

Modeling Timed Concurrent Systems

Xiaojun Liu¹, Eleftherios Matsikoudis², and Edward A. Lee²

¹ Sun Microsystems, Inc., xiaojun.liu@sun.com

² University of California, Berkeley, [ematsi,eal@eecs.berkeley.edu](mailto:ematsi@eecs.berkeley.edu)

Abstract. Timed concurrent systems are widely used in concurrent and distributed real-time software, modeling of hybrid systems, design of hardware systems (using hardware description languages), discrete-event simulation, and modeling of communication networks. They consist of concurrent components that communicate using timed signals, that is, sets of (semantically) time-stamped events. The denotational semantics of such systems is traditionally formulated in a metric space, wherein causal components are modeled as contracting functions. We show that this formulation excessively restricts the models of time that can be used. In particular, it cannot handle super-dense time, commonly used in hardware description languages and hybrid systems modeling, finite time lines, and time with no origin. Moreover, if we admit continuous-time and mixed signals (essential for hybrid systems modeling) or certain Zeno signals, then causality is no longer equivalent to its formalization in terms of contracting functions. In this paper, we offer an alternative semantic framework using a generalized ultrametric that overcomes these limitations.

1 Introduction

This paper focuses on timed concurrent systems modeling. Timed concurrent systems are collections of concurrent components that communicate by use of timed signals. We will define this formally, but intuitively timed signals are functions of a globally defined time. They consist of either continuously evolving values or discrete events or some combination of the two. Semantically, time is a globally shared concept, and causality is intrinsically bound to chronological ordering.

Timed concurrent systems have a wide range of application. They are used in concurrent and distributed real-time software, modeling of hybrid systems, design of hardware systems (using hardware description languages), discrete-event simulation, and modeling of communication networks.

The importance of precise mathematical models for such systems cannot be overemphasized. In short, they establish canonical denotational definitions of timed programming languages, thereby providing the means for reasoning about the correctness of individual implementations, as well as allowing hidden commonalities of seemingly different timed systems to emerge.

Unfortunately, timed systems are not amenable to standard order-theoretic denotational semantic approaches as they may realize non-monotonic functions

over the sequences of observable actions [25]. Yet interesting results have been obtained by imposing a fixed lower bound on the reaction time of the involved components, effectively precluding Zeno behavior, where an infinite number of actions takes place over a finite interval of time. This has permitted the successful employment of traditional metric-space theory in the construction of well-defined mathematical models for these constrained classes of timed concurrent systems [23, 25, 11, 2, 4, 6].

In this paper, we expose a number of limitations in the traditional metric-space approach that hinder generalization to broader classes of timed concurrent systems. We then proceed to develop the fundamentals of a semantic framework for timed concurrent systems that is more broadly applicable. The underlying assumption is that a timed concurrent system can be modeled as a single system function, and that the behavior of the system corresponds to a fixed point of that function. In practice, to obtain this function, we have to be concerned about composition. That is, given the functions for the interconnected components, we need to be able to compose them to obtain the system function. The techniques given in this paper facilitate such composition.

This paper begins with a brief review of metric spaces, a definition of timed signals, and a review of a metric-space semantics for timed concurrent systems. During this review, we point out several limitations in this traditional approach. We then develop an alternative based on generalized ultrametric spaces, and discuss how it overcomes these limitations.

2 Mathematical Preliminaries

A *metric space* (X, d) is a set X with a *metric distance function* $d: X \times X \rightarrow \mathbb{R}_0$ such that for all $x, y, z \in X$,

1. $d(x, y) = 0$ if and only if $x = y$,
2. $d(x, y) = d(y, x)$, and
3. $d(x, z) \leq d(x, y) + d(y, z)$.

If the metric distance function d also satisfies

4. $d(x, z) \leq \max(d(x, y), d(y, z))$

for all $x, y, z \in X$, then (X, d) is an *ultrametric space* and d an *ultrametric distance function*.

The value $d(x, y)$ quantifies how closely x approximates y . An element $x \in X$ is the *limit* of a sequence $\{x_k\}_{k \in \mathbb{N}}$, where \mathbb{N} is the set of all natural numbers $\{0, 1, 2, \dots\}$, iff for all $\epsilon > 0$, there exists $n \in \mathbb{N}$ such that for all $k \geq n$, $d(x_k, x) < \epsilon$. The sequence is then said to *converge* to x , denoted by $x_k \rightarrow x$. A sequence $\{x_k\}_{k \in \mathbb{N}}$ is *Cauchy* iff for all $\epsilon > 0$, there exists $n \in \mathbb{N}$ such that for all $k, l \geq n$, $d(x_k, x_l) < \epsilon$. A metric space (X, d) is *complete* iff every Cauchy sequence converges to some $x \in X$.

If $B_\delta(x)$ is the set $\{y \in X \mid d(y, x) < \delta\}$, then the collection of such sets $\{B_\delta(x) \mid x \in X, \delta \in \mathbb{R}_+\}$ is a basis of a *topology* on X . This topology is called the *metric topology* induced by d .

Let (X, d) and (X', d') be metric spaces. A function $f: X \rightarrow X'$ is *continuous* iff $x_k \rightarrow x$ implies $f(x_k) \rightarrow f(x)$. It is *contracting* iff for all $x, y \in X$,

$$d'(f(x), f(y)) \leq d(x, y).$$

It is *strictly contracting* iff for all $x, y \in X$,

$$x \neq y \implies d'(f(x), f(y)) < d(x, y).$$

It is a δ -*contraction* iff there exists $\delta \in (0, 1)$ such that for all $x, y \in X$,

$$d'(f(x), f(y)) \leq \delta d(x, y).$$

From the theory of metric spaces, the key result used in programming language semantics is the Banach fixed-point theorem [7].

Theorem 1 (Banach). *Let (X, d) be a complete metric space. If the function $f: (X, d) \rightarrow (X, d)$ is a δ -contraction, then f has a unique fixed point in X , denoted by $\text{fix } f$, and for all $x \in X$, $f^k(x) \rightarrow \text{fix } f$.*

3 Timed Signals

In this paper, we are interested in concurrent components that communicate via *timed signals*. We model these using the tagged-signal model [12], where the communication between two components is represented by a set of events. Formally, let \mathcal{T} be a non-empty set of *tags*, and \mathcal{V} a non-empty set of *values*. An *event* is a pair (t, v) in $\mathcal{T} \times \mathcal{V}$. A *signal* is a set of events that typically represents the sum total of the communication between two components along some communication path. For the systems we are interested in, these sets are very likely infinite. Most applications of the tagged-signal model impose structure on the tag set \mathcal{T} and study the consequences of that structure. For example, \mathcal{T} might represent causality properties, time, or activation orders.

3.1 Models of Time

In general, in the tagged signal model, \mathcal{T} is a partially ordered set. In this paper, \mathcal{T} represents time. Our framework admits several models of time, but in all cases, \mathcal{T} will be totally ordered.

Perhaps the most natural choice for \mathcal{T} is the set of non-negative real numbers \mathbb{R}_0 , reflecting a Newtonian physical view of time. The fact that we include only the non-negative reals implies that our timed concurrent networks have a starting point.

A more interesting model of time is the *super dense time* (SDT) model [19], where $\mathcal{T} = \mathbb{R}_0 \times \mathbb{N}$ equipped with the lexicographic order, that is, for all $(r_1, n_1), (r_2, n_2) \in \mathbb{R}_0 \times \mathbb{N}$,

$$(r_1, n_1) \leq (r_2, n_2) \iff r_1 < r_2 \text{ or } (r_1 = r_2 \ \& \ n_1 \leq n_2).$$

This is a total order. SDT can be similarly defined as $\mathcal{T} = I \times \mathbb{N}$, with I being any interval of real numbers. SDT has been used in studying the semantics of hybrid systems [10, 14, 18]. Its subset $\mathbb{N} \times \mathbb{N}$ is used as the model of time in the hardware description languages Verilog and VHDL. SDT is in a sense strictly richer than \mathbb{R}_0 as a model of time, in that there is no order-embedding from $\mathbb{R}_0 \times \mathbb{N}$ into \mathbb{R}_0 , as may be easily verified.

3.2 Signals

A signal in the tagged-signal model is a set of events, or equivalently, a relation with domain some subset of \mathcal{T} and range some subset of \mathcal{V} . In this paper, we constrain such relations to be single-valued, and thus commit to the following definition.

Definition 1 (Signal). *A set s is a signal if and only if $s \in (\mathcal{T} \rightarrow \mathcal{V})$.*³

We denote the set of all signals with tag set \mathcal{T} and value set \mathcal{V} by $\mathcal{S}(\mathcal{T}, \mathcal{V})$, that is, $\mathcal{S}(\mathcal{T}, \mathcal{V}) = (\mathcal{T} \rightarrow \mathcal{V})$. We adopt common practice in modern set theory and identify a function with its graph. The events of a signal s then are precisely the members of s . And the domain of the signal $\text{dom } s$ is the set of all tags where events of the signal s are present. The signal with no events is simply the empty set \emptyset .

For notational convenience, we will write $s_1(t) \simeq s_2(t)$ iff the signals s_1 and s_2 are either both defined, or both undefined at tag t , and if defined $s_1(t) = s_2(t)$.

The following examples in $\mathcal{S}(\mathbb{R}_0, \mathbb{R})$ are sketched in Fig. 1:

$$\begin{aligned} \text{const}_1 &\stackrel{\text{def}}{=} \{(t, 1) \mid t \in \mathbb{R}_0\}, \\ \text{clock}_1 &\stackrel{\text{def}}{=} \{(k, 1) \mid k \in \mathbb{N}\}, \text{ and} \\ \text{zeno} &\stackrel{\text{def}}{=} \{(1 - 1/2^k, 1) \mid k \in \mathbb{N}\}. \end{aligned}$$

The *zeno* example is particularly interesting. It is really the timed systems version of Zeno's paradox [1], where an infinite number of events may take place before some finite instance of time. Although not physically realizable, such signals may easily crop up in simulation and modeling environments where time is represented as an actual program variable. In Section 4 we discuss some of the related subtleties that in the past have compelled researchers to invariably impose certain conditions on their systems that effectively preclude Zeno behavior [23, 25, 11].

4 The Cantor Metric and its Limitations

The Cantor metric is a distance function that may be defined on streams or sequences [5]. The same metric is called the Baire-distance in [6]. The focus here

³ We denote the set of all functions f with $\text{dom } f \subseteq A$ and $\text{ran } f \subseteq B$ by $(A \rightarrow B)$.




Fig. 1. Examples of timed signals: (a) $const_1$, (b) $clock_1$, (c) $zeno$.

is on the Cantor metric for timed signals, a typical choice in the metric-space approach to the denotational semantics of timed concurrent systems [11, 15, 23].

Under the assumption that $\mathcal{T} = \mathbb{R}_0$, the Cantor metric for timed signals is a function $d_{\text{cantor}}: \mathcal{S}(\mathcal{T}, \mathcal{V}) \times \mathcal{S}(\mathcal{T}, \mathcal{V}) \rightarrow \mathbb{R}_0$ such that for all $s_1, s_2 \in \mathcal{S}(\mathcal{T}, \mathcal{V})$,

$$d_{\text{cantor}}(s_1, s_2) \stackrel{\text{def}}{=} 2^{-\sup \{t \in \mathcal{T} \mid (\forall \tau \leq t)(s_1(\tau) \simeq s_2(\tau))\}}. \quad (1)$$

It is understood that since $\mathcal{T} = \mathbb{R}_0$, $2^{-\sup \emptyset} = 1$, and $2^{-\sup \mathbb{R}_0} = 0$. That is, two signals that differ at their start have distance one, and two signals that are everywhere identical have distance zero. It is easy to show that $(\mathcal{S}(\mathbb{R}_0, \mathcal{V}), d_{\text{cantor}})$ is a complete (ultra)metric space [16].

In the remainder of this section we discuss a number of limitations in the use of the Cantor metric in a semantic framework for timed concurrent systems, thus demonstrating our motivation for turning to the theory of generalized ultrametric spaces. This is not to say that the Cantor metric is the only metric applicable to timed signals. For example, any Hausdorff distance function, such as the Skorohod J_2 metric in [8], may be used to endow the set of timed signals with the structure of a metric space. In the special case of discrete-event signals, even the Baire-distance on sequences, or other ad-hoc distance functions, as in [24], may be viable alternatives. However, all these alternatives seem to more or less suffer from the same limitations. And for timed concurrent systems modeling, the Cantor metric seems, at least to the authors, as the most prominent metric candidate, rendering causal components contracting functions on timed signals, and yielding a convergence process that closely resembles the actual operation of timed systems.

4.1 Convergence in the Cantor Metric Space

A sequence $\{s_k\}_{k \in \mathbb{N}}$ of signals is said to converge to a signal s if for any $\epsilon > 0$, there is an $n \in \mathbb{N}$ such that for all $k > n$,

$$d_{\text{cantor}}(s_k, s) < \epsilon.$$

Such convergence gives us a theory of approximation that enables assigning semantics to timed systems with infinite executions. If partial executions yield signals s_k , and these signals converge to some signal s , then in an operational semantics, s is the semantics of the signal. Of course, we would expect that this s be the same as the signal delivered by our denotational semantics. The main obstacle to achieving this in timed concurrent systems is the potential for Zeno conditions.

Consider, for example, finite approximations s_k to the signal $zeno \in \mathcal{S}(\mathbb{R}_0, \mathbb{R})$. Let $\{s_k\}_{k \in \mathbb{N}}$ be the sequence of signals where s_k is defined by

$$s_k = \{(1 - 1/2^n, 1) \mid n \in \{0, \dots, k\}\}. \quad (2)$$

That is, s_k is the prefix of $zeno$ that contains the first k events. Intuitively, the sequence $\{s_k\}_{k \in \mathbb{N}}$ converges to $zeno$, but it does not converge in the Cantor metric. It is easy to see that for all k ,

$$d_{\text{cantor}}(s_k, zeno) > 1/2.$$

It is also easy to see that for any k and k' such that $k \neq k'$,

$$d_{\text{cantor}}(s_k, s_{k'}) > 1/2.$$

So this sequence is not Cauchy. Although we have no mathematical contradiction, the Cantor metric has failed to provide us with a framework where we can consider the sequence of signals $\{s_k\}_{k \in \mathbb{N}}$ to be a sequence of finite approximations to the signal $zeno$.

On a related note, systems that give rise to Zeno signals cannot be modeled as δ -contractions in the Cantor metric space [25]. It is thus impossible to utilize the Banach fixed-point theorem for reasoning about the behavior of such systems. The generalized ultrametric space that we construct in Section 5 allows us to use a variant of the Banach fixed-point theorem that is less restrictive in terms of how contractive the system functions are.

4.2 Causal Components versus Contracting Functions

Causality is the relationship between causes and effects. If a component in a timed concurrent system models a physical or computational process, the time of an effect cannot be earlier than the time of the corresponding cause. It is common to relate this intuitive notion of causality with contracting functions in the Cantor metric space [23, 11]. However, this does not quite work with arbitrary timed signals.

Consider a deterministic component that accepts as input a timed signal and produces as output a timed signal. We can model such a component as a function $F: \mathcal{S}(\mathcal{T}, \mathcal{V}) \rightarrow \mathcal{S}(\mathcal{T}, \mathcal{V})$. If $\mathcal{T} = \mathbb{R}_0$, then the domain and range of this function are complete metric spaces under the Cantor metric. If F is contracting, then from the definition of the Cantor metric, we can see that if two possible inputs are identical up to some time, then the corresponding outputs are identical up

to that same time. This motivates some authors to model causal components as contracting input-output functions.

Consider the following example. Let $u_1, u_2 \in \mathcal{S}(\mathbb{R}_0, \mathbb{R})$ be such that for all $t \in \mathbb{R}_0$,

$$\begin{aligned} u_1(t) &= 0, \\ u_2(t) &= \begin{cases} 0 & \text{if } t \in [0, 1], \\ 1 & \text{if } t > 1. \end{cases} \end{aligned}$$

Suppose a component with one input and one output is modeled by the function $F: \mathcal{S}(\mathbb{R}_0, \mathbb{R}) \rightarrow \mathcal{S}(\mathbb{R}_0, \mathbb{R})$, where for any input $s \in \mathcal{S}(\mathbb{R}_0, \mathbb{R})$ and any $t \in \mathbb{R}_0$,

$$F(s)(t) = \begin{cases} \lim_{r \rightarrow t^+} s(r) & \text{if the right limit exists,} \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

The function F is contracting. However, $F(u_1) = u_1$, and

$$F(u_2)(t) = \begin{cases} 0 & \text{if } t \in [0, 1), \\ 1 & \text{if } t \geq 1. \end{cases}$$

The input signals u_1 and u_2 are equal over $[0, 1]$, whereas the output signals $F(u_1)$ and $F(u_2)$ are equal only over $[0, 1)$. Consequently, the component is not causal.

The fact that contracting functions in the Cantor metric space are not necessarily causal is certainly disturbing. We can avoid this discrepancy by restricting our attention to a certain class of signals and components. The following definition comes from [11].

Definition 2 (DE Signals). *A timed signal $s \in \mathcal{S}(\mathcal{T}, \mathcal{V})$ is a discrete-event (DE) signal if and only if there exists an order-embedding from $\text{dom } s$ into \mathbb{N} .*

It is not hard to show that when equipped with the Cantor metric, the set of all DE signals yields a complete ultrametric space [16]. If we consider only DE signals and components that operate only on DE input signals yielding DE signals as outputs, then our informal notion of causality actually coincides with its classic formalization in terms of contracting functions, as the reader is invited to verify.

In Section 5 we define a generalized ultrametric that enables us to accurately formalize our informal notion of causality in terms of contracting functions in the respective generalized ultrametric space, and thus have a unified framework of causal systems that arbitrarily combine discrete (e.g. software) components with continuous (e.g. physical) processes.

4.3 Tag-Set Choices

When using the Cantor metric, the choice of tag set has a profound impact. We would like to be able to use any totally ordered set to model time but we cannot do this with the Cantor metric.




Fig. 2. Timed signals with tag set \mathbb{R} .

For example, we would like to be able to use super dense time as our model of time, an indispensable choice when studying hybrid systems. However, the fact that there is no order-embedding from $\mathbb{R}_0 \times \mathbb{N}$ into \mathbb{R}_0 makes this impossible.

Similarly, it is tempting to restrict the tag set to the interval $[0, 1)$ so as to have the signal *zeno* extend over the whole time line. By the definition of Cantor metric in (1), for all signals $s_1, s_2 \in \mathcal{S}([0, 1), \mathcal{V})$,

$$s_1 \neq s_2 \implies d_{\text{cantor}}(s_1, s_2) > \frac{1}{2}.$$

Hence, the metric topology induced by d_{cantor} on $\mathcal{S}([0, 1), \mathcal{V})$ is the *discrete topology*. In other words, the Cantor metric does not provide any useful structure on $\mathcal{S}([0, 1), \mathcal{V})$.

Another interesting case is to take \mathbb{R} as the tag set. The Cantor distance between two signals in $\mathcal{S}(\mathbb{R}, \mathcal{V})$ may be infinity. For example, let

$$\begin{aligned} s_{\text{clk}} &\stackrel{\text{def}}{=} \{(k, 1) \mid k \in \mathbb{Z}\}, \\ s_{\text{alt}} &\stackrel{\text{def}}{=} \{(2k, 1) \mid k \in \mathbb{Z}\}, \end{aligned}$$

where \mathbb{Z} is the set of all integers. These signals, illustrated in Fig. 2, have an empty common prefix, and hence

$$\{t \in \mathbb{R} \mid (\forall \tau \leq t)(s_1(\tau) \simeq s_2(\tau))\} = \emptyset.$$

With \mathbb{R} as the tag set, it is understood that $2^{-\sup \emptyset} = \infty$, so the Cantor distance between these signals is infinite.

$(\mathcal{S}(\mathbb{R}, \mathcal{V}), d_{\text{cantor}})$ is not a metric space⁴, because a metric space is defined as a function into \mathbb{R}_0 . We have a function onto $\mathbb{R}_0 \cup \{\infty\}$. Consequently, we cannot immediately use the Banach fixed point theorem, and in fact we can show that it does not hold in this space.

Consider a component Delay_d that shifts every event in its input signal by $d \in \mathbb{R}_0$ into the future. That is, $\text{Delay}_d: \mathcal{S}(\mathbb{R}, \mathcal{V}) \rightarrow \mathcal{S}(\mathbb{R}, \mathcal{V})$ such that for any $s \in \mathcal{S}(\mathbb{R}, \mathcal{V})$ and any $t \in \mathbb{R}$,

$$\text{Delay}_d(s)(t) \stackrel{\text{def}}{\simeq} \begin{cases} s(t-d) & \text{if } t-d \in \text{dom } s, \\ \text{undefined} & \text{otherwise.} \end{cases} \quad (4)$$

⁴ It is erroneously assumed to be one in [20].

It is easy to show that the function $Delay_d$ is a δ -contraction for any $d > 0$. In $\mathcal{S}(\mathbb{R}, \mathcal{V})$, this function has more than one fixed point. In particular, note that with $d = 1$, both \emptyset and s_{clk} are fixed points. Although the Banach fixed point theorem assures us that a δ -contraction has only one fixed point, there is no contradiction because $\mathcal{S}(\mathbb{R}, \mathcal{V})$ is not a metric space.

$(\mathcal{S}(\mathbb{R}, \mathcal{V}), d_{\text{cantor}})$ is an *extended metric space* [3]. This extended metric space can be divided into a set of complete metric spaces. Let relation R on $\mathcal{S}(\mathbb{R}, \mathcal{V})$ be defined by

$$(s_1, s_2) \in R \iff d_{\text{cantor}}(s_1, s_2) < \infty.$$

An equivalent definition of relation R is that s_1 and s_2 have a non-empty common prefix—going back in time, s_1 and s_2 are eventually the same.

It is straightforward to show that R is an *equivalence relation*. For any signal $s \in \mathcal{S}(\mathbb{R}, \mathcal{V})$, let

$$E_s \stackrel{\text{def}}{=} \{s' \in \mathcal{S}(\mathbb{R}, \mathcal{V}) \mid d_{\text{cantor}}(s', s) < \infty\}$$

be the equivalence class containing s . (E_s, d_{cantor}) is a complete (ultra)metric space. The extended metric space $(\mathcal{S}(\mathbb{R}, \mathcal{V}), d_{\text{cantor}})$ thus contains an infinite number of complete (ultra)metric spaces.

Consider $Delay_1$ as a function of the form $Delay_1: E_{s_{\text{clk}}} \rightarrow E_{s_{\text{clk}}}$. This is legitimate, because given any $s \in E_{s_{\text{clk}}}$, $Delay_1(s) \in E_{s_{\text{clk}}}$. This function has a unique fixed point in $E_{s_{\text{clk}}}$, namely s_{clk} .

Consider $Delay_1$ as a function of the form $Delay_1: E_{\emptyset} \rightarrow E_{\emptyset}$. This is legitimate, because given any $s \in E_{\emptyset}$, $Delay_1(s) \in E_{\emptyset}$. This function has a unique fixed point in E_{\emptyset} , namely \emptyset .

Notice that it is not possible to consider $Delay_1$ as a function of the form $Delay_1: E_{s_{\text{alt}}} \rightarrow E_{s_{\text{alt}}}$, because in fact $Delay_1(s_{\text{alt}}) \notin E_{s_{\text{alt}}}$.

To borrow an analogy from cosmology, the equivalence classes of R partition $\mathcal{S}(\mathbb{R}, \mathcal{V})$ into parallel universes, and all signals in an equivalence class originate from the same “Big Bang.” If a (contracting) component can take us from one universe to another, then putting it in a feedback loop can yield multiple parallel behaviors.

We next develop a framework, based on the theory of generalized ultrametric spaces, that is far more admissive with respect to the tag-set choices. The generalized ultrametric space that we construct renders $Delay_d$ not strictly contracting.

5 Generalized Ultrametrics and their Application

We have seen that the traditional approach to metric space semantics has a number of limitations for timed concurrent systems. Several restrictions have to be applied in order for it to be useful. It effectively rules out models of time that are used in practice (such as super dense time) or are interesting in theory (such as \mathbb{R} , which has no least time, or bounded intervals of \mathbb{R}). Moreover, even when time is modeled using \mathbb{R}_0 , Zeno conditions can render the Banach fixed-point

theorem irrelevant. Finally, the equivalence between our informal understanding of causality and its formalization in terms of contracting functions breaks down when continuous-time and mixed signals (essential for hybrid systems modeling) or certain Zeno signals are allowed.

In this section we resort to the theory of generalized ultrametric spaces and define a generalize ultrametric distance function on timed signals that eliminates the abovementioned problems.

5.1 Generalized Ultrametric Spaces

While the codomain of a metric distance function is required to be the set of all non-negative real numbers \mathbb{R}_0 , the codomain of a generalized ultrametric distance function [21] may be chosen as any partially ordered set with a minimum element.

Definition 3 (Generalized Ultrametric Space). *Let X be a set, Γ a partially ordered set with a minimum element 0_Γ . Then (X, d, Γ) is a generalized ultrametric space iff $d: X \times X \rightarrow \Gamma$ is a function such that for all $x, y, z \in X$ and $\gamma \in \Gamma$,*

1. $d(x, y) = 0_\Gamma$ if and only if $x = y$,
2. $d(x, y) = d(y, x)$, and
3. if $d(x, y) \leq \gamma$ and $d(y, z) \leq \gamma$, then $d(x, z) \leq \gamma$.

A function $d: X \times X \rightarrow \Gamma$ that adheres to the above definition is called a *generalized ultrametric distance function*.

If (X, d, Γ) is a generalized ultrametric space, then for any $\gamma \in \Gamma \setminus \{0_\Gamma\}$ and $a \in X$, the set

$$B_\gamma(a) = \{x \in X \mid d(x, a) \leq_\Gamma \gamma\}$$

is called the *ball* with *center* a and *radius* γ . It is easy to verify that for all $x, y \in X$ and $\alpha, \beta \in \Gamma$, if $0_\Gamma < \alpha \leq \beta$ and $x \in B_\beta(y)$, then $B_\alpha(x) \subseteq B_\beta(y)$; every point in a ball is also its center.

The usual notion of completeness for metric spaces extends naturally to the case of generalized ultrametric spaces. However, it is the stronger notion of spherical completeness that most interesting results in the theory of generalized ultrametric spaces rely on.

Definition 4 (Spherical Completeness). *A generalized ultrametric space is spherically complete iff every chain of balls (ordered by inclusion) has a non-empty intersection.*

Let (X, d, Γ) and (X', d', Γ) be generalized ultrametric spaces. A function $f: X \rightarrow X'$ is *contracting* iff for all $x, y \in X$,

$$d'(f(x), f(y)) \leq d(x, y).$$

It is *strictly contracting* iff for all $x, y \in X$,

$$x \neq y \implies d'(f(x), f(y)) < d(x, y).$$

The above definitions are evidently identical to those introduced in the case of metric spaces. The notion of δ -contraction, however, has no immediate counterpart in the context of generalized ultrametric spaces. There is nevertheless an analogue, in some sense, to the Banach fixed-point theorem, namely the Priess-Crampe & Ribenboim fixed-point theorem. There are several variants of this theorem [21, 22]. The following is from section 5.2 of [21].

Theorem 2 (Priess-Crampe & Ribenboim). *Let (X, d, Γ) is a spherically complete generalized ultrametric space. If the function $f: X \rightarrow X$ is strictly contracting, then f has a unique fixed point.*

We note that the proof of this theorem relies on the Axiom of Choice and is thus inherently non-constructive.

5.2 Generalized Ultrametrics on Timed Signals

Let s_1 and s_2 be signals in $\mathcal{S}(\mathbb{R}_0, \mathcal{V})$. The Cantor metric in essence maps the set $\{t \in \mathbb{R}_0 \mid (\forall \tau \leq t)(s(\tau) \simeq s_2(\tau))\}$, namely the largest down set⁵ of \mathbb{R}_0 on which the signals s_1 and s_2 coincide, to an element of \mathbb{R}_0 such that for all $s'_1, s'_2 \in \mathcal{S}(\mathbb{R}_0, \mathcal{V})$,

$$\begin{aligned} \{t \in \mathbb{R}_0 \mid (\forall \tau \leq t)(s_1(\tau) \simeq s_2(\tau))\} &\supseteq \{t \in \mathbb{R}_0 \mid (\forall \tau \leq t)(s'_1(\tau) \simeq s'_2(\tau))\} \\ &\implies \\ d_{\text{cantor}}(s_1, s_2) &\leq d_{\text{cantor}}(s'_1, s'_2). \end{aligned}$$

The inverse implication is not generally true, which is the reason that a contracting process is not necessarily causal.

Let $\mathcal{D}(\mathbb{R}_0)$ denote the set of all down sets of \mathbb{R}_0 . We can define a totally ordered set $(\mathcal{D}(\mathbb{R}_0), \supseteq)$ whose order relation is reverse set containment \supseteq . It is easy to show that there is no order-embedding from $(\mathcal{D}(\mathbb{R}_0), \supseteq)$ into \mathbb{R}_0 . Hence, it is impossible to define a metric d on $\mathcal{S}(\mathbb{R}_0, \mathcal{V})$ such that for all $s_1, s_2, s'_1, s'_2 \in \mathcal{S}(\mathbb{R}_0, \mathcal{V})$,

$$\begin{aligned} \{t \in \mathbb{R}_0 \mid (\forall \tau \leq t)(s_1(\tau) \simeq s_2(\tau))\} &\supseteq \{t \in \mathbb{R}_0 \mid (\forall \tau \leq t)(s'_1(\tau) \simeq s'_2(\tau))\} \\ &\iff \\ d_{\text{cantor}}(s_1, s_2) &\leq d_{\text{cantor}}(s'_1, s'_2). \end{aligned}$$

However, we can easily define a generalized ultrametric that satisfies this equivalence.

For any tag set \mathcal{T} , let the set of generalized ultrametric distances, $\Gamma_{\mathcal{T}}$, be the partially ordered set

$$\Gamma_{\mathcal{T}} \stackrel{\text{def}}{=} (\mathcal{D}(\mathcal{T}), \supseteq).$$

If we use the notation \leq_{Γ} for the order relation of $\Gamma_{\mathcal{T}}$, then for any two down sets $D, D' \in \mathcal{D}(\mathcal{T})$, $D \leq_{\Gamma_{\mathcal{T}}} D'$ if and only if $D \supseteq D'$. That is, the order is reverse

⁵ A subset D of a partially ordered set (P, \leq) is a *down set* of P iff for all $p, p' \in P$, if $p' \in D$ and $p \leq p'$, then $p \in D$.

set containment. \mathcal{T} is the minimum element of $\Gamma_{\mathcal{T}}$. That is, for all $D \in \mathcal{D}(\mathcal{T})$, $\mathcal{T} \leq_{\Gamma_{\mathcal{T}}} D$ because $\mathcal{T} \supseteq D$. Similarly, the maximum element is \emptyset , the empty set. It is easy to show that for any tag set \mathcal{T} , the partially ordered set $(\mathcal{D}(\mathcal{T}), \supseteq)$ is a complete lattice.

For any tag set \mathcal{T} and any value set \mathcal{V} , we define the function $d_{ds}: \mathcal{S}(\mathcal{T}, \mathcal{V}) \times \mathcal{S}(\mathcal{T}, \mathcal{V}) \rightarrow \Gamma_{\mathcal{T}}$ such that for all $s_1, s_2 \in \mathcal{S}(\mathcal{T}, \mathcal{V})$,

$$d_{ds}(s_1, s_2) \stackrel{\text{def}}{=} \{t \in \mathcal{T} \mid (\forall \tau \leq t)(s_1(\tau) \simeq s_2(\tau))\}.$$

The following lemma establishes that d_{ds} is in fact a generalized ultrametric on timed signals [17].

Lemma 1 (Generalized Ultrametric on Timed Signals). *For any tag set \mathcal{T} and any value set \mathcal{V} , d_{ds} is a generalized ultrametric distance function on $\mathcal{S}(\mathcal{T}, \mathcal{V})$.*

Stated differently, Lemma 1 shows that for any tag set \mathcal{T} and any value set \mathcal{V} , $(\mathcal{S}(\mathcal{T}, \mathcal{V}), d_{ds}, \Gamma_{\mathcal{T}})$ is a generalized ultrametric space. In particular, we are no longer restricted to $\mathcal{T} = \mathbb{R}_0$. We can choose any totally ordered set to model time, including super dense time, time lines with no origin, and bounded intervals.

We remark here that if we choose \mathcal{T} to be the interval of reals $[0, 1)$, then the sequence $\{s_k\}_{k \in \mathbb{N}}$ as specified in (2) actually converges to *zeno* in the respective generalized ultrametric space, whereas it failed to converge in the Cantor metric space.

On a similar note, observe that while the function $Delay_d \in \mathcal{S}(\mathbb{R}_0, \mathcal{V})$ as defined in (4) is a δ -contraction with respect to the Cantor metric, it is not strictly contracting in the respective generalized ultrametric space (consider any two $s_1, s_2 \in \mathcal{S}(\mathbb{R}, \mathcal{V})$ such that for all $t \in \mathbb{R}$, $s_1(t) \not\approx s_2(t)$). Hence, and in accordance with intuition, we cannot apply the Priess-Crampe & Ribenboim fixed-point theorem to establish the existence of a unique fixed point for the function $Delay_d$.

Finally, notice that the function F defined in (3) is not contracting in the respective generalized ultrametric space. It can therefore no longer serve as a counterexample to the equivalence between the informal notion of causality and its formalization in terms of contracting functions. In fact, it is no longer possible to find such a counterexample. Contracting functions with respect to generalized ultrametrics on timed signals accurately capture the chronological precedence relationship between causes and effects. The following formal definition of causality is essentially equivalent to the respective definitions in [13] and [20].

Definition 5 (Causal Function). *A function $f: \mathcal{S}(\mathcal{T}, \mathcal{V}) \rightarrow \mathcal{S}(\mathcal{T}, \mathcal{V})$ is causal iff it is contracting in the generalized ultrametric space $(\mathcal{S}(\mathcal{T}, \mathcal{V}), d_{ds}, \Gamma_{\mathcal{T}})$.*

Causal functions represent system components that are non-anticipative, in the sense that the output does not anticipate future events of the input. But non-anticipative components may still react instantaneously to input stimuli. The concept of strict causality is thus introduced in order to further assert, when appropriate, the impossibility of instantaneous reaction.

Definition 6 (Strictly Causal Function). *A function $f: \mathcal{S}(\mathcal{T}, \mathcal{V}) \rightarrow \mathcal{S}(\mathcal{T}, \mathcal{V})$ is strictly causal iff it is strictly contracting in the generalized ultrametric space $(\mathcal{S}(\mathcal{T}, \mathcal{V}), d_{ds}, \Gamma_{\mathcal{T}})$.*

Again, the above formal definition of strict causality is equivalent to the respective definition in [20]. We remark here that δ -contractions in the Cantor metric space are strictly contracting functions in the respective generalized ultrametric space. The converse does not hold in general.

The next lemma ensures the applicability of the Priess-Crampe & Ribenboim fixed-point theorem in the context of timed concurrent systems [17].

Lemma 2. *For any tag set \mathcal{T} and any value set \mathcal{V} , the generalized ultrametric space $(\mathcal{S}(\mathcal{T}, \mathcal{V}), d_{ds}, \Gamma_{\mathcal{T}})$ is spherically complete.*

As a simple demonstration, we can immediately apply the Priess-Crampe & Ribenboim fixed-point theorem to establish the following result, first obtained by Naundorf in [20].

Theorem 3. *For any tag set \mathcal{T} and any value set \mathcal{V} , if $f: \mathcal{S}(\mathcal{T}, \mathcal{V}) \rightarrow \mathcal{S}(\mathcal{T}, \mathcal{V})$ is a strictly causal function, then f has a unique fixed point in $\mathcal{S}(\mathcal{T}, \mathcal{V})$.*

We remark that the above results can be specialized to certain classes of timed signals, including discrete-event signals [16].

The relative advantage of the approach taken here over that in [20] relates to the general formulation of the problem. The development of a semantic framework based on generalized ultrametric spaces makes it possible to apply off-the-shelf results from the theory of generalized ultrametric spaces, and share relevant findings with seemingly irrelevant research communities such as the programming logic community [9].

We conclude with the observation that the proof of Theorem 3 here (a trivial application of the Priess-Crampe & Ribenboim fixed-point theorem), as well as in [20], is non-constructive. At the moment, we can only guarantee the existence of a unique fixed point for strictly causal functions. This certainly limits our ability to reason about the behavior of strictly causal systems, a rather broad and interesting class of timed systems. It may nevertheless prove possible to construct a denotational semantics that accurately reflects the actual execution of such systems. In any case, a constructive proof of Theorem 3 is of both practical and theoretical interest, and is the subject of future work.

6 Conclusions

Timed concurrent systems are aggregations of components that communicate by use of timed signals. Such systems have been traditionally modeled in a semantic framework that leverages the theory of metric spaces, representing causal components as contracting functions. This approach has a number of limitations. In particular, it rules out models of time that are used in practice (such as super dense time) or are interesting in theory (such as finite time lines

and time lines with no origin). Moreover, even when more conventional models of time are used, Zeno conditions can render key results from the theory of metric spaces (such as the Banach fixed-point theorem) irrelevant. Finally, the equivalence between the informal notion of causality and its formalization in terms of contracting functions breaks down when continuous-time and mixed signals or certain Zeno signals are allowed.

With these considerations in mind, we have introduced an alternative semantic framework for timed concurrent systems that relies on the theory of generalized ultrametric spaces. We defined an appropriate generalized ultrametric on timed signals that eliminates the aforementioned limitations, yielding a spherically complete generalized ultrametric space under any model of time. We also presented an elegant and formal definition of causality that is exactly equivalent to its informal counterpart. The resultant mathematical structure allows us to apply off-the-shelf results from the theory of generalized ultrametric spaces to reason about the behavior of timed concurrent systems. And as evidence for this thesis, we applied the Priess-Crampe & Ribenboim fixed-point theorem to trivially establish the fact that strictly causal functions have unique fixed points.

7 Acknowledgements

Thanks to Adam Cataldo for many helpful suggestions. 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, the State of California Micro Program, and the following companies: Agilent, DGIST, General Motors, Hewlett Packard, Infineon, Microsoft, and Toyota.

References

1. M. Abadi and L. Lamport. An old-fashioned recipe for real time. *ACM Trans. Program. Lang. Syst.*, 16(5):1543–1571, 1994.
2. A. Arnold and M. Nivat. Metric interpretations of infinite trees and semantics of non deterministic recursive programs. *Fundamenta Informaticae*, 11(2):181–205, 1980.
3. N. Aronszajn and P. Panitchpakdi. Extension of uniformly continuous transformations and hyperconvex metric spaces. *Pacific Journal of Mathematics*, 6(3):405–439, 1956.
4. C. Baier and M. E. Majster-Cederbaum. Denotational semantics in the cpo and metric approach. *Theoretical Computer Science*, 135(2):171–220, 1994.
5. C. S. Calude, S. Marcus, and L. Staiger. A topological characterization of random sequences. *Information Processing Letters*, 88(5):245–250, 2003.
6. J. W. de Bakker and E. P. de Vink. Denotational models for programming languages: Applications of Banach’s fixed point theorem. *Topology and its Applications*, 85:35–52, 1998.
7. A. Granas and J. Dugundji. *Fixed Point Theory*. Springer-Verlag, 2003.

8. V. Gupta, R. Jagadeesan, and P. Panangaden. Approximate reasoning for real-time probabilistic processes. In *Proceedings of the First International Conference on the Quantitative Evaluation of Systems (QEST'04)*, pages 304–313, Sep. 2004.
9. P. Hitzler and A. K. Seda. Generalized metrics and uniquely determined logic programs. *Theoretical Computer Science*, 305(1-3):187–219, 2003.
10. A. Kapur. *Interval and Point-Based Approaches to Hybrid Systems Verification*. Ph.d., Stanford University, 1997. Uses super dense time (super-dense, superdense).
11. E. A. Lee. Modeling concurrent real-time processes using discrete events. *Annals of Software Engineering*, 7:25–45, 1999.
12. E. A. Lee and A. Sangiovanni-Vincentelli. A framework for comparing models of computation. *IEEE Transactions on CAD*, 17(12), 1998.
13. E. A. Lee and P. Varaiya. *Structure and Interpretation of Signals and Systems*. Addison Wesley, 2003.
14. E. A. Lee and H. Zheng. Operational semantics of hybrid systems. In M. Morari and L. Thiele, editors, *Hybrid Systems: Computation and Control (HSCC)*, volume LNCS 3414, pages pp. 25–53, Zurich, Switzerland, 2005. Springer-Verlag.
15. J. Liu and E. A. Lee. On the causality of mixed-signal and hybrid models. In *6th International Workshop on Hybrid Systems: Computation and Control (HSCC '03)*, Prague, Czech Republic, 2003.
16. X. Liu. Semantic foundation of the tagged signal model. Phd thesis, EECS Department, University of California, December 20 2005.
17. X. Liu, E. Matsikoudis, and E. A. Lee. Modeling timed concurrent systems using generalized ultrametrics. Technical Report UCB/EECS-2006-45, EECS Department, University of California, Berkeley, May 1 2006.
18. O. Maler, Z. Manna, and A. Pnueli. From timed to hybrid systems. In *Real-Time: Theory and Practice, REX Workshop*, pages 447–484. Springer-Verlag, 1992.
19. Z. Manna and A. Pnueli. Verifying hybrid systems. *Hybrid Systems*, pages 4–35, 1992.
20. H. Naundorf. Strictly causal functions have a unique fixed point. *Theoretical Computer Science*, 238(1-2):483–488, 2000.
21. S. Priess-Crampe and P. Ribenboim. Logic programming and ultrametric spaces. *Rendiconti di Matematica, Serie VII*, 19:155–176, 1999.
22. S. Priess-Crampe and P. Ribenboim. Fixed point and attractor theorems for ultrametric spaces. *Forum Mathematicum*, 12:53–64, 2000.
23. G. M. Reed and A. W. Roscoe. Metric spaces as models for real-time concurrency. In *3rd Workshop on Mathematical Foundations of Programming Language Semantics*, pages 331–343, London, UK, 1988.
24. F. van Breugel. Comparative semantics for a real-time programming language with integration. In *TAPSOFT '91*, pages 397–411, New York, NY, USA, 1991. Springer-Verlag New York, Inc.
25. R. K. Yates. Networks of real-time processes. In E. Best, editor, *Proc. of the 4th Int. Conf. on Concurrency Theory (CONCUR)*, volume LNCS 715. Springer-Verlag, 1993.