



Universiteit
Leiden

Master Computer Science

A coalgebraic modal logic for cellular automata

Name: Lulof Pirée
Student ID: s3406512
Date: 15/07/2024
Specialisation: Artificial Intelligence
1st supervisor: Dr. Henning Basold
2nd supervisor: Dr. Chase Ford

Master's Thesis in Computer Science

Leiden Institute of Advanced Computer Science (LIACS)
Leiden University
Einsteinweg 55
2333 CC Leiden
The Netherlands

Abstract

Dynamic patterns of states in the evolution of a cellular automaton can be described in the simple syntax of our new modal logic. Not only can the logic express local properties such as the local state-transition rule, but also global properties such as periodicity and nilpotency. We present the semantics of the logic via a novel coalgebraic model of cellular automata that explicitly separates cells from connections between them. This model is sufficiently general to describe any graph shape, and generalises to non-uniform cellular automata in which different cells may admit different local rules. Finally, we define a coalgebraic definition of similarity that includes a map between the parameters of the type functor. A generalised Hennessey-Milner theorem proves that two cellular automata are similar if and only if every modal logic formula valid on the one can be translated into a formula valid on the other.

Contents

1	Introduction	3
1.1	Related work	5
2	Preliminaries	6
2.1	Sets and functions	7
2.2	Cellular automata and monoids	8
2.3	Category theory	11
2.4	Modal logic	13
3	Uniform cellular automata	15
3.1	Graph as coalgebras	15
3.2	Extending graphs with local rules	19
3.3	Uniform cellular automata as coalgebras	19
3.4	Modelling irregular and non-uniform CA	25
3.5	Summary of uniform CA	27
4	General cellular automata	27
4.1	Classes of general CA	27
4.2	General CA functor	34
4.3	Neighbourhood- and extension-uniform CA	36
4.3.1	The local rule as a natural transformation	37
4.4	Summary of general CAs	38
5	Coalgebra homomorphisms and behavioural equivalence	38
5.1	Characterising coalgebra homomorphisms	39
5.1.1	First property: commutativity with cowriter-comonad coalgebras	39
5.1.2	Second property: extension of local rules	41
5.2	Illustration of behavioural equivalence	42
5.3	The terminal coalgebra	43
5.3.1	Coalgebra contractions	44
5.3.2	Construction of the terminal coalgebra	51
5.3.3	Morphisms between embeddings	52
5.3.4	Cardinality of union of contractions	56
5.4	Summary of coalgebra homomorphisms and behavioural equivalence	58
6	Logic for computation traces	59
6.1	Space modalities	59
6.1.1	Grammar	60
6.1.2	Semantics	60

6.1.3	Expressivity of space modalities	61
6.2	Time modalities	62
6.2.1	Expressivity of time modalities	63
6.3	Global variables and uniform substitution	64
6.4	Logical reasoning rules	67
6.5	Summary of basic modal logic	70
7	Logic extension: infinitary disjunctions	71
7.1	Intuition and motivation	71
7.2	Formal definitions of infinitary disjunctions	72
7.3	Examples of expressible properties	74
7.4	Summary of infinitary disjunctions	79
8	Correspondence semantical and logical bisimulation	79
8.1	Examples of similarity	80
8.1.1	A basic example of similarity	80
8.1.2	Motivation extending similarity with monoid homomorphisms	80
8.1.3	Motivation extending similarity with state mappings	82
8.2	Basic Hennessy-Milner theorem	82
8.3	Generalised Hennessy-Milner theorem	88
8.3.1	General CA morphisms	88
8.3.2	Generalised notions of similarity	89
8.3.3	Generalised Hennessy-Milner theorem	91
8.3.4	Example applications	95
8.4	Summary of the Hennessy-Milner theorem	97
9	Concluding remarks	98
9.1	Discussion of the coalgebraic model	98
9.2	Discussion of the new modal logic	98
9.3	Conclusion	99
	References	100

1 Introduction

Cells in a human body, particles in a gas cloud, ants in an ant-hill, components in electronic circuits; the world is full of systems consisting of a collection of simple objects (cells, particles, ants, components) that interact only locally with similar objects in their surrounding. Cellular automata (CA) are computational models for studying such systems mathematically and/or simulating them on a computer. In practise, CA are employed in the empirical sciences to for example biological [10] and physical systems [32][33], but also as tools in areas such as game design [47][17][44], music generation [3] and in and cryptography [28][49][27][46]. In the theory of computation, CA are a very general computational model: many other models such

such as Turing Machines [42][24], have been described as specific CA.

A CA consists of a graph where every vertex (usually called a *cell*) is inhabited by a deterministic automaton. These cells hold a single symbol, their *state*, and in discrete time steps they synchronously take their neighbour's states as inputs and update their own state according to a deterministic *local rule*. This defines a global transition function between assignments of states to the cells (*configurations*). From a given initial configuration, one can recursively compute the *trace sequence* of successor configurations using this global function. Despite the simplicity of individual cells, consisting of no more than a single state and a deterministic local rule, the global trace can exhibit complex behaviour.

Local spatial and/or temporal patterns in the trace of CA are often of great importance in their analysis and design. Such patterns are, for example, data structures or signals between isolated groups of actively interacting cells. Currently, there is a deficit in the supply of universal tools to describe, present, compare and reason about such patterns. Visualisations are commonly employed, but are informal and describe only limited portions of space and time. Furthermore, some graphs (such as four-dimensional grids) can be visualised. The aim of this study is to provide a symbolic logic capable of describing spatio-temporal patterns with a uniform syntax for all graphs.

Modal logics are logical languages that have successfully been employed to describe and reason about local properties of various computational models, such as labelled transition systems [12]. Essentially, modal logic is propositional (Boolean) logic that uses a coloured graph of *worlds*, which are assignments of Boolean values to variables, instead of a single such assignment. Formulas can reason about truths in a given world, but also evaluate subformulas in worlds related via a coloured edge. It is this graph-based nature that makes modal logic an inviting choice for a logic of CA: spatially, CA have cells arranged in a graph, and temporally, CA represent their computation via their trace sequences. Hence the research question of the present work:

Research question: *Is it possible to define a modal logic on CA, capable of describing spatial and temporal relations of cells' states in a trace sequence?*

Many dynamical systems have been modelled in the form of *coalgebras* [39], a concept from category theory that has been proposed to be a universal framework for modal logic [6][23]. In particular, here exists a standard framework of *predicate-lifting* developed by Pattinson [31][30], Schröder [41] and Klin [22] for automatically generating a modal logic with an expressivity guarantee for coalgebras.

Hence a natural starting point for the present study was to model CA coalgebraically as to employ this framework. This line of thought led to a natural generalisation of CA to *non-uniform* cellular automata, in which distinct cells are allowed to have different local rules. However, we show that the predicate-lifting framework would generate a modal logic that is only capable of describing static properties of the graph structure of a CA, and not of properties of computation traces.

Fortunately, the present coalgebraic model provided a suitable foundation for the semantics of a tailor made modal logic that is capable of describing patterns in traces. The logic is especially expressive when extending it with infinitary disjunctions; in some cases can even describe certain global asymptotic properties of traces.

Finally, we define a notion of *mutual similarity* on cells in two CA, based on the ability to simulate each other. In a Hennessy-Milner style theorem we establish that CA are mutually similar if and only if they have a similar logic. In fact, we prove the stronger statement that a cell x can be simulated by another cell if and only if x 's logic can be embedded in y 's logic, using a translation of incompatible logical symbols.

Summary of contributions The contributions of this thesis can be summarised as follows:

1. A new coalgebraic model of uniform and non-uniform cellular automata based on a **Set**-valued comonad, which explicitly separates cells from connections between cells. We propose different classes of non-uniform CA that arise as full subcategories of our coalgebras.

2. When using the logic generated by the predicate lifting framework, two cells would have the same logic if they are *behaviourally equivalent*. We show how CA can be *contracted* to the minimal CA to which it is behaviourally equivalent, and how all contractions can be embedded in a *terminal coalgebra*. The outcome of this study is that many CA are BE, even when their trace sequences on the same initial configuration are different. Hence the predicate lifting’s logic would be unable to describe the details of the trace sequences, which our design goal requires.
3. A tailor made modal logic for describing spatio-temporal patterns in traces of non-uniform CA: an extension of propositional logic with unary modalities for connections between cells and for timesteps in a trace sequence. We show how it can be extended with infinitary disjunctions, and give examples properties of CA that the logic can express.
4. A Hennessy-Milner style theorem establishing the equivalence of semantical simulation and logical simulation.

Thesis overview In Section 2 we gather the required background concepts and we set notational conventions used in this thesis. The main ideas behind our coalgebraic model are introduced in Section 3, which also proves a correspondence between the classical and coalgebraic definitions of CA. Thereafter, Section 4 refines those ideas, and simultaneously generalises them to non-uniform CA; in particular it shows how definitions of two subclasses of non-uniform CA can naturally be formulated in terms of coalgebras. Section 5 analyses morphisms between our coalgebras, and shows that the standard framework for coalgebraic modal logic does not provide the desired logic. We explain the basic version of our manually designed logic in Section 6, and extend it with infinitary disjunctions in Section 7. The latter section also showcases examples of what the logic can express. In Section 8 we define simulations and logical simulations between cells, and prove their relation in a Hennessy-Milner style theorem. The concluding remarks in Section 9.3 finish the thesis.

1.1 Related work

Coalgebraic modal logic The semantics of various modal logics can be defined in terms of coalgebras, and general results for coalgebraic model logics have been published; see the surveys by Cîrstea et al. [6] and by Kupke and Pattinson [23] for an overview. Both these surveys argue that coalgebras provide the ideal abstract framework to reason about modal logics. We have attempted to use the *predicate-lifting* framework which authors such as Pattinson [31][30], Schröder [41] and Klin [22] develop for automatically creating an expressive modal logic for endofunctors on **Set**, **Set** and **Set**-like categories respectively, but the output did not match our design goals. Our manually designed modal logic instead uses a coalgebraic model of CA *indirectly* to define the semantics, which is incompatible with the standard framework. Most of the existing meta-theory therefore not apply to our logic.

Bisimulations We prove a CA analogue of the Hennessy-Milner (HM) theorem [12]. The original Hennessy-Milner theorem states that two labelled transition systems (LTS) satisfy the same model logic formulas if and only if they are *bisimilar*. This latter notion of bisimilarity is a special case of many coalgebraic definitions of bisimilarity (see Staton [43] for an overview). Our HM theorem does not use the standard definition of bisimilarity, but a novel CA definition of *simulation*. Our definition extends a one-way embedding correspondence (a *simulation* in [40]) with requiring the preservation of a cell’s behaviour dynamics. It is unclear if our definition is a special case of the simulations defined by Hughes and Jacobs in [13]. Unlike the original HM theorem, we first prove that one-directional simulation is equivalent to embedding of the logic, and the two-way case follows by applying this in both directions. We also allow simulations and logical equivalence to include a injection on the underlying signatures (of the coalgebras and logical formulas, respectively).

Path logic of concurrent systems Joyal, Nielsen and Winskel [18] also define a uniform modal logic for paths (analogous to CA traces) for various concurrent systems, such as Labelled Transition Systems and Event Structures, and prove a HM theorem for a generalised definition of simulation. Our work is orthogonal to this, since there seems to be no direct correspondence between CA and the considered concurrent systems.

Temporal Description Logic for CA Delivorias et al. [7] construct a decidable Temporal Description Logic for describing the evolution of states of a CA on \mathbb{Z}^2 , with semantics given via knowledge bases. Like ours, their language has atomic formulas for states, includes time modalities and an “until” operator, and they give similar examples of expressible properties. The main differences are: (1) we use coalgebraic semantics that apply to any graph (not only to \mathbb{Z}^2) and to non-uniform CA, (2) our logic is undecidable, and (3) our language has additional space modalities.

Simulations between CA A CA \mathcal{B} *simulates* a CA \mathcal{A} if the configurations of \mathcal{A} can be injectively embedded into the configurations of \mathcal{B} , in such a way that the dynamics are preserved. A CA \mathcal{B} is *intrinsically universal* if it can simulate every CA in a given class. Various authors use different choices of the precise definition of *simulation*. These choices vary in the constraints on the embedding and how they define the preservation of the dynamics. We only give a partial account, and focus on the most general definitions. A historic overview of the study of simulation and universality is given by Ollinger in [29], and §5.2 therein gives a very a general definition of *direct simulation* (given for CA on grids of the form \mathbb{Z}^d for some $d \in \mathbb{N}$, but it applies to all CA). Simulations between CA on Cayley graphs have been studied by Roká [36]. For a discussion on different definitions of intrinsic universality, see [9]. Also non-injective simulations have been studied, for example in [14].

The definition of *simulation* used in this thesis (Definition 8.15) is in some sense more general than the definitions cited above, as applies to many classes of CA (all non-uniform CA, which includes uniform CA on \mathbb{Z}^d and all CA on Cayley graphs), but it is more restricted in what kinds of simulations it permits. In particular, our notation of simulation does imply Ollinger’s definition of direct simulation, but not vice versa.

Categorical models of CA The literature provides a few other categorical models of CA, all different from the current model. Capobianco and Uustalu have modelled CA as morphisms between coalgebras of a comonad on the category of uniform spaces [4], as well as morphisms between coalgebras of a *graded* comonad on **Set** [5]. We have attempted to adapt this latter construction to the standard coalgebraic modal logic framework, which unfortunately only yielded a trivial logic. In the present work, CA are modelled as coalgebra objects themselves and not as morphisms between coalgebras.

Widemann and Hauhs [48] have published a bialgebraic model specific to two-dimensional CA. Their focus lies on the theoretical relation between different programming styles for implementing CA.

Non-uniform CA In hindsight we discovered that non-uniform CA have been well-studied: see [1] for an overview.

2 Preliminaries

This section fixes the notations used in this thesis and recites necessary background concepts. Many topics are too large to cover fully, but we try to provide references for further reading where possible. We assume that the reader is familiar with basic category theory. For a general introduction to category theory, see for example one of the books [26], [35] or [25].

2.1 Sets and functions

We fix the following notational conventions for set-theoretical concepts.

For a function $f: A \rightarrow B$ and a subset $C \subseteq A$, the function $f|_C: C \rightarrow B$ is the domain restriction of f , $f[C] \stackrel{\text{def}}{=} \{f(c) \mid c \in C\}$ is the f -image of C . The preimage of f on C is partitioning of C in sets of elements that f maps to the same point in B :

$$\text{Prelm}(f) \stackrel{\text{def}}{=} \{\{c \in C \mid f(c) = b\} \mid b \in f[C]\}.$$

For functions $f: A \setminus B \rightarrow C$ and $g: B \rightarrow C$, we write $f \cup g: A \rightarrow C$ to denote the function

$$x \mapsto \begin{cases} f(x) & x \in A \setminus B \\ g(x) & x \in B. \end{cases}$$

We will sometimes use λ -notation to define a function. As an example of λ -notation, $\lambda n . n + 1$ denotes the map $\mathbb{N} \rightarrow \mathbb{N}$ defined as $n \mapsto n + 1$. Sometimes we explicitly annotate the domain of variables, e.g., as in $\lambda n \in \mathbb{N} . n + 1$. Functions $A \rightarrow B$ that always output the same element $b \in B$ (i.e., those of the form $\lambda a . b$) will also be denoted as const_b .

The set $\mathbb{1}$ denotes the set with exactly one element. The powerset of a set A is denoted as $\wp(A)$. Set difference is denoted with \setminus , defined as:

$$A \setminus B \stackrel{\text{def}}{=} \{a \in A \mid a \notin B\}.$$

Partition orderings Similar to how a complete pizza can be cut into slices or a continent be divided into countries, a set can *partitioned* into mutually disjoint subsets that together form the whole set:

Definition 2.1: Partitioning

A *partitioning* of a set $A \in \mathbf{Set}$ is a set $B \subseteq \wp(A)$ such that every $a \in A$ is included in *exactly one* set in B (so $\bigcup B = A$), and $\emptyset \notin B$. Elements of B are called *partitions*. For every $a \in A$, we denote the partition in B containing a by $|a|_B$.

The *set of all partitionings* of A is denoted by

$$\text{Part}(A) \stackrel{\text{def}}{=} \{B \in \wp(B) \mid \bigcup B = A, \emptyset \notin B, |n|_B, |m|_B \in B \Rightarrow (|n|_B = |m|_B) \vee (|n|_B \cap |m|_B = \emptyset)\}.$$


We define the following partial ordering on $\text{Part}(A)$: for partitionings B and C , we have $B \leq C$ if and only if C is further “cut up” partitioning of A than B is: every partition in C can be obtained as a subset of a partition in B . This is equivalent to saying that $a, a' \in |a|_C$ only occurs if $a, a' \in |a|_B$. Formally:

Definition 2.2: Partitionings ordering

The *partitionings ordering* of a set A , denoted as $\text{PartOrd} = (\text{Part}(A), \leq)$, is the poset where \leq is defined for $B, C \in \text{Part}(A)$ as:

$$B \leq C \Leftrightarrow \bigvee_{a \in A} |a|_C \subseteq |a|_B.$$

For example, the Hasse-diagram of $\text{PartOrd}(\{a, b, c\})$ is:



2.2 Cellular automata and monoids

A cellular automaton (CA) is a computational model whose static structure consists of a directed graph with labelled directed edges whose vertices, called *cells*.

In the dynamics, every cell is assigned a state (often visualised as a colour) from a chosen set of states. Such an assignment of states to cells is called a *configuration*. Every CA comes equipped with (1) a subset of edge labels called a *neighbourhood*, such that every cell has an outgoing edge for every label in the neighbourhood, and (2) a *local rule* that tells how a cell updates its state depending on the states it can see in its local neighbourhood (so the local rule maps an assignment of states to neighbourhood labels to a single state). The local rules define a *global rule* that maps configurations to configurations as follows: the cells synchronously read the states of their neighbours according to the input configuration, feed this to the local rule, and use the output as their state in the output configuration. By iterating the global rule on a given *initial configuration*, one obtains a *trace* of successor configurations. This can be seen as a computation with the initial configuration as input. For general introductions to CA, see for example the books [8] and [37] or the survey [20].

A classical example is Conway’s Game of Life (CGOL) [11][9]. This CA uses the regular 2D grid \mathbb{Z}^2 as a graph, and as neighbourhood the 3×3 region around a cell (including the central cell itself). It uses only two states: “dead” (white) and “alive” (black). The local rule has an intuitive description:

If you are dead, you will become alive the next timestep if you have 3 or 4 alive neighbours. If you are alive, you will stay alive the next timestep if you have 2 or 3 alive neighbours. In all other cases you will be dead the next timestep.




Figure 1: Example of one of CGOL’s rules: a dead cell (the central cell, with relative coordinates $(0, 0)$) with three alive neighbours will become alive next timestep.




Figure 2: First three configurations in a trace on CGOL. Live cells are visualised as black, dead cells as white. Only parts of the configurations are shown, the other cells are assumed to be white.

To give a precise description of the classical definition of a CA, we fix the convention of representing grids by monoids.

Definition 2.3: Monoid

A monoid $(\mathcal{M}, \bullet, \mathbf{1})$ is a set of elements \mathcal{M} together with an associative binary *multiplication* operation $\bullet: \mathcal{M} \times \mathcal{M} \rightarrow \mathcal{M}$ and a distinguished *identity element* $\mathbf{1} \in \mathcal{M}$. The multiplication operation satisfies, for all $m, n, \ell \in \mathcal{M}$, the *monoid associativity law*:

$$m \bullet (n \bullet \ell) = (m \bullet n) \bullet \ell$$

and the *monoid unit laws*:

$$\mathbf{1} \bullet m = m = m \bullet \mathbf{1}.$$

Remark 2.4. We typeset monoids in a thin calligraphic font, e.g., \mathcal{M} . Often we use the same symbol for both the entire monoid as well as the underlying set of elements. Expressions such as $X^{\mathcal{M}}$, $m \in \mathcal{M}$ and $N \subseteq \mathcal{M}$ *always* refer to the underlying set of elements.

Example 2.5. The natural numbers with addition $(\mathbb{N}, +, 0)$ form a monoid.

Definition 2.6: Freely generated monoid

For any set X , let X^* be the set of finite sequences of elements from X (called *strings*):

$$X^* \stackrel{\text{def}}{=} \{[x_1, x_2, \dots, x_k] \mid k \in \mathbb{N}, x_1, \dots, x_k \in X\}.$$

Let $\varepsilon = []$ be the empty string, and let $\bullet_X: X^* \times X^* \rightarrow X^*$ be string concatenation:

$$[x_1, x_2, \dots, x_k] \bullet_X [y_1, y_2, \dots, y_\ell] \stackrel{\text{def}}{=} [x_1, x_2, \dots, x_k, y_1, y_2, \dots, y_\ell]$$

(for all $k, \ell \in \mathbb{N}$ and $x_1, \dots, x_k, y_1, \dots, y_\ell \in X$). Then $(X^*, \bullet_X, \varepsilon)$ is a monoid called the *monoid freely generated over X* .

Example 2.7. If we encode $0 = \varepsilon$, $1 = S$, $2 = SS$, $3 = SSS$, etc., and denote string concatenation with $+$, then \mathbb{N} is freely generated over $\{S\}$.

Definition 2.8: Generated monoid

A *monoid congruence* \sim on a monoid $(\mathcal{M}, \bullet, \mathbf{1})$ is an equivalence relation $\sim \subseteq \mathcal{M} \times \mathcal{M}$ such that for all $m, n, n' \in \mathcal{M}$, if $n \sim n'$ then also $m \bullet n \sim m \bullet n'$ and $n \bullet m \sim n' \bullet m$.

A *monoid generated* over a set X and a relation $R \subseteq X^* \times X^*$ is the monoid $(X^*/\sim_R, \bullet_R, |\varepsilon|_{\sim_R})$ where

- $(X^*, \bullet, \varepsilon)$ is the monoid freely generated over X .
- $\sim_R \subseteq X^* \times X^*$ is the monoid congruence

$$\sim_R \stackrel{\text{def}}{=} \cap \{ \sim \mid R \subseteq \sim, \sim \text{ is a monoid congruence on } (X^*, \bullet, \varepsilon) \}.$$

- $|n|_{\sim_R}$ for $n \in X^*$ denotes the equivalence class of n under \sim_R .
- The multiplication \bullet_R is:

$$|n|_{\sim_R} \bullet_R |m|_{\sim_R} \stackrel{\text{def}}{=} |n \bullet m|_{\sim_R}.$$

This is well-defined since \sim_R is a monoid congruence.

Remark 2.9. Relations are often written as equalities. E.g., if $R = \{(AB, C)\}$ then we write $R = \{AB = C\}$.

Example 2.10. Encode $0 = \varepsilon$, $1 = S$, $2 = SS$, $3 = SSS$, etc., and $-1 = P$, $-2 = PP$, $-3 = PPP$, etc. Then the integers \mathbb{Z} are generated over $\{S, P\}$ and the relation $R = \{SP = \varepsilon\}$. Addition is string concatenation modulo \sim_R .

A homomorphism between monoids is a map of elements that preserves monoid multiplication.

Definition 2.11: Monoid homomorphism

A *monoid homomorphism* $h: (\mathcal{M}, \bullet, \mathbf{1}) \rightarrow (\mathcal{M}', \bullet', \mathbf{1}')$ between two monoids $(\mathcal{M}, \bullet, \mathbf{1})$ and $(\mathcal{M}', \bullet', \mathbf{1}')$ is a function $h: \mathcal{M} \rightarrow \mathcal{M}'$ such that

$$h(\mathbf{1}) = \mathbf{1}' \tag{2}$$

and

$$h(m \bullet n) = h(m) \bullet' h(n)$$

for all $m, n \in \mathcal{M}$.

Definition 2.12: Category of monoids

The category **Mon** has as objects monoids and as morphisms monoid homomorphisms.

We can now state the definition of the static structure of a CA:

Definition 2.13: Cellular automaton (classical)

A *cellular automaton* (CA) \mathfrak{A} is a 4-tuple $\mathfrak{A} = (\mathcal{M}, N, S, \gamma)$, where:

- $(\mathcal{M}, \bullet, \mathbf{1})$ is a monoid that represents the graph;
- $N \subseteq \mathcal{M}$ is the *neighbourhood*. We require that $\mathbf{1} \in N$;
- S is a set of *states*;

- $\gamma: S^N \rightarrow S$ is a *local rule*.

We interpret elements $x \in \mathcal{M}$ as cells. But an element $m \in \mathcal{M}$ also represents an edge label for the edge between x and $m \bullet x$ (for every $x \in \mathcal{M}$).

Remark 2.14. This is a very general definition of a CA. Many authors limit the allowed monoids to relatively regular grids. For example, only monoids generated over a finite set. Or, even stronger, only those of the form \mathbb{Z}^d with $d \in \mathbb{N}$. The present definition allows monoids that are not grid-like at all, such as binary trees.

Remark 2.15. Not all authors require that $\mathbf{1} \in N$. Ultimately this makes little difference; one can always define γ to ignore the state at $\mathbf{1}$ (i.e., $\gamma(z) = \gamma(z')$ for all $z, z': N \rightarrow S$ that only differ on input $\mathbf{1}$).

The static structure of a CA determines the dynamics given by the global rule as follows:

Definition 2.16: Global rule of a CA (classical)

An assignment of states to cells $c: \mathcal{M} \rightarrow S$ is called a *configuration*. Every CA \mathfrak{A} induces a map $G_{\mathfrak{A}}: S^{\mathcal{M}} \rightarrow S^{\mathcal{M}}$ between configurations called the *global rule*. For every cell $m \in \mathcal{M}$, it is defined as follows:

$$G_{\mathfrak{A}}(c)(m) \stackrel{\text{def}}{=} \gamma(\lambda n \in N . c(n \bullet m)). \quad (3)$$

An initial configuration $c_0: \mathcal{M} \rightarrow S$ induces an sequence of configurations $(c_i)_{i \in \mathbb{N}}$ called the *trace sequence*, which is inductively defined as $c_i \stackrel{\text{def}}{=} G_{\mathfrak{A}}^i(c_0)$.

Example 2.17 (Conway’s Game Of Life (CGOL)). We can now describe CGOL formally. The Game of Life is the CA

$$\text{CGOL} \stackrel{\text{def}}{=} (\mathbb{Z}^2, N, \{\bullet, \circ\}, \gamma), \quad (4)$$

where $N = \{(i, j) \mid -1 \leq i, j \leq 1\}$ is the 3×3 rectangle surrounding a cell, and where the state set $\{\bullet, \circ\}$ has only a “dead” (\circ) and “alive” (\bullet) state. For every input $k: N \rightarrow \{\bullet, \circ\}$ of the local rule, denote $A(k) \stackrel{\text{def}}{=} \text{card}\{(i, j) \in N \mid k((i, j)) = \bullet\}$. Then the local rule $\gamma: \{\bullet, \circ\}^N \rightarrow \{\bullet, \circ\}$ can be defined as:

$$\gamma(k) \stackrel{\text{def}}{=} \begin{cases} \circ & A(k) \leq 2 \\ \bullet & A(k) = 3 \\ k((0, 0)) & A(k) = 4 \\ \circ & A(k) \geq 5 \end{cases} \quad (5)$$

2.3 Category theory

We briefly review concepts from category theory, to start with coalgebras and comonads. For general introduction into coalgebras, see for example [15], [39], [38] or [16]. More background on comonads can be found in textbooks such as [26, Ch6] or [35, Ch5].

Definition 2.18: Coalgebra

A *coalgebra* of an endofunctor $F: \mathcal{C} \rightarrow \mathcal{C}$ on a category \mathcal{C} an F -coalgebra is a pair $\mathcal{X} = (X, d)$ where $X \in \mathcal{C}$ is an object (referred to as the *carrier*) and $d: X \rightarrow FX$ a morphism in \mathcal{C} (referred to as the *dynamics*, or (if clear from context) as the *coalgebra*).

A *coalgebra homomorphism* between two F -coalgebras (X, d) and (Y, e) is a morphism $h: X \rightarrow Y$ in \mathcal{C}

such that commutes:

$$\begin{array}{ccc} X & \xrightarrow{h} & Y \\ d \downarrow & & \downarrow e \\ FX & \xrightarrow{Fh} & FY \end{array} \quad (6)$$

This diagram will be called the *coalgebra homomorphism law*.

The coalgebras of F together with the coalgebra homomorphisms form a category $\text{Coalg}(F)$.

We typeset coalgebras in a thick calligraphic font, e.g., \mathcal{X} . If a coalgebra name contains multiple symbols we use overarrows instead, e.g., $\overrightarrow{\text{CGOL}}$.

If the endofunctor F is a *comonad*, then one can construct the *coEilenberg-Moore* category $\text{coEM}(F)$ of comonad-coalgebras, which is a subcategory of $\text{Coalg}(F)$.

Definition 2.19: Comonad

A *comonad* is an endofunctor $F: \mathcal{C} \rightarrow \mathcal{C}$ equipped with two natural transformation $\varepsilon_X: FX \rightarrow X$ (called the *counit*) and $\delta_X: FX \rightarrow FFX$ (called the *comultiplication*) such that the following diagrams commute for all $X \in \mathcal{C}$:

$$\begin{array}{ccc} FX & \xrightarrow{\delta_X} & FFX \\ \delta_X \downarrow & & \downarrow F\delta_X \\ FFX & \xrightarrow{\delta_{FX}} & FFFX \end{array} \quad (7)$$

$$\begin{array}{ccccc} & & FX & & \\ & & \downarrow \delta_X & & \\ FX & \xleftarrow{\delta_{FX}} & FFX & \xrightarrow{F\delta_X} & FC \end{array} \quad (8)$$

which will be referred to as the *comonad coassociativity law* and *comonad counit law* respectively.

Definition 2.20: coEilenberg-Moore coalgebras

The category $\text{coEM}(F)$ of comonad $F: \mathcal{C} \rightarrow \mathcal{C}$ is the subcategory of $\text{Coalg}(F)$ containing only the objects (X, d) such that the following diagrams commute:

$$\begin{array}{ccc} X & \xrightarrow{d} & FX \\ d \downarrow & & \downarrow Fd \\ FX & \xrightarrow{\delta_X} & FFX \end{array} \quad (9)$$

$$\begin{array}{ccc} X & \xrightarrow{d} & FX \\ & \searrow & \downarrow \varepsilon_X \\ & & X \end{array} \quad (10)$$

These diagrams will be referred to as the *comonad coalgebra coassociativity law* and the *comonad coalgebra counit law*, respectively.

For coalgebras on endofunctors of **Set** there exists a natural notion of similarity between elements of their carriers, known as *behavioural equivalence*. See Staton [43] for an overview of related definitions generalised beyond **Set**.

Definition 2.21: Behavioural equivalence

Let $\mathcal{X} = (X, d)$ and $\mathcal{Y} = (Y, e)$ be F -coalgebras. Let $x \in X$ and $y \in Y$. Then x is *behaviourally equivalent* to y , denoted $x \overset{\text{BE}}{\approx} y$, if there exist a coalgebra (Z, f) and coalgebra homomorphism $h: (X, d) \rightarrow (Z, f)$ and $k: (Y, e) \rightarrow (Z, f)$ such that $h(x) = k(y)$.

Remark 2.22. For any coalgebra \mathcal{Y} there always exists an identity endohomomorphism $\text{Id}_{\mathcal{Y}}: \mathcal{Y} \rightarrow \mathcal{Y}$. As a result, for any coalgebra homomorphism $h: \mathcal{X} \rightarrow \mathcal{Y}$ it automatically holds that $x \overset{\text{BE}}{\approx} h(x)$ for all $x \in X$ (this uses $k := \text{Id}_{\mathcal{Y}}$).

In this thesis we use functors defined in terms of dependent sums in **Set**, for which we use the following definition:

Definition 2.23: Dependent sums of sets

Given a set A and a A -indexed family of sets B the dependent sum $\sum_{a \in A} B(a)$ is defined as the set of tuples (a, b) such that $b \in B(a)$ (the set $B(a)$ from which the second element is drawn depends on the first element a).

Remark 2.24. The set $\sum_a B(A)$ is the coproduct in **Set** of the diagram

$$\begin{aligned} D: \Delta_A &\rightarrow \mathbf{Set} \\ D: \Delta_A &\mapsto \{a\} \times B(a) \end{aligned}$$

where Δ_A is the small discrete category with elements of A as objects and with only identity morphisms.

2.4 Modal logic

We review only the basic definitions and relevant results of modal logic. We refer the reader to the textbook by Blackburn, De Rijke and Venema [2] for a comprehensive introduction.

A Kripke modal logic is an extension of propositional (Boolean) logic with unary operators \diamond_i known as *modal operators*.

Definition 2.25: Grammar of modal logic

A *signature* for modal logics is a pair (τ, Φ) where τ is a set of unary operators $(\diamond_i)_{i \in \tau}$ and $\Phi = \{p, q, r, \dots\}$ a set of propositional letters (equivalent to Boolean variables).

The set of well-formed formulas $\text{Form}(\tau, \Phi)$ on a signature (τ, Φ) is:

$$\text{Form}(\tau, \Phi) ::= \neg\psi \mid \psi_1 \vee \psi_2 \mid \perp \mid p \mid \diamond_i\psi$$

where $p \in \Phi$ and $i \in \tau$. We abbreviate $\neg \diamond_i \neg\psi$ as $\square_i\psi$.

One can generalise the diamonds to be n -ary operators for any $n \in \mathbb{N}$, but this generalisation is not relevant for the present work.

Modal formulas are interpreted at points in relational structures, also known as *Kripke models*.

Definition 2.26: Kripke frame and models

Let (τ, Φ) be a signature for a Kripke modal logic.

A *Kripke frame* $\mathbb{F} = (W, R_i)_{i \in \tau}$ consists of a set W *points* (also called *worlds*), and a τ -indexed set of

binary relations $(R_i \subseteq W \times W)_{i \in \tau}$.

A *Kripke model* (\mathbb{F}, V) consists of a Kripke frame $\mathbb{F} = (W, R_i)_{i \in \tau}$ together with a *valuation* $V: W \rightarrow \wp(\Phi)$.

A formula $\psi \in \text{Form}(\tau, \Phi)$ is either true (*satisfied*) or false (*unsatisfied*) when interpreted at a specific point $w \in W$ of a model $((W, R_i)_{i \in \tau}, V)$. Propositional letters are true at w when they occur in $V(w)$. However, the truth of ψ at w does not only depend on $V(w)$: the modal operators $\diamond_i \psi'$ allow to evaluate the truth a subformula ψ' at a R_i -neighbour v (i.e., $(w, v) \in R_i$). In particular, $\diamond_i \psi'$ is true at w as soon as there exist *any* R_i neighbour of w where ψ' is true. Dually, $\Box_i \psi'$ is true at w if ψ' is true at *all* R_i -neighbours of w .

Definition 2.27: Kripke semantics modal logic

Given a Kripke model $M = (\mathbb{F} = (W, R_i)_{i \in \tau}, V)$ defined on a signature (τ, Φ) , we define for each $w \in W$ and formula $\psi \in \text{Form}(\tau, \Phi)$ the proposition $M, w \models \psi$ to either be true or false (synonyms: ψ is either *satisfied* or *unsatisfied* at w in M). This truth is defined by structural induction on ψ as follows:

$$\begin{aligned} M, w \models \neg\psi & \quad \text{iff it is false that } M, w \models \psi \\ M, w \models \psi_1 \vee \psi_2 & \quad \text{iff } (M, w \models \psi_1 \text{ or } M, w \models \psi_2) \\ M, w \models \perp & \quad \text{is always false} \\ M, w \models p & \quad \text{iff } p \in V(w) \\ M, w \models \diamond_i \psi & \quad \text{iff there exists a } v \in W \text{ such that } (w, v) \in R_i \text{ and } M, v \models \psi. \end{aligned}$$

A formula ψ is:

- *globally satisfied in a model* M if $M, w \models \psi$ for all $w \in W$.
- *valid in a frame* \mathbb{F} if $(\mathbb{F}, V), w \models \psi$ for all $V: W \rightarrow \wp(\Phi)$ and $w \in W$.
- *valid at a point* $w \in W$ if $(\mathbb{F}, V), w \models \psi$ for all $V: W \rightarrow \wp(\Phi)$.

Remark 2.28. The semantics of formulas without modal operators are the same as in propositional logic, which justifies the following familiar notational shorthands (for all $\psi_1, \psi_2 \in \text{Form}(\tau, \Phi)$):

$$\begin{aligned} \psi_1 \wedge \psi_2 & \stackrel{\text{def}}{=} \neg((\neg\psi_1) \vee (\neg\psi_2)) \\ \psi_1 \rightarrow \psi_2 & \stackrel{\text{def}}{=} (\neg\psi_1) \vee \psi_2 \\ \psi_1 \leftrightarrow \psi_2 & \stackrel{\text{def}}{=} (\psi_1 \rightarrow \psi_2) \wedge (\psi_2 \rightarrow \psi_1) \\ \top & \stackrel{\text{def}}{=} \neg\perp. \end{aligned}$$

An important result of modal logic is the Hennessy-Milner theorem [12]. Intuitively, the theorem asserts that two points $w, v \in W$ of a model satisfy the same set of formulas if and only if they have a similar successor structure (by recursively following outgoing relations R_i) in the model. The latter property is called *bisimilarity*, and the precise definition is as follows:

Definition 2.29: Bisimulation and logical equivalence

Let (τ, Φ) be a modal logic signature. Let $M = (W, (R_i)_{i \in \tau}, V)$ and $M' = (W', (R'_i)_{i \in \tau}, V')$ be Kripke models. A bisimulation $Z: M \rightleftharpoons M'$ is a binary relation $Z \subseteq W \times W'$ that satisfies the following three conditions:

1. If uZu' then $V(u) = V'(u')$ (for all $u \in W$ and $u' \in W'$).
2. If $uR_i v$ and uZu' then there exists a $v' \in W$ such that $u'R_i v'$ (for all $u, v \in W$, $u' \in W'$ and $i \in \tau$).
3. If $u'R_i v'$ and uZu' then there exists a $v \in W$ such that $uR_i v$ (for all $u \in W$, $u', v' \in W'$ and $i \in \tau$).

Points $w \in W$ and $w' \in W'$ are *bisimilar*, denoted $w \Leftrightarrow w'$, if there exists a bisimulation $Z: M \Leftrightarrow M'$ such that wZw' . The points are *logically equivalent*, denoted $w \rightsquigarrow w'$, if for all $\psi \in \text{Form}(\tau, \Phi)$ it holds that

$$(M, w \models \psi) \quad \text{iff} \quad (M', w' \models \psi).$$

Definition 2.30: Image finite frame

A Kripke frame $\mathbb{F} = (W, (R_i)_{i \in \tau})$ is *image-finite* if for all relations R_i and all points $w \in W$ the set

$$R_i(w) \stackrel{\text{def}}{=} \{v \in W \mid wR_i v\}$$

is finite.

Theorem 2.31: Hennessy and Milner, 1984, [12]

Let (τ, Φ) be a Kripke modal logic signature, $M = (W, (R_i)_{i \in \tau}, V)$ and $M' = (W', (R'_i)_{i \in \tau}, V')$ be image-finite Kripke models and $w \in W$, $w' \in W'$ be points. Then

$$w \Leftrightarrow w' \quad \text{iff} \quad w \rightsquigarrow w'.$$

3 Uniform cellular automata

In this section we exhibit the class of CA as a subcategory $\mathbf{CCA}_{\mathbb{U}}^+$ of the coalgebras of an endofunctor $C_{\mathbb{U}}: \mathbf{Set} \rightarrow \mathbf{Set}$. The main result, Theorem 3.17, states that there is an essentially surjective functor from classical CA c.f. Definition 2.13 to $\mathbf{CCA}_{\mathbb{U}}^+$. This coalgebraic formalisation will be refined in section 4, and thereafter feature in the semantics of the modal logic described in Sections 6 and 7.

3.1 Graph as coalgebras

We first sketch the intuition of our coalgebraic model of CA. Unlike the classical definition of a CA (Definition 2.13) and Capobianco's and Uustalu's [5][4] model, we will distinguish between cells and paths between cells. To describe how cells are interconnected, we use the same *cowriter comonad* as used by Capobianco and Uustalu (in a limited form described for CA before by Piponi [34]). To model paths, we use a monoid $(\mathcal{M}, \bullet, \mathbf{1})$, and interpret the monoid operation $m \bullet n$ as path composition (so traversing $m \bullet n$ is equivalent to first traversing n and m thereafter).

The comonad works as follows: a coalgebra (X, χ) consists of a carrier X , the set of cells, and a mapping $\chi: X \rightarrow X^{\mathcal{M}}$. Unlike Capobianco and Uustalu, we interpret X as the set of cells instead of the set of states. We take the convention to write χ as an infix operator \otimes . For every path $m \in \mathcal{M}$ and cell $x \in X$, we compute the cell one arrives at by traversing the path m from the initial cell x as $m \otimes x$.

We will show that the comonad structure ensures that the graph structure is coherent. First of all, traversing the empty path $\mathbf{1} \in \mathcal{M}$ should not lead to a different cell. In particular, we require that the following identity holds for all $x \in X$:

$$\mathbf{1} \otimes x = x. \quad (11)$$

We further require that \otimes respects path composition. Suppose that m_1 is a path from x to x' , and that m_2 is a path from x' to x'' . Then in the \bullet -is-path-composition interpretation it should hold that $m_2 \bullet m_1$ is a path from x to x'' .

$$\begin{array}{ccccc} & & m_2 \bullet m_1 & & \\ & & \curvearrowright & & \\ x & \xrightarrow{m_1} & x' & \xrightarrow{m_2} & x'' \end{array}$$

This means that we require (for all $m_1, m_2 \in \mathcal{M}$ and all $x \in X$):

$$m_2 \otimes (m_1 \otimes x) = (m_2 \bullet m_1) \otimes x. \quad (12)$$

(From now on, we will usually omit parenthesis in expressions of the form $m_2 \otimes m_1 \otimes x$).

Requiring χ to adhere to conditions Eq. (11) and (12) turns out to be equivalent to requiring (X, χ) to be a coalgebra in the coEilenberg-Moore category (Definition 2.20) of the following cowriter comonad:

Definition 3.1: Cowriter functor

The *cowriter functor* $P = P: \mathbf{Set} \rightarrow \mathbf{Set}$ is the functor whose action on objects $X \in \mathbf{Set}$ is

$$P: X \mapsto X^{\mathcal{M}} \quad (13)$$

and whose action on functions $f: X \rightarrow Y$ is:

$$Pf \stackrel{\text{def}}{=} \lambda k . f \circ k: X^{\mathcal{M}} \rightarrow Y^{\mathcal{M}} \quad (14)$$

for all $k: \mathcal{M} \rightarrow X$.

To make P a comonad, we define the associated natural transformations:

Definition 3.2: Comonad structure of the cowriter functor

The counit ε and comultiplication δ of P are:

$$\begin{array}{ll} \varepsilon: P \Rightarrow \text{Id} & \delta: P \Rightarrow PP \\ \varepsilon_X: X^{\mathcal{M}} \rightarrow X & \delta_X: X^{\mathcal{M}} \rightarrow (X^{\mathcal{M}})^{\mathcal{M}} \\ \varepsilon_X(k) \stackrel{\text{def}}{=} k(\mathbf{1}) & \delta_X(k) \stackrel{\text{def}}{=} \lambda m . \lambda n . k(n \bullet m) \end{array} \quad (15)$$

for all $k: \mathcal{M} \rightarrow X$ and where $m, n \in \mathcal{M}$.

Remark 3.3. Alternatively, we could also define comultiplication δ as

$$\delta_X: k \mapsto \lambda m . \lambda n . k(m \bullet n). \quad (16)$$

In this case, we should interpret $m \bullet n$ as the path obtained from first traversing m and thereafter n (for $n, m \in \mathcal{M}$), and it would be more convenient to write $x \otimes m \otimes n$ instead of $m \otimes n \otimes x$. We do not follow this alternative approach, but it showcases that the interpretation of the order in path composition is arbitrary.

Before we prove that this is indeed a comonad, we check that the comonad coassociativity and counit laws, (9) and (10) respectively, indeed correspond to Eq. (12) and (11) respectively. The first law (9) expresses

that the following diagram commutes:

$$\begin{array}{ccc} X & \xrightarrow{\chi} & X^{\mathcal{M}} \\ x \downarrow & & \downarrow P\chi \\ X^{\mathcal{M}} & \xrightarrow{\delta_x} & (X^{\mathcal{M}})^{\mathcal{M}} \end{array}$$

On inputs $x \in X$ and $m, n \in \mathcal{M}$, the upper path outputs:

$$\begin{aligned} (P\chi \circ \chi)(x, m, n) &= (P\chi(\chi(x)))(m, n) \\ &= (\chi \circ (\chi(x)))(m, n) && // \text{Definition 3.1.} \\ &= (\chi(m \otimes x))(n) \\ &= n \otimes (m \otimes x). \end{aligned}$$

The lower path outputs:

$$\begin{aligned} (\delta_X \circ \chi)(x, m, n) &= (\delta_X(\chi(x)))(m, n) \\ &= (\lambda m' \cdot \lambda n' \cdot (n' \bullet m') \otimes x)(m, n) && // \text{Definition 3.2.} \\ &= (n \bullet m) \otimes x. \end{aligned}$$

Equality of these paths is indeed exactly requirement (12).

The second law (10) states that the following diagram commutes:

$$\begin{array}{ccc} X & \xrightarrow{\chi} & X^{\mathcal{M}} \\ & \searrow & \downarrow \varepsilon_x \\ & & X \end{array}$$

Then, for every $x \in X$, we have:

$$x = (\varepsilon_X \circ \chi)(x) = \varepsilon_X(\chi(x)) = \chi(x)(\mathbf{1}) = \mathbf{1} \otimes x,$$

as required by (11).

We now show that Definition 3.2 indeed a comonad structure on the cowriter functor P . This result is not new, but we present the details for the reader's convenience:

Lemma 3.4

(P, ε, δ) is a comonad.

Proof. Referring to Definition 3.2, we need to show that $\varepsilon: P \Rightarrow \text{Id}$ and $\delta: P \Rightarrow PP$ are indeed natural in their arguments, and that the comonad coassociativity law (7) and comonad counit law (8) are satisfied.

As for the naturality of ε , let $f: X \rightarrow Y$, and it is to show that commutes:

$$\begin{array}{ccc} X^{\mathcal{M}} & \xrightarrow{Pf} & Y^{\mathcal{M}} \\ \varepsilon_X \downarrow & & \downarrow \varepsilon_Y \\ X & \xrightarrow{f} & Y \end{array}$$

Indeed, for any $k: \mathcal{M} \rightarrow X$ we observe that

$$(f \circ \varepsilon_X)(k) = f(k(\mathbf{1})) = (f \circ k)(\mathbf{1}) = (\varepsilon_X \circ Pf)(k).$$

The naturality condition for δ for a function $f: X \rightarrow Y$ is

$$\begin{array}{ccc} X^{\mathcal{M}} & \xrightarrow{Pf} & Y^{\mathcal{M}} \\ \delta_X \downarrow & & \downarrow \delta_Y \\ (X^{\mathcal{M}})^{\mathcal{M}} & \xrightarrow{PPf} & (Y^{\mathcal{M}})^{\mathcal{M}} \end{array}$$

We can indeed calculate that for any input $k: \mathcal{M} \rightarrow X$ we have:

$$\begin{aligned} (PPf) \circ \delta_X(k) &= PPf(\lambda m . \lambda n . k(n \bullet m)) && // \text{Definition } \delta \text{ (Definition 3.2).} \\ &= \lambda m . Pf \circ \lambda n . k(n \bullet m) && // \text{Definition outermost } P \text{ (Definition 3.1).} \\ &= \lambda m . \lambda n . (f \circ k)(n \bullet m) && // \text{Definition } P \text{ (Definition 3.1).} \\ &= (\delta_Y \circ (Pf))(k). && // \text{Definition } \delta \text{ (Definition 3.2).} \end{aligned}$$

The comonad coassociativity law (7) requires that commutes:

$$\begin{array}{ccc} X^{\mathcal{M}} & \xrightarrow{\delta_X} & (X^{\mathcal{M}})^{\mathcal{M}} \\ \delta_X \downarrow & & \downarrow P\delta_X \\ (X^{\mathcal{M}})^{\mathcal{M}} & \xrightarrow{\delta_{X^{\mathcal{M}}}} & ((X^{\mathcal{M}})^{\mathcal{M}})^{\mathcal{M}} \end{array} \quad (17)$$

For an input $k: \mathcal{M} \rightarrow X$, the upper path evaluates to:

$$\begin{aligned} (\delta_{X^{\mathcal{M}}} \circ \delta_X)(k) &= \lambda u . \lambda v . \delta_X(k)(v \bullet u) && // \text{Definition } \delta \text{ (Definition 3.2).} \\ &= \lambda u . \lambda v . \lambda m . \lambda n . k(n \bullet m)(v \bullet u) && // \text{Definition } \delta \text{ (Definition 3.2).} \\ &= \lambda u . \lambda v . \lambda n . k(n \bullet v \bullet u). \end{aligned}$$

This indeed agrees with the lower path:

$$\begin{aligned} (P\delta_X \circ \delta_X)(k) &= P\delta_X[\lambda u . \lambda w . k(w \bullet u)] \\ &= \lambda u . \delta_X(\lambda u . \lambda w . k(w \bullet u)) && // \text{Definition } P \text{ (Definition 3.1).} \\ &= \lambda u . \lambda v . \lambda n . (\lambda u . \lambda w . k(w \bullet u))(n \bullet v) && // \text{Definition } \delta \text{ (Definition 3.2).} \\ &= \lambda u . \lambda v . \lambda n . k(n \bullet v \bullet u), \end{aligned}$$

which shows that both paths in the diagram indeed agree.

Finally, we need to show that the comonad counit law (8) holds:

$$\begin{array}{ccc} & X^{\mathcal{M}} & \\ & \parallel & \\ X^{\mathcal{M}} & \xrightarrow{\varepsilon_{(X^{\mathcal{M}})}} & (X^{\mathcal{M}})^{\mathcal{M}} \xrightarrow{P\varepsilon_X} X^{\mathcal{M}} \\ & \parallel & \\ & X^{\mathcal{M}} & \end{array} \quad (18)$$

To this end, we show for both inner triangles separately that they commute on any input $k: \mathcal{M} \rightarrow X$. The following computation shows that the right triangle commutes:

$$\begin{aligned} P\varepsilon_X(\delta_X(k)) &= P\varepsilon_X(\lambda m . \lambda n . k(n \bullet m)) \\ &= \lambda m . \varepsilon_X(\lambda n . k(n \bullet m)) && // \text{Definition } P \text{ (Definition 3.1).} \\ &= \lambda m . k(\mathbf{1} \bullet m) && // \text{Definition } \varepsilon \text{ (Definition 3.2).} \\ &= k. \end{aligned}$$

Commutativity of the left triangle follows by a similar computation, using the definitions of δ and ε (Definition 3.2):

$$\varepsilon_{(X^{\mathcal{M}})}(\delta_X(k)) = \varepsilon_{(X^{\mathcal{M}})}(\lambda m . \lambda n . k(n \bullet m)) = \lambda n . k(n \bullet \mathbf{1}) = k.$$

□

3.2 Extending graphs with local rules

The coalgebras of the paths functor P are cells with a graph structure. These structured need to be complemented with local rules before they become CA, and we will build the local rules directly into the coalgebra type functor. To this end, we will modify P into a “uniform CA” functor C_U with an additional component for specifying local rules. We will use C_U to construct a category of coalgebras that correspond to uniform CA, i.e., CA where all cells use the same local rule, as in Definition 2.13 (the subscript u stands for *uniform*). Section 4 will construct an endofunctor C_G and a category of *non-uniform* coalgebras in which cells are allowed to have different local rules.

In the following, we assume an additional parameter next to \mathcal{M} : a state set $S \in \mathbf{Set}$.

Definition 3.5

The functor $C_U: \mathbf{Set} \rightarrow \mathbf{Set}$ is given by the following action on every set $X \in \mathbf{Set}$:

$$C_U(X) \stackrel{\text{def}}{=} X^{\mathcal{M}} \times S^{S^{\mathcal{M}}}$$

and the action on functions $f: X \rightarrow Y$ is:

$$C_U f \stackrel{\text{def}}{=} \lambda k . (f \circ k) \times \text{Id}$$

where $k: \mathcal{M} \rightarrow X$ and $(f \circ k): \mathcal{M} \rightarrow Y$.

Let (X, d) be a C_U -coalgebra and $x \in X$. Then $d(x) \in X^{\mathcal{M}} \times S^{S^{\mathcal{M}}}$ has two components $\pi_1(d(x))$ and $\pi_2(d(x))$, where π_i denotes the projection to the i^{th} component. All our applications will use these components separately, and we use the following notational convention to ease this: if we define $\chi \stackrel{\text{def}}{=} \pi_1 \circ d$ and $\gamma \stackrel{\text{def}}{=} \pi_2 d$, then we will write $\langle \chi, \gamma \rangle$ for d .

Observe that first component $\chi: X \rightarrow X^{\mathcal{M}}$ has the same type as a cowriter coalgebra as discussed in Section 3.1. The second component has type

$$\gamma: X \rightarrow S^{S^{\mathcal{M}}}. \quad (19)$$

Thus γ assigns to every cell $x \in X$ a local rule $\gamma_x: S^{\mathcal{M}} \rightarrow S$, with neighbourhood \mathcal{M} . This is a more general collection of local rules than the local rules in a classical CA according to Definition 2.13: (1) it is possible that $\gamma_x \neg \gamma_{x'}$ for $x, x' \in X$, and (2) the neighbourhood is the entire monoid \mathcal{M} , while a CA may have a smaller neighbourhood $N \subseteq \mathcal{M}$.

Example 3.6 (A coalgebraic model of Conway’s Game of Life). It is straightforward to describe CGOL (Example 2.17) as a C_U -coalgebra $\overrightarrow{\text{CGOL}} \stackrel{\text{def}}{=} (\mathbb{Z}^2, \langle +, \gamma \rangle)$ where $(\mathbb{Z}^2, +, (0, 0))$ is the underlying monoid and the used cowriter comonad coalgebra “+” is component-wise addition: $(i, j) \oplus (k, \ell) \stackrel{\text{def}}{=} (i + k, j + \ell)$. The neighbourhood N and local rule γ are exactly as in the classical definition of Example 2.17.

3.3 Uniform cellular automata as coalgebras

We will now construct the category of C_U -coalgebras that corresponds to classical CA (Definition 2.13). Different CA may have different underlying monoids \mathcal{M} , neighbourhoods $N \subseteq \mathcal{M}$ and state sets S . Therefore

we will need to combine coalgebras of $C_{\mathbb{U}}$ for different choices of the parameters \mathcal{M} and S . To avoid ambiguity, we write $C_{\mathbb{U}}^{(\mathcal{M}, S)}$ when using $C_{\mathbb{U}}$ with parameters \mathcal{M} and S .

We will call the parameters of a CA a *signature*:

Definition 3.7: Signature

A *signature* is a triplet (\mathcal{M}, N, S) where $(\mathcal{M}, \bullet, \mathbf{1})$ is a monoid, $N \subseteq \mathcal{M}$ with $\mathbf{1} \in \mathcal{M}$ a neighbourhood, and $S \in \mathbf{Set}$ a state set.

After an auxiliary definition (of *rooted coalgebras*), we will specify the subcategory $\mathbf{CCA}_{\mathbb{U}}^{(\mathcal{M}, N, S)}$ of $\mathbf{Coalg}(C_{\mathbb{U}}^{(\mathcal{M}, S)})$ of coalgebras that correspond to CA defined on this signature (\mathcal{M}, N, S) . Thereafter we will combine these subcategories (over all choices of signatures) in one category $\mathbf{CCA}_{\mathbb{U}}^+$, and show that the latter exactly corresponds to the collection of classical CA conform Definition 2.13.

For any $M \subseteq \mathcal{M}$, a coalgebra $(X, \langle \chi, \gamma \rangle)$ and an $x \in X$, we will denote the M -image of $\chi(x)$ as:

$$M \otimes x \stackrel{\text{def}}{=} \{m \otimes r \mid m \in M\}. \quad (20)$$

We can now define a *rooted coalgebra*:

Definition 3.8: Rooted coalgebra

A coalgebra $\mathcal{X} = (X, \langle \chi, \gamma \rangle)$ is *rooted* in a distinguished root $r \in X$ if $\mathcal{M} \otimes r = X$.

Remark 3.9. The choice of r may not be unique, and not every coalgebra may have a root.

Definition 3.10

Let (\mathcal{M}, N, S) be a signature, then $\mathbf{CCA}_{\mathbb{U}}^{(\mathcal{M}, N, S)}$ is the subcategory of $\mathbf{Coalg}(C_{\mathbb{U}}^{(\mathcal{M}, S)})$ whose objects $\mathcal{X} = (X, \langle \chi, \gamma \rangle)$ satisfy the following constraints:

1. There exists an *essential local rule* $\hat{\gamma}: S^N \rightarrow S$ such that $\gamma_x(k) = \hat{\gamma}(k \upharpoonright N)$ for all $x \in X$ and all $k: \mathcal{M} \rightarrow S$.
2. \mathcal{X} is rooted in a chosen root $r \in X$.
3. $(-)\otimes r: \mathcal{M} \rightarrow X$ is an injection.
4. (X, χ) is a cowriter comonad coalgebra (of the functor P , Definition 3.1 and 3.2).

We usually make the essential local rule $\hat{\gamma}$ and the root $r \in X$ explicit, in which case we denote a coalgebra in $\mathbf{CCA}_{\mathbb{U}}$ as $\mathcal{X} = (X, \langle \chi, \gamma \rangle, \hat{\gamma}, r)$.

The morphisms in $\mathbf{CCA}_{\mathbb{U}}^{(\mathcal{M}, N, S)}$ are root-preserving coalgebra homomorphisms.

Remark 3.11. Item 1 ensures that the CA has exactly one local rule for all cell (uniformity), and that this rule is also actually local (its output only depends on the states in an N -neighbourhood around x).

Example 3.12. One-dimensional CA are CA on the monoid $(\mathbb{Z}, +, 0)$ with as cells \mathbb{Z} and as cowriter comonad coalgebra the usual addition. Any cell $n \in \mathbb{Z}$ can be chosen as the root: every number $m \in \mathbb{N}$ can be uniquely written as $m = (m - n) + n$, thus $(-) + n$ has \mathbb{Z} as image, and is injective.

Now we combine the $\mathbf{CCA}_{\mathbb{U}}^{(\mathcal{M}, N, S)}$ over all choices of signature into one category $\mathbf{CCA}_{\mathbb{U}}^+$. A technical detail required to prove the correspondence between classical CA and $\mathbf{CCA}_{\mathbb{U}}^+$ (in Theorem 3.17) is that we need to

introduce additional arrows aside from existing the coalgebra homomorphisms. Intuitively, these extra arrows are isomorphisms between coalgebras that are identical up to a renaming of the elements of the monoid they are defined on.

Definition 3.13

The category $\mathbf{CCA}_{\mathbb{U}}^+$ is the coproduct of the categories $\mathbf{CCA}_{\mathbb{U}}^{(\mathcal{M}, N, S)}$ over all signatures (\mathcal{M}, N, S) , extended with the following arrows: for every pair of signatures (\mathcal{M}, N, S) and (\mathcal{M}', N', S) with the same set of states, and for every monoid isomorphism $\phi: \mathcal{M} \xrightarrow{\sim} \mathcal{M}'$, and for all pairs of coalgebras

$$\begin{aligned} \mathcal{X} &= (X, \langle \chi, \gamma \rangle, \hat{\gamma}, r_X) \in \mathbf{CCA}_{\mathbb{U}}^{(\mathcal{M}, N, S)} \\ \mathcal{Y} &= (Y, \langle \xi, \delta \rangle, \hat{\delta}, r_Y) \in \mathbf{CCA}_{\mathbb{U}}^{(\mathcal{M}', N', S)}, \end{aligned}$$

if

1. $X = Y$ and $r_X = r_Y$,
2. $m \otimes x = \phi(m) \otimes x$ for all $x \in X$ and $m \in \mathcal{M}$,
3. and $\hat{\delta} = \lambda k \cdot \hat{\gamma}(k \circ \phi)$,

then there is an arrow $\phi_{\mathcal{X}, \mathcal{Y}} \in \mathbf{CCA}_{\mathbb{U}}^+(\mathcal{X}, \mathcal{Y})$ whose action on cells is the identity map. The inverse of this arrow is $(\phi^{-1})_{\mathcal{Y}, \mathcal{X}} \in \mathbf{CCA}_{\mathbb{U}}^+(\mathcal{Y}, \mathcal{X})$.

Remark 3.14. The isomorphism $\phi_{\mathcal{X}, \mathcal{Y}} \in \mathbf{CCA}_{\mathbb{U}}^+(\mathcal{X} \rightarrow \mathcal{Y})$ witnesses that the following diagram commutes:

$$\begin{array}{ccc} X & \xlongequal{\text{Id}} & Y \\ \langle \gamma, \chi \rangle \downarrow & & \downarrow \langle \xi, \delta \rangle \\ X^{\mathcal{M}} \times S^{S^{\mathcal{M}}} & \xrightarrow[\text{Id}']{} & Y^{\mathcal{M}'} \times S^{S^{\mathcal{M}'}} \end{array} \quad (21)$$

where

$$\text{Id}' \stackrel{\text{def}}{=} (\lambda f \cdot \text{Id} \circ (f \circ \phi^{-1})) \times (\lambda g \cdot \lambda k \cdot g(k \circ \phi)).$$

This is a condition similar to that of a coalgebra homomorphism (compare with diagram (6)), except that it has been extended with the monoid isomorphism ϕ to ensure type consistency. In the special case where $\phi = \text{Id}_{\mathcal{M}}$, diagram 21 simplifies to the coalgebra homomorphism law on $\text{Id}_X: \mathcal{X} \rightarrow \mathcal{X}$, since

$$(\lambda f \cdot \text{Id}_X \circ (f \circ \text{Id}_{\mathcal{M}}^{-1})) \times (\lambda g \cdot \lambda k \cdot g(k \circ \text{Id}_{\mathcal{M}})) = (\text{Id}_X \circ f) \times \text{Id}_{(S^{S^{\mathcal{M}}})} \stackrel{\text{def}}{=} C_{\mathbb{U}}(\text{Id}_X).$$

To prove that $\mathbf{CCA}_{\mathbb{U}}^+$ corresponds to the classical CA, we view the classical CA as category, and construct an essentially surjective functor. To our knowledge, there is no clear standard definition for morphisms between classical CA, so we formulate the classical CA as a discrete category:

Definition 3.15: Classical CA as a category

The category **classical** has as objects classical CA according to Definition 2.13, i.e., the collection of quadruples $(\mathcal{M}, N, S, \hat{\gamma})$ where \mathcal{M} is a monoid, $N \subseteq \mathcal{M}$ with $\mathbf{1} \in N$, $S \in \mathbf{Set}$ and $\hat{\gamma}: S^N \rightarrow S$. As morphisms it has only identity morphisms.

We will need the following auxiliary lemma in order to establish essential surjectivity:

Lemma 3.16

For every signature (\mathcal{M}, N, S) there exists a functor

$$G^{(\mathcal{M}, N, S)} : \mathbf{CCA}_{\mathbb{U}}^{(\mathcal{M}, N, S)} \rightarrow \mathbf{Mon}$$

such that for every $\mathcal{X} \in \mathbf{CCA}_{\mathbb{U}}^{(\mathcal{M}, N, S)}$ there exist a monoid isomorphism

$$\psi_{\mathcal{X}} : G^{(\mathcal{M}, N, S)}(\mathcal{X}) \xrightarrow{\sim} \mathcal{M}.$$

Proof. We abbreviate $G^{(\mathcal{M}, N, S)}$ as G . We first define the action of G on objects, so let $\mathcal{X} = (X, \langle \chi, \gamma \rangle, \hat{\gamma}, r) \in \mathbf{CCA}_{\mathbb{U}}^{(\mathcal{M}, N, S)}$ be a coalgebra rooted in r . Note that, by definition of a root (Definition 3.8) and the fact that $(-)\otimes r$ is an injection (Definition 3.10), every $x \in X$ can be written in the form $m_x \otimes r$ for a unique $m_x \in \mathcal{M}$. Then we define $G(\mathcal{X})$ to be the monoid (X, \bullet_X, r) . Here \bullet_X is defined for every $x = m_x \otimes r \in X$ and $x' = m_{x'} \otimes r \in X$ as

$$(m_x \otimes r) \bullet_X (m_{x'} \otimes r) \stackrel{\text{def}}{=} (m_x \bullet m_{x'}) \otimes r.$$

The monoid identity laws indeed hold for any $m_x \otimes r \in X$:

$$r \bullet_X (m_x \otimes r) = (\mathbf{1} \bullet m_x) \otimes r = m_x \otimes r = (m_x \bullet \mathbf{1}) \otimes r = (m_x \otimes r) \bullet_X r,$$

where we used the fact that $r = \mathbf{1} \otimes r$. The monoid associativity law follows from the associativity of \bullet : let $m_x \otimes r, m_y \otimes r, m_z \otimes r \in X$, then

$$\begin{aligned} (m_x \otimes r) \bullet_X ((m_y \otimes r) \bullet_X (m_z \otimes r)) &= (m_x \bullet (m_y \bullet m_z)) \otimes r && // \text{Definition } \bullet_X \text{ (used twice).} \\ &= ((m_x \bullet m_y) \bullet m_z) \otimes r && // \text{Associativity of } \bullet. \\ &= ((m_x \otimes r) \bullet_X (m_y \otimes r)) \bullet_X (m_z \otimes r). && // \text{Definition } \bullet_X \text{ (used twice).} \end{aligned}$$

Thus $G(\mathcal{X})$ is a well-defined monoid. To define G on morphisms, let $\mathcal{X} = (X, \langle \chi, \gamma \rangle, r_X, \hat{\gamma})$ and $\mathcal{Y} = (Y, \langle \xi, \delta \rangle, r_Y, \hat{\delta})$ be coalgebras in $\mathbf{CCA}_{\mathbb{U}}^{(\mathcal{M}, N, S)}$ and let $h : \mathcal{X} \rightarrow \mathcal{Y}$ by a coalgebra homomorphism. We claim that we can define $Gh \stackrel{\text{def}}{=} h$, so we show that $h : G(\mathcal{X}) \rightarrow G(\mathcal{Y})$ is a monoid homomorphism. The first condition of a monoid homomorphism (Definition 2.11) requires preservation of identity elements, but indeed $h(r_X) = r_Y$ by definition of $\mathbf{CCA}_{\mathbb{U}}^{(\mathcal{M}, N, S)}$. It remains to prove the other condition of a monoid homomorphism. To this end, we first observe that for any $x \in X$, the coalgebra homomorphism law (diagram (6) in Definition 2.18) on h evaluates to:

$$\xi(h(x)) \times (\delta_{h(x)}) = (h \circ \chi(x)) \times (h \circ \gamma_x)$$

Taking the first projections of both sides of the equation gives

$$\xi(h(x)) = h \circ \chi(x),$$

which on input $m \in \mathcal{M}$ evaluates (in infix notation) to:

$$m \otimes h(x) = h(m \otimes x). \tag{22}$$

This property allows to prove the second condition of a monoid homomorphism, that $h(x \bullet_X x') = h(x) \bullet_Y h(x')$ for all $x, x' \in X$. Write $x = m_x \otimes r_X$ and $x' = m_{x'} \otimes r_X \in X$, then it indeed holds that:

$$\begin{aligned} h((m_x \otimes r_X) \bullet_X (m_{x'} \otimes r_X)) &= h((m_x \bullet m_{x'}) \otimes r_X) && // \text{Definition } \bullet_X. \\ &= (m_x \bullet m_{x'}) \otimes h(r_X) && // \text{Eq. (22).} \\ &= (m_x \otimes h(r_X)) \bullet_Y (m_{x'} \otimes h(r_X)) && // \text{Definition } \bullet_Y. \\ &= h(m_x \otimes r_X) \bullet_Y h(m_{x'} \otimes r_X). && // \text{Eq. (22).} \end{aligned}$$

That G preserves composition and identity morphisms is trivial, thus we established that G is a functor.

It remains to give the monoid isomorphism $\psi_{\mathcal{X}}: G(\mathcal{X}) \xrightarrow{\sim} \mathcal{M}$. In the following, we abbreviate $\psi_{\mathcal{X}}$ by ψ . For all $m \otimes r \in X$, we define:

$$\psi_{\mathcal{X}}(m \otimes r) \stackrel{\text{def}}{=} m$$

This is well-defined since $(-)\otimes r$ is an injection (by Definition 3.10). The monoid homomorphism conditions (Definition 2.11) are straightforward to check:

$$\begin{aligned} \psi(r) &= \psi(\mathbf{1} \otimes r) = \mathbf{1} \\ \psi((m_x \otimes r) \bullet_X (m_{x'} \otimes r)) &= \psi((m_x \bullet m_{x'}) \otimes r) \\ &= m_x \bullet m_{x'} \\ &= \psi(m_x \otimes r) \bullet \psi(m_{x'} \otimes r), \end{aligned}$$

for all $m_x \otimes r, m_{x'} \otimes r \in X$. The inverse of ψ is $(-)\otimes r$, which also satisfies the monoid homomorphism conditions:

$$\begin{aligned} \mathbf{1} \otimes r &= r \\ (m \bullet n) \otimes r &= (m \otimes r) \bullet_X (n \otimes r). \end{aligned}$$

That $((-)\otimes r) \circ \psi = \text{Id}_{G(\mathcal{X})}$ and $\psi \circ ((-)\otimes r) = \text{Id}_{\mathcal{M}}$ are, respectively, easy to see:

$$\begin{aligned} \psi(m \otimes r) \otimes r &= m \otimes r \\ \psi(m \otimes r) &= m. \end{aligned}$$

□

Theorem 3.17

There exists an essentially surjective functor $F: \text{classical} \rightarrow \mathbf{CCA}_{\mathbb{J}}^+$.

Proof. We first construct a functor $F: \text{classical} \rightarrow \mathbf{CCA}_{\mathbb{J}}^+$, and thereafter show that it is essentially surjective. Given a classical CA $\mathfrak{A} = (\mathcal{M}, N, S, \hat{\gamma}) \in \text{classical}$, we define $F(\mathfrak{A})$ to be the coalgebra $(\mathcal{M}, \langle \chi, \gamma \rangle, \mathbf{1}, \hat{\gamma})$, which uses the monoid elements as cells and $\mathbf{1}$ as root. It remains to define χ and γ , and that $F(\mathfrak{A}) \in \mathbf{CCA}_{\mathbb{J}}^+$. For the latter, it suffices to show that $F(\mathfrak{A}) \in \mathbf{CCA}_{\mathbb{J}}^{(\mathcal{M}, N, S)}$, thus we will show that $F(\mathfrak{A})$ satisfies the conditions of $\mathbf{CCA}_{\mathbb{J}}^{(\mathcal{M}, N, S)}$ as given in Definition 3.10.

For all $m, x \in \mathcal{M}$, we define:

$$m \otimes x \stackrel{\text{def}}{=} m \bullet x.$$

We can immediately confirm that condition 4 of Definition 3.10 is satisfied: the properties of a comonad coalgebra, Eq. (11) and (12), follow directly from the monoid identity and associativity laws on \bullet (see Definition 2.3).

For all $x \in \mathcal{M}$ and $k: \mathcal{M} \rightarrow S$, we define

$$\gamma_x(k) \stackrel{\text{def}}{=} \hat{\gamma}(k \upharpoonright N).$$

Condition 1 of Definition 3.10 holds by construction. Condition 2, which requires that that $\mathbf{1}$ is indeed the root of the CA, follows from the monoid identity law (see Definition 2.3): $m = m \bullet \mathbf{1}$ for all $m \in \mathcal{M}$. Finally, Condition 3, i.e., injectivity of $(-)\otimes \mathbf{1}$, also follows from the monoid identity law since $\otimes = \bullet$ (for any $m, n \in \mathcal{M}$, if $m \bullet \mathbf{1} = n \bullet \mathbf{1}$ then the monoid identity law gives $m = n$).

This established F as a well-typed functor. As for essential surjectivity, assume any arbitrary $\mathcal{X} = (X, \langle \chi, \gamma \rangle, \dot{\gamma}, r) \in \mathbf{CCA}_{\cup}^+$. It is to show that there exists an $\mathfrak{A} \in \mathbf{classical}$ such that $F(\mathfrak{A}) \cong \mathcal{X}$. Without loss of generality, we may assume that $\mathcal{X} \in \mathbf{CCA}_{\cup}^{(\mathcal{M}, N, S)}$ for a signature (\mathcal{M}, N, S) .

From the definition of F we can already see that \mathfrak{A} needs to be defined on a monoid with the set X as elements. Using the functor $G^{(\mathcal{M}, N, S)}$ of Lemma 3.16, we can choose

$$\mathfrak{A} \stackrel{\text{def}}{=} (G^{(\mathcal{M}, N, S)}(\mathcal{X}), N \otimes r, S, \dot{\gamma}') \in \mathbf{classical},$$

where

$$\dot{\gamma}' \stackrel{\text{def}}{=} \lambda k . \dot{\gamma}(k \circ \psi_{\mathcal{X}})$$

(here $\psi_{\mathcal{X}}: G^{(\mathcal{M}, N, S)}(\mathcal{X}) \xrightarrow{\sim} \mathcal{M}$ is the isomorphism from the proof of Lemma 3.16 and the type of k is $(N \otimes r) \rightarrow S$). The F -image of \mathfrak{A} is

$$F(\mathfrak{A}) = (X, \langle \bullet_X, \gamma' \rangle, \dot{\gamma}', r) \in \mathbf{CCA}_{\cup}^{(G(\mathcal{X}), N \otimes r, S)}$$

where

$$\gamma'_x = \lambda k . \dot{\gamma}'(k) = \lambda k . \dot{\gamma}(k \circ \psi_{\mathcal{X}}). \quad (23)$$

It remains to show that $F(\mathfrak{A})$ is isomorphic to \mathcal{X} . The coalgebras $F(\mathfrak{A})$ and \mathcal{X} are defined on different signatures but with the same state set, so the only possible isomorphism in one of the extra arrows in \mathbf{CCA}_{\cup}^+ (Definition 3.13) that is defined by a monoid isomorphism. And indeed, the monoid isomorphism $\psi_{\mathcal{X}}$ of the proof of Lemma 3.16 satisfies the conditions listed in Definition 3.13:

1. The first condition is immediate.
2. The second condition requires showing that $x \bullet_X r = \psi_{\mathcal{X}}(x) \otimes r$ for all $x \in X$. Recall from the proof of Lemma 3.16 that we can write $x = m \otimes r$ for a unique $m \in \mathcal{M}$, and observe that $r = \mathbf{1} \otimes r$. Hence we can indeed compute that:

$$\begin{aligned} x \bullet_X r &= (m \otimes x) \bullet_X (\mathbf{1} \otimes r) \\ &= (m \bullet \mathbf{1}) \otimes r && // \text{Definition } \bullet_X \text{ (see the proof of Lemma 3.16).} \\ &= m \otimes r \\ &= \psi_{\mathcal{X}}(m \otimes r) \otimes r && // \text{Definition } \psi_{\mathcal{X}} \text{ (see the proof of Lemma 3.16).} \\ &= \psi_{\mathcal{X}}(x) \otimes r. \end{aligned}$$

3. The last condition is exactly Eq. (23).

Hence we indeed have an isomorphism $(\psi_{\mathcal{X}})_{F(\mathfrak{A}), \mathcal{X}}: F(\mathfrak{A}) \xrightarrow{\sim} \mathcal{X}$ by Definition 3.13. \square

We can also define global rules for coalgebras, in such a way that the global rules for coalgebras in \mathbf{CCA}_{\cup} also correspond to the classical global rules.

Definition 3.18: Global rule (coalgebraic)

Let (\mathcal{M}, N, S) be a signature and let $\mathcal{X} = (X, \langle \chi, \gamma \rangle) \in \mathbf{Coalg}(C_{\cup}^{(\mathcal{M}, S)})$. A *configuration* of \mathcal{X} is a function $c: X \rightarrow S$. The *global rule* of \mathcal{X} is a map between configurations defined as

$$\begin{aligned} G_{\mathcal{X}}: S^X &\rightarrow S^X \\ G_{\mathcal{X}}(c)(x) &\stackrel{\text{def}}{=} \gamma_x(\lambda m \in \mathcal{M} . c(m \otimes x)). \end{aligned} \quad (24)$$

Note that for coalgebras $\mathcal{X} = (X, \langle \chi, \gamma \rangle, \overset{\circ}{\gamma}, r) \in \mathbf{CCA}_{\mathbf{U}}^{(\mathcal{M}, N, S)}$ this definition evaluates to:

$$G_{\mathcal{X}}(c)(x) = \overset{\circ}{\gamma}(\lambda n \in N . c(n \otimes x)).$$

Every initial configuration $c_0: X \rightarrow S$ induces a *trace sequence* $(c_i)_{i \in \mathbb{N}}$ defined as $c_i \stackrel{\text{def}}{=} G_{\mathcal{X}}^i(c_0)$.

The functor F of Theorem 3.17 preserves global rules exactly:

Lemma 3.19

Let $\mathfrak{A} = (\mathcal{M}, N, S, \overset{\circ}{\gamma})$ be a classical CA (as in Definition 2.13), and let $F(\mathfrak{A}) = (\mathcal{M}, \langle \chi, \gamma \rangle, \overset{\circ}{\gamma}, \mathbf{1})$ be the output of the functor constructed in Theorem 3.17. Then: $G_{\mathfrak{A}} = G_{F(\mathfrak{A})}$.

Proof. Take any arbitrary configuration $c: \mathcal{M} \rightarrow S$ and any cell $m \in \mathcal{M}$. Then the classical definition of a global rule (Definition 2.16) shows that

$$G_{\mathfrak{A}}(c)(m) \stackrel{\text{def}}{=} \overset{\circ}{\gamma}(\lambda n . c(n \bullet m)).$$

The coalgebraic global rule (Definition 3.18) of $F(\mathfrak{A})$ is

$$G_{F(\mathfrak{A})}(c)(m) \stackrel{\text{def}}{=} \overset{\circ}{\gamma}(\lambda n . c(n \otimes m)),$$

but by definition of F (see the proof of Theorem 3.17) it holds that $n \otimes m = n \bullet m$, thus the global rules are indeed equal. \square

3.4 Modelling irregular and non-uniform CA

We finish this section by providing arguments for further refining the functor $C_{\mathbf{U}}$ in the next section in order to better model CA whose cells have distinct local behaviour.

There are two ways in which the local behaviour of cells can be different: (1) they have a different neighbourhood structure, and (2) they can have different local rules. To avoid confusion, I will call the former condition *irregularity* and the latter *non-uniformity*.

Definition 3.20: Irregular CA

A CA is *irregular* if it has two cells x and x' and two paths in the neighbourhood $n_1, n_2 \in N$ such that n_1 and n_2 lead to different cells when traversing them from initial cell x , but to the same cell when traversed from x' .

In case the CA is given as a coalgebra $\mathcal{X} = (X, \langle \chi, \gamma \rangle) \in \mathbf{CCA}_{\mathbf{U}}^{(\mathcal{M}, N, S)}$, then this means that $n_1 \otimes x \neq n_2 \otimes x$ but $n_1 \otimes x' = n_2 \otimes x'$.

Example 3.21 (An irregular CA). Let $\mathcal{X} = (\mathbb{Z}, \langle \chi, \gamma \rangle) \in \mathbf{CCA}_{\mathbf{U}}^{(\mathbb{Z}^2, N, S)}$ be the CA with:

- as underlying monoid $(\mathbb{Z}^2, +, (0, 0))$,
- neighbourhood $N = \{(i, j) \mid -1 \leq i, j \leq 1\}$,
- and with the following cowriter comonad coalgebra χ :

$$(a, b) \otimes x \stackrel{\text{def}}{=} a + x.$$

Then \mathcal{X} is irregular, since $(0, 1), (0, -1) \in N$ but for any $x \in \mathbb{Z}$ it holds that $(0, 1) \otimes x = x = (0, -1) \otimes x$.

Irregular CA always have *incomplete* cells:

Definition 3.22: Incomplete cell

An *incomplete* cell is a cell that does not have a distinct neighbour for every $n \in N$.

Remark 3.23. The existence of incomplete cells does not imply irregularity: it is possible that $n_1, n_2 \in N$ lead to the same neighbour for *all* cells in a CA.

Definition 3.24: Non-uniform CA

A CA with neighbourhood N is *non-uniform* if there exist two cells x and x' such that the local rule $\gamma_x: S^N \rightarrow S$ used to compute the next state of x in a trace sequence is different from the local rule $\gamma_{x'}: S^N \rightarrow S$ used at x' .

Remark 3.25. We defined *non-uniformity* in such a way that the neighbourhood N is still the same for all cells. This is not restrictive, since one can model a CA where different cells have different neighbourhoods as an *irregular* CA, and take N to be the union of all neighbourhoods.

Both classical definition of a CA (Definition 2.13) and \mathbf{CCA}_{\cup}^+ allow irregular CA but not non-uniform CA. In particular, all cells in a coalgebra in \mathbf{CCA}_{\cup}^+ use the same local rule $\hat{\gamma}: S^N \rightarrow S$. However, this does *not* imply that all cells behave identically. Furthermore, a uniform CA may have an incomplete cell x , not all inputs of the type S^N to the local rule at x can occur in the computation of a trace sequence. This is because configurations assign states to cells and not to paths. The next example gives a concrete instance of this phenomenon.

Example 3.26. Consider the monoid $\mathcal{M} = (\{\mathbf{1}, A, B\}, \bullet, \varepsilon)$ where

$$A \bullet A = A \quad B \bullet B = B \quad A \bullet B = B \bullet A = \mathbf{1}.$$

Define the neighbourhood $N = \{\mathbf{1}, A, B\}$ and the CA $\mathcal{X} = (X, \langle \chi, \gamma \rangle, \hat{\gamma}, r) \in \mathbf{CCA}_{\cup}^{(\mathcal{M}, N, S)}$ where $X = \{r, x, y\}$ as in Figure 3.




Figure 3: An irregular CA $\mathcal{X} = (\{r, x, y\}, \langle \chi, \gamma \rangle)$ defined on the monoid \mathcal{M} of Example 3.26. Arrows show the action of χ ; arrows for the monoid identity element have been omitted. Under the assumption that γ is well-behaved, we see that \mathcal{X} is in \mathbf{CCA}_{\cup}^+ , since it satisfies the conditions of Definition 3.10: it is rooted in r , and $(-)\otimes r$ is an injection since r has 3 distinct neighbours (including itself) while \mathcal{M} has only 3 elements.

Observe that x has only a B -neighbour that is not itself, and y only an A -neighbour. In any configuration $c_0: X \rightarrow S$, it will always hold that $c_0(x) = c_0(A \otimes x)$ and $c_0(y) = c_0(B \otimes y)$. When using the global rule $G_{\mathcal{X}}: S^X \rightarrow S^X$ to compute the next state for x in a trace, the inputs $k: N \rightarrow S$ to the local rule $\hat{\gamma}$ are always of the form $\lambda n \cdot c_0(n \otimes x)$. These will therefore never have $k(\mathbf{1}) \neq k(A)$, but $k(\mathbf{1}) \neq k(B)$ is possible. Similarly for y , the inputs will never have $k(\mathbf{1}) \neq k(B)$, while $k(\mathbf{1}) \neq k(A)$ is possible. Thus only a strict subset of the domain S^N of $\hat{\gamma}$ will be used at x , which means that the local rule at x contains superfluous data. Even stronger, the reachable part of the domain S^N of $\hat{\gamma}$ at y is different from that at x , and the intersection are only the constant inputs. That is, x and y behave almost independently, except that they behave the same in the case where their only neighbour has the same state as themselves.

The next section will show how to construct CA such that the local rule at every cell does not contain superfluous data. We will adapt the type of the local rule at a cell x to S^{N_x} , where N_x has only one element for every *distinct* neighbour of x . A side effect of allowing different cells to have different local rules is that it becomes possible to formulate constraints on how the local rules should relate to each other, and the next section will discuss two choices of such constraints in depth.

We only required CA in \mathbf{CCA}_U^+ to be rooted in order to prove Theorem 3.17: otherwise the functor $\text{classical} \rightarrow \mathbf{CCA}_U^+$ would not be essentially surjective. The latter worry is irrelevant when allowing non-uniform CA, since this implies intentionally including coalgebras not corresponding to classical CA as described in Definition 2.13. Note that the root is also unrelated to the local behaviour of a CA.

3.5 Summary of uniform CA

We have described a coalgebraic model of CA, that exploits the idea that cells and connections between cells can be separate mathematical objects. Concretely, we first modelled a CA's graph structure as comonad coalgebras of the cowriter comonad P (Definition 3.1), parametrised by a monoid of “paths” between cells. We interpreted the carrier sets of those coalgebras as the set of cells. To endow coalgebras with local rules, we further modified the cowriter comonad P to a “uniform CA” functor C_U . From the coalgebras of C_U we constructed the category \mathbf{CCA}_U^+ , whose objects correspond to the classical, uniform CA (Definition 2.13) as coalgebras: Theorem 3.17 shows that there exists an essentially surjective functor from the classical (uniform) CA to \mathbf{CCA}_U^+ . Finally, we have discussed irregular and non-uniform CA, and sketched why we will refine C_U in the next section in order to model such CA more naturally.

4 General cellular automata

In this section we will refine C_U to a functor C_G , whose coalgebras can also be non-uniform CA. We will define the class of *general CA*, contains not only the class of classical CA but also non-uniform CA and CA without a root. Thereafter we show how the class of general CA and two subclasses thereof can be specified as subcategories of $\text{Coalg}(C_G)$.

Aside from capturing a more general definition of CA, the C_G -coalgebras will also have more appropriately typed local rules for irregular CA. Recall from Section 3.4 (and in particular, Example 3.26) that coalgebras of irregular CA in \mathbf{CCA}_U^+ have redundant data in their local rules, as these local rules take assignments $N \rightarrow S$ of states to paths in the neighbourhood N . In the computation of a trace sequence, states are assigned to cells instead, so not all assignments $N \rightarrow S$ can occur in this computation at an incomplete cell x . The functor C_G instead will ensure, at type level, that local rule of a coalgebra at a cell x takes as input an assignment of one state to each *distinct* neighbour of x .

4.1 Classes of general CA

When allowing cells in a CA to have different local rules, the question arises whether there should be a correspondence between the distinct local rules of distinct cells, if any at all. There is no unique answer to this question, and different applications of CA may require different degrees of uniformity. In particular, we consider the following hierarchy classes of CA with various degrees of uniformity, from most general to most specific:

1. **All general CA:** the largest class allows all cells to have unrelated local rules.
2. **Neighbourhood uniform (general) CA:** cells with the same neighbourhood structure have the same local rule. This will be more clearly defined in Definition 4.8.

3. **Extension uniform (general) CA:** the local rule of a cell is an extension of the local rule of cells whose neighbourhood is smaller. Cells with the same neighbourhood structure have the same local rule. This will be more clearly defined in Definition 4.8.
4. **Uniform general CA:** all cells have essentially the same local rule $\dot{\gamma}: S^N \rightarrow S$, but only with a local modification to the local neighbourhood of the cell. This is simple to implement: given an assignment $k: (N \otimes x) \rightarrow S$ of states to the neighbours of x , one can simply use $\gamma_x(k) = \dot{\gamma}(\lambda n \in N . n \otimes x)$. Hence this class is similar to \mathbf{CCA}_U^+ (Definition 3.10), except for (1) the domain of the local rule at a cell x is adapted the neighbourhood structure of x (2) not requiring the existence of a root r nor injectivity of $(-) \otimes r$.

One may need an infinite amount of data to describe an arbitrary general CA, even when using a finite state set and a finite neighbourhood, since the local rule of every individual cell must be specified separately. However, both in the neighbourhood-uniform and extension-uniform classes, the number of distinct local rules on a neighbourhood N is bounded by $2^{\text{card}(N)}$. Furthermore, the appropriate local rule of a cell does not need to be given individually for every cell: it can be determined locally from observing the neighbourhood structure of the cell.

More classes can be defined, but are not considered in this thesis. For example, there is a class of CA whose local rules to agree on inputs that assign every neighbour the same state.

The next step is to formally define the class of general CA, as well as the subclasses of neighbourhood-uniform CA and extension-uniform CA. Before we can give the definitions, we first need to clarify the notion of neighbourhood structures and how they can be compared. After defining these three classes of CA, we will continue this section with showing how each can be modelled as a subcategory of the coalgebras of a “general CA” functor C_G , which is a modified version of C_U . To this end, we will describe how the relations that neighbourhood-uniformity and extension-uniformity enforce between local rules of a coalgebra can be formulated as naturality conditions on the local rule component γ of the coalgebra’s dynamics.

Neighbourhood quotients The definition of a general CA takes a more abstract view on the neighbourhood of a cell x , not depending on the neighbours $x' \in N \otimes x$ as cells, but on rather on the set N_x of sets of paths from x to distinct neighbours. This perspective makes it possible to compare different cells that have essentially the same neighbourhood lay-out, even when the actual the neighbours (as elements of X) are distinct.

Definition 4.1: Neighbourhood quotients

For a cell $x \in X$, define the *neighbourhood quotient* $N_x \stackrel{\text{def}}{=} N / \sim_x$ where $\sim_x \subseteq N \times N$ is the relation defined for all $n, m \in N$ as:

$$(n \sim_x m) \Leftrightarrow (n \otimes x = m \otimes x). \quad (25)$$

We denote the equivalence class (with respect to x) of a path $n \in N$ by $|n|_x \in N_x$.

We typically use the abbreviation *neighbourhood* for *neighbourhood quotient*.

Observe that N_x is always a partitioning of N (Definition 2.1). Consequently, the poset structure on partitionings (Definition 2.2) provides a means of comparing neighbourhoods:

Definition 4.2: Compatible neighbourhood-configuration

Let $N_x, N_{x'} \in \text{PartOrd}(N)$ be the neighbourhoods of cells $x, y \in X$. Then N_x is *N_y -compatible* if $N_y \leq N_x$.

A neighbourhood configuration $z: N \rightarrow S$ is N_x -compatible if

$$\forall_{n,m \in N} [|n|_x = |m|_x \Rightarrow z(n) = z(m)]. \quad (26)$$

Intuitively: z is N_x -compatible if for all paths n and m (starting from x) that lead to the same cell y , then z assigns these paths the same state.

Remark 4.3. The condition $N_{x'} \leq N_x$ implies

$$\bigvee_{n,m \in N} [|n|_x = |m|_x \Rightarrow |n|_{x'} = |m|_{x'}].$$

It follows that any $N_{x'}$ -compatible $z: N \rightarrow S$ is also N_x compatible when $N_{x'} \leq N_x$ (since we would have $|n|_x = |m|_x \Rightarrow |n|_{x'} = |m|_{x'} \Rightarrow z(n) = z(m)$).

The main purpose of the concept of compatibility is allowing us to “downscale” neighbourhood configurations $N \rightarrow S$ into the shape of the neighbourhood N_x of a particular cell x :

Definition 4.4: Downscaling

Given an N_x compatible neighbourhood configuration $z: N \rightarrow S$, define the x -downscaled function $z \setminus x: N_x \rightarrow S$ as:

$$(z \setminus x)(|n|_x) \stackrel{\text{def}}{=} z(n).$$

Note that this definition is well-defined (independent of the choice of the representative of $|n|_x$): this follows directly from the definition of N_x -compatibility (Definition 4.2).

Example 4.5 (Irregular neighbourhood structures). The CA of Figure 4 has four cells each with a different neighbourhood:

- For x , the three neighbourhood paths in $N = \{\varepsilon, L, D\}$ all lead to a different cell. So $N_x = \{|\varepsilon|, |L|, |D|\}$.
- For y , both ε and L lead back to itself, but D leads to z . Thus $N_y = \{|\varepsilon|, |L|, |D|\}$.
- For w , both ε and D lead back to itself, but L to z . Therefore $N_w = \{|\varepsilon|, |D|, |L|\}$.
- For z , all paths in N lead back to itself. So $N_z = \{N\}$.

(here we denoted entire equivalence classes as a set delimited with $|(-)|$, a convention we will continue to use when explicitly giving neighbourhood-quotient equivalence classes). In particular, note that y and w have the same number of neighbours, but still a different neighbourhood structure. In fact, N_y and N_w are incomparable in $\text{PartOrd}(N)$. However, we do have the relations $N_z \leq N_x$, $N_y \leq N_x$, $N_w \leq N_x$, $N_z \leq N_y$ and $N_z \leq N_w$. Note that all neighbourhoods are N_z -compatible.




Figure 4: Example of a non-uniform CA with four cells $X = \{w, x, y, z\}$ on the monoid freely generated over the alphabet $\{L, D\}$ and neighbourhood $N = \{\varepsilon, L, D\}$, where ε (the empty string) is the monoid identity. Note that all cells have a different neighbourhood structure.

Suppose we have state set $S = \{\bullet, \circ\}$ and a neighbourhood configuration $k: N \rightarrow S$ such that $k(\varepsilon) = k(L) = \circ$ and $k(D) = \bullet$. Then k is compatible with N_x and N_y , but not with N_z or N_w .

Classes of generalised CA The definition of a general CA is very similar to that of a classical CA (Definition 2.13), except that it assigns every cell its own local rule (allowing non-uniformity) and allows CA without a root. The subclasses of neighbourhood-uniform and extension-uniform CA follow the same definition, but with additional constraints on the local rule. The presented form of the definition of a general CA is designed to facilitate modelling it coalgebraically:

Definition 4.6: General CA

Static structure

A general CA is a tuple $\mathbb{X} = (\mathcal{M}, N, S, X, \chi, \gamma)$ where:

- (\mathcal{M}, N, S) is a signature (Definition 3.7).
- $X \in \mathbf{Set}$ is a set of *cells*.
- (X, χ) is a cwriter-comonad coalgebra (see Definition 3.1 and 3.2). Equivalently, $\chi: X \rightarrow X^{\mathcal{M}}$ is an action of \mathcal{M} that satisfies Eq. (11) and (12).
- γ is a family of maps $\gamma_x: S^{N_x} \rightarrow S$ indexed by $x \in X$ (called the *local rule*).

Dynamics

A *configuration* of \mathbb{X} is an assignment $c: X \rightarrow S$. Given an *initial configuration* c_0 , the *trace sequence* of \mathbb{X} on c_0 is the sequence of configurations $(c_i)_{i \in \mathbb{N}}$, where for every $i \in \mathbb{N}$, c_i is computed inductively as:

$$c_{t+1} \stackrel{\text{def}}{=} G_{\mathbb{X}}^{t+1}(c_0). \quad (27)$$

Here the *global rule* $G_{\mathbb{X}}: S^X \rightarrow S^X$ is defined as:

$$G_{\mathbb{X}}(c_t)(x) \stackrel{\text{def}}{=} \gamma_x(\lambda |n|_x \in N_x . c_t(n \otimes x)). \quad (28)$$

The class of neighbourhood-uniform CA is the subclass of general CA (Definition 4.6) of CA with the following property: cells with the same neighbourhood use the same local rule. The subclass of extension-uniform CA only contains general CA with the following property: the local rules of cells with large neighbourhoods “extend” the behaviour of cells with smaller neighbourhoods. Extensions are defined as follows:

Definition 4.7: Extension

Let X and Y be sets of cells, and let $\gamma: X \rightarrow S^{S^N}$ and $\delta: Y \rightarrow S^{S^N}$ be local rules on the same signature (\mathcal{M}, N, S) . A component $\gamma_x: S^{N_x} \rightarrow S$ at a cell $x \in X$ *extends* the component $\delta_y: S^{N_y} \rightarrow S$ at a cell y if $N_y \leq N_x$ and for all N_y -compatible $z: N \rightarrow S$ it holds that

$$\gamma_x(z \parallel x) = \delta_y(z \parallel y).$$

Definition 4.8: Subclasses of general CA

Let $\mathbb{X} = (\mathcal{M}, N, S, X, \chi, \gamma)$ be a general CA.

Then \mathbb{X} is:

- *neighbourhood-uniform* if it has the property that $\gamma_x = \gamma_{x'}$ when $N_x = N_{x'}$ (for all $x, x' \in X$).
- *extension-uniform* if if for all $x, x' \in X$ with $N_{x'} \leq N_x$ it holds that γ_x extends $\gamma_{x'}$ (according to Definition 4.7).

Remark 4.9. The class of extension-universal is a subclass of the class of neighbourhood-universal CA, which can be seen as follows. Suppose that \mathbb{X} is extension-uniform, then \mathbb{X} is also neighbourhood-uniform if and only if for all $x, x' \in X$ such that $N_x = N_{x'}$ it holds that $\gamma_x = \gamma_{x'}$. But if $N_x = N_{x'}$, then also $N_{x'} \leq N_x$, thus by (4.7) we find

$$\gamma_x(z \parallel x) = \gamma_{x'}(z \parallel x') \quad (29)$$

for all $z: N \rightarrow S$ that are $N_{x'}$ -compatible. But z is $N_{x'}$ -compatible if and only if it is N_x -compatible, so all inputs to γ_x and $\gamma_{x'}$ are of the form $z \parallel x$ and $z \parallel x'$, respectively, for N_x -compatible $z: N \rightarrow S$. Therefore (29) fully specifies γ_x and $\gamma_{x'}$, and we indeed conclude that $\gamma_x = \gamma_{x'}$.

The next two examples illustrate the consequences of neighbourhood- and extension-uniformity on two specific CA graphs.

Example 4.10 (Forktree). Define the “forktree” \mathcal{F} as the tree where the branching factor of nodes at an even depth is 2, but 1 for nodes at an odd depth. More precisely, \mathcal{F} is the monoid generated (Definition 2.8) over the alphabet $\{L, R\}$ and the following equalities between elements:

$$sL = sR \iff \bigcap_{n \in \mathbb{N}} [s \in \{L, R\}^{2n+1}], \quad (30)$$

that is, the L - and R -successors of odd-length strings are the same element.

See Figure 5 for a visualisation of \mathcal{F} .




Figure 5: The forktree monoid. The top point is the empty string ε . Labels of arrows denote the neighbourhood element that need to be right-multiplied with the tail point to obtain the head point; the notation “ L/R ” is used in case multiplication with L and R lead to the same result. Arrows with label ε (which are all self-loops) have been omitted. Observe that some points have the same L - and R -successor, while other points have two distinct successors. Only the generating elements $\{L, R\}$ are shown as arrows, composite arrows are omitted.

In this example we consider general CA of the form $\mathbb{X} = (\mathcal{F}, N, S, X, \langle \chi, \gamma \rangle)$ whose graph is exactly the forktree, and whose neighbourhood is the direct-successor relation:

- $N = \{\varepsilon, L, R\}$.
- $X = \mathcal{F}$.

- \otimes is the forktree's multiplication.

Observe that this CA is irregular: there are two classes of cells based on their neighbourhood quotient. The first class $X_1 \subseteq X$ are the cells at an odd depth in the tree (containing cells such as R, L, RRR, LLL , etc.). By the equalities of the forktree Eq. (30), these cells have branching factor 1 and therefore have the same L - and R -successor. Their neighbourhood quotient is:

$$N_R = \{|\varepsilon|, |L, R|\}.$$

The other class, X_2 , consists of cells at even depth with branching factor 2: they have distinct L - and R -successors. These cells have neighbourhood quotient

$$N_\varepsilon = \{|\varepsilon|, |L|, |R|\} \cong N.$$

This class contains cells such as $\varepsilon, LL, RR, LLRR$, etc.

Neighbourhood-uniformity

If we want \mathbb{X} to be *neighbourhood-uniform*, then all cells in X_1 need to have the same local rule, and all cells in X_2 need to have the same local rule as well. Thus we would be free to independently choose two local rules $\gamma_1, \gamma_2: X \rightarrow S^{S^N}$ and set, for every $x \in X$:

$$\gamma_x = \begin{cases} \gamma_1 & x \in X_1 \\ \gamma_2 & x \in X_2 \end{cases}$$

For example, we may take $S = \{\bullet, \circ\}$ and define γ the following rule set:




Figure 6: Example definitions for γ_1 (left two rules) and γ_2 (right four rules). The \otimes denotes that a copy of the rule holds for every possible value at that cell.

An example of a possible trace sequence with this local rule is the following:




Figure 7: Example trace sequence on the forktree graph \mathcal{F} using the local rules of Figure 6 and an arbitrary initial configuration $\mathcal{F} \rightarrow \{\bullet, \circ\}$ (the leftmost configuration). Cells that are not drawn are all in state \circ .

Extension-uniformity

If we instead want \mathbb{X} to be *extension-uniform*, then we need to take the fact that $N_R \leq N_\varepsilon$ into account.

This means that we can choose γ_1 freely, but γ_2 must extend γ_1 : for every neighbourhood configurations $z: N \rightarrow S$ such that $z(L) = z(R)$, we need $\gamma_2(z \setminus \varepsilon) = \gamma_1(z \setminus R)$. The outputs of γ_2 on inputs with $z(L) = z(R)$ can still be chosen independently of γ_1 .

The local rule γ presented above in Figure 6 is not extension-uniform. In particular, the following combination of local rules violates the extension property, since the right neighbourhood configuration $N_\varepsilon \rightarrow S$ is an “upscaled” version of the left configuration $N_R \rightarrow S$, and should therefore give the same output (which is \circ):



Example 4.11 (Three-forktree). Define the three-forktree monoid $3\mathcal{F}$ as the monoid generated over the alphabet $\{L, M, R\}$ and with equalities:

$$\begin{aligned} s_1 M s_2 R &= s_1 M s_2 M \\ s_1 M s_2 L &= s_1 M s_2 M \\ s_1 R s_2 L &= s_1 R s_2 M \\ s_1 L s_2 R &= s_1 L s_2 M \end{aligned}$$

for all $s_1, s_2 \in \{L, M, R\}^*$. Figure 8 sketches this monoid.




Figure 8: Visualisation of (a part of) $3\mathcal{F}$ with the points $a = \varepsilon$, $x = L$, $y = RR$ and $w = M$ annotated.

As in Example 4.10, we consider CA $\mathbb{X} = (\mathcal{F}, N, S, X, \langle \chi, \gamma \rangle)$ whose graph structure correspond exactly to the monoid structure of $3\mathcal{F}$, and whose neighbourhood is the direct-successor relation:

- The neighbourhood is $N = \{\varepsilon, L, M, R\}$.
- $X = 3\mathcal{F}$.
- \otimes is $3\mathcal{F}$'s multiplication.

Now consider the following three neighbourhood configurations $z_i: N \rightarrow S$ with $N \stackrel{\text{def}}{=} \{\varepsilon, L, M, R\}$ and $S \stackrel{\text{def}}{=} \{s_0, s_1, s_2, s_3\}$:




Figure 9: Three neighbourhood configurations $z_1, z_2, z_3: N \rightarrow S$ defined on $3\mathcal{F}$ (see Figure 8) with neighbourhood $N \stackrel{\text{def}}{=} \{\varepsilon, L, M, R\}$ and state-set $S \stackrel{\text{def}}{=} \{s_0, s_1, s_2, s_3\}$. All three map ε to s_0 , as indicated by the central nodes.

From Figure 8 it becomes clear that there are four classes of cells with the same neighbourhood quotient:

- $X_\varepsilon = \{\varepsilon\}$: ε is the unique cell with neighbourhood $N_\varepsilon \cong N$.
- X_L : the cells $\{L, LL, LLL, \dots\}$ with neighbourhood $N_L = \{|\varepsilon|, |L|, |M, R|\}$.
- X_R : the cells $\{R, RR, RRR, \dots\}$ with neighbourhood $N_R = \{|\varepsilon|, |L, M|, |R|\}$.
- X_M : all other cells (so this includes M, MM, RM, LMM , etc.). These all have neighbourhood $N_M = \{|\varepsilon|, |L, M, R|\}$.

Note that $N_L \neq N_R$, despite being isomorphic as sets. Even stronger, $N_L \not\leq N_R$ and $N_R \not\leq N_L$: the neighbourhood quotients are incomparable in $\text{PartOrd}(N)$.

Neighbourhood-uniformity

If we want \mathbb{X} to be *neighbourhood-uniform*, then we can independently define a local rule γ_i for every choice of $i \in \{\varepsilon, L, R, M\}$, with the requirement that $\gamma_x = \gamma_i$ whenever $x \in X_i$.

Extension-uniformity

If instead we want \mathbb{X} to be *extension-uniform*, then we need to take into account that

$$N_M \leq N_L \leq N_\varepsilon$$

and

$$N_M \leq N_R \leq N_\varepsilon.$$

Therefore γ_L and γ_R only need to extend γ_M and not need to extend each other (since N_L and N_R are incomparable). However, γ_ε needs to extend γ_L and γ_R (and γ_M , but any extension of γ_L or γ_R always is an extension of γ_M).

One may wonder if the extension requirement of γ_ε restricts the choices of γ_L and γ_R : after all, γ_ε cannot be an extension of two disagreeing local rules at the same time. The answer is negative: we can choose γ_L independently from γ_R , this will not cause problems when choosing γ_ε . To see this, we observe that any $z: N \rightarrow S$ that is N_L -compatible but not N_M -compatible assigns L and M different states, and is therefore not N_R -compatible. Similarly, if z is N_R -compatible but not N_M -compatible, then z is not N_L compatible. So the only inputs on which both γ_L and γ_R are defined are N_M -compatible, on which γ_L and γ_R agree.

4.2 General CA functor

We will now define the “general CA” functor C_G , which is a modification of C_U (Definition 3.5) that ensures that the local rule components of coalgebras adapt to the neighbourhoods of cells. The idea is simple: where the γ component of a C_U -coalgebra had type $\gamma_x: S^{\mathcal{M}} \rightarrow S$ for every cell $x \in X$, the γ component of a C_G -coalgebra instead has type $\gamma_x: S^{N_x} \rightarrow S$. Hence it takes inputs of type $N_x \rightarrow S$ that assign exactly one state to every distinct neighbour of x .

Coalgebraic models of general CA arise as a full subcategory \mathbf{CCA}_G of $\mathbf{Coalg}(C_G)$. Since neighbourhood-uniform CA and extension-uniform CA are subclasses of general CA, we in turn construct coalgebraic models of these subclasses as subcategories of \mathbf{CCA}_G .

The following definition defines C_G . It is not directly obvious that it satisfies the functor laws, but this will be proven in Lemma 4.15.

Definition 4.12

For every signature (\mathcal{M}, N, S) , the functor $C_G: \mathbf{Set} \rightarrow \mathbf{Set}$ is defined on objects $X \in \mathbf{Set}$ as:

$$C_G(X) \stackrel{\text{def}}{=} \sum_{f \in X^{\mathcal{M}}} S^{S^{\text{Prelm}_N(f)}}. \quad (31)$$

(where \sum denotes the dependent sum of Definition 2.23).

The action of C_G on functions $h: X \rightarrow Y$ is defined as:

$$\begin{aligned} C_G h: \sum_{f \in X^{\mathcal{M}}} S^{S^{\text{Prelm}_N(f)}} &\rightarrow \sum_{f' \in Y^{\mathcal{M}}} S^{S^{\text{Prelm}_N(f')}} \\ C_G h: (f, g) &\mapsto (h \circ f, \lambda k . g(\lambda |n|_f . k(|n|_{h \circ f}))). \end{aligned} \quad (32)$$

where $|n|_f$ gives the equivalence class of $n \in N$ under the relation $n' \sim n''$ iff $f(n') = f(n'')$ (and analogous for $|n|_{h \circ f}$).

Remark 4.13. As for C_U -coalgebras (Definition 3.5), C_G -coalgebras (X, d) have two components $\chi(x) \stackrel{\text{def}}{=} \pi_1(d(x))$ and $\gamma_x \stackrel{\text{def}}{=} \pi_2(d(x))$ for every $x \in X$. We continue the convention of omitting the projections and denote the coalgebra directly as $(X, \langle \chi, \gamma \rangle)$.

Remark 4.14. There are three important insights not directly obvious from the definition of C_G :

1. For a coalgebra $(X, \langle \chi, \gamma \rangle) \in \mathbf{Coalg}(C_G)$ it holds that $\text{Prelm}_N((-) \otimes x) = N_x$ (see Definition 4.1) for every cell $x \in X$. This means that $\gamma_x: S^{N_x} \rightarrow S$, as desired.
2. The element in $\text{preim}_N((-) \otimes x)$ are equivalence classes of the form $|n|_{(-) \otimes x}$ for some $n \in N$. By the previous remark, this class is the same as $|n|_x$.
3. For $h: X \rightarrow Y$, the map $C_G h$ is well-defined only if $|n|_{h \circ f}$ is independent of the choice of the representative of $|n|_f$. This is indeed true: if $|n|_f = |n'|_f$, then $f(n) = f(n')$, which implies $h(f(n)) = h(f(n'))$. The latter implies $|n|_{h \circ f} = |n'|_{h \circ f}$. Note that the type of the bound variable k is $k: \text{Prelm}_N(h \circ f) \rightarrow S$.

Lemma 4.15

The action of C_G (Definition 4.12) on morphisms is functorial.

Proof. Let $X \in \mathbf{Set}$ and $(f, g) \in C_G X = \sum_{f \in X^{\mathcal{M}}} S^{S^{\text{Prelm}_N(f)}}$, then identity preservation is simple to verify:

$$\begin{aligned} C_G \text{Id}_X(f, d) &= (\text{Id}_X \circ f, \lambda k . d(\lambda x . k(\text{Id}_X(x)))) \\ &= (f, \lambda k . d(\lambda x . k(x))) \\ &= (f, \lambda k . d(k)) \\ &= (f, d). \end{aligned}$$

Composition preservation follows from evaluating λ -expressions. Take any pair of functions:

$$X \xrightarrow{h} Y \xrightarrow{\ell} Z$$

and an input $(f, g) \in C_G X = \sum_{f \in X^{\mathcal{M}}} S^{S^{\text{Prelm}_N(f)}}$. Then we compute:

$$\begin{aligned} (C_G \ell \circ C_G h)(f, g) &= C_G \ell(h \circ f, \lambda c \cdot g(\lambda |n|_f \cdot c(|n|_{h \circ f}))) && // \text{ Evaluate } C_G h. \\ &= (\ell \circ h \circ f, \lambda k \cdot (g(\lambda |n|_f \cdot c(|n|_{h \circ f}))) \left(\lambda |n'|_{h \circ f} \cdot k(|n'|_{\ell \circ h \circ f}) \right)) && // \text{ Evaluate } C_G \ell. \\ &= (\ell \circ h \circ f, \lambda k \cdot (g(\lambda |n|_f \cdot (\lambda |n'|_{h \circ f} \cdot k(|n'|_{\ell \circ h \circ f}))) (|n|_{h \circ f}))) && // \text{ Substitute } \lambda |n'|_{h \circ f} \cdot k(|n'|_{\ell \circ h \circ f}) \text{ for } c. \\ &= (\ell \circ h \circ f, \lambda k \cdot (g(\lambda |n|_f \cdot k(|n|_{\ell \circ h \circ f})))) && // \text{ Substitute } |n|_{h \circ f} \text{ for } |n'|_{h \circ f}. \\ &\stackrel{\text{def}}{=} C_G(\ell \circ h)(f, g) && // \text{ Definition } C_G(\ell \circ h). \end{aligned}$$

as required. Note that the bound variables have types

$$c: \text{Prelm}_N(h \circ f) \rightarrow S$$

and

$$k: \text{Prelm}_N(\ell \circ h \circ f) \rightarrow S.$$

□

The type of coalgebras of $C_G^{(\mathcal{M}, N, S)}$ together with Remark 4.14.1 make it clear that the subcategory of $C_G^{(\mathcal{M}, N, S)}$ -coalgebras whose first component is a cowriter-comonad coalgebra are exactly the general CA on signature (\mathcal{M}, N, S) (Definition 4.6). Thus we can define general CA coalgebraically as:

Definition 4.16: Category of general CA coalgebras

For every signature (\mathcal{M}, N, S) , let $\mathbf{CCA}_G^{(\mathcal{M}, N, S)}$ be the subcategory of $\mathbf{Coalg}(C_d^{(\mathcal{M}, N, S)})$ of coalgebras whose first component is a cowriter-comonad coalgebra (Definition 3.1 and 3.2). Let the *category of general CA coalgebras*, denoted \mathbf{CCA}_G , be the coproduct of $\mathbf{CCA}_G^{(\mathcal{M}, N, S)}$ over all signatures (\mathcal{M}, N, S) .

4.3 Neighbourhood- and extension-uniform CA

We claimed that the coalgebras $\mathcal{X} = (X, \langle \chi, \gamma \rangle) \in \mathbf{CCA}_G$ that are neighbourhood-uniform or extension-uniform CA respectively, are exactly the coalgebras whose second projection (the local rule γ) satisfies a respective naturality constraint. We will show this via backward reasoning: first we show that neighbourhood- and extension-uniformity can be stated as an commuting diagram involving γ , and thereafter we construct the categories and functors that show that this diagram is a naturality condition.

In the remainder of this section, we fix a general CA $\mathcal{X} = (X, \langle \chi, \gamma \rangle) \in \mathbf{CCA}_G$ on a signature (\mathcal{M}, N, S) .

Extension-uniformity First consider extension-uniformity (Definition 4.8), which requires that for any pair of cells $x, x' \in X$ with $N_{x'} \leq N_x$ it holds that γ_x *extends* $\gamma_{x'}$ (extensions were defined in Definition 4.7, and \leq is from the poset $\text{PartOrd}(N)$ of Definition 2.2). That is, for any $N_{x'}$ -compatible $z: N \rightarrow S$ it holds that

$$\gamma_x(z \parallel x) = \gamma_{x'}(z \parallel x'),$$

where $(z \ll x) \in S^{N_x}$ and $(z \ll x') \in S^{N_{x'}}$. We can construct a mapping $S^{N_{x'}} \rightarrow S^{N_x}$ that sends $z \ll x'$ to $z \ll x$ as follows:

$$\begin{aligned} Q_{x,x'}: S^{N_{x'}} &\rightarrow S^{N_x} \\ Q_{x,x'}(k) &= \lambda |n|_x \cdot k(|n|_{x'}). \end{aligned}$$

This mapping is well-defined (independent of the choice of representative of $|n|_x$), this follows from the definition of \leq on $N_{x'} \leq N_x$. Now indeed we can check that

$$\begin{aligned} Q_{x,x'}(z \ll x') &\stackrel{\text{def}}{=} \lambda |n|_x \cdot (z \ll x')(|n|_{x'}) && // \text{Definition } Q_{x,x'}. \\ &= \lambda |n|_x \cdot z(n) && // \text{Definition } \ll \text{ (Definition 4.4)}. \\ &= \lambda |n|_x \cdot (z \ll x)(|n|_x) && // \text{Definition } \ll \text{ (Definition 4.4)}. \\ &= z \ll x, \end{aligned}$$

as we desired.

Thus, γ is extension-uniform if, for all $x, x' \in X$ with $N_{x'} \leq N_x$, the following diagram commutes:

$$\begin{array}{ccc} S^{N_{x'}} & \xrightarrow{\gamma_{x'}} & S \\ Q_{x,x'} \downarrow & & \parallel \\ S^{N_x} & \xrightarrow{\gamma_x} & S \end{array} \quad (33)$$

Neighbourhood-uniformity The reasoning for neighbourhood-uniformity is the same, except with the condition $N_{x'} \leq N_x$ replaced by $N_{x'} = N_x$. To see this, recall from the discussion in Remark (4.9) that, under the assumption $N_{x'} = N_x$, the condition $\gamma_x = \gamma_{x'}$ is equivalent to the condition that γ_x extending $\gamma_{x'}$ (which is equivalent to $\gamma_{x'}$ extending γ_x). Thus, γ is neighbourhood-uniform if and only if Diagram (33) commutes for any pair of cells $x, x' \in X$ such that $N_x = N_{x'}$.

4.3.1 The local rule as a natural transformation

The previous paragraphs showed that γ is neighbourhood- and extension-uniform, respectively, if Diagram (33) commutes whenever $N_{x'} \leq N_x$ or $N_{x'} = N_x$, respectively. To show that this is a naturality condition on γ , we define appropriate categories and functors such that Diagram (33) is a naturality square. As domain of the functors we use the following thin category, using the poset structure of $\text{PartOrd}(N)$:

Definition 4.17: Neighbourhood category

The *neighbourhood category* of \mathcal{X} , denoted as $\mathbf{N}_{\mathcal{X}}$, is the category with as objects the neighbourhood quotients N_x of every $x \in X$, and a unique morphism $N_{x'} \rightarrow N_x$ whenever $N_{x'} \leq N_x$.

The *neighbourhood-iso-category* of \mathcal{X} , denoted as $\mathbf{N}_{\mathcal{X}}^{\overline{=}}$, has the same objects as $\mathbf{N}_{\mathcal{X}}$ but only identity morphism.

For extension-uniformity, the first functor is $Q: \mathbf{N}_{\mathcal{X}} \rightarrow \mathbf{Set}$ that sends every object N_x to S^{N_x} . On a morphism $f: N_{x'} \leq N_x$ we define $Qf \stackrel{\text{def}}{=} Q_{x,x'}$, as required by Diagram (33). The second functor is Δ_S that maps every object to the set S , and every morphism to Id_S .

The condition that γ is a natural transformation $\gamma: Q \Rightarrow \Delta_S$ means that, for every morphism $f: N_{x'} \leq N_x$, Diagram (33) commutes. As argued above, this is exactly the extension-uniformity requirement.

As for neighbourhood-uniformity, we can use the functor $Q^= : \mathbf{N}_{\mathcal{X}}^= \rightarrow \mathbf{Set}$, which is defined the same way as Q . The naturality condition $\gamma : Q^= \Rightarrow \Delta_S$ means that Diagram 33 need to commute only for every morphism witnessing $f : N_{x'} = N_x$ (by definition of the morphisms in $\mathbf{N}_{\mathcal{X}}^=$), i.e., whenever $N_{x'} = N_x$. Again, we have already argued that this is exactly the neighbourhood-uniformity condition.

In the remainder of this thesis, we will denote the subcategory of $\mathbf{CCA}_{\mathcal{G}}$ of coalgebras where γ is a natural transformation $Q^= \Rightarrow \Delta_S$ (neighbourhood-uniform) and $Q \Rightarrow \Delta_S$ (extension-uniform) by \mathbf{CCA}_{NU} and \mathbf{CCA}_{EU} respectively.

4.4 Summary of general CAs

This section defined the class of *general CA* (Definition 4.6): a generalisation of CA (in comparison to Definition 2.13) that also allows non-uniform CA and CA without root. We constructed a functor $C_{\mathcal{G}}$ and a category $\mathbf{CCA}_{\mathcal{G}}$ of $C_{\mathcal{G}}$ -coalgebras that are exactly the general CA.

We also defined two subclasses of general CA, namely the *neighbourhood-uniform CA* and the *extension-uniform CA*. These classes consist of CA in which different cells may have different local rules, but where CA with similar neighbourhood structures must have related local rules. We showed that these arise coalgebraically as the subcategories \mathbf{CCA}_{NU} and \mathbf{CCA}_{EU} , respectively, of coalgebras in $\mathbf{CCA}_{\mathcal{G}}$ whose local rule satisfies a naturality condition.

In the remainder of this thesis we will investigate the options of defining a modal logic over coalgebraic CA. We will take the coalgebras in the category $\mathbf{CCA}_{\mathcal{G}}$ as our working definition of “CA”.

5 Coalgebra homomorphisms and behavioural equivalence

An important theoretical tool that can automatically construct modal logic languages for coalgebras is the *predicate lifting* framework developed by Pattinson [30, 31] and Schröder [41] for coalgebras of endofunctors on \mathbf{Set} . See Klin [22] for a treatment of base categories beyond \mathbf{Set} . This framework uses *behavioural equivalence* (BE, Definition 2.21) as the main notion of similarity. In particular, the language that this framework provides ensures that cells $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ satisfy the same formulas iff $x \overset{\text{BE}}{\approx} y$. A consequence that any differences between cells $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ cannot be described by the logic as soon as $x \overset{\text{BE}}{\approx} y$. We will see that this unfortunately means that the framework is inappropriate to describe spacial-temporal patterns in trace sequences of CA, since certain pairs of CA with very different dynamics turn out to be behaviourally equivalent. Phrased differently: to fulfill our research question, a model logic must apparently be able to distinguish behaviourally equivalent coalgebras.

Before we can justify this conclusion, we need to develop a clear overview of what behavioural equivalence means in context of CA. To this end, we will first analyse homomorphisms between CA. This will reveal that coalgebra homomorphisms are characterised two properties: (1) they commute with cowriter-comonad coalgebras and (2) they witness an extension condition similar to the one of the class of extension-uniform CA (Definition 4.7). Thereafter we construct two methods of characterising behaviourally equivalent cells: (1) contracting a CA such that it has no distinct behaviourally equivalent cells anymore, and (2) constructing a terminal coalgebra in $\mathbf{CCA}_{\mathcal{G}}^{\Sigma}$ that has exactly one cell for every equivalence class of behaviourally equivalent cells in $\mathbf{CCA}_{\mathcal{G}}^{\Sigma}$ (for a any signature Σ).

Throughout this section, we assume a fixed signature $\Sigma = (\mathcal{M}, N, S)$. Analogously to $\mathbf{CCA}_{\mathcal{G}}^{\Sigma}$ (Definition 4.16), we define $\mathbf{CCA}_{\text{NU}}^{\Sigma}$ to be the full subcategory of \mathbf{CCA}_{NU} of coalgebras defined of Σ (equivalently, it is the full subcategory of $\mathbf{CCA}_{\mathcal{G}}^{\Sigma}$ of neighbourhood-uniform CA).

5.1 Characterising coalgebra homomorphisms

We will begin the investigation of behavioural equivalence by characterising the properties of coalgebra homomorphisms in \mathbf{CCA}_G^Σ . These homomorphisms have two properties, because the output of the object map of the functor C_G (Definition 4.12) is a Cartesian product between two sets. The first property states that coalgebra homomorphisms commute with the cowriter coalgebras, and the second property states an extension relation between local rules of cells in the domain and image, similar to the extension relation between local rules in extension-uniform CA (Definition 4.8).

In this subsection, we assume two fixed general CA $\mathcal{X} = (X, \langle \chi, \gamma \rangle)$, $\mathcal{Y} = (Y, \langle \xi, \delta \rangle) \in \mathbf{CCA}_G^\Sigma$.

Both properties follow from the coalgebra homomorphism law (Eq. (6)). For a coalgebra homomorphism $h: \mathcal{X} \rightarrow \mathcal{Y}$ this law asserts that the following diagram commutes:

$$\begin{array}{ccc}
 X & \xrightarrow{h} & Y \\
 \langle \chi, \gamma \rangle \downarrow & & \downarrow \langle \xi, \delta \rangle \\
 \sum_{f \in X^{\mathcal{M}}} S^{S^{\text{Prelm}_N(f)}} & \xrightarrow{C_G h} & \sum_{f' \in Y^{\mathcal{M}}} S^{S^{\text{Prelm}_N(f')}}
 \end{array} \tag{34}$$

Both paths from X (upper left) to $\sum_{f' \in Y^{\mathcal{M}}} S^{S^{\text{Prelm}_N(f')}}$ (lower right) in Diagram (34) output a tuple of the form $((-) \otimes y, \delta_y)$ on any input $x \in X$. The diagram commutes for every coalgebra homomorphism h , which means that both the upper-left path and the right-lower path output the same tuple. The upper-left path is the morphism:

$$\langle \chi, \delta \rangle \circ h$$

and the right-lower path is:

$$C_G h \circ \langle \chi, \gamma \rangle = (h \circ ((-) \otimes \chi), \lambda x . \lambda k . \gamma_x(\lambda |n|_x . k(|n|_{h(x)}))$$

by definition of C_G (Definition 4.12), where the type of the bound variable k is $N_{h(x)} \rightarrow S$.

Remark 5.1. The discussion above used the notational simplifications indicated in Remark 4.14. We wrote N_x for $\text{Prelm}_N((-) \otimes x) = \text{Prelm}_N(\chi(x))$, $N_{h(x)}$ for $\text{Prelm}_N(\xi(h(x)))$ and $|n|_x$ for $|n|_{\chi(x)}$. Finally, we also used the following simplification:

$$\begin{aligned}
 |n|_{h(x)} &= |n|_{(-) \otimes h(x)} \\
 &= |n|_{\xi(h(x))} \\
 &= |n|_{(\xi \circ h)(x)}.
 \end{aligned}$$

5.1.1 First property: commutativity with cowriter-comonad coalgebras

We first consider the equality of the first elements of the tuples that the paths in Diagram (34) output. This equality asserts that h is a \mathcal{M} - χ - ξ -homomorphism, defined as follows:

Definition 5.2: \mathcal{M} - χ - ξ -homomorphism

A \mathcal{M} - χ - ξ -homomorphism between the cowriter-comonad coalgebras (X, χ) and (Y, ξ) is a function $h: X \rightarrow Y$ with the following property: for all $x \in X$ and $m \in \mathcal{M}$, it holds that

$$h(m \otimes x) = m \otimes h(x). \tag{35}$$

A \mathcal{M} - χ - ξ -homomorphism h is an \mathcal{M} -restricted surjection (injection/isomorphism respectively) if, for

every $x \in X$, the restriction

$$h \upharpoonright (\mathcal{M} \otimes x): (\mathcal{M} \otimes x) \rightarrow (\mathcal{M} \otimes h(x)) \quad (36)$$

is a surjection (injection/bijection respectively).

Lemma 5.3

Every coalgebra homomorphism $h: \mathcal{X} \rightarrow \mathcal{Y}$ is a \mathcal{M} - χ - ξ -homomorphism.

Proof. The first projections of the composite paths in Diagram (34) gives $\xi \circ h = (h \circ (-)) \circ \chi$, i.e.,

$$\xi \circ h = \lambda x . \lambda m . h(\chi(x)(m)).$$

On inputs $x \in X$ and $m \in \mathcal{M}$ this evaluates, in infix notation, to:

$$m \otimes h(x) = h(m \otimes x).$$

□

An important property of M - χ - ξ -homomorphisms is that they are always \mathcal{M} -restricted surjective.

Lemma 5.4

Every M - χ - ξ -homomorphism $h: (X, \chi) \rightarrow (Y, \xi)$ is an \mathcal{M} -restricted surjection.

Proof. Take any $x \in X$, then we need to show that $\mathcal{M} \otimes h(x)$ is contained in the image of h under $\mathcal{M} \otimes x$. To this end, take an arbitrary $m \in M$, the property of an M - χ - ξ -homomorphism asserts that:

$$h(m \otimes x) = m \otimes h(x),$$

and indeed $(m \otimes x) \in (\mathcal{M} \otimes x)$. □

Remark 5.5. It might still be possible that there exist an $y \in Y$ for which there exist no $m \in \mathcal{M}$ such that $y = m \otimes h(x)$. In this case y may not be in the image of h under $\mathcal{M} \otimes x$.

We provide two examples of coalgebra homomorphisms: the first one is \mathcal{M} -restricted injective, the second not.

Example 5.6 (Coalgebra isomorphism). Recall Example 4.10, which showed that we can construct a neighbourhood-uniform CA on the forktree monoid \mathcal{F} whose graph structure is the same as the forktree. Assume such a neighbourhood-uniform CA $\mathcal{X} \in \mathbf{CCA}_{\text{NU}}^{\Sigma}$, and we will construct a coalgebra homomorphism $\mathcal{X} \rightarrow \mathcal{X}$. Let h be the mapping $h: \mathcal{F} \rightarrow \mathcal{F}$ that appends the suffix LL to every string, that is,

$$h: x \mapsto xLL.$$

It sends the root ε to LL , but the subtree rooted in LL is isomorphic to \mathcal{F} (as seen in Figure 5). Furthermore, it preserves the structure of the forktree: it sends cells to cells with the same neighbourhood. For example, h sends the children of the root L and R to LLL and RLL , respectively, and the only child of L , which is $LR = LL$ to the only child of LLL , which is $LLLL = LRLL$, etc. Therefore h is a M - χ - χ -homomorphism, the first property of a coalgebra homomorphism by Theorem 5.9.

Neighbourhood-uniformity ensures that cells with the same neighbourhood have the same local rule, which implies that the extension property is satisfied. Thus from Theorem 5.9 it follows that h is indeed a coalgebra homomorphism.

Example 5.7 (Non-injective coalgebra homomorphism). Example 4.11 showed that we can construct an extension-uniform CA on the three-forktree monoid $3\mathcal{F}$ whose graph structure is exactly the three-forktree. Let $\mathcal{X} \in \mathbf{CCA}_{\text{EU}}$ be such an extension-uniform CA. We will construct a non-injective coalgebra homomorphism $h: \mathcal{X} \rightarrow \mathcal{X}$.

Recall that all cells whose string starts with the character M have the same neighbourhood N_M , and therefore the same local rule γ_M . Furthermore, we saw that the local rules of all other cells need to extend γ_M . Therefore the extension-property of Theorem 5.9 is always satisfied if we map cells to cells whose strings starts with the character M . Thus we conjecture that we can define h as follows: it sends every string $s \in 3\mathcal{F}$ to Ms' where s' is s but with all occurrences of L and R replaced by M . For example, $h(MRR) = MMMM$, $h(\varepsilon) = M$ and $h(MLRML) = MMMMMM$. Note that this mapping is independent of the choice of representative of equivalence classes under congruence, since by definition of $3\mathcal{F}$, all generating equalities only equate strings of the same length (see Example 4.11, it can also be seen in Figure 8). Put visually, h maps every cell of depth k in the tree to the cell M^{k+1} of depth $k+1$ in the central branch of Figure 8.

It follows that for all monoid elements $m \in 3\mathcal{F}$ and every cell $x \in X = 3\mathcal{F}$ that $h(m \otimes x) = M^{k_m+k_x}$ where k_m and k_x are the lengths of the strings m and x respectively. But we can write $m = s_0s_1s_2 \dots s_{k_m}$ for $s_i \in \{L, M, R\}$ and $\ell \in \mathbb{N}$. Then the generating equalities of $3\mathcal{F}$ and the cowriter-comonad property of \otimes (Eq. (12)) ensure that:

$$\begin{aligned}
 m \otimes x &= s_{k_m} \otimes s_{k_m-1} \otimes \dots \otimes s_1 \otimes s_0 \otimes h(x) \\
 &= s_{k_m} \otimes s_{k_m-1} \otimes \dots \otimes s_1 \otimes s_0 \otimes M^{k_x} \\
 &= M^{k_x} s_0 s_1 \dots s_{k_m-1} s_{k_m} \\
 &= M^{k_x} M s_1 \dots s_{k_m-1} s_{k_m} && // \text{Generating equalities of } 3\mathcal{F}. \\
 &= \dots \\
 &= M^{k_x} M^{k_m-1} s_{k_m} && // \text{Generating equalities of } 3\mathcal{F}. \\
 &= M^{k_x} M^{k_m} && // \text{Generating equalities of } 3\mathcal{F}. \\
 &= M^{k_x+k_m} \\
 &= h(m \otimes x),
 \end{aligned}$$

which shows that h is also a M - χ - χ -homomorphism. Thus h has all properties of a coalgebra homomorphism as required by Theorem 5.9. Now h is not injective, as it collapses the three-forktree into a tree with branching factor 1 (e.g., it maps L and R , which are distinct cells, both to MM).

5.1.2 Second property: extension of local rules

Now we analyse the implication of equality of the second elements of the tuples that the paths in Diagram (34) output. It asserts that the local rule of a cell in the image of a coalgebra homomorphism extends (Definition 4.7) the local rule of the corresponding cell in the domain.

Lemma 5.8

A \mathcal{M} - χ - ξ -homomorphism $h: (X, \chi) \rightarrow (Y, \xi)$ is a coalgebra homomorphism $h: \mathcal{X} \rightarrow \mathcal{Y}$ if and only if it has the property that γ_x extends $\delta_{h(x)}$ for all $x \in X$.

Proof. Unfolding the definition of an extension (Definition 4.7), we need to show that for all $x \in X$ and all $N_{h(x)}$ -compatible $z: N \rightarrow S$ (Definition 4.2) it holds that:

$$\delta_{h(x)}(z \parallel h(x)) = \gamma_x(z \parallel x).$$

The second projection of the output of the paths in Diagram 34 asserts that for all $x \in X$ and $k \in S^{N_{h(x)}}$ it holds that

$$(\delta \circ h)(x)(k) = (C_G h \circ \gamma)(x)(k) \quad (37)$$

But note that $N_{h(x)}$ -compatible elements of S^N are in bijection with $S^{N_{h(x)}}$ (with the bijection given by $(-)\|h(x)$). Thus the above statement is equivalent to the statement that for every $N_{h(x)}$ -compatible $z: N \rightarrow S$ it holds that:

$$(\delta \circ h)(x)(z\|h(x)) = (C_G h \circ \gamma)(x)(z\|h(x)).$$

Simplifying both sides gives

$$\begin{aligned} \delta_{h(x)}(z\|h(x)) &= (C_G h(\gamma_x))(z\|h(x)) \\ &= \gamma_x(\lambda n_x \cdot (z\|h(x))(|n|_{h(x)})) && // \text{Definition } C_G h \text{ (Definition 4.12).} \\ &= \gamma_x(\lambda n_x \cdot z(n)) && // \text{Definition } \|h(x) \text{ (Definition 4.4).} \\ &= \gamma_x(\lambda n_x \cdot (z\|x)(|n|_x)) = \gamma_x(z\|x), && // \text{Definition } \|x \text{ (Definition 4.4).} \end{aligned}$$

as required. \square

We can summarise the characterisation of coalgebra homomorphisms as follows: convenient notation as follows:

Theorem 5.9: Characterisation of CCA_G^Σ -morphisms

The coalgebra homomorphisms $h: \mathcal{X} \rightarrow \mathcal{Y}$ in CCA_G^Σ are exactly the M - χ - ξ -homomorphism with the following extension property: extension property: for all $x \in X$ and $N_{h(x)}$ -compatible $z: N \rightarrow S$ it holds that:

$$\delta_{h(x)}(z\|h(x)) = \gamma_x(z\|x). \quad (38)$$

5.2 Illustration of behavioural equivalence

Theorem 5.9 characterises the coalgebra homomorphisms, but on a very abstract level. We will now show several examples and more specific results to obtain an intuition when cells are behaviourally equivalent.

We begin with a basic example of finding a coalgebra homomorphism (in this case even an isomorphism) via Theorem 5.9:

Example 5.10. Define the data:

$$\begin{aligned} \mathcal{M} &= \mathbb{Z} \\ N &= \{-1, 0, 1\} \\ \mathcal{X} &= \{\mathbb{Z}^+ = \{m^+\}_{m \in \mathbb{Z}}, \langle \chi, \gamma^+ \rangle\} \\ &\quad \text{with } n \otimes m \stackrel{\text{def}}{=} n + m \\ \mathcal{Z} &= \{\mathbb{Z}^- = \{m^-\}_{m \in \mathbb{Z}}, \langle \xi, \gamma^- \rangle\} \\ &\quad \text{with } n \odot m \stackrel{\text{def}}{=} m - n \end{aligned}$$

where $\gamma_{n^+}^+ = \gamma_{n^-}^-$ is an arbitrary local rule on an arbitrary state set S . Note that all cells have $N_{n^+} = N_{n^-} \cong N$, so with a minor abuse of notation we will write all neighbourhood quotients as N .

Now define $h: \mathbb{Z}^+ \rightarrow \mathbb{Z}^-$ as negation: $h(n) \stackrel{\text{def}}{=} -n$. We can verify that h is a coalgebra homomorphism by checking it has the two required properties (see Theorem 5.9).

To check that h is a \mathcal{M} - χ - ξ -homomorphism, observe that:

$$h(n \otimes m) = h(n + m) = -(n + m) = -n - m = (-m) - n = n \otimes h(m).$$

To check that h satisfies the extension relation Eq. (38), take any $z: N \rightarrow S$, observe that:

$$(z \setminus h(n^+))(|m|_{h(n^+)}) = z(m) = (z \setminus n^+)(|m|_{n^+})$$

thus when $\gamma_{h(n^+)}^-$ queries $m \in N$ to the input configuration $z \setminus h(n^+)$, it obtains the same state in S as when $\gamma_{n^+}^+$ queries m to its own input $z \setminus n^+$. Since $\gamma_{n^+}^+ = \gamma_{n^-}$, they must therefore give the same outputs on those inputs.

Note that the m -neighbour of $h(n^+)$ is $m \otimes -n^- = -n^- - m$ whereas the m -neighbour of n^+ is $n^+ + m$. This does not matter for the extension relation, which only depends on whether paths in N lead to distinct cells.

Lemma 5.11

Every cell in a uniform CA $\mathcal{X} = (X, \langle \chi, \gamma \rangle) \in \mathbf{CCA}_{\mathcal{G}}^{\Sigma}$ (i.e., $\gamma_x = \gamma_{x'}$ for all $x, x' \in X$) is behavioural equivalent to the cell of a single-celled CA.

Proof. Let $\mathcal{X} = (X, \langle \chi, \gamma \rangle)$ be a uniform CA on the signature (\mathcal{M}, N, S) . If $X = \emptyset$ the lemma holds vacuously, so assume that X is not empty.

We construct a single-celled CA $\mathcal{U} = (\{*\}, \langle \xi, \delta \rangle) \in \mathbf{CCA}_{\mathcal{G}}^{\Sigma}$ on the same signature. The cowriter-comonad coalgebra of \mathcal{U} is trivial: $(-) \otimes * \stackrel{\text{def}}{=} \lambda m . *$.

As for the local rule δ , first note that the neighbourhood quotient $N_* = \{N\}$ is a singleton set, so all neighbourhood configurations $z: N_* \rightarrow S$ are of the form $\text{const}_s = \lambda |n|_* . s: N_* \rightarrow S$ for some $s \in S$. Since \mathcal{U} has only one cell, δ has only one component $\delta_*: N_* \rightarrow S$ that we need to define. Since X is not empty, there exists an $x \in X$, and uniformity asserts that γ_x is the same for any choice of x . Therefore we can define δ_* unambiguously as follows:

$$\delta_*(\text{const}_s) \stackrel{\text{def}}{=} \gamma_x(\lambda |x| . s).$$

Now $h: \mathcal{X} \rightarrow \mathcal{U}$ defined as $h = \text{const}_*$ satisfies all properties of a coalgebra homomorphism (Theorem 5.9), and indeed sends all cells of \mathcal{X} to a CA with a single cell. \square

Example 5.12 (CGOL is behavioural equivalent to a constant dead CA). The proof of Lemma 5.11 shows that two cells $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ of two uniform CA are behaviourally equivalent as soon as γ_x and δ_y agree on constant neighbourhood configurations (inputs in which all neighbours have the same state).

For CGOL (Examples 2.17 and 3.6) this means that all cells are behavioural equivalent to a rather trivial single-celled CA whose single cell always becomes \circ the next timestep. Indeed, we observe that CGOL's local rule (Eq. (5)) always outputs \circ on constant neighbourhood configurations $z: N \rightarrow \{\bullet, \circ\}$ (for $z = \text{const}_\circ$, the target cell and all neighbours are dead, so it stays dead. For $z = \text{const}_\bullet$, the cell has too many alive neighbours to stay alive).

This observation is quite significant, since CGOL is universal (both Turing universal, as well as universal for 2D CA [9]), while the dynamics of the constant-dead-single-cell CA are completely trivial.

5.3 The terminal coalgebra

With a practical characterisation of coalgebra homomorphisms at hand (theorem 5.9), we can construct a more complete overview of the behavioural equivalence relation. The terminal object of $\mathbf{CCA}_{\mathcal{G}}^{\Sigma}$ provides a succinct characterisation of behavioural equivalent cells: all pairs of behavioural equivalent cells are mapped to the same cell in this terminal coalgebra. In this subsection we construct this terminal coalgebra.

5.3.1 Coalgebra contractions

We will construct the terminal coalgebra by (1) *contracting* every CA $\mathcal{X} \in \mathbf{CCA}_{\mathbb{G}}^{\Sigma}$ into a canonical form $\|\mathcal{X}\| \in \mathbf{CCA}_{\mathbb{G}}^{\Sigma}$, such that there exist exactly one $\iota_{\mathcal{X}}: \mathcal{X} \rightarrow \|\mathcal{X}\|$, and (2) combining all contractions in one coalgebra. The second step requires some additional care to avoid quantifying over a large collection that is not a set.

Definition 5.13: Contraction

Given a coalgebra $\mathcal{X} = (X, \langle \chi, \gamma \rangle) \in \mathbf{CCA}_{\mathbb{G}}^{\Sigma}$, define the *contraction* $\|\mathcal{X}\|$ of \mathcal{X} to be the $C_{\mathbb{G}}$ -coalgebra $(\|X\|, \langle \|\chi\|, \|\gamma\| \rangle)$ where:

- $\|X\|$ are the equivalence classes $\|x\|$ (with $x \in X$) under the transitive closure of $\overset{\text{BE}}{\approx}$.
- $m \|\otimes\| \|x\| \stackrel{\text{def}}{=} \|m \otimes x\|$.
- $\|\gamma\|_{\|x\|}: S^{N_{\|x\|}} \rightarrow S$ is defined to be $\gamma_x \circ u$ where

$$u: S^{N_{\|x\|}} \rightarrow S^{N_x}$$

$$u(k)(|n|_x) \stackrel{\text{def}}{=} k(|n|_{\|x\|}),$$

which is well defined only if $N_{\|x\|} \leq N_x$ for all $x \in \|x\|$. This indeed holds; see Theorem 5.14 below.

Lemma 5.18 will show that $\|\mathcal{X}\| \in \mathbf{CCA}_{\mathbb{G}}^{\Sigma}$.

Theorem 5.14

For all $x \in X$ it holds that:

$$N_{\|x\|} \leq \bigwedge_{x \in \|x\|} N_x \leq N_x.$$

This theorem is easiest to prove with the help of a new definition and two lemmas. For convenience, denote $N_{\wedge} \stackrel{\text{def}}{=} \bigwedge_{x \in \|x\|} N_x$ and the classes in N_{\wedge} by $|n|_{\wedge}$.

Definition 5.15: Domino relation

Let $U \in \mathbf{Set}$ and let $\{A_i \subseteq U\}_{i \in I}$ be a set of subsets of U . Then there exists a *domino path* from $a \in U$ to $b \in U$ if there exists an $n \in \mathbb{N}$ and a sequence $(x_i, A_i)_{1 \leq i \leq n+1}$ such that $a \in A_0$, $b \in A_{n+1}$ (x_{n+1} may differ from b) and

$$x_{i+1} \in A_i \cap A_{i+1} \quad \text{for all } 0 \leq i \leq n.$$

This relation is an equivalence relation $\boxplus \subseteq U \times U$:

- The singleton sequence $((a, A))$ proves the reflexive case.
- Symmetry follows from the sequence (x_{n+1-i}, A_{n+1-i}) , which gives a domino path from b to a .
- Transitivity is simply concatenating sequences.




Figure 10: Example of a domino path proving $1 \boxplus 5$ in the collection $U = \{\{1, 6\}, \{6, 3\}, \{3, 2\}, \{2, 4\}, \{4, 5\}\}$. Note that we can interpret the elements of U as domino pieces in the actual game, thus justifying the name of the relation.

Lemma 5.16

Consider the domino relation over the N -partitions $\{N_x \mid x \in \|x\|\}$, then

$$|m|_\lambda = |m'|_\lambda \text{ iff } m \boxplus m'. \quad (39)$$

Proof. We show both directions, to begin with right-to-left. Let $(n_i, |n_i|_\lambda)_{1 \leq i \leq k}$ be the domino sequence proving $m \boxplus m'$. All sets in N_λ are mutually disjoint, but $n_{i+1} \in |n_i|_\lambda$ and $n_{i+1} \in |n_{i+1}|_\lambda$. Therefore it must be that $|n_i|_\lambda = |n_{i+1}|_\lambda$, which implies all sets in the sequence are equal, and in particular $|m|_\lambda = |m'|_\lambda$.

We prove the other direction via contraposition: assume that there exists no domino path from m to m' , but still $|m|_\lambda = |m'|_\lambda$. We will derive a contradiction by constructing an $N_\Delta \in \text{PartOrd}(N)$ such that $N_\lambda \lesssim N_\Delta \leq N_{x'}$ for all $x' \in \|x\|$, which cannot exist since N_λ is the greatest lower bound.

Let $N_\Delta = \{|n|_\Delta\}_{n \in N}$ where:

$$|n|_\Delta \stackrel{\text{def}}{=} \{n' \in N \mid n \boxplus n'\}.$$

Since we are using domino paths over $\{N_{x'} \mid x' \in \|x'\|\}$ it holds that $\bigcup_{x' \in \|x'\|} |m|_{x'} \subseteq |m|_\Delta$ (since clearly $m \boxplus n$ for all $n \in |m|_{x'}$ via a length-1 path, so $|m|_{x'} \subseteq |m|_\Delta$). Hence $N_\Delta \leq N_x$.

It remains to show that $N_\lambda \lesssim N_\Delta$. We first show $N_\lambda \leq n_\Delta$. Take any $n \in N$, then it is to show that $|n|_\Delta \subseteq |n|_\lambda$. For any $n' \in N$ we have $n \boxplus n'$ iff it holds that $|n|_\Delta = |n'|_\Delta$. But we have already shown the right-to-left direction of the current lemma, thus we also know that $n \boxplus n'$ implies $|n|_\lambda = |n'|_\lambda$. Hence $|n|_\Delta = |n'|_\Delta$ implies $|n|_\lambda = |n'|_\lambda$, i.e., $|N|_\Delta \subseteq |n|_\lambda$, as required.

The last step is to show that $N_\lambda \neq N_\Delta$. By our assumption, $m \not\boxplus m'$, and thus $|m|_\Delta \neq |m'|_\Delta$, while we assumed $|m|_\lambda = |m'|_\lambda$. Thus we conclude $N_\lambda \lesssim N_\Delta \leq N_x$ for all $x \in \|x\|$, which is our sought contradiction. \square

Lemma 5.17

For $x \in X$ and $y \in Y$ (given coalgebras $\mathcal{X} = (X, \langle \chi, \gamma \rangle)$ and $\mathcal{Y} = (Y, \langle \xi, \delta \rangle)$), if $x \stackrel{\text{BE}}{\approx} y$ then $(m \otimes x) \stackrel{\text{BE}}{\approx} (m \odot y)$.

Proof. If $x \stackrel{\text{BE}}{\approx} y$ then there exists a $\mathcal{Z} = (Z, \langle \zeta, \alpha \rangle)$ and $h: \mathcal{X} \rightarrow \mathcal{Z}$ and a $k: \mathcal{Y} \rightarrow \mathcal{Z}$ such that $h(x) = k(y)$. From the \mathcal{M} - χ - ζ - and \mathcal{M} - ξ - ζ -homomorphism properties of h and k respectively it follows that

$$h(m \otimes x) = m \odot h(x) = m \odot k(y) = k(m \odot y),$$

which indeed proves $(m \otimes x) \stackrel{\text{BE}}{\approx} (m \odot y)$. \square

We can now prove Theorem 5.14:

Proof. For convenience, denote the classes $N_{\|x\|} = \{|n|_{\bullet} \mid n \in N\}$, and take any $n \in N$; then it is to show that $|n|_{\wedge} \subseteq |n|_{\bullet}$. To this end, we pick arbitrary $m, m' \in |n|_{\wedge}$ and show that $|m|_{\bullet} = |m'|_{\bullet}$. (If we do this for all $m \in |n|_{\wedge}$, then it follows that $|n|_{\bullet} = |m|_{\bullet}$ for all such m , so $m \in |n|_{\bullet}$ and it follows that indeed $|n|_{\wedge} \subseteq |m|_{\bullet}$.) There either exists or not exists an $x_0 \in \|x\|$ such that $y := m \otimes x_0 = m' \otimes x_0 =: y'$. If such an x_0 exists, then the identity coalgebra homomorphism shows $y \stackrel{\text{BE}}{\approx} y'$, which implies $\|y\| = \|y'\|$. Hence both m and m' have that

$$m \|\otimes\| \|x_0\| = \|m \otimes x_0\| = \|y\| = \|y'\| = \|m' \otimes x_0\| = m' \|\otimes\| \|x_0\|,$$

which proves that m and m' lead from $\|x_0\|$ to the same cell $\|y\|$ via $\|\otimes\|$, thus $|m|_{\bullet} = |m'|_{\bullet}$ by definition of neighbourhood quotients.

The case where no such x_0 exists remains. From Lemma 5.16 it still follows that $m \boxplus m'$. So we obtain a domino sequence $(n_i, |n_i|_{x_i})_{1 \leq i \leq k}$ such that $x_i \in \|x\|$, $n_{i+1} \in |n_i|_{x_i} \cap |n_{i+1}|_{x_{i+1}}$, $m = n_0$ and $m' \in |n_{k+1}|_{x_{k+1}}$. From the definition of neighbourhood quotients and the fact that $n_{i+1} \in |n_i|_{x_i}$ it follows that

$$n_i \otimes x_i = n_{i+1} \otimes x_i.$$

Therefore the identity coalgebra homomorphism proves that $n_i \otimes x_i \stackrel{\text{BE}}{\approx} n_{i+1} \otimes x_i$, so by definition of $\|(-)\|$ we find $\|n_i \otimes x_i\| = \|n_{i+1} \otimes x_i\|$. But by definition of $\|\otimes\|$ this means that $n_i \|\otimes\| \|x_i\| = n_{i+1} \|\otimes\| \|x_i\|$. Thus from the definition of neighbourhood quotients $|(-)|_{\bullet} = |(-)|_{\|x\|}$ (now in the contracted coalgebra) it follows that $|n_i|_{\bullet} = |n_{i+1}|_{\bullet}$. This holds for all $0 \leq i \leq k$, thus by transitivity we find the desired $|m|_{\bullet} = |m'|_{\bullet}$. \square

The following lemma shows that the output of a contraction is again in $\mathbf{CCA}_{\mathcal{G}}^{\Sigma}$, i.e., the first component is again a cowriter-comonad coalgebra:

Lemma 5.18

If $\mathcal{X} \in \mathbf{CCA}_{\mathcal{G}}^{\Sigma}$ then $\|\mathcal{X}\| \in \mathbf{CCA}_{\mathcal{G}}^{\Sigma}$. That is, $\|\otimes\|$ is a cowriter-comonad coalgebra that satisfies Eq. (11) and (12).

Proof. Take any $n, m \in \mathcal{M}$ and $\|x\| \in \|X\|$ (for some $x \in X$), then we verify that Eq. (11) and (12) are satisfied as follows:

$$\begin{aligned} \mathbf{1} \|\otimes\| \|x\| &\stackrel{\text{def}}{=} \mathbf{1} \otimes x = \|x\| \\ n \|\otimes\| (m \|\otimes\| \|x\|) &\stackrel{\text{def}}{=} n \|\otimes\| \|m \otimes x\| \stackrel{\text{def}}{=} \|(n \bullet m) \otimes x\| \stackrel{\text{def}}{=} (n \bullet m) \|\otimes\| \|x\|. \end{aligned}$$

\square

Example 5.19 (Example of a contraction). Figure 11 visualises the effect of a contraction of a general CA. We assume that $\gamma_a = \gamma_{a'}$ and $\gamma_b = \gamma_{b'}$, which implies that $a \stackrel{\text{BE}}{\approx} a'$ and $b \stackrel{\text{BE}}{\approx} b'$. We also assume that γ_a and γ_b disagree on constant configurations, which ensures that $a \not\stackrel{\text{BE}}{\approx} b$ and $a \not\stackrel{\text{BE}}{\approx} b'$. If they would happen to agree on constant configurations, then the CA would be contracted to a single cell.




Figure 11: **Left:** a small CA on the monoid freely generated over $\{L, M, R, A\}$ with the empty string ε as identity element, neighbourhood $N = \{\varepsilon, L, M, R, A\}$ and four cells $X = \{a, a', b, b'\}$. Only paths in N are shown, with labels are indicated in gray; omitted paths (in N) implicitly denote paths to the cell itself. **Right:** contraction of the CA on the left, assuming that γ_a and γ_b disagree on constant configurations, but that $\gamma_a = \gamma_{a'}$ and $\gamma_b = \gamma_{b'}$.

Remark 5.20. It is possible that $N_{\|x\|} \lesssim \bigwedge_{x' \in \|x\|} N_{x'}$. Take for example the neighbourhood-uniform CA $\mathcal{X} \in \mathbf{CCA}_{\text{NU}}^\Sigma$ of Figure 12 below:




Figure 12: The neighbourhood-uniform CA $\mathcal{X} \in \mathbf{CCA}_{\text{NU}}^\Sigma$ with $X = \{a, b, c\}$. The underlying monoid \mathcal{M} is freely generated over the alphabet $\{L, M, R\}$ and has the empty string ε as identity element. The neighbourhood is $N = \{\varepsilon, L, M, R\}$. Only arrows in N are shown (but this is a generating set of \mathcal{M}), and the notation L/R means that both the paths L and R lead to the same neighbour.

The neighbourhood quotients of its cells are:

$$N_c = \{|\varepsilon, M|, |L, R|\}$$

$$N_a = N_b = \{|\varepsilon, L, R|, |M|\}.$$

These neighbourhoods are incompatible, thus γ_c can be chosen independently from γ_a and γ_b . Thus we can choose γ_c such that it disagrees with γ_a on constant inputs (e.g., pick $S = \{\bullet, \circ\}$, $\gamma_c = \text{const}_\circ$ and $\gamma_a = \gamma_b = \text{const}_\bullet$). In this case, c will not be contracted to a in $\|\mathcal{X}\|$, but a and b will both be contracted to $\|a\|$. We would obtain the contraction of Figure 13 below:




Figure 13: Contraction $\|\mathcal{X}\|$ of the CA \mathcal{X} Figure 12 (in the case that the local rule at cell c disagrees with the local rule at cell a on constant configurations; otherwise the contraction would be a single cell).

In this contraction, it trivially holds that $\bigwedge_{x \in \|\mathcal{X}\|} N_x = N_c$ (the meet of a singleton set), but $N_{\|\mathcal{X}\|} = \{|\varepsilon, M|, |L, R|\}$. Thus $N_{|c|} \lesssim \bigwedge_{x \in \|\mathcal{X}\|} N_x$.

Before we can prove the uniqueness of $\iota_{\mathcal{X}}$, we need to show several other properties of contractions. But we can already start with defining the map itself:

Lemma 5.21

There exist a coalgebra homomorphism $\iota_{\mathcal{X}}: \mathcal{X} \rightarrow \|\mathcal{X}\|$.

Proof. The most obvious definition works:

$$\iota_{\mathcal{X}}(x) \stackrel{\text{def}}{=} \|x\| \quad (40)$$

for all $x \in X$. Applying the characterisation of coalgebra homomorphisms given by Theorem 5.9, we need to show that $\iota_{\mathcal{X}}$ (1) is a $\mathcal{M}\text{-}\chi\text{-}\|\chi\|$ -homomorphisms and (2) that $(z \setminus\!\! \setminus x) = (z \setminus\!\! \setminus \|x\|)$ for all $N_{\|\mathcal{X}\|}$ -compatible $z: N \rightarrow S$. This last requirement follows directly from the definition of $\|\gamma\|$ (Definition 5.13).

To show $\iota_{\mathcal{X}}$ is an $\mathcal{M}\text{-}\chi\text{-}\|\chi\|$ -homomorphism, simply observe that:

$$\iota_{\mathcal{X}}(m \otimes x) \stackrel{\text{def}}{=} \|m \otimes x\| = m \|\otimes\| \|x\| \stackrel{\text{def}}{=} m \|\otimes\| \iota_{\mathcal{X}}(x).$$

□

The next observation is rather trivial, but of great importance:

Lemma 5.22

There exist no distinct $u, v \in \|\mathcal{X}\|$ such that $u \stackrel{\text{BE}}{\approx} v$.

Proof. By definition of a contraction, we can write $u = \|x\|$ and $v = \|x'\|$ for some $x, x' \in X$. If there would exist $h: \|\mathcal{X}\| \rightarrow \mathcal{Y}$ and $\ell: \|\mathcal{X}\| \rightarrow \mathcal{Y}$ such that $h(u) = \ell(v)$, then $(h \circ \iota_{\mathcal{X}})(x) = (\ell \circ \iota_{\mathcal{X}})(x')$, which would imply $x \stackrel{\text{BE}}{\approx} x'$. But then $u = \|x\| = \|x'\| = v$. □

Lemma 5.23

$\|(-)\|$ can be extended to an endofunctor $\mathbf{CCA}_{\mathbb{G}}^{\Sigma} \rightarrow \mathbf{CCA}_{\mathbb{G}}^{\Sigma}$.

Proof. Lemma 5.18 asserts that the object map of $\|(-)\|$ indeed outputs another object in \mathbf{CCA}_G^Σ .

We define the action of $\|(-)\|$ on homomorphisms $h: \mathcal{X} \rightarrow \mathcal{Y}$ as follows:

$$\|h\|(\|x\|) \stackrel{\text{def}}{=} \|h(x)\| \quad (41)$$

for all $\|x\| \in \|x\|$.

Functoriality is immediate, but it is obvious not that $\|h\|$ is a coalgebra homomorphism; we prove this via the characterisation of coalgebra homomorphisms (Theorem 5.9). The first property we need to show is that $\|h\|$ is a \mathcal{M} - $\|\chi\|$ - $\|\xi\|$ -homomorphism. This can indeed be seen as follows (for arbitrary $\|x\| \in \|X\|$):

$$\begin{aligned} \|h\|(\|m \|\otimes\| \|x\|) &= \|h\|(\|m \otimes x\|) && // \text{Definition } \|\otimes\|. \\ &= \|h(m \otimes x)\| && // \text{Definition } \|h\|. \\ &= \|m \otimes h(x)\| && // h \text{ is a } \mathcal{M}\text{-}\chi\text{-}\xi\text{-homomorphism.} \\ &= m \|\otimes\| \|h(x)\| && // \text{Definition } \|\otimes\|. \\ &= m \|\otimes\| \|h\|(\|x\|). && // \text{Definition } \|h\|. \end{aligned}$$

The other required property states that, for all $N_{\|h(x)\|}$ -compatible $z: N \rightarrow S$, it holds that:

$$\|\gamma\|_{\|x\|}(z \|\backslash\| \|x\|) = \|\delta\|_{\|h(x)\|}(z \|\backslash\| \|h(x)\|).$$

But we know h is a coalgebra homomorphism, so

$$\gamma_x(z' \|\backslash\| x) = \delta_{h(x)}(z' \|\backslash\| h(x)) \quad (42)$$

for all $N_{h(x)}$ -compatible $z': N \rightarrow S$. Since both $\iota_{\mathcal{Y}}$ and h are N -restricted surjections, it holds that $N_{\|h(x)\|} \leq N_{h(x)} \leq N_x$, so z is also $N_{h(x)}$ - and N_x -compatible. Therefore:

$$\|\gamma\|_{\|x\|}(z \|\backslash\| \|x\|) \stackrel{\text{def}}{=} \gamma_x(z \|\backslash\| x) \stackrel{(42)}{=} \delta_{h(x)}(z \|\backslash\| h(x)) \stackrel{\text{def}}{=} \|\delta\|_{\|h(x)\|}(z \|\backslash\| \|h(x)\|).$$

□

Corollary 5.24

$\iota: \mathbf{Id}_{\mathbf{Coalg}(C_G)} \Rightarrow \|(-)\|$ is a natural transformation.

Proof. Directly from the definitions of $\iota_{\mathcal{X}}$ and $\|(-)\|$ on morphisms (see the proof of Lemma 5.23:

$$\|h\|(\iota_{\mathcal{X}}(x)) \stackrel{\text{def}}{=} \|h\|(\|x\|) \stackrel{\text{def}}{=} \|h(x)\| \stackrel{\text{def}}{=} \iota_{\mathcal{Y}}(h(x))$$

(for all $h: \mathcal{X} \rightarrow \mathcal{Y}$ and $x \in X$). Hence the naturality square commutes:

$$\begin{array}{ccc} \mathcal{X} & \xrightarrow{h} & \mathcal{Y} \\ \downarrow \iota_{\mathcal{X}} & & \downarrow \iota_{\mathcal{Y}} \\ \|\mathcal{X}\| & \xrightarrow{\|h\|} & \|\mathcal{Y}\| \end{array} \quad (43)$$

□

Lemma 5.25

For any coalgebra $\mathcal{X} \in \mathbf{CCA}_{\mathcal{C}}^{\Sigma}$ it holds that $\|\mathcal{X}\| \cong \|\|\mathcal{X}\|\|$ with the isomorphism being $\iota_{\|\mathcal{X}\|}$.

Proof. No two distinct cells of $\|\mathcal{X}\|$ are behaviourally equivalent, so any coalgebra homomorphism $h: \|\mathcal{X}\| \rightarrow \|\|\mathcal{X}\|\|$ must be injective. Contracting never adds cells, thus h must also be surjective. Hence h must be an isomorphism; and $\iota_{\|\mathcal{X}\|}$ witnesses the existence of such a h . \square

Theorem 5.26

The coalgebra homomorphism $\iota_{\mathcal{X}}: \mathcal{X} \rightarrow \|\mathcal{X}\|$ is the unique morphism between those coalgebras.

Proof. Suppose any other $h: \mathcal{X} \rightarrow \|\mathcal{X}\|$, and take any $x \in X$. Then there exists an $x' \in X$ such that $h(x) = \|x'\|$. From the proof of Lemma 5.23 it follows that the action of $\|h\|: \|\mathcal{X}\| \rightarrow \|\|\mathcal{X}\|\|$ maps $\|x\|$ to $\|h(x)\| = \|\|x'\|\|$ (recall that $\|\|X\|\|$ has as elements singleton sets $\|\|x\|\| \stackrel{\text{def}}{=} \{\|x\|\}$ for every $\|x\| \in \|\mathcal{X}\|$). See the proof of Lemma 5.25).

From Lemma 5.25 it follows that $\iota_{\|\mathcal{X}\|}$ is an isomorphism, thus $\iota_{\|\mathcal{X}\|}^{-1} \circ \|h\|: \|\mathcal{X}\| \rightarrow \|\mathcal{X}\|$ is a coalgebra homomorphism that sends $\|x\|$ to $\|x'\|$. Thus $\|x\| \stackrel{\text{BE}}{\approx} \|x'\|$ in $\|\mathcal{X}\|$; but a contraction has no distinct BE cells (Lemma 5.22), thus

$$\iota_{\mathcal{X}}(x) \stackrel{\text{def}}{=} \|x\| = \|x'\| = h(x).$$

Since x was arbitrary it follows that $\iota_{\mathcal{X}} = h$. \square

It is important to realise that contractions may *not* preserve uniformity, neighbourhood-uniformity and extension-uniformity. Note that $\mathbf{CCA}_{\text{NU}}^{\Sigma}$ includes all CA that have one of these three forms of uniformity, therefore we prove the following: if $\mathcal{X} \in \mathbf{CCA}_{\text{NU}}^{\Sigma}$, then $\|\mathcal{X}\|$ might not be in $\mathbf{CCA}_{\text{NU}}^{\Sigma}$.

Lemma 5.27


There exists a CA $\mathcal{X} \in \mathbf{CCA}_{\text{NU}}^{\Sigma}$ such that $\|\mathcal{X}\| \notin \mathbf{CCA}_{\text{NU}}^{\Sigma}$.

Proof. We construct a CA such that its contraction has two cells with the same neighbourhood quotients, but with different local rules. By definition of neighbourhood-uniformity (Definition 4.8, any CA in $\mathbf{CCA}_{\text{NU}}^{\Sigma}$ has the property that all cells with the same neighbourhood have the same local rule.

Define the CA $\mathcal{X} = (X, \langle \chi, \gamma \rangle)$ with $X = \{x_i\}_{i \in \mathbb{Z}} \cup \{y_i\}_{i \in \mathbb{Z}}$, on the monoid \mathcal{M} freely generated over the alphabet $\{L, M, R, A\}$ with the empty string ε as the unit element and neighbourhood $N = \{L, M, R, A\}$. We use the two-state alphabet $S = \{\bullet, \circ\}$, local rules $\gamma_{x_i} = \text{const}_{\bullet}$ and $\gamma_{y_i} = \text{const}_{\circ}$ for all $i \in \mathbb{Z}$, and cells connections χ defined as:

$$\begin{aligned} L \otimes x_i &= M \otimes x_i = x_{i-1} \\ R \otimes x_i &= x_{i+1} \\ A \otimes x_i &= y_i \\ L \otimes y_i &= y_{i-1} \\ M \otimes y_i &= R \otimes y_i = y_{i+1} \\ A \otimes y_i &= x_i, \end{aligned}$$


which is visualised in Figure 5.27.

Figure 14: Fragments of the CA \mathcal{X} of the proof of Lemma 5.27.

There exists a “shifting”-endohomomorphism $h_n: \mathcal{X} \rightarrow \mathcal{X}$ for every $n \in \mathbb{Z}$ defined as:

$$\begin{aligned} h_n(x_i) &\stackrel{\text{def}}{=} x_{i+n} \\ h_n(y_i) &\stackrel{\text{def}}{=} y_{i+n} \end{aligned}$$

for all $i \in \mathbb{Z}$. Thus $x_i \stackrel{\text{BE}}{\approx} x_{i+n}$ and $y_i \stackrel{\text{BE}}{\approx} y_{i+n}$ for all $i, n \in \mathbb{Z}$. Note that $x_i \stackrel{\text{BE}}{\not\approx} y_j$ for all $i, j \in \mathbb{Z}$ since γ_{x_i} and γ_{y_j} disagree on all inputs.

Figure 15: Contraction $\|\mathcal{X}\|$ of the CA \mathcal{X} of Figure 14 and the proof of Lemma 5.27.

Thus the contraction of \mathcal{X} (see Figure 15) has two classes, $\|x_0\|$ for the x_i 's and $\|y_0\|$ for the y_i 's, with the same local rules $\|\gamma\|_{\|x_0\|} = \text{const}_\bullet$ and $\|\gamma\|_{\|y_0\|} = \text{const}_\circ$ respectively. However, these two classes have the same neighbourhood quotients:

$$N_{\|x_0\|} = N_{\|y_0\|} = \{|\varepsilon, L, M, R|, |A|\},$$

and thus both the conditions of neighbourhood-uniformity and extension-uniformity would require $\|x_0\|$ and $\|y_0\|$ to have the same local rule if $\|X\|$ were to be a CA. \square

5.3.2 Construction of the terminal coalgebra

The previous discussion showed that there is a unique homomorphism from every $\mathcal{X} \in \mathbf{CCA}_{\mathbb{C}}^{\Sigma}$ to $\|\mathcal{X}\|$. The strategy is now to construct the terminal coalgebra by putting all contractions of CA together in one coalgebra. To do so, two remaining challenges need to be solved:

1. There may (and in fact do) exist morphisms between contractions, and hence different morphisms from \mathcal{X} to the union of all contractions. This is shown in Section 5.3.3.
2. The disjointed union of coalgebras $\bigsqcup_{\mathcal{X} \in \mathbf{CCA}_{\mathbb{C}}^{\Sigma}} \|\mathcal{X}\|$ quantifies over a collection greater than \mathbf{Set} ; the resulting collection is therefore not a set, and hence not the carrier of a coalgebra (since $C_{\mathbb{C}}: \mathbf{Set} \rightarrow \mathbf{Set}$). We will instead only work with contractions that are unique up to isomorphism, i.e., $\bigsqcup_{\mathcal{X} \in \mathbf{CCA}_{\mathbb{C}}^{\Sigma}/\cong} \|\mathcal{X}\|$, but it is still not obvious that this is a set. However, Section 5.3.4 will show that this union has a cardinality.

5.3.3 Morphisms between embeddings

To analyse challenge 1, we will show that the image of $\|(-)\|$ under \mathbf{CCA}_G^Σ (denoted by $\|\mathbf{CCA}_G^\Sigma\|$) is a thin category containing only embeddings as morphisms (i.e., there exists at most one embedding $\|\mathcal{X}\| \hookrightarrow \|\mathcal{Y}\|$), but also that not all morphisms are isomorphisms. Hence even after removing isomorphic contractions via a quotient, the union $\mathcal{T} = \bigcup_{\|\mathcal{X}\| \in \|\mathbf{CCA}_G^\Sigma\| / \cong} \|\mathcal{X}\|$ of all (up to isomorphism unique) contractions still contains distinct BE cells; but – provided T it is a set – we can contract all the morphisms by using $\|\mathcal{T}\|$ to obtain a coalgebra towards which every CA has a unique morphism. The following four lemmas and theorem will justify those claims.

Lemma 5.28

For $\|\mathcal{X}\|, \|\mathcal{Y}\| \in \|\mathbf{CCA}_G^\Sigma\|$, every $e: \|\mathcal{X}\| \rightarrow \|\mathcal{Y}\|$ is an embedding.

Proof. If h is not injective, then there would exist $\|x\| \neq \|x'\|$ such that $h(\|x\|) = h(\|x'\|)$. But then $\|x\| \stackrel{\text{BE}}{\approx} \|x'\|$, contradicting Lemma 5.22. \square

The next lemma is needed to show that there exists at most one embedding between contractions. The intuition of this next lemma is as follows: given two CA cells z_1 and z_2 that are essentially the same, and also $\mathcal{M} \odot z_1$ is identical to $\mathcal{M} \odot z_2$, then we can exchange all cells in $\mathcal{M} \odot z_1$ with the corresponding cells $\mathcal{M} \odot z_2$ without any essential change to the CA.

More generally, the next lemma will show that for a coalgebra $\mathcal{Y} = (Y, \langle \xi, \delta \rangle)$ containing two isomorphic subcoalgebras with isomorphism $\psi: \mathcal{Z}_1 \xrightarrow{\sim} \mathcal{Z}_2$, we can construct a coalgebra \mathcal{Y}' with \mathcal{Z}_1 and \mathcal{Z}_2 exchanged, and “lift” ψ to:

$$\psi^+: \mathcal{Y} \rightarrow \mathcal{Y}'$$

$$\psi^+(y) \stackrel{\text{def}}{=} \begin{cases} \psi(z_1) \in Z_2 & u = z_1 \in Z_1 \\ \psi^{-1}(z_2) \in Z_1 & u = z_2 \in Z_2 \\ y & u = y \in Y \setminus (Z_1 \cup Z_2) \end{cases}. \quad (44)$$




Figure 16: If \mathcal{Y} has two isomorphic subcoalgebras \mathcal{Z}_1 and \mathcal{Z}_2 , then there exists an isomorphism $\psi^+: \mathcal{Y} \cong \mathcal{Y}'$ that exchanges these subcoalgebras.

Lemma 5.29

The image $\mathcal{Y}' \stackrel{\text{def}}{=} \psi^+[\mathcal{Y}]$ is a coalgebra and $\psi^+ : \mathcal{Y} \xrightarrow{\sim} \mathcal{Y}'$ an isomorphism.

Proof. We first show that $\mathcal{Y}' = (Y', \langle \xi', \delta' \rangle)$ is a well-defined C_G -coalgebra.

Carrier: On the level of sets, it is clear that ψ^+ is a bijection $Y \xrightarrow{\sim} Y'$, so the carrier Y' of \mathcal{Y}' is simply Y .

ξ' of dynamics: As for the ξ' of \mathcal{Y}' , note that every element of Y' can be written in the form $\psi^+(u)$ for an $u \in Y$, so define for all $m \in \mathcal{M}$:

$$\begin{aligned} m \otimes' (-) &: Y' \rightarrow Y' \\ m \otimes' \psi^+(u) &\stackrel{\text{def}}{=} \psi^+(m \otimes u). \end{aligned} \tag{45}$$

It indeed satisfies the path-comonad laws (Eq. (11) and (12)):

$$\begin{aligned} \mathbf{1} \otimes' \psi^+(u) &\stackrel{\text{def}}{=} \psi^+(\mathbf{1} \otimes u) = \psi^+(u). \\ (m \bullet n) \otimes' \psi^+(u) &\stackrel{\text{def}}{=} \psi^+((m \bullet n) \otimes u) = \psi^+(m \otimes n \otimes u) \stackrel{\text{def}}{=} m \otimes' \psi^+(n \otimes u) \stackrel{\text{def}}{=} m \otimes' n \otimes' \psi^+(u). \end{aligned}$$

δ' of dynamics: Being an isomorphism, we have $N_{z_1} = N_{\psi(z_1)}$ for all $z_1 \in Z_1$, thus the extension property (Theorem 5.9) gives $\delta_{z_1} = \delta_{\psi(z_1)}$, and by definition of ψ^+ , the latter indeed equals $\delta_{\psi^+(z_1)}$. Analogously $\delta_{z_2} = \delta_{\psi^{-1}(z_2)} = \delta_{\psi^+(z_2)}$ for all $z_2 \in Z_2$. Trivially also $\delta_y = \delta_{\psi^+(y)}$ for all $y \in Y \setminus (Z_1 \cup Z_2)$ (since $\psi^+(y) \stackrel{\text{def}}{=} y$), thus the δ' of \mathcal{Y}' is just δ .

ψ^+ is a coalgebra isomorphism: That ψ^+ is a \mathcal{M} - ξ - ξ' -homomorphism follows directly from the definitions of δ' and ξ' . To show it is an isomorphism, we construct its inverse by lifting ψ^{-1} analogously:

$$\begin{aligned} (\psi^{-1})^+ &: \mathcal{Y}' \rightarrow \mathcal{Y} \\ (\psi^{-1})^+(u) &\stackrel{\text{def}}{=} \begin{cases} \psi^{-1}(z_1) \in Z_2 & u = z_1 \in Z_1 \\ \psi(z_2) \in Z_1 & u = z_2 \in Z_2 \\ y & u = y \in Y \setminus (Z_1 \cup Z_2) \end{cases}. \end{aligned}$$

Before we prove that $(\psi^{-1})^+$ has the properties of a coalgebra homomorphism, we first show that it is indeed the inverse of ψ^+ on cells. For any $u \in Y$ we have:

$$\begin{aligned} (\psi^{-1})^+(\psi^+(u)) &\stackrel{\text{def}}{=} \begin{cases} u & u \in Y \setminus (Z_1 \cup Z_2) \\ \psi^{-1}(\psi(u)) & u \in Z_1 \\ \psi(\psi^{-1}(u)) & u \in Z_2 \end{cases} \\ &= u. \end{aligned}$$

Analogously we find $\psi^+((\psi^{-1})^+(u)) = u$ for all $u \in Y'$.

As usual, to prove that $(\psi^{-1})^+$ is a coalgebra homomorphism we will use Theorem 5.9. To see that $(\psi^{-1})^+$ is a \mathcal{M} - ξ' - ξ -homomorphism, again note that every element of Y' can be written in the form $\psi^+(u)$ for some $u \in Y$. Thus it suffices to take arbitrary $m \in \mathcal{M}$ and $u \in Y$ and observe that:

$$\begin{aligned} (\psi^{-1})^+(m \otimes' \psi^+(u)) &= (\psi^{-1})^+(\psi^+(m \otimes u)) \\ &= m \otimes u \quad // (\psi^{-1})^+ \text{ and } \psi^+ \text{ are inverses, so also on the cell } m \otimes u. \\ &= m \otimes (\psi^{-1})^+(\psi^+(u)). \quad // \text{Idem but for the cell } u. \end{aligned}$$

Writing $u' := \psi^+(u)$ this can be summarised as:

$$(\psi^{-1})^+(m_{\otimes} 'u') = m_{\otimes} (\psi^{-1})^+(u'),$$

exactly the required relation.

It remains to show the extension property of $(\psi^{-1})^+$. But indeed, for an arbitrary $\psi^+(u) \in Y'$:

$$\delta_{(\psi^{-1})^+(\psi^+(u))} = \delta_u = \delta_{\psi^+(u)}$$

where the first equation follows from the fact that $(\psi^{-1})^+$ and ψ^+ are inverses on cells, and the second is the earlier observed extension property of ψ^+ . \square

Lemma 5.30

For $\|\mathcal{X}\|, \|\mathcal{Y}\| \in \|\mathbf{CCA}_{\mathbb{G}}^{\Sigma}\|$, there exist *at most one* embedding $e: \|\mathcal{X}\| \hookrightarrow \|\mathcal{Y}\|$.

Proof. Let $e_1, e_2: \|\mathcal{X}\| \hookrightarrow \|\mathcal{Y}\|$, then we need to show that $e_1 = e_2$.

Since both e_i are surjective on $\mathcal{M} \otimes e_i(\|x\|)$ for all $\|x\| \in \|\mathcal{X}\|$ (by Lemma 5.4), they are isomorphisms $e_i: \|\mathcal{X}\| \xrightarrow{\sim} e_i[\|\mathcal{X}\|]$ (in $\mathbf{Coalg}(C_{\mathbb{G}})$). Surjectivity also means that no cell in $\|\mathcal{Y}\| \setminus e_i[\|\mathcal{X}\|]$ is reachable from a cell in $e_i[\|\mathcal{X}\|]$. Hence $e_i[\|\mathcal{X}\|]$ is a subcoalgebra of $\|\mathcal{Y}\|$ with only *incoming* paths from cells in $\|\mathcal{Y}\| \setminus e_i[\|\mathcal{X}\|]$.

In particular we have $e_1[\|\mathcal{X}\|] \cong \|\mathcal{X}\| \cong e_2[\|\mathcal{X}\|]$. Denote this isomorphism by $\psi: e_1[\|\mathcal{X}\|] \xrightarrow{\sim} e_2[\|\mathcal{X}\|]$. Using Lemma 5.29 we lift ψ to $\psi^+: \|\mathcal{Y}\| \xrightarrow{\sim} \|\mathcal{Y}\|'$, where $\|\mathcal{Y}\|'$ is $\|\mathcal{Y}\|$ but with subcoalgebras $e_1[\|\mathcal{X}\|]$ and $e_2[\|\mathcal{X}\|]$ exchanged. The construction of lemma 5.29 shows that $\psi^+ \upharpoonright e_1[\|\mathcal{X}\|]$ is just the isomorphism ψ with image $e_2[\|\mathcal{X}\|]$ in $\|\mathcal{Y}\|'$, which allows us to define the composite isomorphism $\phi \stackrel{\text{def}}{=} e_2^{-1} \circ (\psi^+ \upharpoonright e_1[\|\mathcal{X}\|]) \circ e_1$.

$$\begin{array}{ccc} \|\mathcal{X}\| & \xrightarrow{\phi} & \|\mathcal{X}\| \\ e_1 \downarrow & & \uparrow e_2^{-1} \\ e_1[\|\mathcal{X}\|] & \xrightarrow{\psi^+ \upharpoonright e_1[\|\mathcal{X}\|]} & e_2[\|\mathcal{X}\|] \end{array}$$

But $\|\mathcal{X}\|$ has no distinct BE cells (Lemma 5.22), thus all endohomomorphism must send every cell to itself. Hence $\phi = \text{Id}_{\|\mathcal{X}\|}$.

Now take any $\|x\| \in \|\mathcal{X}\|$. Then

$$e_2^{-1}(\psi^+(e_1(\|x\|))) = \phi(\|x\|) = \|x\|.$$

Applying e_2 to both sides gives

$$\psi^+(e_1(\|x\|)) = e_2(\|x\|),$$

which is an equality of elements in the subcoalgebra $e_2[\|\mathcal{X}\|]$ of $\|\mathcal{Y}\|'$. Thus these elements are also equal in $\|\mathcal{Y}\|'$. But since ψ^+ is an isomorphism (this is the observation for which Lemma 5.29 is really used) ψ^+ has an inverse, and applying the inverse to both sides gives:

$$(\psi^+)^{-1}(\psi^+(e_1(\|x\|))) = (\psi^+)^{-1}(e_2(\|x\|)),$$

which is an equality in $\|\mathcal{Y}\|$. Choosing $h := (\psi^+)^{-1} \circ \psi^+$ and $k := (\psi^+)^{-1}$ shows that $h(e_1(\|x\|)) = k(e_2(\|x\|))$, thus $e_1(\|x\|) \stackrel{\text{BE}}{\approx} e_2(\|x\|)$. But $\|\mathcal{Y}\|$ is a contraction, and therefore has no distinct BE cells. Thus $e_1(\|x\|) = e_2(\|x\|)$. Since $\|x\|$ was arbitrary, this proves $e_1 = e_2$. \square

The next lemma will justify the need of contracting \mathcal{T} , since it may have non-trivial endohomomorphisms otherwise. It would be possible to postcompose such endohomomorphisms with arbitrary homomorphism $h: \mathcal{X} \rightarrow \mathcal{T}$, which would show that h is not the unique homomorphism from \mathcal{X} to \mathcal{T} .

Lemma 5.31

There exist $\|\mathcal{X}\|, \|\mathcal{Y}\| \in \|\mathbf{CCA}_{\mathbb{C}}^{\Sigma}\|$ and $e: \|\mathcal{X}\| \hookrightarrow \|\mathcal{Y}\|$ such that e is not an isomorphism.

Proof. Let the monoid \mathcal{M} be freely generated over the alphabet $\{L, M, R, U\}$, and consider the contracted coalgebras $\|\mathcal{X}\|$ and $\|\mathcal{Y}\|$ as depicted in Figure 17. We can check that these are indeed contractions for a suitable choice of the local rules: all cells have different neighbourhood quotients that are incomparable under the $\text{PartOrd}(N)$ ordering (Definition 2.2), and hence mutually incompatible (Definition 4.2).

$$\begin{aligned} N_{\|x_0\|} &= N_{\|y_0\|} = \{L, U, |\varepsilon, M, R\} \\ N_{\|x_1\|} &= N_{\|y_1\|} = \{\varepsilon, L, M, |R, U\} \\ N_{\|y_2\|} &= \{\varepsilon, |L, |M, U, |R\} \\ N_{\|y_3\|} &= \{\varepsilon, L, R, |M, U\} \end{aligned}$$

So the local rules at every cell can be chosen independently and differently. We choose them as such; then no pair of cells within $\|\mathcal{X}\|$ or within $\|\mathcal{Y}\|$ is behavioural equivalent to another cell in the same coalgebra.

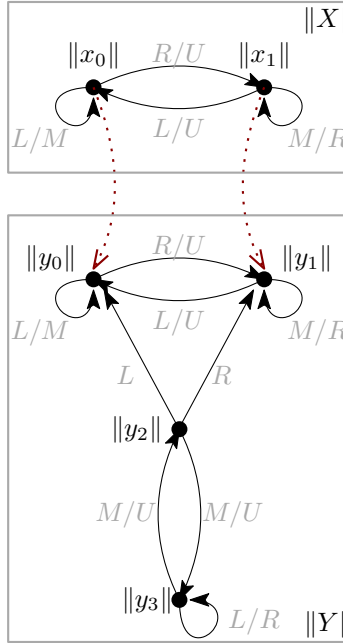


Figure 17: Embedding of a contracted coalgebra $\|\mathcal{X}\|$ (top) into a larger contracted coalgebra $\|\mathcal{Y}\|$ (bottom). The dotted arrows indicate the cell-map of the embedding. The underlying monoid \mathcal{M} is freely generated over $\{L, M, R, U\}$. Only paths in the neighbourhood $N = \{\varepsilon, L, M, R, U\}$ are shown, but this is a generating set of \mathcal{M} . The notation M/U denotes that paths $M, U \in N$ lead to the same cell.

□

Remark 5.32. Any embedding $e: \|\mathcal{X}\| \hookrightarrow \|\mathcal{Y}\|$ has in its image all cells of $\mathcal{M} \circledast e(\|x\|)$ (for all $\|x\| \in \|\mathcal{X}\|$), because of Lemma 5.4. Thus the embedding $e[\|\mathcal{X}\|]$ has exactly the same form as $\|\mathcal{X}\|$ and no outgoing arrows to other cells in $\|\mathcal{Y}\|$, but as the proof of Lemma 5.31 shows, it may have incoming arrows.

The next theorem will be used to prove the uniqueness of the homomorphism towards $\|\mathcal{T}\|$:

Theorem 5.33

There exists at most one unique coalgebra homomorphism $h: \mathcal{X} \rightarrow \|\mathcal{Y}\|$ when $\mathcal{X}, \mathcal{Y} \in \mathbf{CCA}_{\mathcal{G}}^{\Sigma}$.

Proof. It suffices to factor h via $\iota_{\mathcal{X}}: \mathcal{X} \rightarrow \|\mathcal{X}\|$ and an embedding $e: \|\mathcal{X}\| \hookrightarrow \|\mathcal{Y}\|$, since by Theorem 5.26 and Lemma 5.30 these two maps are necessary unique.

By naturality of ι w.r.t. the functor $\|(-)\|$ (Corollary 5.24) the following solid square commutes:

$$\begin{array}{ccc} \mathcal{X} & \xrightarrow{h} & \|\mathcal{Y}\| \\ \iota_{\mathcal{X}} \downarrow & & \downarrow \iota_{\|\mathcal{Y}\|} \\ \|\mathcal{X}\| & \xrightarrow{\|h\|} & \|\|\mathcal{Y}\|\| \end{array}$$

(Note: A dotted arrow labeled $\tilde{\iota}_{\|\mathcal{Y}\|}$ also points from $\|\mathcal{X}\|$ to $\|\|\mathcal{Y}\|\|$.)

Lemma 5.25 tells us that $\iota_{\|\mathcal{Y}\|}$ is an isomorphism, thus it has an inverse (the dotted arrow in the above diagram). Hence

$$\iota_{\|\mathcal{Y}\|}^{-1} \circ \|h\| \circ \iota_{\mathcal{X}} = h,$$

and therefore we can take $e := \iota_{\|\mathcal{Y}\|}^{-1} \circ \|h\|$. □

5.3.4 Cardinality of union of contractions

The next lemmas will establish that $\|\mathbf{CCA}_{\mathcal{G}}^{\Sigma}\|/\cong$ is a set, and thus that $\biguplus_{\|\mathcal{X}\| \in \|\mathbf{CCA}_{\mathcal{G}}^{\Sigma}\|/\cong} \|\mathcal{X}\|$ is also a set.

The first lemma will be used to bound the number of behaviourally distinct (as in not-behaviourally-equivalent) cells in a contraction: contractions have no two cells with essentially the same successors.

Lemma 5.34

Let $\mathcal{X} = (X, \langle \chi, \gamma \rangle) \in \mathbf{CCA}_{\mathcal{G}}^{\Sigma}$. Let $x, x' \in X$ be such that for all $m \in \mathcal{M}$,

$$N_m \otimes x = N_m \otimes x'$$

$$\gamma_m \otimes x = \gamma_m \otimes x'$$

then $x \stackrel{\text{BE}}{\approx} x'$.

Proof. We construct a coalgebra \mathcal{X}' by using a construction similar to a contraction (c.f. Definition 5.13), but we only contract pairs of cells of the form $m \otimes x$ and $m \otimes x'$ into an equivalence class (for every $m \in \mathcal{M}$).

So let $\mathcal{X}' \stackrel{\text{def}}{=} (X/\sim, \langle \chi', \gamma' \rangle)$ where $\sim \subseteq X \times X$ is the reflexive and symmetric closure of the relation $m \otimes x \sim m \otimes x'$ (for all $m \in \mathcal{M}$). Denote equivalence classes of \sim by $|y|$ (for all $y \in X$), and note that equivalence classes contain either 1 or 2 cells.

Define χ' as follows:

$$n \otimes' |y| \stackrel{\text{def}}{=} |n \otimes y|$$

for all $n \in \mathcal{M}$ and $y \in X$. To see that this definition is independent of the representative of $|y|$, there are only two cases to consider:

1. $|y| = \{y\}$, in which independence of the choice of representative is trivial.

2. $|y| = |m \otimes x|$ for some $m \in \mathcal{M}$. Then well-definedness follows from the following observation:

$$n \otimes' |m \otimes x| = |n \otimes m \otimes x| = |(n \bullet m) \otimes x| = |(n \bullet m) \otimes x'| = |n \otimes m \otimes x'| = n \otimes' |m \otimes x'|.$$

To see that χ' is a cowriter-comonad-coalgebra, observe that:

$$n \otimes' \ell \otimes' |y| \stackrel{\text{def}}{=} n \otimes' |\ell \otimes y| \stackrel{\text{def}}{=} |n \otimes \ell \otimes y| = |(n \bullet \ell) \otimes y| \stackrel{\text{def}}{=} (n \bullet \ell) \otimes |y|.$$

The definition of γ' follows the same approach as $\|\gamma\|$:

$$\gamma'_{|y|} \stackrel{\text{def}}{=} \gamma_y$$

where:

$$u: S^{N_{|y|}} \rightarrow S^{N_y}$$

$$u(k)(m) \stackrel{\text{def}}{=} k(|m|).$$

Note that u is only well-defined when $N_{|y|} \leq N_y$ (in $\text{PartOrd}(N)$, Definition 2.2). That is, if $|m|_y = |n|_y$ implies $|m|_{|y|} = |n|_{|y|}$. To prove this, assume $|m|_y = |n|_y$, that is, $m \otimes y = n \otimes y$. Then applying $|(-)|$ to both sides gives $|m \otimes y| = |n \otimes y|$, which by definition of χ' is equivalent to $m \otimes' |y| = n \otimes' |y|$, which is equivalent to $|m|_{|y|} = |n|_{|y|}$.

Finally, we let $\phi: \mathcal{X} \rightarrow \mathcal{X}'$ be the quotient map $\phi: y \mapsto |y|$. That this is a well-defined coalgebra homomorphism follows directly from the definitions of χ' and γ' (and Theorem 5.9). We further observe that $\phi(x) = \phi(x')$, thus we conclude that $x \stackrel{\text{BE}}{\approx} x'$. \square

We proceed now by counting cardinalities. Fix an arbitrary monoid \mathcal{M} , neighbourhood $N \subseteq \mathcal{M}$ and state set S , and denote their cardinalities as:

$$\mu = \text{card}(\mathcal{M})$$

$$\nu = \text{card}(N)$$

$$\sigma = \text{card}(S).$$

The following lemma gives a (very coarse) upper bound on the number of behaviourally distinct cells a CA may have.

Lemma 5.35

For a CA $\mathcal{X} = (X, \langle \chi, \gamma \rangle) \in \mathbf{CCA}_{\mathcal{G}}^{\Sigma}$, the number of distinct equivalence classes of X under the BE-relation is upper bounded by

$$\kappa = 2^{\mu} \mu(2^{2^{\nu}})(\sigma^{\sigma^{\nu}}).$$

Proof. From Lemma 5.34 it follows that a sufficient condition for cells $x, x' \in X$ to be BE is having:

- The same number of \mathcal{M} -reachable successors.
- The same $N_m \otimes x = N_m \otimes x'$ and $\gamma_m \otimes x = \gamma_m \otimes x'$ for all $m \in \mathcal{M}$.

The number of disjoint partitions of N is upper bounded by $2^{2^{\nu}}$, the number of subsets of subsets of N . Hence, there are at most $2^{2^{\nu}}$ choices for N_x for a given cell $x \in X$.

Recall that $\gamma_x \in S^{S^{N_x}}$, and since $\text{card}(N_x) \leq \nu$, the number of choices of γ_x is bounded by $\sigma^{\sigma^{\nu}}$.

From the above two observations, it follows that there are at most $(2^{2^\nu})(\sigma^{\sigma^\nu})$ distinct choices of pairs (N_x, γ_x) . A cell x has at least 1 successor (itself) and at most μ successors. For a chosen number $1 \leq \alpha \leq \mu$ of successors of x , there are $\alpha(2^{2^\nu})(\sigma^{\sigma^\nu})$ ways to choose the data inside the successors. But by the cowriter-comonad structure (Eq. (12)), the successors of successors of x are also successors of x , hence this fixes the successors of all cells $m \otimes x$ automatically as well. Since there are 2^μ choices for α , and since $\alpha \leq \mu$, this means that the total number of ways to configure x and all reachable successors is bounded by

$$\kappa \stackrel{\text{def}}{=} 2^\mu \mu (2^{2^\nu})(\sigma^{\sigma^\nu}).$$

□

Note that a cell x may still have a myriad of (possibly independent) parents that it cannot see. The trick is that those parents themselves are also among those κ choices.

Corollary 5.36

If $\mathcal{X} \in \mathbf{CCA}_G^\Sigma$ then $\|\mathcal{X}\|$ has at most κ cells.

Proof. $\|\mathcal{X}\|$ has no distinct behaviourally equivalent cells, and there are at most κ behaviourally distinct cells by Lemma 5.35. □

Corollary 5.37

There are, up to isomorphism, at most 2^κ distinct contractions of coalgebras in \mathbf{CCA}_G^Σ .

Proof. This is the number of subsets of κ , the upper bound of the total number of possible cells. □

Let $A = \{\|\mathcal{X}\| \mid \mathcal{X} \in \mathbf{CCA}_G^\Sigma\} / \cong$.

Theorem 5.38

$T \stackrel{\text{def}}{=} \uplus_{\|\mathcal{X}\| \in A} \|\mathcal{X}\|$ is a set, and so is $\|T\|$. Therefore $\|\mathcal{T}\|$ is the terminal coalgebra in \mathbf{CCA}_G^Σ .

Proof. The fact that T is a set follows from Corollary 5.37, which gives an upper bound on the cardinality of T .

Every $\mathcal{X} \in \mathbf{CCA}_G^\Sigma$ has a unique morphism to $\|\mathcal{X}\|$ (Theorem 5.26), and clearly there is the inclusion map $i: \|\mathcal{X}\| \hookrightarrow T$. Thus we have the composition:

$$\mathcal{X} \xrightarrow{\iota_{\mathcal{X}}} \|\mathcal{X}\| \xrightarrow{\iota_{\|\mathcal{X}\|}} \|\|\mathcal{X}\|\| \xrightarrow{\|i\|} \|T\|$$

Theorem 5.33 asserts that this coalgebra homomorphism is unique as well. □

5.4 Summary of coalgebra homomorphisms and behavioural equivalence

The concept of behavioural equivalence seems appropriate for describing the distinct atomic building blocks of a non-uniform CA (uniform CA have only one such building block), but not for reasoning about the global dynamics of a CA.

Using Theorem 5.38 as a practical characterisation of coalgebra homomorphisms, we have developed the concept of the *contraction* $\|\mathcal{X}\|$ of a CA \mathcal{X} . In a contraction, all BE cells are merged into one cell. There

is a unique map from a CA \mathcal{X} to any contracted CA (Theorem 5.33), and also at most one map between contractions (Lemma 5.30). The reader acquainted with Homotopy Type Theory may see an analogy with propositional truncations: for every type A there is a unique map to its propositional truncation $\|A\|$, and there is at most one map between two propositional truncations (as they are propositions).

We used contractions to construct the terminal coalgebra $\|\mathcal{T}\|$ (Theorem 5.38). This coalgebra encodes a complete picture of all equivalence classes of BE cells of CA (defined for a given signature (\mathcal{M}, N, S)), as it contains exactly one cell for each class.

A modal logic obtained via the predicate-lifting framework for coalgebras cannot logically distinguish BE cells, and therefore it cannot distinguish a CA from its contraction. However, contractions preserve few of the global dynamics of a CA. The contraction of a uniform CA is particularly trivial, as it is a single cell (Lemma 5.11). For example, Example 5.12 shows that CGOL is contracted to the single-celled constant-dead-state CA; the logic would not be able to differentiate between the computationally universal dynamics of GCOL and a trivial constant state.

The conclusion is that a logic capable of describing spatio-temporal patterns in the trace sequence of a CA is necessarily also capable of distinguishing behaviourally equivalent cells. A logic generated by the predicate lifting framework is not able to do the latter. This justifies manually defining of a modal logic with semantics based on trace sequences, which will be the topic for the remainder of this thesis.

6 Logic for computation traces

Section 5 argued that the standard framework for defining modal logics for coalgebras due to Schröder [41] and Klin [22] is not appropriate for describing the global dynamics of a CA.

This section will manually construct an elementary modal logic for describing the states of cells in a trace of a CA. Like how Kripke modal logics are evaluated from the perspective of points, our logic will be evaluated locally from the perspective of cells. Given a coalgebra $\mathcal{X} = (X, \langle \chi, \gamma \rangle)$ on a monoid \mathcal{M} , we can assign a state to cells via a configuration $c_0: X \rightarrow S$. Therefore it seems natural to take states as out analog of propositional letters and traces $[c_0, c_1, \dots]$ (which are entirely determined by c_0) as analogs to valuations. Hence our atomic formulas will be individual states $s \in S$ and models will be a coalgebras \mathcal{X} paired with an initial configuration c_0 .

In Kripke modal logics, when evaluating a formula at a point x , it is possible to specify that a subformula should be evaluated at a related point y via a diamond \diamond_R if xRy . In the context of CA, there are two relation-like structures we can define diamonds for: the CA’s connections between cells are encoded in χ , as well as the chain-shaped trace $c_0 \rightarrow c_1 \rightarrow c_2 \rightarrow c_3 \rightarrow \dots$. Therefore we will introduce “space” diamonds \diamond_m for $m \in \mathcal{M}$ to traverse the CA’s graph, as well as a “time” diamond \varkappa to evaluate a subformula at the next configuration in a trace. In combination with the familiar Boolean logic connectives \wedge, \vee, \neg and \rightarrow , we can encode statements such as “If my $m \in \mathcal{M}$ neighbour has state s_0 , then I will have state s_1 in the next timestep”:

$$(\diamond_m s_0) \rightarrow \varkappa s_1.$$

This section will incrementally build a basic modal language \mathcal{L}_{stg} that implements the ideas below. We will first construct the space diamonds, then add the time diamonds. Thereafter we discuss the possibilities of performing uniform substitution, and this section concludes with proving a collection of practical reasoning tools.

6.1 Space modalities

A configuration of a CA $\mathcal{X} = (X, \langle \chi, \gamma \rangle) \in \mathbf{CCA}_{\mathbf{C}}$ is an assignment $c: X \rightarrow S$ of each cell $x \in X$ to a state $s \in S$. This is essentially a Kripke model:

- The propositional letters are S , the worlds are X .
- There is a successor relation R_m for every $m \in \mathcal{M}$, namely xR_mx' iff $x' = m \otimes x$.
- The valuation $V: X \rightarrow \wp(S)$ is just the configuration: $V(x) \stackrel{\text{def}}{=} \{c(x)\}$. Note that that every world is assigned *exactly one* letter in a valuation c .

This perspective allows us to introduce a propositional “space modalities” modal logic \mathcal{L}_s with modalities \diamond_m for every $m \in \mathcal{M}$, that can express “space” properties of a CA. This logic cannot yet express the dynamics of a CA (which follow from the local rules γ_x), but it is a foundation to which we will incrementally extend with more symbols in subsequent subsections.

6.1.1 Grammar

The grammar of \mathcal{L}_s is as follows:

$$\Psi ::= \neg\psi \mid \psi_1 \vee \psi_2 \mid \perp \mid s \mid \diamond_m$$

where $m \in \mathcal{M}$ and $s \in S$. The notational abbreviations \wedge , \rightarrow , \leftrightarrow and \top are defined as usual. The states s act as logic constants; the intuition is that a formula “ s ” is satisfied at a cell x in configuration c iff x is assigned state s (i.e., iff $c(x) = s$).

6.1.2 Semantics

As in Kripke semantics, we distinguish between static *frames* and *models*, which are assignments of values to “variables”. In the context of CA, the natural choice of these variables are the initial states of the cells.

Definition 6.1: Frame

A *frame* $\mathbb{X} = (\mathcal{M}, N, S, X, \langle \chi, \gamma \rangle)$ is a coalgebra $\mathcal{X} = (X, \langle \chi, \gamma \rangle) \in \mathbf{CCA}_{\mathbf{G}}$ together with the monoid \mathcal{M} , neighbourhood $N \subseteq \mathcal{M}$ and state set S on which \mathcal{X} is defined.

Definition 6.2: Model

A *model* (\mathbb{X}, c_0) is a frame $\mathbb{X} = (\mathcal{M}, N, S, \mathcal{X})$ together with an initial configuration $c_0: X \rightarrow S$.

Definition 6.3: Space modalities semantics

Define *modal satisfaction* of a formula $\psi \in \Psi$ in a model (\mathbb{X}, c_0) at a cell $x \in X$, to be the proposition (which is strictly either true or false):

$$\mathbb{X}, c_0, x \models \psi,$$

whose truth is given recursively as follows:

$$\begin{aligned} \mathbb{X}, c_0, x \models \perp & \quad \text{is always false} \\ \mathbb{X}, c_0, x \models s & \quad \text{iff } c_0(x) = s \\ \mathbb{X}, c_0, x \models \neg\psi & \quad \text{iff } \mathbb{X}, c_0, x \not\models \psi \\ \mathbb{X}, c_0, x \models \psi_1 \vee \psi_2 & \quad \text{iff } (\mathbb{X}, c_0, x \models \psi_1 \text{ or } \mathbb{X}, c_0, x \models \psi_2) \\ \mathbb{X}, c_0, x \models \diamond_m\psi & \quad \text{iff } \mathbb{X}, c_0, m \otimes x \models \psi \end{aligned}$$

where $s \in S$, $m \in \mathcal{M}$, $\psi, \psi_1, \psi_2 \in \Psi$ and “ $\mathbb{X}, c_0, x \not\models \psi$ ” denotes “*it is false that* $\mathbb{X}, c_0, x \models \psi$ ”.

A formula that holds in all choices of the initial configuration is called *valid*:

Definition 6.4: Validity of a formula

A formula ψ is *valid* in frame \mathbb{X} at a cell $x \in X$, denoted $\mathbb{X}, x \models \psi$, if $\mathbb{X}, c_0, x \models \psi$ for all $c_0: X \rightarrow S$. A formula ψ is *valid* in \mathbb{X} itself, denoted $\mathbb{X} \models \psi$, if ψ is valid at all cells in X .

The *logic of a frame*, denoted $\text{Log}(\mathbb{F})$, is the set of all formulas that a frame validates. The *logic of x* for a cell $x \in X$, denoted $\text{Log}(x)$, is the set of formulas that are valid at x .

Stacking space diamonds, e.g., $\diamond_m \diamond_n$, semantically corresponds to monoid multiplication $n \bullet m$. An example usage of this property is that a logic for finitely generated monoids only need a diamond for every generating element. Formally, the property is:

Lemma 6.5

For all $\psi \in \Psi$, the formula $\diamond_m \diamond_n \psi$ is semantically equivalent to $\diamond_{n \bullet m} \psi$.

Proof. Directly from the definition of the semantics:

$$\begin{aligned} \mathbb{X}, c_0, x \models \diamond_m \diamond_n \psi &\text{ iff } \mathbb{X}, c_0, m \otimes x \models \diamond_n \psi \\ &\text{ iff } \mathbb{X}, c_0, n \otimes m \otimes x \models \psi \\ &\text{ iff } \mathbb{X}, c_0, (n \bullet m) \otimes x \models \psi \\ &\text{ iff } \mathbb{X}, c_0, x \models \diamond_{n \bullet m} \psi \end{aligned}$$

□

6.1.3 Expressivity of space modalities

The logic \mathcal{L}_s can express both basic properties of the input configuration $c: X \rightarrow S$ and some properties of χ . As for a basic example:

Example 6.6. Consider the 2D grid (i.e., the monoid \mathbb{Z}^2), and denote the modalities corresponding to $(1, 0), (-1, 0), (0, 1), (0, -1) \in \mathbb{Z}^2$ with $\diamondright, \diamondleft, \diamondup, \diamonddown$ respectively. As cells, simply use $X = \mathbb{Z}^2$ as well. Use the colours green, orange and purple as states (i.e., $S = \{s_{\text{or}}, s_{\text{gr}}, s_{\text{pu}}\}$).




Figure 18: Part of a configuration $S \rightarrow \mathbb{Z}^2$ using the three-colour state set $S = \{s_{\text{or}}, s_{\text{gr}}, s_{\text{pu}}\}$.

In the configuration c of Figure 18, and the frame $\mathbb{F} = (\mathbb{Z}^2, S, \overrightarrow{\mathbb{Z}^2})$, the point $x = (0, 0)$ satisfies and unsatisfies

respectively the following formulas:

$$\begin{aligned}
& \mathbb{F}, c, x \vDash s_{or} \\
& \mathbb{F}, c, x \not\vDash s_{gr} \\
& \mathbb{F}, c, x \vDash \diamond \diamond s_{or} \\
& \mathbb{F}, c, x \not\vDash \diamond s_{pu} \\
& \mathbb{F}, x, x \vDash (\diamond s_{gr}) \wedge ((\diamond s_{pu}) \vee (\neg \diamond s_{gr})).
\end{aligned}$$

By using the notion of *validity*, the logic \mathcal{L}_s can already separate cells that are BE, implying \mathcal{L}_s is more expressive than the notion of BE. A frame \mathbb{F} *validates* $\psi \in \Psi$, denoted as $\mathbb{F} \vDash \psi$, if $\mathbb{F}, c, x \vDash \psi$ for all $c: X \rightarrow S$ and $x \in X$. The *logic of a frame*, denoted $\text{Log}(\mathbb{F})$, is the set of formulas that the frame validates. Analogously a cell $x \in X$ may validate a ψ if $\mathbb{F}, c, x \vDash \psi$ for all $c: X \rightarrow S$, which in turn defines the logic of a cell $\text{Log}(x)$.

Lemma 6.7

There exist $\mathcal{X}, \mathcal{Y} \in \mathbf{CCA}_{\mathbf{G}}$ and $x \in X$ and $y \in Y$ such that $x \overset{\text{BE}}{\approx} y$ but $\text{Log}(x) \neq \text{Log}(y)$.

Proof. Refer back to Example 5.12, which showed that all cells in CGOL are BE to the CA with a single cell and the local rule that always outputs the “dead” state \circ . Denote this single-cell-constant-dead CA by $\mathbb{1}_{\circ}$, with the single cell $*$, state set $\{\bullet, \circ\}$ and $\gamma_* = \text{const}_{\circ}$. The CA $\mathbb{1}_{\circ}$ uses the same signature $(\mathbb{Z}^2, N, S = \{\bullet, \circ\})$ as $\overrightarrow{\text{CGOL}}$ (coalgebra homomorphisms only exist between CA defined on the same signature) where $N = \{(i, j) \mid i, j \in \{-1, 0, 1\}\}$. Now a configuration on $\mathbb{1}_{\circ}$ is of type $c: \{*\} \rightarrow \{\bullet, \circ\}$ and all neighbours of $*$ (which is only $*$ itself!) all receive the same state in any configuration. Hence

$$\circ \rightarrow \diamond \circ \in \text{Log}(*) .$$

In CGOL this formula is not *valid* at any cell $(x, y) \in \mathbb{Z}^2$, seen (for example) by the configuration $c(x, y) = \circ$ and $c(x, y + 1) = \bullet$. (Note that the formula is *satisfied* in the *model* of $\overrightarrow{\text{CGOL}}$'s frame with a configuration c s.t. $c(x, y) = c(x, y + 1)$). But this does not hold for all configurations on $\overrightarrow{\text{CGOL}}$. \square

Example 6.8 (Neighbourhood quotients). The logic \mathcal{L}_s can express properties of the neighbourhood quotients. For example the formula

$$\diamond_n s_0 \leftrightarrow \diamond_m s_0 \tag{46}$$

is valid at a cell x iff $|n|_x = |m|_x$ (assuming $S \not\cong \mathbb{1}$). This can be seen as follows: if $n \otimes x \neq m \otimes x$, then one can define a configuration such that $c(n \otimes x) \neq c(m \otimes x)$, in which case (46) is not satisfied.

6.2 Time modalities

The logic \mathcal{L}_s cannot express any properties about the dynamics of a CA; neither local dynamic properties such as the local rules γ_x (the semantics c.f. Definition 6.3 are independent of γ_x) nor properties about the global dynamics of a CA. Recall from the definition of a general CA (Definition 4.6) that the global dynamics are a sequence of configurations $(c_i)_{i \in \mathbb{N}}$ (where $c_i: X \rightarrow S$) computed by recursively applying the global rule $G_{\mathbb{X}}$ (Eq. (28)) to an configuration c_0 .

In order to express the evolution of the configurations over time we introduce a new *time modality* \varkappa to the logic \mathcal{L}_s , obtaining the “space & time modalities” logic \mathcal{L}_{st} , with semantics:

Definition 6.9: Space & time modalities semantics

Define *modal satisfaction* of a formula $\psi \in \Psi$ in a model (\mathbb{X}, c_0) at a cell $x \in X$, to be the proposition (which is strictly either true or false):

$$\mathbb{X}, c_0, x \models \psi.$$

Unlike in the semantics of \mathcal{L}_s , modal satisfaction depends recursively on the truth in models $(\mathbb{X}, G_{\mathbb{X}}^t(c_0))$. Thus we define the truth of propositions

$$\mathbb{X}, G_{\mathbb{X}}^t(c_0), x \models \psi \quad (47)$$

for all $t \in \mathbb{N}$. Since c_0 is usually clear from context, we ease the notation by writing c_t for $G_{\mathbb{X}}^t(c_0)$. Now the truth of (47) can be given recursively as:

$$\begin{aligned} \mathbb{X}, c_t, x \models \perp & \quad \text{is always false} \\ \mathbb{X}, c_t, x \models s & \quad \text{iff } c_t(x) = s \\ \mathbb{X}, c_t, x \models \neg\psi & \quad \text{iff } \mathbb{X}, c_t, x \not\models \psi \\ \mathbb{X}, c_t, x \models \psi_1 \vee \psi_2 & \quad \text{iff } (\mathbb{X}, c_t, x \models \psi_1 \text{ or } \mathbb{X}, c_t, x \models \psi_2) \\ \mathbb{X}, c_t, x \models \diamond_m \psi & \quad \text{iff } \mathbb{X}, c_t, m \otimes x \models \psi \\ \mathbb{X}, c_t, x \models \varepsilon \psi & \quad \text{iff } \mathbb{X}, c_{t+1}, m \otimes x \models \psi \end{aligned}$$

The only difference between the semantics of \mathcal{L}_{st} and \mathcal{L}_s (Definition 6.3) are the ε case, and generalising from c_0 to all c_t .

6.2.1 Expressivity of time modalities

The logic \mathcal{L}_{st} gives a positive answer to our research question, but only for finite patterns:

Example 6.10. The logic \mathcal{L}_{st} can already express any finite spatio-temporal pattern of a trace of a rooted (Definition 3.8) CA. To see this, let $\mathbb{X} = (\mathcal{M}, N, S, X, \langle \chi, \gamma \rangle) \in \mathbf{CCA}_{\mathbf{G}}$ be a CA rooted in $r \in X$. For any cell $x \in X$ and state $s \in S$, we can express that x has state s in the $(t-1)^{\text{th}}$ configuration of a trace with initial configuration $c_0: X \rightarrow S$ as follows. First we rewrite $x = m \otimes r$ (for some $m \in \mathcal{M}$, which is possible since r is the root), and then we note that $G_{\mathbb{X}}^t(x) = s$ holds iff $\mathbb{X}, c_0, r \models \varepsilon^t \diamond_m s$. For patterns involving multiple cells and/or multiple points in time we simply use conjunctions.

We will now show examples of properties related to the dynamics of a CA that \mathcal{L}_{st} can express but \mathcal{L}_s cannot.

Example 6.11. \mathcal{L}_{st} can logically distinguish a CA whose local rules are the constant- \circ -output from a CA whose local rules are the constant- \bullet -output; the latter validates $\varepsilon\bullet$ but not $\varepsilon\circ$, for the former it is exactly the other way around.

Lemma 6.12

If N is finite, then \mathcal{L}_{st} can express the mapping of γ_x piecewise for every input $z: N_x \rightarrow S$.
If S is also finite, then \mathcal{L}_{st} can express γ_x entirely in one formula.

Proof. The formula for $\gamma_x(z)$ is:

$$\left(\bigwedge_{n \in N} \diamond_n z|n|_x \right) \rightarrow \varepsilon \gamma_x(z). \quad (48)$$

The reading is as follows: $\bigwedge_{n \in N} \diamond_n z(|n|_x)$ is satisfied in \mathbb{X}, c_t, x iff $c_t \upharpoonright (N \otimes x) = z \circ |(-)|_x$, so it encodes *if the neighbours of x have the states as encoded in z* . The conclusion $\exists \gamma_x(z)$ encodes that $\mathbb{X}, c_{t+1}, x \models \gamma_x(z)$, i.e., that x will have state $\gamma_x(z)$ in the next timestep.

If S is finite, then S^{N_x} is also finite, so one can take the conjunction over all $z \in S^{N_x}$ of equations (48). \square

Example 6.13 (Local rule). Suppose we have a CA \mathcal{X} on the monoid freely generated over $\{L, R\}$ with the finite neighbourhood $N = \{\mathbf{1}, L, R\}$ and state set S . Given a cell x with a full neighbourhood $N_x = \{N\}$, we can express γ_x on an input $z: N \rightarrow S$ as:

$$\bigwedge_{z: N \rightarrow S} (z(\mathbf{1}) \wedge \diamond_L z(L) \wedge \diamond_R z(R)) \rightarrow \exists \gamma_x(z \setminus x).$$

6.3 Global variables and uniform substitution

In the context of Kripke frames, a key strength of normal modal logics is the ability to characterise a class \mathcal{C} of frames by a finite set of formulae $\Psi_{\mathcal{C}}$ such that $\mathbb{X} \in \mathcal{C}$ iff $\mathbb{X} \models \psi$ for all $\psi \in \Psi_{\mathcal{C}}$. For example, $p \rightarrow \diamond p$ characterises frames based on a reflexive relation. This is possible because of the propositional letters $p \in Phi$ in Kripke semantics: these act as global variables and can be *uniformly substituted* by arbitrary subformulas without changing the validity of the outer formula.

This raises the question whether one can also introduce a form of uniform substitution in the logic for CA, as this would be helpful to finitely classify classes of CA.

The answer is positive, but we cannot simply use the states $s \in S$ as propositional letters. To see this, observe that Kripke propositional letters are indistinguishable and whose semantics are entirely determined by the valuation $V: W \rightarrow \wp(\Phi)$ added to a frame $\mathbb{X} = (W, R \subseteq W \times W)$ to construct a model (\mathbb{X}, V) . However, states $s \in S$ are not necessarily indistinguishable and may partially have a constant value that the choice of c_0 (our analog of V) cannot change. The simplest example is a constant local rule $\gamma_x = \text{const}_{s_0}$ for some $s_0 \in S$. A model may freely choose c_0 , but $c_t + 1(x) = s_0$ for all $t \in \mathbb{N}$. Hence x validates $\exists s_0$. Obviously we cannot substitute s_0 by any arbitrary formula, such as “ s_1 ” for a state $s_1 \neq s_0 \in S$!

The solution is straightforward: only consider the states in c_0 , which do not have any constant value. We introduce new 0-ary symbols v_s for all $s \in S$ and extend the semantics with the rule:

$$\mathbb{X}, G_{\mathbb{X}}^t(c_0), x \models v_s \text{ iff } c_0(x) = s. \quad (49)$$

When c_0 is clear from context, this will be written as $\mathbb{X}, c_t, x \models v_s$ iff $c_0(x) = s$. We denote extension of \mathcal{L}_{st} with global variables as \mathcal{L}_{stg} .

One complication remains: a normal modal logic formula ψ in a Kripke frame becomes either true or false after giving a valuation V . A formula can be seen as a function $\wp(\Phi)^W \rightarrow \mathbb{2}$ that maps a valuation to a true/false truth value. But for CA, the truth value of a formula ψ might also depend on the point in time in which it is evaluated. It follows that one cannot substitute a variable v_s (which does not depend on the point in time in which it is evaluated, i.e., $\mathbb{X}, c_0, x \models v_s$ iff $\mathbb{X}, c_t, x \models v_s$ for all $t \in \mathbb{N}$) by a formula ψ whose truth does change over time: that is, if $\mathbb{X}, c_{t_1}, x \models \psi$ but $\mathbb{X}, c_{t_2}, x \not\models \psi$ for some $t_1, t_2 \in \mathbb{N}$. However, when restricting to time invariant formulas, uniform substitution *does* remain sound for CA.

Definition 6.14: Time invariant formula

A formula $\psi \in \Psi$ is *time invariant* if, for all $c_0: X \rightarrow S$:

$$\mathbb{X}, x \models \psi \text{ iff for all } t \in \mathbb{N}: \mathbb{X}, x \models \exists^t \psi.$$

Denote the time invariant formulas in Ψ of a cell by $\text{TI}_\Psi(x)$, and the formulas that are time invariant in all cells in a CA \mathcal{X} by $\text{TI}_\Psi(\mathcal{X})$.

For a frame \mathbb{X} based on \mathcal{X} we will also write $\text{TI}_\Psi(\mathbb{X})$ to denote the same set as $\text{TI}_\Psi(\mathcal{X})$.

Lemma 6.15: Valid formulas are time invariant

If $\mathbb{X}, x \models \psi$ then ψ is time-invariant.

Proof. Take any $c_0: X \rightarrow S$ and any $t \in \mathbb{N}$. We need to show that $\mathbb{X}, c_0, x \models z^t \psi$, but this holds iff $\mathbb{X}, c_t, x \models \psi$, which indeed follows from the assumption $\mathbb{X}, x \models \psi$ (since $c_t = G_{\mathbb{X}}^t(c_0)$ is just another choice of model). \square

Note that we cannot restrict the grammar to only time invariant formulas, since time invariance depends on the CA. For example, recall **CGOL** and $\mathbb{1}_\circ$ from the proof of Lemma 6.7 and Example 5.12. The formula $\circ \rightarrow z\circ$ is not valid in **CGOL** but it is if valid (and hence time invariant) in $\mathbb{1}_\circ$.

Further note that the absence of constants $s \in S$ in a formula ψ is not required for time invariance. For example, $s \vee \top$ is always valid and $s \wedge \perp$ is always invalid, and hence both are time invariant.

Lemma 6.16: Uniform substitution of time invariant formulas

If $\mathbb{X} \models \psi$ (and $S \not\cong \mathbb{1}$), then for all $s \in S$ and $\phi \in \text{TI}_\Psi(\mathbb{X})$ it holds also that $\mathbb{X} \models \psi [\phi/v_S]$.

It is most convenient to first prove the following auxiliary lemma:

Lemma 6.17

For all $\theta \in \Psi$ containing the term v_S , $s \neq s' \in S$, $\phi \in \text{TI}_\Psi(\mathbb{X})$ and $c_0: X \rightarrow S$, if we define

$$c'_0 \stackrel{\text{def}}{=} \begin{cases} s & \mathbb{X}, c_0, x \models \phi \\ s' & \text{otherwise} \end{cases},$$

then for all $t \in \mathbb{N}$ and $x \in X$:

$$\mathbb{X}, c'_t, x \models \theta \text{ iff } \mathbb{X}, c_t, x \models \theta [\phi/v_S].$$

where we abbreviated $G_{\mathbb{X}}^t(c'_0)$ as c'_t .

Proof. Proceed by structural induction on θ . We assumed that θ contains v_S , so the cases $\theta = s \in S$ and $\theta = \perp$ do not apply.

Case $\theta = v_S$: By the semantics of v_S , $\mathbb{X}, c'_t, x \models v_S$ iff $c'_0(x) = s$. By definition of c'_0 , this holds iff $\mathbb{X}, c_0, x \models \phi$. The latter holds iff $\mathbb{X}, c_t, x \models v_S [\phi/v_S]$, as required.

Note that this first case is the reason for introducing v_S . If we would use the atoms $s \in S$ instead as propositional letters to be substituted, then we would only be able to show the equality for $t = 0$ in the $\theta = s$ case. However, we need it to show for all $t \in \mathbb{N}$ in order to have a sufficiently strong induction hypothesis for the $\theta = z\psi$ case.

Case $\theta = \psi_1 \vee \psi_2$:

Induction hypothesis: for all $x \in X$, $t \in \mathbb{N}$ and $j \in \{1, 2\}$ it holds that:

$$\mathbb{X}, c'_t, x \models \psi_j \text{ iff } \mathbb{X}, c_t, x \models \psi_j [\phi/v_S].$$

Observe that $(\psi_1 \vee \psi_2) [a/b] = (\psi_1 [a/b]) \vee (\psi_2 [a/b])$ for all strings a and b . Therefore we can compute:

$$\begin{aligned} \mathbb{X}, c'_t, x \models \psi_1 \vee \psi_2 & \\ \text{iff } (\mathbb{X}, c'_t, x \models \psi_1 \text{ or } \mathbb{X}, c'_t, x \models \psi_2) & \\ \text{iff } (\mathbb{X}, c_t, x \models \psi_1 [\phi/v_s] \text{ or } \mathbb{X}, c_t, x \models \psi_2 [\phi/v_s]) & \quad // \text{ By the induction hypothesis.} \\ \text{iff } \mathbb{X}, c_t, x \models (\psi_1 \vee \psi_2) [\phi/v_s]. & \end{aligned}$$

Case $\theta = \neg\psi$:

Induction hypothesis: for all $x \in X$ and $t \in \mathbb{N}$ it holds that:

$$\mathbb{X}, c'_t, x \models \psi \text{ iff } \mathbb{X}, c_t, x \models \psi [\phi/v_s].$$

Now we compute:

$$\begin{aligned} \mathbb{X}, c'_t, x \models \neg\psi \text{ iff } \mathbb{X}, c'_t, x \not\models \psi & \\ \text{iff } \mathbb{X}, c_t, x \not\models \psi [\phi/v_s] & \\ // \text{ By the induction hypothesis and the fact that a formula is either satisfied or not.} & \\ \text{iff } \mathbb{X}, c_t, x \models \neg(\psi [\phi/v_s]) & \\ \text{iff } \mathbb{X}, c_t, x \models (\neg\psi) [\phi/v_s]. & \end{aligned}$$

Case $\theta = \exists\psi$:

The induction hypothesis is as in the \neg case.

$$\begin{aligned} \mathbb{X}, c'_t, x \models \exists\psi \text{ iff } \mathbb{X}, c'_{t+1}, x \models \psi & \\ \text{iff } \mathbb{X}, c_{t+1}, x \models \psi [\phi/v_s] & \quad // \text{ By the induction hypothesis.} \\ \text{iff } \mathbb{X}, c_t, x \models \exists(\psi [\phi/v_s]) & \\ \text{iff } \mathbb{X}, c_t, x \models (\exists\psi) [\phi/v_s]. & \end{aligned}$$

Case $\theta = \diamond_m\psi$ ($m \in \mathcal{M}$):

The induction hypothesis is again as in the \neg case.

$$\begin{aligned} \mathbb{X}, c'_t, x \models \diamond_m\psi \text{ iff } \mathbb{X}, c'_t, m \otimes x \models \psi & \\ \text{iff } \mathbb{X}, c_t, m \otimes x \models \psi [\phi/v_s] & \quad // \text{ By the induction hypothesis.} \\ \text{iff } \mathbb{X}, c_t, x \models \diamond_m(\psi [\phi/v_s]) & \\ \text{iff } \mathbb{X}, c_t, x \models (\diamond_m\psi) [\phi/v_s]. & \end{aligned}$$

□

Proof of Lemma 6.16. Assume that $\mathbb{X} \models \psi$ and that $\phi \in \text{TI}_\Psi(\mathcal{X})$. It is to show that $\mathbb{X} \models \psi [\phi/v_s]$. To this end, take any arbitrary $c_0: X \rightarrow S$ and $x \in X$, then it suffices to show that $\mathbb{X}, c_0, x \models \psi [\phi/v_s]$, i.e., that

$$\mathbb{X}, c_0, x \models \psi [\phi/v_s] \tag{50}$$

But define c'_0 as in Lemma 6.17, then $\mathbb{X} \models \psi$ implies that $\mathbb{X}, c'_0, x \models \psi$. Now Eq. (50) follows directly from Lemma 6.17. □

We conclude with an example application of global variables and time-invariant uniform substitution.

Example 6.18. Given a CA $\mathbb{X} = (\mathcal{M}, N, S, \mathcal{X})$ with at least two distinct states $s_0, s_1 \in S$, the formula

$$\mathbf{v}_{s_{s_0}} \rightarrow \diamond_m \mathbf{v}_{s_{s_0}} \quad (51)$$

is valid at a cell $x \in X$ iff $x = m \otimes x$ (otherwise we could unsatisfy with a configuration c_0 s.t. $c(x) = s_0$ and $c(m \otimes x) = s_1$). Uniform substitution of time-invariant formulas allows to substitute $\mathbf{v}_{s_{s_0}}$ by any other time invariant formula. This is not surprising: if x satisfies ψ and $x = m \otimes x$, then $m \otimes x$ must clearly also satisfy ψ !

6.4 Logical reasoning rules

One of the main applications of defining a logic is the possibility to perform symbolic reasoning. Logical reasoning can, for example, be used to prove properties about the mathematical structures that form the semantics. This requires reasoning rules that are *sound*, i.e., that are consistent the semantics. We will present several such rules in the style of natural deduction, many of which will be used in proofs in later sections.

In this section, we use $\mathbb{X} = (\mathcal{M}, N, S, X, \langle \chi, \gamma \rangle)$ as an arbitrary frame. A rule about validity of a formula in \mathbb{X} will be presented as follows:

Definition 6.19: Sound rule

A rule with n -premises $\mathbb{X} \models \phi_1, \mathbb{X} \models \phi_2, \dots, \mathbb{X} \models \phi_n$ and conclusion $\mathbb{X} \models \psi$,

denoted $\frac{\mathbb{X} \models \phi_1 \quad \mathbb{X} \models \phi_2 \quad \dots \quad \mathbb{X} \models \phi_n}{\mathbb{X} \models \psi}$, is *sound* if $\mathbb{X} \models \psi$ whenever $\mathbb{X} \models \phi_i$ for all $1 \leq i \leq n$.

Double bars, as in $\frac{A}{B}$, denote the pair of rules $\frac{A}{B}$ and $\frac{B}{A}$.

A rule may apply both to space diamonds \diamond_m (for all choices of $m \in \mathcal{M}$) as well as for time diamonds \diamond , in which case the notation \diamond will be used. All instances of \diamond within a rule must be substituted with the same diamond.

We will also write metamathematical statements (such as $s \in S$) as premises in rules. Technically this means that we define a family of rules over all choices of the metavariables in the metamathematical premises.

Be aware that we give no guarantee of completeness: there may exist valid formulas that cannot be derived from the rules in this section.

The first two rules are very versatile; they show that we can perform basic propositional logic in \mathcal{L}_{stg} .

Lemma 6.20

For any propositional logical tautology τ in propositional variables x_1, x_2, \dots, x_n , we have the rule:

$$\frac{\psi_1, \psi_2, \dots, \psi_n \in \Psi}{\mathbb{X} \models \tau [\psi_1/x_1] [\psi_2/x_2] \dots [\psi_n/x_n]} \text{Taut}$$

Proof. The semantics for \mathcal{L}_{stg} concerning the symbols \neg, \vee and \perp (and hence also \rightarrow, \wedge and \top) are equivalent to those of propositional logic.

In any model given by a configuration $c_0: X \rightarrow S$, the semantics of the formulas ψ_i evaluate to **tt** or to **ff**. So the semantics in c_0 of a formula ϕ containing ψ_i are (in c_0) the same as $\phi [\phi_i/\top]$ or $\phi [\phi_i/\perp]$ respectively. Since τ is a tautology, it is true for any substitution of the variables x_1, x_2, \dots, x_n by \top and \perp . \square

Lemma 6.21: Modus Ponens

The following rule is sound:

$$\frac{\mathbb{X} \models \psi \rightarrow \phi \quad \mathbb{X} \models \psi \quad \psi, \phi \in \Psi}{\mathbb{X} \models \phi} \text{MP}$$

Proof. Take any configuration $c_0: X \rightarrow S$. Then for all $x \in X$ we have (unfolding the definition of \rightarrow) $\mathbb{X}, c_0, x \models \neg\psi \vee \phi$ and $\mathbb{X}, c_0, x \models \psi$. This combination is only possible if $\mathbb{X}, c_0, x \models \phi$. Since c_0 and x were arbitrary we find $\mathbb{X} \models \phi$. \square

Remark 6.22. Many common reasoning strategies can be obtained as a combination of a propositional tautology and MP. For example, \wedge -elimination can be encoded as a tautology $A \wedge B \rightarrow A$. Thus, if one has derived that $\mathbb{X} \models \phi_A \wedge \phi_B$ for some $\phi_A, \phi_B \in \Psi$, we first obtain the tautology $\mathbb{X} \models \phi_A \wedge \phi_B \rightarrow \phi_A$ via Taut (Lemma 6.20), and conclude $\mathbb{X} \models \phi_A$ via MP (Lemma 6.21).

MP gives a tool for implication elimination, but implication introduction is also often useful. The following lemma shows that we can introduce implications by the usual method of propositional logic.

Lemma 6.23: Implication introduction

If one can prove $\mathbb{X}, c_0, x \models \phi$ when assuming $\mathbb{X}, c_0, x \models \psi$, then $\mathbb{X}, c_0, x \models \psi \rightarrow \phi$ holds in general.

Proof. Assume we can derive $\mathbb{X}, c_0, x \models \phi$ from $\mathbb{X}, c_0, x \models \psi$.

By definition of “ \rightarrow ” we know $\mathbb{X}, c_0, x \models \psi \rightarrow \phi$ holds iff $\mathbb{X}, c_0, x \models \neg\psi \vee \phi$. Assume $\mathbb{X}, c_0, x \not\models \neg\psi$, then we need to show that $\mathbb{X}, c_0, x \models \phi$. But $\mathbb{X}, c_0, x \models \psi$ holds iff $\mathbb{X}, c_0, x \models \psi$ (by Lemma 6.29). By the premise of the lemma, this implies $\mathbb{X}, c_0, x \models \phi$. The latter was to be shown. \square

The next two rules follow from the fact that, in any configuration, a cell has exactly one state.

Lemma 6.24: Single state

The following rule is sound:

$$\frac{s_0, s_1 \in S \quad s_0 \neq s_1}{\mathbb{X} \models s_0 \rightarrow \neg s_1} \text{Det1}$$

Proof. Take any $x \in X$ and $c_0: X \rightarrow S$. Suppose $s_0 \neq s_1 \in S$ and that $\mathbb{X}, c_0, c \models s_0$. Then, by the semantics of the formula s_0 , it must be that $c_0(x) = s_0$. Thus $c_0(x) \neq s_1$, which implies $\mathbb{X}, c_0, x \not\models s_1$. Hence $\mathbb{X}, c_0, x \models s_0 \rightarrow \neg s_1$. \square

The following rule shows that we cannot distinguish \diamond from $\neg\diamond\neg$; this is unlike in Kripke modal logics, where $\square \stackrel{\text{def}}{=} \neg\diamond\neg$ does have different semantics than \diamond .

Lemma 6.25: Diamonds are boxes

For any $\psi \in \Psi$ the following rule is sound:

$$\frac{\mathbb{X} \models \diamond\psi}{\mathbb{X} \models \neg\diamond\neg\psi} \text{Det2}$$

Proof. For any $c_0: X \rightarrow S$ and $x \in X$ it holds:

$$\begin{aligned} \mathbb{X}, c_0, x \models \varepsilon\psi &\text{ iff } \mathbb{X}, G_{\mathbb{X}}(c_0), x \models \psi \\ &\text{ iff } \mathbb{X}, G_{\mathbb{X}}(c_0), x \not\models \neg\psi \\ &\text{ iff } \mathbb{X}, c_0, x \not\models \varepsilon\neg\psi \\ &\text{ iff } \mathbb{X}, c_0, x \models \neg\varepsilon\neg\psi. \end{aligned}$$

(Note that, by definitions of the semantics of ε , $(\mathbb{X}, G_{\mathbb{X}}(c_0), x \models \neg\psi)$ and $(\mathbb{X}, c_0, x \models \varepsilon\neg\psi)$ have the same truth value).

The case where $\diamond = \diamond_m$ for some $m \in \mathcal{M}$ is analogous (but changing x to $m \otimes x$ instead of c_0 to $G_{\mathbb{X}}(c_0)$). \square

Lemma 6.26: Generalisation

The following rule is sound:

$$\frac{\mathbb{X} \models \psi \quad \psi \in \Psi}{\mathbb{X} \models \diamond\psi} \text{Gen}$$

Proof. In case $\diamond = \varepsilon$, take any $c_0: X \rightarrow S$ and $x \in X$. We need to show that $\mathbb{X}, c_0, x \models \varepsilon\psi$, i.e., that $\mathbb{X}, G_{\mathbb{X}}(c_0), x \models \psi$. But $G_{\mathbb{X}}(c_0)$ is just a configuration, so this follows from $\mathbb{X} \models \psi$. \square

Lemma 6.27: K-axiom

The following rule is sound:

$$\frac{s_0, s_1 \in S}{\mathbb{X} \models \diamond(\mathbf{v}_{s_0} \rightarrow \mathbf{v}_{s_1}) \rightarrow ((\diamond\mathbf{v}_{s_0}) \rightarrow (\diamond\mathbf{v}_{s_1}))} \text{K}$$

Proof. We first consider the case where $\diamond = \diamond_m$ for some $m \in \mathcal{M}$. Assume any $s_0, s_1 \in S$, $c_0: X \rightarrow S$ and $x \in X$. Applying Lemma 6.23, we assume that $\mathbb{X}, c_0, x \models \diamond_m(\mathbf{v}_{s_0} \rightarrow \mathbf{v}_o)$ and we need to show that $\mathbb{X}, c_0, x \models (\diamond_m\mathbf{v}_{s_0}) \rightarrow (\diamond_m\mathbf{v}_{s_1})$. We apply Lemma 6.23 again, and assume that $\mathbb{X}, c_0, x \models \diamond_m\mathbf{v}_{s_0}$; it remains to show that $\mathbb{X}, c_0, x \models \diamond_m\mathbf{v}_{s_1}$. Our two assumptions are equivalent to:

$$\mathbb{X}, c_0, m \otimes x \models \mathbf{v}_{s_0} \rightarrow \mathbf{v}_{s_1}$$

and

$$\mathbb{X}, c_0, m \otimes x \models \mathbf{v}_{s_0}.$$

respectively. We can apply MP on these latter forms to obtain $\mathbb{X}, c_0, m \otimes x \models \mathbf{v}_{s_1}$. This is equivalent to $\mathbb{X}, c_0, x \models \diamond_m\mathbf{v}_{s_1}$, which was to be shown.

The case $\diamond = \varepsilon$ is analogous, except that we replace c_0 by $G_{\mathbb{X}}(c_0)$ instead of x by $m \otimes x$. \square

Space-time independence Suppose it is 11:55 AM. Then you could (1) walk downstairs and wait there for 5 minutes, and (2) wait for 5 minutes and thereafter go downstairs. But both options make you end up on the ground floor at 12:00. The key insight is that the stairs lead to the same place at any point in time.

Lemma 6.28: Space-Time Independence (STI)

For any $m \in \mathcal{M}$ and $\psi \in \Psi$, the following rules are sound:

$$\frac{\mathbb{X} \vDash \diamond_m \varkappa \psi}{\mathbb{X} \vDash \varkappa \diamond_m \psi} \text{STI}$$

Proof. Directly from the semantics:

$$\begin{aligned} \mathbb{X}, c_0, x \vDash \diamond_m \varkappa \psi &\text{ iff } \mathbb{X}, c_0, m \otimes x \vDash \varkappa \psi \\ &\text{ iff } \mathbb{X}, G_{\mathbb{X}}(c_0), m \otimes x \vDash \psi \\ &\text{ iff } \mathbb{X}, G_{\mathbb{X}}(c_0), x \vDash \diamond_m \psi \\ &\text{ iff } \mathbb{X}, c_0, x \vDash \varkappa \diamond_m \psi. \end{aligned}$$

□

Lemma 6.29: Double Negation

For any $\psi \in \Psi$, the following rule is sound:

$$\frac{}{\mathbb{X} \vDash \neg \neg \psi \leftrightarrow \psi} \text{DN}$$

Proof. This follows from the definition of the semantics: a formula in a model (\mathbb{X}, c_0) at a cell $x \in X$ is either satisfied or not. Thus, if $\mathbb{X}, c_0, x \not\vDash \phi$ is *not* true (for any $\phi \in \Phi$), then the only remaining possibility is that $\mathbb{X}, c_0, x \vDash \phi$. So we observe that $\mathbb{X}, c_0, x \vDash \neg \neg \psi$ iff $\mathbb{X}, c_0, x \not\vDash \neg \psi$. Now $\mathbb{X}, c_0, x \vDash \neg \psi$ holds iff $\mathbb{X}, c_0, x \not\vDash \psi$, thus $\mathbb{X}, c_0, x \not\vDash \neg \psi$ holds iff $\mathbb{X}, c_0, x \vDash \psi$. □

6.5 Summary of basic modal logic

This section constructed the basic modal logic \mathcal{L}_{stg} , a cellular automaton analogue of Kripke modal logics. The frames $\mathbb{X} = (\mathcal{M}, N, S, \mathcal{X})$ in this logic are CA coalgebras in $\mathbf{CCA}_{\mathcal{G}}$, together with the signature (\mathcal{M}, N, S) on which they are defined. Models are frames paired with an initial configuration $c_0: X \rightarrow S$, which define a trace $[c_0, c_1, c_2, \dots]$ that assigns a state to every cell at every point in time. These states $s \in S$ act as atomic formulas, which are true at a cell $x \in X$ in a configuration c_0 iff c_0 assigns state s to x (i.e., $c_0(x) = s$). Paired with the Boolean connectives \wedge, \vee, \neg and \rightarrow , these atomic formulas form the basic building blocks of the language \mathcal{L}_{stg} . But by encapsulating subformulas in space \diamond_m and time \varkappa diamonds, we allow to evaluate subformulas at cells $m \otimes x$ for $m \in \mathcal{M}$, as well as in a the next configuration c_1 in the trace. Diamonds can be nested, which allow to specify any reachable point in the CA's graph and any future configuration c_t .

Unlike in Kripke modal logics, the atomic formulas $s \in S$ cannot be freely substituted by other formulas. We introduced a global variable v_s for every $s \in S$ that can be freely substituted by *time-invariant* formulas, while preserving formula satisfaction. Luckily, many other reasoning tools familiar from Kripke modal logics, such as the K -axiom, modus ponens, generalisation of valid formulas and the validity of all Boolean tautologies do carry over to \mathcal{L}_{stg} .

The language \mathcal{L}_{stg} can express certain properties of a CA, such as the local rule (provided that the neighbourhood N is finite). But it also has its limits. Every diamond connects a cell in a specific point in the trace to a single other specific cell in a specific point in the trace. This makes it impossible to express properties such as “for all $m \in \mathcal{M}$, it holds that ...” or “for all c_t with $t \in \mathbb{N}$ it holds that ...”. The next section will extend the language with infinitary disjunctions and (derived) infinitary conjunctions, which can encode such quantifications and greatly increases the expressivity of the language.

7 Logic extension: infinitary disjunctions

The expressivity of \mathcal{L}_{stg} can significantly be improved after extending it with infinitary disjunctions, which are practically existential quantifiers.

7.1 Intuition and motivation

We first motivate the need of this extension with three examples. Thereafter we present the formal extensions to the grammar and semantics.

Example 7.1 (Expressing all spatio-temporal patterns in a trace). Example 6.10 showed how \mathcal{L}_{st} can express that a given cell has a certain state in a certain point in time in a trace. With the help of a finite number of conjunctions we could describe any finite patterns this way. With infinitary conjunctions (which automatically are as the duals $\bigwedge = \neg \bigvee \neg$ of infinitary disjunctions) we can describe any spatio-temporal pattern in a trace of a rooted CA.

Example 7.2 (Expressing “forever after” and “eventually”). There are many examples of properties that can be expressed with the notions of *forever* and *eventually*. For example, a CA may have a “dead” state $s_d \in S$ that a cell cannot leave once entered. This state satisfies the statement “*if in state s_d , then forever after in state s_d* ”. Logically, “*forever after ψ holds*” means that $\mathbb{z}^n \psi$ holds for all choices of $n \in \mathbb{N}$ repetitions of the time diamond. Thus our statement would be:

$$s_d \rightarrow \bigwedge_{n \in \mathbb{N}} \mathbb{z}^n s_d. \quad (52)$$

As for another example, some CA have dynamics such that every trace will converge to a fixpoint in which all cells have a “dead” (quiescent) state s_d . However, the exact number of timesteps before a cell dies may vary with the initial configuration chosen. We only know there exists a choice of a finite number of timesteps before this occurs, which is an infinite disjunction over all number of timesteps. Hence we can express “*eventually state s_d will be attained*” as:

$$\bigvee_{\ell \in \mathbb{N}} \mathbb{z}^\ell s_d. \quad (53)$$

This does not yet express that the cell will remain in s_d . However, combining these two examples, we can express “*every cell will eventually obtain state s_d , and once in s_d , it will retain this state forever after*”:

$$\bigvee_{\ell \in \mathbb{N}} \mathbb{z}^\ell \left(\bigwedge_{n \in \mathbb{N}} \mathbb{z}^n s_d \right). \quad (54)$$

Example 7.3 (Quiescent states). A *quiescent state* is a state $s_q \in S$ such that if a cell and all its neighbours have state s_q , then the cell will stay in state s_q the next timestep. Quiescent states are useful for simulating a CA with in infinite grid on a real computer: this is possible if the number of non-quiescent cells is finite, since the quiescent cells do not need to be stored or computed explicitly.

If we allow infinitary disjunctions, then we obtain infinitary conjunctions as their dual ($\bigwedge = \neg \bigvee \neg$). With these conjunctions we can express “*all neighbours have state s_q* ” in a formula as $\bigwedge_{n \in \mathbb{N}} \diamond_n s_q$. The statement “*the cell will stay in state s_q* ” means that state s_q will occur the next timestep, i.e., the formula $\mathbb{z}s_q$. So quiescence of s_q can be expressed with

$$\left(\bigwedge_{n \in \mathbb{N}} \diamond_n s_q \right) \rightarrow \mathbb{z}s_q. \quad (55)$$

Finally, “*there exists a quiescent state*” holds as soon as (55) is valid for some choice of $s_q \in S$. This is essentially a disjunction over all states:

$$\bigvee_{s \in S} \left(\left(\bigwedge_{n \in N} \diamond_n s \right) \rightarrow \exists s \right). \quad (56)$$

7.2 Formal definitions of infinitary disjunctions

We extend the language \mathcal{L}_{stg} with a new constructor $\bigvee_{i \in I} \psi_i$ for every index set I and set of subformulas $\{\psi_i\}_{i \in I}$. We interpret $\bigvee_{i \in I} \psi_i$ the same way as the (possible infinite) formula obtained as the disjunction of ψ_i over all $i \in I$. While results in very large and very general language \mathcal{L}'_{CA} , we will primarily continue to work with a subset $\mathcal{L}_{\text{CA}} \subseteq \mathcal{L}'_{\text{CA}}$ of a few specific choices of sets of formulas $\{\psi_i\}_{i \in I}$.

Definition 7.4: Grammar of \mathcal{L}'_{CA}

The well-formed formulas Ψ that form the grammar of \mathcal{L}'_{CA} are given by:

$$\Psi ::= \perp \mid \neg\psi \mid \psi_1 \vee \psi_2 \mid \diamond_m \psi \mid \exists\psi \mid s \mid \mathbf{v}_s \mid \bigvee_{i \in I} \psi_i$$

where $\psi, \psi_1, \psi_2 \in \Psi$, $m \in \mathcal{M}$ and $\{\psi_i\}_{i \in I} \subseteq \Psi$.

Definition 7.5: Semantics of \mathcal{L}'_{CA}

The semantics of \mathcal{L}'_{CA} are as for \mathcal{L}_{stg} (see Definition 6.9 and Eq. (49)), extended with the rule:

$$\mathbb{X}, G_{\mathbb{X}}^t(c_0), x \models \bigvee_{i \in I} \psi_i \text{ iff there exists an } i \in I \text{ such that } \mathbb{X}, G_{\mathbb{X}}^t(c_0), x \models \psi_i.$$

Remark 7.6. We will write $\bigwedge_{i \in I} \psi_i$ to denote $\neg \bigvee_{i \in I} \neg \psi_i$; observe we can deduce the familiar semantics:

$$\mathbb{X}, G_{\mathbb{X}}^t(c_0), x \models \bigwedge_{i \in I} \psi_i \text{ iff for all } i \in I \text{ it holds that } \mathbb{X}, G_{\mathbb{X}}^t(c_0), x \models \psi_i.$$

Remark 7.7. All the deduction rules of Section 6.4 are also sound for \mathcal{L}'_{CA} . This is straightforward to see: the given proofs also work for \mathcal{L}'_{CA} .

Remark 7.8. Uniform substitution of time-invariant variables (Lemma 6.16 and Lemma 6.17) also holds for \mathcal{L}'_{CA} . The only required modification in the proofs is adding the $\theta = \bigvee_{i \in I} \psi_i$ case to the structural induction on θ in Lemma 6.17 (with the induction hypothesis assumed for each ψ_i). However, this case is analogous to the binary disjunction $\theta = \psi_1 \vee \psi_2$ case since we have a similar distribution of the substitution: $(\bigvee_{i \in I} \psi_i) [\phi/\mathbf{v}_s] = \bigvee_{i \in I} (\psi_i [\phi/\mathbf{v}_s])$.

We will be working with \mathcal{L}_{CA} , which uses the following infinitary disjunctions:

Definition 7.9: Subset \mathcal{L}_{CA} of \mathcal{L}'_{CA}

\mathcal{L}_{CA} is \mathcal{L}'_{CA} but with restricted choices of the sets of subformulas used in infinitary disjunctions.

The first few of those sets are “decorated” version of a fixed formula $\psi \in \Psi$, and defined for each chose of ψ . Let $A \stackrel{\text{def}}{=} \{\mathcal{M}, \mathcal{M} \setminus \mathbf{1}, N, N \setminus \mathbf{1}\}$, then the sets of decorations for ψ over signature (\mathcal{M}, N, S) in \mathcal{L}_{CA} are:

- The set of steps in an arbitrary direction:

$$\{\diamond_m \mid m \in Z\} \quad (57)$$

for any choice of $Z \in A$. Disjunctions over this set will be denoted as $\bigvee_{m \in Z} \diamond_m \psi$.

- The set of paths over Z for any choice of $Z \in A$ and $\psi \in \Psi$:

$$\text{Paths}(Z) \stackrel{\text{def}}{=} \{\diamond_{m_1} \diamond_{m_2} \dots \diamond_{m_k} \mid m_i \in Z, k \in \mathbb{N}\}. \quad (58)$$

Notation: $\bigvee_{\pi \in \text{Paths}(Z)} \pi \psi$.

- For any choice of Z and $m \in Z$ the set of repetitions of \diamond_m :

$$\{\diamond_m^k \mid k \in \mathbb{N}\}. \quad (59)$$

Notation: $\bigvee_{k \in \mathbb{N}} \diamond_m^k \psi$.

- The set of repetitions of the time modality:

$$\{\varepsilon^k \mid k \in \mathbb{N}\}. \quad (60)$$

Notation: $\bigvee_{k \in \mathbb{N}} \varepsilon^k \psi$.

\mathcal{L}_{CA} may also include disjunctions over set of states $s \in S$ and the set of global variables v_s with $s \in S$. This are sets of atomic formulas (and do not depend on ψ).

Finally, \mathcal{L}_{CA} also has all variants of the following *until modality*:

$$\phi \mathbf{U} \psi \stackrel{\text{def}}{=} \bigvee_{k \in \mathbb{N}} (\varepsilon^k \psi) \wedge \bigwedge_{0 \leq j \leq k-1} \varepsilon^j \psi. \quad (61)$$

(Intuitively: “ ψ will eventually hold after k timesteps, and if $k > 0$ then ϕ holds after any $0, 1, \dots, k-1$ timesteps.”)

Remark 7.10 (Cardinality of formulas in \mathcal{L}_{CA}). The cardinality on the size of a set of formulas in infinitary disjunctions in \mathcal{L}_{CA} is bounded by the cardinality κ of the signature (\mathcal{M}, N, S) . The effective length of formulas (when unfolding the infinitary disjunctions into actual infinite sequences of disjunctions) can be any finite power of κ since we can nest infinitary disjunctions (e.g., $\bigvee_{m \in \mathcal{M}} \bigvee_{n \in \mathcal{M}} \bigvee_{\ell \in \mathcal{M}} \diamond_m \diamond_n \diamond_\ell \psi$ is effectively a disjunction over $\text{card}(\mathcal{M})^3$ subformulas).

Remark 7.11 (Until modality). Many variants of the “until modality” are possible: e.g. one could analogously define it for repetitions of a space diamonds \diamond_m , or drop the constraint that ψ must eventually hold (then it is possible that ϕ holds forever and ψ never).

Also interesting is to define the following “until modality” for paths. Paths work a bit different than mere repetition, so we define $\phi \mathbf{U}_{\text{Paths}(Z)} \psi$ to denote that every path first visits a sequence of cells where ϕ hold, then a cell where ψ holds, and continues without further constraints. Since we do not both control the length and direction of the path, it may keep circling in the ϕ -region; the solution is simply to use the “until”-variant where the end condition ψ is not guaranteed to hold.

Formally, let \sqsubset denote the is-a-strict-prefix relation (e.g., $\diamond_{m_1} \diamond_{m_2} \diamond_{m_3} \sqsubset \diamond_{m_1} \diamond_{m_2} \diamond_{m_3} \diamond_{m_4}$), then we can define:

$$\phi \mathbf{U}_{\text{Paths}(Z)} \psi \stackrel{\text{def}}{=} \bigwedge_{\pi \in \text{Paths } Z} \left(\left(\pi \phi \wedge \bigwedge_{\theta \sqsubset \pi} \theta \phi \right) \vee \bigvee_{\theta \sqsubset \pi} \left(\theta \psi \wedge \bigwedge_{\omega \sqsubset \theta} \omega \phi \right) \right).$$

The interpretation is easiest to explain when we take the CA on the grid \mathbb{Z}^2 with two colours $S = \{\circ, \bullet\}$ and pick $\phi := \bullet$ and $\psi := \circ$. Then $\pi \circ \wedge \bigwedge_{\theta \sqsubseteq \pi} \theta \bullet$ means “all cells on the path π are \circ ”, and $\bigvee_{\theta \sqsubseteq \pi} (\theta \bullet \wedge \bigwedge_{\omega \sqsubseteq \theta} \omega \circ)$ means “there exists a prefix of π that first only visits \circ cells and ends in a cell in state \bullet .” Thus $\circ \cup_{\text{Paths}(\mathbb{Z}^2)} \bullet$ reads “I am in a blob of \circ cells, and this blob is surrounded by \bullet cells.” See Figure 19 for a visualisation.

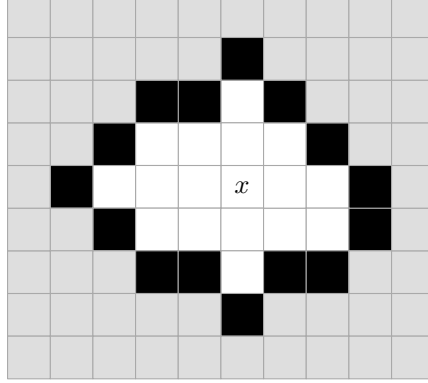


Figure 19: Example configuration of black and white states ($\{\circ, \bullet\}$) on the grid \mathbb{Z}^2 , the cell x satisfies $\circ \cup_{\text{Paths}(\mathbb{Z}^2)} \bullet$. The formula does not concern the states of gray cells, which can therefore be arbitrary.

7.3 Examples of expressible properties

The expressivity of \mathcal{L}_{CA} is best illustrated by concrete examples. In particular, we will explore certain formulas that are valid on a frame iff that frame has a specific property.

Most examples only apply to *rooted* (recall from Definition 3.8) frames with at least two states ($\text{card}(S) \geq 2$). We will refer to these CA as *nice* CA.

Rootedness ensures that there is a cell that can describe properties of every other cell, i.e., global properties. At least two states are needed because a “CA” with $\text{card}(S) = 1$ has no dynamics at all, and the only non- \perp atomic formulas are states; hence all cells satisfy the same formulas at any timestep and cannot be told apart.

Local rules Lemma 6.12 showed how to express local rules for finite neighbourhoods N as formulas; with infinitary conjunctions, the same proof generalises to local rules on infinite N .

Eventually everyone We can express that, after a given number of timesteps, all cells satisfy a given formula ψ .

Lemma 7.12

A nice CA \mathbb{X} validates

$$\text{EvEv}(\psi) \stackrel{\text{def}}{=} \bigvee_{k \in \mathbb{N}} \mathbb{X}^k \bigwedge_{m \in \mathcal{M}} \diamond_m \psi \quad (62)$$

iff for every initial configuration $c_0: X \rightarrow S$ there exists a $t \in \mathbb{N}$ s.t. $\mathbb{X}, G_{\mathbb{X}}^t, x \models \psi$ for all $x \in X$.

Proof. Let $r \in X$ be the root. Then in a given model c_0 we have that $\mathbb{X}, c_0, r \models \text{EvEv}(\psi)$ iff there exists a $t \in \mathbb{N}$ s.t. $\mathbb{X}, G_{\mathbb{X}}^t(c_0), r \models \bigwedge_{m \in \mathcal{M}} \diamond_m \psi$. The latter holds iff $\mathbb{X}, G_{\mathbb{X}}^t(c_0), m \otimes r \models \psi$ for all $m \in \mathcal{M}$. Since r is the root, $\mathcal{M} \otimes r = X$, thus this holds iff all cells $x \in X$ satisfy ψ at time t . \square

Example 7.13 (Eventually everyone is dead). Suppose we have a “dead” state $s_d \in S$. Then $\bigwedge_{t \in \mathbb{N}} \mathbb{Z}^t s_d$ expresses that a cell is dead and stays dead forever after. Hence

$$\text{EvEv}(\bigwedge_{t \in \mathbb{N}} \mathbb{Z}^t s_d) = \bigvee_{k \in \mathbb{N}} \bigwedge_{m \in \mathcal{M}} \bigwedge_{t \in \mathbb{N}} \diamond_m \mathbb{Z}^t s_d \quad (63)$$

expresses (in a nice CA) that “*eventually, every cell will be forever dead*”.

The trace will become periodic A CA $\mathbb{X} = (\mathcal{M}, N, S, \mathcal{X})$ might be such that for every initial configuration $c_0: X \rightarrow S$, the resulting trace $[c_0, c_1, c_2, c_3, \dots]$ (computed c.f. (28)) becomes periodic. More precisely, there exists (for every c_0) a starting timesteps $k \in \mathbb{N}$ and a period $p \in \mathbb{N}$ such that $c_{k+i} = c_{k+i+p} = c_{k+i+2p} = c_{k+i+3p} = \dots$ for all $i \in \mathbb{N}$. An intuitive example is a traffic light: it is a periodic 3-cell CA with states {off, green, orange, red}.

Lemma 7.14

A nice CA \mathbb{X} validates

$$\bigvee_{p \in \mathbb{N}} \text{EvEv}(\psi(p)) \quad (64)$$

where

$$\psi(p) \stackrel{\text{def}}{=} \bigwedge_{k \in \mathbb{N}} \bigwedge_{s \in S} \mathbb{Z}^k (s \rightarrow \mathbb{Z}^p s) \quad (65)$$

iff it all its traces eventually become periodic.

Proof. A trace is periodic with period p iff every cell has a sequence of states that repeats with period p . Thus, by Lemma 7.12, it suffices to show that $\psi(p)$ expresses “*my states are repeating periodically with period p* ”.

Assume an arbitrary initial configuration $c_0: X \rightarrow S$ and write c_i for $G_{\mathbb{X}}^i(c_0)$. We need to show that $\mathbb{X}, c_i, x \models \psi(p)$ iff x ’s states in the suffix $[c_i, c_{i+1}, c_{i+2}, \dots]$ of the trace repeat with period p .

For the first direction, assume that

$$\mathbb{X}, c_i, x \models \psi(p). \quad (66)$$

Take an arbitrary point $0 \leq z \leq p - 1$ in the period and let $s_0 := c_{i+z}(x)$, then we need to show that $\mathbb{X}, c_{i+z+np}, x \models s_0$ for all $n \in \mathbb{N}$. We proceed with induction on n . Note that the base case holds trivially by definition of s_0 .

For the inductive step, assume the induction hypothesis $\mathbb{X}, c_{i+z+np}, x \models s_0$. If we choose $k := z + np$ and $s := s_0$ in (66) we obtain $\mathbb{X}, c_{i+z+pn}, x \models s_0 \rightarrow \mathbb{Z}^p s_0$. Using the inductive hypothesis and modus ponens (Lemma 6.21) we find $\mathbb{X}, c_{i+z+pn}, x \models \mathbb{Z}^p s_0$. This holds iff $\mathbb{X}, c_{i+z+p(n+1)}, x \models s_0$, which was the goal of the inductive step.

For the reverse direction, assume $c_{i+z} = c_{i+z+np}$ for all $0 \leq z \leq p - 1$ and $n \in \mathbb{N}$. It is to show that $\mathbb{X}, c_i, x \models \bigwedge_{k \in \mathbb{N}} \bigwedge_{s \in S} \mathbb{Z}^k (s \rightarrow \mathbb{Z}^p s)$. It suffices to show that, for an arbitrary $k \in \mathbb{N}$ and $s_0 \in S$ it holds that $\mathbb{X}, c_{i+k}, x \models s_0 \rightarrow \mathbb{Z}^p s_0$. We use the method from propositional logic of “ \rightarrow ”-introduction (Lemma 6.23). So assume $\mathbb{X}, c_{i+k}, x \models s_0$; this holds iff $c_{i+k}(x) = s_0$. But by our periodicity assumption we know $c_{i+k} = c_{i+k+p}$, so $c_{i+k+p}(x) = s_0$ as well. The latter is equivalent to $\mathbb{X}, c_{i+k+p}, x \models s_0$, which holds iff $\mathbb{X}, c_{i+k+p}, x \models \mathbb{Z}^p s_0$. \square

Undecidability Two of the previous examples are actually already show that \mathcal{L}_{CA} can encode undecidable properties of CA. On the one hand, this highlights the expressivity of the language; on the other hand, it seems not practically possible to compute whether a formula is valid on a given CA.

Corollary 7.15

Frame validity of a formula in \mathcal{L}_{CA} is undecidable. (I.e., there exist a frame \mathbb{X} and a formula ψ in \mathcal{L}_{CA} such that it is undecidable whether $\mathbb{X} \models \psi$).

Proof. We give two proofs that use a different classical result.

A CA \mathbb{X} is *nilpotent* if there exists a $t \in \mathbb{N}$ and a configuration $c_d: X \rightarrow S$ such that c_d is a fixpoint of $G_{\mathbb{X}}$, and $G_{\mathbb{X}}^t(c_0) = c_d$ for any initial configuration $c_0: X \rightarrow S$. Kari proved in 1992 [19] that it is undecidable for a CA on \mathbb{Z} whether it is nilpotent. This cited paper also shows that in a nilpotent CA, c_d necessarily assigns all cells the same state, say $s_d \in S$. Example 7.13 exactly expresses that the configuration const_{s_d} will be reached; we only need to add “there exists a $s_d \in S$ such that eventually all cells will be in state s_d ”. Hence a CA is nilpotent iff it validates:

$$\bigvee_{s_d \in S} \text{EvEv} \left(\bigwedge_{k \in \mathbb{N}} \mathbb{z}^k s_d \right) \stackrel{\text{def}}{=} \bigvee_{s_d \in S} \bigvee_{t \in \mathbb{N}} \mathbb{z}^t \left(\bigwedge_{m \in \mathcal{M}} \left(\diamond_m \bigwedge_{k \in \mathbb{N}} \mathbb{z}^k s_d \right) \right).$$

A CA is *periodic* if there exists an $p \in \mathbb{N}$ such that $G_{\mathbb{X}}^p = \text{Id}$. Since the trace are just the outputs of $G_{\mathbb{X}}^t(c_0)$ for all t , this definition is equivalent to periodicity of the trace. Lemma 7.14 shows how to express that the trace will *eventually* become periodic. Removing the disjunction $\bigvee_{k \in \mathbb{N}} \mathbb{z}^k$ over the starting point in time of the periodicity (it should start immediately), we obtain the formula that expresses periodicity:

$$\bigvee_{p \in \mathbb{N}} \bigwedge_{m \in \mathcal{M}} \diamond_m \left(\bigwedge_{t \in \mathbb{N}} \bigwedge_{s \in S} \mathbb{z}^t (s \rightarrow \mathbb{z}^p s) \right).$$

Whether a CA on \mathbb{Z} is periodic is undecidable (proven by Kari and Ollinger in 2008 [21]).

Note that the classical undecidability results remain when restricting to nice CA. \mathbb{Z} is a rooted grid, and both problems are trivial to decide when a CA has only a single state (hence it is not possible that only the subclass of CA on \mathbb{Z} with $\text{card}(S) = 1$ is undecidable). \square

Steady states A cell x may, in a given model (\mathbb{X}, c_0) , eventually reach a steady state. That is, $c_i(x) = c_{i+k}(x)$ for some $i \in \mathbb{N}$ and all $k \in \mathbb{N}$. This is simply expressed as $\mathbb{X}, c_0, x \models \bigvee_{i \in \mathbb{N}} \bigvee_{s \in S} \mathbb{z}^i (\bigwedge_{k \in \mathbb{N}} \mathbb{z}^k s)$. More interestingly, we can also show that a cell (or even stronger, no cell) will never reach a steady state.

Lemma 7.16: Never reach a steady state

A cell $x \in X$ will never reach a steady state in model (\mathbb{X}, c_0) iff

$$\mathbb{X}, c_0, x \models \bigwedge_{s \in S} \left(s \rightarrow \bigvee_{i \in \mathbb{N}} \mathbb{z}^i \neg s \right). \quad (67)$$

It follows that no cell in any model on a nice CA will reach a steady state iff

$$\mathbb{X} \models \bigwedge_{m \in \mathcal{M}} \bigwedge_{s \in S} \left(s \rightarrow \bigvee_{i \in \mathbb{N}} \mathbb{z}^i \neg s \right). \quad (68)$$

Proof sketch. A cell x will reach a steady iff it will reach some state s s.t. all future states of x are also s . Equivalently, a state is not steady iff there exist a future state of x that is different. Thus it suffices to show that x has no steady state iff every state will eventually be succeeded by another state, which is exactly expressed by (67). \square

The local rule of a CA may ignore the state of a given neighbour at path $n \in N$. I.e., $\gamma_x(z) = \gamma_x(z')$ for all $z, z': \mathcal{M} \otimes x \rightarrow S$ that only differ on input $n \otimes x$.

Lemma 7.17

Suppose $N = \{n_0, n_1, \dots, n_k\}$ is a finite neighbourhood. The local rule γ of a frame \mathbb{X} always ignores neighbour $n_0 \in N$ iff

$$\mathbb{X} \models \bigwedge_{s_0 \in S} \bigwedge_{s_1 \in S} \dots \bigwedge_{s_k \in S} \bigwedge_{s' \in S} [(((\diamond_{n_0} s_0) \wedge (\diamond_{n_1} s_1) \wedge \dots \wedge (\diamond_{n_k} s_k)) \rightarrow \Sigma s') \rightarrow [((\diamond_{n_1} s_1) \wedge \dots \wedge (\diamond_{n_k} s_k)) \rightarrow \Sigma s']] \quad (69)$$

Information flow graph In configuration c_i in a trace of a model (\mathbb{X}, c_i) , the state of a cell x in the next timestep (i.e., $c_{i+1}(x)$) is entirely determined by (1) γ_x and (2) the states of x 's neighbours in c_i . But the neighbours' states depend on the states of their neighbours in c_{i-1} , which in turn depend on their neighbours' states in c_{i-2} , and so on, until c_0 . Thus, the transitive closure of the is-a-neighbour-of relation relates a cell x to all the cells x' that may influence the states occurring at x in the trace. If x is not related to a cell x'' this way, then no information coming from x'' can influence x . Hence, we can construct the graph that depicts how information can “flow” in a CA:

Definition 7.18: Information flow graph

Given a frame $\mathbb{X} = (\mathcal{M}, N, S, \mathcal{X})$, define the *information flow graph* $\mathfrak{G}^N = (V, E)$ to be the directed graph where:

- The vertices are the cells: $V = X$.
- The edges $E \subseteq V \times V$ are the is-a-neighbour-of-relation: $(x, x') \in E$ iff $x' \in N \otimes x$.

In many cases, \mathfrak{G}^N has the same shape as “the grid” (which is the graph obtained by $(x, x') \in E$ iff $x' \in \mathcal{M} \otimes x$), such as for the forktree (Figure 5) and the \mathbb{Z}^2 grid in CGOL. An interesting exception is the CA \mathbb{Y} of Example 8.28: this one has an hexagonal information graph.

There are many important properties of \mathfrak{G}^N that can be expressed as \mathcal{L}_{CA} formulas. A first example is essential undirectedness of \mathfrak{G}^N :

Definition 7.19: Essential undirected graph

A graph (V, E) is *essentially undirected* if E is a symmetric relation (i.e., $(x, x') \in E$ iff $(x', x) \in E$).

Lemma 7.20

The information graph \mathfrak{G}^N of a nice CA \mathbb{X} is essentially undirected iff

$$\mathbb{X} \models \mathbf{v}_{s_0} \rightarrow \bigwedge_{m \in N} \bigvee_{n \in N} \diamond_m \diamond_n \mathbf{v}_{s_0} \quad (70)$$

where $s_0 \neq s_1 \in S$ are distinct states.

Proof. Suppose \mathbb{X} validates (70) and $(x, x') \in E$. Choose states $s_0, s_1 \in S$ such that $s_0 \neq s_1$, and pick $c_0: X \rightarrow S_X$ as $c_0(x'') = \begin{cases} s_0 & x'' = x \\ s_1 & \text{otherwise} \end{cases}$. Then $\mathbb{X}, c_0, x \models \mathbf{v}_{s_0}$, and by modus ponens we find that $\mathbb{X}, c_0, x \models$

$\bigwedge_{m \in N} \bigvee_{n \in N} \diamond_m \diamond_n v_{s_0}$. Since $(x, x') \in E$, there exists an $m \in N$ such that $x' = m \otimes x$. Using this m , by \wedge -elimination we find $\mathbb{X}, c_0, x \models \bigvee_{n \in N} \diamond_m \diamond_n v_{s_0}$. Thus there exists an $n \in N$ such that $\mathbb{X}, c_0, x \models \diamond_m \diamond_n v_{s_0}$, which holds iff $\mathbb{X}, c_0, x' \models \diamond_n v_{s_0}$, which holds iff $\mathbb{X}, c_0, n \otimes x' \models v_{s_0}$. So x' has a n -neighbour with state s_0 ; but x is the only cell with state s_0 in c_0 , so $n \otimes x' = x$. Hence n witnesses that $(x', x) \in E$ as well.

For the other direction, suppose E is symmetric and take any $x \in X$ and $c_0: X \rightarrow S_X$. By Lemma 6.23 it suffices to assume $\mathbb{X}, c_0, x \models v_{s_0}$ and to show $\mathbb{X}, c_0, x \models \bigwedge_{m \in N} \bigvee_{n \in N} \diamond_m \diamond_n v_{s_0}$. So take an arbitrary $m \in N$. By definition $(x, m \otimes x) \in E$, and since E is symmetric, this means $(m \otimes x, x) \in E$. The latter implies the existence of an $n \in N$ s.t. $n \otimes m \otimes x = x$. Rewriting this into $\mathbb{X}, c_0, x \models v_{s_0}$ gives $\mathbb{X}, c_0, n \otimes m \otimes x \models v_{s_0}$, i.e., $\mathbb{X}, c_0, x \models \diamond_m \diamond_n v_{s_0}$, which implies $\mathbb{X}, c_0, x \models \bigvee_{n \in N} \diamond_m \diamond_n v_{s_0}$. As m was arbitrary, we obtain the outer the infinitary conjunction. \square

Other interesting properties of \mathfrak{G}^N in \mathcal{L}_{CA} as well. The proofs are similar as the essential undirectedness formula, and therefore omitted.

Commuting information flow graph Call a \mathfrak{G}^N *commuting* if first following the step m and then n leads to the same cell as first following n and then m .

Lemma 7.21

The information graph \mathfrak{G}^N of a nice CA \mathbb{X} is *commuting* iff

$$\mathbb{X} \models \bigwedge_{m \in N} \bigwedge_{n \in N} (\diamond_m \diamond_n v_{s_0} \leftrightarrow \diamond_n \diamond_m v_{s_0}). \quad (71)$$

Many regular grids such as \mathbb{Z}^2 are commuting. In fact, \mathfrak{G}^N will commute as soon as \mathcal{M} is a group; this ensures that \bullet is commuting and from Lemma 6.5) it follows that $\diamond_m \diamond_n = \diamond_{n \bullet m} = \diamond_{m \bullet n} = \diamond_n \diamond_m$.

Reachability implies dependence It is interesting to know whether the existence of a path $m \in \mathcal{M}$ from x to a cell $m \otimes x$ implies that information in $m \otimes x$ can flow to x . That is, $(x, m \otimes x) \in E$ for all $m \in \mathcal{M}$. Many CA have this property, but for example not in \mathbb{Z} when using the neighbourhood $N = \{-1, 0\}$; then the state of x is never influenced by a past state of $x + 1$.

Lemma 7.22

The information graph \mathfrak{G}^N of nice CA \mathbb{X} has the property $(x, m \otimes x) \in E$ for all $x \in X$ and $m \in \otimes$ iff

$$\mathbb{X} \models \bigvee_{m \in \mathcal{M}} \left((\diamond_m v_{s_0}) \rightarrow \bigvee_{\pi \in \text{Paths}(N)} \pi v_{s_0} \right). \quad (72)$$

One-way Cellular Automata A One-way Cellular Automaton (OCA) is a CA s.t. information can “flow in only one direction”. The typical example is \mathbb{Z} with $N = -1, 0$; in such a CA, cells do not depend on their right neighbour, thus information can only flow from left to right. There are different ways to generalise one-wayness to other CA, but the following expresses it in terms of \mathfrak{G}^N :

Definition 7.23: One-way Cellular Automaton (OCA)

A *One-way Cellular Automaton* is a CA whose information graph \mathfrak{G}^N has no cycles except for self-loops $(x, x) \in E$.

The property “*is an OCA*” cannot be expressed via validity of a \mathcal{L}'_{CA} formula; we will prove this in Lemma 8.30 using the results of Section 8. However, \mathcal{L}_{CA} can express that a CA is *not* an OCA.

Lemma 7.24

A nice CA \mathbb{X} is *not* an OCA if

$$\mathbb{X} \models v_{s_0} \rightarrow \bigvee_{\pi \in \text{Paths}(N \setminus \{1\})} \pi v_{s_0}. \quad (73)$$

This is an “*if*” statement, and not an “*iff*”; indeed, if \mathbb{X} does not validate (73), then that only means that a counterexample exists. We have such a counterexample as soon as there exists *some* cell that does not lie on a nontrivial cycle in \mathfrak{G}^N ; it does not imply that *no* cell has such a nontrivial cycle!

7.4 Summary of infinitary disjunctions

This section showed that \mathcal{L}_{stg} can be extended with disjunctions over arbitrary (and possibly infinite) sets to obtain the logical language \mathcal{L}'_{CA} . The fragment \mathcal{L}_{CA} , which only takes disjunctions over sets directly in or related to the sets in a frame $(\mathcal{M}, N, S, \mathcal{X})$, can already express many properties of CA that \mathcal{L}_{stg} cannot. This includes properties of the global dynamics, such as the property of the trace converging to a fixpoint with all cells in a permanent “dead” state.

8 Correspondence semantical and logical bisimulation

There are two ways to compare two cells $x \in \mathcal{X}$ and $y \in \mathcal{Y}$: we can compare the sets $\text{Log}(x)$ and $\text{Log}(y)$ of logical formulas they validate, or we can directly compare them categorically as elements of coalgebras in $\mathbf{CCA}_{\mathbb{G}}$. This situation also occurs in modal logics for Kripke frames: the Hennessy-Milner theorem [12] states that points $x \in \mathbb{F}$ and $y \in \mathbb{G}$ are semantically similar (*bisimilar*, denoted $x \simeq y$) iff they are logically equivalent ($x \rightsquigarrow y$) (provided that the frames are *image finite*: every point has only finitely many successors for each given edge colour).

Since the semantics of the logic are entirely defined in the dynamics of a coalgebra, it is clear that similar coalgebras should indeed validate a similar logic. Conversely, the logic $\text{Log}(x)$ of a cell x in a CA frame $\mathbb{X} = (\mathcal{M}, N, S, X, \langle \chi, \gamma \rangle)$ can express the essential data of the dynamics at x , such as the local rule γ_x (Lemma 6.12) and the neighbourhood layout N_x (Example 6.8), thus we expect a similar logic to imply similar cells.

This section will first develop an intuition for the right CA definitions of *similarity* and *logical equivalence* with a few examples. We first show an example of similarity between CA with the same monoid and state set; Section 8.2 will prove a HM theorem for this setting. Thereafter we illustrate that a more general notion of similarity is needed to capture many pairs of CA that can directly simulate each other. The solution is to define a general notion of simulation, which is roughly a coalgebra homomorphism extended with a monoid homomorphism and a state set mapping. We also need a corresponding logical embedding that translates occurrences of monoid elements and states in formulas. A HM theorem proving that a cell x can be simulated by y iff the logic of x can be embedded in y can be found in Section 8.3. Note that this implies that x and y can mutually simulate each other (*mutual similarity*, our generalisation of *similarity*) iff their logics can be embedded into each other (*mutual logical simulation*). Example 8.28 concludes this section with a less trivial mutual similarity from the literature.

The results in this section apply to all the languages \mathcal{L}_{stg} , \mathcal{L}'_{CA} and \mathcal{L}_{CA} . However, for \mathcal{L}_{stg} there often is the additional condition that the neighbourhood N is finite (the apparent analog to image-finiteness condition of the Kripke logic HM theorem).

8.1 Examples of similarity

The first example shows similarity between CA with the same monoid and state set. This specific form of similarity can be directly phrased in terms of coalgebra homomorphisms.

8.1.1 A basic example of similarity

A starting point is to define logical equivalence $x \rightsquigarrow y$ simply as $\text{Log}(x) = \text{Log}(y)$, for cells $x \in \mathcal{X} = (X, \langle \chi, \gamma \rangle)$ and $y \in \mathcal{Y} = (Y, \langle \xi, \delta \rangle)$. As for coalgebraic similarity $x \simeq y$, called *similarity*, we should demand that $N_x = N_y$ (similar neighbourhood layout) and that $\gamma_x = \delta_y$ (same dynamics). This demand is satisfied if there exist homomorphisms

$$\overrightarrow{\mathcal{M} \otimes x} \xrightleftharpoons[k]{h} \overrightarrow{\mathcal{M} \otimes y}$$

such that $h(x) = y$ and $k(y) = x$, because (1) h and k are surjective on $\mathcal{M} \otimes y$ and $\mathcal{M} \otimes x$ respectively (Lemma 5.4), which already implies $N_x = N_y$, and (2) Theorem 5.9 guarantees that $\gamma_x = \delta_y$.

Example 8.1. Take the binary tree as monoid, i.e., the monoid $\mathcal{T} = (\{L, R\}^*, \bullet, \varepsilon)$ freely generated over $\{L, R\}$, and the neighbourhood $N = \{\varepsilon, L, R\}$ (i.e., the node itself and its two children). Now consider a frame $\mathbb{X} = (\mathcal{T}, N, S, \mathcal{T}, \langle \chi, \gamma \rangle)$ where χ is simply string concatenation, and the states S and local rules γ can be chosen freely.

Now every cell has the same successor structure: it has two children, each of which has two children, who both have two children in turn, and so on. So $\mathcal{T} \otimes x$ is isomorphic to $\mathcal{T} \otimes x'$ for any $x, x' \in \mathcal{T}$, which is also a coalgebra isomorphism $\overrightarrow{\mathcal{T} \otimes x} \rightarrow \overrightarrow{\mathcal{T} \otimes x'}$ since all cells have the same local rule. The only difference between x and x' is the number of ancestors they may have; but modal formulas can only refer to cells reachable via paths in \mathcal{T} , and there are no paths from x to its ancestors. Validity of modal formulas at x can only depend on cells reachable from x (this is proven in Lemma 8.7), i.e., only on $\mathcal{T} \otimes x$. Hence we conclude that *all* cells in \mathcal{T} validate exactly the same formulas.




Figure 20: Visualisation of the freely generated binary tree monoid $\mathcal{T} = (\{L, R\}, \bullet, \varepsilon)$.

A similar observation holds for the forktree CA \mathcal{F} (Figure 5), but now $x \rightsquigarrow x'$ iff $N_x = N_{x'}$. Hence there are two classes of logically equivalent cells on \mathcal{F} .

8.1.2 Motivation extending similarity with monoid homomorphisms

Example 8.2 (90°rotation). Let $\mathbb{X} = (\mathbb{Z}, N, S, \mathcal{X})$ where $\mathcal{X} = (X, \langle \chi, \gamma \rangle)$ be a simple 1D CA frame, in particular with the 110 Wolfram number rule¹ (the numbers were introduced by Wolfram in [50], but see [45,

§2.4] for this specific rule).

$$\begin{array}{ll}
X = \{x_i\}_{i \in \mathbb{Z}} & N = \{-1, 0, +1\} \\
n \otimes x_i = x_{i+n} & S = \{\bullet, \circ\} \\
\gamma(\circ, \circ, \circ) = \circ & \gamma(\circ, \circ, \bullet) = \bullet \\
\gamma(\circ, \bullet, \circ) = \bullet & \gamma(\circ, \bullet, \bullet) = \bullet \\
\gamma(\bullet, \circ, \circ) = \circ & \gamma(\bullet, \circ, \bullet) = \bullet \\
\gamma(\bullet, \bullet, \circ) = \bullet & \gamma(\bullet, \bullet, \bullet) = \circ
\end{array}$$

Now make a copy \mathcal{X}^j for every $j \in \mathbb{Z}$, and “stack up” these copies along the vertical axis of \mathbb{Z}^2 to obtain the CA $\mathcal{H} = (\{x_i^j\}_{i,j \in \mathbb{Z}}, \langle \chi^H, \{\gamma^j\}_{j \in \mathbb{Z}} \rangle)$ with $N^H = ((-1, 0), (0, 0), (1, 0))$ and $(a, b) \otimes^H x_{i,j} = x_{i+a, j+b}$. This CA is simply a parallel execution of independent copies of \mathcal{X} .

Now we can repeat the same trick, but rotated by 90° counter clockwise (\odot): let $\mathcal{V} = (\{x_i^j\}_{i,j \in \mathbb{Z}}, \langle \chi^V, \{\gamma^j\}_{j \in \mathbb{Z}} \rangle)$ with $N^V = ((0, -1), (0, 0), (0, 1))$ and $(a, b) \chi^V x_i^j = x_{i+b}^{j+a}$. We obtain the same parallel execution of the copies of \mathcal{X} , but with the whole graph rotated 90° .

Intuitively, \mathcal{H} and \mathcal{V} are “the same”, but they are not similar in the sense defined above ($N^H \neq N^V$, so coalgebra homomorphisms $\mathcal{H} \rightarrow \mathcal{V}$ are not even defined). Nor do they satisfy the same logic, as evident from the fact that \mathcal{H} validates

$$\psi = ((\diamond_{(-1,0)} \bullet) \wedge (\diamond_{(0,0)} \circ) \wedge (\diamond_{(1,0)} \bullet)) \rightarrow \exists \bullet, \quad (74)$$

i.e., the encoding of the rule “ $\gamma(\bullet, \circ, \bullet) = \bullet$ ”. This formula is *not* valid in \mathcal{V} , since the configuration $c_0: V \rightarrow \{\bullet, \circ\}$ with around the origin is:

$$c_0 = \begin{array}{c} 1 \\ 0 \\ -1 \\ - \\ 0 \\ 1 \\ 1 \end{array} \begin{array}{|c|c|c|} \hline \circ & \circ & \circ \\ \hline \bullet & \circ & \bullet \\ \hline \circ & \circ & \circ \\ \hline \end{array}$$

the next timestep will result in

$$c_1 = \begin{array}{c} 1 \\ 0 \\ -1 \\ - \\ 0 \\ 1 \\ 1 \end{array} \begin{array}{|c|c|c|} \hline & & \\ \hline \bullet & \circ & \bullet \\ \hline & & \\ \hline \end{array}$$

(states that we cannot compute from the partial configuration are omitted).

Since c_0 satisfies the premise of ψ but not the consequence we conclude that $\forall, c_0, x_0^0 \not\models \psi$.

The problem is that the neighbours of x_0^0 are above and below it, but left and right in \mathcal{H} and in ψ . A simple fix is to apply a 90° counterclockwise rotation

$$\begin{aligned}
\rho: \mathbb{Z}^2 &\xrightarrow{\sim} \mathbb{Z}^2 \\
\rho: (a, b) &\mapsto (-b, a)
\end{aligned}$$

to all modalities. This gives

$$\rho[\psi] = ((\diamond_{(0,-1)} \bullet) \wedge (\diamond_{(0,0)} \circ) \wedge (\diamond_{(0,1)} \bullet)) \rightarrow \exists \bullet, \quad (75)$$

which *is* valid in \mathcal{V} .

¹Here we slightly abuse the notation $\gamma(\circ, \bullet, \circ) = \bullet$ to denote $\gamma_{x_i}(z) = \bullet$ where $z(|-1|_x) = z(|1|_x) = \circ$ and $z(|0|_x) = \bullet$.

On the semantical side, ρ can also be extended to a variant of a coalgebra homomorphism $h: H \rightarrow V$, that rotates cell coordinates 90° : $h(x_i^j) \stackrel{\text{def}}{=} x_{-j}^i$. This mapping satisfies a property very similar to the \mathbb{Z}^2 - \otimes^H - \otimes^V -homomorphism property (Definition 5.2) that coalgebra homomorphisms satisfy (as proven in Lemma 5.3):

$$h((a, b) \otimes^H x_i^j) = h(x_{i+1}^{j+b}) \stackrel{\text{def}}{=} x_{-j-b}^{i+a} = (-b, a) \otimes^V x_{-j}^i = \rho(a, b) \otimes^V h(x_i^j)$$

(For a real \mathbb{Z}^2 - \otimes^H - \otimes^V -homomorphism $k: \mathcal{Y} \rightarrow \mathcal{Z}$, the property is of the form $k(m \otimes y) = m \otimes k(y)$, but here the RHS is of the form $\rho(m) \otimes k(y)$.) Observe that h preserves the dynamics of the CA; the $(-1, 0)$ neighbour of a cell x_i^j , which is x_{i-1}^j is mapped to the $(0, -1)$ neighbour of x_{-j}^i , which is x_{-j}^{i-1} . Indeed, the left neighbour (horizontal layout of \mathcal{X} in \mathbb{Z}) is mapped to the up neighbour (vertical layout), and the local rules are the same. Also, h has an inverse: rotating 90° clockwise.

Conclusion: we retain a similarity-to-logical-equivalence correspondence if (1) we generalise the definition of coalgebra homomorphisms to be paired with a monoid homomorphism, and (2) allow logical equivalence up to a monoid homomorphism on space modalities. Note that this is indeed a generalisation, since the old definition is recovered by taking the identity monoid homomorphism.

8.1.3 Motivation extending similarity with state mappings

A CA \mathbb{Y} might be able to simulate a CA \mathbb{X} , while using a different state $s' \in S_Y$ to represent a certain simulated state $s \in S_X$.

As for a simple example, take the single-cell CA $\mathbb{X} = (\mathbb{1}, \mathbb{1}, S_X = \mathbb{N}, (X = \mathbb{1}, \langle \chi, \gamma \rangle))$ where the graph is trivial $\mathbb{1} \otimes \mathbb{1} = \mathbb{1}$ but the state becomes increment every timestep $\gamma_1(\text{const}_n) = n + 1$. Let \mathbb{Y} be the same CA, except for using $S_Y = \mathbb{N}_{\geq 1}$. Clearly, \mathbb{X} becomes \mathbb{Y} after incrementing all states by 1. This also holds for the logic, for example, $\mathbb{X}, \text{const}_0, \mathbf{1} \models \text{zz}2$ iff $\mathbb{Y}, \text{const}_1, \mathbf{1} \models \text{zz}3$. Note that const_1 is $i \circ \text{const}_0$ where $i: n \mapsto n + 1$, and that $\text{zz}3$ is $\text{zz}2$ with i applied to all occurrences of states.

A reverse simulation, of \mathbb{Y} by \mathbb{X} , is obviously given by decrementing states by 1. But also incrementing states by any other fixed $n \in \mathbb{N}$ works; if $i_{1000}: n \mapsto n + 1000$, then $\mathbb{Y}, \text{const}_1, \mathbf{1} \models \psi$ iff $\mathbb{X}, \text{const}_{1001}, \mathbf{1} \models i_{1000}[\psi]$.

Conclusion: we obtain more bisimilarity-to-logical-equivalences if we extend coalgebra homomorphisms (in this case the identity on cells) with a function between state sets $S_X \rightarrow S_Y$, and define logical equivalence also up to a map between states.

8.2 Basic Hennessy-Milner theorem

The logic \mathcal{L}_{stg} admits a cellular automaton variant of the Hennessy-Milner theorem of Kripke modal logics. The Hennessy-Milner theorem states that two worlds in two image-finite Kripke models satisfy the same formulas iff they are “bisimilar” [12]. We will define an alternative notion of similarity for CA cells, and prove two cells *validate* (i.e., satisfy on all models) the same formulas iff they are similar. The upcoming definitions assume that the compared CA have the same monoid; Section 8.3 will generalise the definitions to distinct monoids.

Definition 8.3: Similar cells

Let $\mathcal{X} = (X, \langle \chi, \gamma \rangle)$ and $\mathcal{Y} = (Y, \langle \xi, \delta \rangle)$ be CA in $\mathbf{CCA}_{\mathbf{G}}$ based on the same signature (\mathcal{M}, N, S) , and let $x \in X$ and $y \in Y$. Write $\overrightarrow{\mathcal{M} \otimes x}$ and $\overrightarrow{\mathcal{M} \otimes y}$ for the subcoalgebras of \mathcal{X} and \mathcal{Y} respectively containing the cells $\mathcal{M} \otimes x = \{m \otimes x \mid m \in \mathcal{M}\}$ and $\mathcal{M} \otimes y = \{m \otimes y \mid m \in \mathcal{M}\}$.

Then x is similar to y , denoted $x \simeq y$, iff there exist coalgebra homomorphisms

$$\overrightarrow{\mathcal{M} \otimes x} \begin{array}{c} \xrightarrow{h} \\ \xleftarrow{k} \end{array} \overrightarrow{\mathcal{M} \otimes y}$$

such that $h(x) = y$ and $k(y) = x$.

Remark 8.4. This definition of similarity is very strict: observe that cells are already non-similar if they have a different number of neighbours (h and k are both surjective on the set of reachable cells, by Lemma 5.4) or if they have a different local rule.

Remark 8.5. The standard coalgebraic definition of bisimilarity states that cells x and y are bisimilar in \mathcal{Z} if there exist coalgebra homomorphism $h: \mathcal{X} \rightarrow \mathcal{Z}$ and $k: \mathcal{Y} \rightarrow \mathcal{Z}$ such that $h(x) = k(y)$. Similarity according to Definition 8.3 hence means that x and y are both bisimilar in \mathcal{X} and in \mathcal{Y} .

Definition 8.6: Logically equivalent

Let x be a cell in the frame $\mathbb{X} = (\mathcal{M}, N, \mathcal{X})$ and y be a cell in the frame $\mathbb{Y} = (\mathcal{M}, N, \mathcal{Y})$. Then x and y are logically equivalent, denoted as $x \leftrightarrow y$, when for all $\psi \in \Psi$, $\mathbb{X}, x \models \psi$ iff $\mathbb{Y}, y \models \psi$.

Note that logical equivalence is about *validity* of a formula at a cell, not about *satisfiability*, since we are not using an initial configuration $c_0: X \rightarrow S$.

Whether a formula is satisfied in a cell x depends only on the cells that are reachable from a path from x . This is intuitively to be expected, since (1) the local rule of x depends only on cells that x can reach (since $N \subseteq \mathcal{M}$), and (2) we only have space modalities that follow existing paths in \mathcal{M} .²

Lemma 8.7

Let x be a cell in a frame $\mathbb{X} = (\mathcal{M}, N, X, \langle \chi, \gamma \rangle)$. Let $c, c': X \rightarrow S$ s.t. If $c \upharpoonright (\mathcal{M} \otimes x) = c' \upharpoonright (\mathcal{M} \otimes x)$. Then for all $\psi \in \Psi$ it holds: $\mathbb{X}, c, x \models \psi$ iff $\mathbb{X}, c', x \models \psi$.

Proof. Prove the stronger claim: for all $t \in \mathbb{N}$, $m \in \mathcal{M}$ it holds that:

$$\mathbb{X}, c, m \otimes x \models \psi \text{ iff } \mathbb{X}, c', m \otimes x \models \psi$$

by induction on ψ . □

Lemma 8.8

If there exist coalgebra homomorphisms

$$\overrightarrow{\mathcal{M} \otimes x} \xrightleftharpoons[k]{h} \overrightarrow{\mathcal{M} \otimes y}$$

such that $h(x) = y$ and $k(y) = x$, then h and k are inverses of each other.

Proof. For all $m \in \mathcal{M}$, it holds that:

$$\begin{aligned} k(h(m \otimes x)) &= k(m \otimes h(x)) \\ &= m \otimes k(h(x)) \\ &= m \otimes k(y) \\ &= m \otimes x. \end{aligned}$$

An analogous argument shows that $h(k(m \otimes y)) = m \otimes y$. □

2

Corollary 8.9

If there exist coalgebra homomorphisms

$$\overrightarrow{\mathcal{M} \otimes x} \xrightarrow{h} \overrightarrow{\mathcal{M} \otimes y} \xleftarrow{k}$$

such that $h(x) = y$ and $k(y) = x$, then for all $m \in \mathcal{M}$:

$$\begin{aligned} N_m \otimes x &= N_{h(m \otimes x)} \\ N_{k(m \otimes y)} &= N_m \otimes y \end{aligned}$$

Proof. The statements are symmetric, so we only prove $N_m \otimes x = N_{h(m \otimes x)}$ explicitly.

Let $x' = m \otimes x$. We have $|n_1|_{x'} = |n_2|_{x'} \in N_{x'}$ iff $n_1 \otimes x = n_2 \otimes x$. Since h is a bijection (Lemma 8.8), this holds iff $n_1 \otimes h(x') = h(n_1 \otimes x') = h(n_2 \otimes x') = n_2 \otimes h(x')$, which holds iff $|n_1|_{h(x')} = |n_2|_{h(x')}$. \square

The following lemma is (a slightly strengthened variant) of one direction of the Hennessy-Milner theorem (the other direction is proven in Theorem 8.12).

Lemma 8.10

If $x \Leftrightarrow y$ then $x \rightsquigarrow y$.

Proof. We will prove the following two claims:

- For all $t \in \mathbb{N}$, $c_0: X \rightarrow S$, $m \in \mathcal{M}$ and $\psi \in \Psi$:

$$\mathbb{X}, G_{\mathbb{X}}^t(c_0), m \otimes x \models \phi \text{ iff } \mathbb{Y}, G_{\mathbb{Y}}^t(c_0 \circ k), h(m \otimes x) \models \phi \quad (\text{H1})$$

- For all $t \in \mathbb{N}$, $c_0: Y \rightarrow S$, $m \in \mathcal{M}$ and $\psi \in \Psi$:

$$\mathbb{Y}, G_{\mathbb{Y}}^t(c_0), m \otimes y \models \psi \text{ iff } \mathbb{X}, G_{\mathbb{X}}^t(c_0 \circ h), k(m \otimes y) \models \psi \quad (\text{H2})$$

To see why these claims imply $x \rightsquigarrow y$, suppose $\mathbb{X}, x \models \psi$ and take any arbitrary $c_0: Y \rightarrow S$.

Satisfiability of ψ at y only depends on $c_0 \upharpoonright (\mathcal{M} \otimes y)$ (Lemma 8.7), we have $\mathbb{Y}, c_0, y \models \psi$ iff $\mathbb{Y}, c'_0 \circ k, y \models \psi$ for all $c'_0: X \rightarrow S$ such that $c'_0 \circ k \upharpoonright (\mathcal{M} \otimes y) = c_0 \upharpoonright (\mathcal{M} \otimes y)$. From (H1) it follows that $\mathbb{Y}, c'_0 \circ k, y \models \psi$ holds iff $\mathbb{X}, c'_0, x \models \psi$, but that is known to hold (we assumed $\mathbb{X}, x \models \psi$). Since h and k are bijections (Lemma 8.8, we do have exactly such an c'_0 , namely $c'_0 := c_0 \circ h$ (now $c'_0 \circ k = c_0$, so the requirement is satisfied trivially).

It remains to prove (H1) and (H2). Because of symmetry we only need to prove (H1) explicitly (the proof of (H2) is exactly the same after swapping the symbols $h \leftrightarrow k$, $X \leftrightarrow Y$, $x \leftrightarrow y$, $\chi \leftrightarrow \xi$ and $\gamma \leftrightarrow \delta$).

We proceed by induction on ψ . As a shorthand for induction hypotheses, write $H(\phi)$ to denote:

$$H(\phi) \stackrel{\text{def}}{=} \bigvee_{\substack{c_0: X \rightarrow S \\ t \in \mathbb{N} \\ m \in \mathcal{M}}} \mathbb{X}, G_{\mathbb{X}}^t(c_0), m \otimes x \models \phi \text{ iff } \mathbb{Y}, G_{\mathbb{Y}}^t(c_0 \circ k), h(m \otimes x) \models \phi \quad (76)$$

Case $\psi = \phi_1 \vee \phi_2$

Induction hypotheses: $H(\phi_1)$ and $H(\phi_2)$.

We observe that, for all t , c_0 and $x' = m \otimes x \in (\mathcal{M} \otimes x)$, it holds that:

$$\begin{aligned} \mathbb{X}, G_{\mathbb{X}}^t(c_0), x' \models \phi_1 \vee \phi_2 & \\ \text{iff } (\mathbb{X}, G_{\mathbb{X}}^t(c_0), x' \models \phi_1 \text{ or } \mathbb{X}, G_{\mathbb{X}}^t(c_0), x' \models \phi_2) & \quad // \text{ Semantics of } \vee. \\ \text{iff } (\mathbb{Y}, G_{\mathbb{Y}}^t(c_0 \circ k), h(x') \models \phi_1 \text{ or } \mathbb{Y}, G_{\mathbb{Y}}^t(c_0 \circ k), h(x') \models \phi_2) & \quad // \text{ Induction hypothesis.} \\ \text{iff } \mathbb{Y}, G_{\mathbb{Y}}^t(c_0 \circ k), h(x') \models \phi_1 \vee \phi_2 & \quad // \text{ Semantics of } \vee. \end{aligned}$$

Case $\psi = \bigvee_{i \in I} \phi_i$

Induction hypotheses: $H(\phi_i)$ for all $i \in I$.

As in the binary disjunction case, for all t , c_0 and $x' = m \otimes x$ it holds that:

$$\begin{aligned} \mathbb{X}, G_{\mathbb{X}}^t(c_0), x' \models \bigvee_{i \in I} \psi_i & \\ \text{iff exists } i \in I \text{ such that } \mathbb{X}, G_{\mathbb{X}}^t(c_0), x' \models \psi_i & \quad // \text{ Semantics of infinitary disjunctions.} \\ \text{iff exists } i \in I \text{ such that } \mathbb{Y}, G_{\mathbb{Y}}^t(c_0 \circ k), h(x') \models \psi_i & \quad // \text{ Induction hypothesis.} \\ \text{iff } \mathbb{Y}, G_{\mathbb{Y}}^t(c_0 \circ k), h(x') \models \bigvee_{i \in I} \psi_i. & \quad // \text{ Semantics of } \vee. \end{aligned}$$

Case $\psi = \perp$

It follows immediately from the semantics of \perp that, for all t , c_0 and $x' = m \otimes x$, that both $\mathbb{X}, G_{\mathbb{X}}^t(c_0), x' \models \perp$ and $\mathbb{Y}, G_{\mathbb{Y}}^t(c_0 \circ k), h(x') \models \perp$ are always false.

Case $\psi = \neg\phi$

Induction hypothesis: $H(\phi)$.

For all t , c_0 and $x' = m \otimes x$ it holds that:

$$\begin{aligned} \mathbb{X}, G_{\mathbb{X}}^t(c_0), x' \models \neg\phi & \text{ iff } \mathbb{X}, G_{\mathbb{X}}^t(c_0), x' \not\models \phi \\ & \text{ iff } \mathbb{Y}, G_{\mathbb{Y}}^t(c_0 \circ k), h(x') \not\models \phi \quad // \text{ Induction hypothesis.} \\ & \text{ iff } \mathbb{Y}, G_{\mathbb{Y}}^t(c_0 \circ k), h(x') \models \neg\phi \end{aligned}$$

Case $\psi = v_S$ for some $s \in S$

This case is still quite straightforward, since the semantics of v_S do not depend on t . In particular, for all t , c_0 and $x' = m \otimes x$, it holds that:

$$\begin{aligned} \mathbb{X}, G_{\mathbb{X}}^t(c_0), x' \models v_S & \text{ iff } c_0(x') = s \\ & \text{ iff } c_0(k(h(x'))) = s \quad // k(h(x')) = x' \text{ by Lemma 8.8.} \\ & \text{ iff } \mathbb{Y}, G_{\mathbb{Y}}^t(c_0 \circ k), h(x') \models v_S \end{aligned}$$

Case $\psi = s$ for some $s \in S$

This case is considerably more complicated than the v_S case, since the truth value of s *does* depend on the point in time. Inductively we need to show that $c_t(x') = (c_t \circ k)(h(x'))$ (where $c_t = G_{\mathbb{X}}^t(c_0)$) for all $x' \in \mathcal{M} \otimes x$ and $t \in \mathbb{N}$, which is done in Lemma 8.11 below. With this lemma available, we can show (for all t , c_0 and $x' = m \otimes x$) that:

$$\begin{aligned} \mathbb{X}, G_{\mathbb{X}}^t(c_0), x' \models s & \\ \text{iff } G_{\mathbb{X}}^t(c_0)(x') = s & \quad // \text{ Semantics of } s. \\ \text{iff } G_{\mathbb{Y}}^t(c_0 \circ k)(h(x')) = s & \quad // \text{ Lemma 8.11.} \\ \text{iff } \mathbb{Y}, G_{\mathbb{Y}}^t(c_0 \circ k), h(x') \models s. & \quad // \text{ Semantics of } s. \end{aligned}$$

Case $\psi = \diamond_n \phi$ for some $n \in \mathcal{M}$

Induction hypothesis: $H(\phi)$.

For all t, c_0 and m it holds that:

$$\begin{aligned}
\mathbb{X}, G_{\mathbb{X}}^t(c_0), m \otimes x \models \diamond_n \phi & \\
& \text{iff } \mathbb{X}, G_{\mathbb{X}}^t(c_0), n \otimes m \otimes x \models \phi & // \text{Semantics } \diamond_n. \\
& \text{iff } \mathbb{X}, G_{\mathbb{X}}^t(c_0), (n \bullet m) \otimes x \models \phi \\
& \text{iff } \mathbb{Y}, G_{\mathbb{Y}}^t(c_0 \circ k), h((n \bullet m) \otimes x) \models \phi \\
& // \text{Induction hypothesis (note that } (n \bullet m) \in \mathcal{M}, \text{ so } (n \bullet m) \otimes x \in (\mathcal{M} \otimes x)). \\
& \text{iff } \mathbb{Y}, G_{\mathbb{Y}}^t(c_0 \circ k), h(n \otimes m \otimes x) \models \phi \\
& \text{iff } \mathbb{Y}, G_{\mathbb{Y}}^t(c_0 \circ k), n \otimes h(m \otimes x) \models \phi \\
& \text{iff } \mathbb{Y}, G_{\mathbb{Y}}^t(c_0 \circ k), h(m \otimes x) \models \diamond_n \phi. & // \text{Semantics } \diamond_n.
\end{aligned}$$

Case $\psi = \varepsilon \phi$

Induction hypothesis: $H(\phi)$.

This is the case for which the quantification over time in (H1) was needed; for all t, c_0 and $x' = m \otimes x$ it now holds that:

$$\begin{aligned}
\mathbb{X}, G_{\mathbb{X}}^t(c_0), x' \models \varepsilon \phi & \\
& \text{iff } \mathbb{X}, G_{\mathbb{X}}^{t+1}(c_0), x' \models \phi & // \text{Semantics } \varepsilon. \\
& \text{iff } \mathbb{Y}, G_{\mathbb{Y}}^{t+1}(c_0 \circ k), h(x') \models \phi & // \text{Induction hypothesis.} \\
& \text{iff } \mathbb{Y}, G_{\mathbb{Y}}^t(c_0 \circ k), h(x') \models \varepsilon \phi. & // \text{Semantics } \varepsilon.
\end{aligned}$$

□

Lemma 8.11

If there exist coalgebra homomorphisms

$$\overrightarrow{\mathcal{M} \otimes x} \xrightleftharpoons[k]{h} \overrightarrow{\mathcal{M} \otimes y}$$

such that $h(x) = y$ and $k(y) = x$, then it holds for all $t \in \mathbb{N}$ and $c_0: X \rightarrow S$ that:

$$G_{\mathbb{X}}^t(c_0) \upharpoonright (\mathcal{M} \otimes x) = (G_{\mathbb{Y}}^t(c_0 \circ k) \circ h) \upharpoonright (\mathcal{M} \otimes x). \tag{77}$$

Proof. Induction on $t \in \mathbb{N}$. The base case, when $t = 0$, follows directly from Lemma 8.8 and the fact that $G^0 = \text{Id}$:

$$c_0(m \otimes x) = c_0(k(h(m \otimes x))) = ((c_0 \circ k) \circ h)(m \otimes x)$$

holds for all $m \otimes x \in \mathcal{M} \otimes x$.

For the inductive case, assume (77) holds for a given $t \in \mathbb{N}$. We need to show it also holds for $t + 1$. Take an

arbitrary $x' = m \otimes x \in \mathcal{M} \otimes x$, then we compute:

$$\begin{aligned}
G_{\mathbb{X}}^{t+1}(c_0)(x') &= G_{\mathbb{X}}(G_{\mathbb{X}}^t(c_0))(x') \\
&= \gamma_{x'}(\lambda |n|_{x'} \cdot G_{\mathbb{X}}^t(c_0)(n \otimes x')) && // \text{Definition of global rule (Eq. (28)).} \\
&= \delta_{h(x')}(\lambda |n|_{h(x')} \cdot G_{\mathbb{X}}^t(c_0)(h(n \otimes x'))) && // \text{By Corollary 8.9, } |n|_{x'} = |n|_{h(x')}. \\
&\quad // \text{Thus by the extension property of coalgebra homomorphisms (Theorem 5.9).} \\
&= \delta_{h(x')}(\lambda |n|_{h(x')} \cdot G_{\mathbb{X}}^t(c_0)(n \otimes h(x'))) \\
&= \delta_{h(x')}(\lambda |n|_{h(x')} \cdot G_{\mathbb{Y}}^t(c_0 \circ k)(n \otimes h(x'))) && // \text{Induction hypothesis.} \\
&= G_{\mathbb{Y}}(G_{\mathbb{Y}}^t(c_0 \circ k)(h(x'))) && // \text{Definition of global rule (Eq. (28)).} \\
&= G_{\mathbb{Y}}^{t+1}(c_0 \circ k)(h(x')).
\end{aligned}$$

□

The converse of Lemma 8.8 also holds, but in \mathcal{L}_{stg} only if neighbourhoods N and M are finite:

Theorem 8.12

Let x be a cell in the frame $\mathbb{X} = (\mathcal{M}, N, \mathcal{X})$ and y be a cell in the frame $\mathbb{Y} = (\mathcal{M}, N, \mathcal{Y})$. then $x \Leftrightarrow y$ iff $x \rightsquigarrow y$. (This holds in \mathcal{L}'_{CA} and \mathcal{L}_{CA} , and also in \mathcal{L}_{stg} if N and M are finite).

Proof. The left-to-right direction follows from Lemma 8.8. For the other direction, assume $x \rightsquigarrow y$. We now need to construct coalgebra homomorphisms

$$\begin{array}{ccc}
\longrightarrow & \xrightarrow{h} & \longrightarrow \\
\mathcal{M} \otimes x & \xleftrightarrow[k]{h} & \mathcal{M} \otimes y
\end{array}$$

such that $h(x) = y$ and $k(y) = x$. By Theorem 5.9, these need to be \mathcal{M} - χ - ξ -homomorphisms. Hence we have no choice but to define (for all $m \in \mathcal{M}$):

$$\begin{aligned}
h &: x \mapsto y, \\
h &: m \otimes x \mapsto m \otimes h(x) = m \otimes y, \\
k &: y \mapsto x, \\
k &: m \otimes y \mapsto m \otimes k(y) = m \otimes x.
\end{aligned}$$

Note that $h = k^{-1}$.

It remains to prove the extension property. Because of symmetry, we only show this for h ; so it remains to prove that for all $x' = m \otimes x \in \mathcal{M} \otimes x$ it holds that:

$$\gamma_{x'}(z \parallel x') = \delta_{h(x')}(z \parallel h(x'))$$

for all $z: N \rightarrow S$ that are $N_{h(x')}$ compatible. Note that $N_{x'} = N_{h(x')}$ since h is an isomorphism (isomorphism implies $n_1 \otimes x' = n_2 \otimes x'$ iff $n_1 \otimes h(x') = h(n_1 \otimes x') = h(n_2 \otimes x') = n_2 \otimes h(x')$, i.e., $|n_1|_{x'} = |n_2|_{x'}$ iff $|n_1|_{h(x')} = |n_2|_{h(x')}$).

Further observe that

$$\text{Log}(m \otimes x) = \text{Log}(h(m \otimes x)). \quad (78)$$

This follows from the fact that (omitting frames) $m \otimes x \vDash \phi$ iff $x \vDash \diamond_m \phi$ iff (because $x \rightsquigarrow y$) $y \vDash \diamond_m \phi$ iff $m \otimes y \vDash \phi$ and $m \otimes y = m \otimes h(x) = h(m \otimes x)$ (these statements can easily be proven via the semantics of \diamond_m).

We can use Lemma 6.12 to express $\gamma_{x'}(z \parallel x')$ as a formula (in \mathcal{L}'_{CA} and \mathcal{L}_{CA} this is possible for all N , in \mathcal{L}_{stg} only if N is finite). In particular, $\gamma_{x'}(z \parallel x') = s_0$ holds iff

$$\mathbb{X}, x' \models \left(\bigwedge_{n \in N} \diamond_n z(|n|_{x'}) \right) \rightarrow \varkappa s_0,$$

which holds iff (by Eq. (78), and since $|n|_{h(x')} = |n|_{x'}$)

$$\mathbb{Y}, h(x') \models \left(\bigwedge_{n \in N} \diamond_n z(|n|_{h(x')}) \right) \rightarrow \varkappa s_0,$$

which (again by Lemma 6.12) holds iff $\delta_{h(x')}(z \parallel h(x')) = s_0$. Since z and s_0 were arbitrary, the desired conclusion also holds: $\gamma_{x'}(z \parallel x') = \delta_{h(x')}(z \parallel h(x'))$ for all $N_{x'} = N_{h(x')}$ compatible configuration $z: N \rightarrow S$. \square

8.3 Generalised Hennessy-Milner theorem

Theorem 8.12 gives a variant of the Hennessy-Milner theorem that allows to compare similarity of two CA \mathcal{X} and \mathcal{Y} that are based on the same monoid \mathcal{M} and use the same set of states S . We will now formalise the generalised definitions of similarity and logical equivalence as sketched in Sections 8.1.2 and 8.1.3, and prove a corresponding Hennessy-Milner theorem.

8.3.1 General CA morphisms

We will now generalise the definition of coalgebra homomorphisms to include morphisms between CA defined on different signatures. Recall that $\mathbf{CCA}_{\mathcal{G}}$ (Definition 4.16) is the coproduct of $\mathbf{CCA}_{\mathcal{G}}^{\Sigma}$ over all signatures Σ . As such, it has no morphisms between CA defined on different signatures.

We will extend $\mathbf{CCA}_{\mathcal{G}}$ to the category $\mathbf{CCA}_{\mathcal{G}}^+$, which has the same objects, but whose morphisms are *general CA morphisms*, which have the characteristic properties of coalgebra homomorphisms (Theorem 5.9) but also include a monoid homomorphism and a function between state sets. Coalgebra homomorphisms are special cases of general CA morphisms (where the monoid homomorphism and state set maps are the identities), so $\mathbf{CCA}_{\mathcal{G}}$ is a subcategory of $\mathbf{CCA}_{\mathcal{G}}^+$.

To define general CA morphisms we first need the following auxiliary definition:

Definition 8.13: Graph-neighbourhood morphism

Let $\mathcal{M}, \mathcal{N} \in \mathbf{Mon}$ be monoids and let $M \subseteq \mathcal{M}$ and $N \subseteq \mathcal{N}$. A *graph-neighbourhood morphism* $h_{\mathcal{G}}: (\mathcal{M}, M) \rightarrow (\mathcal{N}, N)$ is a monoid homomorphism (in \mathbf{Mon}) such that $h_{\mathcal{G}}[M] \subseteq N$.

Definition 8.14: General CA morphism

A *general CA morphism* $h: \mathbb{X} \rightarrow \mathbb{Y}$ from $\mathbb{X} = (\mathcal{M}, M, S_X, X, \langle \chi, \gamma \rangle)$ to $\mathbb{Y} = (\mathcal{N}, N, S_Y, Y, \langle \xi, \delta \rangle)$ is triple of functions $(h_{\mathcal{G}}, h_C, h_S)$ where

- $h_{\mathcal{G}}: (\mathcal{M}, M) \rightarrow (\mathcal{N}, N)$ is a graph-neighbourhood morphism.
- $h_C: X \rightarrow Y$ is a function between cells. We usually omit the subscript and write $h(x)$ for $h_C(x)$.
- $h_S: S_X \rightarrow S_Y$ is a function between states.

Furthermore, this triple adheres to the following constraints:

1. **(Module property)**

$$h_C(m \otimes x) = h_G(m) \otimes h_C(x)$$

for all $x \in X$ and $m \in \mathcal{M}$.

2. h_S is injective.

3. For any $x \in X$ and $y \in Y$ s.t. $h_C(x) = y$, the restriction $h_C \upharpoonright (\mathcal{M} \otimes x): (\mathcal{M} \otimes x) \rightarrow (\mathcal{N} \otimes y)$ is an injection.

4. **(Extension property)**

$$h_S(\gamma_x(z)) = \delta_{h_C(x)}(z')$$

for every $x \in X$ and every $z: M_x \rightarrow S_X$, where $z': N_{h_C(x)} \rightarrow S_Y$ is any extension of $h_S \circ z \circ (h_G^{M_x})^{-1}$.

Here $h_G^{M_x}: M_x \rightarrow N_{h_C(x)}$ is defined as $h_G^{M_x}: |m|_x \mapsto |h_G(m)|_{h_C(x)}$ ³.

As a commuting diagram this is:

$$\begin{array}{ccc} S_X^{M_x} & \xrightarrow{h_S \circ (-) \circ (h_G^{M_x})^{-1}} & S_Y^{h_G^{M_x}[M_x]} & \xrightarrow{\text{extend}} & S_Y^{N_{h_C(x)}} \\ \gamma \downarrow & & & & \downarrow \delta_{h_C(x)} \\ S_X & \xrightarrow{h_S} & & & S_Y \end{array}$$

A coarse shorthand to capture the essence of this law is:

$$h_S \gamma = \delta h_S.$$

Note that coalgebra homomorphisms are special cases of general CA homomorphisms.

8.3.2 Generalised notions of similarity

Definition 8.15: Simulation and mutual similarity

Let $\mathbb{X} = (\mathcal{M}, M, S_X, X, \langle \chi, \gamma \rangle)$ and $\mathbb{Y} = (\mathcal{N}, N, S_Y, Y, \langle \xi, \delta \rangle)$ be frames in \mathbf{CCA}_G^+ . A cell $x \in X$ is simulated by a cell $y \in Y$, denoted as $x \Rightarrow y$, if there exists a general CA morphism $h: \mathcal{M} \otimes x \rightarrow \mathcal{M} \otimes y$ such that $h_C(x) = y$. If also $y \Rightarrow x$ (i.e., if there also exists a $k: \mathcal{M} \otimes y \rightarrow \mathcal{M} \otimes x$ such that $k_C(y) = x$), then x is mutually similar to y , denoted as $x \Leftrightarrow y$.

Remark 8.16. The definition of a simulation $h: x \Rightarrow y$ tells directly that the local rules γ and δ commute with the embedding given by h . We will later see (in Lemma 8.24) that h also commutes with the global rules $G_{\mathbb{X}}$ and $G_{\mathbb{Y}}$.

Remark 8.17. The definition of mutual similarity does not require h and k to be each other's inverse.

The general definition of logical equivalence is less straightforward than Definition 8.3, since the two CA may have different states and different path monoids, and hence different symbols. We propose the following “translation” of formulas between incompatible CA:

Definition 8.18: Logical simulation

Let $\Psi(\mathcal{M}, M, S)$ denote the \mathcal{L}_{stg} -formulas (idem when using \mathcal{L}'_{CA} or \mathcal{L}_{CA}) defined on the monoid \mathcal{M} ,

³It may not be obvious that $h_G^{M_x}$ is well-defined, since $|h_G(m)|_{h_C(x)}$ is ill-defined when $h_C(m) \notin N$. The latter cannot occur, since the $h_G[M] \subseteq N$ condition of a graph-neighbourhood morphism (Definition 8.13) ensures that $h_C(m) \in N$ when $m \in M$.

neighbourhood $M \subseteq \mathcal{M}$ and state set S . Then a *logical simulation* $h_L: \Psi(\mathcal{M}, M, S) \rightarrow \Psi(\mathcal{N}, N, S')$ is a pair (h_G, h_S) where $h_G: (\mathcal{M}, M) \rightarrow (\mathcal{N}, N)$ is a graph-neighbourhood morphism (Definition 8.13) and $h_S: S \rightarrow S'$ a function.

The computation of h_L on a formula in $\Psi(\mathcal{M}, S)$ is defined by straightforward recursion:

$$\begin{aligned} h_L(\perp) &\stackrel{\text{def}}{=} \perp \\ h_L(s_i) &\stackrel{\text{def}}{=} h_S(s_i) \\ h_L(\mathbf{v}_{s_i}) &\stackrel{\text{def}}{=} \mathbf{v}_{h_S(s_i)} \\ h_L(\neg\psi) &\stackrel{\text{def}}{=} \neg h_L(\psi) \\ h_L(\psi_1 \vee \psi_2) &\stackrel{\text{def}}{=} h_L(\psi_1) \vee h_L(\psi_2) \\ h_L(\bigvee_{i \in I} \psi_i) &\stackrel{\text{def}}{=} \bigvee_{i \in I} h_L(\psi_i) \\ h_L(\diamond_m \psi) &\stackrel{\text{def}}{=} \diamond_{h_G(m)} h_L(\psi) \\ h_L(\mathbf{z}\psi) &\stackrel{\text{def}}{=} \mathbf{z}h_L(\psi) \end{aligned}$$

We define the *image configuration* of h_L on \mathbb{Y} as the set

$$\text{ImC}(h_L) \stackrel{\text{def}}{=} \{c_0: Y \rightarrow S_Y \mid c_0(h_G(x)) = h_S(s) \text{ for some } s \in S_X, \text{ for all } x \in X\}. \quad (79)$$

This are, intuitively, all possible extensions of the (h_G, h_S) images of \mathbb{X} configurations.

To define logical equivalence, we first define logical simulation. Intuitively, $x \rightsquigarrow y$ if $\text{Log}(x)$ can injectively be translated into a fragment of $\text{Log}(y)$. It is possible that \mathbb{Y} has more states or more paths than \mathbb{X} does, but these are then simply not part of the fragment in which $\text{Log}(x)$ is embedded. We call x and y if their logics can be embedded into one each other's. The formal definitions are as follows:

Definition 8.19: Logical simulation

Let $\mathbb{X} = (\mathcal{M}, M, S_X, X, \langle \chi, \gamma \rangle)$ and $\mathbb{Y} = (\mathcal{N}, N, S_Y, Y, \langle \xi, \delta \rangle)$ be frames in \mathbf{CCA}_G^+ . A cell $x \in X$ is *logically simulated* by a cell $y \in Y$, denoted as $x \rightsquigarrow y$, if there exists a logical simulation $h_L: \Psi(\mathcal{M}, M, S_X) \rightarrow \Psi(\mathcal{N}, N, S_Y)$ such that for all $c_0: Y \rightarrow S_Y \in \text{ImC}(h_L)$:

$$\mathbb{X}, h_S^{-1} \circ c_0 \circ h_C, x \models \psi \text{ iff } \mathbb{Y}, c_0, y \models h_L(\psi) \quad (80)$$

for all $\psi \in \Psi(\mathcal{M}, S_X)$, where

$$\begin{aligned} h_C: (\mathcal{M} \otimes x) &\rightsquigarrow (h_G[\mathcal{M}] \otimes y), \\ h_C: m \otimes x &\mapsto h_G(m) \otimes y. \end{aligned}$$

If both $x \rightsquigarrow y$ and $y \rightsquigarrow x$, then x and y are *mutually logically similar*, which is denoted by $x \rightsquigarrow\!\!\!\!\!\leftrightarrow y$.

Remark 8.20. Eq. (80) does apply to all models on the frame \mathbb{X} : Lemma 8.22 will show that h_C is injective, so every $c'_0: X \rightarrow S_X$ can be written in the form $c'_0 = h_S^{-1} \circ c_0 \circ h_C$ for some $c_0 \in \text{ImC}(h_L)$. Note that this is well defined (despite h_S not being a bijection), as by definition of $\text{ImC}(h_L)$, all cells in the image of h_C do have a state in the image of h_S .

Time invariance (Definition 6.14) is preserved under logical simulations, for configurations using only states in $h_S[S_X]$:

Lemma 8.21

If $h_L: x \rightsquigarrow y$ and if ψ is time invariant for x then $h_L(\psi)$ is time invariant for y on models using configurations $c_0: Y \rightarrow h_S[S_X]$.

Proof. Let $c_0: Y \rightarrow h_S[S_X]$, then (since h_S is injective) it is of the form $h_S \circ c'_0$ for some $c'_0: Y \rightarrow S_X$. Hence $\mathbb{Y}, c_0, y \models h_L(\psi)$ iff $\mathbb{X}, c'_0 \circ h_C, x \models \psi$. Since ψ is time invariant for x , the latter holds iff $\mathbb{X}, c'_0 \circ h_C, x \models z^t \psi$ for all $t \in \mathbb{N}$. Since $x \rightsquigarrow y$, this in turn holds iff $\mathbb{Y}, c_0, x \models z^t h_L(\psi)$ for all $t \in \mathbb{N}$ (note that $h_L(z^t \psi) = z^t h_L(\psi)$). Hence $\mathbb{Y}, c_0, y \models h_L(\psi)$ iff $\mathbb{Y}, c_0, x \models z^t h_L(\psi)$ for all $t \in \mathbb{N}$, which is exactly the property of time invariance. \square

8.3.3 Generalised Hennessy-Milner theorem

We will now show that $x \rightleftharpoons y$ iff $x \rightsquigarrow y$. First we prove the left-to-right direction. Although I state it as a lemma, in reality I first performed the proof, and used the result to find the correct definition for $x \rightleftharpoons y$:

Lemma 8.22

If $x \rightsquigarrow y$ then $x \rightleftharpoons y$, under the conditions that that $\text{card}(S_X) \geq 2$ and $\text{card}(S_Y) \geq 2$ (and in case of \mathcal{L}_{stg} , also that M and N are finite).

Proof. The assumption $x \rightsquigarrow y$ gives us a logical simulation $h_L = (h_G, h_S): \Psi(\mathcal{M}, M, S_X) \rightarrow \Psi(\mathcal{N}, N, S_Y)$. We will extend (h_G, h_S) to a general CA morphism $h = (h_G, h_C, h_S): \mathcal{M} \otimes x \rightarrow \mathcal{M} \otimes y$. The definition of $h_C: (\mathcal{M} \otimes x) \rightarrow (\mathcal{N} \otimes y)$ is straightforward:

$$\begin{aligned} h_C: x &\mapsto y \\ h_C: (m \otimes x) &\mapsto h_G(m) \otimes y \end{aligned}$$

This definition already implies that property 1 of a general CA morphism (Definition 8.14) is satisfied.

As for the injectivity of h_S (property 2), suppose for contradiction that there exist two states $s_0 \neq s_1 \in S_X$ such that $h_S(s_0) = h_S(s_1)$. In any configuration, x can have at most one state, so $\mathbb{X}, x \models s_0 \rightarrow \neg s_1$ (see Lemma 6.24). From $x \rightsquigarrow y$ it follows that $\mathbb{Y}, c_0, y \models h_S(s_0) \rightarrow \neg h_S(s_1)$ (for any $c_0: \text{ImC}(h_L)$, see Remark 8.20). Now write $s' = h_S(s_0)$, then this is $\mathbb{Y}, c_0, y \models s' \rightarrow \neg s'$. Now pick the particular $c'_0 \in \text{ImC}(h_L)$ such that $c'_0(y) = s'$ (e.g., $c'_0 = \text{const}_{s'}$). Then $\mathbb{Y}, c'_0, y \models s'$, and from modus ponens (Lemma 6.21) it follows that $\mathbb{Y}, c'_0, y \models \neg s'$. Thus $s' = c'_0(y) \neq s'$; a contradiction.

We show injectivity of h_C (property 3) via a similar argument. First observe that, for any $x' = m \otimes x \in \mathcal{M} \otimes x$:

$$\text{if } \mathbb{X}, x' \models \phi \text{ then } \mathbb{X}, x \models \diamond_m \phi. \quad (81)$$

. Thus $\mathbb{Y}, c_0, y \models \diamond_{h_G(m)} h_L(\phi)$ (for all $c_0 \in \text{ImC}(h_L)$), which is equivalent to $\mathbb{Y}, c_0, h_G(m) \otimes y \models h_L(\phi)$. But by definition of h_C , it holds that $h_C(x') = h_G(m) \otimes y$, so $\mathbb{Y}, c_0, h_C(x') \models h_L(\phi)$.

Now suppose h_C is not injective, and that $h_C(m_1 \otimes x) = h_C(m_2 \otimes x)$ while $m_1 \otimes x \neq m_2 \otimes x$. Then $h_G(m_1) \otimes y$ and $h_G(m_2) \otimes y$ are the same cell, so they have always the same state. Hence $\mathbb{Y}, c_0, y \models (\diamond_{h_G(m_1)} h_S(s)) \leftrightarrow (\diamond_{h_G(m_2)} h_S(s))$ for any $s \in S_X$ and all $c_0 \in \text{ImC}(h_L)$. By definition of $x \rightsquigarrow y$, this means that for all $s \in S$:

$$\mathbb{X}, x \models (\diamond_{m_1} s) \leftrightarrow (\diamond_{m_2} s) \quad (82)$$

(by Remark 8.20 we indeed obtain satisfaction in all models, i.e., validity). Since S_X has at least two states, so there exists $s_1 \neq s_2 \in S_X$. Recall that $m_1 \otimes x \neq m_2 \otimes x$ are different cells, we can pick $c'_0: X \rightarrow S_X$ such that $c'_0(m_1 \otimes x) = s_1$ and $c'_0(m_2 \otimes x) = s_2$. Then $\mathbb{X}, c'_0, x \models \diamond_{m_1} s_1$ by the semantics $\diamond_{m_1} s_1$. But \wedge -elimination

and MP^4 we obtain $\mathbb{X}, c'_0, x \vDash \diamond_{m_2} s_1$, which holds iff $\mathbb{X}, c'_0, m_2 \otimes x \vDash s_1$, i.e., iff $c'_0(m_2 \otimes x) = s_1$; the latter is indeed a contradiction!

It remains to prove the extension property (property 4). To this end, take any arbitrary $x' = m \otimes x \in \mathcal{M} \otimes x$ and any $z: M_{x'} \rightarrow S_X$.

Let $z': N_{h_C(x')} \rightarrow S_Y$ be any arbitrary extension of $h_S \circ z \circ (h_G^{M_{x'}})^{-1}: h_G^{M_{x'}}[M_{x'}] \rightarrow S_Y$. Using Lemma 6.12, we can express the local rule $\gamma_{x'}$ on input z as the formula:

$$\mathbb{X}, x' \vDash \left(\bigwedge_{|m_i| \in M_{x'}} \diamond_{m_i} z(|m_i|) \right) \rightarrow \varkappa \gamma_{x'}(z). \quad (83)$$

We can do the same for $h_C(x')$ and $\delta_{h_C(x')}$:

$$\mathbb{Y}, h_C(x') \vDash \left(\bigwedge_{|n| \in N_{h_C(x')}} \diamond_n z'(|n|) \right) \rightarrow \varkappa \delta_{h_C(x')}(z'). \quad (84)$$

But, similarly as we did for Eq. (81), we can translate Eq. (83) under the logical simulation to:

$$\mathbb{Y}, c_0, h_C(x') \vDash \left(\bigwedge_{|m_i| \in M_{x'}} \diamond_{h_G(m_i)} h_S(z(|m_i|)) \right) \rightarrow \varkappa h_S(\gamma_{x'}(z)), \quad (85)$$

for all $c_0 \in \text{ImC}(h_L)$. The restriction to configurations in $\text{ImC}(h_L)$ is not problematic, since the extension property only concerns the action of δ on local neighbourhood configurations z' where all cells in the image of h_G obtain a state in the image of h_S (i.e., extensions of $h_S \circ z \circ (h_G^{M_{x'}})^{-1}$ for some $z: X \rightarrow S_X$). All such neighbourhood configurations indeed arise as a subconfiguration in $\text{ImC}(h_L)$. Hence there exist enough choices for c_0 to use (85) to completely specify δ on inputs z' .

Since $h_G[M] \subseteq N$ (Definition 8.13), it follows that $h_G^{M_{x'}}[M_{x'}] \subseteq N_{h_C(x')}$. Hence the premise of (85) contains a subset of the conjuncts of the premise of (84). Abstracting to propositional logic, we have two formulas:

$$\begin{aligned} (A_1 \wedge A_2 \wedge \dots \wedge A_\ell) &\rightarrow B_1, \\ (A_1 \wedge A_2 \wedge \dots \wedge A_\ell \wedge C_1 \wedge \dots \wedge C_q) &\rightarrow B_2, \end{aligned}$$

These can be further abstracted to

$$\begin{aligned} A &\rightarrow B_1, \\ A \wedge C &\rightarrow B_2, \end{aligned}$$

Using the logical tautology $(A \rightarrow B_1) \rightarrow (A \wedge C \rightarrow B_1)$ and modus ponens (Lemmas 6.20 and 6.21) we find that $A \wedge C \rightarrow B_1$, i.e.,

$$\mathbb{Y}, c_0, h_C(x') \vDash \left(\bigwedge_{|n| \in N_{h_C(x')}} \diamond_n z'(|n|) \right) \rightarrow \varkappa h_S(\gamma_{h_C(x')}(z)). \quad (86)$$

So, given that the neighbours reachable from $h_C(x')$ via paths $n \in h_G[M]$ have states $z'(|n|_{h_C(x')})$, the state of $h_C(x')$ in the next timestep in the trace is $h_S(\gamma_{h_C(x')}(z))$. This holds regardless of the states z' assigns to cells at paths $N \setminus h_G[M]$. Since $\delta_{h_C(x')}$ also gives the next state of $h_C(x')$, and since $h_C(x')$ has at most one state per timestep, these states must coincide (i.e., semantically we know $B_1 \leftrightarrow B_2$) (i.e., varying the states of cells at paths $N \setminus h_G[M]$ does not change the output of $\delta_{h_C(x')}$ as long as the states of the paths in $h_G[M]$ are given by z'). Hence the desired conclusion follows: $h_S(\gamma_{x'}(z)) = \delta_{h_C(x')}(z)$. \square

⁴More precisely, we use Lemma 6.20 with the logical tautology $\alpha \wedge \beta \rightarrow \alpha$, with $\alpha := \diamond_{m_1} s_1 \rightarrow \diamond_{m_2} s_1$ and $\beta := \diamond_{m_2} s_1 \rightarrow \diamond_{m_1} s_1$. Now Eq. (82) gives $\alpha \wedge \beta$, thus by modus ponens (Lemma 6.21) we obtain α . Finally apply modus ponens on α , since the premise $\mathbb{X}, c'_0, x \vDash \diamond_{m_1} s_1$ is known to hold.

Corollary 8.23

If $x \rightsquigarrow y$ then $x \rightleftharpoons y$.

Proof. Apply Lemma 8.22 twice. □

We will now prove the converse of Lemma 8.22 and Corollary 8.23, but this first requires an auxiliary lemma:

Lemma 8.24

Given $h: x \Rightarrow y$, then for all $c_0 \in \text{ImC}(h_L)$:

$$h_S \circ G_X(h_S^{-1} \circ c_0 \circ h_C) \circ h_C^{-1} = G_Y(c_0) \upharpoonright h_C[\mathcal{M} \otimes x] \quad (87)$$

are the same configuration $h_C[\mathcal{M} \otimes x] \rightarrow S_Y$.

$$\begin{array}{ccc} S_X^X & \xleftarrow{h_S^{-1} \circ (-) \circ h_C} & \text{ImC}(h_L) \\ G_X \downarrow & & \downarrow G_Y \\ S_X^X & \xrightarrow{h_S \circ (-) \circ h_C^{-1}} & \text{ImC}(h_L) \end{array} \quad (88)$$

It follows that

$$G_X(h_S^{-1} \circ c_0 \circ h_C) = h_S^{-1} \circ G_Y(c_0) \circ h_C \quad (89)$$

(and that $h_S^{-1} \circ G_Y(c_0) \circ h_C$ is well-defined).

Proof. Write h for h_C and take any $h_C(x') \in h[\mathcal{M} \otimes x]$. Then:

$$\begin{aligned} (h_S \circ G_X(h_S^{-1} \circ c_0 \circ h_C) \circ h_C^{-1})(h_C(x')) &= h_S(G_X(h_S^{-1} \circ c_0 \circ h_C)(x')) \\ &= h_S(\gamma_{x'}(\lambda | m| \in M_{x'} \cdot (h_S^{-1} \circ c_0 \circ h_C)(m \otimes x'))) \\ &\quad // \text{Definition global rule } G_X \text{ (Eq. (28)).} \\ &= h_S(\gamma_{x'}(z)) \quad // \text{Let } z := \lambda | m| \in M_{x'} \cdot (h_S^{-1} \circ c_0 \circ h_C)(m \otimes x'): \\ &= \delta_{h_C(x')}(z') \quad // \text{Extension property (Definition 8.14.4).} \end{aligned}$$

for any extension $z': Y \rightarrow S_Y$ of $h_S \circ z \circ (h_G^{M_{x'}})^{-1}: h_G^{M_{x'}}[M_{x'}] \rightarrow S_Y$. The latter equals

$$\begin{aligned} h_S \circ z \circ (h_G^{M_{x'}})^{-1} &= \lambda | m| \in ((h_G^{M_{x'}})^{-1} \circ h_G^{M_{x'}})[M_{x'}] \cdot (h_S \circ h_S^{-1} \circ c_0 \circ h_C)(m \otimes x') \\ &= \lambda | m| \in M_{x'} \cdot (c_0 \circ h_C)(m \otimes x') \\ &= \lambda | m| \in M_{x'} \cdot c_0(h_G(m) \otimes h_C(x')) \\ &= \lambda | n| \in h_G^{M_{x'}}[M_{x'}] \cdot c_0(n \otimes h_C(x')). \quad // \text{Since } h_G^{M_{x'}} \text{ is an injection.} \end{aligned}$$

The extension property guarantees that $\delta_{h_C(x')}(z')$ behaves the same on any extension z' of ($//$ Since $h_G^{M_{x'}}$ is an injection.), including the intuitively natural extension:

$$z' := \lambda | n| \in N_{h_C(x')} \cdot c_0(n \otimes h_C(x')).$$

Continuing where we left off in ($//$ Extension property (Definition 8.14.4).) with this choice of z' , we find:

$$\begin{aligned} &= \delta_{h_C(x')}(z') \\ &= \delta_{h_C(x')}(\lambda | n| \in N_{h_C(x')} \cdot c_0(n \otimes h_C(x')).) \\ &\stackrel{\text{def}}{=} G_Y(c_0)(h_C(x')), \quad // \text{Definition global rule } G_Y \text{ (Eq. (28))} \end{aligned}$$

which was to be shown. \square

Lemma 8.25

If $x \Rightarrow y$ then $x \rightsquigarrow y$.

Proof. We need to prove that, for all $c_0 \in \text{ImC}(h_L)$ it holds that

$$\mathbb{X}, h_S^{-1} \circ c_0 \circ h_C, x \vDash \phi \text{ iff } \mathbb{Y}, c_0, y \vDash h_L(\phi).$$

We will prove the stronger claim that for all $k_0 \in \text{ImC}(h_L)$ and all $x' = m \otimes x \in \mathcal{M} \otimes x$:

$$\mathbb{X}, h_S^{-1} \circ k_0 \circ h_C, x' \vDash \phi \text{ iff } \mathbb{Y}, k_0, h_C(x') \vDash h_L(\phi). \quad (90)$$

For convenience, we omit explicitly notating the frames \mathbb{X} and \mathbb{Y} .

Case $\phi = \exists \psi$:

This is the only difficult case. Observe the following:

$$\begin{aligned} & \bigvee_{c_0 \in \text{ImC}(h_L)} h_S^{-1} \circ c_0 \circ h_C, x' \vDash \exists \psi \\ \text{iff } & \bigvee_{c_0} G_{\mathbb{X}}(h_S^{-1} \circ c_0 \circ h_C), x' \vDash \psi && // \text{Semantics } \exists. \\ \text{iff } & \bigvee_{c_0} h_S^{-1} \circ G_{\mathbb{Y}}(c_0) \circ h_C, x' \vDash \psi && // \text{Lemma 8.24 Eq. (89)}. \\ \text{iff } & \bigvee_{c_0} G_{\mathbb{Y}}(c_0), h_C(x') \vDash h_L(\psi) && // \text{Induction hypothesis (with } k_0 := G_{\mathbb{Y}}(c_0)). \\ \text{iff } & \bigvee_{c_0} G_{\mathbb{Y}}(c_0), h_C(x') \vDash \exists h_L(\psi) && // \text{Semantics } \exists. \\ \text{iff } & \bigvee_{c_0} G_{\mathbb{Y}}(c_0), h_C(x') \vDash h_L(\exists \psi). \end{aligned}$$

Case $\phi = s \in S_X$ (case v_S is analogous):

Fix any arbitrary $c_0 \in \text{ImC}(h_L)$, then:

$$\begin{aligned} & \mathbb{X}, h_S^{-1} \circ c_0 \circ h_L, x' \vDash s \\ \text{iff } & h_S^{-1}(c_0(h_C(x'))) = s \\ \text{iff } & h_S(h_S^{-1}(c_0(h_C(x')))) = h_S(s) && // h_S \text{ is an injection.} \\ \text{iff } & c_0(h_C(x')) = h_S(s) \\ \text{iff } & c_0, h_C(x') \vDash h_S(s), \end{aligned}$$

and indeed $h_L(\phi) = h_S(s)$ in this case.

Case $\phi = \diamond_m \psi$ (for some $m \in \mathcal{M}$):

Again fix any arbitrary $c_0 \in \text{ImC}(h_L)$ and observe:

$$\begin{aligned} & \mathbb{X}, h_S^{-1} \circ c_0 \circ h_C, x' \vDash \diamond_m \psi \\ \text{iff } & \mathbb{X}, h_S^{-1} \circ c_0 \circ h_C, m \otimes x' \vDash \psi && // \text{Semantics } \diamond_m. \\ \text{iff } & \mathbb{Y}, c_0, h_C(m \otimes x') \vDash h_L(\psi) && // \text{Induction hypothesis.} \\ \text{iff } & \mathbb{Y}, c_0, h_G(m) \circledast h_C(x') \vDash h_L(\psi) \\ \text{iff } & \mathbb{Y}, c_0, h_C(x') \vDash \diamond_{h_G(m)} h_L(\psi) && // \text{Semantics } \diamond_{h_G(m)}. \\ \text{iff } & \mathbb{Y}, c_0, h_C(x') \vDash h_L(\diamond_m \psi). \end{aligned}$$

Remaining cases:

The $\phi = \perp$ case is trivial. The $\phi = \neg\psi$, $\phi = \psi_1 \vee \psi_2$ and $\phi = \bigvee_{i \in I} \psi_i$ cases for almost immediately from the induction hypothesis. \square

Corollary 8.26

If $x \rightleftharpoons y$ then $x \rightsquigarrow y$.

Proof. Apply Lemma 8.25 twice. \square

We now summarise the work done above as:

Theorem 8.27

$$x \Rightarrow y \quad \text{iff} \quad x \rightsquigarrow y$$

and

$$x \rightleftharpoons y \quad \text{iff} \quad x \rightsquigarrow y$$

(given and $\text{card}(S_X) \geq 2$ and $\text{card}(S_Y) \geq 2$, and in case of \mathcal{L}_{stg} , also that M and N are finite).

8.3.4 Example applications

The following example shows a mutual similarity between two CA on very different graphs. It is based on an example from the literature [36, Prop. 2, 2nd proof], and demonstrates how the current approach, which separated the set of cells and the paths monoid, fits well with the description of the simulation:

Example 8.28 (Hexagonal to 3D). Let the hexagonal paths grid \mathcal{H} be the Abelian group generated over the set $\{A, B, C\}$ with equalities $ABC = \mathbf{1}$, $AB = BA$, $AC = CA$ and $BC = CB$ ⁵.

As neighbourhood we use $M = \{\mathbf{1}, A, B, C, A^{-1}, B^{-1}, C^{-1}\}$. Since this group is Abelian, we can draw it as the following grid:




Figure 21: Right: a piece of the grid encoded by \mathcal{H} , not showing the inverse group elements. Right: the equality $ABC = \mathbf{1}$.

As cells of the CA frame $\mathbb{X} = (\mathcal{H}, N, S, X, \langle \chi, \gamma \rangle)$ we take $X = \mathcal{H}$, and the neighbours-mapping $\textcircled{\times}$ is the multiplication on \mathcal{H} . The state set S and local rule $\gamma: (M_X \rightarrow S)^X$ can be arbitrary (as long as $\text{card}(S) \geq 2$). Since γ_x is the same for all $x \in X$, we can simply write $\gamma_x = \gamma_*$. Since all M neighbours of every cell are distinct, we will (with a minor abuse of notation) also write M for M_x .

⁵The equalities $AC = CA$ and $BC = CB$ are superfluous and can be derived from the pair $ABC = \mathbf{1}$ and $AB = BA$.

We can simulate this CA on \mathbb{Z}^3 by using the CA

$$\mathbb{Y} = (\mathbb{Z}^3, N = \{(i, j, k) \mid |i| + |j| + |k| \leq 1\}, S, Y = \mathbb{Z}^3 + (1, 1, 1)\mathbb{Z}^3, \langle \xi, \delta \rangle),$$

such that \oplus is addition modulo $(1, 1, 1)$ (so $(a, b, c) \oplus ((p, q, r) + (1, 1, 1)\mathbb{Z}^3) \stackrel{\text{def}}{=} (p+a, q+b, r+c) + (1, 1, 1)\mathbb{Z}^3$). Again we will write δ_* for $\delta_{\vec{y}}$ and N for $N_{\vec{y}}$. It remains to define δ , but this is easier after a few more observations.

The intuitive layout of Y in \mathbb{Z}^3 is as follows: if we fix the value 0 for the third coordinate, we obtain a plane containing every cell of Y exactly once (for every $(i, j, k) \in \mathbb{Z}^3 + (1, 1, 1)\mathbb{Z}^3$, we have that $(i - k, j - k, 0) + (1, 1, 1)\mathbb{Z}^3$ is the same cell, and the only occurrence of this cell at third coordinate 0). The plane where the third coordinate is 1 is the same, but with the origin at $(1, 1, 1)$ instead of $(0, 0, 0)$. We see that every cell is repeated endlessly along diagonal lines with slope $(1, 1, 1)$.

To define δ , observe that every element of \mathcal{H} can be written as $A^i B^j C^k$ because \mathcal{H} is Abelian. Since $ABC = \mathbf{1}$, we may assume a canonical representation of the cells in \mathcal{H} (and hence X) in which at least one of i, j and k is zero. This already gives an isomorphism from \mathcal{H} to $\mathbb{Z}^3 + (1, 1, 1)\mathbb{Z}^3$:

$$h: A^i B^j C^k \mapsto (i, j, k).$$

Intuitively, we identify A, B and C with the axes of the 3D grid. The A and A^{-1} neighbours become the left and right neighbours, the B and B^{-1} the forth and back neighbours, and C and C^{-1} the up and down neighbours. Thus $h: M \cong N$, and we define $\delta_* \stackrel{\text{def}}{=} \gamma_* \circ h_G^M: S^N \rightarrow S$.




Figure 22: Left: the neighbourhood M on the hexagonal grid. Right: the isomorphic neighbourhood N in \mathbb{Z}^3 .

After taking $h_S \stackrel{\text{def}}{=} \text{Id}_S$, we see that h trivially satisfies extension property of a general CA morphism (Definition 8.14.4), and the other requirements are also straightforward to check. Since h also has an inverse, we obtain a mutual similarity $A^i B^j C^k \Leftrightarrow (i, j, k) + (1, 1, 1)\mathbb{Z}^3$ for every $(i, j, k) \in \mathbb{Z}^3$. Furthermore, from Theorem 8.27 it follows that (up to an isomorphism) $\text{Log}(A^i B^j C^k)$ and $\text{Log}((i, j, k))$ are the same.

The Hennessy-Milner theorems can also be used to prove what kind of simulations are impossible. For example, one cannot simulate a CA with a quiescent state on an CA without a quiescent state:

Lemma 8.29

Consider frames \mathbb{X} and \mathbb{Y} such that \mathbb{X} has a quiescent state. If $x \Rightarrow y$ for some $x \in X$ and $y \in Y$, then \mathbb{Y} has a quiescent state.

Proof. Let ψ be the has-a-quiescent-state-formula (Eq. (56)), then $\mathbb{X}, x \models \psi$. By Lemma 8.25 it holds that

$$\mathbb{Y}, c_0, y \models h_L(\psi) \tag{91}$$

for all $c_0: Y \rightarrow h_S[S_X]$. Now it might be that $h_C[M \otimes x] \subsetneq N \otimes y$, thus (91) does not completely specify δ_y on the constant neighbourhood configuration $\text{const}_{h_S(s)}: (N \otimes y) \rightarrow S_Y$ yet. This is not an issue, since we can pick a c_0 such that all neighbours of y (even those not in the image of X under h_C) have state $h_S(s)$, in which case (91) does imply that $\delta_y(\text{const}_{h_S(s)}) = h_S(s)$. The latter is exactly the quiescence condition. \square

The converse may not hold when h_S is not surjective; in this case \mathbb{Y} might have a quiescent state in $S_Y \setminus h_S[S_X]$ which is not of a concern in the definition of $x \rightsquigarrow y$.

Finally, the Hennessy-Milner theorem allows to prove that certain properties cannot be axiomatised in \mathcal{L}'_{CA} . One example is the property of being an OCA (Definition 7.23 of Section 7.3).

Lemma 8.30

There exists no $\psi \in \mathcal{L}'_{CA}$ such that $\mathbb{X}, x \models \psi$ iff there exist no nontrivial path from x to x in \mathfrak{G}^N .

Proof. Let $\mathbb{X} = (\mathbb{Z}, \{0\}, S, \mathbb{Z}, \langle \chi, \gamma \rangle)$ and $\mathbb{Y} = (\mathbb{Z}, \{-1, 0, 1\}, S, \mathbb{Z}, \langle \chi, \delta \rangle)$ where $S = \{\bullet, \circ\}$, $\gamma_z = \text{const}_\circ: \{z\} \rightarrow S$ and $\delta_z = \text{const}_\circ: \{z-1, z, z+1\} \rightarrow S$ for all $z \in \mathbb{Z}$. Using $h_C = \text{ld}_{\mathbb{Z}}$, $h_S = \text{ld}_S$ and $h_G = \text{ld}_{\mathbb{Z}}$ we obtain a mutual similarity $h: z \Rightarrow z$ for all $z \in \mathbb{Z}$ (between z in \mathbb{X} to z in \mathbb{Y}). Thus by Theorem 8.27 also $z \rightsquigarrow z$, since $\text{ImC}(h_L) = S^{\mathbb{Z}}$ are all configurations on \mathbb{Y} , this implies $\mathbb{X}, z \models \psi$ iff $\mathbb{Y}, z \models \psi$.

Now suppose there exists a ψ that is validated in a point z iff there exist no nontrivial path in \mathfrak{G}^N from z to z . But \mathbb{X} is an OCA, so $\mathbb{X}, z \models \psi$ holds for all $z \in \mathbb{Z}$. Hence by the previous observation, also $\mathbb{Y}, z \models \psi$ for all $z \in \mathbb{Z}$. The latter implies that \mathbb{Y} is an OCA; but this is a contradiction since \mathbb{Y} 's information graph is the essentially undirected (c.f. Definition 7.19) Cayley graph of \mathbb{Z} . \square

8.4 Summary of the Hennessy-Milner theorem

This section proved a Hennessy-Milner theorem: a direct link between the logic of CA and the ability of CA to semantically simulate each other. We first saw, in Theorem 8.12, that CA \mathcal{X} and \mathcal{Y} defined over the same signature (\mathcal{M}, N, S) validate exactly the same logic iff there exist coalgebra homomorphisms $k: \mathcal{X} \rightarrow \mathcal{Y}$ and $h: \mathcal{Y} \rightarrow \mathcal{X}$.

Thereafter we generalised coalgebra homomorphisms to *general CA morphisms* between frames. In order to translate signatures, these CA morphisms are coalgebra homomorphisms between cells extended with an additional monoid homomorphism and a state-set map. Intuitively, exists a general CA morphism $h: \mathbb{X} \rightarrow \mathbb{Y}$ when \mathbb{X} can be simulated by \mathbb{Y} .

We likewise generalised *logical equivalence* to *logical simulation* in order to compare logics of frames on different signatures. This is implemented by allowing a monoid homomorphism and state-set map to translate the signature-dependent symbols in the logic.

Finally, we proved a generalised Hennessy-Milner theorem (Theorem 8.27): frames can simulate each other via general CA morphisms if and only if their logics can be embedded into one each other's.

9 Concluding remarks

We finish this thesis by reflecting on the advantages and disadvantages of the presented constructions, outlining directions for potential future work, and a final summary.

9.1 Discussion of the coalgebraic model

The main particularity of the present coalgebraic model of CA is the separation between the set of cells X and the monoid \mathcal{M} of paths between cells; most existing works identify those sets. The separation seems to offer new flexibility. For example, we showed in Example 8.28 how to embed the hexagonal grid into \mathbb{Z}^3 , which was straightforward since we only use the cells of the hexagonal CA and adjust the paths accordingly (intuitively, we identify multiple grid points with the same cell). This should be compared to the non-coalgebraic description by Roká [36], which had to explain how to translate a configuration on the hexagonal grid to a configuration on \mathbb{Z}^3 by copying the state of a source cell to all \mathbb{Z}^3 grid-points with which it is identified.

Classes of general CA We proposed several classes of general CA with different levels of uniformity between local rules in Section 4.1, None of these is the obvious “correct” class: this is probably application dependent. However, our coalgebras are sufficiently expressive that any coalgebraic categories of CA according to of the desired degree of uniformity simply arise as full subcategories of $\mathbf{CCA}_{\mathbf{G}}$, the category of general CA.

The standard coalgebra framework The present model does not fit in well with existing approaches of describing dynamical systems coalgebraically. The strongest discrepancies are Section 5, which showed that the concept of *behavioural equivalence* for our coalgebras has little to do with the actual behaviour of CA. But the issue runs deeper: traditionally, the elements of carrier sets of coalgebras are called “states”. In our model, this are the *cells*, and part of the static structure rather than the dynamics of the system. The question whether there exist a coalgebraic model for CA that is better aligned with existing work remains open for future work.

Alternative base categories Our coalgebraic model uses an endofunctor on \mathbf{Set} , but future work could investigate other base categories with more structure. For example, we can “lift” the functor $C_{\mathbf{G}}$ to an endofunctor on the category of posets, in such a way that $C_{\mathbf{G}}$ is recovered by precomposing with the forgetful functor. This would result in orderings on cells and states, and monotone general CA morphisms and local rules. Alternatively, we can consider the base category of small categories and use the morphisms of a category as connections between cells, instead of a separate monoid.

9.2 Discussion of the new modal logic

The presented modal logic $\mathcal{L}_{\mathbf{CA}}$ is quite expressive; aside from any finite spatio-temporal pattern in the trace of a CA, it can sometimes even describe asymptotic behaviour of the trace (recall Examples 7.13 and Lemma 7.14). However, infinitary disjunctions come at the cost of making the logic undecidable (Corollary 7.15). We conjecture that $\mathcal{L}_{\mathbf{stg}}$ is decidable, at least for CA with a finite state sets and neighbourhoods, but a proof of this remains open for future work.

One-way cellular automata Lemma 8.30 shows that $\mathcal{L}'_{\mathbf{CA}}$ does not have a formula that expresses that a CA is a one-way CA. This is not in conflict with the design goal of the language, since the property of

being an OCA is not visible in the trace. Indeed, the proof of Lemma 8.30 constructs two CA with the same traces, but one is an OCA and the other not.

This example shows that it is debatable whether the definition of an OCA is “correct”, and more generally, whether the definition of the information flow graph (Definition 7.18) is “correct”. The alternative is to define one-directedness based on the actual dynamics of the CA: we could for example define *essential one-directedness* as follows: a CA is essentially one-directed if for any cell x in a configuration c_t of which it is known that x 's state in configuration c_{t+1} is s_{t+1} , for any $k \in \mathbb{N}$, the state of x in c_{t+1+k} is the same for any state of x in the current configuration c_t . While somewhat technical, this expresses that information about x 's past states cannot flow back to influence its future states (other than directly influence the direct successor state of course). One could generalise the definition of the information flow graph analogously. While semantically more intuitive, it may not solve the inexpressivity issue: we have not found a formula to express *essential one-directedness*.

A more general definition of simulation? Theorem 8.27, our generalised Hennessy-Milner theorem, is based on our definition of a *simulation* and CA being *mutually similar* (Definition 8.15). It is straightforward to see that a simulation $h: x \Rightarrow y$ indeed embeds the dynamics of \mathbb{X} visible from x into y , but note that this is a simulation that always simulated 1 cell by 1 cell, 1 state by 1 state, and 1 timestep by 1 timestep. There are many examples of more general definitions of simulations between CA found in the literature, such as allowing groups of cell to simulate groups of cells, and allowing a constant time delay between the simulated and simulating CA, or even mapping entire configurations to configurations (see e.g., [36] for simulations of CA on Cayley graphs, and [29, §5.2] for a very general definition for graphs of the form \mathbb{Z}^d). It seems possible to generalise our definition of \Rightarrow accordingly, and it might also be possible to extend the definition of logical equivalence with a search-and-replace operation that multiplies occurrences of \varkappa (to express a delay), replaces occurrences of atomic states by a conjunction of states (when a single state of the simulated CA can be represented by multiple states in the simulating CA) or even replaces occurrences of atomic states $s \in S$ by spatially arranged “blocks” of states (e.g., when s is encoded in the simulating CA by multiple cells in the states s_1, \dots, s_k arranged in the pattern given by m_1, \dots, m_k then one would replace s by $\diamond_{m_1} s_1 \wedge \diamond_{m_2} s_2 \wedge \dots \wedge \diamond_{m_k} s_k$). This possibility remains open for future work.

Uniform substitution We showed in Section 6.3 that the global variables of our logic, which semantically represent the states in the initial configuration, allow a limited form of uniform substitution: only substitution of time-invariant formulas. In future work we may investigate the alternative option of assigning propositional letters to cells (as in the Kripke semantics), in which case full uniform substitution likely becomes possible.

9.3 Conclusion

We have successfully constructed a modal logic for describing spatio-temporal patterns in the trace of a CA (see Example 7.1). Aside from these patterns, the logic can also express some global properties of a CA, such as nilpotency (Example 7.13) and periodicity (Example 7.14). To implement the semantics of the logic, we constructed a novel coalgebraic model of CA that separates cells from paths between cells. We proposed two variants of a generalisation of CA in which different cells can have different local behaviour. Our coalgebraic model is sufficiently flexible to also represent both *irregular* and *non-uniform* CA, to which the logic also applies.

The logic is not based on the predicate-lifting framework for coalgebraic modal logic, since we showed that this framework cannot express dynamic properties for our coalgebraic representation of CA.

The syntax of the logic is uniform for all CA, but formulas may contain state labels or paths labels specific to a CA. However, we gave a generalised Hennessy-Milner theorem stating that the formulas validated by a CA can be translated into formulas validated by another CA if and only if the latter CA can directly simulate the

former. This result simplifies to a direct logical equivalence in the special case where both CA have exactly the same states and paths.

References

- [1] Kamalika Bhattacharjee, Nazma Naskar, Souvik Roy, and Sukanta Das. “A survey of cellular automata: types, dynamics, non-uniformity and applications”. In: *Natural Computing* 19.2 (July 2018), pp. 433–461. ISSN: 1572-9796. DOI: [10.1007/s11047-018-9696-8](https://doi.org/10.1007/s11047-018-9696-8). URL: <http://dx.doi.org/10.1007/s11047-018-9696-8>.
- [2] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*. Cambridge University Press, 2001. URL: www.cambridge.org/9780521527149.
- [3] Dave Burraston, Ernest A. Edmonds, Dan Livingstone, and Eduardo Reck Miranda. “Cellular Automata in MIDI based Computer Music”. In: *International Conference on Mathematics and Computing*. 2004. URL: <http://hdl.handle.net/2027/spo.bbp2372.2004.047>.
- [4] Silvio Capobianco and Tarmo Uustalu. “A categorical outlook on cellular automata”. In: *Journées Automates Cellulaires* (2010).
- [5] Silvio Capobianco and Tarmo Uustalu. “Additive Cellular Automata Graded-Monadically”. In: *International Symposium on Principles and Practice of Declarative Programming (PPDP 2023)* (2023).
- [6] Corina Cîrstea, Alexander Kurz, Dirk Pattinson, Lutz Schröder, and Yde Venema. “Modal Logics are Coalgebraic1”. In: *The Computer Journal* 54.1 (Feb. 2009), pp. 31–41. ISSN: 0010-4620. DOI: [10.1093/comjnl/bxp004](https://doi.org/10.1093/comjnl/bxp004). eprint: <https://academic.oup.com/comjnl/article-pdf/54/1/31/963631/bxp004.pdf>. URL: <https://doi.org/10.1093/comjnl/bxp004>.
- [7] Stathis Delivourias, Haralampos Hatzikirou, Rafael Peñaloza, and Dirk Walther. “Detecting Emergent Phenomena in Cellular Automata Using Temporal Description Logics”. In: *Cellular Automata*. Ed. by Jarosaw Ws, Georgios Ch. Sirakoulis, and Stefania Bandini. Cham: Springer International Publishing, 2014, pp. 357–366. ISBN: 978-3-319-11520-7.
- [8] M. Delorme and J. Mazoyer. *Cellular Automata A Parallel Model*. Springer Science+Business Media Dordrecht, 1999.
- [9] Bruno Durand and Zsuzsanna Róka. “The game of life: universality revisited”. In: *Cellular Automata A Parallel Model*. Springer Science+Business Media Dordrecht, 1999.
- [10] G. Bard Ermentrout and Leah Edelstein-Keshet. “Cellular Automata Approaches to Biological Modeling”. In: *Journal of Theoretical Biology* 160.1 (1993), pp. 97–133. ISSN: 0022-5193. DOI: <https://doi.org/10.1006/jtbi.1993.1007>. URL: <https://www.sciencedirect.com/science/article/pii/S0022519383710076>.
- [11] Martin Gardner. “MATHEMATICAL GAMES”. In: *Scientific American* 223.4 (1970), pp. 120–123. ISSN: 00368733, 19467087. URL: <http://www.jstor.org/stable/24927642> (visited on 06/27/2024).
- [12] Matthew Hennessy and Robin Milner. “Algebraic laws for nondeterminism and concurrency”. In: *J. ACM* 32.1 (Jan. 1985), pp. 137–161. ISSN: 0004-5411. DOI: [10.1145/2455.2460](https://doi.org/10.1145/2455.2460). URL: <https://doi.org/10.1145/2455.2460>.

- [13] Jesse Hughes and Bart Jacobs. “Simulations in coalgebra”.
In: *Theoretical Computer Science* 327.1 (2004). Selected Papers of CMCS '03, pp. 71–108.
ISSN: 0304-3975. DOI: <https://doi.org/10.1016/j.tcs.2004.07.022>.
URL: <https://www.sciencedirect.com/science/article/pii/S030439750400444X>.
- [14] Navot Israeli and Nigel Goldenfeld.
“Computational Irreducibility and the Predictability of Complex Physical Systems”.
In: *Physical Review Letters* 92.7 (Feb. 2004). DOI: [10.1103/physrevlett.92.074105](https://doi.org/10.1103/physrevlett.92.074105).
URL: <https://doi.org/10.1103/physrevlett.92.074105>.
- [15] Bart Jacobs. *Introduction to Coalgebra: Towards Mathematics of States and Observation*.
Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2016.
- [16] Bart Jacobs and Jan Rutten. “An introduction to (co)algebra and (co)induction”. In:
Advanced Topics in Bisimulation and Coinduction. Ed. by Davide Sangiorgi and Jan Rutten.
Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2011, pp. 38–99.
- [17] Lawrence Johnson, Georgios N. Yannakakis, and Julian Togelius.
“Cellular automata for real-time generation of infinite cave levels”.
In: *Proceedings of the 2010 Workshop on Procedural Content Generation in Games*. PCGames '10.
Monterey, California: Association for Computing Machinery, 2010. ISBN: 9781450300230.
DOI: [10.1145/1814256.1814266](https://doi.org/10.1145/1814256.1814266). URL: <https://doi.org/10.1145/1814256.1814266>.
- [18] André Joyal, Mogens Nielsen, and Glynn Winskel. “Bisimulation from Open Maps”.
In: *Information and Computation* 127 (1996).
- [19] Jarkko Kari. “The Nilpotency Problem of One-Dimensional Cellular Automata”.
In: *SIAM Journal on Computing* 21.3 (1992), pp. 571–586. DOI: [10.1137/0221036](https://doi.org/10.1137/0221036).
eprint: <https://doi.org/10.1137/0221036>. URL: <https://doi.org/10.1137/0221036>.
- [20] Jarkko Kari. “Theory of cellular automata: A survey”.
In: *Theoretical Computer Science* 334.1 (2005), pp. 3–33. ISSN: 0304-3975.
DOI: <https://doi.org/10.1016/j.tcs.2004.11.021>.
URL: <https://www.sciencedirect.com/science/article/pii/S030439750500054X>.
- [21] Jarkko Kari and Nicolas Ollinger. “Periodicity and Immortality in Reversible Computing”.
In: *Mathematical Foundations of Computer Science 2008*.
Ed. by Jerzy Ochmasi Edward and Tyszkiewicz. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008,
pp. 419–430. ISBN: 978-3-540-85238-4.
- [22] Bartek Klin. “Coalgebraic Modal Logic Beyond Sets”.
In: *Electronic Notes in Theoretical Computer Science* (173 2007).
- [23] Clemens Kupke and Dirk Pattinson. “Coalgebraic semantics of modal logics: An overview”.
In: *Theoretical Computer Science* 412.38 (2011). CMCS Tenth Anniversary Meeting, pp. 5070–5094.
ISSN: 0304-3975. DOI: <https://doi.org/10.1016/j.tcs.2011.04.023>.
URL: <https://www.sciencedirect.com/science/article/pii/S0304397511003215>.
- [24] Petr Krka. “On topological dynamics of Turing machines”.
In: *Theoretical Computer Science* 174.1 (1997), pp. 203–216. ISSN: 0304-3975.
DOI: [https://doi.org/10.1016/S0304-3975\(96\)00025-4](https://doi.org/10.1016/S0304-3975(96)00025-4).
URL: <https://www.sciencedirect.com/science/article/pii/S0304397596000254>.
- [25] Tom Leinster. *Basic Category Theory*. Vol. 143. Cambridge Studies in Advanced Mathematics.
Cambridge University Press, July 2014. URL: <http://www.cambridge.org/9781107044241>.
- [26] Saunders MacLane. *Categories for the Working Mathematician*.
Graduate Texts in Mathematics, Vol. 5. New York: Springer-Verlag, 1971, pp. ix+262.

- [27] Surendra Kumar Nanda, Suneeta Mohanty, and Prasant Kumar Pattnaik. “A Survey on Cellular Automata-Based Security Systems for Information Security and Optimized Performance”. In: *Proceedings of 2nd International Conference on Smart Computing and Cyber Security*. Ed. by Prasant Kumar Pattnaik, Mangal Sain, and Ahmed A. Al-Absi. Singapore: Springer Nature Singapore, 2022, pp. 235–242. ISBN: 978-981-16-9480-6.
- [28] S. Nandi, B.K. Kar, and P. Pal Chaudhuri. “Theory and applications of cellular automata in cryptography”. In: *IEEE Transactions on Computers* 43.12 (1994), pp. 1346–1357. DOI: [10.1109/12.338094](https://doi.org/10.1109/12.338094).
- [29] Nicolas Ollinger. “Universalities in Cellular Automata”. In: Grzegorz Rozenberg, Thomas Bäck, and Joost N. Kok. *Handbook of Natural Computation*. Springer-Verlag Berlin Heidelberg, 2012.
- [30] Dirk Pattinson. “Coalgebraic modal logic: soundness, completeness and decidability of local consequence”. In: *Theoretical Computer Science* 309 (2003). DOI: [10.1016/S0304-3975\(03\)00201-9](https://doi.org/10.1016/S0304-3975(03)00201-9).
- [31] Dirk Pattinson. “Semantical Principles in the Modal Logic of Coalgebras”. In: *STACS 2001*. Ed. by Afonso Ferreira and Horst Reichel. Berlin, Heidelberg: Springer Berlin Heidelberg, 2001, pp. 514–526. ISBN: 978-3-540-44693-4.
- [32] *Physica D: Nonlinear Phenomena*. Vol. 10. 1984. URL: <https://www.sciencedirect.com/journal/physica-d-nonlinear-phenomena/vol/10/issue/1>.
- [33] *Physica D: Nonlinear Phenomena*. Vol. 45. 1990. URL: <https://www.sciencedirect.com/journal/physica-d-nonlinear-phenomena/vol/45/issue/1>.
- [34] Dan (sigfpe) Piponi. *Evaluating cellular automata is comonadic*. Accessed 4 July 2024. 2006. URL: <http://blog.sigfpe.com/2006/12/evaluating-cellular-automata-is.html>.
- [35] Emily Riehl. *Category Theory in Context*. Dover, 2016. URL: <http://www.math.jhu.edu/~5C-%7B%7Deriehl/context.pdf>.
- [36] Zsuzsanna Róka. “Simulations between cellular automata on Cayley graphs”. In: *Theoretical Computer Science* 225 (1999).
- [37] Grzegorz Rozenberg, Thomas Bäck, and Joost N. Kok. *Handbook of Natural Computation*. Springer-Verlag Berlin Heidelberg, 2012.
- [38] Jan J. M. M. Rutten. “The Method of Coalgebra: exercises in coinduction”. In: *World Congress on Formal Methods*. 2019. URL: <https://api.semanticscholar.org/CorpusID:86847919>.
- [39] Jan J.M.M. Rutten. “Universal coalgebra: a theory of systems”. In: *Theoretical Computer Science* 249.1 (2000). Modern Algebra, pp. 3–80. ISSN: 0304-3975. DOI: [https://doi.org/10.1016/S0304-3975\(00\)00056-6](https://doi.org/10.1016/S0304-3975(00)00056-6). URL: <https://www.sciencedirect.com/science/article/pii/S0304397500000566>.
- [40] Davide Sangiorgi. *Introduction to Bisimulation and Coinduction*. Cambridge University Press, 2011.
- [41] Lutz Schröder. “Expressivity of coalgebraic modal logic: The limits and beyond”. In: *Theoretical Computer Science* (390 2008).
- [42] Alvy Ray Smith. “Simple Computation-Universal Cellular Spaces”. In: *J. ACM* 18.3 (July 1971), pp. 339–353. ISSN: 0004-5411. DOI: [10.1145/321650.321652](https://doi.org/10.1145/321650.321652). URL: <https://doi.org/10.1145/321650.321652>.
- [43] Sam Staton. “Relating coalgebraic notions of bisimulation”. In: *Logical Methods in Computer Science* 7.(1:13) (2011). DOI: [10.2168/LMCS-7](https://doi.org/10.2168/LMCS-7). URL: <https://doi.org/10.2168/LMCS-7>.

- [44] Penelope Sweetser and Janet Wiles. “Combining Influence Maps and Cellular Automata for Reactive Game Agents”. In: *Intelligent Data Engineering and Automated Learning - IDEAL 2005*. Ed. by Marcus Gallagher, James P. Hogan, and Frederic Maire. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 524–531. ISBN: 978-3-540-31693-0.
- [45] Véronique Terrier. “Language Recognition by Cellular Automata”. In: Grzegorz Rozenberg, Thomas Bäck, and Joost N. Kok. *Handbook of Natural Computation*. Springer-Verlag Berlin Heidelberg, 2012.
- [46] Oleksii O Vodka and Mariia I Shapovalova. “Exploration of cellular automata: a comprehensive review of dynamic modeling across biology, computer and materials science”. In: *Computer Methods in Materials Science* 23.4 (2023), p. 58. URL: <https://doi.org/10.7494/cmms.2023.4.0820>.
- [47] Gabriel Wainer, Qi Liu, Olivier Dalle, and Bernard P. Zeigler. “Applying Cellular Automata and DEVS Methodologies to Digital Games: A Survey”. In: *Simulation & Gaming* 41.6 (2010), pp. 796–823. DOI: [10.1177/1046878110378708](https://doi.org/10.1177/1046878110378708). eprint: <https://doi.org/10.1177/1046878110378708>. URL: <https://doi.org/10.1177/1046878110378708>.
- [48] Baltasar Trancón y Widemann and Michael Hauhs. “Distributive-Law Semantics for Cellular Automata and Agent-Based Models”. In: *Algebra and Coalgebra in Computer Science*. Ed. by Andrea Corradini, Bartek Klin, and Corina Cirstea. Springer-Verlag Berlin Heidelberg, 2011. DOI: [10.1007/978-3-642-22944-2](https://doi.org/10.1007/978-3-642-22944-2).
- [49] Stephen Wolfram. “Cryptography with Cellular Automata”. In: *Advances in Cryptology — CRYPTO ’85 Proceedings*. Ed. by Hugh C. Williams. Berlin, Heidelberg: Springer Berlin Heidelberg, 1986, pp. 429–432. ISBN: 978-3-540-39799-1.
- [50] Stephen Wolfram. “Statistical mechanics of cellular automata”. In: *Rev. Mod. Phys.* 55 (3 July 1983), pp. 601–644. DOI: [10.1103/RevModPhys.55.601](https://doi.org/10.1103/RevModPhys.55.601). URL: <https://link.aps.org/doi/10.1103/RevModPhys.55.601>.