

Copyright © 2005, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

**DISCRETE-EVENT SYSTEMS:
GENERALIZING METRIC SPACES
AND FIXED-POINT SEMANTICS**

by

Adam Cataldo, Edward A. Lee, Xiaojun Liu,
Eleftherios Matsikoudis and Haiyang Zheng

Memorandum No. UCB/ERL M05/12

8 April 2005

**DISCRETE-EVENT SYSTEMS:
GENERALIZING METRIC SPACES
AND FIXED-POINT SEMANTICS**

by

Adam Cataldo, Edward A. Lee, Xiaojun Liu,
Eleftherios Matsikoudis and Haiyang Zheng

Memorandum No. UCB/ERL M05/12

8 April 2005

ELECTRONICS RESEARCH LABORATORY

College of Engineering
University of California, Berkeley
94720

Discrete-Event Systems: Generalizing Metric Spaces and Fixed-Point Semantics

Adam Cataldo, Edward Lee, Xiaojun Liu, Eleftherios Matsikoudis, and
Haiyang Zheng

Center for Hybrid and Embedded Software Systems (CHESS)
University of California, Berkeley, 94720, USA
{acataldo, eal, liuxj, ematsi, hyzheng}@eecs.berkeley.edu

Abstract. This paper studies the semantics of discrete-event systems as a concurrent model of computation. The classical approach, which is based on metric spaces, does not handle well multiplicities of simultaneous events, yet such simultaneity is a common property of discrete-event models and modeling languages. (Consider, for example, delta time in VHDL.) In this paper, we develop a semantics using an extended notion of time. We give a generalization of metric spaces that we call tetric spaces. (A tetric functions like a metric, but its value is an element of a totally-ordered monoid rather than an element of the non-negative reals.) A straightforward generalization of the Banach fixed point theorem to tetric spaces supports the definition of a fixed-point semantics and generalizations of well-known sufficient conditions for avoidance of Zeno conditions.

1 Introduction

Discrete-event (DE) systems are widely used in modeling and simulation (e.g. [1]) and in circuit design (e.g. [2]). Historically, distributed and parallel implementations of DE systems have been constructed to achieve faster simulation (e.g. [3, 4]). Recently, however, DE principles are getting applied to intrinsically distributed systems, where the focus is not on faster simulation but rather on a timed coordination mechanism. For example, the TeaTime protocol in Croquet, a shared 3-D immersion environment [5], is a distributed DE system that combines the concept of optimistic computation [6, 4] with distributed consensus [7]. The emergence of high-precision network time synchronization (e.g. the IEEE 1588 standard) also creates compelling new possibilities for the use of DE principles in distributed embedded software for applications such as industrial automation and instrumentation [8].

This paper studies the semantics of DE systems as a concurrent model of computation. The classical approach to this semantics is based on metric spaces [9, 10]. We show in this paper that these semantics has some serious limitations that can be overcome by using the notion of superdense time [11, 12] and a generalization of metric spaces that we call *tetric spaces*. Whereas a metric is a function that yields a non-negative real number, a tetric is a function that yields

an element of a totally ordered monoid. The classical semantics uses a fixed-point whose uniqueness is assured by the Banach fixed point theorem. Our semantics uses a fixed point whose uniqueness is assured by a straightforward generalization of the Banach fixed-point theorem to tetric spaces. Many of the classical results also generalize in a straightforward way, including sufficient conditions for the avoidance of Zeno conditions.

In DE systems, concurrent objects (which we call processes) interact via *signals* consisting of *events*, where an event has a *time* and a *value*. As a concurrency model, DE, at first, seems straightforward and easily understood by system designers: events are processed in chronological order, much as in the physical world. The semantics can become quite subtle, however, when one considers events that are arbitrarily close in time or simultaneous.

One approach is to avoid these subtleties by assuming them away. For example, in [9], Yates assumes a minimum separation in time between any two events in a signal. However, in the design of practical discrete-event languages, such as VHDL (which is widely used to specify highly concurrent systems, namely digital circuits), such separations are impractical, and would seriously weaken the discrete-event abstraction. In VHDL, time is given by a natural number that is interpreted as a multiple of a minimum time resolution. A signal, however, may have multiple events at the same time, in which case they are semantically distinguished by a second natural number, called the delta time, which gives simultaneous events an ordering.

Even for discrete-event languages where time is a real number, simultaneity is a useful concept. A common use of DE languages is to model mixed physical and software systems. For example, OPNET Modeler, a commercial tool from OPNET Technologies, and NS-2, a widely used research tool, are discrete-event languages for modeling computer networks. Such systems mix models of the physical world with software systems. In the physical world, if we ignore quantum effects, it is arguable propagation delays and arbitrary precision of time make simultaneity at least unlikely. However, in the software world, we can make no such argument. There is no time in software semantics, only an ordering of events. Mapping such semantics onto a time line without simultaneity is at best an artifice.




Fig. 1. The Newton's Cradle office toy.

Even in models of the physical world though, simultaneity proves to be a useful concept. Consider, for example, the Newton’s cradle office toy [13], which consists of three pendulums whose balls touch when at rest. See Figure 1. If we pull the first ball from the rest position and drop it, it will collide with the other two balls. The third ball will then fly off, while the first and second balls remain at rest. If we had started with the first two balls removed from rest, the second and third ball would fly off while the first ball remains at rest. To explain this, when we pull only the first ball, it collides with the second and comes to rest; then the second ball collides with the third and then comes to rest. We model these two collisions occurring at the same time, but in a well-defined order. In the two-ball case, the second ball collides with the third and comes to rest, then the first ball collides with the second and comes to rest, all at the same time. Here time is important in detecting when the collisions occur, but it is irrelevant in computing the sequence of collisions that occur at the same time.

This example illustrates a second subtlety, which is that even if events are not simultaneous, the time gap between them may get arbitrarily small. If we include friction in the model, then the time between collisions will decrease monotonically, and we will observe Zeno behavior, where the number of events in a finite time becomes infinite.

In this paper, we generalize the semantics of discrete-event systems to handle simultaneous events well. We generalize classical results about uniqueness of behavior and freedom from Zeno conditions. To do this, we first adopt a model of time that supports both events distributed over time and simultaneous events that are ordered. We then generalize the classical metric-space approach so that it works with the *superdense* model of time. We then show that many of the classical results generalize in a straightforward way to this new mathematical framework, including uniqueness and conditions for avoidance of Zeno behavior.

2 Discrete-Event Semantics

Our discrete-event semantics follows the tagged-signal framework of Lee and Sangiovanni-Vincentelli [14]. Time is represented by an element from a set of tags. To support simultaneous events, our tags have the following structure:

Definition 1 (Tags). Let $T = \mathbb{R}_+ \times \mathbb{N}_0$ be the set of tags¹.

Notice that this is different than the tag set $T = \mathbb{R}_+$ defined in [9, 10]. When $T = \mathbb{R}_+$, each $t \in T$ represents a time. When $T = \mathbb{R}_+ \times \mathbb{N}_0$, for each $t = (\tau, n) \in T$, τ represents a time, and n represents an index, which give us an ordering on the events at time τ . This notion of time is called *superdense time* in [11] and introduced in [12]. In [11], the authors claim that this model of time makes verification of hybrid systems difficult. We will show that this is the right

¹ In this paper, \mathbb{R} is the set of real numbers, \mathbb{N}_0 is the set of natural numbers (beginning with zero), \mathbb{Z} is the set of integers, and \mathbb{Q} is the set of rational numbers. The non-negative reals are denoted by \mathbb{R}_+ .

notion of time for our semantic model. Note that T is totally ordered under the lexicographic order² \preceq .

Definition 2 (Discrete). A set $D \subseteq T$ of tags is discrete³ if there exists an injective, order-preserving map $f : D \rightarrow \mathbb{N}_0$.

Definition 3 (Zeno). A discrete set $Z \subseteq T$ is Zeno if Z is infinite and there exists a $t \in T$ such that Z is bounded above by t .

Note that T itself is not discrete. Any discrete set is countable, but not every countable set is discrete, for example, $\{0, 1\} \times \mathbb{N}$ is not discrete. The discrete sets $Z_1 = \{0\} \times \mathbb{N}_0$, and $Z_2 = \{1/n \mid n \in \mathbb{N}\} \times \{0\}$ are Zeno. The set $D = \mathbb{N}_0 \times \{0\}$ is a non-Zeno-discrete set.

Definition 4 (Values). We let V be some arbitrary set of values.

These are the values a signal can take on. Since we are interested in the time behavior of discrete events, the structure of V is irrelevant to us.

Definition 5 (Signal). A partial function $s : T \rightarrow V$ is a signal.

We will denote the set of tags at which s is defined by

$$\text{Tag}(s) = \{t \in T \mid s(t) \text{ is defined}\}, \quad (1)$$

and the set of times at which s is defined by

$$\text{Time}(s) = \{\tau \in \mathbb{R}_+ \mid \exists n \in \mathbb{N}_0, (\tau, n) \in \text{Tag}(s)\}, \quad (2)$$

When s is not defined at t , we will say $s(t) = \perp$, for convenience. Here we assume $\perp \notin V$. In this sense s is a total function from T to $V \cup \{\perp\}$. We say that $s : T \rightarrow V$ is a *discrete-event signal* if $\text{Tag}(s)$ is discrete and that it is a *Zeno signal* if $\text{Tag}(s)$ is Zeno. This is consistent with the definition in hybrid systems [16]. We let $S = [T \rightarrow V]$ be the set of all signals. Given $n \in \mathbb{N}$, S^n is the set of all n -tuples of signals. If $\mathbf{s} \in S^n$ and $t \in T$, then, with slight abuse of notation, we say $\mathbf{s}(t) = (s_1(t), \dots, s_n(t))$. In this sense, we can group n signals together, so

$$\mathbf{s} : T \rightarrow (V \cup \{\perp\})^n \quad (3)$$

is simply another signal. As in [10], $S^0 := \{\sigma\}$, a singleton set with element σ . In general, if I is any finite index set, then S^I is the set of functions from I to S , or set of I -tuples of signals. In this sense, S^n is just syntactic sugar for $S^{\{1, \dots, n\}}$. For a signal $\mathbf{s} \in S^I$ and a subset $K = \{k_1, \dots, k_m\}$ of I , we define the projection onto K by

$$\pi_K(\mathbf{s}) = (s_{k_1}, \dots, s_{k_m}). \quad (4)$$

² $(\tau_1, n_1) \preceq (\tau_2, n_2)$ if and only if $\tau_1 < \tau_2$, or $\tau_1 = \tau_2$ and $n_1 \leq n_2$.

³ Our definition of “discrete” is equivalent to, but simpler than that given by Mazurkiewicz [15].

Definition 6 (Process). Given two finite sets I and J , a process⁴ is any function $F : S^I \rightarrow S^J$.

The most basic process is the identity process $F : s \mapsto s$. The delay process delays the input by time $\tau' \in \mathbb{R}_+$:

$$\forall (\tau, n) \in T, F(s)(\tau + \tau', n) = s(\tau, n). \quad (5)$$

We call any process $F : S^0 \rightarrow S^J$ a *source* and any process $F : S^I \rightarrow S^0$ a *sink*.

Definition 7 (Composite Process). Given processes $F_1 : S^{I_1} \rightarrow S^{J_1}$ and $F_2 : S^{I_2} \rightarrow S^{J_2}$ and $K \subset J_1 \cap I_2$, for each $s \in S^{I_2}$, let

$$F_2(\pi_{I_2/K}(s), \pi_{I_2 \cap K}(s)) := F_2(s). \quad (6)$$

Then we define the composite process $F_K : S^{I_1} \times S^{I_2/K} \rightarrow S^{J_1} \times S^{J_2}$ as

$$F_K : (s_1, s_2) \mapsto (F_1(s_1), F_2(s_2, F_1(s_1))). \quad (7)$$

When F_1 and F_2 are the same process, we call this *feedback composition*. If F_1 and F_2 are distinct and $K = \emptyset$, we have *parallel composition*. Otherwise we have *series composition*. In the composite process, a signal is an output if it is an output of either F_1 or F_2 . Otherwise, it is an input. This is consistent with the definition of inputs and outputs in [9]. Compositions are easy to visualize using block diagrams. See the example in Figure 2. Repeated composition of processes




Fig. 2. In this case, $s_3 = F(s_1, G(\sigma))$.

allows us to build networks of processes.

Consider the identity process composed with itself. Then our composite process $F_K : S^0 \rightarrow S$ could be defined by $F_K(\sigma) = s$ for any $s \in S$, since $F(s) = s$. In this case, there is no unique solution to the feedback composition of F with itself. If we disallow feedback, we severely limit the types of systems we can construct. We will develop conditions that give a network of processes a unique output for each input. Before we can address this, we must introduce some topological concepts which we will use to reason about the meaning.

⁴ This is called a functional process in [14].

3 An Extension of Metric Spaces

Recall that the Banach fixed-point theorem gives us a method to prove the existence of a unique solution to an ordinary differential equation under certain conditions. Moreover, it gives us a way to construct the solution starting from any guess of the solution. We are interested in when there exists a unique solution to a network of processes with feedback. In this section we will extend the concept of a metric space and extend the Banach fixed point theorem to this generalization. We will then use this new fixed point theorem to show when there exists a unique solution to a network of processes with feedback.

3.1 Tetric Spaces

Recall the definition of a monoid from group theory.

Definition 8 (Monoid⁵). A set M combined with a binary operation $+$ defined on M is a monoid iff for all $a, b, c \in M$:

1. *Closure:* $a + b \in M$,
2. *Associativity:* $a + (b + c) = (a + b) + c$,
3. *Identity:* $\exists 0 \in M, a + 0 = 0 + a = a$.

A monoid $(M, +)$ is *commutative* iff $a + b = b + a$ for all $a, b \in M$. As an example, $(\mathbb{R}_+, +)$, with $+$ being the standard addition operator, is a commutative monoid. Another example is the structure (\mathbb{N}_0, \vee) , with \vee such that for all $a, b \in \mathbb{N}_0$:

$$a \vee b := \max\{a, b\} \tag{8}$$

We now define a total order.

Definition 9 (Total Order). A structure (M, \leq) , with $\leq \subseteq M \times M$, is a total order iff for all $a, b, c \in M$:

1. *Reflexivity:* $a \leq a$.
2. *Weak antisymmetry:* $a \leq b$ and $b \leq a \Rightarrow a = b$.
3. *Transitivity:* $a \leq b$ and $b \leq c \Rightarrow a \leq c$.
4. *Comparability:* either $a \leq b$ or $b \leq a$.

As an example, (\mathbb{R}, \leq) is a total order under the standard denotation of the operator \leq .

The following two definitions are adopted from [17]:

Definition 10 (Tomonoid). $(M, +, \leq)$ is a tomonoid iff:

- $(M, +)$ is a monoid,
- (M, \leq) is a total order,

⁵ Note that the closure property is not always included in the definition of a monoid. We include it here since all the monoids we are interested in exhibit closure, and we wish to exploit this property.

– *Translation invariance*: $\forall a, b, c \in M, a \leq b \Rightarrow a + c \leq b + c$.

A tomonoid $(M, +, \leq)$ is *positive* iff the identity element 0 of the monoid $(M, +)$ is the minimum element of the total order (M, \leq) . Note that $(\mathbb{R}_+, +, \leq)$ is a positive tomonoid, whereas $(\mathbb{R}, +, \leq)$ is not. A less trivial positive tomonoid is $(\mathbb{N}_0, \vee, \leq)$. We say that the tomonoid is *commutative* iff the underlying monoid is commutative.

For any monoid $(M, +)$, any sequence $\langle m_i \rangle \in M^\omega$, and any $n \in \mathbb{N}$, we can define the sum

$$\sum_{i=1}^n m_i = m_1 + m_2 + \dots + m_n. \quad (9)$$

Definition 11 (Summable). We call a positive tomonoid $(M, +, \leq)$ *summable* if whenever there is an $m \in M$ such that for all $n \in \mathbb{N}$

$$\sum_{i=1}^n m_i \leq m, \quad (10)$$

$$\sum_{i=1}^{\infty} m_i := \lim_{n \rightarrow \infty} \left(\sum_{k=1}^n m_k \right) \in M. \quad (11)$$

Observe that $(\mathbb{R}_+, +, \leq)$ and $(\mathbb{R}_+, \vee, \leq)$ are both summable while $(\mathbb{Q}_+, +, \leq)$ is not. For example, the natural number e can be written as a sum of rational numbers.

Using these definitions, we can define an extension of a metric. We call this extension *tetric*, short for tomonoid metric.

Definition 12 (Tetric). Given a set X and a positive, commutative, summable tomonoid $(M, +, \leq)$, a function $d : X \times X \rightarrow M$ is a *tetric* iff for all $a, b, c \in X$:

- *Identity of Indiscernibles*: $d(a, b) = 0 \Leftrightarrow a = b$,
- *Symmetry*: $d(a, b) = d(b, a)$,
- *Triangle Inequality*: $d(a, c) \leq d(a, b) + d(b, c)$.

If $M = \mathbb{R}_+$ and $+$ and \leq denote the standard operators, then our tetric becomes a metric. If (M, \leq) is any total order with a minimum element 0, and $+$ is \vee with the semantics of equation (8), then our tetric becomes a generalized ultrametric, in the sense of [18], restricted to a totally ordered set M .

Given a totally ordered set M and any $a, b \in M$, we can define $(a, b) = \{x \in M \mid a < x < b\}$. We similarly define $(a, b]$, $[a, b]$, $[a, b)$, where the closed bracket means to replace $<$ with \leq . We can define the following topology⁶ on M :

⁶ A *topology* on a set X is a collection \mathcal{T} of subsets of X that include X itself and \emptyset . A topology must be closed under finite set intersection and arbitrary set union. An element of the topology is called an *open set* and its complement is a *closed set*. A sequence $\langle x_i \rangle$ in X^ω *converges* to $x \in X$ if for every open set U containing x , there is some $n \in \mathbb{N}$ such that $i > n$ implies $x_i \in U$. For two topological spaces (X, \mathcal{T}) and (Y, \mathcal{S}) , we say that a function $f : X \rightarrow Y$ is continuous if for all $S \in \mathcal{S}$, $f^{-1}(S) \in \mathcal{T}$.

Definition 13 (Order Topology). Given a total order (M, \leq) , let

$$\mathcal{B} := \{(a, b) \subseteq M \mid a, b \in M\}. \quad (12)$$

If M has a minimum element 0 , include all sets of the form $[0, a)$ in \mathcal{B} . If M has a maximal element ∞ , include all sets of the form $(a, \infty]$ in \mathcal{B} . Let \mathcal{T} be the collection of all arbitrary unions of elements in \mathcal{B} . Then \mathcal{T} is a topology on M , which we call the order topology.

If $M = \mathbb{R}$, and we use the standard \leq operator, then this is the standard topology of the real numbers. In an order topology a sequence may converge to only one point.

Given a tetric $d : X \times X \rightarrow M$ and any $x \in X$, $\epsilon \in M$, let

$$B_d(x, \epsilon) := \{y \in X \mid d(x, y) < \epsilon\} \quad (13)$$

be the ϵ -ball centered at x . We can now define a tetric space.

Definition 14 (Tetric Space). Given a tetric $d : X \times X \rightarrow M$, let \mathcal{B} be the collection of ϵ -balls in X . Let \mathcal{T} , be the collection of all arbitrary unions of elements of \mathcal{B} . Then \mathcal{T} is a topology, and we call (X, \mathcal{T}) a tetric space.

Any metric space is a tetric space, with $M = \mathbb{R}_+$. An open research question is whether ever tetric space is a metric space under some suitable metric. We are interested in the convergence of functions with tetrizable⁷ domains:

Theorem 1. Let $f : X \rightarrow Y$ be a continuous function (in the topological sense). If a sequence $\langle x_i \rangle$ converges to x in X , then the sequence $\langle f(x_i) \rangle$ converges to $f(x)$ in Y . If X is tetrizable, then the converse is also true.

We omit the proof of this theorem here, as there is no novelty involved (e.g. replace metric with tetric in the proof of Theorem 21.3 in [19]).

3.2 The Tetric Fixed-Point Theorem

Given a function $f : X \rightarrow X$, we let $f^n(x)$ be the function applied n times to x . We let $f^0(x) := x$. We use this to define a type of contraction map on a tetric space:

Definition 15 (Additive Contraction). Given a tetric space (X, \mathcal{T}) , a function $f : X \rightarrow X$ is an additive contraction iff for all $a, b \in X$:

1. the sequence $\langle t_i \rangle$, with

$$t_i = d(f^i(a), f^i(b)), \quad (14)$$

is strictly decreasing and converges to 0 in M , and

⁷ As in metric spaces, we call a space (X, \mathcal{T}) tetrizable if there exists a tetric $d : X \times X \rightarrow M$ and corresponding $(M, +, \leq)$ such that (X, \mathcal{T}) is the tetric space induced by d .

2. the sequence $\langle s_i \rangle$, with

$$s_i = \sum_{k=i}^{\infty} t_k, \quad (15)$$

converges to 0 in M .

As an example, a δ -contraction defined over a metric space is an additive contraction. In a δ -contraction we have some $\delta \in (0, 1)$ such that for all $a, b \in X$, $d(a, b) \leq \delta \cdot d(f(a), f(b))$. In the case of an ultrametric $d : X \times X \rightarrow M$, for some positive tomonoid (M, \vee, \leq) , we only require that $\langle t_i \rangle$ is strictly decreasing and converges to 0, since:

$$s_i = \sum_{k=i}^{\infty} t_k = \max_{k \in \{i, i+1, \dots\}} \{t_i, t_{i+1}, \dots\} = t_i. \quad (16)$$

Definition 16 (Cauchy Sequence). A sequence $\langle x_i \rangle$ in a tetric space (X, T) is Cauchy iff for all $\epsilon \in M$, such that $\epsilon > 0$, there exists some $k \in \mathbb{N}_0$, such that for all $n, m > k$:

$$d(x_n, x_m) < \epsilon. \quad (17)$$

Definition 17 (Completeness). A tetric space (X, T) is complete iff every Cauchy sequence converges to some limit in X .

As an example, the standard real metric space is a complete tetric space. The subspace of rational numbers with this metric is an incomplete tetric space. We now present the main theoretical result of this paper:

Theorem 2 (Tetric Fixed-Point Theorem). Given a complete tetric space (X, T) , an additive contraction $f : X \rightarrow X$ has a unique fixed point. That is, there is a unique $x \in X$ such that $f(x) = x$.

Proof. We prove this theorem in three parts.

1. f is continuous.

Given $x \in X$ and some open set V containing $f(x)$, there exists some $\epsilon \in M$ such that $B_d(f(x), \epsilon) \subseteq V$. Since f is an additive contraction, $d(x, a) \geq d(f(x), f(a))$ for all $a \in X$, and thus $f(B_d(x, \epsilon)) \subseteq B_d(f(x), \epsilon) \subseteq V$, so f is continuous.

2. For all $x \in X$, the sequence $\langle f^i(x) \rangle$ converges to a fixed point x_* .

Letting $a = x$ and $b = f(x)$, and applying the triangle inequality and the definition of additive contraction, it is easy to see that for all $\epsilon > 0$, there exists a $k \in \mathbb{N}$ such that for all $n, m > k$ with $m > n$:

$$d(f^n(x), f^m(x)) \leq \sum_{p=n}^{m-1} d(f^p(x), f^{p+1}(x)) \leq \sum_{p=n}^{\infty} d(f^p(x), f^{p+1}(x)) < \epsilon. \quad (18)$$

Since this is a Cauchy sequence in a complete tetric space, it converges to some point x_* . Because f is continuous, by theorem 1, $\langle f^i(x) \rangle$ converges to x_* implies $\langle f^{i+1}(x) \rangle$ converges to $f(x_*)$. The limit of a sequence is unique if it exists, so $f(x_*) = x_*$.

3. The fixed point x_* is unique.
Suppose y_* is a different fixed point. Then

$$d(y_*, x_*) > d(f(y_*), f(x_*)) = d(y_*, x_*). \quad (19)$$

This contradiction can only be resolved if $y_* = x_*$.

Application of this theorem to metric spaces with δ -contractions yields the classic Banach fixed-point theorem. The following corollary deals with the application to ultrametric spaces, and is of particular importance in our discrete-event systems study.

Corollary 1. *Given a complete ultrametric space, a function satisfying the first condition of Definition 15 has a unique fixed point.*

4 Feedback Semantics

We now consider when a processes $F : S \rightarrow S$ has a unique feedback signal $s = F(s)$. Let $M = \{(m_1, m_2) \in \mathbb{R}_+^2 \mid (m_1 = 0) \Rightarrow (m_2 = 0)\}$. Observe that M , as a subset of \mathbb{R}_+^2 , is totally ordered under the relation \preceq . Using the definition of \vee given in Equation 8, (M, \vee, \preceq) is a positive, commutative, summable tomonoid. Given $s_1, s_2 \in S$, let

$$\Delta(s_1, s_2) := \{\tau \in \mathbb{R}_+ \mid \exists n \in \mathbb{N}_0. s_1(\tau, n) \neq s_2(\tau, n)\}. \quad (20)$$

We can then define the tetric $d : S \times S \rightarrow M$ as

$$d(s_1, s_2) = \begin{cases} (0, 0), & s_1 = s_2. \\ \left(\frac{1}{2^\tau}, 0\right), & s_1 \neq s_2. \tau = \inf \Delta(s_1, s_2) \notin \Delta(s_1, s_2) \\ \left(\frac{1}{2^\tau}, \frac{1}{2^n}\right), & s_1 \neq s_2. \tau = \inf \Delta(s_1, s_2) \in \Delta(s_1, s_2), \\ & n = \min \{n \in \mathbb{N}_0 \mid s_1(\tau, n) \neq s_2(\tau, n)\} \end{cases} \quad (21)$$

Note that if s_1 and s_2 never contain multiple events at a given time, that is, $\forall \tau \in \mathbb{R}_+, \forall n > 0, s_1(\tau, n) = s_2(\tau, n) = \perp$, then this is equivalent to the Cantor metric of [20], where we simply ignore the second element of $d(s_1, s_2)$. Note the following observation about (S, d) :

Lemma 1 (Completeness). *(S, d) is a complete tetric space.*

Proof. Let $\langle s_i \rangle$ be a Cauchy sequence in S^ω . Then for any $(\tau, n) \in T$, there exists a $k \in \mathbb{N}_0$ where $k_1, k_2 \geq k$ implies

$$d(s_{k_1}, s_{k_2}) \prec \left(\frac{1}{2^\tau}, \frac{1}{2^n}\right). \quad (22)$$

This in turn implies that for all $k_1, k_2 \geq k$ and for all $(\tau', n') \preceq (\tau, n)$,

$$s_{k_1}(\tau', n') = s_{k_2}(\tau', n'). \quad (23)$$

We then let $s(\tau', n') := s_k(\tau', n')$ for all $(\tau', n') \preceq (\tau, n)$ and all n' . We can extend the values of s for all time by choosing large enough (τ, n) . Thus, for any (τ, n) , we can find a $k \in \mathbb{N}_0$ where $k' \geq k$ implies

$$d(s_{k'}, s) \preceq \left(\frac{1}{2^{\tau}}, \frac{1}{2^n} \right), \quad (24)$$

and (s_i) converges to s .

Definition 18 (Delta Causal). A process $F : S \rightarrow S$ is delta causal if there exists a $\delta > 0$ and a map $N : \mathbb{R}_+ \rightarrow \mathbb{N}_0$, such that for all s_1 and s_2 with $d(s_1, s_2) = \left(\frac{1}{2^\tau}, \frac{1}{2^n} \right)$:

$$d(F(s_1), F(s_2)) \preceq \begin{cases} \left(\frac{1}{2^\tau}, \frac{1}{2^{n+1}} \right), & n < N(\tau) \\ \left(\frac{1}{2^{\tau+\delta}}, 1 \right), & n \geq N(\tau). \end{cases} \quad (25)$$

If we only allow one event at each time, then this is equivalent to delta causal as in [10]. Given a delta causal process and two inputs which agree through tag (τ, n) , their outputs will agree through tag $(\tau, n+1)$. If $n > N(\tau)$, their outputs will agree through tag $(\tau + \delta, 0)$. Note that the identity process is not delta causal. For the identity process any signal is a fixed point. A delay process, $\forall(\tau, n), F(s)(\tau+1, n) = s(\tau, n)$ with some initial value $F(s)(0, 0) := v_i$, is delta causal. If $F : S^I \rightarrow S^J$, we will say F is delta causal with respect to index $k \in I \cap J$ if for all inputs $s \in S^{I \setminus \{k\}}$, for all $s_1, s_2 \in S$

$$d(F(s, s_1), F(s, s_2)) \preceq \begin{cases} \left(\frac{1}{2^\tau}, \frac{1}{2^{n+1}} \right), & n < N(\tau) \\ \left(\frac{1}{2^{\tau+\delta}}, 1 \right), & n \geq N(\tau). \end{cases} \quad (26)$$

We have the following:

Proposition 1 (Fixed Points of Delta Causal Processes). A process $F : S \rightarrow S$ which is delta causal has a unique fixed point.

Proof. Given any $s_1, s_2 \in S$, define the sequence $\langle t_i \rangle \in M^\omega$ by

$$t_i = d(F^i(s_1), F^i(s_2)). \quad (27)$$

From the definition of delta causal, it follows that this sequence is monotonically decreasing. Suppose $d(s_1, s_2) = \left(\frac{1}{2^\tau}, \frac{1}{2^n} \right)$. Then we can find a subsequence, indexed by $n_1 < n_2 < \dots$ such that

$$t_{n_i} \leq \left(\frac{1}{2^{i\delta}}, \frac{1}{2^\tau}, 1 \right) \quad (28)$$

has the first component converging to 0. Since $(0, 0)$ is the only element (m_1, m_2) in M with $m_1 = 0$, the subsequence, and thus the sequence, converges to $(0, 0)$. Applying the result of Corollary 1, we see that F must admit a unique fixed point.

We can apply this result as follows: If we make any guess of the fixed point, s_0 , then the sequence $\langle F^i(s_0) \rangle$ converges to the fixed point. Note that the fixed point might be a Zeno signal and might not even be a discrete event signal. As an example, the process, $F : s \mapsto s_c$ is delta causal for any s_c , whether s_c is a non-Zeno discrete-event signal or not.

4.1 Alternative Metrics

It turns out that the Cantor metric of [20] can be extended to a metric over our tag set. Define $d_C : S \times S \rightarrow \mathbb{R}_+$ as

$$d_C(s_1, s_2) = \begin{cases} 0, & s_1 = s_2 \\ \frac{1}{2^\tau}, & s_1 \neq s_2, \tau = \inf \Delta(s_1, s_2) \end{cases} \quad (29)$$

d_C is an ultrametric on S . Using a proof similar to Lemma 1 we can show that (S, d_C) is a complete metric space.

Lemma 2. *The topology of the tetric space (S, d) , where d is defined in Equation 21, and the topology induced by the metric d_C are the same.*

Proof. Let T_C denote the topology induced by the metric d_C .

We first show that any open set O of the tetric space (S, d) is an open set of T_C . For any $s \in O$, there exists $(\tau, n) \in T$ such that for all $s' \in S$,

$$d(s', s) < \left(\frac{1}{2^\tau}, \frac{1}{2^n} \right) \quad (30)$$

implies $s' \in O$. By the definition of d and d_C , $d_C(s', s) < \frac{1}{2^{\tau+1}}$ implies 30,

$$\left\{ s' \in S \mid d(s', s) < \frac{1}{2^{\tau+1}} \right\} \subseteq O \quad (31)$$

O is an open set of T_C .

Showing that any open set O' of T_C is an open set of (S, d) is similar.

Because the two topologies are the same, they define the same notion of convergence and the same set of continuous functions. Their interesting differences are in defining contracting mappings and in the applicability of fixed point theorems. Using the Banach fixed point theorem on the metric space (S, d_C) requires that a process be δ -causal. With the tetric d defined in 21 we can apply the tetric fixed point theorem to a larger set of processes.

4.2 Non-Zeno Fixed Points

Definition 19 (Simple). *A process $F : S^I \rightarrow S^J$ is simple if $\text{Tag}(s)$ is a discrete, non-Zeno set implies that $\text{Tag}(F(s))$ is a discrete, non-Zeno set.*

The identity process is simple, even though it admits Zeno signals as fixed points. The constant output process, $\forall t \in T, F(s)(t) = c \neq \perp$ is not. It is trivial to show that parallel and series composition of simple processes are simple. We have the following result about feedback composition of simple processes:

Proposition 2 (Non-Zeno Fixed Points). *If $F : S \rightarrow S$ is a simple, delta causal process, then its fixed-point solution $s = F(s)$ is a non-Zeno discrete-event signal.*

Proof. Let s be the fixed point. We prove that over any finite interval $I \subset \mathbb{R}_+$, $\text{Time}(s) \cap I$ is a finite set. Suppose not. Then there is no minimum distance $\delta' > 0$ between any arbitrary pair of event times. Given the $\delta > 0$ and $N : \mathbb{R}_+ \rightarrow \mathbb{N}_0$ of our delta causal process, we let

$$\tau_1 := \inf \left\{ \tau \in \mathbb{R}_+ \mid ([\tau, \tau + \delta) \cap \text{Time}(s)) \text{ infinite} \right\}. \quad (32)$$

$$s_1(\tau, n) := \begin{cases} s(\tau, n), & (\tau, n) \preceq (\tau_1, N(\tau_1)), \\ \perp, & (\tau, n) \succ (\tau_1, N(\tau_1)). \end{cases} \quad (33)$$

Now $\text{Time}(s_1)$ is a finite set. If $\text{Tag}(s_1)$ is a non-Zeno, discrete set, then

$$d(F(s_1), F(s)) = d(F(s_1), s) \preceq \left(\frac{1}{2^{\tau_1 + \delta}}, 1 \right), \quad (34)$$

which implies $\text{Time}(F(s_1)) \cap [0, \tau_1 + \delta]$ is an infinite set. Since F is simple, $\text{Tag}(s_1)$ is a Zeno set. Let

$$\tau_2 := \min \left\{ \tau \in \mathbb{R}_+ \mid \left(\left[(0, 0), \left(\frac{1}{2^\tau}, \frac{1}{2^{N(\tau)}} \right) \right] \cap \text{Tag}(s_1) \right) \text{ finite} \right\}. \quad (35)$$

$$s_2(\tau, n) := \begin{cases} s_1(\tau, n), & (\tau, n) \preceq (\tau_2, N(\tau_2)), \\ \perp, & (\tau, n) \succ (\tau_2, N(\tau_2)). \end{cases} \quad (36)$$

Now $\text{Tag}(s_2)$ is finite, and s_1 has an infinite number of events at time τ_2 . Thus

$$d(F(s_2), F(s_1)) \preceq d(F(s_2), F(s)) \vee d(F(s), F(s_1)) \preceq \left(\frac{1}{2^{\tau_2 + \delta}}, 1 \right). \quad (37)$$

Thus $\{(\tau, n) \in \text{Tag}(F(s_2)) \mid \tau \leq \delta\}$ is infinite, but since F is not simple, we must conclude that $\text{Time}(s) \cap I$ is finite when I is finite. We can similarly show that $\text{Tag}(s) \cap I$ is finite when I is finite, so our fixed-point is a non-Zeno, discrete-event signal.

Whenever we have a non-Zeno discrete-event fixed point solution $s = F(s)$, then if we can approximate the solution arbitrarily closely with only a finite number of events, even if the solution has an infinite number of events. For a Zeno solution, we lose this ability. This theorem gives us a condition under which we can be sure that the solution is non-Zeno. We will call a process $F : S^I \rightarrow S^J$ causal if for all signals $s_1, s_2 \in S^I$, $d(F(s_1), F(s_2)) \leq d(s_1, s_2)$. It is hard, if not impossible, to physically build a non-causal system. We now show when we can compose simple, causal processes:

Proposition 3. *Given a network of simple processes, if in every cycle there is a delta causal process, then the composite system is simple.*

Proof. Since parallel and series composition of simple processes yield simple processes, we need only show that as long as every cycle contains a delta causal process, the composite is simple. We restrict ourselves to processes which map S to S , as extending to the case of multiple inputs and outputs is mechanical. From Proposition 2, a cycle with one delta causal, simple process is simple. Suppose we have n simple, causal processes $F_i : S \rightarrow S$, of which at least one is delta causal. Then for each $i \in \{0, \dots, n-1\}$, the composite $(F_{i \bmod n} \circ F_{(i+1) \bmod n} \circ \dots \circ F_{(i+n) \bmod n})$ is delta causal and simple. Thus any signal between F_k and F_{k+1} must be a non-Zeno discrete-event signal. Finally, every such composition of simple processes is simple.

5 Conclusions

Discrete-event systems offer an attractive model of computation that has proven effective for large-scale concurrent system design (such as digital circuits). Moreover, there is growing interest in the use of such a model of computation for distributed software systems. Practical uses, however, require the semantics to admit simultaneous events. Moreover, if time is dense, then the semantics allows for the possibility of Zeno behaviors.

In this paper, we first broaden the notion of time to support simultaneous events. We then generalize a classical approach to discrete-event semantics that uses fixed point theorems on a metric space to allow for simultaneous events. We first generalize the notion of a metric to define what we call a tetric, which is a function that yields an element of a totally ordered monoid, rather than a non-negative real number as done by a metric. We then generalize the classical Banach fixed point theorem to tetric spaces and apply this generalization to the semantics of discrete-event systems with superdense time. We give conditions for uniqueness in a fixed-point semantics and for avoidance of Zeno conditions that are straightforward generalizations of the corresponding classical notions.

6 Acknowledgments

We would like to acknowledge useful discussions with Aaron Ames (about Zeno phenomena) and Oded Maler (who pointed out previous work on superdense semantics).

This paper describes work that is part of the Ptolemy project, which is supported by the National Science Foundation (NSF award number CCR-00225610), and Chess (the Center for Hybrid and Embedded Software Systems), which receives support from NSF and the following companies: Agilent, General Motors, Hewlett-Packard, Honeywell, Infineon, Samsung, and Toyota.

References

1. Zeigler, B.P., Praehofer, H., Kim, T.G.: Theory of Modeling and Simulation. 2nd edn. Academic Press (2000)
2. Armstrong, J.R., Gray, F.G.: VHDL Design Representation and Synthesis. Second edition edn. Prentice-Hall (2000)
3. Chandy, K.M., Misra, J.: Parallel Program Design: A Foundation. Addison Wesley (1988)
4. Jefferson, D.: Virtual time. *ACM Trans. Programming Languages and Systems* **7** (1985) 404–425 Virtual time here means simulated time, as used in discrete event systems. Describes time warp as an implementation mechanism.
5. Smith, D.A., Kay, A., Raab, A., Reed, D.P.: Croquet: A collaboration system architecture. White paper (2003)
6. Reed, D.P.: Naming and Synchronization in a Decentralized Computer System. Ph.d. thesis, MIT (1978)
7. Lamport, L.: The part-time parliament. *ACM Transactions on Computer Systems* **16** (1998) 133–169
8. Lee, E.A.: Embedded software. In Zelkowitz, M., ed.: *Advances in Computers*. Volume 56. Academic Press (2002)
9. Yates, R.K.: Networks of real-time processes. In Best, E., ed.: *Proc. of the 4th Int. Conf. on Concurrency Theory (CONCUR)*. Volume LNCS 715.. Springer-Verlag (1993)
10. Lee, E.A.: Modeling concurrent real-time processes using discrete events. *Annals of Software Engineering* **7** (1999) 25–45
11. Manna, Z., Pnueli, A.: Verifying hybrid systems. In: *Hybrid Systems*. Volume 736 of *Lecture Notes in Computer Science*. Springer-Verlag (1993) 4–35
12. Maler, O., Manna, Z., Pnueli, A.: From timed to hybrid systems. In: *REX workshop Real-Time: Theory in Practice*. *Lecture Notes in Computer Science* (1992) 447–48
13. Mosterman, P.J.: An overview of hybrid simulation phenomena and their support by simulation packages. In Vaandrager, F.W., van Schuppen, J.H., eds.: *Hybrid Systems: Computation and Control*, Springer-Verlag (1999) 165–177
14. Lee, E., Sangiovanni-Vincentelli, A.: A framework for comparing models of computation. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **17** (1998) 1217–1279
15. Mazurkiewicz, A.: Traces, histories, graphs: Instances of a process monoid. In Chytil, M.P., Koubek, V., eds.: *Proc. Conf. on Mathematical Foundations of Computer Science*. Volume LNCS 176.. Springer-Verlag (1984)
16. Ames, A.D., Sastry, S.: Characterization of zeno behavior in hybrid systems using homological methods. In: To appear at 24th American Control Conference 2005 (ACC05), Portland, OR (2005)
17. Evans, K., Konikoff, M., Madden, J.J., Mathis, R., Whipple, G.: Totally ordered commutative monoids. *Semigroup Forum* **62** (2001) 249 – 278
18. Priess-Crampe, S., Ribenboim, P.: Generalized ultrametric spaces, I. In: *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*. Volume 66. (1996) 55–73
19. Munkres, J.R.: *Topology*. Prentice Hall (2000)
20. Roscoe, B., Reed, G.: Metric spaces as models for real-time concurrency. In: *Proceedings of the Third Workshop on the Mathematical Foundations of Programming Language Semantics*. Volume 298.. Springer LNCS (1988) 331–343