



Centrum voor Wiskunde en Informatica

REPORTRAPPORT

Discrete Time Process Algebra and the Semantics of SDL

J.A. Bergstra, C.A. Middelburg, Y.S. Usenko

Software Engineering (SEN)

SEN-R9809 June 30, 1998

Report SEN-R9809
ISSN 1386-369X

CWI
P.O. Box 94079
1090 GB Amsterdam
The Netherlands

CWI is the National Research Institute for Mathematics and Computer Science. CWI is part of the Stichting Mathematisch Centrum (SMC), the Dutch foundation for promotion of mathematics and computer science and their applications.

SMC is sponsored by the Netherlands Organization for Scientific Research (NWO). CWI is a member of ERCIM, the European Research Consortium for Informatics and Mathematics.

Copyright © Stichting Mathematisch Centrum
P.O. Box 94079, 1090 GB Amsterdam (NL)
Kruislaan 413, 1098 SJ Amsterdam (NL)
Telephone +31 20 592 9333
Telefax +31 20 592 4199

Discrete Time Process Algebra and the Semantics of SDL

J.A. Bergstra

*Programming Research Group, University of Amsterdam
P.O. Box 41882, 1009 DB Amsterdam, The Netherlands*

and

*Department of Philosophy, Utrecht University
P.O. Box 80126, 3508 TC Utrecht, The Netherlands*

Email: janb@wins.uva.nl

C.A. Middelburg

CWI

P.O. Box 94079, 1090 GB Amsterdam, The Netherlands

and

*Department of Philosophy, Utrecht University
P.O. Box 80126, 3508 TC Utrecht, The Netherlands*

Email: cam@cw.nl

Y.S. Usenko

*Department of Theoretical Cybernetics, Taras Shevchenko University
2/6, Glushkov Avenue, Kiev, Ukraine*

Email: yarick@ysu.kiev.ua

ABSTRACT

We present an extension of discrete time process algebra with relative timing where recursion, propositional signals and conditions, a counting process creation operator, and the state operator are combined. Except the counting process creation operator, which subsumes the original process creation operator, these features have been developed earlier as largely separate extensions of time free process algebra. The change to the discrete time case and the combination of the features turn out to be far from trivial. We also propose a semantics for a simplified version of SDL, using this extension of discrete time process algebra to describe the meaning of the language constructs. This version covers all behavioural aspects of SDL, except for communication via delaying channels – which can easily be modelled. The semantics presented here facilitates the generation of finitely branching transition systems for SDL specifications and thus it enables validation.

1991 Mathematics Subject Classification: 68Q55, 68Q60

1991 Computing Reviews Classification System: D.2.1, D.3.1, F.3.1, F.3.2

Keywords and Phrases: process algebra, ACP, discrete time, relative timing, semantics, specification language, SDL, asynchronous communication, timers

Note: The work presented in this report has been partly carried out while the second and third author were at UNU/IIST (United Nations University, International Institute for Software Technology).

Note: A slightly adapted version of this report will appear as a chapter in the forthcoming Handbook of Process Algebra (editors J.A. Bergstra, A. Ponse and S.A. Smolka) to be published by Elsevier.

1. INTRODUCTION

In this chapter, we present an extension of discrete time process algebra with relative timing where recursion, propositional signals and conditions, a counting process creation operator, and the state operator are combined. We also propose a semantics for a simplified version of SDL, called φ -SDL, using this extension of discrete time process algebra with relative timing to describe the meaning of the language constructs. The choice of a process algebra in the style of ACP [6, 7] as the basis of the presented semantics is obvious. This algebraic approach to concurrency represents a large body of relevant theory. In particular, many features of φ -SDL are related to topics that have been studied extensively in the framework of ACP. Besides, the axiom system and operational semantics of an ACP-style process algebra facilitate advances in the areas of validation and verification.

We take the remainder of this introductory section to introduce φ -SDL, to motivate the choices made in the selection of this dramatically simplified version of SDL, and to describe its close connection with full SDL. We also explain the need for a semantics that deals properly with the time related aspects of SDL in case one intends to validate SDL specifications, or to justify design steps made using SDL by formal verification.

1.1 Background

At present, SDL [8, 16] is widely used in telecommunications for describing structure and behaviour of generally complex systems at different levels of abstraction. It originated from an informal graphical description technique already commonly used in the telecommunications field at the time of the first computer controlled telephone switches. Our starting-point is the version of SDL defined in [30], the ITU-T Recommendation Z.100 published in 1994. There, a subset of SDL, called Basic SDL, is identified and used to describe the meaning of the language constructs of SDL that are not in Basic SDL. This subset is still fairly complicated.

φ -SDL is a simplified version of Basic SDL.¹ The following simplifications have been made:

- blocks and channels are removed;
- all variables are revealed and they can be viewed freely;
- timer setting is regarded as just a special use of signals;
- timer setting is based on discrete time.

Besides, φ -SDL does not deal with the specification of abstract data types. An algebraic specification of all data types used in a φ -SDL specification is assumed as well as an initial algebra semantics for it. The pre-defined data types **Boolean** and **Natural**, with the obvious interpretation, should be included.

We decided to focus in φ -SDL on the behavioural aspects of SDL. We did so for the following two reasons. Firstly, the structural aspects of SDL are mostly of a static nature and therefore not very relevant from a semantic point of view. Secondly, the part of SDL that deals with the specification of abstract data types is well understood – besides, it can easily be isolated and treated as a parameter.² For practical reasons, we also chose not to include initially procedures, syntypes with a range condition and process types with a bound on the number of instances that may exist simultaneously. Similarly, the **any** expression is omitted as well. Services are not supported by φ -SDL for the following reasons: the semantics of services is hard to understand, ETSI forbids for this reason their use in European telecommunication standards (see [28]), and the SDL community discusses its usefulness.

In [12], we introduced a simplified version of SDL, called φ SDL, which covers all behavioural aspects of SDL, including communication via delaying channels. φ -SDL is φ SDL without communication via delaying channels. The process algebra semantics of φ SDL proposed in [12] made clear that φ SDL specifications can always be transformed to semantically equivalent ones in φ -SDL. Apart from the

¹This subset is called φ -SDL, where φ stands for flat, as it does not cover the structural aspects of SDL, and $-$ indicates that delaying channels are left out.

²The following is also worth noticing: (1) ETSI discourages the use of abstract data types other than the pre-defined ones in European telecommunication standards (see [28]); (2) ASN.1 [29] is widely used for data type specification in the telecommunications field, and there is an ITU-T Recommendation, Z.105, for combining SDL and ASN.1 (see [33]).

data type definitions, SDL specifications can be transformed to φ SDL specifications, and hence to φ^- SDL specifications, provided that no use is made of facilities that are not included initially. The transformation from SDL to φ SDL has, apart from some minor adaptations, already been given. The first part of the transformation is the mapping for the shorthand notations of SDL which is given informally in the ITU-T Recommendation Z.100 [30] and defined in a fully precise manner in its Annex F.2 [31]. The second and final part is essentially the mapping *extract-dict* defined in its Annex F.3 [32].

The semantics of φ^- SDL agrees with the semantics of SDL as far as reasonably possible. This means in the first place that obvious errors in [32] have not been taken over. For example, the intended effect of SDL's create and output actions may sometimes be reached with interruption according to [32] – allowing amongst other things that a process ceases to exist while a signal is sent to it instantaneously. Secondly, the way of dealing with time is considered to be unnecessarily complex and inadequate in SDL and has been adapted as explained below.

In SDL, real numbers are used for times and durations. So when a timer is set, its expiration time is given by a real number. However, the time considered is the system time which proceeds actually in a discrete manner: the system receives ticks from the environment which increase the system time with a certain amount (how much real time they represent is left open). Therefore, the timer is considered to expire when the system receives the first tick that indicates that its expiration time has passed. So nothing is lost by adopting in φ^- SDL a discrete time approach, using natural numbers for times and durations, where the time unit can be viewed as the time between two ticks but does not really rely upon the environment. This much simpler approach also allows us to remove the original inadequacy to relate the time used with timer setting to the time involved in waiting for signals by processes.

We generally had to make our own choices with respect to the time related aspects of SDL, because they are virtually left out completely in the ITU-T Recommendation Z.100. Our choices were based on communications with various practitioners from the telecommunications field using SDL, in particular the communications with Leonard Pruitt [24]. They provided convincing practical justification for the premise of our current choices: communication with the environment takes a good deal of time, whereas internal processing takes a negligible deal of time. Ease of adaptation to other viewpoints on time in SDL is guaranteed relatively well by using a discrete time process algebra, PA_{drt}^- (see [4]) without immediate deadlock, as the basis of the presented semantics.

In the telecommunications field, SDL is increasingly used for describing generally complex telecommunications systems, including switching systems, services and protocols, at different levels of abstraction – from initial specification till implementation. Initial specification of systems is done with the intention to analyse the behavioural properties of these systems and thus to validate the specification. There is also a growing need to verify whether the properties represented by one specification are preserved in another, more concrete, specification and thus to justify design steps. However, neither SDL nor the tools and techniques that are used in conjunction with SDL provide appropriate support for validation of SDL specifications and verification of design steps made using SDL. The main reason is that the semantics of SDL according to the ITU-T Recommendation Z.100 is at some points inadequate for advanced validation and formal verification. In particular, the semantics of time related features, such as timers and delaying channels, is insufficiently precise. Moreover, the semantics is at some other points unnecessarily complex. Consequently, rules of logical reasoning, indispensable for formal verification, have not yet been developed and most existing analysis tools, e.g. GEODE [1] and SDT [34], offer at best a limited kind of model checking for validation.

Prerequisites for advanced validation and formal verification is a dramatically simplified version of SDL and an adequate semantics for it. Only after that possibilities for advanced analysis can be elaborated and proof rules for formal verification devised. The language φ^- SDL and the presented semantics for it are primarily intended to come up to these prerequisites.

1.2 Organization of this chapter

The structure of this chapter is as follows. In Section 2, we present the extension of discrete time process algebra with relative timing that is used for the process algebra semantics of φ -SDL proposed in Section 4. An overview of φ -SDL is given in Section 3. Following the overview, in Section 4, we present the proposed semantics of φ -SDL in two steps. First, a semantics of φ -SDL process definitions, which are the main elements of φ -SDL specifications, is given. This semantics abstracts from dynamic aspects of process behaviour such as process creation and process execution in a state. A semantics of φ -SDL system definitions, i.e. complete φ -SDL specifications, is then given in terms of the semantics of φ -SDL process definitions using the counting process creation operator and the state operator. In Section 5, we give an overview of related work and we explain how the semantics presented in Section 4 can be used to transform φ -SDL specifications to transition systems that can be used for advanced validation. There are appendices about notational conventions used (Appendix A) and details concerning the contexts used to model scope in the presented semantics (Appendix B). Small examples of specification in φ -SDL and the meaning of the process definitions being found in these examples are also presented (in Section 3 and Section 4, respectively).

Acknowledgements

For one month all three authors were at UNU/IIST (United Nations University, International Institute for Software Technology) in Macau. During that period, some of the more important corrections and technical changes of the material presented here were made. We thank Dines Bjørner, the former director of UNU/IIST, for bringing us together in Macau. The third author thanks Radu Şoricuţ and Bogdan Warinschi for their helpful comments and discussions.

2. PROCESS ALGEBRA

2.1 Introduction

In this section, we present an extension of discrete time process algebra with relative timing where recursion, propositional signals and conditions, a counting process creation operator, and the state operator are combined. Its signature, axioms and a structured operational semantics are given and it is shown that strong bisimulation equivalence is a congruence for all operations. Except the counting process creation operator, which subsumes the original process creation operator, these features have been developed earlier as largely separate extensions of time free process algebra. However, both the change to the discrete time case and the combination of the features turn out to be far from trivial. Besides, some of the features are slightly adapted versions of the original ones in order to meet the needs of the semantics of φ -SDL.

In Section 2.2, we present discrete relative time process algebra without immediate deadlock and delayable actions (PA_{drt}^- -ID). In Section 2.3, we add propositional signals and conditions to PA_{drt}^- -ID. In the discrete relative time case, this addition requires some axioms to be refined. We introduce a new guarded command operator that yields a deadlock in the current time slice if the condition does not hold at the start, i.e. waiting is no option if the condition does not hold. In Section 2.4, we add recursion to the extension presented in Section 2.3. The main definitions related to recursion, such as the definitions of recursive specification, solution and guardedness, are given here for the case with relative timing in discrete time as well as propositional signals and conditions.

In Section 2.5 and 2.6, we describe the counting process creation operator and the state operator, respectively, for the discrete relative time case in the presence of propositional signals and conditions. The counting process creation operator is a straightforward extension of the original process creation operator that allows to assign a unique “process identification value” to each process created. The state operator presented here allows to deal with conditions whose truth depends on the state and with state changes due to progress of time to the next time slice.

The main reference to discrete time process algebra in the style of ACP is [4]. The features with which it is combined here are discussed as separate extensions of time free process algebra in [5] (propositional signals and conditions, state operator), [6] (recursion) and [9] (process creation). Our

discussion of axioms is concentrated on the crucial axioms for the discrete time case and each of these features, and on the alterations and additions needed if all this is combined. For a systematic introduction to process algebra in the style of ACP, the reader is referred to [6] and [7].

2.2 Discrete relative time process algebra

In this subsection, we present discrete relative time process algebra without immediate deadlock and delayable actions. The term discrete time is used here to indicate that time is divided into time slices and timing of actions is done with respect to the time slices in which they are performed – within a time slice there is only the order in which actions are performed. Additionally, performance of actions and passage to the next time slice are separated here. This corresponds to the two-phase functioning scheme for modeling timed processes [23]. Note that it means that processes are supposed to be capable of performing certain actions, like in time free process algebra, as well as passing to the next time slice.

First we treat the basic discrete relative time process algebra $\text{BPA}_{\text{drt}}^-$ -ID. Then we treat PA_{drt}^- -ID, the extension of $\text{BPA}_{\text{drt}}^-$ -ID with parallel composition in which no communication between processes is involved. $\text{ACP}_{\text{drt}}^-$ -ID, the extension of $\text{BPA}_{\text{drt}}^-$ -ID with parallel composition in which synchronous communication between processes is involved, will not be treated. $\text{BPA}_{\text{drt}}^-$ -ID, PA_{drt}^- -ID and $\text{ACP}_{\text{drt}}^-$ -ID are presented in detail in [25]. We also present the extension of PA_{drt}^- -ID with encapsulation, described before for the discrete relative time case without immediate deadlock in [3].

2.2.1 Basic process algebra In $\text{BPA}_{\text{drt}}^-$ -ID, we have the sort \mathbf{P} of processes, the constants \underline{a} (one for each action a) and $\underline{\delta}$, the unary operator σ_{rel} (time unit delay), and the binary operators \cdot (sequential composition) and $+$ (alternative composition). The constants \underline{a} stand for a in the current time slice. Similarly, the constant $\underline{\delta}$ stands for a deadlock in the current time slice. The process $\sigma_{\text{rel}}(x)$ is the process x delayed till the next time slice. The process $x \cdot y$ is the process x followed after successful termination by the process y . The process $x + y$ is the process that proceeds with either the process x or the process y , but not both. We also have the auxiliary unary operator ν_{rel} (now) in $\text{BPA}_{\text{drt}}^-$ -ID. This operator makes an elegant axiomatization of PA_{drt}^- -ID possible. The process $\nu_{\text{rel}}(x)$ is the part of x that is not delayed till the next time slice.

It is assumed that a fixed but arbitrary set A of *actions* has been given.

Signature of $\text{BPA}_{\text{drt}}^-$ -ID The signature of $\text{BPA}_{\text{drt}}^-$ -ID consists of the *undelayable action* constants $\underline{a} : \mathbf{P}$ (for each $a \in A$), the *undelayable deadlock* constant $\underline{\delta} : \mathbf{P}$, the *alternative composition* operator $+: \mathbf{P} \times \mathbf{P} \rightarrow \mathbf{P}$, the *sequential composition* operator $\cdot : \mathbf{P} \times \mathbf{P} \rightarrow \mathbf{P}$, the *time unit delay* operator $\sigma_{\text{rel}} : \mathbf{P} \rightarrow \mathbf{P}$, and the *now* operator $\nu_{\text{rel}} : \mathbf{P} \rightarrow \mathbf{P}$.

We assume that an infinite set of variables (of sort \mathbf{P}) has been given. Given the signature of $\text{BPA}_{\text{drt}}^-$ -ID, terms of $\text{BPA}_{\text{drt}}^-$ -ID, often referred to as process expressions, are constructed in the usual way. The need to use parentheses is reduced by ranking the precedence of the operators. Throughout this chapter we adhere to the following precedence rules: (i) all unary operators have the same precedence, (ii) unary operators have a higher precedence than binary operators, (iii) the operator \cdot has the highest precedence amongst the binary operators, (iv) the operator $+$ has the lowest precedence amongst the binary operators, and (v) all other binary operators have the same precedence. We will also use the following abbreviation. Let $(p_i)_{i \in \mathcal{I}}$ be an indexed set of terms of $\text{BPA}_{\text{drt}}^-$ -ID where $\mathcal{I} = \{i_1, \dots, i_n\}$. Then we write $\sum_{i \in \mathcal{I}} p_i$ for $p_{i_1} + \dots + p_{i_n}$. We further use the convention that $\sum_{i \in \mathcal{I}} p_i$ stands for $\underline{\delta}$ if $\mathcal{I} = \emptyset$.

We denote variables by x, x', y, y', \dots . An important convention is that a, a', b, b', \dots denote elements of $A \cup \{\delta\}$ in the context of an equation, but elements of A in the context of an operational semantics rule. Furthermore, H denotes a subset of A .

Axioms of $\text{BPA}_{\text{drt}}^-$ -ID The axiom system of $\text{BPA}_{\text{drt}}^-$ -ID consists of the equations A1-A5, DRT1-DRT4A and DCS1-DCS4 given in Table 1.

$x + y = y + x$	A1	$\sigma_{\text{rel}}(x) + \sigma_{\text{rel}}(y) = \sigma_{\text{rel}}(x + y)$	DRT1
$(x + y) + z = x + (y + z)$	A2	$\sigma_{\text{rel}}(x) \cdot y = \sigma_{\text{rel}}(x \cdot y)$	DRT2
$x + x = x$	A3	$\underline{\underline{\delta}} \cdot x = \underline{\underline{\delta}}$	DRT3
$(x + y) \cdot z = x \cdot z + y \cdot z$	A4	$x + \underline{\underline{\delta}} = x$	DRT4A
$(x \cdot y) \cdot z = x \cdot (y \cdot z)$	A5	$\nu_{\text{rel}}(\underline{\underline{a}}) = \underline{\underline{a}}$	DCS1
		$\nu_{\text{rel}}(x + y) = \nu_{\text{rel}}(x) + \nu_{\text{rel}}(y)$	DCS2
		$\nu_{\text{rel}}(x \cdot y) = \nu_{\text{rel}}(x) \cdot y$	DCS3
		$\nu_{\text{rel}}(\sigma_{\text{rel}}(x)) = \underline{\underline{\delta}}$	DCS4

Table 1: Axioms of $\text{BPA}_{\text{drt}}^-$ -ID

Axioms DRT1 and DRT2 represent the interaction of time unit delay with alternative composition and sequential composition, respectively. Axiom DRT1, called the time factorization axiom, expresses that passage to the next time slice by itself can not determine a choice. Axiom DRT2 expresses that timing is relative to the performance of the previous action.

In [25], a structured operational semantics of $\text{BPA}_{\text{drt}}^-$ -ID is presented and proofs are given of the soundness and completeness of the axiom system of $\text{BPA}_{\text{drt}}^-$ -ID for the set of closed $\text{BPA}_{\text{drt}}^-$ -ID terms modulo (strong) *bisimulation* equivalence. This notion is precisely defined in [6]. Roughly, bisimilarity of two processes means that if one process is capable of doing a certain step, i.e. performing some action or passing to the next time slice, and next going on as a certain follow-up process then the other process is capable of doing the same step and next going on as a process bisimilar to the follow-up process.

2.2.2 Parallel composition In PA_{drt}^- -ID, we have, in addition to sequential and alternative composition, parallel composition of processes. In PA_{drt}^- -ID, unlike in $\text{ACP}_{\text{drt}}^-$ -ID, parallel composition does not involve communication between processes. The parallel composition operator \parallel of PA_{drt}^- -ID is called free merge to indicate that no communication is involved. The process $x \parallel y$ is the process that proceeds simultaneously with the processes x and y . In order to get a finite axiomatization, we also have the auxiliary operator \ll (left merge) in PA_{drt}^- -ID. The processes $x \ll y$ and $x \parallel y$ are the same except that $x \ll y$ must start with a step of x .

Signature of PA_{drt}^- -ID The signature of PA_{drt}^- -ID is the signature of $\text{BPA}_{\text{drt}}^-$ -ID extended with the *free merge* operator $\parallel: \mathbb{P} \times \mathbb{P} \rightarrow \mathbb{P}$ and the *left merge* operator $\ll: \mathbb{P} \times \mathbb{P} \rightarrow \mathbb{P}$.

We will use the following abbreviation. Let $(p_i)_{i \in \mathcal{I}}$ be an indexed set of terms of PA_{drt}^- -ID where $\mathcal{I} = \{i_1, \dots, i_n\}$. Then, we write $\parallel_{i \in \mathcal{I}} p_i$ for $p_{i_1} \parallel \dots \parallel p_{i_n}$.

Axioms of PA_{drt}^- -ID The axiom system of PA_{drt}^- -ID consists of the axioms of $\text{BPA}_{\text{drt}}^-$ -ID and the equations DRTM1-DRTM6 given in Table 2.

Axioms DRTM5 and DRTM6 represent the interaction between time unit delay and left merge. These axioms express that passage to the next time slice of parallel processes must synchronize.

In [25], a structured operational semantics of PA_{drt}^- -ID is presented and proofs are given of the soundness and completeness of the axiom system of PA_{drt}^- -ID for the set of closed PA_{drt}^- -ID terms modulo (strong) bisimulation equivalence.

$x \parallel y = x \parallel y + y \parallel x$	DRTM1
$\underline{a} \parallel x = \underline{a} \cdot x$	DRTM2
$\underline{a} \cdot x \parallel y = \underline{a} \cdot (x \parallel y)$	DRTM3
$(x + y) \parallel z = x \parallel z + y \parallel z$	DRTM4
$\sigma_{\text{rel}}(x) \parallel \nu_{\text{rel}}(y) = \underline{\delta}$	DRTM5
$\sigma_{\text{rel}}(x) \parallel (\nu_{\text{rel}}(y) + \sigma_{\text{rel}}(z)) = \sigma_{\text{rel}}(x \parallel z)$	DRTM6

Table 2: Additional axioms for PA_{drt}^- -ID

2.2.3 Encapsulation We extend the signature of PA_{drt}^- -ID with the encapsulation operator $\partial_H : \mathbf{P} \rightarrow \mathbf{P}$. This operator turns all undelayable actions \underline{a} , where a in $H \subseteq A$, into undelayable deadlock. The encapsulation operator is defined by the equations DRTD1-DRTD5 given in Table 3.

$\partial_H(\underline{a}) = \underline{a}$ if $a \notin H$	DRTD1
$\partial_H(\underline{a}) = \underline{\delta}$ if $a \in H$	DRTD2
$\partial_H(x \cdot y) = \partial_H(x) \cdot \partial_H(y)$	DRTD3
$\partial_H(x + y) = \partial_H(x) + \partial_H(y)$	DRTD4
$\partial_H(\sigma_{\text{rel}}(x)) = \sigma_{\text{rel}}(\partial_H(x))$	DRTD5

Table 3: Axioms for encapsulation

An operational semantics of encapsulation is presented in [3].

2.3 Propositional signals and conditions

In [5], process algebra with propositional signals and conditions is introduced for the time free case. In this subsection, we adapt it for discrete relative time. The result is referred to by $\text{PA}_{\text{drt}}^{\text{psc}}$. In later sections, we will call propositional signals “propositions” in order to avoid ambiguity with signals in φ -SDL.

In process algebra with propositional signals and conditions, propositions are used both as signals that are emitted by processes and as conditions that are imposed on processes to proceed. Condition testing is looked upon as signal inspection. The intuition is that the signal emitted by a process, as well as each of its logical consequences, holds at the start of the process. The signal emitted by a process is called its root signal.

Like in the time free case we have \perp (non-existence) as additional constant and $\overset{\curvearrowright}{\rightarrow}$ (root signal emission) and $:\rightarrow$ (guarded command) as additional operators. Like the constant $\underline{\delta}$, the constant \perp stands for a process that is incapable of doing any step and incapable of terminating successfully. In addition, going on as \perp after performing an action is impossible. The process $\phi \overset{\curvearrowright}{\rightarrow} x$ is the process x where the proposition ϕ holds at its start. Broadly speaking, the process $\phi : \rightarrow x$ is the process that may proceed as the process $\nu_{\text{rel}}(x)$ if the proposition ϕ holds at its start, but may also proceed as the process $\sigma_{\text{rel}}(\phi : \rightarrow y)$ in case $x = \nu_{\text{rel}}(x) + \sigma_{\text{rel}}(y)$. In other words, with the guarded command operator $:\rightarrow$, it is possible to wait till a proposition holds. This agrees with the original intention of the operator to make actions conditional.

We also have a non-waiting version of the guarded command operator, namely the operator $\dot{:\rightarrow}$ (strict guarded command). The process $\phi \dot{:\rightarrow} x$ is the process that proceeds as the process x if the proposition ϕ holds at its start, and otherwise yields a deadlock in the current time slice. Both guarded

commands agree with the one in the time free case for processes that are not capable of passing to the next time slice.

Lifting propositional signals and conditions to the discrete time case requires the axioms for left merge to be adapted. The root signal of $x \parallel y$ is now the conjunction of the root signals of the processes x and y . Thus, different from the time free case, the process $x \parallel y$ is neither capable of performing an action nor capable of passing to the next time slice if the root signal of y is equal to f . This difference is not relevant to the free merge operator because the root signal of $x \parallel y$ is the conjunction of the root signals of x and y anyhow.

It is assumed that a fixed but arbitrary set B_{at} of *atomic propositions* has been given. From now on we have, in addition to the sort P of processes, the sort B of propositions over B_{at} ; with constants t, f (true, false) and operators $\neg, \vee, \wedge, \rightarrow, \leftrightarrow$ (negation, disjunction, conjunction, implication, bi-implication). In case B_{at} is empty, B represents the boolean algebra over the set $\mathbb{B} = \{t, f\}$.

We denote propositions by ϕ, ψ, \dots . In derivations we may always use logical equivalences of (classical) propositional logic. So we are actually using equivalence classes of formulas, with respect to logical equivalence, instead of the formulas themselves.

A *valuation* v of atomic formulas is a function $v : B_{\text{at}} \rightarrow \mathbb{B}$. Any valuation v can be extended to B in the usual homomorphic way, i.e.:

$$\begin{aligned} v(\kappa) &= \kappa \text{ for the constants } \kappa \in \{t, f\}, \\ v(\neg\phi) &= \neg v(\phi), \\ v(\phi \circ \psi) &= v(\phi) \circ v(\psi) \text{ for the binary operators } \circ \in \{\vee, \wedge, \rightarrow, \leftrightarrow\}. \end{aligned}$$

We will use the same name for a valuation v and its extension to B . If a proposition ϕ is satisfied by a valuation v ($v(\phi) = t$), we write $v \models \phi$ to indicate this.

Signature of $\text{PA}_{\text{drt}}^{\text{psc}}$ The signature of $\text{PA}_{\text{drt}}^{\text{psc}}$ is the signature of the PA_{drt}^- -ID extended with the encapsulation operator, the *non-existence* constant $\perp : P$, the *strict guarded command* operator $\text{!} : B \times P \rightarrow P$, the *root signal emission* operator $\overset{\sim}{\rightarrow} : B \times P \rightarrow P$ and the *weak guarded command* operator $\text{!} : B \times P \rightarrow P$.

Axioms of $\text{PA}_{\text{drt}}^{\text{psc}}$ The axiom system of $\text{PA}_{\text{drt}}^{\text{psc}}$ consists of the axioms of $\text{BPA}_{\text{drt}}^-$ -ID and the equations DRTM1, DRTM4 from Table 2, the equations DRTM2', DRTM3', DRTM5', DRTM6' from Table 4, the equations NE1-NE3 from Table 5, the equations SGC1-SGC6, MSGC from Table 6, the equations SRSE1-SRSE7, MSRSE from Table 7, the equations DCS5, DCS6, DRTM7 from Table 8, the equations DGC1-DGC8 from Table 9, the equations DRTD1-DRTD5 from Table 3, and the equations PD1-PD3 from Table 10.

$\underline{a} \parallel x = \underline{a} \cdot x + \partial_A(\nu_{\text{rel}}(x))$	DRTM2'
$\underline{a} \cdot x \parallel y = \underline{a} \cdot (x \parallel y) + \partial_A(\nu_{\text{rel}}(y))$	DRTM3'
$\sigma_{\text{rel}}(x) \parallel \nu_{\text{rel}}(y) = \partial_A(\nu_{\text{rel}}(y))$	DRTM5'
$\sigma_{\text{rel}}(x) \parallel (\nu_{\text{rel}}(y) + \sigma_{\text{rel}}(z)) = \sigma_{\text{rel}}(x \parallel z) + \partial_A(\nu_{\text{rel}}(y))$	DRTM6'

Table 4: Adapted axioms for left merge

The axioms A3, DRT3 and DRT4A of $\text{BPA}_{\text{drt}}^-$ -ID (Table 1) are derivable from the axioms SGC1, SGC2, SGC4 and SGC6 (Table 6).

The axioms DRTM2', DRTM3', DRTM5' and DRTM6' (Table 4) are the axioms DRTM2, DRTM3, DRTM5 and DRTM6 of PA_{drt}^- -ID (Table 2) where a summand is added to the right hand side of the axioms, viz. $\partial_A(\nu_{\text{rel}}(x))$ in case of DRTM2 and $\partial_A(\nu_{\text{rel}}(y))$ in case of the other axioms. Thus is

$x + \perp = \perp$	NE1
$\perp \cdot x = \perp$	NE2
$\underline{a} \cdot \perp = \underline{\delta}$	NE3

Table 5: Axioms for non-existence

$\mathbf{t} \dot{\rightarrow} x = x$	SGC1
$\mathbf{f} \dot{\rightarrow} x = \underline{\delta}$	SGC2
$\phi \dot{\rightarrow} (x + y) = (\phi \dot{\rightarrow} x) + (\phi \dot{\rightarrow} y)$	SGC3
$(\phi \vee \psi) \dot{\rightarrow} x = (\phi \dot{\rightarrow} x) + (\psi \dot{\rightarrow} x)$	SGC4
$\phi \dot{\rightarrow} (\psi \dot{\rightarrow} x) = (\phi \wedge \psi) \dot{\rightarrow} x$	SGC5
$\phi \dot{\rightarrow} (x \cdot y) = (\phi \dot{\rightarrow} x) \cdot y$	SGC6
$(\phi \dot{\rightarrow} x) \parallel y = \phi \dot{\rightarrow} (x \parallel y)$	MSGC

Table 6: Axioms for strict guarded command

$(\phi \overset{\frown}{\rightarrow} x) \cdot y = \phi \overset{\frown}{\rightarrow} (x \cdot y)$	SRSE1
$(\phi \overset{\frown}{\rightarrow} x) + y = \phi \overset{\frown}{\rightarrow} (x + y)$	SRSE2
$\phi \overset{\frown}{\rightarrow} (\psi \overset{\frown}{\rightarrow} x) = (\phi \wedge \psi) \overset{\frown}{\rightarrow} x$	SRSE3
$\mathbf{t} \overset{\frown}{\rightarrow} x = x$	SRSE4
$\mathbf{f} \overset{\frown}{\rightarrow} x = \perp$	SRSE5
$\phi \dot{\rightarrow} (\psi \overset{\frown}{\rightarrow} x) = (\phi \rightarrow \psi) \overset{\frown}{\rightarrow} (\phi \dot{\rightarrow} x)$	SRSE6
$\phi \overset{\frown}{\rightarrow} (\phi \dot{\rightarrow} x) = \phi \overset{\frown}{\rightarrow} x$	SRSE7
$(\phi \overset{\frown}{\rightarrow} x) \parallel y = \phi \overset{\frown}{\rightarrow} (x \parallel y)$	MSRSE

Table 7: Axioms for root signal emission

$\nu_{\text{rel}}(\phi \dot{\rightarrow} x) = \phi \dot{\rightarrow} \nu_{\text{rel}}(x)$	DCS5
$\nu_{\text{rel}}(\phi \overset{\frown}{\rightarrow} x) = \phi \overset{\frown}{\rightarrow} \nu_{\text{rel}}(x)$	DCS6
$\sigma_{\text{rel}}(x) \parallel ((\phi \dot{\rightarrow} y) + z) = (\phi \dot{\rightarrow} (\sigma_{\text{rel}}(x) \parallel (y + z))) + (\neg\phi \dot{\rightarrow} (\sigma_{\text{rel}}(x) \parallel z))$	DRTM7

Table 8: Additional axioms for time unit delay and now operator

expressed that the root signal of the left merge of two processes is always the conjunction of the root signals of both processes.

The axioms NE1-NE3, SGC1-SGC6, SRSE1-SRSE7, MSGC and MSRSE are straightforward reformulations of corresponding axioms for the time free case, i.e. axioms of PA_{ps} , given in [5]. The constants a and the constant δ have been replaced by the constants \underline{a} and the constant $\underline{\delta}$, respectively; and the operator $\dot{\rightarrow}$ has been replaced by the operator $\overset{\frown}{\rightarrow}$. The axioms NE1 and NE2 (Table 5) are

$\phi : \rightarrow \underline{\underline{\delta}} = \underline{\underline{\delta}}$	DGC1
$\phi : \rightarrow \underline{\underline{a}} = \phi \dot{\rightarrow} \underline{\underline{a}}$	DGC2
$\phi : \rightarrow (\underline{\underline{a}} \cdot x) = (\phi \dot{\rightarrow} \underline{\underline{a}}) \cdot x$	DGC3
$\phi : \rightarrow (\psi \dot{\rightarrow} x) = \psi \dot{\rightarrow} (\phi : \rightarrow x)$	DGC4
$\phi : \rightarrow (x + y) = (\phi \dot{\rightarrow} x) + (\phi \dot{\rightarrow} y)$	DGC5
$\phi : \rightarrow \sigma_{\text{rel}}(x) = \sigma_{\text{rel}}(\phi \dot{\rightarrow} x)$	DGC6
$\phi : \rightarrow \nu_{\text{rel}}(x) = \nu_{\text{rel}}(\phi \dot{\rightarrow} x)$	DGC7
$\phi : \rightarrow (\psi \overset{\curvearrowright}{\rightarrow} x) = \psi \overset{\curvearrowright}{\rightarrow} (\phi \dot{\rightarrow} x)$	DGC8

Table 9: Axioms for weak guarded command

$\partial_H(\phi \dot{\rightarrow} x) = \phi \dot{\rightarrow} \partial_H(x)$	PD1
$\partial_H(\phi \dot{\rightarrow} x) = \phi \dot{\rightarrow} \partial_H(x)$	PD2
$\partial_H(\phi \overset{\curvearrowright}{\rightarrow} x) = \phi \overset{\curvearrowright}{\rightarrow} \partial_H(x)$	PD3

Table 10: Additional axioms for encapsulation

derivable from the axiom A1 of $\text{BPA}_{\text{drt}}^-$ -ID and the axioms SRSE1, SRSE2 and SRSE5 (Table 7). Axiom NE3 expresses that going on as \perp after performing an action is impossible. Axiom SRSE5 expresses that a process where falsity holds at its start is non-existent. The crucial axioms are SRSE6 and SRSE7 which represent the interaction between the root signal emission operator and the strict guarded command. Axiom SRSE6 expresses that if a proposition holds at the start of a process and that process is guarded by another proposition then at the start of the whole the former proposition holds or the latter proposition does not hold. Axiom SRSE7 expresses that it is superfluous to guard a process by a proposition if the proposition holds at the start of the whole.

The additional axioms DCS5, DCS6 and DRTM7 (Table 8) are needed because propositional signals and conditions are lifted to the discrete time case. The strict guarded command is non-waiting, i.e. we do not have $\phi \dot{\rightarrow} \sigma_{\text{rel}}(x) = \sigma_{\text{rel}}(\phi \dot{\rightarrow} x)$. Root signal emission is non-persistent, i.e. we do not have $\phi \overset{\curvearrowright}{\rightarrow} \sigma_{\text{rel}}(x) = \sigma_{\text{rel}}(\phi \overset{\curvearrowright}{\rightarrow} x)$. Axiom DRTM7 is necessary for the elimination of parallel composition. This axiom expresses that if processes are capable of passing to the next time slice conditionally then their parallel composition can do so if all conditions concerned hold. Note that from DRTM7 we can derive $\sigma_{\text{rel}}(x) \parallel (\phi \dot{\rightarrow} y) = \phi \dot{\rightarrow} (\sigma_{\text{rel}}(x) \parallel y)$.

The axioms DGC1-DGC8 define the weak guarded command with which, unlike with the strict guarded command, it is possible to wait till a proposition holds. From these axioms we can derive

$$\begin{aligned} x = \nu_{\text{rel}}(x) &\Rightarrow \phi \dot{\rightarrow} x = (\phi \dot{\rightarrow} x) + \partial_A(\nu_{\text{rel}}(x)) \\ x = \nu_{\text{rel}}(x) + \sigma_{\text{rel}}(y) &\Rightarrow \phi \dot{\rightarrow} x = (\phi \dot{\rightarrow} \nu_{\text{rel}}(x)) + \sigma_{\text{rel}}(\phi \dot{\rightarrow} y) + \partial_A(\nu_{\text{rel}}(x)) \end{aligned}$$

which gives a full picture of the differences between the two guarded commands.

Semantics of $\text{PA}_{\text{drt}}^{\text{psc}}$ We shall give a structured operational semantics for $\text{PA}_{\text{drt}}^{\text{psc}}$ using rules in the style of Plotkin to define the following unary and binary relations on the closed terms of $\text{PA}_{\text{drt}}^{\text{psc}}$:

$$\begin{aligned} \text{a unary relation} &- \frac{v, a}{\sqrt{\quad}} && \text{for each valuation } v \text{ and } a \in A, \\ \text{a binary relation} &- \frac{v, a}{_} && \text{for each valuation } v \text{ and } a \in A, \\ \text{a binary relation} &- \frac{v, \sigma}{_} && \text{for each valuation } v, \\ \text{a unary relation} &- v \in [\mathfrak{s}_\rho(_)] && \text{for each valuation } v. \end{aligned}$$

These relations can be explained as follows:

- $t \xrightarrow{v,a} \surd$: under valuation v , t is capable of first performing a in the current time slice and then terminating successfully;
- $t \xrightarrow{v,a} t'$: under valuation v , t is capable of first performing a in the current time slice and then proceeding as t' ;
- $t \xrightarrow{v,\sigma} t'$: under valuation v , t is capable of first passing to the next time slice and then proceeding as t' ;
- $v \in [s_\rho(t)]$: v makes the root signal of t true.

The rules have the form $\frac{p_1, \dots, p_m}{c_1, \dots, c_n} s$, where s is optional. They are to be read as “if p_1 and \dots and p_m then c_1 and \dots and c_n , provided s ”. As usual, p_1, \dots, p_m and c_1, \dots, c_n are called the premises and the conclusions, respectively. The conclusions are positive formulas of the form $t \xrightarrow{v,a} \surd$, $t \xrightarrow{v,a} t'$, $t \xrightarrow{v,\sigma} t'$ or $v \in [s_\rho(t)]$, where t and t' are open terms of $\text{PA}_{\text{drt}}^{\text{psc}}$. The premises are positive formulas of the above forms or negative formulas of the form $t \not\xrightarrow{v,\sigma}$. A negative formula $t \not\xrightarrow{v,\sigma}$ means that for all closed terms t' of $\text{PA}_{\text{drt}}^{\text{psc}}$ not $t \xrightarrow{v,\sigma} t'$. The rules are actually rule schemas. The optional s is a side-condition restricting the valuations over which v ranges, the actions over which a ranges, the propositions over which ϕ ranges, and the sets of actions over which H ranges. If $m = 0$ and there is no side-condition, the horizontal bar is left out.

The signature of $\text{PA}_{\text{drt}}^{\text{psc}}$ together with the rules that will be given constitute a term deduction system in *panth* format as defined in [6]. It is known from [6] that if a term deduction system in *panth* format is *stratifiable*, (strong) bisimulation equivalence is a congruence for the operators in the signature concerned.

Let T be a term deduction system and $\text{PF}(T)$ be the set of positive formulas occurring in the rules of T . Then a mapping $S : \text{PF}(T) \rightarrow \alpha$ for an ordinal α is called a *stratification* for T if for all rules $\frac{P}{C}$ of T , formulas c in C , and closed substitutions σ the following conditions hold:

- for all positive formulas p in P , $S(\sigma(p)) \leq S(\sigma(c))$;
- for all negative formulas $t \neg R$ in P , $S(\sigma(tRt')) < S(\sigma(c))$ for all closed terms t' ;
- for all negative formulas $\neg Pt$ in P , $S(\sigma(Pt)) < S(\sigma(c))$.

Recall that the rules that will be given are actually rule schemas. Within the framework of term deduction systems, the instances of the rule schemas that satisfy the stated side-conditions should be taken as the rules under consideration. For the rest, we continue to use the word rule in the broader sense.

A structured operational semantics of $\text{PA}_{\text{drt}}^{\text{psc}}$ is described by the rules given in Tables 11, 12, 13 and 14. Note that we write $t \not\xrightarrow{v,\sigma}$ instead of $t \neg \xrightarrow{v,\sigma}$.

All rules are in *panth* format. In order to prove the fact that strong bisimulation is a congruence, we only have to find a stratification. We define a stratification S as follows:

$$S(t \xrightarrow{v,\sigma} t') = n_+(t) \text{ and } S(t \xrightarrow{v,a} \surd) = S(t \xrightarrow{v,a} t') = S(v \in [s_\rho(t)]) = 0,$$

where $n_+(t)$ stands for the number of occurrences of $+$ in t . So $S(F)$ is the number of occurrences of $+$ in the terms t occurring as the left-hand side of the formulas in F that have the form $t \xrightarrow{v,\sigma} t'$. It is straightforward to prove that the mapping S is a stratification. We have to check all rules. This is trivial except for the only rule with a negative formula in its premises, viz. the last rule of Table 11. The number of occurrences of $+$ in the conclusion of that rule is strictly greater than the number of occurrences of $+$ in the negative formula in the premises.

Note that the two rules for alternative composition concerning passage to the next time slice (Table 11) have complementary conditions. Together they enforce that the choice between two processes

$\underline{a} \xrightarrow{v,a} \surd$	$\sigma_{\text{rel}}(x) \xrightarrow{v,\sigma} x$	
$\frac{x \xrightarrow{v,a} \surd}{\phi \dot{\rightarrow} x \xrightarrow{v,a} \surd} v \models \phi$	$\frac{x \xrightarrow{v,a} x'}{\phi \dot{\rightarrow} x \xrightarrow{v,a} x'} v \models \phi$	$\frac{x \xrightarrow{v,\sigma} x'}{\phi \dot{\rightarrow} x \xrightarrow{v,\sigma} x'} v \models \phi$
$\frac{x \xrightarrow{v,a} \surd}{\phi \dot{\rightarrow} x \xrightarrow{v,a} \surd} v \models \phi$	$\frac{x \xrightarrow{v,a} x'}{\phi \dot{\rightarrow} x \xrightarrow{v,a} x'} v \models \phi$	$\frac{x \xrightarrow{v,\sigma} x'}{\phi \dot{\rightarrow} x \xrightarrow{v,\sigma} \phi \dot{\rightarrow} x'}$
$\frac{x \xrightarrow{v,a} \surd}{\phi \overset{\curvearrowright}{\rightarrow} x \xrightarrow{v,a} \surd} v \models \phi$	$\frac{x \xrightarrow{v,a} x'}{\phi \overset{\curvearrowright}{\rightarrow} x \xrightarrow{v,a} x'} v \models \phi$	$\frac{x \xrightarrow{v,\sigma} x'}{\phi \overset{\curvearrowright}{\rightarrow} x \xrightarrow{v,\sigma} x'} v \models \phi$
$\frac{x \xrightarrow{v,a} \surd, w \in [s_\rho(y)]}{x \cdot y \xrightarrow{v,a} y}$	$\frac{x \xrightarrow{v,a} x'}{x \cdot y \xrightarrow{v,a} x' \cdot y}$	$\frac{x \xrightarrow{v,\sigma} x'}{x \cdot y \xrightarrow{v,\sigma} x' \cdot y}$
$\frac{x \xrightarrow{v,a} \surd, v \in [s_\rho(y)]}{x + y \xrightarrow{v,a} \surd, y + x \xrightarrow{v,a} \surd}$	$\frac{x \xrightarrow{v,a} x', v \in [s_\rho(y)]}{x + y \xrightarrow{v,a} x', y + x \xrightarrow{v,a} x'}$	
$\frac{x \xrightarrow{v,\sigma} x', y \xrightarrow{v,\sigma} y'}{x + y \xrightarrow{v,\sigma} x' + y'}$	$\frac{x \xrightarrow{v,\sigma} x', y \xrightarrow{v,\sigma} y', v \in [s_\rho(y)]}{x + y \xrightarrow{v,\sigma} x', y + x \xrightarrow{v,\sigma} x'}$	

Table 11: Rules for basic operators of $\text{PA}_{\text{drt}}^{\text{psc}}$

$\frac{x \xrightarrow{v,a} \surd, v \in [s_\rho(y)]}{x \parallel y \xrightarrow{v,a} y, y \parallel x \xrightarrow{v,a} y, x \parallel y \xrightarrow{v,a} y}$
$\frac{x \xrightarrow{v,a} x', v \in [s_\rho(y)], w \in [s_\rho(x')], w \in [s_\rho(y)]}{x \parallel y \xrightarrow{v,a} x' \parallel y, y \parallel x \xrightarrow{v,a} y \parallel x', x \parallel y \xrightarrow{v,a} x' \parallel y}$
$\frac{x \xrightarrow{v,\sigma} x', y \xrightarrow{v,\sigma} y'}{x \parallel y \xrightarrow{v,\sigma} x' \parallel y', x \parallel y \xrightarrow{v,\sigma} x' \parallel y'}$

Table 12: Rules for parallel composition

$v \in [s_\rho(\underline{a})]$	$v \in [s_\rho(\underline{\delta})]$	$v \in [s_\rho(\sigma_{\text{rel}}(x))]$
$\frac{v \in [s_\rho(x)]}{v \in [s_\rho(\phi \dot{\rightarrow} x)]}$	$\frac{}{v \in [s_\rho(\phi \dot{\rightarrow} x)]} v \not\models \phi$	
$\frac{v \in [s_\rho(x)]}{v \in [s_\rho(\phi \overset{\curvearrowright}{\rightarrow} x)]} v \models \phi$	$\frac{v \in [s_\rho(x)]}{v \in [s_\rho(\phi \overset{\curvearrowright}{\rightarrow} x)]} v \models \phi$	
$\frac{v \in [s_\rho(x)]}{v \in [s_\rho(x \cdot y)]}$	$\frac{v \in [s_\rho(x)], v \in [s_\rho(y)]}{v \in [s_\rho(x + y)], v \in [s_\rho(x \parallel y)], v \in [s_\rho(x \parallel y)]}$	

Table 13: Rules for $v \in [s_\rho(-)]$

$\frac{x \xrightarrow{v,a} \surd}{\nu_{\text{rel}}(x) \xrightarrow{v,a} \surd}$	$\frac{x \xrightarrow{v,a} x'}{\nu_{\text{rel}}(x) \xrightarrow{v,a} x'}$	$\frac{v \in [s_\rho(x)]}{v \in [s_\rho(\nu_{\text{rel}}(x))]}$
$\frac{x \xrightarrow{v,a} \surd}{\partial_H(x) \xrightarrow{v,a} \surd} a \notin H$	$\frac{x \xrightarrow{v,a} x'}{\partial_H(x) \xrightarrow{v,a} \partial_H(x')} a \notin H$	
$\frac{x \xrightarrow{v,\sigma} x'}{\partial_H(x) \xrightarrow{v,\sigma} \partial_H(x')}$	$\frac{v \in [s_\rho(x)]}{v \in [s_\rho(\partial_H(x))]}$	

Table 14: Rules for now operator and encapsulation

that both can pass to the next time slice is postponed till after the passage to the next time slice. This corresponds to the property reflected by the axiom DRT1 of $\text{BPA}_{\text{drt}}^-$ -ID (Table 1).

In order to rule out processes that are capable of performing an action and then going on as \perp , there are premises in the first rule for sequential composition (Table 11) and the second rule for parallel composition (Table 12) concerning the existence of valuations that makes the root signal of certain processes true. This corresponds to the property reflected by the axiom NE3 of $\text{PA}_{\text{drt}}^{\text{psc}}$ (Table 5).

The rule for time unit delay (Table 11), shows that $\sigma_{\text{rel}}(t)$ is capable of passing to the next time slice under all valuations instead of only the valuations under which t is capable of doing things. This excludes persistency of root signal emission and waiting of the strict guarded command because all rules for these operators restrict the valuations under which the resulting processes are capable of doing things.

2.4 Recursion

In this subsection, we add recursion to $\text{PA}_{\text{drt}}^{\text{psc}}$. Recursive specification, solution, guardedness, etc. are defined in a similar way as for BPA in [6].

Let V be a set of variables (of sort \mathbf{P}). A *recursive specification* $E = E(V)$ in $\text{PA}_{\text{drt}}^{\text{psc}}$ is a set of equations $E = \{X = s_X \mid X \in V\}$ where each s_X is a $\text{PA}_{\text{drt}}^{\text{psc}}$ term that only contains variables from V . We shall use X, X', Y, Y', \dots for variables bound in a recursive specification. A *solution* of a recursive specification $E(V)$ is a set of processes $\{\langle X|E \rangle \mid X \in V\}$ in some model of $\text{PA}_{\text{drt}}^{\text{psc}}$ such that the equations of $E(V)$ hold if, for all $X \in V$, X stands for $\langle X|E \rangle$. Mostly, we are interested in one particular variable $X \in V$.

We can now introduce the equational theory of $\text{PA}_{\text{drt}}^{\text{psc}}\text{rec}$.

Signature of $\text{PA}_{\text{drt}}^{\text{psc}}\text{rec}$ The signature of $\text{PA}_{\text{drt}}^{\text{psc}}\text{rec}$ consists of the signature of $\text{PA}_{\text{drt}}^{\text{psc}}$ extended with a constant $\langle X|E \rangle : \mathbf{P}$ for each $X \in V$ and each recursive specification $E(V)$.

Let t be an open term in $\text{PA}_{\text{drt}}^{\text{psc}}$ and $E = E(V)$ be a recursive specification. Then we write $\langle t|E \rangle$ for t with, for all $X \in V$, all occurrences of X in t replaced by $\langle X|E \rangle$.

Axioms of $\text{PA}_{\text{drt}}^{\text{psc}}\text{rec}$ The axiom system of $\text{PA}_{\text{drt}}^{\text{psc}}\text{rec}$ consists of the axioms of $\text{PA}_{\text{drt}}^{\text{psc}}$ and an equation $\langle X|E \rangle = \langle s_X|E \rangle$ for each $X \in V$ and each recursive specification $E(V)$.

Let t be a term of $\text{PA}_{\text{drt}}^{\text{psc}}$ containing a variable X . We call an occurrence of X in t *guarded* if t has a subterm of the form $\underline{a} \cdot t'$ or $\sigma_{\text{rel}}(t')$ with t' a $\text{PA}_{\text{drt}}^{\text{psc}}$ term containing this occurrence of X . We call a recursive specification *guarded* if all occurrences of all its variables in the right-hand sides of all its equations are guarded or it can be rewritten to such a recursive specification using the axioms of $\text{PA}_{\text{drt}}^{\text{psc}}$ and its equations. An interesting form of guarded recursive specification is linear recursive specification. We call a recursive specification $E(V)$ *linear* if each equation in E has the form

$$X = \sum_{i < n} \phi_i \dot{\rightarrow} \underline{a}_i \cdot X_i + \sum_{i < m} \psi_i \dot{\rightarrow} \underline{b}_i + \sum_{i < k} \chi_i \dot{\rightarrow} \sigma_{\text{rel}}(X'_i) + \xi \dot{\rightarrow} \underline{\delta}$$

for certain actions a_i and b_i , propositions ϕ_i , ψ_i , χ_i and ξ , and variables $X, X_i, X'_i \in V$. Note that, without loss of generality we can assume that for all i and j such that $i \neq j$: $\underline{a_i} \cdot X_i \not\equiv \underline{a_j} \cdot X_j$, $\underline{b_i} \not\equiv \underline{b_j}$, $X'_i \not\equiv X'_j$ and $\chi_i \not\equiv \chi_j$. We can also assume that ϕ_i , ψ_i and χ_i are not f.

Principles of $\text{PA}_{\text{drt}}^{\text{psc}}\text{rec}$ The (*restricted*) *recursive definition principle* ($\text{RDP}^{(-)}$) is the assumption that every (guarded) recursive specification has a solution. The *recursive specification principle* (RSP) is the assumption that every guarded recursive specification has at most one solution.

Note that the axioms $\langle X|E \rangle = \langle s_X|E \rangle$ for a fixed E express that the constants $\langle X|E \rangle$ make up a solution of E , i.e. RDP holds for any model of $\text{PA}_{\text{drt}}^{\text{psc}}\text{rec}$. The conditional equations $E \Rightarrow X = \langle X|E \rangle$ express that this solution is the only one. So RSP can be described by means of conditional equations – as already mentioned in [26].

Semantics of $\text{PA}_{\text{drt}}^{\text{psc}}\text{rec}$ A structured operational semantics of $\text{PA}_{\text{drt}}^{\text{psc}}\text{rec}$ is described by the rules for $\text{PA}_{\text{drt}}^{\text{psc}}$ and the rules given in Table 15.

$\frac{\langle s_X E \rangle \xrightarrow{v,a} \surd}{\langle X E \rangle \xrightarrow{v,a} \surd}$	$\frac{\langle s_X E \rangle \xrightarrow{v,a} y}{\langle X E \rangle \xrightarrow{v,a} y}$
$\frac{\langle s_X E \rangle \xrightarrow{v,\sigma} y}{\langle X E \rangle \xrightarrow{v,\sigma} y}$	$\frac{v \in [\mathfrak{s}_\rho(\langle s_X E \rangle)]}{v \in [\mathfrak{s}_\rho(\langle X E \rangle)]}$

Table 15: Rules for recursion

The rules added for recursion are also in panth format. We define a stratification S as follows:

$$S(t \xrightarrow{v,\sigma} t') = \omega \cdot n_{\text{sol}}(t) + n_+(t) \text{ and } S(t \xrightarrow{v,a} \surd) = S(t \xrightarrow{v,a} t') = S(v \in [\mathfrak{s}_\rho(t)]) = 0,$$

where $n_{\text{sol}}(t)$ stands for the number of unguarded occurrences of constants $\langle X|E \rangle$ in t and $n_+(t)$ stands for the number of occurrences of $+$ in t . The addition of the summand $\omega \cdot n_{\text{sol}}(t)$ for formulas $t \xrightarrow{v,\sigma} t'$ solves the problem that the conclusion of the rule for recursion concerning passage to the next time slice does not contain occurrences of $+$.

Let $E = \{X = s_X \mid X \in V\}$ be a recursive specification. Then roughly, the rules for recursion come down to looking upon $\langle X|E \rangle$ as the process s_X with, for all $X' \in V$, all occurrences of X' in s_X replaced by $\langle X'|E \rangle$.

2.5 Counting process creation

In this subsection, we introduce the counting process creation operator E_Φ^n that is used for the semantics of φ -SDL. This operator subsumes the original process creation operator introduced in [9]. The latter process creation operator was used in [12] for the semantics of φ SDL. But the approach used there does not guarantee that a unique process identification is assigned to each created process.

It is assumed that a fixed but arbitrary set D of *data* has been given together with a function $\Phi : \mathbb{N} \times D \rightarrow \mathbf{P}$, and that there exist actions $cr(d)$ and $\overline{cr}(n, d)$ for all $d \in D$ and $n \in \mathbb{N}$. The process creation operator E_Φ^n allows, given the function Φ , the use of actions $cr(d)$ to create processes $\Phi(n, d)$.

The counting process creation operator $E_\Phi^n : \mathbf{P} \rightarrow \mathbf{P}$ is defined by the equations PCR1-PCR8 given in Table 16. The crucial axiom is PCR4 . It expresses that counting process creation applied to a process, with the counter set to n , leaves the action $\overline{cr}(n, d)$ as a trace and starts a process $\Phi(n, d)$ in parallel with the remaining process when it comes across an undelayable process creation action $\underline{cr}(d)$. Besides, it increases the counter by one. The counting process creation operator E_Φ^n is an extension of the process creation operator E_ϕ from [9]. We can write $E_\Phi^n = E_\phi$ if $\Phi(n, d) = \phi(d)$ for all $n \in \mathbb{N}$ and $d \in D$.

$E_{\Phi}^n(\underline{a}) = \underline{a}$ if $a \neq cr(d)$ for $d \in D$	PRCR1
$E_{\Phi}^n(\underline{cr}(d)) = \underline{cr}(n, d) \cdot E_{\Phi}^{n+1}(\Phi(n, d))$	PRCR2
$E_{\Phi}^n(\underline{a} \cdot x) = \underline{a} \cdot E_{\Phi}^n(x)$ if $a \neq cr(d)$ for $d \in D$	PRCR3
$E_{\Phi}^n(\underline{cr}(d) \cdot x) = \underline{cr}(n, d) \cdot E_{\Phi}^{n+1}(\Phi(n, d) \parallel x)$	PRCR4
$E_{\Phi}^n(x + y) = E_{\Phi}^n(x) + E_{\Phi}^n(y)$	PRCR5
$E_{\Phi}^n(\sigma_{rel}(x)) = \sigma_{rel}(E_{\Phi}^n(x))$	PRCR6
$E_{\Phi}^n(\phi \dot{\rightarrow} x) = \phi \dot{\rightarrow} E_{\Phi}^n(x)$	PRCR7
$E_{\Phi}^n(\phi \overset{\sim}{\rightarrow} x) = \phi \overset{\sim}{\rightarrow} E_{\Phi}^n(x)$	PRCR8

Table 16: Axioms for counting process creation

A structured operational semantics for the counting process creation operator is described by the rules given in Table 17. The stratification introduced for $\text{PA}_{\text{drt}}^{\text{psc}}\text{rec}$ still works if we add the rules for

$\frac{x \xrightarrow{v,a} \surd}{E_{\Phi}^n(x) \xrightarrow{v,a} \surd} \quad a \neq cr(d)$	$\frac{x \xrightarrow{v,cr(d)} \surd, w \in [s_{\rho}(\Phi(n, d))]}{E_{\Phi}^n(x) \xrightarrow{v, \underline{cr}(n, d)} E_{\Phi}^{n+1}(\Phi(n, d))}$
$\frac{x \xrightarrow{v,a} x'}{E_{\Phi}^n(x) \xrightarrow{v,a} E_{\Phi}^n(x')} \quad a \neq cr(d)$	$\frac{x \xrightarrow{v,cr(d)} x', w \in [s_{\rho}(\Phi(n, d))], w \in [s_{\rho}(x')]}{E_{\Phi}^n(x) \xrightarrow{v, \underline{cr}(n, d)} E_{\Phi}^{n+1}(\Phi(n, d) \parallel x')}$
$\frac{x \xrightarrow{v,\sigma} x'}{E_{\Phi}^n(x) \xrightarrow{v,\sigma} E_{\Phi}^n(x')}$	$\frac{v \in [s_{\rho}(x)]}{v \in [s_{\rho}(E_{\Phi}^n(x))]}$

Table 17: Rules for counting process creation

the counting process creation operator.

2.6 State operator

In this subsection, we introduce a state operator for $\text{PA}_{\text{drt}}^{\text{psc}}$. It generalizes and extends the state operator for ACP_{ps} proposed in [5]: the truth value of propositional signals and conditions may depend upon the state and passage to the next time slice may have an effect on the state. The possibility of non-deterministic behaviour as the result of applying the operator is included as well, like for the extended state operator Λ (see e.g. [6]).

It is assumed that a fixed but arbitrary set S of states has been given, together with functions:

$$\begin{aligned}
act &: A \times S \rightarrow \mathcal{P}_{fn}(A) \\
eff &: A \times S \times A \rightarrow S \\
eff_{\sigma} &: S \rightarrow \mathcal{P}_{fn}(S) \setminus \{\emptyset\} \\
sig &: S \rightarrow \mathbb{B} \\
val &: B_{\text{at}} \times S \rightarrow \mathbb{B}
\end{aligned}$$

The state operator λ_s ($s \in S$) allows, given these functions, processes to interact with a state. The process $\lambda_s(x)$ is the process x executed in a state s . The function act gives, for each action a and state s , the set of actions that may be performed if a is executed in state s . The function eff gives, for each action a , state s and action a' , the state that results when a is executed in state s and a' is the

action that is actually performed as the result of the execution. The function eff_σ gives, for each state s , the set of states that may result when time passes to the next time slice in state s . The function sig gives, for each state s , the propositional signal that holds at the start of any process executed in state s . The function val gives, for each state s , the valuation $val(_, s)$ of the atomic propositions in state s . The valuation $val(_, s)$ can be extended to all propositions in the usual homomorphic way as any other valuation. We will use the notation $val(_, s)$ to refer to the extension as well.

The state operator $\lambda_s : \mathbf{P} \rightarrow \mathbf{P}$ is defined by the equations SO1-SO6 given in Table 18. The axioms

$\lambda_s(\underline{a}) = sig(s) \overset{\curvearrowright}{\sim} \sum_{a' \in act(a,s)} \underline{a}'$	SO1
$\lambda_s(\underline{a} \cdot x) = sig(s) \overset{\curvearrowright}{\sim} \sum_{a' \in act(a,s)} (\underline{a}' \cdot \lambda_{eff(a,s,a')}(x))$	SO2
$\lambda_s(x + y) = \lambda_s(x) + \lambda_s(y)$	SO3
$\lambda_s(\sigma_{rel}(x)) = sig(s) \overset{\curvearrowright}{\sim} \sigma_{rel}(\sum_{s' \in eff_\sigma(s)} \lambda_{s'}(x))$	SO4
$\lambda_s(\phi \dot{\rightarrow} x) = sig(s) \overset{\curvearrowright}{\sim} (val(\phi, s) \dot{\rightarrow} \lambda_s(x))$	SO5
$\lambda_s(\phi \overset{\curvearrowright}{\sim} x) = val(\phi, s) \overset{\curvearrowright}{\sim} \lambda_s(x)$	SO6

Table 18: Axioms for state operator

SO1-SO3 are straightforward reformulations – in the same way as for $\text{PA}_{\text{drt}}^{\text{psc}}$ – of corresponding axioms given in [5] for the time free case. The additional axiom SO4 expresses how passage to the next time slice has influence on the execution of a process in a state. The axioms SO5 and SO6 are also reformulations of corresponding axioms given in [5]. In these axioms the proposition ϕ has been replaced by $val(\phi, s)$. Thus the case is covered where the truth value of propositional signals and conditions may depend upon the state.

A structured operational semantics for the state operator is described by the rules given in Table 19. Note that the rules added for the state operator have a common side-condition given at the bottom of the table. The stratification introduced for $\text{PA}_{\text{drt}}^{\text{psc}} \text{rec}$ still works if we add the rules for the state

$\frac{x \xrightarrow{v,a} \surd}{\lambda_s(x) \xrightarrow{v',a'} \surd} \quad a' \in act(a,s)$
$\frac{x \xrightarrow{v,a} x', w \in [s_\rho(x')]}{\lambda_s(x) \xrightarrow{v',a'} \lambda_{eff(a,s,a')}(x')} \quad a' \in act(a,s) \wedge w \models sig(eff(a,s,a'))$
$\frac{x \xrightarrow{v,\sigma} x'}{\lambda_s(x) \xrightarrow{v',\sigma} \lambda_{s'}(x')} \quad s' \in eff_\sigma(s) \quad \frac{v \in [s_\rho(x)]}{v' \in [s_\rho(\lambda_s(x))]}$

for all v, v' and s such that $v' \models sig(s) \wedge \forall \phi (v \models \phi \leftrightarrow v' \models val(\phi, s))$

Table 19: Rules for state operator

operator.

3. OVERVIEW OF φ -SDL

3.1 Introduction

In this section, we give an overview of φ -SDL, i.e. φ SDL without delaying channels. φ SDL is a small subset of SDL, introduced in [12], which covers all behavioural aspects of SDL, including communication via delaying channels, timing and process creation. Leaving out delaying channels simplifies

the presentation. Besides work on the process algebra semantics of φ SDL made clear that φ SDL specifications can always be transformed to semantically equivalent ones in φ -SDL. At the end of Section 4 is shown how to model a delaying channel by means of a φ -SDL process.

In Sections 3.2, 3.3 and 3.4, the syntax of φ -SDL is described by means of production rules in the form of an extended BNF grammar (the extensions are explained in Appendix A). The meaning of the language constructs of the various forms distinguished by these rules is explained informally. Some peculiar details of the semantics, inherited from full SDL, are left out to improve the comprehensibility of the overview. These details will, however, be taken into account in Section 4, where the process algebra semantics of φ -SDL is presented. In Section 3.5, some remarks are made about the context-sensitive conditions for syntactic correctness of φ -SDL specifications. The syntactic differences with full SDL are summarized in this section as well. In Section 3.6 some examples of φ -SDL specifications are given.

In line with full SDL, we can define a graphical representation for φ -SDL specifications. We pay no attention to this practically important point because it is not relevant to the subject of this chapter.

3.2 System definition

First of all, the φ -SDL view of a system is explained in broad outline.

Basically, a system consists of *processes* which communicate with each other and the environment by sending and receiving *signals* via *signal routes*. A process proceeds in parallel with the other processes in the system and communicates with these processes in an asynchronous manner. This means that a process sending a signal does not wait until the receiving process consumes it, but it proceeds immediately. A process may also use local *variables* for storage of values. A variable is associated with a value that may change by assigning a new value to it. A variable can only be assigned new values by the process to which it is local, but it may be viewed by other processes. Processes can be distinguished by unique addresses, called *pid values* (process identification values), which they get with their creation.

A signal can be sent from the environment to a process, from a process to the environment or from a process to a process. A signal may carry values to be passed from the sender to the receiver; on consumption of the signal, these values are assigned to local variables of the receiver. A signal route is a unidirectional communication path for sending signals from the environment to a process, from one process to another process or from a process to the environment. If a signal is sent to a process via a signal route it will be instantaneously delivered to that process.

Syntax:

```

<system definition> ::=
  system <system nm> ; { <definition> }+ endsystem ;

<definition> ::=
  dcl <variable nm> <sort nm> ;
  | signal <signal nm> [ ( <sort nm> { , <sort nm> } * ) ] ;
  | signalroute <signalroute nm>
    from { <process nm> | env } to { <process nm> | env }
    with <signal nm> { , <signal nm> } * ;
  | process <process nm> ( <natural ground expr> ) ;
    [ fpar <variable nm> { , <variable nm> } * ; ]
    start ; <transition> { <state def> } *
  endprocess ;

```

A system definition consists of definitions of the types of processes present in the system, the local variables used by the processes for storage of values, the types of signals used by the processes for communication and the signal routes via which the signals are conveyed.

A variable definition **dcl** $v T$; defines a variable v that may be assigned values of sort T .

A signal definition **signal** $s(T_1, \dots, T_n)$; defines a type of signals s of which the instances carry values of the sorts T_1, \dots, T_n . If (T_1, \dots, T_n) is absent, the signals of type s do not carry any value.

A signal route definition **signalroute** r **from** X_1 **to** X_2 **with** s_1, \dots, s_n ; defines a signal route r that delivers without delay signals sent by processes of type X_1 to processes of type X_2 , for signals of the types s_1, \dots, s_n . The process types X_1 and X_2 are called the sender type of r and the receiver type of r , respectively. A signal route from the environment can be defined by replacing **from** X_1 by **from env**. A signal route to the environment can be defined analogously.

A process definition **process** $X(k)$; **fpar** v_1, \dots, v_m ; **start**; $tr d_1 \dots d_n$ **endprocess**; defines a type of processes X of which k instances will be created during the start-up of the system. On creation of a process of type X after the start-up, the creating process passes values to it which are assigned to the local variables v_1, \dots, v_m . If **fpar** v_1, \dots, v_m ; is absent, no values are passed on creation. The process body **start**; $tr d_1 \dots d_n$ describes the behaviour of the processes of type X in terms of states and transitions (see further Section 3.3). Each process will start by making the transition tr , called its start transition, to enter one of its states. The state definitions d_1, \dots, d_n define all the states in which the process may come while it proceeds.

3.3 Process behaviour

First of all, the φ -SDL view of a process is briefly explained.

To begin with, a process is either in a *state* or making a *transition* to another state. Besides, when a signal arrives at a process, it is put into the unique *input queue* associated with the process until it is consumed by the process. The states of a process are the points in its behaviour where a signal may be consumed. However, a state may have signals that have to be saved, i.e. withhold from being consumed in that state. The signal consumed in a state of a process is the first one in its input queue that has not to be saved for that state. If there is no signal to consume, the process waits until there is a signal to consume. So if a process is in a state, it is either waiting to consume a signal or consuming a signal.

A transition from a state of a process is initiated by the consumption of a signal, unless it is a spontaneous transition. The start transition is not initiated by the consumption of a signal either. A transition is made by performing certain actions: signals may be sent, variables may be assigned new values, new processes may be created and *timers* may be set and reset. A transition may at some stage also take one of a number of branches, but it will eventually come to an end and bring the process to a next state or to its termination.

A timer can be set which sends at its expiration time a signal to the process setting it. A timer is identified with the type and carried values of the signal it sends on expiration. Thus an active timer can be set to a new time or reset; if this is done between the sending of the signal noticing expiration and its consumption, the signal is removed from the input queue concerned. A timer is de-activated when it is reset or the signal it sends on expiration is consumed.

Syntax:

```

<state def> ::=
  state <state nm> ;
  [ save <signal nm> {, <signal nm>}* ; ] {<transition alt>}*

<transition alt> ::=
  {<input guard> | input none;} <transition>

<input guard> ::=
  input <signal nm> [ ( (<variable nm> {, <variable nm>}*) ) ] ;

<transition> ::=
  {<action>}* {nextstate <state nm> | stop | <decision>} ;

```

```

<action> ::=
  output <signal nm> [ ( <expr> { , <expr> } * ) ]
    [ to <pid expr> ] via <signalroute nm> { , <signalroute nm> } * ;
  | set ( <time expr> , <signal nm> [ ( <expr> { , <expr> } * ) ] ) ;
  | reset ( <signal nm> [ ( <expr> { , <expr> } * ) ] ) ;
  | task <variable nm> := <expr> ;
  | create <process nm> [ ( <expr> { , <expr> } * ) ] ;

<decision> ::=
  decision { <expr> | any } ;
  ( [ <ground expr> ] ) : <transition>
  { ( [ <ground expr> ] ) : <transition> } +
  enddecision

```

A state definition **state** st ; **save** s_1, \dots, s_m ; $alt_1 \dots alt_n$ defines a state st . The signals of the types s_1, \dots, s_m are saved for the state. The input guard of each of the transition alternatives alt_1, \dots, alt_n gives a type of signals that may be consumed in the state; the corresponding transition is the one that is initiated on consumption of a signal of that type. The alternatives with **input none**; instead of an input guard are the spontaneous transitions that may be made from the state. No signals are saved for the state if **save** s_1, \dots, s_m ; is absent.

An input guard **input** $s(v_1, \dots, v_n)$; may consume a signal of type s and, on consumption, it assigns the carried values to the variables v_1, \dots, v_n . If the signals of type s carry no value, (v_1, \dots, v_n) is left out.

A transition $a_1 \dots a_n$ **nextstate** st ; performs the actions a_1, \dots, a_n in sequential order and ends with entering the state st . Replacing **nextstate** st by the keyword **stop** yields a transition ending with process termination. Replacing it by the decision dec leads instead to transfer of control to one of two or more transition branches.

An output action **output** $s(e_1, \dots, e_n)$ **to** e **via** r_1, \dots, r_m ; sends a signal of type s carrying the current values of the expressions e_1, \dots, e_n to the process with the current (pid) value of the expression e as its address, via one of the signal routes r_1, \dots, r_m . If the signals of type s carry no value, (e_1, \dots, e_n) is left out. If **to** e is absent, the signal is sent via one of the signal routes r_1, \dots, r_m to an arbitrary process of its receiver type. The output action is called an output action with explicit addressing if **to** e is present. Otherwise, it is called an output action with implicit addressing.

A set action **set** $(e, s(e_1, \dots, e_n))$; sets a timer that expires, unless it is set again or reset, at the current (time) value of the expression e with sending a signal of type s that carries the current values of the expressions e_1, \dots, e_n .

A reset action **reset** $(s(e_1, \dots, e_n))$; de-activates the timer identified with the signal type s and the current values of the expressions e_1, \dots, e_n .

An assignment task action **task** $v := e$; assigns the current value of the expression e to the local variable v .

A create action **create** $X(e_1, \dots, e_n)$; creates a process of type X and passes the current values of the expressions e_1, \dots, e_n to the newly created process. If no values are passed on creation of processes of type X , (e_1, \dots, e_n) is left out.

A decision **decision** $e; (e_1); tr_1 \dots (e_n); tr_n$ **enddecision** transfers control to the transition branch tr_i ($1 \leq i \leq n$) for which the value of the expression e_i equals the current value of the expression e . Non-existence of such a branch results in an error. A non-deterministic choice can be obtained by replacing the expression e by the keyword **any** and removing all the expressions e_i .

3.4 Values

The value of expressions in φ -SDL may vary according to the last values assigned to variables, including local variables of other processes. It may also depend on timers being active, the system

time, etc.

Syntax:

```

<expr> ::=
  <operator nm> [ ( <expr> { , <expr> } * ) ]
  | if <boolean expr> then <expr> else <expr> fi
  | <variable nm>
  | view ( <variable nm> , <pid expr> )
  | active ( <signal nm> [ ( <expr> { , <expr> } * ) ] )
  | now | self | parent | offspring | sender

```

An operator application $op(e_1, \dots, e_n)$ evaluates to the value yielded by applying the operation op to the current values of the expressions e_1, \dots, e_n .

A conditional expression **if** e_1 **then** e_2 **else** e_3 **fi** evaluates to the current value of the expression e_2 if the current (Boolean) value of the expression e_1 is true, and the current value of the expression e_3 otherwise.

A variable access v evaluates to the current value of the local variable v of the process evaluating the expression.

A view expression **view** (v, e) evaluates to the current value of the local variable v of the process with the current (pid) value of the expression e as its address.

An active expression **active** ($s(e_1, \dots, e_n)$) evaluates to the Boolean value true if the timer identified with the signal type s and the current values of the expressions e_1, \dots, e_n is currently active, and false otherwise.

The expression **now** evaluates to the current system time.

The expressions **self**, **parent**, **offspring** and **sender** evaluate to the pid values of the process evaluating the expression, the process by which it was created, the last process created by it, and the sender of the last signal consumed by it. Natural numbers are used as pid values. The pid value 0 is a special pid value that never refers to any existing process – in full SDL this pid value is denoted by **null** – and the pid value 1 is reserved for the environment. The expressions **parent**, **offspring** and **sender** evaluate to 0 in case there exists no parent, offspring and sender, respectively.

3.5 Miscellaneous issues

3.5.1 Context-sensitive syntactic rules We remain loose about the context-sensitive conditions for syntactic correctness of φ -SDL specifications. For the most part, they are as usual: only defined names may be used, use of names must agree with their definitions, types of expressions must be in accordance with the type information in the definitions, etc. There is one condition that needs particular attention: signals of the same type may not be used for both signal sending and timer setting/resetting. All φ -SDL specifications that are obtained by semantics preserving transformations of syntactically correct specifications in full SDL will be syntactically correct φ -SDL specifications as well.

3.5.2 Syntactic differences with SDL Syntactically, φ -SDL is not exactly a subset of SDL. The syntactic differences are as follows:

- variable definitions occur at the system level instead of inside process definitions;
- signal route definitions and process definitions occur at the system level instead of inside block definitions;
- formal parameters in process definitions are variable names instead of pairs of variable names and sort names;
- signal names are used as timer names.

These differences are all due to the simplifications mentioned in Section 1.

3.6 Examples

We give three small examples to illustrate how systems are specified in φ -SDL. The examples concern simple repeaters and routers.

3.6.1 Repeater The first example concerns a simple repeater, i.e. a system that simply passes on what it receives. The system, called `Repeater`, consists of only one process, viz. `rep`, which communicates signals `s` with the environment via the signal routes `fromenv` and `toenv`. The process has only one state.

```

system Repeater
  signal s;

  signalroute fromenv from env to rep with s;
  signalroute toenv from rep to env with s;

  process rep(1);
    start;
    nextstate pass;
    state pass;
    input s;
    output s via toenv;
    nextstate pass;
  endprocess;
endsystem;

```

3.6.2 Address driven router This example concerns address driven routing of data. The system, called `AddrRouter`, consists of three processes, one instance of `rtr` and two instances of `rep`. The latter two processes are created by the former process. The process `rtr` consumes signals `s(a)`, delivered via signal route `fromenv`, and passes them to one of the instances of `rep` (via signal route `rs`) depending on the value `a`. The instances of `rep` then pass the signals received from `rtr` to the environment via the signal route `toenv`.

```

system AddrRouter
  signal s(Bool);

  signalroute fromenv from env to rtr with s;
  signalroute rs from rtr to rep with s;
  signalroute toenv from rep to env with s;

  dcl a Bool; dcl rep1 Nat; dcl rep2 Nat;

  process rtr(1);
    start;
    create rep; task rep1 := offspring;
    create rep; task rep2 := offspring;
    nextstate route;
  state route;
  input s(a);
  decision a;
    (False):
      output s(a) to rep1 via rs;
      nextstate route;
    (True):
      output s(a) to rep2 via rs;
      nextstate route;
  endprocess;
endsystem;

```

```

        enddecision;
    endprocess;

    process rep(0);
    start;
        nextstate pass;
    state pass;
        input s(a);
        output s(a) via toenv;
        nextstate pass;
    endprocess;
endsystem;

```

3.6.3 Load driven router The last example concerns load driven routing of data. The system, called *LoadRouter*, consists of three processes, one instance of *rtr* and two instances of *trep*. The latter two processes are created by the former process. The process *rtr* consumes signals *s*, delivered via signal route *fromenv*, and passes them to one of the instances of *trep* (via signal route *rs*) depending on their load. The instances of *trep* then pass the signals received from *rtr* to the environment via the signal routes *toenv*. Either delivers the data consumed after a fixed time *delay*. One repeater delivers twice as fast as the other one. Only two load factors are considered for each of the two repeaters: one indicating its idleness and one indicating the opposite.

```

system LoadRouter
    signal s; signal t;

    signalroute fromenv from env to rtr with s;
    signalroute rs from rtr to trep with s;
    signalroute toenv from trep to env with s;

    dcl idle Bool; dcl delay Nat;
    dcl rep1 Nat; dcl rep2 Nat;

    process rtr(1);
    start;
        create trep(10); task rep1 := offspring;
        create trep(20); task rep2 := offspring;
        nextstate route;
    state route;
        input s;
        decision view(idle,rep1) <-> view(idle,rep2);
            (True):
                output s via rs; nextstate route;
            (False):
                decision view(idle,rep1)
                    (True):
                        output s to rep1 via rs;
                        nextstate route;
                    (False):
                        output s to rep2 via rs;
                        nextstate route;
                enddecision
            enddecision
        endprocess;

    process trep(0) fpar delay;

```



```

start;
  task idle := True;
  nextstate get;
state get;
  input s;
  task idle := False;
  set(now + delay, t);
  nextstate put;
state put;
  save s;
  input t;
  output s via toenv;
  task idle := True;
  nextstate get;
endprocess;
endsystem;

```

Most features of φ -SDL are used in this example. Of the main features, only spontaneous transitions, i.e. transition alternatives with **input none**; instead of an input guard, are missing. In Section 4.6.3, use is made of this feature to define a process modeling a delaying channel.

4. SEMANTICS OF φ -SDL

4.1 Introduction

In this section, we propose a process algebra semantics for φ -SDL. This semantics is presented in two steps.

In Section 4.3, a semantics for φ -SDL process definitions is given. This semantics abstracts from dynamic aspects of process behaviour such as process creation and process execution in a state. It describes the meaning of φ -SDL process definitions by means of finite guarded recursive specifications in $\text{BPA}_{\text{drt}}^{\text{psc}}$. The counting process creation operator and the state operator are not needed for this semantics. Preceding, in Section 4.2, the actions and atomic conditions used are introduced. These actions and conditions are parametrized by expressions with values that depend on the state in which the action or condition concerned is executed.

In Section 4.6, a semantics of φ -SDL system definitions is given in terms of the semantics for φ -SDL process definitions using the counting process creation operator and the state operator. Preceding, in Section 4.5, all the details of the instance of the state operator needed for this semantics are given. Included are the definitions of the state space and the functions that describe how the actions and conditions used for the semantics of φ -SDL process definitions interact with the state when this instance of the state operator is applied. The interaction with the environment is another aspect covered by the semantics of φ -SDL system definitions. For this purpose an environment process is introduced as well.

The semantics of φ -SDL is described by means of a set of equations recursively defining interpretation functions for all syntactic categories. The special notation used is explained in Appendix A. We will be lazy about specifying the range of each interpretation function, since this is usually clear from the context. Many of the interpretations are functions from natural numbers to process expressions or equations. These process expressions and equations will simply be written in their display form. If an optional clause represents a sequence, its absence is always taken to stand for an empty sequence. Otherwise, it is treated as a separate case. The semantics is defined using contextual information κ extracted from the φ -SDL specification concerned. This is further described in Appendix B.

The special notation used for parametrized actions, conditions and propositions is explained in Appendix A, and so is the uncustomary notation as regards sets, functions and sequences. The words action and process are used in connection with both ACP-style process algebras and versions of SDL, but with slightly different meanings. In case it is not clear from the context, these words will be preceded by ACP if they should be taken in the ACP sense, and by SDL otherwise.

Data types We mentioned before that φ -SDL does not deal with the specification of abstract data types. We assume a fixed algebraic specification covering all data types used and an initial algebra semantics, denoted by \mathcal{A} , for it. The data types **Boolean** and **Natural**, with the obvious interpretation, must be included. We will write $Sort_{\mathcal{A}}$ and $Op_{\mathcal{A}}$ for the set of all sort names and the set of all operation names, respectively, in the signature of \mathcal{A} . We additionally assume a constant name, i.e. a nullary operation name, in $Op_{\mathcal{A}}$, referred to as \underline{n} , for each $n \in \mathbb{N}$. We will write U for $\bigcup_{T \in Sort_{\mathcal{A}}} T^{\mathcal{A}}$, where $T^{\mathcal{A}}$ is the interpretation of the sort name T in \mathcal{A} . We have that $\mathbb{B}, \mathbb{N} \subset U$ because of our earlier requirement that **Boolean**, **Natural** $\in Sort_{\mathcal{A}}$. We assume that $\text{nil} \notin U$. In the sequel, we will use for each $op \in Op_{\mathcal{A}}$ an extension to $U \cup \{\text{nil}\}$, also denoted by op , such that $op(t_1, \dots, t_n) = \text{nil}$ if not all of the t_i s are of the appropriate sort. Thus, we can change over from the many-sorted case to the one-sorted case for the description of the meaning of the language constructs of φ -SDL. We can do so without loss of generality, because it can (and should) be statically checked that only terms of appropriate sorts occur.

4.2 Actions and conditions

In the semantics of φ -SDL process definitions, which will be presented in Section 4.3, ACP actions and conditions are used. Here, we introduce the actions and atomic conditions concerned.

The actions and atomic conditions, used to describe the meaning of φ -SDL process definitions, are parametrized by various domains. These domains depend upon the specific variable names, signal names and process names introduced in the system definition concerned. These sets of names largely make up the context ascribed to the system definition by means of the function $\{[-]\}$ defined in Appendix B. For convenience, we define these sets for arbitrary contexts κ (the notation concerning contexts introduced in Appendix B is used):

$$\begin{aligned} V_{\kappa} &= \text{vars}(\kappa) \cup \{\text{parent, offspring, sender}\} \\ S_{\kappa} &= \text{sigs}(\kappa) \\ P_{\kappa} &= \text{procs}(\kappa) \end{aligned}$$

Most arguments of the parametrized actions and conditions introduced here are expressions that originate in φ -SDL expressions or objects that are somehow composed of such expressions and variable, signal and process names. The reason of this is that the value of the original φ -SDL expressions may vary according to the last values assigned to the variables referred to, the status of the timers referred to, etc. In other words, these expressions stand for values that are not known until the action or condition concerned is executed in the appropriate state.

The syntax of the expressions concerned, called value expressions, is as follows:

$$\begin{aligned} \langle \text{value expr} \rangle ::= & \\ & \langle \text{operator nm} \rangle [(\langle \text{value expr} \rangle \{ , \langle \text{value expr} \rangle \}^*)] \\ & | \text{cond} (\langle \text{boolean value expr} \rangle , \langle \text{value expr} \rangle , \langle \text{value expr} \rangle) \\ & | \text{value} (\langle \text{variable nm} \rangle , \langle \text{pid value expr} \rangle) \\ & | \text{active} (\langle \text{signal nm} \rangle [(\langle \text{value expr} \rangle \{ , \langle \text{value expr} \rangle \}^*)] , \langle \text{pid value expr} \rangle) \\ & | \text{now} \end{aligned}$$

where the terminal productions of $\langle \text{operator nm} \rangle$, $\langle \text{variable nm} \rangle$ and $\langle \text{signal nm} \rangle$ are assumed to yield the sets $Op_{\mathcal{A}}$, V_{κ} and S_{κ} , respectively. $ValE_{\kappa}$ denotes the set of all syntactically correct value expressions. The forms of value expressions distinguished above correspond to operator applications, conditional expressions, view expressions, active expressions, and the expression **now**, respectively, in φ -SDL.

We define now the set $SigP_{\kappa}$ of signal patterns, the set $SigE_{\kappa}$ of signals expressions, the set $SaveSet_{\kappa}$ of save sets and the set $PrCrD_{\kappa}$ of process creation data. Further on, we will look at a signal as an object that consists of the name of the signal and the sequence of values that it carries. Signal patterns and signal expressions are like signals, but variables and value expressions, respectively, are

used instead of values. A save set is just a finite set of signal names. A process creation datum consists of the name of the process to be created, its formal parameters, expressions denoting its actual parameters and the pid value of its creator.

$$\begin{aligned}
\text{Sig}P_\kappa &= \{(s, \langle v_1, \dots, v_n \rangle) \mid (s, \langle T_1, \dots, T_n \rangle) \in \text{sigds}(\kappa), v_1, \dots, v_n \in V_\kappa\} \\
\text{Sig}E_\kappa &= \{(s, \langle e_1, \dots, e_n \rangle) \mid (s, \langle T_1, \dots, T_n \rangle) \in \text{sigds}(\kappa), e_1, \dots, e_n \in \text{Val}E_\kappa\} \\
\text{SaveSet}_\kappa &= \mathcal{P}_{\text{fin}}(S_\kappa) \\
\text{PrCr}D_\kappa &= \{(X, \langle v_1, \dots, v_n \rangle, \langle e_1, \dots, e_n \rangle, e) \mid \\
&\quad X \in P_\kappa, v_1, \dots, v_n \in V_\kappa, e_1, \dots, e_n, e \in \text{Val}E_\kappa\}
\end{aligned}$$

We write $\text{pnm}(d)$, where $d = (X, vs, es, e) \in \text{PrCr}D_\kappa$, for X . Each process creation datum contains the formal parameters for the process type concerned. The alternative would be to make the association between process types and their formal parameters itself a parameter of the state operator, which is very unattractive.

The following actions are used:

$$\begin{aligned}
\text{input} &: \text{Sig}P_\kappa \times \text{SaveSet}_\kappa \times \text{Val}E_\kappa \\
\text{outpute} &: \text{Sig}E_\kappa \times \text{Val}E_\kappa \times \text{Val}E_\kappa \\
\text{outputi} &: \text{Sig}E_\kappa \times \text{Val}E_\kappa \times (P_\kappa \cup \{\mathbf{env}\}) \\
\text{set} &: \text{Val}E_\kappa \times \text{Sig}E_\kappa \times \text{Val}E_\kappa \\
\text{reset} &: \text{Sig}E_\kappa \times \text{Val}E_\kappa \\
\text{ass} &: V_\kappa \times \text{Val}E_\kappa \times \text{Val}E_\kappa \\
\text{cr} &: \text{PrCr}D_\kappa \\
\text{stop} &: \text{Val}E_\kappa \\
\text{inisport} &: \text{Val}E_\kappa \\
\# &:
\end{aligned}$$

Cr_κ denotes the set of all cr actions, and Act_κ^- denotes the set of all input , outpute , outputi , set , reset , ass , stop and inisport actions. The $\#$ action is a special action with no observable effect whatsoever. The other actions correspond to the input guards, the SDL actions, the terminator **stop** and the void guard **input none**. The last argument of each action is the pid value of the process from which the action originates, except for the outpute and outputi actions where the pid value concerned is available as the second argument. The outpute and outputi actions correspond to **output** actions with explicit addressing and implicit addressing, respectively, in φ -SDL. The last argument of these actions refers to the receiver. With an outpute action, the receiver is fully determined by the pid value given as the last argument. With an outputi action, the receiver is not fully determined; it is an arbitrary process of the given type.

The conditions used are built from the following atomic conditions:

$$\begin{aligned}
\text{waiting} &: \text{SaveSet}_\kappa \times \text{Val}E_\kappa \\
\text{type} &: \text{Val}E_\kappa \times (P_\kappa \cup \{\mathbf{env}\}) \\
\text{hasinst} &: P_\kappa \cup \{\mathbf{env}\} \\
\text{eq} &: \text{Val}E_\kappa \times \text{Val}E_\kappa
\end{aligned}$$

AtCond_κ denotes the set of all waiting , type , hasinst and eq conditions. A condition $\text{waiting}(ss, e)$ is used to test whether the process with the pid value denoted by e has to wait for a signal if the signals with names in ss are withhold from being consumed. A condition $\text{type}(e, X)$ is used to test whether X is the type of the process with the pid value denoted by e . A condition $\text{hasinst}(X)$ is used to test whether there exists a process of type X . A condition $\text{eq}(e_1, e_2)$ is used to test whether the expressions e_1 and e_2 denote the same value. These conditions are used to give meaning to the state definitions, output actions and decisions of φ -SDL.

4.3 Semantics of process definitions

We describe now the meaning of φ -SDL process definitions and their constituents. The meaning of the process definitions occurring in the examples from Section 3.6 is presented in Section 4.4.

4.3.1 Process definitions The meaning of each process definition occurring in a system definition consists of the process name introduced and a family of processes, one for each possible pid value, which are described by means of finite guarded recursive specifications in $\text{BPA}_{\text{drt}}^{\text{psc}}$. The meaning of a process definition is expressed in terms of the meaning of its start transition and its state definitions.

$$\begin{aligned} & \llbracket \text{process } X(k); \text{fpar } v_1, \dots, v_m; \text{start}; \text{tr } d_1 \dots d_n \text{endprocess}; \rrbracket^\kappa := \\ & (X, \{i \mapsto \langle X \mid \{X = \llbracket \text{tr} \rrbracket_i^{\kappa'}, \llbracket d_1 \rrbracket_i^{\kappa'}, \dots, \llbracket d_n \rrbracket_i^{\kappa'} \} \mid i \in \mathbb{N}\} \rangle) \\ & \text{where } \kappa' = \text{updscopeunit}(\kappa, X) \end{aligned}$$

The recursive specification of the process of type X with pid value i describes how it behaves at its start (the equation $X = \llbracket \text{tr} \rrbracket_i^{\kappa'}$) and how it behaves from each of the n states in which it may come while it proceeds (the equations $\llbracket d_1 \rrbracket_i^{\kappa'}, \dots, \llbracket d_n \rrbracket_i^{\kappa'}$).

In the remainder of this section, we will be loose in the explanation of the meaning of the constituents of process definitions about the fact that there is always a family of meanings, one for each possible pid value.

4.3.2 States and transitions The meaning of a state definition, occurring in the scope of a process definition, is a process described by an equation of the form $Z = s_Z$ where Z is a variable corresponding to the state and s_Z is a $\text{BPA}_{\text{drt}}^{\text{psc}}$ term that only contains variables corresponding to states introduced in the process definition concerned. The equation describes how a process of the type being defined behaves from the state. The meaning of the state definition is expressed in terms of the meaning of its transition alternatives, which are process expressions describing the behaviour from the state being defined for the individual signal types of which instances may be consumed and, in addition, possibly for some spontaneous transitions. The meaning of each transition alternative is in turn expressed in terms of the meaning of its input guard, if the alternative is not a spontaneous transition, and its transition.

$$\begin{aligned} & \llbracket \text{state } st; \text{save } s_1, \dots, s_m; \text{alt } t_1 \dots \text{alt } t_n \rrbracket_i^\kappa := \\ & X_{st} = \underline{\underline{t}} \cdot (\llbracket \text{alt } t_1 \rrbracket_i^{\kappa'} + \dots + \llbracket \text{alt } t_n \rrbracket_i^{\kappa'} + \text{waiting}(\{s_1, \dots, s_m\}, \underline{\underline{t}}) \mapsto \sigma_{\text{rel}}(X_{st})) \\ & \text{where } X = \text{scopeunit}(\kappa), \\ & \quad \kappa' = \text{updsaveset}(\kappa, \{s_1, \dots, s_m\}) \end{aligned}$$

$$\begin{aligned} & \llbracket \text{input } s(v_1, \dots, v_n); \text{tr} \rrbracket_i^\kappa := \underline{\underline{\text{input}}}((s, \langle v_1, \dots, v_n \rangle), ss, \underline{\underline{t}}) \cdot \llbracket \text{tr} \rrbracket_i^\kappa \\ & \text{where } ss = \text{saveset}(\kappa) \end{aligned}$$

$$\llbracket \text{input none}; \text{tr} \rrbracket_i^\kappa := \underline{\underline{\text{inispont}}}(\underline{\underline{t}}) \cdot \llbracket \text{tr} \rrbracket_i^\kappa$$

The equation that corresponds to a state definition describes that a process of type X behaves from the state st as one of the given transition alternatives, and that this behaviour is possibly delayed till the first future time slice in which there is a signal to consume if there are no more signals to consume in the current time slice. Entering a state is supposed to take place by way of some internal action – thus it is precluded that a process is in more than one state. We use process names with state name subscripts, such as X_{st} above, as variables. Notice that, in the absence of spontaneous transitions, a delay becomes inescapable if there are no more signals to consume in the current time slice. The process expression that corresponds to a guarded transition alternative expresses that the transition tr is initiated on consumption of a signal of type s . In case of an unguarded transition alternative, the process expression expresses that the transition tr is initiated spontaneously, i.e. without a preceding signal consumption – with **sender** set to the value of **self** (see Section 4.5).

The meaning of a transition is described by a process expression – a $\text{BPA}_{\text{drt}}^{\text{psc}}$ term to be precise. It is expressed in terms of the meaning of its actions and its transition terminator.

$$\begin{aligned} \llbracket a_1 \dots a_n \text{ nextstate } st; \rrbracket_i^\kappa &:= \llbracket a_1 \rrbracket_i^\kappa \cdot \dots \cdot \llbracket a_n \rrbracket_i^\kappa \cdot X_{st} \\ \text{where } X &= \text{scopeunit}(\kappa) \\ \llbracket a_1 \dots a_n \text{ stop}; \rrbracket_i^\kappa &:= \llbracket a_1 \rrbracket_i^\kappa \cdot \dots \cdot \llbracket a_n \rrbracket_i^\kappa \cdot \underline{\underline{\text{stop}}}(\dot{\iota}) \\ \llbracket a_1 \dots a_n \text{ dec}; \rrbracket_i^\kappa &:= \llbracket a_1 \rrbracket_i^\kappa \cdot \dots \cdot \llbracket a_n \rrbracket_i^\kappa \cdot \llbracket \text{dec} \rrbracket_i^\kappa \end{aligned}$$

The process expression that corresponds to a transition terminated by **nextstate** st expresses that the transition performs the actions a_1, \dots, a_n in sequential order and ends with entering state st – i.e. goes on behaving as defined for state st . In case of termination by **stop**, the process expression expresses that it ends with ceasing to exist; and in case of termination by a decision dec , that it goes on behaving as described by decision dec .

The meaning of a decision is described by a process expression as well. It is expressed in terms of the meaning of its expressions and transitions.

$$\begin{aligned} \llbracket \text{decision } e; (e_1):tr_1 \dots (e_n):tr_n \text{ enddecision} \rrbracket_i^\kappa &:= \\ \text{eq}(\llbracket e \rrbracket_i, \llbracket e_1 \rrbracket_i) \dot{\rightarrow} \llbracket tr_1 \rrbracket_i^\kappa + \dots + \text{eq}(\llbracket e \rrbracket_i, \llbracket e_n \rrbracket_i) \dot{\rightarrow} \llbracket tr_n \rrbracket_i^\kappa \\ \llbracket \text{decision any}; () : tr_1 \dots () : tr_n \text{ enddecision} \rrbracket_i^\kappa &:= \llbracket tr_1 \rrbracket_i^\kappa + \dots + \llbracket tr_n \rrbracket_i^\kappa \end{aligned}$$

The process expression that corresponds to a decision with a question expression e expresses that the decision transfers control to the transition tr_i for which the value of e equals the value of e_i . In case of a decision with **any** instead, the process expression expresses that the decision transfers non-deterministically control to one of the transitions tr_1, \dots, tr_n .

4.3.3 Actions The meaning of an SDL action is described by a process expression, of course. It is expressed in terms of the meaning of the expressions occurring in it. It also depends on the occurring names (variable names, signal names, signal route names and process names – dependent on the kind of action).

$$\begin{aligned} \llbracket \text{output } s(e_1, \dots, e_n) \text{ to } e \text{ via } r_1, \dots, r_m; \rrbracket_i^\kappa &:= \\ \text{type}(\llbracket e \rrbracket_i, X_1) \vee \dots \vee \text{type}(\llbracket e \rrbracket_i, X_m) \dot{\rightarrow} \underline{\underline{\text{outpute}}}((s, \langle \llbracket e_1 \rrbracket_i, \dots, \llbracket e_n \rrbracket_i \rangle), \dot{\iota}, \llbracket e \rrbracket_i) + \\ \neg(\text{type}(\llbracket e \rrbracket_i, X_1) \vee \dots \vee \text{type}(\llbracket e \rrbracket_i, X_m)) \dot{\rightarrow} \underline{\underline{\text{t}}} \\ \text{where for } 1 \leq j \leq m: X_j &= \text{rcv}(\kappa, r_j) \\ \llbracket \text{output } s(e_1, \dots, e_n) \text{ via } r_1, \dots, r_m; \rrbracket_i^\kappa &:= \\ \underline{\underline{\text{outputi}}}((s, \langle \llbracket e_1 \rrbracket_i, \dots, \llbracket e_n \rrbracket_i \rangle), \dot{\iota}, X_1)) + \dots + \underline{\underline{\text{outputi}}}((s, \langle \llbracket e_1 \rrbracket_i, \dots, \llbracket e_n \rrbracket_i \rangle), \dot{\iota}, X_m)) + \\ \neg(\text{hasinst}(X_1) \wedge \dots \wedge \text{hasinst}(X_m)) \dot{\rightarrow} \underline{\underline{\text{t}}} \\ \text{where for } 1 \leq j \leq m: X_j &= \text{rcv}(\kappa, r_j) \\ \llbracket \text{set } (e, s(e_1, \dots, e_n)); \rrbracket_i^\kappa &:= \underline{\underline{\text{set}}}(\llbracket e \rrbracket_i, (s, \langle \llbracket e_1 \rrbracket_i, \dots, \llbracket e_n \rrbracket_i \rangle), \dot{\iota}) \\ \llbracket \text{reset } (s(e_1, \dots, e_n)); \rrbracket_i^\kappa &:= \underline{\underline{\text{reset}}}((s, \langle \llbracket e_1 \rrbracket_i, \dots, \llbracket e_n \rrbracket_i \rangle), \dot{\iota}) \\ \llbracket \text{task } v := e; \rrbracket_i^\kappa &:= \underline{\underline{\text{ass}}}(v, \llbracket e \rrbracket_i, \dot{\iota}) \\ \llbracket \text{create } X(e_1, \dots, e_n); \rrbracket_i^\kappa &:= \underline{\underline{\text{cr}}}((X, \text{fpars}(\kappa, X), \langle \llbracket e_1 \rrbracket_i, \dots, \llbracket e_n \rrbracket_i \rangle), \dot{\iota}) \end{aligned}$$

All cases except the ones for output actions are straightforward. The cases of output actions need further explanation. In the case of an output action with explicit addressing, the process with pid value e must be of the receiver type associated with one of the signal routes r_1, \dots, r_m . Therefore, the condition $\text{type}(\llbracket e \rrbracket_i, X_1) \vee \dots \vee \text{type}(\llbracket e \rrbracket_i, X_m)$ is used. If the process with pid value e is not of the receiver type associated with any of the signal routes r_1, \dots, r_m , or a process with that pid value does not exist, the signal is simply discarded and no error occurs. This is expressed by the summand

$\neg(\text{type}(\llbracket e \rrbracket_i, X_1) \vee \dots \vee \text{type}(\llbracket e \rrbracket_i, X_m)) \dot{\rightarrow} \underline{\#}$. In the case of an output action with implicit addressing, first an arbitrary choice from the signal routes r_1, \dots, r_m is made and thereafter an arbitrary choice from the existing processes of the receiver type for the chosen signal route is made. Therefore, there is a summand for the receiver type of each signal route. If a process of the receiver type for the chosen signal route does not exist, the signal is simply discarded and no error occurs. This is expressed by the summand $\neg(\text{hasinst}(X_1) \wedge \dots \wedge \text{hasinst}(X_m)) \dot{\rightarrow} \underline{\#}$. Note that the signal may already be discarded if there is one signal route for which there exists no process of its receiver type.

4.3.4 Values The meaning of a φ^- SDL expression is given by a translation to a value expression of the same kind. There is a close correspondence between the φ^- SDL expressions and their translations. Essential of the translation is that i is added where the local states of different processes need to be distinguished. Consequently, a variable access v is just treated as a view expression **view** (v , **self**). For convenience, the expressions **parent**, **offspring** and **sender** are also regarded as variable accesses.

$$\begin{aligned}
\llbracket op(e_1, \dots, e_n) \rrbracket_i &:= op(\llbracket e_1 \rrbracket_i, \dots, \llbracket e_n \rrbracket_i) \\
\llbracket \text{if } e_1 \text{ then } e_2 \text{ else } e_3 \text{ fi} \rrbracket_i &:= \text{cond}(\llbracket e_1 \rrbracket_i, \llbracket e_2 \rrbracket_i, \llbracket e_3 \rrbracket_i) \\
\llbracket v \rrbracket_i &:= \text{value}(v, i) \\
\llbracket \text{view}(v, e) \rrbracket_i &:= \text{value}(v, \llbracket e \rrbracket_i) \\
\llbracket \text{active}(s(e_1, \dots, e_n)) \rrbracket_i &:= \text{active}((s, \langle \llbracket e_1 \rrbracket_i, \dots, \llbracket e_n \rrbracket_i \rangle), i) \\
\llbracket \text{now} \rrbracket_i &:= \text{now} \\
\llbracket \text{self} \rrbracket_i &:= i \\
\llbracket \text{parent} \rrbracket_i &:= \text{value}(\text{parent}, i) \\
\llbracket \text{offspring} \rrbracket_i &:= \text{value}(\text{offspring}, i) \\
\llbracket \text{sender} \rrbracket_i &:= \text{value}(\text{sender}, i)
\end{aligned}$$

4.4 Examples

We present the meaning of the process definitions occurring in the examples from Section 3.6. To be more precise, we give for each process definition a constant of the form $\langle X | E \rangle$ that stands for the process of the type concerned with pid value i .

It is clear that there are many similarities with the original process definitions in φ^- SDL. There is an equation for each state, the right hand side of each equation has a summand for each transition alternative of the corresponding state, etc. In addition, there is always a summand in which the time unit delay operator appears; this summand allows a delay to a future time slice to occur if there is no input to be read from the input queue of the process concerned. The main difference between the φ^- SDL process definitions and the description of their meaning in $\text{BPA}_{\text{drt}}^{\text{psc}} \text{rec}$ is that the latter can be subjected to equational reasoning using the axioms of $\text{BPA}_{\text{drt}}^{\text{psc}} \text{rec}$.

4.4.1 Repeater

The **rep** process with pid value i is

$$\begin{aligned}
&\langle \text{rep} | \\
&\quad \{ \text{rep} = \text{rep}_{\text{pass}}, \\
&\quad \quad \text{rep}_{\text{pass}} = \underline{\#} \cdot (\text{input}((s, \langle \rangle), \emptyset, i) \cdot \\
&\quad \quad \quad \text{output}((s, \langle \rangle), i, \text{env}) + \neg \text{hasinst}(\text{env}) \dot{\rightarrow} \underline{\#}) \cdot \text{rep}_{\text{pass}} + \\
&\quad \quad \quad \text{waiting}(\emptyset, i) \dot{\rightarrow} \sigma_{\text{rel}}(\text{rep}_{\text{pass}}) \\
&\quad \} \\
&\rangle
\end{aligned}$$

4.4.2 Address driven router

The \mathbf{rtr} process with pid value i is

$$\begin{aligned}
& \langle \mathbf{rtr} | \\
& \{ \mathbf{rtr} = \underline{cr}((\mathbf{rep}, \langle \rangle, \langle \rangle, i)) \cdot \underline{ass}(\mathbf{rep1}, \mathit{value}(\mathbf{offspring}, i), i) \cdot \\
& \quad \underline{cr}((\mathbf{rep}, \langle \rangle, \langle \rangle, i)) \cdot \underline{ass}(\mathbf{rep2}, \mathit{value}(\mathbf{offspring}, i), i) \cdot \mathbf{rtr}_{\mathbf{route}}, \\
& \quad \mathbf{rtr}_{\mathbf{route}} = \underline{\underline{t}} \cdot \underline{input}((s, \langle a \rangle), \emptyset, i) \cdot \\
& \quad \quad (eq(\mathit{value}(a, i), \mathbf{False}) \rightarrow (type(\mathit{value}(\mathbf{rep1}, i), \mathbf{rep}) \rightarrow \\
& \quad \quad \underline{outpute}((s, \langle \mathit{value}(a, i) \rangle), i, \mathit{value}(\mathbf{rep1}, i)) + \\
& \quad \quad \neg type(\mathit{value}(\mathbf{rep1}, i), \mathbf{rep}) \rightarrow \underline{\underline{t}}) \cdot \mathbf{rtr}_{\mathbf{route}} + \\
& \quad \quad eq(\mathit{value}(a, i), \mathbf{True}) \rightarrow (type(\mathit{value}(\mathbf{rep2}, i), \mathbf{rep}) \rightarrow \\
& \quad \quad \underline{outpute}((s, \langle \mathit{value}(a, i) \rangle), i, \mathit{value}(\mathbf{rep2}, i)) + \\
& \quad \quad \neg type(\mathit{value}(\mathbf{rep2}, i), \mathbf{rep}) \rightarrow \underline{\underline{t}}) \cdot \mathbf{rtr}_{\mathbf{route}}) + \\
& \quad \quad \mathit{waiting}(\emptyset, i) \rightarrow \sigma_{\mathit{rel}}(\mathbf{rtr}_{\mathbf{route}})) \\
& \} \\
& \rangle
\end{aligned}$$

The \mathbf{rep} process with pid value i is

$$\begin{aligned}
& \langle \mathbf{rep} | \\
& \{ \mathbf{rep} = \mathbf{rep}_{\mathbf{pass}}, \\
& \quad \mathbf{rep}_{\mathbf{pass}} = \underline{\underline{t}} \cdot \underline{input}((s, \langle a \rangle), \emptyset, i) \cdot \\
& \quad \quad (\underline{outputi}((s, \langle a \rangle), i, \mathbf{env}) + \neg \mathit{hasinst}(\mathbf{env}) \rightarrow \underline{\underline{t}}) \cdot \mathbf{rep}_{\mathbf{pass}} + \\
& \quad \quad \mathit{waiting}(\emptyset, i) \rightarrow \sigma_{\mathit{rel}}(\mathbf{rep}_{\mathbf{pass}})) \\
& \} \\
& \rangle
\end{aligned}$$

4.4.3 Load driven router

The \mathbf{rtr} process with pid value i is

$$\begin{aligned}
& \langle \mathbf{rtr} | \\
& \{ \mathbf{rtr} = \underline{cr}((\mathbf{trep}, \langle \mathit{delay} \rangle, \langle 10 \rangle, i)) \cdot \underline{ass}(\mathbf{rep1}, \mathit{value}(\mathbf{offspring}, i), i) \cdot \\
& \quad \underline{cr}((\mathbf{trep}, \langle \mathit{delay} \rangle, \langle 20 \rangle, i)) \cdot \underline{ass}(\mathbf{rep2}, \mathit{value}(\mathbf{offspring}, i), i) \cdot \mathbf{rtr}_{\mathbf{route}}, \\
& \quad \mathbf{rtr}_{\mathbf{route}} = \underline{\underline{t}} \cdot \underline{input}((s, \langle \rangle), \emptyset, i) \cdot \\
& \quad \quad (eq(\mathit{value}(\mathbf{idle}, \mathit{value}(\mathbf{rep1}, i)) \leftrightarrow \mathit{value}(\mathbf{idle}, \mathit{value}(\mathbf{rep2}, i)), \mathbf{True}) \rightarrow \\
& \quad \quad \underline{outputi}((s, \langle \rangle), i, \mathbf{trep}) + \neg \mathit{hasinst}(\mathbf{trep}) \rightarrow \underline{\underline{t}}) \cdot \mathbf{rtr}_{\mathbf{route}} + \\
& \quad \quad eq(\mathit{value}(\mathbf{idle}, \mathit{value}(\mathbf{rep1}, i)) \leftrightarrow \mathit{value}(\mathbf{idle}, \mathit{value}(\mathbf{rep2}, i)), \mathbf{False}) \rightarrow \\
& \quad \quad (eq(\mathit{value}(\mathbf{idle}, \mathit{value}(\mathbf{rep1}, i)), \mathbf{True}) \rightarrow \\
& \quad \quad \quad (type(\mathit{value}(\mathbf{rep1}, i), \mathbf{rep}) \rightarrow \underline{outpute}((s, \langle \rangle), i, \mathit{value}(\mathbf{rep1}, i)) + \\
& \quad \quad \quad \neg type(\mathit{value}(\mathbf{rep1}, i), \mathbf{rep}) \rightarrow \underline{\underline{t}}) \cdot \mathbf{rtr}_{\mathbf{route}} + \\
& \quad \quad \quad eq(\mathit{value}(\mathbf{idle}, \mathit{value}(\mathbf{rep1}, i)), \mathbf{False}) \rightarrow \\
& \quad \quad \quad (type(\mathit{value}(\mathbf{rep2}, i), \mathbf{rep}) \rightarrow \underline{outpute}((s, \langle \rangle), i, \mathit{value}(\mathbf{rep2}, i)) + \\
& \quad \quad \quad \neg type(\mathit{value}(\mathbf{rep2}, i), \mathbf{rep}) \rightarrow \underline{\underline{t}}) \cdot \mathbf{rtr}_{\mathbf{route}})) + \\
& \quad \quad \mathit{waiting}(\emptyset, i) \rightarrow \sigma_{\mathit{rel}}(\mathbf{rtr}_{\mathbf{route}})) \\
& \} \\
& \rangle
\end{aligned}$$

The \mathbf{trep} process with pid value i is

$$\begin{aligned}
& \langle \text{trep} | \\
& \quad \{ \text{trep} = \underline{\text{ass}}(\text{idle}, \text{True}) \cdot \text{trep}_{\text{get}}, \\
& \quad \quad \text{trep}_{\text{get}} = \underline{\underline{t}} \cdot (\underline{\text{input}}((\underline{s}, \langle \rangle), \emptyset, \underline{i}) \cdot \\
& \quad \quad \quad \underline{\text{ass}}(\text{idle}, \text{False}) \cdot \underline{\text{set}}(\text{now} + \text{delay}, \underline{t}) \cdot \text{trep}_{\text{put}} + \\
& \quad \quad \quad \underline{\text{waiting}}(\emptyset, \underline{i}) \dot{\rightarrow} \sigma_{\text{rel}}(\text{trep}_{\text{get}})), \\
& \quad \quad \text{trep}_{\text{put}} = \underline{\underline{t}} \cdot (\underline{\text{input}}((\underline{t}, \langle \rangle), \{\underline{s}\}, \underline{i}) \cdot \\
& \quad \quad \quad (\underline{\text{output}}((\underline{s}, \langle \rangle), \underline{i}, \text{env}) + \neg \text{hasinst}(\text{env}) \dot{\rightarrow} \underline{\underline{t}}) \cdot \\
& \quad \quad \quad \underline{\text{ass}}(\text{idle}, \text{True}) \cdot \text{trep}_{\text{get}} + \\
& \quad \quad \quad \underline{\text{waiting}}(\{\underline{s}\}, \underline{i}) \dot{\rightarrow} \sigma_{\text{rel}}(\text{trep}_{\text{put}})) \\
& \quad \} \\
& \rangle
\end{aligned}$$

4.5 Interaction with states

In the semantic of φ -SDL process definitions, ACP actions are used to give meaning to input guards, SDL actions, the terminator **stop** and the void guard **input none**. Thus, the facilities for storage, communication, timing and process creation offered by these language constructs are not fully covered; the ACP actions are meant to interact with a system state. In the semantics of φ -SDL system definitions, which will be presented in Section 4.6, the state operator mentioned in Section 2 is used to describe this. First, we will describe the state space, the actions that may appear as the result of executing the above-mentioned actions in a state, and the result of executing processes, built up from these actions, in a state from the state space.

4.5.1 State space, actions and propositional signals The state space, used to describe the meaning of system definitions, depends upon the specific variable names, signal names and process names introduced in the system definition concerned. That is, the sets V_κ , S_κ and P_κ are used here as well.

First, we define the set Sig_κ of signals and the set ExtSig_κ of extended signals. A signal consist of the name of the signal and the sequence of values that it carries. An extended signal contains, in addition to a signal, the pid values of its sender and receiver.

$$\begin{aligned}
\text{Sig}_\kappa &= \{(s, \langle u_1, \dots, u_n \rangle) \mid (s, \langle T_1, \dots, T_n \rangle) \in \text{sigds}(\kappa), u_1, \dots, u_n \in U\} \\
\text{ExtSig}_\kappa &= \text{Sig}_\kappa \times \mathbb{N}_1 \times \mathbb{N}_1
\end{aligned}$$

We write $\text{snm}(\text{sig})$ and $\text{vals}(\text{sig})$, where $\text{sig} = (s, \text{vs}) \in \text{Sig}_\kappa$, for s and vs , respectively. We write $\text{sig}(\text{xsig})$, $\text{snd}(\text{xsig})$ and $\text{rcv}(\text{xsig})$, where $\text{xsig} = (\text{sig}, i, i') \in \text{ExtSig}_\kappa$, for sig , i and i' , respectively. Note that 0 is excluded as pid value of the sender or receiver of a signal because it is a special pid value that never refers to any existing process.

The local state of a process includes a storage which associates local variables with the values assigned to them, an input queue where delivered signals are kept until they are consumed, and a component keeping track of the expiration times of active timers. We define the set Stg_κ of storages, the set InpQ_κ of input queues and the set Timers_κ of timers as follows:

$$\begin{aligned}
\text{Stg}_\kappa &= \bigcup_{V \subseteq V_\kappa} (V \rightarrow U) \\
\text{InpQ}_\kappa &= \text{ExtSig}_\kappa^* \\
\text{Timers}_\kappa &= \bigcup_{T \in \mathcal{P}_{\text{fin}}(\text{Sig}_\kappa)} (T \rightarrow (\mathbb{N} \cup \{\text{nil}\}))
\end{aligned}$$

We will follow the convention that the domain of a function from Stg_κ does not contain variables with which no value is associated because a value has never been assigned to them. We will also follow the convention that the domain of a function from Timers_κ contains precisely the active timers. While an expired timer is still active, its former expiration time will be replaced by nil. The basic operations on Stg_κ and Timers_κ are general operations on functions: function application, overriding (\oplus) and domain subtraction (\Leftarrow). Overriding and domain subtraction are defined in Appendix A. In so far as the facilities for communication are concerned, the basic operations on InpQ_κ are the functions

$$\begin{aligned}
\mathit{getnext} & : \mathit{InpQ}_\kappa \times \mathit{SaveSet}_\kappa \rightarrow \mathit{ExtSig}_\kappa \cup \{\mathit{nil}\}, \\
\mathit{rmvfirst} & : \mathit{InpQ}_\kappa \times \mathit{Sig}_\kappa \rightarrow \mathit{InpQ}_\kappa, \\
\mathit{merge} & : \mathcal{P}_{\mathit{fin}}(\mathit{ExtSig}_\kappa) \rightarrow \mathcal{P}_{\mathit{fin}}(\mathit{InpQ}_\kappa)
\end{aligned}$$

defined below. The value of $\mathit{getnext}(\sigma, ss)$ is the first (extended) signal in σ with a name different from the ones in ss . The value of $\mathit{rmvfirst}(\sigma, sig)$ is the input queue σ from which the first occurrence of the signal sig has been removed. Both functions are used to describe the consumption of signals by SDL processes. The function $\mathit{getnext}$ is recursively defined by

$$\begin{aligned}
\mathit{getnext}(\langle \rangle, ss) & = \mathit{nil} \\
\mathit{getnext}((sig, i, i') \& \sigma, ss) & = (sig, i, i') \quad \text{if } \mathit{snm}(sig) \notin ss \\
\mathit{getnext}((sig, i, i') \& \sigma, ss) & = \mathit{getnext}(\sigma, ss) \quad \text{if } \mathit{snm}(sig) \in ss
\end{aligned}$$

and the function $\mathit{rmvfirst}$ is recursively defined by

$$\begin{aligned}
\mathit{rmvfirst}(\langle \rangle, sig) & = \langle \rangle \\
\mathit{rmvfirst}((sig, i, i') \& \sigma, sig) & = \sigma \\
\mathit{rmvfirst}((sig, i, i') \& \sigma, sig') & = (sig, i, i') \& \mathit{rmvfirst}(\sigma, sig') \quad \text{if } sig \neq sig'
\end{aligned}$$

For each process, signals noticing timer expiration have to be merged when time progresses to the next time slice. The function merge is used to describe this precisely. For a given set of extended signals it gives the set of all possible sequences of them. It is inductively defined by

$$\begin{aligned}
\langle \rangle & \in \mathit{merge}(\emptyset) \\
(sig, i, i') \notin \Sigma \wedge \sigma \in \mathit{merge}(\Sigma) & \Rightarrow (sig, i, i') \& \sigma \in \mathit{merge}(\{(sig, i, i')\} \cup \Sigma)
\end{aligned}$$

We define now the set \mathcal{L}_κ of local states. The local state of a process contains, in addition to the above-mentioned components, the name of the process. Thus, the type of the process concerned will not get lost. This is important, because a signal may be sent to an arbitrary process of a process type.

$$\mathcal{L}_\kappa = \mathit{Stg}_\kappa \times \mathit{InpQ}_\kappa \times \mathit{Timers}_\kappa \times P_\kappa$$

We write $\mathit{stg}(L)$, $\mathit{inpq}(L)$, $\mathit{timers}(L)$ and $\mathit{ptype}(L)$, where $L = (\rho, \sigma, \theta, X) \in \mathcal{L}_\kappa$, for ρ , σ , θ and X , respectively.

The global state of a system contains, besides a local state for each existing process, a component keeping track of the system time. To keep track of the system time, natural numbers suffice. We define the state space \mathcal{G}_κ of global states as follows:

$$\mathcal{G}_\kappa = \mathbb{N} \times \bigcup_{I \in \mathcal{P}_{\mathit{fin}}(\mathbb{N}_2)} (I \rightarrow \mathcal{L}_\kappa)$$

We write $\mathit{now}(G)$ and $\mathit{lsts}(G)$, where $G = (n, \Sigma) \in \mathcal{G}_\kappa$, for n and Σ , respectively. We write $\mathit{exists}(i, G)$, where $i \in \mathbb{N}$ and $G \in \mathcal{G}_\kappa$, for $i \in \mathit{dom}(\mathit{lsts}(G))$. Note that 1 is excluded from being used as pid value of an existing process of the system because it is a special pid value that is reserved for the environment.

Every state from \mathcal{G}_κ produces a proposition which is considered to hold in the state concerned. In this way, the state of a process is made partly visible. These propositions are built from the following atomic propositions:

$$\begin{aligned}
\mathit{value} & : V_\kappa \times U \times \mathbb{N}_2 \\
\mathit{active} & : \mathit{Sig}_\kappa \times \mathbb{N}_2
\end{aligned}$$

AtProp_κ denotes the set of all value and active propositions. We write Prop_κ for the set of all propositions that can be built from AtProp_κ . An atomic proposition of the form $\mathit{value}(v, u, i)$ is intended to indicate that u is the value of the local variable v of the process with pid value i . An atomic proposition of the form $\mathit{active}(sig, i)$ is intended to indicate that the timer of the process with pid value i identified with signal sig is active. By using only atomic propositions of these forms,

the state of a process can not be made fully visible via the proposition produced. The proposition produced by each state, given by the function *sig* defined in Section 4.5.4, makes only visible the value of all local variables and the set of active timers for all existing processes.

Below, we introduce the additional actions that are used for the semantics of φ^- SDL system definitions. Like most of the actions used for the semantics of φ^- SDL process definitions, these actions are parametrized. The following additional actions are used:

$$\begin{aligned} \overline{cr} & : \mathbb{N} \times PrCrD_\kappa \\ input' & : ExtSig_\kappa \\ output' & : ExtSig_\kappa \\ set' & : \mathbb{N} \times Sig_\kappa \times \mathbb{N}_2 \\ reset' & : Sig_\kappa \times \mathbb{N}_2 \end{aligned}$$

\overline{Cr}_κ denotes the set of all \overline{cr} actions; and Act'_κ denotes the set of all $input'$, $output'$, set' and $reset'$ actions. The \overline{cr} actions appears as the result of applying the process creation operator E_Φ^p to cr actions. The $input'$, $output'$, set' and $reset'$ actions appear as the result of applying the state operator λ_G to $input$, $output/outputi$, set and $reset$ actions, respectively (see Section 4.5.4).

4.5.2 State transformers and observers In the process algebra semantics of φ^- SDL process definitions, presented in Section 4.3, ACP actions are used to describe the meaning of input guards, SDL actions, the terminator **stop** and the void guard **input none**. These ACP actions are meant to interact with a state from \mathcal{G}_κ . Later on, we will define the result of executing a process, built up from these actions, in a state from \mathcal{G}_κ . That is, we will define the relevant state operator. This will partly boil down to describing how the actions, and the progress of time (modeled by the time unit delay operator σ_{rel}), transform states. For the sake of comprehensibility, we will first define matching state transforming operations.

In addition, we will define some state observing operations. Two of the state observing operations are used directly to define the action and effect function of the state operator and three others are used to define the valuation function of the state operator. The remaining two are used to define the signal function of the state operator as well as an auxiliary evaluation function needed for the value expressions that occur as arguments, and as components of arguments, of the above-mentioned actions and conditions (see Section 4.2).

State transformers: In general, the state transformers change one or two components of the local state of one process. The notable exception is *csmsig*, which is defined first. It may change all components except, of course, the process type. This is a consequence of the fact that the facilities for storage, communication and timing are rather intertwined on the consumption of signals in φ^- SDL. For each state transformer it holds that everything remains unchanged if an attempt is made to transform the local state of a non-existing process. This will not be explicitly mentioned in the explanations given below.

The function $csmsig : ExtSig_\kappa \times V_\kappa^* \times \mathcal{G}_\kappa \rightarrow \mathcal{G}_\kappa$ is used to describe how the ACP actions corresponding to the input guards of φ^- SDL transform states.

$$csmsig((sig, i, i'), \langle v_1, \dots, v_n \rangle, G) = \begin{array}{ll} (now(G), lsts(G) \oplus \{i' \mapsto (\rho, \sigma, \theta, X)\}) & \text{if } exists(i', G) \\ G & \text{otherwise} \end{array}$$

$$\begin{aligned} \text{where } \rho & = stg(lsts(G)_{i'}) \oplus \{v_1 \mapsto vals(sig)_1, \dots, v_n \mapsto vals(sig)_n, \mathbf{sender} \mapsto i\}, \\ \sigma & = rmvfirst(inpq(lsts(G)_{i'}, sig), \\ \theta & = \{sig\} \triangleleft timers(lsts(G)_{i'}), \\ X & = ptype(lsts(G)_{i'}) \end{aligned}$$

$csmsig((sig, i, i'), \langle v_1, \dots, v_n \rangle, G)$ deals with the consumption of signal *sig* by the process with pid

value i' . It transforms the local state of the process with pid value i' , the process by which the signal is consumed, as follows:

- the values carried by sig are assigned to the local variables v_1, \dots, v_n of the consuming process and the sender's pid value (i) is assigned to **sender**;
- the first occurrence of sig in the input queue of the consuming process is removed;
- if sig is a timer signal, it is removed from the active timers.

Everything else is left unchanged.

The function $sndsig : ExtSig_\kappa \times \mathcal{G}_\kappa \rightarrow \mathcal{G}_\kappa$ is used to describe how the ACP actions corresponding to the output actions of φ -SDL transform states.

$$sndsig((sig, i, i'), G) = \begin{array}{ll} (now(G), lsts(G) \oplus \{i' \mapsto (\rho, \sigma, \theta, X)\}) & \text{if } exists(i', G) \\ G & \text{otherwise} \end{array}$$

$$\text{where } \begin{array}{ll} \rho & = stg(lsts(G)_{i'}), \\ \sigma & = inpq(lsts(G)_{i'}) \frown \langle (sig, i, i') \rangle, \\ \theta & = timers(lsts(G)_{i'}), \\ X & = ptype(lsts(G)_{i'}) \end{array}$$

$sndsig((sig, i, i'), G)$ deals with passing signal sig from the process with pid value i to the process with pid value i' . It transforms the local state of the process with pid value i' , the process to which the signal is passed, as follows:

- sig is put into the input queue of the process to which the signal is passed (unless $i' = 1$, indicating that the signal is passed to the environment).

Everything else is left unchanged.

The function $settimer : \mathbb{N} \times Sig_\kappa \times \mathbb{N}_2 \times \mathcal{G}_\kappa \rightarrow \mathcal{G}_\kappa$ is used to describe how the ACP actions corresponding to the set actions of φ -SDL transform states.

$$settimer(t, sig, i, G) = \begin{array}{ll} (now(G), lsts(G) \oplus \{i \mapsto (\rho, \sigma, \theta, X)\}) & \text{if } exists(i, G) \\ G & \text{otherwise} \end{array}$$

$$\text{where } \begin{array}{ll} \rho & = stg(lsts(G)_i), \\ \sigma & = \begin{array}{ll} rmvfirst(inpq(lsts(G)_i), sig) & \text{if } t \geq now(G) \\ rmvfirst(inpq(lsts(G)_i), sig) \frown \langle (sig, i, i) \rangle & \text{otherwise,} \end{array} \\ \theta & = \begin{array}{ll} timers(lsts(G)_i) \oplus \{sig \mapsto t\} & \text{if } t \geq now(G) \\ timers(lsts(G)_i) \oplus \{sig \mapsto nil\} & \text{otherwise,} \end{array} \\ X & = ptype(lsts(G)_i) \end{array}$$

$settimer(t, sig, i, G)$ deals with setting a timer, identified with signal sig , to time t . If t has not yet passed, it transforms the local state of the process with pid value i , the process to be notified of the timer's expiration, as follows:

- the occurrence of sig in the input queue originating from an earlier setting, if any, is removed;
- sig is included among the active timers with expiration time t ; thus overriding an earlier setting, if any.

Otherwise, it transforms the local state of the process with pid value i as follows:

- sig is put into the input queue after removal of its occurrence originating from an earlier setting, if any;

- sig is included among the active timers without expiration time.

Everything else is left unchanged.

The function $resettimer : Sig_\kappa \times \mathbb{N}_2 \times \mathcal{G}_\kappa \rightarrow \mathcal{G}_\kappa$ is used to describe how the ACP actions corresponding to the reset actions of φ -SDL transform states.

$$resettimer(sig, i, G) = \begin{array}{ll} (now(G), lsts(G) \oplus \{i \mapsto (\rho, \sigma, \theta, X)\}) & \text{if } exists(i, G) \\ G & \text{otherwise} \end{array}$$

$$\begin{array}{ll} \text{where } \rho & = stg(lsts(G)_i), \\ \sigma & = rmvfirst(inpq(lsts(G)_i), sig), \\ \theta & = \{sig\} \triangleleft timers(lsts(G)_i), \\ X & = ptype(lsts(G)_i) \end{array}$$

$resettimer(sig, i, G)$ deals with resetting a timer, identified with signal sig . It transforms the local state of the process with pid value i , the process that would otherwise have been notified of the timer's expiration, as follows:

- the occurrence of sig in the input queue originating from an earlier setting, if any, is removed;
- if sig is an active timer, it is removed from the active timers.

Everything else is left unchanged.

Notice that $settimer(t, sig, i, G)$ and $settimer(t, sig, i, resettimer(sig, i, G))$ have the same effect. In other words, $settimer$ resets implicitly. In this way, at most one signal from the same timer will ever occur in an input queue. Furthermore, the context-sensitive conditions for syntactic correctness of φ -SDL specifications imply that timer signals and other signals are kept apart: not a single signal can originate from both timer setting and customary signal sending. Thus, resetting, either explicitly or implicitly, will solely remove signals from input queues that originate from timer setting.

The function $assignvar : V_\kappa \times U \times \mathbb{N}_2 \times \mathcal{G}_\kappa \rightarrow \mathcal{G}_\kappa$ is used to describe how the ACP actions corresponding to the assignment task actions of φ -SDL transform states.

$$assignvar(v, u, i, G) = \begin{array}{ll} (now(G), lsts(G) \oplus \{i \mapsto (\rho, \sigma, \theta, X)\}) & \text{if } exists(i, G) \\ G & \text{otherwise} \end{array}$$

$$\begin{array}{ll} \text{where } \rho & = stg(lsts(G)_i) \oplus \{v \mapsto u\}, \\ \sigma & = inpq(lsts(G)_i), \\ \theta & = timers(lsts(G)_i), \\ X & = ptype(lsts(G)_i) \end{array}$$

$assignvar(v, u, i, G)$ deals with assigning value u to variable v . It transforms the local state of the process with pid value i , the process to which the variable is local, as follows:

- u is assigned to the local variable v , i.e. v is included among the variables in the storage with value u ; thus overriding an earlier assignment, if any.

Everything else is left unchanged.

The function $createproc : \mathbb{N}_2 \times PrCrD'_\kappa \times \mathcal{G}_\kappa \rightarrow \mathcal{G}_\kappa$ is used to describe how the ACP actions corresponding to the create actions of φ -SDL transform states. The elements of $PrCrD'_\kappa$ are like process creation data, i.e. elements of $PrCrD_\kappa$, but values are used instead of value expressions: $PrCrD'_\kappa = \{(X, \langle v_1, \dots, v_n \rangle, \langle u_1, \dots, u_n \rangle, i) \mid X \in P_\kappa, v_1, \dots, v_n \in V_\kappa, u_1, \dots, u_n \in U, i \in \mathbb{N}_2\}$.

$$createproc(i', (X', \langle v_1, \dots, v_n \rangle, \langle u_1, \dots, u_n \rangle, i), G) = \begin{array}{ll} (now(G), lsts(G) \oplus \{i \mapsto (\rho, \sigma, \theta, X), i' \mapsto (\rho', \sigma', \theta', X')\}) & \text{if } exists(i, G) \\ G & \text{otherwise} \end{array}$$

$$\begin{aligned}
\text{where } \rho &= stg(lsts(G)_i) \oplus \{\mathbf{offspring} \mapsto i'\}, \\
\sigma &= inpq(lsts(G)_i), \\
\theta &= timers(lsts(G)_i), \\
X &= ptype(lsts(G)_i), \\
\rho' &= \{v_1 \mapsto u_1, \dots, v_n \mapsto u_n, \mathbf{parent} \mapsto i, \mathbf{offspring} \mapsto 0, \mathbf{sender} \mapsto 0\}, \\
\sigma' &= \langle \rangle, \\
\theta' &= \{ \}
\end{aligned}$$

$createproc(i', (X', \langle v_1, \dots, v_n \rangle, \langle u_1, \dots, u_n \rangle, i), G)$ deals with creating a process of type X' . It transforms the local state of the process with pid value i , the parent of the created process, as follows:

- the pid value of the created process (i') is assigned to **offspring**.

Besides, it creates a new local state for the created process which is initiated as follows:

- the values u_1, \dots, u_n are assigned to the local variables v_1, \dots, v_n of the created process and the parent's pid value (i) is assigned to **parent**;
- X' is made the process type.

Everything else is left unchanged.

The function $stopproc : \mathbb{N}_2 \times \mathcal{G}_\kappa \rightarrow \mathcal{G}_\kappa$ is used to describe how the ACP actions corresponding to **stop** in φ -SDL transform states.

$$stopproc(i, G) = (now(G), \{i\} \triangleleft lsts(G))$$

$stopproc(i, G)$ deals with terminating the process with pid value i . It disposes of the local state of the process with pid value i . Everything else is left unchanged.

The function $inispont : \mathbb{N}_2 \times \mathcal{G}_\kappa \rightarrow \mathcal{G}_\kappa$ is used to describe how the ACP actions used to initiate spontaneous transitions transform states.

$$\begin{aligned}
inispont(i, G) &= \\
&\begin{array}{ll}
(now(G), lsts(G) \oplus \{i \mapsto (\rho, \sigma, \theta, X)\}) & \text{if } exists(i, G) \\
G & \text{otherwise}
\end{array}
\end{aligned}$$

$$\begin{aligned}
\text{where } \rho &= stg(lsts(G)_i) \oplus \{\mathbf{sender} \mapsto i\}, \\
\sigma &= inpq(lsts(G)_i), \\
\theta &= timers(lsts(G)_i), \\
X &= ptype(lsts(G)_i)
\end{aligned}$$

$inispont(i, G)$ deals with initiating spontaneous transitions. It transforms the local state of the process with pid value i , the process for which a spontaneous transition is initiated, by assigning i to **sender**. Everything else is left unchanged.

The function $unitdelay : \mathcal{G}_\kappa \rightarrow \mathcal{P}_{fin}(\mathcal{G}_\kappa)$ is used to describe how progress of time transforms states. In general, these transformations are non-deterministic – how signals from expiring timers enter input queues is not uniquely determined. Therefore, this function yields for each state a set of possible states.

$$\begin{aligned}
G' \in unitdelay(G) &\leftrightarrow \\
now(G') &= now(G) + 1 \wedge \\
\forall i \in dom(lsts(G)) &\cdot \\
stg(lsts(G')_i) &= stg(lsts(G)_i) \wedge \\
(\exists \sigma \in InpQ &\cdot \\
inpq(lsts(G')_i) &= inpq(lsts(G)_i) \frown \sigma \wedge \\
\sigma \in merge(\{ \{sig, i, i\} \mid timers(lsts(G)_i)(sig) \leq now(G) \}) &\wedge \\
timers(lsts(G')_i) &= \\
timers(lsts(G)_i) \oplus \{sig \mapsto nil \mid timers(lsts(G)_i)(sig) \leq now(G)\} &\wedge \\
ptype(lsts(G')_i) &= ptype(lsts(G)_i)
\end{aligned}$$

$unitdelay(G)$ transforms the global state as follows:

- the system time is incremented by one unit;
- for the local state of each process:
 - its storage is left unchanged;
 - the signals that correspond to expiring timers are put into the input queue in a non-deterministic way;
 - for each of the expiring timers, the expiration time is removed;
 - its process type is left unchanged.

State observers: In general, the state observers examine one component of the local state of some process. The only exceptions are *has-instance* and *instances*, which may even examine the process type component of all processes. If an attempt is made to observe the local state of a non-existing process, each non-boolean-valued state observer yields nil and each boolean-valued state observer yields f . This will not be explicitly mentioned in the explanations given below.

The function $contents : V_\kappa \times \mathbb{N}_2 \times \mathcal{G}_\kappa \rightarrow U \cup \{\text{nil}\}$ is used to describe the value of expressions of the form $value(v, e)$, which correspond to the variable accesses and view expressions of φ -SDL.

$$contents(v, i, G) = \begin{array}{ll} \rho(v) & \text{if } exists(i, G) \wedge v \in dom(\rho) \\ \text{nil} & \text{otherwise} \end{array}$$

$$\text{where } \rho = stg(lsts(G)_i)$$

$contents(v, i, G)$ yields the current value of the variable v that is local to the process with pid value i .

The function $is-active : Sig_\kappa \times \mathbb{N}_2 \times \mathcal{G}_\kappa \rightarrow \mathbb{B}$ is used to describe the value of expressions of the form $active(sig, e)$, which correspond to the active expressions of φ -SDL.

$$is-active(sig, i, G) = \begin{array}{ll} \text{t} & \text{if } exists(i, G) \wedge sig \in dom(\theta) \\ \text{f} & \text{otherwise} \end{array}$$

$$\text{where } \theta = timers(lsts(G)_i)$$

$is-active(sig, i, G)$ yields true iff sig is an active timer signal of the process with pid value i .

The function $is-waiting : SaveSet_\kappa \times \mathbb{N}_2 \times \mathcal{G}_\kappa \rightarrow \mathbb{B}$ is used to describe the value of conditions of the form $waiting(\{s_1, \dots, s_n\}, e)$, which are used to give meaning to the state definitions of φ -SDL.

$$is-waiting(ss, i, G) = \begin{array}{ll} \text{t} & \text{if } exists(i, G) \wedge getnxt(\sigma, ss) = \text{nil} \\ \text{f} & \text{otherwise} \end{array}$$

$$\text{where } \sigma = inpq(lsts(G)_i)$$

$is-waiting(ss, i, G)$ yields true iff there is no signal in the input queue of the process with pid value i that has a name different from the ones in ss .

The function $type : \mathbb{N}_1 \times \mathcal{G}_\kappa \rightarrow P_\kappa \cup \{\mathbf{env}\} \cup \{\text{nil}\}$ is used to describe the value of conditions of the form $type(e, X)$, which are used to give meaning to the output actions with explicit addressing of φ -SDL.

$$type(i, G) = \begin{array}{ll} X & \text{if } exists(i, G) \\ \mathbf{env} & \text{if } i = 1 \\ \text{nil} & \text{otherwise} \end{array}$$

$$\text{where } X = ptype(lsts(G)_i)$$

$type(i, G)$ yields the type of the process with pid value i . Different from the other state observers, it yields a result if $i = 1$ as well, viz. \mathbf{env} .

The function $has-instance : (P_\kappa \cup \{\mathbf{env}\}) \times \mathcal{G}_\kappa \rightarrow \mathbb{B}$ is used to describe the value of conditions of the form $hasinst(X)$, which are used to give meaning to the output actions with implicit addressing of φ -SDL.

$$\begin{aligned} \text{has-instance}(X, G) = & \text{t} \quad \text{if } \exists i \in \mathbb{N}_1 \cdot (i = 1 \vee \text{exists}(i, G)) \wedge \text{type}(i, G) = X \\ & \text{f} \quad \text{otherwise} \end{aligned}$$

$\text{has-instance}(X, G)$ yields true iff there exists a process of type X .

The function $\text{nextsig} : \text{SaveSet}_\kappa \times \mathbb{N}_2 \times \mathcal{G}_\kappa \rightarrow \text{ExtSig}_\kappa \cup \{\text{nil}\}$ is used to describe the result of executing *input* actions, which are used to give meaning to the input guards of φ -SDL.

$$\begin{aligned} \text{nextsig}(ss, i, G) = & \text{getnext}(\sigma, ss) \quad \text{if } \text{exists}(i, G) \\ & \text{nil} \quad \text{otherwise} \end{aligned}$$

where $\sigma = \text{inpq}(\text{lsts}(G)_i)$

$\text{nextsig}(ss, i, G)$ yields the first signal in the input queue of the process with pid value i that has a name different from the ones in ss .

The function $\text{instances} : (P_\kappa \cup \{\mathbf{env}\}) \times \mathcal{G}_\kappa \rightarrow \mathcal{P}_{\text{fin}}(\mathbb{N}_2) \cup \{\{1\}\}$ is used to describe the result of executing *outputi* actions, which correspond to the output actions with implicit addressing of φ -SDL.

$$\begin{aligned} \text{instances}(X, G) = & \{i \in \text{dom}(\text{lsts}(G)) \mid \text{type}(i, G) = X\} \quad \text{if } X \neq \mathbf{env} \\ & \{1\} \quad \text{otherwise} \end{aligned}$$

$\text{instances}(X, G)$ yields the set of pid values of all existing processes of type X if $X \neq \mathbf{env}$ and the singleton set $\{1\}$ otherwise.

4.5.3 Evaluation of value expressions The function eval_G is used to evaluate value expressions in a state G . The state observers *contents* and *is-active* defined in Section 4.5.2 are used to define the evaluation function.

$$\begin{aligned} \text{eval}_G(\text{op}(e_1, \dots, e_n)) = & \\ \text{op}(\text{eval}_G(e_1), \dots, \text{eval}_G(e_n)) & \quad \text{if } \text{eval}_G(e_1) \neq \text{nil} \wedge \dots \wedge \text{eval}_G(e_n) \neq \text{nil} \\ \text{nil} & \quad \text{otherwise} \end{aligned}$$

$$\begin{aligned} \text{eval}_G(\text{cond}(e_1, e_2, e_3)) = & \text{eval}_G(e_2) \quad \text{if } \text{eval}_G(e_1) = \text{t} \\ & \text{eval}_G(e_3) \quad \text{if } \text{eval}_G(e_1) = \text{f} \\ \text{nil} & \quad \text{otherwise} \end{aligned}$$

$$\begin{aligned} \text{eval}_G(\text{value}(v, e)) = & \text{contents}(v, \text{eval}_G(e), G) \quad \text{if } \text{eval}_G(e) \in \mathbb{N}_2 \\ \text{nil} & \quad \text{otherwise} \end{aligned}$$

$$\begin{aligned} \text{eval}_G(\text{active}(s(e_1, \dots, e_n), e)) = & \\ \text{is-active}(\text{sig}, \text{eval}_G(e), G) & \quad \text{if } \text{eval}_G(e_1) \neq \text{nil} \wedge \dots \wedge \text{eval}_G(e_n) \neq \text{nil} \wedge \\ & \text{eval}_G(e) \in \mathbb{N}_2 \\ \text{nil} & \quad \text{otherwise} \end{aligned}$$

where $\text{sig} = (s, \langle \text{eval}_G(e_1), \dots, \text{eval}_G(e_n) \rangle)$

$$\text{eval}_G(\text{now}) = \text{now}(G)$$

In all of these cases, if the value of at least one of the subexpressions occurring in an expression is undefined in the state concerned, the expression will be undefined, i.e. yield *nil*.

We extend eval_G to sequences of value expressions and signal expressions as follows:

$$\begin{aligned} \text{eval}_G(\langle e_1, \dots, e_n \rangle) = & \\ \langle \text{eval}_G(e_1), \dots, \text{eval}_G(e_n) \rangle & \quad \text{if } \text{eval}_G(e_1) \neq \text{nil} \wedge \dots \wedge \text{eval}_G(e_n) \neq \text{nil} \\ \text{nil} & \quad \text{otherwise} \end{aligned}$$

$$\begin{aligned} \text{eval}_G((s, es)) = & (s, \text{eval}_G(es)) \quad \text{if } \text{eval}_G(es) \neq \text{nil} \\ \text{nil} & \quad \text{otherwise} \end{aligned}$$

4.5.4 Definition of the state operator In this subsection, we define the functions act , eff , eff_σ , sig and val in completion of the definition of the state operator used to describe the meaning of φ -SDL system definitions. Recall that for this state operator $S = \mathcal{G}_\kappa$.

Action and effect functions: In the current application of $\text{PA}_{\text{drt}}^{\text{psc}}$, $A = Cr_\kappa \cup Act_\kappa \cup Act'_\kappa \cup \{\#\}$, where $Act_\kappa = Act_\kappa^- \cup \overline{Cr}_\kappa$. The actions in Act_κ are actions that may change the state in which they are executed. The actions in Act'_κ are actions that are performed as the result of the execution of actions in Act_κ in a state. The actions in Cr_κ and \overline{Cr}_κ are the process creation actions and the actions that are left as a trace of the process creations that occur, respectively – note that we use the data set $D = PrCrD_\kappa$ for process creation. The action and effect functions are trivial for actions that are not in Act_κ .

In order to keep the definitions comprehensible, we will use the following abbreviations. Let e be a value expression, let es be a sequence of value expressions, and let se be a signal expression. Then we write e' for $eval_G(e)$, es' for $eval_G(es)$ and se' for $eval_G(se)$.

$$\begin{aligned}
act(input((s, \langle v_1, \dots, v_n \rangle), ss, e), G) &= \\
&\{input'((sig, i, j) \mid snm(sig) = s, natsig(ss, e', G) = (sig, i, j), e' \in \mathbb{N}_2)\} \\
eff(input((s, \langle v_1, \dots, v_n \rangle), ss, e), G, a) &= \\
&csmsig(natsig(ss, e', G), \langle v_1, \dots, v_n \rangle, G) \text{ if } a \in act(input((s, \langle v_1, \dots, v_n \rangle), ss, e), G) \\
&G \qquad \qquad \qquad \text{otherwise} \\
act(output(se, e_1, e_2), G) &= \\
&\{output'((se', e'_1, e'_2) \mid se' \neq \text{nil}, e'_1 \in \mathbb{N}_2, e'_2 \in \mathbb{N}_1)\} \\
eff(output(se, e_1, e_2), G, a) &= \\
&sndsig((se', e'_1, e'_2), G) \text{ if } a \in act(output(se, e_1, e_2), G) \\
&G \qquad \qquad \qquad \text{otherwise} \\
act(outputi(se, e, X), G) &= \\
&\{output'((se', e', i) \mid se' \neq \text{nil}, e' \in \mathbb{N}_2, i \in instances(X, G))\} \\
eff(outputi(se, e, X), G, a) &= \\
&sndsig((se', e', i), G) \text{ if } a \in act(outputi(se, e, X), G) \wedge a = output'((se', e', i)) \\
&G \qquad \qquad \qquad \text{otherwise} \\
act(set(e_1, se, e_2), G) &= \{set'(e'_1, se', e'_2) \mid e'_1 \in \mathbb{N}, se' \neq \text{nil}, e'_2 \in \mathbb{N}_2\} \\
eff(set(e_1, se, e_2), G, a) &= settimer(e'_1, se', e'_2, G) \text{ if } a \in act(set(e_1, se, e_2), G) \\
&G \qquad \qquad \qquad \text{otherwise} \\
act(reset(se, e), G) &= \{reset'(se', e') \mid se' \neq \text{nil}, e' \in \mathbb{N}_2\} \\
eff(reset(se, e), G, a) &= resettimer(se', e', G) \text{ if } a \in act(reset(se, e), G) \\
&G \qquad \qquad \qquad \text{otherwise} \\
act(ass(v, e_1, e_2), G) &= \{\# \mid e'_1 \neq \text{nil}, e'_2 \in \mathbb{N}_2\} \\
eff(ass(v, e_1, e_2), G, a) &= assignvar(v, e'_1, e'_2, G) \text{ if } a \in act(ass(v, e_1, e_2), G) \\
&G \qquad \qquad \qquad \text{otherwise} \\
act(\overline{cr}(i, (X, vs, es, e)), G) &= \{\# \mid i \in \mathbb{N}_2, es' \neq \text{nil}, e' \in \mathbb{N}_2\} \\
eff(\overline{cr}(i, (X, vs, es, e)), G, a) &= \\
&createproc(i, (X, vs, es', e'), G) \text{ if } a \in act(\overline{cr}(i, (X, vs, es, e)), G) \\
&G \qquad \qquad \qquad \text{otherwise}
\end{aligned}$$

$$\begin{aligned}
act(stop(e), G) &= \{\mathbf{t} \mid e' \in \mathbb{N}_2\} \\
eff(stop(e), G, a) &= \begin{array}{l} stopproc(e', G) \text{ if } a \in act(stop(e), G) \\ G \text{ otherwise} \end{array} \\
act(inispont(e), G) &= \{\mathbf{t} \mid e' \in \mathbb{N}_2\} \\
eff(inispont(e), G, a) &= \begin{array}{l} inispont(e', G) \text{ if } a \in act(inispont(e), G) \\ G \text{ otherwise} \end{array}
\end{aligned}$$

For all actions $a \in Cr_\kappa \cup Act'_\kappa \cup \{\mathbf{t}\}$:

$$\begin{aligned}
act(a, G) &= a \\
eff(a, G, a') &= G
\end{aligned}$$

The effect of applying the state operator to a process of the form $\sigma_{rel}(x)$ is described by mean of the function eff_σ .

$$eff_\sigma(G) = unitdelay(G)$$

Signal function: First, we define the function $atoms : \mathcal{G}_\kappa \rightarrow \mathcal{P}_{fin}(AtProp_\kappa)$ giving for each state the set of atomic propositions that hold in that state. It is inductively defined by

$$\begin{aligned}
contents(v, i, G) = u &\Rightarrow value(v, u, i) \in atoms(G) \\
is-active(sig, i, G) = \mathbf{t} &\Rightarrow active(sig, i) \in atoms(G)
\end{aligned}$$

We now define the function $sig : \mathcal{G}_\kappa \rightarrow Prop_\kappa$, giving the propositions produced by the states. as follows:

$$sig(G) = \bigwedge_{\phi \in atoms(G)} \phi$$

So $sig(G)$ is the conjunction of all atomic propositions that hold in state G .

Valuation function: In the current application of PA_{drt}^{psc} , $B_{at} = AtCond_\kappa \cup AtProp_\kappa$.

The function val is defined in the following way:

$$\begin{aligned}
val(eq(e_1, e_2), G) &= \mathbf{t} \text{ if } e'_1 = e'_2 \wedge e'_1 \neq nil \wedge e'_2 \neq nil \\
&\mathbf{f} \text{ otherwise}
\end{aligned}$$

$$val(waiting(ss, e), G) = is-waiting(ss, e', G) \wedge e' \in \mathbb{N}_2$$

$$\begin{aligned}
val(type(e, X), G) &= \mathbf{t} \text{ if } type(e', G) = X \wedge e' \in \mathbb{N}_1 \\
&\mathbf{f} \text{ otherwise}
\end{aligned}$$

$$val(hasinst(X), G) = has-instance(X, G)$$

$$\begin{aligned}
val(value(v, u, i), G) &= \mathbf{t} \text{ if } value(v, u, i) \in atoms(G) \\
&\mathbf{f} \text{ otherwise}
\end{aligned}$$

$$\begin{aligned}
val(active(sig, i), G) &= \mathbf{t} \text{ if } active(sig, i) \in atoms(G) \\
&\mathbf{f} \text{ otherwise}
\end{aligned}$$

Note that the sets $AtCond_\kappa$ and $AtProp_\kappa$ are disjoint. The elements of $AtCond_\kappa$ are used as conditions with a truth value that depends upon the state. The elements of $AtProp_\kappa$ are not used as conditions and their valuation in a state is not needed for the semantics of φ^- SDL.

4.6 Semantics of system definitions

In this subsection, we present a semantics for φ -SDL system definitions. It relies heavily upon the specifics of the state operator defined in Section 4.5.

According to the semantics presented here, the meaning of a φ -SDL system definition is a process described by a process expression – a term of $\text{PA}_{\text{drt}}^{\text{psc}}$ extended with the counting process creation operator and the state operator to be precise. It is given in terms of the semantics of φ -SDL process definitions presented in Section 4.3 and the contextual information extracted by means of the function $\{\llbracket - \rrbracket\}$ defined in Appendix B.

4.6.1 System definition The semantics of a φ -SDL system definition depends on a parameter Env representing the environment of the system. This environment Env has to be described by a $\text{PA}_{\text{drt}}^{\text{psc}}$ term.

$$\begin{aligned} \llbracket \text{system } S; D_1 \dots D_n \text{ endsystem}; \rrbracket(Env) &:= \lambda_{G_0} (E_{\Phi}^{n_0+2}(P) \parallel Env) \\ \text{where } P &= \parallel_{i \in \{1, \dots, n_0\}} \Phi(i+1, (pt(i+1), \langle \rangle, \langle \rangle, 0)), \\ \Phi &= \{(i, d) \mapsto \Psi(i) \mid \exists D \in \{D_1, \dots, D_n\} \cdot \llbracket D \rrbracket^{\kappa} = (pnm(d), \Psi)\}, \\ G_0 &= (0, \{i+1 \mapsto L_0(i+1) \mid i \in \{1, \dots, n_0\}\}), \\ L_0(i) &= (\{\mathbf{parent} \mapsto 0, \mathbf{offspring} \mapsto 0, \mathbf{sender} \mapsto 0\}, \langle \rangle, \{ \}, pt(i)), \\ n_0 &= \sum_{X \in \text{procs}(\kappa)} \text{init}(\kappa, X), \\ \kappa &= \llbracket \text{system } S; D_1 \dots D_n \text{ endsystem}; \rrbracket \\ \text{and } pt &: \{1+1, \dots, n_0+1\} \rightarrow \text{procs}(\kappa) \text{ is such that} \\ &\forall X \in \text{procs}(\kappa) \cdot \text{card}(pt^{-1}(X)) = \text{init}(\kappa, X). \end{aligned}$$

The function pt is used to assign pid values for the processes created during system start-up. It maps a set of pid values to the types of the processes with these pid values.

The process expression that corresponds to a system definition expresses that, for each of the process types defined, the given initial number of processes are created and these processes are executed in parallel, starting in the state G_0 , while they receive signals via signal routes from the environment Env . G_0 is the state in which the system time is zero and there is a local state for each of the processes that is created during system start-up. Recall that the pid value 1 is reserved for the environment. The mapping Φ from pid values and process creation data to process expressions is derived from the meaning of the process definitions occurring in the system definition. This mapping is used by the counting process creation operator, which is needed for process creation after system start-up.

4.6.2 Environments The semantics of φ -SDL system definitions describes the meaning of a system definition for an arbitrary process Env that represents the behaviour of the environment. Here we explain how the environment's behaviour is described by a $\text{PA}_{\text{drt}}^{\text{psc}}$ process.

Some general assumptions have to be made about the behaviour of the environment of any system described using φ -SDL. Further assumptions may be made about the behaviour of the environment of a specific system described using φ -SDL, including ones that facilitate analysis of the system concerned. The general assumptions made about environments are:

- the environment can only send signals that are defined in the system definition concerned;
- the environment can only send signals to processes to which it is connected by signal routes;
- the environment can send only a finite number of signals during a time slice.

Besides, the viewpoint is taken that the processes that constitute a system are not observable to its environment. This leads to the use of output actions with implicit addressing in representing the environment's behaviour.

The set $EnvSig_{\kappa}$ of possible environment signals is determined by the specific types of signals and signal routes introduced in the system definition concerned. It can be obtained from the environment

signal description yielded by applying the function *envsigd* (defined in Appendix B) to the context ascribed to the system definition. For an arbitrary context κ , the set of environment signals is obtained as follows:

$$EnvSig_\kappa = \bigcup_{((s, \langle T_1, \dots, T_n \rangle), X) \in envsigd(\kappa)} \{((s, \langle t_1, \dots, t_n \rangle), 1, X) \mid t_1 \in T_1^A, \dots, t_n \in T_n^A\}$$

It is clear that in general the set $EnvSig_\kappa$ is infinite because signals can carry values from infinite domains. Besides, the environment can send an arbitrary signal from $EnvSig_\kappa$. To represent this we need the alternative composition of an infinite number of alternatives. This can be done using the operator $\sum_{d:D}$ of μ CRL [20]. However, the combination of this operator with the extension of discrete time process algebra we are using has not been investigated thoroughly. Besides, the potential unbounded non-determinism introduced by this operator does not allow to use conventional validation techniques.

A process that satisfies the three above-mentioned assumptions can be defined in the following way:

$$\begin{aligned} Env_\kappa &= \sum_{n:\mathbb{N}} Env_n \\ Env_0 &= \sigma_{rel}(Env_\kappa) \\ Env_{n+1} &= Env_n + \sum_{osig \in EnvSig_\kappa} \underline{outputi(osig)} \cdot Env_n \end{aligned}$$

In order to use an environment process as a parameter of the presented semantics of φ -SDL, it has to be given as a process in PA_{drt}^{psc} . Below we define such an environment process. We call it a standard environment process for the semantics of φ -SDL. It is determined by two restrictions:

- the set of signals which the environment can send to the system is restricted to a finite subset $ES \subseteq EnvSig_\kappa$;
- the maximal number of signals which can be sent in one time slice is bounded by a natural number N_s .

A standard environment is defined as follows:

$$\begin{aligned} Env_\kappa^{st} &= \sum_{n \in \{0, \dots, N_s\}} Env'_n \\ Env'_0 &= \sigma_{rel}(Env_\kappa^{st}) \\ Env'_{n+1} &= Env'_n + \sum_{osig \in ES} \underline{outputi(osig)} \cdot Env'_n \end{aligned}$$

One may also define another environment process for a specific system. In any case the process representing the system's environment must satisfy the above-mentioned assumptions.

4.6.3 Delaying channels The process algebra semantics of φ -SDL makes clear how to model a delaying channel by means of a φ -SDL process. Below a φ -SDL process definition of such a process, called **ch**, is given. It is assumed that the process can only receive signals of type **es**.

The process consumes signals **es(rcv, v1, ..., vn)** and passes them on after an arbitrary delay, possibly zero, with **rcv** replaced by **snd**. Each signal consumed carries the pid value of the ultimate receiver, and this pid value is replaced by the one of the original sender when the signal is passed on. This is needed because the original sender and ultimate receiver have now an intermediate receiver and intermediate sender, respectively. The **decision** construct is used to find out whether the original sender used implicit or explicit addressing. Note that we write **Null** for 0, i.e. the special pid value that never refers to any existing process.

```
process ch(1);
  start;
```

```

    nextstate receive;
state receive;
  input es(rcv,v1,...,vn);
    task snd := sender;
    nextstate deliver;
state deliver;
  save es;
  input none;
  decision rcv = Null;
    (True):
      output es(snd,v1,...,vn) via sr_out;
      nextstate receive;
    (False):
      output es(snd,v1,...,vn) to rcv via sr_out;
      nextstate receive;
  enddecision;
endprocess;

```

In state `deliver`, there will never be signals to consume because all signals are withheld from being consumed by means of `save es`. This means that the behaviour from this state may be delayed till any future time slice. The total lack of signals to consume does not preclude the process to proceed, because the only transition alternative is a spontaneous transition, i.e. it is not initiated by the consumption of a signal.

5. CLOSING REMARKS

In this section, we sum up what has been achieved. We also describe in broad outline what is anticipated to be achieved more, thus hinting at topics for future research. But, to begin with, we present an overview of related work.

5.1 Related work

In [15], a foundation for the semantics of SDL, based on streams and stream processing functions, has been proposed. This proposal indicates that the SDL view of a system gives an interesting type of asynchronous dataflow networks, but the treatment of time in the proposal is however too sketchy to be used as a starting point for a semantics of the time related features of SDL. Besides, process creation is not covered. In [13], we present a process algebra model of asynchronous dataflow networks as a semantic foundation for SDL. The model is close to the concepts around which SDL has been set up. However, we are not able to cover process creation too.

An operational semantics for a subset of SDL, which covers timing, has been given in [19]. Many relevant details are worked out, but it is not quite clear whether time is treated satisfactory. This is largely due to the intricacy of the operational semantics. At the outset, we have also tried shortly to give an operational semantics for φ -SDL, but we found that it is very difficult. Our experience is that the main motivations for the rules describing an operational semantics are unavoidably intuitive. This may already lead to an undesirable semantics in relatively simple cases. For example, working on $\text{PA}_{\text{drt}}^{\text{psc}}$, we have seriously considered to have the premise $w \in [s_p(x)]$ in the rule for time unit delay (see Table 11). However, this plausible option would render all delayed processes bisimilar to deadlock in the current time slice, i.e. $\sigma_{\text{rel}}(x) = \underline{\delta}$ would hold.

It is likely that, if we had taken parallel composition with communication, we would have been able to use special processes, put in parallel with the other ones, instead of the counting process creation operator and the state operator. The approach to use special processes is followed in [32]. There, it is largely responsible for the exceptional intricacy of the resulting semantics, which, by the way, has kept various obvious errors unnoticed for a long time. Amongst other things for this reason, we have chosen to use the counting process creation operator and the state operator instead. Besides, the approach to use special processes brings along a lot of internal communication that is irrelevant

from a semantic point of view. Of course, in case an ACP-style process algebra is used as the basis of the presented semantics, there is the option to use an abstraction operator to abstract from the internal communication (for abstraction in process algebra, see e.g. [7]). However, it seems far from easy to elaborate the addition of abstraction to $\text{PA}_{\text{drt}}^{\text{psc}}$ or its adaptation to parallel composition with communication.

For a subset of SDL, called μSDL , both an operational semantics and a process algebra semantics has been given in [18]. The operational semantics of μSDL is related to the one presented in [19], but time is not treated. The basis of the corresponding process algebra semantics is a time free process algebra, essentially μCRL [20] extended with the state operator. Like in [32], special processes are used for channels and input queues although there is no need for that with the state operator at one's disposal. Interesting is that first the intended meaning of the language constructs is laid down in an operational semantics so that it can be reasonably checked later whether the process algebra semantics reflects the intentions. However, the SDL facilities for storage, timing and process creation are not available at all in μSDL ; and the facilities for communication are only partially available.

In [17], BSDL is proposed as a basis for the semantics of SDL. BSDL is developed from scratch, using Object-Z, but it does not seem to fit in very well with SDL. In [22], it is proposed to use Duration Calculus [27] to describe the meaning of the language constructs of SDL. Thus an interesting semantic link is made between SDL and Duration Calculus, but it seems a little bit odd to view the main semantic objects used, viz. traces, as phases of system behaviour, called state variables in Duration Calculus, of which the duration is the principal attribute.

5.2 Conclusions and future work

We have presented an extension of discrete time process algebra with relative timing and we have proposed a semantics for the core of SDL, using this extension to describe the meaning of the language constructs. The operational semantics and axiom system of this ACP-style process algebra facilitates advances in the areas of validation of SDL specifications and verification of design steps made using SDL, respectively. At present, we focus on validation. We do so because practically useful advanced tools and techniques are within reach now while there is a tremendous need for them.

For validation purposes, we have to transform φ -SDL specifications to transition systems in accordance with the process algebra semantics. Generating transition systems from finite linear recursive specifications in $\text{BPA}_{\text{drt}}^{\text{psc}}$ (see Section 2.4) is straightforward. In Section 4, the meaning of a φ -SDL process definition is described by a finite guarded recursive specification in $\text{BPA}_{\text{drt}}^{\text{psc}}$ that can definitely be brought into linear form. The meaning of a φ -SDL system definition is given in terms of the meaning of the occurring process definitions using the parallel composition operator, the counting process creation operator and the state operator. An obvious direction is to devise, for each of these operators on processes, a corresponding syntactic operator on finite linear recursive specifications that, under certain conditions, yields a linear recursive specification of the process that results from applying the operator on processes to the process(es) defined by the recursive specification(s) to be operated on. Of course, we look for non-restrictive conditions, but finiteness of the state space, a finite upper bound on the total number of process creations and a finite upper bound on the number of signals per time slice originating from the environment are inescapable. It goes without saying that we have to show the correctness of these syntactic operators. For that purpose, we have available the axioms presented in Section 2 and RSP (see Section 2.4).

In [10], timed frames, which are closely related to the kind of transition systems presented by the operational semantics of, for example, $\text{BPA}_{\text{drt}}^-$ -ID and PA_{drt}^- -ID, are studied in a general algebraic setting and results concerning the connection between timed frames and discrete time processes with relative timing are given. In [14], a model for $\text{BPA}_{\text{drt}}^-$ -IDlin ($\text{BPA}_{\text{drt}}^-$ -ID with finite linear recursive specifications) is presented that gives an interpretation of its constants and operators on timed frames; and it is shown that the bisimulation model induced by the original structured operational semantics is isomorphic to the timed frame model. It is plausible that these results can be extended to $\text{BPA}_{\text{drt}}^{\text{psc}}$ lin – timed frames support propositional signals. This would mean that we can transform φ -SDL speci-

fications to timed frames. In that case, we can basically check whether a system described in φ -SDL satisfies a property expressed in TFL [11], an expressive first-order logic proposed for timed frames. We are considering to look for a fragment of TFL that is suitable to serve as a logic for φ -SDL and to adapt an existing model checker to φ -SDL and this logic – and thus to automate the checking. In particular, the model checker MEC [2] seems suited for this purpose – at least for small-scale systems. A fragment of Duration Calculus may be considered as well, since in [21] validity for Duration Calculus formulas in timed frames is defined.

The extension of discrete time process algebra with relative timing, used to describe the meaning of the language constructs of φ -SDL, is fairly large and rather intricate. Theoretically interesting general properties, such as elimination, conservativity, completeness, etc. have yet to be established. We think that we are near the limit of what can be made provably free from defect. Still, owing to the nontrivial state space taken for the state operator, the presented semantics for φ -SDL uses an excursion outside the realm of process algebra that is not negligible. We wish to have abstraction added to the process algebra used in order to provide a more abstract semantics for φ -SDL, but right now we consider the consequences of this addition too difficult to grasp. All this suggests the option to develop a special process algebra that is closer to the concepts around which SDL has been set up. Of course, there is also the alternative to simplify SDL by removing SDL features that introduce semantic complexities but do not serve any practical purpose. The presented semantics of φ -SDL may assist in identifying such cases.

REFERENCES

1. B. Algayres, Y. Lejeune, F. Hugonnet, and F. Hantz. The AVALON project: A validation environment for SDL/MSD descriptions. In O. Færgemand and A. Sarma, editors, *SDL '93: Using Objects*, pages 221–235. North-Holland, 1993.
2. A. Arnold. MEC, a system for constructing and analysing transition systems. In J. Sifakis, editor, *Automatic Verification Methods for Finite State Systems*, pages 117–132. LNCS 407, Springer Verlag, 1990.
3. J.C.M. Baeten and J.A. Bergstra. Discrete time process algebra (extended abstract). In W.R. Cleaveland, editor, *CONCUR'92*, pages 401–420. LNCS 630, Springer-Verlag, 1992. Full version: Report P9208b, Programming Research Group, University of Amsterdam.
4. J.C.M. Baeten and J.A. Bergstra. Discrete time process algebra. *Formal Aspects of Computing*, 8:188–208, 1996.
5. J.C.M. Baeten and J.A. Bergstra. Process algebra with propositional signals. *Theoretical Computer Science*, 177:381–405, 1997.
6. J.C.M. Baeten and C. Verhoef. Concrete process algebra. In S. Abramsky, D. Gabbay, and T.S.E. Maibaum, editors, *Handbook of Logic in Computer Science, Volume IV*, pages 149–268. Oxford University Press, 1995.
7. J.C.M. Baeten and W.P. Weijland. *Process Algebra*. Cambridge Tracts in Theoretical Computer Science 18, Cambridge University Press, 1990.
8. F. Belina, D. Hogrefe, and A. Sarma. *SDL with Applications from Protocol Specification*. Prentice-Hall, 1991.
9. J.A. Bergstra. A process creation mechanism in process algebra. In J.C.M. Baeten, editor, *Applications of Process Algebra*, pages 81–88. Cambridge Tracts in Theoretical Computer Science 17, Cambridge University Press, 1990.
10. J.A. Bergstra, W.J. Fokkink, and C.A. Middelburg. Algebra of timed frames. *International Journal of Computer Mathematics*, 61:227–255, 1996.
11. J.A. Bergstra, W.J. Fokkink, and C.A. Middelburg. A logic for signal inserted timed frames. Logic Group Preprint Series 155, Utrecht University, Department of Philosophy, January 1996.

12. J.A. Bergstra and C.A. Middelburg. Process algebra semantics of φ SDL. Logic Group Preprint Series 129, Utrecht University, Department of Philosophy, March 1995.
13. J.A. Bergstra, C.A. Middelburg, and R. Şoricuţ. Discrete time network algebra for a semantic foundation of SDL. Research Report 98, United Nations University, International Institute for Software Technology, October 1997.
14. J.A. Bergstra, C.A. Middelburg, and B. Warinschi. Timed frame models for discrete time process algebras. Research Report 100, United Nations University, International Institute for Software Technology, October 1997.
15. M. Broy. Towards a formal foundation of the specification and description language SDL. *Formal Aspects of Computing*, 3:21–57, 1991.
16. J. Ellsberger, D. Hogrefe, and A. Sarma. *SDL, Formal Object-oriented Language for Communicating Systems*. Prentice-Hall, 1997.
17. J. Fischer, S. Lau, and A. Prinz. A short note about BSDL. *SDL Newsletter*, 18:15–21, 1995.
18. A. Gammelgaard and J.E. Kristensen. A correctness proof of a translation from SDL to CRL. In O. Færgemand and A. Sarma, editors, *SDL '93: Using Objects*, pages 205–219. Elsevier (North-Holland), 1993. Full version: Report TFL RR 1992-4, Tele Danmark Research.
19. J.C. Godskesen. An operational semantics model for Basic SDL (extended abstract). In O. Færgemand and R. Reed, editors, *SDL '91: Evolving Methods*, pages 15–22. Elsevier (North-Holland), 1991. Full version: Report TFL RR 1991-2, Tele Danmark Research.
20. J.F. Groote and A. Ponse. The syntax and semantics of μ CRL. In A. Ponse, C. Verhoef, and S.F.M. van Vlijmen, editors, *Algebra of Communicating Processes 1994*, pages 26–62. Workshop in Computing Series, Springer-Verlag, 1995.
21. C.A. Middelburg. Truth of duration calculus formulae in timed frames. Research Report 82, United Nations University, International Institute for Software Technology, September 1996.
22. S. Mørk, J.C. Godskesen, M.R. Hansen, and R. Sharp. A timed semantics for SDL. In R. Gotzhein and J. Brederke, editors, *FORTE IX: Theory, Application and Tools*, pages 295–309. Chapman & Hall, 1997.
23. X. Nicollin and J. Sifakis. The algebra of timed processes ATP: Theory and application. *Information and Computation*, 114:131–178, 1994.
24. L. Pruitt, 1994. Personal Communications.
25. M.A. Reniers and J.J. Vereijken. Completeness in discrete-time process algebra. Computing Science Report 96-15, Eindhoven University of Technology, Department of Mathematics and Computing Science, 1996.
26. F.W. Vaandrager. *Algebraic Techniques for Concurrency and their Applications*. PhD thesis, Centre for Mathematics and Computer Science, Amsterdam, 1989.
27. Zhou Chaochen, C.A.R. Hoare, and A.P. Ravn. A calculus of durations. *Information Processing Letters*, 40:269–276, 1991.
28. Methods for Testing and Specification (MTS); use of SDL in European telecommunications standards rules for testability and facilitating validation. ETSI Document ETS 300 414, 1995.
29. Specification of Abstract Syntax Notation One (ASN.1). ITU-T Recommendation X.208, 1988.
30. Specification and Description Language (SDL). ITU-T Recommendation Z.100, 1994.
31. Specification and Description Language (SDL) – SDL formal definition: Static semantics. ITU-T Recommendation Z.100 F2, 1994. Annex F.2 to Recommendation Z.100.
32. Specification and Description Language (SDL) – SDL formal definition: Dynamic semantics. ITU-T Recommendation Z.100 F3, 1994. Annex F.3 to Recommendation Z.100.

33. SDL combined with ASN.1 (SDL/ASN.1). ITU-T Recommendation Z.105, 1995.

34. SDT user's guide. TeleLOGIC AB, Sweden, 1992.

A. NOTATIONAL CONVENTIONS

Meta-language for syntax: The syntax of φ -SDL is described by means of production rules in the form of an *extended* BNF grammar. The curly brackets “{” and “}” are used for grouping. The asterisk “*” and the plus sign “+” are used for zero or more repetitions and one or more repetitions, respectively, of curly bracketed groups. The square brackets “[” and “]” are also used for grouping, but indicate that the group is optional. An underlined part included in a nonterminal symbol does not belong to the context-free syntax; it describes a context-sensitive condition.

Meta-language for semantics: The semantics of φ -SDL is described by means of a set of equations recursively defining interpretation functions for all syntactic categories. For each syntactic category, the corresponding interpretation function gives a meaning to each language construct c belonging to the category. We use the notation $\llbracket c \rrbracket$ or $\llbracket c \rrbracket^\kappa$ for applications of all interpretation functions. The exact interpretation function is always clear from the context. If contextual information κ is needed for the interpretation, it is provided by an additional argument and the notation $\llbracket c \rrbracket^\kappa$ is used.

Special action, condition and proposition notation: We write $a : D_1 \times \dots \times D_n$ to indicate that a is an action parametrized by $D_1 \times \dots \times D_n$. This means that there is an action, referred to as $a(d_1, \dots, d_n)$, for each $(d_1, \dots, d_n) \in D_1 \times \dots \times D_n$. For atomic conditions and propositions, we use analogous notations.

Special set, function and sequence notation: We write $\mathcal{P}(A)$ for the set of all subsets of A , and we write $\mathcal{P}_{fin}(A)$ for the set of all finite subsets of A . We use abbreviations \mathbb{N}_1 and \mathbb{N}_2 for the sets $\mathbb{N} \setminus \{0\}$ and $\mathbb{N} \setminus \{0, 1\}$, respectively.

We write $f : A \rightarrow B$ to indicate that f is a total function from A to B , that is $f \subseteq A \times B \wedge \forall x \in A \cdot \exists_1 y \in B \cdot (x, y) \in f$. We write $dom(f)$, where $f : A \rightarrow B$, for A . We also write $A \rightarrow B$ for the set of all functions from A to B . For an (ordered) pair (x, y) , where x and y are intended for argument and value of some function, we use the notation $x \mapsto y$ to emphasize this intention. The binary operators \triangleleft (domain subtraction) and \oplus (overriding) on functions are defined by

$$\begin{aligned} A \triangleleft f &= \{x \mapsto y \mid x \in dom(f) \wedge x \notin A \wedge f(x) = y\} \\ f \oplus g &= (dom(g) \triangleleft f) \cup g \end{aligned}$$

For a function $f : A \rightarrow B$ presenting a family B indexed by A , we use the notation f_i (for $i \in A$) instead of $f(i)$.

Functions are also used to present sequences; as usual we write $\langle x_1, \dots, x_n \rangle$ for the sequence presented by the function $\{1 \mapsto x_1, \dots, n \mapsto x_n\}$. The unary operators *hd* and *tl* stand for selection of head and tail, respectively, of sequences. The binary operator $\hat{\ }^{\ } stands for concatenation of sequences. We write $x \& t$ for $\langle x \rangle \hat{\ } t$.$

B. CONTEXTUAL INFORMATION

The meaning of a language construct of φ -SDL generally depends on the definitions in the scope in which it occurs. Contexts are primarily intended for modeling the scope. The context that is ascribed to a complete φ -SDL specification is also used to define the state space used to describe its meaning. The context of a language construct contains all names introduced by the definitions of variables, signal types, signal routes and process types occurring in the specification on hand and additionally:

- if the language construct occurs in the scope of a process definition, the name introduced by that process definition, called the *scope unit*;

- if the language construct occurs in the scope of a state definition, the set of names occurring in the **save** part of that state definition, called the *save set*.

The names introduced by the definitions are in addition connected with their static attributes. For example, a name of a variable is connected with the name of the sort of the values that may be assigned to it; and a name of a process type is connected with the names of the variables that are its formal parameters and the number of processes of this type that have to be created during the start-up of the system.

$$\text{Context} = \mathcal{P}_{\text{fin}}(\text{VarD}) \times \mathcal{P}_{\text{fin}}(\text{SigD}) \times \mathcal{P}_{\text{fin}}(\text{RouteD}) \times \mathcal{P}_{\text{fin}}(\text{ProcD}) \times (\text{ProcId} \cup \{\text{nil}\}) \times \mathcal{P}_{\text{fin}}(\text{SigId})$$

where

$$\begin{aligned} \text{VarD} &= \text{VarId} \times \text{SortId} \\ \text{SigD} &= \text{SigId} \times \text{SortId}^* \\ \text{RouteD} &= \text{RouteId} \times (\text{ProcId} \cup \{\mathbf{env}\}) \times (\text{ProcId} \cup \{\mathbf{env}\}) \times \mathcal{P}_{\text{fin}}(\text{SigId}) \\ \text{ProcD} &= \text{ProcId} \times \text{VarId}^* \times \mathbb{N} \end{aligned}$$

For language constructs that do not occur in a process definition, the absence of a scope unit will be represented by nil and, for language constructs that do not occur in a state definition, the absence of a save set will be represented by \emptyset . We write $\text{varids}(\kappa)$, $\text{sigds}(\kappa)$, $\text{routed}(\kappa)$, $\text{procds}(\kappa)$, $\text{scopeunit}(\kappa)$ and $\text{saveset}(\kappa)$, where $\kappa = (V, S, R, P, X, ss) \in \text{Context}$, for V, S, R, P, X and ss , respectively. We write $\text{vars}(\kappa)$ for $\{v \mid \exists T. (v, T) \in \text{varids}(\kappa)\}$. The abbreviations $\text{sigs}(\kappa)$ and $\text{procs}(\kappa)$ are used analogously.

We make use of the following functions on *Context*:

$$\begin{aligned} \text{rcv} &: \text{Context} \times \text{RouteId} \rightarrow \text{ProcId} \cup \{\mathbf{env}\}, \\ \text{fpars} &: \text{Context} \times \text{ProcId} \rightarrow \text{VarId}^*, \\ \text{init} &: \text{Context} \times \text{ProcId} \rightarrow \mathbb{N}, \\ \text{updscopeunit} &: \text{Context} \times \text{ProcId} \rightarrow \text{Context}, \\ \text{updsaveset} &: \text{Context} \times \mathcal{P}_{\text{fin}}(\text{SigId}) \rightarrow \text{Context}, \\ \text{envsigd} &: \text{Context} \rightarrow \mathcal{P}_{\text{fin}}(\text{SigD} \times \text{ProcId}) \end{aligned}$$

The first four functions are partial functions, but they will only be applied in cases where the result is defined. The function rcv is used to extract the receiver type of a given signal route from the context. This function is inductively defined by

$$(r, X_1, X_2, ss) \in \text{routed}(\kappa) \Rightarrow \text{rcv}(\kappa, r) = X_2$$

The functions fpars and init are used to extract the formal parameters and the initial number of processes, respectively, of a given process type from the context. These functions are inductively defined by

$$\begin{aligned} (X, vs, k) \in \text{procds}(\kappa) &\Rightarrow \text{fpars}(\kappa, X) = vs, \\ (X, vs, k) \in \text{procds}(\kappa) &\Rightarrow \text{init}(\kappa, X) = k \end{aligned}$$

The functions updscopeunit and updsaveset are used to update the scope unit and the save set, respectively, of the context. These functions are inductively defined by

$$\begin{aligned} \kappa = (V, S, R, P, X, ss) &\Rightarrow \text{updscopeunit}(\kappa, X') = (V, S, R, P, X', ss), \\ \kappa = (V, S, R, P, X, ss) &\Rightarrow \text{updsaveset}(\kappa, ss') = (V, S, R, P, X, ss') \end{aligned}$$

The function envsigd is used to determine the possible environment signals, i.e. signals that the system can receive via signal routes from the environment. It is inductively defined by

$$\begin{aligned} s \in ss \wedge (s, \langle T_1, \dots, T_n \rangle) \in \text{sigds}(\kappa) \wedge (r, \mathbf{env}, X_2, ss) \in \text{routed}(\kappa) &\Rightarrow \\ ((s, \langle T_1, \dots, T_n \rangle), X_2) \in \text{envsigd}(\kappa) & \end{aligned}$$

The context ascribed to a complete φ -SDL specification is a minimal context in the sense that the contextual information available in it is common to all contexts on which language constructs occurring in it depend. The additional information that may be available applies to the scope unit for language constructs occurring in a process definition and the save set for language constructs occurring in a state definition. The context ascribed to a complete specification is obtained by taking the union of the corresponding components of the partial contexts contributed by all definitions occurring in it, except for the scope unit and the save set which are permanently the same – nil and \emptyset , respectively.

$$\begin{aligned} \{\{\mathbf{system} \ S; D_1 \dots D_n \ \mathbf{endsystem};\}\} &:= \\ &(\mathit{vars}(\{\{D_1\}\}) \cup \dots \cup \mathit{vars}(\{\{D_n\}\}), \\ &\mathit{sigds}(\{\{D_1\}\}) \cup \dots \cup \mathit{sigds}(\{\{D_n\}\}), \\ &\mathit{routed}(\{\{D_1\}\}) \cup \dots \cup \mathit{routed}(\{\{D_n\}\}), \\ &\mathit{procds}(\{\{D_1\}\}) \cup \dots \cup \mathit{procds}(\{\{D_n\}\}), \\ &\text{nil}, \emptyset) \end{aligned}$$

$$\{\{\mathbf{dcl} \ v \ T;\}\} := (\{(v, T)\}, \emptyset, \emptyset, \emptyset, \text{nil}, \emptyset)$$

$$\{\{\mathbf{signal} \ s(T_1, \dots, T_n);\}\} := (\emptyset, \{(s, \langle T_1, \dots, T_n \rangle)\}, \emptyset, \emptyset, \text{nil}, \emptyset)$$

$$\begin{aligned} \{\{\mathbf{signalroute} \ r \ \mathbf{from} \ X_1 \ \mathbf{to} \ X_2 \ \mathbf{with} \ s_1, \dots, s_n;\}\} &:= \\ &(\emptyset, \emptyset, \{(r, X_1, X_2, \{s_1, \dots, s_n\})\}, \emptyset, \text{nil}, \emptyset) \end{aligned}$$

$$\begin{aligned} \{\{\mathbf{process} \ X(k); \ \mathbf{fpar} \ v_1, \dots, v_m; \ \mathbf{start}; \ \mathit{tr} \ d_1 \dots d_n \ \mathbf{endprocess};\}\} &:= \\ &(\emptyset, \emptyset, \emptyset, \{(X, \langle v_1, \dots, v_m \rangle, k)\}, \text{nil}, \emptyset) \end{aligned}$$