Finite configurations

Let $s \in S$ be an arbitrary state. The **s-support** of a configuration $c \in S^{\mathbb{Z}^d}$ is the set

$$supp_s(c) = \{\vec{n} \in \mathbb{Z}^d \mid c(\vec{n}) \neq s\}$$

of cells not in state s.

Configuration c is **s-finite** if $supp_s(c)$ is a finite set: all but a finite number of cells are in state s.

Let us denote

$$\mathcal{F}_s(d, S) = \{c \in S^{\mathbb{Z}^d} \mid \ c \text{ is } s\text{-finite } \}.$$

Note that $\mathcal{F}_s(d, S)$ is **countably infinite** while $S^{\mathbb{Z}^d}$ is uncountable.

Sometimes one state $q \in S$ is identified as **quiescent**. The quiescent state q must satisfy

$$f(q, q, \ldots, q) = q.$$

(A cell whose neighbors are all quiescent becomes quiescent.)

If a quiescent state q is identified and fixed then the q-support of c is called simply the **support** of c and denoted by $supp(c)$. Also, q-finite configurations are called **finite**, and the set of finite configurations in $S^{\mathbb{Z}^d}$ is denoted by $\mathcal{F}(d, S)$, or simply by $\mathcal{F}$ when d and S are clear from the context.

The configuration in which every cell is in state q is called the **quiescent configuration**.

If c is s-finite then $G(c)$ is t-finite where $t = f(s, s \ldots, s)$.

In particular, in the presence of **quiescent state** q, finite configurations are mapped into finite configurations. In this case we denote by

$$G_F : \mathcal{F} \longrightarrow \mathcal{F}$$

the **restriction of G on finite configurations**.

Examples. In **XOR** CA we can name state 0 quiescent, in which case the space-time diagram

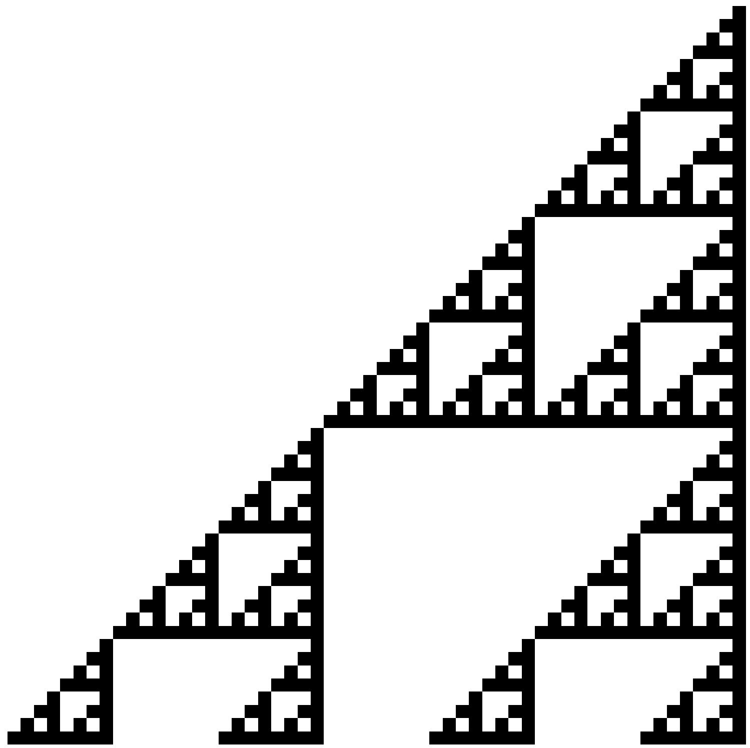

depicts a time-evolution according to G_F.

In **Game-of-life** the **dead** state is taken as the quiescent state.

Spatially periodic configurations

Let the dimension d and the state set S be fixed.

The **translation** $\tau_{\vec{r}}$ by vector $\vec{r} \in \mathbb{Z}^d$ is the function

$$\tau_{\vec{r}} : S^{\mathbb{Z}^d} \longrightarrow S^{\mathbb{Z}^d}$$

that maps $c \mapsto c'$ where $c'(\vec{n}) = c(\vec{n} + \vec{r})$ for all $\vec{n} \in \mathbb{Z}^d$.

(It is the global transition function of the CA whose neighborhood contains only $\vec{r}$ and whose local rule is the identity function.)

Clearly for all $\vec{r}, \vec{s} \in \mathbb{Z}^d$ and $k \in \mathbb{Z}$ we have

$$\tau_{\vec{r}} \circ \tau_{\vec{s}} = \tau_{\vec{r}+\vec{s}},$$

$$\tau_{\vec{r}}^{-1} = \tau_{-\vec{r}}, \text{ and}$$

$$\tau_{\vec{r}}^{k} = \tau_{k\vec{r}}.$$

For each dimension $i = 1, 2, \ldots, d$ we call the translation by one cell down in dimension i a **shift** and denote it by σ_i. In other words: $\sigma_i = \tau_{\vec{e}_i}$ for the i'th coordinate unit vector

$$\vec{e}_i = (0, \ldots, 0, 1, 0, \ldots 0).$$

Every translation is a composition of shifts and their inverses.

In the one-dimensional case the only shift σ_1 is called the **left shift** and we denote it simply by σ.

The following proposition states an elementary but important property of cellular automata, based on the fact that all cells use the same local update rule:

Proposition. Let G be an arbitrary CA function and τ a translation. Functions G and τ commute, i.e., $G \circ \tau = \tau \circ G$:

$$\begin{array}{ccc} S^{\mathbb{Z}^d} & \xrightarrow{\quad G \quad} & S^{\mathbb{Z}^d} \\ \tau \big\downarrow & & \big\downarrow \tau \\ S^{\mathbb{Z}^d} & \xrightarrow{\quad G \quad} & S^{\mathbb{Z}^d} \end{array}$$

Proof.

A configuration $c \in S^{\mathbb{Z}^d}$ is **$\vec{r}$-periodic** if

$$c = \tau_{\vec{r}}(c),$$

i.e., c is **invariant under the translation** by $\vec{r}$.

In other words,

$$c(\vec{n}) = c(\vec{n} + \vec{r}) \text{ for all } \vec{n} \in \mathbb{Z}^d.$$

A configuration is **spatially periodic** if it is $\vec{r}$-periodic for some $\vec{r} \neq \vec{0}$.

A d-dimensional configuration is **strongly periodic** if it is $\vec{r}_i$-periodic for some linearly independent $\vec{r}_1, \vec{r}_2, \ldots, \vec{r}_d \in \mathbb{Z}^d$.

A strongly periodic configuration consists of a $k \times k \times \cdots \times k$ pattern, for some $k \geq 1$, that is repeated periodically in each of the d-dimensions of the space:

Proposition. If c is strongly periodic then there exists $k \in \mathbb{Z}_+$ such that c is σ_i^k-periodic for all $1 \leq i \leq d$.

Proof.

Let $\mathcal{P}(d, S)$ denote the set of strongly periodic elements of $S^{\mathbb{Z}^d}$, or if d and S are clear from the context, we may simply use $\mathcal{P}$.

Set $\mathcal{P}(d, S)$ is **countably infinite**.

In the one-dimensional case there is no difference between spatial periodicity and strong periodicity. In two- and higher dimensional spaces there is a difference.

Example. A two-dimensional configuration (infinite horizontal stripe) that is $\vec{e}_1$-periodic but not strongly periodic:

Example. A strongly periodic configuration (infinite checker board):

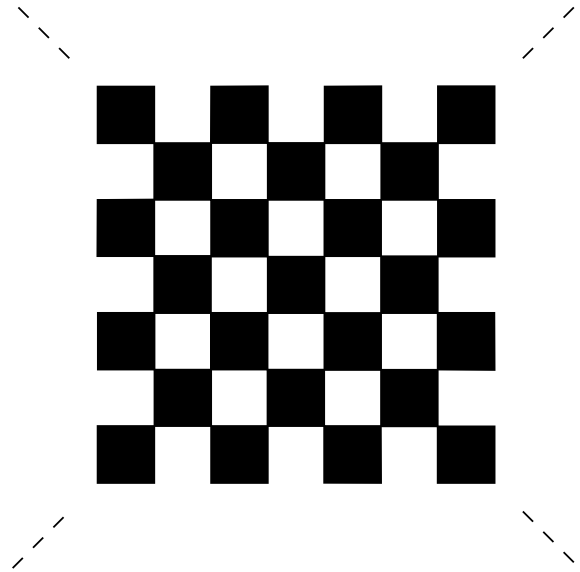

Cellular automata preserve spatial periods of configurations:

Proposition. If G is a CA and c is an $\vec{r}$-periodic configuration then $G(c)$ is also $\vec{r}$-periodic.

Proof.

Cellular automata preserve spatial periods of configurations:

Proposition. If G is a CA and c is an $\vec{r}$-periodic configuration then $G(c)$ is also $\vec{r}$-periodic.

In particular, if c is strongly periodic then also $G(c)$ is strongly periodic. We denote by

$$G_P : \mathcal{P} \longrightarrow \mathcal{P}$$

the restriction of G on strongly periodic configurations.

Finite configurations and periodic configurations are used in effective **simulations** of cellular automata on computers.

Periodic configurations are often referred to as the **periodic boundary conditions** on a finite cellular array.

Remark: Periodicity of a configuration may refer to temporal or spatial periodicity. I try to be careful and make clear which of the two I am talking about.

Compactness

Let $c_1, c_2, \ldots$ be a **sequence** of configurations in $S^{\mathbb{Z}^d}$. The sequence **converges** to a **limit** configuration c if

$$(\forall \vec{n} \in \mathbb{Z}^d)\ (\exists k \in \mathbb{N})\ (\forall i \geq k)\ :\ c_i(\vec{n}) = c(\vec{n}).$$

Such a limit (if it exists) is unique and we denote

$$c = \lim_{i \to \infty} c_i.$$

In other words: if we look at any cell $\vec{n}$ and scan $c_1, c_2, \ldots$ then from some moment on we always see the same state $c(\vec{n})$ in position $\vec{n}$.

Later we give the set $S^{\mathbb{Z}^d}$ of configurations a **metric**. The convergence of sequences under this metric is exactly this convergence concept.

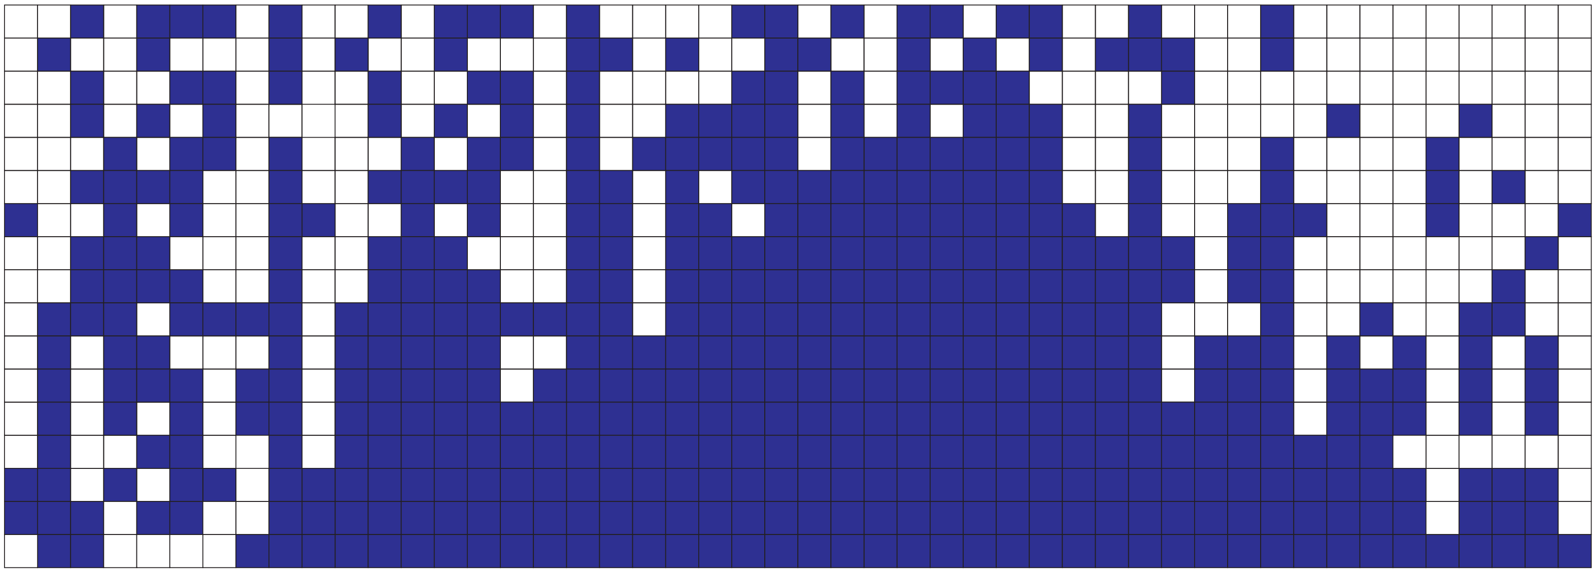

A **subsequence** of $c_1, c_2, \ldots$ is a sequence

$$c_{i_1}, c_{i_2}, \ldots$$

where $i_1 < i_2 < \ldots$.

(So a subsequence is obtained by picking infinitely many elements of the sequence, preserving their relative order.)

Obviously every subsequence of a converging sequence also converges and has the same limit.

The following theorem states the compactness of the configuration space:

Proposition. Every sequence of configurations has a converging subsequence.

The following theorem states the compactness of the configuration space:

Proposition. Every sequence of configurations has a converging subsequence.

Proof. In the proof we use the following simple special case of **König's infinity lemma**. Consider an infinite directed rooted tree where each node has a finite number of children:

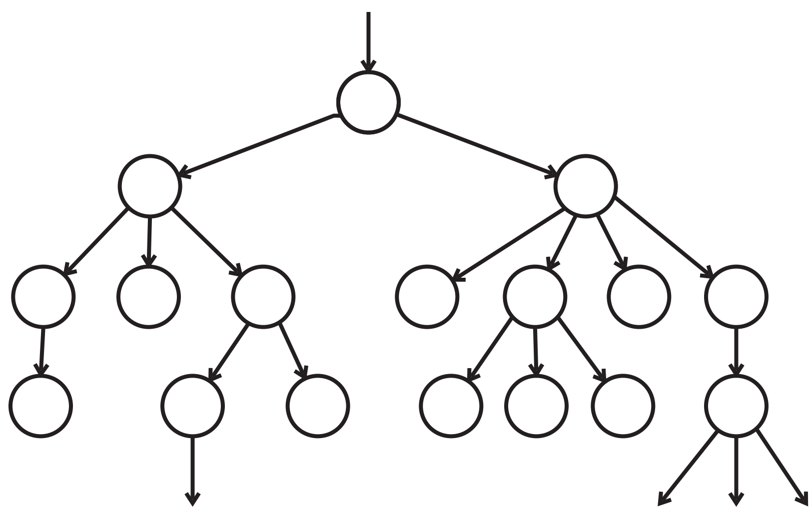

Lemma: The tree contains an infinite path down from the root.

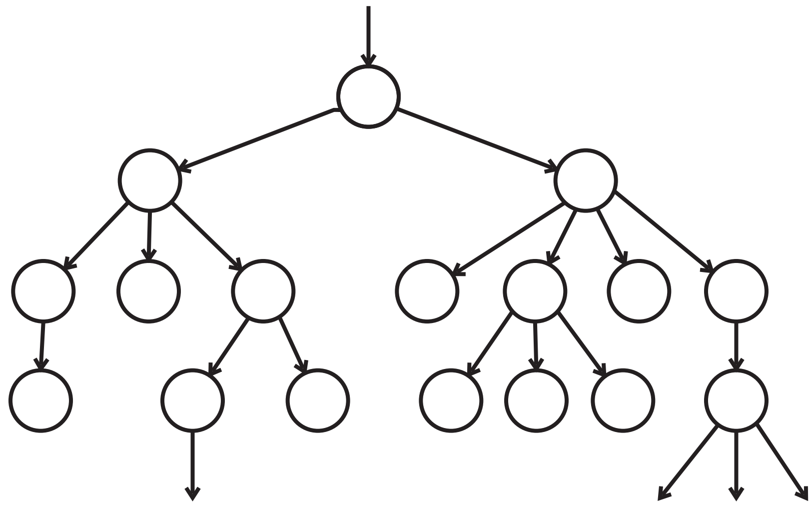

Lemma: The tree contains an infinite path down from the root.

Proof of the lemma: If a node is the root of an infinite subtree, then it has a child that is also the root of an infinite subtree (because the node has just finite number of children).

So starting from the root one can move down the tree by always moving to a child that is the root of an infinite subtree. This path is never blocked so the path follows an infinite branch of the tree. □

Now we can prove our compactness proposition:

Proposition. Every sequence of configurations has a converging subsequence.

Proof. Let $c_1, c_2, \ldots$ be an arbitrary sequence in $S^{\mathbb{Z}^d}$.

Fix one $t \in S$ such that $c_i(\vec{0}) = t$ for infinitely many i. (Such t exists since S is finite.) Remove from the sequence all c_i such that $c_i(\vec{0}) \neq t$. Thus we may assume all c_i satisfy $c_i(\vec{0}) = t$.

For $n = 0, 1, \ldots,$ let

$$D_n = \{-n, \ldots, n\}^d,$$

that is, the $(2n+1) \times (2n+1) \times \cdots \times (2n+1)$ size hypercube centered at the origin $\vec{0}$.

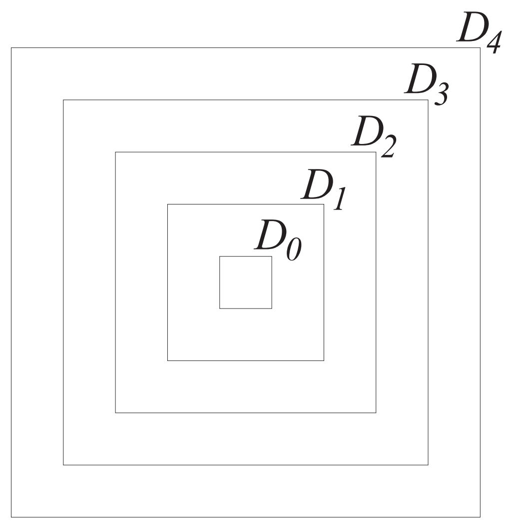

Define the following infinite tree:

- Nodes at level n are patterns $p : D_n \longrightarrow S$ such that for infinitely many i we have
$$c_i|_{D_n} = p.$$
(Patterns that appear centered around origin in infinitely many elements of the sequence.)

- The parent of a pattern $x : D_n \longrightarrow A$ is the pattern $x|_{D_{n-1}}$. Thus children extend the parent pattern to a larger domain:

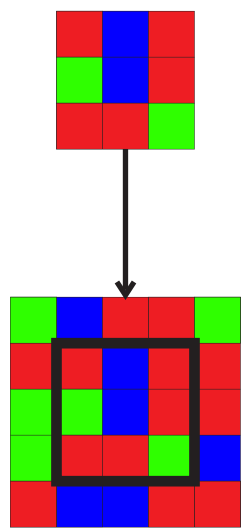

- The root is the unique node of level 0. (It is the pattern with state t at cell $\vec{0}$.)

The tree is infinite so it has an **infinite branch** b. The infinite branch defines a configuration c where the state in any cell $\vec{n}$ is the unique state put in that position by patterns of the branch b.

There is a subsequence that converges to c: we pick indices $i_1, i_2, \ldots$ such that for every n

- $i_n < i_{n+1}$, and
- $c_{i_n}|_{D_n}$ is the level n node in branch b.

We then have that for all $k \geq n$

$$c_{i_k}|_{D_n} = c_{i_n}|_{D_n} = c|_{D_n},$$

hence the subsequence converges to c. □

Our next proposition essentially states that all CA functions are **continuous**.

Proposition. Let G be a CA function and $c_1, c_2, \ldots$ a converging sequence of configurations with limit c. Then also the sequence $G(c_1), G(c_2), \ldots$ converges and

$$\lim_{i\to\infty} G(c_i) = G(c).$$

Proof.

Our last proposition states that the sets of finite and strongly periodic configurations are **dense**:

Proposition. Let $c \in S^{\mathbb{Z}^d}$ and $s \in S$. There exist sequences

(a) $c_1, c_2, \ldots$ of s-finite configurations $c_i \in \mathcal{F}_s(d, S)$, and

(b) $p_1, p_2, \ldots$ of strongly periodic configurations $p_i \in \mathcal{P}(d, S)$

such that $c = \lim_{i \to \infty} c_i = \lim_{i \to \infty} p_i$.

Proof.