On Fixed Points of Strictly Causal Functions*

Eleftherios Matsikoudis and Edward A. Lee

University of California, Berkeley

Abstract. We ask whether strictly causal components form well defined systems when arranged in feedback configurations. The standard interpretation for such configurations induces a fixed-point constraint on the function modelling the component involved. We define strictly causal functions formally, and show that the corresponding fixed-point problem does not always have a well defined solution. We examine the relationship between these functions and the functions that are strictly contracting with respect to a generalized distance function on signals, and argue that these strictly contracting functions are actually the functions that one ought to be interested in. We prove a constructive fixed-point theorem for these functions, and introduce a corresponding induction principle.

1 Introduction

This work is part of a larger effort aimed at the construction of well defined mathematical models that will inform the design of programming languages and model-based design tools for timed systems. We use the term "timed" rather liberally here to refer to any system that will determinately order its events relative to some physical or logical clock. But our emphasis is on timed computation, with examples ranging from concurrent and distributed real-time software to hardware design, and from discrete-event simulation to continuous-time and hybrid modelling, spanning the entire development process of what we would nowadays refer to as cyber-physical systems. Our hope is that our work will lend insight into the design and application of the many languages and tools that have and will increasingly come into use for the design, simulation, and analysis of such systems. Existing languages and tools to which this work applies, to varying degrees, include hardware description languages such as VHDL and SystemC, modeling and simulation tools such as Simulink and LabVIEW, network simulation tools such as ns-2/ns-3 and OPNET, and general-purpose simulation formalisms such as DEVS (e.g., see [1]).

* This work was supported in part by the Center for Hybrid and Embedded Software Systems (CHESS) at UC Berkeley, which receives support from the National Science Foundation (NSF awards #0720882 (CSR-EHS: PRET), #0931843 (CPS: Large: ActionWebs), and #1035672 (CPS: Medium: Ptides)), the Naval Research Laboratory (NRL #N0013-12-1-G015), and the following companies: Bosch, National Instruments, and Toyota.

V. Braberman and L. Fribourg (Eds.): FORMATS 2013, LNCS 8053, pp. 183–197, 2013.

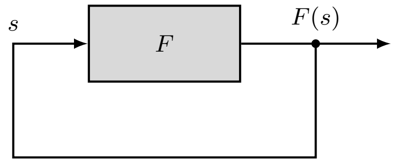

Fig. 1. Block-diagram of a functional component F in feedback. The input signal s and the output signal $F(s)$ are but the same signal; that is, $s = F(s)$.

Considering the breadth of our informal definition for timed systems, we cannot hope for a comprehensive formalism or syntax for such systems at a granularity finer than that of a network of components. We will thus ignore any internal structure or state, and think of any particular component as an opaque flow transformer. Formally, we will model such components as functions, and use a suitably generalized concept of signal as flow. This point of view is consistent with the one presented by most of the languages and tools mentioned above.

The greatest challenge in the construction of such a model is, by and large, the interpretation of feedback. Feedback is an extremely useful control mechanism, present in all but the most trivial systems. But it makes systems self-referential, with one signal depending on another, and vice versa (see Fig. 1). Mathematically, this notion of self-reference manifests itself in the form of a fixed-point problem, as illustrated by the simple block-diagram of Fig. 1: the input signal s and the output signal $F(s)$ are but the same signal transmitted over the feedback wire of the system; unless F has a fixed point, the system has no model; unless F has a unique or otherwise canonical fixed point, the model is not uniquely determined; unless we can construct the unique or otherwise canonical fixed point of F, we cannot know what the model is.

In this work, we consider the fixed-point problem of strictly causal functions, namely functions modelling components whose output at any time depends only on past values of the input. After a careful, precise formalization of this folklore, yet universally accepted, definition, we show that the property of strict causality is by itself too weak to accommodate a uniform fixed-point theory. We then consider functions that are strictly contracting with respect to a suitably defined generalized distance function on signals, and study their relationship to strictly causal ones, providing strong evidence that the former are actually the functions that one ought to be interested in. Finally, we consider the fixed-point problem of these functions, which is not amenable to classical methods (see [2, thm. A.2 and thm. A.4]), and prove a constructive fixed-point theorem, what has resisted previous efforts (e.g., see [3], [4]). We also introduce a corresponding induction principle, and discuss its relationship to the standard ones afforded by the fixed-point theories of order-preserving functions and contraction mappings.

For lack of space, we omit all proofs; they can be found in [2]. We also assume a certain level of familiarity with the theory of generalized ultrametric spaces (e.g., see [5]) and order theory (e.g., see [6]).

2 Signals

We postulate a non-empty set T of *tags*, and a total order relation $\preceq$ on T. We use T to represent our time domain. The order relation $\preceq$ is meant to play the role of a chronological precedence relation, and therefore, it is reasonable to require that $\preceq$ be a total order. We note, however, that such a requirement is often unnecessary, and in fact, all results of Sect. 5 remain valid when $\langle \mathrm{T}, \preceq \rangle$ is only partially ordered.

We would like to define signals as functions over an independent variable ranging over T. But being primarily concerned with computational systems, we should expect our definition to accommodate the representation of variations that may be undefined for some instances or even periods of time. In fact, we think of such instances and periods of time as part of the variational information. Such considerations lead directly to the concept of partial function.

We postulate a non-empty set V of *values*.

Definition 1. *An* event *is an ordered pair* $\langle \tau, v \rangle \in \mathrm{T} \times \mathrm{V}$.

Definition 2. *A* signal *is a single-valued set of events.*

We write S for the set of all signals.

Our concept of signal is based on [7]. But here, unlike in [7], we restrict signals to be single-valued.

Notice that the empty set is vacuously single-valued, and hence, a signal, which we call the *empty signal*.

We adopt common practice in modern set theory and identify a function with its graph. A signal is then a function with domain some subset of T, and range some subset of V, or in other words, a partial function from T to V.

For every $s_1, s_2 \in \mathrm{S}$ and $\tau \in \mathrm{T}$, we write $s_1(\tau) \simeq s_2(\tau)$ if and only if one of the following is true:

1. $\tau \notin \mathsf{dom}\, s_1$ and $\tau \notin \mathsf{dom}\, s_2$;
2. $\tau \in \mathsf{dom}\, s_1$, $\tau \in \mathsf{dom}\, s_2$, and $s_1(\tau) = s_2(\tau)$.

There is a natural, if abstract, notion of distance between any two signals, corresponding to the largest segment of time closed under time precedence, and over which the two signals agree; the larger the segment, the closer the two signals. Under certain conditions, this can be couched in the language of metric spaces (e.g., see [7], [8], [9]). All one needs is a map from such segments of time to non-negative real numbers. But this step of indirection excessively restricts the kind of ordered sets that one can use as models of time (see [4]), and can be avoided as long as one is willing to think about the notion of distance in more abstract terms, and use the language of generalized ultrametric spaces instead (see [5]).

We write d for a map from $\mathrm{S} \times \mathrm{S}$ to $\mathscr{L}\langle \mathrm{T}, \preceq \rangle$ such that for every $s_1, s_2 \in \mathrm{S}$,[1]

$$\mathrm{d}(s_1, s_2) = \{\tau \mid \tau \in \mathrm{T}, \text{ and for every } \tau' \preceq \tau,\ s_1(\tau') \simeq s_2(\tau')\} \ .$$

Proposition 1. $\langle \mathrm{S}, \mathscr{L}\langle \mathrm{T}, \preceq \rangle, \supseteq, \mathrm{T}, \mathrm{d} \rangle$ *is a generalized ultrametric space.*

The following is immediate, and indeed, equivalent:

Proposition 2. *For every* $s_1, s_2, s_3 \in \mathrm{S}$*, the following are true:*

1. $\mathrm{d}(s_1, s_2) = \mathrm{T}$ *if and only if* $s_1 = s_2$*;*
2. $\mathrm{d}(s_1, s_2) = \mathrm{d}(s_2, s_1)$*;*
3. $\mathrm{d}(s_1, s_2) \supseteq \mathrm{d}(s_1, s_3) \cap \mathrm{d}(s_3, s_2)$.

We refer to clause 1 as the *identity of indiscernibles*, clause 2 as *symmetry*, and clause 3 as the *generalized ultrametric inequality*.

Proposition 3. $\langle \mathrm{S}, \mathscr{L}\langle \mathrm{T}, \preceq \rangle, \supseteq, \mathrm{T}, \mathrm{d} \rangle$ *is spherically complete.*

Spherical completeness implies Cauchy-completeness, but the converse is not true in general (see [10, prop. 10], [2, exam. 2.6]). The importance of spherical completeness will become clear in Section 4 (see Theorem 2).

There is also a natural order relation on signals, namely the prefix relation on signals.

We write $\sqsubseteq$ for a binary relation on S such that for every $s_1, s_2 \in \mathrm{S}$,

$$s_1 \sqsubseteq s_2 \iff \text{for every } \tau, \tau' \in \mathrm{T}, \text{ if } \tau \in \mathsf{dom}\, s_1 \text{ and } \tau' \preceq \tau, \text{ then } s_1(\tau') \simeq s_2(\tau').$$

We say that s_1 is a *prefix* of s_2 if and only if $s_1 \sqsubseteq s_2$.

It is easy to see that for every $s_1, s_2 \in \mathrm{S}$, $s_1 \sqsubseteq s_2$ if and only if there is $L \in \mathscr{L}\langle \mathrm{T}, \preceq \rangle$ such that $s_1 = s_2 \restriction L$,[3] and in particular, $s_1 \sqsubseteq s_2$ if and only if $s_1 = s_2 \restriction \mathrm{d}(s_1, s_2)$ (see [2, prop. 2.12 and thm. 2.13]).

Notice that for every $s \in \mathrm{S}$, $\emptyset \sqsubseteq s$; that is, the empty signal is a prefix of every signal.

Proposition 4. $\langle \mathrm{S}, \sqsubseteq \rangle$ *is a complete semilattice.*

For every $C \subseteq \mathrm{S}$ such that C is consistent in $\langle \mathrm{S}, \sqsubseteq \rangle$, we write $\bigsqcup C$ for the least upper bound of C in $\langle \mathrm{S}, \sqsubseteq \rangle$.

For every $s_1, s_2 \in \mathrm{S}$, we write $s_1 \sqcap s_2$ for the greatest lower bound of s_1 and s_2 in $\langle \mathrm{S}, \sqsubseteq \rangle$.

[1] For every ordered set $\langle P, \leqslant \rangle$, we write $\mathscr{L}\langle P, \leqslant \rangle$ for the set of all lower sets[2] of $\langle P, \leqslant \rangle$.

[2] For every ordered set $\langle P, \leqslant \rangle$, and every $L \subseteq P$, L is a *lower set* (also called a *down-set* or an *order ideal*) of $\langle P, \leqslant \rangle$ if and only if for any $p_1, p_2 \in P$, if $p_1 \leqslant p_2$ and $p_2 \in L$, then $p_1 \in L$.

[3] For every function f and every set A, we write $f \restriction A$ for the *restriction* of f to A, namely the function $\{\langle a, b \rangle \mid \langle a, b \rangle \in f \text{ and } a \in A\}$.

3 Causal and Strictly Causal Functions

Causality is a concept of fundamental importance in the study of timed systems. Informally, it represents the constraint that, at any time instance, the output events of a component do not depend on its future input events. This is only natural for components that model or simulate physical processes, or realize online algorithms; an effect cannot precede its cause.

Assume a non-empty partial function F on S.

We say that F is *causal* if and only if there is a partial function f such that for every $s \in \mathsf{dom}\, F$ and every $\tau \in \mathrm{T}$,

$$F(s)(\tau) \simeq f(s \upharpoonright \{\tau' \mid \tau' \preceq \tau\}, \tau) \ .$$

Example 1. Suppose that $\mathrm{T} = \mathbb{R}$, and $\preceq$ is the standard order on $\mathbb{R}$. Let p be a positive real number, and F a function on S such that for every $s \in \mathrm{S}$ and every $\tau \in \mathrm{T}$,

$$F(s)(\tau) \simeq \begin{cases} s(\tau) & \text{if there is } i \in \mathbb{Z} \text{ such that } \tau = p \cdot i; \\ \text{undefined} & \text{otherwise.} \end{cases}$$

Clearly, F is causal.

The function of Example 1 models a simple sampling process.

Now, as explained in Section 1, due to its relevance to the interpretation of feedback, of special interest is whether any particular partial function F on S has a fixed point, that is, whether there is $s \in \mathrm{S}$ such that $s = F(s)$ (see Fig. 1), and whether that fixed point is unique. The function of Example 1 has uncountably many fixed points, namely every $s \in \mathrm{S}$ such that

$$\mathsf{dom}\, s \subseteq \{\tau \mid \text{there is } i \in \mathbb{Z} \text{ such that } \tau = p \cdot i\} \ .$$

And it is easy to construct a causal function that does not have a fixed point.

Example 2. Let τ be a tag in T, v a value in V, and F a function on S such that for every $s \in \mathrm{S}$,

$$F(s) = \begin{cases} \emptyset & \text{if } \tau \in \mathsf{dom}\, s; \\ \{\langle \tau, v \rangle\} & \text{otherwise.} \end{cases}$$

It is easy to verify that F is causal. But F has no fixed point; for $F(\{\langle \tau, v \rangle\}) = \emptyset$, whereas $F(\emptyset) = \{\langle \tau, v \rangle\}$.

The function of Example 2 models a component whose behaviour at τ resembles a logical inverter, turning presence of event into absence, and vice versa.

Strict causality is causality bar instantaneous reaction. Informally, it is the constraint that, at any time instance, the output events of a component do not depend on its present or future input events. This operational definition has its origins in natural philosophy, and is of course inspired by physical reality: every physical system, at least, in the classical sense, is a strictly causal system.

We say that F is *strictly causal* if and only if there is a partial function f such that for every $s \in \mathsf{dom}\, F$ and every $\tau \in \mathrm{T}$,

$$F(s)(\tau) \simeq f(s \upharpoonright \{\tau' \mid \tau' \prec \tau\}, \tau) \ .$$

Proposition 5. *If F is strictly causal, then F is causal.*

Of course, the converse is false. For example, the sampling function of Example 1 is causal but not strictly causal.

Example 3. Suppose that $\mathrm{T} = \mathbb{R}$, and $\preceq$ is the standard order on $\mathbb{R}$. Let F be a function on S such that for every $s \in \mathrm{S}$ and every $\tau \in \mathrm{T}$,

$$F(s)(\tau) \simeq s(\tau - 1) \ .$$

Clearly, F is stricty causal.

The function of Example 3 models a simple constant-delay component. It is in fact a "delta causal" function, as defined in [7] and [8], and it is not hard to see that every such function is strictly causal (as is every "Δ-causal" function, as defined in [11] and [12]). The function of our next example models a variable reaction-time component, and is a strictly causal function that is not "delta causal" (nor "Δ-causal").

Example 4. Suppose that $\mathrm{T} = [0, \infty)$, $\preceq$ is the standard order on $[0, \infty)$, and $\mathrm{V} = (0, \infty)$. Let F be a function on S such that for every $s \in \mathrm{S}$ and any $\tau \in \mathrm{T}$,

$$F(s)(\tau) \simeq \begin{cases} 1 & \text{if there is } \tau' \in \mathsf{dom}\, s \text{ such that } \tau = \tau' + s(\tau'); \\ \text{undefined} & \text{otherwise.} \end{cases}$$

Clearly, F is strictly causal.

Now, the function of Example 3 has uncountably many fixed-points, namely every $s \in \mathrm{S}$ such that for every $\tau \in \mathbb{R}$,

$$s(\tau) \simeq s(\tau + 1) \ .$$

And the function of Example 4 has exactly one fixed point, namely the empty signal. And having ruled out instantaneous reaction, the reason behind the lack of fixed point in Example 2, one might expect that every strictly causal function has a fixed point. But this is not the case.

Example 5. Suppose that $\mathrm{T} = \mathbb{Z}$, $\preceq$ is the standard order on $\mathbb{Z}$, and $\mathrm{V} = \mathbb{N}$. Let F be a function on S such that for every $s \in \mathrm{S}$ and every $\tau \in \mathrm{T}$,

$$F(s)(\tau) \simeq \begin{cases} s(\tau - 1) + 1 & \text{if } \tau - 1 \in \mathsf{dom}\, s; \\ 0 & \text{otherwise.} \end{cases}$$

Clearly, F is strictly causal. However, F does not have a fixed point; any fixed point of F would be an order-embedding from the integers into the natural numbers, which is of course impossible.

Example 5 alone is enough to suggest that the classical notion of strictly causality is by itself too general to support a useful theory of timed systems.

4 Contracting and Strictly Contracting Functions

There is another, intuitively equivalent way to articulate the property of causality: a component is causal just as long as any two possible output signals differ no earlier than the input signals that produced them (see [8, p. 36], [13, p. 11], and [14, p. 383]). And this can be very elegantly expressed using the generalized distance function of Section 2.

We say that F is *contracting* if and only if for every $s_1, s_2 \in \mathsf{dom}\, F$,

$$\mathrm{d}(F(s_1), F(s_2)) \supseteq \mathrm{d}(s_1, s_2) \ .$$

In other words, a function is contracting just as long as the distance between any two signals in the range of the function is smaller than or equal to that between the signals in the domain of the function that map to them.

In [4, def. 5], causal functions were defined to be the contracting functions. Here, we prove that indeed they are.

Theorem 1. *F is causal if and only if F is contracting.*

Following the same line of reasoning, one might expect that a component is strictly causal just as long as any two possible output signals differ later, if at all, than the signals that produced them (see [8, p. 36]).

We say that F is *strictly contracting* if and only if for every $s_1, s_2 \in \mathsf{dom}\, F$ such that $s_1 \neq s_2$,

$$\mathrm{d}(F(s_1), F(s_2)) \supset \mathrm{d}(s_1, s_2) \ .$$

Proposition 6. *If F is strictly contracting, then F is contracting.*

In [3, p. 484], Naundorf defined strictly causal functions as the functions that we here call strictly contracting, and in [4, def. 6], that definition was rephrased using the generalized distance function to explicitly identify strictly causal functions with the strictly contracting functions. But the relationship between the proposed definition and the classical notion of strict causality was never formally examined. The next theorem, which is a direct application of the fixed-point theorem of Priess-Crampe and Ribenboim for strictly contracting functions on spherically complete generalized ultrametric spaces (see [5, thm. 1]), implies that, in fact, the two are not the same.

Assume non-empty $X \subseteq \mathrm{S}$.

Theorem 2. *If $\langle X, \mathscr{L}\langle \mathrm{T}, \preceq \rangle, \supseteq, \mathrm{T}, \mathrm{d} \rangle$ is spherically complete, then every strictly contracting function on X has exactly one fixed point.*

By Theorem 2, the function of Example 5 is not strictly contracting. What is then the use, if any, of strictly contracting functions in a fixed-point theory for strictly causal functions? The next couple of theorems are key in answering this question.

Theorem 3. *If F is strictly contracting, then F is strictly causal.*

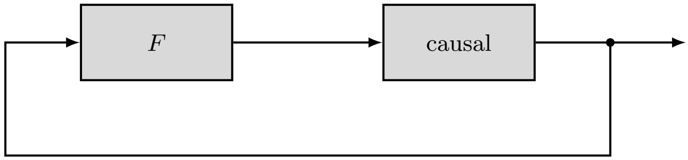

Fig. 2. A functional component realizes a strictly contracting function F if and only if the cascade of the component and any arbitrary causal component has a unique, well defined behaviour when arranged in a feedback configuration.

Theorem 4. *If* $\langle \mathsf{dom}\, F, \mathscr{L}\langle \mathrm{T}, \preceq\rangle, \supseteq, \mathrm{T}, \mathrm{d}\rangle$ *is spherically complete, then* F *is strictly contracting if and only if for every causal function* F' *from* $\mathsf{ran}\, F$ *to* $\mathsf{dom}\, F$*,* $F' \circ F$ *has a fixed point.*

An informal but informative way of reading Theorem 4 is the following: a functional component realizes a strictly contracting function if and only if the cascade of the component and any arbitrary causal filter has a unique, well defined behaviour when arranged in a feedback configuration (see Fig. 2); that is, the components that realize strictly contracting functions are those functional components that maintain the consistency of the feedback loop no matter how we chose to filter the signal transmitted over the feedback wire, as long as we do so in a causal way.

Theorem 4 completely characterizes strictly contracting functions in terms of the classical notion of causality, identifying the class of all such functions as the largest class of functions that have a fixed point not by some fortuitous coincidence, but as a direct consequence of their causality properties.

The implication of Theorem 3 and 4, we believe, is that the class of strictly contracting functions is the largest class of strictly causal functions that one can reasonably hope to attain a uniform fixed-point theory for.

Finally, if we further require that the domain of any signal in the domain of a function is well ordered under $\prec$, then the difference between a strictly causal function and a strictly contracting one vanishes.

Theorem 5. *If for every* $s \in \mathsf{dom}\, F$*,* $\langle \mathsf{dom}\, s, \prec\rangle$ *is well ordered, then* F *is strictly causal if and only if* F *is strictly contracting.*

For example, the variable-reaction-time function of Example 4 is not strictly contracting, as can be witnessed by the signals $\{\langle \frac{1}{2n+1}, \frac{2}{4n^2-1}\rangle \mid n \in \mathbb{N}\}$ and $\{\langle \frac{1}{2n}, \frac{1}{2n(n-1)}\rangle \mid n \in \mathbb{N} \text{ and } n \geq 2\}$, but its restriction to the set of all discrete-event signals, for instance, is.

Theorem 5, immediately applicable in the case of discrete-event systems, is most pleasing considering our emphasis on timed computation. It implies that for all kinds of computational timed systems, where components are expected to operate on discretely generated signals, including all programming languages and model-based design tools mentioned in the beginning of the introduction, the strictly contracting functions are exactly the strictly causal ones.

5 Fixed-Point Theory

We henceforth concentrate on the strictly contracting functions, and begin to develop the rudiments of a constructive fixed-point theory for such functions.

For every partial endofunction F on S, and any $s \in \mathsf{dom}\, F$, we say that s is a *post-fixed point* of F if and only if $s \sqsubseteq F(s)$.

Assume non-empty $X \subseteq$ S.

Lemma 1. *If $\langle X, \sqsubseteq \rangle$ is a subsemilattice of $\langle \mathrm{S}, \sqsubseteq \rangle$, then for every contracting function F on X, and every $s \in X$, the following are true:*

1. *$F(s) \sqcap F(F(s))$ is a post-fixed point of F;*
2. *if s is a post-fixed point of F, then $s \sqsubseteq F(s) \sqcap F(F(s))$.*

Lemma 2. *If $\langle X, \sqsubseteq \rangle$ is a subsemilattice of $\langle \mathrm{S}, \sqsubseteq \rangle$, then for every contracting function F on X, and any set P of post-fixed points of F, if P has a least upper bound in $\langle X, \sqsubseteq \rangle$, then $\bigsqcup_X P$ is a post-fixed point of F.*

Lemma 1 and 2 contain nearly all the ingredients of a transfinite recursion facilitating the construction of a chain that will converge to the desired fixed point. We may start with any arbitrary post-fixed point of the function F, and iterate through the function $\lambda x : X \,.\, F(x) \sqcap F(F(x))$ to form an ascending chain of such points. Every so often, we may take the supremum in $\langle X, \sqsubseteq \rangle$ of all signals theretofore constructed, and resume the process therefrom, until no further progress can be made. Of course, the phrase "every so often" is to be interpreted rather liberally here, and certain groundwork is required before we can formalize its transfinite intent.

We henceforth assume some familiarity with transfinite set theory, and in particular, ordinal numbers. The unversed reader may refer to any introductory textbook on set theory for details (e.g., see [15]).

Assume a subsemilattice $\langle X, \sqsubseteq \rangle$ of $\langle \mathrm{S}, \sqsubseteq \rangle$, and a function F on X.

We write $\mathsf{1m2}\, F$ for a function on X such that for any $s \in X$,

$$(\mathsf{1m2}\, F)(s) = F(s) \sqcap F(F(s)) \enspace .$$

In other words, $\mathsf{1m2}\, F$ is the function $\lambda x : X \,.\, F(x) \sqcap F(F(x))$.

Assume a directed-complete subsemilattice $\langle X, \sqsubseteq \rangle$ of $\langle \mathrm{S}, \sqsubseteq \rangle$, a contracting function F on X, and a post-fixed point s of F.

We let

$$(\mathsf{1m2}\, F)^0(s) = s \enspace ,$$

for every ordinal α,

$$(\mathsf{1m2}\, F)^{\alpha+1}(s) = (\mathsf{1m2}\, F)((\mathsf{1m2}\, F)^{\alpha}(s)) \enspace ,$$

and for every limit ordinal λ,

$$(\mathsf{1m2}\, F)^{\lambda}(s) = \bigsqcup\nolimits_X \{(\mathsf{1m2}\, F)^{\alpha}(s) \mid \alpha \in \lambda\} \enspace .$$

The following implies that for every ordinal α, $(\mathsf{1m2}\, F)^{\alpha}(s)$ is well defined:

Lemma 3. *If $\langle X, \sqsubseteq \rangle$ is a directed-complete subsemilattice of $\langle \mathrm{S}, \sqsubseteq \rangle$, then for every contracting function F on X, any post-fixed point s of F, and every ordinal α,*

1. $(\mathsf{1m2}\, F)^{\alpha}(s) \sqsubseteq F((\mathsf{1m2}\, F)^{\alpha}(s))$;
2. *for any* $\beta \in \alpha$, $(\mathsf{1m2}\, F)^{\beta}(s) \sqsubseteq (\mathsf{1m2}\, F)^{\alpha}(s)$.

By Lemma 3.2, and a simple cardinality argument, there is an ordinal α such that for every ordinal β such that $\alpha \in \beta$, $(\mathsf{1m2}\, F)^{\beta}(s) = (\mathsf{1m2}\, F)^{\alpha}(s)$. In fact, for every directed-complete subsemilattice $\langle X, \sqsubseteq \rangle$ of $\langle \mathrm{S}, \sqsubseteq \rangle$, there is a least ordinal α such that for every contracting function F on X, any post-fixed point s of F, and every ordinal β such that $\alpha \in \beta$, $(\mathsf{1m2}\, F)^{\beta}(s) = (\mathsf{1m2}\, F)^{\alpha}(s)$.

We write $\mathsf{oh}\, \langle X, \sqsubseteq \rangle$ for the least ordinal α such that there is no function φ from α to X such that for every $\beta, \gamma \in \alpha$, if $\beta \in \gamma$, then $\varphi(\beta) \sqsubset \varphi(\gamma)$.

In other words, $\mathsf{oh}\, \langle X, \sqsubseteq \rangle$ is the least ordinal that cannot be orderly embedded in $\langle X, \sqsubseteq \rangle$, which we may think of as the *ordinal height* of $\langle X, \sqsubseteq \rangle$. Notice that the Hartogs number of X is an ordinal that cannot be orderly embedded in $\langle X, \sqsubseteq \rangle$, and thus, $\mathsf{oh}\, \langle X, \sqsubseteq \rangle$ is well defined, and in particular, smaller than or equal to the Hartogs number of X.

Lemma 4. *If $\langle X, \sqsubseteq \rangle$ is a directed-complete subsemilattice $\langle X, \sqsubseteq \rangle$ of $\langle \mathrm{S}, \sqsubseteq \rangle$, then for every contracting function F on X, any post-fixed point s of F, and every ordinal α, if $(\mathsf{1m2}\, F)^{\alpha}(s)$ is not a fixed point of $\mathsf{1m2}\, F$, then $\alpha + 2 \in \mathsf{oh}\, \langle X, \sqsubseteq \rangle$.*

By Lemma 4, $(\mathsf{1m2}\, F)^{\mathsf{oh}\, \langle X, \sqsubseteq \rangle}(s)$ is a fixed point of $\mathsf{1m2}\, F$. Nevertheless, $(\mathsf{1m2}\, F)^{\mathsf{oh}\, \langle X, \sqsubseteq \rangle}(s)$ need not be a fixed point of F as intended. For example, if F is the function of Example 2, then for every ordinal α, $(\mathsf{1m2}\, F)^{\alpha}(\emptyset) = \emptyset$, even though $\emptyset$ is not a fixed point of F. This rather trivial example demonstrates how the recursion process might start stuttering at points that are not fixed under the function in question. If the function is strictly contracting, however, progress at such points is guaranteed.

Lemma 5. *If $\langle X, \sqsubseteq \rangle$ is a subsemilattice of $\langle \mathrm{S}, \sqsubseteq \rangle$, then for every strictly contracting function F on X, s is a fixed point of F if and only if s is a fixed point of $\mathsf{1m2}\, F$.*

We may at last put all the pieces together to obtain a constructive fixed-point theorem for strictly contracting functions on directed-complete subsemilattices of $\langle \mathrm{S}, \sqsubseteq \rangle$.

Theorem 6. *If $\langle X, \sqsubseteq \rangle$ is a directed-complete subsemilattice of $\langle \mathrm{S}, \sqsubseteq \rangle$, then for every strictly contracting function F on X, and any post-fixed point s of F,*

$$\mathsf{fix}\, F = (\mathsf{1m2}\, F)^{\mathsf{oh}\, \langle X, \sqsubseteq \rangle}(s) \ .$$

To be pedantic, Theorem 6 does not directly prove that F has a fixed point; unless there is a post-fixed point of F, the theorem is true vacuously. But by Lemma 1.1, for every $s \in X$, $(\mathsf{1m2}\, F)(s)$ is a post-fixed point of F.

This construction of fixed points as "limits of stationary transfinite iteration sequences" is very similar to the construction of extremal fixed points of monotone operators in [16] and references therein, where the function iterated is not $\mathsf{1m2}\,F$, but F itself. Notice, however, that if F preserves the prefix relation, then for any post-fixed point of F, $(\mathsf{1m2}\,F)(s) = F(s)$.

The astute reader will at this point anticipate the following:

Theorem 7. *If $\langle X, \sqsubseteq \rangle$ is a non-empty, directed-complete subsemilattice of $\langle \mathrm{S}, \sqsubseteq \rangle$, then for every strictly contracting function F on X,*

$$\mathsf{fix}\,F = \bigsqcup\nolimits_X \{s \mid s \in X \text{ and } s \sqsubseteq F(s)\} \ .$$

The construction of Theorem 7 is identical in form to Tarski's well known construction of greatest fixed points of order-preserving functions on complete lattices (see [17, thm. 1]).

We note here that $\mathsf{1m2}\,F$ is not, in general, order-preserving under the above premises (see [2, exam. 5.15]), as might be suspected, and thus, our fixed-point theorem is not a reduction to a standard order-theoretic one.

Now, having used transfinite recursion to construct fixed points, we may use transfinite induction to prove properties of them. And in particular, we may use Theorem 6 to establish a special proof rule.

Assume $P \subseteq \mathrm{S}$.

We say that P is *strictly inductive* if and only if every non-empty chain in $\langle P, \sqsubseteq \rangle$ has a least upper bound in $\langle P, \sqsubseteq \rangle$.

Note that P is strictly inductive if and only if $\langle P, \sqsubseteq \rangle$ is directed-complete (see [18, cor. 2]).

Theorem 8. *If $\langle X, \sqsubseteq \rangle$ is a non-empty, directed-complete subsemilattice of $\langle \mathrm{S}, \sqsubseteq \rangle$, then for every strictly contracting function F on X, and every non-empty, strictly inductive $P \subseteq X$, if for every $s \in P$, $(\mathsf{1m2}\,F)(s) \in P$, then $\mathsf{fix}\,F \in P$.*

Theorem 8 is an induction principle that one may use to prove properties of fixed points of strictly contracting endofunctions. We think of properties extensionally here; that is, a property is a set of signals. And the properties that are admissible for use with this principle are those that are non-empty and strictly inductive.

It is interesting to compare this principle with the fixed-point induction principle for order-preserving functions on complete partial orders (see [19]), which we will here refer to as *Scott-de Bakker induction*, and the fixed-point induction principle for contraction mappings on complete metric spaces (see [20]), which we will here refer to as *Reed-Roscoe induction* (see also [21], [22], [23]).

For a comparison between our principle and Scott-de Bakker induction, let F be a function of the most general kind of function to which both our principle and Scott-de Bakker induction apply, namely an order-preserving, strictly contracting function on a pointed, directed-complete subsemilattice $\langle X, \sqsubseteq \rangle$ of $\langle \mathrm{S}, \sqsubseteq \rangle$. Now assume a property $P \subseteq X$. If P is admissible for use with Scott-de Bakker induction, that is, closed under suprema in $\langle X, \sqsubseteq \rangle$ of arbitrary chains in

$\langle P, \sqsubseteq \rangle$, then $\{s \mid s \in P \text{ and } s \sqsubseteq F(s)\}$ is non-empty and strictly inductive. And if P is closed under F, then $\{s \mid s \in P \text{ and } s \sqsubseteq F(s)\}$ is trivially closed under $\mathsf{1m2}\,F$. Therefore, given any reasonable property-specification logic, our principle is at least as strong a proof rule as Scott-de Bakker induction. At the same time, the often inconvenient requirement that a property P that is admissible for use with Scott-de Bakker induction contain the least upper bound in $\langle X, \sqsubseteq \rangle$ of the empty chain, namely the least element in $\langle X, \sqsubseteq \rangle$, and the insistence that the least upper bound of every non-empty chain in $\langle P, \sqsubseteq \rangle$ be the same as in $\langle X, \sqsubseteq \rangle$ make it less likely that every property true of $\mathsf{fix}\,F$ that can be proved using our principle can also be proved using Scott-de Bakker induction. For this reason, we are inclined to say that, given any reasonable property-specification logic, our principle is a strictly stronger proof rule than Scott-de Bakker induction, in the case, of course, where both apply.

The relationship between our principle and Reed-Roscoe induction is less clear. $\langle \mathrm{S}, \mathscr{L}\,\langle \mathrm{T}, \preceq \rangle, \supseteq, \mathrm{T}, \mathrm{d} \rangle$ being a generalized ultrametric space rather than a metric one, it might even seem that there can be no common ground for a meaningful comparison between the two. Nevertheless, it is possible to generalize Reed-Roscoe induction in a way that extends its applicability to the present case, while preserving its essence. According to the generalized principle, then, for every strictly contracting function F on any Cauchy complete, non-empty, directed-complete subsemilattice of $\langle \mathrm{S}, \sqsubseteq \rangle$ such that every orbit under F is a Cauchy sequence, every non-empty property closed under limits of Cauchy sequences that is preserved by F is true of $\mathsf{fix}\,F$. One similarity between this principle and our own, and an interesting difference from Scott-de Bakker induction, is the lack of an explicit basis for the induction; as long as the property in question is non-empty, there is some basis available. In terms of closure and preservation of admissible properties, however, the two principles look rather divergent from one another. For example, the property of a signal having only a finite number of events in any finite interval of time is Cauchy complete, but not strictly inductive. On the other hand, by Lemma 1.1 and Theorem 7, our principle is better fit for proving properties that are closed under prefixes, such as, for example, the property of a signal having at most one event in any time interval of a certain fixed size. And for this reason, we suspect that, although complimentary to the generalized Read-Roscoe induction principle in theory, our principle might turn out to be more useful in practice, what can of course only be evaluated empirically.

6 Related Work

Fixed points have been used extensively in the construction of mathematical models in computer science. In most cases, ordered sets and monotone functions have been the more natural choice. But in the case of timed computation, metric spaces and contraction mappings have proved a better fit, and Tarski's fixed-point theorem and its variants have given place to Banach's contraction principle.

Common to all approaches using this kind of modeling framework that we know of is the requirement of a positive lower bound on the reaction time of each component in a system (e.g., see [20], [24], [12], [25], [7], [8], [9]). This constraint is used to guarantee that the functions modelling these components are actually contraction mappings with respect to the defined metrics. The motivation is of course the ability to use Banach's fixed-point theorem in the interpretation of feedback, but a notable consequence is the absence of non-trivial Zeno phenomena, what has always been considered a precondition for realism in the real-time systems community. And yet in modelling and simulation, where time is represented as an ordinary program variable, Zeno behaviours are not only realizable, but occasionally desirable as well. Simulating the dynamics of a bouncing ball, for example, will naturally give rise to a Zeno behaviour, and the mathematical model used to study or even define the semantics of the simulation environment should allow for that behaviour. This is impossible with the kind of metric spaces found in [20], [24], [12], [25], [7], [8], and [9] (see [4, sec. 4.1]).

Another limiting factor in the applicability of the existing approaches based on metric spaces is the choice of tag set. The latter is typically some unbounded subset of the real numbers, excluding other interesting choices, such as, for example, that of superdense time (see [4, sec. 4.3]).

Naundorf was the first to address these issues, abolishing the bounded reaction time constraint, and allowing for arbitrary tag sets (see [3]). He defined strictly causal functions as the functions that we here call strictly contracting, and used an ad hoc, non-constructive argument to prove the existence of a unique fixed point for every such function. Unlike that in [7], [8], and [9], Naundorf's definition of strict causality was at least sound (see Theorem 3), but nevertheless incomplete (e.g., see Example 5). It was rephrased in [4] using the generalized distance function to explicitly identify strictly causal functions with the strictly contracting ones. This provided access to the fixed-point theory of generalized ultrametric spaces, which, however, proved less useful than one might have hoped. The main fixed-point theorem of Priess-Crampe and Ribenboim for strictly contracting endofunctions offered little more than another non-constructive proof of Naundorf's theorem, improving only marginally on the latter by allowing the domain of the function to be any arbitrary spherically complete set of signals, and the few constructive fixed-point theorems that we know to be of any relevance (e.g., see *Proof of Theorem 9 for ordinal distances* in [10], [26, thm. 43]) were of limited applicability.

There have also been a few attempts to use complete partial orders and least fixed points in the study of timed systems. In [11], Yates and Gao reduced the fixed-point problem related to a system of so-called "Δ-causal" components to that of a suitably constructed Scott-continuous function, transferring the Kahn principle to networks of real-time processes, but once more, under the usual bounded reaction-time constraint. A more direct application of the principle in the context of timed systems was put forward in [27]. But the proposed definition of strict causality was still incomplete, unable to accommodate components with more arbitrarily varying reaction-times, such as the one modelled by the function of Example 4 restricted to the set of all discrete-event signals, and ultimately, systems with non-trivial Zeno behaviours. Finally, a more naive approach was proposed in [28], where components were modelled as Scott-continuous functions with respect to the prefix relation on signals, creating, of course, all kinds of causality problems, which, however, seem to have gone largely unnoticed.

7 Conclusion

Strictly contracting functions form the largest class of strictly causal functions for which fixed points exist uniformly and canonically. More importantly, they coincide with strictly causal function in the case of discrete-event computation. The constructive fixed-point theorem for such functions presented in this work is, we believe, a leap forward in our understanding of functional timed systems. But a complete treatment of such systems should allow for arbitrarily complex networks of components. What we need, then, is an abstract characterization of a class of structures that will support the development of the theory, and stay closed under the construction of products and function spaces of interest, enabling the treatment of arbitrary, even higher-order composition in a uniform and canonical way. This is the subject of future work.

References

1. Zeigler, B.P.: Theory of Modeling and Simulation. John Wiley & Sons (1976)
2. Matsikoudis, E., Lee, E.A.: The fixed-point theory of strictly causal functions. Technical Report UCB/EECS-2013-122, EECS Department, University of California, Berkeley (June 2013)
3. Naundorf, H.: Strictly causal functions have a unique fixed point. Theoretical Computer Science 238(1-2), 483–488 (2000)
4. Liu, X., Matsikoudis, E., Lee, E.A.: Modeling timed concurrent systems. In: Baier, C., Hermanns, H. (eds.) CONCUR 2006. LNCS, vol. 4137, pp. 1–15. Springer, Heidelberg (2006)
5. Priess-Crampe, S., Ribenboim, P.: Fixed points, combs and generalized power series. Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg 63(1), 227–244 (1993)
6. Davey, B.A., Priestley, H.A.: Introduction to Lattices and Order, 2nd edn. Cambridge University Press (2002)
7. Lee, E.A., Sangiovanni-Vincentelli, A.: A framework for comparing models of computation. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems 17(12), 1217–1229 (1998)
8. Lee, E.A.: Modeling concurrent real-time processes using discrete events. Annals of Software Engineering 7(1), 25–45 (1999)
9. Liu, J., Lee, E.A.: On the causality of mixed-signal and hybrid models. In: Maler, O., Pnueli, A. (eds.) HSCC 2003. LNCS, vol. 2623, pp. 328–342. Springer, Heidelberg (2003)
10. Hitzler, P., Seda, A.K.: Generalized metrics and uniquely determined logic programs. Theoretical Computer Science 305(1-3), 187–219 (2003)
11. Yates, R.K., Gao, G.R.: A Kahn principle for networks of nonmonotonic real-time processes. In: Bode, A., Reeve, M., Wolf, G. (eds.) PARLE 1993. LNCS, vol. 694, pp. 209–227. Springer, Heidelberg (1993)

12. Yates, R.K.: Networks of real-time processes. In: Best, E. (ed.) CONCUR 1993. LNCS, vol. 715, pp. 384–397. Springer, Heidelberg (1993)
13. Broy, M.: Refinement of time. Theoretical Computer Science 253(1), 3–26 (2001)
14. Lee, E.A., Varaiya, P.: Structure and Interpretation of Signals and Systems, 2nd edn. (2011), `http://LeeVariaya.org`
15. Enderton, H.B.: Elements of Set Theory. Academic Press (1977)
16. Cousot, P., Cousot, R.: Constructive versions of Tarski's fixed point theorems. Pacific J. of Math. 82(1), 43–57 (1979)
17. Tarski, A.: A lattice-theoretical fixpoint theorem and its applications. Pacific J. of Math. 5(2), 285–309 (1955)
18. Markowsky, G.: Chain-complete posets and directed sets with applications. Algebra Universalis 6(1), 53–68 (1976)
19. Scott, D.S., de Bakker, J.W.: A theory of programs. Unpublished notes, Seminar on Programming, IBM Research Center, Vienna, Austria (1969)
20. Reed, G.M., Roscoe, A.W.: A timed model for communicating sequential processes. In: Kott, L. (ed.) ICALP 1986. LNCS, vol. 226, pp. 314–323. Springer, Heidelberg (1986)
21. Rounds, W.C.: Applications of topology to semantics of communicating processes. In: Brookes, S., Roscoe, A.W., Winskel, G. (eds.) Seminar on Concurrency. LNCS, vol. 197, pp. 360–372. Springer, Heidelberg (1985)
22. Roscoe, A.W.: Topology, computer science, and the mathematics of convergence. In: Reed, G.M., Roscoe, A.W., Wachter, R.F. (eds.) Topology and Category Theory in Computer Science, pp. 1–27. Oxford University Press, Inc., New York (1991)
23. Kozen, D., Ruozzi, N.: Applications of metric coinduction. In: Mossakowski, T., Montanari, U., Haveraaen, M. (eds.) CALCO 2007. LNCS, vol. 4624, pp. 327–341. Springer, Heidelberg (2007)
24. Reed, G.M., Roscoe, A.W.: Metric spaces as models for real-time concurrency. In: Main, M.G., Mislove, M.W., Melton, A.C., Schmidt, D. (eds.) MFPS 1987. LNCS, vol. 298, pp. 331–343. Springer, Heidelberg (1988)
25. Müller, O., Scholz, P.: Functional specification of real-time and hybrid systems. In: Maler, O. (ed.) HART 1997. LNCS, vol. 1201, pp. 273–285. Springer, Heidelberg (1997)
26. Krötzsch, M.: Generalized ultrametric spaces in quantitative domain theory. Theoretical Computer Science 368(1-2), 30–49 (2006)
27. Liu, X., Lee, E.A.: CPO semantics of timed interactive actor networks. Theoretical Computer Science 409(1), 110–125 (2008)
28. Caspi, P., Benveniste, A., Lublinerman, R., Tripakis, S.: Actors without directors: A Kahnian view of heterogeneous systems. In: Majumdar, R., Tabuada, P. (eds.) HSCC 2009. LNCS, vol. 5469, pp. 46–60. Springer, Heidelberg (2009)