

On the Behaviours Produced by Instruction Sequences under Execution

J.A. Bergstra, C.A. Middelburg*

Informatics Institute, University of Amsterdam

Science Park 904, 1098 XH Amsterdam, the Netherlands

{J.A.Bergstra,C.A.Middelburg}@uva.nl

Abstract. We study several aspects of the behaviours produced by instruction sequences under execution in the setting of the algebraic theory of processes known as ACP. We use ACP to describe the behaviours produced by instruction sequences under execution and to describe two protocols implementing these behaviours in the case where the processing of instructions takes place remotely. We also show that all finite-state behaviours considered in ACP can be produced by instruction sequences under execution.

Keywords: instruction sequence, remote instruction processing, instruction sequence producible process

1. Introduction

The concept of an instruction sequence is a very primitive concept in computing. It has always played a central role in computing because of the fact that execution of instruction sequences underlies virtually all past and current generations of computers. It happens that, given a precise definition of an appropriate notion of an instruction sequence, many issues in computer science can be clearly explained in terms of instruction sequences. A simple yet interesting example is that a program can simply be defined as a text that denotes an instruction sequence. Such a definition corresponds to an empirical perspective found among practitioners.

*Address for correspondence: Informatics Institute, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, the Netherlands

In theoretical computer science, the meaning of programs usually plays a prominent part in the explanation of many issues concerning programs. Moreover, what is taken for the meaning of programs is mathematical by nature. On the other hand, it is customary that practitioners do not fall back on the mathematical meaning of programs in case explanation of issues concerning programs is needed. They phrase their explanations from an empirical perspective. An empirical perspective that we consider appealing is the perspective that a program is in essence an instruction sequence and an instruction sequence under execution produces a behaviour that is controlled by its execution environment in the sense that each step performed actuates the processing of an instruction by the execution environment and a reply returned at completion of the processing determines how the behaviour proceeds.

This paper concerns the behaviours produced by instruction sequences under execution as such and two issues relating to the behaviours produced by instruction sequences under execution, namely the issue of implementing these behaviours in the case where the processing of instructions takes place remotely and the issue of the extent to which the behaviours considered in process algebra can be produced by instruction sequences under execution.

Remote instruction processing means that a stream of instructions to be processed arises at one place and the processing of that stream of instructions is handled at another place. This phenomenon is increasingly encountered. It is found if loading the instruction sequence to be executed as a whole is impracticable. For instance, the storage capacity of the execution unit is too small or the execution unit is too far away. Remote instruction processing requires special attention because the transmission time of the messages involved in remote instruction processing makes it hard to keep the execution unit busy without intermission.

In the literature on computer architecture, hardly anything can be found that contributes to a sound understanding of the phenomenon of remote instruction processing. As a first step towards such an understanding, we give rigorous descriptions of two protocols for remote instruction processing at a level of abstraction that captures the underlying essence of the protocols. One protocol is very simple, but makes no effort keep the execution unit busy without intermission. The other protocol is more complex and is directed towards keeping the execution unit busy without intermission. It is reminiscent of an instruction pre-fetching mechanism as found in pipelined processors (see e.g. [26]), but its range of application is not restricted to pipelined instruction processing.

The work presented in this paper belongs to a line of research which started with an attempt to approach the semantics of programming languages from the perspective mentioned above. The first published paper on this approach is [7]. That paper is superseded by [8] with regard to the groundwork for the approach: program algebra, an algebraic theory of single-pass instruction sequences, and basic thread algebra, an algebraic theory of mathematical objects that represent in a direct way the behaviours produced by instruction sequences under execution.¹ The main advantages of the approach are that it does not require a lot of mathematical background and that it is more appealing to practitioners than the main approaches to programming language semantics: the operational approach, the denotational approach and the axiomatic approach. For an overview of these approaches, see e.g. [32].

The work presented in this paper is based on the instruction sequences considered in program algebra and the representation of the behaviours produced by instruction sequences under execution considered in basic thread algebra. It is rather awkward to describe and analyse the behaviours of this kind using algebraic theories of processes such as ACP [3, 6], CCS [27, 31] and CSP [23, 29]. However, the objects

¹In [8], basic thread algebra is introduced under the name basic polarized process algebra.

considered in basic thread algebra can be viewed as representations of processes as considered in process algebra. This allows for the protocols for remote instruction processing to be described using ACP or rather ACP^τ , an extension of ACP which supports abstraction from internal actions.

Process algebra is an area of the study of concurrency which is considered relevant to computer science, as is witnessed by the extent of the work on algebraic theories of processes such as ACP, CCS and CSP in theoretical computer science. This strongly hints that there are programmed systems whose behaviours can be taken for processes as considered in process algebra. Therefore, it is interesting to know to which extent the behaviours considered in process algebra can be produced by programs under execution, starting from the perception of a program as an instruction sequence. In this paper, we will show that, by apposite choice of instructions, all finite-state processes can be produced by instruction sequences (provided that the cluster fair abstraction rule, see e.g. Section 5.6 of [24], is valid).

The instruction sequences considered in program algebra are single-pass instruction sequences, i.e. finite or infinite sequences of instructions of which each instruction is executed at most once and can be dropped after it has been executed or jumped over. Program algebra does not provide a notation for programs that is intended for actual programming: programs written in an assembly language are finite instruction sequences for which single-pass execution is usually not possible. We will also show that all finite-state processes can as well be produced by programs written in a program notation which is close to existing assembly languages.

Instruction sequences under execution may make use of services provided by their execution environment such as counters, stacks and Turing tapes. The use operators added to basic thread algebra in e.g. [12] can be used to describe the behaviours produced by instruction sequences under execution that make use of services. Interesting is that instruction sequences under execution that make use of services may produce infinite-state processes. On that account, we will make precise what processes are produced by instruction sequences under execution that make use of services provided by their execution environment.

As a continuation of the work on a new approach to programming language semantics mentioned above, the notion of an instruction sequence was subjected to systematic and precise analysis using the groundwork laid earlier. This led among other things to expressiveness results about the instruction sequences considered and variations of the instruction sequences considered (see e.g. [12, 18, 20, 21, 36]). Instruction sequences are under discussion for many years in diverse work on computer architecture, as witnessed by e.g. [4, 22, 25, 30, 33, 34, 35, 39, 41], but the notion of an instruction sequence has never been subjected to any precise analysis before. As another continuation of the work on a new approach to programming language semantics mentioned above, selected issues relating to well-known subjects from the theory of computation and the area of computer architecture were rigorously investigated thinking in terms of instruction sequences (see e.g. [14, 15, 16, 17, 19]). The subjects from the theory of computation, namely the halting problem and non-uniform computational complexity, are usually investigated thinking in terms of a common model of computation such as Turing machines and Boolean circuits (see e.g. [1, 28, 38]). The subjects from the area of computer architecture, namely instruction sequence performance, instruction set architectures and remote instruction processing, are usually not investigated in a rigorous way at all. The general aim of the work in both continuations mentioned is to bring instruction sequences as a theme in computer science better into the picture. The work presented in this paper forms a part of the last mentioned continuation.

This paper is organized as follows. The body of the paper consists of three parts. The first part (Sections 2–7) concerns the behaviours produced by instruction sequences under execution as such and

includes surveys of program algebra, basic thread algebra and the algebraic theory of processes known as ACP. The second part (Sections 8–10) concerns the issue of implementing these behaviours in the case where the processing of instructions takes place remotely and includes rigorous descriptions of two protocols for remote instruction processing. The third part (Sections 11–14) concerns the issue of the extent to which the behaviours considered in process algebra can be produced by instruction sequences under execution and includes the result that, by apposite choice of instructions, all finite-state processes can be produced by instruction sequences.

This paper consolidates material from [11, 13, 14].

2. Program Algebra

In this section, we review PGA (ProGram Algebra). The starting-point of program algebra is the perception of a program as a single-pass instruction sequence. The concepts underlying the primitives of program algebra are common in programming, but the particular form of the primitives is not common. The predominant concern in the design of program algebra has been to achieve simple syntax and semantics, while maintaining the expressive power of arbitrary finite control.

In PGA, it is assumed that a fixed but arbitrary set \mathcal{A} of *basic instructions* has been given. The intuition is that the execution of a basic instruction may modify a state and produces a reply at its completion. The possible replies are the Boolean values T and F.

PGA has the following *primitive instructions*:

- for each $a \in \mathcal{A}$, a *plain basic instruction* a ;
- for each $a \in \mathcal{A}$, a *positive test instruction* $+a$;
- for each $a \in \mathcal{A}$, a *negative test instruction* $-a$;
- for each $l \in \mathbb{N}$, a *forward jump instruction* $\#l$;
- a *termination instruction* $!$.

We write \mathcal{I} for the set of all primitive instructions of PGA. On execution of an instruction sequence, these primitive instructions have the following effects:

- the effect of a positive test instruction $+a$ is that basic instruction a is executed and execution proceeds with the next primitive instruction if T is produced and otherwise the next primitive instruction is skipped and execution proceeds with the primitive instruction following the skipped one — if there is no primitive instructions to proceed with, inaction occurs;
- the effect of a negative test instruction $-a$ is the same as the effect of $+a$, but with the role of the value produced reversed;
- the effect of a plain basic instruction a is the same as the effect of $+a$, but execution always proceeds as if T is produced;
- the effect of a forward jump instruction $\#l$ is that execution proceeds with the l -th next instruction of the program concerned — if l equals 0 or there is no primitive instructions to proceed with, inaction occurs;

Table 1. Axioms of PGA

$(x ; y) ; z = x ; (y ; z)$	PGA1
$(x^n)^\omega = x^\omega$	PGA2
$x^\omega ; y = x^\omega$	PGA3
$(x ; y)^\omega = x ; (y ; x)^\omega$	PGA4

- the effect of the termination instruction ! is that execution terminates.

PGA has the following constants and operators:

- for each $u \in \mathcal{I}$, an *instruction* constant u ;
- the binary *concatenation* operator ; ;
- the unary *repetition* operator $^\omega$.

We assume that there is a countably infinite set of variables which includes x, y, z . Terms are built as usual. We use infix notation for concatenation and postfix notation for repetition.

A closed PGA term is considered to denote a non-empty, finite or eventually periodic infinite sequence of primitive instructions.² The instruction sequence denoted by a closed term of the form $t ; t'$ is the instruction sequence denoted by t concatenated with the instruction sequence denoted by t' . The instruction sequence denoted by a closed term of the form t^ω is the instruction sequence denoted by t concatenated infinitely many times with itself. Some simple examples of closed PGA terms are

$$a ; b ; c , \quad +a ; \#2 ; \#3 ; b ; ! , \quad a ; (b ; c)^\omega .$$

On execution of the instruction sequence denoted by the first term, the basic instructions a, b and c are executed in that order and after that inaction occurs. On execution of the instruction sequence denoted by the second term, the basic instruction a is executed first, if the execution of a produces the reply T, the basic instruction b is executed next and after that execution terminates, and if the execution of a produces the reply F, inaction occurs. On execution of the instruction sequence denoted by the third term, the basic instruction a is executed first, and after that the basic instructions b and c are executed in that order repeatedly forever.

Closed PGA terms are considered equal if they represent the same instruction sequence. The axioms for instruction sequence equivalence are given in Table 1. In this table, n stands for an arbitrary positive natural number. The term t^n , where t is a PGA term, is defined by induction on n as follows: $t^1 = t$ and $t^{n+1} = t ; t^n$. The *unfolding* equation $x^\omega = x ; x^\omega$ is derived as follows:

$$\begin{aligned} x^\omega &= (x ; x)^\omega && \text{by PGA2} \\ &= x ; (x ; x)^\omega && \text{by PGA4} \\ &= x ; x^\omega && \text{by PGA2 .} \end{aligned}$$

Each closed PGA term is derivably equal to a term in *canonical form*, i.e. a term of the form t or $t ; t'^\omega$, where t and t' are closed PGA terms in which the repetition operator does not occur. For example:

$$\begin{aligned} (a ; b)^\omega ; c ; ! &= a ; (b ; a)^\omega , \\ +a ; (\#4 ; b ; (-c ; \#5 ; !)^\omega)^\omega &= +a ; \#4 ; b ; (-c ; \#5 ; !)^\omega . \end{aligned}$$

²An eventually periodic infinite sequence is an infinite sequence with only finitely many distinct suffixes.

The initial models of PGA are considered its standard models. Henceforth, we restrict ourselves to the initial model \mathcal{I}_{PGA} of PGA in which:

- the domain is the set of all non-empty, finite and eventually periodic infinite sequences over the set \mathcal{I} of primitive instructions;
- the operation associated with \cdot is concatenation;
- the operation associated with $^\omega$ is the operation $^\omega$ defined as follows:
 - if F is a finite sequence over \mathcal{I} , then F^ω is the unique eventually periodic infinite sequence F' such that F concatenated n times with itself is a proper prefix of F' for each $n \in \mathbb{N}$;
 - if F is an eventually periodic infinite sequence over \mathcal{I} , then F^ω is F .

In the sequel, we use the term *instruction sequence* for the elements of the domain of \mathcal{I}_{PGA} , and we denote the interpretations of the constants and operators of PGA in \mathcal{I}_{PGA} by the constants and operators themselves. \mathcal{I}_{PGA} is loosely called *the* initial model of PGA because all initial models of PGA are isomorphic, i.e. there exist bijective homomorphism between them (see e.g. [37, 40]).

3. Basic Thread Algebra

In this section, we review BTA (Basic Thread Algebra). BTA is an algebraic theory of mathematical objects that represent in a direct way the behaviours produced by instruction sequences under execution. The objects concerned are called threads.

In BTA, it is assumed that a fixed but arbitrary set \mathcal{A} of *basic actions*, with $\text{tau} \notin \mathcal{A}$, has been given. Besides, tau is a special basic action. We write \mathcal{A}_{tau} for $\mathcal{A} \cup \{\text{tau}\}$. A thread performs basic actions in a sequential fashion. Upon each basic action performed, a reply from an execution environment determines how it proceeds. The possible replies are the Boolean values T and F. Performing tau , which is considered performing an internal action, always leads to the reply T.

Although BTA is one-sorted, we make this sort explicit. The reason for this is that we will extend BTA with an additional sort in Section 13.

BTA has one sort: the sort \mathbf{T} of *threads*. To build terms of sort \mathbf{T} , it has the following constants and operators:

- the *inaction* constant $D : \mathbf{T}$;
- the *termination* constant $S : \mathbf{T}$;
- for each $a \in \mathcal{A}_{\text{tau}}$, the binary *postconditional composition* operator $\triangleleft a \triangleright : \mathbf{T} \times \mathbf{T} \rightarrow \mathbf{T}$.

We assume that there are infinitely many variables of sort \mathbf{T} , including x, y, z . Terms of sort \mathbf{T} are built as usual. We use infix notation for the postconditional composition operators. We introduce *basic action prefixing* as an abbreviation: $a \circ t$, where $a \in \mathcal{A}_{\text{tau}}$ and t is a term of sort \mathbf{T} , abbreviates $t \triangleleft a \triangleright t$.

The thread denoted by a closed term of the form $t \triangleleft a \triangleright t'$ will first perform a , and then proceed as the thread denoted by t if the reply from the execution environment is T and proceed as the thread

Table 2. Axiom of BTA

$$\underline{x \triangleleft \text{tau} \triangleright y = x \triangleleft \text{tau} \triangleright x} \quad \text{T1}$$

Table 3. RDP, RSP and AIP

$\langle X E \rangle = \langle t_X E \rangle$ if $X = t_X \in E$	RDP	$\pi_0(x) = \text{D}$	P0
$E \Rightarrow X = \langle X E \rangle$ if $X \in V(E)$	RSP	$\pi_{n+1}(\text{S}) = \text{S}$	P1
		$\pi_{n+1}(\text{D}) = \text{D}$	P2
$\bigwedge_{n \geq 0} \pi_n(x) = \pi_n(y) \Rightarrow x = y$	AIP	$\pi_{n+1}(x \triangleleft a \triangleright y) = \pi_n(x) \triangleleft a \triangleright \pi_n(y)$	P3

denoted by t' if the reply from the execution environment is F. The threads denoted by D and S will become inactive and terminate, respectively. Some simple examples of closed BTA terms are

$$a \circ (\text{S} \triangleleft b \triangleright \text{D}), \quad (b \circ \text{S}) \triangleleft a \triangleright \text{D}.$$

The first term denotes the thread that first performs basic action a , next performs basic action b , if the reply from the execution environment on performing b is T, after that terminates, and if the reply from the execution environment on performing b is F, after that becomes inactive. The second term denotes the thread that first performs basic action a , if the reply from the execution environment on performing a is T, next performs the basic action b and after that terminates, and if the reply from the execution environment on performing a is F, next becomes inactive.

BTA has only one axiom. This axiom is given in Table 2. Using the abbreviation introduced above, axiom T1 can be written as follows: $x \triangleleft \text{tau} \triangleright y = \text{tau} \circ x$.

Notice that each closed BTA term denotes a thread that will become inactive or terminate after it has performed finitely many actions. Infinite threads can be described by guarded recursion.

A *guarded recursive specification* over BTA is a set of recursion equations $E = \{X = t_X \mid X \in V\}$, where V is a set of variables of sort **T** and each t_X is a BTA term of the form D , S or $t \triangleleft a \triangleright t'$ with t and t' that contain only variables from V . We write $V(E)$ for the set of all variables that occur in E . We are only interested in models of BTA in which guarded recursive specifications have unique solutions, such as the projective limit model of BTA presented in [5].

A simple example of a guarded recursive specification is the one consisting of following two equations:

$$x = x \triangleleft a \triangleright y, \quad y = y \triangleleft b \triangleright \text{S}.$$

The x -component of the solution of this guarded recursive specification is the thread that first performs basic action a repeatedly until the reply from the execution environment on performing a is F, next performs basic action b repeatedly until the reply from the execution environment on performing b is F, and after that terminates.

For each guarded recursive specification E and each $X \in V(E)$, we introduce a constant $\langle X|E \rangle$ of sort **T** standing for the X -component of the unique solution of E . We write $\langle t_X|E \rangle$ for t_X with, for all $Y \in V(E)$, all occurrences of Y in t_X replaced by $\langle Y|E \rangle$. The axioms for the constants for the components of the solutions of guarded recursive specifications are RDP (Recursive Definition Principle) and RSP (Recursive Specification Principle), which are given in Table 3. RDP and RSP are actually axiom schemas in which X stands for an arbitrary variable, t_X stands for an arbitrary BTA term, and E

stands for an arbitrary guarded recursive specification over BTA. Side conditions are added to restrict what X , t_X and E stand for. The equations $\langle X|E \rangle = \langle t_X|E \rangle$ for a fixed E express that the constants $\langle X|E \rangle$ make up a solution of E . The conditional equations $E \Rightarrow X = \langle X|E \rangle$ express that this solution is the only one.

RDP and RSP are means to prove closed terms that denote the same infinite thread equal. We introduce AIP (Approximation Induction Principle) as an additional means to prove closed terms that denote the same infinite thread equal. AIP is based on the view that two threads are identical if their approximations up to any finite depth are identical. The approximation up to depth n of a thread is obtained by cutting it off after it has performed n actions. AIP is also given in Table 3. Here, approximation up to depth n is phrased in terms of the unary *projection* operator $\pi_n : \mathbf{T} \rightarrow \mathbf{T}$. The axioms for the projection operators are axioms P0–P3 in Table 3. P1–P3 are actually axiom schemas in which a stands for arbitrary basic action and n stands for an arbitrary natural number.

We write BTA+REC for BTA extended with the constants for the components of the solutions of guarded recursive specifications, the projection operators and the axioms RDP, RSP, AIP and P0–P3.

The minimal models of BTA+REC are considered its standard models.³ Recall that a model of an algebraic theory is minimal iff all elements of the domains associated with the sorts of the theory can be denoted by closed terms. Henceforth, we restrict ourselves to the minimal models of BTA+REC. We assume that a minimal model $\mathcal{M}_{\text{BTA+REC}}$ of BTA+REC has been given.

In the sequel, we use the term *thread* for the elements of the domain of $\mathcal{M}_{\text{BTA+REC}}$, and we denote the interpretations of constants and operators in $\mathcal{M}_{\text{BTA+REC}}$ by the constants and operators themselves.

Let T be a thread. Then the set of *states* or *residual threads* of T , written $\text{Res}(T)$, is inductively defined as follows:

- $T \in \text{Res}(T)$;
- if $T' \trianglelefteq a \triangleright T'' \in \text{Res}(T)$, then $T' \in \text{Res}(T)$ and $T'' \in \text{Res}(T)$.

Let T be a thread and let $\mathcal{A}' \subseteq \mathcal{A}_{\text{tau}}$. Then T is *regular over* \mathcal{A}' if the following conditions are satisfied:

- $\text{Res}(T)$ is finite;
- for all $T', T'' \in \text{Res}(T)$ and $a \in \mathcal{A}_{\text{tau}}$, $T' \trianglelefteq a \triangleright T'' \in \text{Res}(T)$ implies $a \in \mathcal{A}'$.

We say that T is *regular* if T is regular over \mathcal{A}_{tau} .

For example, the x -component of the solution of the guarded recursive specification consisting of the following two equations:

$$x = a \circ y, \quad y = (c \circ y) \trianglelefteq b \triangleright (x \trianglelefteq d \triangleright S)$$

has five states and is regular over any $\mathcal{A}' \subseteq \mathcal{A}_{\text{tau}}$ for which $\{a, b, c, d\} \subseteq \mathcal{A}'$.

We will make use of the fact that being a regular thread coincides with being a component of the solution of a finite guarded recursive specification in which the right-hand sides of the recursion equations are of a restricted form.

A *linear recursive specification* over BTA is a guarded recursive specification $E = \{X = t_X \mid X \in V\}$ over BTA, where each t_X is a term of the form D, S or $Y \trianglelefteq a \triangleright Z$ with $Y, Z \in V$.

³A minimal model of an algebraic theory is a model of which no proper subalgebra is a model as well.

Table 4. Defining equations for thread extraction operation

$ a = a \circ D$	$ \#l = D$
$ a ; F = a \circ F $	$ \#0 ; F = D$
$ +a = a \circ D$	$ \#1 ; F = F $
$ +a ; F = F \trianglelefteq a \triangleright \#2 ; F $	$ \#l + 2 ; u = D$
$ -a = a \circ D$	$ \#l + 2 ; u ; F = \#l + 1 ; F $
$ -a ; F = \#2 ; F \trianglelefteq a \triangleright F $	$ \! = S$
	$ \! ; F = S$

Proposition 3.1. Let T be a thread and let $\mathcal{A}' \subseteq \mathcal{A}_{\text{tau}}$. Then T is regular over \mathcal{A}' iff there exists a finite linear recursive specification E over BTA in which only basic actions from \mathcal{A}' occur such that T is a component of the solution of E .

Proof:

The implication from left to right is proved as follows. Because T is regular, $\text{Res}(T)$ is finite. Hence, there are finitely many threads T_1, \dots, T_n , with $T = T_1$, such that $\text{Res}(T) = \{T_1, \dots, T_n\}$. Now T is the x_1 -component of the solution of the linear recursive specification consisting of the following equations:

$$x_i = \begin{cases} S & \text{if } T_i = S \\ D & \text{if } T_i = D \\ x_j \trianglelefteq a \triangleright x_k & \text{if } T_i = T_j \trianglelefteq a \triangleright T_k \end{cases} \quad \text{for all } i \in [1, n].$$

Because T is regular over \mathcal{A}' , only basic actions from \mathcal{A}' occur in the linear recursive specification constructed in this way.

The implication from right to left is proved as follows. Thread T is a component of the unique solution of a finite linear specification in which only basic actions from \mathcal{A}' occur. This means that there are finitely many threads T_1, \dots, T_n , with $T = T_1$, such that for every $i \in [1, n]$, $T_i = S$, $T_i = D$ or $T_i = T_j \trianglelefteq a \triangleright T_k$ for some $j, k \in [1, n]$ and $a \in \mathcal{A}'$. Consequently, $T' \in \text{Res}(T)$ iff $T' = T_i$ for some $i \in [1, n]$ and moreover $T' \trianglelefteq a \triangleright T'' \in \text{Res}(T)$ only if $a \in \mathcal{A}'$. Hence, $\text{Res}(T)$ is finite and T is regular over \mathcal{A}' . \square

Remark 3.2. A structural operational semantics of BTA+REC and a bisimulation equivalence based on it can be found in e.g. [10]. The quotient algebra of the algebra of closed terms of BTA+REC by this bisimulation equivalence is one of the minimal models of BTA+REC.

4. Thread Extraction

In this short section, we use BTA+REC to make mathematically precise which threads are produced by instruction sequences under execution.

For that purpose, \mathcal{A} is taken such that $\mathcal{A} \supseteq \mathfrak{A}$ is satisfied.

The *thread extraction* operation $|_|_$ assigns a thread to each instruction sequence. The thread extraction operation is defined by the equations given in Table 4 (for $a \in \mathfrak{A}$, $l \in \mathbb{N}$, and $u \in \mathfrak{J}$) and the rule that $|\#l ; F| = D$ if $\#l$ is the beginning of an infinite jump chain. This rule is formalized in e.g. [12].

Let F be an instruction sequence and T be a thread. Then we say that F produces T if $|F| = T$. For example,

$$\begin{array}{ll} a ; b ; c & \text{produces } a \circ b \circ c \circ D , \\ +a ; \#2 ; \#3 ; b ; ! & \text{produces } (b \circ S) \trianglelefteq a \triangleright D , \\ +a ; -b ; c ; ! & \text{produces } (S \trianglelefteq b \triangleright (c \circ S)) \trianglelefteq a \triangleright (c \circ S) , \\ +a ; \#2 ; (b ; \#2 ; c ; \#2)^\omega & \text{produces } D \trianglelefteq a \triangleright (b \circ D) . \end{array}$$

In the case of instruction sequences that are not finite, the produced threads can be described as components of the solution of a guarded recursive specification. For example, the infinite instruction sequence

$$(a ; +b)^\omega$$

produces the x -component of the solution of the guarded recursive specification consisting of following two equations:

$$x = a \circ y , \quad y = x \trianglelefteq b \triangleright y$$

and the infinite instruction sequence

$$a ; (+b ; \#2 ; \#3 ; c ; \#4 ; -d ; ! ; a)^\omega$$

produces the x -component of the solution of the guarded recursive specification consisting of following two equations:

$$x = a \circ y , \quad y = (c \circ y) \trianglelefteq b \triangleright (x \trianglelefteq d \triangleright S) .$$

5. Algebra of Communicating Processes

In this section, we review ACP^τ (Algebra of Communicating Processes with abstraction). This algebraic theory of processes will among other things be used to make precise what processes are produced by the threads denoted by closed terms of BTA+REC. For a comprehensive overview of ACP^τ , the reader is referred to [3, 24].

In ACP^τ , it is assumed that a fixed but arbitrary set A of *atomic actions*, with $\tau, \delta \notin A$, and a fixed but arbitrary commutative and associative function $| : A \cup \{\tau\} \times A \cup \{\tau\} \rightarrow A \cup \{\delta\}$, with $\tau | e = \delta$ for all $e \in A \cup \{\tau\}$, have been given. The function $|$ is regarded to give the result of synchronously performing any two atomic actions for which this is possible, and to give δ otherwise. In ACP^τ , τ is a special atomic action, called the *silent step*. The act of performing the silent step is considered unobservable. Because it would otherwise be observable, the silent step is considered an atomic action that cannot be performed synchronously with other atomic actions. We write A_τ for $A \cup \{\tau\}$.

ACP^τ has the following constants and operators:

- for each $e \in A$, the *atomic action* constant e ;
- the *silent step* constant τ ;
- the *inaction* constant δ ;
- the binary *alternative composition* operator $+$;

- the binary *sequential composition* operator \cdot ;
- the binary *parallel composition* operator \parallel ;
- the binary *left merge* operator \ll ;
- the binary *communication merge* operator $|$;
- for each $H \subseteq A$, the unary *encapsulation* operator ∂_H ;
- for each $I \subseteq A$, the unary *abstraction* operator τ_I .

We assume that there are infinitely many variables, including x, y, z . Terms are built as usual. We use infix notation for the binary operators. The precedence conventions used with respect to the operators of ACP^τ are as follows: $+$ binds weaker than all others, \cdot binds stronger than all others, and the remaining operators bind equally strong.

Let t and t' be closed ACP^τ terms, $e \in A$, and $H, I \subseteq A$. Intuitively, the constants and operators to build ACP^τ terms can be explained as follows:

- the process denoted by e first performs atomic action e and next terminates successfully;
- the process denoted by τ performs an unobservable atomic action and next terminates successfully;
- the process denoted by δ can neither perform an atomic action nor terminate successfully;
- the process denoted by $t + t'$ behaves either as the process denoted by t or as the process denoted by t' , but not both;
- the process denoted by $t \cdot t'$ first behaves as the process denoted by t and on successful termination of that process it next behaves as the process denoted by t' ;
- the process denoted by $t \parallel t'$ behaves as the process that proceeds with the processes denoted by t and t' in parallel;
- the process denoted by $t \ll t'$ behaves the same as the process denoted by $t \parallel t'$, except that it starts with performing an atomic action of the process denoted by t ;
- the process denoted by $t | t'$ behaves the same as the process denoted by $t \parallel t'$, except that it starts with performing an atomic action of the process denoted by t and an atomic action of the process denoted by t' synchronously;
- the process denoted by $\partial_H(t)$ behaves the same as the process denoted by t , except that atomic actions from H are blocked;
- the process denoted by $\tau_I(t)$ behaves the same as the process denoted by t , except that atomic actions from I are turned into unobservable atomic actions.

Table 5. Axioms of ACP^τ

$x + y = y + x$	A1	$x \cdot \tau = x$	B1
$(x + y) + z = x + (y + z)$	A2	$x \cdot (\tau \cdot (y + z) + y) = x \cdot (y + z)$	B2
$x + x = x$	A3		
$(x + y) \cdot z = x \cdot z + y \cdot z$	A4	$\partial_H(a) = a$	if $a \notin H$ D1
$(x \cdot y) \cdot z = x \cdot (y \cdot z)$	A5	$\partial_H(a) = \delta$	if $a \in H$ D2
$x + \delta = x$	A6	$\partial_H(x + y) = \partial_H(x) + \partial_H(y)$	D3
$\delta \cdot x = \delta$	A7	$\partial_H(x \cdot y) = \partial_H(x) \cdot \partial_H(y)$	D4
$x \parallel y = x \parallel y + y \parallel x + x \mid y$	CM1	$\tau_I(a) = a$	if $a \notin I$ TI1
$a \parallel x = a \cdot x$	CM2	$\tau_I(a) = \tau$	if $a \in I$ TI2
$a \cdot x \parallel y = a \cdot (x \parallel y)$	CM3	$\tau_I(x + y) = \tau_I(x) + \tau_I(y)$	TI3
$(x + y) \parallel z = x \parallel z + y \parallel z$	CM4	$\tau_I(x \cdot y) = \tau_I(x) \cdot \tau_I(y)$	TI4
$a \cdot x \mid b = (a \mid b) \cdot x$	CM5		
$a \mid b \cdot x = (a \mid b) \cdot x$	CM6	$a \mid b = b \mid a$	C1
$a \cdot x \mid b \cdot y = (a \mid b) \cdot (x \parallel y)$	CM7	$(a \mid b) \mid c = a \mid (b \mid c)$	C2
$(x + y) \mid z = x \mid z + y \mid z$	CM8	$\delta \mid a = \delta$	C3
$x \mid (y + z) = x \mid y + x \mid z$	CM9	$\tau \mid a = \delta$	C4

The operators \parallel and \mid are of an auxiliary nature. They are needed to axiomatize ACP^τ .

The axioms of ACP^τ are given in Table 5. CM2–CM3, CM5–CM7, C1–C4, D1–D4 and TI1–TI4 are actually axiom schemas in which a , b and c stand for arbitrary constants of ACP^τ , and H and I stand for arbitrary subsets of A . ACP^τ is extended with guarded recursion like BTA.

A *recursive specification* over ACP^τ is a set of recursion equations $E = \{X = t_X \mid X \in V\}$, where V is a set of variables and each t_X is an ACP^τ term containing only variables from V . We write $V(E)$ for the set of all variables that occur in E . Let t be an ACP^τ term without occurrences of abstraction operators containing a variable X . Then an occurrence of X in t is *guarded* if t has a subterm of the form $e \cdot t'$ where $e \in A$ and t' is a term containing this occurrence of X . Let E be a recursive specification over ACP^τ . Then E is a *guarded recursive specification* if, in each equation $X = t_X \in E$: (i) abstraction operators do not occur in t_X and (ii) all occurrences of variables in t_X are guarded or t_X can be rewritten to such a term using the axioms of ACP^τ in either direction and/or the equations in E except the equation $X = t_X$ from left to right. We are only interested models of ACP^τ in which guarded recursive specifications have unique solutions, such as the models of ACP^τ presented in [3].

For each guarded recursive specification E and each $X \in V(E)$, we introduce a constant $\langle X|E \rangle$ standing for the X -component of the unique solution of E . We write $\langle t_X|E \rangle$ for t_X with, for all $Y \in V(E)$, all occurrences of Y in t_X replaced by $\langle Y|E \rangle$. The axioms for the constants for the components of the solutions of guarded recursive specifications are RDP and RSP, which are given in Table 6. RDP and RSP are actually axiom schemas in which X stands for an arbitrary variable, t_X stands for an arbitrary ACP^τ term, and E stands for an arbitrary guarded recursive specification over ACP^τ . Side conditions are added to restrict what X , t_X and E stand for.

Closed terms of ACP^τ extended with constants for the components of the solutions of guarded recursive specifications that denote the same process cannot always be proved equal by means of the axioms of ACP^τ together with RDP and RSP. We introduce AIP to remedy this. AIP is based on the

Table 6. RDP, RSP and AIP

$\langle X E \rangle = \langle t_X E \rangle$ if $X = t_X \in E$	RDP	$\pi_0(a) = \delta$	PR1
$E \Rightarrow X = \langle X E \rangle$ if $X \in V(E)$	RSP	$\pi_{n+1}(a) = a$	PR2
		$\pi_0(a \cdot x) = \delta$	PR3
$\bigwedge_{n \geq 0} \pi_n(x) = \pi_n(y) \Rightarrow x = y$	AIP	$\pi_{n+1}(a \cdot x) = a \cdot \pi_n(x)$	PR4
		$\pi_n(x + y) = \pi_n(x) + \pi_n(y)$	PR5
		$\pi_n(\tau) = \tau$	PR6
		$\pi_n(\tau \cdot x) = \tau \cdot \pi_n(x)$	PR7

view that two processes are identical if their approximations up to any finite depth are identical. The approximation up to depth n of a process behaves the same as that process, except that it cannot perform any further atomic action after n atomic actions have been performed. AIP is given in Table 6. Here, approximation up to depth n is phrased in terms of a unary *projection* operator π_n . The axioms for the projection operators are axioms PR1–PR7 in Table 6. PR1–PR7 are actually axiom schemas in which a stands for arbitrary constants of ACP^τ different from τ and n stands for an arbitrary natural number.

We write $ACP^\tau + REC$ for ACP^τ extended with the constants for the components of the solutions of guarded recursive specifications, the projection operators, and the axioms RDP, RSP, AIP and PR1–PR7.

The minimal models of $ACP^\tau + REC$ are considered its standard models. Henceforth, we restrict ourselves to the minimal models of $ACP^\tau + REC$. We assume that a fixed but arbitrary minimal model $\mathcal{M}_{ACP^\tau + REC}$ of $ACP^\tau + REC$ has been given.

From Section 12, we will sometimes assume that CFAR (Cluster Fair Abstraction Rule) is valid in $\mathcal{M}_{ACP^\tau + REC}$. CFAR says that a cluster of silent steps that has exits can be eliminated if all exits are reachable from everywhere in the cluster. A precise formulation of CFAR can be found in [24].

We use the term *process* for the elements from the domain of $\mathcal{M}_{ACP^\tau + REC}$, and we denote the interpretations of constants and operators in $\mathcal{M}_{ACP^\tau + REC}$ by the constants and operators themselves.

Let P be a process. Then the set of *states* or *subprocesses* of P , written $Sub(P)$, is inductively defined as follows:

- $P \in Sub(P)$;
- if $e \cdot P' \in Sub(P)$, then $P' \in Sub(P)$;
- if $e \cdot P' + P'' \in Sub(P)$, then $P' \in Sub(P)$.

Let P be a process and let $A' \subseteq A_\tau$. Then P is *regular over* A' if the following conditions are satisfied:

- $Sub(P)$ is finite;
- for all $P' \in Sub(P)$ and $e \in A_\tau$, $e \cdot P' \in Sub(P)$ implies $e \in A'$;
- for all $P', P'' \in Sub(P)$ and $e \in A_\tau$, $e \cdot P' + P'' \in Sub(P)$ implies $e \in A'$.

We say that P is *regular* if P is regular over A_τ .

We will make use of the fact that being a regular process over A coincides with being a component of the solution of a finite guarded recursive specification in which the right-hand sides of the recursion equations are linear terms. *Linearity* of terms is inductively defined as follows:

- δ is linear;
- if $e \in A_\tau$, then e is linear;
- if $e \in A_\tau$ and X is a variable, then $e \cdot X$ is linear;
- if t and t' are linear, then $t + t'$ is linear.

A *linear recursive specification* over ACP^τ is a guarded recursive specification $E = \{X = t_X \mid X \in V\}$ over ACP^τ , where each t_X is linear.

Proposition 5.1. Let P be a process and let $A' \subseteq A$. Then P is regular over A' iff there exists a finite linear recursive specification E over ACP^τ in which only atomic actions from A' occur such that P is a component of the solution of E .

Proof:

The proof follows the same line as the proof of Proposition 3.1. □

Remark 5.2. Proposition 5.1 is concerned with processes that are regular over A . We can also prove that being a regular process over A_τ coincides with being a component of the solution of a finite linear recursive specification over ACP^τ if we assume that the cluster fair abstraction rule [24] holds in the model $\mathcal{M}_{ACP^\tau+REC}$. However, we do not need this more general result.

We will write $\sum_{i \in S} t_i$, where $S = \{i_1, \dots, i_n\}$ and t_{i_1}, \dots, t_{i_n} are ACP^τ terms, for $t_{i_1} + \dots + t_{i_n}$. The convention is that $\sum_{i \in S} t_i$ stands for δ if $S = \emptyset$. We will often write X for $\langle X | E \rangle$ if E is clear from the context. It should be borne in mind that, in such cases, we use X as a constant.

6. Program-Service Interaction Instructions

Recall that, in PGA, it is assumed that a fixed but arbitrary set \mathfrak{A} of basic instructions has been given. In the sequel, we will make use a version of PGA in which the following additional assumptions relating to \mathfrak{A} are made:

- a fixed but arbitrary finite set \mathcal{F} of *foci* has been given;
- a fixed but arbitrary finite set \mathcal{M} of *methods* has been given;
- $\mathfrak{A} = \{f.m \mid f \in \mathcal{F}, m \in \mathcal{M}\}$.

Each focus plays the role of a name of some service provided by an execution environment that can be requested to process a command. Each method plays the role of a command proper. Executing a basic instruction of the form $f.m$ is taken as making a request to the service named f to process command m .

A basic instruction of the form $f.m$ is called a *program-service interaction instruction*. Recall that, in BTA, it is assumed that a fixed but arbitrary set \mathcal{A} of basic actions has been given. In the sequel, we will make use of a version of BTA in which $\mathcal{A} = \mathfrak{A}$. A basic action of the form $f.m$ is called a *thread-service interaction action*.

The intuition concerning program-service interaction instructions given above will be made fully precise in Section 7, using ACP.

Table 7. Defining equations for process extraction operation

$ S ^c = \text{stop}$
$ D ^c = \delta$
$ T \triangleleft \text{tau} \triangleright T' ^c = i \cdot i \cdot T ^c$
$ T \triangleleft f.m \triangleright T' ^c = s_f(m) \cdot (r_f(T) \cdot T ^c + r_f(F) \cdot T' ^c)$

7. Process Extraction

In this section, we use $ACP^\tau + \text{REC}$ to make mathematically precise which processes are produced by threads.

For that purpose, A and $| \cdot |$ are taken such that the following conditions are satisfied:⁴

$$A \supseteq \{s_f(d) \mid f \in \mathcal{F}, d \in \mathcal{M} \cup \mathbb{B}\} \cup \{r_f(d) \mid f \in \mathcal{F}, d \in \mathcal{M} \cup \mathbb{B}\} \cup \{\text{stop}, i\}$$

and for all $f \in \mathcal{F}$, $d \in \mathcal{M} \cup \mathbb{B}$, and $e \in A$:

$$\begin{aligned} s_f(d) \mid r_f(d) &= i, \\ s_f(d) \mid e &= \delta & \text{if } e \neq r_f(d), & & \text{stop} \mid e &= \delta & \text{if } e \neq \text{stop}, \\ e \mid r_f(d) &= \delta & \text{if } e \neq s_f(d), & & i \mid e &= \delta. \end{aligned}$$

Actions of the forms $s_f(d)$ and $r_f(d)$ are send and receive actions, respectively, stop is an explicit termination action, and i is a concrete internal action.

The *process extraction operation* $| \cdot |$ assigns a process to each thread. The process extraction operation $| \cdot |$ is defined by $|T| = \tau_{\{\text{stop}\}}(|T|^c)$, where $| \cdot |^c$ is defined by the equations given in Table 7 (for $f \in \mathcal{F}$ and $m \in \mathcal{M}$).

Let P be a process, T be a thread, and F be an instruction sequence. Then we say that T *produces* P if $\tau \cdot \tau_I(|T|) = \tau \cdot P$ for some $I \subseteq A$, and we say that F *produces* P if $|F|$ produces P .

Notice that two atomic actions are involved in performing a basic action of the form $f.m$: one for sending a request to process command m to the service named f and another for receiving a reply from that service upon completion of the processing. Notice also that, for each thread T , $|T|^c$ is a process that in the event of termination performs a special termination action just before termination. Abstraction from this termination action yields the process denoted by $|T|$.

The process extraction operation preserves the axioms of $\text{BTA} + \text{REC}$. Before we make this fully precise, we have a closer look at the axioms of $\text{BTA} + \text{REC}$.

A proper axiom is an equation or a conditional equation. In Table 3, we do not find proper axioms. Instead of proper axioms, we find axiom schemas without side conditions and axiom schemas with side conditions. The axioms of $\text{BTA} + \text{REC}$ are obtained by replacing each axiom schema by all its instances.

Henceforth, we write α^* , where α is a valuation of variables in $\mathcal{M}_{\text{BTA} + \text{REC}}$, for the unique homomorphic extension of α to terms of $\text{BTA} + \text{REC}$. Moreover, we identify $t_1 = t_2$ and $\emptyset \Rightarrow t_1 = t_2$.

Proposition 7.1. Let $E \Rightarrow t_1 = t_2$ be an axiom of $\text{BTA} + \text{REC}$, and let α be a valuation of variables in $\mathcal{M}_{\text{BTA} + \text{REC}}$. Then $|\alpha^*(t_1)| = |\alpha^*(t_2)|$ if $|\alpha^*(t'_1)| = |\alpha^*(t'_2)|$ for all $t'_1 = t'_2 \in E$.

⁴As usual, we will write \mathbb{B} for the set $\{T, F\}$.

Proof:

The proof is trivial for the axiom of BTA and the axioms RDP and RSP. Using the equation $|\pi_n(T)|^c = \pi_{2n}(|T|^c)$, the proof is also trivial for the axioms AIP and P0–P3. This equation is easily proved by induction on n and case distinction on the structure of T in both the basis step and the inductive step. \square

Remark 7.2. Proposition 7.1 would go through if no abstraction of the above-mentioned special termination action was made. Notice further that ACP^τ without the silent step constant and the abstraction operator, better known as ACP, would suffice if no abstraction of the special termination action was made.

8. A Simple Protocol for Remote Instruction Processing

In this section and the next section, we consider two protocols for remote instruction processing. The simple protocol described in this section is presumably the most straightforward protocol for remote instruction processing that can be achieved. Therefore, we consider it a suitable starting-point for the design of more advanced protocols for remote instruction processing – such as the one described in the next section. Before this simple protocol is described, an extension of ACP is introduced to simplify the description of the protocols.

The following extension of ACP from [2] will be used: the non-branching conditional operator $:\rightarrow$ over \mathbb{B} . The expression $b : \rightarrow p$, is to be read as if b then p else δ . The additional axioms for the non-branching conditional operator are

$$T : \rightarrow x = x \quad \text{and} \quad F : \rightarrow x = \delta .$$

In the sequel, we will use expressions whose evaluation yields Boolean values instead of the constants T and F. Because the evaluation of the expressions concerned are not dependent on the processes denoted by the terms in which they occur, we will identify each such expression with the constant for the Boolean value that its evaluation yields. Further justification of this can be found in [9, Section 9].

The protocols concern systems whose main components are an *instruction stream generator* and an *instruction stream execution unit*. The instruction stream generator generates different instruction streams for different threads. This is accomplished by starting it in different states. The general idea of the protocols is that:

- the instruction stream generator generating an instruction stream for a thread $T \triangleleft a \triangleright T'$ sends a to the instruction stream execution unit;
- on receipt of a , the instruction stream execution unit gets the execution of a done and sends the reply produced to the instruction stream generator;
- on receipt of the reply, the instruction stream generator proceeds with generating an instruction stream for T if the reply is T and for T' otherwise.

In the case where the thread is S or D, the instruction stream generator sends a special instruction (stop or dead) and the instruction stream execution unit does not send back a reply.

In this section, we consider a very simple protocol for remote instruction processing that makes no effort to keep the execution unit busy without intermission.

In the protocols, the generation of an instruction stream start from the thread produced by an instruction sequence under execution instead of the instruction sequence itself. It follows immediately from the definition of the thread extraction operation that the threads produced by instruction sequences under execution are regular threads. Therefore, we restrict ourselves to regular threads.

We write \mathcal{I} for the set $\mathcal{A} \cup \{\text{stop}, \text{dead}\}$. Elements from \mathcal{I} will loosely be called instructions. The restriction of the domain of $\mathcal{M}_{\text{BTA}+\text{REC}}$ to the regular threads will be denoted by \mathcal{RT} .

The functions act , $thrt$, and $thrf$ defined below give, for each thread T different from S and D, the basic action that T will perform first, the thread with which it will proceed if the reply from the execution environment is T, and the thread with which it will proceed if the reply from the execution environment is F, respectively. The functions $act : \mathcal{RT} \rightarrow \mathcal{I}$, $thrt : \mathcal{RT} \rightarrow \mathcal{RT}$, and $thrf : \mathcal{RT} \rightarrow \mathcal{RT}$ are defined as follows:

$$\begin{aligned} act(S) &= \text{stop} , & thrt(S) &= D , & thrf(S) &= D , \\ act(D) &= \text{dead} , & thrt(D) &= D , & thrf(D) &= D , \\ act(T \trianglelefteq a \triangleright T') &= a , & thrt(T \trianglelefteq a \triangleright T') &= T , & thrf(T \trianglelefteq a \triangleright T') &= T' . \end{aligned}$$

The function $next^0$ defined below is used by the instruction stream generator to distinguish when it starts with handling the instruction to be executed next between the different instructions that it may be. The function $next^0 : \mathcal{I} \times \mathcal{RT} \rightarrow \mathbb{B}$ is defined as follows:

$$next^0(a, T) = \begin{cases} T & \text{if } act(T) = a \\ F & \text{if } act(T) \neq a . \end{cases}$$

For the purpose of describing the simple protocol outlined above in ACP^τ , A and | are taken such that, in addition to the conditions mentioned at the beginning of Section 7, the following conditions are satisfied:

$$\begin{aligned} A \supseteq & \{s_i(d) \mid i \in \{1, 2\}, d \in \mathcal{I}\} \cup \{r_i(d) \mid i \in \{1, 2\}, d \in \mathcal{I}\} \\ & \cup \{s_i(r) \mid i \in \{3, 4\}, r \in \mathbb{B}\} \cup \{r_i(r) \mid i \in \{3, 4\}, r \in \mathbb{B}\} \cup \{j\} \end{aligned}$$

and for all $i \in \{1, 2\}$, $j \in \{3, 4\}$, $d \in \mathcal{I}$, $r \in \mathbb{B}$, and $e \in A$:

$$\begin{aligned} s_i(d) \mid r_i(d) &= j , & s_j(r) \mid r_j(r) &= j , \\ s_i(d) \mid e = \delta & \quad \text{if } e \neq r_i(d) , & s_j(r) \mid e = \delta & \quad \text{if } e \neq r_j(r) , \\ e \mid r_i(d) = \delta & \quad \text{if } e \neq s_i(d) , & e \mid r_j(r) = \delta & \quad \text{if } e \neq s_j(r) , \\ j \mid e &= \delta . \end{aligned}$$

Notice that the set \mathbb{B} is the set of replies.

Let $T \in \mathcal{RT}$. Then the process representing the simple protocol for remote instruction processing with regard to thread T is described by

$$\partial_H(ISG_T^0 \parallel IMTC^0 \parallel RTC^0 \parallel ISEU^0) ,$$

where the process ISG_T^0 is recursively specified by the following equation:

$$\begin{aligned} ISG_T^0 &= \sum_{f.m \in \mathcal{A}} next^0(f.m, T) : \rightarrow \\ & \quad s_1(f.m) \cdot (r_4(T) \cdot ISG_{thrt(T)}^0 + r_4(F) \cdot ISG_{thrf(T)}^0) \\ & \quad + next^0(\text{stop}, T) : \rightarrow s_1(\text{stop}) + next^0(\text{dead}, T) : \rightarrow s_1(\text{dead}) , \end{aligned}$$

the process $IMTC^0$ is recursively specified by the following equation:

$$IMTC^0 = \sum_{d \in \mathcal{I}} r_1(d) \cdot s_2(d) \cdot IMTC^0 ,$$

the process RTC^0 is recursively specified by the following equation:

$$RTC^0 = \sum_{r \in \mathbb{B}} r_3(r) \cdot s_4(r) \cdot RTC^0 ,$$

the process $ISEU^0$ is recursively specified by the following equation:

$$\begin{aligned} ISEU^0 = & \sum_{f.m \in \mathcal{A}} r_2(f.m) \cdot s_f(m) \cdot (r_f(\mathbb{T}) \cdot s_3(\mathbb{T}) + r_f(\mathbb{F}) \cdot s_3(\mathbb{F})) \cdot ISEU^0 \\ & + r_2(\text{stop}) + r_2(\text{dead}) \cdot \delta \end{aligned}$$

and

$$\begin{aligned} H = & \{s_i(d) \mid i \in \{1, 2\}, d \in \mathcal{I}\} \cup \{r_i(d) \mid i \in \{1, 2\}, d \in \mathcal{I}\} \\ & \cup \{s_i(r) \mid i \in \{3, 4\}, r \in \mathbb{B}\} \cup \{r_i(r) \mid i \in \{3, 4\}, r \in \mathbb{B}\} . \end{aligned}$$

ISG_T^0 is the instruction stream generator for thread T , $IMTC^0$ is the transmission channel for messages containing instructions, RTC^0 is the transmission channel for replies, and $ISEU^0$ is the instruction stream execution unit.

If we abstract from all communications via the transmission channels, then the process denoted by $\partial_H(ISG_T^0 \parallel IMTC^0 \parallel RTC^0 \parallel ISEU^0)$ and the process $|T|$ are equal modulo an initial silent step.

Theorem 8.1. For each $T \in \mathcal{RT}$, $\tau \cdot \tau_{\{j\}}(\partial_H(ISG_T^0 \parallel IMTC^0 \parallel RTC^0 \parallel ISEU^0))$ denotes the process $\tau \cdot |T|$.

Proof:

Let $T \in \mathcal{RT}$. Moreover, let E be a finite linear recursive specification over ACP^τ with $X \in V(E)$ such that $|T|$ is the X -component of the solution of E in $\mathcal{M}_{ACP^\tau + REC}$. By Proposition 5.1 and the definition of the process extraction operation, it is sufficient to prove that

$$\tau \cdot \tau_{\{j\}}(\partial_H(ISG_T^0 \parallel IMTC^0 \parallel RTC^0 \parallel ISEU^0)) = \tau \cdot \langle X|E \rangle .$$

By AIP, it is sufficient to prove that for all $n \geq 0$:

$$\pi_n(\tau \cdot \tau_{\{j\}}(\partial_H(ISG_T^0 \parallel IMTC^0 \parallel RTC^0 \parallel ISEU^0))) = \pi_n(\tau \cdot \langle X|E \rangle) .$$

This is easily proved by induction on n and in the inductive step by case distinction on the structure of T , using the axioms of ACP^τ and RDP and in addition the fact that $|T'| \in Sub(|T|)$ for all $T' \in Res(T)$ and the fact that there exists a bijection between $Sub(|T|)$ and $V(E)$. \square

9. A More Complex Protocol

In this section, we consider a more complex protocol for remote instruction processing that makes an effort to keep the execution unit busy without intermission.

The specifics of the more complex protocol considered here are that:

- the instruction stream generator may run ahead of the instruction stream execution unit by not waiting for the receipt of the replies resulting from the execution of instructions that it has sent earlier;
- to ensure that the instruction stream execution unit can handle the run-ahead, each instruction sent by the instruction stream generator is accompanied with the sequence of replies after which the instruction must be executed;
- to correct for replies that have not yet reached the instruction stream generator, each instruction sent is also accompanied with the number of replies received since the last sending of an instruction.

This protocol is reminiscent of an instruction pre-fetching mechanism as found in pipelined processors (see e.g. [26]), but its range of application is not restricted to pipelined instruction processing.

We write $\mathbb{B}^{\leq n}$, where $n \in \mathbb{N}$, for the set $\{u \in \mathbb{B}^* \mid \text{len}(u) \leq n\}$.⁵

It is assumed that a natural number ℓ has been given. The number ℓ is taken for the maximal number of steps that the instruction stream generator may run ahead of the instruction stream execution unit. Whether the execution unit can be kept busy without intermission with the given ℓ depends on the actual execution times of instructions and the actual transmission times over the transmission channels involved. If the execution unit can be kept busy without intermission with the given ℓ , then it is useless to increase ℓ .

The set \mathcal{IM} of *instruction messages* is defined as follows:

$$\mathcal{IM} = [0, \ell] \times \mathbb{B}^{\leq \ell} \times \mathcal{I} .$$

In an instruction message $(n, u, a) \in \mathcal{IM}$:

- n is the number of replies that are acknowledged by the message;
- u is the sequence of replies after which the instruction that is part of the message must be executed;
- a is the instruction that is part of the message.

The instruction stream generator sends instruction messages via an instruction message transmission channel to the instruction stream execution unit. We refer to a succession of transmitted instruction messages as an *instruction stream*. An instruction stream is dynamic by nature, in contradistinction with an instruction sequence.

The set \mathcal{S}_{ISG} of *instruction stream generator states* is defined as follows:

$$\mathcal{S}_{\text{ISG}} = [0, \ell] \times \mathcal{P}(\mathbb{B}^{\leq \ell+1} \times \mathcal{RT}) .$$

In an instruction stream generator state $(n, R) \in \mathcal{S}_{\text{ISG}}$:

⁵As usual, we write D^* for the set of all finite sequences with elements from set D and $\text{len}(\sigma)$ for the length of finite sequence σ . Moreover, we write ϵ for the empty sequence, d for the sequence having d as sole element, $\sigma\sigma'$ for the concatenation of finite sequences σ and σ' , and $\text{tl}(\sigma)$ for the tail of finite sequence σ .

- n is the number of replies that has been received by the instruction stream generator since the last acknowledgement of received replies;
- in each $(u, T) \in R$, u is the sequence of replies after which the thread T must be performed.

The functions $updpm$ and $updcr$ defined below are used to model the updates of the instruction stream generator state on producing a message and consuming a reply, respectively. The function $updpm : (\mathbb{B}^{\leq \ell} \times \mathcal{RT}) \times \mathcal{S}_{\text{ISG}} \rightarrow \mathcal{S}_{\text{ISG}}$ is defined as follows:

$$\begin{aligned} updpm((u, T), (n, R)) = \\ \begin{cases} (0, (R \setminus \{(u, T)\}) \cup \{(u\top, thrt(T)), (u\text{F}, thrf(T))\}) & \text{if } act(T) \in \mathcal{A} \\ (0, (R \setminus \{(u, T)\})) & \text{if } act(T) \notin \mathcal{A}. \end{cases} \end{aligned}$$

The function $updcr : \mathbb{B} \times \mathcal{S}_{\text{ISG}} \rightarrow \mathcal{S}_{\text{ISG}}$ is defined as follows:

$$updcr(r, (n, R)) = (n + 1, \{(u, T) \mid (ru, T) \in R\}).$$

The function sel defined below is used to model the selection of the sequence of replies and the instruction that will be part of the next message produced by the instruction stream generator. The function $sel : \mathcal{P}(\mathbb{B}^{\leq \ell} \times \mathcal{RT}) \rightarrow \mathcal{P}(\mathbb{B}^{\leq \ell} \times \mathcal{RT})$ is defined as follows:

$$sel(R) = \{(u, T) \in R \mid \forall (v, T') \in R \bullet len(u) \leq len(v)\}.$$

Notice that $(u, T) \in sel(R)$ and $(v, T') \in R$ only if $len(u) \leq len(v)$. By that breadth-first run-ahead is enforced. The performance of the protocol would change considerably if breadth-first run-ahead was not enforced.

The set $\mathcal{S}_{\text{ISEU}}$ of *instruction stream execution unit states* is defined as follows:

$$\mathcal{S}_{\text{ISEU}} = [0, \ell] \times \mathcal{P}(\mathbb{B}^{\leq \ell} \times \mathcal{I}).$$

In an instruction stream execution unit state $(n, S) \in \mathcal{S}_{\text{ISEU}}$:

- n is the number of replies for which the instruction stream execution unit still has to receive an acknowledgement;
- in each $(u, a) \in S$, u is the sequence of replies after which the instruction a must be executed.

The functions $updcm$ and $updpr$ defined below are used to model the updates of the instruction stream execution unit state on consuming a message and producing a reply, respectively. The function $updcm : \mathcal{IM} \times \mathcal{S}_{\text{ISEU}} \rightarrow \mathcal{S}_{\text{ISEU}}$ is defined as follows:

$$updcm((k, u, a), (n, S)) = (n \dot{-} k, S \cup \{(tl^{n \dot{-} k}(u), a)\}).⁶$$

The function $updpr : \mathbb{B} \times \mathcal{S}_{\text{ISEU}} \rightarrow \mathcal{S}_{\text{ISEU}}$ is defined as follows:

$$updpr(r, (n, S)) = (n + 1, \{(u, a) \mid (ru, a) \in S\}).$$

⁶As usual, we write $i \dot{-} j$ for the monus of i and j , i.e. $i \dot{-} j = i - j$ if $i \geq j$ and $i \dot{-} j = 0$ otherwise. As usual, $tl^n(u)$ is defined by induction on n as follows: $tl^0(u) = u$ and $tl^{n+1}(u) = tl(tl^n(u))$.

The function $next$ defined below is used by the instruction stream execution unit to distinguish when it starts with handling the instruction to be executed next between the different instructions that it may be. The function $next : \mathcal{I} \times \mathcal{P}(\mathbb{B}^{\leq \ell} \times \mathcal{I}) \rightarrow \mathbb{B}$ is defined as follows:

$$next(a, S) = \begin{cases} \text{T} & \text{if } (\epsilon, a) \in S \\ \text{F} & \text{if } (\epsilon, a) \notin S . \end{cases}$$

The instruction stream execution unit sends replies via a reply transmission channel to the instruction stream generator. We refer to a succession of transmitted replies as a *reply stream*.

For the purpose of describing the transmission protocol in ACP^T , A and $|$ are taken such that, in addition to the conditions mentioned at the beginning of Section 7, the following conditions are satisfied:

$$A \supseteq \{s_i(d) \mid i \in \{1, 2\}, d \in \mathcal{IM}\} \cup \{r_i(d) \mid i \in \{1, 2\}, d \in \mathcal{IM}\} \\ \cup \{s_i(r) \mid i \in \{3, 4\}, r \in \mathbb{B}\} \cup \{r_i(r) \mid i \in \{3, 4\}, r \in \mathbb{B}\} \cup \{j\}$$

and for all $i \in \{1, 2\}, j \in \{3, 4\}, d \in \mathcal{IM}, r \in \mathbb{B}$, and $e \in A$:

$$\begin{aligned} s_i(d) \mid r_i(d) = j, & & s_j(r) \mid r_j(r) = j, \\ s_i(d) \mid e = \delta & \quad \text{if } e \neq r_i(d), & s_j(r) \mid e = \delta & \quad \text{if } e \neq r_j(r), \\ e \mid r_i(d) = \delta & \quad \text{if } e \neq s_i(d), & e \mid r_j(r) = \delta & \quad \text{if } e \neq s_j(r), \\ j \mid e = \delta . & & & \end{aligned}$$

Let $T \in \mathcal{RT}$. Then the process representing the more complex protocol for remote instruction processing with regard to thread T is described by

$$\partial_H(ISG_T \parallel IMTC \parallel RTC \parallel ISEU) ,$$

where the process ISG_T is recursively specified by the following equations:

$$\begin{aligned} ISG_T &= ISG'_{(0, \{\epsilon, T\})} , \\ ISG'_{(n, R)} &= \sum_{(u, T) \in sel(R)} s_1((n, u, act(T))) \cdot ISG'_{updpm((u, T), (n, R))} \\ &\quad + \sum_{r \in \mathbb{B}} r_4(r) \cdot ISG'_{updcr(r, (n, R))} \\ &\text{(for every } (n, R) \in \mathcal{S}_{ISG} \text{ with } R \neq \emptyset \text{),} \\ ISG'_{(n, \emptyset)} &= j \\ &\text{(for every } (n, \emptyset) \in \mathcal{S}_{ISG} \text{),} \end{aligned}$$

the process $IMTC$ is recursively specified by the following equation:

$$IMTC = \sum_{d \in \mathcal{IM}} r_1(d) \cdot s_2(d) \cdot IMTC ,$$

the process RTC is recursively specified by the following equation:

$$RTC = \sum_{r \in \mathbb{B}} r_3(r) \cdot s_4(r) \cdot RTC ,$$

the process $ISEU$ is recursively specified by the following equations:

$$\begin{aligned} ISEU &= ISEU'_{(0, \emptyset)} , \\ ISEU'_{(n, S)} &= \sum_{d \in \mathcal{IM}} r_2(d) \cdot ISEU'_{updcm(d, (n, S))} \\ &\quad + \sum_{f.m \in \mathcal{A}} nxt(f.m, S) : \rightarrow s_f(m) \cdot ISEU''_{(f, (n, S))} \\ &\quad + nxt(\text{stop}, S) : \rightarrow j + nxt(\text{dead}, S) : \rightarrow \delta \\ &\text{(for every } (n, S) \in \mathcal{S}_{ISEU} \text{)} , \\ ISEU''_{(f, (n, S))} &= \sum_{r \in \mathbb{B}} r_f(r) \cdot s_3(r) \cdot ISEU'_{updpr(r, (n, S))} \\ &\quad + \sum_{d \in \mathcal{IM}} r_2(d) \cdot ISEU''_{(f, updcm(d, (n, S)))} \\ &\text{(for every } (f, (n, S)) \in \mathcal{F} \times \mathcal{S}_{ISEU} \text{)} , \end{aligned}$$

and

$$\begin{aligned} H &= \{s_i(d) \mid i \in \{1, 2\}, d \in \mathcal{IM}\} \cup \{r_i(d) \mid i \in \{1, 2\}, d \in \mathcal{IM}\} \\ &\quad \cup \{s_i(r) \mid i \in \{3, 4\}, r \in \mathbb{B}\} \cup \{r_i(r) \mid i \in \{3, 4\}, r \in \mathbb{B}\} . \end{aligned}$$

ISG_T is the instruction stream generator for thread T , $IMTC$ is the transmission channel for instruction messages, RTC is the transmission channel for replies, and $ISEU$ is the instruction stream execution unit.

The protocol described above has been designed such that, for each $T \in \mathcal{RT}$, $\tau \cdot \tau_{\{j\}}(\partial_H(ISG_T \parallel IMTC \parallel RTC \parallel ISEU))$ denotes the process $\tau \cdot |T|$. We refrain from presenting a proof of the claim that the protocol satisfies this because this paper is first and foremost a conceptual paper and the proof is straightforward but tedious.

The transmission channels $IMTC$ and RTC can keep one instruction message and one reply, respectively. The protocol has been designed in such a way that the protocol will also work properly if these channels are replaced by channels with larger capacity and even by channels with unbounded capacity.

Suppose that the transmission times over the transmission channels are small compared with the execution times of instructions. Even then the protocol described in Section 8 will always have to idle for a short time after the execution of an instruction, whereas after an initial phase the protocol described above will never have to idle after the execution of an instruction if the instruction stream generator may run a few steps ahead of the instruction stream execution unit.

10. Adaptations of the Protocol

In this section, we discuss some conceivable adaptations of the protocol described in Section 9. While we were thinking through the details of that protocol, various variations suggested themselves. The

variations discussed below are among the most salient ones. We think they deserve mention. However, their discussion is not in depth. The reason for this is that these variations have not yet been investigated thoroughly.

Consider the case where, for each instruction, it is known what the probability is with which its execution leads to the reply T. This might give reason to adapt the protocol described in Section 9. Suppose that the instruction stream generator states do not only keep the sequences of replies after which threads must be performed, but also the sequences of instructions involved in producing those sequences of replies. Then the probability with which the sequences of replies will happen can be calculated and several conceivable adaptations of the protocol to this probabilistic knowledge are possible by mere changes in the selection of the sequence of replies and the instruction that will be part of the next instruction message produced by the instruction stream generator. Among those adaptations are:

- restricting the instruction messages that are produced ahead to the ones where the sequence of replies after which the instruction must be executed will happen with a probability ≥ 0.50 , but sticking to breadth-first run-ahead;
- restricting the instruction messages that are produced ahead to the ones where the sequence of replies after which the instruction must be executed will happen with a probability ≥ 0.95 , but not sticking to breadth-first run-ahead.

At first sight, these adaptations are reminiscent of combinations of an instruction pre-fetching mechanism and a branch prediction mechanism as found in pipelined processors (see e.g. [26]). However, usually branch prediction mechanisms make use of statistics based on recently processed instructions instead of probabilistic knowledge of the kind used in the protocols sketched above.

Regular threads can be represented in such a way that it is effectively decidable whether the two threads with which a thread may proceed after performing its first action are identical. Consider the case where threads are represented in the instruction stream generator states in such a way. Then the protocol can be adapted such that no duplication of instruction messages takes place in the cases where the two threads with which a thread possibly proceeds after performing its first action are identical. This can be accomplished by using sequences of elements from $\mathbb{B} \cup \{*\}$, instead of sequences of elements from \mathbb{B} , in instruction messages, instruction stream generator states, and instruction stream execution unit states. The occurrence of $*$ at position i in a sequence indicates that the i th reply may be either T or F. The impact of this change on the updates of instruction stream generator states and instruction stream execution unit states is minor. This adaptation is reminiscent of an instruction pre-fetch mechanism as found in pipelined processors that prevents instruction pre-fetches that are superfluous due to identity of branches.

11. Alternative Choice Instructions

Process algebra is an area of the study of concurrency which is considered relevant to computer science, as is witnessed by the extent of the work on algebraic theories of processes such as ACP, CCS and CSP in theoretical computer science. This strongly hints that there are programmed systems whose behaviours can be taken for processes as considered in process algebra. Therefore, it is interesting to know to which extent the behaviours considered in process algebra can be produced by programs under execution, starting from the perception of a program as an instruction sequence. In coming sections, we

Table 8. Additional defining equation for process extraction operation

$$\overline{|T \triangleleft \text{ac}(e, e') \triangleright T'|^c = e \cdot |T|^c + e' \cdot |T'|^c}$$

will establish results concerning the processes as considered in ACP that can be produced by instruction sequences under execution.

For the purpose of producing processes as considered in ACP, we need a version of PGA with special basic instructions to deal with the non-deterministic choice between alternatives that stems from the alternative composition of processes. Recall that, in PGA, it is assumed that a fixed but arbitrary set \mathfrak{A} of basic instructions has been given. In the coming sections, we will make use a version of PGA in which the following additional assumptions relating to \mathfrak{A} are made:

- a fixed but arbitrary finite set \mathcal{F} of *foci* has been given;
- a fixed but arbitrary finite set \mathcal{M} of *methods* has been given;
- a fixed but arbitrary set \mathcal{AA} of *atomic actions*, with $t \notin \mathcal{AA}$, has been given;
- $\mathfrak{A} = \{f.m \mid f \in \mathcal{F}, m \in \mathcal{M}\} \cup \{\text{ac}(e_1, e_2) \mid e_1, e_2 \in \mathcal{AA} \cup \{t\}\}$.

On execution of a basic instruction $\text{ac}(e_1, e_2)$, first a non-deterministic choice between the atomic actions e_1 and e_2 is made and then the chosen atomic action is performed. The reply T is produced if e_1 is performed and the reply F is produced if e_2 is performed. Basic instructions of this kind are material to produce all regular processes by means of instruction sequences. A basic instruction of the form $\text{ac}(e_1, e_2)$ is called an *alternative choice instruction*. Henceforth, we will write PGA_{ac} for the version of PGA with alternative choice instructions.

The intuition concerning alternative choice instructions given above will be made fully precise at the end of this section, using ACP^τ . It will not be made fully precise using an extension of BTA because it is considered a basic property of threads that they are deterministic behaviours.

Recall that we make use of a version of BTA in which $\mathcal{A} = \mathfrak{A}$. A basic action of the form $\text{ac}(e_1, e_2)$ is called an *alternative choice action*. Henceforth, we will write BTA^{ac} for the version of BTA with alternative choice actions.

For the purpose of making precise what processes are produced by the threads denoted by closed terms of $\text{BTA}^{\text{ac}} + \text{REC}$, A and $|$ are taken such that, in addition to the conditions mentioned at the beginning of Section 7, the following conditions are satisfied:

$$A \supseteq \mathcal{AA} \cup \{t\}$$

and for all $e, e' \in A$:

$$e' | e = \delta \text{ if } e' \in \mathcal{AA} \cup \{t\}.$$

The process extraction operation for BTA^{ac} has as defining equations the equations given in Table 7 and in addition the equation given in Table 8.

Proposition 7.1 goes through for BTA^{ac} .

12. Instruction Sequence Producible Processes

It follows immediately from the definitions of the thread extraction and process extraction operations that the instruction sequences considered in PGA produce regular processes. The question is whether all regular processes are producible by these instruction sequences. In this section, we show that all regular processes can be produced by the instruction sequences with alternative choice instructions.

We will make use of the fact that all regular threads over \mathcal{A} can be produced by the single-pass instruction sequences considered in PGA.

Proposition 12.1. For each thread T that is regular over \mathcal{A} , there exists a PGA instruction sequence F such that F produces T , i.e. $|F| = T$.

Proof:

By Proposition 3.1, T is a component of the solution of some finite linear recursive specification E over BTA. There occur finitely many variables X_0, \dots, X_n in E . Assume that T is the X_0 -component of the solution of E . Let F be the PGA instruction sequence $(F_0 ; \dots ; F_n)^\omega$, where F_i is defined as follows ($0 \leq i \leq n$):

$$F_i = \begin{cases} ! ; ! ; ! & \text{if } X_i = S \in E \\ \#0 ; \#0 ; \#0 & \text{if } X_i = D \in E \\ +a ; \#3 \cdot (j-i) - 1 ; \#3 \cdot (k-i) - 2 & \text{if } X_i = X_j \triangleleft a \triangleright X_k \in E \wedge i < j \wedge i < k \\ +a ; \#3 \cdot (j-i) - 1 ; \#3 \cdot (n+1-(i-k)) - 2 & \text{if } X_i = X_j \triangleleft a \triangleright X_k \in E \wedge i < j \wedge i \geq k \\ +a ; \#3 \cdot (n+1-(i-j)) - 1 ; \#3 \cdot (k-i) - 2 & \text{if } X_i = X_j \triangleleft a \triangleright X_k \in E \wedge i \geq j \wedge i < k \\ +a ; \#3 \cdot (n+1-(i-j)) - 1 ; \#3 \cdot (n+1-(i-k)) - 2 & \text{if } X_i = X_j \triangleleft a \triangleright X_k \in E \wedge i \geq j \wedge i \geq k. \end{cases}$$

Then F is a PGA instruction sequence such that the interpretation of $|F| = T$. □

All regular processes over \mathcal{AA} can be produced by the instruction sequences considered in PGA_{ac} .

Theorem 12.2. Assume that CFAR is valid in $\mathcal{M}_{\text{ACP}^\tau + \text{REC}}$. Then, for each process P that is regular over \mathcal{AA} , there exists an instruction sequence F in which only basic instructions of the form $\text{ac}(e, t)$ occur such that F produces P , i.e. $\tau \cdot \tau_{\{t\}}(|F|) = \tau \cdot P$.

Proof:

By Propositions 3.1, 5.1 and 12.1, it is sufficient to show that, for each finite linear recursive specification E over ACP^τ in which only atomic actions from \mathcal{AA} occur, there exists a finite linear recursive specification E' over BTA^{ac} in which only basic actions of the form $\text{ac}(e, t)$ occur such that $\tau \cdot \langle X | E \rangle = \tau \cdot \tau_{\{t\}}(|\langle X | E' \rangle|)$ for all $X \in V(E)$.

Take the finite linear recursive specification E over ACP^τ that consists of the recursion equations

$$X_i = e_{i1} \cdot X_{i1} + \dots + e_{ik_i} \cdot X_{ik_i} + e'_{i1} + \dots + e'_{il_i},$$

where $e_{i1}, \dots, e_{ik_i}, e'_{i1}, \dots, e'_{il_i} \in \mathcal{AA}$, for $i \in \{1, \dots, n\}$. Then construct the finite linear recursive specification E' over BTA^{ac} that consists of the recursion equations

$$X_i = X_{i1} \triangleleft \text{ac}(e_{i1}, t) \triangleright (\dots (X_{ik_i} \triangleleft \text{ac}(e_{ik_i}, t) \triangleright (\text{S} \triangleleft \text{ac}(e'_{i1}, t) \triangleright (\dots (\text{S} \triangleleft \text{ac}(e'_{il_i}, t) \triangleright X_i) \dots))) \dots)$$

for $i \in \{1, \dots, n\}$; and the finite linear recursive specification E'' over ACP^τ that consists of the recursion equations

$$\begin{array}{ll} X_i &= e_{i1} \cdot X_{i1} + t \cdot Y_{i2} , & Z_{i1} &= e'_{i1} + t \cdot Z_{i2} , \\ Y_{i2} &= e_{i2} \cdot X_{i2} + t \cdot Y_{i3} , & Z_{i2} &= e'_{i2} + t \cdot Z_{i3} , \\ &\vdots & &\vdots \\ Y_{ik_i} &= e_{ik_i} \cdot X_{ik_i} + t \cdot Z_{i1} , & Z_{il_i} &= e'_{il_i} + t \cdot X_i , \end{array}$$

where $Y_{i2}, \dots, Y_{ik_i}, Z_{i1}, \dots, Z_{il_i}$ are fresh variables, for $i \in \{1, \dots, n\}$. It follows immediately from the definition of the process extraction operation that $|\langle X|E' \rangle| = \langle X|E'' \rangle$ for all $X \in V(E)$. Moreover, it follows from CFAR that $\tau \cdot \langle X|E \rangle = \tau \cdot \tau_{\{t\}}(\langle X|E'' \rangle)$ for all $X \in V(E)$. Hence, $\tau \cdot \langle X|E \rangle = \tau \cdot \tau_{\{t\}}(|\langle X|E' \rangle|)$ for all $X \in V(E)$. \square

For example, assuming that CFAR is valid, the instruction sequence

$$(+ac(r_3(T), t) ; \#4 ; +ac(r_3(F), t) ; \#5 ; \#7 ; +ac(s_4(T), t) ; \#5 ; \#9 ; +ac(s_4(F), t) ; \#2 ; \#9)^\omega$$

produces the reply transmission channel process RTC of which a guarded recursive specification is given in Section 9.

Remark 12.3. Theorem 12.2 with “ $\tau \cdot \tau_{\{t\}}(|F|) = \tau \cdot P$ ” replaced by “ $|F| = P$ ” can be established if PGA is extended with multiple-reply test instructions, see [11]. In that case, the assumption that CFAR is valid is superfluous.

13. Services and Use Operators

An instruction sequence under execution may make use of services. That is, certain instructions may be executed for the purpose of having the behaviour produced by the instruction sequence affected by a service that takes those instructions as commands to be processed. Likewise, a thread may perform certain actions for the purpose of having itself affected by a service that takes those actions as commands to be processed. The processing of an action may involve a change of state of the service and at completion of the processing of the action the service returns a reply value to the thread. The reply value determines how the thread proceeds. The use operators can be used in combination with the thread extraction operation from Section 4 to describe the behaviour produced by instruction sequences that make use of services. In this section, we first review the use operators, which are concerned with threads making such use of services, and then extend the process extraction operation to the use operators.

A *service* H consists of

- a set S of *states*;
- an *effect* function $eff : \mathcal{M} \times S \rightarrow S$;
- a *yield* function $yld : \mathcal{M} \times S \rightarrow \mathbb{B} \cup \{\mathbf{B}\}$;
- an *initial state* $s_0 \in S$;

Table 9. Axioms for use operators

$S /_f H = S$	U1
$D /_f H = D$	U2
$(x \trianglelefteq \text{tau} \trianglerighteq y) /_f H = (x /_f H) \trianglelefteq \text{tau} \trianglerighteq (y /_f H)$	U3
$(x \trianglelefteq g.m \trianglerighteq y) /_f H = (x /_f H) \trianglelefteq g.m \trianglerighteq (y /_f H)$ if $f \neq g$	U4
$(x \trianglelefteq f.m \trianglerighteq y) /_f H = \text{tau} \circ (x /_f \frac{\partial}{\partial m} H)$	if $H(m) = \mathbf{T}$ U5
$(x \trianglelefteq f.m \trianglerighteq y) /_f H = \text{tau} \circ (y /_f \frac{\partial}{\partial m} H)$	if $H(m) = \mathbf{F}$ U6
$(x \trianglelefteq f.m \trianglerighteq y) /_f H = \text{tau} \circ D$	if $H(m) = \mathbf{B}$ U7
$(x \trianglelefteq \text{ac}(e_1, e_2) \trianglerighteq y) /_f H = (x /_f H) \trianglelefteq \text{ac}(e_1, e_2) \trianglerighteq (y /_f H)$	U8
$\pi_n(x /_f H) = \pi_n(\pi_n(x) /_f H)$	U9

satisfying the following condition:

$$\forall m \in \mathcal{M}, s \in S \bullet (yld(m, s) = \mathbf{B} \Rightarrow \forall m' \in \mathcal{M} \bullet yld(m', eff(m, s)) = \mathbf{B}) .$$

The set S contains the states in which the service may be, and the functions eff and yld give, for each method m and state s , the state and reply, respectively, that result from processing m in state s . By the condition imposed on services, once the service has returned \mathbf{B} as reply, it keeps returning \mathbf{B} as reply.

Let $H = (S, eff, yld, s_0)$ be a service and let $m \in \mathcal{M}$. Then the *derived service* of H after processing m , written $\frac{\partial}{\partial m} H$, is the service $(S, eff, yld, eff(m, s_0))$; and the *reply* of H after processing m , written $H(m)$, is $yld(m, s_0)$.

When a thread makes a request to the service to process m :

- if $H(m) \neq \mathbf{B}$, then the request is accepted, the reply is $H(m)$, and the service proceeds as $\frac{\partial}{\partial m} H$;
- if $H(m) = \mathbf{B}$, then the request is rejected and the service proceeds as a service that rejects any request.

We introduce the sort \mathbf{S} of *services*. However, we will not introduce constants and operators to build terms of this sort. The sort \mathbf{S} , standing for the set of all services, is considered a parameter of the extension of BTA being presented. Moreover, we introduce, for each $f \in \mathcal{F}$, the binary *use operator* $/_f : \mathbf{T} \times \mathbf{S} \rightarrow \mathbf{T}$. The axioms for these operators are given in Table 9. Intuitively, $T /_f H$ is the thread that results from processing all actions performed by thread T that are of the form $f.m$ by service H . When a basic action of the form $f.m$ performed by thread T is processed by service H , it is turned into the basic action tau and postconditional composition is removed in favour of basic action prefixing on the basis of the reply value produced.

We add the use operators to PGA_{ac} as well. We will only use the extension in combination with the thread extraction operation $|_$ and define $|F /_f H| = |F| /_f H$. Hence, $|F /_f H|$ denotes the thread produced by F if F makes use of H . If H is a service such as an unbounded counter, an unbounded stack or a Turing tape, then a non-regular thread may be produced.

In order to extend the process extraction operation to the use operators, we need an extension of ACP^τ with action renaming operators ρ_h , where $h : A_\tau \rightarrow A_\tau$ such that $h(\tau) = \tau$. The axioms for action renaming are given in [24]. Intuitively, $\rho_h(P)$ behaves as P with each atomic action replaced according to h . We write $\rho_{e' \mapsto e''}$ for the renaming operator ρ_h with h defined by $h(e') = e''$ and $h(e) = e$ if $e \neq e'$.

Table 10. Adapted and additional defining equations for process extraction operation

$$\begin{aligned} |T \trianglelefteq f.m \triangleright T'|^c &= s_f(m) \cdot (r_f(T) \cdot |T|^c + r_f(F) \cdot |T'|^c + r_f(B) \cdot \delta) \\ |T /_f H|^c &= \rho_{\text{stop}^* \mapsto \text{stop}}(\partial_{\{\text{stop}\}}(\partial_{A_f}(|T|^c \parallel \rho_{h_f}(|H|^c)))) \end{aligned}$$

For the purpose of extending the process extraction operation to the use operators, A and $|$ are taken such that, in addition to the conditions mentioned at the beginning of Section 7, with everywhere \mathbb{B} replaced by $\mathbb{B} \cup \{\mathbf{B}\}$, and the conditions mentioned at the end of Section 11, the following conditions are satisfied:

$$A \supseteq \{s_{\text{serv}}(r) \mid r \in \mathbb{B} \cup \{\mathbf{B}\}\} \cup \{r_{\text{serv}}(m) \mid m \in \mathcal{M}\} \cup \{\text{stop}^*\}$$

and for all $e \in A$, $m \in \mathcal{M}$, and $r \in \mathbb{B} \cup \{\mathbf{B}\}$:

$$\begin{aligned} s_{\text{serv}}(r) \mid e &= \delta, & \text{stop} \mid \text{stop} &= \text{stop}^*, \\ e \mid r_{\text{serv}}(m) &= \delta, & \text{stop}^* \mid e &= \delta. \end{aligned}$$

We also need to define a set $A_f \subseteq A$ and a function $h_f : A_\tau \rightarrow A_\tau$ for each $f \in \mathcal{F}$:

$$A_f = \{s_f(d) \mid d \in \mathcal{M} \cup \mathbb{B} \cup \{\mathbf{B}\}\} \cup \{r_f(d) \mid d \in \mathcal{M} \cup \mathbb{B} \cup \{\mathbf{B}\}\};$$

for all $e \in A_\tau$, $m \in \mathcal{M}$ and $r \in \mathbb{B} \cup \{\mathbf{B}\}$:

$$\begin{aligned} h_f(s_{\text{serv}}(r)) &= s_f(r), \\ h_f(r_{\text{serv}}(m)) &= r_f(m), \\ h_f(e) &= e & \text{if } \bigwedge_{r' \in \mathbb{N}} e \neq s_{\text{serv}}(r') \wedge \bigwedge_{m' \in \mathcal{M}} e \neq r_{\text{serv}}(m'). \end{aligned}$$

To extend the process extraction operation to the use operators, the defining equation concerning the postconditional composition operators has to be adapted and a new defining equation concerning the use operators has to be added. These two equations are given in Table 10, where $|H|^c$ is the X_H -component of the solution of

$$\{X_{H'} = \sum_{m \in \mathcal{M}} r_{\text{serv}}(m) \cdot s_{\text{serv}}(H'(m)) \cdot X_{\frac{\partial}{\partial m} H'} + \text{stop} \mid H' \in \Delta(H)\},$$

where $\Delta(H)$ is inductively defined as follows:

- $H \in \Delta(H)$;
- if $m \in \mathcal{M}$ and $H' \in \Delta(H)$, then $\frac{\partial}{\partial m} H' \in \Delta(H)$.

The extended process extraction operation preserves the axioms for the use operators. Owing to the presence of axiom schemas with semantic side conditions in Table 9, the axioms for the use operators include proper axioms, which are all of the form $t_1 = t_2$, and axioms that have a semantic side condition, which are all of the form $t_1 = t_2$ if $H(m) = r$. By that, the precise formulation of the preservation result is somewhat complicated.

Proposition 13.1.

1. Let $t_1 = t_2$ be a proper axiom for the use operators, and let α be a valuation of variables in $\mathcal{M}_{\text{BTA}+\text{REC}}$. Then $|\alpha^*(t_1)| = |\alpha^*(t_2)|$.
2. Let $t_1 = t_2$ if $H(m) = r$ be an axiom with semantic side condition for the use operators, and let α be a valuation of variables in $\mathcal{M}_{\text{BTA}+\text{REC}}$. Then $|\alpha^*(t_1)| = |\alpha^*(t_2)|$ if $H(m) = r$.

Proof:

The proof is straightforward. We sketch the proof for axiom U5. By the definition of the process extraction operation, it is sufficient to show that $|(T \trianglelefteq f.m \triangleright T') /_f H|^c = |\text{tau} \circ (T /_f \frac{\partial}{\partial m} H)|^c$ if $H(m) = \text{T}$. In outline, this goes as follows:

$$\begin{aligned}
 & |(T \trianglelefteq f.m \triangleright T') /_f H|^c \\
 &= \rho_{\text{stop}^* \mapsto \text{stop}} \\
 &\quad (\partial_{\{\text{stop}\}}(\partial_{A_f}(\text{s}_f(m) \cdot (\text{r}_f(\text{T}) \cdot |T|^c + \text{r}_f(\text{F}) \cdot |T'|^c + \text{r}_f(\text{B}) \cdot \delta) \parallel \rho_{h_f}(|H|^c)))) \\
 &= \text{i} \cdot \text{i} \cdot \rho_{\text{stop}^* \mapsto \text{stop}}(\partial_{\{\text{stop}\}}(\partial_{A_f}(|T|^c \parallel \rho_{h_f}(|\frac{\partial}{\partial m} H|^c)))) \\
 &= |\text{tau} \circ (T /_f \frac{\partial}{\partial m} H)|^c.
 \end{aligned}$$

In the first and third step, we apply defining equations of $|-|^c$. In the second step, we apply axioms of $\text{ACP}^\tau + \text{REC}$ with action renaming, and use that $H(m) = \text{T}$. □

Remark 13.2. Let F be a PGA_{ac} instruction sequence and H be a service. Then $\|F /_f H\|$ is the process produced by F if F makes use of H . Instruction sequences that make use of services such as unbounded counters, unbounded stacks or Turing tapes are interesting because they may produce non-regular processes.

14. PGLD Programs and the Use of Boolean Registers

In this section, we show that all regular processes can also be produced by programs written in a program notation which is close to existing assembly languages, and even by programs in which no atomic action occurs more than once in an alternative choice instruction. The latter result requires programs that make use of Boolean registers.

A hierarchy of program notations rooted in PGA is introduced in [8]. One program notation that belongs to this hierarchy is PGLD, a very simple program notation which is close to existing assembly languages. It has absolute jump instructions and no explicit termination instruction.

In PGLD, like in PGA, it is assumed that there is a fixed but arbitrary finite set of *basic instructions* \mathcal{Q} . The primitive instructions of PGLD differ from the primitive instructions of PGA as follows: for each $l \in \mathbb{N}$, there is an *absolute jump instruction* $\#\#l$ instead of a forward jump instruction $\#l$. PGLD programs have the form $u_1; \dots; u_k$, where u_1, \dots, u_k are primitive instructions of PGLD.

The effects of all instructions in common with PGA are as in PGA with one difference: if there is no next instruction to be executed, termination occurs. The effect of an absolute jump instruction $\#\#l$ is that execution proceeds with the l -th instruction of the program concerned. If $\#\#l$ is itself the l -th instruction, then inaction occurs. If l equals 0 or l is greater than the length of the program, then termination occurs.

We define the meaning of PGLD programs by means of a function pgld2pga from the set of all PGLD programs to the set of all closed PGA terms. This function is defined by

$$\text{pgld2pga}(u_1; \dots; u_k) = (\phi_1(u_1); \dots; \phi_k(u_k); !; !)^{\omega},$$

where the auxiliary functions ϕ_j from the set of all primitive instructions of PGLD to the set of all primitive instructions of PGA are defined as follows ($1 \leq j \leq k$):

$$\begin{aligned} \phi_j(\#\#l) &= \#l - j && \text{if } j \leq l \leq k, \\ \phi_j(\#\#l) &= \#k + 2 - (j - l) && \text{if } 0 < l < j, \\ \phi_j(\#\#l) &= ! && \text{if } l = 0 \vee l > k, \\ \phi_j(u) &= u && \text{if } u \text{ is not a jump instruction.} \end{aligned}$$

PGLD is as expressive as PGA. Before we make this fully precise, we introduce a useful notation.

Let α is a valuation of variables in \mathcal{I}_{PGA} , and let α^* be the unique homomorphic extension of α to terms of PGA. Then $\alpha^*(t)$ is independent of α if t is a closed term, i.e. $\alpha^*(t)$ is uniquely determined by \mathcal{I}_{PGA} . Therefore, we write $t^{\mathcal{I}_{\text{PGA}}}$ for $\alpha^*(t)$ if t is a closed term.

Proposition 14.1. For each closed PGA term t , there exists a PGLD program p such that $|t^{\mathcal{I}_{\text{PGA}}}| = |\text{pgld2pga}(p)^{\mathcal{I}_{\text{PGA}}}|$.

Proof:

In [8], a number of functions (called embeddings in that paper) are defined, whose composition gives, for each closed PGA term t , a PGLD program p such that $|t^{\mathcal{I}_{\text{PGA}}}| = |\text{pgld2pga}(p)^{\mathcal{I}_{\text{PGA}}}|$. \square

Let p be a PGLD program and P be a process. Then we say that p *produces* P if $|\text{pgld2pga}(p)^{\mathcal{I}_{\text{PGA}}}|$ produces P .

Below, we will write PGLD_{ac} for the version of PGLD in which the additional assumptions relating to \mathfrak{A} mentioned in Section 11 are made. As a corollary of Theorem 12.2 and Proposition 14.1, we have that all regular processes over \mathcal{AA} can be produced by PGLD_{ac} programs.

Corollary 14.2. Assume that CFAR is valid in $\mathcal{M}_{\text{ACP}\tau+\text{REC}}$. Then, for each process P that is regular over \mathcal{AA} , there exists a PGLD_{ac} program p such that p produces P .

We switch to the use of Boolean registers now. First, we describe services that make up Boolean registers.

A Boolean register service accepts the following methods:

- a *set to true method* set:T ;
- a *set to false method* set:F ;
- a *get method* get .

We write \mathcal{M}_{BR} for the set $\{\text{set:T}, \text{set:F}, \text{get}\}$. It is assumed that $\mathcal{M}_{\text{BR}} \subseteq \mathcal{M}$.

The methods accepted by Boolean register services can be explained as follows:

- set:T : the contents of the Boolean register becomes T and the reply is T;

- set:F : the contents of the Boolean register becomes F and the reply is F;
- get : nothing changes and the reply is the contents of the Boolean register.

Let $s \in \mathbb{B} \cup \{\mathbf{B}\}$. Then the *Boolean register service* with initial state s , written BR_s , is the service $(\mathbb{B} \cup \{\mathbf{B}\}, \text{eff}, \text{eff}, s)$, where the function eff is defined as follows ($b \in \mathbb{B}$):

$$\begin{aligned} \text{eff}(\text{set:T}, b) &= \mathbf{T}, & \text{eff}(m, b) &= \mathbf{B} \text{ if } m \notin \mathcal{M}_{\text{BR}}, \\ \text{eff}(\text{set:F}, b) &= \mathbf{F}, & \text{eff}(m, \mathbf{B}) &= \mathbf{B}. \\ \text{eff}(\text{get}, b) &= b, \end{aligned}$$

Notice that the effect and yield functions of a Boolean register service are the same.

Let p be a PGLD program and P be a process. Then we say that p produces P using Boolean registers if $(\dots(|\text{pgld2pga}(p)^{\mathcal{I}_{\text{PGA}}}| /_{\text{br:1}} BR_{\mathbf{F}}) \dots /_{\text{br:k}} BR_{\mathbf{F}})$ produces P for some $k \in \mathbb{N}^+$.

We have that PGLD_{ac} programs in which no atomic action from \mathcal{AA} occurs more than once in an alternative choice instruction can produce all regular processes over \mathcal{AA} using Boolean registers.

Theorem 14.3. Assume that CFAR is valid in $\mathcal{M}_{\text{ACP}\tau+\text{REC}}$. Then, for each process P that is regular over \mathcal{AA} , there exists a PGLD_{ac} program p in which each atomic action from \mathcal{AA} occurs no more than once in an alternative choice instruction such that p produces P using Boolean registers.

Proof:

By the proof of Theorem 12.2 given in Section 12, it is sufficient to show that, for each thread T that is regular over \mathcal{A} , there exist a PGLD program p in which each basic action from \mathcal{A} occurs no more than once and a $k \in \mathbb{N}^+$ such that $(\dots(|\text{pgld2pga}(p)^{\mathcal{I}_{\text{PGA}}}| /_{\text{br:1}} BR_{\mathbf{F}}) \dots /_{\text{br:k}} BR_{\mathbf{F}}) = T$.

Let T be a thread that is regular over \mathcal{A} . We may assume that T is produced by a PGLD program p' of the following form:

$$\begin{aligned} &+a_1 ; \#\#(3 \cdot k_1 + 1) ; \#\#(3 \cdot k'_1 + 1) ; \\ &\quad \vdots \\ &+a_n ; \#\#(3 \cdot k_n + 1) ; \#\#(3 \cdot k'_n + 1) ; \\ &\#\#0 ; \#\#0 ; \#\#0 ; \#\#(3 \cdot n + 4) , \end{aligned}$$

where, for each $i \in [1, n]$, $k_i, k'_i \in [0, n - 1]$ (cf. the proof of Proposition 2 from [36]). It is easy to see that the PGLD program p that we are looking for can be obtained by transforming p' : by making use of n Boolean registers, p can distinguish between different occurrences of the same basic instruction in p' , and in that way simulate p' . □

15. Conclusions

Using the algebraic theory of processes known as ACP, we have described two protocols to deal with the phenomenon that, on execution of an instruction sequence, a stream of instructions to be processed arises at one place and the processing of that stream of instructions is handled at another place. The more complex protocol is directed towards keeping the execution unit busy. In this way, we have brought the phenomenon better into the picture and have ascribed a sense to the term instruction stream which makes

clear that an instruction stream is dynamic by nature, in contradistinction with an instruction sequence. We have also discussed some conceivable adaptations of the more complex protocol.

The description of the protocols start from the behaviours produced by instruction sequences under execution. By that we abstract from the instruction sequences which produce those behaviours. How instruction streams can be generated efficiently from instruction sequences is a matter that obviously requires investigations at a less abstract level. The investigations in question are an option for future work.

We believe that the more complex protocol described in this paper provides a setting in which basic techniques aimed at increasing processor performance, such as pre-fetching and branch prediction, can be studied at a more abstract level than usual (cf. [26]). In particular, we think that the protocol can serve as a starting-point for the development of a model with which trade-offs encountered in the design of processor architectures can be clarified. We consider investigations into this matter an interesting option for future work.

The fact that process algebra is an area of the study of concurrency which is considered relevant to computer science, strongly hints that there are programmed systems whose behaviours are taken for processes as considered in process algebra. In that light, we have investigated the connections between programs and the processes that they produce, starting from the perception of a program as an instruction sequence. We have shown that, by apposite choice of basic instructions, all regular processes can be produced by means of instruction sequences as considered in PGA.

We have also made precise what processes are produced by instruction sequences under execution that make use of services. The reason for this is that instruction sequences under execution are regular threads and regular threads that make use of services such as unbounded counters, unbounded stacks or Turing tapes may produce non-regular processes. An option for future work is to characterize the classes of processes that can be produced by single-pass instruction sequences that make use of such services.

References

- [1] Arora, S., Barak, B.: *Computational Complexity: A Modern Approach*, Cambridge University Press, Cambridge, 2009.
- [2] Baeten, J. C. M., Bergstra, J. A.: Process Algebra with Signals and Conditions, *Programming and Mathematical Methods* (M. Broy, Ed.), vol. F88 of NATO ASI Series, Springer-Verlag, 1992, 273–323.
- [3] Baeten, J. C. M., Weijland, W. P.: *Process Algebra*, vol. 18 of *Cambridge Tracts in Theoretical Computer Science*, Cambridge University Press, Cambridge, 1990.
- [4] Baker, H. G.: Precise Instruction Scheduling without a Precise Machine Model, *SIGARCH Computer Architecture News*, **19**(6), 1991, 4–8.
- [5] Bergstra, J. A., Bethke, I.: Polarized Process Algebra and Program Equivalence, *Proceedings 30th ICALP* (J. C. M. Baeten, J. K. Lenstra, J. Parrow, G. J. Woeginger, Eds.), vol. 2719 of *Lecture Notes in Computer Science*, Springer-Verlag, 2003, 1–21.
- [6] Bergstra, J. A., Klop, J. W.: Process Algebra for Synchronous Communication, *Information and Control*, **60**(1–3), 1984, 109–137.
- [7] Bergstra, J. A., Loots, M. E.: Program Algebra for Component Code, *Formal Aspects of Computing*, **12**(1), 2000, 1–17.

- [8] Bergstra, J. A., Loots, M. E.: Program Algebra for Sequential Code, *Journal of Logic and Algebraic Programming*, **51**(2), 2002, 125–156.
- [9] Bergstra, J. A., Middelburg, C. A.: Splitting Bisimulations and Retrospective Conditions, *Information and Computation*, **204**(7), 2006, 1083–1138.
- [10] Bergstra, J. A., Middelburg, C. A.: Maurer Computers with Single-Thread Control, *Fundamenta Informaticae*, **80**(4), 2007, 333–362.
- [11] Bergstra, J. A., Middelburg, C. A.: Instruction Sequences for the Production Of Processes, [arXiv:0811.0436v2 \[cs.PL\]](#), November 2008.
- [12] Bergstra, J. A., Middelburg, C. A.: Program Algebra with a Jump-Shift Instruction, *Journal of Applied Logic*, **6**(4), 2008, 553–563.
- [13] Bergstra, J. A., Middelburg, C. A.: A Protocol for Instruction Stream Processing, [arXiv:0905.2257v1 \[cs.PL\]](#), May 2009.
- [14] Bergstra, J. A., Middelburg, C. A.: Transmission Protocols for Instruction Streams, *ICTAC 2009* (M. Leucker, C. Morgan, Eds.), vol. 5684 of Lecture Notes in Computer Science, Springer-Verlag, 2009, 127–139.
- [15] Bergstra, J. A., Middelburg, C. A.: Instruction Sequences and Non-uniform Complexity Theory, [arXiv:0809.0352v3 \[cs.CC\]](#), July 2010.
- [16] Bergstra, J. A., Middelburg, C. A.: On the Operating Unit Size of Load/Store Architectures, *Mathematical Structures in Computer Science*, **20**(3), 2010, 395–417.
- [17] Bergstra, J. A., Middelburg, C. A.: Indirect Jumps Improve Instruction Sequence Performance, [arXiv:0909.2089v2 \[cs.PL\]](#), December 2011. To appear in *Scientific Annals of Computer Science*, 2012.
- [18] Bergstra, J. A., Middelburg, C. A.: Thread Extraction for Polyadic Instruction Sequences, *Scientific Annals of Computer Science*, **21**(2), 2011, 283–310.
- [19] Bergstra, J. A., Middelburg, C. A.: Instruction Sequence Processing Operators, *Acta Informatica*, **49**(3), 2012, 139–172.
- [20] Bergstra, J. A., Middelburg, C. A.: On the Expressiveness of Single-Pass Instruction Sequences, *Theory of Computing Systems*, **50**(2), 2012, 313–328.
- [21] Bergstra, J. A., Ponse, A.: An Instruction Sequence Semigroup with Involutive Anti-Automorphisms, *Scientific Annals of Computer Science*, **19**, 2009, 57–92.
- [22] Brock, C., Hunt, W. A.: Formally Specifying and Mechanically Verifying Programs for the Motorola Complex Arithmetic Processor DSP, *ICCD '97*, 1997, 31–36.
- [23] Brookes, S. D., Hoare, C. A. R., Roscoe, A. W.: A Theory of Communicating Sequential Processes, *Journal of the ACM*, **31**(3), 1984, 560–599.
- [24] Fokkink, W. J.: *Introduction to Process Algebra*, Texts in Theoretical Computer Science, An EATCS Series, Springer-Verlag, Berlin, 2000.
- [25] Hennessy, J., Jouppi, N., Przybylski, S., Rowen, C., Gross, T., Baskett, F., Gill, J.: MIPS: A Microprocessor Architecture, *MICRO '82*, 1982, 17–22.
- [26] Hennessy, J. L., Patterson, D. A.: *Computer Architecture: A Quantitative Approach*, Third edition, Morgan Kaufmann, San Francisco, 2003.
- [27] Hennessy, M., Milner, R.: Algebraic Laws for Non-determinism and Concurrency, *Journal of the ACM*, **32**(1), 1985, 137–161.

- [28] Hermes, H.: *Enumerability, Decidability, Computability*, Springer-Verlag, Berlin, 1965.
- [29] Hoare, C. A. R.: *Communicating Sequential Processes*, Prentice-Hall, Englewood Cliffs, 1985.
- [30] Lunde, A.: Empirical Evaluation of Some Features of Instruction Set Processor Architectures, *Communications of the ACM*, **20**(3), 1977, 143–153.
- [31] Milner, R.: *Communication and Concurrency*, Prentice-Hall, Englewood Cliffs, 1989.
- [32] Mosses, P. D.: Formal Semantics of Programming Languages — An Overview, *Electronic Notes in Theoretical Computer Science*, **148**, 2006, 41–73.
- [33] Nair, R., Hopkins, M. E.: Exploiting Instruction Level Parallelism in Processors by Caching Scheduled Groups, *SIGARCH Computer Architecture News*, **25**(2), 1997, 13–25.
- [34] Ofelt, D., Hennessy, J. L.: Efficient Performance Prediction for Modern Microprocessors, *SIGMETRICS '00*, 2000, 229–239.
- [35] Patterson, D. A., Ditzel, D. R.: The Case for the Reduced Instruction Set Computer, *SIGARCH Computer Architecture News*, **8**(6), 1980, 25–33.
- [36] Ponse, A., van der Zwaag, M. B.: An Introduction to Program and Thread Algebra, *CiE 2006* (A. Beckmann, et al., Eds.), vol. 3988 of Lecture Notes in Computer Science, Springer-Verlag, 2006, 445–458.
- [37] Sannella, D., Tarlecki, A.: Algebraic Preliminaries, in: *Algebraic Foundations of Systems Specification* (E. Astesiano, H.-J. Kreowski, B. Krieg-Brückner, Eds.), Springer-Verlag, Berlin, 1999, 13–30.
- [38] Sipser, M.: *Introduction to the Theory of Computation*, Second edition, Thomson, Boston, MA, 2006.
- [39] Tennenhouse, D. L., Wetherall, D. J.: Towards an Active Network Architecture, *SIGCOMM Computer Communication Review*, **37**(5), 2007, 81–94.
- [40] Wirsing, M.: Algebraic Specification, in: *Handbook of Theoretical Computer Science* (J. van Leeuwen, Ed.), vol. B, Elsevier, Amsterdam, 1990, 675–788.
- [41] Xia, C., Torrellas, J.: Instruction Prefetching of Systems Codes with Layout Optimized for Reduced Cache Misses, *ISCA '96*, 1996, 271–282.