A Domain for Concurrent Termination: A Generalization of Mazurkiewicz traces*

Volker Diekert
Universität Stuttgart
Institut für Informatik
Breitwiesenstr. 20-22
D-70565 Stuttgart

Paul Gastin
Institut Blaise Pascal
LITP
4, place Jussieu
F-75252 Paris Cedex 05

November 10, 1994

Abstract

This paper generalizes the concept of Mazurkiewicz traces to a *fuzzy* description of a concurrent process, where a known prefix is given in a first component and a second alphabetic component yields necessary information about future actions. This allows to define a good semantic domain where the concatenation is continuous with respect to the Scott- and to the Lawson topology. For this we define the notion of $\alpha-$ and of δ-trace. We show various mathematical results proving thereby the soundness of our approach. Our theory is a proper generalization of the theory of finite and infinite words (with explicit termination) and of the theory of finite and infinite (real and complex) traces. We make use of trace theory, domain theory, and topology.

1 Introduction

The theory of Mazurkiewicz traces has been recognized as an important tool for investigating concurrent systems, overviews are in [Maz87, AR88, Per89, Die90, DR95]. The underlying idea is that for a given alphabet Σ of actions (or events) we fix an irreflexive and symmetric independence relation $I \subseteq \Sigma \times \Sigma$. The semantics is that independent actions can be executed concurrently; and finite processes are described by elements of the quotient monoid $\mathbb{M}(\Sigma, D) = \Sigma^*/\{ab = ba \mid (a,b) \in I\}$. The complement $D = \Sigma \times \Sigma \setminus I$ is called the *dependence relation.*

*This research has been supported by the ESPRIT Basic Research Actions No. 6317 ASMICS II and No. 6067 CALIBAN.

The elements of $\mathbb{M}(\Sigma, D)$ are called *traces*. The concatenation of traces combines by a single operation sequential and parallel composition. If $D = \Sigma \times \Sigma$ is full, then the concatenation is the sequential composition of words. If s, t are traces such that the actions of s and t are pairwise independent, then st can be viewed as the parallel composition $s\|t$.

Another important situation is given when the dependency D is transitive. Then we have $\mathbb{M}(\Sigma, D) = \Sigma_1^* \times \cdots \times \Sigma_k^*$ for some partition $\Sigma = \Sigma_1 \dot\cup \cdots \dot\cup \Sigma_k$. In this case we describe a system with k independent components, where each component behaves sequentially. Let us continue with this example. A process p is then described as a tuple of words $p = (v_1, \ldots, v_k)$. Assume that p' is a prefix of p describing the part of p which is accomplished, and let $q = (w_1, \ldots, w_k)$ be a second process, which should be executed after p. Assume that $p' = (v_1', \ldots, v_k')$ has some components which are terminated, $v_i' = v_i$, and some components, where $v_i' < v_i$ is a proper prefix. Using the symbol $\surd$ for non-termination, we may assume $p' = (v_1'\surd, \ldots, v_i'\surd, v_{i+1}, \ldots, v_k)$. We want the following law:

$$(v_1'\surd, \ldots, v_i'\surd, v_{i+1}, \ldots, v_k)(w_1, \ldots, w_k) = (v_1'\surd, \ldots, v_i'\surd, v_{i+1}w_{i+1}, \ldots, v_kw_k)$$

This law expresses the fact that if we know termination on a subset of components $K \subseteq \{1, \ldots, k\}$ (e.g., as above $K = \{i+1, \ldots, k\}$), then the execution of $q = (w_1, \ldots, w_k)$ can be started on these components without waiting for global termination. On the other components we have no further information than $v_i'\surd$.

This basic idea generalizes perfectly as follows in the frame of Mazurkiewicz traces. We describe a process p as a pair (t, A) where $t \in \mathbb{M}(\Sigma, D)$ is trace and $A \subseteq \Sigma$ is a subset of actions which may still be performed before termination. If $A = \emptyset$ is the empty set then this explicitly means global termination. One should note the difference with failure (or ready) semantics in process algebra like CCS or CSP [Hoa85, Mil80], where usually the behavior in the next step is specified. Here the finite alphabetic information A is given for the whole future and for the future only.

However, alphabetic information about the future for composition laws is obviously useful for process algebra too. Some recent papers about this subject have been published, see e.g. [RW94].

If we are uncertain about the future then, we have possibly to allow everything. Thus in our formalism (t, Σ) is always an approximation of (t, A) for all $A \subseteq \Sigma$. More precisely, if $A \subseteq B$ then (t, A) imposes more restrictions on the process than (t, B) does. Hence for $A \subseteq B$, we view (t, A) as a better approximation than (t, B) and therefore $(t, B) \sqsubseteq (t, A)$. In the following we will call these pairs *α-traces*. In many cases, it is convenient to abstract from the second component to the set $D(A) = \{b \in \Sigma \mid (a, b) \in D\}$, since, as we will see, this is the only information we need for concatenation. This leads to the notion of *δ-trace*, which identifies (t, A) with (t, A'), if $(t, D(A)) = (t, D(A'))$. Although our motivation stems from the description of finite processes, we are able to cope with infinite behaviors, too. In fact, it is convenient to deal with finite and infinite behaviors simultaneously since, if the process runs on different independent components, then it is very natural that on some components there is termination whereas on other components it runs infinitely. The precise mathematical formalism is given below. We will develop a theory which results in an algebraic cpo (Scott domain), where the concatenation is continuous. In particular, it is monotone with respect to both components. Moreover, we will define a natural metric on the space of α-traces (δ-traces resp.) such that the induced topological space becomes compact (by the Lawson topology [Law88]). The concatenation is in fact uniformly continuous in this topology. These mathematical results underline the correctness of our approach. We think that it is an important contribution to an algebraic understanding of concurrent processes. Our theory generalizes the theory of finite and infinite words (with explicit termination) and the theory of infinite (real and complex) traces. Some of the following ideas were proposed in the invited lecture [Die93a].

2 Preliminaries from poset theory

Let $(Z, \leq)$ be a poset, i.e., a partially ordered set. A subset $Y \subseteq Z$ is *coherent* if for all $x, y \in Y$, there exists $z \in Z$ such that $x \leq z$ and $y \leq z$. A subset $Y \subseteq Z$ is *directed* if it is non-empty and for all $x, y \in Y$, there exists $z \in Y$ such that $x \leq z$ and $y \leq z$. A subset $Y \subseteq Z$ is *bounded* if there exists $z \in Z$ such that $x \leq z$ for all $x \in Y$. The poset $(Z, \leq)$ is a *coherently complete partial order* (a *ccpo* for short) if every coherent subset of Z has a least upper bound in Z. Similarly, the poset $(Z, \leq)$ is a *directed complete partial order* (a *dcpo* for short) if every directed subset of Z has a least upper bound in Z. Finally, the poset $(Z, \leq)$ is a *bounded complete partial order* (a *bcpo* for short) if every bounded subset of Z has a least upper bound in Z. Note that a ccpo is indeed both a dcpo and a bcpo. Whenever they exist, the least upper bound and the greatest lower bound of a subset $Y \subseteq Z$ will be denoted by $\sqcup Y$ and $\sqcap Y$ respectively.

Let $(Z, \leq)$ be a ccpo. An element $x \in Z$ is *prime* if for all *coherent* subset $Y \subseteq Z$, $x \leq \bigsqcup Y$ implies $x \leq y$ for some $y \in Y$. For all $z \in Z$, we denote by $P(z) = \{x \in Z \mid x \leq z \text{ and } x \text{ is prime}\}$ the set of prime elements of Z less than or equal to z. A ccpo $(Z, \leq)$ is *prime algebraic* if $z = \bigsqcup P(z)$ for all $z \in Z$.

Similarly, let $(Z, \leq)$ be a dcpo. An element $x \in Z$ is *compact* if for all *directed* subset $Y \subseteq Z$, $x \leq \bigsqcup Y$ implies $x \leq y$ for some $y \in Y$. For $z \in Z$, we denote by $K(z) = \{x \in Z \mid x \leq z \text{ and } x \text{ is compact}\}$ the set of compact elements less than or equal to z. A dcpo $(Z, \leq)$ is *algebraic* if $K(z)$ is directed and $z = \bigsqcup K(z)$ for all $z \in Z$. Note that, if any pair of elements of Z which is bounded above admits a least upper bound (in particular if $(Z, \leq)$ is a ccpo or a bcpo) then $K(z)$ is directed for all $z \in Z$. Finally, an algebraic dcpo is called a *(Scott) domain.* For a poset $(Z, \leq)$ and $z \in Z$ we use the following standard notations:

$$\downarrow z \;=\; \{y \in Z \mid y \leq z\} \quad \text{and} \quad \uparrow z \;=\; \{y \in Z \mid y \geq z\}.$$

Let $(X, \leq)$ and $(Y, \leq)$ be two posets. A mapping $f : X \longrightarrow Y$ is monotone if for all $x \in X$ and $y \in Y$, $x \leq y$ implies $f(x) \leq f(y)$. Let $(X, \leq)$ and $(Y, \leq)$ be two dcpos. A monotone mapping $f : X \longrightarrow Y$ is continuous if $f(\bigsqcup Z) = \bigsqcup f(Z)$ for all directed set $Z \subseteq X$.

3 Preliminaries from trace theory

We assume that the reader is familiar with the basic concepts of trace theory. A detailed presentation of traces and dependence graphs appears in [GP95]. Let Σ be a finite alphabet with a symmetric and irreflexive independence relation $I \subseteq \Sigma\times\Sigma$. The complement $D = \Sigma\times\Sigma\setminus I$ is called the dependence relation, and (Σ, D) is called a dependence alphabet. A dependence alphabet is represented as an undirected graph where self-loops are omitted. A dependence graph (over (Σ, D)) is (an isomorphism class of) a directed acyclic node-labelled graph $[V, E, \lambda]$ where V is a countable set of vertices, $\lambda : V \longrightarrow \Sigma$ is the node-labelling, and where the edge relation $E \subseteq V \times V$ satisfies the following restriction: $E \cup E^{-1} \cup \mathrm{id}_V = \lambda^{-1}(D)$, i.e., edges are between dependent vertices, only. Furthermore, for technical reasons we demand that the partial order, which is induced by the reflexive and transitive closure E^*, is well-founded. The set of dependence graphs $\mathbb{G}(\Sigma, D)$ is a monoid with the following concatenation $[V_1, E_1, \lambda_1] \cdot [V_2, E_2, \lambda_2] = [V, E, \lambda]$ where $V = V_1 \dot\cup V_2$, $\lambda = \lambda_1 \dot\cup \lambda_2$ and $E = E_1 \dot\cup E_2 \dot\cup (V_1 \times V_2 \cap \lambda^{-1}(D))$. This means that we take the disjoint union and introduce new arcs from V_1 to V_2 whenever vertices have dependent labels. The neutral element is the empty graph, denoted by $1 = [\emptyset, \emptyset, \emptyset]$.

For $x \in V$ we denote by $\downarrow x$ the dependence graph which is induced by all elements below x, i.e., $\downarrow x$ is the restriction of $[V, E, \lambda]$ to the set $\{x' \in V \mid (x', x) \in E^*\}$. If the set V is finite, then $[V, E, \lambda]$ has a canonical identification with a finite trace of $\mathbb{M}(\Sigma, D) = \Sigma^*/\{ab = ba \mid (a, b) \in I\}$. If all vertices of V have a finite past, i.e., if $\downarrow x$ is finite for all $x \in V$, then $[V, E, \lambda]$ is called a *real trace*. Thus, a real trace is a finite or infinite trace in the sense of [Maz87, Gas90, Kwi90]. The set of real traces is denoted by $\mathbb{R}(\Sigma, D)$. Note that $\mathbb{R}(\Sigma, D)$ is not a submonoid of $\mathbb{G}(\Sigma, D)$, e.g. the concatenation of a^ω with itself yields a dependence graph with a non-empty transfinite part, i.e., some vertices do not have a finite past.

For dependence graphs we have a natural notion of prefix (here the assumption of well-foundedness is used). For $g, h \in \mathbb{G}(\Sigma, D)$ we write $g \leq h$ if $gf = h$ for some dependence graph f. More generally, we denote this *prefix relation* by $\leq$ for any monoid. Note that $\leq$ is not necessarily antisymmetric. The subset $\mathbb{R}(\Sigma, D)$ with the induced prefix ordering forms a prime algebraic coherently complete domain, [GR93]. The same is *almost* true for $\mathbb{G}(\Sigma, D)$, [GP92]. We do not go into details what "almost" means, but the basic reason that $\mathbb{G}(\Sigma, D)$ can not form a complete partial order results from set theoretical considerations: The least upper bound of all countable ordinals is not countable anymore. For $g = [V, E, \lambda] \in \mathbb{G}(\Sigma, D)$, we let the *real part* $\mathrm{Re}(g)$ of g be the restriction of g to the set $\{x \in V \mid \downarrow x \text{ is finite}\}$. The real part $\mathrm{Re}(g) \in \mathbb{R}(\Sigma, D)$ is the maximal real prefix of g. We also define the *alphabet* and *alphabet at infinity* for $g = [V, E, \lambda] \in \mathbb{G}(\Sigma, D)$ by:

$$\begin{aligned} \mathrm{alph}(g) &= \lambda(V) \\ \mathrm{alphinf}(g) &= \lambda(\{x \in V \mid \downarrow x \text{ is infinite or } \lambda^{-1}\lambda(x) \text{ is infinite}\}) \end{aligned}$$

For $A \subseteq \Sigma$ and $g = [V, E, \lambda] \in \mathbb{G}(\Sigma, D)$, the maximal real prefix of g which is independent of A, denoted by $\mu_A(g)$, is the restriction of g to the set $\{\, x \in V \mid \downarrow x \text{ is finite and } \mathrm{alph}(\downarrow x) \times A \subseteq I\,\}$. Note that $\mu_A(g) \in \mathbb{R}(\Sigma, D)$ is a real trace and that $\mu_\emptyset(g) = \mathrm{Re}(g)$.

Since $\mu_A(g)$ is a prefix of g and $\mathbb{G}(\Sigma, D)$ is left-cancellative, we can define the dependence graph $\mu_A(g)^{-1}g$ by the equation $\mu_A(g)(\mu_A(g)^{-1}g) = g$. In the following we are only interested in the alphabet of this graph. As a matter of fact, it is more convenient to include the set $\mathrm{alphinf}(g)$ in this alphabet. Thus, we define

$$\sigma_A(g) = \mathrm{alph}(\mu_A(g)^{-1}g) \cup \mathrm{alphinf}(g)$$

4 The calculus of α- and δ-traces

Definition 4.1 *An α-trace (δ-trace resp.) is a pair $x = (r, A)$ (a pair $x = (r, D(A))$ resp.) where $r \in \mathbb{R}(\Sigma, D)$ is a real trace, and $A \subseteq \Sigma$ is a subset of Σ such that* $\mathrm{alphinf}(r) \subseteq A$. *The trace r is also denoted by* $\mathrm{Re}(x)$. *The set A (the set $D(A)$ resp.) is denoted by* $\mathrm{Im}(x)$ *and it is called the* imaginary part *of x. The set of α-traces is denoted by $\mathbb{F}_\alpha(\Sigma, D)$, the set of δ-traces is denoted by $\mathbb{F}_\delta(\Sigma, D)$. Both sets are monoids with $(1, \emptyset)$ as neutral element and the following concatenation*

$$\begin{aligned} (r, A) \cdot (s, B) &= (r\mu_A(s), \sigma_A(s) \cup A \cup B) && \textit{for } \alpha\textit{-traces} \\ (r, D(A)) \cdot (s, D(B)) &= (r\mu_A(s), D(\sigma_A(s) \cup A \cup B)) && \textit{for } \delta\textit{-traces} \end{aligned}$$

Note that the concatenation of δ-traces is well-defined since $\mu_A(g) = \mu_{A'}(g)$ and $\sigma_A(g) = \sigma_{A'}(g)$ for all $g \in G$ and $A, A' \subseteq \Sigma$ such that $D(A) = D(A')$.

We can think of an α- or δ-trace as a *fuzzy* description or an approximation of a process. The first component is a real trace which describes a prefix of the actual process; the second part is an alphabetic information about possible future actions in the α-case or about possible future dependencies (which is sufficient in many cases) in the δ-case. This interpretation leads to the following definition of a natural approximation ordering.

Definition 4.2 *For α-traces (r, A), (s, B) we define $(r, A) \sqsubseteq (s, B)$, if $r \leq s$ and* $\mathrm{alph}(r^{-1}s) \cup B \subseteq A$. *Similarly, for δ-traces $(r, D(A))$, $(s, D(B))$ we define $(r, D(A)) \sqsubseteq (s, D(B))$, if $r \leq s$ and* $D(\mathrm{alph}(r^{-1}s) \cup B) \subseteq D(A)$.

Remark 4.3
i) The natural mapping $\mathbb{F}_\alpha(\Sigma, D) \longrightarrow \mathbb{F}_\delta(\Sigma, D)$, $(r, A) \longmapsto (r, D(A))$ is a surjective and monotone homomorphism.
ii) The prefix relation $\leq$ of $\mathbb{F}_\alpha$ ($\mathbb{F}_\delta$ resp.) is antisymmetric, whence a partial order. Moreover, the prefix ordering $\leq$ and the approximation ordering $\sqsubseteq$ are disjoint in the following sense: If $(r, A) \leq (s, B)$ and $(r, A) \sqsubseteq (s, B)$ then we have $(r, A) = (s, B)$.
iii) If $D = \Sigma \times \Sigma$ is the full dependence relation then the consideration of $\mathbb{F}_\delta$ is equivalent to the formalism of Σ^∞ together with the non-termination symbol $\surd$ with $\surd \notin \Sigma$. In fact, for $D(A)$ we only have two possibilities: either $D(A) = \Sigma$, which means non-termination, or $D(A) = \emptyset$, which means explicit termination. A δ-trace in this special case either is a finite non-terminated word $v\surd \in \Sigma^*\surd$, either it is a finite terminated word $v \in \Sigma^*$, or it is an infinite word $v \in \Sigma^\omega$ (where the non-termination symbol is redundant and therefore it may be omitted). The calculus is as usual such that $v\surd$ or infinite words are right-absorbent. Furthermore, we have $u\surd \sqsubseteq v\surd \sqsubseteq v$ if and only if $u \leq v$, $u, v \in \Sigma^\infty$.
iv) The bottom element with respect to $\sqsubseteq$ is the pair $(1, \Sigma)$. It signifies that nothing has happened and everything is possible. Note that $(1, \Sigma)$ is right-absorbent with respect to the concatenation and we have $(r, A) \cdot (1, \Sigma) = (r, \Sigma)$.

We have a sequence of canonical mappings

$$\mathbb{M}(\Sigma, D) \hookrightarrow \mathbb{R}(\Sigma, D) \hookrightarrow \mathbb{G}(\Sigma, D) \xrightarrow{\varphi} \mathbb{F}_\alpha(\Sigma, D) \twoheadrightarrow \mathbb{F}_\delta(\Sigma, D)$$

The first two mappings are inclusions, the mapping $\varphi : \mathbb{G}(\Sigma, D) \to \mathbb{F}_\alpha(\Sigma, D)$, $g \mapsto (\mathrm{Re}(g), \mathrm{alphinf}(g))$ is a homomorphism, but it is not surjective (for $\Sigma \neq \emptyset$), and $\mathbb{F}_\alpha(\Sigma, D) \to \mathbb{F}_\delta(\Sigma, D)$ is the surjective homomorphism defined above. The image of $\mathbb{G}(\Sigma, D)$ in $\mathbb{F}_\alpha(\Sigma, D)$ is in fact the monoid of α-complex traces $\mathbb{C}_\alpha(\Sigma, D)$, the image of $\mathbb{G}(\Sigma, D)$ in $\mathbb{F}_\delta(\Sigma, D)$ is the monoid of complex traces $\mathbb{C}(\Sigma, D)$; both monoids are defined in [Die93b]. The mapping $\mathbb{R}(\Sigma, D) \to \mathbb{F}_\delta(\Sigma, D)$, $r \mapsto (r, D(\mathrm{alphinf}(r)))$ is injective, hence $\mathbb{R}(\Sigma, D) \subseteq \mathbb{C}(\Sigma, D) \subseteq \mathbb{F}_\delta(\Sigma, D)$. Similarly, the mapping $\mathbb{R}(\Sigma, D) \to \mathbb{F}_\alpha(\Sigma, D)$, $r \mapsto (r, \mathrm{alphinf}(r))$ is injective, hence $\mathbb{R}(\Sigma, D) \subseteq \mathbb{C}_\alpha(\Sigma, D) \subseteq \mathbb{F}_\alpha(\Sigma, D)$.

In the following we focus our attention to the set of α-traces $\mathbb{F}_\alpha(\Sigma, D)$. Analogous results can be obtained for the slightly more abstract model of δ-traces. This is left to the reader.

The homomorphism $\mathbb{M}(\Sigma, D) \hookrightarrow \mathbb{F}_\alpha(\Sigma, D)$, $t \mapsto (t, \emptyset)$ is monotone with respect to the prefix ordering $\leq$. It maps a finite trace t to a finite terminated process $(t, \emptyset)$. If $s \neq t$ then $(s, \emptyset)$ and $(t, \emptyset)$ are incomparable with respect to the approximation ordering $\sqsubseteq$. This is a special case of Rem. 4.3, ii).

The prefix ordering of the monoid $\mathbb{F}_\alpha(\Sigma, D)$ has the following nice property.

Proposition 4.4 *The set $(\mathbb{F}_\alpha(\Sigma, D), \leq)$ is a coherently complete partial order. If $X \subseteq \mathbb{F}_\alpha(\Sigma, D)$ is $\leq$-coherent, then we have $\bigsqcup_\leq X = (r, A)$ where $r = \bigsqcup_\leq \mathrm{Re}(X)$ and $A = \mathrm{alphinf}(r) \cup \bigcup_{x\in X} \mathrm{Im}(x)$.*

The proposition above allows in a convenient way to define an ω-product of α-traces $x_1 x_2 \cdots$ for $x_i \in \mathbb{F}_\alpha(\Sigma, D)$, $i \geq 1$. The ω-product is defined as the least upper bound $\sqcup_{i<\omega}(x_1 \ldots x_i)$.

The main drawback of the prefix ordering is that it is not monotone with respect to the concatenation. If $x' \leq x$ and $y' \leq y$ then we cannot infer $x'y' \leq xy$. Overcoming this difficulty is a basic motivation to study $\mathbb{F}_\alpha(\Sigma, D)$ with the approximation ordering $\sqsubseteq$.

5 Poset properties of $\mathbb{F}_\alpha(\Sigma, D)$ with the approximation ordering

Proposition 5.1 *The concatenation of α-traces is monotone for the approximation order: Let $x, x', y, y' \in \mathbb{F}_\alpha(\Sigma, D)$ be α-traces. If $x' \sqsubseteq x$ and $y' \sqsubseteq y$ then $x' \cdot y' \sqsubseteq x \cdot y$.*

Proof: Let $x' = (r', A')$, $x = (r, A)$, $y' = (s', B')$ and $y = (s, B)$ such that $r' \leq r$, $A' \supseteq \mathrm{alph}(r'^{-1}r) \cup A$ and $s' \leq s$, $B' \supseteq \mathrm{alph}(s'^{-1}s) \cup B$. We have $x'y' = (r'\mu_{A'}(s'), \sigma_{A'}(s') \cup A' \cup B')$. Note that $\mu_{A'}(s') \leq \mu_A(s') \leq \mu_A(s)$. Furthermore, since $A' \supseteq \mathrm{alph}(r'^{-1}r)$, we see that $r'\mu_{A'}(s') \leq r\mu_A(s)$. Therefore it remains to show that $\sigma_{A'}(s') \cup A' \cup B' \supseteq \sigma_A(s) \cup A \cup B \cup \mathrm{alph}((r'\mu_{A'}(s'))^{-1}(r\mu_A(s)))$. Note that $(r'\mu_{A'}(s'))^{-1}(r\mu_A(s))(\mu_A(s)^{-1}s) = (r'^{-1}r)(\mu_{A'}(s')^{-1}s')(s'^{-1}s)$. Since $A' \supseteq A \cup \mathrm{alph}(r'^{-1}r)$ and $B' \supseteq B \cup \mathrm{alph}(s'^{-1}s)$, the result follows from $\sigma_{A'}(s') \supseteq \sigma_A(s)$. □

The following generalization remains to the reader.

Proposition 5.2 *Let $(x_i)_{i\geq 1}$, $(y_i)_{i\geq 1}$ be two infinite sequences of α-traces such that $x_i \sqsubseteq y_i$ for all $i \geq 1$. Then we have for the infinite ω-product*

$$x_1 x_2 \ldots \sqsubseteq y_1 y_2 \ldots \ .$$

The following investigation shows that $(\mathbb{F}_\alpha(\Sigma, D), \sqsubseteq)$ is a good semantic domain with many desirable completeness properties. Moreover, based on Prop. 5.1, we will see that the concatenation behaves well with respect to these properties.

Theorem 5.3 *The poset $(\mathbb{F}_\alpha(\Sigma, D), \sqsubseteq)$ is a coherently complete partial order. If $X \subseteq \mathbb{F}_\alpha(\Sigma, D)$ is a coherent set, then we have*

$$\sqcup X = (\sqcup_\leq \mathrm{Re}(X), \cap_{x\in X} \mathrm{Im}(x)) \ .$$

Proof: Clearly $\mathrm{Re}(X)$ is a coherent subset of $(\mathbb{R}(\Sigma, D), \leq)$. Hence $r = \sqcup_{\leq} \mathrm{Re}(X) \in \mathbb{R}(\Sigma, D)$ exists. The next observation is that if $x \in X$ then $\mathrm{Im}(x) \supseteq \mathrm{alphinf}(r)$. This follows from the fact that $\mathrm{Im}(y) \supseteq \mathrm{alphinf}(\mathrm{Re}(y))$ for all α-trace y and that X is coherent. Once we have this, we see that $(r, \cap_{x \in X} \mathrm{Im}(x))$ is an α-trace. An easy reflection then shows the relation $x \sqsubseteq (r, \cap_{x \in X} \mathrm{Im}(x))$ for all $x \in X$. Finally let $x \sqsubseteq (s, B)$ for all $x \in X$. Then $r \leq s$ and $\mathrm{Im}(x) \supseteq \mathrm{alph}(r^{-1}s) \cup B$ for all $x \in X$. Hence we get the result. $\square$

Corollary 5.4 *Let $Y \subseteq \mathbb{F}(\Sigma, D)$ be any non-empty set of α-traces. Then the greatest lower bound $\sqcap Y$ exists and we have $\sqcap Y = (r, A)$ where*

$$r = \sqcap \mathrm{Re}(Y) \quad \textit{and} \quad A = \bigcup_{y \in Y} \left(\mathrm{Im}(y) \cup \mathrm{alph}(r^{-1} \mathrm{Re}(y)) \right)$$

Proof: The existence of a greatest lower bound is a formal consequence of Thm. 5.3. However, the calculation of $\sqcap Y$ has to be done. Choose any $y \in Y$ then $\mathrm{Im}(y) \supseteq \mathrm{alphinf}(\mathrm{Re}(y)) \supseteq \mathrm{alphinf}(r)$. Hence (r, A) is an α-trace. Clearly it is a lower bound. Assume $(s, B) \sqsubseteq y$ for all $y \in Y$. Then we have $s \leq r$ and $B \supseteq \mathrm{Im}(y) \cup \mathrm{alph}(s^{-1} \mathrm{Re}(y))$ for all $y \in Y$. Since $s \leq r \leq \mathrm{Re}(y)$, we have $\mathrm{alph}(s^{-1} \mathrm{Re}(y)) \supseteq \mathrm{alph}(r^{-1} \mathrm{Re}(y))$ and hence the corollary. $\square$

Theorem 5.5 *An α-trace $(t, A) \in \mathbb{F}_\alpha(\Sigma, D)$ is compact if and only if t is a finite trace. Moreover, the poset $(\mathbb{F}_\alpha(\Sigma, D), \sqsubseteq)$ is an algebraic coherently complete partial order, in particular it is a domain.*

Proof: Let $x = (r, A)$ be any α-trace. It is easy to see that the set of finite α-traces

$$\{ (p, \mathrm{alph}(p^{-1}r) \cup A) \mid p \leq r \text{ and } p \text{ is finite} \}$$

is directed. Moreover for some $p \leq r$ which is large enough we have $\mathrm{alph}(p^{-1}r) = \mathrm{alphinf}(r)$. Since $\mathrm{alphinf}(r) \subseteq A$ by definition, we see from the formula in Thm. 5.3 and [GR93, Prop. 3.3] that x is the least upper bound of this directed set of finite α-traces. Hence compact elements must be finite. To see that finite α-traces are compact we consider any directed set $Y \subseteq \mathbb{F}_\alpha(\Sigma, D)$ and $x = (t, A) \sqsubseteq \sqcup Y$ such that t is finite. Since $\mathrm{Re}(Y)$ is a directed set of $(\mathbb{R}(\Sigma, D), \leq)$ and t is finite, we have $t \leq \mathrm{Re}(y)$ for some $y \in Y$. Since Y is directed we may in fact assume that $t \leq \mathrm{Re}(y)$ for all $y \in Y$. Next we replace Y by the set $Y' = \{(t, \mathrm{alph}(t^{-1} \mathrm{Re}(y)) \cup \mathrm{Im}(y)) \mid y \in Y\}$. The new set Y' is still directed and we have $x \sqsubseteq \sqcup Y'$. Since for all $y' \in Y'$ there exists some $y \in Y$ with $y' \sqsubseteq y$, it is enough to find some $y' \in Y'$ such that $x \sqsubseteq y'$. Now, the set Y' is finite, hence $\sqcup Y' \in Y'$ and the claim follows. Since we have just seen that every α-trace is approximated by finite (compact) elements, we conclude that $(\mathbb{F}_\alpha(\Sigma, D), \sqsubseteq)$ is algebraic. $\square$

For a finite trace $t \in \mathbb{M}(\Sigma, D)$ we denote by $\max(t)$ the (pairwise independent) set of labels of maximal elements in the dependence graph of t. Following Viennot [Vie86] we call a finite trace with exactly one maximal element a *pyramid*.

The next lemma characterizes prime elements. Recall that prime elements are compact. Hence only finite α-traces can be prime.

Proposition 5.6
i) *The α-trace $(1, A) \in \mathbb{F}_\alpha(\Sigma, D)$, $A \subseteq \Sigma$ is prime if and only if $|\Sigma \setminus A| = 1$.*
ii) *Let $(t, A) \in \mathbb{F}_\alpha(\Sigma, D)$, $t \neq 1$, be an α-trace. Then (t, A) is prime if and only if t is a pyramid (hence finite) and $\Sigma = \max(t) \cup A$.*

Proof:
i) Let $x = (1, A)$ with $A \subseteq \Sigma$. If $|\Sigma \setminus A| \geq 2$, then choose $a, b \in \Sigma \setminus A$, $a \neq b$. We have $(1, A) = \sqcup\{(1, A \cup \{a\}), (1, A \cup \{b\})\}$, but neither $(1, A) \sqsubseteq (1, A \cup \{a\})$ nor $(1, A) \sqsubseteq (1, A \cup \{b\})$. Hence $(1, A)$ is not prime in this case. If $|\Sigma \setminus A| = 0$, then $A = \Sigma$ and $(1, \Sigma) = \sqcup\emptyset$ is not prime. Finally, if $|\Sigma \setminus A| = 1$ and $x \sqsubseteq \sqcup X$ for some coherent set X, then, from the formula in Thm. 5.3, there exists some $y \in X$ with $A \supseteq \mathrm{Im}(y)$. It follows that $x \sqsubseteq y$. Hence x is prime.
ii) Let $x = (t, A)$ with $t \in \mathbb{M}(\Sigma, D)$, $t \neq 1$ and $A \subseteq \Sigma$. Consider first the case where $A = \Sigma$. Then we have $x \sqsubseteq \sqcup X$ if and only if $t \leq \sqcup_{\leq} \mathrm{Re}(X) \in \mathbb{R}(\Sigma, D)$. It follows (e.g. from [GR93]) that (t, Σ) is prime if and only if t is a pyramid. If $|\Sigma \setminus A| \geq 2$, then a similar argument as in i) applies showing that (t, A) is not prime. Therefore we may assume that $|\Sigma \setminus A| = 1$ and $A \cup \{a\} = \Sigma$ for some $a \in \Sigma$.

Now, let first t be a pyramid such that $\max(t) = \{a\}$. Consider a coherent set $X \subseteq \mathbb{F}_\alpha(\Sigma, D)$ such that $(w, E) = \sqcup X$ and $(t, A) \sqsubseteq (w, E)$. Knowing $a \notin E$ there has to be some $(s, C) \in X$ such that $a \notin C$. If we have $t \leq s$ then $A \supseteq \mathrm{alph}(t^{-1}w) \cup C \supseteq \mathrm{alph}(t^{-1}s) \cup C$ and hence $(t, A) \sqsubseteq (s, C)$. So assume that $a \notin C$ and that t is not a prefix of s. Applying the well-known generalization of Levi's Lemma, see e.g. [Die90], we can factorize $t = pu$, $s = pv$ and $w = puvq$ such that $\mathrm{alph}(u) \times \mathrm{alph}(v) \subseteq I$. In addition, since t is a pyramid, we have $a \in \mathrm{alph}(u)$. Since $(s, C) \sqsubseteq (w, E) = (suq, E)$ we have in particular $C \supseteq \mathrm{alph}(uq) \ni a$ which is a contradiction. Hence, if t is a pyramid and $A \cup \max(t) = \Sigma$, then (t, A) is prime.

There remain two cases: either t is not a pyramid or $A \cup \max(t) \neq \Sigma$. In both cases it is easy to see that (t, A) is not prime. This is left to the reader. $\square$

The following theorem summarizes the results above. It is the main theorem of this section. The proof uses the characterization of prime α-traces, Prop. 5.6. It is left to the reader.

Theorem 5.7 *The poset $(\mathbb{F}_\alpha(\Sigma, D), \sqsubseteq)$ is a prime algebraic coherently complete domain.*

Let us point out that δ-traces do not form a prime algebraic domain in general. We have the following counterexample

Example 5.8 *Let* $(\Sigma, D) =$ 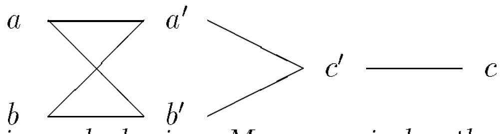 *. Then, the poset* $(\mathbb{F}_\delta(\Sigma, D), \sqsubseteq)$ *is not prime algebraic. More precisely, the* δ*-trace* $x = (ab, D(c)) \in \mathbb{F}_\delta(\Sigma, D)$ *is not the least upper bound of its prime approximations.*

Although we have not yet been able to characterize those dependence alphabets where $\mathbb{F}_\delta(\Sigma, D)$ is prime algebraic, we can state the following theorem.

Theorem 5.9 *The poset* $\mathbb{F}_\delta(\Sigma, D)$ *of* δ*-traces is an algebraic coherently complete domain. In general it is not prime algebraic.*

The next theorem states that in both domains the concatenation is continuous with respect to the approximation order. Its proof can be based on Prop. 5.1. Later we will show an even stronger result, therefore the proof can be omitted.

Theorem 5.10 *Let* $X, Y \subseteq \mathbb{F}_\alpha(\Sigma, D)$ *(*$\mathbb{F}_\delta(\Sigma, D)$ *resp.) be two directed sets of* α*-traces (*δ*-traces resp.). Then* $X \cdot Y$ *is directed and we have* $\sqcup(X \cdot Y) = (\sqcup X) \cdot (\sqcup Y)$*.*

6 Topology

In this section we define a natural ultrametric of $\mathbb{F}_\alpha(\Sigma, D)$ and we show that it induces the Lawson topology. The length of an α-trace $x = (t, A)$ is the length of its real part. It is denoted by $|x|$. The same notation is used for real and δ-traces, too. The definition below follows the following scheme. We define a function l to $\mathbb{N} \cup \{\infty\}$ and putting $d(x, y) = 2^{-l(x,y)}$ as resulting ultrametric.

Definition 6.1 *Let* x*,* y *be real (*α*- resp.,* δ*- resp.) traces. Define*

$$\begin{aligned}
l_{\mathbb{R}}(x, y) &= \sup \{\, n \mid \forall p \in \mathbb{R}(\Sigma, D), |p| \leq n : p \leq x \Leftrightarrow p \leq y \,\} \\
l_\alpha(x, y) &= \sup \{\, n \mid \forall p \in \mathbb{F}_\alpha(\Sigma, D), |p| < n : p \sqsubseteq x \Leftrightarrow p \sqsubseteq y \,\} \\
l_\delta(x, y) &= \sup \{\, n \mid \forall p \in \mathbb{F}_\delta(\Sigma, D), |p| < n : p \sqsubseteq x \Leftrightarrow p \sqsubseteq y \,\}
\end{aligned}$$

The resulting ultrametrics are denoted by $d_{\mathbb{R}}$*,* d_α *and* d_δ*.*

Remark 6.2 The metric $d_{\mathbb{R}}$ for real traces is the same as in [Kwi90]. The metrics d_α and d_δ have been studied for α-complex and complex traces in [Die93b].

In the following we study only d_α for α-traces. Analogous results for $d_{\mathbb{R}}$ and real traces are well known. The extension to δ-traces remains to the reader.

Definition 6.3 *Let $x \in \mathbb{F}_\alpha(\Sigma, D)$ be an α-trace. The n-th* open ball $B(x,n)$ *with center x is defined by*

$$B(x,n) \;=\; \{\, y \in \mathbb{F}_\alpha(\Sigma, D) \mid l_\alpha(x,y) > n \,\} \;.$$

The n-th approximation $x[n]$ *of x is defined by*

$$x[n] \;=\; \sqcup \{\, y \in \mathbb{F}_\alpha(\Sigma, D) \mid y \sqsubseteq x \text{ and } |y| \le n \,\} \;.$$

Lemma 6.4 *Let $x, y \in \mathbb{F}_\alpha(\Sigma, D)$, then we have:*
i) $B(x,n) = B(x[n],n)$
ii) $B(x,n) \cap B(y,n) \neq \emptyset$, *if and only if* $x[n] = y[n]$
iii) *The open ball $B(x,n)$ is closed, too.*

Proposition 6.5 *The set of finite α-traces $\{(t,A) \mid t \in \mathbb{M}(\Sigma, D), A \subseteq \Sigma\}$ is an open, discrete and dense subset of $(\mathbb{F}_\alpha(\Sigma, D), d_\alpha)$.*

Proof: It follows by Lem. 6.4 that the set of finite α-traces is dense. To see that it is open and discrete observe that $B((t,A),n) = \{(t,A)\}$ is a one-point set for all $n > |t|$. □

The Lawson topology is a refinement of the Scott topology for a domain such that on one side the topology becomes Hausdorff, on the other side (for bounded complete domains) it is still compact, see [GHK+80, Chap. III, Thm. 1.10 and Exercises].

Let us first recall the definition of the Lawson topology for a domain $(Z, \le)$. In this case, the family

$$\mathcal{S} = \{\, \uparrow k \mid k \text{ compact} \,\} \cup \{\, Z \setminus \uparrow k \mid k \text{ compact} \,\}$$

is a sub-basis for the Lawson topology on $(Z, \le)$. This means the open sets are precisely the unions of finite intersections of the sub-basis.

Proposition 6.6 *The metric d_α induces the Lawson topology on $\mathbb{F}_\alpha(\Sigma, D)$, which is defined with respect to the approximation order $\sqsubseteq$.*

Proof: Let $B(x,n)$ be an open ball. It is easy to verify that $B(x,n) = \uparrow x[n] \setminus \uparrow \{z[n] \mid z[n] \not\sqsubseteq x[n]\}$. Since finite α-traces are compact, it follows that every d_α-open set is Lawson open.

For the other direction let $k \in \mathbb{F}_\alpha(\Sigma, D)$ be compact. Then, for all open ball $B(x,n)$, either $B(x,n) \subseteq \uparrow k$ or $B(x,n) \cap \uparrow k = \emptyset$. Therefore, $\uparrow k$ is both d_α-open and d_α-closed. Hence, all Lawson open sets are open in $(\mathbb{F}_\alpha(\Sigma, D), d_\alpha)$. □

Since $\mathbb{F}_\alpha(\Sigma, D)$ is coherently complete, it is bounded complete. Therefore, the Lawson topology yields a compact space. From general facts about the Lawson topology we deduce the following:

Theorem 6.7
i) *The metric space* $(\mathbb{F}_\alpha(\Sigma, D), d_\alpha)$ *is compact and complete.*
ii) *Let* $(x_n)_{n\geq 0} \subseteq \mathbb{F}_\alpha(\Sigma, D)$ *be a* $\sqsubseteq$*-increasing sequence. Then we have*

$$\lim_{n\to\infty} x_n \;=\; \sqcup\{\, x_n \mid n \geq 0 \,\}$$

Our final result shows that the concatenation of α-traces is uniformly continuous with respect to the metric d_α. By Prop. 5.1 and Thm. 6.7 this is a stronger result than Thm. 5.10.

Theorem 6.8 *The concatenation of* α*-traces is uniformly continuous. More precisely, let* $x, x', y, y' \in \mathbb{F}_\alpha(\Sigma, D)$ *be* α*-traces. Then*

$$d_\alpha(xy, x'y') \;\leq\; \max(d_\alpha(x, x'), d_\alpha(y, y'))\ .$$

Proof: For the proof it is enough to show that $(x \cdot y)[n] = (x[n] \cdot y[n])[n]$ for all $x, y \in \mathbb{F}_\alpha(\Sigma, D)$, $n \geq 0$. Let $x = (r, A)$, $y = (s, B)$, and $z = (t, C)$ with $|t| < n$. We have $xy = (r\mu_A(s), \sigma_A(s) \cup A \cup B)$.

First, let $z \sqsubseteq xy$. Then, by the generalization of Levi's Lemma (see e.g. [Die90]), we have $t = pv$, $r = pu$, $\mu_A(s) = vq$ such that $\mathrm{alph}(u) \times \mathrm{alph}(v) \subseteq I$, $C \supseteq \mathrm{alph}(u) \cup \mathrm{alph}(q) \cup \sigma_A(s) \cup A \cup B$. Note that $(p, \mathrm{alph}(u) \cup A) \sqsubseteq x[n]$ and $(v, \mathrm{alph}(q) \cup \sigma_A(s) \cup B) \sqsubseteq y[n]$. We obtain by the calculus of α-traces and Prop. 5.1 the following formula

$$\begin{aligned} z = (t, C) \;&\sqsubseteq\; (t, \mathrm{alph}(u) \cup \mathrm{alph}(q) \cup \sigma_A(s) \cup A \cup B) \\ &=\; (p, \mathrm{alph}(u) \cup A) \cdot (v, \mathrm{alph}(q) \cup \sigma_A(s) \cup B) \\ &\sqsubseteq\; x[n] \cdot y[n]\ . \end{aligned}$$

Since $|t| < n$, we have $z \sqsubseteq (x[n] \cdot y[n])[n]$. It follows $(x \cdot y)[n] \sqsubseteq (x[n] \cdot y[n])[n]$.

The other direction is trivial, since again by monotonicity we have $x[n] \cdot y[n] \sqsubseteq x \cdot y$ and it follows $(x[n] \cdot y[n])[n] \sqsubseteq (x \cdot y)[n]$. $\square$

References

[AR88] IJ.J. Aalbersberg and G. Rozenberg. Theory of traces. *Theoretical Computer Science*, 60:1–82, 1988.

[Die90] V. Diekert. *Combinatorics on Traces.* Number 454 in Lecture Notes in Computer Science. Springer, Berlin-Heidelberg-New York, 1990.

[Die93a] V. Diekert. Complex and complex-like traces. In A. Borzyszkowski et al., editor, *Proceedings of the 18th Mathematical Foundations of Computer Science (MFCS 93), Gdansk (Polen) 1993*, number 711 in Lecture Notes in Computer Science, pages 68–82, Berlin-Heidelberg-New York, 1993. Springer. Invited Lecture.

[Die93b] V. Diekert. On the concatenation of infinite traces. *Theoretical Computer Science*, 113:35–54, 1993. Special issue STACS 91.

[DR95] V. Diekert and G. Rozenberg, editors. *The Book of Traces.* World Scientific, Singapore, 1995. to appear.

[Gas90] P. Gastin. Infinite traces. In I. Guessarian, editor, *Proceedings of the Spring School of Theoretical Computer Science on Semantics of Systems of Concurrent Processes*, number 469 in Lecture Notes in Computer Science, pages 277–308, Berlin-Heidelberg-New York, 1990. Springer.

[GHK+80] G. Gierz, K.H. Hofmann, K. Keimel, Lawson J.D., M. Mislove, and D. Scott. *A Compendium of Continuous Lattices.* Springer, 1980.

[GP92] Paul Gastin and Antoine Petit. Poset properties of complex traces. In I. M. Havel and V. Koubek, editors, *Proceedings of the 17th Symposium on Mathematical Foundations of Computer Science (MFCS'92), Prague, (Czechoslovakia), 1992*, number 629 in Lecture Notes in Computer Science, pages 255–263, Berlin-Heidelberg-New York, 1992. Springer.

[GP95] P. Gastin and A. Petit. Infinite traces. In G. Rozenberg and V. Diekert, editors, *The Book of Traces*, pages 393–486. World Scientific, Singapore, 1995.

[GR93] P. Gastin and B. Rozoy. The poset of infinitary traces. *Theoretical Computer Science*, 120:101–121, 1993.

[Hoa85] C. A. R. Hoare. *Communicating Sequential Processes.* Prentice-Hall International, London, 1985.

[Kwi90] M. Kwiatkowska. A metric for traces. *Information Processing Letters*, 35:129–135, 1990.

[Law88] Jimmie D. Lawson. The versatile continuous order. In M. Main, A. Melton, M. Mislove, and D. Schmidt, editors, *Proceedings of the Workshop on Mathematical Foundations of Programming Language Semantics*, number 298 in Lecture Notes in Computer Science, pages 134–160, Berlin-Heidelberg-New York, 1988. Springer.

[Maz87] A. Mazurkiewicz. Trace theory. In W. Brauer et al., editors, *Petri Nets, Applications and Relationship to other Models of Concurrency*, number 255 in Lecture Notes in Computer Science, pages 279–324, Berlin-Heidelberg-New York, 1987. Springer.

[Mil80] R. Milner. *A Calculus of Communicating Systems.* Number 92 in Lecture Notes in Computer Science. Springer, Berlin-Heidelberg-New York, 1980.

[Per89] D. Perrin. Partial commutations. In *Proceedings of the 16th International Colloquium on Automata, Languages and Programming (ICALP'89), Stresa (Italy) 1989*, number 372 in Lecture Notes in Computer Science, pages 637–651, Berlin-Heidelberg-New York, 1989. Springer.

[RW94] Arend Rensink and Heike Wehrheim. Weak sequential composition in process algebras. In Bengt Jonssen and Joachim Parrow, editors, *Proceedings of the Fifth International Conference on Concurrency Theory CONCUR'94, Uppsala (Sweden)*, number 836 in Lecture Notes in Computer Science, pages 226–241, Berlin-Heidelberg-New York, 1994. Springer.

[Vie86] X. G. Viennot. Heaps of pieces I: Basic definitions and combinatorial lemmas. In G. Labelle et al., editors, *Proceedings Combinatoire énumerative, Montréal, Québec (Canada) 1985*, number 1234 in Lecture Notes in Mathematics, pages 321–350, Berlin-Heidelberg-New York, 1986. Springer.