

# Does universal controllability of physical systems prohibit thermodynamic cycles?

Dominik Janzing

Max Planck Institute for Intelligent Systems

Max-Planck-Ring 4

72076 Tübingen, Germany

Email: [dominik.janzing@tuebingen.mpg.de](mailto:dominik.janzing@tuebingen.mpg.de)

Pawel Wocjan

Department of Computer Science, University of Central Florida

4328 Scorpius Street

Orlando, FL 32816, USA

Email: [wocjan@cs.ucf.edu](mailto:wocjan@cs.ucf.edu)

March 28, 2018

## Abstract

Here we study the thermodynamic cost of computation and control using 'physically universal' cellular automata (CAs) or Hamiltonians. The latter were previously defined as systems that admit the implementation of any desired transformation on a finite target region by first initializing the state of the surrounding and then letting the system evolve according to its autonomous dynamics. This way, one obtains a model of control where each region can play both roles, the controller or the system to be controlled. In physically universal systems every degree of freedom is indirectly accessible by operating on the remaining degrees of freedom.

In a nutshell, the thermodynamic cost of an operation is then given by the size of the region around the target region that needs to be initialized. In the meantime, physically universal CAs have been constructed by Schaeffer (in two dimensions) and Salo & Törmä (in one dimension). Here we show that in Schaeffer's CA the cost for implementing  $n$  operations grows linearly in  $n$ , while operating in a thermodynamic cycle requires sublinear growth to ensure zero cost *per operation* in the limit  $n \rightarrow \infty$ . Although this particular result need not hold for general physically universal CAs, this strong notion of universality does imply a certain kind of instability of information, which could result in lower bounds on the cost of protecting information from its noisy environment.

The technical results of the paper are sparse and quite simple. The contribution of the paper is mainly conceptual and consists in illustrating the type of thermodynamic questions raised by models of control that rely on the concept of physical universality.

# 1 Why thermodynamics of computation and control requires new models

## 1.1 The debate on thermodynamics of computation since the 1960s

The question of whether there are fundamental lower bounds on the energy consumption of computing devices has attracted the attention of researchers since the 1960s. Landauer [1] realized that logically irreversible operations like erasure of memory space necessarily require to transfer the energy  $\ln 2kT$  per bit to the environment (with  $k$  denoting Boltzmann's constant and  $T$  the temperature of the environment) due to the second law of thermodynamics.<sup>1</sup> Bennett [6] clarified that computation can be performed without logically irreversible operations and thus Landauer's argument does not prove any fundamental lower bound for the energy needed by computation tasks without further specification. Ref. [7] argues that physical models of reversible computation should include the clocking mechanism (that control the implementation of logical gates) because otherwise one neglects the question of how to implement clocking in a thermodynamically reversible way (after all, if both gates and clocking device are described as quantum systems then the influence of the latter on the former would, to some extent, also imply an influence of the former on the latter [8]).

## 1.2 External clocking and control signals as loopholes

To motivate this work step by step we first discuss the thermodynamics of clocking and synchronization briefly which is a sophisticated problem [9, 10, 11, 12]. Ref. [11], for instance, studies some synchronization protocols that suggest that thermodynamically reversible synchronization requires to exchange *quantum* information, which links the a priori different tasks of reversible computation and quantum computing.<sup>2</sup>

Going beyond the question of whether implementing reversible logical operations is possible in a thermodynamically reversible way, we ask whether implementing unitary operations on some quantum system is possible in a thermodynamically reversible way. Regardless of how we call the physical devices controlling the implementation (we called it 'clock' in the case of computation processes), also the implementation of a unitary  $U$  requires to 'change Hamiltonians' – except for the special case where  $U = e^{-iHt}$  with  $H$  being the free Hamiltonian of the system of consideration. However, do we really have appropriate models for discussing the thermodynamic cost of 'changing a system's Hamiltonian'? After all, describing a control field in classical terms is only a valid approximation if it can be considered *macroscopic*. For instance, a 'macroscopic' number of electrons, sufficiently distant from some probe particle under consideration, could create such a 'classical' field. It is hard, however, to imagine a macroscopic controller whose energy consumption does not exceed the energy content of the microscopic target system. This suggests that discussing potential thermodynamic limitations requires *microscopic* models of control.

For both tasks, computation and control, we are criticizing basically the same issue: as long as the device controlling or triggering the operations (regardless of whether we call it 'clock' or 'controller') is not included in our microscopic description, we are skeptical about the claim that the operation could 'in principle' be implemented in a thermodynamic cycle without any energy cost.

---

<sup>1</sup>In [2] we have argued that the energy requirements for *reliable* erasure are even larger than Landauer's bound when the state of the energy source is noisy, for instance if it is given by two thermodynamic reservoirs of different temperatures. For further different perspectives on Landauer's principle see, e.g., [3, 4, 5].

<sup>2</sup>Here, the formal distinction between quantum and classical clock signals as well as the conversion of time information between them is based on the rather general framework introduced in [13].

These remarks raise the following two questions: (1) What are appropriate models for discussing resource requirements of computation and control? Given such a model, we need to ask (2) how to define *resource requirements* within the model.

To discuss the cost of ‘changing Hamiltonians’ we first recall that changing ‘effective Hamiltonians’ is what is actually done: Let the target system, for instance, be a single particle. Changing control fields actually means to change the quantum state of the physical systems surrounding the particle. In a certain mean-field limit, this state change amounts to the change of a classical field. Thus, the particles interact according to a *fixed* Hamiltonian. Taking this perspective seriously, we are looking for a model where control operations are implemented by a fixed interaction Hamiltonian if the states of the surrounding quantum systems are tuned in an appropriate way. Ref. [14] also studies thermodynamic laws in a scenario where system, controller, and baths are coupled by a fixed time-independent Hamiltonian, while [15] also considers autonomous dynamics of *open* systems. Although the goal of the present paper is also to study thermodynamics in a scenario with autonomous time evolution, we consider a model that is nevertheless general enough to enable controlling controllers by ‘meta’-controllers and so on. This, in turn, requires to couple the target system considered in the first place to an infinite system that is not just a ‘heat bath’ as it is often assumed but something that can be controlled and, further, act as a controller at the same time.

### 1.3 Spin lattice Hamiltonians as autonomous models of computation

As models for reversible computing, Hamiltonians on spin lattices have been constructed that are able to perform computation [16] by their autonomous evolution. This addresses the above criticism in the sense that these models do not require any external clocking. Instead, synchronization is achieved by the fixed and spatially homogeneous interaction Hamiltonian itself. Refs. [17, 18] go one step further and describe Hamiltonians on spin lattices for which the result of the computation need not be read out within a certain time interval because the *time average state* encodes the result. This solves the more subtle problem that otherwise the readout required an external clock.

There are several properties that make spin lattices attractive as physical toy models of the world (and not only as model for a computing device): the discrete lattice symmetry represents spatial homogeneity of the physical laws and the constant Hamiltonian the homogeneity in time. By looking at lattices as discrete approximations of a field theoretical description of the physical world, even the presence and absence of matter can be seen as just being different states of the lattice. Accordingly, one can argue that spin lattices allow for a quite principled way of studying thermodynamics of computation and control because they model not only the computing device itself but also its interaction with the environment: to this end, just consider some region in the lattice as the *computing device* and the complement of that region as the *environment*.

### 1.4 Why we propose to add physical universality

For the purpose of developing our ‘toy thermodynamics of computing and control’ we propose to consider spin lattices or cellular automata (as their discrete analog) that satisfy the additional condition of *physical universality* introduced in [19]. This property will be explained and motivated on an informal level in the following section. Roughly speaking, physical universality means that the autonomous time evolution of the system is able to implement any mathematically possible process on an arbitrarily large finite region after the complement of the region is prepared to an appropriate initial state. To appreciate that physical universality is a strong requirement it is essential to emphasize that the autonomous dynamics of the system is fixed and not adapted to the particular process that is supposed to be implemented. In the case of quantum systems, we mean by ‘mathematically possible’ the set of completely positive trace preserving maps. In the classical case, we refer to the set of stochastic maps. Given that one

believes in the hypothesis that real physical systems admit in principle the implementation of any mathematically possible process<sup>3</sup>, it is natural to demand that the interaction at hand itself is able to implement the transformation. Otherwise, the interaction does not fully describe the interface between system and its environment. For the purpose of our thermodynamic considerations, however, we want to study systems whose interface is completely described by the interaction under consideration rather than relying on control operations that come as additional, external, ingredients.

The paper is structured as follows. Section 2 briefly motivates the notion of physical universality introduced in [19] for both Hamiltonians and cellular automata<sup>4</sup>, although we focus on the latter for sake of simplicity. Section 3 introduces the condition of physical universality formally and describes and discusses the notion of resource requirements introduced in [19], which is also the basis of this paper. Further, we raise the question of whether the resource requirements of repeating a certain operation can grow sublinear in the number of repetitions (which we argue to be necessary to justify the term 'thermodynamic cycle'). Section 4 explains why CAs that are not physically universal may admit thermodynamic cycles in our sense. This is because they admit initializations of a finite region that ensure the implementation of endless repetitions of the same control operation. Section 5 explains why this simple construction is impossible in *physically universal* CAs and shows that Schaeffer's CA does not admit sublinear growth. Whether no physically universal CA admits sublinear growth has to be left to the future.

## 2 Physical universality: informal description and possible consequences

### 2.1 Physically universal systems as consistent models of control

Ref. [19] introduces the notion of physical universality for three types of systems:

- (1) translationally invariant finite-range interaction Hamiltonians on infinite spin lattices,
- (2) quantum cellular automata, and
- (3) classical cellular automata.

While (1) is the model that is closest to physics, (2) and (3) describe increasing abstractions that are useful for our purpose. Essentially, (2) is just the discrete time version of (1). We will restrict the attention to (3) because it turns out that the problem is already difficult enough for this case.

At an abstract level, the definition of physical universality coincides for all three cases: a system is called physically universal if any transformation on any desired target region (of arbitrary but finite size) can be implemented by first initializing the (infinite) complement of that region to an appropriate state and then letting the system evolve according to its autonomous

---

<sup>3</sup>For critical remarks on this postulate see [20], Chapter 7: here doubts are raised that every self-adjoint operator in a multi-particle system can be measured in practice. However, there exists always a unitary transformation that reduces the observable to an observable that is diagonal in the tensor product basis, i.e., measurements of every single particle. Given that one believes that these individual measurements are always possible even for multi-partite systems, the doubts thus question the implementation of arbitrary unitaries. Further, Ref. [21] discusses the concept of physical universality for an understanding of life and also proposes to weaken physical universality – just to mention a second critical point of view.

<sup>4</sup>Note that this paper contains several ideas that already appear in the preprint [19], but often less explicit than here. Since [19] will not be published because its main purpose had been to state a question that has been solved in the meantime, we do not care about this overlap.

dynamics for a certain ‘waiting time’  $t$ . For the cases (2) and (3),  $t$  is a positive integer while it is a positive real number for the case (1). Since cases (1) and (2) refer to *quantum* systems the set of possible transformations (completely positive trace preserving maps) is uncountably infinite, we should only demand that one can get *arbitrarily close* to the desired transformation via appropriate initializations and waiting times instead of being able to implement the desired transformation exactly. For physically universal quantum cellular automata, we demand that quantum transformation of the target region be implemented by initializing the complement of the target region in a basis state. This is reasonable because the initialization should be simple. Otherwise, the complexity of implementing a desired quantum transformation could be “hidden” in the preparation of a complex quantum state. We also refer the reader to [22, Section 5], where it is explained that the results on programmable quantum gate arrays imply that nothing is gained by allowing superpositions.

**Shifting the boundary between target and controller** Physically universal systems are intriguing because they provide a model class where every physical degree of freedom is indirectly accessible by operating on the remaining degrees of freedom in the ‘world’ and then letting the joint system evolve. In other words, the complement of the target region acts as the controller of the target region so that any part of the world can become the controller or the system to be controlled. This is in contrast to some physical models of computation, e.g., [17], for which data and program registers are represented by different types of physical degrees of freedom. These systems are able to perform any desired transformation on the *data register* by appropriate initialization of the *program register*. The question of how to act on the program register cannot be addressed within the model. In physically universal systems, on the other hand, the preparation of *any* region can be achieved by operating on its complement. This reduces the question of how to act on some target region to the question of how to act on some ‘controller’ region around it. In turn, this controller region can be prepared by acting on some ‘meta-controller’ region around it. Although this does not solve the problem it shows at least that the boundary between controller and target region can be arbitrarily shifted.

**Analogy to the quantum measurement problem** The above shift of the boundary is similar to the quantum measurement problem where the boundary between the measurement apparatus and the quantum system to be measured (the famous ‘Heisenberg cut’) can be arbitrarily shifted as long as the quantum description is considered appropriate: the transition from a pure superposition to the corresponding mixture can be explained by entanglement between the target system and its measurement apparatus [23] (for simplicity, one may define ‘measurement apparatus’ as all parts of the environment that carry information about the result). The resulting joint superposition of measurement apparatus and target system can be transferred to a mixture by entanglement with a ‘meta’ measurement apparatus and so on.

## 2.2 Potential thermodynamic implications

Physical universality can have important thermodynamic consequences because it excludes the ability to completely protect information. Physically universality means that any system can be controlled by its surrounding. Therefore, the unknown state of the surrounding will eventually cause the state of the system to change. In contrast, in systems such as [17] the state of the program register never changes during the autonomous time evolution because of the strict separation between data and program registers. Here, we do not want to accept the latter class of models as physical models of computation because in the real world program registers are also physical systems that can be somehow accessed by actions on their environment. In other words, the information of the ‘program’ register is only safe because the model fails to describe

how to act on that part of the system using the given interactions (these actions are external to the theory).

**Trade-off between stability and controllability** Physical universality thus gives rise to a thermodynamic theory in which the *inability to protect* information is the result of the *ability to control* every degree of freedom. On the one hand, the target system needs to interact with its environment otherwise we would not be able to control it. On the other hand, this interaction makes entropy leak from the surrounding into the target system. Ref. [19] defines the model class of physically universal systems for the purpose of studying this conflict on an abstract level. Here, we restrict the attention to discrete time dynamics on *classical cellular automata*. In the long run, one should certainly address our thermodynamic questions using *continuous time dynamics on quantum systems*. As a first approach, however, it is convenient to simplify the problem by restricting oneself to classical CAs. Another reason for considering classical CAs is also to make this problem more accessible to the computer science community.<sup>5</sup> After all, it is one of the lessons learned from quantum information theory [25] that translating physics into computer scientific language can provide a new perspective and new paradigms. Indeed, the past two decades have shown that understanding thermodynamics via computer scientific models is also promising.<sup>6</sup> On the microscopic level one can hardly tell apart computing devices from thermodynamic machines in the conventional sense.<sup>7</sup> As part of this oversimplification, we will define the thermodynamic cost of an operation simply by the size of the region in the surrounding of the target system that needs to be initialized. This will be partly justified in Section 3.2.

### 3 The formal setting

#### 3.1 Notation and terminology

For the basic notation we mainly follow [30]. The cells of our CA in  $d$  dimensions are located at lattice points in  $\Omega := \mathbb{Z}^d$ . The *state* of each cell is given by an element of the alphabet  $\Sigma$ . For any subset  $X \subset \Omega$ , a *configuration*  $\gamma_X$  of  $X$  is a map  $X \rightarrow \Sigma$ . Let  $\Sigma^X$  denote the set of all configurations of  $X$ . The dynamics of the CA is given by a map  $\alpha : \Sigma^\Omega \rightarrow \Sigma^\Omega$  that is local (i.e. the state of each cell is only influenced by the state of cells in a fixed neighborhood) and spatially homogeneous (i.e., it commutes with all lattice translations).

Later, we will often consider a class of CAs in dimension  $d = 2$  where the state of a cell one time step later only depends on the state of the cell itself and its 8 surrounding neighbors, the so-called Moore neighborhood, and refer to this class as ‘Moore CAs’.

If  $\gamma' := \alpha(\gamma)$  for any  $\gamma \in \Sigma^\Omega$ , we also write  $\gamma \rightarrow \gamma'$  to indicate that the configuration  $\gamma$  evolves to  $\gamma'$  in one time step and  $\gamma \xrightarrow{n} \gamma' = \alpha^n(\gamma)$  means that  $\gamma$  evolves to  $\gamma'$  in  $n$  time steps.

**Definition 1** (implementation of a function). *Let  $X, Y \subset \Omega$  be finite sets and  $f : \Sigma^X \rightarrow \Sigma^Y$  be an arbitrary function. Then we say a configuration  $\phi \in \Sigma^X$  implements  $f$  in time  $t$  if for every  $x \in \Sigma^X$*

$$\phi \oplus x \xrightarrow{t} \psi_x \oplus f(x),$$

*holds for some  $\psi_x \in \Sigma^Y$ . Here, the sign  $\oplus$  denotes merging configurations of disjoint regions to a configuration of the union.*

---

<sup>5</sup>Note, further, that already von Neumann’s self-reproducing automata [24] follows the principle to study physical or biological universality properties using CAs.

<sup>6</sup>For instance, the principle of cooling devices [26, 27] and heat engines [28] can be illustrated using an  $n$ -bit register represented by  $n$  two-level systems or other simple discrete systems. For this model class, the relation between physics and information is most obvious.

<sup>7</sup>See also the adaptive heat engine in Ref. [29].

For physical universality, we follow Schaeffer’s modified definition [30], which is equivalent to our original one, and also his definition of *efficiently* physically universal:

**Definition 2** (physical universality). *We say a cellular automaton is physically universal if for all finite regions  $X, Y$  and all transformations  $f : \Sigma^X \rightarrow \Sigma^Y$ , there exists a configuration  $\phi$  of the complement of  $X$  and a natural number  $t \in \mathbb{N}$  such that  $\phi$  implements  $f$  in time  $t$ .*

*We say the CA is efficiently physically universal if the implementation runs in time  $t_0$ , where  $t_0$  is polynomial in*

- *the diameter of  $X$  (i.e., the width of the smallest hypercube containing the set) and diameter of  $Y$ ,*
- *the distance between  $X$  and  $Y$ , and*
- *the computational complexity of  $f$  under some appropriate model of computation (e.g., the number of logical gates in a circuit for  $f$ ).*

For simplicity, we will often consider only the case  $Y = X$ . Since every signal in our CA propagates only one cite per time step, at most a margin of thickness  $t$  around  $X$  matters for what happens after  $t$  time steps. Depending on the dynamical law and the desired operation on the target region, the relevant part of the state can be significantly less. To explore the resource requirements of an ‘implementation’ we phrase the notion of an implementation formally in a way that is explicit about which parts of the surrounding cells matter to achieve the desired operation:

**Definition 3** (device for implementing  $f$ ). *A device for implementing  $f : \Sigma^X \rightarrow \Sigma^Y$  is a triple  $(Z, \phi_Z, t)$  such that  $\phi_Z \oplus \phi'$  implements  $f$  in  $t$  time steps for all  $\phi' \in \Sigma^{\bar{Z} \cap \bar{X}}$ . Here,  $X$  and  $Y$  are called the ‘source region’ and ‘target region’, respectively, and  $Z \subset \bar{X}$  is called the ‘relevant region’,  $\phi_Z \in \Sigma^Z$  the state of this region, and  $t \in \mathbb{N}$  the ‘implementation time’. Then, the ‘size’ of the device is the size of  $W := Z \cup X \cup Y$ . The ‘range’ of the device is the side length of the smallest  $d$ -dimensional hypercube containing  $W$ .*

Note that the relevant region may overlap with the target region while it needs to be disjoint with the source region. Further, note that the definition of a device does not imply that the relevant region has been chosen in a minimal way. Accordingly, future theorems on the resource requirements of implementations may read ‘the relevant region consists of at least  $\dots$  cells.’ The range can be seen as the size of the apparatus. Assume, for instance, that  $W$  consists of a small number  $n$  of single cells spread over a hypercube of side length  $k \gg n$ . Then we would still call this a ‘large’ apparatus even if  $n$  is small.

So far, we have only considered the ability to implement *one* specific transformation once. We also want to be able to study processes where one desired operation is performed after time  $t_1$ , a second one after time  $t_2 + t_1$ , and so on. Assume, for instance, that we want to achieve that the information content of a certain cell  $c_1 \in \Omega$  is shifted to cell  $c_2$  after some time  $t_1$  and then shifted to cell  $c_3$  at some later time  $t_2 + t_1$ . Then the entire process consisting of both subprocesses should be performed by one initialization rather than demanding re-preparing the system after each transformation. To this end, we define devices for implementing concatenations of transformations as generalization of Definition 3:

**Definition 4** (device for implementing a sequence of transformations). *Let  $X_1, \dots, X_{n+1}$  be finite regions and  $f_1, \dots, f_n$  be functions with  $f_j : \Sigma^{X_j} \rightarrow \Sigma^{X_{j+1}}$  for  $j = 1, \dots, n$ . In other words, the target region of  $f_j$  is the source region of  $f_{j+1}$ . A device for implementing the sequence  $f_1, f_2, \dots, f_n$  is an  $n + 2$ -tuple  $(Z, \phi_Z, t_1, \dots, t_n)$  with  $t_j > 0$ , where  $Z \subset \bar{X}_1$  is called the ‘relevant region’ and  $\phi_Z \in \Sigma^Z$  is a configuration such that  $\phi_Z \oplus \phi'$  implements  $f_j \circ f_{j-1} \circ \dots \circ f_1$  in  $\sum_{i=1}^j t_i$  time steps for all  $\phi' \in \Sigma^{\bar{Z} \cap \bar{X}_1}$ . The size of the device is the size of  $W := Z \cup (\cup_{i=1}^{n+1} X_i)$  and its range is the side length of the smallest  $d$ -dimensional hypercube containing  $W$ .*

The idea of Definition 4 is that the CA implements the transformation  $f_j$  within  $t_j$  time steps. Note, however, that the definition only specifies that any initial configuration  $x_1$  of  $X_1$

undergoes the state transitions

$$x_1 \xrightarrow{t_1} f_1(x_1) \xrightarrow{t_2} f_2(f_1(x_1)) \xrightarrow{t_3} \dots \xrightarrow{t_n} f_n(\dots(f_1(x_1))\dots).$$

It does *not* specify what happens if we disturb the entire process by an external intervention. For instance, after changing  $f_1(x_1)$  at time  $t_1$  to some state  $y$  other than  $f_1(x_1)$ , we need not necessarily obtain  $f_2(y)$  after  $t_2$  steps.<sup>8</sup>

A priori it is not obvious that physical universality entails the ability of implementing sequences with  $n > 1$ . The following result shows that this is the case:

**Theorem 1** (ability to implement sequences). *In every physically universal CA there is a device for any sequence of transformations.*

*Proof.* We provide a proof by induction over  $n$ . The base case  $n = 1$  follows from physical universality. For the induction hypothesis assume that sequences of  $n$  arbitrary functions can be implemented.

For the induction step, let  $f_1, \dots, f_n, f_{n+1}$  be a sequence of  $n + 1$  arbitrary transformations with

$$f_j : \Sigma^{X_j} \rightarrow \Sigma^{X_{j+1}}$$

for  $j = 1, \dots, n + 1$ .

By physical universality there exists a device  $(Z_{n+1}, \phi_{Z_{n+1}}, t_{n+1})$  with  $Z_{n+1} \subset \bar{X}_{n+1}$  that implements the last function  $f_{n+1}$  of the above sequence. Using this device we define the following augmented version  $\hat{f}_n$  of the second last function  $f_n$  of the above sequence by setting

$$\hat{f}_n : \begin{cases} \Sigma^{X_n} & \rightarrow \Sigma^{X_{n+1}} \cup \Sigma^{Z_{n+1}} \\ y & \mapsto f_n(y) \oplus \phi_{Z_{n+1}} \end{cases}$$

for all  $y \in \Sigma^{X_n}$ . In words, the output of the augmented function  $\hat{f}_n$  consists of the output of original function  $f_n$  on the region  $X_n$  and the constant output  $\phi_{Z_{n+1}}$  on the region  $Z_{n+1}$ .

By induction hypothesis there exists a device  $(Z, \phi, t_1, \dots, t_{n-1}, t_n)$  that implements the sequence  $f_1, \dots, f_{n-1}, \hat{f}_n$ . The special form of the output of the augmented function  $\hat{f}_n$  ensures that the device  $(Z, \phi, t_1, \dots, t_{n-1}, t_n, t_{n+1})$  also implements the sequence  $f_1, \dots, f_{n-1}, f_n, f_{n+1}$ . This is because after  $t_1 + \dots + t_n$  times steps the output is

$$\hat{f}_n(y) = f_n(y) \oplus \phi_{Z_{n+1}} \in \Sigma_{X_{n+1}} \cup \Sigma_{Z_{n+1}} \quad \text{where } y = f_{n-1}(\dots(f_1(x))\dots),$$

so that after  $t_{n+1}$  additional time steps the final output is

$$f_{n+1}(f_n(y)) \in \Sigma^{X_{n+2}},$$

as desired. □

To mention a simple example of the kind of sequences we are interested in, consider a CA with binary alphabet  $\Sigma = \{0, 1\}$ . Assume the task is to implement a NOT gate on the same bit  $n$  times on some target bit. Then the desired functions read  $f_j = \text{NOT}$  and the numbers  $t_j$  specify the time instants for which the autonomous dynamics has implemented another NOT gate on our target bit, given that some region  $Z$  has been initialized to the state  $\phi_Z$ .

---

<sup>8</sup>Rephrased in causal language [31], if we denote the state of  $X_j$  at time  $\sum_{i=1}^j t_i$  by  $x_j$ , then the equation

$$x_j = f_{j-1}(x_{j-1}), \tag{1}$$

is not a ‘structural equation’, since the latter describes, by definition, also the impact that interventions on the input variable on the right hand side have for the variable on the left hand side.



### 3.2 Formalizing ‘thermodynamic cost’ of operations

Here we will consider the size of the relevant region as the thermodynamic cost of an implementation. This first approximation is justified by the following idea: a priori, the state of each cell is unknown, i.e., we assume uniform distribution over  $\Sigma$ . According to Landauer’s principle it then requires the energy  $kT \ln |\Sigma|$  to initialize one cell to the desired state. This way, the thermodynamic cost of the initialization process is simply proportional to the number of cells to be initialized. This view will be further discussed at the end of this subsection.

Note that the size of the relevant region can only grow with  $O(t^d)$  if  $t$  is the running time for an implementation since a signal can only proceed a constant number of cells per time step. Therefore, the thermodynamic cost scales only polynomial in the computational complexity if a CA is efficiently physically universal. This statement, however, is too weak for our purpose. To phrase the main questions of this paper (which looks for stronger statements) we need the following terminology:

**Definition 5** (zero cost per operation). *Given a function  $f : \Sigma^X \rightarrow \Sigma^X$ , a physically universal CA is said to admit the implementation of  $f$  at zero cost per operation, if there are devices  $(Z_n, \phi_{Z_n}, t_1, \dots, t_n)$  for every  $n \in \mathbb{N}$ , each implementing the sequence  $\underbrace{f, \dots, f}_n$ , such that*

$$\lim_{n \rightarrow \infty} \frac{|Z_n|}{n} = 0.$$

Note that this definition does not require that the implementation of  $f$  stops after the time  $t_n$ . Likewise, we define:

**Definition 6** (zero cost of information storage per time). *For some region  $X$ , a physically universal CA is said to admit zero cost of information storage per time on  $X$  if there are devices  $(Z_n, \phi_{Z_n}, t_n)$  for every  $n \in \mathbb{N}$  with  $t_n \rightarrow \infty$  that implement the identity on  $X$  after the time  $t_n$  such that*

$$\lim_{n \rightarrow \infty} \frac{|Z_n|}{t_n} = 0.$$

In a certain sense, the two tasks are complementary: while Definition 5 asks for the resource requirements of *actively* implementing transformations, Definition 6 asks for the cost of passively keeping information over time. It is possible that there is a fundamental trade-off between the two different resource requirements. If a CA does ‘nothing’ on the target region for the large majority of states of the target’s complement, it may require a large amount of resources to implement non-trivial transformations because only few states of the complement yield non-trivial actions on the target region. We are now able to phrase our main questions:

- **Question 1:** Does there exist a physically universal CA that admits zero cost per operation for any / for all functions  $f$ ?
- **Question 2:** Does there exist a physically universal CA that admits zero cost for information storage per time for any / for all finite regions  $X$ ?

If we recall that the state of the CA may also encode the presence or absence of matter, our definition of implementation cost also includes the aspect of hardware deterioration. Assume one has built some microscopic control device that degrades after performing an operation some large number  $n_0$  of times, a device for implementing the operation  $n > n_0$  times includes a ‘meta’ device repairing the original one.<sup>9</sup>

---

<sup>9</sup>Thermodynamic considerations that account also for reproduction processes are certainly related to thermodynamics of life [32].

On the one hand, we will show that the answers to both questions above are negative for Schaeffer’s CA [30]. On the other hand, we will show that there exist physically *non*-universal CAs for which both answers are positive. We leave it as an open question whether physical universality precludes the ability to achieve zero cost. However, we give some intuitive arguments that suggest that physical universality makes it at least more difficult to achieve zero implementation cost per operation or zero cost for information storage per time.

**Discussion of the above formalization of thermodynamic cost** It is certainly an oversimplification to identify the size of the region that needs to be initialized with the thermodynamic cost of an implementation. Consider, for instance, a physical many particle system where each cell is a physical system that is weakly interacting with its neighbours. This ensures that the total energy of the composed system is approximately given by the sum of the energy of the individual systems. Assume, furthermore, that the state  $0 \in \Sigma$  corresponds to the ground state, that is, the state of lowest energy. In the limit of low temperatures, this state has probability close to 1, which implies that initializing the lattice to the all-zero state does not require significant free energy resources. In this case, however, it requires significant free energy resources to set a cell to *any state other than 0* and the resource requirement then depends on the number of cells that need to be in a non-zero state (which may correspond to the number of particles in physics).

On the other hand, identifying the number of cells to be initialized with the thermodynamic cost, can also be justified from the following point of view: assume we are not interested in the amount of free energy that is required for one specific transformation. Instead, we only ask whether the amount increases sublinearly or not. Assuming, in the above physical picture, non-zero temperature (although it may still be low, which favors the state 0), initializing  $n$  states to 0 *with certainty* yet requires an amount of free energy *of the order*  $n$ . This way, the asymptotic behavior of resource requirements is unaffected by the details of the physical hardware assumptions.

## 4 Cost of operations in Turing complete CAs

The ability to write finite computer programs that implement endless loops is a crucial feature of any computing device that has the same computational capabilities as a universal Turing machine. The finiteness of the program code may suggest that such a machine (when choosing one that is ‘reversible’ in the sense of [33]) is able to implement a cycle at zero thermodynamic cost. However, we now discuss in what sense this idea is misleading because physical universality turns out to be an issue for thermodynamically reversibility. As a simple toy example, we consider the control task of repeatedly turning on and off a target bit without ever stopping. Intuitively, this process already reminds us of a program with an infinite loop:

**Example 1** (infinite bit switching).

```

a := 0
while (1) do
  a := 1 ⊕ a // bit XOR
end while

```

Every Turing-complete CA is capable of implementing the above program. We recall that a CA is called Turing-complete if there exists a finite configuration that allows the CA to simulate any universal Turing machine, where the concepts of ‘finite configuration’ and ‘halting’ are defined as follows. ‘Finite configuration’ means that only finitely many cells are in a non-zero state, where a single element of the alphabet  $\Sigma$  is chosen to be zero, denoted by 0. ‘Halting’ is defined as the event of a single previously selected cell becoming non-zero.




Figure 1: A simple configuration in Conway’s Game of Life that yields a dynamical behavior with period 2. The system changes between the two configurations on the left and the right hand side, respectively. ‘Alive’ and ‘dead’ cells are indicated by gray and white, respectively.

It is important to observe that *finite configuration* does not imply *finite resources* in our sense. ‘Finite configuration’ means that all but a finite number of cells are in the *zero* state, whereas ‘finite resources’ means that all but a finite number of cells are in an *unknown* state.

Consider the following situation: the simulation of a universal Turing machine by a CA could *require* that all but a finite number of cells be zero because otherwise the non-zero cells would eventually perturb the simulation. This would mean *infinite* resources in our sense. However, as long as we do not demand physical universality, we can easily modify Turing complete CAs such that they are able to implement an infinite loop with finite resources, as will be discussed in the following two subsections.

#### 4.1 Conway’s Game of Life

We first consider the implementation of our target operation ‘infinite bit switch’ in a well-known cellular automaton, namely Conway’s Game of Life. It is a CA in two dimensions, each cell being ‘alive’ or ‘dead’, i.e., formally each cell is just one bit. The rules are [34]:

- (1) Any live cell with fewer than two live neighbours dies, as if caused by under-population.
- (2) Any live cell with two or three live neighbours lives on to the next generation.
- (3) Any live cell with more than three live neighbours dies, as if by over-population.
- (4) Any dead cell with exactly three live neighbours becomes a live cell, as if by reproduction.

To implement the bit flip, as desired, we find simple oscillating patterns in [34]: The ‘Blinker’ has period 2, as shown in Figure 1. We now focus on the space requirements of this 2-cycle and recall that space requirements in our sense refer to the amount of space that needs to be initialized to a specific value. For the Blinker to work, it is essential that there are no ‘particles’ in the direct neighborhood that disturb the patterns. Whenever there is a region outside which the state is not known at all, this complementary region contains with some probability a pattern that moves towards the blinker and disturbs its cycle. It is therefore possible, that, without having some control about the entire space, we cannot guarantee that the blinker works forever.

#### 4.2 Modified Game of Life with impenetrable walls

There is, however, a simple modification of the Game of Life for which we can ensure that the blinker works forever although we only control the state of a finite region. To this end, we augment each cell by an additional third state ‘brick’ ■, indicated by black color, that prevents any diffusion into the finite region from its surrounding. The transition rule of the new CA now consist of the following rules:

- (0) a cell being in the state ■ remains there forever. (1)-(4) as before, with the convention that the brick ■ counts as □ for its neighbors.




Figure 2: The blinker surrounded by a wall of ‘bricks’, which protect it from uncontrolled perturbations from its environment.

The idea of bricks is that they can form a ‘wall’ around our blinker that protects it from the influence of its surrounding (which can be in an unknown state). In physical terms, the wall protects the blinker from the heat of the environment, as shown in Figure 2.

### 4.3 Reversible CA: Margolus’ billiard ball model

To get one step closer to physics and account for the bijectivity of microscopic dynamics in the physical world, we now consider reversible CAs, i.e., CAs in which every state has a unique predecessor, which is not the case for Game of Life. We now show that even reversible CAs exist that admit perfect protection of an implementation of an infinite loop, which results in zero cost per operation.

*Margolus’ billiard ball model CA* [35] is a CA in 2 dimensions whose update rules are defined on Margolus neighborhoods, i.e., there are two partitions of the grid into blocks of  $2 \times 2$  cells describing the updates at even and odd time instants: At even time instances, the update is done on the blocks  $\{(2i, 2j), (2i, 2j + 1), (2i + 1, 2j), (2i + 1, 2j + 1)\}$ , at odd times it is done on the blocks  $\{(2i - 1, 2j - 1), (2i - 1, 2j), (2i, 2j - 1), (2i, 2j)\}$ , as visualized by the black and the red grid in Figure 3, right. For each such block, the update rules are shown in Figure 3, left. To interpret such a CA with Margolus neighbourhood as a so-called Moore CA where the update rules do not change between even and odd time steps (see Subsection 3.1), we consider two time steps in the Margolus CA as one time step of a Moore CA. To ensure that the update of a cell of the Moore CA only depends on its surrounding neighbors (which is convenient for some purposes) one may consider each  $2 \times 2$  block of the Margolus CA as one cell of the Moore CA.

As noted in [30], the billiard ball CA is not physically universal since it allows for impenetrable walls [35]. We will use such walls to implement a bit switching process that continues forever although only a finite region has been initialized. A simple example is shown in Figure 4. In the sense of the present paper, this CA implements the NOT operation in a thermodynamic cycle since there are no resource requirements per operation because there is no need to initialize the cells outside the wall.




Figure 3: Left: Transition rules of Margolus' billiard ball model CA. Right: the two different partitions are indicated by the black and the red grid.




Figure 4: Configuration of Margolus billiard ball CA that implements bit switching forever: applying Rule (3) to the black partitioning takes the configuration on the left hand side to the one on the right hand side. Then, an update according to the red partitioning leaves the state unchanged due to Rule (5). Applying Rule (3) to the black partitioning takes the configuration on the right hand side back to the one on the left hand side. Again, updating according to the red partitioning has no effect.




Figure 5: Transition rules of Schaeffer's physically universal CA. Further rules are given by rotation invariance.




Figure 6: Free particle propagation in Schaeffer's physically universal CA: the configuration on the left turns into the one in the middle by applying Rule (2) to the red partitioning. The middle configuration turns into the right one by applying the same rule to the black partitioning.

## 5 Cost of operations in physically universal CAs

### 5.1 Schaeffer's physically universal CA

Schaeffer [30], see also [36], constructed an efficiently physically universal CA that is close to Margolus' billiard ball model CA. The update rules are shown in Figure 5. Here, physical universality refers to the Moore CA whose update rule consists of two time steps of the Margolus CA (following the remarks in Subsection 4.3). We now discuss a rather primitive solution of implementing our bit switching task in Schaeffer's CA. Its resource requirements grow at least linearly in  $n$ , which at first appears to be suboptimal. Yet, we will later show that linear growth is optimal. We first observe that the CA admits free particle propagation in diagonal direction, a fact that is heavily used in the proof for physical universality [30]. Figure 6 visualizes this motion. We now use a 'beam of particles' in diagonal direction in which a particle and a hole alternate, as shown in Figure 7. Then choose a target bit along the diagonal, as indicated by the blue square in Figure 7. Just by waiting, this bit is turned on and off when particles and holes appear, respectively. The resource requirements of this implementation are large: not only does it require to correctly locate particles and holes, it also requires to keep the space around the beam empty to protect the beam from collisions.

**Remark 1** (complexity aspect of preparation). *Apart from being costly from the thermodynamic point of view, the implementation is also 'not nice' in other respects: compared to the simplicity of our control problem, the initialization is rather complex. Assume, for comparison, the following general control task: given some arbitrary binary string  $b$  of length  $2n$ , the target bit is supposed to attain the value  $b_j$  at time  $j$ . Then, the above beam solves this task for the special case where  $b = 101010 \cdots 10$ . The general task can be obviously solved by the same procedure as above: just locate particles and holes according to  $b$ . The fact that the solution of the simple special case is based on the same principle suggests that it is a 'bad' solution; it is*




Figure 7: Beam of 3 propagating particles which implement the turning on and off of the blue target cell 3 times.

*inappropriately complex compared to the simplicity of the task. In a way, it reduces a simple control operation to one that seems more complex. This raises the question of what one wants to call a ‘solution’ of a control task.*

To return to the thermodynamic question, one may wonder if there exist smarter implementations of the bit switch process where the resource requirements do not grow linearly in  $n$ . We can easily show that the *range* of the implementation of the  $n$ -fold bit switch grows linearly in  $n$ . To this end, we first need the Diffusion Theorem of [30]:

**Theorem 2** (Diffusion Theorem). *Let  $S \subset \mathbb{Z}^2$  be an arbitrary square of side length  $s$  in the Moore CA and  $\phi$  an arbitrary configuration that is empty on  $\bar{S}$ . Then  $\alpha^t(\phi)$  is empty on  $S$  for all  $t \geq s$ .*

We then have:

**Theorem 3** (range of device restoring a region after  $t$  time steps). *Let  $f : \Sigma^X \rightarrow \Sigma^X$  denote an arbitrary bijection for some region  $X \subset \mathbb{Z}^2$ . Assume that  $(Z, \phi_Z, t)$  is a device for implementing  $f$ . Then its range is at least  $t$ .*

*Proof.* Let  $S$  be the smallest square containing  $W = Z \cup X$  and  $s$  its side length. We must have  $s \geq t$ . Assume to the contrary that  $s \leq t - 1$ . Then by the diffusion theorem any configuration  $\phi$  that is empty on  $\bar{S}$  evolves in  $t$  time steps to a configuration that is empty on  $S$ . This contradicts the assumption there is a configuration  $\phi_Z$  that implements a bijection on site  $X$ .  $\square$

An important special case of the above theorem is when  $f$  is the identity function ID. Moreover, we have:

**Theorem 4** (range of device implementing  $n$  powers of a transformation). *Let  $f : \Sigma^X \rightarrow \Sigma^X$  be an arbitrary bijection for some region  $X \subset \mathbb{Z}^2$ . Assume that  $(Z, \phi_Z, t_1, t_2, \dots, t_n)$  be a device for implementing of the sequence  $f, f, \dots, f$ . Then its range is at least  $n$ .*

*Proof.* The proof is very similar to the proof of Theorem 3. If  $W = Z \cup X$  were contained in a square of side length  $s \leq t_n - 1$  the configuration after  $t_n$  would be empty on  $X$ . Thus  $s \geq t_n$ . The result follows since we must have  $t_n \geq n$ .  $\square$

**Remark 2** (resources requirements for 1D physically universal CA). *We make some comments on resource requirements of the one-dimensional physically universal CA in [37]. This CA uses interacting particles that propagate with different speeds, namely  $\pm 1$  or  $\pm 2$  sites per time step. Similar results to Theorem 3 and Theorem 4 hold for this CA as well.*

*[37, Lemma 2] is also a kind of diffusion theorem similar to Theorem 2. We reformulate its statements slightly. Let  $S$  be an interval of length  $s$  and  $\phi$  a configuration that is empty on  $\bar{S}$ . Then, after  $t(s) \in O(s)$  time steps all configurations  $\phi'$  that arise from  $\phi$  under the autonomous time evolution are empty on  $S$ . It is convenient to rephrase  $t(s) \in O(s)$  as follows: there exist two constants  $s_0$  and  $\kappa$  such that  $t(s) \leq \kappa s$  for all  $s \geq s_0$ .*

*Using the same arguments but now with the one-dimensional diffusion theorem, we may conclude that for the one-dimensional CA the ranges must be at least  $t/\kappa$  and  $n/\kappa$  in Theorem 3 and Theorem 4, respectively, provided that  $X$  is sufficiently large. The latter condition on  $X$  is necessary because the diffusion theorem only applies for intervals of length at least  $s_0$ .*

The range is a rather crude measure of the resource requirements. A finer measure is the size, that is, the number of cells of the relevant region. We focus on the elementary control task of restoring a bit  $n$  times and derive a lower bound on the size of the corresponding device.

**Theorem 5** (size of device restoring a bit  $n$  times). *Let ID denote the identity on some cell of  $\mathbb{Z}^2$  in the Moore CA corresponding to Schaeffer's construction. Assume that  $(Z, \phi_Z, t_1, t_2, \dots, t_n)$  is a device for implementing ID, ID,  $\dots$ , ID. Then  $Z$  contains at least  $n/4 - 1$  cells (also counted in the Moore CA).*

*Proof.* Below, the term 'cell' refers to a cell in the Margolus CA (containing just one bit), not the  $2 \times 2$  block defining the cell of the corresponding Moore CA. Let  $X$  denote the source/target  $2 \times 2$  block. Since  $Z$  consists, by definition, of cells of the Moore CA, it consists only of complete  $2 \times 2$  blocks in the Margolus CA.

We now rely on the techniques developed in the proof of Theorem 4 in [30]. We also consider an 'abstract' CA that consists of three states  $\{0, 1, \top\}$ , where  $\top$  denotes a 'wild card' that stands for an uncertain state. The purpose of the abstract CA is merely to keep track of how uncertain states propagate in the concrete CA. Ref. [30] describes a rather simple set of update rules for the abstract CA, whose details are not needed. The essential observation that we adopt is that  $\top$  particles exhibit free particle propagation as long as the following 'forbidden' patterns



and their rotated versions do not occur. Here a grey box indicates the  $\top$  state, which stands for either the 0 state (white) or the 1 state (black). It is important that these forbidden patterns will never occur during the dynamical evolution of the abstract CA if the initial configuration does not contain any forbidden patterns [30].

First, we assign  $\top$  to all cells in  $W := Z \cup X$  representing the fact that their states are unknown or could be arbitrary (because we do not know what the correct  $\phi_Z$  looks like and the source block  $X$  could also be in any state). Second, we assign 0 to all cells in the complement of  $W$ . This is possible because cells outside  $W$  do not matter for the correct implementation.

This way, the forbidden patterns do not occur in the initial configuration and the dynamics of the abstract CA can be described by free propagation of  $\top$ -particles: each  $\top$ -particle moves to the diagonally opposite cell, that is, in either NE, NW, SE, SW direction. Consequently, any cell can attain  $\top$  and, in particular 1, at most  $4|W|$  times.

Assume one of the cells in the source region  $X$  is in the state 1 at  $t = 0$ . Consequently, it must be in the state 1 at least  $n$  times during the interval  $1, 2, \dots, t_n$  to ensure the correct implementation of the  $n$ -fold repetition of ID. By combining these two arguments together we conclude that  $W$  consists of at least  $n$  cells of the Margolus CA. Hence, it consists of at least



$n/4$  cells of the Moore CA. Since  $Z$  differs from  $W$  by only one cell we finally obtain the lower bound  $|Z| \geq n/4 - 1$ .  $\square$

Theorem 5 can easily be applied to our task of  $n$ -fold NOT since the latter amounts to implementing the identity for all  $t_j$  with even  $j$ .

It is unclear whether some of these insights apply to a general physically universal CA. The question whether there exist physically universal CAs that do not satisfy the Diffusion Theorem has already been raised by Schaeffer [30], which seems related to our thermodynamic questions since diffusion is what makes information so extremely unstable.

It is, however, clear that in any physically universal CA a configuration of a finite region is unstable in the following sense:

**Theorem 6** (instability of patterns). *For some physically universal CA, let  $Z \subset \mathbb{Z}^d$  be a finite region that is initialized to the state  $\phi_Z$ . Assume that the states of all cells of  $\bar{Z}$  are unknown and described by some probability distribution  $P$  that assigns a non-zero probability to every possible state in  $\Sigma^{\bar{Z}}$ . Then, for any configuration  $\phi'_Z$  of  $Z$  there is a time  $t$  such that  $\phi_Z$  evolves to  $\phi'_Z$  with non-zero probability.*

*Proof.* Choose a function  $f : \Sigma^{\bar{Z}} \rightarrow \Sigma^{\bar{Z}}$  with  $f(\phi_Z) = \phi'_Z$ . By physical universality, there is a configuration of the complement of  $Z$  implementing  $f$  for some  $t$ . Since only the restriction of the configuration to a finite region matters (cells that are further away than  $t$  sites do not matter) the set of all configurations implementing  $f$  has a non-zero probability.  $\square$

The absence of impenetrable walls in physically universal automata is only the most intuitive consequence of this observation. Less obvious consequences remain to be discovered in the future.

## 6 Conclusions

Common discussions on thermodynamic irreversibility of operations often focus on entropy generation while they substantially differ with respect to the underlying notion of entropy (e.g. Boltzmann entropy, Shannon respective von Neumann entropy, or Kolmogorov complexity [33, 38, 39]). Given these different notions of entropy, entropy generation is explained by coarse graining [40], because complexity also contributes to physical entropy by definition [33, 38], or because entropy leaks into the system from its environment.

Irreversibility in physically universal CAs or Hamiltonian systems is not due to entropy production – at least not in any obvious sense. Instead, every evolution is to some extent irreversible simply because one has no access to the evolution, the autonomous time evolution of the system just continues forever. Therefore, simulating the inverse evolution on some target system involves sophisticated initialization of a large number of cells in the surrounding (acting as the controller). Since this initialization is typically destroyed by the autonomous evolution of the system, restoring the state of the joint system of target *and* its controller involves a sophisticated initialization of a ‘meta-controller’, which, in turn, will then be destroyed by the evolution. The question of how to reverse the dynamics of one system without disturbing the state of its surrounding thus raises the same question for an even larger system.

The idea that control operations, even when they are unitary, imply heat generation in the controlling device, is certainly not new. However, physically universal CAs and Hamiltonians may allow us to look at the idea from a new perspective because they admit to describe target, controller, meta-controller and so on, in a unified way since all of them are just regions of cells. Moreover, physically universal CAs formalize the conflict between controllability and isolability of a system in a principled way. This is because physical universality, which formalizes the ability to control subsystems, implies instability of information, although quantitative results

have to be left to the future. Here we have shown that in the existing constructions of physically universal cellular automata information is extremely unstable – for instance, in the sense that the resources required for protecting information grows linearly in time.

The intention of this article is to inspire other researchers to explore implications of physical universality rather than exploring properties of specific constructions of CAs. Here we have discussed properties of Schaeffer’s construction only to illustrate how to work with our notion of resource requirements in the context of a physically universal CA.

**Acknowledgements:** We would like to thank Scott Aaronson and Luke Schaeffer for interesting discussions on related questions and Armen Allahverdyan for comments on an earlier version of this manuscript.

## References

- [1] R. Landauer. Irreversibility and heat generation in the computing process. *IBM J. Res. Develop.*, 5:183–191, 1961.
- [2] D. Janzing, P. Wocjan, R. Zeier, R. Geiss, and Th. Beth. Thermodynamic cost of reliability and low temperatures : Tightening Landauer’s principle and the Second Law. *Int. Jour. Theor. Phys.*, 39(12):2217–2753, 2000.
- [3] O. J. E. Maroney. Generalizing landauer’s principle. *Phys. Rev. E*, 79:031105, 2009.
- [4] Takahiro S. Thermodynamic and logical reversibilities revisited. *Journal of Statistical Mechanics: Theory and Experiment*, 2014(3):P03025, 2014.
- [5] D. H. Wolpert. Extending Landauer’s Bound from Bit Erasure to Arbitrary Computation. [arxiv:1508.05319](https://arxiv.org/abs/1508.05319), 2015.
- [6] C. H. Bennett. Logical reversibility of computation. *IBM J. Res. Develop.*, 17:525–532, 1973.
- [7] D. Janzing and Th. Beth. Are there quantum bounds on the recyclability of clock signals in low power computers? In *Proceedings of the DFG-Kolloquium VIVA*, Chemnitz, 2002. See also preprint [arXiv:quant-ph/0202059](https://arxiv.org/abs/quant-ph/0202059).
- [8] D. Janzing and T. Decker. How much is a quantum controller controlled by the controlled system? *Applicable Algebra in Engineering, Communication and Computing*, 19(3):241–258, 2008.
- [9] D. Janzing and B. Steudel. Quantum broadcasting problem in classical low power signal processing. *Phys. Rev.*, A(75):022309, 2007.
- [10] D. Janzing and T. Beth. Bounds on the entropy generated when timing information is extracted from microscopic systems. preprint [arXiv:quant-ph/0301125](https://arxiv.org/abs/quant-ph/0301125), 2003.
- [11] D. Janzing and T. Beth. Synchronizing quantum clocks with classical one-way communication: Bounds on the generated entropy. preprint [arXiv:quant-ph/0306023v1](https://arxiv.org/abs/quant-ph/0306023v1), 2003.
- [12] P. Erker, M. Mitchison, R. Silva, M. Wood, Brunner M., and M. Huber. Autonomous quantum clocks: how thermodynamics limits our ability to measure time. preprint [arXiv:1609.06704](https://arxiv.org/abs/1609.06704), 2016.
- [13] D. Janzing and T. Beth. Quasi-order of clocks and their synchronism and quantum bounds for copying timing information. *IEEE Transactions Information Theory*, 49(1):230–240, 2003.
- [14] S. Deffner and C. Jarzynski. Information processing and the second law of thermodynamics: An inclusive, hamiltonian approach. *Phys. Rev. X*, 3:041003, 2013.

- [15] J. Horowitz and M. Esposito. Thermodynamics with continuous information flow. *Phys. Rev. X*, 4:031015, 2014.
- [16] N. Margolus. Parallel quantum computation. In W. Zurek, editor, *Complexity, Entropy, and the Physics of Information*. Addison Wesley Longman, 1990.
- [17] D. Janzing and P. Wocjan. Ergodic quantum computing. *Quant. Inf. Process.*, 4(2):129–158, 2005.
- [18] D. Janzing. Spin-1/2 particles moving on a 2D lattice with nearest-neighbor interactions can realize an autonomous quantum computer. *Phys. Rev. A*, 75:012307, 2007.
- [19] D. Janzing. Is there a physically universal cellular automaton or Hamiltonian? preprint [arXiv:1009.1720](https://arxiv.org/abs/1009.1720).
- [20] Omnès, R. *The interpretation of quantum mechanics*. Princeton Series in Physics. Princeton University Press, 1994.
- [21] A. Adams, A. Berner, P. daview, and S. Walker. Physical universality, state-dependent dynamical laws and open-ended novelty. *Entropy*, 19(461):1–20, 2017.
- [22] L. Schaeffer. A physically universal quantum cellular automaton. In J. Kari, editor, *Proceedings of 21st IFIP WG 1.5 International Workshop, AUTOMATA 2015*, pages 46–58, Heidelberg, 2015. Springer.
- [23] D. Giulini, E. Joos, C. Kiefer, J. Kupsch, I.-O. Stamatescu, and H.D. Zeh. *Decoherence and the Appearance of a Classical World in Quantum Theory*. Springer, Berlin, 1996.
- [24] J. von Neumann. *Theory of self-reproducing automata*. University of Illinois Press, Champaign, 1966.
- [25] M. Nielsen and I. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [26] J. Fernandez, S. Lloyd, T. Mor, and V. Roychowdhury. Algorithmic cooling of spins: A practicable method for increasing polarization. *Int. Journ. Quant. Inf.*, 2(4):461–467, 2004.
- [27] A. Allahverdyan, K. Hovhannisyanyan, D. Janzing, and G. Mahler. Thermodynamic limits of dynamic cooling. *Phys. Rev. E*, E(48):041109, 2011.
- [28] D. Janzing. On the computational power of molecular heat engines. *J. Stat. Phys.*, 122(3):531–556, 2006.
- [29] A. E. Allahverdyan, S. G. Babajanyan, N. H. Martirosyan, and A. V. Melkikh. Adaptive heat engine. *Phys. Rev. Lett.*, 117:030601, 2016.
- [30] L. Schaeffer. A physically universal cellular automaton. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science, ITCS '15*, pages 237–246, New York, NY, USA, 2015. ACM.
- [31] J. Pearl. *Causality: Models, reasoning, and inference*. Cambridge University Press, 2000.
- [32] C. P. Kempes, D. Wolpert, Z. Cohen, and J. Pérez-Mercader. The thermodynamic efficiency of computations made in cells across the range of life. *Philosophical Transactions of the Royal Society of London Series A*, 375:20160343, 2017.
- [33] C. H. Bennett. Time/space trade-offs for reversible computation. *SIAM J. Computing*, 18(4):766–776, 1989.
- [34] Game of life. Wikipedia. accessed 22 May 2016.
- [35] N. Margolus. Physics-like models of computation. *Physica D: Nonlinear Phenomena*, 10(1):81–95, 1984.
- [36] L. Schaeffer. A physically universal cellular automaton (website). <http://web.mit.edu/lrs/www/physCA/>, accessed February 2016.

- [37] V. Salo and I. Törmä. A one-dimensional physically universal cellular automaton. preprint [arXiv:1501.03988](https://arxiv.org/abs/1501.03988).
- [38] W. Zurek. Algorithmic randomness and physical entropy. *Phys Rev A*, 40(8):4731–4751, 1989.
- [39] D. Janzing, R. Chaves, and B. Schölkopf. Algorithmic independence of initial condition and dynamical law in thermodynamics and causal inference. *New Journal of Physics*, 18(093052):1–13, 2016.
- [40] R. Balian. *From microphysics to macrophysics*, volume 1. Springer, 2007.