Complex Systems **6** (1992) 287–300

Effect of Noise on Long-term Memory in Cellular Automata with Asynchronous Delays between the Processors

Reza Gharavi
Venkat Anantharam
School of Electrical Engineering, Cornell University, Ithaca, NY 14853, USA

Abstract. We consider monotonic binary cellular automata on the lattice $\mathbf{Z}^p$ as models of systems capable of long-term memory, that is, capable of admitting multiple invariant configurations. Long-term memory in cellular automata can be robust in the presence of noise, so that in a noisy environment the automaton may admit more than one stationary distribution on configurations. We examine the effects of asynchronous communication delays between the processors on long-term memory in cellular automata and describe when asynchronism can cause the erasure of long-term memory in the presence of noise. Our main result is a simple generalization to the asynchronous computation model of a deep result of Toom characterizing the invariant configurations of monotonic binary cellular automata that are robust to noise. Several qualitative consequences of asynchronism are illustrated through examples.

1. Introduction

Cellular automata are simple computational models that are capable of exhibiting a wide range of complex dynamical behavior (see [7]). The computation is considered as proceeding synchronously via identical processors at each site on a regular lattice—usually $\mathbf{Z}^p$—and the computational rule is assumed to be spatially homogenous. In other words, at every time step the state at a site is updated as a function of certain previous states of some of the neighboring sites with a spatially homogenous updating rule. The interest in studying such automata comes from several points of view. For example, there is a belief that the complex dynamical behavior exhibited by these automata is a good model for the natural statistical behavior of physical systems such as gases, which consist of large numbers of interacting elements. Another powerful source of reawakened interest in cellular automata has been the development of parallel computational systems that employ different types of regular architectures (e.g., [1]).

A processor at a site can be in one of a finite set of states, and the entire collection of states at the sites is called a *configuration.* One of the remarkable properties of certain cellular automaton updating rules is that they can admit more than one invariant configuration, representing the ability to maintain long-term memory. Further, it is known [6] that there are automata whose long-term memory persists under noise. Even in environments in which the state of a site may change due to noise or computational error, the evolution of the automaton is such that one can give a test that distinguishes between initial conditions even after arbitrarily long periods of time. This ability is particularly important from the point of view of the automaton as a computational device, where the initial configuration is the input on which the processors perform their calculations (see [2]). For the operation to be reliable in a noisy or unreliable environment, we would require the system to remember enough relevant information about its initial configuration over arbitrarily long periods of time. For a general discussion of the subject of reliable computation, see the survey paper [5].

Our interest in this paper is in the persistence of long-term memory in automata operating in an unreliable or noisy environment when there is also unreliability in the data transfer between the processors. We examine this question in an asynchronous computational model, where the computations are synchronous but there are unknown delays in the transmission of data between processors. In other words, a processor carrying out a computation may only have available to it delayed versions of the data on which it depends. Such a situation may arise when data is lost or delayed in transit between the processors. We assume that the delays are bounded by some integer d; in other words, the systems we consider are partially asynchronous (see [1]). We are particularly interested in discussing how large a delay in data transfer can be tolerated before the ability of the automaton to remember its initial configuration breaks down.

Our results are for a class of automata called monotonic binary tessellations (MBTs), defined in section 2. In an MBT the state of each processor is 1 or 0, and the all-zero and the all-one configurations are invariant under the updating rule. Our main result is a simple generalization to the partially asynchronous computational model of a deep result of Toom [6], which gives necessary and sufficient conditions for the stability of the all-zero configuration of an MBT to noise. Toom's characterization is widely acknowledged to be one of the deepest results to date about cellular automata—it is a simple geometric characterization based on the geometry of the data set that a processor depends upon for its computation. We discuss this in section 2 after giving a formal definition of MBTs. Our result is stated and proved as Theorem 2 in section 3. Our characterization of the stability of the states under asynchronous computation allows us to discuss the effects of asynchronism on the stability of states. Several qualitative aspects of this are brought out through examples in section 4.

2. Problem definition

Before defining an MBT, we will need the following items:

1. A function $U : \mathbf{Z}^{p+1} \mapsto \mathbf{Z}^{(p+1)r}$ defined by

$$U(v) \stackrel{\text{def}}{=} (v + (u_1, t_1), \ldots, v + (u_r, t_r))$$

   where r and p are positive integers and $u_i \in \mathbf{Z}^p$, $t_i < 0$ for $i = 1, \ldots, r$. We denote $t_W = \max_{i=1,\ldots,r} |t_i|$.

2. A set $V = \{(s,t) \in \mathbf{Z}^{p+1} : s \in \mathbf{Z}^p,\ t \geq -t_W\}$. Each *point* $v \in V$ is denoted by the pair (s,t), where $s \in \mathbf{Z}^p$ is the *space* coordinate and $t \in \mathbf{Z}$ is the *time* coordinate, and we say v is *located* at s at time t. In this section, all pairs (s,t) will be assumed to be of this form. We also define $W \stackrel{\text{def}}{=} \{(s,t) \in V : -t_W \leq t < 0\}$.

3. A *space-time configuration* $x \in \{0,1\}^V$. The state of a point $v \in V$ is denoted by x_v and for any $A \subset V$, by $x_A \in \{0,1\}^A$ we will mean a tuple, components of which are x_a, $a \in A$. We term a configuration x_W of W a *base*.

4. A *monotonic* function $\phi : \{0,1\}^r \mapsto \{0,1\}$, which means that if for some $z, z' \in \{0,1\}^r$ we have $z_i \leq z'_i$ for $i = 1, \ldots, r$, then $\phi(z) \leq \phi(z')$. To avoid trivialities, we will assume that $\phi(0, \ldots, 0) = 0$ and $\phi(1, \ldots, 1) = 1$. By $\phi(y_{U(v)})$ we will mean $\phi(y_{v+(u_1,t_1)}, \ldots, y_{v+(u_r,t_r)})$.

We call any pair (U, ϕ) satisfying the above conditions a (p-*dimensional*) monotonic binary tessellation (MBT). Note that specifying U uniquely determines V and W. For an example of an MBT, see figures 1(a) and 1(b). All of the following definitions are in relation to some fixed (U, ϕ).

Definition. Given any base x_W, we say a configuration y is a *trajectory* if $y_W = x_W$, and if for all $v \in V \setminus W$ we have $y_v = \phi(y_{U(v)})$. We denote such a configuration y by $T(x_W)$. By our assumptions on ϕ, it is clear that a configuration x with $x_v = 0$ or 1 for all $v \in V$ is a trajectory. We term these two special trajectories the *zero* and the *one* trajectories, respectively.

Definition. A trajectory y is termed *attractive* if the following holds for any trajectory x: if the set $\{a \in W : x_a \neq y_a\}$ is finite, then the set $\{a \in V : x_a \neq y_a\}$ is finite. Defining operator I by $I(x_A) \stackrel{\text{def}}{=} \{v \in A : x_v = 1\}$ for $A \subset V$, we see that the zero trajectory is attractive if for any trajectory x such that $|I(x_W)| < \infty$, we have $|I(T(x_W))| < \infty$.

Definition. Given $\epsilon \in (0,1)$, let M_ϵ be a set of probability measures on the σ-algebra generated by the cylinder subsets of $\{0,1\}^V$, such that a measure $\mu \in M_\epsilon$ if and only if for any finite $A \subset V$,

$$\mu(x_v \neq \phi(x_{U(v)})\, \forall\, v \in A) \leq \epsilon^{|A|}.$$

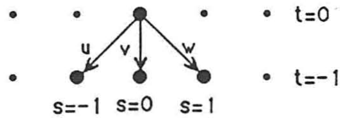

(a)

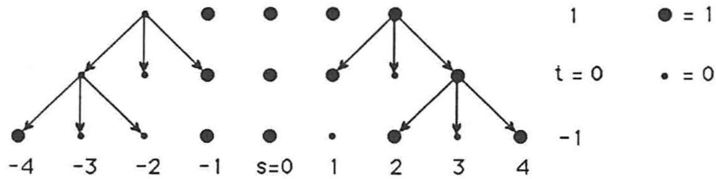

(b)

Figure 1: An example of a one-dimensional MBT, where $(u_1, t_1) = (-1,-1)$, $(u_2, t_2) = (0,-1)$, and $(u_3, t_3) = (1,-1)$ are represented respectively by vectors $\mathbf{u}$, $\mathbf{v}$, and $\mathbf{w}$ in part (a), and where ϕ is the majority rule. A section of a trajectory of this MBT is given in part (b). The minimal zero sets of this example are $\{(-1,-1),(0,-1)\}$, $\{(-1,-1),(1,-1)\}$, and $\{(0,-1),(1,-1)\}$.

Given a base x_W, $M_\epsilon(x_W)$ denotes the set of measures in M_ϵ whose projection to W is a δ measure concentrated on x_W. We term a trajectory y *stable* if

$$\lim_{\epsilon\to 0} \sup_{\substack{\mu\in M_\epsilon(y_W)\\ v\in V}} \mu(x_v \neq y_v) = 0. \tag{1}$$

Definition. A set $A \subset \{(u_1, t_1), \ldots, (u_r, t_r)\}$ is called a *zero set* if the condition $(\forall\, a \in A : x_a = 0)$ implies $\phi(x_{U(0)}) = 0$. A zero set is *minimal* if it contains no strictly smaller zero sets. Henceforth, $Z_1, \ldots, Z_Q$ will denote all of the distinct minimal zero sets of the (U, ϕ) under discussion.

The lattice V in $\mathbf{Z}^{p+1}$ will now be immersed in the real space $\mathbf{R}^{p+1}$ with the same coordinates; for example, Z_i's now become collections of points in $\mathbf{R}^{p+1}$. Given an MBT (U, ϕ), we define σ as follows:

$$\sigma = -\bigcap_{q=1}^{Q} \bigcup_{\{\alpha\in\mathbf{R}:\alpha\geq 0\}} \{\alpha v : v \in \mathrm{conv}(Z_q)\},$$

where $\mathrm{conv}(Z_q)$ is the convex hull of the set Z_q in $\mathbf{R}^{p+1}$.

Theorem 1. *(Theorems 6 and 7 in [6].) The following are equivalent:*

- (a) $\sigma = \{0\}$.
- (b) *The zero trajectory is attractive.*
- (c) *The zero trajectory is stable.*

Condition (a) will be called Toom's criterion.

3. Asynchronous MBTs

We now introduce the possibility of delays between the processors. When constructing a trajectory for an MBT, it is of course expected that a point $(s,t) \in V$ has available to it the state of the points located at $s+s_1, \dots, s+s_r$ at times $t+t_1, \dots, t+t_r$. (Recall that t_i's are negative quantities.) When there are communication delays, (s,t) may have available to it the state of the points located at $s+s_1, \dots, s+s_r$, but only at times previous to $t+t_1, \dots, t+t_r$. Loosely speaking, by an "asynchronous MBT" we mean an MBT with such delays. ("Asynchronous MBT" is a misleading term because such a system is not necessarily an MBT. We will later use the term "scenario" instead.) By *depth d* of an asynchronous MBT we will mean the maximum allowable delay. Note that for a given MBT and a given $d > 0$, there are infinitely many such asynchronous MBTs.

We now define this concept formally. Given d, we let S denote the set of functions $\{\tau : V \setminus W \mapsto \{0,1,2,\dots,d\}^r\}$. As was done for an MBT, we first need to define the following items:

1. the sets $V' \stackrel{\text{def}}{=} \{(s,t) \in \mathbf{Z}^{p+1} : -t_W - d \le t\}$ and $W' \stackrel{\text{def}}{=} \{(s,t) \in V' : t < 0\}$;

2. for all $\tau \in S$, a function $U^\tau : V' \setminus W' \mapsto V'$ defined by

$$U^\tau(v) \stackrel{\text{def}}{=} (v + (u_1, t_1 - \tau_1(v)), \dots, v + (u_r, t_r - \tau_r(v)))$$

where $\tau_i(v)$ denotes the ith component of $\tau(v)$.

Given a $\tau \in S$, we say the pair (U^τ, ϕ) is the *scenario* τ of the MBT (U, ϕ), or simply the scenario τ whenever (U, ϕ) is given. If $\tau(v)$ is a constant r-tuple for all $v \in V' \setminus W'$, then τ will be called a *homogenous* scenario. Note that all homogenous scenarios are MBTs. Moreover, if that r-tuple is $(0, \dots, 0)$, then (U^τ, ϕ) is the same as (U, ϕ). Loosely speaking, $\tau(v)$ are the delays at v with which the data relevant to the computation at v is received. In the following discussion it is assumed that some (U, ϕ) is given.

All the definitions related to a configuration $x \in \{0,1\}^{V'}$ of a scenario τ are analogous to those of an MBT. For completeness we will mention some of them. Given a base $x_{W'}$ and a $\tau \in S$, we say a configuration y is a trajectory for scenario τ, or simply a τ *trajectory*, if $y_{W'} = x_{W'}$, and if for all $v \in V' \setminus W'$ we have $y_v = \phi(y_{U^\tau(v)})$. We denote such a configuration by $T^\tau(x_{W'})$. Note that the zero and the one configurations are τ trajectories for all $\tau \in S$. By replacing V with V', W with W', and T with T^τ in the definitions of attractive and M_ϵ of (U, ϕ), we get definitions for *attractive* and M_ϵ^τ for a scenario τ. If in addition we replace M_ϵ with M_ϵ^τ, we get the definition of a *stable* τ trajectory. Finally, we define the operator C by

$$C(A) \stackrel{\text{def}}{=} \{(s, t-i) \in V' : i = 0, \dots, d; (s,t) \in A\}$$

for any $A \subset V$, and we define $\sigma' \subset \mathbf{R}^{(p+1)}$ by

$$\sigma' \stackrel{\text{def}}{=} -\bigcap_{q=1}^{Q} \bigcup_{\alpha \geq 0} \{\alpha v : v \in \text{conv}(C(Z_q))\}.$$

Our main result is a characterization of when the zero trajectory in an asynchronous MBT is stable under all scenarios.

Theorem 2. *The following are equivalent:*

(*a*) $\sigma' = 0$.

(*b*) *The zero* τ *trajectory is attractive for all* $\tau \in S$.

(*c*) *The zero* τ *trajectory is stable for all* $\tau \in S$.

Before proving the theorem, we introduce the *dominant* MBT of (U, ϕ), denoted by (U', ϕ'), which is also an MBT. We define $U' : V' \setminus W' \mapsto (V')^{rd}$ and $\phi' : \{0,1\}^{rd} \mapsto \{0,1\}$ as follows:

$$U'(v) \stackrel{\text{def}}{=} (v + (u_1, t_1), v + (u_1, t_1 - 1), \ldots, v + (u_1, t_1 - d), \ldots, v + (u_r, t_r - d)),$$

$$\phi'(x_{U'(v)}) \stackrel{\text{def}}{=} \phi\left(\bigvee_{i=0}^{d} x_{v+(u_1,t_1-i)}, \ldots, \bigvee_{i=0}^{d} x_{v+(u_r,t_r-i)}\right).$$

Clearly, ϕ' is monotonic. We denote a trajectory of (U', ϕ') with base $x_{W'}$ by $T'(x_{W'})$ and call it a *dominant* trajectory. In order to apply Theorem 1 to (U', ϕ'), we will need to know its minimal zero sets.

Lemma 1. *If $Z_1, \ldots, Z_Q$ are all of the minimal zero sets of (U, ϕ), then $C(Z_1), \ldots, C(Z_Q)$ are all of the minimal zero sets of (U', ϕ').*

Proof. By virtue of the $\vee$'s in the definition of ϕ', it is clear that the $C(Z_i)$'s are zero sets, and furthermore that they are indeed minimal. Now let A' be a minimal zero set of (U', ϕ'). Since it is minimal, it must have the form $C(A)$ for some $A \in W$. It is then easy to see that A must be a minimal zero set of (U, ϕ). ∎

Lemma 2. *Given any base $x_{W'}$, there exists a $\tau \in S$ such that $T'(x_{W'}) = T^\tau(x_{W'})$.*

Proof. Let $x = T'(x_{W'})$ and $y = T^\tau(x_{W'})$. We will construct the desired τ inductively. Assume that we know the values of $\tau(v)$ for all $v \in \{(s,t) : t \leq t^0\}$ for some $t^0 \geq 0$. Letting $a = (s, t^0 + 1)$, where $s \in \mathbf{Z}^p$, we define $\tau(a)$ as follows. If $\bigvee_{i=0}^{d} x_{a+(u_j,t_j-i)} = 0$, then $\tau_j(a)$ can be arbitrary. Else, if the expression equals one, then for some k, $x_{v+(u_j,t_j-k)} = 1$, and we let $\tau(a) = k$. Simple verification shows that $\phi'(x_{U'(a)}) = \phi(x_{U^\tau(a)})$. ∎

Proof of Theorem 2. We will first show that (a) implies (b). Let $x_{W'}$ be any base such that $|I(x_{W'})|$ is finite. By Theorem 1, (a) implies that the zero state is attractive for the dominant MBT, that is, $|I(T'(x_{W'}))|$ is finite. Now $(T'(x_{W'}))_v \geq (T^\tau(x_{W'}))_v$ for all $v \in V'$ and $\tau \in S$, a fact easily established by induction. Thus, $|I(T'(x_{W'}))| \geq |I(T^\tau(x_{W'}))|$ for all $\tau \in S$, which implies (b).

Now we will prove that (a) implies (c). Let z be the zero trajectory and M'_ϵ be the M_ϵ set of the dominant MBT. Again by Theorem 1, (a) implies that

$$\lim_{\epsilon \to 0} \sup_{\substack{\mu \in M'_\epsilon(z_W) \\ v \in V}} \mu(x_v = 1) = 0.$$

Therefore, given $\mu^\tau \in M^\tau_\epsilon(z_{W'})$ (where $\tau \in S$ and $\epsilon \in (0,1)$ are arbitrary), if we can find a $\mu' \in M'_\epsilon(z_{W'})$ such that $\mu'(x_v = 1) \geq \mu^\tau(x_v = 1)$ for all $v \in V' \setminus W'$, we will be done. Let μ^τ and ϵ be fixed.

We now define a $\{0,1\}^{V'}$ valued random variable e on an auxiliary space $\Omega = \{0,1\}^{V'}$ such that $P(e_v = 1 \, \forall \, v \in A) = \mu^\tau(x_v \neq \phi(x_{U^\tau(v)}) \, \forall \, v \in A)$ for all finite A. We know that the last term is not greater than $\epsilon^{|A|}$. Then for each $\omega \in \Omega$, we construct random states as follows. For all $v \in W'$, $x'_v = x^\tau_v = 0$, and $v \in V' \setminus W'$,

$$x^\tau_v = \begin{cases} \phi^\tau(x_{U^\tau(v)}), & \text{if } e_v = 0, \\ \overline{\phi^\tau(x_{U^\tau(v)})}, & \text{if } e_v = 1, \end{cases} \qquad x'_v = \begin{cases} \phi'(x_{U'(v)}), & \text{if } e_v = 0, \\ 1, & \text{if } e_v = 1, \end{cases}$$

where $\overline{(\cdot)}$ is the logical negation operator. Let μ' be the measure induced by x'; that is, for any state y, we will have $\mu'(x_v = y_v \, \forall \, v \in A) = P(x'_v = y_v \, \forall \, v \in A)$ for all finite A. Note that for all $\omega \in \Omega$, $x'_v \geq x^\tau_v$ for all $v \in V' \setminus W'$, which implies $P(x'_v = 1) \geq P(x^\tau_v = 1)$. Thus

$$\begin{aligned} \mu'(x_v = 1) &= P(x'_v = 1) \\ &\geq P(x^\tau_v = 1) \\ &= \mu^\tau(x_v = 1). \end{aligned}$$

On the other hand, for any finite A,

$$\begin{aligned} \mu'(x_v \neq \phi'(x_{U'(v)}) \, \forall \, v \in A) &\leq P(e_v = 1 \, \forall \, v \in A) \\ &\leq \epsilon^{|A|}, \end{aligned}$$

which implies $\mu' \in M'_\epsilon$.

We will now show by contradiction that (b) implies (a). Assume $\sigma' \neq 0$. Then by Theorem 1 the zero trajectory is not attractive for (U', ϕ'), that is, there exists an $x_{W'}$ such that $|I(x_{W'})|$ is finite but $|I(T'(x_{W'}))|$ is infinite. By Lemma 2 there exists a scenario τ such that $I(T'(x_{W'})) = I(T^\tau(x_{W'}))$. Thus for this τ the zero state is not an attractive τ trajectory.

Finally, we will show that (c) implies (a), also by contradiction. Again assume $\sigma' \neq 0$. Then there exists a ray $\beta \subset \sigma'$ originating from the origin. By Theorem 8 and Lemma 11 in [6], we know that there exists a finite set

$A \subset \{(s,t) \in \mathbf{R}^{p+1} : t < 0\}$ such that the set $P \stackrel{\text{def}}{=} (A+\beta) \cap V'$ possesses the property that $P \cap W' \subset I(x_{W'})$ implies $P \subset I(T'(x_{W'}))$. For example, A can be defined as

$$A = -(p+1)\sum_{q=1}^{Q} \operatorname{conv}(Z_q) + B + b,$$

where B is a unit cube in $\mathbf{R}^{p+1}$ that is added to make $A + \sigma'$ contain points $(s,t) \in V'$ for all $t \geq 0$; and b is a vector that is added to make $(s,t) \in A \implies t < 0$ [6]. Other choices can also work.

To facilitate the rest of the proof, we define the sets $D_n \subset V'$ for $n \geq 0$ by

$$D_n \stackrel{\text{def}}{=} \{(s,t) \in A+\beta : n \leq t < n + d + t_W\} \cap V'.$$

We also let y be the state

$$y_v = \begin{cases} 1, & \text{if } v \in P; \\ 0, & \text{otherwise.} \end{cases}$$

We claim that if $v \in \{(s,t) \in P : t \geq 0\}$, then $\phi'(y_{U^\tau(v)}) = 1$. To see this, note that $(a,\ldots,a) - U'(a)$ is a constant for all $a \in V' \setminus W'$. Therefore, by Theorem 8 in [1], for any dominant trajectory x such that $D_n \subset I(x)$, we have $D_{n+1} \subset I(x)$, for any non-negative integer n.

We now construct a scenario τ as follows. If $v \notin P$, then $\tau(v)$ can be arbitrary. If $v \in P$, then as we argued $\phi'(y_{U'(v)}) = 1$, and as was done in the proof of Lemma 2, we can define $\tau(v)$ such that $\phi(y_{U^\tau(v)}) = 1$. With this τ, simple inductive arguments show that for any τ trajectory x such that $D_n \subset I(x)$, we have $D_m \subset I(x)$ for all $m \geq n$.

We now construct a measure $\mu_\epsilon \in M_\epsilon^\tau$. Let $\omega \in \{0,1\}^{V'\setminus W'}$ and define a map $\Gamma : \{0,1\}^{V'\setminus W'} \mapsto \{0,1\}^{V'}$ by

$$x_a \stackrel{\text{def}}{=} \begin{cases} 0, & \text{if } a \in W', \\ \max(\phi(x_{U^\tau(a)}), \omega_a), & \text{if } a \in V' \setminus W'. \end{cases}$$

Let $\tilde{\mu}_\epsilon$ be a Bernoulli measure on $\{0,1\}^{V'\setminus W'}$ such that for any finite $A \subset V'$

$$\tilde{\mu}_\epsilon(\omega_a = 1 \, \forall \, a \in A) = \epsilon^{|A|}.$$

Let μ_ϵ designate the measure $\{0,1\}^{V'}$ induced by the measure $\tilde{\mu}_\epsilon$ on $\{0,1\}^{V'\setminus W'}$ with the map Γ. Clearly $\mu_\epsilon \in M_\epsilon^\tau(z_{W'})$, where z is the zero state.

For any t, let s_t be such that $(s_t, t) \in P$. Now, by the arguments in the previous paragraph, if $\omega_v = 1$ for all $v \in D_n$ for some $n \geq 0$, then $(\Gamma(\omega))_a = 1$ for all $a \in D_m$ with $m \geq n$. Thus the condition $(\Gamma(\omega))_{(s_t,t)} = 1$ is ensured if there is some $n \leq t$ such that $\omega_v = 1$ for all $v \in D_n$. Let $n_i \stackrel{\text{def}}{=} i(t_W + d)$. Then, clearly, the events "$\omega_a = 1$ for all $a \in D_{n_i}$" are independent of each other. The probability of every quoted event is at least ϵ^N, where $N \stackrel{\text{def}}{=} \sup_{n\geq 0} D_n$, a finite integer. Thus

$$\mu_\epsilon(x_{(s_t,t)} = 1) \geq 1 - (1-\epsilon^N)^{[t/(t_W+d)]} \stackrel{t\to\infty}{\longrightarrow} 1.$$

In other words, the zero state is not a stable τ trajectory. ∎

Thus, in order to determine if the zero trajectory of an MBT is stable under asynchronism, it is enough to apply Toom's criterion to the dominant MBT.

Remark 1. Dual arguments can be used to study the stability of the one trajectory under asynchronism. We are particularly interested in the effects of asynchronism on MBTs where both the zero and the one trajectories are stable. We will discuss these effects through a series of examples in section 4.

Remark 2. The concept of stability used in Theorems 1 and 2 is rather strong, namely, an MBT is called stable only if equation (1) holds with the supremum over all $\mu \in M_\epsilon$. In particular, this requires stability under asymmetric noise where 0's can become 1's, but not vice versa. It would be interesting to characterize stability under more symmetric noise models (e.g., with positive probabilities for each kind of error). See, however, Example 3 of section 4, which shows that asynchronism can cause MBTs with multiple stable trajectories to become ergodic even under symmetric noise.

4. Examples

We conclude our discussion with several examples and a theorem.

Example 1. Here we note that there are MBTs such that for a given d, the zero trajectory is attractive for all of the homogenous scenarios but not attractive for all scenarios. For example, consider a one-dimensional (U, Φ) with $(u_1, t_1) = (2, -1), (u_2, t_2) = (3, -1)$, and $\phi(a, b) = a \wedge b$. Let $d = 2$. By Theorem 1 the zero trajectory is attractive for all four possible homogenous scenarios, whereas by Theorem 2 it is not so for all scenarios. The reason for this becomes clear when we try to construct a scenario τ according to the procedure in the proof. For this example, we can take

$$P = \{(s, t) \in V' : -\tfrac{1}{2}s - 2 \leq t \leq -\tfrac{1}{2}s\}.$$

Then, for instance, for the point $(-4, 0)$ we get $\tau((-4, 0)) = (0, 0)$, while for the point $(0, 0)$ we get $\tau((0, 0)) = (0, 1)$ or $(1, 1)$. (See figure 2.)

Example 2. The next two examples show how asynchronism may affect the zero and the one trajectory differently. Let $a_1 = (-1, 2, -1)$, $a_2 = (-1, 0, -1)$, $a_3 = (1, 0, -1)$, and $a_4 = (0, 2, -3)$. Define U by the assignment $(u_i, t_i) = a_i$, $i = 1, 2, 3, 4$, and define ϕ by zero sets $\{a_1, a_2\}$, $\{a_2, a_3\}$, $\{a_3, a_4\}$, $\{a_4, a_1\}$, and one sets $\{a_1, a_3\}$, $\{a_2, a_4\}$. Such sets do determine a unique monotonic ϕ.

By Theorem 1, the zero and the one trajectories are both stable for this (U, ϕ). However, by Theorem 2, the one trajectory becomes unstable for some scenarios when $d \geq 1$, while the zero trajectory remains stable for all d. In the next example we will show how both the one and the zero trajectories may become unstable for different values of d. (See figure 3.)

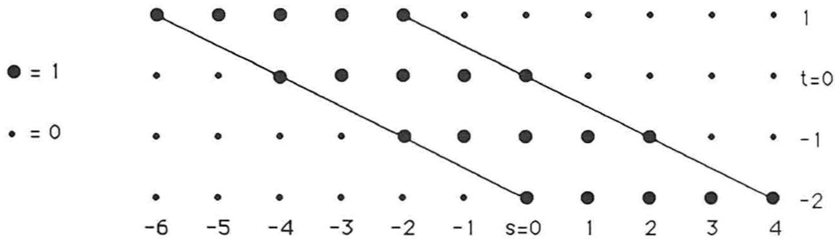

Figure 2: Constructing a scenario as prescribed by Example 1. Given the base, we want a τ that guarantees a state of 1 for all points in the slab P. The two lines in the figure are the boundary of P. For example, in order to get a state 1 for the point $(-4, 0)$, we cannot have any delays (i.e., $\tau(-4, 0) = (0, 0)$), whereas to get a state 1 at point $(0, 0)$, we must have $\tau_2(0, 0) = 1$. Thus τ cannot be homogenous.

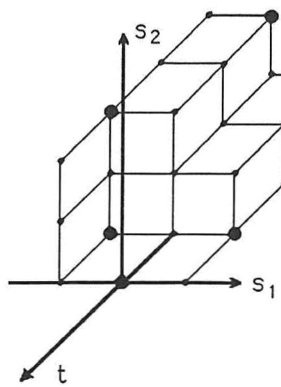

Figure 3: The origin and $U(0)$ of the MBT in Example 2. By Theorem 1, both the zero and the one trajectories are stable for this MBT; however, by Theorem 2, the one trajectory may become unstable for any $d \geq 1$, whereas the zero trajectory remains stable for all d.

Example 3. Here we indicate that the depth of asynchronism needed to erase the stability of the zero trajectory can be different from that for the one trajectory. Let $a_1 = (3, 2, -2)$, $a_2 = (3, 1, -2)$, $a_3 = (9, 2, -2)$, and $a_4 = (9, 3, -4)$. Define U by the assignment $(u_i, t_i) = a_i$, $i = 1, 2, 3, 4$, and define ϕ (uniquely) by zero sets $\{a_1, a_2\}$, $\{a_2, a_3\}$, $\{a_3, a_4\}$, $\{a_4, a_1\}$, and one sets $\{a_1, a_3\}$, $\{a_2, a_4\}$.

Now Theorem 1 tells us that the zero and the one trajectories are stable for this (U, ϕ). Theorem 2 and simple geometric arguments show that the one trajectory becomes unstable for some scenarios when $d = 1$, whereas the zero trajectory becomes unstable when $d = 2$. Thus the zero and the one trajectories may behave differently under asynchronism. (See figure 4.)

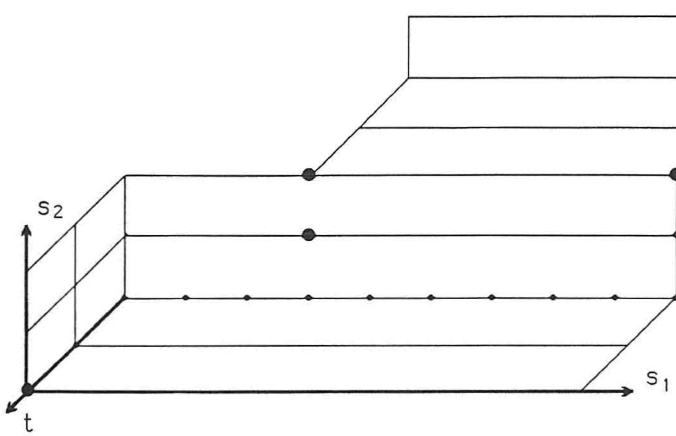

Figure 4: The origin and $U(0)$ of the MBT in Example 3. By Theorem 1, both the zero and the one trajectories are stable for this MBT; however, by Theorem 2, the one trajectory becomes unstable for some scenarios when $d = 1$, whereas the zero trajectory becomes unstable when $d = 2$.

Example 4. This example demonstrates that asynchronism may cause a non-ergodic MBT (that is, one with several stable trajectories) to become ergodic. In fact, this can happen under symmetric noise models (see Remark 2). Consider the two-dimensional MBT (U, ϕ), where $(u_1, t_1) = (2, 0, -1)$, $(u_2, t_2) = (1, 1, -2)$, $(u_3, t_3) = (0, 2, -2)$, and ϕ is the majority voter rule. (See figure 5(a).) We also define $\mu_\epsilon \in M_\epsilon$ by the relation

$$\mu_\epsilon(x_v \neq \phi(x_{U(v)}) \, \forall \, v \in A) = \epsilon^{|A|}$$

for any finite $A \subset V$. In other words, the probability of an error at any point is independent of other points and is equal to ϵ. Now let μ_ϵ^0 and μ_ϵ^1 be such measures in $M_\epsilon(x_W^0)$ and $M_\epsilon(x_W^1)$, respectively, where x^0 is the zero trajectory and x^1 is the one trajectory. In view of Theorem 1, we have

$$\limsup_{\epsilon \to 0} \sup_{v \in V} \mu_\epsilon^0(x \neq 0) = 0 \qquad \text{and} \qquad \limsup_{\epsilon \to 1} \sup_{v \in V} \mu_\epsilon^1(x \neq 1) = 1.$$

Now consider the homogenous scenario τ, where $\tau(v) = (1, 0, 0)$ for all v. This MBT is actually composed of infinitely many non-interacting, one-dimensional MBTs that are independent of each other. To see this, take π to be the plane in V' that goes through the origin and has a normal $(1, 1, 1)$ (figure 5(b)). Intersection of π with V' is a lattice isomorphic to $\{(s, t) \in \mathbf{Z}^2 : t \geq -3\}$, which we will call V''. Consider the one-dimensional MBT (U'', ϕ), where $(u_1, t_1) = (-1, -2)$, $(u_2, t_2) = (0, -2)$, and $(u_3, t_3) = (1, -2)$. Define μ_ϵ^0 and μ_ϵ^1 for (U'', ϕ), as was done previously. It has recently been shown that this one-dimensional automaton is ergodic under symmetric noise

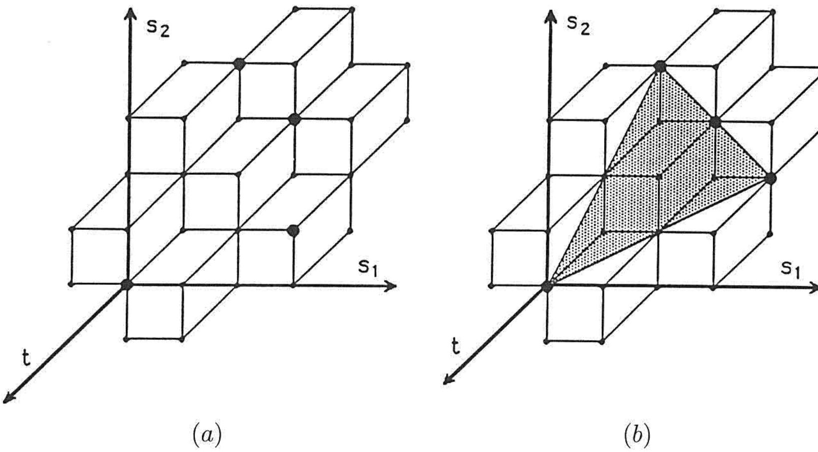

Figure 5: The origin and $U(0)$ of the original MBT of Example 4 are indicated in part (a). Once we let $\tau(v) = (1, 0, 0)$ for all v, the origin and $U^\tau(0)$ will lie on the plane that has a section of it shaded in part (b). It is then easy to see that the MBT (U^τ, ϕ) is composed of infinitely many non-interacting, one-dimensional MBTs that reside on planes parallel to the mentioned plane.

of arbitrary small error probability ϵ (see [3]). In particular, defining μ^0_ϵ and μ^1_ϵ for (U'', ϕ), as was done previously, we have

$$\lim_{\epsilon \to 0} \sup_{v \in V''} \mu^0_\epsilon(x \neq 0) = 1/2 \qquad \text{and} \qquad \lim_{\epsilon \to 0} \sup_{v \in V''} \mu^1_\epsilon(x \neq 1) = 1/2.$$

Once we see that (U^τ, ϕ) is composed of infinitely many such one-dimensional MBTs operating on planes parallel to the one mentioned, the ergodicity of (U^τ, ϕ) follows.

As a final remark, we conclude with a discussion of the possibility of the constructing one-dimensional examples in which asynchronism has a nontrivial effect. This is marginally related to the famous *positive rates conjecture* [2, 4].

It is clearly possible to construct one-dimensional MBTs in which the zero and the one configurations are trajectories. It is then natural to ask if they can both be stable. If this were so it would be easy to construct scenarios in which asynchronism erases memory by considering homogenous scenarios where the automaton breaks up into zero-dimensional automata, which are clearly ergodic.

The following theorem shows that in a one-dimensional MBT the zero and the one trajectories cannot both be stable.

Theorem 3. No one-dimensional MBT (U, ϕ) can have both a stable zero trajectory and a stable one trajectory.

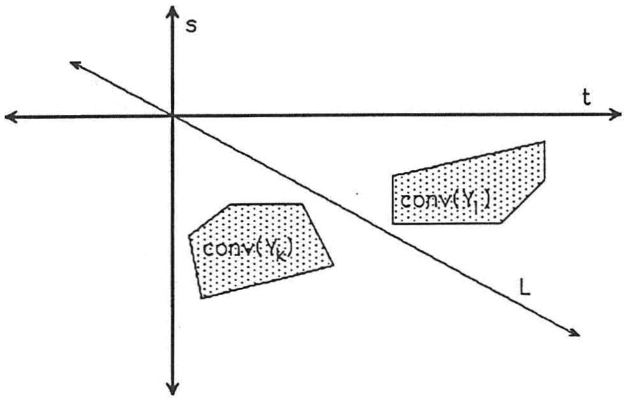

Figure 6: As mentioned in the Proof of Theorem 3, $\sigma_0 = \{0\}$ implies the existence of a line L through the origin that properly separates some Y_k from some Y_l.

Proof of Theorem 3. Let $Y_1, \ldots, Y_n$ be an enumeration of the minimal one sets of ϕ and $Z_1, \ldots, Z_m$ be an enumeration of the minimal zero sets of ϕ (see Remark 1). Since Z_j is a minimal zero set of ϕ, even if $x = 0$ on $U(0) \setminus Z_j$, we will have $\phi(x_{U(0)}) = 1$ if $x_a = 1$ for all $a \in Z_j$. Thus none of the Y_i's can be contained in $U(0) \setminus Z_j$, that is, each Y_i meets every Z_j.

Let σ_0 be the σ set of (U, ϕ) as before. We define an analogous set σ_1 for the minimal one sets of (U, ϕ). In view of Remark 1, the one trajectory is stable (attractive) if and only if $\sigma_1 = \{0\}$. Now assume that the zero trajectory is stable. Since this implies that $\sigma_0 = \{0\}$, there must exist a line L through the origin that properly separates some $\text{conv}(Y_k)$ from some $\text{conv}(Y_l)$, and thus Y_k from Y_l (figure 6). Now since each Z_j meets both Y_k and Y_l, L must meet each $\text{conv}(Z_j)$. Thus σ_1 contains $\{(s,t) \in \text{L} : t \geq 0\}$, which implies that the one trajectory is not stable. ∎

Acknowledgments

Research supported by NSF under PYI award NCR 8857731 and IRI 9005849, by IBM under a Faculty Development award, and by BellCore Inc.

References

[1] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Computation* (Englewood Cliffs, NJ, Prentice Hall, 1989).

[2] P. Gács, "Nonergodic One-dimensional Media and Reliable Computation," in *Particle Systems, Random Media and Large Deviations*, edited by R. Durrett (Providence, RI, American Mathematical Society, 1985).

[3] L. F. Gray, "The Behavior of Processes with Statistical Mechanical Properties," in *Percolation Theory and Ergodic Theory*, edited by H. Kesten (New York, Springer Verlag, 1987).

[4] T. M. Liggett, *Interacting Particle Systems* (New York, Springer Verlag, 1985).

[5] N. Pippenger, "Developments in the Synthesis of Reliable Organisms from Unreliable Components," in *The Legacy of John von Neumann*, edited by J. Glimm, J. Impagliazzo, and I. Singer (Providence, RI, American Mathematical Society, 1990).

[6] A. L. Toom, "Stable and Attractive Trajectories in Multicomponent Systems," in *Multicomponent Random Systems*, edited by R. L. Dobrushin and Y. G. Sinai (New York, Marcel Dekker, 1980).

[7] S. Wolfram, "Universality and Complexity in Cellular Automata," *Physica*, **10D** (1984) 1–35.