Resource Traces: A Domain for Processes Sharing Exclusive Resources [*]

Paul Gastin and Dan Teodosiu

Institut Blaise Pascal, LITP
2, place Jussieu, F-75252 Paris Cedex 05
{teodosiu,gastin}@litp.ibp.fr

Abstract

The theory of finite and infinite words with explicit termination is commonly used to give denotational semantics for algebras of sequential processes. In this paper we generalize this theory to the case of partially terminated processes synchronizing on a fixed set of shared exclusive resources.

We start with an alphabet of atomic actions, a set of resources, and a resource map assigning to each action the non-empty subset of resources it uses. Actions that do not share common resources are considered to be independent.

The specification we use for partially terminated processes is similar to failure semantics. It consists of two components: an already observed part represented as an action-labeled event structure (Mazurkiewicz trace), and a guard set containing the resources granted to the process for its further development. A process concatenation is then defined such that independent actions can be dispatched concurrently. Specification refinement leads to the definition of a natural approximation ordering between processes which generates a coherently complete prime algebraic Scott domain, where process concatenation is continuous in both arguments. Furthermore, we define a natural ultrametric on processes based on prefix information. The induced topology is shown to be equivalent to the compact Lawson topology generated by the approximation ordering. Moreover, process concatenation is shown to be uniformly continuous with respect to the defined ultrametric.

We develop a mathematical theory which perfectly extends the central properties of the domain of finite and infinte words with explicit termination and the domain of finite and infinite Mazurkiewicz traces. Its natural semantics is well suited to the design of modular specifications for processes sharing exclusive resources such as instruction sets using some set of registers (PRAM) or concurrent sequential processes sharing exclusive communication channels (CSP).

1 Introduction

The theory of Mazurkiewicz traces has been recognized over the last decade as an expressive tool for investigating concurrent systems according to the inter-

---

[*] This research has been supported by the PROCOPE programme No 95143 ASTRA

leaving viewpoint (overviews can be found in [Maz87, AR88, Per89, Die90] and in the recent monograph [DR95]). The idea is to start with a finite alphabet Σ of actions and an explicit symmetric and irreflexive independence relation $I \subseteq \Sigma \times \Sigma$ specifying those pairs of actions that can be run concurrently. The complement relation $D = \Sigma \times \Sigma \setminus I$ is called dependence relation. The intended semantics is that performing two independent actions in any order should lead to the same result. Therefore, finite sequences of actions which can be mutualy transformed into each other by commuting consecutive independent actions are considered to represent one and the same concurrent computation. Thus, finite concurrent processes are mathematically described as elements of the quotient monoid $\mathbb{M}(\Sigma, D) = \Sigma^*/\{\, ab = ba \mid (a, b) \in I \,\}$, called the monoid of Mazurkiewicz traces over the dependence alphabet (Σ, D). The concatenation operation on traces can be regarded as a weak sequential composition: only dependent actions are synchronized whereas at the process level no explicit synchronization is performed, thus retaining the maximal possible concurrency for the resulting process. This accounts in fact very well for both sequential composition, in the word case of full dependency $D = \Sigma \times \Sigma$, and parallel composition, in the vector-addition case of total independency $I = \Sigma \times \Sigma \setminus \mathbb{1}_\Sigma$.

Natural frames for defining dependence alphabets mostly come about when considering actions that share exclusive resources from some fixed set $\mathcal{R}$. This can be expressed by a resource map $\mathrm{res} : \Sigma \to \mathcal{P}(\mathcal{R}) \setminus \{\emptyset\}$ assigning to each action the non-empty set of resources it uses. The induced dependence relation is then implicitly defined by setting two actions dependent if they share some common resource, i.e. $a \ \mathrm{D} \ b$ iff $\mathrm{res}(a) \cap \mathrm{res}(b) \neq \emptyset$. Typical examples are e.g. instruction sets using some set of registers (PRAM), concurrent sequential processes sharing exclusive communication channels (CSP) and string rewriting systems over a set of symbols. Obviously, a dependence relation can have different representations by resource maps. Since the information we need in order to perform concatenation restricts to dependency, one normally abstracts away the exact representation and retains only the induced dependency.

In order to use the concatenation of processes as a modular operation for specifying concurrent systems, we should be able to use some refinement technique compatible with process concatenation, such that refining the processes in a concatenation should lead to a refinement of the result. The main technique used to this end is to go over to partially terminated processes consisting of an already observed part and some restriction on the continuation of the process, which guarantees the continuity of the concatenation. This technique has been applied in the setting of trace theory in [Die93a, DG95], where the models of α- and δ-traces are proposed as generalizations of the classical completion for sequential processes. The approach is in fact similar to the failure (or ready) semantics used for process algebras like CCS or CSP [Hoa85, Mil80]. The difference is that whereas in failure semantics the behaviour is specified merely for the next step, in the trace models the specification applies to the whole future of the process.

The model of resource traces proposed in this paper follows the general line of thought developed in [DG95]. However, it uses in an essential way the representation of dependence alphabets in terms of resources for defining the approximation ordering between processes, corresponding to specification refinement. This approximation ordering is shown to have good completeness and algebraicity properties, thus defining an appropriate domain structure. Using the approximation ordering we furthermore define an ultrametric distance between processes. The induced topology is shown to coincide with the compact Lawson topology [Law88] associated with the approximation ordering. In addition, the concatenation operation proves to be continuous with respect to the approximation ordering, and even uniformly continuous with respect to the defined ultrametric, which enables its use as a modular operation for specifying concurrent systems.

We use a rich mathematical arsenal comprising techniques from algebra, domain theory, and topology in order to present a generalisation of the theory of finite and infinite words which proves to be adequate from an algebraic as well as from an order-theoretic perspective.

2 Domains

In this section we briefly introduce some standard terminology of domain theory in order to lay down the context. Subsequently, we prove two general domain-theoretic propositions which will substantially clarify our further investigation.

A *partial order* $(X, \leq)$ is a set X equiped with a reflexive, antisymmetric and transitive binary relation $\leq$. If $x \leq y$, we shall say that x is *less than* or *below* y and that y is *greater than* or *above* x. The *lower set* (*upper set* resp.) of a subset $Y \subseteq X$ is $\downarrow Y = \{x \in X \mid x \leq y \text{ for some } y \in Y\}$ ($\uparrow Y = \{x \in X \mid x \geq y \text{ for some } y \in Y\}$ resp.). A subset $Y \subseteq X$ is *downwards closed* (*upwards closed* resp.) iff $\downarrow Y = Y$ ($\uparrow Y = Y$ resp.).

A subset $Y \subseteq X$ is *upper bounded* (*lower bounded* resp.) by an element $x \in X$ iff $y \leq x$ ($x \leq y$ resp.) for all $y \in Y$. It is *directed* (*coherent* resp.) iff it is non-empty and for all $x, y \in X$ there exists $z \in Y$ ($z \in X$ resp.) such that $x \leq z$ and $y \leq z$. Whenever they exist the *least upper bound* and the *greatest lower bound* of a subset $Y \subseteq X$ will be denoted by $\bigvee Y$ and $\bigwedge Y$, respectively.

The partial order $(X, \leq)$ is a *coherently complete partial order* (CCPO), *directed complete partial order* (DCPO), *bounded complete partial order* (BCPO) iff every coherent subset, directed subset, bounded subset, respectively, has a least upper bound.

A BCPO $(X, \leq)$ is *bounded distributive* iff for all bounded subsets $Y \cup \{x\} \subseteq X$, we have $x \wedge (\bigvee Y) = \bigvee (x \wedge Y)$ where $x \wedge Y = \{x \wedge y \mid y \in Y\}$.

An element $x \in X$ is *irreducible* (*prime* resp.) iff for all subsets $Y \subseteq X$ having a least upper bound, $x = \bigvee Y$ ($x \leq \bigvee Y$ resp.) implies $x \leq y$ for some $y \in Y$. It is *compact*, iff for all directed subsets $Y \subseteq X$ having a least upper bound, $x \leq \bigvee Y$ implies $x \leq y$ for some $y \in Y$. The set of all irreducible, prime, compact elements in X will be denoted by Irr, Prm, Kmp, respectively, the sets of those below $x \in X$ by $\mathrm{Irr}(x)$, $\mathrm{Prm}(x)$, $\mathrm{Kmp}(x)$, respectively.

The partial order $(X, \leq)$ is *i-algebraic* (*p-algebraic* resp.) iff $x = \bigvee \mathrm{Irr}(x)$ ($x = \bigvee \mathrm{Prm}(x)$ resp.) for all $x \in X$. It is *k-algebraic* iff $\mathrm{Kmp}(x)$ is directed and $x = \bigvee \mathrm{Kmp}(x)$ for all $x \in X$. A k-algebraic DCPO is called a Scott domain.

A mapping $F : (X, \leq) \to (X', \leq')$ between two partial orders is *monotone* iff for all $x, y \in X$, $x \leq y$ implies $F(x) \leq' F(y)$. It is *continuous* iff it is monotone and for all directed sets $Y \subseteq X$, such that $\bigvee Y$ exists, $\bigvee' F(Y)$ exists, and $F(\bigvee Y) = \bigvee' F(Y)$. A pair of mappings $E : (X, \leq) \to (X', \leq')$, $P : (X', \leq) \to (X, \leq')$ between two partial orders is an *embedding-projection pair* iff both mappings are monotone and $P \circ E = \mathbb{1}_X$, $E \circ P \leq' \mathbb{1}_{X'}$. E is then called the *embedding* and P the *projection.*

Let us stop for an instant and comment on the above definitions. It follows directly from the definitions that any directed or bounded set is coherent. Therefore, CCPOs are DCPOs and BCPOs at the same time. Furthermore, we see that primes are irreducible and compact, so p-algebraicity implies i-algebraicity and, in the case of BCPOs, k-algebraicity. The later statement is a consequence of the fact that in a p-algebraic BCPO, compacts are precisely finite suprema of primes and for all $x \in X$, $\mathrm{Kmp}(x) = \{\bigvee P \mid P \text{ finite and } P \subseteq \mathrm{Prm}(x)\}$ which is indeed directed.

Our models turn out to be p-algebraic coherently complete partial orders (CCPOs), hence enjoy the most powerful properties of completeness and algebraicity stated above. Whereas coherent completeness is fairly easy to prove, p-algebraicity is rather difficult. We shall, therefore, exploit to this end another property shared by all our models through their very construction, namely well-foundedness. A partial order is *well-founded* iff every descending sequence of elements is finite or, equivalently, iff every non-empty subset has a minimal element. As we shall see in a moment, i-algebraicity is automatically satisfied under this assumption (Proposition 2.1), so we are left with showing that all irreducibles are prime in order to conclude p-algebraicity. This implication is then further reduced to checking bounded distributivity (Proposition 2.2). The latter task can finally be easily accomplished within our models, by using the form of the $\bigwedge$ and $\bigvee$ operators on bounded subsets. We now proceed to the details.

Proposition 2.1 *Every well-founded partial order is i-algebraic.*

Proof: Let $(X, \leq)$ be a well-founded partial order. We have to prove that every element $x \in X$ is a supremum of irreducible elements. We use well-founded induction. Assume that every element strictly below x is a supremum of irreducibles. We distinguish two cases. Either x itself is irreducible, in which case the one-element set $\{x\}$ will do, or otherwise x is not irreducible. Then, there must exist by definition a subset $Y \subseteq X$ having as supremum $\bigvee Y = x$ such that $x \notin Y$. But then all elements $y \in Y$ are necessarily strictly below x. Applying now the induction hypothesis, we obtain for every $y \in Y$ a subset I_y of irreducibles having as supremum $\bigvee I_y = y$. Taking the union of all these subsets ranging over $y \in Y$ yields a subset of irreducibles whose supremum is x, thus concluding the induction proof. □

It is well-known that a BCPO is allways bounded distributive if it is p-algebraic. The next proposition shows that additionally assuming well-foundedness makes the two conditions equivalent. As a matter of fact, this is also derivable from previous work of Winskel on prime event structures. We do not go here into details and choose instead to present the following simple proof.

Proposition 2.2 *A well-founded BCPO is p-algebraic iff it is bounded distributive.*

Proof: Let $(X, \leq)$ be a well-founded BCPO.

We first prove the easy direct implication. Consider any $x \in X$, $Y \subseteq X$ such that $\{x\} \cup Y$ is bounded. First, for all $y \in Y$ we have $x \wedge y \leq x \wedge (\bigvee Y)$, which implies $\bigvee(x \wedge Y) \leq x \wedge (\bigvee Y)$. Secondly, for any prime $p \leq x \wedge (\bigvee Y)$ we have $p \leq x$ and $p \leq \bigvee Y$, hence $p \leq y$ for some $y \in Y$. Therefore, $p \leq x \wedge y \leq \bigvee(x \wedge Y)$. Since $(X, \leq)$ is p-algebraic it follows that $x \wedge (\bigvee Y) = \bigvee \mathrm{Prm}(x \wedge (\bigvee Y))$ $\leq \bigvee(x \wedge Y)$.

We now prove the converse implication. In view of the preceeding proposition, all we have to show is that irreducibles are prime. So let $x \in X$ be irreducible and suppose that $x \leq \bigvee Y$ for some set $Y \subseteq X$. Then $\bigvee Y$ is a common upper bound to x and all elements of Y. But then $x = x \wedge (\bigvee Y) = \bigvee(x \wedge Y)$ by bounded distributivity. Since x is irreducible, it follows that there exists $y \in Y$, such that $x = x \wedge y$, which in turn implies $x \leq y \in Y$, showing thus that x must be prime. $\square$

3 Traces

In this section we review the basics of trace theory and stress upon some results necessary to the subsequent presentation of our model. We assume the reader to be familiar with the notions involved and, therefore, only sketch the general ideas. Extensive treatments of the subject can be found in [Maz87, AR88, Per89, Die90, DR95].

We fix in the following a finite *alphabet of actions* Σ and a symmetric and irreflexive relation I on Σ called *independence relation*. The symmetric and reflexive complement of I in $\Sigma \times \Sigma$ is denoted by D and called *dependence relation*. The pairs (Σ, D) and (Σ, I) are called *dependence alphabet* and *independence alphabet*, respectively. For $A \subseteq \Sigma$ we denote by $\mathrm{D}(A)$ the set of actions dependent on some action in A and by $\mathrm{I}(A)$ the set of actions independent from all actions in A.

A *trace* over (Σ, D) is (the isomorphism class of) a labeled directed graph $x = [V, E, \lambda]$, where V is a countable set of *events*, E is a well-founded *synchronization relation* on V, and $\lambda : V \to \Sigma$ is a *node-labelling* satisfying the *well-synchronization condition*

$$(\lambda(p), \lambda(q)) \in D \Longleftrightarrow (p, q) \in E \text{ or } (q, p) \in E \text{ or } p = q$$

for any $p, q \in V$. The *length* of the trace x is $|x| = |V|$. A trace is *finite* iff its length is finite. The *set of traces* over (Σ, D) is denoted by $\mathbb{G}(\Sigma, D)$ or simply $\mathbb{G}$, the *set of finite traces* by $\mathbb{M}(\Sigma, D)$ or simply $\mathbb{M}$.

As E is assumed to be well-founded, it will necessarily be acyclic. Hence, its reflexive-transitive closure E^* defines a well-founded *partial ordering on the* *events* of the trace x. The *past* $\downarrow v$ of the event $v \in V$ is the set of all events below v with respect to the ordering E^*. The *set of maximal actions* of the trace x, denoted by $\max(x)$, is the set of labels of maximal events in x with respect to the ordering E^*. We define the *alphabet* of the trace x by $\mathrm{alph}(x) = \lambda(V)$ and set $\mathrm{D}(x) = \mathrm{D}(\mathrm{alph}(x))$ and $\mathrm{I}(x) = \mathrm{I}(\mathrm{alph}(x))$.

The set of traces $\mathbb{G}$ forms a monoid together with the *concatenation* defined by $[V_1, E_1, \lambda_1] \cdot [V_2, E_2, \lambda_2] = [V, E, \lambda]$, where $V = V_1 \dot\cup V_2$, $\lambda = \lambda_1 \dot\cup \lambda_2$, $E = E_1 \dot\cup E_2 \dot\cup (V_1 \times V_2 \cap \lambda^{-1}(D))$, and the *empty trace* $1 = [\emptyset, \emptyset, \emptyset]$ as neutral element. The set of finite traces $\mathbb{M}$ is a countable submonoid of $\mathbb{G}$. It is actually isomorphic to the factor monoid obtained from Σ^* by setting two words equivalent, if they can be mutually transformed into each other by commuting consecutive independent actions, that is $\mathbb{M}(\Sigma, D) = \Sigma^*/\{\, ab = ba \mid (a, b) \in I \,\}$.

The monoid $\mathbb{G}$ is left-cancellative (i.e. $xy = xz$ implies $y = z$) and group free (i.e. $xy = 1$ implies $x = y = 1$). The reflexive and transitive *prefix relation*, defined by $x \leq y$ iff there exists z such that $xz = y$, is therefore antisymmetric, hence defines a partial order on $\mathbb{G}$. Moreover, for any traces $x \leq y$ we can define $x^{-1}y$ by the equation $x(x^{-1}y) = y$. For all $x \leq y \leq z$ we then have $x^{-1}z = (x^{-1}y)(y^{-1}z)$, which is the standard usage of this notation. Furthermore, as a consequence of the well-known Levi Lemma, we have $(x \wedge y)^{-1}y = x^{-1}(x \vee y)$ for any traces x and y having a common upper bound.

One can easily see that the prefix ordering must be well-founded. The properties of $\mathbb{M}$ and $\mathbb{G}$ with respect to the prefix ordering have been investigated in [Maz87, Gas90, Kwi90, GR93, GP95], where it has been shown that $(\mathbb{M}, \leq)$ and $(\mathbb{G}, \leq)$ are p-algebraic BCPOs. Since any prefix of a finite trace is finite, $\mathbb{M}$ is a downwards closed subset of $\mathbb{G}$. Moreover, all finite traces can be shown to be compact. Neither of the two monoids is directed complete with respect to the prefix ordering. The reason is in fact purely set-theoretical: in both cases the set of ordinal numbers corresponding to the ordered sets of underlying events does not contain its limit points.

We can nevertheless consider for any trace $x \in \mathbb{G}$ the (bounded) set of all its finite prefixes and extract its *real part*

$$\mathrm{Re}(x) = \bigvee \{\, k \mid k \in \mathbb{M}, k \leq x \,\}.$$

This defines a mapping $\mathrm{Re} : \mathbb{G} \to \mathbb{G}, x \mapsto \mathrm{Re}(x)$. A trace $x \in \mathbb{G}$ is *real* iff it equals its real part. This is equivalent to demanding of all events in the trace to have finite past. The *set of real traces*, denoted by $\mathbb{R}(\Sigma, D)$ or simply $\mathbb{R}$, is a directed complete, downwards closed subset of $\mathbb{G}$ containing all finite traces. The compact real traces are precisely the finite ones whereas the prime real traces are the non-empty finite traces having exactly one maximal vertex. Moreover, $\mathbb{R}$ is a p-algebraic CCPO isomorphic to the ideal completion of $\mathbb{M}$, which shows that $\mathbb{R}$ is a good semantic domain from an order-theoretic perspective. Unfortunately, we completely loose the algebraic structure in passing over to $\mathbb{R}$, as concatenation is not an internal operation on $\mathbb{R}$.

An even more penalizing deficiency, from a modular point of view, is caused by the fact that the algebraic and order structures of $\mathbb{M}$ and $\mathbb{G}$ lack good interconnection properties, since concatenation is not monotone with respect to the prefix ordering. This is even true of the classical word case, when all actions are considered to be dependent. The direct consequence of this circumstance is that we can not use the prefix ordering in modular judgements about trace compositions, since increasing the factor traces will not necessary lead to an increase of their composition.

The model we present in the next sections provides a solution to these problems. It refines the approach proposed in [Die93a, DG95], meanwhile avoiding some of its inconveniences which we discuss in the last section.

4 The Calculus of Resource Traces

A solution for making concatenation monotonous (continuous) in the word case has been known for quite a while. It consists in augmenting the alphabet Σ with an explicit termination symbol $\top$ and non-termination symbol $\bot$ and considering partially terminated (infinite) words, defined as elements of $\Sigma^\infty\bot \cup \Sigma^*\top$, where $\Sigma^\infty = \Sigma^* \cup \Sigma^\omega$. Words ending with $\top$ are finite and called *terminated*, those ending with $\bot$ are possibly infinite and called *non-terminated*. The *concatenation* is defined by $u\bot \cdot w = u\bot$ and $v\top \cdot w = vw$ for all $u \in \Sigma^\infty$ $v \in \Sigma^*$ and $w \in \Sigma^\infty\bot \cup \Sigma^*\top$. This is formally equivalent to setting the non-termination symbol $\bot$ to be right-absorbent and the termination symbol $\top$ to be a left unit. $\Sigma^\infty\bot \cup \Sigma^*\top$ is a monoid with respect to the concatenation and $\top$ as unit. An *approximation order* can be introduced by putting $u\bot \sqsubseteq uv$ for any $u \in \Sigma^*$ and $v \in \Sigma^\infty\bot \cup \Sigma^*\top$, which means that only finite non-terminated words can grow. $\Sigma^\infty\bot \cup \Sigma^*\top$ is a p-algebraic coherently complete domain with respect to the approximation order. The compact elements are exactly the finite words, that is the elements of $\Sigma^*\{\bot, \top\}$. Moreover, concatenation can be shown to be monotonous (continuous) with respect to the approximation order. This formalism has been used as a foundation for specifying denotational semantics for process calculi such as CCS and CSP. The model we present tries to extend this line of thought to the trace model of computation.

Let us fix in the following a finite *alphabet of actions* Σ, a finite *set of resources (registers)* $\mathcal{R}$ and a *resource mapping* $\mathrm{res} : \Sigma \to \mathcal{P}(\mathcal{R}) \setminus \{\emptyset\}$ assigning to every action in Σ a non-empty set of resources from $\mathcal{R}$ of which it makes use. The resource mapping res induces an implicit dependence relation D on Σ by setting two actions dependent if they use some common resource, i.e. $a \mathrel{\mathrm{D}} b$ iff $\mathrm{res}(a) \cap \mathrm{res}(b) \neq \emptyset$ for all $a, b \in \Sigma$. The resource mapping res will be called a *representation* of the induced dependence relation D.

We extend the resource mapping to $\mathcal{P}(\Sigma)$ and $\mathbb{G}(\Sigma, D)$ by $\mathrm{res} : \mathcal{P}(\Sigma) \to \mathcal{P}(\mathcal{R})$, $A \mapsto \mathrm{res}(A) = \bigcup_{a \in A} res(a)$ and $\mathrm{res} : \mathbb{G}(\Sigma, D) \to \mathcal{P}(\mathcal{R})$, $x \mapsto \mathrm{res}(x) = \mathrm{res}(\mathrm{alph}(x))$. Furthermore, we define the *resources at infinity* of a trace by

$$\mathrm{resinf}(x) = \bigcap\{\, \mathrm{res}(k^{-1}x) \mid k \in \mathbb{M}, k \leq x \,\}$$

This defines a mapping $\mathbb{G} \to \mathcal{P}(\mathcal{R})$, $x \mapsto \mathrm{resinf}(x)$. For real traces, the resources at infinity of a trace can be identified as being the set of resources used infinitely often by action labellings of events in the trace. For non-real traces, it additionally contains the resources of the non-real part, hence $\mathrm{resinf}(x) =$

$\operatorname{resinf}(\operatorname{Re}(x)) \cup \operatorname{res}(\operatorname{Re}(x)^{-1}x)$. Furthermore, as can be easily shown, we have $\operatorname{resinf}(x) \subseteq \operatorname{resinf}(y) \subseteq \operatorname{resinf}(x) \cup \operatorname{res}(x^{-1}y)$ whenever $x \leq y$.

Definition 4.1 *A* resource trace, *ρ-*trace *for short, over $(\Sigma, \mathcal{R}, \operatorname{res})$ is a pair $x = (r, R)$, where $r \in \mathbb{R}(\Sigma, D)$ is a real trace over (Σ, D), and $R \subseteq \mathcal{R}$ is a subset of resources, such that* $\operatorname{resinf}(r) \subseteq R$. *The trace r is denoted by* $\operatorname{Re}(x)$ *and called the* real part *of x. The length $|x|$ of x is the length $|r|$ of its real part. The set R is denoted by* $\operatorname{Im}(x)$ *and called the* imaginary part *of x. The* set of resource traces *over $(\Sigma, \mathcal{R}, \operatorname{res})$ is denoted by $\mathbb{F}_\rho(\Sigma, \mathcal{R}, \operatorname{res})$, or simply $\mathbb{F}_\rho$.*

We should think of a ρ-trace as an approximation of a terminated process. The first component is a real trace describing the part of the process which is certain. The second component is the set of resources claimed by the process for its continuation. The process needs not use all resources claimed but it cannot use a resource which has not been claimed. Thus, ρ-traces can be viewed as specifications over traces. The ρ-trace (r, A) represents in this view the set of all real traces starting with r and continuing with actions using only the resources belonging to R.

This interpretation of ρ-traces leads in a natural way to the following *approximation ordering*, essentially corresponding to specification refinement

$$(r, R) \sqsubseteq (s, S) \qquad \text{iff} \qquad r \leq s \quad \text{and} \quad \operatorname{res}(r^{-1}s) \cup S \subseteq R$$

The underlying idea is that we increase the information about a process by letting grow its real part with actions using only the specified resources ($r \leq s$, $\operatorname{res}(r^{-1}s) \subseteq R$) and by possibly reducing the resources claimed for the continuation of the process ($S \subseteq R$). The bottom element with respect to $\sqsubseteq$ is $\perp = (1, \mathcal{R})$. It signifies that nothing has been observed and everything is still possible. If the imaginary part is the empty set, the process is called *terminated*, otherwise *non-terminated.*

Let us see how a meaningful concatenation of two ρ-traces (r, R) and (s, S) should be defined. Since the information we have about the continuation of the first process restricts to the set of resources which can still be used, we can only let execute actions of the second process which do not use these resources. This motivates us to define for all $R \subseteq \mathcal{R}$ and $s \in \mathbb{G}$, the *maximal real prefix of s independent of R* and the corresponding *resource suffix* by

$$\begin{aligned} \mu_R(s) &= \bigvee \{\, k \mid k \in \mathbb{M}, k \leq s, \operatorname{res}(k) \cap R = \emptyset \,\}, \\ \sigma_R(s) &= \operatorname{res}(\mu_R(s)^{-1}s) \end{aligned}$$

They define two mappings $\mu_R : \mathbb{G} \to \mathbb{R}$ and $\sigma_R : \mathbb{G} \to \mathcal{P}(\mathcal{R})$. One can easily see that $\mu_R(s)$ is in fact the restriction of s to the set of events having a past which is finite and labelled with actions not using resources from R. Note that for $R' \subseteq R$ and $s \leq s'$ we have $\mu_R(s) \leq \mu_{R'}(s')$ and

$$\begin{array}{rclclclcl} 1 & = & \mu_{\mathcal{R}}(s) & \leq & \mu_R(s) & \leq & \mu_\emptyset(s) & = & \operatorname{Re}(s) \\ \operatorname{res}(s) & = & \sigma_{\mathcal{R}}(s) & \supseteq & \sigma_R(s) & \supseteq & \sigma_\emptyset(s) & = & \emptyset \end{array}$$

Furthermore, we have

$$\sigma_R(s) \cup \mathrm{resinf}(s) = \bigcap \{ \, \mathrm{res}(k^{-1}s) \mid k \in \mathbb{M}, k \leq s, \mathrm{res}(k) \cap R = \emptyset \, \}.$$

Using the notations above, the following *concatenation* can be defined

$$(r, R) \cdot (s, S) \quad = \quad (r\,\mu_R(s), R \cup \sigma_R(s) \cup S).$$

When concatenating two processes, we let pass into the real component the maximal real prefix of the second process not using resources reserved by the first process for its future and adjoin the resources used by the remaining suffix to the cumulated imaginary parts. As can be easily checked, $\mathbb{F}_\rho$ forms a monoid together with the concatenation defined above and $\top = (1, \emptyset)$ as the neutral element. Note furthermore that the bottom element $\bot = (1, \mathcal{R})$ is right-absorbent with respect to concatenation, i.e. $\bot \cdot x = \bot$ for all $x \in \mathbb{F}_\rho$.

The concatenation in $\mathbb{F}_\rho$ determines as usual a reflexive and transitive prefix relation $\leq$ on $\mathbb{F}_\rho$. Actually, one can easily check that

$$(r, R) \leq (s, S) \quad \text{iff} \quad r \leq s,\ R \subseteq S, \mathrm{res}(r^{-1}s) \cap R = \emptyset.$$

This shows in particular that the prefix relation is antisymmetric, hence it defines a partial order on $\mathbb{F}_\rho$. Comparing the prefix ordering $\leq$ and the approximation ordering $\sqsubseteq$ on $\mathbb{F}_\rho$, we see that they are disjoint in the following sense: If $x \leq y$ and $x \sqsubseteq y$, then $x = y$. This means that a partial computation cannot be at the same time the extension and the refinement of another partial computation.

We have a chain of canonical (partial) monoid homomorphisms (note that the concatenation of two real traces is not necessarily a real trace)

$$\mathbb{M} \hookrightarrow \mathbb{R} \hookrightarrow \mathbb{G} \overset{\varphi_\rho}{\to} \mathbb{F}_\rho$$

where the first two mappings are inclusions and $\varphi_\rho(g) = (\mathrm{Re}(g), \mathrm{resinf}(g))$. Since the mappings are homomorphisms, they are also monotone with respect to the associated prefix orderings. The injection $\mathbb{R} \hookrightarrow \mathbb{F}_\rho$, $r \mapsto (r, \mathrm{resinf}(r))$ is therefore monotone with respect to the prefix ordering $\leq$. It maps every finite trace to a terminated ρ-trace. This canonical embedding allows us to identify $\mathbb{R}$ with its image in $\mathbb{F}_\rho$ and set $\mathbb{R} \subseteq \mathbb{F}_\rho$. Likewise, we have a canonical embedding $\mathcal{P}(\mathcal{R}) \hookrightarrow \mathbb{F}_\rho$, $R \mapsto (1, R)$, which allows us to set $\mathcal{P}(\mathcal{R}) \subseteq \mathbb{F}_\rho$. Using these embeddings and the definition of the concatenation, we see that $x = \mathrm{Re}(x) \cdot \mathrm{Im}(x)$ for all $x \in \mathbb{F}_\rho$. Furthermore, we also have an injection $\mathbb{R} \hookrightarrow \mathbb{F}_\rho$, $r \mapsto (r, \mathcal{R})$ which is monotone with respect to the approximation ordering $\sqsubseteq$. It maps every finite trace to a ρ-trace whose continuation is bound to no restriction.

Given an arbitrary dependence alphabet we can choose conveniently a set of resources and a representing resource mapping depending on the structure of the dependence alphabet. In the word case for instance, where we have full dependency between all actions, we recover the classical completion by restricting to a single resource on which all actions synchronize. We then have on the second component of any ρ-trace either the empty set, corresponding thus to termination, or otherwise the chosen resource as singleton, acting as non-termination symbol. A similar case can be made for a disjoint union of full dependencies by taking a separate resource for each full dependency. In general, we could use as resources the cliques of any covering of the dependence alphabet and assign to every action the cliques of the covering to which it belongs. This representation of dependence alphabets by resources (registers) is similar to the one used in constructing asynchronous automata recognizing regular trace languages [DR95].

5 Properties of the Approximation Ordering

The approximation order on $\mathbb{F}_\rho$ is evidently a refinement of the product order of $(\mathbb{R}, \leq)$ and $(\mathcal{P}(\mathcal{R}), \supseteq)$. Since both orders are well-founded, this implies in turn the well-foundedness of the approximation order. The next proposition shows that for ρ-traces having a common upper bound the approximation order reduces to the product order. Thus, on coherent sets the approximation order is equivalent to the product order.

Proposition 5.1
(1) (r, R) *and* (s, S) *have a common upper bound in* $\mathbb{F}_\rho$ *iff* r *and* s *have a common upper bound in* $\mathbb{R}$ *and moreover* $\mathrm{res}((r \wedge s)^{-1}s) = \mathrm{res}(r^{-1}(r \vee s)) \subseteq R$ *and* $\mathrm{res}((r \wedge s)^{-1}r) = \mathrm{res}(s^{-1}(r \vee s)) \subseteq S$.
(2) If (r, R) *and* (s, S) *have a common upper bound in* $\mathbb{F}_\rho$*, then* $(r, R) \sqsubseteq (s, S)$ *iff* $r \leq s$ *and* $R \supseteq S$.

Proof:
(1) Let (t, T) be a common upper bound of (r, R) and (s, S). It follows from the definition that $r \leq t$ and $s \leq t$, hence $r \vee s \leq t$. Moreover, $\mathrm{res}(r^{-1}(r \vee s)) \subseteq \mathrm{res}(r^{-1}t) \subseteq R$ and similarly, $\mathrm{res}(s^{-1}(r \vee s)) \subseteq S$. Conversely, $\mathrm{resinf}(r \vee s) \subseteq \mathrm{resinf}(r) \cup \mathrm{res}(r^{-1}(r \vee s)) \subseteq R$ and similarly, $\mathrm{resinf}(r \vee s) \subseteq S$. Hence $(r \vee s, R \cap S)$ is a ρ-trace which is clearly a common upper bound of (r, R) and (s, S).
(2) Clear from (1). □

The following proposition shows that $\mathbb{F}_\rho$ is a coherently complete partial order (CCPO) and gives the computation formula for the supremum operator. It also implies the continuity of the mappings $\mathrm{Re} : (\mathbb{F}_\rho, \sqsubseteq) \to (\mathbb{R}, \leq)$ and $\mathrm{Im} : (\mathbb{F}_\rho, \sqsubseteq) \to (\mathcal{P}(\mathcal{R}), \supseteq)$. The proof uses

Lemma 5.2 *Let* $x \in \mathbb{R}$ *and* Re_x *be the set of real traces coherent with* x*. The mapping* $f_x : (\mathrm{Re}_x, \leq) \longrightarrow (\mathbb{R}, \leq)$ *defined by* $y \mapsto x^{-1}(x \vee y)$ *is monotone. Moreover, for all coherent sets* $Y \subseteq \mathrm{Re}_x$*,* $f_x(Y)$ *is coherent and* $f_x(\bigvee Y) = \bigvee f_x(Y)$.

Proof: From the computation of the least upper bound for real traces, it is easy to see that f_x is monotone. Hence, if $Y \subseteq \mathbb{R}$ is coherent, so is $f_x(Y)$. Now remark that if $Z \subseteq \mathbb{R}$ is coherent then so is $x \cdot Z = \{x \cdot z \mid z \in Z\}$ and $x \cdot \bigvee Z = \bigvee x \cdot Z$. Therefore, $x \cdot \bigvee f_x(Y) = \bigvee x \cdot f_x(Y) = \bigvee_{y \in Y} x \cdot x^{-1}(x \vee y) = \bigvee_{y \in Y}(x \vee y) = x \vee \bigvee Y$ and since the concatenation in $\mathbb{R}$ is left-cancellative, we deduce that $\bigvee f_x(Y) = x^{-1}(x \vee \bigvee Y) = f_x(\bigvee Y)$. □

Proposition 5.3 *$\mathbb{F}_\rho$ is a CCPO. In particular it is also a DCPO. The supremum of a coherent set $X \subseteq \mathbb{F}_\rho$ is*

$$\bigsqcup X = (\bigvee_{x\in X} \mathrm{Re}(x), \bigcap_{x\in X} \mathrm{Im}(x))$$

Proof: Let X be coherent in $\mathbb{F}_\rho$. Then $\mathrm{Re}(X)$ must be coherent in $\mathbb{R}$, which is a CCPO, so there exists $r = \bigvee_{x\in X} \mathrm{Re}(x)$. Let $R = \bigcap_{x\in X} \mathrm{Im}(x)$.

We first check that $(r, R) \in \mathbb{F}_\rho$ and that (r, R) is an upper bound of X. For all $x \in X$ we obtain using Lemma 5.2

$$\mathrm{Re}(x)^{-1} r = \mathrm{Re}(x)^{-1}(\mathrm{Re}(x) \vee \bigvee_{y\in X} \mathrm{Re}(y)) = \bigvee_{y\in X} \mathrm{Re}(x)^{-1}(\mathrm{Re}(x) \vee \mathrm{Re}(y)).$$

Using Proposition 5.1 (1), we deduce

$$\mathrm{res}(\mathrm{Re}(x)^{-1} r) = \bigcup_{y\in X} \mathrm{res}(\mathrm{Re}(x)^{-1}(\mathrm{Re}(x) \vee \mathrm{Re}(y))) \subseteq \mathrm{Im}(x).$$

Therefore, $\mathrm{resinf}(r) \subseteq \mathrm{resinf}(\mathrm{Re}(x)) \cup \mathrm{res}(\mathrm{Re}(x)^{-1} r) \subseteq \mathrm{Im}(x)$. Hence $\mathrm{resinf}(r) \subseteq \bigcap_{x\in X} \mathrm{Im}(x) = R$ which shows that $(r, R) \in \mathbb{F}_\rho$. Moreover (r, R) is an upper bound to all $x \in X$ since $\mathrm{res}(\mathrm{Re}(x)^{-1} r) \subseteq \mathrm{Im}(x)$ and evidently $\mathrm{Re}(x) \le r$ and $\mathrm{Im}(x) \supseteq R$.

We now show that (r, R) is the minimal upper bound of X. Let z be any upper bound of X. For all $x \in X$ we then have $\mathrm{Re}(x) \le \mathrm{Re}(z)$ and, therefore, $r = \bigvee \mathrm{Re}(X) \le \mathrm{Re}(z)$. Furthermore, $\mathrm{Im}(x) \supseteq \mathrm{Im}(z) \cup \mathrm{res}(\mathrm{Re}(x)^{-1}\mathrm{Re}(z)) \supseteq \mathrm{Im}(z) \cup \mathrm{res}(r^{-1}\mathrm{Re}(z))$, hence $R = \bigcap_{x\in X} \mathrm{Im}(x) \supseteq \mathrm{Im}(z) \cup \mathrm{res}(r^{-1}\mathrm{Re}(z))$. This proves $(r, R) \sqsubseteq z$. □

Analogously, we can prove the following proposition concerning the infimum operator, which is defined on all non-empty subsets. Note the substantial simplification of the formula for the greatest lower bound in the presence of coherence due to

Lemma 5.4 *Let $x \in \mathbb{R}$. The mapping $g_x : (\mathbb{R}, \le) \longrightarrow (\mathcal{P}(\Sigma), \supseteq)$ defined by $y \mapsto \mathrm{alph}((y \wedge x)^{-1} x)$ is monotone and continuous. Moreover, for all non-empty sets $Y \subseteq \mathbb{R}$ we have $g_x(\bigwedge Y) = \bigcup_{y\in Y} g_x(y)$.*

Proof: If $y \le z$ then $y \wedge x \le z \wedge x \le x$. Hence $(z \wedge x)^{-1} x$ is a suffix of $(y \wedge x)^{-1} x$ and $g_x(y) \supseteq g_x(z)$ which shows that g_x is monotone.

Let $Y \subseteq \mathbb{R}$ be non-empty. For all $y \in Y$, we have $\bigwedge Y \le y$. Since g_x is monotone it follows that $g_x(\bigwedge Y) \supseteq \bigcup_{y\in Y} g_x(y)$. Conversely, note that $a \in g_x(y)$ if and only if there exists a prime real trace z such that $z \le x$, $z \not\le y$ and $\max(z) = \{a\}$. Hence, if $a \in g_x(\bigwedge Y)$, let z be as above. Since $z \not\le \bigwedge Y$, there exists $y \in Y$ such that $z \not\le y$ and it follows that $a \in g_x(y) \subseteq \bigcup_{y\in Y} g_x(y)$. □

Proposition 5.5 *The infimum of a non-empty set $X \subseteq \mathbb{F}_\rho$ is*

$$\sqcap X = (r, R) \text{ where } r = \bigwedge_{x\in X} \mathrm{Re}(x) \text{ and } R = \bigcup_{x\in X} (\mathrm{Im}(x) \cup \mathrm{res}(r^{-1}\mathrm{Re}(x)))$$

In particular, if X is coherent, then

$$\sqcap X = (\bigwedge \mathrm{Re}(X), \bigcup \mathrm{Im}(X))$$

Proof: Since $X \neq \emptyset$, there exists $r = \bigwedge \mathrm{Re}(X)$. Let $R = \bigcup_{x \in X}(\mathrm{Im}(x) \cup \mathrm{res}(r^{-1}\,\mathrm{Re}(x)))$. Choosing some $x \in X$, we have $r \leq \mathrm{Re}(x)$, hence $\mathrm{resinf}(r) \subseteq \mathrm{resinf}(\mathrm{Re}(x)) \subseteq \mathrm{Im}(x) \subseteq R$ and, therefore, $(r, R) \in \mathbb{F}_\rho$. Moreover, (r, R) is clearly a lower bound of X.

Let now (s, S) be any lower bound of X. Then $s \leq \mathrm{Re}(x)$ for all $x \in X$, hence $s \leq r$. Furthermore, for all $x \in X$, $S \supseteq \mathrm{Im}(x) \cup \mathrm{res}(s^{-1}\,\mathrm{Re}(x)) = \mathrm{Im}(x) \cup \mathrm{res}(s^{-1}r) \cup \mathrm{res}(r^{-1}\,\mathrm{Re}(x))$. Therefore $S \supseteq R \cup \mathrm{res}(s^{-1}r)$. We conclude that $(s, S) \sqsubseteq (r, R)$. This shows that (r, R) is the greatest lower bound of X.

If X is coherent, then we see that for all $x \in X$ we have

$$\begin{array}{rcl|l}
\mathrm{res}(r^{-1}\,\mathrm{Re}(x)) & = & \mathrm{res}((r \wedge \mathrm{Re}(x))^{-1}\,\mathrm{Re}(x)) & \\
& = & \bigcup\limits_{y \in X} \mathrm{res}((\mathrm{Re}(y) \wedge \mathrm{Re}(x))^{-1}\,\mathrm{Re}(x)) & \text{by Lemma 5.4} \\
& \subseteq & \bigcup\limits_{y \in X} \mathrm{Im}(y) & \text{by Prop. 5.1.(1) .}
\end{array}$$

Therefore, $\bigcup_{x \in X} \mathrm{res}(r^{-1}\,\mathrm{Re}(x)) \subseteq \bigcup_{x \in X} \mathrm{Im}(x)$, which finally implies that $R = \bigcup_{x \in X}(\mathrm{Im}(x) \cup \mathrm{res}(r^{-1}\,\mathrm{Re}(x))) = \bigcup_{x \in X} \mathrm{Im}(x)$. $\square$

Proposition 5.6 $(\mathbb{F}_\rho, \sqsubseteq)$ *is bounded distributive.*

Proof: By Propositions 5.3 and 5.5 the supremum and infimum of bounded sets can be computed componentwise. Using the bounded distributivity in $\mathbb{R}$ and $\mathcal{P}(\mathcal{R})$ we can then easily check the bounded distributivity in $\mathbb{F}_\rho$. $\square$

As a consequence of Propositions 5.3, 5.6, and the well-foundedness of $(\mathbb{F}_\rho, \sqsubseteq)$ we are now able to state, due to Proposition 2.2, that $(\mathbb{F}_\rho, \sqsubseteq)$ is a p-algebraic CCPO, thus assuring its adequacy as a semantic domain for modelling distributed processes.

Theorem 5.7 $(\mathbb{F}_\rho, \sqsubseteq)$ *is a p-algebraic CCPO. In particular it is also a Scott domain.*

The next theorem identifies the compact elements in $(\mathbb{F}_\rho, \sqsubseteq)$. They are in fact exactly those having a compact and thus finite real part.

Theorem 5.8 *An element* $x \in \mathbb{F}_\rho$ *is compact iff* $\mathrm{Re}(x)$ *is compact (finite).*

Proof: Let x be a compact element of $\mathbb{F}_\rho$ and consider the set $Y = \{\, (k, \mathrm{res}(k^{-1}\,\mathrm{Re}(x)) \cup \mathrm{Im}(x)) \mid k \in \mathbb{M}, k \leq \mathrm{Re}(x) \,\}$. We will first prove that Y is directed and $\bigsqcup Y = x$. Recall that $\cap\{\mathrm{res}(k^{-1}\,\mathrm{Re}(x)) \mid k \in \mathbb{M}, k \leq \mathrm{Re}(x)\} = \mathrm{resinf}(\mathrm{Re}(x)) \subseteq \mathrm{Im}(x)$. Since $\mathbb{R}$ is k-algebraic and the compacts are precisely the finite traces, we have $\mathrm{Re}(x) = \bigvee\{k \in \mathbb{M} \mid k \leq \mathrm{Re}(x)\}$. It follows from Proposition 5.3 that $\bigsqcup Y = x$. Furthermore, since for any $k \leq k' \leq \mathrm{Re}(x)$ we have $(k, \mathrm{res}(k^{-1}\,\mathrm{Re}(x)) \cup \mathrm{Im}(x)) \sqsubseteq (k', \mathrm{res}(k'^{-1}\,\mathrm{Re}(x)) \cup \mathrm{Im}(x))$ we deduce that Y is directed. Now, since x is compact and $x = \bigsqcup Y$, we have $x \in Y$ which implies that $\mathrm{Re}(x) \in \mathbb{M}$ is compact.

Suppose now, conversely, that $\mathrm{Re}(x) \in \mathbb{R}$ is compact and $x \sqsubseteq \bigsqcup Y$ for some directed set $Y \subseteq \mathbb{F}_\rho$. Then $\mathrm{Re}(x) \leq \bigvee \mathrm{Re}(Y)$ and $\mathrm{Im}(x) \supseteq \bigcap \mathrm{Im}(Y)$, whereby both suprema are directed. We can therefore extract from Y two elements y_1 and y_2, such that $\mathrm{Re}(x) \leq \mathrm{Re}(y_1)$ and $\mathrm{Im}(x) \supseteq \mathrm{Im}(y_2)$. Since Y is directed, there exists moreover an upper bound $y \in Y$ to y_1 and y_2. Consequently, we have $\mathrm{Re}(x) \leq \mathrm{Re}(y_1) \leq \mathrm{Re}(y)$ and $\mathrm{Im}(x) \supseteq \mathrm{Im}(y_2) \supseteq \mathrm{Im}(y)$. As $\bigsqcup Y$ is a common upper bound of x and y, this implies by Proposition 5.1.(2) that $x \sqsubseteq y \in Y$ is compact. □

Let us now identify the form of the primes in $(\mathbb{F}_\rho, \sqsubseteq)$. First, one can easily see that the primes in $(\mathcal{P}(\mathcal{R}), \supseteq)$ are the sets missing a single element. Secondly, the primes in $(\mathbb{R}, \leq)$ are so-called *pyramids* (see [GR93]), that is traces having a single maximal event. Now, the primes of the direct product of any two domains are exactly those pairs having a prime on the first component and the bottom element on the second component or vice versa. Recall that the bottom element of $(\mathbb{R}, \leq)$ is 1 and the bottom element of $(\mathcal{P}(\mathcal{R}), \supseteq)$ is $\mathcal{R}$. As the order structure of $(\mathbb{F}_\rho, \sqsubseteq)$ is finer than the direct product order of $(\mathbb{R}, \leq)$ and $(\mathcal{P}(\mathcal{R}), \supseteq)$, we expect to have additional primes corresponding to the constraint imposed on actions increasing the real part to use only those resources specified by the imaginary part. We, therefore, obtain three forms of primes, as stated in the following

Theorem 5.9 *An element $(p, R) \in \mathbb{F}_\rho$ is prime in the following three cases*

(1) p is prime (i.e. $|\max(p)| = 1$) and $R = \mathcal{R}$

(2) $p = 1$ and R is prime (i.e. $|\bar{R}| = 1$)

(3) p and R are both prime and $\mathrm{res}(\max(p)) \cup R = \mathcal{R}$.

Proof: We first show that any prime $(p, R) \in \mathbb{F}_\rho$ has one of the three forms. Let (p, R) be prime, hence necessarily compact and irreducible. Since $(1, \mathcal{R}) = \bigsqcup \emptyset$ is not irreducible, we must have $|\max(p)| \geq 1$ or $|\bar{R}| \geq 1$.

Suppose first that $|\bar{R}| \geq 2$. Then, choosing $\alpha, \beta \in \bar{R}, \alpha \neq \beta$ we would have $(p, R) = (p, R \cup \{\alpha\}) \sqcup (p, R \cup \{\beta\})$. On the other hand, neither $(p, R) = (p, R \cup \{\alpha\})$ nor $(p, R) = (p, R \cup \{\beta\})$. This implies $|\bar{R}| \leq 1$.

Suppose now that $|\max(p)| \geq 2$. Then $p = qab = qba$ for some $a, b \in \Sigma$ with $(a, b) \in I$ and consequently $\mathrm{res}(a) \cap \mathrm{res}(b) = \emptyset$ and $(p, R) = (qab, R) = (qb, R \cup \mathrm{res}(a)) \sqcup (qa, R \cup \mathrm{res}(b))$, but again neither $(p, R) = (qb, R \cup \mathrm{res}(a))$ nor $(p, R) = (qa, R \cup \mathrm{res}(b))$. This implies $|\max(p)| \leq 1$.

Putting these observations together we obtain for (p, R) either one of the first two cases stated or otherwise $|\max(p)| = 1$ and $|\bar{R}| = 1$. Let then $\max(p) = \{a\}$ and $\bar{R} = \{\alpha\}$. Assuming $\alpha \notin \mathrm{res}(a)$ we would have $(p, R) = (qa, R) = (q, R) \sqcup (qa, \mathcal{R})$, but once again neither $(p, R) = (q, R)$ nor $(p, R) = (qa, \mathcal{R})$. We conclude that $\bar{R} = \{\alpha\} \subseteq \mathrm{res}(a) = \mathrm{res}(\max(p))$, which shows that (p, A) satisfies the third case stated.

Let us now show that, conversely, any $(p, R) \in \mathbb{F}_\rho$ of one of the three forms is prime. Let $X \subseteq \mathbb{F}_\rho$ be such that $(p, R) \sqsubseteq \bigsqcup_{x \in X} (p_x, R_x) = (\bigvee_{x \in X} p_x, \bigcap_{x \in X} R_x)$. Then $p \leq \bigvee_{x \in X} p_x$ and $R \supseteq \bigcap_{x \in X} R_x$.

In the first case $R = \mathcal{R}$ and $p \in \mathbb{R}$ is prime, hence there exists $x \in X$, such that $p \leq p_x$ and therefore $(p, R) = (p, \mathcal{R}) \sqsubseteq (p_x, R_x)$.

Similarly, in the second case $p = 1$ and $R \subseteq \mathcal{R}$ is prime, hence there exists $x \in X$ such that $R \supseteq R_x$. Since $p = 1 \leq p_x$, we have by Proposition 5.1 (2) that $(p, R) \sqsubseteq (p_x, R_x)$.

Finally, in the third case $p \in \mathbb{R}$ and $R \subseteq \mathcal{R}$ are both prime so $\max(p) = \{a\}$, $\bar{R} = \{\alpha\}$ and, furthermore, $\bar{R} \subseteq \mathrm{res}(\max(p))$ so $\alpha \in \mathrm{res}(a)$. In particular, since R is prime, there exists $x \in X$ such that $R \supseteq R_x$. As (p, R) and (p_x, R_x) are coherent, we have by Proposition 5.1.(2) $R \supseteq R_x \supseteq \mathrm{res}(p_x^{-1}(p_x \vee p))$. Hence, $\alpha \notin \mathrm{res}(p_x^{-1}(p_x \vee p))$ and, therefore, $a \notin \mathrm{alph}(p_x^{-1}(p_x \vee p)) \supseteq \max(p_x^{-1}(p_x \vee p))$. Since $\max(p_x^{-1}(p_x \vee p)) \subseteq \max(p) = \{a\}$ we must have $\max(p_x^{-1}(p_x \vee p)) = \emptyset$. But then $p_x^{-1}(p_x \vee p) = 1$ which implies $p_x = p_x \vee p$ and finally $p \leq p_x$. Together with $R \supseteq R_x$ we again conclude using Proposition 5.1.(2) that $(p, R) \sqsubseteq (p_x, R_x)$.
□

As expected from the word case, the prefix ordering $\leq$ on $\mathbb{F}_\rho$ is not monotone with respect to concatenation. On the other hand, as stated in the next proposition, the approximation ordering $\sqsubseteq$ does have this property. This shows that the approximation ordering is well suited for refining modular compositions of processes since increasing the knowledge about the components leads in turn to an increase in the knowledge about their composition.

Proposition 5.10 *The concatenation in $\mathbb{F}_\rho$ is monotone with respect to the approximation order, i.e. $x \sqsubseteq x'$ and $y \sqsubseteq y'$ implies $x \cdot y \sqsubseteq x' \cdot y'$ for all $x, x', y, y' \in \mathbb{F}_\rho$.*

Proof: Let $x = (r, R)$, $y = (s, S)$, $x' = (r', R')$, $y' = (s', S')$ be such that $x \sqsubseteq x'$ and $y \sqsubseteq y'$. Then $r \leq r'$, $R \supseteq \mathrm{res}(r^{-1}r') \cup R'$ and $s \leq s'$, $S \supseteq \mathrm{res}(s^{-1}s') \cup S'$. By definition, we have $xy = (r\,\mu_R(s), R \cup \sigma_R(s) \cup S)$ and $x'y' = (r'\,\mu_{R'}(s'), R' \cup \sigma_{R'}(s') \cup S')$. We have to prove the two inequalities

$$\begin{array}{rcl} r\,\mu_R(s) & \leq & r'\,\mu_{R'}(s') \\ R \cup \sigma_R(s) \cup S & \supseteq & \mathrm{res}((r\,\mu_R(s))^{-1}(r'\,\mu_{R'}(s'))) \cup R' \cup \sigma_{R'}(s') \cup S'. \end{array}$$

Since $\mathrm{res}(r^{-1}r') \subseteq R$ and $\mathrm{res}(\mu_R(s)) \cap R = \emptyset$ we have $(r^{-1}r')\ I\ \mu_R(s)$, hence $(r^{-1}r') \cdot \mu_R(s) = \mu_R(s) \cdot (r^{-1}r')$. Furthermore, since $R \supseteq R'$ and $s \leq s'$ we have $\mu_R(s) \leq \mu_{R'}(s')$.

To check the first inequality, we use

$$\begin{array}{rcl} r' \cdot \mu_{R'}(s') & = & r \cdot (r^{-1}r') \cdot \mu_R(s) \cdot (\mu_R(s)^{-1}\,\mu_{R'}(s')) \\ & = & r \cdot \mu_R(s) \cdot (r^{-1}r') \cdot (\mu_R(s)^{-1}\,\mu_{R'}(s')). \end{array}$$

To prove the second inequality we proceed analogously. We have $\mu_R(s) \leq \mu_{R'}(s') \leq s'$ and $\mu_R(s) \leq s \leq s'$ hence

$$\mu_R(s)^{-1}s' = (\mu_R(s)^{-1}\,\mu_{R'}(s'))(\mu_{R'}(s')^{-1}s') = (\mu_R(s)^{-1}s)(s^{-1}s').$$

Therefore,

$$\begin{array}{cl} & \mathrm{res}((r\,\mu_R(s))^{-1}(r'\,\mu_{R'}(s'))) \cup \sigma_{R'}(s') \\ = & \mathrm{res}(r^{-1}r') \cup \mathrm{res}(\mu_R(s)^{-1}\,\mu_{R'}(s')) \cup \mathrm{res}(\mu_{R'}(s')^{-1}s') \\ = & \mathrm{res}(r^{-1}r') \cup \mathrm{res}(\mu_R(s)^{-1}s) \cup \mathrm{res}(s^{-1}s') \\ \subseteq & R \cup \sigma_R(s) \cup S \end{array}$$

which together with $R' \subseteq R$ and $S' \subseteq S$ shows the validity of the second condition. $\Box$

The next theorem states that the concatenation is continuous with respect to the approximation order. The proof uses the following Proposition.

Proposition 5.11 *Let $x, y \in \mathbb{F}_\rho$ be ρ-traces. Then,*

$$\mathrm{Kmp}(x \cdot y) = \downarrow(\mathrm{Kmp}(x) \cdot \mathrm{Kmp}(y)).$$

Proof: From Proposition 5.10 it is clear that $\downarrow(\mathrm{Kmp}(x) \cdot \mathrm{Kmp}(y)) \subseteq \mathrm{Kmp}(x \cdot y)$. Conversely, let $x = (r, R)$, $y = (s, S)$ be ρ-traces and consider any $z = (t, T) \in \mathrm{Kmp}(x \cdot y)$. Let $u = r \wedge t$. Then $x' = (u, R \cup \mathrm{res}(u^{-1}r)) \in \mathrm{Kmp}(x)$. Let furthermore $v = u^{-1}t$. Then $v \leq \mu_R(s) \leq s$, hence, $y' = (v, S \cup \mathrm{res}(v^{-1}s)) \in \mathrm{Kmp}(y)$. Now, $\mathrm{res}(v) \cap (\mathrm{res}(u^{-1}r) \cup R) = \emptyset$, consequently

$$\begin{aligned} x' \cdot y' &= (uv, R \cup S \cup \mathrm{res}(u^{-1}r) \cup \mathrm{res}(v^{-1}s)) \\ &= (t, R \cup S \cup \mathrm{res}(u^{-1}r) \cup \mathrm{res}(v^{-1}\,\mu_R(s)) \cup \sigma_R(s)) \\ &= (t, R \cup S \cup \sigma_R(s)) \cup \mathrm{res}(t^{-1}r\,\mu_R(s))) \end{aligned}$$

Since $z \sqsubseteq x \cdot y$ we have $T \supseteq R \cup S \cup \sigma_R(s) \cup \mathrm{res}(t^{-1}r\,\mu_R(s))$. Therefore, $z \sqsubseteq x' \cdot y' \in \mathrm{Kmp}(x) \cdot \mathrm{Kmp}(y)$. $\Box$

Theorem 5.12 *The concatenation in $\mathbb{F}_\rho$ is continuous with respect to the approximation order, i.e. if $X, Y \in \mathbb{F}_\rho$ are directed, then $X \cdot Y$ is directed and $\bigsqcup(X \cdot Y) = (\bigsqcup X) \cdot (\bigsqcup Y)$.*

Proof: Let X and Y be directed in $\mathbb{F}_\rho$. Using Proposition 5.10 one can easily check that $X \cdot Y$ is directed. Moreover, since for all $x \in X$ and $y \in Y$ we have $x \sqsubseteq \bigsqcup X$ and $y \sqsubseteq \bigsqcup Y$, again by Proposition 5.10 we see that $x \cdot y \sqsubseteq (\bigsqcup X) \cdot (\bigsqcup Y)$. Therefore, $\bigsqcup(X \cdot Y) \sqsubseteq (\bigsqcup X) \cdot (\bigsqcup Y)$.

Let us now show the opposite direction. Since X and Y are directed and the concatenation is monotone we have $\mathrm{Kmp}(\bigsqcup X) \cdot \mathrm{Kmp}(\bigsqcup Y) \subseteq \downarrow(X \cdot Y)$. Using Proposition 5.11 we deduce $\bigsqcup X \cdot \bigsqcup Y = \bigsqcup \mathrm{Kmp}(\bigsqcup X \cdot \bigsqcup Y) = \bigsqcup(\mathrm{Kmp}(\bigsqcup X) \cdot \mathrm{Kmp}(\bigsqcup Y)) \sqsubseteq \bigsqcup(X \cdot Y)$. $\Box$

6 Topology

In this section, we introduce a very natural ultra-metric d on ρ-traces. We show that the topology induced by this ultra-metric is precisely the Lawson topology [Law88] induced by the approximation order. Using this fact and the properties of the approximation ordering, we deduce that the metric space $(\mathbb{F}_\rho, d)$ is compact, hence complete. Moreover, the set of finite ρ-traces is countable, discrete and dense in $(\mathbb{F}_\rho, d)$. Finally we prove that the concatenation is uniformly continuous with respect to the defined ultra-metric.

We start by recalling some well-known toplogical definitions. Let $(X, \mathcal{O})$ be a topological space. A family $\mathcal{B} \subseteq \mathcal{O}$ is called a *basis* of the topology iff any open set is an arbitrary union of sets in $\mathcal{B}$. A family $\mathcal{S} \subseteq \mathcal{O}$ is called a *subbasis* of the topology iff the family of finite intersections of sets in $\mathcal{S}$ is a basis of the topology.

A function $d : X \times X \to \mathbb{R}$ is called an *ultrametric* on X iff for any $x, y, z \in X$ it is *non-degenerate*, i.e. $d(x, y) = 0 \Leftrightarrow x = y$, *symmetric*, i.e. $d(x, y) = d(y, x)$ and satisfies the *ultrametric inequality*, i.e. $d(x, z) \leq \max(d(x, y), d(y, z))$.

Finally, let $\mathrm{Kmp}_n = \{ k \in \mathrm{Kmp} \mid |k| \leq n \}$ and $\mathrm{Kmp}_n(x) = \mathrm{Kmp}_n \cap \mathrm{Kmp}(x)$ for any $n \in \mathbb{N}$, $x \in \mathbb{F}_\rho$.

Definition 6.1 *Let $x, y \in \mathbb{F}_\rho$. The* distance *between x and y is defined by $d(x, y) = 2^{-l(x,y)}$ where*

$$l(x, y) = \inf\{ n \in \mathbb{N} \mid \mathrm{Kmp}_n(x) \neq \mathrm{Kmp}_n(y) \} \in \mathbb{N} \cup \{\infty\}$$

and by convention, $\inf(\emptyset) = \infty$ *and* $2^{-\infty} = 0$.

Let $x \in \mathbb{F}_\rho$ and $n \in \mathbb{N}$. The n^{th} open ball centred in x is

$$B(x, n) = \{ y \in \mathbb{F}_\rho \mid l(x, y) > n \}.$$

For all $x \in \mathbb{F}_\rho$ we have by convention $d(x, x) = 0$. Conversely, if $x, y \in \mathbb{F}_\rho$ are such that $d(x, y) = 0$ then $\mathrm{Kmp}_n(x) = \mathrm{Kmp}_n(y)$ for all $n \in \mathbb{N}$, hence, $\mathrm{Kmp}(x) = \cup_{n\in\mathbb{N}} \mathrm{Kmp}_n(x) = \cup_{n\in\mathbb{N}} \mathrm{Kmp}_n(y) = \mathrm{Kmp}(y)$, and using the algebraicity of $(\mathbb{F}_\rho, \sqsubseteq)$ it follows that $x = \bigsqcup \mathrm{Kmp}(x) = \bigsqcup \mathrm{Kmp}(y) = y$. The distance function is clearly symmetric by definition and, as can be easily checked, satisfies the ultrametric inequality. Therefore, d is an ultra-metric on $\mathbb{F}_\rho$. Moreover, the family of open balls

$$\mathcal{B} = \{B(x, n) \mid x \in \mathbb{F} \text{ and } n \in \mathbb{N}\}$$

is a basis of the topology induced by the metric d on $\mathbb{F}_\rho$.

The Lawson topology can be defined in a quite general setting, namely for continuous complete semi-lattices (see [GHK+80] for a thorough treatment of the subject). Here we are mainly interested in the properties of the Lawson topology for domains (i.e. k-algebraic DCPOs) $(X, \leq)$. In this case the family of sets

$$\mathcal{S} = \{\uparrow k \mid k \text{ is compact}\} \cup \{\overline{\uparrow k} \mid k \text{ is compact}\}$$

can be shown to build a subbasis of the Lawson topology on $(X, \leq)$. We use this fact in order to prove the equivalence between the d-topology and the $\sqsubseteq$-Lawson topology on $\mathbb{F}_\rho$.

Proposition 6.2 *The topology on $\mathbb{F}_\rho$ induced by the ultra-metric d coincides with the Lawson topology on $\mathbb{F}_\rho$ induced by the approximation order $\sqsubseteq$.*

Proof: Consider any compact element $k \in \mathbb{F}_\rho$. For all $x, y \in \mathbb{F}_\rho$ such that $l(x, y) > |k|$ we have $k \sqsubseteq x \Longleftrightarrow k \sqsubseteq y$ or stated equivalently $x \in \uparrow k \Longleftrightarrow y \in \uparrow k$. Therefore, if $x \in \uparrow k$ then $B(x, |k|) \subseteq \uparrow k$ and if $x \in \overline{\uparrow k}$ then $B(x, |k|) \subseteq \overline{\uparrow k}$. This shows that $\uparrow k$ and $\overline{\uparrow k}$ are both d-open (and d-closed). Thus, the d-topology is finer than the $\sqsubseteq$-Lawson topology.

We now show that, conversely, for any $x \in \mathbb{F}_\rho$, $n \in \mathbb{N}$ the open ball $B(x, n)$ is Lawson-open. For $k \in \mathrm{Kmp}$, let $O_x(k) = \uparrow k$ if $k \sqsubseteq x$ and $O_x(k) = \overline{\uparrow k}$ otherwise. Note that in either case $O_x(k)$ belongs to the subbasis of the Lawson topology. Now, for any $y \in \mathbb{F}_\rho$ we have $y \in B(x, n)$ iff $l(x, y) > n$ iff $\mathrm{Kmp}_n(x) = \mathrm{Kmp}_n(y)$. It is easy to check that the later statement is equivalent to $y \in O_x(k)$ for all $k \in \mathrm{Kmp}_n$. Therefore, $B(x, n) = \bigcap_{k \in \mathrm{Kmp}_n} O_x(k)$. Since Kmp_n is a finite set, $B(x, n)$ is a finite intersection of sets belonging to the subbasis of the Lawson topology, hence Lawson-open itself. Therefore, the $\sqsubseteq$-Lawson topology is finer than the d-topology, thus concluding the equivalence proof between the two topologies. □

From Proposition 6.2, we infer the following property which is a general fact for the Lawson topology (see [GHK$^+$80]). Essentially it states that $\sqsubseteq$-convergence is stronger then d-convergence and that the two limit notions, when defined in both settings, coincide.

Proposition 6.3 *Let $(x_n)_{n \geq 0} \subseteq \mathbb{F}_\rho$ be an $\sqsubseteq$-increasing sequence. Then, the sequence $(x_n)_{n \geq 0}$ is d-convergent and furthermore*

$$\lim_{n \to \infty} x_n = \bigsqcup_{n \geq 0} x_n$$

For any partial order $(X, \leq)$ which is both bounded complete and directed complete the Lawson topology is quasi-compact and T_1. If $(X, \leq)$ is in addition k-algebraic, then the Lawson topology becomes T_2 (Hausdorff), whence compact, and the set of compact elements of $(X, \leq)$ is dense in X. For all these statements, the reader is referred to [GHK$^+$80] where they are proved for the more general case of continuous complete semi-lattices. Using these properties we now give the most important corollary of Proposition 6.2.

Theorem 6.4 *The metric space $(\mathbb{F}_\rho, d)$ is compact (hence complete). Furthermore, the set of finite ρ-traces is countable, dense, discrete and open in $(\mathbb{F}_\rho, d)$.*

Proof: By Theorem 5.7, the partial order $(\mathbb{F}_\rho, \sqsubseteq)$ is a bounded complete domain. Since the d-topology is the $\sqsubseteq$-Lawson topology, it follows that $(\mathbb{F}_\rho, d)$ is a compact (hence complete) metric space. Moreover, since the compact ρ-traces are by Theorem 5.8 precisely the finite ones, it follows that the countable set of finite ρ-traces is dense in $(\mathbb{F}_\rho, d)$. There remains for us to show that the set of finite ρ-traces is discrete and open (which is no general fact for the Lawson topology of bounded complete domains).

To this end we prove that $B(k, |k| + 1) = \{k\}$ for any finite ρ-trace k. Consider, therefore, any $x \in B(k, |k| + 1)$. Then $l(k, x) > |k| + 1$, hence, for any finite ρ-trace h with $|h| \leq |k| + 1$ we have $h \sqsubseteq k \iff h \sqsubseteq x$. In particular, taking $h = k$, it follows that $k \sqsubseteq x$, hence, $\mathrm{Re}(k) \leq \mathrm{Re}(x)$ and $|k| \leq |x|$.

Assume that $\mathrm{Re}(k) < \mathrm{Re}(x)$. Then we can find a finite real trace l such that $\mathrm{Re}(k) < l \leq \mathrm{Re}(x)$ with $|l| = |k| + 1$. Taking now $h = l \sqcap x$ we have by Proposition 5.5 $\mathrm{Re}(h) = l \wedge \mathrm{Re}(x) = l$, whence $|h| = |l| = |k| + 1$. Since furthermore $h \sqsubseteq x$, we infer $h \sqsubseteq k$ which contradicts $\mathrm{Re}(k) < l = \mathrm{Re}(h)$.

Therefore, $\mathrm{Re}(k) = \mathrm{Re}(x)$ and $|x| = |k|$. Taking $h = x$, we then deduce $x \sqsubseteq k$, which together with $k \sqsubseteq x$ allows us to conclude $x = k$. $\Box$

We have seen in the previous section (Proposition 5.12) that the concatenation of ρ-traces is continuous with respect to the approximation order. We prove here the following stronger result.

Proposition 6.5 *The concatenation on ρ-traces is uniformly continuous. Moreover, for any $x, x', y, y' \in \mathbb{F}_\rho$*

$$d(xy, x'y') \leq \max(d(x, x'), d(y, y')).$$

Proof: Let $x, x', y, y' \in \mathbb{F}_\rho$. We have to show that for all $n \in \mathbb{N}$ we have $l(x, x') > n$ and $l(y, y') > n \Rightarrow l(xy, x'y') > n$ or, stated equivalently, $\mathrm{Kmp}_n(x) = \mathrm{Kmp}_n(x')$ and $\mathrm{Kmp}_n(y) = \mathrm{Kmp}_n(y') \Rightarrow \mathrm{Kmp}_n(xy) = \mathrm{Kmp}_n(x'y')$. The later statement is a direct consequence of the following claim: for any $x, y \in \mathbb{F}_\rho$, $n \in \mathbb{N}$

$$\mathrm{Kmp}_n(xy) = \mathrm{Kmp}_n \cap {\downarrow}(\mathrm{Kmp}_n(x) \cdot \mathrm{Kmp}_n(y)).$$

Note that this claim strengthens Proposition 5.11 and has exactly the same proof. For the non trivial inclusion, we only have to note that in the proof of Proposition 5.11 we have $|x'| \leq |z|$ and $|y'| \leq |z|$. $\Box$

Note that Proposition 5.12 is in fact a corollary of Proposition 6.5. Indeed, a Lawson-continuous mapping between two domains is necessarily $\sqsubseteq$-continuous. However, the converse is false in general. Hence, we cannot infer the uniform continuity of the concatenation from its $\sqsubseteq$-continuity and the compactness of $(\mathbb{F}_\rho, d)$.

7 The Calculus of α- and δ-Traces

The calculus of α- and δ-traces has been previously introduced in [Die93a, DG95] as means to cope with the problem of non-monotonicity of trace concatenation. We discuss in short some of their strong and weak points, thus providing an a posteriori motivation for the model we have put forward in this paper.

Let the *alphabet at infinity* of a trace $x \in \mathbb{G}$ be defined by

$$\mathrm{alphinf}(x) = \bigcap \{\, \mathrm{alph}(k^{-1}x) \mid k \in \mathbb{M}, k \leq x \,\}$$

It determines a mapping $\mathbb{G} \to \mathcal{P}(\Sigma)$, $x \mapsto \mathrm{alphinf}(x)$. For real traces, the alphabet at infinity represents the set of actions occurring infinitely many times as event labels in the trace. For non-real traces, it additionally contains the alphabet of the non-real part, hence $\mathrm{alphinf}(x) = \mathrm{alphinf}(\mathrm{Re}(x)) \cup \mathrm{alph}(\mathrm{Re}(x)^{-1}x)$.

Definition 7.1 *An α-trace (δ-trace respectively) over the dependence alphabet (Σ, D) is a pair $x = (r, A)$ ($x = (r, \mathrm{D}(A))$ respectively), where $r \in \mathbb{R}(\Sigma, D)$ is a real trace, and $A \subseteq \Sigma$ is a subset of actions such that* $\mathrm{alphinf}(r) \subseteq A$. *The trace r is denoted by* $\mathrm{Re}(x)$ *and called the* real part *of x. The set A ($\mathrm{D}(A)$ respectively) is denoted by* $\mathrm{Im}(x)$ *and called the* imaginary part *of x. The set of α-traces (δ-traces respectively) over (Σ, D) is denoted by $\mathbb{F}_\alpha$ ($\mathbb{F}_\delta$ respectively).*

The informal semantics of α-traces is much the same as for ρ-traces, except that the restriction imposed by the second component on the continuation of the process applies directly to the actions that can still be performed rather than to the resources that can still be used by the process. The interpretation is similar for δ-traces.

Specification refinement leads as expected to the following natural *approximation orderings* for $\mathbb{F}_\alpha$ and $\mathbb{F}_\delta$

$$\begin{array}{lrcl} \alpha\text{-traces:} & (r, A) \sqsubseteq (s, B) & \text{iff} & r \leq s, \mathrm{alph}(r^{-1}s) \cup B \subseteq A \\ \delta\text{-traces:} & (r, \mathrm{D}(A)) \sqsubseteq (s, \mathrm{D}(B)) & \text{iff} & r \leq s, \mathrm{D}(\mathrm{alph}(r^{-1}s) \cup B) \subseteq \mathrm{D}(A). \end{array}$$

Furthermore, $\mathbb{F}_\alpha$ and $\mathbb{F}_\delta$ are monoids with $(1, \emptyset)$ as neutral element for the following *concatenations*

$$\begin{array}{lrcl} \alpha\text{-traces:} & (r, A) \cdot (s, B) & = & (r\,\mu_A(s), A \cup \sigma_A(s) \cup B) \\ \delta\text{-traces:} & (r, \mathrm{D}(A)) \cdot (s, \mathrm{D}(B)) & = & (r\,\mu_A(s), \mathrm{D}(A \cup \sigma_A(s) \cup B)) \end{array}$$

where

$$\begin{array}{rcl} \mu_A(s) & = & \bigvee\{\, k \mid k \in \mathbb{M}, k \leq s, \mathrm{alph}(k) \cap \mathrm{D}(A) = \emptyset \,\}, \\ \sigma_A(s) & = & \mathrm{alph}(\mu_A(s)^{-1}s) \end{array}$$

are the *maximal real prefix of s independent of A* and the corresponding *alphabetic suffix*. Note that concatenation is well-defined in the δ-case since $\mu_A(s)$ and $\sigma_A(s)$ depend on $\mathrm{D}(A)$ only, which is actually the main reason for introducing the δ-model. The canonical mapping $\varphi_{\alpha\delta} : \mathbb{F}_\alpha \to \mathbb{F}_\delta$, $\varphi_{\alpha\delta}(r, A) = (r, \mathrm{D}(A))$ is at the same time a monoid homomorphism and a projection between $\mathbb{F}_\alpha$ and $\mathbb{F}_\delta$. The associated embedding is given by $\psi_{\delta\alpha} : \mathbb{F}_\delta \to \mathbb{F}_\alpha$, $\psi_{\delta\alpha}(r, \mathrm{D}(A)) = (r, \mathrm{I}(\mathrm{I}(A)))$. Therefore, $\mathbb{F}_\delta$ can be regarded as being a more abstract view of $\mathbb{F}_\alpha$, in both algebraical and order-theoretical senses.

As shown in [DG95], the approximation orders $(\mathbb{F}_\alpha, \sqsubseteq)$ and $(\mathbb{F}_\delta, \sqsubseteq)$ are CCPOs. Moreover, the concatenation operations are monotone and even continuous with respect to the approximation orderings defined in each case. We thus have in both cases compatible algebraic and order structures. This is perhaps the most pleasing aspect of these models and actually the original motivation for their investigation.

Now, $(\mathbb{F}_\alpha, \sqsubseteq)$ is shown in [DG95] to be p-algebraic for any dependence alphabet, hence makes a good candidate as a process algebra. The main drawback of $(\mathbb{F}_\alpha, \sqsubseteq)$ is that we do not recover in the sequential case the classical completion using one termination symbol and one non-termination symbol, since we allow for any subset of actions to be specified on the second component. This information about the future of the process is highly redundant in the sequential case and reveals, as a matter of fact, that the α-model is not abstract enough.

The case of $(\mathbb{F}_\delta, \sqsubseteq)$ is in a certain sense exactly opposite. We do recover the classical completion in the sequential case, since the alphabetic information about the future actions restricts to dependency and, therefore, delivers either $\top = \emptyset$, or $\bot = \Sigma$, in all cases. However, $(\mathbb{F}_\delta, \sqsubseteq)$ is only k-algebraic, but fails to be p-algebraic in general. One can easily check this even on small dependence alphabets such as the square or the line of four actions [DG95]. This indicates that the δ-model is too abstract, in the sense that some of the (atomic) information necessary for p-algebraicity has been abstracted away.

The next propositions show that the cause for this lack of p-algebraicity comes indeed from the fact that the family of sets appearing on the second component, $\mathrm{D}(\mathcal{P}(\Sigma)) = \{\, \mathrm{D}(A) \subseteq \Sigma \mid A \subseteq \Sigma \,\}$, is not p-algebraic with respect to inclusion. The p-algebraicity of $(\mathrm{D}(\mathcal{P}(\Sigma)), \subseteq)$ is furthermore shown to be equivalent to a representation condition for the dependence alphabet (Σ, D). To this purpose we need the following definitions.

Let us say that two actions *share* an action α iff they are both dependent on α. An action $\alpha \in \Sigma$ is called a *resource* (*register*) iff any two actions sharing it are mutually dependent. This is tantamount to saying that $\mathrm{D}(\alpha)$ forms a *clique* of the dependence alphabet (Σ, D).

Proposition 7.2
(1) $\alpha \in \Sigma$ *is a resource iff* $\alpha \in \mathrm{D}(a)$ *implies* $\mathrm{D}(\alpha) \subseteq \mathrm{D}(a)$ *for all* $a \in \Sigma$.
(2) The prime elements of the lattice $(\mathrm{D}(\mathcal{P}(\Sigma)), \subseteq)$ *are the subsets* $\mathrm{D}(\alpha)$*, where* α *is a resource.*

Proof:
(1) Suppose α is a resource and $\alpha \in \mathrm{D}(a)$. Then $\mathrm{D}(\alpha)$ is a clique and $a \in \mathrm{D}(\alpha)$, therefore, $\mathrm{D}(\alpha) \subseteq \mathrm{D}(a)$. Conversely, let $a, b \in \mathrm{D}(\alpha)$. Then $\alpha \in \mathrm{D}(a)$, hence $b \in \mathrm{D}(\alpha) \subseteq \mathrm{D}(a)$, which shows that $\mathrm{D}(\alpha)$ is a clique.
(2) The family of sets $(\mathrm{D}(\mathcal{P}(\Sigma)), \subseteq)$ is closed under unions, hence the supremum of any set of elements in $(\mathrm{D}(\mathcal{P}(\Sigma)), \subseteq)$ is given by their union. Furthermore, since $\mathrm{D}(A) = \bigcup_{\alpha \in A} \mathrm{D}(\alpha)$ for all $A \subseteq \Sigma$, we can easily check that any prime in $(\mathrm{D}(\mathcal{P}(\Sigma)), \subseteq)$ must be of the form $\mathrm{D}(\alpha)$ for some $\alpha \in \Sigma$. We show that $\mathrm{D}(\alpha)$ is prime iff α is a resource.

Let first α be a resource and assume that $\mathrm{D}(\alpha) \subseteq \bigcup_{a \in A} \mathrm{D}(a)$. Since $\alpha \in \mathrm{D}(\alpha)$, there exists some $a \in A$, such that $\alpha \in \mathrm{D}(a)$. By (1) this implies $\mathrm{D}(\alpha) \subseteq \mathrm{D}(a)$, showing thus that $\mathrm{D}(\alpha)$ is prime.

Let now $\mathrm{D}(\alpha)$ be prime. We show that $\alpha \in \mathrm{D}(a)$ implies $\mathrm{D}(\alpha) \subseteq \mathrm{D}(a)$ for all $a \in \Sigma$, which by (1) proves that α is a resource. Assume that $\alpha \in D(a)$. We have $\mathrm{D}(\alpha) \subseteq \mathrm{D}(a) \cup \mathrm{D}(\alpha) = \mathrm{D}(a) \cup (\mathrm{D}(\alpha) \setminus \mathrm{D}(a)) \subseteq \mathrm{D}(a) \cup \mathrm{D}(\mathrm{D}(\alpha) \setminus \mathrm{D}(a))$. As $\mathrm{D}(\alpha)$ is prime, we must have either the case that $\mathrm{D}(\alpha) \subseteq \mathrm{D}(a)$, or else that $\mathrm{D}(\alpha) \subseteq \mathrm{D}(\mathrm{D}(\alpha) \setminus \mathrm{D}(a))$. Now, $\mathrm{D}(a) \cap (\mathrm{D}(\alpha) \setminus \mathrm{D}(a)) = \emptyset$ hence $a \notin \mathrm{D}(\mathrm{D}(\alpha) \setminus \mathrm{D}(a))$. Since $a \in \mathrm{D}(\alpha)$ this implies that $\mathrm{D}(\alpha) \not\subseteq \mathrm{D}(\mathrm{D}(\alpha) \setminus \mathrm{D}(a))$, which excludes the second case. We are therefore left with the first case, namely, $\mathrm{D}(\alpha) \subseteq \mathrm{D}(a)$. □

Let $\mathcal{R} \subseteq \Sigma$ denote the set of resources of (Σ, D). The function $\mathrm{res} : \mathcal{P}(\Sigma) \to \mathcal{P}(\mathcal{R})$, $\mathrm{res}(A) = \mathcal{R} \cap \mathrm{D}(A)$ assigns to every subset $A \subseteq \Sigma$ the set of resources on which it depends. Note that $\mathrm{res}(A \cup B) = \mathrm{res}(A) \cup \mathrm{res}(B)$ for all $A, B \subseteq \Sigma$. In particular, $\mathrm{res}(A) = \bigcup_{a \in A} \mathrm{res}(a)$, and using Proposition 7.2.(1), we see that $\mathrm{res}(A) = \{\, \alpha \in \mathcal{R} \mid \mathrm{D}(\alpha) \subseteq \mathrm{D}(A) \,\}$, hence $\mathrm{D}(\mathrm{res}(A)) \subseteq \mathrm{D}(A)$ for all $A \subseteq \Sigma$. By Proposition 7.2.(2), $\mathrm{res}(A)$ denotes therefore all actions α, such that $\mathrm{D}(\alpha)$ is a prime below $\mathrm{D}(A)$ in $(\mathrm{D}(\mathcal{P}(\Sigma)), \subseteq)$.

The very definition of resources guarantees that two actions are mutually dependent, whenever they share some common resource, i.e. $\mathrm{res}(a) \cap \mathrm{res}(b) \neq \emptyset$ implies $a \mathrel{\mathrm{D}} b$ for all $a, b \in \Sigma$. A dependence alphabet is said to be *resource representable* iff the converse implication is also true, that is, iff two actions share some common resource, whenever they are mutually dependent, i.e. $a \mathrel{\mathrm{D}} b$ implies $\mathrm{res}(a) \cap \mathrm{res}(b) \neq \emptyset$ for all $a, b \in \Sigma$.

Proposition 7.3 *The following statements are equivalent for any dependence alphabet* (Σ, D)*:*

(1) $(\mathbb{F}_\delta, \sqsubseteq)$ *is p-algebraic.*

(2) $(\mathrm{D}(\mathcal{P}(\Sigma)), \supseteq)$ *is a p-algebraic (distributive) lattice.*

(3) $\mathrm{D}(A) = \mathrm{D}(\mathrm{res}(A))$ *for all* $A \subseteq \Sigma$*.*

(4) (Σ, D) *is resource representable.*

Proof:
(1) ⇒ (2): Let $(\mathbb{F}_\delta, \sqsubseteq)$ be p-algebraic. The set of δ-traces $\{\, (1, \mathrm{D}(A) \mid A \in \Sigma \,\}$ is a downwards closed subset of $(\mathbb{F}_\delta, \sqsubseteq)$. Hence, as can be easily checked, the restriction of the approximation ordering to this subset must be p-algebraic. But this suborder is evidently isomorphic to the lattice $(\mathrm{D}(\mathcal{P}(\Sigma)), \supseteq)$, which is therefore p-algebraic (distributive).
(2) ⇒ (3): If $(\mathrm{D}(\mathcal{P}(\Sigma)), \supseteq)$ is a p-algebraic lattice, it follows from general lattice-theoretical arguments that its dual $(\mathrm{D}(\mathcal{P}(\Sigma)), \subseteq)$ is a p-algebraic lattice, too. Therefore, every element of $(\mathrm{D}(\mathcal{P}(\Sigma)), \subseteq)$ is the supremum of the primes below itself. Using Proposition 7.2.(2), we see that $\mathrm{D}(A) = \bigcup_{\alpha \in \mathrm{res}(A)} \mathrm{D}(\alpha) = \mathrm{D}(\mathrm{res}(A))$ for all $A \subseteq \Sigma$.
(3) ⇒ (4): For any $a, b \in \Sigma$, such that $a \mathrel{\mathrm{D}} b$, we have $b \in \mathrm{D}(a) = \mathrm{D}(\mathrm{res}(a))$. Consequently, there exists $\alpha \in \mathrm{res}(a)$ with $b \in \mathrm{D}(\alpha)$. But then $\alpha \in \mathcal{R} \cap \mathrm{D}(b) = \mathrm{res}(b)$, thus showing that $\alpha \in \mathrm{res}(a) \cap \mathrm{res}(b) \neq \emptyset$.
(4) ⇒ (1): We will show that $\mathbb{F}_\delta(\Sigma, D) \simeq \mathbb{F}_\rho(\Sigma, \mathcal{R}', \mathrm{res}')$ for a suitable resource alphabet $(\Sigma, \mathcal{R}', \mathrm{res}')$. Let (Σ, D) be resource representable. This implies directly from the definition that the restriction of D to the set of resources $\mathcal{R}$ is an equivalence relation. Let $\mathcal{R}' \subseteq \mathcal{R}$ denote a set of representants with respect to this equivalence relation and define the resource map $\mathrm{res}' : \Sigma \to \mathcal{P}(\mathcal{R}') \setminus \{\emptyset\}$, $\mathrm{res}'(a) = \mathrm{D}(a) \cap \mathcal{R}'$, which we extend in the canonical way to $\mathcal{P}(\Sigma)$. Obviously, $\mathcal{R}'$ is composed of pairwise independent actions, hence, for any $R \subseteq \mathcal{R}'$

$$\mathrm{res}'(R) = R. \tag{1}$$

For all $a, b \in \mathcal{R}$ we have by definition that $a \mathrel{\mathrm{D}} b$ iff there exists $\alpha \in \mathcal{R}$ such that $a, b \in \mathrm{D}(\alpha)$. Since α is a resource and for the representant $\alpha' \in \mathcal{R}'$ of α we have $\alpha' \in \mathrm{D}(\alpha)$, it follows that $a, b \in D(\alpha')$. Hence, $(a, b) \in D$ if and only if there exists $\alpha' \in \mathcal{R}'$ such that $a, b \in \mathrm{D}(\alpha')$, which can be rewritten as $\mathrm{res}'(a) \cap \mathrm{res}'(b) \neq \emptyset$. This shows that D is the dependence relation induced by the resource mapping res'.

For all $b \in \Sigma$ and $A \subseteq \Sigma$ we then have that $b \in \mathrm{D}(A)$ iff $\mathrm{res}'(b) \cap \mathrm{res}'(A) \neq \emptyset$ iff $\mathrm{D}(b) \cap \mathrm{res}'(A) \neq \emptyset$ iff $b \in \mathrm{D}(\mathrm{res}'(A))$. Therefore, for all $A \subseteq \Sigma$,

$$\mathrm{D}(A) = \mathrm{D}(\mathrm{res}'(A)). \tag{2}$$

Together with the defining equation $\mathrm{res}'(A) = \mathrm{D}(A) \cap \mathcal{R}'$ we obtain for any $A, B \subseteq \Sigma$ the equivalence

$$\mathrm{D}(A) \subseteq \mathrm{D}(B) \quad \text{iff} \quad \mathrm{res}'(A) \subseteq \mathrm{res}'(B). \tag{3}$$

Let now $\mathbb{F}_\rho(\Sigma, \mathcal{R}', \mathrm{res}')$ be the set of ρ-traces associated with the resource alphabet $(\Sigma, \mathcal{R}', \mathrm{res}')$ and consider the mapping $\psi : (\mathbb{F}_\rho(\Sigma, \mathcal{R}', \mathrm{res}'), \sqsubseteq_\rho) \to (\mathbb{F}_\delta(\Sigma, D), \sqsubseteq_\delta)$ defined by $(r, R) \mapsto (r, D(R))$.

First, from equation (2) we see that the mapping ψ is surjective.

Secondly, using equation (1) and equation (3) we obtain the following chain of equivalences for any $x, y \in \mathbb{F}_\rho(\Sigma, \mathcal{R}', \mathrm{res}')$, $x = (r, R)$, $y = (s, S)$

$$\begin{aligned}
x \sqsubseteq_\rho y \quad &\Longleftrightarrow \quad r \leq s \text{ and } \mathrm{res}'(r^{-1}s) \cup S \subseteq R \\
&\Longleftrightarrow \quad r \leq s \text{ and } \mathrm{res}'(r^{-1}s) \cup \mathrm{res}'(S) \subseteq \mathrm{res}'(R) \\
&\Longleftrightarrow \quad r \leq s \text{ and } \mathrm{D}(r^{-1}s) \cup \mathrm{D}(S) \subseteq \mathrm{D}(R) \\
&\Longleftrightarrow \quad (r, \mathrm{D}(R)) \sqsubseteq_\delta (s, D(S)) \\
&\Longleftrightarrow \quad \psi(x) \sqsubseteq_\delta \psi(y)
\end{aligned}$$

This proves that the mapping ψ is injective and, furthermore, that ψ and ψ^{-1} are monotone. Thus ψ is an order isomorphism. Since by Theorem 5.7 $(\mathbb{F}_\rho(\Sigma, \mathcal{R}', \mathrm{res}'), \sqsubseteq_\rho)$ is p-algebraic, we finally conclude that $(\mathbb{F}_\delta(\Sigma, D), \sqsubseteq_\delta)$ is p-algebraic. Let us furthermore note that the mapping ψ is also a monoid morphism between $(\mathbb{F}_\rho(\Sigma, \mathcal{R}', \mathrm{res}'), \cdot_\rho)$ and $(\mathbb{F}_\delta(\Sigma, D), \cdot_\delta)$. ☐

The proposition above provides for an efficient algorithm to test the p-algebraicity of $\mathbb{F}_\delta$ by checking the resource representability of the dependence graph. One can easily check by hand that this condition is not fulfilled by the square or the line of four actions. Furthermore, all dependence alphabets not containing these two configurations as induced subgraphs can be shown to satisfy resource representability. Indeed, it is well-known that this class of graphs coincides with the so-called *transitive forests*, that is graphs obtained by taking the transitive closure of the descendent relation in a disjoint union of rooted trees. One can easily see that in a transitive forest, two vertices are connected by an edge if and only if they lie on a same branch. Hence, all leaves are resources and two actions are dependent if and only if they depend on a same leaf. Therefore, the dependence alphabet is resource representable. Moreover, the set of leaves of the transitive forest is a set of representatives of the equivalence classes of resources defined in the proof of Proposition 7.3.

On the other hand, it is clear that the class of resource representable graphs is not definable by forbidden induced subgraphs, since every graph can be completed to a resource representable graph, by extending the vertex set with the elements of any clique covering of the graph and setting every clique dependent on the vertices it contains.

The proof of the theorem also shows the remarkable fact that in the case of representable dependence alphabets the model of δ-traces is isomorphic to a particular model of resource traces where the set of resources is chosen as a subset of the action set. The idea of dropping this additional restriction meanwhile retaining the equivalence between dependence and resource sharing lead us in fact to the investigation of the model of resource traces which we have presented in this paper.

8 Conclusion

Much of the theory presented above can be substantially simplified by using stable embedding-projection pairs. We intend to do this as the next step, in order to clarify the mechanism by which a *minimal completion*, having compatible algebraic and order structures, can be achieved for trace monoids.

There also remains some non-trivial work to be done, which consists in identifying for all dependence alphabets a *canonical set of resources.* Since the resource (register) representation we use is similar to the one employed in constructing asynchronous automata, we think that Zielonka's time-stamping function could provide for such a standard representation.

Acknowledgement

We would like to thank V. Diekert for fruitful discussions and for pointing out that transitive forests satisfy our resource representability condition.

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 ins fmputer CSpceV.