5 Chapter 5: Topological Dynamics of Cellular Automata

In this chapter, we study the dynamical properties of global maps of cellular automata. To this end, we endow the space of configurations with a natural topology, in which two configurations are close when they agree on a large finite region around the origin. This allows us to view cellular automata as self maps on a compact topological space and leads to a fundamental characterization: cellular automata are precisely the continuous maps that commute with all shifts.

Once this framework is in place, we present several classical results about the global dynamics of cellular automata. In particular, we discuss the interplay between injectivity, surjectivity and reversibility, and present the famous Garden-of-Eden theorem, which gives a precise relationship between surjectivity and a weaker form of injectivity.

We then discuss some standard topological notions related to the sensitivity to small perturbations of initial configurations and very briefly show how this leads to a topological classification of cellular automata in one dimension.

Finally, we introduce subsystems, factors, and embeddings as canonical ways to study the richness of a cellular automaton's dynamics. These notions will play an important role in the next chapter, where we focus on the computational capacity of cellular automata.

This chapter is mainly based on the excellent notes by Jarkko Kari [2] and on a more in-depth textbook by Petr Kůrka [3].

5.1 The Topology of CA Configuration Space

We can endow the space $S^{\mathbb{Z}^d}$ with a metric as follows. We define the distance between two configurations $c_1, c_2 \in S^{\mathbb{Z}^d}$ as:

$$\mathrm{d}(c_1, c_2) = \begin{cases} 0 & \text{if } c_1 = c_2 \\ 2^{-\min\{\|v\| \mid c_1(v) \neq c_2(v)\}} & \text{if } c_1 \neq c_2 \end{cases}.$$

Thus, two configurations are close if they agree on a large finite region around the origin 0.

Exercise 5.1 (Metric Space of Configurations)**.** *Prove that the function defined above is a metric. I.e., show that* d *satisfies for all* $c_1, c_2, c_3 \in S^{\mathbb{Z}^d}$ *the following:*

(a) $\mathrm{d}(c_1, c_2) = 0 \iff c_1 = c_2$ *and* $\mathrm{d}(c_1, c_2) \geq 0$.

(b) $\mathrm{d}(c_1, c_2) = \mathrm{d}(c_2, c_1)$

(c) $\mathrm{d}(c_1, c_3) \leq \mathrm{d}(c_1, c_2) + \mathrm{d}(c_2, c_3)$.

For $r \in \mathbb{N}$ we define the ball of radius 2^{-r} centred at $c \in S^{\mathbb{Z}^d}$ as

$$B_r(c) = \{c' \in S^{\mathbb{Z}^d} \mid c'(v) = c(v) \text{ for all } v \in \mathbb{Z}^d \text{ such that } \|v\| \leq r\}.$$

Thus, $B_r(c)$ consists of all configurations that agree with c on the finite box $[-r, r]^d = \{v \in \mathbb{Z}^d \mid \|v\| \leq r\}$. The collection of all such balls forms a basis for the topology induced by the metric d. Fixing $r \in \mathbb{N}$, we notice there are only finitely many unique balls $B_r(c)$, $c \in S^{\mathbb{Z}^d}$. Therefore, this basis is countable.

Example 5.2 (Example of an open set)**.** *Let* $d = 1$ *and* $S = \{0, 1\}$. *Consider the set*

$$U = \{c \in S^{\mathbb{Z}} \mid c(-1) = 1,\ c(2) = 0\}.$$

This is an open set: if $c \in U$, then every configuration c' that agrees with c on $[-2, 2]$ also lies in U. In other words, $B_2(c) \subseteq U$. It is easy to see that U is a union of finitely many balls of the form $B_2(x)$.

The example above is an instance of a cylinder set: a basic open set defined by prescribing finitely many values of a configuration.

Definition 5.3. *Let $D \subseteq \mathbb{Z}^d$ be finite and let $c \in S^{\mathbb{Z}^d}$. The* cylinder set *determined by c on D is*

$$\mathrm{cyl}(c, D) = \{c' \in S^{\mathbb{Z}^d} \mid c'(v) = c(v) \text{ for all } v \in D\}.$$

Example 5.4. *Every cylinder set is open. Indeed, let $D \subseteq \mathbb{Z}^d$ be finite and $c' \in \mathrm{cyl}(c, D)$. Choose r so large that $D \subseteq [-r, r]^d$. Then every configuration $x \in B_r(c')$ agrees with c', and hence with c, on D. Therefore*

$$B_r(c') \subseteq \mathrm{cyl}(c, D),$$

which shows that $\mathrm{cyl}(c, D)$ is open.

Conversely, every basic open set $B_r(c)$ is itself a cylinder set, namely

$$B_r(c) = \mathrm{cyl}\big(c, [-r, r]^d\big).$$

So the cylinder sets form a basis for the topology on $S^{\mathbb{Z}^d}$.

The following observation gives a convenient characterization of convergence in this metric space.

Observation 5.5. *A sequence $(c^i)_{i=0}^{\infty}$ in $S^{\mathbb{Z}^d}$ converges to $c^* \in S^{\mathbb{Z}^d}$ if and only if for every finite set $D \subseteq \mathbb{Z}^d$, the restrictions $c^i|_D$ eventually agree with $c^*|_D$. Equivalently, for every $v \in \mathbb{Z}^d$, the sequence*

$$c^1(v),\ c^2(v),\ c^3(v),\ \ldots$$

is eventually constant with value $c^*(v)$.

Exercise 5.6. *Work out a proof for the above observation.*

It should now not be difficult to see that the configuration space together with the metric we have defined is compact, as is stated in the exercise below.

Exercise 5.7 (Compactness of the Configuration Space)**.** *Show that the metric space $(S^{\mathbb{Z}^d}, \mathrm{d})$ is compact. That is, show that every sequence of configurations $(c^i)_{i=0}^{\infty}$ has a convergent subsequence.* Hint: *Enumerate $\mathbb{Z}^d = \{v_1, v_2, v_3, \ldots\}$. Since S is finite, one can repeatedly pass to subsequences so that the values at v_1, then at v_2, then at v_3, and so on, eventually stabilize.*

Remark 5.8. *There is another natural way to define the topology on $S^{\mathbb{Z}^d}$. Since S is finite, we equip it with the discrete topology, in which every subset of S is open. The product topology on $S^{\mathbb{Z}^d}$ is then generated by cylinder sets. Hence, the product topology and the topology induced by* d *coincide.*

In particular, since every finite topological space is compact, the compactness of $S^{\mathbb{Z}^d}$ also follows from Tychonoff's theorem.

This topology is designed so that agreement on a large finite region characterizes closeness of two configurations. As a result, continuity of a map $F : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ expresses exactly the idea that finite information about the image is determined by finite information about the input. To be explicit, a map $F : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ is continuous if for every $c \in S^{\mathbb{Z}^d}$ and every $\varepsilon > 0$, there exists $\delta > 0$ such that

$$\mathrm{d}(c, c') < \delta \implies \mathrm{d}\big(F(c), F(c')\big) < \varepsilon.$$

In particular, if F is continuous and $c \in S^{\mathbb{Z}^d}$, then there exists $r \in \mathbb{N}$ such that whenever c' agrees with c on $[-r, r]^d$, we have

$$F(c')(0) = F(c)(0).$$

Thus, continuity says that the value at the origin is determined by finitely many input cells, although at this stage the required finite region may still depend on the configuration c.

It is beyond the scope of these notes to give a general introduction to compact metric spaces. Below, we list a few standard properties of compact metric spaces that we will use in later proofs. For readers who would like a brief overview, we recommend Section 5.1 of [2] or Appendix A of [3].

Proposition 5.9. *Let (X, d) be a compact metric space. Then, the following statements hold.*

1. *Let $F : X \to X$ be a continuous mapping. Then, F is uniformly continuous. I.e., for every $\varepsilon > 0$ there exists a $\delta > 0$ such that for all $c, c' \in X$:*

$$\mathrm{d}(c, c') < \delta \implies \mathrm{d}\left(F(c), F(c')\right) < \varepsilon.$$

2. *Let $K_1 \supseteq K_2 \supseteq K_3 \supseteq \dots$ be a sequence of non-empty closed subsets of X. Then, $\bigcap_{i=1}^{\infty} K_i \neq \emptyset$.*
3. *Let $F : X \to X$ be a continuous mapping which is bijective. Then, F^{-1} is also continuous.*

We are now ready to prove a fundamental characterization of cellular automata: their global maps are exactly the continuous self-maps of $S^{\mathbb{Z}^d}$ that commute with all shifts.

Theorem 5.10 (Curtis–Hedlund–Lyndon Theorem)**.** *Let S be a finite set and let $d \in \mathbb{N}$. A function $F : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ is the global rule of a cellular automaton if and only if it is continuous and commutes with every shift operator.*

Proof. $\implies$ Suppose $F : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ is the global rule of a cellular automaton with radius r and local rule f. We have already observed that F commutes with every shift operator, so it remains to prove continuity.

Let $n \in \mathbb{N}$ and let $c \in S^{\mathbb{Z}^d}$. We claim that if c' agrees with c on the box $[-(n+r), n+r]^d$, then $F(c')$ agrees with $F(c)$ on $[-n, n]^d$. Indeed, let $v \in \mathbb{Z}^d$ satisfy $\|v\| \le n$. To compute $F(c)(v)$, the local rule only inspects the values of c on cells $v + u$ with $\|u\| \le r$. For such u, we have

$$\|v + u\| \le \|v\| + \|u\| \le n + r,$$

so c and c' agree on all cells needed to compute the value at v. Hence $F(c)(v) = F(c')(v)$. Therefore, for all $n \in \mathbb{N}$

$$\mathrm{d}(c, c') < 2^{-(n+r)} \implies \mathrm{d}\left(F(c), F(c')\right) < 2^{-n},$$

which shows that F is continuous.

$\impliedby$ Now suppose that $F\colon S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ is continuous and commutes with all shifts. Since $S^{\mathbb{Z}^d}$ is compact, Proposition 5.9 implies that F is uniformly continuous.

Apply uniform continuity with $\varepsilon = 1$. Then there exists $r \in \mathbb{N}$ such that for all $c, c' \in S^{\mathbb{Z}^d}$,

$$\mathrm{d}(c, c') < 2^{-r} \quad \Longrightarrow \quad \mathrm{d}\big(F(c), F(c')\big) < 1.$$

Since $\mathrm{d}(x, y) < 1$ holds if and only if $x(0) = y(0)$, it follows that whenever c and c' agree on the box $[-r, r]^d$, we have

$$F(c)(0) = F(c')(0).$$

Thus, the value $F(c)(0)$ depends only on the ordered tuple of states $\big(c(n_1), \ldots, c(n_k)\big)$, where $N_r = (n_1, \ldots, n_k)$ is a fixed ordering of the sites in the box $[-r, r]^d$. We may therefore define a local rule $f : S^k \to S$ by

$$f(s_1, \ldots, s_k) = F(c)(0),$$

where c is any configuration satisfying $c(n_i) = s_i$ for all $1 \le i \le k$. The previous paragraph shows that this is well-defined.

Finally, let $c \in S^{\mathbb{Z}^d}$ and $v \in \mathbb{Z}^d$. Since F commutes with shifts, we have

$$F(c)(v) = \sigma_v(F(c))(0) = F(\sigma_v(c))(0).$$

The value on the right depends only on the restriction of $\sigma_v(c)$ to N_r, that is, only on the values $(c(v + n_1), \ldots, c(v + n_k))$. Hence

$$F(c)(v) = f\big(c(v + n_1), \ldots, c(v + n_k)\big),$$

so F is the global rule of a cellular automaton with neighbourhood N_r and local rule f. □

Exercise 5.11 (Both assumptions in the Curtis–Hedlund–Lyndon Theorem are essential)**.** *Construct an example of a map $F : \mathbf{2}^{\mathbb{Z}} \to \mathbf{2}^{\mathbb{Z}}$ such that:*

(a) F is continuous, but is not a global rule of any cellular automaton.

(b) F commutes with every shift operator, but is not a global rule of any cellular automaton.

As a first application of the Curtis–Hedlund–Lyndon theorem, we consider reversibility. Reversible cellular automata are of particular interest because many microscopic laws of physics are reversible. They therefore provide a natural class of discrete models for systems in which the evolution does not destroy information.

Definition 5.12 (Reversible Cellular Automata)**.** *We say that a cellular automaton $\mathcal{A} = (S^{\mathbb{Z}^d}, F)$ is* reversible *if $F : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ is bijective and $(S^{\mathbb{Z}^d}, F^{-1})$ is also a cellular automaton.*

Proposition 5.13. *Every bijective cellular automaton is reversible.*

Proof. Let $\mathcal{A} = (S^{\mathbb{Z}^d}, F)$ be a CA such that $F : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ is bijective. All we have to do is to verify that $F^{-1} : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ is a continuous map that commutes with all the shifts. Since F is continuous and $S^{\mathbb{Z}^d}$ is a compact space, it follows immediately that F^{-1} is also continuous. We check the commuting property now. Let $v \in \mathbb{Z}^d$ and $c \in S^{\mathbb{Z}^d}$. Since F is bijective, there exists some $c' \in S^{\mathbb{Z}^d}$ such that $F(c') = c$. Then

$$F^{-1}\left(\sigma_v(c)\right) = F^{-1}\left(\sigma_v\left(F(c')\right)\right) = F^{-1}\left(F\left(\sigma_v(c')\right)\right) = \sigma_v(c') = \sigma_v\left(F^{-1}(c)\right).$$

So F^{-1} commutes with every shift. By the Curtis–Hedlund–Lyndon theorem, F^{-1} is the global rule of a cellular automaton. □

5.2 Injectivity, Surjectivity, and the Garden-of-Eden Theorem

When considering the global map of a cellular automaton, several basic questions naturally arise: when is the global map injective, surjective, or reversible, and what are the relationships between these properties? Injectivity, surjectivity, and reversibility are among the most fundamental global properties of a cellular automaton, and their study leads to one of the central classical results in the theory: the Garden-of-Eden theorem. For the rest of this section, let us fix a CA $\mathcal{A} = (S^{\mathbb{Z}^d}, F)$ with radius r and local rule f.

Definition 5.14 (Injective CAs, surjective CAs, Garden-of-Eden configurations)**.** *A cellular automaton $\mathcal{A} = (S^{\mathbb{Z}^d}, F)$ is* injective *if its global rule $F : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ is injective. It is* surjective *if $F : S^{\mathbb{Z}^d} \to S^{\mathbb{Z}^d}$ is surjective.*

A configuration $c \in S^{\mathbb{Z}^d}$ is called a Garden-of-Eden configuration *if $F^{-1}(c) = \emptyset$.*

Clearly, a cellular automaton is surjective if and only if it has no Garden-of-Eden configurations. At first sight, surjectivity is a genuinely global property: to verify it, one has to know that every configuration has a preimage. A remarkable feature of cellular automata is that this question can already be detected locally. To make this precise, we first introduce the notion of a finite pattern.

Definition 5.15 (Patterns)**.** *Let $D \subseteq \mathbb{Z}^d$ be a finite set. A* pattern *is a pair $\mathcal{P} = (p, D)$, where $p : D \to S$ assigns a state to each site in D. We say that a configuration $c \in S^{\mathbb{Z}^d}$* contains *the pattern $\mathcal{P}$ if there exists $v \in \mathbb{Z}^d$ such that $c(v + u) = p(u)$ for all $u \in D$. In other words, c contains $\mathcal{P}$ if $\mathcal{P}$ appears in some translation of c.*

Definition 5.16 (Orphans)**.** *A finite pattern $\mathcal{P}$ is called an* orphan *if every configuration containing $\mathcal{P}$ is a Garden-of-Eden configuration.*

Thus, an orphan is a finite witness of non-surjectivity: once such a pattern appears in a configuration, that configuration cannot have a preimage. An important consequence of the compactness of $S^{\mathbb{Z}^d}$ is that every Garden-of-Eden configuration contains an orphan. We prove this below.

Proposition 5.17. *Let $\mathcal{A} = (S^{\mathbb{Z}^d}, F)$ be a CA with radius r and local rule f and suppose that $c \in S^{\mathbb{Z}^d}$ is a Garden-of-Eden configuration. Then, c contains an orphan.*

Proof. We will suppose that c contains no orphan, and use this to construct a preimage of c under F.

For each $n \in \mathbb{N}$, let $C_n = B_n(c)$. Therefore, $C_1 \supseteq C_2 \supseteq C_3 \supseteq \dots$ is a sequence of balls with shrinking radii. Moreover, as each ball is both open and closed, we have that each C_n is closed.

Clearly, since for any $n \in \mathbb{N}$ the pattern given by $c\big|_{[-n,n]^d}$ cannot be an orphan, we have that for all $n \in \mathbb{N}$

$$D_n = F^{-1}(C_n) \neq \emptyset.$$

Since F is continuous and each C_n is closed, D_n is also closed. Clearly, it holds that

$$D_1 \supseteq D_2 \supseteq D_3 \supseteq \dots.$$

Again, using the compactness of the space, we have that there exists

$$c' \in \bigcap_{n=1}^{\infty} D_n.$$

We claim that c' is a preimage of c. Indeed, for every $n \in \mathbb{N}$ we have that $F(c') \in C_n$, and therefore

$$F(c')\big|_{[-n,n]^d} = c\big|_{[-n,n]^d}.$$

Since this holds for every $n \in \mathbb{N}$, it follows that $F(c') = c$. $\square$

The notion of orphans brings us one step closer to understanding which patterns are constructible by a self-replicator. Indeed, an orphan cannot be produced by the dynamics of the cellular automaton, since any configuration containing it has no preimage. Thus, orphans define a hard limit on what a self-replicator capable of universal construction can build.

Despite knowing that each non-surjective CA admits a finite witness, it is not easy to prove that a given CA is non-surjective as the search for orphans can be computationally very demanding. Indeed, Game of Life is not surjective, yet the smallest known orphan (due to Steven Eker, 2016) measured by its bounding-box area has a rectangular shape of size 8×12 cells. In what follows, we show that surjectivity is equivalent to other criteria that can prove very useful when showing a given CA is not surjective. The first one we discuss is the balancedness of the CA. We first illustrate this on an example.

Example 5.18 (Imbalance of Rule 110, Example from Kari's scripts)**.** *Consider Rule 110. Among the eight possible neighbourhood configurations, three are mapped to state* 0 *and five are mapped to state* 1*. Let us demonstrate how this imbalance automatically implies the existence of an orphan. The idea is straightforward: we amplify this imbalance by creating a family of patterns with many repeated* 0*'s for which the number of preimage candidates is strictly smaller.*

Let k be a positive integer, and consider a finite pattern u with domain $[0, 3k-3]$ in which

$$u(0) = u(3) = \cdots = u(3k-3) = 0.$$

See Figure 1 for an illustration. There are $2^{2(k-1)} = 4^{k-1}$ possible choices for the missing states between the 0*'s in u (shown as $*$ in Figure 1).*

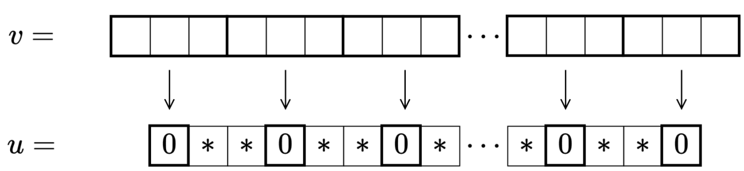

Figure 1: Illustration of the pattern c and its pre-image b.

*Now let w be a preimage pattern of u with domain $[-1, 3k-2]$ (i.e., once we apply f to every consecutive triple of w, we obtain u). Then, for every $i = 0, 3, \ldots, 3k-3$, the length-*3 *segment $w(i-1)w(i)w(i+1)$ must be mapped to state* 0 *by the local rule f. Since $|f^{-1}(0)| = 3$, each such segment has exactly three possible choices. Moreover, these k segments are pairwise disjoint, so the choices are independent. Therefore, there are exactly 3^k possibilities for these segments.*

If k is sufficiently large, then $3^k < 4^{k-1}$. Hence, there must exist some choice of the intermediate states in u for which no corresponding preimage pattern w exists. Therefore, u is an orphan and Rule 110 is not surjective.

However, it does not suffice to characterize surjectivity in terms of the balanced rule table itself, as illustrated by the following example. The correct notion of balancedness concerns not only the local rule table, but also all of its "unravelled forms on finite patterns" as we will discuss below.

Example 5.19 ("Traffic Cellular Automaton")**.** *Consider ECA Rule 184 with its rule table and dynamics illustrated in Figure 2.*

Rule 184

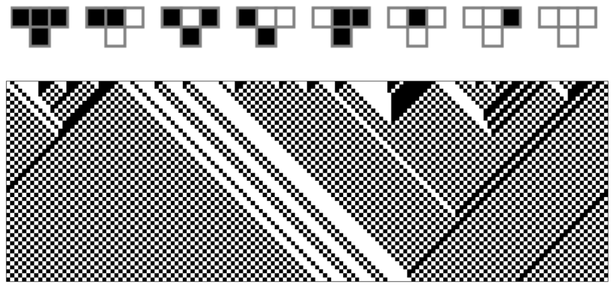

Figure 2: Rule 184 rule table and space-time diagram.

This CA can be viewed as a simplistic model of a traffic jam: each black cell can be interpreted as a presence of a car, each white square as its absence. The car moves right whenever there is a free space, otherwise it stays in place. Clearly, the rule table is balanced; however, this CA is not surjective. Verify that 1100 is its orphan.

Below we formally define the unravelling of a CA's local rule. Though it may sound slightly cumbersome, the notion is very intuitive: the CA's local rule is defined on patterns spanning the box $[-r,r]^d$. We can generalize this to a map $\tilde{f}$ which works for any finite pattern with domain $[-r,k+r-1]^d$ for an arbitrary k simply by applying the local rule f synchronously to each subdomain of the form $[-r,r]^d$.

Definition 5.20 (Unravelling a local function, Balanced CA). *Let $\mathcal{A} = (S^{\mathbb{Z}^d}, F)$ be a CA with radius r and local rule $f : S^{(2r+1)^d} \to S$. We define its unravelled form $\tilde{f}$ which operates on finite patterns as follows. Let $k \in \mathbb{N}^+$ and consider the domains $D = [0,k-1]^d$ and $D' = [-r,k+r-1]^d$. Then*

$$\tilde{f} : S^{D'} \to S^D$$

is simply defined as

$$\tilde{f}(c\big|_{D'}) = F(c)\big|_D$$

for any $c \in S^{\mathbb{Z}^d}$. This is indeed well-defined since for any two configurations c, c' that agree on D' it holds that $F(c)\big|_D = F(c')\big|_D$.

For each $k \in \mathbb{N}^+$ the function $\tilde{f}$ maps a set of size $|S|^{(k+2r)^d}$ into a set of size $|S|^{k^d}$. Therefore, on average every pattern with domain $D = [0,k-1]^d$ has $|S|^{(k+2r)^d - k^d}$ preimages under $\tilde{f}$. We say that a CA is balanced *if every pattern with domain of the form $D = [0,k-1]^d$ has exactly $|S|^{(k+2r)^d - k^d}$ preimages for each $k \in \mathbb{N}^+$.*

In dimension one, the situation simplifies as we can naturally identify every pattern $p : [0,k-1] \to S$ with a word $p(0)p(1)\cdots p(k-1) \in S^k$. Then, for each $k \in \mathbb{N}^+$ we can interpret $\tilde{f} : S^{k+2r} \to S^k$. Thus, a balanced CA has exactly $|S|^{2r}$ preimages of every finite pattern $p : [0,k-1] \to S$ under $\tilde{f}$. Notice that in dimension 1 the number of preimages does not depend on the domain's size. We illustrate this in Figure 3.

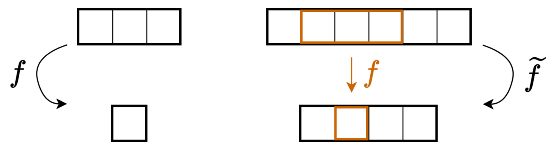

Figure 3: Unravelled local rule of a one-dimensional CA with radius 1.

Proposition 5.21. *Let $\mathcal{A} = (S^{\mathbb{Z}^d}, F)$ be a CA with radius r and local rule f. Then, $\mathcal{A}$ is balanced if and only if it is surjective.*

Proof. We show the proof in dimension one as the notation becomes much simpler there, and refer the reader to [2, p. 20–24] for the general result. Thus, in this proof we identify patterns on intervals with words, and for each $k \in \mathbb{N}^+$ we regard $\tilde{f}$ as the map

$$\tilde{f} : S^{k+2r} \to S^k.$$

Assume first that $\mathcal{A}$ is not balanced. Then there exists some $n \in \mathbb{N}^+$ for which not every word in S^n has exactly $|S|^{2r}$ preimages under $\tilde{f} : S^{n+2r} \to S^n$. Since the average number of preimages is $|S|^{2r}$, there must exist a word $u \in S^n$ such that

$$m = |\tilde{f}^{-1}(u)| < |S|^{2r}.$$

We again use the amplifying principle from Example 5.18. For each $\ell \in \mathbb{N}^+$, consider the set

$$U_\ell = \{u\, a_1\, u\, a_2 \cdots a_{\ell-1} u \mid a_i \in S^{2r} \text{ for all } i\} \subseteq S^{\ell n + 2r(\ell-1)}.$$

Clearly,

$$|U_\ell| = |S|^{2r(\ell-1)}.$$

Now let $w \in U_\ell$ and suppose that $b \in S^{\ell(n+2r)}$ satisfies $\tilde{f}(b) = w$. Since the copies of u in w are separated by words of length $2r$, the ℓ windows of b of length $n+2r$ that determine these copies of u are pairwise disjoint; in fact, they partition b into ℓ consecutive blocks of length $n+2r$. Each of these blocks must belong to $\tilde{f}^{-1}(u)$. Therefore, there are at most m^ℓ possible preimages of words from U_ℓ.

Since $m < |S|^{2r}$, there exists $\ell \in \mathbb{N}^+$ large enough such that

$$m^\ell < |S|^{2r(\ell-1)} = |U_\ell|.$$

Hence there exists a word $w \in U_\ell$ with no preimage under $\tilde{f}$. Such a word is an orphan and therefore, $\mathcal{A}$ is not surjective.

Conversely, assume that $\mathcal{A}$ is balanced. Then every word on a finite interval has a preimage under the corresponding unravelled map $\tilde{f}$. Hence, there is no orphan whose domain is an interval. In dimension one, this already implies that there are no orphans at all: if there were an orphan on an arbitrary finite domain, then any completion of it to the smallest interval containing its domain would again be an orphan. Thus, $\mathcal{A}$ has no orphans. By Proposition 5.17, $\mathcal{A}$ has no Garden-of-Eden configurations and is therefore surjective. □

Below, we establish yet another property characterizing surjectivity: the so-called pre-injectivity, which is a certain weaker form of injectivity. This characterization is exactly the content of the Garden-of-Eden Theorem. In particular, this result implies that for cellular automata, injectivity is a stronger property than surjectivity.

Let $\mathcal{A} = (S^{\mathbb{Z}^d}, F)$ be a CA with local rule f. We say that a state $q \in S$ is *quiescent* if it holds that $f(q, q, \ldots, q) = q$. Fixing a quiescent state, we say that a configuration $c \in S^{\mathbb{Z}^d}$ is *finite* if $\{v \in \mathbb{Z}^d \mid c(v) \neq q\}$ is finite. Let Fin $\subset S^{\mathbb{Z}^d}$ denote the set of all finite configurations. Clearly, $F\big|_{\text{Fin}} : \text{Fin} \to \text{Fin}$.

Definition 5.22 (Pre-injective Cellular Automata)**.** *A CA $\mathcal{A} = (S^{\mathbb{Z}^d}, F)$ is pre-injective, if for every pair of distinct configurations $c, c' \in S^{\mathbb{Z}^d}$ such that $\{v \in \mathbb{Z}^d \mid c(v) \neq c'(v)\}$ is finite, it holds that $F(c) \neq F(c')$. We call such configurations that differ only in finitely many cells* asymptotic configurations.

Observation 5.23. *Injective $\implies$ pre-injective.*

The following proposition shows that for cellular automata with a quiescent state, pre-injectivity of a CA global rule F is in fact equivalent to injectivity on the set of its finite configurations Fin.

Proposition 5.24. *Let $\mathcal{A} = (S^{\mathbb{Z}^d}, F)$ be a CA with a quiescent state that we fix. Then, $\mathcal{A}$ is pre-injective if and only if $F\big|_{\text{Fin}} : \text{Fin} \to \text{Fin}$ is injective.*

Proof. Clearly, pre-injectivity implies the injectivity of F when restricted to the set of its finite configurations, since any two configurations from Fin are asymptotic.

Now, assume that $\mathcal{A}$ is not pre-injective. Then, there are two asymptotic configurations $c_1, c_2 \in S^{\mathbb{Z}^d}$ such that $F(c_1) = F(c_2)$. We will transform them into two finite configurations $c_1', c_2' \in$ Fin such that $F(c_1') = F(c_2')$. Since c_1 and c_2 are asymptotic, there exists a $k \in \mathbb{N}^+$ such that for all $v \in \mathbb{Z}^d$ with $\|v\| \geq k$ we have that $c_1(v) = c_2(v)$. We finish this transformation by simply replacing the original states with q on all positions "far enough" to preserve the collision of the configurations. We illustrate this in the one-dimensional case in Figure 4. Formally, let r be the radius of $\mathcal{A}$. We

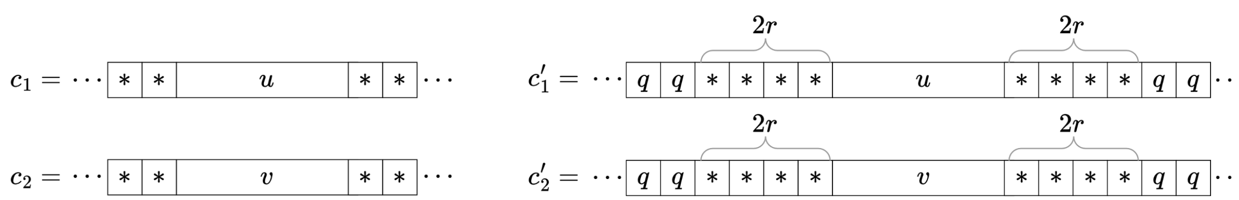

Figure 4: Substituting identical states in c_1 and c_2 with the quiescent state while preserving the collision.

define $c_1' \in$ Fin so that for all $v \in \mathbb{Z}^d$ such that $\|v\| \leq k + 2r$ we have that $c_1'(v) = c_1(v)$ and otherwise is equal to q. We define c_2' analogously. It is straightforward to verify that:

- for all $v \in \mathbb{Z}^d$ such that $\|v\| \leq k + r$: $F(c_1')(v) = F(c_1)(v) = F(c_2)(v) = F(c_2')(v)$.
- for all $v \in \mathbb{Z}^d$ such that $\|v\| > k + r$: $F(c_1')(v) = F(c_2')(v)$ since in this region, c_1' and c_2' are identical.

Thus, $F\big|_{\text{Fin}}$ is not injective. □

Finally, we present the famous Garden-of-Eden Theorem.

Theorem 5.25 (Garden-of-Eden)**.** *A cellular automaton $\mathcal{A} = (S^{\mathbb{Z}^d}, F)$ is surjective if and only if it is pre-injective.*

Proof. The proofs in both directions follow a counting argument and a certain variant of the "amplification" strategy. Again, for simplicity of notation, we only show the proof for dimension one; for the general case we refer the reader to [2]. Let r denote the CA's radius and $q \in S$ the fixed quiescent state.

pre-injectivity $\implies$ surjectivity: Suppose that $\mathcal{A}$ is not surjective. Then there exists a Garden-of-Eden configuration, and by Proposition 5.17 it contains an orphan. In dimension one, by extending this pattern to the smallest interval containing its domain, we may assume that the orphan is a word $u \in S^n$ of length n. Let $q \in S$ be an arbitrary state and let $q' = f(q, q, \ldots, q)$.

For each $k \in \mathbb{N}^+$ such that $kn > 2r$, consider the set

$$C_k = \{c \in S^{\mathbb{Z}} \mid c_i = q \text{ for all } i < r \text{ and } i \geq kn - r\}.$$

Each configuration from C_k is determined by its values on the interval $[r, kn - r - 1]$, so

$$|C_k| = |S|^{kn-2r}.$$

Every configuration in $F(C_k)$ is equal to q' outside the interval $[0, kn - 1]$. We illustrate this situation in Figure 5. Hence, an element of $F(C_k)$ is completely determined by the k consecutive blocks of length n that partition this interval. None of these blocks can be equal to the orphan u, because every configuration in $F(C_k)$ has a preimage. Therefore,

$$|F(C_k)| \leq (|S|^n - 1)^k.$$

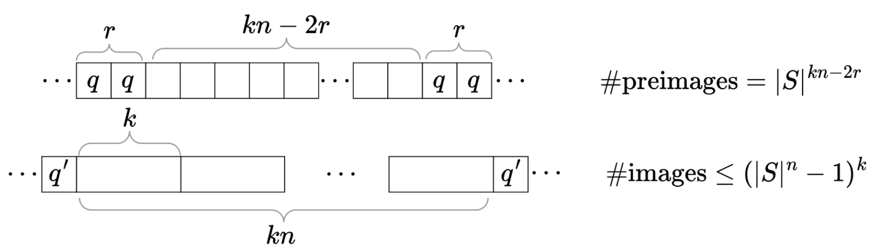

Figure 5: Constructing two colliding asymptotic configurations, using the fact that there is an orphan of size n.

By Lemma 5.26 for the case $d = 1$, there exists $k \in \mathbb{N}^+$ such that

$$(|S|^n - 1)^k < |S|^{kn-2r} = |C_k|.$$

Thus, $|F(C_k)| < |C_k|$. Hence there exist distinct configurations $c_1, c_2 \in C_k$ such that

$$F(c_1) = F(c_2).$$

Since both c_1 and c_2 are equal to q outside the finite interval $[r, kn - r - 1]$, they are asymptotic. This contradicts pre-injectivity.

surjectivity $\implies$ pre-injectivity: Suppose $\mathcal{A}$ is not pre-injective and there exist two asymptotic configurations $c_1, c_2 \in S^{\mathbb{Z}^d}$ such that $F(c_1) = F(c_2)$. We use this fact to show that $\mathcal{A}$ has an orphan. Let $D \subset \mathbb{Z}$ be a large enough interval such that c_1 and c_2 agree outside of it, as well as on the first and last $2r$ positions in D. Let $n = |D|$ and let

$$p_1 = c_1\big|_D \in S^n, \qquad p_2 = c_2\big|_D \in S^n.$$

Then $p_1 \neq p_2$, the words p_1 and p_2 agree on the first and last $2r$ positions, and since $F(c_1) = F(c_2)$, the output on the interior of D is the same for both configurations. Therefore,

$$\tilde{f}(p_1) = \tilde{f}(p_2) \in S^{n-2r}.$$

Now let $k \in \mathbb{N}^+$ and define

$$C_k = \{u_1 \cdots u_k \mid u_i \in S^n \setminus \{p_1\} \text{ for all } i\} \subseteq S^{kn}.$$

Clearly,

$$|C_k| = (|S|^n - 1)^k.$$

We claim that every word $w \in S^{kn-2r}$ that has a preimage under $\tilde{f} : S^{kn} \to S^{kn-2r}$ has a preimage in the set C_k. Indeed, let

$$b = u_1 \cdots u_k \in S^{kn}$$

satisfy $\tilde{f}(b) = w$. Whenever some block u_i is equal to p_1, replace it by p_2. Because p_1 and p_2 agree on the first and last $2r$ positions, all output cells whose neighbourhood crosses the boundary of this block remain unchanged. The output cells whose neighbourhood is fully contained inside the block form the word $\tilde{f}(p_1) = \tilde{f}(p_2)$, so they also remain unchanged. Hence this replacement does not change the image under $\tilde{f}$. Repeating the procedure for every occurrence of p_1 yields a preimage of w that belongs to C_k.

Therefore, the number of words in S^{kn-2r} that have a preimage is at most $|C_k| = (|S|^n - 1)^k$. By Lemma 5.26 for the case $d = 1$, there exists $k \in \mathbb{N}^+$ such that

$$(|S|^n - 1)^k < |S|^{kn-2r}.$$

Hence there exists a word in S^{kn-2r} with no preimage under $\tilde{f}$. Such a word is an orphan, and therefore $\mathcal{A}$ is not surjective. □

Lemma 5.26. *For all $d, n, s, r \in \mathbb{N}^+$, there exists $k \in \mathbb{N}^+$ such that*

$$(s^{n^d} - 1)^{k^d} < s^{(kn-2r)^d}.$$

Proof. If $s = 1$, then $(s^{n^d} - 1)^{k^d} = (1-1)^{k^d} = 0$ and $s^{(kn-2r)^d} = 1$ so the inequality holds for every k. Now assume $s \geq 2$ and apply $\log_s$ to both sides. The original inequality holds if and only if

$$k^d \log_s(s^{n^d} - 1) < (kn - 2r)^d.$$

And this is equivalent to

$$\log_s(s^{n^d} - 1) < (n - \frac{2r}{k})^d.$$

Now, for the left-hand side we have that $\log_s(s^{n^d} - 1) < n^d$ while for the right-hand side $(n - \frac{2r}{k})^d \to n^d$ as $k \to \infty$. Therefore, there must exist a $k \in \mathbb{N}^+$ large enough for the inequality to hold. □

Corollary 5.27. *Every injective cellular automaton is also surjective. In particular, we have the following relationship:*

$$\begin{array}{ccccc} \textit{injective} & \Longleftrightarrow & \textit{bijective} & \Longleftrightarrow & \textit{reversible} \\ & & \Downarrow & & \\ & & \textit{surjective} & & \end{array}$$

However, the converse is not true, as illustrated by the following example.

Example 5.28 (Rule 90)**.** *Elementary CA Rule 90 is surjective but not injective. In particular, every configuration has exactly four preimages. This comes from the fact that Rule 90 is both* left-permutive *and* right-permutive*. We say that a 1D CA with local rule f and radius r is left-permutive if $f(*, s_1, \dots, s_{2r}) : S \to S$ is a bijection for every possible combination of inputs $s_1, \dots, s_{2r} \in S$. Similarly, we say it is right-permutive if $f(s_1, \dots, s_{2r}, *) : S \to S$ is a bijection for every possible combination of inputs $s_1, \dots, s_{2r} \in S$. It is easy to verify that since Rule 90 satisfies $f(x, y, z) = x + z \bmod 2$ it is indeed both left- and right-permutive. Let $c \in \mathbf{2}^{\mathbb{Z}}$ be an arbitrary configuration and choose any $x, y \in \mathbf{2}$. We show that there is exactly one configuration $b \in \mathbf{2}^{\mathbb{Z}}$ such that $F(b) = c$ and $b_0 = x, b_1 = y$. The process is straightforward from the illustration in Figure 6.*

$$b = \cdots \begin{array}{|c|c|c|c|c|c|c|} \hline b_{-3} & b_{-2} & b_{-1} & x & y & b_2 & b_3 \\ \hline \end{array} \cdots$$

$$c = \cdots \begin{array}{|c|c|c|c|c|c|c|} \hline c_{-3} & c_{-2} & c_{-1} & c_0 & c_1 & c_2 & c_3 \\ \hline \end{array} \cdots$$

Figure 6: Constructing preimages under Rule 90.

Using left-permutivity, we can progressively uniquely determine the values of b to the "left of xy": given c_0 and x, y we can determine the unique b_{-1} satisfying $f(b_{-1}xy) = c_0$, etc. Analogously, we use the right-permutivity to uniquely prolong this partial preimage to the right. Since we picked $x, y \in \mathbf{2}$ arbitrarily, it is clear that c has exactly four preimages.

The equivalence provided by the Garden-of-Eden Theorem proves useful when showing that a particular CA is not surjective. One simply needs to find two asymptotic configurations that collide under the CA's single iteration. This often proves to be easier than finding a counterexample to the balance condition. We show this on the example of Game of Life.

Example 5.29 (Game of Life is not surjective)**.** *It is very easy to verify that the two asymptotic configurations surrounded by zeros shown in Figure 7 have the same image under the global update rule of Game of Life. We often call such two asymptotic configurations that collide in the next time-step the* twins*. This immediately shows that Game of Life is not surjective.*

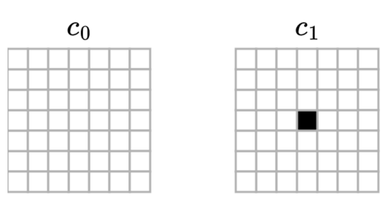

Figure 7: Example of Game of Life twins.

We note that in one-dimension there is an algorithm which decides whether a given CA is surjective, which runs in polynomial time with respect to the size of its local rule description. The algorithm is based on a particular representation of the rule table through a directed graph called the *de Bruijn diagram* that is properly defined in the exercises. However, in dimensions two and higher, we have the following result due to Kari that we state without a proof.

Theorem 5.30 (Kari [1])**.** *In two dimensions, the problem whether a CA is surjective is undecidable.*

This result is a first example out of many that we will see of a CA's global property that is undecidable. It also illustrates the important fact that dimensionality of cellular automata plays an important role.

5.3 Topological Classification of Cellular Automata in One Dimension

We now return to studying the dynamical properties of cellular automata. In particular, we discuss several classical notions from topological dynamics that express how sensitive a CA is to small perturbations of its initial configurations. Our main goal is to present a clean picture for one-dimensional cellular automata: they split into two classes, namely almost equicontinuous and sensitive systems. This gives one possible mathematical formalization of the rough distinction between Wolfram Classes 1+2 and Classes 3+4.

Let us a fix a d-dimensional CA $\mathcal{A} = (S^{\mathbb{Z}^d}, F)$. Below, we present some of the standard notions related to CA's sensitivity.

Equicontinuity The strongest form of stability is that small perturbations of initial configurations never produce a visible difference near the origin, no matter how long we iterate the CA.

Definition 5.31 (Equicontinuity point, Equicontinuous CA)**.** *Let $c \in S^{\mathbb{Z}^d}$. We say that c is an* equicontinuity point *of F if for every $\varepsilon > 0$ there exists $\delta > 0$ such that for every $c' \in S^{\mathbb{Z}}$,*

$$\mathrm{d}(c, c') < \delta \implies \mathrm{d}(F^t(c), F^t(c')) < \varepsilon \quad \textit{for all } t \in \mathbb{N}.$$

We say that $\mathcal{A}$ is equicontinuous *if every configuration is an equicontinuity point.*

Intuitively, if we know the initial configuration with sufficiently high precision, then we can predict the future evolution in any fixed finite window for all times.

Almost equicontinuity A natural weakening of equicontinuity is to require stable behaviour not everywhere, but on a topologically large set of initial configurations. In topology, such sets are often called *residual*: this means they contain a countable intersection of dense open sets.

Definition 5.32 (Almost Equicontinuous CA)**.** *We say that $\mathcal{A}$ is* almost equicontinuous *if the set of equicontinuity points of F is residual.*

Remark 5.33. *For one-dimensional cellular automata, almost equicontinuity admits a very useful simplification: it is equivalent to the existence of a single equicontinuity point. However, this equivalence does not hold in higher dimensions, and there are known counterexamples.*

Exercise 5.34 (Finding equicontinuity points)**.** *Consider the ECA Rule 128 whose local rule $f : \mathbf{2}^3 \to \mathbf{2}$ can be expressed as $f(a, b, c) = abc \bmod 2$. Find an equicontinuity point of this CA. Is the CA equicontinuous?*

Sensitivity The opposite type of behaviour is that for every configuration we can carefully choose a perturbation that will eventually become visible.

Definition 5.35 (Sensitivity)**.** *We say that $\mathcal{A}$ is* sensitive to initial conditions *if there exists $\varepsilon > 0$ such that for every $c \in S^{\mathbb{Z}^d}$ and every $\delta > 0$, there exist $c' \in S^{\mathbb{Z}^d}$ and $t \in \mathbb{N}$ such that*

$$\mathrm{d}(c, c') < \delta \quad \textit{and} \quad \mathrm{d}(F^t(c), F^t(c')) > \varepsilon.$$

Since d only takes values among $0, 1, 2^{-1}, 2^{-2}, \ldots$, one may equivalently assume that $\varepsilon = 2^{-m}$ for some $m \in \mathbb{N}$. Then sensitivity means that there exists a fixed finite window $[-m, m]$ such that however far away we place a perturbation, after some number of iterations the two evolutions will differ somewhere inside this window.

Example 5.36. *The shift map is sensitive: a perturbation placed far away is simply transported toward the origin until it becomes visible there.*

Remark 5.37 (Chaos)**.** *There are several stronger notions of chaotic behaviour in topological dynamics, for instance Devaney chaos. We will not discuss them in these notes, since for our purposes the dichotomy between almost equicontinuity and sensitivity in dimension one already provides a useful and reasonably clean topological picture.*

Below, we present an important result for 1D CAs without a proof; this can be found in [3].

Theorem 5.38 (One-dimensional dichotomy)**.** *Let $\mathcal{A} = (S^{\mathbb{Z}}, F)$ be a one-dimensional CA. Then the following are equivalent:*

(1) $\mathcal{A}$ is not sensitive,

(2) $\mathcal{A}$ has an equicontinuity point,

(3) $\mathcal{A}$ is almost equicontinuous.

In particular, every one-dimensional CA is either almost equicontinuous or sensitive.

This theorem gives a nice topological picture: almost equicontinuous CAs are, in a rough sense, the more regular and predictable ones, while sensitive CAs amplify local perturbations and their dynamics is difficult to predict up to a finite precision. This could provide a formal split between Wolfram Classes 1+2 and Classes 3+4. However, one should be careful not to identify the two classifications completely. For example, the shift map is sensitive, so it is chaotic in the topological sense, but it is not usually regarded as computationally rich in the sense associated with typical Class 4 behaviour. Moreover, some cellular automata classically assumed in Class 2 are indeed sensitive, such as the elementary CA Rule 2.

Remark 5.39 (Higher dimensions)**.** *The previous theorem is special to the one-dimensional setting. In dimension $d \geq 2$, the dichotomy between almost equicontinuity and sensitivity fails, and it is no longer true that the existence of a single equicontinuity point implies almost equicontinuity.*

Summary For cellular automata, equicontinuity means complete stability under small perturbations, almost equicontinuity means that such stable behaviour occurs on a residual set of configurations, and sensitivity means that arbitrarily small perturbations can eventually become visible. The key theorem says that every one-dimensional CA is either almost equicontinuous or sensitive. This gives a useful topological analogue of the rough Wolfram split between Classes 1+2 and 3+4, though the correspondence is not perfect. However, the result does not generalize to higher dimensions.

5.4 Subsystems, Embeddings, and Factors of Cellular Automata

It is often useful to compare two cellular automata and ask whether the dynamics of one is somehow contained in the dynamics of the other. This will become a central question in the next chapter when we start studying the computational capacity of a CA by assessing how many other CAs can it "simulate". The two crucial ingredients for such comparisons are the embeddings and factors of a CA.

Before we define these notions, we first introduce an important class of topological subspaces that are shift-invariant and admit finite descriptions. We will also encounter them in the next section as they provide a good notion of "computationally feasible subsystems".

Subshifts of Finite Type Let $S^{\mathbb{Z}^d}$ be a configuration space. Each finite set L of finite patterns in $S^{\mathbb{Z}^d}$ defines a *subshift of finite type (SFT)* $\Sigma_L \subseteq S^{\mathbb{Z}^d}$ as follows:

$$\Sigma_L = \{c \in S^{\mathbb{Z}^d} \mid c \text{ does not contain any pattern from } L\}.$$

Thus, an SFT is obtained by imposing finitely many local constraints on configurations. Subshifts of finite type provide natural configuration subspaces on which one can study shift dynamics and, more generally, the action of cellular automata.

As the next exercise shows, every subshift of finite type is closed in $S^{\mathbb{Z}^d}$ and invariant under all the shift maps.

Exercise 5.40 (Subshifts of finite type are closed and invariant under shifts)**.** *Let $\Sigma_L \subseteq S^{\mathbb{Z}^d}$ be a subshift of finite type defined by some finite set L of finite patterns in $S^{\mathbb{Z}^d}$. Show that Σ_L is closed and invariant under every shift map.*

Subsystems, Embeddings and Factors Let (X, F) and (Y, G) be topological dynamical systems; i.e., dynamical systems endowed with a topology.

Definition 5.41 (Subsystem)**.** *We say that (Y, G) is a* subsystem *of (X, F) if $Y \subseteq X$ is closed, $F(Y) \subseteq Y$, and $G = F\big|_Y$.*

Thus, a subsystem is obtained by restricting the dynamics to a closed subset that is preserved by the map.

Example 5.42 (SFTs are natural subsystems of the shift map)**.** *Consider the shift CA $\mathcal{A} = (\mathbf{2}^{\mathbb{Z}}, \sigma)$. Then, every SFT induces a subsystem of $\mathcal{A}$, since it is closed and invariant under σ.*

In practice, one often studies such subsystems only up to conjugacy. This leads to the notion of embedding.

Definition 5.43 (Topological Conjugacy)**.** *We say that (X, F) and (Y, G) are* topologically conjugate *if there exists a homeomorphism $\varphi : X \to Y$ such that*

$$\varphi \circ F = G \circ \varphi.$$

Definition 5.44 (Embedding)**.** *We say that (X, F)* embeds into *(Y, G) if there exists a subsystem $(Z, G\big|_Z)$ of (Y, G) which is topologically conjugate to (X, F).*

In the case of cellular automata, there is an equivalent way to define systems that embed into a cellular automaton.

Observation 5.45. *Suppose that (X, F) and (Y, G) are both dynamical systems on compact metric spaces. Then, (X, F) embeds into (Y, G) if and only if there exists a continuous injective map $\iota : X \hookrightarrow Y$ satisfying*

$$\iota \circ F = G \circ \iota.$$

We often call ι the embedding*.*

The dual notion of an embedded system is that of a factor.

Definition 5.46 (Factor)**.** *We say that (X, F) is a* factor *of (Y, G) if there exists a continuous surjective map $\pi : Y \twoheadrightarrow X$ satisfying*

$$\pi \circ G = F \circ \pi.$$

In that case, we call π the factor map*.*

Hence, a factor is a dynamical system obtained from (Y, G) by identifying configurations in a way that is compatible with the dynamics.

Example 5.47 (Trivial Factor)**.** *Let $\mathcal{B} = (T^{\mathbb{Z}^d}, G)$ be an arbitrary CA and $\mathcal{A} = (\mathbf{1}^{\mathbb{Z}^d}, F)$ be a CA with a single state. Then, $\mathcal{A}$ is a factor of $\mathcal{B}$ witnessed by the constant map $\pi : T^{\mathbb{Z}^d} \to \mathbf{1}^{\mathbb{Z}^d}$. It is straightforward to verify that $\pi \circ G = F \circ \pi$ as both sides are simply equal to π.*

Remark 5.48 (Shift-commuting versions)**.** *Since cellular automata live on configuration spaces equipped with shift maps, one often strengthens the notions above by additionally requiring the witnessing maps (i.e., the conjugacy, the embedding, and the factor map) to commute with all shifts. Such shift-commuting versions are often called* strong conjugacies/strong embeddings/strong factor maps*. This stronger condition leads to a useful fact: whenever a dynamical system $(S^{\mathbb{Z}^d}, F)$ is topologically conjugate to, embeds into, or is a factor of a cellular automaton $\mathcal{B} = (T^{\mathbb{Z}^d}, G)$ via a shift-commuting map, then $(S^{\mathbb{Z}^d}, F)$ is itself a cellular automaton. Indeed, one obtains continuity and shift-commutation of F from the intertwining relation and then applies the Curtis–Hedlund–Lyndon theorem.*

This is not true for the weaker notions we have defined in this section: as Example 5.49 shows, a cellular automaton can be topologically conjugate to a system which is not a CA. However, the shift-commuting versions exclude certain natural examples, such as one obtained by regrouping cells into blocks; we illustrate this in Example 5.50. These will be important in the next chapter when we discuss global simulation. Therefore, in this course we will work with the weaker notions where the witnessing maps need not commute with shifts.

Example 5.49 (Conjugacy need not preserve cellular automata)**.** *Consider the one-dimensional cellular automaton given by shift map $\mathcal{A} = (\mathbf{2}^{\mathbb{Z}}, \sigma)$. Define a homeomorphism $h \colon \mathbf{2}^{\mathbb{Z}} \to \mathbf{2}^{\mathbb{Z}}$ which the coordinates 0 and 1:*

$$(h(x))_0 = x_1, \qquad (h(x))_1 = x_0, \qquad (h(x))_n = x_n \quad \text{for all } n \neq 0, 1.$$

Clearly, h is a homeomorphism and $h^{-1} = h$.

Consider the map

$$F = h \circ \sigma \circ h.$$

Then $(\mathbf{2}^{\mathbb{Z}}, F)$ is topologically conjugate to the cellular automaton $\mathcal{A}$, with conjugacy h. We claim, however, that F is not itself a cellular automaton. Indeed, a direct computation shows that

$$(F(x))_n = \begin{cases} x_1, & \text{if } n = -1, \\ x_2, & \text{if } n = 0, \\ x_0, & \text{if } n = 1, \\ x_{n+1}, & \text{otherwise.} \end{cases}$$

Which now allows us to see that F does not commute with the shift:

$$(\sigma \circ F)(x)_0 = (F(x))_1 = x_0,$$

whereas

$$(F \circ \sigma)(x)_0 = (\sigma(x))_2 = x_3,$$

and these are not equal in general. Hence

$$\sigma \circ F \neq F \circ \sigma.$$

Therefore F is not a cellular automaton.

The following example shows two cellular automata which are conjugate via a map that does not commute with the shifts. This example illustrates an important way to rescale the CA grid to obtain a CA with equivalent dynamics but a richer state space.

Example 5.50 (Strong conjugacy is restricting)**.** *Let $\mathcal{A} = (\mathbf{2}^{\mathbb{Z}}, \sigma^2)$ be the CA which shifts every configuration by two cells to the left and consider the state space*

$$\mathbf{2}^2 = \{00, 01, 10, 11\},$$

whose symbols we think of as blocks of length 2 over $\{0,1\}$. Define a map

$$o\colon \mathbf{2}^{\mathbb{Z}} \to (\mathbf{2}^2)^{\mathbb{Z}}$$

by grouping each configuration into consecutive pairs:

$$o(x)_n = x_{2n}x_{2n+1}.$$

In other words, the symbol at coordinate n in $o(x)$ records the two bits of x at coordinates $2n$ and $2n+1$. For example, if

$$x = \dots x_{-2}x_{-1}x_0x_1x_2x_3\dots,$$

then

$$o(x) = \dots(x_{-2}x_{-1})(x_0x_1)(x_2x_3)\dots.$$

The map o is a homeomorphism. And it is straightforward to check it is a conjugacy between $(\mathbf{2}^{\mathbb{Z}}, \sigma^2)$ and the CA $((\mathbf{2}^2)^{\mathbb{Z}}, \sigma)$ which now operates on the richer state space and shifts every configuration by one cell to the left.

However, the map o is not a strong conjugacy, since it does not commute with the shift. Indeed,

$$o \circ \sigma^2 = \sigma \circ o,$$

but in general

$$o \circ \sigma \neq \sigma \circ o.$$

This example shows that strong conjugacy is more restrictive than ordinary topological conjugacy: it excludes natural conjugacies obtained by grouping cells into blocks.

Whenever we have a non-surjective CA, we can use the existence of an orphan to create non-trivial subsystems and factors of such a CA. This is illustrated by the following exercise.

Exercise 5.51. *Let $\mathcal{A} = (S^{\mathbb{Z}^d}, F)$ be a non-surjective CA.*

a) Construct a non-trivial SFT $\Sigma \subsetneq S^{\mathbb{Z}^d}$ which is a subsystem of $\mathcal{A}$.

b) Let $\mathcal{O} = (\mathbf{2}^{\mathbb{Z}^d}, O)$ be the constant CA mapping every configuration to $0^{\mathbb{Z}^d}$. Show that $\mathcal{O}$ is a factor of $\mathcal{A}$.

Chapter Summary

In this chapter, we equipped the configuration space $S^{\mathbb{Z}^d}$ with its natural topology and observed that it is compact. This allowed us to view a cellular automaton as a continuous self-map of a compact metric space. The Curtis–Hedlund–Lyndon theorem then gave a fundamental characterization: the global rules of cellular automata are precisely the continuous maps on $S^{\mathbb{Z}^d}$ that commute with all shifts.

We next studied the global properties of injectivity, surjectivity and reversibility. We introduced Garden-of-Eden configurations and orphans, showed that every Garden-of-Eden configuration contains an orphan, and presented two classical characterizations of surjectivity: balancedness and pre-injectivity. In particular, the Garden-of-Eden theorem implies that every injective cellular automaton is surjective, although the converse fails in general.

We also defined equicontinuity, almost equicontinuity and sensitivity and discussed that every one-dimensional cellular automaton is either almost equicontinuous or sensitive. This dichotomy gives a useful topological perspective on the qualitative behaviour of cellular automata, although it should not be identified too literally with Wolfram's empirical classification.

Finally, we introduced subsystems, embeddings and factors as basic ways to compare dynamics of two systems. These notions will play an important role in the next chapter, where we study the computational capacity of cellular automata.

References

[1] Kari, J. Reversibility and surjectivity problems of cellular automata. *Journal of Computer and System Sciences 48*, 1 (1994), 149–182.

[2] Kari, J. Cellular automata, 2011. Lecture notes.

[3] Kůrka, P. *Topological and Symbolic Dynamics*, vol. 11 of *Cours Spécialisés*. SMF, 2003.