierarchy are tractable in the one-dimensional case, its global properties turn out to be undecidable. In particular, we prove that it is undecidable whether this hierarchy is trivial or unbounded. We point out some interconnections between the structure of this hierarchy and the dynamical properties of the CA. In particular, we show that positively expansive CA do not have non-trivial real-valued conservation laws.

Keywords Cellular automata · Conservation laws · Energy · Reversibility · Undecidability · Dynamical systems · Chaos

E. Formenti · S. Taati
Laboratoire I3S, Université de Nice Sophia Antipolis, 2000 route des Lucioles, Les Algorithmes – Bât. Euclide B, BP 121, 06903 Sophia Antipolis Cedex, France

J. Kari (✉)
Department of Mathematics, University of Turku, 20014 Turku, Finland
e-mail: jkari@utu.fi

Published online: 05 November 2010

1 Introduction

Conservation laws in physics provide numerical invariants of the dynamics of a system. In cellular automata (CA), a similar concept has already been defined and studied (see e.g. Hattori and Takesue 1991; Boccara and Fukś 1998; Pivato 2002; Durand et al. 2003; Formenti and Grange 2003; Moreira et al. 2004). 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. Usually the CA has a designated non-active state and the "energy" of the uniformly non-active pattern is chosen to be zero. The total "energy" of a configuration is obtained by sliding the window all over the lattice and adding up the energy values of the local patterns we see. The total energy is not always meaningful, but it is so for any finite configuration (i.e., a configuration with only finitely many active cells). We have a conservation law for that energy provided the evolution of the CA preserves the total energy of each finite 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.

The study of conservation laws in cellular automata was initiated in connection with cellular automata models of physical phenomena (see e.g. Hardy et al. 1976; Pomeau 1984; Takesue 1987). Hattori and Takesue were first to consider such laws in a general abstract setting and address the problem of finding and verifying them. Various mathematical and algorithmic characterizations of conservation laws in cellular automata are investigated in details (see e.g. Hattori and Takesue 1991; Pivato 2002; Durand et al. 2003). Certain natural conservation laws are also used to define a rich class of cellular automata (see e.g. Boccara and Fukś 1998; Fukś 2000; Durand et al. 2003; Formenti and Grange 2003; Moreira et al. 2004). Local representation of conservation laws using flows and particles is particularly challenging (see e.g. Fukś 2000; Moreira et al. 2004; Pivato 2002).

We study what happens if instead of mere real numbers, we allow energy valuations from a commutative group or semigroup. A remarkable (but trivial) 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 the most general law 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. We also provide an example of a reversible one-dimensional CA that has no non-trivial real-valued conservation laws, refuting our conjecture in Formenti et al. (2008).

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. Biryukov 1967). 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 in the most general energy valuation of a given

range, 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. This hierarchy may be trivial, non-trivial but bounded or unbounded. In the bounded case there is a window that provides the *absolutely most general* conservation law amongst all. We prove that it is undecidable (even in the one-dimensional case) whether a CA has any non-trivial conservation law at all, or whether the hierarchy is unbounded. These properties are undecidable regardless of whether we consider the semigroup-, group- or real-valued conservation laws. In fact, there is no algorithm to separate the CA without any semigroup-valued conservation laws from those that have an unbounded hierarchy of real-valued conservation laws.

We note that the concept of the most general conservation law is closely related to Conway's tiling group (Conway and Lagarias 1990; Thurston 1990), 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.

The paper is organized as follows. We start with basic definitions related to cellular automata (Sect. 2.1) and conservation laws, including the definitions related to the most general semigroup- and group-valued conservation laws (Sect. 2.2). We continue with three examples in Sect. 3: first two examples show that moving from real numbers to more general groups, or from groups to semigroups, can improve the fidelity of conservation laws. The third example is a simple one-dimensional reversible CA that has no real-valued conservation laws of any interaction range.

In Sect. 4 we study algorithmic questions concerning semigroup-valued conservation laws of a fixed interaction range. It is shown (Theorem 1) that it is undecidable whether a given two-dimensional CA conserves a given energy valuation, even in the very restricted setup where the energy values are from a two-element semigroup and the values are assigned to cells without neighborhood interaction. In contrast, the one-dimensional case is effective: one can algorithmically construct the most general conservation law for any given interaction range (Theorem 2). In Sect. 5 we investigate algorithmic questions in the setup when the interaction window of the energy valuation is not fixed. In this case, it is undecidable if any non-trivial conservation laws exist, or whether their hierarchy is unbounded (Theorem 4). This undecidability holds in semigroup-, group- and real-valued cases. The main results of Sects. 4 and 5 were already reported in Formenti et al. (2008) without detailed proofs.

In Sect. 6 we show an example of what purely dynamical properties of a CA can tell about their conservation laws. We show that strong transitivity and positive expansivity of a CA guarantee that the CA can have no real-valued conservation laws (Theorem 5).

2 Preliminaries

2.1 Cellular automata

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* that 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 $s \in S$ and each configuration $c \in S^{\mathbb{Z}^d}$ we denote by

$$supp_s(c) \triangleq \left\{ i \in \mathbb{Z}^d | c[i] \neq s \right\}$$

the set of those cells where c does not take value s, and call it the *s-support* of c. An *s-uniform* configuration is a configuration whose s-support is empty, and an *s-finite* configuration is one whose s-support is finite. The set of all s-finite configurations is denoted by $C_s[S]$. For any two configurations c and e we denote by

$$diff(c, e) \triangleq \left\{ i \in \mathbb{Z}^d | c[i] \neq e[i] \right\}$$

the set of cells where c and e take different values. Configurations c and e are called *asymptotic*, denoted $c \sim e$, if *diff* (c, e) is a finite set.

A *pattern* over a set $A \subseteq \mathbb{Z}^d$ is an assignment $p : A \to S$. If A is finite, we say that p is a *finite* pattern. Note that every configuration is a pattern with domain $\mathbb{Z}$. 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 *translation* of a pattern $p : A \to S$ by $a \in \mathbb{Z}^d$ is denoted by $\sigma^a p$ and is defined by $(\sigma^a p)[i] \triangleq p[a + i]$. When $d = 1$, we may write σ for σ^1. Finite patterns are often considered modulo translations; we do not distinguish between a finite pattern and its translations.

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 $\diamond \in S$ such that F maps the $\diamond$-uniform configuration to itself; i.e., $f(\diamond^N) = \diamond$. If $\diamond$ is a quiescent state, the image of every $\diamond$-finite configuration is also $\diamond$-finite. More generally, if $c \sim e$ then $F(c) \sim F(e)$.

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 $e \in (S^A)^{\mathbb{Z}^d}$ 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 (up to translations) using edge-labeled De Bruijn graphs, which we are going to exploit in Sect. 4. 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 $\lambda : E \to S^{k+1-(l+r)}$ which is defined as follows. For every edge $u_0 u_1 \ldots u_k \in S^{k+1}$, let $\lambda(u_0 u_1 \ldots u_k) = v_l v_{l+1} \ldots v_{k-r}$ 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.

Cellular automata are often studied as topological dynamical systems. This will be the point of view in Sect. 6 of this article. The configuration space $S^{\mathbb{Z}^d}$ can be naturally topologized with the product topology. The convergence in this topology is defined cell-wise: a sequence $c_1, c_2, \ldots$ of configurations converges to a configuration c if and only if for every cell i there is $n \geq 1$ such that $c_k[i] = c[i]$ for all $k \geq n$. The product topology on the configuration space is compact and metric. A simple distance function for this topology is defined by

$$\rho(c, e) \triangleq 2^{-\kappa(c,e)}$$

for every two configurations c and e, where $\kappa(c, e)$ is the smallest non-negative integer i (∞ if none exists) such that c and e differ on the central hypercube $[-i, i]^d$. All CA global mappings are continuous with respect to the product topology. In fact, the CA global mappings are exactly those continuous mappings that are translation symmetric (Hedlund 1969).

A continuous mapping $F : X \to X$ on a compact metric space X defines a topological *dynamical system*. Two dynamical systems (X, F) and (Y, G) are *isomorphic* if there is a homeomorphism $\varphi : X \to Y$ such that $\varphi \circ F = G \circ \varphi$. Topological dynamics is the study of the dynamical properties that can be expressed topologically and hence are invariant under isomorphisms. These include sensitivity properties (whether small perturbations lead to large deviations) and mixing properties (whether each orbit covers the whole space) that are used to describe chaos in dynamical systems. See Kůrka (2003) for more information.

2.2 Conservation laws in cellular automata

We formulate conservation laws in a general setting in that the energy values can be chosen from an arbitrary semigroup. This makes the possible relation between different conservation laws and the information embedded in them transparent, and organizes all conservation laws in a hierarchy. However, the material of Sects. 5 and 6, as well as Example 3 can be followed without reference to this algebraic point of view, by simply choosing the set of real numbers $\mathbb{R}$ as the semigroup.

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 $\diamond$. 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 $\hat{\mu}(c)$ of a configuration $c \in S^{\mathbb{Z}^d}$ as the sum

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

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 $\mu(\diamond^W) = 0$. Then the above sum will have a natural meaning for all $\diamond$-finite

configurations—only finite number of the terms have values other than 0. For all other configurations we leave the total energy undefined.[1]

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

$$\hat{\mu}(F(c)) = \hat{\mu}(c) \tag{1}$$

for every $\diamond$-finite configuration c. Then $\hat{\mu}$ is a Φ-valued conserved energy defined using window W, and (1) is a *conservation law* for the CA. More concisely, the conservation law can be identified by the pair (Φ, μ). If Φ is in fact a group (or if $\Phi = \mathbb{R}$ is the set of real numbers), we may speak of a group-valued (resp. real-valued) energy or conservation law.

Notice that the values of $\hat{\mu}$ do not depend on the displacement of the window. That is, for every translation σ^a, $\hat{\mu} \circ \sigma^a = \hat{\mu}$. The conservation of $\hat{\mu}$ by a CA F means further that its value is constant over the orbits of F; i.e., $\hat{\mu} \circ F = \hat{\mu}$. Therefore, the kernel of $\hat{\mu}$ 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 $\hat{\mu} : C_\diamond[S] \to \Phi$ 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 $\hat{\Phi} \triangleq \hat{\mu}(C_\diamond[S])$ of Φ. So (Φ, μ) is trivial, if and only if, the realizable sub-monoid $\hat{\Phi}$ is trivial.

Let us now fix a CA $F : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ with a quiescent state $\diamond$, and a finite window $W \subseteq \mathbb{Z}^d$. Every conserved energy valuation μ for F satisfies (1) for every $\diamond$-finite configuration c. An important observation is that the largest semigroup generated by the (formal) values of μ for which (1) holds for every $\diamond$-finite configuration c provides the *most general* conservation law for F defined using window W.

Put it precisely, let $\Sigma \triangleq S^W - \{\diamond^W\}$ 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 $\mu_\circ(\diamond^W) = 0$, and $\mu_\circ(p) = p$ for any other pattern p over W, and let $\hat{\mu}_\circ : C_\diamond[S] \to \mathbb{N}^\Sigma$ be the corresponding total energy. Let $\cong \subseteq \mathbb{N}^\Sigma \times \mathbb{N}^\Sigma$ be the smallest monoid congruence where $\hat{\mu}_\circ(F(c)) \cong \hat{\mu}_\circ(c)$ for every $\diamond$-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 $\mu_F : S^W \to \Phi_F$ with $\mu_F = h_\cong \circ \mu_\circ$. Clearly the pair (Φ_F, μ_F) identifies a conservation law, because

$$\hat{\mu}_F(F(c)) = h_\cong(\hat{\mu}_\circ(F(c))) = h_\cong(\hat{\mu}_\circ(c)) = \hat{\mu}_F(c).$$

If (Φ, μ) is any other conservation law with window W, then for any $\diamond$-finite configurations c and e,

$$\hat{\mu}_F(c) = \hat{\mu}_F(e) \Longrightarrow \hat{\mu}_\circ(c) \cong \hat{\mu}_\circ(e) \Longrightarrow \hat{\mu}(c) = \hat{\mu}(e).$$

The second implication follows from the fact that, by (1), the kernel of the homomorphism $\mathbb{N}^\Sigma \to \Phi$ determined by $\mu_\circ(p) \mapsto \mu(p)$ is coarser than $\cong$. We see that μ_F defines a finer partitioning of $C_\diamond[S]$ than any other conservation law with window W, i.e., it is the most

[1] This is not the only possible approach, but it sounds most natural to us. For the real-valued energies, this and several other approaches lead to equivalent definitions (see Durand et al. 2003; Hattori and Takesue 1991; Pivato 2002).

general conservation law with window W. More specifically, there is a monoid homomorphism $h : \Phi_F \to \Phi$ satisfying $\mu = h \circ \mu_F$. Then, schematically,

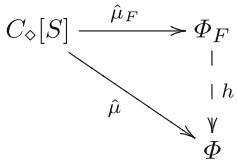

(2)

One can verify that using $\hat{\Phi}_F$ 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., $\left(\Phi_F^{(W)}, \mu_F^{(W)}\right)$ is *more general* than $\left(\Phi_F^{(W')}, \mu_F^{(W')}\right)$. 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, μ_F) based on W for F. Likewise, we define the realizable subgroup $\hat{G}_F \triangleq \hat{\mu}_F(C_\diamond[S])$ of G_F, where $\hat{\mu}_F$ is the total energy mapping corresponding to μ_F. The conservation law $\left(\hat{G}_F, \mu_F\right)$ satisfies a similar universal property (2) among group-valued conservation laws. The usual real-valued conservation laws are obtained when the energies are in the additive group of real numbers.

In the group-valued case we have a more elegant formulation of energy conservation that does not require a quiescent state but uses asymptotic configurations instead. For any group-valued energy valuation μ with window W we define the *potential difference*

$$\Delta(c, e) \triangleq \sum_{i \in \mathbb{Z}^d} [\mu(e[i + W]) - \mu(c[i + W])] \tag{3}$$

between asymptotic configurations c and e. The sum is defined as it contains only a finite number of non-zero terms. Note the additivity of Δ: for any three asymptotic configurations c_1, c_2 and c_3 we have

$$\Delta(c_1, c_2) + \Delta(c_2, c_3) = \Delta(c_1, c_3).$$

Now, the conservation of μ by a CA F can be stated as the requirement that

$$\Delta(F(c), F(e)) = \Delta(c, e) \tag{4}$$

holds for all asymptotic c and e. In the presence of a quiescent state $\diamond$, this is easily seen equivalent to (1).

3 Examples

The following example shows that semigroup-valued conservation laws can extract more information than group-valued ones.

Example 1 (Spreading 1's) Consider the one-dimensional CA F with binary state set $\mathbb{Z}_2 = \{0, 1\}$, 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

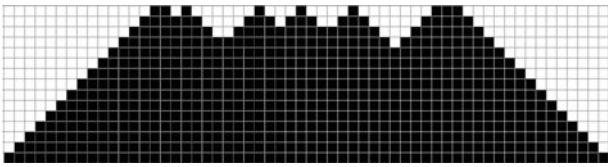

Fig. 1 A space–time snapshot from the CA in Ex. 1

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 configuration has total energy 0, while every other 0-finite configuration has total energy 1. This energy valuation is easily seen to be the only non-trivial conservation law for this CA, so in this case the hierarchy of semi-group valued conservation laws is non-trivial but bounded.

Our second example shows that group-valued conservation laws, in turn, can define a finer partitioning than real-valued conservation laws.

Example 2 (XOR) This example is again a one-dimensional CA F with binary state set, and neighborhood $\{-1, 0, 1\}$. The local rule is the XOR rule $f(a, b, c) = a \oplus b \oplus c$ where we denote $x \oplus y = x + y \pmod 2$. 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. Robison (1987)). This implies that F cannot have any non-trivial real-valued conservation law. On the other hand, F preserves the parity of the configurations. Let $G = \mathbb{Z}_2$ be the binary cyclic group, and consider the identity energy function $\mu : \{0, 1\} \to \mathbb{Z}_2$ on window $\{0\}$. The total energy $\hat{\mu}(c)$ is simply the parity of the number of 1's in c, and is preserved by F.

The following example is a *reversible* one-dimensional CA that has no non-trivial real-valued conservation laws of any range, resolving the question we asked in Formenti et al. (2008). Recall that a CA F is called reversible if F is injective (in which case it is known to be a bijection, and the inverse function is given by the *inverse CA*, see e.g. Kari (2005)).

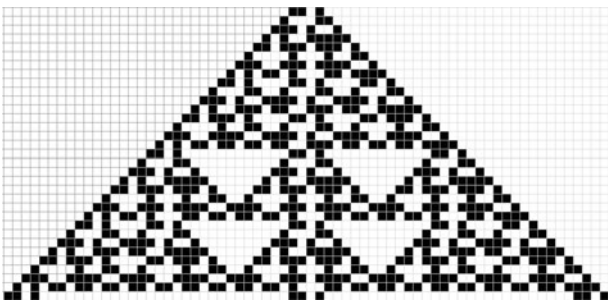

Fig. 2 A space–time snapshot from the CA in Ex. 2

Analogously to Example 2, the CA in this example creates new copies of any finite initial pattern at times 2^n, for all sufficiently large n.

Example 3 (Replicating reversible CA) The state set is $\mathbb{Z}_2 \times \mathbb{Z}_2$, so each state is a pair (a, b) of bits (Fig. 3). The neighborhood is $\{-1, 0\}$ and the local rule is

$$f((a, b), (c, d)) = (b, c \oplus d).$$

For this CA, any (0, 0)-finite configuration is conveniently represented as the column vector $((p(x), q(x))^T$ of two Laurent polynomials $p(x)$ and $q(x)$ over the ring $\mathbb{Z}_2$, where the coefficients of any x^k in $p(x)$ and $q(x)$ are the two bits in cell k. See Kari (2000) for more details on using generating functions to analyze linear CA over rings $\mathbb{Z}_m^n$.

If the local rule is expressed as the 2×2 matrix

$$M = \begin{pmatrix} 0 & x \\ 1 & 1 \end{pmatrix}$$

of Laurent polynomials (as described in Kari 2000) then the configuration $(p(x), q(x))^T$ becomes in one time step the configuration

$$M \begin{pmatrix} p(x) \\ q(x) \end{pmatrix}.$$

Because matrix M is invertible, the CA it defines is reversible. The inverse CA is given by the inverse matrix

$$M^{-1} = \begin{pmatrix} x^{-1} & 1 \\ x^{-1} & 0 \end{pmatrix},$$

that is, the inverse CA uses the neighborhood $\{0, 1\}$, and the local update rule is $((a, b), (c, d)) \mapsto (b \oplus c, c)$.

The t'th iteration of the CA corresponds to multiplying with the matrix M^t. For $t = 2^n$ we easily see that

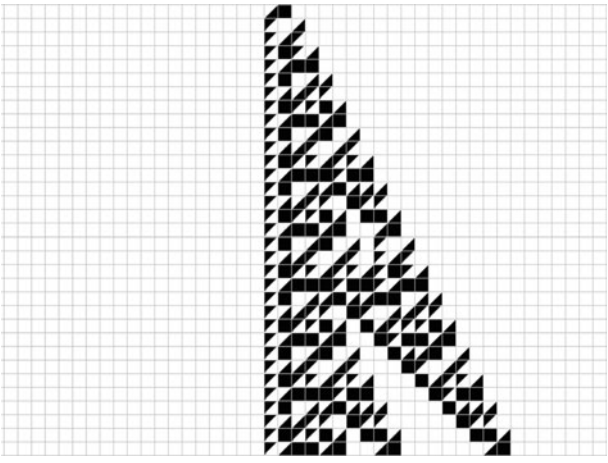

Fig. 3 A space–time snapshot from the CA in Ex. 3. The *black triangles* represent states (0, 1) and (1, 0), while the *solid white* and *black* indicate (0, 0) and (1, 1), respectively.

$$M^{2^n} = \begin{pmatrix} \alpha_n(x) & x \\ 1 & 1+\alpha_n(x) \end{pmatrix},$$

where

$$\alpha_n(x) = x + x^2 + x^4 + \cdots + x^{2^{n-1}}.$$

Indeed, this is true for $n = 1$, and assuming it for n, we have for $n+1$ that

$$M^{2^{n+1}} = M^{2^n} M^{2^n} = \begin{pmatrix} \alpha_n(x) & x \\ 1 & 1+\alpha_n(x) \end{pmatrix}^2 = \begin{pmatrix} \alpha_n(x)^2 + x & x \\ 1 & \alpha_n(x)^2 + x + 1 \end{pmatrix}.$$

Now it is enough to notice that

$$\alpha_n(x)^2 = \left(x + x^2 + x^4 + \cdots + x^{2^{n-1}}\right)^2 = x^2 + x^4 + \cdots + x^{2^n} = \alpha_{n+1}(x) + x$$

in the ring $\mathbb{Z}_2$.

Consider any finite initial configuration $(p(x), q(x))^T$, and let us denote its 2^nth successor $M^{2^n}(p(x), q(x))^T$ by $(p_n(x), q_n(x))^T$. Because

$$M^{2^{n+1}} - M^{2^n} = x^{2^n} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix},$$

we immediately see that

$$\begin{pmatrix} p_{n+1}(x) \\ q_{n+1}(x) \end{pmatrix} = x^{2^n} \begin{pmatrix} p(x) \\ q(x) \end{pmatrix} + \begin{pmatrix} p_n(x) \\ q_n(x) \end{pmatrix},$$

so a new copy of the initial configuration $(p(x), q(x))^T$ is added at the position 2^n during the time interval $2^n \longrightarrow 2^{n+1}$. By choosing large n, an arbitrarily long block of 0's can be left between the supports of the 2^nth successor $(p_n(x), q_n(x))^T$ and the new copy of $(p(x), q(x))^T$ in the 2^{n+1}st successor. Consequently, the contribution of the initial contribution $(p(x), q(x))^T$ under any real-valued energy valuation μ must be zero, showing that the CA has no non-trivial real-valued conservation law of any range. □

4 Semigroup-valued conservation laws

The definition of the most general conservation law of a certain range, given in Sect. 2.2, 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. Grillet 1995). 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 $\diamond \in S$. Clearly, $(\diamond, \diamond)$ is a quiescent state for the product $F \times F$. Let $\Phi_B = \{0, 1\}$ be the Boolean semigroup with $a + b \triangleq a \vee b$, for any $a, b \in \Phi_B$. Define an energy valuation μ with window $\{0\}$ and values from Φ_B by

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

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 (Moore 1962; Myhill 1963), 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 (Kari 1994). 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* $\mu : S \to \Phi_B$ *into the semigroup* $\Phi_B = (\{0,1\}, \vee)$*, 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 $\diamond$ 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 $k = l + r + m - 1$, consider the k'th order De Bruijn representation $(B_k[S], \lambda)$ of F. This has a vertex $\diamond^k$, with a loop edge $\diamond^{k+1}$ which is labeled by $\diamond^m$. Any path corresponding to a $\diamond$-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 $u_0 u_1 \ldots u_k \in S^{k+1}$ let us assign two elements

$$\alpha(u_0 u_1 \ldots u_k) \triangleq \mu(u_0 u_1 \ldots u_{m-1}) \tag{5}$$

and

$$\beta(u_0 u_1 \ldots u_k) \triangleq \mu(v_l v_{l+1} \ldots v_{l+m-1}) \tag{6}$$

from Φ, where $v_l v_{l+1} \ldots v_{k-r} = \lambda(u_0 u_1 \ldots u_k)$ is the label of $u_0 u_1 \ldots u_k$. The total energy of a $\diamond$-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 $\diamond^{k+1}$ do not contribute to the total energy, because $\mu(\diamond^m) = 0$. For any path $p = p_1 p_2 \ldots p_n$ (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.,

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

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 $\diamond^k$, $\alpha(p) = \beta(p)$. There are infinitely many such paths p, but it turns out that it is sufficient to verify $\alpha(p) = \beta(p)$ for a suitably selected finite collection of paths p.

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*

$$\alpha(p) = \beta(p) \tag{7}$$

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.

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 Eq. (7) 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 Eq. (7) holds by the assumption.

Suppose that the Eq. (7) 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 Eq. (7) 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

$$\alpha(\tilde{x}) + \alpha(\tilde{y}) = \beta(\tilde{x}) + \beta(\tilde{y}) \tag{8}$$

$$\alpha(\tilde{x}) + \alpha(c) + \alpha(\tilde{y}) = \beta(\tilde{x}) + \beta(c) + \beta(\tilde{y}) \tag{9}$$

$$\alpha(x) + \alpha(y) = \beta(x) + \beta(y) \tag{10}$$

from which we obtain that

$$\begin{aligned}
\alpha(p) &= \alpha(x) + \alpha(c) + \alpha(y) \\
&= \gamma_x + \gamma_y + \alpha(\tilde{x}) + \alpha(c) + \alpha(\tilde{y}) \\
&= \gamma_x + \gamma_y + \beta(\tilde{x}) + \beta(c) + \beta(\tilde{y}) \quad \text{(by (9))} \\
&= \gamma_x + \gamma_y + \alpha(\tilde{x}) + \alpha(\tilde{y}) + \beta(c) \quad \text{(by (8))} \\
&= \alpha(x) + \alpha(y) + \beta(c) \\
&= \beta(x) + \beta(y) + \beta(c) \quad \text{(by (10))} \\
&= \beta(p)
\end{aligned}$$

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. □

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.*

Proof In Prop. 1 choose as G the de Bruijn graph $B_k[S]$, and let $A = B = \{\diamond^k\}$. Let α and β be defined by (5) and (6) with $\mu = \mu_o$. Then $\Phi = \Phi_F$. □

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

5 The existence problem

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 μ. (This is similar to the real-valued case; see e.g. (Kari 2005)). In particular, for any fixed window W one can effectively construct the most general group-valued conservation law $\left(G_F^{(W)}, \mu_F^{(W)}\right)$ based on W.

The next challenge would be to analyze the hierarchy of conservation laws, obtained by increasing the size of their window. For example, given a CA, is it possible to decide if it has any conservation law at all? Can one determine if the hierarchy of conservation laws is bounded, i.e., whether there is a window size that provides the absolutely most general conservation law? In this section we prove that the answer is negative: these questions are undecidable even among one-dimensional CA.

We first establish a simple condition that guarantees that no non-trivial conservation law exists. Let $\diamond$ be a quiescent state for CA F. We say that F is $\diamond$-*nilpotent* if all $\diamond$-finite configurations eventually become $\diamond$-uniform. In this case, any energy valuation μ that is conserved by F must assign the same value to all $\diamond$-finite configurations, and hence no non-trivial conservation law exists:

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

Our next lemma deals with the other extreme: it gives a condition that guarantees an unbounded hierarchy of conservation laws. Let F be a CA. An energy valuation μ with window W is termed *F-invariant* if for every configuration c we have

$$\mu(c[W]) = \mu(F(c)[W]).$$

In other words, the energy values are preserved "at the same position". Note that this definition does not require the values of μ to have any algebraic structure. However, if the values are in a semigroup or group, it is clear that F-invariance of μ implies that F conserves μ. Let us say that μ is trivial if all patterns are given the same value.

Lemma 2 *Let F be a cellular automaton and μ a non-trivial F-invariant valuation. Then F has an unbounded hierarchy of real-valued conservation laws. In fact, for every semigroup-valued conservation law μ' of F there exists two configurations c and e, and a real-valued conservation law μ'' such that $\hat{\mu}'(c) = \hat{\mu}'(e)$ but $\hat{\mu}''(c) \neq \hat{\mu}''(e)$.*

Proof Let $\diamond$ be an arbitrary state. Since μ is non-trivial, there is a pattern p over domain W such that $\mu(p) \neq \mu(\diamond^W)$. For every $i \in \mathbb{Z}^d$, we define the real-valued energy valuation μ_i with window $W \cup (i + W)$ as follows: For every pattern $q : W \cup (i + W) \to S$ let $q_1 \triangleq q[W]$ and $q_2 \triangleq q[i + W]$ and define

$$\mu_i(q) = \begin{cases} 1, & \text{if } \mu(q_1) = \mu(q_2) = \mu(p), \\ 0, & \text{otherwise.} \end{cases}$$

In other words, value 1 is assigned at a cell if and only if μ gives value $\mu(p)$ to both patterns around that cell and around the cell offset by i. Since μ is F-invariant, each μ_i is also F-invariant and hence conserved by F.

For all $i \in \mathbb{Z}^d$ such that $W \cap (i + W) = \emptyset$, let us denote by c_i the configuration that contains pattern p in domains W and $i + W$, and state $\diamond$ outside of $W \cup (i + W)$. Clearly, $\hat{\mu}_i(c_i) > 0$.

Let μ' be an arbitrary energy valuation with window A. It is clear that μ' gives the same total energy to all but a finite number of the configurations c_i: When i is sufficiently large, no translate of A intersects both W and $i + W$. However, for all sufficiently large $i \in \mathbb{Z}^d$, all but a finite number of $j \in \mathbb{Z}^d$ satisfy $\hat{\mu}_i(c_j) = 0$, so there are configurations c_i and c_j that are given the same value by μ' but different values by μ_i. □

In the following, we show that the situations indicated in Lemmas 1 and 2 are recursively inseparable, even for given one-dimensional CA F. In the proof, we use a standard simulation of two-counter machines by one-dimensional cellular automata, and exploit the following theorem about two-counter machines. A *two-counter machine* is a finite automaton equipped with two unbounded *counters*, each storing a natural number. The automaton can increase or decrease the value of each counter, and can test if either has value zero. Two-counter machines are known to be equivalent in power with Turing machines—any algorithm can be implemented on a two-counter machine (see e.g. Minsky 1967). Blondel, Cassaigne and Nichitiu have shown that the presence of periodic orbits in two-counter machines cannot be decided algorithmically:

Theorem 3 (Blondel et al. 2002) *Given a deterministic two-counter machine A with no halting state, it is undecidable whether A has a periodic orbit.*

Consider the following two extreme cases concerning the hierarchy of conservation laws of a given CA F:

I. F has no non-trivial semigroup-valued conservation law.
II. For every semigroup-valued conservation law μ' of F there exists a real-valued conservation law μ'' of F such that $\hat{\mu}'(c) = \hat{\mu}'(e)$ but $\hat{\mu}''(c) \neq \hat{\mu}''(e)$ for some configurations c and e.

In particular, condition II means that the hierarchy of (real-valued/group-valued/semigroup-valued) conservation laws for F is unbounded.

Theorem 4 *Cases I and II above are recursively inseparable for a given one-dimensional CA F.*

Proof We show how to reduce the problem of whether a given counter machine has a periodic orbit to the problem of separating I and II. Since the former is undecidable, we conclude the recursive inseparability of I and II.

Let A be a two-counter machine with state set Q, two registers x_1 and x_2, and transition function $\delta : Q \times \{0, 1\}^2 \to Q \times \{1, 2\} \times \{-, 0, +\}$. A *configuration* of the machine consists of its current state $q \in Q$, and the current value of its two registers $x_1, x_2 \in \mathbb{N}$. The transition rule reads the current state, checks whether either of the registers contains zero (1 means the value of the register is zero, and 0 means otherwise), decides the new state, chooses one of the registers (1 or 2 as the index of the chosen register), and instructs whether the chosen register must be decreased ($-$), left unchanged (0), or increased ($+$).

We construct a CA F with a designated quiescent state $\diamond$ such that

(a) if A has a periodic orbit, F has a non-trivial F-invariant valuation μ (and hence, by Lemma 2, it satisfies condition II), while
(b) if A has no periodic orbit, F is $\diamond$-nilpotent (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 $\diamond$. 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 $\diamond$-finite configuration eventually goes quiescent; i.e., F is $\diamond$-nilpotent.

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 $B = Lb_1b_2 \ldots b_{n-1}R$. Let us denote by $\mathcal{B}$ the set of all simulation blocks $B' \in S^{n+1}$ that eventually turn into B. Since no new end-marker is ever created by F, and since the information cannot pass through the end-markers, we can argue that the set $\mathcal{B}$ is "stable"; once we know that a block in $\mathcal{B}$ occurs in a certain position at a certain time during the evolution of the CA, we also know that in any other time a block in $\mathcal{B}$ occurs in that same position. In other words, the valuation $\mu : S^{[0,n]} \to \{0, 1\}$ defined by

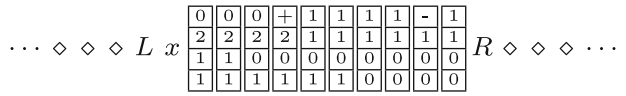

Fig. 4 A syntactically correct simulation block in the CA described in the proof of Theorem 4. The simulated machine is in state $x \in Q$. The first counter contains 2, while the second counter has 6. There is a signal, moving to the right, commanding the second counter to increase

$$\mu(x) \triangleq \begin{cases} 1 & \text{if } x[[0,n]] \in \mathcal{B}, \\ 0 & \text{otherwise} \end{cases}$$

is invariant under F.

Let us now give the precise construction. Let $E = \{L, R\}$, $X = \{0, 1\}$, $K = \{1, 2\}$ and $C = \{0, +, -, 1\}$. The state set of F is $S = \{\diamond\} \cup E \cup Q \cup (X \times X \times K \times C)$, 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 to 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. A sample syntactically correct simulation block is depicted in Fig. 4.

To increase the value of the counter x_b (for $b \in K$), a right-traveling signal is initiated in the form of a + tailed by a sequence of 0's in the C component, accompanied by a sequence of b's in the K component. The segment between the head of the signal and the right end of the simulation block is irrelevant and can take arbitrary values in the C and K components. Once the signal reaches the end of the unary representation of the counter x_b, it increases the value of the counter by attaching an extra 1 to the end of its representation. The signal is then destroyed and an acknowledgement signal is created. The acknowledgement signal consists only of a 1 in the C component that travels to the left, eating the sequence of 0's. During the propagation of the + signal and the return of the acknowledgement, the automaton is kept on hold. Once an acknowledgement reaches the left end of the block, the automaton proceeds to its next step. The protocol for decreasing a counter is similar. The local rule $f : S^3 \to S$ of the CA is presented in Table 1. □

Corollary 2 *It is undecidable whether a given one-dimensional CA has any non-trivial (semigroup/group/real-valued) conservation laws. It is also undecidable whether it has an unbounded hierarchy of such conservation laws.*

6 Restrictions on the dynamics

In this section, we argue that the structure of the hierarchy of conservation laws in a CA may severely restrict its dynamical properties (or vice versa). A rather trivial example of an interconnection between the dynamical properties of a CA and the structure of its hierarchy of conservation laws is already implied by Lemmas 1 and 2 of the previous section. A dynamical system (X, F) is said to be *equicontinuous* if the orbit of every point $x \in X$ is

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

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

stable, in the sense that for every $\varepsilon > 0$, there is a $\delta > 0$ such that for any other point $y \in X$ with $\rho(x, y) < \delta$ we have $\rho(F^i(x), F^i(y)) < \varepsilon$ for all $i \geq 0$. For instance, every nilpotent CA is equicontinuous. It is well-known that the equicontinuous CA are precisely those CA that are eventually periodic (Kůrka 1997).

It is easy to see that every eventually periodic CA that is not nilpotent has a non-trial invariant energy valuation. Therefore, while according to Lemma 1, nilpotent CA have no non-trivial conservation law, Lemma 2 implies that every non-nilpotent equicontinuous CA has an unbounded hierarchy of conservation laws.

We are now going to present a less trivial result about the interconnection between dynamics and conservation laws in cellular automata. We show that a CA with a strong

chaotic behavior (positive expansivity or strong transitivity) cannot have any non-trivial real-valued conservation laws. In the following we use the potential difference formulation (4) for the conservation of energy valuation.

Let μ be a real-valued energy function that uses window W, and let Δ be the potential difference defined by (3) between asymptotic pairs of configurations. Let us call a configuration c a *ground* configuration if for every configuration e asymptotic to c, we have $\Delta(c,e) \geq 0$. Let us first show that ground configurations always exist.

Proposition 3 *Every real-valued local potential difference has at least one ground configuration.*

Proof Let $\diamond \in S$ be arbitrary and let $\diamondsuit$ be the $\diamond$-uniform configuration. Let us fix a growing sequence $D_1 \subseteq D_2 \subseteq \cdots$ of finite subsets of $\mathbb{Z}^d$ whose union is $\mathbb{Z}^d$. For example, we can choose $D_n \triangleq [-n,n]^d$.

For every $n \in \mathbb{N}$, let us choose a configuration c_n with minimum potential difference $\Delta(\diamondsuit, c_n)$ among all configurations with $supp_\diamond(c_n) \subseteq D_n$. There is a finite number of such configurations so a minimum exists. Due to the compactness of the configuration space, there exists a configuration c that is the limit of a subsequence of $c_1, c_2, \ldots$ We claim that c is a ground configuration.

Let e be any configuration that is asymptotic to c, and let $D \triangleq diff\,(c,e)$ be the set of cells on which c and e differ. We denote

$$W(D) \triangleq D + W + (-W) = \{k + i - j : k \in D, i,j \in W\}$$

so that $\Delta(c,e) = \Delta(c',e')$ holds for any c' and e' that agree with each other outside of $W(D)$, and with c and e, respectively, inside of $W(D)$. Let $n \in \mathbb{N}$ be such that

(i) $W(D) \subseteq D_n$, and
(ii) c_n and c agree on $W(D)$.

Let x be the $\diamond$-finite configuration which agrees with e on $W(D)$ and with c_n outside of $W(D)$. In particular, x takes value $\diamond$ outside of D_n. We have

$$\Delta(c,e) = \Delta(c_n,x) = \Delta(\diamondsuit,x) - \Delta(\diamondsuit,c_n) \geq 0.$$

Therefore, c is a ground configuration. □

Let us denote the set of ground configurations of μ by G_μ. Next we see that ground configurations form a subshift.

Proposition 4 *The set G_μ is closed in the product topology.*

Proof Let $g_1, g_2, \ldots$ be a converging sequence of ground configurations, with limit c. It is enough to show that $c \in G_\mu$. Let e be an arbitrary configuration that is asymptotic to c. Let $D \triangleq diff\,(c,e)$ and, as in the proof above, we denote $W(D) \triangleq D + W + (-W)$. There exists $n \in \mathbb{N}$ such that g_n and c agree inside $W(D)$. Define the configuration x which agrees with e inside $W(D)$ and with g_n outside of $W(D)$. Then

$$\Delta(c,e) = \Delta(g_n,x) \geq 0.$$

We conclude that c is a ground configuration. □

Lemma 3 *Suppose μ is conserved by CA F. If $c \in F^{-1}(g)$ for a ground configuration g, then also c is a ground configuration.*

Proof Let e be any configuration asymptotic to c. Then

$$\Delta(c, e) = \Delta(F(c), F(e)) = \Delta(g, F(e)) \geq 0.$$

□

Now we are ready to prove the main result of this section. A dynamical system is *strongly transitive* if the orbits of the points in every open ball (no matter how small) cover the whole space (Manzini and Margara 1999). In other words, a dynamical system (X, F) is strongly transitive if for every point $x \in X$, the set $\bigcup_{t>0} F^{-t}(x)$ is dense.

Theorem 5 *A CA with a non-trivial real-valued conservation law cannot be strongly transitive.*

Proof Let F be a strongly transitive CA, and let μ be a real-valued energy function which is conserved by F. By Proposition 3 there exists a ground configuration c, so by strong transitivity and Lemma 3 the set G_μ of ground configurations is dense. It follows then from Proposition 4 that all configurations are ground configurations, so μ must be trivial. □

A characteristic property of chaotic dynamical systems is their sensitivity to initial conditions. Positively expansive dynamical systems are systems with a high degree of sensitivity to initial conditions: A system (X, F) is said to be *positively expansive* if there exists a constant $\varepsilon > 0$ such that any two distinct points $x, y \in X$ (no matter how close) are eventually separated by distance at least ε. Cellular automata in higher than one dimensions cannot be positively expansive (Finelli et al. 1998). In one dimension, however, positively expansive CA exist and constitute a curious class of CA. Intuitively, they are those CA in which any information on the lattice propagates indefinitely in both directions. Every positively expansive CA is strongly transitive. In fact, every positively expansive CA is known to be isomorphic to a so-called mixing one-sided shift space of finite type (Kůrka 1997; Nasu 1995) and every such system is strongly transitive.

Corollary 3 *Positively expansive CA have no non-trivial real-valued conservation laws.*

The XOR cellular automaton in Example 2 is positively expansive, and as we already knew, has no non-trivial real-valued conservation law.

7 Conclusion

We have examined a number of algorithmic problems that arise from studying algebraic conservation laws for cellular automata. Semigroup-valued conservation laws are highly expressive, still not so tractable in two- and higher dimensional spaces. 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? Example 3 shows that there exist reversible CA without any real-valued conservation laws so the question has both positive and negative instances. Our proof of Theorem 4 takes advantage of the existence of very long transients to construct CA whose conservation laws need very large windows, so the method cannot be applied in the reversible case.

Acknowledgments This work has been supported by the French ANR Blanc Project "EMC" and the Academy of Finland grant 131558.

References

Biryukov AP (1967) Some algorithmic problems for finitely defined commutative semigroups. Sib Math J 8:384–391

Blondel VD, Cassaigne J, Nichitiu C (2002) On the presence of periodic configurations in turing machines and in counter machines. Theor Comput Sci 289:573–590

Boccara N, Fukś H (1998) Cellular automaton rules conserving the number of active sites. J Phys A 31:6007–6018

Conway JH, Lagarias JC (1990) Tiling with polyominoes and combinatorial group theory. J Comb Theory A 53:183–208

Durand B, Formenti E, Róka Z (2003) Number conserving cellular automata I: decidability. Theor Comput Sci 299:523–535

Finelli M, Manzini G, Margara K (1998) Lyapunov exponents versus expansivity and sensitivity in cellular automata. J Complex 14:210–233

Formenti E, Grange A (2003) Number conserving cellular automata II: dynamics. Theor Comput Sci 304:269–290

Formenti E, Kari J, Taati S (2008) The most general conservation law for a cellular automaton. In: Hirsch EA, Razborov AA, Semenov AL, Slissenko A (eds) CSR vol 5010 of Lecture notes in computer science. Springer, New York, pp 194–203

Fukś H (2000) A class of cellular automata equivalent to deterministic particle systems. In: Feng S, Lawniczak AT, Varadhan SRS (eds) Hydrodynamic limits and related topics, vol 27 of Fields institute communications. American Mathematical Society, Providence, pp 57–69

Grillet PA (1995) Semigroups: an introduction to the structure theory. Dekker, New York

Hardy J, de Pazzis O, Pomeau Y (1976) Molecular dynamics of a classical lattice gas: transport properties and time correlation functions. Phys Rev A 13:1949–1961

Hattori T, Takesue S (1991) Additive conserved quantities in discrete-time lattice dynamical systems. Physica D 49:295–322

Hedlund GA (1969) Endomorphisms and automorphisms of the shift dynamical system. Math Syst Theory 3:320–375

Kari J (1994) Reversibility and surjectivity problems of cellular automata. J Comput Syst Sci 48:149–182

Kari J (2000) Linear cellular automata with multiple state variables. In: Reichel H, Tison S (eds) STACS vol 1770 of Lecture notes in computer science. Springer, New York, pp 110–121

Kari J (2005) Theory of cellular automata: a survey. Theor Comput Sci 334:3–33

Kůrka P (1997) Languages, equicontinuity and attractors in cellular automata. Ergod Theory Dyn Syst 17:417–433

Kůrka P (2003) Topological and symbolic dynamics, vol 11 of Cours Spécialisés, Société Mathématique de France

Manzini G, Margara L (1999) A complete and efficiently computable topological classification of d-dimensional linear cellular automata over $\mathbb{Z}_m$. Theor Comput Sci 221:157–177

Minsky M (1967) Computation: finite and infinite machines. Prentice-Hall, Englewood Cliffs

Moore EF (1962) Machine models of self-reproduction. In: Proceedings of symposia in applied mathematics. AMS, Providence, pp 17–33

Moreira A, Boccara N, Goles E (2004) On conservative and monotone one-dimensional cellular automata and their particle representation. Theor Comput Sci 325:285–316

Myhill J (1963) The converse of Moore's garden-of-Eden theorem. Proc Am Math Soc 14:685–686

Nasu M (1995) Textile systems for endomorphisms and automorphisms of the shift. Mem Am Math Soc 114, no. 546, pp viii + 215

Pivato M (2002) Conservation laws in cellular automata. Nonlinearity 15:1781–1793

Pomeau Y (1984) Invariant in cellular automata. J Phys A 17:L415–L418

Robison AD (1987) Fast computation of additive cellular automata. Complex Syst 1:211–216

Takesue S (1987) Reversible cellular automata and statistical mechanics. Phys Rev Lett 59:2499–2502

Thurston WP (1990) Conway's tiling groups. Am Math Mon 97:757–773