

A Max-Plus Model of Asynchronous Cellular Automata

Ebrahim L. Patel

*Mathematical Institute, University of Oxford, Woodstock Road
Oxford OX2 6GG, UK*

David Broomhead

*School of Mathematics, University of Manchester, Oxford Road
Manchester M13 9PL, UK*

This paper presents a new framework for asynchrony. This has its origins in our attempts to better harness the internal decision-making process of cellular automata (CA). Thus, we show that a max-plus algebraic model of asynchrony arises naturally from the CA requirement that a cell receive the state of each neighbor before updating. The significant result is the existence of a bijective mapping between the asynchronous system and the synchronous system classically used to update CA. Consequently, although the CA outputs look qualitatively different, when surveyed on “contours” of real time, the asynchronous CA replicates the synchronous CA. Moreover, this type of asynchrony is simple—it is characterized by the underlying network structure of the cells, and long-term behavior is deterministic and periodic due to the linearity of max-plus algebra. The findings lead us to proffer max-plus algebra as: (i) a more accurate and efficient underlying timing mechanism for models of patterns seen in nature; and (ii) a foundation for promising extensions and applications.

1. Introduction

A cellular automaton (or CA, where we also abbreviate the plural “cellular automata” to CA, allowing the context to save confusion) is a discrete dynamical system, consisting of an array of identical cells, each possessing a state. The states evolve, according to some local rule, in discrete time steps. The first CA models were synchronous, where all cells update once on each time step, and the foundations of the study of these CA were laid by Wolfram in the 1980s [1, 2]. A popular application of such CA is the construction of models of natural growth processes such as seashell patterns and snowflakes [3, 4]. Figure 1 (right) shows a CA pattern typically examined by Wolfram inscribed on a seashell; the similarity to the real seashell pattern (on the left) is evident. A natural extension is the introduction of asynchronous update times. Indeed, in terms of seashell patterns, Gunji

demonstrated that different asynchronous update methods yield different patterns, leading to the conjecture that asynchrony is intrinsic to living systems [5].




Figure 1. Seashell patterns: left is real; right is generated by a CA [3].

A preliminary observation of asynchronous CA was made in [6], where the authors compared the properties of synchronous CA with two types of CA that iterate asynchronously. Subsequent studies employed the methods of [6] as special cases to conduct specific studies into asynchronous CA [7–11]. Many of these authors attest to asynchrony as being stochastic in nature. This is a general viewpoint in light of their applications: such asynchrony relies on continuous time [11] and is also likely to be more robust [7, 9, 11], thereby aiding a better description of biological phenomena. For example, given a system of coupled cells, the update of cell states depends on a predefined probability [6, 11]. This consequently also led Schönfish and de Roos to conjecture that while synchronous updating can produce periodic orbits, asynchronous systems will only yield patterns that converge to a fixed point or patterns that are chaotic [11].

The argument for asynchronous updating's being stochastic has been challenged by authors such as Cornforth et al. [12]. They claim that such probabilistic updating schemes are used because of the oversimplification of biologically inspired models. They further argue for mimicking appropriate aspects of nature more closely to create better computational models. Thus, the authors have drawn attention to a large class of behaviors of natural processes, in which the updating is asynchronous but not stochastic [13].

Moreover, underlying a synchronous update scheme is the notion of a “global clock,” in the sense that it assigns the same update time

to all cells. In recent years, the disadvantages of synchrony in this context have been exposed [14]. On the other hand, a remedy has been presented from the perspective of parallel computing devices, wherein such a distribution of a global signal proves costly. Thus, the proposal is to allow cells to determine their own update time through local interactions (see [14] and the references therein, including [15]). Although the storage cost is higher, this scheme points toward a more natural form of computation.

In [1], Wolfram explored synchronous CA on a one-dimensional lattice, where cells take the Boolean states 1 or 0. The CA state of cell i was dependent on the states of three connected cells, called the *neighborhood* of i . These three-cell neighborhood CA were termed elementary cellular automata (ECA).

The ECA may be regarded as special cases of random Boolean networks. The different types of these networks were first classified by Gershenson in 2002 [16]. The types considered included asynchronous random Boolean networks (ARBNs), in which nodes are selected to be randomly updated at each time step, and deterministic ARBNs (DARBNs), where the node to be updated is selected deterministically. Gershenson talked of DARBNs as being more advantageous because of their modeling capabilities, which are more straightforward than ARBNs that rely on the stochasticity of asynchronous phenomena. Gershenson further proposed DARBNs as better representations of the famous genetic regulatory networks of Kauffman [17], as they are asynchronous but do not rely on stochastic methods.

Following on from Gershenson's idea of using determinism as a more "model-friendly" form for asynchrony, a goal of this paper is to exploit this avenue by presenting a new, deterministic framework for asynchrony. No matter how well they match the real system, we claim that the essence of many interesting and important asynchronous processes is lost by using probabilistic updating schemes.

As a real example, consider Figure 1 again, showing a seashell pattern; it is interesting to see that the same pattern (and many other such seashell patterns, as well as growth processes such as snowflakes [4]) may be replicated quite accurately by a CA model. Traditionally, some difference between the two versions—real and CA—would be ascribed to a fault, random or otherwise, such that a better approximation may be obtained by adding stochastic asynchrony in the CA rule. As alluded to by Cornforth et al. [12], such asynchrony has tended to simplify these dynamics into a probability (or the like) of cells updating their states. Inspired by the ideas of "local clocks over global clocks" in [14] and [15], we will get into the heart of the matter and study the pattern of exchanges that take place locally, that is, before any data is transferred between cells.

1.1 Network Description of the Cellular Automaton Lattice

Our work views each cell in a one-dimensional lattice as a processor that receives input from its neighborhood. Having received this input, the processor computes its new state (as a function of the input states), then sends a corresponding output to its connected neighbors. This type of information exchange can be represented by the digraph in Figure 2, where each node represents a processor, and directed arcs between nodes indicate the direction of information transfer.

Figure 2 shows three arcs pointing to each node, indicating that there are three processors—therefore three neighbors (including i itself)—sending information to each process i . Thus, the neighborhood size n of each node is $n = 3$; we also refer to such a neighborhood as an n -neighborhood (or n -nbhd). The figure particularly shows that each processor sends output to itself as well as to its left and right neighbors. We refer to this type of network as a *regular n -nbhd network* or simply a regular network if n is understood. Thus, the regular 3-nbhd network describes the lattice for the ECA, where cells are depicted by nodes. For this reason, we use the terms *cell* and *node* to mean the same thing.




Figure 2. Digraph representing a regular 3-nbhd network.

It is also assumed that there is a processing delay associated with each nodal computation of CA state. Moreover, we incorporate a transmission delay, which is the time taken for a CA state to be transmitted to other cells that require it. These two parameters are the means by which we obtain asynchrony: a divergence from classical (synchronous) CA models, since there the computations are assumed to occur instantaneously.

1.2 Contents

In Section 2, we review the ECA, a classical synchronous model of CA, which is later used to present our asynchronous model. In Section 3, we show how max-plus algebra provides a natural way to mathematically model asynchronous CA. By covering related graph theoretical techniques and known proven results, we also show that this asynchronous system is periodic, characterized by the connectivity of the underlying network. It will be seen from the depth of theory covered that a max-plus algebraic model of asynchrony is more bespoke and addresses *all* the intricacies of the internal dynamics within

a cell (not just the external). Nevertheless, those already familiar with max-plus algebra may skim over Sections 3.1 and 3.2, noting the few places where we mention the link to our system. The “contour plot” is introduced as a framework for this asynchrony in Section 3. This is followed by Section 4, which shows the effect of the max-plus asynchrony on cellular automata. We finish with concluding remarks.

2. Synchronous Cellular Automata

Let s_i denote the state of cell i . The index i denotes the position of the cell in the one-dimensional Euclidean plane, so that the state of the CA at a given time $t \in \mathbb{R}$ is represented by the string $s_1(t) s_2(t) \dots s_N(t)$ or the vector $s(t) = (s_1(t), s_2(t), \dots, s_N(t))$, where N is the size of the lattice. The ECA assigns a symmetrical neighborhood of three nodes, where node i is contained in its own neighborhood (as in Figure 2). Here, the CA rule is a function $f: \{0, 1\}^3 \rightarrow \{0, 1\}$ given by $s_i(t + 1) = f(s_{i-1}(t), s_i(t), s_{i+1}(t))$. An example of such a CA rule is the following:

$$s_i(t + 1) = \sum_{j=i-1}^{j=i+1} s_j(t) \pmod 2. \tag{1}$$

The rule in equation (1) was named ECA rule 150 by Wolfram [2], and this is how we refer to it throughout the paper.

Consider a regular 3-nbhd network of 20 cells (connected in a ring such that end cells are adjacent). If each cell is depicted by a square—colored if $s_i = 1$, clear if $s_i = 0$ —then the output produced after each of 30 iterations of equation (1) is shown in Figure 3. For fixed t , the CA state $s_1(t) \dots s_{20}(t)$ represents a horizontal line of cell states, and the initial CA state is

$$s_i(0) = \begin{cases} 1 & \text{if } i = 10 \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

Such an output as Figure 3 is referred to as a *space-time pattern* (or plot). It is evident that the absence of transmission and processing delays means that the t^{th} update of each cell occurs at the same time as the t^{th} update of every other cell, that is, synchronously.

2.1 State Transition Graph

Figure 3 displays periodic behavior. In other words, for each time step t , the CA yields a state $s(t)$, which is seen again after a few more time




Figure 3. CA pattern produced by ECA rule 150. Time travels vertically down; cells are labeled on the x axis.

steps. Given the initial CA state $s(0)$ and the CA rule f , an *orbit* of $s(0)$ is the sequence of states obtained by applying f on $s(0)$ sequentially. If f is applied k times, we represent this as

$$f^k(s(0)) = \underbrace{f(f(\dots f(s(0))))}_{k \text{ times}}.$$

We define periodic behavior as follows.

Definition 1. Consider the CA rule f and network size N . Let $s(k) = f^k(s(0))$ for all $k \geq 0$, where $s(k)$ is a CA state represented by a $1 \times N$ vector. For some $t \geq 0$, if there exists a finite number $p \in \mathbb{N}$ such that $s(t+p) = s(t)$, then the set of states

$$\{s(t), s(t+1), \dots, s(t+p-1)\}$$

is called a periodic CA orbit, where p is the CA period of the orbit.

As an example, consider the underlying network of size 4 as given in Figure 4. We give the CA rule in words: the new state of each cell is the sum of the states (mod 2) of its neighborhood cells on the previous time step, where the neighborhood of cell i comprises those cells whose outgoing arc points to i . (Note that this rule is an extension of ECA rule 150 to arbitrary lattices.) For small N , as is the case here, it is useful to represent each system state as a vertex in a digraph. Thus, there is an arc from CA state s^i to CA state s^{ii} if and only if

$f(s^i) = s^{i+1}$. The digraph is known as a *state transition graph* (STG). For the system in question, the STG is given in Figure 5. Each CA state is shown in string form, where the i^{th} digit represents the CA state of the i^{th} node.




Figure 4. Size 4 network of CA cells.




Figure 5. STG of ECA rule 150 generalized to a lattice of size 4 (namely the network given in Figure 4).

To determine the evolution of the CA, we can follow the arcs in Figure 5. It can be seen that all initial CA states asymptotically evolve into four periodic orbits, represented as circuits in the STG. Two of these (states 0000 and 1001) are period-1 orbits (or *fixed points* in conventional dynamical systems language), and two are period-3 orbits; all other states are transient. Such an STG is an artifact of a synchronous CA.

3. Asynchronous Model

Consider a synchronous CA. Due to the synchrony, it is possible to draw horizontal lines in the corresponding space-time pattern such that each line represents the update times of all cells at some fixed time. The CA is a discrete time, dynamical system, so the horizontal lines may be drawn in sequence, evenly spaced, as in Figure 6(a). We call such a space-time plot a *contour plot*, and each horizontal line is referred to as a *contour*. The contour plot may be thought of as a frame on which the CA states are overlaid and simulated. Now consider altering these contours so that cells do not necessarily update synchronously. The corresponding contour plot may then look like Figure 6(b), which shows the contours having variable shapes—updates occur asynchronously. We shall return to the contour plot after looking at how our asynchronous model may produce it.




Figure 6. (a) Synchronous and (b) asynchronous contour plots. The contours indicate update times of cells and act as a frame on which the CA may be evolved. Both lattices are connected as a regular 3-nbhd network of 20 cells. Time is on the vertical axis, and the horizontal axis represents the cell positions. In (a), contours are horizontal. In (b), update k of all cells is represented by contour k (counting from the top).

We present the asynchronous model as follows. Consider a cell i contained in a regular n -nbhd network of N cells. The cell carries a CA state (1 or 0) that changes with time, depending on the rules that we employ. Thus, we can plot points on the real line corresponding to when these changes occur. The real line represents time, and the points are the *update times* of the CA state. Let $x_i(k)$ denote the k^{th} update time for cell i . We also refer to k as a *cycle number*. Once each cell in the neighborhood of cell i has completed its k^{th} cycle, it sends the updated state to i . The transmission of such a state from cell j to i takes *transmission time* $\tau_{ij}(k)$. The update of cell i takes a *processing*

time, and it is represented in the k^{th} cycle by $\xi_i(k)$. If $n = 3$, we have the following iterative scheme for the $(k + 1)^{\text{th}}$ update time of cell i :

$$x_i(k + 1) = \max \{x_{i-1}(k) + \tau_{i,i-1}(k), x_i(k) + \tau_{i,i}(k), x_{i+1}(k) + \tau_{i,i+1}(k)\} + \xi_i(k + 1). \tag{3}$$

This sequence of interactions yielding a state change is depicted in Figure 7. Notice that we have now expanded on the simpler notion of intercellular communication by focusing more closely on intracellular communication, that is, the communication within a cell itself. This is the key to our study of asynchrony, and it has arisen naturally from the requirement that a CA cell know the states of its neighbors.




Figure 7. The processes internal to the k^{th} state change at cell i . Real time travels vertically down. Arrows indicate the destination of the sent state. Labels on arrows indicate the time taken for the corresponding process.

We refer to these internal processes as *events*. In Figure 7, there are two significant types of events: “receive” and “send.” The three times $x_{i-1}(k)$, $x_i(k)$, and $x_{i+1}(k)$ are send event times (i.e., when the corresponding CA states are sent). The time

$$\max \{x_{i-1}(k) + \tau_{i,i-1}(k), x_i(k) + \tau_{i,i}(k), x_{i+1}(k) + \tau_{i,i+1}(k)\}$$

is when node i receives the aforementioned send states; it is therefore a receive event. Once received, node i processes its new CA state (by applying a CA rule on the received states); this takes time duration $\xi_i(k + 1)$. Once processed, node i sends its state to connected nodes at

time $x_i(k+1)$; this is another send event. Note that we make no distinction between update times and events; both are the same. However, we formally refer to $x_i(k+1)$ as the update time of i , as introduced earlier.

Note 1. This idea of a network depicting events in a space-time diagram is not too dissimilar from Wolfram's *causal network* [25, Chapter 9]. We shall delve into this relationship later.

3.1 Preliminaries: Max-Plus Algebra and Graph Theory

The max operation enables one to interpret equation (3) in max-plus algebra. This is useful because it converts a nonlinear system into a linear system in this new algebra, which subsequently shares many important features with conventional linear algebra. Particularly novel applications include the modeling of railway network timetables [18], manufacturing processes [19], and even cellular protein production [20]. In presenting max-plus algebra, we borrow most notation and terminology from [18].

Define $\varepsilon = -\infty$ and $e = 0$, and denote by \mathbb{R}_{\max} the set $\mathbb{R} \cup \{\varepsilon\}$. For elements $a, b \in \mathbb{R}_{\max}$, define operations \oplus and \otimes by

$$a \oplus b = \max(a, b) \text{ and } a \otimes b = a + b.$$

The set \mathbb{R}_{\max} together with the operations \oplus and \otimes is what we refer to as max-plus algebra and it is denoted by $\mathcal{R}_{\max} = (\mathbb{R}_{\max}, \oplus, \otimes, \varepsilon, e)$. $\varepsilon = -\infty$ is the "zero" (i.e., $\forall x \in \mathbb{R}_{\max}$, $\varepsilon \otimes x = x \otimes \varepsilon = \varepsilon$, and $\varepsilon \oplus x = x \oplus \varepsilon = x$), while $e = 0$ is the "unit" element (i.e., $\forall x \in \mathbb{R}_{\max}$, $e \otimes x = x \otimes e = x$).

\mathcal{R}_{\max} is associative and commutative over both operations \oplus and \otimes , while \otimes is distributive over \oplus . (In addition, \oplus is idempotent in \mathcal{R}_{\max} , so that max-plus algebra is a commutative and idempotent semiring.)

In this paper, we remove the dependence on k of the processing and transmission times, so that equation (3) is written

$$x_i(k+1) = \max\{x_{i-1}(k) + \tau_{i,i-1}, x_i(k) + \tau_{i,i}, x_{i+1}(k) + \tau_{i,i+1}\} + \xi_i. \quad (4)$$

We can now write this in max-plus algebra:

$$x_i(k+1) = \{(\tau_{i,i-1} \otimes x_{i-1}(k)) \oplus (\tau_{i,i} \otimes x_i(k)) \oplus (\tau_{i,i+1} \otimes x_{i+1}(k))\} \oplus \xi_i. \quad (5)$$

We often save space and clarify the presentation by omitting \otimes , much as in conventional algebra. Thus, $x \otimes y \equiv xy$ and equation (5) can be written as

$$x_i(k+1) = \xi_i \{ \tau_{i,i-1} x_{i-1}(k) \oplus \tau_{i,i} x_i(k) \oplus \tau_{i,i+1} x_{i+1}(k) \}. \quad (6)$$

Since \otimes is distributive over \oplus , we can write this as

$$x_i(k + 1) = \xi_i \tau_{i,i-1} x_{i-1}(k) \oplus \xi_i \tau_{i,i} x_i(k) \oplus \xi_i \tau_{i,i+1} x_{i+1}(k). \tag{7}$$

To represent a max-plus power, we follow from the associativity of \otimes and define, for $x \in \mathbb{R}_{\max}$,

$$x^{\otimes n} \stackrel{\text{def}}{=} \underbrace{x \otimes x \otimes \dots \otimes x}_{n \text{ times}} \tag{8}$$

for all $n \in \mathbb{N}$ with $n \neq 0$. For $n = 0$, we define $x^{\otimes 0} = e$ ($= 0$).

Max-plus algebra extends naturally to matrices, and this allows the concurrent modeling of the update times for all nodes. Denote the set of $n \times m$ matrices with underlying max-plus algebra by $\mathbb{R}_{\max}^{n \times m}$. The sum of matrices $A, B \in \mathbb{R}_{\max}^{n \times m}$, denoted by $A \oplus B$, is defined by

$$[A \oplus B]_{ij} = a_{ij} \oplus b_{ij}, \tag{9}$$

where $a_{ij} = [A]_{ij}$ and $b_{ij} = [B]_{ij}$. In the same vein, for matrices $A \in \mathbb{R}_{\max}^{n \times l}$ and $B \in \mathbb{R}_{\max}^{l \times m}$, the matrix product $A \otimes B$ is defined by

$$[A \otimes B]_{ij} = \bigoplus_{k=1}^l a_{ik} \otimes b_{kj} = \max_{k \in \{1, \dots, l\}} \{a_{ik} + b_{kj}\}. \tag{10}$$

For $\alpha \in \mathbb{R}$, the scalar multiple $\alpha \otimes A$ is defined by

$$[\alpha \otimes A]_{ij} = \alpha \otimes a_{ij}. \tag{11}$$

As in classical matrix manipulation, the max-plus matrix addition \oplus is associative and commutative, while the matrix product \otimes is associative and distributive with respect to \oplus ; it is usually not commutative. Similarly, the operation \otimes has priority over \oplus .

The elements of $\mathbb{R}_{\max}^n \stackrel{\text{def}}{=} \mathbb{R}_{\max}^{n \times 1}$ are called *vectors*. A vector is usually written in bold, as in \mathbf{x} , while the vector with all elements equal to e is called the *unit vector* and is denoted by \mathbf{u} .

We are now able to define matrix-vector products. The product $A \otimes \mathbf{x}$, where $A \in \mathbb{R}_{\max}^{n \times m}$ and $\mathbf{x} \in \mathbb{R}_{\max}^m$, is

$$[A \otimes \mathbf{x}]_i = \bigoplus_{k=1}^m a_{ik} \otimes x_k = \max_{k \in \{1, \dots, m\}} \{a_{ik} + x_k\}. \tag{12}$$

Moreover, for the square matrix $A \in \mathbb{R}_{\max}^{n \times n}$, denote the k^{th} power of A by $A^{\otimes k}$, defined by

$$A^{\otimes k} \stackrel{\text{def}}{=} \underbrace{A \otimes A \otimes \dots \otimes A}_{n \text{ times}} \tag{13}$$

for all $k \in \mathbb{N}$ with $k \neq 0$. For $k = 0$, we set $A^{\otimes 0} \stackrel{\text{def}}{=} E(n, n)$, the identity matrix whose diagonal elements equal e and all of whose other elements are ε .

Having established the preliminaries above, a system of N such equations as (7) can now be given in the form

$$\mathbf{x}(k + 1) = P \otimes \mathbf{x}(k), \tag{14}$$

where $\mathbf{x}(k) = (x_1(k), x_2(k), \dots, x_N(k))^T$. P is the $N \times N$ matrix defined by $A_\xi \otimes T$, where

$$A_\xi = \begin{pmatrix} \xi_1 & \varepsilon & \cdots & \varepsilon \\ \varepsilon & \xi_2 & \cdots & \varepsilon \\ \vdots & & \ddots & \vdots \\ \varepsilon & \varepsilon & \cdots & \xi_N \end{pmatrix}$$

and

$$T = \begin{pmatrix} \tau_{11} & \tau_{12} & \varepsilon & \varepsilon & \cdots & \varepsilon & \tau_{1N} \\ \tau_{21} & \tau_{22} & \tau_{23} & \varepsilon & \cdots & \varepsilon & \varepsilon \\ \vdots & & & \ddots & & & \vdots \\ \varepsilon & \varepsilon & \cdots & \varepsilon & \tau_{N-1,N-2} & \tau_{N-1,N-1} & \tau_{N-1,N} \\ \tau_{N,1} & \varepsilon & \cdots & \varepsilon & \varepsilon & \tau_{N,N-1} & \tau_{NN} \end{pmatrix}.$$

A_ξ is referred to as the *processing matrix* and T is the *transmission matrix*. We call equation (14) a *max-plus system* (of dimension N), where the vector $x(k)$ is the *state of the system*. P is called the *timing dependency matrix* of the network of cells. Note that here, the term “state” refers to update time and is not to be confused with “CA state.” Nevertheless, the context should make this distinction clear. The name timing dependency matrix (along with the later *timing dependency graph*) is inspired by its use in another novel application of \mathcal{R}_{\max} to the timing of digital hardware [21].

To a network of cells, we associate a *digraph*, such as in Figure 2. In general, we define a digraph as $\mathcal{G} = (V, E)$, consisting of a set V and a set E of ordered pairs (a, b) of V . (Often, we refer to the digraph simply as a graph, which, in turn, also refers to the “network” of our application.) The elements of V are called *vertices* or *nodes*, and those of E are *arcs*.

An arc (a, b) is also denoted $a b$, and we refer to an arc $a a$ as a *self-loop*. For the arc $a b$, a is the *start node* and b is the *end node*; a is also referred to as the *predecessor* node of b , while b is the *successor* of a .

By assigning real numbers (called *weights*) to the arcs of a graph $\mathcal{G} = (V, E)$, we obtain a *weighted graph*. The *weighted adjacency matrix of \mathcal{G} over \mathbb{R}_{\max}* is the matrix $W \in \mathbb{R}_{\max}^{n \times n}$ whose $(i, j)^{\text{th}}$ entry w_{ij} is nonzero ($\neq \varepsilon$) if and only if j is a predecessor of i . We also refer to W as a *max-plus adjacency matrix*. Given $W \in \mathbb{R}_{\max}^{n \times n}$, we denote the associated network as $\mathcal{G}(W)$. Thus, P in equation (14) is a max-plus adjacency matrix of the regular 3-nbhd network, since the network is exactly that shown in Figure 2 augmented with arc weights $\xi_i \tau_{ij}$. In fact, since our particular max-plus system concerns update times, we call $\mathcal{G}(P)$ the *timing dependency graph* of the system. The neighborhood of a node may now be defined in terms of the adjacency matrix.

Definition 2. Let W be the max-plus adjacency matrix of a digraph of connected cells. The neighborhood of i is $\mathcal{N}_i = \{j \mid w_{ij} \neq \varepsilon\}$.

Definition 3. Let $p = \{a_1, a_2, \dots, a_n\}$ be a sequence of arcs. If there are vertices v_0, v_1, \dots, v_n (not necessarily distinct) such that $a_j = v_{j-1} v_j$ for $j = 1, \dots, n$, then p is called a walk from v_0 to v_n . A walk for which the a_j are distinct is called a path. Such a path is said to consist of the nodes v_0, v_1, \dots, v_n and to have length n , which is denoted $|p| = n$.

If $v_n = v_0$, then the path is called a circuit. If the nodes in the circuit are all distinct (i.e., $v_i \neq v_k$ for $i \neq k$), then it is called an elementary circuit.

We define the *weight* $|p|_w$ of a path p as the sum of the weights of all arcs constituting the path. The *average weight* of p is $|p|_w / |p|$. For a circuit, we refer to this quantity as the *average circuit weight*.

Definition 4. For a graph $\mathcal{G} = (V, E)$, node $j \in V$ is said to be reachable from node $i \in V$, denoted $i \rightarrow j$, if there exists a path from i to j . Graph \mathcal{G} is strongly connected if $i \rightarrow j$ for any two nodes $i, j \in V$.

Moreover, matrix $A \in \mathbb{R}_{\max}^{n \times n}$ is called *irreducible* if $\mathcal{G}(A)$ is strongly connected; if a matrix is not irreducible, it is called *reducible*. Thus, in our max-plus system of equation (14), matrix P is irreducible.

Definition 5. Denote the cyclicity of a graph \mathcal{G} by $\sigma_{\mathcal{G}}$.

- If \mathcal{G} is strongly connected, then $\sigma_{\mathcal{G}}$ equals the greatest common divisor of the lengths of all elementary circuits in \mathcal{G} . If \mathcal{G} consists of only one node without a self-loop, then $\sigma_{\mathcal{G}}$ is defined to be one.
- If \mathcal{G} is not strongly connected, then $\sigma_{\mathcal{G}}$ equals the least common multiple of the cyclicities of all maximal strongly connected subgraphs (MSCSs) of \mathcal{G} . (See Appendix A for definition of MSCSs.)

Definition 6. Let $A \in \mathbb{R}_{\max}^{n \times n}$ be irreducible. The *cyclicity* of A , denoted $\sigma(A)$, is defined as the cyclicity of the critical graph of A .

When the matrix is understood, the cyclicity is also denoted by σ .

3.2 Asymptotic Behavior of the Max-Plus System

Let $\mathbf{x}(0)$ represent the initial state of all cells. Then we can rewrite equation (14) as

$$\mathbf{x}(k+1) = P \otimes P \otimes \cdots \otimes P \otimes \mathbf{x}(0) = P^{\otimes(k+1)} \otimes \mathbf{x}(0) \quad (15)$$

or, equivalently,

$$\mathbf{x}(k) = P^{\otimes k} \otimes \mathbf{x}(0). \quad (16)$$

Given $\mathbf{x}(0)$, the sequence of vectors $\{\mathbf{x}(k) : k \in \mathbb{N}_0\}$, obtained by iterating equation (16), is referred to as the *orbit* of $\mathbf{x}(0)$. A study of such sequences is provided in greater scope in [18, Chapters 3 and 4]. We detail the topics relevant for this work, with some well-known results expanded upon in Appendix B.

Definition 7. Let $A \in \mathbb{R}_{\max}^{n \times n}$. If $\lambda \in \mathbb{R}_{\max}$ is a scalar and $\mathbf{v} \in \mathbb{R}_{\max}^n$ is a vector that contains at least one finite element such that

$$A \otimes \mathbf{v} = \lambda \otimes \mathbf{v}, \quad (17)$$

then λ is called an eigenvalue of A , and \mathbf{v} is an eigenvector of A associated with eigenvalue λ .

For a system with irreducible P such as equation (14) (or (16)), it turns out that there is only one eigenvalue, and it is equal to the maximal average weight of elementary circuits in $\mathcal{G}(P)$ (see Appendix B, Theorem 2). Such circuits with maximal average weight are called *critical*, and the *critical graph* of P is the graph consisting only of critical circuits in $\mathcal{G}(P)$.

Definition 8. Let $A \in \mathbb{R}_{\max}^{n \times n}$. For some $k \geq 0$, consider the set of vectors

$$\mathbf{x}(k), \mathbf{x}(k+1), \mathbf{x}(k+2), \dots \in \mathbb{R}_{\max}^n,$$

where $\mathbf{x}(k) = A^{\otimes k} \otimes \mathbf{x}(0)$ for all $k \geq 0$. The set is called a (periodic) regime if there exists $\mu \in \mathbb{R}_{\max}$ and a finite number $\rho \in \mathbb{N}$ such that

$$\mathbf{x}(k+\rho) = \mu \otimes \mathbf{x}(k).$$

The period of the regime is ρ .

Suppose the initial vector $\mathbf{x}(0)$ is an eigenvector of A . Then $\mathbf{x}(k+1) = \lambda \otimes \mathbf{x}(k)$ for $k \geq 0$, so that the period is one. Thus, larger periods are obtained when the system is not initialized to an eigenvector. The remainder of Section 3 explores this in more detail. We start with the following crucial theorem of max-plus algebra.

Theorem 1. Let $A \in \mathbb{R}_{\max}^{n \times n}$ be an irreducible matrix with eigenvalue λ and cyclicity σ . Then there is a k_\star such that

$$A^{\otimes(k+\sigma)} = \lambda^{\otimes\sigma} \otimes A^{\otimes k}$$

for all $k \geq k_\star$.

Proof. See [18, Theorem 3.9]. \square

Let $\sigma = \sigma(P)$. For an indication of the asymptotic behavior of our system, we apply this theorem to observe the state at epoch $k + \sigma$ for $k \geq k_\star$:

$$\begin{aligned} \mathbf{x}(k + \sigma) &= P^{\otimes(k+\sigma)} \otimes \mathbf{x}(0) = \\ &\lambda^{\otimes\sigma} \otimes P^{\otimes k} \otimes \mathbf{x}(0) = \lambda^{\otimes\sigma} \otimes \mathbf{x}(k), \end{aligned} \tag{18}$$

where $\lambda^{\otimes\sigma}$ is read as $\lambda \times \sigma$ in terms of classical algebra. This guarantees the periodic behavior of the max-plus system, where σ is the upper bound on the period. More specifically, the period ρ is dependent on the choice of $\mathbf{x}(0)$ and, since we have seen that the system must also be periodic with period σ , we have that ρ is a factor of σ . We saw earlier that if $\mathbf{x}(0)$ is not an eigenvector, then $\rho > 1$; we now also know that such a period ρ will not be larger than σ .

We can use the above to show that the vectors $\mathbf{x}(k)$ in a regime turn out to be eigenvectors of $P^{\otimes\sigma}$ associated with eigenvalue $\lambda \times \sigma$:

$$\begin{aligned} \lambda^{\otimes\sigma} \otimes \mathbf{x}(k) &\stackrel{\text{from above}}{=} P^{\otimes(k+\sigma)} \otimes \mathbf{x}(0) = \\ P^{\otimes\sigma} \otimes P^{\otimes k} \otimes \mathbf{x}(0) &= P^{\otimes\sigma} \otimes \mathbf{x}(k). \end{aligned} \tag{19}$$

In fact, given $\mathbf{x}(0)$ and corresponding period ρ , vectors in a regime are also eigenvectors of $P^{\otimes\rho}$; this can be shown in the same way as above.

We now define a measure for the average delay between consecutive event times $x_i(k)$ and $x_i(k + 1)$.

Definition 9. Let $\{x_i(k) : k \in \mathbb{N}\}$ be an orbit of $x_i(0)$ in \mathbb{R}_{\max} . Assuming that it exists, the quantity χ_i , defined by

$$\chi_i = \lim_{k \rightarrow \infty} \frac{x_i(k)}{k},$$

is called the cycle time of i .

For an irreducible system such as ours, the vector $\chi = (x_1, x_2, \dots, x_N)$ of cycle times is unique (see Appendix B, Theorem 3). Moreover, the irreducibility of P ensures that each element of χ is the same—specifically the eigenvalue of P , that is,

$$\lim_{k \rightarrow \infty} \frac{x_i(k)}{k} = \lambda$$

for any initial condition $\mathbf{x}(0) \in \mathbb{R}^n$ (see Appendix B, Lemma 2).

Thus, since χ is independent of the initial condition, we relate it to our timing dependency graph by calling it the *cycle time vector* of P . P is irreducible in this paper, so we let $\chi = \chi_i$ for any i , and therefore refer to the cycle time vector of irreducible P simply as the cycle time of P .

As a compact summary of this subsection, we have shown that asynchrony due to an irreducible max-plus system always leads to periodic behavior, and it is characterized by the circuit(s) in $\mathcal{G}(P)$ with largest average weight. Section 3.3 addresses the impact of the above theory on the asynchronous time framework of the contour plot introduced earlier.

■ 3.3 The Contour Plot

Figure 7 can be seen as a *Hasse diagram* of events. We say that events are *causally related* if they are contained in the same *chain*. For example, the “send” events at times $x_{i-1}(k-1)$, $x_i(k-1)$, and $x_{i+1}(k-1)$ are not causally related since they are not contained in a chain—there is no path of directed arcs connecting any of the three events. Consequently, these three events form an *antichain*. (For a formal definition of a Hasse diagram, including chains and antichains, we refer the reader to Appendix D.) By connecting those elements in the same antichain, we obtain a piecewise linear plot of the vector $\mathbf{x}(k)$, which we define next.

Definition 10. Consider the vector $\mathbf{x}(k)$. A contour is the plot obtained by connecting $(i, x_i(k))$ to $(i+1, x_{i+1}(k))$ with a straight line for each i , ($i = 1, \dots, N$).

Creating a contour for each k gives a pictorial representation of vectors $\mathbf{x}(k)$ as a function of k . We call this a *contour plot*. Figure 6 displays the contour plots of a size 20 system, where the sequence $\{\mathbf{x}(0), \mathbf{x}(1), \mathbf{x}(2), \dots\}$ represents the contours (counting k from the top). For this reason, we interchangeably refer to the vector $\mathbf{x}(k)$ as “the k^{th} contour” from now on. Between successive contours, we can imagine there being drawn the internal processes of those in Figure 7.

Consider an example system with $\mathbf{x}(0) = \mathbf{u}$ that yields the following periodic behavior: $\mathbf{x}(k+1) = 5 \otimes \mathbf{x}(k)$ for $k \geq 3$. The contours for this system would represent vectors in the periodic regime $\{\mathbf{x}(k) \mid \mathbf{x}(k+1) = 5 \otimes \mathbf{x}(k), \mathbf{x}(0) = \mathbf{u}, k \geq 3\}$. The period of a regime and cyclicity are related by $1 \leq \rho \leq \sigma$. Thus, if $\sigma = 1$ in this example, then $\rho = \sigma = 1$, so that no other period can be obtained for all initial states $\mathbf{x}(0)$. Therefore, each contour in the contour plot has the same shape (separated by five time units) as $k \rightarrow \infty$; we call this a *limiting shape* of the contours or a *limiting contour*. For a larger period, we

obtain a different set of limiting contours. In particular, we obtain different limiting contour plots for the cases $\rho = 1, 2, \dots, \sigma$, each dependent on the choice of $\mathbf{x}(0)$. We will formalize this in Section 4.

The idea of a limiting shape in contours suggests a change of coordinates: given the irreducible matrix $P \in \mathbb{R}_{\max}^{N \times N}$ with eigenvalue λ , let

$$\mathbf{x}(k) = \lambda^{\otimes k} \otimes \mathbf{y}(k). \tag{20}$$

We can think of $\lambda^{\otimes k}$ as a diagonal matrix or the product of $\lambda^{\otimes k}$ and the identity matrix $E(N, N)$. The advantage of this is that such a diagonal matrix is invertible, its inverse being the diagonal matrix with diagonal entries equal to $\lambda^{\otimes -k}$. Using this property, we rearrange equation (20) to obtain

$$\mathbf{y}(k) = \lambda^{\otimes -k} \otimes \mathbf{x}(k). \tag{21}$$

In other words, $\mathbf{y}(k)$ is the limit to which the vectors $\mathbf{x}(k) - \lambda k$ tend as $k \rightarrow \infty$. By studying the asymptotic behavior of $\mathbf{y}(k)$ itself, we can deduce the shape of the limiting contour.

The original system follows the recurrence relation $\mathbf{x}(k + 1) = P \otimes \mathbf{x}(k)$ for some $\mathbf{x}(0) \in \mathbb{R}_{\max}^N$. Substitute equation (20) into this to obtain

$$\lambda^{\otimes(k+1)} \otimes \mathbf{y}(k + 1) = P \otimes \lambda^{\otimes k} \otimes \mathbf{y}(k). \tag{22}$$

Interpreting $\lambda^{\otimes k}$ as a diagonal matrix again yields

$$\begin{aligned} \mathbf{y}(k + 1) &= \lambda^{\otimes -(k+1)} \otimes P \otimes \lambda^{\otimes k} \otimes \mathbf{y}(k) = \\ &\lambda^{\otimes -1} \otimes P \otimes \mathbf{y}(k) = \hat{P} \otimes \mathbf{y}(k), \end{aligned} \tag{23}$$

where $\hat{P} = \lambda^{\otimes -1} \otimes P$ represents the *normalized matrix* of P , equivalently obtained by subtracting the eigenvalue of P from each of its entries. The communication graph of \hat{P} is the same as that for P (but with different arc weights), so that \hat{P} is also irreducible. However, the maximum average circuit weight of $\mathcal{G}(\hat{P})$, hence the eigenvalue of \hat{P} , is zero.

Moreover, it can be shown that P and \hat{P} have the same cyclicity (i.e., $\sigma(\hat{P}) = \sigma(P) = \sigma$). Theorem 1 tells us of the asymptotic behavior of the powers of an irreducible matrix. Apply this to \hat{P} to obtain $\hat{P}^{\otimes(k+\sigma)} = 0^{\otimes \sigma} \otimes \hat{P}^{\otimes k} = \hat{P}^{\otimes k}$ for $k \geq k_*$. Thus, using equation (23),

$$\mathbf{y}(k + \sigma) = \hat{P}^{\otimes(k+\sigma)} \otimes \mathbf{y}(0) = \hat{P}^{\otimes k} \otimes \mathbf{y}(0) = \mathbf{y}(k). \tag{24}$$

So the limiting contour $\mathbf{y}(k)$ is periodic with period $\sigma(P)$. In fact, like $\mathbf{x}(k)$, $\mathbf{y}(k)$ is periodic with period ρ , dependent on $\mathbf{x}(0)$, where ρ is a factor of σ . Note that this period now conforms with the traditional dynamical systems definition of a period, in that the sequence $\{\mathbf{y}(k) \mid k \in \mathbb{N}_0\}$ is not monotonically increasing, in contrast to the original sequence $\{\mathbf{x}(k) \mid k \in \mathbb{N}_0\}$.

Loosely speaking, Section 3.3 has shown that there is no unique shape to the limit of a contour plot. This yields an interesting feature of the max-plus asynchronous model: the asynchrony is related not only to the timing dependency graph but also to the system's starting point in time.

■ 3.4 The Eigenspace in Max-Plus Algebra

The set of all eigenvectors of $A \in \mathbb{R}_{\max}^{n \times n}$ associated to eigenvalue λ is the *eigenspace* of A . For max-plus asynchrony, the importance of the eigenspace is demonstrated through its links to the contour plot, as shown next.

Consider the recurrence relation $\mathbf{x}(k+1) = P \otimes \mathbf{x}(k)$ and its corresponding contour plot. We have already established that the vectors in a regime are eigenvectors of $P^{\otimes \sigma(P)}$. In other words, each contour is an eigenvector of $P^{\otimes \sigma(P)}$. Taking linear combinations of eigenvectors enables the construction of the eigenspace of $P^{\otimes \sigma(P)}$ (see Appendix C). This eigenspace is the set of all possible periodic regimes (i.e., of all periods ρ , $1 \leq \rho \leq \sigma(P)$, obtained for all initial states $\mathbf{x}(0)$), which corresponds to the set of all contour plots that can be obtained.

Theorem 4 (Appendix C) gives a method for constructing the eigenspace of an irreducible matrix. Let $B = P^{\otimes \sigma(P)}$. Applying Theorem 4 to irreducible B will yield its eigenspace and consequently all possible contour plots for the system $\mathbf{x}(k+1) = P \otimes \mathbf{x}(k)$.

Note 2. If P is irreducible, then it is not necessarily the case that $B = P^{\otimes \sigma(P)}$ will also be irreducible. Nevertheless, in a highly connected lattice such as the regular network of this paper, it is more likely that B is irreducible. If B is reducible, then methods other than Theorem 4 must be employed. One such method is Howard's policy improvement scheme (see [18] or [22]).

Lemma 1. Consider a max-plus system having irreducible timing dependency matrix P with cyclicity σ . For ρ fixed ($1 \leq \rho \leq \sigma$), period ρ regimes are not necessarily unique.

Proof. Each contour in the periodic regime is an eigenvector of $P^{\otimes \rho}$. By taking linear combinations of eigenvectors, it is possible to construct linearly independent eigenvectors of $P^{\otimes \rho}$ such that the corre-

sponding contours are also linearly independent. Thus, there is no unique periodic regime of period ρ . \square

The top of Figure 6(a) shows the limiting contours of a size 20 system. The contours depict a period 1 regime and therefore have the same shape. Lemma 1 says that if $\sigma > 1$ for that system, then other shapes of limiting contour may be possible for $\rho = 1$ periodic regimes. In other words, the limiting contour plot is not necessarily fixed, despite ρ, χ 's being fixed. The significance of this is that a corresponding CA is asymptotically not unique, but dependent on the initial time $\mathbf{x}(0)$; this can affect the time $x_i(k)$ of the k^{th} update at node i relative to $x_j(k)$ (at node j), even when the “pattern” of consecutive update times $(\mathbf{x}(k), \mathbf{x}(k+1), \dots, \mathbf{x}(k+\rho), \dots)$ is independent of $\mathbf{x}(0)$ (since the cycle time is independent of $\mathbf{x}(0)$).

4. Cellular Automata in Max-Plus Time

We now present the first formalism for implementing a CA asynchronously, such that update times are determined by a max-plus system.

4.1 Contour Plot as a Foundation for Cellular Automata

Let $s_i(k)$ denote the CA state of node i at epoch k . We are concerned with Boolean CA states, so that $s_i(k) \in \{0, 1\}$. The unit k is as used in the max-plus model, which updates the times $\mathbf{x}(k)$. Thus, to be precise, $s_i(k)$ is the CA state of node i at time $x_i(k)$. The CA state of the system is represented by the string $s_1(k) s_2(k) \dots s_N(k)$, which can also be read as the state of all nodes on contour k . As a consequence, just as we represented the vector $\mathbf{x}(k)$ by a contour, we can represent the CA state $\mathbf{s}(k)$ by the same contour but with the addition that the coordinates $(i, x_i(k))$ now display the state $s_i(k)$ (e.g., in colored form, where two different colors are used to distinguish the two states 1 and 0).

Recall the main events that are internal to node i ; they occur within cycle k . These events are grouped in two categories: “receive” and “send.” Once the receive CA states have all arrived, node i applies a CA rule on this set, to obtain the new state $s_i(k)$. If all nodes have neighborhood size n , the applied CA rule is the function $f: \{0, 1\}^n \rightarrow \{0, 1\}$, and the new state $s_i(k)$ is calculated as

$$s_i(k) = f(s_{N_i}(k-1)). \quad (25)$$

4.2 Cellular Automaton Space-Time Plot

The classical one-dimensional CA is synchronous, so that the k^{th} update time of each cell is the same. Consequently, we can think of such a system as having a contour plot that contains only horizontal contours. Updates of the CA state of the system take place every one time unit, thereby giving the synchronous CA a cycle time of 1. The time between contours in this system is thus of duration one, although no such duration is depicted; for example, if $s_i(k) = 1$ for all k , then this is shown as a continuous vertical colored block in position i .

Despite varying contour shapes dictating the varying time gaps between contours, it is simple and intuitive to construct the space-time plot for CA in max-plus time. To illustrate, we use the example of a regular 3-nbhd network with size $N = 10$. Let the positive (diagonal) entries in matrix A_ξ be represented by the vector ξ of processing times. We choose the entries in ξ at random with equal probability from all integers between 1 and 30, while the nonzero entries in T are selected likewise from the integers between 1 and 10. Taking the initial time $\mathbf{x}(0) = \mathbf{u}$, we obtain a contour plot of update times by iterating the max-plus system. We now address the CA state by assigning the depicted space between contours as *memory*: for each node, the CA state remains fixed until the time of update, which corresponds to a contour. For node i , the time t_i that elapses between contours implies that the storing of the CA state in memory can be represented as a vertical block of length t_i (which is colored accordingly, depending on the CA state). Correspondingly, this may be depicted in a space-time plot, the construction of which is shown in three stages in Figure 8.

At this juncture, it is important to distinguish between the variables t and k . The term “time” (or “real time”) now refers to a point $t \in \mathbb{R}^+$; it can be thought of as time as we know it. k maintains its role as a discrete epoch. Thus, node i carries a CA state for every point in real time. If we denote the state of node i at real time t as $s_i^{(t)}$, the CA state of node i can now be understood in two ways: $s_i(k)$ denotes the CA state on contour k , and it is discretely dynamic, while $s_i^{(t)}$ represents the state in a dynamical system with a continuous underlying real time t . Thus, for example, on contour $k = 2$, if $t = 5.8$, we have $s_i(2) = s_i^{(5.8)}$.

Figure 8(d) is for comparison with Figure 8(c), and it shows the classical synchronous CA having the same initial time $\mathbf{x}(0)$, initial CA state $\mathbf{s}(0)$, and CA rule. While sharing initial conditions and CA rule, the states $s^{(t)}$ in both patterns will generally differ. (This can be seen by simply drawing a horizontal line across both patterns at time t and

reading off the state of each node at that time.) The difference in pattern is clearly due to the asynchrony of the max-plus system in Figure 8(c). We show next that we can characterize this difference somewhat and, in fact, map the synchronous CA to the max-plus CA via the contour plot.




Figure 8. Construction of the CA space-time plot in max-plus time. The CA rule is ECA rule 150. The initial CA state on contour 0 is $s_5(0) = 1$, $s_i(0) = 0$ for $i \neq 5$. State 0 is colored dark; state 1 is light. In all figures, the vertical axis denotes real time, traveling down. (a) Contour plot with CA states indicated on each contour. (b) CA states indicated for all time by filling spaces between contours with memory. (c) Contours and space between nodes removed to obtain the CA space-time plot. (d) Classical (synchronous) CA space-time plot.

4.3 Bijection

Recall the state transition graph. In Figure 8(d), the CA period is 6 and the periodic orbit is the following set:

$$\{0000100000, 0001110000, 0010101000, 0110101100, 1000100010, 1101110110\}. \tag{26}$$

If we consider the CA states on only the contours in Figure 8(c) (which is seen better in Figure 8(a)), we see that they are exactly the same as the states in Figure 8(d). This is a consequence of the max-

plus model's requiring all neighborhood states to arrive before processing new CA states. Indeed, this notion has also been mentioned (albeit briefly) in [25, p. 1035] under the heading of "Intrinsic synchronization in cellular automata."

Given the same initial CA state $s(0)$ and CA rule, let \mathcal{S} and \mathcal{M} denote the orbit of $s(0)$ generated in the synchronous system and the max-plus system, respectively. Let $s_{\mathcal{S}}(k)$ denote the CA state after k iterations of the synchronous system; $s_{\mathcal{M}}(k)$ denotes the CA state after k iterations of the max-plus system. The model uses the same CA rule, applied to the same neighborhoods, the only difference being that the time of application of the rule is different. Then, after k iterations of both systems, we clearly have $s_{\mathcal{S}}(k) = s_{\mathcal{M}}(k)$. This defines a one-to-one and onto mapping—a bijection—between \mathcal{S} and \mathcal{M} , and we say that both systems have the same STG (defined as the transitions between states on contours). Thus, the max-plus system need not evolve the CA concurrently, since the CA plot for the max-plus system may be obtained from this mapping.

In summary, the STG provides a deterministic form for predicting the behavior of the CA in max-plus time. Each state $s_{\mathcal{M}}(k)$ in the STG does not necessarily correspond to a state in real time $t \in \mathbb{R}$, due to contours not necessarily being horizontal. Nevertheless, the property of memory can be applied to ascertain such real-time states.

5. Conclusion

We have shown that when modeling discrete asynchronous systems, more attention needs to be paid to the internal processes of a cell. This has resulted in the uncovering of a useful theory—that of max-plus algebra.

Cellular automata (CA) are naturally well-suited to be modeled in max-plus time, because this model requires update on knowledge of all neighbors. This has the additional benefit of cells updating only when they are ready. Thus, whereas a fixed, periodic, global update time (as in the synchronous case) can be slower and less energy efficient, a max-plus asynchronous model consumes only the time and energy that local neighborhoods require.

The classical ECA corresponds to a strongly connected network in our model. Thus, P is irreducible, and this ensures periodic behavior, which is not what is usually associated with the word "asynchrony." We comment that periodic behavior is expected also when P is reducible; in this case, the theory is similar to what we have covered here, but the resultant cycle time vector is not necessarily uniform

[18]. We have therefore seen that not only is it an efficient system for the timing of asynchronous CA, the max-plus system is also a simple, deterministic, asynchronous model.

We can get a visual sense of such periodicity via the contour plot, which couples the asynchronous update times with the CA. The eigenspace of $P^{\otimes \sigma(P)}$ yields the range of contour plots that can be generated due to P ; this enables us to get a sense of the range of the corresponding CA space-time patterns that are possible. By identifying a bijection between synchronous and max-plus CA, we can further narrow the aforementioned range of CA patterns, since CA states on contours can be obtained directly from the synchronous CA. We note that this notion of memory has been suggested previously (see [14, 15, 25] and the references therein); this bijection is captured under the topic of “causal invariance” in [25], wherein the same causal network (i.e., structure of events—past and present) emerges, irrespective of CA states. In this paper, however, we have identified a numerical link between the idea and the theory in the form of max-plus algebra.

Importantly, the contour plot also allows one to see that the system visits many more “interim” states (in memory) in real time—these are transient CA states that illustrate the local dynamics. Thus, although the STG is the same as the synchronous STG, max-plus CA gives more information: we can now visualize exactly how one state in the STG evolves into another. Max-plus algebra offers promising scope for assigning numerical measures to the states in memory; the parameters ξ_i , τ_{ij} , along with the cycle time, are likely to play major roles for this purpose. Such classification scales up to the CA space-time pattern as a whole, and attempts to do this are the focus of further work.

Synchronous CA have previously been held up as models for patterns seen in nature (such as those seen on seashells and the growth of snowflakes). Certain probabilistic CA improved on this to account for random fluctuations in the growth processes. We hypothesize that a max-plus algebraic approach adds further realism because it considers the actual processing and delay times that may be present within the chemical reactions and does not rely on probability. One might conjecture the aforementioned transient states in memory to be precisely the processes observed during the natural construction of such patterns as found on seashells and snowflakes; that is, apparent deviations from “normal” growth may, in fact, be part of a transient phase as opposed to some random fault. Noting the absence of significant literature on the matter, we believe Boolean networks provide a notable avenue to exploit; it would be interesting to see how the various topologies of these networks behave under max-plus time. How does a max-plus model impact applications such as Kauffman’s genetic regulatory network [17]? Moreover, in light of connections to the idea of

causal invariance in [25], might max-plus algebra now provide a numerical measure for the related topics (e.g., of space, time, and relativity) contained therein?

In addition to the above work, this paper has also laid the theoretical groundwork for extending the asynchronous model itself. For instance, one such extension includes the minimum operator and can subsequently be used to conduct further studies of asynchronous CA, particularly to better describe an intended application. Preliminary results of these new “max-min-plus” models are less predictable, particularly manifested by an absence of the bijection described in Section 4.3 [23].

Acknowledgments

The novel idea in this paper was conceived by D. Broomhead, but ill health—and his subsequent death in July, 2014—prevented involvement in its preparation.

This work was initiated as part of the Centre for Interdisciplinary Computational and Dynamical Analysis (CICADA) project, which was funded by the Engineering and Physical Sciences Research Council (EPSRC). E. L. Patel is now funded by the EPSRC Horizon: Digital Economy project (EP/G065802/1).

Appendix

A. Maximal Strongly Connected Subgraph

We say that node i communicates with node j , denoted $i \leftrightarrow j$, if either $i = j$ or $i \rightarrow j$ and $j \rightarrow i$. Note that we allow a solitary node to communicate with itself, even if there is no self-loop attached to it.

It is thus possible to partition the node set V of a graph into disjoint subsets V_i such that $V = V_1 \cup V_2 \cup \dots \cup V_q$, where each subset V_i contains nodes that communicate with each other but not with other nodes of V . By taking V_i together with arc set E_i , each of whose arcs has start node and end node in V_i , we obtain the subgraph $\mathcal{G}_i = (V_i, E_i)$. We call this subgraph a *maximal strongly connected subgraph* (MSCS) of $\mathcal{G} = (V, E)$.

B. Asymptotic Behavior of the Max-Plus System

Theorem 2. Let $A \in \mathbb{R}_{\max}^{n \times n}$ be irreducible. Then A possesses a unique eigenvalue, denoted $\lambda(A)$, which is finite ($\neq \varepsilon$). Moreover, this eigen-

value is equal to the maximal average weight of elementary circuits in $\mathcal{G}(A)$. Let c denote an elementary circuit of $\mathcal{G}(A)$. Denote the set of all elementary circuits of $\mathcal{G}(A)$ by $C(A)$. Then

$$\lambda(A) = \max_{c \in C(A)} \frac{|c|_{w}}{|c|_l} \tag{B.1}$$

Proof. See [18, Theorem 2.9]. \square

The following theorem shows that the cycle time vector

$$\chi = \left(\lim_{k \rightarrow \infty} \frac{x_1(k)}{k}, \lim_{k \rightarrow \infty} \frac{x_2(k)}{k}, \dots, \lim_{k \rightarrow \infty} \frac{x_N(k)}{k} \right)^T \tag{27}$$

is unique.

Theorem 3. Consider the recurrence relation $\mathbf{x}(k+1) = A \otimes \mathbf{x}(k)$ for $k \geq 0$ and $A \in \mathbb{R}_{\max}^{n \times n}$ irreducible. For some $\mathbf{x}_*(0) \in \mathbb{R}_{\max}^n$ whose elements are all finite, if the limit

$$\lim_{k \rightarrow \infty} \frac{A^k \otimes \mathbf{x}_*(0)}{k}$$

exists, then this limit is the same for any initial condition $\mathbf{x}(0) \in \mathbb{R}_{\max}^n$ whose elements are all finite.

Proof. See [18, Theorem 3.11]. \square

In fact, the condition of irreducibility can be relaxed to that of reducibility, as long as each node has at least one predecessor node in $\mathcal{G}(A)$; this corresponds to all rows of A containing at least one nonzero ($\neq -\infty$) element. Thus, in the latter (relaxed) case, we obtain a cycle time vector whose elements may not necessarily be identical. However, for A irreducible, each element of χ turns out to be the same—specifically the eigenvalue of A —as stated in the following lemma.

Lemma 2. For the recurrence relation $\mathbf{x}(k+1) = A \otimes \mathbf{x}(k)$ with $k \geq 0$, let $A \in \mathbb{R}_{\max}^{n \times n}$ be an irreducible matrix having eigenvalue $\lambda \in \mathbb{R}$. Then, for $i = 1, 2, \dots, n$,

$$\lim_{k \rightarrow \infty} \frac{x_i(k)}{k} = \lambda$$

for any initial condition $\mathbf{x}(0) \in \mathbb{R}^n$.

Proof. See [18, Lemma 3.12]. \square

C. Eigenspace of an Irreducible Matrix

As in conventional linear algebra, eigenvectors are not unique in max-plus algebra because they are defined up to scalar multiplication. (It can easily be shown that if v and w are eigenvectors of $A \in \mathbb{R}_{\max}^{n \times n}$ associated with eigenvalue λ , then for $\alpha, \beta \in \mathbb{R}_{\max}$, $\alpha \otimes v \oplus \beta \otimes w$ is also an eigenvector of A .)

Consider the definition of the Kleene star for any $A \in \mathbb{R}_{\max}^{n \times n}$:

$$A^* \stackrel{\text{def}}{=} \bigoplus_{k=0}^{\infty} A^{\otimes k}. \tag{C.1}$$

It is known that if circuit weights in $\mathcal{G}(A)$ are nonpositive, then the Kleene star of a square matrix over \mathbb{R}_{\max} exists [18]. Denote the critical graph of A as $\mathcal{G}^{cr}(A) = (V^{cr}(A), E^{cr}(A))$, and the normalized matrix $\hat{A} = -\lambda \otimes A$.

Theorem 4. Let $A \in \mathbb{R}_{\max}^{n \times n}$ be irreducible and consider \hat{A}^* to be the Kleene star of $\hat{A} = -\lambda \otimes A$.

1. If node i belongs to $\mathcal{G}^{cr}(A)$, then $[\hat{A}^*]_{,i}$ is an eigenvector of A .
2. The eigenspace of A is

$$V(A) = \left\{ v \in \mathbb{R}_{\max}^n \mid v = \bigoplus_{i \in V^{cr}(A)} a_i \otimes [\hat{A}^*]_{,i} \text{ for } a_i \in \mathbb{R}_{\max} \right\}. \tag{C.2}$$

3. For i, j belonging to $\mathcal{G}^{cr}(A)$, there exists $a \in \mathbb{R}$ such that

$$a \otimes [\hat{A}^*]_{,i} = [\hat{A}^*]_{,j} \tag{C.3}$$

if and only if i and j belong to the same MSCS of $\mathcal{G}^{cr}(A)$.

Proof. In [18, Theorem 4.5]. \square

D. Hasse Diagram

With elements taken from [24], we define the Hasse diagram formally. We first require the ‘‘happened before’’ relation, denoted by $<$.

Definition 11. The relation $<$ on a set of events is defined by the following conditions:

1. If the events a and b are processed by the same processor, and a occurs before b , then $a < b$.
2. If a is the sending of a message by processor A , and b is the receipt of the message by another processor B , then $a < b$.
3. If $a < b$ and $b < c$, then $a < c$.

We say that two distinct events a and b are concurrent if $a \not\prec b$ and $b \not\prec a$. We also assume the properties of irreflexivity, that is, $a \not\prec a$, and antisymmetry on the times of events; that is, if a_t and b_t represent the times of events a and b , then $a_t \not\prec b_t, b_t \not\prec a_t \Rightarrow a_t = b_t$.

In fact, Definition 11, along with the properties of irreflexivity and antisymmetry, defines a *partial ordering* on the set X of all events in our system. We say that the set X , along with the relation \prec , forms a partially ordered set (*poset*). In the following, the pair (X, \prec) denotes a poset.

Definition 12. Let x and y be distinct elements of a poset (X, \prec) . y is said to cover x if $x \prec y$, but no element z satisfies $x \prec z \prec y$.

Definition 13. Let the set X_1 of n elements $\{x_1, x_2, \dots, x_n\}$ be a subset of (X, \prec) such that each element may be totally ordered according to \prec as $x_1 \prec x_2 \prec \dots \prec x_n$. Then X_1 is a chain. The subset $X_2 \in X$ is called an antichain if and only if no elements of X_2 may be totally or partially ordered.

These definitions now allow us to formulate the following definition.

Definition 14. The Hasse diagram of a poset (X, \prec) is a graph drawn in the Euclidean plane such that each element of the poset is represented by a unique vertex in the graph. Each covering pair $x \prec y$ is depicted by a directed arc from x to y , where the point representing x is below the point representing y (i.e., it has a smaller Y coordinate).

References

- [1] S. Wolfram, "Statistical Mechanics of Cellular Automata," *Reviews of Modern Physics*, 55(3), 1983 pp. 601–644. doi:10.1103/RevModPhys.55.601.
- [2] S. Wolfram, "Universality and Complexity in Cellular Automata," *Physica D: Nonlinear Phenomena*, 10(1–2), 1984 pp. 1–35. doi:10.1016/0167-2789(84)90245-8.
- [3] H. Meinhardt, P. Prusinkiewicz, and D. R. Fowler, *The Algorithmic Beauty of Sea Shells*, New York: Springer, 1995.
- [4] "Snowflake Generator." (Jun 26, 2014) www.incendia.net/wiki/index.php/Snowflake_Generator.
- [5] Y. Gunji, "Pigment Color Patterns of Molluscs as an Autonomous Process Generated by Asynchronous Automata," *BioSystems*, 23(4), 1990 pp. 317–334. doi:10.1016/0303-2647(90)90014-R.

- [6] T. E. Ingerson and R. L. Buvel, "Structure in Asynchronous Cellular Automata," *Physica D: Nonlinear Phenomena*, 10(1–2), 1984 pp. 59–68. doi:10.1016/0167-2789(84)90249-5.
- [7] H. Bersini and V. Detours, "Asynchrony Induces Stability in Cellular Automata Based Models," in *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems* (R. A. Brooks and P. Maes, eds.), Cambridge, MA: MIT Press, 1994 pp. 382–387.
- [8] N. Fatès, E. Thierry, M. Morvan, and N. Schabanel, "Fully Asynchronous Behavior of Double-Quiescent Elementary Cellular Automata," *Theoretical Computer Science*, 362(1–3), 2006 pp. 1–16. doi:10.1016/j.tcs.2006.05.036.
- [9] N. Fatès and M. Morvan, "An Experimental Study of Robustness to Asynchronism for Elementary Cellular Automata," *Complex Systems*, 16(1), 2005 pp. 1–28. <http://www.complex-systems.com/pdf/16-1-1.pdf>.
- [10] G. Le Caër, "Comparison between Simultaneous and Sequential Updating in $2^{n+1} - 1$ Cellular Automata," *Physica A: Statistical Mechanics and its Applications*, 157(2), 1989 pp. 669–687. doi:10.1016/0378-4371(89)90061-7.
- [11] B. Schönfisch and A. de Roos, "Synchronous and Asynchronous Updating in Cellular Automata," *BioSystems*, 51(3), 1999 pp. 123–143. doi:10.1016/S0303-2647(99)00025-8.
- [12] D. G. Green, D. Newth, D. Cornforth, and M. Kirley, "On Evolutionary Processes in Natural and Artificial Systems," in *Proceedings of the Fifth Australia-Japan Joint Workshop on Intelligent and Evolutionary Systems* (P. Whigham, ed.), The University of Otago, Dunedin, New Zealand, 2001 pp. 1–10.
- [13] D. Cornforth, D. Green, D. Newth, and M. Kirley, "Ordered Asynchronous Processes in Natural and Artificial Systems," in *Proceedings of the Fifth Australia-Japan Joint Workshop on Intelligent and Evolutionary Systems* (P. Whigham, ed.), The University of Otago, Dunedin, New Zealand, 2001 pp. 105–112.
- [14] N. Fatès, "A Guided Tour of Asynchronous Cellular Automata," in *Cellular Automata and Discrete Complex Systems*, New York: Springer, 2013 pp. 15–30.
- [15] C. L. Nehaniv, "Asynchronous Automata Networks Can Emulate Any Synchronous Automata Network," *International Journal of Algebra and Computation*, 14(5–6), 2004 pp. 719–739. doi:10.1142/S0218196704002043.
- [16] C. Gershenson, "Classification of Random Boolean Networks," *ICAL 2003 Proceedings of the Eighth International Conference on Artificial Life* (R. Standish, M. A. Bedau, and H. A. Abbass, eds.), Cambridge, MA: MIT Press, 2003 pp. 1–8.

- [17] S. A. Kauffman, "Metabolic Stability and Epigenesis in Randomly Constructed Genetic Nets," *Journal of Theoretical Biology*, 22(3), 1969 pp. 437–467. doi:10.1016/0022-5193(69)90015-0.
- [18] B. Heidergott, G. J. Olsder, and J. W. van der Woude, *Max Plus at Work: Modeling and Analysis of Synchronized Systems: A Course on Max-Plus Algebra and Its Applications*, Princeton: Princeton University Press, 2006.
- [19] A. Doustmohammadi and E. W. Kamen, "Direct Generation of Event-Timing Equations for Generalized Flow Shop Systems," in *Proceedings of SPIE 2596, Modeling, Simulation, and Control Technologies for Manufacturing*, 1995 pp. 50–62. doi:10.1117/12.227228.
- [20] C. A. Brackley, D. S. Broomhead, M. C. Romano, and M. Thiel, "A Max-Plus Model of Ribosome Dynamics During mRNA Translation," *Journal of Theoretical Biology*, 303, 2012 pp. 128–140. doi:10.1016/j.jtbi.2012.03.007.
- [21] D. Broomhead, S. Furber, and M. Johnson, "Algebraic Approach to Time Borrowing," *IET Computers & Digital Techniques*, 7(1), 2013 pp. 1–10. doi:10.1049/iet-cdt .2012.0103.
- [22] J. Cochet-Terrasson, G. Cohen, S. Gaubert, M. M. McGettrick, and J.-P. Quadrat, "Numerical Computation of Spectral Elements in Max-Plus Algebra," in *System Structure and Control 1998 (SSC'98): A Proceedings Volume from the 5th IFAC Conference*, Nantes, France, July 8–10, 1998 (J. F. LaFay, ed.), New York: Oxford, 1998.
- [23] E. L. Patel, "Maxmin-Plus Models of Asynchronous Computation," Ph.D. thesis, University of Manchester, 2012.
- [24] P. J. Cameron, *Combinatorics: Topics, Techniques, Algorithms*, New York: Cambridge University Press, 1994.
- [25] S. Wolfram, *A New Kind of Science*, Champaign, IL: Wolfram Media, Inc., 2002.