

Model Based Reactive Planning and Prediction for Autonomic Systems*

Invited Paper

Peter H. Deussen
Fraunhofer Institute for Open Communication Systems
Berlin, Germany
peter.deussen@fokus.fraunhofer.de

ABSTRACT

Dynamic adaptation according to situational variety requires the employment of planning algorithms for reactive and proactive dealing with detected or predicted conditions. The complexity of large scaled, distributed systems however prohibits planning strategies acting on the granularity of single components because (a) a local system view from the perspective of singular components is usually not sufficient, and (b) the synthesis (or a-priori definition) of a global perspective is far beyond feasibility. We therefore propose the use of dynamic abstraction mechanisms to generate planning models of a suitable degree of granularity. The framework of Abstract Interpretation is used to define both a reactive planning algorithm and a pro-active prediction.

Categories and Subject Descriptors

F.1.2 [Computation by Abstract Devices]: Modes of Computation—*Parallelism and concurrency, Relativized computation*

Keywords

Models, Autonomic Systems, Pomsets, Pomtrees, Abstraction, Embedding, Zoom, Abstract Interpretation

1. INTRODUCTION

Autonomic Computing [13] and Communication [20] have been introduced as paradigms to cope with the complexity of highly distributed systems and networks where a centralized management approach is doomed to fail. A number of approaches have been presented to perform system operations on the basis of local information, based on the vision of a set of—more or less—simple rules associated with each system component, with the underlying idea that useful behavioral properties “emerge” on the system level.

*This work has been supported by the project “Component-ware for Autonomic, Situation-aware Communications, And Dynamically Adaptable Services (CASCADAS)” (IST-027807) funded by the FET Program of the European Commission.

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

Autonomics October 28-30, Rome, Italy

In this paper, we use a slightly different complementary vision of management that is to be performed using information that goes beyond the system perception of a single component, but does not necessarily involve global system view, leading to an approach that performs management action on a level that is as global as needed and as local as possible.

A key element of our approach is that we assume the existence of behavioral descriptions of system components, subsystems, and systems, and that these *models* can be accessed and processed at system runtime. This model based view opens a variety of possible research topics on system and interaction definition and management, thus for instance:

1. A model description of possible interaction modes (semantically annotated protocols) describing the set of services provided by a system component can be used for automatic service composition, based e. g. on negotiation. Models become a mean of **communication**.
2. Moreover, interaction modes of communities of components can be expressed as suitable combinations of the models of the components involved. Models become a mean to understand **emergent behavior**.
3. Both local component behavior and emergent behavior of communities can be compared with the model descriptions of this behavior, leading to a autonomous self-management approach. Models become a mean to **supervise** systems and system components.

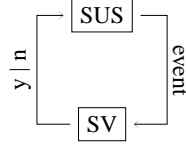
The above list defines a challenging research agenda. In this paper, we are going to provide an approach towards the third topic.

Model-based System Supervision. Supervision may be defined as the ongoing (passive) observation and assessment of the state of a system and the active intervention if an undesired state (such as a failure state) is encountered, and to guide the system back into a more desirable state. Alternatively, intervention can be done proactively to prevent the system from entering a problem state. Both approaches require the availability of planning strategies. A potential supervisor needs to be able to estimate the effect of a supervision action to the overall system state in order to determine whether or not to perform this action. Following the model based approach, we assume that operational system models can be used to estimate effects of system executions.

Clearly, the construction and usage of an overall system model of suitable granularity is—even if this model comprises a finite number of states—far beyond feasibility. Moreover, since autonomic systems are supposed to be open systems able to react on a (probably infinite) set of stimuli from their environment, the finite state

assumption is not longer valid. Thus a model based supervision approach has either to exploit the component structure of the supervised system (compositional approach) or, alternatively, has to work with model abstractions exhibiting a smaller (or even finite) state space, or a combination of both alternatives.

The work presented in this paper is of mathematical nature, thus a number of simplifications is imposed. The first one concerns the actuation model that is used. It bases on the permission of possible actions (events) that the system under supervision (SUS) is able to execute, i. e. if there is a choice between different alternative behaviors, the supervisor (SV) is able to control this choice. With other words, we assume that management capabilities are already represented in the SUS, to be used by the SV. The right hand side figure shows the proposed “architecture” for the most simple case.



Another simplification concerns the notion of a “system”. We use *deterministic labeled transition systems* with probably infinite state space as generic system “syntax”. A partial order semantics based on *partially ordered multisets* [12, 21, 19] is employed to take into account concurrency issues (and to reduce computation complexity that occurs from the explicit representation of interleaving of concurrent actions). *Pomtrees* (also known as *Event Structures* [18]) can be viewed as partial order decision trees. We use pomtrees to represent finite reactive and proactive supervision plans. For the computation of those plans, results from the Petri net theory are exploited, namely the construction of *finite prefixes of branching processes* [16].

Abstract interpretations has been introduced by Cousot & Cousot as a framework for static program analysis [4]. We use a notion of abstraction defined in terms of morphisms between systems and show that the Cousot framework can be applied in this setting: Planning for reactive and pro-active supervision is presented as solutions of fixed point iterations performed in a concrete or abstract domain.

Finally, we explain how to apply the result of this paper in the context of hierarchically organized system models, and discuss how planning for reactive and proactive supervision can be performed on the most concrete (i. e. local) level of a model hierarchy (given a set of abstractions/embeddings).

Methodological Remark. To carry out this work, a combination of a number of approaches from different field has been exploited. We therefore refrain from including the usual section on “Related Work” which could only presented as a seemingly unconnected list of references. We have included a number of remarks that discuss related work in an appropriate context.

We assume that the reader is familiar with the basics of *Lattice theory* (partial orders, (complete) lattices, Tarski’s fixed point theorem); the standard references is [2]. Moreover, the language of *commutating diagrams* used in *Category theory* will prove as useful. Deeper knowledge in this area is not required, the interested reader is referred to [1, 14], as well as to Goldblatts excellent introduction [11, Chapter 3]. Examples are presented using a process algebra with semantics presented in a “structures operational semantics (SOS)” style. Milners book on CCS [17] is probably the best introduction into this field.

Organization of the Paper. Section 2 summarizes the mathematical notions used in this paper, including a brief discussion of Abstract Interpretation. Section 3 introduces labeled transition systems as generic system descriptions, and defines their partial order semantics in terms of pomsets. Moreover, pomtrees are introduced,

and computational issues are discussed. In Section 4, abstractions and embeddings are introduced. Embeddings serve as a tool to describe system composition. Supervision is addressed in Section 5. Both reactive planning and pro-active prediction of problem situations is considered. Section 6 shows how to apply the results presented so far in the setting of hierarchical models. Finally, Section 7 summarizes the paper and gives an outlook on further works.

2. BACKGROUND

Basic Notations. For the sake of brevity we adopt the convention that whenever a structure $A = \langle X, Y, \dots \rangle$ is introduced in a definition, then its components are denoted by X_A, Y_A, \dots . Since we make extensive use of relations, we use some non-standard notations: Let $R \subseteq A \times B$ be a binary relation. Then for some $a \in A$ we let $R(a) =_{\text{df}} \{b \in B \mid a R b\}$. This notion applies also to order relations like \leq . For $C \subseteq A$ we put $R(C) =_{\text{df}} \bigcup_{a \in C} R(a)$. The *inverse* of R is the relation $R^{-1} \subseteq B \times A$ and is defined by: $b R^{-1} a \Leftrightarrow_{\text{df}} a R b$. The identity relation $\text{id}_A \subseteq A \times A$ is defined as $a \text{id}_A a \Leftrightarrow_{\text{df}} a = b$. If $R \subseteq A \times A$ is a relation, we denote by R^+ the smallest transitive relation that contains R , and R^* is $R^+ \cup \text{id}_A$. By $f : A \rightarrow B$ we denote the fact that f is a partial function from A to B (i.e. it might be the case that for some $a \in A$ there is no $b \in B$ such that $\langle a, b \rangle \in f$), while as usual $f : A \rightarrow B$ indicates that f is a total function. Functional restriction is denoted by $f \upharpoonright C$ for $f : A \rightarrow B$ and $C \subseteq A$ and is defined as $f \upharpoonright C =_{\text{df}} f \cap (C \times B)$. If $f : A \rightarrow B$ be a partial function, we use as a further shortcut equations like $f(a) = b$ to indicate both that $f(a)$ is defined *and* that f yields b if applied to a .

\mathbb{N} denotes the set of non-negative integers and \mathbb{O} is the set of finite and transfinite ordinals with first limit ordinal ω .

Let $f : X \rightarrow X$ be a function. A fixed point of f is some $x \in X$ such that $f(x) = x$. If X is partially ordered by \sqsubseteq , then x is a *least fixed point* if $x \sqsubseteq y$ for all other fixed points y of f .

The following is (the first part of) the well-known fixed point theorem of Tarski [22].

THEOREM 1. *Let $\mathcal{X} = \langle X, \leq, \perp, \top, \sqcup, \sqcap \rangle$ be a complete lattice, and $\Phi : X \rightarrow X$ be a monotone function, i. e. $x \sqsubseteq \Phi(x)$ for all $x \in X$. We define:*

$$\begin{aligned} \Phi^0 &=_{\text{df}} \text{id}_X, \\ \Phi^{\delta+1} &=_{\text{df}} \Phi \circ \Phi^\delta \text{ for successor ordinals } \delta + 1 \in \mathbb{O}, \\ \Phi^\lambda &=_{\text{df}} \bigsqcup_{\delta < \lambda} \Phi^\delta \text{ for limit ordinals } \lambda \in \mathbb{O}. \end{aligned}$$

Then $\bigsqcup_{\lambda \in \mathbb{O}} \Phi^\lambda(\perp)$ is the least fixed point of Φ .

Sometimes, sequences of fixed point iterations are considered, each step taking the result of the previous step as starting point. To make the definitions of those chains more readable, we use the notion $\mathbf{lfp}[\Phi, \perp] =_{\text{df}} \bigsqcup_{\lambda \in \mathbb{O}} \Phi^\lambda(\perp)$.

Abstract Interpretation. In this paper, we adopt the framework of *Abstract Interpretation* (AI) introduced by Cousot & Cousot [4]. AI has its roots in static program analysis (strictness analysis for functional programs and ground term analysis for logic programs are classical examples). The basic idea is to “emulate” the execution of a program not in its original domain of variable values but in an abstract, “more simple” version of this domain. A number of variants of the AI framework are possible [5], that mainly differ on the set of operators for abstraction, concretization, and computation that is used, and the consistency criteria between abstract and concrete program executions. In this paper, we use a variant that comprises of abstraction and concretization functions that form *Galois connections*.

Let $\mathcal{C} = \langle X_{\mathcal{C}}, \sqsubseteq_{\mathcal{C}} \rangle$ and $\mathcal{A} = \langle X_{\mathcal{A}}, \sqsubseteq_{\mathcal{A}} \rangle$ be partially ordered sets. A *Galois connection* between \mathcal{C} and \mathcal{A} is a pair of mappings $\alpha : X_{\mathcal{C}} \rightarrow X_{\mathcal{A}}$ and $\gamma : X_{\mathcal{A}} \rightarrow X_{\mathcal{C}}$ such that $\alpha(x) \sqsubseteq_{\mathcal{A}} y \Leftrightarrow x \sqsubseteq_{\mathcal{C}} \gamma(y)$ does hold. We write $\alpha, \gamma : \mathcal{D} \rightleftharpoons \mathcal{A}$ in this case.

Let $\mathcal{C} = \langle X_{\mathcal{C}}, \sqsubseteq_{\mathcal{C}}, \perp_{\mathcal{C}}, \top_{\mathcal{C}}, \sqcup_{\mathcal{C}}, \sqcap_{\mathcal{C}} \rangle$ be a complete lattice, where the domain $X_{\mathcal{C}}$ is interpreted as states of some computation process, and $x \sqsubseteq y$ expresses the fact that the state y is more advanced than the state x , i. e. y occurs in a later stage of the computation process. The computation process itself is expressed by some monotone function $\Phi : X_{\mathcal{C}} \rightarrow X_{\mathcal{C}}$, and the initial state is the bottom element $\perp_{\mathcal{C}}$. Then the complete computation processes is expressed as

$$\perp_{\mathcal{C}}, \Phi(\perp_{\mathcal{C}}), \Phi^2(\perp_{\mathcal{C}}), \dots, \Phi^{\omega}(\perp_{\mathcal{C}}), \Phi^{\omega+1}(\perp_{\mathcal{C}}), \dots$$

which, by Theorem 1, has the limit $\mathbf{lfp}[\Phi, \perp_{\mathcal{C}}]$.¹

Let $\alpha, \gamma : \mathcal{D} \rightleftharpoons \mathcal{A}$ be a Galois connection for complete lattices \mathcal{C} and \mathcal{A} , and let $\Psi : X_{\mathcal{A}} \rightarrow X_{\mathcal{A}}$ be a monotone function. Then we have $x \sqsubseteq_{\mathcal{C}} (\gamma \circ \Psi \circ \alpha)(x)$, i. e. $\Phi(\Phi^{\lambda}(\perp_{\mathcal{C}})) \sqsubseteq_{\mathcal{C}} (\gamma \circ \Psi \circ \alpha)(\Phi^{\lambda}(\perp_{\mathcal{C}}))$, and $\mathbf{lfp}[\Phi, \perp_{\mathcal{C}}] \sqsubseteq_{\mathcal{C}} \gamma(\mathbf{lfp}[\Psi, \perp_{\mathcal{A}}])$. Thus the abstract computation process

$$\perp_{\mathcal{A}}, \Psi(\perp_{\mathcal{A}}), \Psi^2(\perp_{\mathcal{A}}), \dots, \Psi^{\omega}(\perp_{\mathcal{A}}), \Psi^{\omega+1}(\perp_{\mathcal{A}}), \dots$$

has a limit that is an over-approximation of the limit of the concrete computation process, and moreover, each state of the abstract computation process is an over-approximation of its concrete counterpart.

Remark 1. The approach described above can be further improved by means of so-called *widening* and *narrowing* operators. The idea is that widening is used to accelerate the stabilization of the abstract computation process by replacing the various stages of the process by over-approximations so that the resulting process contains only finitely many steps. The price that is paid is that the resulting limit may exceed the limit of the original computation process, i. e. the now finite process reaches a limit $x \sqsupset \mathbf{lfp}[\Psi, \perp]$. A further element of imprecision is added. Narrowing is then used to compute a solution y with $x \sqsupset y \sqsupset \mathbf{lfp}[\Psi, \perp]$ as a corrective measure. Still, the result might be an over-approximation. It is however shown in [6] that Abstract Interpretation based on widening/narrowing is more general than the Galois connection variant.

While computation sequences with widening/narrowing can be defined in a fairly generic way, the concrete definition of these operators seem to require a distinction between program control structure and program data: widening is evidently used to terminate program loops by “jumping” to program data that fulfil the termination condition.

In this paper, we use labeled transition systems to describe computation processes, so no particular syntactically identifiable control structure is available. It is therefore not obvious how to apply the widening/narrowing approach to the notion of systems used in this paper. A possible direction is to exploit the fact that so-called cut-off events (introduced in Section 3) define—in a certain sense—articulation points of loops within sequential processes. This however is subject of further work. \square

¹Usually, countable domains X are used. In this case, the computation process converges at some $\lambda \leq \omega$.

3. SYSTEMS AND THEIR PARTIAL ORDER SEMANTICS

A very generic notion of the term “systems” bases on a notion of the *states* that the system can assume, and a relation on states that models the actions that modify states. We consider furthermore the situation that certain actions can be executed concurrently, i. e. in no definite order, and assume that concurrent executions are reflected in the structure of the system (diamond properties, see below). This section summarizes the necessary definitions and propositions on systems that are needed throughout this paper. Proofs of lemmas and theorems can be found in [7] if not otherwise stated.

Distributed Alphabets and Systems. Let A be a set of symbols, called *actions*, and $I \subseteq A \times A$ be an irreflexive and symmetric *independence relation*. Then $\Sigma = \langle A, I \rangle$ is called a *distributed alphabet*, or, for short, a *d-alphabet*. By $D_{\Sigma} =_{\text{df}} (A \times A) \setminus I$ we denote the *dependence relation* associated with Σ . We moreover stipulate that for each $a \in A$, the set $D(a)$ is finite.

Example 1. Consider a distributed system comprising a number of communicating sequential processes. Each of these processes can be assumed to run at some “location” (e. g. an Internet host, a JAVA thread, etc.). Give a set of *locations* L , and an assignment $l : A \rightarrow \mathcal{P}(L)$ encoding the fact that an action $a \in A$ is executed at the locations $l(a) \subseteq L$, we can define a dependence relation $D \subseteq A \times A$ as $a D b \Leftrightarrow_{\text{df}} l(a) \cap l(b) \neq \emptyset$. Then $\Sigma = \langle A, D \rangle$ forms a d-alphabet. Conversely, if Σ is a d-alphabet, then we may define a D_{Σ} -clique $C \subseteq A_{\Sigma}$ in Σ as a maximal dependent subset, i. e. $C \in \max_{\subseteq} \{C' \subseteq A_{\Sigma} \mid C' \subseteq D_{\Sigma} \times D_{\Sigma}\}$. Let $\mathcal{C}(\Sigma)$ denote the set of all D_{Σ} -cliques in Σ . Then $\mathcal{C}(\Sigma)$ can be used as a set of locations, and the mapping $l : A_{\Sigma} \rightarrow \mathcal{P}(\mathcal{C}(\Sigma))$ is simply given by $l(a) = \{C \in \mathcal{C}(\Sigma) : a \in C\}$. \square

A *labeled transition system* (lts) over a d-alphabet Σ is a pair $\mathcal{S} = \langle S, \delta \rangle$, where S is a (not necessarily finite) set of *states*, and $\delta = \{\delta_a : S \rightarrow S\}_{a \in A}$ is a family of partial transition functions. We usually abbreviate $\delta_a(s)$ by $a_{\mathcal{S}}(s)$. $\mathcal{I} = \langle S, \delta, \varphi \rangle$ is an *interpretation* of Σ if $\langle S, \delta \rangle$ is a lts over Σ and for all $s, s_1, s_2 \in S$ and for all $a, b \in A_{\Sigma}$ with $a I_{\Sigma} b$ the following *diamond properties* are fulfilled whenever the expressions displayed are defined:

$$a_{\mathcal{S}}(s) = s_1 \ \& \ b_{\mathcal{S}}(s) = s_2 \Rightarrow b_{\mathcal{S}}(s_1) = a_{\mathcal{S}}(s_2), \\ a_{\mathcal{S}}(b_{\mathcal{S}}(s)) = b_{\mathcal{S}}(a_{\mathcal{S}}(s)).$$

Further, $\varphi \subseteq S$ is a *state invariant*. By $\text{en}_{\mathcal{S}}(s) =_{\text{df}} \{a \in A_{\Sigma} : a_{\mathcal{S}}(s) \text{ is defined}\}$ we denote the set of actions *enabled* at a state $s \in S$. The class of interpretations of a d-alphabet Σ is denoted by $\mathbf{I}(\Sigma)$.

Example 2. For examples, we need a concrete syntax to express distributed systems. The obvious choice is to use directed graphs, where nodes represent states, and edges are labeled with actions. Another choice would be some state machine syntax, probably equipped with state variables, transition conditions and transitions actions (depending on the values of the state variables), i. e. some type of Extended Finite State Machine. A third option (which would make the definition of d-alphabets quite instructive) would be a Petri-net based formalism.

Graphical formalisms however have the disadvantage that the dynamic creation of system components and their termination requires some extension to the language (e. g. graph grammar rules), since introducing a new component means to modify the graph itself. To avoid concentrating too much on the explanation of the example formalism, which already assumes a significant part of this paper, we therefore use a text based formalism where dynamic component creation can be defined in terms of textual replacement.

$$\begin{aligned}
P &::= \varepsilon \mid a.P \mid x \leftarrow d.P \mid X \mid \mathbf{0} \mid \mathbf{0}_v \mid \sum_{i \in I} P_i \mid P_1 \parallel P_2 \mid X =_{\text{df}} P \mid X^x =_{\text{df}} P \mid X^{x \leftarrow v}; \\
\text{(a) Syntax. } v \in R &\text{ is a process identifier, } d \in D \text{ is a value of domain } D, x \text{ a variable, } X \text{ a process variable, and } a \text{ is an action} \\
&\text{(probably indexed with values } v \text{ or variables } x). \varepsilon \text{ denotes the } \textit{empty process}. \\
&\text{(step)} \frac{}{a.P; R, \sigma \xrightarrow{a} P; R, \sigma}; \quad \text{(ass)} \frac{}{x \leftarrow d.P, R, \sigma \xrightarrow{a} P, R, \sigma[x/d]}; \\
&\text{(rec1)} \frac{X^{x \leftarrow v}; R, \sigma}{P[x \leftarrow v]; R, \sigma} X^x =_{\text{df}} P; \quad \text{(rec2)} \frac{X^x; R, \sigma}{P[x \leftarrow v]; R \setminus \{v\}, \sigma} X^x =_{\text{df}} P \ \& \ v \in R; \quad \text{(rec3)} \frac{X; R, \sigma}{P; R, \sigma} X =_{\text{df}} P \\
&\text{(asyn}_1) \frac{P_1; R, \sigma \xrightarrow{!a} P_2; R, \sigma}{P_1 \parallel P; R, \sigma \xrightarrow{!a} P_2 \parallel P; R, \sigma}; \quad \text{(asyn}_2) \frac{P_1; R, \sigma \xrightarrow{!a} P_2; R, \sigma; P_3; R, \sigma \xrightarrow{?a} P_4; R, \sigma}{P_1 \parallel P_3; R, \sigma \xrightarrow{!a \circ \Sigma ?a} P_2 \parallel P_4; R, \sigma}; \quad \text{(syn)} \frac{P_1; R, \sigma \xrightarrow{a} P_2; R, \sigma; P_3; R, \sigma \xrightarrow{a} P_4; R, \sigma}{P_1 \parallel P_3; R, \sigma \xrightarrow{a} P_2 \parallel P_4; R, \sigma}; \\
&\text{(com)} \frac{P_1 \parallel P_2; R, \sigma}{P_2 \parallel P_1; R, \sigma}; \quad \text{(choice)} \frac{P_i; R, \sigma \xrightarrow{a} P; R, \sigma}{\sum_{i \in I} P_i; R, \sigma \xrightarrow{a} P; R, \sigma} i \in I; \quad \text{(stop}_1) \frac{\mathbf{0}; R, \sigma}{R, \sigma}; \quad \text{(stop}_2) \frac{\mathbf{0}; R, \sigma}{R \cup \{i\}, \sigma}; \quad \text{(cnt)} \frac{C[P; R, \sigma], R}{C[P]; R, \sigma}.
\end{aligned}$$

(b) *Semantics.* $P[x \leftarrow v]$ is obtained by replacing the variable x in the term P by the value v , and $C[P]$ means that the term or state P occurs in the syntactic context C . For $\sigma : V \rightarrow D$, $\sigma[x/d] : V \rightarrow D$ is defined by $\sigma[x/d](y) =_{\text{df}} d$, if $x = y$, and $\sigma[x/d](y) =_{\text{df}} \sigma(y)$ otherwise.

Figure 1: Example process algebra

Figure 1 describes syntax and semantics of a process algebra with variables (taken from a set V) over a data domain D . States of distributed systems described in this language are tuples $\langle P, R, \sigma \rangle$ (written as of $P; R, \sigma$) comprising a *term* P of the algebra, together with a variable state $\sigma : V \rightarrow D$, and a countable infinite set R of unused *process identifiers*. This set is used to create new components, and to keep track of the termination of existing ones. Processes are defined by recursive equations of the form $X^x =_{\text{df}} P$, where X is a process variable and x is a process identifier variable ranging over R . We allow x to be omitted, and to be target of an assignment (in this case we write $X^{x \leftarrow v}$). The transition function δ_a is written in a relational style using arrows, i.e. $P_1; R_1, \sigma_1 \xrightarrow{a} P_2; R_2, \sigma_2 \Leftrightarrow_{\text{df}} \delta_a(\langle P_1, R_1, \sigma_1 \rangle) = \langle P_2, R_2, \sigma_2 \rangle$. An example system is defined in the following paragraphs.

System Actions and Dependencies. Let $I \subseteq R$ be a set of *user identifiers*, J be a set of *scheduler identifiers*, and Q be a set of *service identifiers*. The action set A_Σ of a distributed alphabet Σ is defined by the regular expression $(?+!)(r_{q,i,j} + (g+d+p+e)_{q,i}) + u_{q,i} + x_q \leftarrow d$ for $q \in Q$, $i \in I$, and $j \in J$, and variable symbols x_q for each service $q \in Q$. Moreover, to define the dependence relation D_Σ , we impose the rule that actions a and b (described by the above syntax) are dependent if they share the same user identifier, the same service identifier, or the same scheduler identifier.

Process Equations. Our example system consists of the following processes, defined by a set of recursive equations:

$$\begin{aligned}
G &=_{\text{df}} \sum_{q \in Q} U_q^x \parallel G \\
U_q^x &=_{\text{df}} \sum_{j \in J} !r_{q,x,j} \cdot \{ ?g_{q,x} \cdot ?p_{q,x} \cdot u_{q,x} \cdot !e_{q,x} + ?d_{q,x} \} \cdot \mathbf{0}_x \\
C_j &=_{\text{df}} \sum_{q \in Q} ?r_{q,i,j} \cdot \{ !g_{q,i} \cdot (V_q^{x \leftarrow i} \parallel C_j) + !d_{q,i} \cdot C_j \} \\
V_q^x &=_{\text{df}} x_q \leftarrow x_q + 1 \cdot !p_{q,x} \cdot ?e_{q,x} \cdot x_q \leftarrow x_q - 1 \cdot \mathbf{0},
\end{aligned}$$

where for each service q the variable x_q ranges over the domain $D = \mathbb{N}$. We chose $G \parallel \parallel_{j \in J} C_j; R, \sigma : x_q \mapsto 0$ as an initial state of the system's executions.

The intuition of the example is as follow: A generator process G initiates an arbitrary number of user processes U_q^i which in turn contact one of the scheduler processes C_j by sending a request $!r_{q,i,j}$ for a service q . The scheduler decides whether to grant the request by replying with $!g_{q,i}$, or to deny it by sending $!d_{q,i}$. In the former case, the scheduler initiates a service process V_q^i . This process increases the number of users of the service q stored in the

variable x_q , and sends a $!p_{q,i}$ to the user i to indicate that the service is ready to be used. User i executes the internal action $u_{q,i}$, then it indicates the exit from the service usage by sending $!e_{q,i}$ to the service process which then decreases x_q . Finally, both the user process and the service process terminate.

State Invariant and Error Mode. We now assume that the number of users that are allowed to use a service q is limited by some constant N_q . We establish the state invariant $\varphi =_{\text{df}} \{ \langle P, R, \sigma \rangle : \forall q \in Q. \sigma(x_q) \leq N_q \}$.

The trouble with our example is now that the state invariant is defined in terms that do not have an influence to the execution of the system—a user process will use the service q regardless wether $x_q \leq N_q$ does hold or not. Instead of extending our process language and the example even more, we simply impose new semantic rules:

$$\text{(use)} \frac{u_q.P; R, \sigma}{P; R, \sigma} \sigma(x_q) \leq N_q, \quad \text{(crash)} \frac{u_q.P; R, \sigma}{\varepsilon; R, \sigma[x_q/N_q + 1]} \sigma(x_q) > N_q,$$

i.e. we “crash” a user process that tries to access a service q in case $\sigma(x_q) > N_q$, blocking this service.² \square

Pomsets. A *labeled partial order* (lpo) $u = \langle E, \leq, \lambda \rangle$ over some alphabet A comprises a set of *events* E , a partial order $\leq \subseteq E \times E$ called *causal order* of u , and a *labeling function* $\lambda : E \rightarrow A$. We define $e_1 \leq e_2 =_{\text{df}} \forall e \in e. [e_1 \leq e \leq e_2 \Rightarrow (e = e_1 \vee e = e_2)]$. We usually write \bar{e} instead of $\lambda(e)$, if it is clear from the context which pomset u is meant. For subsets $E \subseteq E_u$ we put $E^\perp =_{\text{df}} \leq_u^{-1}(E)$, and $e^\perp =_{\text{df}} \{e\}^\perp$ for $e \in E_u$, if confusion is not possible. By $\varepsilon =_{\text{df}} \langle \emptyset, \emptyset, \emptyset \rangle$ we denote the *empty lpo*. Moreover, if $a \in A$, we denote the *letter* $\langle \{0\}, \emptyset, \lambda : 0 \mapsto a \rangle$ also by a .

A subset $E \subseteq E_u$ is called *pre-closed* in u if $E^\perp = E$. For pre-closed sets E we define $u[E] =_{\text{df}} \langle E, \leq_u \cap E \times E, \lambda_u \upharpoonright E \rangle$. The *concurrency relation* of u is given by $e_1 \text{ co}_u e_2 =_{\text{df}} \neg(e_1 \leq_u e_2) \ \& \ \neg(e_2 \leq_u e_1)$. A co-set of u is a set $E \subseteq E_u$ such that $E \times E \subseteq \text{co}_u$. A lpo u is called *finitary* if E^\perp is finite for each finite co-set E of u . Lpos u and v are *isomorphic* if there is a bijection $f : E_u \rightarrow E_v$ such that $e_1 \leq_u e_2 \Leftrightarrow f(e_1) \leq_v f(e_2)$ and $\lambda_v \circ f = \lambda_u$. By $[u]$ we denote the equivalence class of u and call $[u]$ a *partially ordered multiset* (pomset).

Remark 2. To get rid of the just introduced notion $[u]$ we do not distinguish between u and its equivalence class $[u]$. From a formal

²The service can actually still be used, as $\sigma(x_q)$ can be decreased by service processes that are active when the “crash” happens.

point of view, this is clearly an oversimplification. But since all concepts that we will discuss throughout this paper are either independent of the concrete choice of a representative $v \in [u]$, or—if not—can be “corrected” by the right choice of v , no harm is done. \square

Let Σ be a d-alphabet and u be a pomset. Assume $\tilde{e}_1 D_\Sigma \tilde{e}_2 \Rightarrow e_1 \leq_u e_2 \vee e_2 \leq_u e_1$ does hold. Then u is called a *pomset* over Σ . By $\mathbf{Ps}(\Sigma)$ we denote the class of pomsets over Σ . Throughout this paper, we assume that all pomsets $u \in \mathbf{Ps}(\Sigma)$ are finitary, and moreover, that E_u is a countable set.

LEMMA 1. *If $u \in \mathbf{Ps}(\Sigma)$ is finitary, then $u = \bigvee \{v \in \mathbf{Ps}(\Sigma) : v \leq u \text{ and } v \text{ is finite}\}$.*

An example where blurring up representatives and equivalence classes requires the right choice of a representative is the following: Define $u \leq v$ if there is some injection $f : E_u \rightarrow E_v$ such that $f(e^\perp) = f(e)^\perp$ and $\lambda_v \circ f = \lambda_u$. If $u \leq v$ holds, then u is called a *prefix* of v . Furthermore, put $u \preceq v$ if there is some bijection $f : E_u \rightarrow E_v$ such that $f(e^\perp) \subseteq f(e)^\perp$ and, as before, $\lambda_v \circ f = \lambda_u$. Then u is called *weaker* than v .

For some $u \in \mathbf{Ps}(\Sigma)$, we define $\langle u \rangle_\Sigma =_{\text{df}} \langle E_u, R^*, \lambda_u \rangle$, with $e_1 R e_2 \Leftrightarrow_{\text{df}} e_1 \leq_u e_2$ & $\tilde{e}_1 D_\Sigma \tilde{e}_2$. $\langle u \rangle_\Sigma$ is called the *weakening* of u w. r. t. Σ . We define $\mathbf{Ps}^w(\Sigma) =_{\text{df}} \{\langle u \rangle_\Sigma : u \in \mathbf{Ps}(\Sigma)\}$

LEMMA 2. *For $u, v \in \mathbf{Ps}(\Sigma)$, (a) $\langle u \rangle_\Sigma \in \mathbf{Ps}(\Sigma)$, (b) $v \preceq \langle u \rangle_\Sigma \Rightarrow v = \langle u \rangle_\Sigma$; (c) $u \preceq v \Rightarrow \langle u \rangle_\Sigma = \langle v \rangle_\Sigma$; (d) $\langle \langle u \rangle_\Sigma \rangle_\Sigma = \langle u \rangle_\Sigma$.*

Furthermore, let us assume $E_u \cap E_v = \emptyset$ for appropriate representatives of $u, v \in \mathbf{Ps}(\Sigma)$. We define $u \circ_\Sigma v =_{\text{df}} \langle E_u \cup E_v, R^*, \lambda_u, \lambda_v \rangle$, with $e_1 R e_2 =_{\text{df}} e_1 \leq_u e_2 \vee e_1 \leq_v e_2 \vee (e_1 \in E_u \text{ \& } e_2 \in E_v \text{ \& } \tilde{e}_1 D_\Sigma \tilde{e}_2)$. $u \circ_\Sigma v$ is called the *weak sequential composition* of u and v w. r. t. Σ .

LEMMA 3. *For all $u, v, w \in \mathbf{Ps}^w(\Sigma)$, (a) $u \circ_\Sigma v \in \mathbf{Ps}^w(\Sigma)$, (b) $\langle u \rangle_\Sigma \circ_\Sigma \langle v \rangle_\Sigma = \langle u \circ_\Sigma v \rangle_\Sigma$; (c) $\varepsilon \circ_\Sigma u = u = u \circ_\Sigma \varepsilon$; (d) $u \circ_\Sigma (v \circ_\Sigma w) = (u \circ_\Sigma v) \circ_\Sigma w$.*

Semantics. By virtue of the operator \circ_Σ , the transition function $\delta_{\mathcal{F}}$ of a system $\mathcal{F} \in \mathbf{I}(\Sigma)$ can be extended to a partial mapping $\delta_u^v : S_{\mathcal{F}} \rightarrow S_{\mathcal{F}}$ for $u \in \mathbf{Ps}^w(\Sigma)$ as follows: $\delta_\varepsilon =_{\text{df}} \text{id}_{S_{\mathcal{F}}}$, $\delta_{u \circ_\Sigma v} =_{\text{df}} \delta_{\mathcal{F}.a} \circ \delta_u$ for finite u . We say that $\delta_u(s)$ is *defined* even for infinite u if $\delta_v(s)$ is defined for all finite v with $v < u$. As before, we abbreviate $\delta_v(s) =_{\text{df}} v_{\mathcal{F}}(s)$. By $\mathbf{L}_{\mathcal{F}}(s) =_{\text{df}} \{v \in \mathbf{Ps}^w(\Sigma) : v_{\mathcal{F}}(s) \text{ is defined}\}$ we denote the *pomset language* of \mathcal{F} at $s \in S_{\mathcal{F}}$.

Example 3. Consider the system from Example 2 and suppose for simplicity $J = \{c\}$, and $Q = \{q\}$. The first steps of an example execution are shown in Fig. 2. Note that the resulting pomset forms a chain as in the example execution no concurrent computations are performed.

Pomsets are more readable if drawn as directed acyclic graphs (without transitive edges). Fig. 3 shows a pomset representing an execution of our example system where two users try to access a service q , the first one succeeds while the second one fails. \square

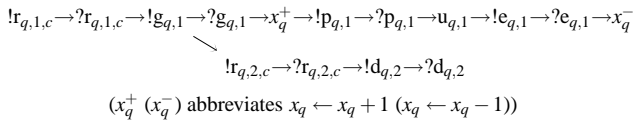


Figure 3: Example execution (pomset).

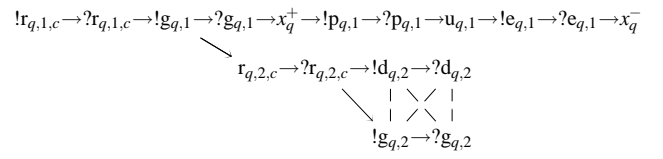


Figure 4: A pomtree.

Pomtrees. The trouble with pomsets is that they do not form a complete lattice which is required to apply the AI framework of Cousot & Cousot, because $u \vee v$ might not be a pomset over Σ . We therefore extend the definition of pomsets to incorporate a notion of “conflict” or “branching behavior”.³

Let Σ be a d-alphabet. A *pomtree* over Σ is a pomset ζ is a pomset over A_Σ together with a *conflict relation* $\# \subseteq E_\zeta \times E_\zeta$ defined to be the smallest irreflexive and symmetric relation satisfying: (a) $\tilde{e}_1 D_\Sigma \tilde{e}_2$ & $e_1 \text{ co}_\zeta e_2 \Rightarrow e_1 \# e_2$, and (b) $e_1 \# e_2$ & $e_1 \leq_\zeta e_3 \Rightarrow e_3 \# e_2$. By $\mathbf{Pt}(\Sigma)$ we denote the class of pomtrees over Σ .

Example 4. Fig. 4 shows an example of a pomtree that expresses two different executions of the system described in Example 2. The conflict relation is represented by a dashed line. \square

A configuration of a pomtree ζ is a prefix $u \leq \zeta$ of ζ such that $u \in \mathbf{Ps}^w(\Sigma)$. By $\mathbf{Cf}(\zeta)$ we denote the class of configurations of ζ . Further, define $\mathbf{Cf}_\vee(\zeta) =_{\text{df}} \max_{\leq} \mathbf{Cf}(\zeta)$

LEMMA 4. *Let $\zeta \in \mathbf{Pt}(\Sigma)$ and let $E \subseteq E_\zeta$ a pre-closed set of event. If $E \times E \cap \#_\zeta = \emptyset$, then $\zeta[E] \in \mathbf{Cf}(\zeta)$.*

LEMMA 5. *For all $\zeta, \xi \in \mathbf{Pt}(\Sigma)$, $\zeta \vee \xi$ is defined and is a least upper bound of ζ and ξ w. r. t. \leq . Moreover, $\zeta = \bigvee \mathbf{Cf}(\zeta)$.*

Obviously, $\mathbf{Cf}(\zeta)$ is prefix closed, in particular, $\varepsilon \in \mathbf{Cf}(\zeta)$. Furthermore, note that $\#_\zeta$ is uniquely defined by ζ and Σ . Thus each pomset $u \in \mathbf{Ps}(A_\Sigma)$ can be viewed as a pomtree over Σ with generated conflict relation $\#_u = \emptyset$. We therefore allow expressions like $u \vee v$ for $u, v \in \mathbf{Ps}^w(\Sigma)$, even if $u \vee v \notin \mathbf{Ps}^w(\Sigma)$, since (using the implicit construction rule for $\#_{u \vee v}$) $u \vee v \in \mathbf{Pt}(\Sigma)$.

THEOREM 2. *Let $\mathcal{F} \in \mathbf{I}(\Sigma)$ and $s \in S_{\mathcal{F}}$. Define $\mathbf{T}_{\mathcal{F}}(s) =_{\text{df}} \{\bigvee \{u : \exists v \in X. u \leq v\} : X \subseteq \mathbf{L}_{\mathcal{F}}(s)\}$ to be the pomtree language of \mathcal{F} at s . Then the structure $(\mathbf{T}_{\mathcal{F}}(s), \leq, \varepsilon, \bigvee \mathbf{T}_{\mathcal{F}}(s), \vee, \wedge)$ forms a complete lattice.*

Remark 3. We occasionally use the notion $\zeta \circ_\Sigma \xi$ for some $\zeta, \xi \in \mathbf{Pt}(\Sigma)$ to denote one of the pomtrees $\bigvee (\{u \circ_\Sigma v : v \in \mathbf{Cf}_\vee(\xi)\} \cup \mathbf{Cf}(\zeta) \setminus \{u\})$ for $u \in \mathbf{Cf}(\zeta)$. Clearly, applied to pomtrees, the operation \circ_Σ is not longer well-defined (it depends on the choice of u). Thus we restrict the usage of expressions of this form to those cases where the concrete result does not matter.

Finite Representations of Pomtrees. Let $\mathcal{F} \in \mathbf{I}(\Sigma)$ be an interpretation of Σ and let $s \in S_{\mathcal{F}}$. In this paragraph, we will show that if the set of states reachable from s if finite that there is a finite pomtree ζ which is *complete* in the sense that each state reachable from s is represented in ζ .

Let $s \in S_{\mathcal{F}}$ and let $\zeta \in \mathbf{T}_{\mathcal{F}}(s)$. We call an event $e \in E_\zeta$ a *cut-off event* if there is another event $e' \in E_\zeta$ with $e' <_\zeta e$ and $\zeta[e^\perp]_{\mathcal{F}}(s) =$

³Alternatively, one can work with sets of pomsets pre-closed under \leq , see Theorem 2.

$$\begin{array}{lcl}
G \parallel C; R, x_q = 0 & \vdash_{\text{rec}_3} & U_q^x \parallel G \parallel C; R, x_q = 0 \\
& \vdash_{\text{rec}_2} & !r_{q,1,c} \cdot \{ ?g_{q,1} \cdot ?p_{q,1} \cdot !e_{q,1} + ?d_{q,1} \} \cdot \mathbf{0}_1 \parallel G \parallel ?r_{q,1,c} \cdot \{ !g_{q,1} \cdot (V_q^{x-1} \parallel C_c) + !d_{q,1} \cdot C_c \}; R \setminus \{1\}, x_q = 0 \\
& \xrightarrow[\text{asyn}_2]{!r_{q,1,c} \circ_{\Sigma} ?r_{q,1,c}} & \{ ?g_{q,1} \cdot ?p_{q,1} \cdot !e_{q,1} + ?d_{q,1} \} \cdot \mathbf{0}_1 \parallel G \parallel \{ !g_{q,1} \cdot (V_q^{x-1} \parallel C_c) + !d_{q,1} \cdot C_c \}; R \setminus \{1\}, x_q = 0 \\
& \xrightarrow[\text{choice, asyn}_2]{\dots \circ_{\Sigma} !g_{q,1} \circ_{\Sigma} ?g_{q,1}} & ?p_{q,1} \cdot !e_{q,1} \cdot \mathbf{0}_1 \parallel G \parallel V_q^{x-1} \parallel C_c; R \setminus \{1\}, x_q = 0 \\
& \vdash_{\text{rec}_1} & ?p_{q,1} \cdot !e_{q,1} \cdot \mathbf{0}_1 \parallel G \parallel x_q \leftarrow x_q + 1 \cdot !p_{q,1} \cdot ?e_{q,1} \cdot x_q \leftarrow x_q - 1 \cdot \mathbf{0} \parallel C_c; R \setminus \{1\}, x_q = 0 \\
& \xrightarrow[\text{ass}]{\dots \circ_{x_q} x_q + 1} & ?p_{q,1} \cdot !e_{q,1} \cdot \mathbf{0}_1 \parallel G \parallel !p_{q,1} \cdot ?e_{q,1} \cdot x_q \leftarrow x_q - 1 \cdot \mathbf{0} \parallel C_c; R \setminus \{1\}, x_q = 1 \quad \dots
\end{array}$$

Figure 2: Example execution.

$\zeta[e^{\downarrow}]_{\mathcal{G}}(s)$. With other words, the state that is reached after executing e has already been reached by the execution of smaller event e' , i. e. is already represented in the history of e . A pomtree $\zeta \in \mathbf{T}_{\mathcal{G}}(s)$ is called *locally complete* at s if

LC₁. each event $e \in \max_{\leq_{\zeta}} E_{\zeta}$ is either a cut-off event or for all $a \in A_{\Sigma}$ it holds that $\zeta[e^{\downarrow}]_{\mathcal{G}} \circ_{\Sigma} a \notin \mathbf{L}_{\mathcal{G}}(s)$, and

LC₂. if $e \in E_{\zeta}$ is not a cut-off event and $\zeta[e^{\downarrow}]_{\mathcal{G}} \circ_{\Sigma} a \in \mathbf{L}_{\mathcal{G}}(s)$ for some $a \in A_{\Sigma}$, then $\zeta[e^{\downarrow}]_{\mathcal{G}} \circ_{\Sigma} a \leq_{\zeta}$.

The pomtree ζ is called *globally complete* at s if for each state $s' \in S_{\mathcal{G}}$ we have: If there is some $w \in \mathbf{L}_{\mathcal{G}}(s)$ such that $w_{\mathcal{G}}(s) = s'$, then there is some configuration $v \in \mathbf{Cf}(\zeta)$ with $v_{\mathcal{G}}(s) = s'$.

We now show that local completeness implies global completeness.

THEOREM 3. *Let $\mathcal{G} \in \mathbf{I}(\Sigma)$, $s \in S_{\mathcal{G}}$ and let $\zeta \in \mathbf{T}_{\mathcal{G}}(s)$. Assume that the set $R_s =_{df} \{s' \in S_{\mathcal{G}} : \exists u \in \mathbf{L}_{\mathcal{G}}(s). s' = u_{\mathcal{G}}(s)\}$ is finite. If ζ is locally complete at s than it is globally complete at s .*

PROOF. (Sketch) Suppose some state $s' \in R_s$ that is not represented in ζ , i. e. $u_{\mathcal{G}}(s) \neq s'$ for all configurations $u \in \mathbf{Cf}(\zeta)$. Since $s' \in R_s$, there is some $w \in \mathbf{L}_{\mathcal{G}}(s)$ such that $w_{\mathcal{G}}(s) = s'$. Let $v \in \mathbf{Cf}(\zeta)$ be a maximum pomset with $v \leq w$. Then there is a cut-off event $e \in E_v$, i. e. there is some $e' <_{\zeta} e$ such that $\zeta[e^{\downarrow}]_{\mathcal{G}}(s) = \zeta[e'^{\downarrow}]_{\mathcal{G}}(s)$. Furthermore, we have $w = \zeta[e^{\downarrow}]_{\mathcal{G}} \circ_{\Sigma} u$ for some pomset u . It follows that $\zeta[e'^{\downarrow}]_{\mathcal{G}} \circ_{\Sigma} u \in \mathbf{L}_{\mathcal{G}}(s)$. But since ζ is local complete, there must be non-empty u_1, u' such $u = u_1 \circ_{\Sigma} u'$ and $\zeta[e'^{\downarrow}]_{\mathcal{G}} \circ_{\Sigma} u_1 \leq_{\zeta}$. Let u_1 be the maximum pomset with this property. Again, u_1 must contain a cut-off event. Thus we can repeat the above construction until the part of u that does not belong to ζ reduces to ε . \square

THEOREM 4. *Let $\mathcal{G} \in \mathbf{I}(\Sigma)$, $s \in S_{\mathcal{G}}$ and let $\zeta \in \mathbf{T}_{\mathcal{G}}(s)$. Assume that the set $R_s =_{df} \{s' \in S_{\mathcal{G}} : \exists u \in \mathbf{L}_{\mathcal{G}}(s). s' = u_{\mathcal{G}}(s)\}$ is finite, moreover, assume that $\text{en}_{\mathcal{G}}(s)$ is finite for each $s \in R_s$. Then there is a finite locally complete pomtree ζ at s .*

PROOF. (Sketch) Since R_s is finite, there is no infinite sequence of pomtrees $\zeta_0 \leq \zeta_1 \leq \dots \leq \zeta_i \leq \dots$ without replicating the states that are represented in ζ_i . The finiteness of $\text{en}_{\mathcal{G}}(s)$ ensures that ζ_i does not contain an infinite co-set. \square

The following theorem explains how to compute locally (and thus globally) complete pomtrees as result of a fixed point iteration. Since the result of each step of this iteration depends on the actual choice of an action $a \in A_{\Sigma}$, the step function Φ_s is defined on sets of pomtrees rather than on pomtrees.

THEOREM 5. *With the assumptions of Theorem 4, a locally complete finite pomtree can be effectively computed as the least*

fixed point of the function $\Phi_s : \mathcal{P}(\mathbf{Pt}(\Sigma)) \rightarrow \mathcal{P}(\mathbf{Pt}(\Sigma))$:

$$\begin{aligned}
\Phi_s(X) &=_{df} \{ \zeta \circ_{\Sigma} a : a \in A_{\Sigma} \ \& \ \zeta \circ_{\Sigma} a \in \mathbf{T}_{\mathcal{G}}(s) \ \& \\
&\quad c(\zeta) \cap c_{<_{\Sigma} a}^{-1}(e_a) = \emptyset \} \cup X \quad (1)
\end{aligned}$$

where we assume that $E_{\zeta \circ_{\Sigma} a} = E_{\zeta} \cup \{e_a\}$, and $e_a \notin E_{\zeta}$ (i. e. e_a is the event that is added to ζ if composed with a), and $c(\zeta)$ denotes the set of cut-off events of ζ . More precisely, there is some $n < \omega$ such that $\Phi_s^n(\{\varepsilon\}) = \mathbf{lfp}[\Phi_s, \{\varepsilon\}] = \{\xi\}$, and ξ is local complete at s .

Since the result of the iteration process is unique, we write $\zeta = \mathbf{lfp}[\Phi_s, \varepsilon]$. Fig. 5 shows an algorithm which computes a local complete pomtree according to Theorem 5.

Remark 4. Local completeness has been elaborated in the context of *safe Petri nets* by McMillan [16]. McMillan uses a more efficient definition of cut-off events that replaces the ordering \leq_{ζ} that is used to relate a cut-off event to its predecessor. In [9] it has been shown that McMillans ordering can be further improved, and that the ordering we use is most inefficient (in terms of the size of the resulting complete pomtree) in a hierarchy of orderings. We use this ordering for the sake of simplicity; the results of [9] can easily be applied in the context of our work. \square

4. MORPHISMS

Effective planning for large scaled distributed systems requires (a) an abstraction mechanism that prevents from taking into account an overwhelming amount of details, and (b) a notion of system composition or, more general, the embedding of subsystems into composite systems.

Proofs of lemmas and theorems of this section (and Section 6) can be found in [8], if not stated otherwise.⁴

Abstractions. Let Σ_1 and Σ_2 be distributed alphabets, and $\mathcal{G}_1 \in \mathbf{I}(\Sigma_1)$, $\mathcal{G}_2 \in \mathbf{I}(\Sigma_2)$. An *abstraction* from \mathcal{G}_1 to \mathcal{G}_2 is a pair of mappings $\alpha : A_{\Sigma_1} \rightarrow A_{\Sigma_2}$ and $\beta : S_{\mathcal{G}_1} \rightarrow S_{\mathcal{G}_2}$ such that (a) $\alpha(a) I_{\Sigma_2} \alpha(b) \Rightarrow a I_{\Sigma_1} b$; (b) if $a_{\mathcal{G}_1}(s_1) = s_2$ then if $\alpha(a)$ is defined then $\alpha(a)_{\mathcal{G}_2}(\beta(s_1)) = \beta(s_2)$ and if $\alpha(a)$ is not defined then $\beta(s_1) = \beta(s_2)$, and (c) $\varphi_{\mathcal{G}_2} \subseteq \beta(\varphi_{\mathcal{G}_1})$. We denote the fact that α, β form an abstraction by $[\alpha, \beta] : \mathcal{G}_1 \rightarrow \mathcal{G}_2$.

Remark 5. One can show that the pair of identity mappings on A_{Σ} and $S_{\mathcal{G}}$ is an abstraction, and that the component-wise composition of abstraction is again an abstraction. Thus, the class of interpretations of d-alphabets together with abstraction arrows forms a category. \square

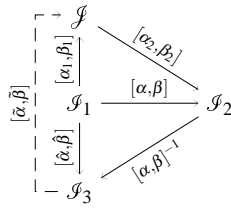
⁴Actually, [8] uses a slightly different version of the notions presented here. These differences do however not jeopardize the results presented here.

Example 5. To obtain an abstraction for our running Example 2, use abstract states of the form $\langle P, R, K \rangle$, with the abstraction mapping $\beta : \langle P, R, \sigma \rangle \mapsto \langle P, R, \sigma, \max_{q \in Q} \sigma(x_q) \rangle$. We furthermore have to redefine the state invariant to $\varphi =_{\text{df}} \{ \langle P, R, \sigma, n \rangle : n < \min_{q \in Q} N_q \}$. We suppose that the concrete and the abstract d-alphabet Σ coincides, and put $\alpha =_{\text{df}} \text{id}_{A_\Sigma}$.

The reader may ask whether the set of abstract states is indeed more “simple” than that of concrete states. The idea here is that we use the state component n as a “quality indicator” to assess the suitability of a state. We furthermore encode the “correlation process” of n (i.e. $n = \max_{q \in Q} \sigma(x_q)$) directly into the abstract states by keeping the variable state σ , because the abstract version of an interpretation based on quality indicators needs to have a way to estimate the effect of a particular (abstract) system action to the values of these indicators. The resulting complicatedness of abstract states is thus due to the simplicity of the running example. \square

Let $[\alpha, \beta] : \mathcal{I}_1 \rightarrow \mathcal{I}_2$ be an abstraction for interpretations $\mathcal{I}_1 \in \mathbf{I}(\Sigma_1)$, $\mathcal{I}_2 \in \mathbf{I}(\Sigma_2)$. A refinement $[\alpha, \beta]^{-1} : \mathcal{I}_2 \rightarrow \mathcal{I}$ for $\mathcal{I} \in \mathbf{I}(\Sigma_1)$ is defined by: (a) $S_{\mathcal{I}} =_{\text{df}} \{ \beta^{-1}(s) : s \in S_{\mathcal{I}_2} \text{ \& } \beta^{-1}(s) \neq \emptyset \}$, (b) $a_{\mathcal{I}}(\beta^{-1}(s_1)) = \beta^{-1}(s_2) \Leftrightarrow \exists a' \in \alpha^{-1}(a). a'_{\mathcal{I}_2}(s'_1) = s'_2$, and (c) $\varphi_{\mathcal{I}} =_{\text{df}} \{ \beta^{-1}(s) : s \in \varphi_{\mathcal{I}_2} \}$.

LEMMA 6. Let $[\alpha, \beta] : \mathcal{I}_1 \rightarrow \mathcal{I}_2$ be an abstraction and let $[\alpha, \beta]^{-1} : \mathcal{I}_2 \rightarrow \mathcal{I}$ be the associated refinement. Then (a) there is an abstraction $[\hat{\alpha}, \hat{\beta}] : \mathcal{I}_1 \rightarrow \mathcal{I}_3$; and moreover, (b) if $[\alpha_1, \beta_1] : \mathcal{I}_1 \rightarrow \mathcal{I}$ and $[\alpha_2, \beta_2] : \mathcal{I} \rightarrow \mathcal{I}_2$ be abstractions, then there is a unique abstraction $[\tilde{\alpha}, \tilde{\beta}] : \mathcal{I}_3 \rightarrow \mathcal{I}$ that makes the diagram displayed right commute.



PROOF. Put $\hat{\alpha} =_{\text{df}} \text{id}_{A_{\Sigma_1}} \upharpoonright \alpha^{-1}(A_{\Sigma_2})$ and $\hat{\beta}(s) =_{\text{df}} \beta^{-1}(s)$. It is an easy exercise to show the required properties. \square

Thus the refined system \mathcal{I}_2 is more abstract than the original system \mathcal{I}_1 , but it is nevertheless the “most concrete” system that can be obtained by any refinement process.

We now are going to investigate the relationship between abstractions and pomset languages of interpretations. For that, we first have to explain how a mapping on action sets of distributed alphabets can be extended to pomsets over these alphabets.

algorithm GENERATE(\mathcal{I}, s_0) is

```

1  input  $\mathcal{I} \in \mathbf{I}(\Sigma), s_0 \in S_{\mathcal{I}}$ ;
2  output a locally complete pomtree  $\zeta \in \mathbf{T}_{\mathcal{I}}(s)$  at  $s$ ;
3   $\zeta \leftarrow \varepsilon; s \leftarrow s_0; c \leftarrow \emptyset; A \leftarrow \text{en}_{\mathcal{I}}(s)$ ;
4  while  $A \neq \emptyset$ 
5  do select  $a \in A$ ;
6     if  $\langle \zeta \circ_{\Sigma} a \rangle(e_a) \cap c = \emptyset$ 
7     then  $\zeta \leftarrow \zeta \circ_{\Sigma} a$ ;
8         select  $u \in \mathbf{Cf}(\zeta)$ ;
9          $\exists b \in A_{\Sigma}. u \circ_{\Sigma} b \in \mathbf{L}_{\mathcal{I}}(s_0) \setminus \mathbf{Cf}(\zeta)$ ;
10        on failure output  $\zeta$ ;
11         $s \leftarrow u_{\mathcal{I}}(s_0); A \leftarrow \text{en}_{\mathcal{I}}(s)$ ;
12        if  $e_a \in c(\zeta)$ 
13        then  $c \leftarrow c \cup \{e_a\}$ ;
14    else  $A \leftarrow A \setminus \{a\}$ ;
15  output  $\zeta$ .
```

Figure 5: Pomset generation procedure

Let Σ_1, Σ_2 be distributed alphabets such that $\alpha : A_{\Sigma_1} \rightarrow A_{\Sigma_2}$ is a mapping. Then α is inductively extended to a mapping $\alpha^* : \mathbf{Ps}(\Sigma_1) \rightarrow \mathbf{Ps}(\Sigma_2)$ as follows:

$$\begin{aligned} \alpha^*(\varepsilon) &=_{\text{df}} \varepsilon; \\ \alpha^*(u \circ_{\Sigma_1} a) &=_{\text{df}} \begin{cases} \alpha^*(u) \circ_{\Sigma_2} \alpha(a), & \text{if } \alpha(a) \text{ is defined;} \\ \alpha^*(u), & \text{otherwise;} \end{cases} \\ &\quad \text{for finite } u; \\ \alpha^*(u) &=_{\text{df}} \bigvee \{ \alpha^*(v) : v \leq u \text{ \& } v \text{ is finite} \} \text{ for infinite } u. \end{aligned}$$

for all $u \in \mathbf{Ps}(\Sigma_2)$ and $a \in A_{\Sigma_2}$.

The following lemma justifies the relationship of the notions of refinement and abstraction on distributed alphabets and interpretations by means of the pomset languages of interpretations.

LEMMA 7. Assume $[\alpha, \beta] : \mathcal{I}_1 \rightarrow \mathcal{I}_2$ is an abstraction. Then $\alpha^*(\mathbf{L}_{\mathcal{I}_1}(s)) \subseteq \mathbf{L}_{\mathcal{I}_2}(\beta(s))$ for all $s \in S_{\mathcal{I}_1}$.

Embedding. The embedding operation provides us with a generic notion on how to understand a system component as part of a larger system, or the relationship of a system and its environment. Note that the composition of a number of system components can be understood as the embedding of each of these components into the composed system. Therefore, embeddings give us a generic notion of composition without assuming any specific composition operation such as synchronous or asynchronous communication.

Let $\mathcal{I}_1 \in \mathbf{I}(\Sigma_1)$ and $\mathcal{I}_2 \in \mathbf{I}(\Sigma_2)$ and let $f : A_{\Sigma_1} \rightarrow A_{\Sigma_2}$ be an injective mapping and $g : S_{\mathcal{I}_2} \rightarrow S_{\mathcal{I}_1}$ be a surjective mappings such that for all $a, b \in A_{\Sigma_1}$ (a) $f(a) I_{\Sigma_2} f(b) \Rightarrow a I_{\Sigma_1} b$, (b) $f(a)_{\mathcal{I}_2}(s_1) = s_2 \Rightarrow a_{\mathcal{I}_1}(g(s_1)) = g(s_2)$, and moreover (c) $\varphi_{\mathcal{I}_1} \subseteq \beta(\varphi_{\mathcal{I}_2})$. We write $(f, g) : \mathcal{I}_1 \rightarrow \mathcal{I}_2$ in this case.

Remark 6. Since id_{A_Σ} and $\text{id}_{S_{\mathcal{I}}}$ form an embedding of $\mathcal{I} \in \mathbf{I}(\Sigma)$ into itself, and moreover the component-wise composition of embeddings forms an embedding, the class of interpretations of d-alphabets together with embeddings as arrows is a category. \square

Example 6. In our example process algebra, $f = \text{id}_{A_\Sigma}$ and $g : \langle U \parallel P, R, \sigma \rangle \mapsto \langle U, R, \sigma \rangle$ form an embedding, where U refers to all the terms “derivable” from U_q^x (compare Figure 2). \square

5. SUPERVISION

Supervision is based on observing the processing of some system, and on interacting if the system enters an undesired state (which could be an error state, a performance bottleneck, etc.). Interaction is done by restricting the possible behaviors of the system under supervision to those alternatives that will lead the system back into a desired state. Restricting behavioral alternatives bases on the disabling of actions of the system under supervision. But of course, not all actions can be controlled in this way. System actions like interrupts, timeouts, are beyond the control of a supervisor. If we consider a component of a larger environment, then from the point of the component, actions of the environment are not controllable (this may differ if the environment is considered).

Hence, if Σ is a d-alphabet, we partition its alphabet set A_Σ into two disjoint sets $A_{\Sigma, c}$ and $A_{\Sigma, uc}$ of *controllable* and *uncontrollable* actions, respectively. Furthermore, we put $\Sigma_c =_{\text{df}} \langle A_{\Sigma, c}, I_\Sigma \cap A_{\Sigma, c} \times A_{\Sigma, c} \rangle$, and $\Sigma_{uc} =_{\text{df}} \langle A_{\Sigma, uc}, I_\Sigma \cap A_{\Sigma, uc} \times A_{\Sigma, uc} \rangle$.

Example 7. In our example, actions $!g_{s,i}, !d_{s,i}, !p_{s,i}$ are controllable. The behavior of the user processes are not controllable, as well receiving events of messages from user processes. \square

Our approach is similar to that of *Supervisory Control Theory* (SCT) [23], with the difference that we do not assume “strong” control in the sense that a supervisor is active all the time, i. e. the supervised system is not able to exhibit undesired alternatives at all. SCT has originally been presented using finite state machines as underlying system models. Recently, extensions to infinite state systems have been considered [3, 15, 10].

Controllability. A pomset $w \in \mathbf{Ps}^w(\Sigma)$ is called *controllable* if $\{\bar{e} : e \in \min_{\leq_w} E_w\} \cap A_{\Sigma,uc} \neq \emptyset$. Note that if w is controllable, then it might have a suffix (i. e. some v such that there is an u with $w = u \circ_{\Sigma} v$) that is not controllable. On the other hand, uncontrollable pomsets may have controllable suffices.

Runs. A pomset $w \in \mathbf{Ps}^w(\Sigma)$ is called a *run* of some $\mathcal{F} \in \mathbf{I}(\Sigma)$ at some $s \in S_{\mathcal{F}}$ if $w \in \mathbf{L}_{\mathcal{F}}(s)$ and either it is infinite or it is finite and moreover, $(w \circ_{\Sigma} a)_{\mathcal{F}}$ is undefined for all $a \in A_{\Sigma}$.

Reactive Supervision Problem. Suppose some state s of a system $\mathcal{F} \in \mathbf{I}(\Sigma)$ that violates the state invariant $\varphi_{\mathcal{F}}$. The *reactive supervision problem* consists of the computation of maximum a pomtree ζ comprising of configurations u that, if executed by \mathcal{F} , will lead the system back into some state $s' \in \varphi_{\mathcal{F}}$, and moreover, if there is a pomset that is also executable at s but does not belong to the configurations of ζ (i. e. does not lead the system back into a desired state), then its execution can be prohibited by a supervisor.

More formally: Let $s \in S_{\mathcal{F}}$ such that $s \notin \varphi_{\mathcal{F}}$. A *reactive plan* for s is a finite pomtree $\zeta \in \mathbf{T}_{\mathcal{F}}(s)$ such that

- R₁. $u_{\mathcal{F}}(s) \in \varphi_{\mathcal{F}}$ for all $u \in \mathbf{Cf}_V(\zeta)$;
- R₂. $u_{\mathcal{F}}(s) \notin \varphi_{\mathcal{F}}$ for all $u \in \mathbf{Cf}(\zeta) \setminus \mathbf{Cf}_V(\zeta)$, and moreover,
- R₃. for all $u \in \mathbf{Cf}(\zeta) \setminus \mathbf{Cf}_V(\zeta)$ it holds: If there is some pomset $v \in \mathbf{L}_{\mathcal{F}}(u_{\mathcal{F}}(s))$ such that $u \circ_{\Sigma} v \notin \mathbf{Cf}(\zeta)$, then v is controllable.

A valid plan ζ is called *most permissive* at s if for each valid plan $\xi \in \mathbf{T}_{\mathcal{F}}(s)$ we have $\zeta \leq \xi \Rightarrow \zeta = \xi$.

Example 8. In order to make the system of Example 2 “supervisable” we have to add some “management functions”: We re-define the following equations:

$$U_q^x = \text{df } \sum_{j \in J} !r_{q,x,j} \cdot \{ ?g_{q,x} \cdot (?p_{q,x} \cdot u_{q,x} \cdot !e_{q,x} + ?d_{q,x}) + ?d_{q,x} \} \cdot \mathbf{0}_x$$

$$V_q^x = \text{df } x_q \leftarrow x_q + 1 \cdot (!p_{q,x} \cdot ?e_{q,x} + !d_{q,x}) \cdot x_q \leftarrow x_q - 1 \cdot \mathbf{0}$$

Let $J = \{c\}$, $Q = \{q\}$, $N_q = 1$, and

$$U_i = \text{df } \{ ?p_{s,i} \cdot u_{s,i} \cdot !e_{s,i} + ?d_{s,i} \} \cdot \mathbf{0}_i \text{ and}$$

$$V_i = \text{df } (!p_{s,i} \cdot ?e_{s,i} + !d_{s,i}) \cdot x_s \leftarrow x_s - 1 \cdot \mathbf{0}$$

for $i = 1, 2, 3$. Consider the concrete state $s = \langle P; R \setminus \{1, 2, 3\}, x_q \mapsto 3 \rangle$ with

$$P = \prod_{i=1}^3 U_i \parallel \prod_{i=1}^3 V_i \parallel C_c \parallel G$$

We have $\beta(s) = \langle P, R \setminus \{1, 2, 3\}, x_q \mapsto 3, 3 \rangle \notin \varphi$. A most permissive valid plan at $\beta(s)$ is

$$\bigvee_i \left(!d_{s,i} \circ_{\Sigma} x_q \leftarrow x_q - 1 \circ_{\Sigma} \bigvee_{j \neq i} (!d_{s,j} \circ_{\Sigma} x_q \leftarrow x_q - 1) \right)$$

for $i, j = 1, 2, 3$. \square

Pro-active Supervision Problem. Dually, we can define the *pro-active supervision problem* as the problem to determining all finite behavioral alternatives that lead to a state violating the state invariant and that cannot be prohibited by a supervisor. In formal terms: Let $s \in S_{\mathcal{F}}$ such that $s \in \varphi_{\mathcal{F}}$. A *violation prediction* for s is a finite pomtree $\zeta \in \mathbf{T}_{\mathcal{F}}(s)$ such that

- P₁. $u_{\mathcal{F}}(s) \notin \varphi_{\mathcal{F}}$ for all $u \in \mathbf{Cf}_V(\zeta)$;
- P₂. $u_{\mathcal{F}}(s) \in \varphi_{\mathcal{F}}$ for all $u \in \mathbf{Cf}(\zeta) \setminus \mathbf{Cf}_V(\zeta)$, and moreover,
- P₃. for all $u \in \mathbf{Cf}(\zeta) \setminus \mathbf{Cf}_V(\zeta)$: If there is some pomset $v \in \mathbf{L}_{\mathcal{F}}(u_{\mathcal{F}}(s))$ such that $u \circ_{\Sigma} v \notin \mathbf{Cf}(\zeta)$, then v is controllable.

A violation prediction ζ is called *complete* at s if for each violation prediction $\xi \in \mathbf{T}_{\mathcal{F}}(s)$ we have $\zeta \leq \xi \Rightarrow \zeta = \xi$.

Solutions. To meet the assumptions of Theorem 5 assume that for a given state s the set $R_s = \text{df } \{s' \in \mathcal{F} : \exists u \in \mathbf{L}_{\mathcal{F}}(s) \cdot u_{\mathcal{F}}(s) = s'\}$ is finite, and moreover $\text{en}_{\mathcal{F}}(s')$ is finite for all $s \in R_s$.

We now define—in addition to the functional Φ_s given by Equation (1)—a number of functionals acting on finite pomtrees:

$$\Psi_{s,\Sigma}(\zeta) = \text{df } \bigvee \{ u[E_u \setminus \leq_u(E)] : \\ a \in A_{\Sigma} \ \& \ u \in \mathbf{Cf}_V(\zeta) \ \& \ u \circ_{\Sigma} a \notin \mathbf{Cf}_V(\zeta) \},$$

where $E = \text{df } \max_{\leq_u} \lambda_u^{-1}(D_{\Sigma}(a))$

$$\Theta_{s,S}(\zeta) = \text{df } \bigvee_{\leq} \min \{ u \in \mathbf{Cf}(\zeta) : v_{\mathcal{F}}(s) \in S \}$$

$$\Omega_{\Sigma,S}(s) = \text{df } \Theta_{s,S}(\mathbf{lfp}[\Psi_{s,\Sigma}, \mathbf{lfp}[\Phi_s, \varepsilon]])$$

THEOREM 6. *If $s \in S_{\mathcal{F}} \setminus \varphi_{\mathcal{F}}$ for some $\mathcal{F} \in \mathbf{I}(\Sigma)$, the $\Omega_{\Sigma, \varphi_{\mathcal{F}}}(s)$ is a most permissive reactive plan for s .*

PROOF. To see that $\Omega_{\Sigma, \varphi_{\mathcal{F}}}(s)$ is a reactive plan, recall from Theorem 4 that $\zeta = \mathbf{lfp}[\Phi_s, \varepsilon]$ is globally complete. $\Psi_{s,\Sigma}(\zeta)$ deletes all maximum uncontrollable events e from ζ including all immediate predecessors of e . Thus a situation in which e is executed cannot longer occur. Thus $\mathbf{lfp}[\Psi_{s,\Sigma}, \zeta]$ fulfills property (R₃). The final application of $\Theta_{s, \varphi_{\mathcal{F}}}$ ensures properties (R₁) and (R₂).

To prove maximum permissiveness of $\zeta = \Omega_{\Sigma, \varphi_{\mathcal{F}}}(s)$ assume there is a valid plan ξ such that $\zeta < \xi$. Since $\mathbf{lfp}[\Phi_s, \varepsilon]$ is state complete, we can assume that $\xi \leq \mathbf{lfp}[\Phi_s, \varepsilon]$ is true. Then we have a configuration $u \in \mathbf{Cf}_V(\zeta)$ such that one of the following is true: (a) There is some $a \in A_{\Sigma}$ such that $u \circ_{\Sigma} a \in \mathbf{Cf}(\xi)$. This however violates property (R₂), and ξ is not a reactive plan. (b) So assume $u \circ_{\Sigma} a \notin \mathbf{Cf}(\xi)$. Then there is some pomset v such that there are $e \in \min_{\leq_v}(E_v)$, $e' \in E_v$, $\bar{e} = a$, and $e <_v e'$, such that both a and e' are controllable. But in this case a (more precisely, the event corresponding to a) would not have been removed from $\mathbf{lfp}[\Phi_s, \varepsilon]$ in the computation process of $\mathbf{lfp}[\Psi_{s,\Sigma}, \mathbf{lfp}[\Phi_s, \varepsilon]]$. \square

THEOREM 7. *If $s \in \varphi_{\mathcal{F}}$ for some $\mathcal{F} \in \mathbf{I}(\Sigma)$, the $\Omega_{\Sigma, S_{\mathcal{F}} \setminus \varphi_{\mathcal{F}}}(s)$ is a complete violation prediction for s .*

PROOF. Dualize the proof of Theorem 6. \square

6. PLANNING WITH ZOOMS

A Galois Connection. In order to employ the the Cousot & Cousot framework of AI, we need to find an appropriate Galois connection:

THEOREM 8. *Let $[\alpha, \beta] : \mathcal{F}_1 \rightarrow \mathcal{F}_2$ be an abstraction. Then $\alpha^*, \gamma : \mathcal{P}(\mathbf{L}_{\mathcal{F}_1}(s)) \Rightarrow \mathcal{P}(\mathbf{L}_{\mathcal{F}_2}(\beta(s)))$ is a Galois connection, where*

$\gamma: \mathbf{L}_{\mathcal{G}_2}(\beta(s)) \rightarrow \mathcal{P}(\mathbf{L}_{\mathcal{G}_1}(s))$, and $\gamma(u)$ is defined to be the smallest set satisfying the following conditions:

$$\begin{aligned} \varepsilon \in \gamma(\varepsilon) \\ u \in \gamma(v) \ \& \ a \in [d']_\alpha \Rightarrow u \circ_{\Sigma_1} a \in \gamma(v \circ_{\Sigma_2} d') \text{ for finite } u, v; \\ \bigvee_{C \in \text{Mc}(X)} C \in \gamma(v) \text{ for infinite } u; \end{aligned}$$

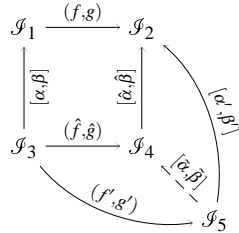
where $X =_{df} \bigcup \{\gamma(w) : w < v \ \& \ w \text{ is finite}\}$, and $\text{Mc}(X)$ denotes the set of all maximum chains in X , that is the set of all maximum subsets $C \subseteq X$ such that $u, v \in C \Rightarrow u \leq v \vee v \leq u$.

PROOF. Assume $\alpha^*(X) \subseteq Y$. By Lemma 6 we have $\gamma(\alpha^*(X)) \subseteq \mathbf{L}_{[\alpha, \beta]^{-1}(\mathcal{G}_2)}(\beta(s))$, hence $X \subseteq \gamma(\alpha^*(X)) \subseteq \gamma(Y)$. Assume $X \subseteq \gamma(Y)$. Then $\alpha^*(X) \subseteq Y$ by Lemma 7. \square

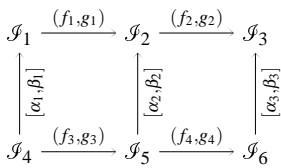
Remark 7. Theorem 8 is formulated in terms of chains because $\gamma(v)$ yields a set of possible concretization of v that are not necessarily comparable by \leq . Thus in ‘‘jumping’’ from finite to infinite pomsets, the sets to build infinite joins need to be made explicit.

Zooms. We now consider the following situation. Let $(f, g) : \mathcal{F}_1 \rightarrow \mathcal{F}_2$ be an embedding and let $[\beta, \alpha] : \mathcal{F}_3 \rightarrow \mathcal{F}_1$ be an abstraction. The question arises whether we can use the embedding process (described by (f, g)) also to embed the refined version of \mathcal{F}_1 into (a refined version) of \mathcal{F}_2 ? Such an operation would provide us with a *magnifying glass*, a notion of local refinement or *zooming*. Of course, we are not looking for an arbitrary zooming operation but for a universal one in the sense that the resulting refined system is the most abstract one that embeds the interpretation \mathcal{F}_3 .

THEOREM 9. Let $\mathcal{F}_i \in \mathbf{I}(\Sigma_i)$ interpretations for $i = 1, 2, 3$ such that $(f, g) : \mathcal{F}_1 \rightarrow \mathcal{F}_2$ is an embedding and $[\alpha, \beta] : \mathcal{F}_3 \rightarrow \mathcal{F}_1$ is an abstraction. Then there is an interpretation $\mathcal{F}_4 \in \mathbf{I}(\Sigma_4)$, an embedding $(\hat{f}, \hat{g}) : \mathcal{F}_3 \rightarrow \mathcal{F}_4$, and an abstraction $[\hat{\alpha}, \hat{\beta}] : \mathcal{F}_4 \rightarrow \mathcal{F}_2$ such that for all interpretations \mathcal{F}_5 and all embeddings $(f', g') : \mathcal{F}_3 \rightarrow \mathcal{F}_5$ and abstractions $[\alpha', \beta'] : \mathcal{F}_5 \rightarrow \mathcal{F}_2$ there is a uniquely defined abstraction $[\tilde{\alpha}, \tilde{\beta}] : \mathcal{F}_5 \rightarrow \mathcal{F}_4$. In other words, $[\tilde{\alpha}, \tilde{\beta}]$ is the only abstraction arrow that makes the diagram displayed right commute.



The following theorem states that zooms can be composed.



THEOREM 10. Consider the diagram show left. If the inner squares form two zooms, then the outer rectangle is also a zoom. If the right hand side inner squares and the outer rectangle are zooms, then the left hand side inner square is also.

Planning with zooms. To prepare the results presented in this paragraph, we have to ensure that controllability remains invariant under abstractions and embeddings. Therefore, if $[\alpha, \beta] : \mathcal{F}_1 \rightarrow \mathcal{F}_2$ is an abstraction for $\mathcal{F}_1 \in \mathbf{I}(\Sigma_1)$ and $\mathcal{F}_2 \in \mathbf{I}(\Sigma_2)$, then we stipulate that $\alpha(A_{\Sigma_1, \text{uc}}) \subseteq A_{\Sigma_2, \text{uc}}$ does hold, since a more concrete event which is not controllable cannot be made controllable by abstraction. If $(f, g) : \mathcal{F}_1 \rightarrow \mathcal{F}_2$ is an embedding, we have to assume

1. If a state $s \notin S_{\mathcal{G}}$ is detected, compute a most permissive valid plan ξ for s .
2. Stop if step (1) does not yield a result.
3. For each localized environment \mathcal{E}_i , compute a prefix $\zeta_i \leq \gamma(\xi)$ such that $\zeta_i \in \mathbf{T}_{\mathcal{G}_i}(s')$, where $s' \in \beta^{-1}(s)$ is the state that is detected on the level of \mathcal{E}_i . Note that $s' \notin \varphi_{\mathcal{E}_i}$. We know that $u_{\mathcal{E}_i}(s') \in \varphi_{\mathcal{E}_i}$ for each configuration $u \in \mathbf{Cf}_v(\zeta_i)$, but ζ_i is not necessarily a valid plan, or most permissive.
4. Stop if step (3) does not yield a result.

Figure 6: Planning algorithm

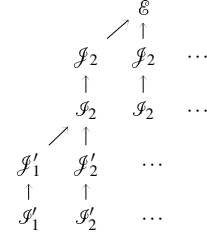
$f(A_{\Sigma_1, c}) \subseteq A_{\Sigma_2, c}$, since actions that are controllable in \mathcal{F}_2 are not necessarily controllable in \mathcal{F}_1 .

Consider again the zoom diagram displayed in Theorem 9. Let $\zeta \in \mathbf{T}_{\mathcal{F}_2}(s)$ be a finite pomtree for some state $s \in S_{\mathcal{F}_2}$. Since the converse of Lemma 7 does not hold, $\gamma(\zeta)$ is not necessary an element of $\mathbf{T}_{\mathcal{F}_4}(s')$ (where γ_i is the adjoint in the Galois connection $\hat{\alpha}_i^*, \gamma_i : \mathbf{L}_{\mathcal{F}_2}(s) \equiv \mathbf{L}_{\mathcal{F}_4}(s)$ and $s' \in [s]_{\hat{\beta}}$). We however have:

THEOREM 11. Assume the notions of Theorem 9. Let $s \in S_{\mathcal{F}_4}$ be a state such that $s \notin \varphi_{\mathcal{F}_4}$ and let $\zeta \in \mathbf{T}_{\mathcal{F}_4}(s)$ be a most permissive valid plan for s . If $\xi \in \mathbf{T}_{\mathcal{F}_2}(\beta(s))$ is a most permissive valid plan for $\beta(s)$, then $\hat{\alpha}^*(\zeta) \leq \xi$.

The proof of the theorem is rather tedious, also the theorem itself is rather obvious. It bases on the idea that if ζ resolves a problem on a concrete level (i. e. in \mathcal{F}_4), then the problem on an abstract level (i. e. in \mathcal{F}_2). But since ξ is most permissive we obtain $\hat{\alpha}^*(\zeta) \leq \xi$.

In practice, we will apply zooms on a number of embedded components, i. e. we will use a families $\{[\alpha_i, \beta_i] : \mathcal{F}_i \rightarrow \mathcal{F}_i\}_{i \in I}$ and $\{(f_i, g_i) : \mathcal{F}_i \rightarrow \mathcal{E}_i\}_{i \in I}$, obtaining families of completions $\{[\hat{\alpha}_i, \hat{\beta}_i] : \mathcal{E}_i \rightarrow \mathcal{E}_i\}_{i \in I}$ and $\{(\hat{f}_i, \hat{g}_i) : \mathcal{F}_i \rightarrow \mathcal{E}_i\}_{i \in I}$ (compare figure right). In this case, we call the interpretation \mathcal{E} a *global environment*, and the interpretations \mathcal{E}_i *localizations* of \mathcal{E} . This situation can be



iterated forming a tree structure of interpretations that are increasingly more concrete—and more localized—towards the leaves of the tree by letting each of the interpretations \mathcal{F}_i the environment of another set of zoom diagrams.

Theorem 11 provides us now (indirectly) with a planning algorithm utilizing zooms as shown in Figure 6. Again, an algorithm for violation prediction with zooms can be obtained by dualizing.

Planning with zooms is done top-down, which has the following implications: The choice of the abstractions $\{[\alpha_i, \beta_i] : \mathcal{F}_i \rightarrow \mathcal{F}_i\}_{i \in I}$ is very sensible. Choosing a version of \mathcal{F}_i that is ‘‘too abstract’’ might render to many abstract action as uncontrollable because some of their concrete counterparts are not controllable, leading to the situation that a valid plan cannot be computed at all. On the other hand, using abstractions that are ‘‘too concrete’’ leads to complexity problems and to non-feasible planning algorithms.

What happens if the execution of the algorithm is terminated in Step (2) because of the fact that too many uncontrollable events exist? If \mathcal{E} is in fact the root of our model tree, we have to conclude that we cannot resolve the problem situation. But an action that is uncontrollable in one component may be controllable in another

one. Thus if there is a model level available *above* \mathcal{E} —say \mathcal{E}' —we can try to repeat the planning algorithm on this level. Since \mathcal{E}' embeds other components besides \mathcal{E} , controllable actions of these components are available to resolve the problem situation on the higher level.

We conclude that a hierarchical structure of the system models that are used for supervision allows to resolve problems on a level that is as local as possible.

For pro-active planing, the situation is similar. Suppose there is a complete violation prediction ζ for some state $s \in S_{\mathcal{E}}$. Again, we have to “concretize” ζ into the localized environments \mathcal{E}_i to determine whether ζ can be executed locally. We omit the details.

7. SUMMARY AND FURTHER WORK

In this paper, we have presented an approach to perform reactive as well as pro-active planning for distributed autonomic systems comprising components that exhibit operational models of themselves. Using the embedding and abstraction morphisms, it had been shown that there is a “most abstract” concretization of an abstracted system component under the embedding into some environment, or—using a more colorful language—it is possible to use component abstractions as “magnification glasses” to zoom into models without inventing details.

Reactive supervision has been defined as the problem to lead a system from an undesired state back into a desired one under the assumption of partial controllability. Pro-active supervision is about the prediction of problem states that cannot be prevented because of the lack of controllability.

We presented algorithms based on the Cousot & Cousot framework of Abstract Interpretation to solve both problems. We further elaborated how these solutions can be applied in the context of a hierarchical system model leading to an iterative planning algorithm that uses the most local view to the system that is possible (under a given set of abstractions/embeddings).

A number of open questions remain, the most immanent are:

1. How to incorporate widening/narrowing operators for Abstract Interpretation (Remark 1)?
2. How to define suitable abstractions (embeddings are usually given by the structure of the SUS)? Is there a way to “adjust” abstractions that are too coarse or too fine?
3. Predictions that incorporate additional probabilities (expressing how likely it is that an uncontrollable action occurs) would be much more useful.

8. REFERENCES

- [1] M. Arbib and E. Manes. *Arrows, Structures, and Functors: The Categorical Imperative*. Academic Press, New York, 1975.
- [2] G. Birkhoff. *Lattice Theory*. American Mathematical Society, Providence, Rhode Island, 3rd edition, 1967.
- [3] Y.-L. Chen and F. Lin. Safety control of discrete event systems using finite state machines with parameters. In *Proc. of American Control Conference 2001*, pages 975–980, Arlington, VA, June 2001.
- [4] P. Cousot and R. Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Conference Record of the Fourth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 238–252, Los Angeles, California, 1977. ACM Press, New York, NY.
- [5] P. Cousot and R. Cousot. Abstract interpretation frameworks. *Journal of Logic and Computation*, 2(4):511–547, 1992.
- [6] P. Cousot and R. Cousot. Comparing the Galois connection and widening/narrowing approaches to abstract interpretation, invited paper. In M. Bruynooghe and M. Wirsing, editors, *Proceedings of the International Workshop Programming Language Implementation and Logic Programming, PLILP '92*, Leuven, Belgium, 13–17 August 1992, Lecture Notes in Computer Science 631, pages 269–295. Springer-Verlag, Berlin, Germany, 1992.
- [7] P. H. Deussen. *Analyse verteilter Systeme mit Hilfe von Prozessautomaten*. PhD thesis, Brandenburg Technical Univ. of Cottbus, 2001. In German.
- [8] P. H. Deussen. A mathematical framework for pervasive supervision. Cascadas project deliverable D2.1, Fraunhofer Institute for Open Communication Systems, 2006.
- [9] J. Esparza, S. Römer, and W. Vogler. An improvement of McMillan’s unfolding algorithm. In *TACAS*, pages 87–106, 1996.
- [10] B. Gaudin and P. H. Deussen. Supervisory control on concurrent discrete event systems with variables. In *Proc. 2007 American Control Conference*, July 2007.
- [11] R. Goldblatt. *Topoi, the Categorical Analysis of Logic*. North Holland Publishing Company, 1979.
- [12] J. Grabowski. On partial languages. *Fund. Inform.*, 4(2):427–498, 1981.
- [13] J. Kephart and D. Chess. The vision of autonomic computing. *IEEE Computer*, 36(1):41–52, 2003.
- [14] S. M. Lane. *Categories for the Working Mathematician*, volume 5 of *Graduate Texts in Mathematics*. Springer Verlag, New York, 1977.
- [15] T. Le Gall, B. Jeannot, and H. Marchand. Supervisory control of infinite symbolic systems using abstract interpretation. In *Proc. of the 44th IEEE Conference on Decision and Control and Control (CDC'05) and European Control Conference (ECC 2005)*, Sevilla (Spain), December 2005.
- [16] K. L. McMillan. A technique of state space search based on unfolding. *Formal Methods in System Design: An International Journal*, 6(1):45–65, January 1995.
- [17] R. Milner. *Communication and Concurrency*. Prentice Hall, 1989.
- [18] M. Nielsen, G. Plotkin, and G. Winskel. Petri nets, event structures and domains, Part I. *Theoretical Computer Science*, 13:85–108, 1981.
- [19] V. Pratt. Modelling concurrency with partial orders. *International Journal of Parallel Programming*, 15(1):33–71, 1986.
- [20] M. Smirnov. Autonomic communication—research agenda for a new communication paradigm. Company whitepaper, Fraunhofer Institute for Open Communication Systems, Berlin, Germany, 2004.
- [21] P. H. Starke. Multiprocessor systems and their concurrency. *J. Inf. Process. Cybern. EIK*, 20(4):207–427, 1984.
- [22] A. Tarski. A lattice-theoretical fixedpoint theorem and its applications. *Pacific Journal of Mathematics*, 5:522–587, 1955.
- [23] W. M. Wonham. Notes on control of discrete-event systems. Tech. rep. ECE 1636F/1637S, Department of Electrical and Computer Engineering, University of Toronto, Toronto, July 2003.