Conservation Laws in Cellular Automata

Siamak Taati

Abstract A conservation law in a cellular automaton is the statement of the invariance of a local and additive energy-like quantity. This chapter reviews the basic theory of conservation laws in cellular automata. A general mathematical framework for formulating conservation laws in cellular automata is presented and several characterizations of them are summarized. Computational problems regarding conservation laws (verification and existence problems) are discussed. Microscopic explanations of the dynamics of the conserved quantities in terms of flows and particle flows are explored. The related concept of dissipating energy-like quantities are also addressed.

Siamak Taati
Department of Mathematics, University of Turku, FI-20014 Turku, Finland
e-mail: staati@utu.fi

1 Introduction

A cellular automaton (CA for short) is an abstract structure, consisting of a d-dimensional checkerboard ($d = 1, 2, 3, \ldots$). Each cell of the board has a state chosen from a finite set of states. The state of each cell changes with time, according to a uniform, deterministic rule, which takes into account the previous state of the cell itself and those in its neighborhood. The changes, however, happen synchronously, and in discrete time steps.

In mathematics and computer science, cellular automata are studied as abstract models of computation, in the same way that Turing machines are (see Chapters **??** and **??**). They are also treated as paradigms of symmetric dynamical systems on the Cantor space (see Chapter **??**). However, the chief reason for the interest in cellular automata comes from their characteristic similarities with nature: they are spatially-extended dynamical systems; they are uniform (the same laws are applied everywhere in the space); the interactions are local (no action-at-a-distance); the amount of information in a finite region of space is finite (cf. [22]). Further characteristics of nature, such as the microscopic reversibility and conservation laws also arise in cellular automaton in a natural way (see Chapter **??**). This makes cellular automata exceptionally suitable for modeling physical and biological phenomena, on the one hand (see Chapter **??**), and as a design framework for natural computing, on the other hand (see e.g. [44]).

This chapter is a survey of the basic known results about local and additive conservation laws in cellular automata. Studying conservation laws in cellular automata may be beneficial from different aspects. When modeling a physical system, conservation laws of the system serve as design constraints that one would like to program in the model (see Chapter **??**). In physically realistic models of computation, conservation laws should naturally be taken into account (see [22, 44]). Moreover, conservation laws in a cellular automaton may provide the same kind of "physical" insight about its dynamics as the insight that conservation laws in physics provide about the physical world.

In the rest of this section, we will illustrate the concept of conservation laws in cellular automata by a number of examples. In Section 2, we will provide a precise mathematical formulation of conservation laws in cellular automata. Section 3 is dedicated to the algorithmic problems arisen from studying conservation laws: how to discover them, and how to verify their validity. In Section 4, we will discuss local explanations of conservation laws in terms of the flow of the conserved quantity. Finally, in Section 5, we will comment on some related issues. Along our exposition, we will also talk about the closely related concept of non-increasing energy-like quantities.

1.1 Few Examples

Conservation laws in cellular automata are obtained in a similar fashion as in physics. A real value is associated to each local pattern of cell states, interpreted as the "energy" (or "electric charge", or ...) of that particular arrangement of states. The total energy of a configuration is the sum of the energy values of the patterns seen in different places on it. Intuitively, a conservation law asserts that the total energy of each configuration remains unchanged under the iterations of the CA.

As an example, consider the well-known *Traffic* CA, which resembles cars moving on a highway. The Traffic CA is a one-dimensional CA, consisting of an infinite number of cells arranged next to each other on a line. Each cell has two possible states: ■ (interpreted as a "car") or □ ("empty space"). At each step, a car moves one cell forward if and only if its front cell is empty. Figure 1 shows a typical space-time diagram of the evolution of the Traffic CA. Not surprisingly, the number of cars on the highway remains constant along the the evolution of the CA. To state

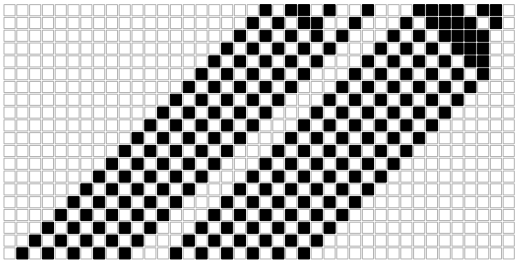

Fig. 1 A typical space-time diagram of the Traffic CA. Time evolves downward. The highway is directed towards the left.

this more precisely, let us index the various positions on the line with integers $i \in \mathbb{Z}$. A configuration of the model is an assignment of values ■ or □ to every position on the line. For each configuration $i \mapsto x[i]$, let us write Fx for the configuration after one step. We can now state the car-conservation law by saying that for any configuration x, the following equality holds:

$$\sum_{i=-\infty}^{+\infty} \theta\left(x[i]\right) = \sum_{i=-\infty}^{+\infty} \theta\left((Fx)[i]\right) \ , \tag{1}$$

where $\theta(\,\blacksquare\,) = 1$ and $\theta(\,\square\,) = 0$. Note that if the the number of cars on a configuration x is infinite, the sum $\sum_i \theta\left(x[i]\right)$ becomes $+\infty$. However, in this case, the configuration Fx has also infinitely many cars on it, and the equality still holds.[1]

Another example is the *Just Gliders* CA. This is also one-dimensional. Each cell can be in either state ◀ (a "particle" moving to the left), or ▶ (a "particle" moving to the right), or · ("empty space"). At each step, each particle moves one cell ahead. Particles moving in opposite directions annihilate when they meet. See Figure 2 for

[1] To learn more about cellular automata models of car traffic, see e.g. [51, 50], and the relevant discussion in Chapter **??**.

a typical space-time diagram. If we define the momentum of a right-moving particle ► to be 1, and the momentum of a left-moving particle ◄ to be -1, one can easily see that the total momentum of a configuration remains constant with time. More

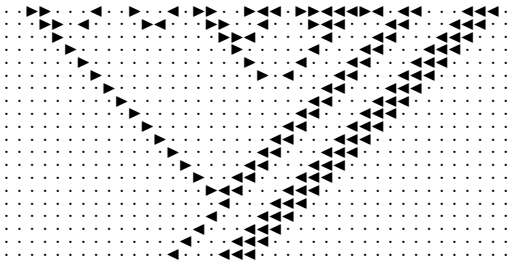

Fig. 2 A typical space-time diagram of the Just Gliders CA. The moving objects annihilate on encounter.

precisely, setting

$$\theta(a) \triangleq \begin{cases} 1 & \text{if } a = \blacktriangleright, \\ -1 & \text{if } a = \blacktriangleleft, \\ 0 & \text{if } a = \cdot. \end{cases} \tag{2}$$

Equation (1) is valid for any configuration x with a finite number of particles. For an infinite configuration (i.e., a configuration with infinitely many particles on it), the sum $\sum_i \theta(x[i])$ is not necessarily meaningful anymore. Therefore, we express the conservation law in terms of finite configurations only.

Alternatively, we could formulate the above conservation law in terms of the average momentum per cell of spatially periodic configurations. Namely, if a configuration x has period $p > 0$ (i.e., if $x[i+p] = x[i]$ for every i), we can say that

$$\frac{\sum_{i=1}^{p} \theta(x[i])}{p} = \frac{\sum_{i=1}^{p} \theta((Fx)[i])}{p} . \tag{3}$$

As we shall discuss later, these two formulations are equivalent in general; a CA conserves the energy of every finite configuration if and only if it preserves the average energy per cell of each spatially periodic configuration.

Next let us talk about few physically interesting examples. The *Ising* model was introduced by Wilhelm Lenz (1888–1957) and Ernst Ising (1900–1998) to explain the phenomenon of phase transition in ferromagnetic materials. It is a stochastic model and is extensively studied in statistical mechanics (see e.g. [38, 57, 26]). Gérard Vichniac has introduced a deterministic CA-like dynamics on it [63] (see also [61, 16]).

In the Ising model, each cell represents a tiny piece a ferromagnetic material having a spin (i.e., a magnetic moment resulting from the angular momentum of the electrons). For simplicity, each spin is approximated by either of two values: ↑ (spin-up) or ↓ (spin-down). Adjacent spins tend to align. This tendency is depicted by assigning an energy 1 to each pair of adjacent spins that are anti-aligned, and energy -1 to those that are aligned.

Vichniac's dynamics is specially designed in such a way as to conserve this energy, hence emulating the regime where there is no heat transfer in and out of the material. The states of the cells are updated in two stages. The cells are colored black and white as on the chess board. At the first stage, all the black cells are updated in the following way: a spin on a black cell is flipped (from $\uparrow$ to $\downarrow$, or from $\downarrow$ to $\uparrow$) if and only if the change does not affect the total energy of the bonds with its adjacent spins. At the second stage, the white cells are updated in a similar fashion. Figure 3 shows few snapshots from a simulation of this CA-like dynamics.

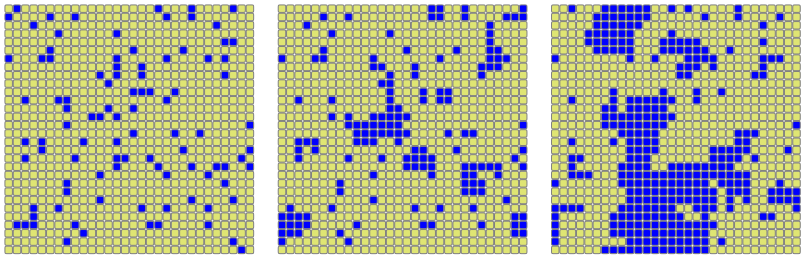

(a) (b) (c)

Fig. 3 Simulation of Vichniac's dynamics on a spatially periodic configuration of the two-dimensional Ising model. Blue represents $\uparrow$. Green represents $\downarrow$. (a) The initial configuration. (b) The configuration at time $t = 10$. (c) The configuration at time $t = 60$.

To put things formally, the set of possible states for each cell is $S \triangleq \{\uparrow, \downarrow\}$. In the d-dimensional model, the cells are indexed by the elements $i = (i_1, i_2, \ldots, i_d)$ in $\mathbb{Z}^d$, the d-dimensional square lattice. Hence a configuration of the model is an assignment $x : \mathbb{Z}^d \to S$. The dynamics is defined in terms of two mappings $x \mapsto F_0 x$ (for updating the black cells) and $x \mapsto F_1 x$ (for updating the white cells). For every cell $i \in \mathbb{Z}^d$, we have

$$(F_0 x)[i] \triangleq \begin{cases} \neg x[i] & \text{if } i \text{ black and } \sum_{j \in N(i)} \varsigma(x[j]) = 0, \\ x[i] & \text{otherwise,} \end{cases} \tag{4}$$

$$(F_1 x)[i] \triangleq \begin{cases} \neg x[i] & \text{if } i \text{ white and } \sum_{j \in N(i)} \varsigma(x[j]) = 0, \\ x[i] & \text{otherwise.} \end{cases} \tag{5}$$

Here we have used few shorthand notations. First, $\neg a$ denotes the reverse of a; that is, $\neg \uparrow \triangleq \downarrow$ and $\neg \downarrow \triangleq \uparrow$. Next, we have used $\varsigma(a)$ to denote the sign of a spin; that is, $\varsigma(\uparrow) \triangleq +1$ and $\varsigma(\downarrow) \triangleq -1$. Finally, $N(i)$ represents the set of immediate neighbors of cell i. In summary, the dynamics is obtained by alternately applying of F_0 and F_1 on a configuration:

$$x_0 \xrightarrow{F_0} x_0' \xrightarrow{F_1} x_1 \xrightarrow{F_0} x_1' \xrightarrow{F_1} x_2 \xrightarrow{F_0} \cdots . \tag{6}$$

The conservation of energy can now be formulated by saying that each of F_0 and F_1 preserves the sum

$$\Theta(x) \triangleq \sum_{\substack{i,j \\ \text{adjacent}}} [1 - \varsigma(x[i])\varsigma(x[j])] , \tag{7}$$

for any configuration x with only finitely many downward spins $\downarrow$, or finitely many upward spins $\uparrow$. A similar statement can be formulated for spatially periodic configurations.

Notice that Vichniac's dynamics is reversible; one can completely regenerate the configuration at time t, knowing the configuration at time $t+1$. In fact, all you need to do is to apply F_0 and F_1 in the reverse order. Reversibility is apparently a fundamental feature of nature. Even though the macroscopic world as we perceive it looks irreversible, every known physical process behaves reversibly in the ultimate microscopic scale (cf. [18]).

The next example, due to Pomeau [55] and Margolus [45], identifies an interesting energy-like invariant in a large class of reversible models. Each cell of the lattice has a state from the finite ring $\mathbb{Z}_m = \{0,1,2,\ldots,m-1\}$. The dynamics is of 2nd order; that is, the configuration at time $t+1$ depends not only on the configuration at time t, but also on the configuration at time $t-1$; that is, $c_{t+1} = F(c_t, c_{t-1})$. The state of a cell $i \in \mathbb{Z}^d$ at time $t+1$ is obtained by a rule of the form

$$c_{t+1}[i] = c_{t-1}[i] + f(c_t[N(i)]) . \tag{8}$$

Here $N(i)$ is a finite set of cells that is called the neighborhood of cell i. The neighborhood is assumed to be uniform; that is, there is a finite set $N \subseteq \mathbb{Z}^d$ such that $N(i) \triangleq \{i+k : k \in N\}$ for every cell i. By $c_t[N(i)]$ we mean the pattern of the states seen on the neighborhood of cell i in configuration c_t. Mathematically, this can be seen as an element of $\mathbb{Z}_m^N$, the set of all possible assignments $p : N \to \mathbb{Z}_m$. The function $f : \mathbb{Z}_m^N \to \mathbb{Z}_m$ assigns a value $f(p) \in \mathbb{Z}_m$ to each neighborhood pattern $p : N \to \mathbb{Z}_m$. Intuitively, in order to update its state, a cell i applies a function f on the current state of its neighbors, and depending on the result, permutes the state it used to have one step before.

Notice that any automaton that is defined this way is automatically reversible; one can retrace an orbit $\ldots, c_{t-1}, c_t, c_{t+1}, \ldots$ backwards using the rule

$$c_{t-1}[i] = c_{t+1}[i] - f(c_t[N(i)]) . \tag{9}$$

Now, suppose that the neighborhood N is symmetric, meaning that $k \in N$ if and only if $-k \in N$. Suppose further that we find a function $g : \mathbb{Z}_m^N \to \mathbb{Z}_m$ of the form

$$g(p) \triangleq \sum_{k \in N} \beta_k p[k] \tag{10}$$

($\beta_k \in \mathbb{Z}_m$) that has the following two properties:

i) it is symmetric; that is, $\beta_{-k} = \beta_k$ for every $k \in N$, and
ii) it is orthogonal to f, in the sense that $f(p)g(p) = 0$ for every $p \in \mathbb{Z}_m^N$.

Let us denote by $\mathbf{0}$ the configuration in which every cell is in state 0. For simplicity, let us assume that $\ldots, \mathbf{0}, \mathbf{0}, \mathbf{0}, \ldots$ is a valid orbit of the automaton. Equivalently, this means that f maps the uniformly-0 pattern into 0.[2] We call a configuration finite if only a finite number of cells have non-zero states in it.

Let $\ldots, c_{t-1}, c_t, c_{t+1}, \ldots$ be an arbitrary orbit consisting of finite configurations. We claim that the value of the sum

$$\Theta(c_{t-1}, c_t) \triangleq \sum_{i \in \mathbb{Z}^d} c_{t-1}[i]\, g\left(c_t[N(i)]\right) \tag{11}$$

is independent of the time t.

From (8) and property (ii) we can write

$$\left(c_{t+1}[i] - c_{t-1}[i]\right) g\left(c_t[N(i)]\right) = 0\,. \tag{12}$$

Summing over all cells i, we obtain

$$\sum_{i \in \mathbb{Z}^d} c_{t-1}[i]\, g\left(c_t[N(i)]\right) = \sum_{i \in \mathbb{Z}^d} c_{t+1}[i]\, g\left(c_t[N(i)]\right)\,. \tag{13}$$

By the symmetry of g (property (i)) we can rewrite the right hand side as follows:

$$\begin{aligned}
\sum_{i \in \mathbb{Z}^d} c_{t+1}[i]\, g\left(c_t[N(i)]\right) &= \sum_{i \in \mathbb{Z}^d} \sum_{k \in N} \beta_k c_{t+1}[i] c_t[i+k] && (14)\\
&= \sum_{i' \in \mathbb{Z}^d} \sum_{k' \in N} \beta_{-k'} c_{t+1}[i'+k'] c_t[i'] && (15)\\
&= \sum_{i' \in \mathbb{Z}^d} \sum_{k' \in N} \beta_{k'} c_{t+1}[i'+k'] c_t[i'] && (16)\\
&= \sum_{i' \in \mathbb{Z}^d} c_t[i']\, g\left(c_{t+1}[N(i')]\right)\,. && (17)
\end{aligned}$$

Therefore, we obtain that

$$\begin{aligned}
\Theta(c_{t-1}, c_t) &= \sum_{i \in \mathbb{Z}^d} c_{t-1}[i]\, g\left(c_t[N(i)]\right) && (18)\\
&= \sum_{i \in \mathbb{Z}^d} c_t[i]\, g\left(c_{t+1}[N(i)]\right) && (19)\\
&= \Theta(c_t, c_{t+1})\,, && (20)
\end{aligned}$$

proving the claim.

Yet another beautiful example is the following discrete model of an excitable medium, due to Greenberg and Hastings [31] (see the relevant part in Chapter **??**,

[2] We could avoid this requirement by formulating our conservation law in terms of spatially periodic configurations.

and also [30, 29]). The CA runs on a two-dimensional board. Each cell is either "at rest" (state □), or "excited" (state ■), or in a "refractory phase" (state ▩). A cell that is at rest remains so unless it is "stimulated" by one or more of its four neighbors (i.e., if at least one of its neighbors is excited). An excited cell undergoes a 1-step refractory phase before going back to rest and starting to respond to stimulations again. Typically, a configuration of the infinite board contains a number of "singularities" with waves continuously swirling around them (Figure 4). The singularities are never created or destroyed. Therefore, the number of such singularities remain constant throughout time. To be more precise, the singularities are the 2×2 blocks of cells with states ▩□, ■■/▩□, or ▩▩/□□, or the rotations or mirror images of these blocks. One can easily verify that a singular 2×2 block remains singular after one step, and a non-singular block remains non-singular.

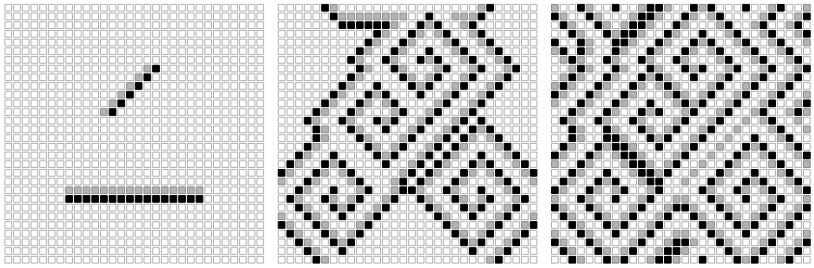

(a) (b) (c)

Fig. 4 Simulation of Greenberg-Hastings model on a spatially periodic configuration. (a) The initial configuration. (b) The configuration at time $t = 10$. (c) The configuration at time $t = 60$.

Sometimes one can find an energy-like function that is not perfectly conserved by the evolution of a CA, yet whose total value is never increased (or never decreased) with time. Physically, such a situation is comparable with a system that is isolated from its environment, except that it may dissipate heat, resulting a decrease in its total energy. Mathematically, a non-increasing energy function may be helpful in studying stability properties of a CA.

As an example, in the Just Gliders CA that we discussed before, the number of left-moving particles is never increased, though it may decrease. Formally, we have

$$\sum_{i=-\infty}^{+\infty} \theta_L(x[i]) \geq \sum_{i=-\infty}^{+\infty} \theta_L((Fx)[i]) \ , \tag{21}$$

where

$$\theta_L(a) \triangleq \begin{cases} 1 & \text{if } a = \blacktriangleleft, \\ 0 & \text{otherwise} \end{cases} \tag{22}$$

for any configuration x with a finite number of particles on it. Equivalently, we can write

$$\frac{\sum_{i=1}^{p} \theta_L(x[i])}{p} \geq \frac{\sum_{i=1}^{p} \theta_L((Fx)[i])}{p}, \tag{23}$$

for every spatially periodic configuration x with period $p > 0$. In the Traffic CA, it is easy to verify that the number of blocks ■■ of two consecutive cars is never increased; that is,

$$\sum_{i=-\infty}^{+\infty} \theta'(x[i,i+1]) \geq \sum_{i=-\infty}^{+\infty} \theta'((Fx)[i,i+1]), \tag{24}$$

where

$$\theta'(ab) \triangleq \begin{cases} 1 & \text{if } ab = \blacksquare\blacksquare, \\ 0 & \text{otherwise.} \end{cases} \tag{25}$$

A more interesting example is the *Sand Pile* model due to Bak, Tang and Wiesenfeld [4, 5]. On each cell of the board, there is a stack of sand grains. We can consider the height of this stack as the state of the cell. So, each cell $i \in \mathbb{Z}^d$ has a state $h[i] \in \{0,1,\ldots,N\}$. (To keep the state set finite, we have assumed that each cell may contain no more than $N > 0$ grains.) If the difference between the height of the stacks in two adjacent cells is more than a threshold $K \geq 4d$, one grain of sand from the higher cell tumbles down onto the lower cell. More precisely, if $h : \mathbb{Z}^d \to \{0,1,\ldots,N\}$ is a configuration of the automaton, its configuration after one step would be $Fh : \mathbb{Z}^d \to \{0,1,\ldots,N\}$, where

$$\begin{aligned}(Fh)[i] = h[i] &- |\{j \in N(i) : h[i] - h[j] \geq K\}| \\ &+ |\{j \in N(i) : h[j] - h[i] \geq K\}|.\end{aligned} \tag{26}$$

Here $N(i)$ denotes the set of immediate neighbors of cell i (Figure 5). Clearly, the

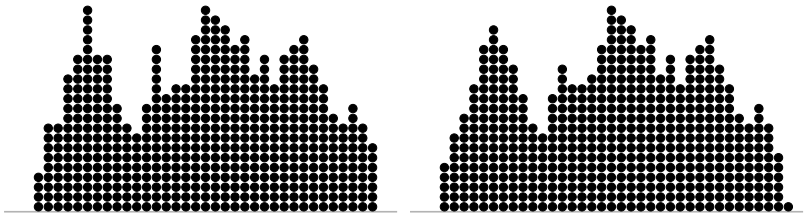

(a) (b)

Fig. 5 Two consecutive configurations of the one-dimensional Sand Pile model with parameter $K = 5$. (a) The initial configuration. (b) The configuration after one step.

total number of sands on a finite configuration is never changed; that is,

$$\sum_{i\in\mathbb{Z}^d} (Fh)[i] = \sum_{i\in\mathbb{Z}^d} h[i]\,. \tag{27}$$

A bit less trivial is the fact that the sum of the squares of the number of sands on the cells is non-increasing:

$$\sum_{i\in\mathbb{Z}^d} ((Fh)[i])^2 \leq \sum_{i\in\mathbb{Z}^d} (h[i])^2\,. \tag{28}$$

To see this, note that the height of a single grain of sand does never increase with time. The sum of the heights of all grains on configuration h is

$$\sum_{i\in\mathbb{Z}^d} \sum_{k=1}^{h[i]} k = \sum_{i\in\mathbb{Z}^d} \frac{1}{2} h[i]\,(h[i]+1) \tag{29}$$

$$= \frac{1}{2} \sum_{i\in\mathbb{Z}^d} (h[i])^2 + \frac{1}{2} \sum_{i\in\mathbb{Z}^d} h[i]\,, \tag{30}$$

and does never increase with time. From (27) we know that the second term in the above sum remains constant with time. Therefore, the fist term cannot increase with time.[3]

2 Mathematical Formulation

In this section, we formulate, in a more precise language, what is exactly meant by a conservation law in a cellular automaton.

2.1 Cellular Automata

We restrict our discussion to cellular automata on the infinite d-dimensional square lattice. Thus the cells of the *lattice* are indexed by the elements of $\mathbb{Z}^d$. The *states* of the cells are chosen from a finite set S that contains at least two elements. A *configuration* of the lattice is a mapping $x : \mathbb{Z}^d \to S$ which assigns a state to each cell on the lattice. By a *pattern* we mean an assignment $p : D \to S$ of states to a subset $D \subseteq \mathbb{Z}^d$ of cells. A *finite* pattern is a pattern that has a finite domain. If $p : D \to S$ is a pattern and $E \subseteq D$, we write $p[E]$ to denote the restriction of p to E; that is, $p[E]$ stands for the pattern $q : E \to S$ seen over the set E under p.

[3] The original Sand Pile automaton of Bak et al. [4, 5] is more elegant in that its cell states store not the height $h[i]$, but the difference $h[i] - h[i+1]$ between the heights of two consecutive cells. This way we can represent sand piles of arbitrary height using a bounded number of cell states. As remarked by Goles [27], a similar non-increasing energy can be found for such CA, provided we allow the energy to depend on the previous configurations, too. Non-uniform non-increasing energies are used in [1, 28] to analyze Sand Pile and similar automata.

For each $k \in \mathbb{Z}^d$, σ^k denotes the *translation* by k. That is, for each pattern $p : D \to S$, $\sigma^k p$ is the pattern with $(\sigma^k p)[i] = p[k+i]$ whenever $k+i \in D$.

To specify a *cellular automaton* (CA), one further needs to specify a *neighorhood* and a *local rule*. The neighborhood is presented by a finite set $0 \in N \subseteq \mathbb{Z}^d$. The neighborhood of a cell i is the set $N(i) \triangleq \{i+k : k \in N\}$. The local rule is a function $f : S^N \to S$ that provides a new state for each cell i by looking at its neighborhood pattern. Hence, for every configuration $x : \mathbb{Z}^d \to S$, one obtains a new configuration $Fx : \mathbb{Z}^d \to S$ where

$$(Fx)[i] = f((\sigma^i x)[N]) \tag{31}$$

$$= f(\sigma^i(x[N(i)])) \tag{32}$$

for each cell i. The dynamics of the CA is realized by iterating the *global mapping* F on an initial configuration x:[4]

$$x \xrightarrow{F} Fx \xrightarrow{F} F^2x \xrightarrow{F} F^3x \xrightarrow{F} \cdots . \tag{33}$$

We often identify a cellular automaton with its global mapping.

It is useful to see the configuration space $S^{\mathbb{Z}^d}$ as a topological space. The product topology on $S^{\mathbb{Z}^d}$ is the smallest topology with respect to which all the projections $x \mapsto x[i]$ are continuous. In this topology, a sequence $\{x_t\}_{t=1}^{\infty}$ converges to a configuration x if and only if for each cell i, $x_t[i] = x[i]$ for all sufficiently large t. The space $S^{\mathbb{Z}^d}$ is compact and metrizable, and is homeomorphic to the Cantor set. A *cylinder* is a set of the form

$$[p] \triangleq \{x \in S^{\mathbb{Z}^d} : x[D] = p\} \tag{34}$$

where $p : D \to S$ is a finite pattern. Cylinders are both open and closed, and form a basis for the product topology. The global mapping of a cellular automaton is continuous with respect to the product topology. Moreover, every translation-invariant continuous mapping on $S^{\mathbb{Z}^d}$ is the global mapping of a cellular automaton [33]. The topological aspects of cellular automata are discussed in Chapter **??**.

The Borel σ-algebra on the space $S^{\mathbb{Z}^d}$ of configurations is the σ-algebra generated by the cylinders. A Borel probability measure π is completely determined by assigning a probability $0 \le \pi([p]) \le 1$ to each cylinder $[p]$ in a consistent way (cf. [53]). The space $\mathscr{M}$ of all Borel probability measures on $S^{\mathbb{Z}^d}$ can be topologized by the vague topology (aka the weak* topology) (see e.g. [64]). This space is also compact and metrizable. A sequence $\{\pi_t\}_{t=0}^{\infty}$ converges to a measure π if and only if for each cylinder $[p]$, the sequence $\{\pi_t([p])\}_{t=0}^{\infty}$ of real numbers converges to $\pi([p])$. A cellular automaton F induces a continuous mapping on $\mathscr{M}$ via $F\pi \triangleq \pi \circ F^{-1}$. A translation-invariant measure is a measure $\pi \in \mathscr{M}$ such that $\sigma^k \pi = \pi$ for every trans-

[4] This definition does not cover some of the examples in the previous section (namely, Vichniac's dynamics on the Ising model, and the 2nd order model). However, those models can be easily transformed into a standard cellular automaton as defined here (cf. [62]).

lation σ^k. The set of all translation-invariant Borel probability measures on $S^{\mathbb{Z}^d}$ is a closed and convex subspace of $\mathscr{M}$ and is denoted by $\mathscr{M}_\sigma$.

2.2 Energy

Let $S^\#$ denote the set of all finite patterns *modulo* translations (i.e., forgetting their exact positions). To be strict, an element of $S^\#$ is a class $\langle p \rangle$ of patterns that can be obtained by translating p. However, unless there is a risk of confusion, we often abuse the notations and identify a class $\langle p \rangle$ with any of its elements. We see the elements of $S^\#$ as generalized words. In particular, when $d = 1$ (i.e., on the one-dimensional lattice $\mathbb{Z}$), we have $S^* \subseteq S^\#$, where S^* stands for the set of finite words on the alphabet S. Let $\varnothing$ be the (unique) pattern with an empty domain.

A (local and additive) *energy* is specified by assigning an *interaction potential* $\theta(p) \in \mathbb{R}$ to each finite pattern $p \in S^\#$. We require that $\theta(\varnothing) = 0$, and that the set $\{p : \theta(p) \neq 0\} \subseteq S^\# \setminus \{\varnothing\}$ is finite. The latter set is called the *support* of θ and is denoted by $\mathrm{supp}(\theta)$. For a configuration $x : \mathbb{Z}^d \to S$ and a finite collection A of cells, the value $\theta(x[A])$ is interpreted as the energy resulting from the interaction of the cells in A. The *total energy* of x (whenever meaningful) is simply the sum

$$\Theta(x) \triangleq \sum_{\substack{A \subseteq \mathbb{Z}^d \\ \text{finite}}} \theta(x[A]) \,. \tag{35}$$

However, the sum (35) typically does not have a well-defined value or is infinite.

There are essentially three ways around this problem. The first approach is to consider only the difference between the energy of two configurations that are only slightly perturbed from each other (more precisely, differ on only a finite number of cells). The second approach is to work with the average or expected energy per cell in a configuration. The third approach is to avoid any global notion of energy and instead describe a conservation law in terms of the local redistribution of energy at each step. Fortunately, all these lead to equivalent concepts of a conservation law. We now discuss the first two approaches and their equivalence. The local approach is discussed later in Section 4.

We say that two configurations $x, y : \mathbb{Z}^d \to S$ are *asymptotic* if they disagree on no more than a finite number of cells. The *difference* between the energy of two asymptotic configurations x and y is defined to be

$$\delta\Theta(x, y) \triangleq \sum_{\substack{A \subseteq \mathbb{Z}^d \\ \text{finite}}} [\theta(y[A]) - \theta(x[A])] \,. \tag{36}$$

Let us remark how the *locality* and *additivity* of energy translate in this framework. The *interaction range* of an energy θ can be identified by a minimal neighborhood $0 \in M \subseteq \mathbb{Z}^d$ such that, for every pattern $p : D \to S$ in the support of θ and

every $i \in D$, we have $D \subseteq M(i)$. Since $\mathrm{supp}(\theta)$ is finite, the interaction range of θ is also finite.

If two patterns $p : D \to S$ and $q : E \to S$ agree on the intersection $D \cap E$ of their domains (in particular, if $D \cap E = \varnothing$), we can merge them together and obtain a pattern $p \vee q : D \cup E \to S$ that agrees with p and q on their domains. The *boundary* of a set A (with respect to the neighborhood M) is the set $\partial M(A) \triangleq M(A) \setminus A$.

Proposition 1 (Locality). *Let $p, q : D \to S$, $r : \partial M(D) \to S$, and $u, v : \mathbb{Z}^d \setminus M(D) \to S$ be arbitrary patterns. Then,*

$$\delta\Theta(p \vee r \vee u, q \vee r \vee u) = \delta\Theta(p \vee r \vee v, q \vee r \vee v) \,. \tag{37}$$

Proposition 2 (Additivity). *Let $D, E \subseteq \mathbb{Z}^d$ be two finite non-empty sets such that $M(D) \cap E = D \cap M(E) = \varnothing$. Let $p_0, p_1 : D \to S$, $q_0, q_1 : E \to S$ and $w : \mathbb{Z}^d \setminus D \setminus E \to S$ be arbitrary patterns. Then,*

$$\begin{aligned} \delta\Theta(p_0 \vee q_0 \vee w, p_1 \vee q_1 \vee w) = {} & \delta\Theta(p_0 \vee q_0 \vee w, p_1 \vee q_0 \vee w) \\ & + \delta\Theta(p_0 \vee q_0 \vee w, p_0 \vee q_1 \vee w) \,. \end{aligned} \tag{38}$$

A *local observable* (or a locally observable property) is a mapping $\mu : S^{\mathbb{Z}^d} \to \Gamma$ (Γ being an arbitrary set) that depends only on the states of a finite number of cells. That is, μ is a local observable if there is a finite neighborhood $W \subseteq \mathbb{Z}^d$ (the observation *window*) and a local rule $g : S^W \to \Gamma$ such that $\mu(x) = g(x[W])$. Observing a configuration x *around* a cell i one gets the value $\mu(\sigma^i x) = g(x[W(i)])$.

Every real-valued local observable μ with local rule $g : S^W \to \mathbb{R}$ defines an interaction potential θ via

$$\theta(p) \triangleq \begin{cases} g(p) & \text{if } p \in S^W \,, \\ 0 & \text{otherwise.} \end{cases} \tag{39}$$

The energy difference $\delta\Theta(x, y)$ can then be calculated using

$$\delta\Theta(x, y) = \sum_{i \in \mathbb{Z}^d} \left[\mu(\sigma^i y) - \mu(\sigma^i x) \right] \,. \tag{40}$$

Conversely, given an interaction potential θ, we can construct a real-valued local observable μ that generates $\delta\Theta$ via (40). For example, we can choose the interaction range M of θ as the observation window and define

$$\mu_\theta(x) \triangleq \sum_{0 \in A \subseteq M} \frac{1}{|A|} \theta(x[A]) \,. \tag{41}$$

Hence, one can equivalently specify an energy using a real-valued local observable.

Consider an energy which is formalized using a local observable μ. For every $n \geq 0$, let $I_n \triangleq [-n, n]^d$ be the centered hyper-cube of size $(2n+1)^d$ on the lattice. The *average* energy per cell in a configuration x is obtained by taking the limit of the finite averages

$$\frac{\sum_{i\in I_n}\mu(\sigma^i x)}{|I_n|} \tag{42}$$

when $n\to\infty$. Since the limit does not always exist, we use the *upper* or *lower* average energy per cell

$$\overline{\mu}(x) \triangleq \limsup_{n\to\infty} \frac{\sum_{i\in I_n}\mu(\sigma^i x)}{|I_n|}\,, \tag{43}$$

$$\underline{\mu}(x) \triangleq \liminf_{n\to\infty} \frac{\sum_{i\in I_n}\mu(\sigma^i x)}{|I_n|}\,. \tag{44}$$

For every translation-invariant Borel probability measure $\pi\in\mathscr{M}_\sigma$ on $S^{\mathbb{Z}^d}$, we can also define the *expected* energy per cell

$$\pi(\mu) \triangleq \int \mu \mathrm{d}\pi \triangleq \sum_{p:W\to S} g(p)\pi([p])\,. \tag{45}$$

The mapping $\pi\mapsto\pi(\mu)$ is uniformly continuous. According to the pointwise ergodic theorem (see e.g. [64]), for every ergodic measure $\pi\in\mathscr{M}_\sigma$ (i.e., ergodic with respect to σ) and π-almost every configuration $x\in S^{\mathbb{Z}^d}$, we have $\overline{\mu}(x)=\underline{\mu}(x)=\pi(\mu)$. Ergodic measures are the extremal points of the convex set $\mathscr{M}_\sigma$.

2.3 Conservation of Energy

We can now formulate conservation laws in several different, but equivalent ways. We say that a cellular automaton F *conserves* an energy defined in terms of an interaction potential θ or an observable μ if any of the following equivalent conditions hold.

Theorem 1 ([32, 14, 9, 54, 10, 40, 17, 48]). *Let $F:S^{\mathbb{Z}^d}\to S^{\mathbb{Z}^d}$ be a cellular automaton and $\theta:S^{\#}\to\mathbb{R}$ an interaction potential. Let $\mu:S^{\mathbb{Z}^d}\to\mathbb{R}$ be a local observable that generates the same energy as θ. Let $\diamondsuit:\mathbb{Z}^d\to S$ be an arbitrary configuration. The following conditions are equivalent.*

C-1 $\delta\Theta(Fx,Fy)=\delta\Theta(x,y)$ *for every two asymptotic configurations x and y.*
C-2 $\delta\Theta(F\diamondsuit,Fx)=\delta\Theta(\diamondsuit,x)$ *for every configuration x asymptotic to $\diamondsuit$.*
C-3 $\overline{\mu}(Fx)=\overline{\mu}(x)$ *for every configuration x.*
C-4 $\overline{\mu}(Fx)=\overline{\mu}(x)$ *for every periodic configuration x.*
C-5 $(F\pi)(\mu)=\pi(\mu)$ *for every translation-invariant probability measure π.*
C-6 $(F\pi)(\mu)=\pi(\mu)$ *for every ergodic translation-invariant probability measure π.*

Condition (C-1) states that the CA preserves the difference between the energy of every two asymptotic configurations, while condition (C-2) requires that the energy of a configuration relative to a fixed orbit remains constant. Condition (C-3)

means that the (upper) average energy per cell of each configuration remains unchanged with time. Condition (C-3) only requires the CA to preserve the average energy per cell of the periodic configurations. Conditions (C-5) and (C-6) express a conservation law in terms of the expected energy per cell.

The equivalence of the conditions (C-1) and (C-2) is an immediate consequence of the locality of energy. To prove that (C-1) implies (C-3), one first shows that $\overline{\mu}(Fz) = \overline{\mu}(z)$ in the special case that z is a uniform configuration. An arbitrary configuration x can then be approximated by configurations x_n that agree with x on larger and larger centered hyper-cubes I_n and have a fixed state s on all the other cells. Similarly, the implication (C-4) $\Rightarrow$ (C-1) follows from the locality of energy, by considering periodic configurations $\hat{x}$ and $\hat{y}$ that agree with x and y on a sufficiently large hyper-cube I_n. That (C-3) implies (C-6) follows from the ergodic theorem. The implication (C-6) $\Rightarrow$ (C-5) is a consequence of the fact that every translation-invariant measure π is a limit of convex combinations of ergodic measures (cf. [64]). To see that (C-5) implies (C-4), for every periodic configuration x, consider a measure π_x whose probability mass is concentrated on the σ-orbit of x. More details can be found, for example, in [58]. Yet other characterizations of conservation laws can be found in [54, 17].

A uniform configuration $\diamondsuit$ for which $F\diamondsuit = \diamondsuit$ is called a *quiescent* configuration. For example, in the Traffic CA, the configuration with no car is quiescent; so is the configuration with cars on every cell. There is often a distinguished quiescent configuration $\diamondsuit$, which is seen as the "background" or the "vacuum", and one is interested in the evolution of the configurations that are asymptotic to $\diamondsuit$. A configuration that is asymptotic to $\diamondsuit$ is then called a *finite* configuration. The image of a finite configuration under F is obviously finite.

Choosing the interaction potential θ properly, one can ensure that this distinguished quiescent configuration $\diamondsuit$ has zero total energy. Then every finite configuration x would have a finite well-defined total energy $\Theta(x) = \delta\Theta(\diamondsuit, x)$. Hence, we can express the conservation law in terms of the total energy of the finite configurations: the CA F conserves θ if and only if

C-7 $\Theta(Fx) = \Theta(x)$ for every finite configuration x.

2.4 Dissipation of Energy

Recall, from the previous section, the Sand Pile example, in which we could find an energy that was not conserved, but its value was never increased. We now formalize such *dissipation laws*. We say that an energy θ is *dissipative* (or non-increasing) under a CA F if any of the following equivalent conditions hold.

Theorem 2 ([40, 48, 7]). *Let $F : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ be a cellular automaton and $\theta : S^{\#} \to \mathbb{R}$ an interaction potential. Let $\mu : S^{\mathbb{Z}^d} \to \mathbb{R}$ be a local observable that generates the same energy as θ. The following conditions are equivalent.*

D-1 $\overline{\mu}(Fx) \leq \overline{\mu}(x)$ for every configuration x.

D-2 $\overline{\mu}(Fx) \leq \overline{\mu}(x)$ *for every periodic configuration* x.
D-3 $(F\pi)(\mu) \leq \pi(\mu)$ *for every translation-invariant probability measure* π.
D-4 $(F\pi)(\mu) \leq \pi(\mu)$ *for every ergodic translation-invariant probability measure* π.

Furthermore, if $\diamondsuit : \mathbb{Z}^d \to S$ *is a quiescent configuration, the following condition is also equivalent to the above.*

D-5 $\delta\Theta(\diamondsuit, Fx) \leq \delta\Theta(\diamondsuit, x)$ *for every finite configuration* x.

The proofs are similar to those for conservation laws. In case the quiescent configuration $\diamondsuit$ has total energy zero, the condition (C-5) takes the following concise form:

D-6 $\Theta(Fx) \leq \Theta(x)$ for every finite configuration x.

3 Algorithmics

In this section, we discuss two algorithmic questions related to conservation laws in cellular automata. We first show how one can verify the validity of a conservation law algorithmically. Next, we discuss the more difficult question of finding conservation laws in a cellular automaton. Though one can enumerate all the possible candidates and verify their validity one by one, there is no algorithm to tell beforehand, whether a cellular automaton has any non-trivial conservation law or not, even when restricted to one-dimensional CA.

Dissipative energies are more complicated to recognize. Fortunately, in one dimension, one can still decide whether a given energy is dissipative or not. The question is, however, undecidable in the higher-dimensional case.

3.1 Conservation Laws

Suppose that we are given a cellular automaton F and an interaction potential θ, and we are asked whether F conserves the energy defined by θ. Verifying either of the characterizations (C-1)–(C-6) requires establishing infinitely many equalities. However, as was first noticed by Hatori and Takesue [32], it suffices to verify the equality

$$\delta\Theta(Fx, Fy) = \delta\Theta(x, y) \tag{46}$$

for all x and y that differ on exactly one cell.

Indeed, suppose that x and y are two configurations that disagree only on a finite set D of cells. One can then find a sequence

$$x = x_0, x_1, x_2, \ldots, x_n = y \tag{47}$$

of configurations, where $n = |D|$, such that x_i and x_{i+1} disagree on exactly one cell. If (46) holds whenever two configurations differ on a single cell, we can write

$$\delta\Theta(x,y) = \sum_{i=0}^{n-1} \delta\Theta(x_i, x_{i+1}) \tag{48}$$

$$= \sum_{i=0}^{n-1} \delta\Theta(Fx_i, Fx_{i+1}) \tag{49}$$

$$= \delta\Theta(Fx, Fy)\ . \tag{50}$$

(Note that $\delta\Theta(a,c) = \delta\Theta(a,b) + \delta\Theta(b,c)$ for every a,b,c.)

Proposition 3 ([32]). *Let $F : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ be a cellular automaton and $\theta : S^{\#} \to \mathbb{R}$ an interaction potential. The following conditions are equivalent.*

C-1 $\delta\Theta(Fx,Fy) = \delta\Theta(x,y)$ *for every two asymptotic configurations x and y.*

C-8 $\delta\Theta(Fx,Fy) = \delta\Theta(x,y)$ *for every two configurations x and y that disagree on exactly one cell.*

Since θ (and hence $\delta\Theta$) is translation-invariant, we can, in fact, consider only configurations x and y that disagree on cell 0.

Now let $0 \in M \subseteq \mathbb{Z}^d$ be the interaction range of θ, and let $0 \in N \subseteq \mathbb{Z}^d$ be the neighborhood of F. For every two configurations x and y that disagree only on cell 0, the value $\delta\Theta(x,y)$ depends only on the state of the cells in x and y that are in the neighborhood $M(0)$. Similarly, since Fx and Fy may disagree only on the set $N^{-1}(0)$ (i.e., the set of cells that have cell 0 as neighbor) the value $\delta\Theta(Fx,Fy)$ depends on the state of the cells in Fx and Fy that are in the set $M(N^{-1}(0))$. Therefore, condition (C-8) reduces to a finite number of equalities, which can each be verified in a finite time.

The following algorithm exploits the above discussion to answer whether F conserves θ. Let $f : S^N \to S$ be the local rule of F. There is a natural way to extend the application of f to any finite pattern $p : N(A) \to S$. Namely, the image of p under f is a pattern $f(p) : A \to S$, where $f(p)[i] \triangleq f((\sigma^i p)[N])$ for every $i \in A$. For every finite pattern $p : D \to S$, let us use the shorthand

$$\Theta(p) \triangleq \sum_{A \subseteq D} \theta(p[A])\ . \tag{51}$$

let $W = M(N^{-1})$;
for every pattern $p : N(W) \to S$ and every state $s \in S$,
    let $q : N(W) \to S$ and
        set $q[0] = s$ and $q[i] = p[i]$ for $i \neq 0$;
    let $p', q' : W \to S$ and
        set $p' = f(p)$ and $q' = f(q)$;
    set $\delta\Theta = \Theta(q) - \Theta(p)$;

```
        set δΘ′ = Θ(q′) − Θ(p′);
        if δΘ′ ≠ δΘ,
                return “no” and halt;
return “yes”;
```

There are few ways to improve the efficiency of the above algorithm. Let us order the cells on the lattice lexicographically. Namely, in this ordering a cell $i = (i_1, i_2, \ldots, i_d)$ precedes a cell $j = (j_1, j_2, \ldots, j_d)$, written $i \prec j$, if there is a $1 \leq k \leq d$ such that $i_k < j_k$ and $i_l = j_l$ for all $1 \leq l < k$. Let $\diamond \in S$ be a fixed state. Using a more clever argument, we can restrict the main loop of the algorithm to all patterns $p : N(W) \to S$ for which $p[i] = \diamond$ for all cells $i \succeq 0$.

The set of all interaction potentials θ that are conserved by F is a linear space $\mathscr{W}_F$. It is interesting to note that, given a finite set $P \subseteq S^{\#}$, one can use a variant of the above algorithm to construct the space $\mathscr{W}_F[P]$ of all those interaction potentials conserved by F that have support P. To do so, one considers $\theta : P \to \mathbb{R}$ as a vector of unknowns, and collects all the equations $\delta\Theta' = \delta\Theta$ given by the algorithm. The space $\mathscr{W}_F[P]$ is simply the solution space of this system of equations.

In [17], a different approach has lead to another efficient algorithm for verifying the validity of conservation laws.

As we already mentioned at the beginning of this section, given a CA F, we can enumerate all the candidate energies θ and verify whether they are conserved by F.[5] Though every conservation law for F is eventually discovered by this algorithm, the algorithm does never end, and there is no way to predict whether, continuing to run the algorithm, we are going to find any new conservation law. In fact, it is undecidable whether a given CA has any non-trivial conservation law or not.

Proposition 4. *There is no algorithm that, given a cellular automaton, decides whether it has any non-trivial conservation law or not.*

This is an immediate consequence of the undecidability of the finite tiling problem. The finite tiling problem is a variant of the tiling problem, which asks whether one can tile the entire plane using decorated tiles chosen from a finite number of given types (see Chapter **??**).

A *Wang tile*, called after the Chinese-American mathematician Hao Wang (1921–1995), is a unit square with colored edges. Two Wang tiles can be placed next to each other if their abutting edges have the same color. A set of Wang tiles is given by a finite set T, and four mappings $n, w, s, e : T \to C$ that identify the colors of the northern, western, southern and eastern edges of the tiles. A *valid tiling* of the

[5] Strictly speaking, one cannot enumerate all the energies, since the set of real-valued energies is uncountable. However, note that for each energy θ, the set of possible values that $\delta\Theta$ can take is a finitely generated subgroup of $\mathbb{R}$, and hence isomorphic to $\mathbb{Z}^m$ for some $m > 0$. Our discussion merely uses the group structure of the energy values. Therefore, we can equivalently work with energies whose values are in $\mathbb{Z}^m$. Such energies can indeed be enumerated.

plane is a configuration $c : \mathbb{Z}^2 \to T$ that respects the tiling rule; that is, $n(c[i, j]) = s(c[i, j+1])$ and $w(c[i, j]) = e(c[i-1, j])$ for every $i, j \in \mathbb{Z}$.

The finite tiling problem asks, given a set T of Wang tiles and a designated *blank* tile $\diamond \in T$, whether there is a non-trivial *finite* valid tiling of the plane; that is, a valid tiling in which all but a finite number of tiles are blank. Here by a *non-trivial* tiling we mean a tiling that uses at least one non-blank tile. Reducing the halting problem of Turing machines to this problem, one can show that the finite tiling problem is undecidable.

Given a tile set T and designated blank tile $\diamond \in T$, let us construct a two-dimensional cellular automaton F as follows: the state set of F is T. On every configuration $c : \mathbb{Z}^2 \to T$, the CA looks at each cell i and its four immediate neighbors. If there is a tiling error on cell i (i.e., if the tile on position i does not match with at least one of its adjacent tiles), the CA changes the state of i to blank. Otherwise, the state of cell i is kept unchanged.

Let $\diamondsuit$ be the uniformly blank configuration. If the tile set T admits no non-trivial finite valid tiling, every finite configuration of F is eventually turned into $\diamondsuit$. Therefore, following the characterization (C-2), in this case F has no non-trivial conservation law. On the other hand, suppose that T admits a non-trivial finite valid tiling c. Let p be the finite pattern consisting of all the non-blank cells of c, as well as a margin of blank cells around them. Clearly F preserves the number of the occurrences of p on any configuration c. In summary, T admits a non-trivial finite valid tiling if and only if F has a non-trivial conservation law. Since no algorithm can decide whether T admits a non-trivial finite valid tiling, no algorithm can either tell whether F has a non-trivial conservation law.

Note that the above argument does not rule out the existence of an algorithm that solves the problem only for one-dimensional cellular automata. Unfortunately, even when restricted to one-dimensional CA, the question remains undecidable.

Theorem 3 ([21]). *There is no algorithm that, given a one-dimensional cellular automaton, decides whether it has any non-trivial conservation law or not.*

The proof of this fact relies on the undecidability of the existence of periodic orbits in 2-counter machines [8]. A 2-counter machine is a finite automaton equipped with two unbounded counters. At each step, the automaton can increase or decrease either of the counters and check whether its value is zero. It is well-known that 2-counter machines are equivalent, in power, with Turing machines (see e.g. [46]); any algorithm can be implemented by a suitable 2-counter machine. As it is proven in [8], there is no algorithm that decides whether a given 2-counter machine has a configuration whose orbit is periodic.

To prove Theorem 3, one (algorithmically) transforms the problem of deciding whether a given 2-counter machine has a periodic configuration to the problem of deciding whether a given cellular automaton has a non-trivial conservation law. Since any algorithm solving the latter problem would lead to an algorithm solving the former, we conclude that no algorithm can solve the latter problem.

The idea of the transformation is simple: given a 2-counter machine M, one constructs a one-dimensional cellular automaton F with a distinguished quiescent configuration $\diamondsuit$ with the property that

i) if M has no periodic orbit, F eventually transforms each finite configuration to the quiescent configuration $\diamondsuit$, while
ii) if M has a periodic orbit, F has a non-trivial conservation law.

Let $\diamond$ be the state of the cells in $\diamondsuit$. Every finite configuration in F is uniquely partitioned into disjoint segments. Each segment is either syntactically valid, in which case it simulates M on some configuration, or it is not syntactically valid, in which case the CA gradually turns all its cells into $\diamond$. The size of a valid segment, however, remains unchanged. So if at some point the simulation requires more cells than available in the segment, the segment overflows and becomes syntactically invalid.

Now, if the machine M does not have any periodic orbit, every valid segment eventually overflows. Hence, every finite configuration eventually changes into $\diamondsuit$. On the other hand, if M does have a periodic configuration, a segment simulating M on such a configuration does never overflow. Therefore, one can construct a non-trivial energy that counts the number of such simulation segments, and that is conserved by F. See [21, 58] for the details.

3.2 Dissipation Laws

Next, let us discuss the problem of verifying whether a given energy is non-increasing under the iteration of a given cellular automaton. While there is an algorithm that solves the problem for one-dimensional CA [40] (see also [48, 7, 6]), the problem is undecidable in higher dimensions [7].

One-dimensional CA have a convenient representation (up to composition with 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 $V \triangleq S^k$ and edge set $E \triangleq S^{k+1}$, where for every $a, b \in S$ 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 every $k \geq l + r$, the CA can be represented on the De Bruijn graph $B_k[S]$ with labeling $\lambda : E \to S^{k-(l+r)}$ that 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 $v_i = f(u_{i-l} u_{i-l+1} \cdots u_{i+r})$.

The edge sequence $p = \{p[i]\}_{i \in \mathbb{Z}}$ of a bi-infinite walk on $B_k[S]$ represents a unique (up to translation) configuration $c : \mathbb{Z} \to S$, while its label sequence $\lambda(p) \triangleq \{\lambda(p[i])\}_{i \in \mathbb{Z}}$ represents Fc. Conversely, every configuration $c : \mathbb{Z} \to S$ corresponds to a unique bi-infinite walk on $B_k[S]$.

Let us now consider an energy defined in terms of a local observable $\mu : S^{\mathbb{Z}} \to \mathbb{R}$. Without loss of generality, we can assume that μ has observation window $[0, m] = \{0, 1, \ldots, m\}$, for some $m \geq 0$, and local rule $g : S^{[0,m]} \to \mathbb{R}$. If $k \geq l + m + r$, we can also represent the energy μ on $B_k[S]$ in a suitable way. To each edge $u_0 u_1 \cdots u_k \in$

S^{k+1}, let us assign two real numbers

$$\alpha(u_0u_1\cdots u_k) \triangleq g(u_0u_1\cdots u_m) \tag{52}$$

and

$$\beta(u_0u_1\cdots u_k) \triangleq g(v_lv_{l+1}\ldots v_{l+m})\,, \tag{53}$$

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 average energy per cell of a periodic configuration c can be calculated by averaging the value of α over the edges of the closed walk p corresponding to c on the graph. The average of β over the same closed walk is the average energy per cell of the configuration Fc. Therefore, the energy μ is non-increasing under F if and only if

$$\frac{\beta(p_1)+\beta(p_2)+\cdots+\beta(p_n)}{n} \leq \frac{\alpha(p_1)+\alpha(p_2)+\cdots+\alpha(p_n)}{n} \tag{54}$$

for every closed walk $p_1p_2\cdots p_n$. It is not difficult to verify that in order to test the above property, one only needs to test (54) for every p that is a cycle (i.e., a closed walk with no repeating vertices). Since the number of cycles on $B_k[S]$ is finite, this gives rise to the following simple algorithm for testing whether the energy μ is non-increasing or not.

```
for every cycle p_1 p_2 ··· p_n on B_k[S],
    if β(p_1) + β(p_2) + ··· + β(p_n) > α(p_1) + α(p_2) + ··· + α(p_n),
        return "no" and halt;
return "yes";
```

For higher-dimensional CA, there is no algorithmic way to verify whether a given energy is dissipative. Intuitively, this means that in order for an energy to be non-increasing, cells that are very far from each other may need to collaborate. Any local increase in energy is compensated by a local decrease somewhere on the lattice, but there is no computable upper bound on the distance we may need to look to see such a decrease.

Theorem 4 ([7]). *There is no algorithm that, given a (two-dimensional) cellular automaton F and an energy θ, decides whether θ is non-increasing under F.*

The proof is via a reduction from the finite tiling problem, which was introduced after Proposition 4. Let $(T,\diamond)$ be an instance of the finite tiling problem, where T is a set of Wang tiles and $\diamond \in T$ is the distinguished blank tile. First we use the tile set P in Figure 6 to construct a tile set $S \subseteq P \times T$. The tile set S consists of all tiles (p,t) satisfying the following two conditions:

- If $p \in \{$, , ,  ,  ,  $\}$, then $t = \diamond$.
- If $p =$, then $t \neq \diamond$.

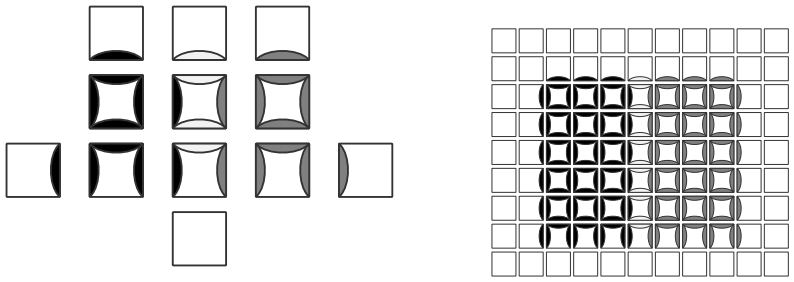

(a) (b)

Fig. 6 (a) Tile set P. (b) A typical finite valid tiling using tile set P.

The color of an edge of tile (p,t) is simply the combination of the colors of that edge in the components p and t; that is, $\alpha(p,t) \triangleq (\alpha(p),\alpha(t))$ for $\alpha = n,e,w,s$. The tile $(\square,\diamond)$ is designated as the blank tile for S.

It is easy to see that S admits a non-trivial finite valid tiling if and only if T admits a non-trivial finite valid tiling. Also, note that every non-trivial finite valid tiling using the tiles in S contains at least one occurrence of $\sqcap$. On the other hand, one can verify that if T (and hence S) does not admit a non-trivial finite valid tiling, on every finite configuration $c : \mathbb{Z}^2 \to S$, the number of occurrences of $\sqcap$ cannot exceed the number of tiling errors (i.e., those positions $i \in \mathbb{Z}^2$ where the tile on i does not match with at least one of its adjacent tiles).

Next, let us construct a two-dimensional CA F and an interaction potential θ. The CA F has S as state set. On each configuration c, the CA looks at each cell i and its four immediate neighbors. If there is a tiling error on cell i, the CA changes the state of i to blank. On the other hand, if there is no tiling error on cell i, the CA keeps the state of i unchanged unless i has a state of the form $(\sqcap,t)$ with $t \in T$. In the latter case, the state of i is turned into an arbitrary value (p',t') with $p' \notin \{\square,\sqcap\}$. The energy $\theta : S^{\#} \to \mathbb{R}$ is, in fact, context-free; that is, it assigns non-zero potentials only to the single-cell patterns. For every state $s = (p,t)$ we define

$$\theta(s) \triangleq \begin{cases} 0 & \text{if } (p,t) = (\square,\diamond), \\ 1 & \text{if } p = \sqcap, \\ 2 & \text{otherwise.} \end{cases} \tag{55}$$

For an arbitrary finite configuration c, let $\gamma(c) \triangleq \Theta(Fc) - \Theta(c)$ denote the change in the total energy of c after one step. By construction, every tiling error contributes either -1 or -2 to γ, while every correctly tiled occurrence of $\sqcap$ contributes $+1$ to γ. If T does not admit any non-trivial finite valid tiling, the number of tiling errors on any configuration c is greater than or equal to the number of occurrences of $\sqcap$ on c. Therefore, for every configuration c, $\gamma(c) \leq 0$. On the other hand, suppose that T admits a non-trivial finite valid tiling. Then, S also admits a non-trivial finite valid

tiling c. While there is no tiling error on c, there is at least one occurrence of  . Hence, $\gamma(c) > 0$.

Since there is no general algorithm for deciding whether T admits a non-trivial finite valid tiling, we conclude that no algorithm can decide whether θ is non-increasing under F. See [7, 6].

4 Flows and Particles

A conservation law, as discussed in the previous sections, is a global property of a CA. It asserts that certain local additive quantity, which we call energy, is globally preserved. It does not, however, provide any microscopic mechanism behind this. Namely, it does not elaborate how the energy is manipulated locally so that its global quantity remains intact. Microscopic explanations of conservation laws can be given in terms of "flow" of energy from one cell to another. In fact, every conservation law in cellular automata has such flow explanations, as it was first noted in [32]. However these flows are not unique.

4.1 Local Conservation Laws

In physics, every local and additive conserved quantity (such as energy, momentum, electric charge, ...) is *locally* conserved. In fact, any conservation law in physics must inevitably be expressible in a local form, in order to be compatible with the principle of relativity (cf. [18]).

In cellular automata, local conservation laws are formalized in terms of the concept of flow. The amount of *flow* from cell i to cell j on a configuration x is specified by a real number $\Phi_{i\to j}(x)$. The mapping $x, i, j \mapsto \Phi_{i\to j}(x) \in \mathbb{R}$ is required to satisfy the following natural conditions:

i) For every $i, j \in \mathbb{Z}^d$, the mapping $x \mapsto \Phi_{i\to j}(x)$ is a local observable.
ii) For every configuration x, all cells $i, j \in \mathbb{Z}^d$, and any displacement $a \in \mathbb{Z}^d$,

$$\Phi_{a+i\to a+j}(x) = \Phi_{i\to j}(\sigma^a x)\ . \tag{56}$$

iii) There is a finite set $I \subseteq \mathbb{Z}^d$ such that $\Phi_{i\to j} = 0$ unless $i - j \in I$.

Equivalently, a mapping $x, i, j \mapsto \Phi_{i\to j}(x)$ is called a flow if there exist finite sets $K, I \subseteq \mathbb{Z}^d$ and a rule $\varphi : S^K \times I \to \mathbb{R}$ such that

$$\Phi_{i\to j}(x) = \begin{cases} \varphi\left(x[K(j)], i - j\right) & \text{if } i - j \in I, \\ 0 & \text{otherwise} \end{cases} \tag{57}$$

for every $x : \mathbb{Z}^d \to S$ and $i, j \in \mathbb{Z}^d$. In summary, the amount of flow to each cell is decided locally, by looking at a finite neighborhood K of that cell. The set I is the set of directions from which energy flows to a cell.

Let $F : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ be a cellular automaton and $\mu : S^{\mathbb{Z}^d} \to \mathbb{R}$ a local observable defining an energy. We say that a flow Φ is *compatible* with μ and F (or, Φ is a flow for μ under the dynamics of F) if the following *continuity equations* hold (see Figure 7(a)):

a) For every configuration x and every cell a,

$$\mu(\sigma^a x) = \sum_{j \in \mathbb{Z}^d} \Phi_{a \to j}(x) \,. \tag{58}$$

b) For every configuration x and every cell a,

$$\sum_{i \in \mathbb{Z}^d} \Phi_{i \to a}(x) = \mu(\sigma^a F x) \,. \tag{59}$$

An energy μ is *locally conserved* by F if it has a flow under F. Conservation laws and local conservation laws are equivalent concepts in cellular automata [32]:

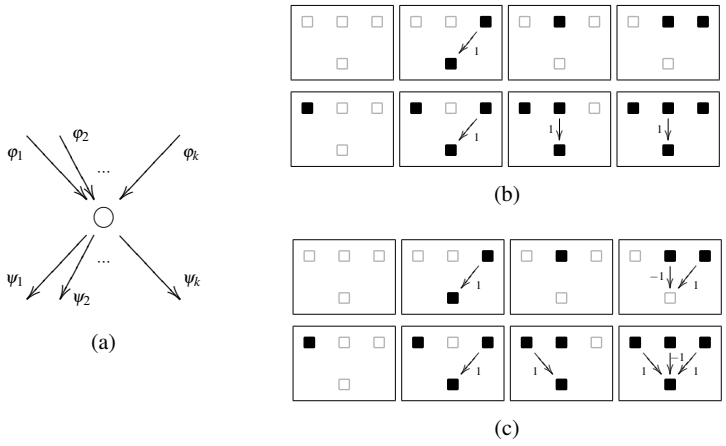

Fig. 7 (a) Continuity of the flow: $\sum_i \varphi_i = \mu = \sum_j \psi_j$. (b, c) Two different flows for the car conservation law in the Traffic CA.

Theorem 5. *Let $F : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ be a cellular automaton and $\mu : S^{\mathbb{Z}^d} \to \mathbb{R}$ a local observable. There is a flow Φ compatible with μ and F if and only if F conserves the energy generated by μ.*

That the existence of a flow compatible with μ and F implies that F conserves μ is easy to see. To prove the converse, one needs to construct, for any CA F and any

energy μ conserved by F, a flow Φ compatible with μ and F. One such construction can be found, for example, in [58]. We now illustrate the idea by an example, namely, the car conservation law in the Traffic CA.

Let x be an arbitrary configuration consisting of a finite number of cars on the highway. Let us start start from the empty highway and place the cars, one by one from left to right, on the highway to obtain x. At each step, let us identify the effect of placing one new car on position i on the configuration of the highway after one time step. We express this effect by flows from cell i to the neighboring cells (see Figure 8(a)). Note that placing a new car on the cell i may only change the state of the cells i and $i-1$ in the following configuration. Furthermore, these changes depend only on whether there has already been a car on the cell $i-1$ or not. Therefore, the effect of each car can be identified locally as depicted in Firgure 7(b). However, note that this is not the only way to identify the effect of each car. For example, if we place the cars from right to left, we obtain a different flow as illustrated in Figures 8(b) and 7(c).

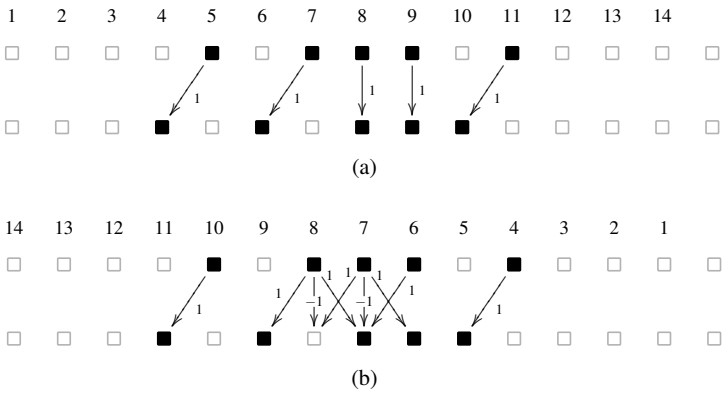

Fig. 8 The identified effect of each car on the configuration after one step. (a) Effects identified from left to right. (a) Effects identified from right to left.

Even though flows provide intuitive ways to think about conservation laws, their non-uniqueness is not plausible. Hence, one may want to identify a flow that is the most natural in some sense.

4.2 Particle Flows

Consider the following game played with pebbles on configurations of a cellular automaton. Let $x : \mathbb{Z}^d \to S$ be an arbitrary configuration. On each cell of the lattice, a number of pebbles have been placed. The number of pebbles on cell i depends only on the state of cell i and is denoted by $\eta(x[i])$. At each iteration of the CA, the

state of each cell i changes to its new value $(Fx)[i]$. The goal is to redistribute the pebbles on the lattice so that the number of pebbles on each cell matches with its new state; i.e., each cell i obtains $\eta((Fx)[i])$ pebbles. Is there a local and uniform strategy to do this? By "local" we mean that each pebble can be moved within a bounded distance of its original position, and we are only allowed to look at the states of a bounded number of cells around it in order to decide where a pebble should be moved. By "uniform" we mean that the same strategy should be used for redistributing the pebbles on every cell and over every configuration.

The assignment $\eta : S \to \mathbb{N}$ ($\mathbb{N} \triangleq \{0,1,2,\ldots\}$) defines an energy in our usual sense. A pebble can be interpreted as the "quantum" of energy; i.e., the tiniest bit of energy, which is indecomposable. The desired strategy for the game is simply a flow Φ that takes its values in the set $\mathbb{N}$ of non-negative integers, and that is compatible with η and F. For example, the flow depicted in Figures 7(b) and 8(a) defines a valid strategy, while the flow in Figures 7(c) and 8(b) does not.

Let us call a flow Φ a *particle flow* if its values are from the set of non-negative integers. Any function $\eta : S \to \mathbb{N}$ is called a *particle assignment*. A necessary condition for the existence of particle flow compatible with a particle assignment η and a cellular automaton F is, of course, that F conserves the energy generated by η. For one-dimensional CA, this condition is also known to be sufficient; Fukś [23] and Pivato [54] have shown that any particle assignment η conserved by a one-dimensional cellular automaton F has a particle flow (see also [48]). For higher-dimensional CA, however, the question is open. See [37] for a partial solution in two dimensions.

Theorem 6 ([23, 54]). *Let $F : S^{\mathbb{Z}} \to S^{\mathbb{Z}}$ be a one-dimensional cellular automaton and $\eta : S \to \mathbb{N}$ a particle assignment. There is a particle flow Φ compatible with η and F if and only if F conserves η.*

Theorem 7 ([37]). *Let $F : S^{\mathbb{Z}^2} \to S^{\mathbb{Z}^2}$ be a two-dimensional cellular automaton with radius-$\frac{1}{2}$ neighborhood, and let $\eta : S \to \mathbb{N}$ be a particle assignment. There is a particle flow Φ compatible with η and F if and only if F conserves η.*

The radius-$\frac{1}{2}$ neighborhood refers to the four-element neighborhood

$$\{(0,0),(0,1),(1,0),(1,1)\} \ . \tag{60}$$

Particle flows, if exist, are not unique; a particle flow compatible with a particle assignment and a CA can be modified in infinitely many ways to obtain new particle flows compatible with the same assignment and CA. In one dimension, however, there are natural criteria that ensure the uniqueness. For example, it is shown in [48] that for each particle assignment conserved by a one-dimensional CA, there is a unique particle flow that preserves the order of the particles.

Let us now mention an argument due to Pivato [54] that strongly suggests that Theorem 6 should be valid also for higher-dimensional CA.

Theorem 8 ([54]). *Let $F : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ be a cellular automaton with neighborhood $0 \in N \subseteq \mathbb{Z}^d$. A particle assignment $\eta : S \to \mathbb{N}$ is conserved by F if and only if*

$$\sum_{i\in A} \eta\left(x[i]\right) \leq \sum_{i\in N^{-1}(A)} \eta\left((Fx)[i]\right) \tag{61}$$

and

$$\sum_{i\in A} \eta\left((Fx)[i]\right) \leq \sum_{i\in N(A)} \eta\left(x[i]\right) \tag{62}$$

for every configuration x and every finite set $A \subseteq \mathbb{Z}^d$.

For every two consecutive configurations x and $y = Fx$, let us construct a bipartite graph $G_N[\eta, x, y] = (U, V, E)$ as follows. For every particle on x, the graph has a vertex in U, and for every particle on y, there is a vertex in V. There is an edge between a particle $u \in U$ coming from a cell i and a particle $v \in V$ coming from a cell j if and only if i is a neighbor of j; that is, if and only if $i - j \in N$.

A perfect matching in graph $G_N[\eta, x, y]$ is a way of identifying particles on x with particles on y in such a way that the position of each particle on x is a neighbor of its position on y. A necessary and sufficient condition for a (possibly infinite, but locally finite) bipartite graph to have a perfect matching is given by Hall's Marriage Theorem (see e.g. [43]): a bipartite graph $G = (U, V, E)$ has a matching that covers U if and only if for every finite set $A \subseteq U$, the number of vertices in V that are adjacent to A is at least $|A|$. If G has a matching that covers U and a matching that covers V, G also has a perfect matching.

It follows from Theorem 8 that if a CA F with neighborhood N conserves a particle assignment η, then for every configuration x, the graph $G_N[\eta, x, Fx]$ must have a perfect matching. Therefore, if F conserves η, the particles on any configuration x can be identified with those on Fx in such a way that the two positions of each particle are within bounded distance from each other. In remains open whether in general such identification can be done in a local and uniform fashion.

4.3 Flows for Dissipative Energies

In the light of Theorem 5, it is natural to ask whether non-increasing energies can also be explained locally. A flow explanation of a non-increasing energy may be possible by relaxing the continuity equations (58) and (59). In one dimension, this is always possible. In higher dimensions, however, the undecidability result of Theorem 4 imposes an obstacle to the general existence of such flows.

Let $F : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ be a cellular automaton and $\mu : S^{\mathbb{Z}^d} \to \mathbb{R}$ a local observable defining an energy. Let us say that a flow Φ is *semi-compatible* with μ and F if

a) For every configuration x and every cell a,

$$\mu(\sigma^a x) \geq \sum_{j\in\mathbb{Z}^d} \Phi_{a\to j}(x)\ . \tag{63}$$

b) For every configuration x and every cell a,

$$\sum_{i\in\mathbb{Z}^d} \Phi_{i\to a}(x) \geq \mu(\sigma^a Fx)\ . \tag{64}$$

Proposition 5. *There exists a two-dimensional cellular automaton $F : S^{\mathbb{Z}^2} \to S^{\mathbb{Z}^2}$ and a local observable $\mu : S^{\mathbb{Z}^2} \to \mathbb{R}$ such that the energy generated by μ is non-increasing under F, but there is no flow semi-compatible with μ and F.*

As mentioned above, this is a consequence of Theorem 4. First, note that there is a semi-algorithm[6] that recognizes those energies that are not non-increasing: given an energy μ and a CA F, such a semi-algorithm tests the inequality $\overline{\mu}(Fx) \leq \overline{\mu}(x)$ for periodic configurations x with larger and larger periods. According to the characterization (D-2), if μ is not non-increasing, the inequality fails on some periodic configuration, which will eventually be found. Now, suppose that every non-increasing energy μ has a semi-compatible flow. Then, one can construct a semi-algorithm that recognizes those energies that are non-increasing: given an energy μ and a CA F, one simply enumerates all the possible flows and verify equations (63) and (63) for them one by one. If μ is non-increasing, the process will eventually find a semi-compatible flow. Running these two semi-algorithms in parallel one obtains an algorithm for deciding whether a given energy is non-increasing; hence contradicting Theorem 4.

The above argument does not provide any explicit example of a non-increasing energy with no semi-compatible flow. An explicit construction has been obtained by A. Rumyantsev (see [6]).

In one-dimensional case, every non-increasing energy has a flow explanation.

Theorem 9. *Let $F : S^{\mathbb{Z}} \to S^{\mathbb{Z}}$ be a one-dimensional cellular automaton and $\mu : S^{\mathbb{Z}} \to \mathbb{R}$ a local observable. There is a flow Φ semi-compatible with μ and F if and only if the energy generated by μ is non-increasing under F.*

A proof in a rather different setting is given in Chapter **??**. In one dimension, in addition, any non-increasing energy defined using a particle assignment has a semi-compatible particle flow.

Theorem 10 ([48]). *Let $F : S^{\mathbb{Z}} \to S^{\mathbb{Z}}$ be a one-dimensional cellular automaton and $\eta : S \to \mathbb{N}$ a particle assignment. If η is non-increasing under F, there is a particle flow Φ semi-compatible with η and F.*

5 Further Topics

Let us conclude this chapter with some miscellaneous remarks.

A large body of research done about conservation laws is concentrated on number-conserving automata (see e.g. [9, 23, 10, 17, 19, 48, 7, 11, 24]). In a *number-conserving* CA, the state of a cell represents the number of particles in that cell, and

[6] By a semi-algorithm we mean an algorithmic process that does not necessarily halt on every input.

the dynamics is such that the total number of particles is preserved. Most results about number-conserving CA have counterparts in the general setting of conservation laws, which can be proven in more or less similar manners.

There are, however, results that are specific to the framework of number-conserving CA. For example, in [47], Moreira has shown that every one-dimensional CA can be simulated by a number-conserving one. In particular, this implies that there are number-conserving one-dimensional CA that are intrinsically universal (cf. Chapter **??**). Morita and Imai have constructed computationally universal reversible CA that conserve a natural quantity associated to each cell [49].

An interesting issue, which needs further investigation, is the possible connections between the conservation laws of a CA and its dynamical properties. Specifically, it is interesting to know what kind of restrictions a conservation law may impose on the dynamics of a CA? Dynamical properties of number-conserving CA are studied in [19]. There, among other things, it has been proved that in every surjective number-conserving CA, the set of temporally periodic configurations is dense. The same holds for any surjective one-dimensional CA that has a conserved energy with a unique ground configuration [20, 58]. A ground configuration for an energy θ is a configuration x such that $\delta\Theta(x,y) \geq 0$ for any configuration y asymptotic to x. There is a long-standing open question, asking whether every surjective CA has a dense set of temporally periodic configurations (see e.g. [35, 15]). It is also shown that positively expansive CA cannot have any non-trivial conservation laws [20, 58]. See Chapter **??** or [39, 41] for information on cellular automata as dynamical systems.

Non-increasing energies have a more established connection with the dynamics, as they define Lyapunov functions for the measure dynamical system $(\mathscr{M}_\sigma, F)$. See Chapter **??** for more on this point of view.

Another trend of research has been to find counterparts of Noether's theorem for cellular automata (see e.g. [44, 3, 12, 13, 52]). Noether's theorem — after the German mathematician Emmy Noether (1882–1935) — establishes a correspondence between the conservation laws of a dynamical system defined in terms of a Hamiltonian or a Lagrangian and its symmetries (see e.g. [42, 2]).

Finally, let us mention the interesting connection between conservation laws and invariant Gibbs measures in reversible cellular automata. Gibbs measures are especially important in statistical mechanics, as they characterize the equilibrium state of lattice Hamiltonian systems (see e.g. [25, 56]). In fact, early study of conservation laws in cellular automata was motivated by the issue of ergodicity in statistical mechanical systems [59].

In reversible cellular automata, there is a one-to-one correspondence between conservation laws and invariant Gibbsian specifications [36, 58]. As a special case, a reversible CA conserves a context-free energy $\eta : S \to \mathbb{R}$ if and only if it preserves the Bernouli measure with probabilities proportional to $2^{-\eta}$. This is at least partially valid also for surjective CA. Similar results, though in different settings, are obtained in [34] (see also [60]). It is an interesting open issue to investigate the questions of statistical mechanics in the framework of (deterministic) reversible cellular automata.

References

1. Anderson, R., Lovász, L., Shor, P., Spencer, J., Tardos, E., Winograd, S.: Disks, balls, and walls: Analysis of a combinatorial game. The American Mathematical Monthly **96**(6), 481–493 (1989)
2. Arnold, V.I.: Mathematical Methods of Classical Mechanics, 2nd edn. Springer-Verlag (1989). English Translation
3. Baez, J.C., Gilliam, J.: An algebraic approach to discrete mechanics. Letters in Mathematical Physics **31**, 205–212 (1994)
4. Bak, P., Tang, C., Wiesenfeld, K.: Self-organized criticality: An explanation of $1/f$ noise. Physical Review Letters **59**(4), 381–384 (1987)
5. Bak, P., Tang, C., Wiesenfeld, K.: Self-organized criticality. Physical Review A **38**(1), 364–374 (1988)
6. Bernardi, V.: Lois de conservation sur automates cellulaires. Ph.D. thesis, Université de Provence (2007)
7. Bernardi, V., Durand, B., Formenti, E., Kari, J.: A new dimension sensitive property for cellular automata. Theoretical Computer Science **345**, 235–247 (2005)
8. Blondel, V.D., Cassaigne, J., Nichitiu, C.: On the presence of periodic configurations in Turing machines and in counter machines. Theoretical Computer Science **289**, 573–590 (2002)
9. Boccara, N., Fukś, H.: Cellular automaton rules conserving the number of active sites. Journal of Physics A: Mathematical and General **31**(28), 6007–6018 (1998)
10. Boccara, N., Fukś, H.: Number-conserving cellular automaton rules. Fundamenta Informaticae **52**, 1–13 (2002)
11. Boccara, N., Fukś, H.: Motion representation of one-dimensional cellular automata rules. International Journal of Modern Physics C **17**, 1605–1611 (2006)
12. Boykett, T.: Towards a Noether-like conservation law theorem for one dimensional reversible cellular automata (2005). `arXiv:nlin.CG/0312003`
13. Boykett, T., Kari, J., Taati, S.: Conservation laws in rectangular CA. Journal of Cellular Automata **3**(2), 115–122 (2008)
14. Boykett, T., Moore, C.: Conserved quantities in one-dimensional cellular automata (1998). Unpublished manuscript
15. Boyle, M.: Open problems in symbolic dynamics. In: K. Burns, D. Dolgopyat, Y. Pesin (eds.) Geometric and Probabilistic Structures in Dynamics, Contemporary Mathematics. American Mathematical Society (2008)
16. Chopard, B., Droz, M.: Cellular Automata Modeling of Physical Systems. Cambridge University Press (1998)
17. Durand, B., Formenti, E., Róka, Z.: Number conserving cellular automata I: decidability. Theoretical Computer Science **299**, 523–535 (2003)
18. Feynman, R.: The Character of Physical Law. MIT Press (1965)
19. Formenti, E., Grange, A.: Number conserving cellular automata II: dynamics. Theoretical Computer Science **304**, 269–290 (2003)
20. Formenti, E., Kari, J., Taati, S.: In preparation
21. Formenti, E., Kari, J., Taati, S.: The most general conservation law for a cellular automaton. In: E.A. Hirsch, A.A. Razborov, A. Semenov, A. Slissenko (eds.) Proceedings of CSR 2008, *LNCS*, vol. 5010, pp. 194–203 (2008)
22. Fredkin, E., Toffoli, T.: Conservative logic. International Journal of Theoretical Physics **21**, 219–253 (1982)
23. Fukś, H.: A class of cellular automata equivalent to deterministic particle systems. In: S. Feng, A.T. Lawniczak, S.R.S. Varadhan (eds.) Hydrodynamic Limits and Related Topics, *Fields Institute Communications*, vol. 27, pp. 57–69. American Mathematical Society (2000)
24. Fukś, H., Sullivan, K.: Enumeration of number-conserving cellular automata rules with two inputs. Journal of Cellular Automata **2**, 141–148 (2007)
25. Georgii, H.O.: Gibbs Measures and Phase Transitions. Walter de Gruyter (1988)

26. Georgii, H.O., Häggström, O., Maes, C.: The random geometry of equilibrium phases. In: C. Domb, J. Lebowitz (eds.) Phase Transitions and Critical Phenomena, vol. 18, pp. 1–142. Academic Press (2000)
27. Goles, E.: Sand pile automata. Annales de l'Institut Henri Poincaré (A) Physique théorique **59**(1), 75–90 (1992)
28. Goles, E., Kiwi, M.A.: Games on line graphs and sand piles. Theoretical Computer Science **115**, 321–349 (1993)
29. Greenberg, J., Greene, C., Hastings, S.: A combinatorial problem arising in the study of reaction-diffusion equations. SIAM Journal on Algebraic and Discrete Methods **1**(1) (1980)
30. Greenberg, J.M., Hassard, B.D., Hastings, S.P.: Pattern formation and periodic structures in systems modeled by reaction-diffusion equations. Bulletin of the American Mathematical Society **84**(6) (1978)
31. Greenberg, J.M., Hastings, S.P.: Spatial patterns for discrete models of diffusion in excitable media. SIAM Journal on Applied Mathematics **34**(3), 515–523 (1978)
32. Hattori, T., Takesue, S.: Additive conserved quantities in discrete-time lattice dynamical systems. Physica D **49**, 295–322 (1991)
33. Hedlund, G.A.: Endomorphisms and automorphisms of the shift dynamical system. Mathematical System Theory **3**, 320–375 (1969)
34. Helvik, T., Lindgren, K., Nordahl, M.G.: Continuity of information transport in surjective cellular automata. Communications in Mathematical Physics **272**, 53–74 (2007)
35. Kari, J.: Theory of cellular automata: A survey. Theoretical Computer Science **334**, 3–33 (2005)
36. Kari, J., Taati, S.: In preparation
37. Kari, J., Taati, S.: A particle displacement representation for conservation laws in two-dimensional cellular automata. In: B. Durand (ed.) Proceedings of JAC 2008, pp. 65–73 (2008)
38. Kindermann, R., Snell, J.L.: Markov Random Fields and Their Applications. American Mathematical Society (1980)
39. Kůrka, P.: Topological dynamics of cellular automata. In: Codes, Systems and Graphical Models, pp. 447–486. Springer-Verlag (2001)
40. Kůrka, P.: Cellular automata with vanishing particles. Fundamenta Informaticae **58**, 1–19 (2003)
41. Kůrka, P.: Topological and Symbolic Dynamics, *Cours Spécialisés*, vol. 11. Société Mathématique de France (2003)
42. Landau, L.D., Lifshitz, E.M.: Course of Theoretical Physics: Mechanics, third edn. Elsevier (1976). English Translation
43. van Lint, J.H., Wilson, R.M.: A Course in Combinatorics. Cambridge University Press (1992)
44. Margolus, N.: Physics-like models of computation. Physica D **10**, 81–95 (1984)
45. Margolus, N.: Physics and computation. Ph.D. thesis, Massachusetts Institute of Technology (1987)
46. Minsky, M.: Computation: Finite and Infinite Machines. Prentice-Hall (1967)
47. Moreira, A.: Universality and decidability of number-conserving cellular automata. Theoretical Computer Science **292**, 711–723 (2003)
48. Moreira, A., Boccara, N., Goles, E.: On conservative and monotone one-dimensional cellular automata and their particle representation. Theoretical Computer Science **325**, 285–316 (2004)
49. Morita, K., Imai, K.: Number-conserving reversible cellular automata and their computation-universality. Theoretical Informatics and Applications **35**, 239–258 (2001)
50. Nagel, K., Herrmann, H.J.: Deterministic models for traffic jams. Physica A **199**, 254–269 (1993)
51. Nagel, K., Schreckenberg, M.: A cellular automaton model for freeway traffic. Journal de Physique I **2**(12), 2221–2229 (1992)
52. Ollinger, N.: Personal communication (2008)
53. Parthasarathy, K.R.: Probability Measures on Metric Spaces. Academic Press (1967)

54. Pivato, M.: Conservation laws in cellular automata. Nonlinearity **15**, 1781–1793 (2002)
55. Pomeau, Y.: Invariant in cellular automata. Journal of Physics A: Mathematical and General **17**(8), L415–L418 (1984)
56. Ruelle, D.: Thermodynamic Formalism, second edn. Cambridge University Press (2004)
57. Sinai, Y.G.: Theory of Phase Transitions: Rigorous Results. Pergamon Press (1982)
58. Taati, S.: Conservation laws in cellular automata. Ph.D. thesis, University of Turku (2009)
59. Takesue, S.: Reversible cellular automata and statistical mechanics. Physical Review Letters **59**(22), 2499–2502 (1987)
60. Toffoli, T.: Information transport obeying the continuity equation. IBM Journal of Research and Development **32**(1), 29–36 (1988)
61. Toffoli, T., Margolus, N.: Cellular Automata Machines: A New Environment for Modeling. MIT Press (1987)
62. Toffoli, T., Margolus, N.: Invertible cellular automata: A review. Physica D **45**, 229–253 (1990)
63. Vichniac, G.Y.: Simulating physics with cellular automata. Physica D **10**, 96–116 (1984)
64. Walters, P.: An Introduction to Ergodic Theory. Springer-Verlag (1982)

Index