

DEPARTAMENTO DE SISTEMAS INFORMÁTICOS Y COMPUTACIÓN
UNIVERSIDAD POLITÉCNICA DE VALENCIA

P.O. Box: 22012 E-46071 Valencia (SPAIN)



Informe Técnico / Technical Report

Ref. No.: DSIC-II/03/10	Pages: 30
Title: A Tracking Semantics for CSP (Extended Version)	
Author(s): M. Llorens, J. Oliver, J. Silva and S. Tamarit	
Date: March, 2010	
Keywords: Concurrency, CSP, Semantics, Trace	

Vº Bº
Leader of research Group

Author(s)

A Tracking Semantics for CSP^{*} (Extended Version)

Marisa Llorens, Javier Oliver, Josep Silva, and Salvador Tamarit

Universidad Politécnic de Valencia,
Camino de Vera S/N, E-46022 Valencia, Spain
{mllorens,fjoliver,jsilva,stamarit}@dsic.upv.es

Abstract. CSP is a powerful language for specifying complex concurrent systems. Due to the non-deterministic execution order of processes and to the restrictions imposed on this order by synchronizations, many analyses such as deadlock analysis, reliability analysis, and program slicing try to predict properties of the specification which can guarantee the quality of the final system. These analyses often rely on the use of CSP's traces. In this work, we introduce the theoretical basis for tracking concurrent and explicitly synchronized computations in process algebras such as CSP. Tracking computations is a difficult task due to the subtleties of the underlying operational semantics which combines concurrency, non-determinism and non-termination. We define an instrumented operational semantics that generates as a side-effect an appropriate data structure (a track) which can be used to track computations at an adequate level of abstraction. Given an execution, its track is more informative than its trace since the former not only contains a lot of information about original program structures but also explicitly relates the sequence of events with the parts of the specification that caused these events. Formal definition of a tracking semantics improves the understanding of the tracking process, but also, it allows to formally prove the correctness of the computed tracks.

Key words: Concurrent Programming, CSP, Semantics, Tracking.

1 Introduction

One of the most important techniques for program understanding and debugging is tracing [3]. A trace gives the user access to otherwise invisible information about a computation. In the context of concurrent languages, computations are particularly complex due to the non-deterministic execution order of processes and to the restrictions imposed on this order by synchronizations; and thus, a

* This work has been partially supported by the Spanish *Ministerio de Ciencia e Innovación* under grant TIN2008-06622-C03-02, by the *Generalitat Valenciana* under grant ACOMP/2009/017, and by the *Universidad Politécnic de Valencia* (Programs PAID-05-08 and PAID-06-08). Salvador Tamarit was partially supported by the Spanish MICINN under FPI grant BES-2009-015019.

tracer is a powerful tool to explore, understand and debug concurrent computations.

One of the most widespread concurrent specification languages is the *Communicating Sequential Processes* (CSP) [7, 16] whose operational semantics allows the combination of parallel, non-deterministic and non-terminating processes. The study and transformation of CSP specifications often uses different analyses such as deadlock analysis [10], reliability analysis [8] and program slicing [18] which are based on a data structure able to represent computations.

In CSP a trace is a sequence of events. Concretely, the operational semantics of CSP is an event-based semantics in which the occurrence of events fires the rules of the semantics. Hence, the final trace of the computation is the sequence of events occurred (see Chapter 8 of [16] for a detailed study of this kind of traces). In this work we introduce an essentially different notion of trace [3] called track. In our setting, a track is a data structure which represents the sequence of expressions that have been evaluated during the computation, and moreover, this data structure is labelled with the location of these expressions in the specification. Therefore, a CSP track is much more informative than a CSP trace since the former not only contains a lot of information about original program structures but also explicitly relates the sequence of events with the parts of the specification that caused these events.

Example 1. Consider the following CSP specification:¹

$$\begin{aligned} \text{MAIN} &= \text{CASINO} \parallel \text{GAMBLING} \\ \text{CASINO} &= (\text{PLAYER} \parallel \text{ROULETTE}) \parallel \text{CROUPIER} \\ &\quad \{ \text{betred, red, black, prize} \} \\ \text{PLAYER} &= \text{betred} \rightarrow (\text{prize} \rightarrow \text{STOP} \sqcap \text{noprize} \rightarrow \text{STOP}) \\ \text{ROULETTE} &= \text{red} \rightarrow \text{STOP} \sqcap \text{black} \rightarrow \text{STOP} \\ \text{CROUPIER} &= (\text{betred} \rightarrow \text{red} \rightarrow \text{prize} \rightarrow \text{STOP}) \\ &\quad \sqcap (\text{betred} \rightarrow \text{black} \rightarrow \text{prize} \rightarrow \text{STOP}) \\ &\quad \sqcap (\text{betblack} \rightarrow \text{black} \rightarrow \text{prize} \rightarrow \text{STOP}) \\ &\quad \sqcap (\text{betblack} \rightarrow \text{red} \rightarrow \text{getmoney} \rightarrow \text{STOP}) \\ \text{GAMBLING} &= \text{Complex Composite Processes} \end{aligned}$$

This specification models several gambling activities running in parallel and modelled by process `GAMBLING`. One of the games is the casino. A `CASINO` is modelled as the interaction of three parallel processes, namely a `PLAYER`, a `ROULETTE`, and a `CROUPIER`. The player bets for red, and she can win a prize or not. The roulette simply takes a color (either red or black); and the croupier checks the

¹ We refer those readers non familiar with CSP syntax to Section 2 where we provide a brief introduction to CSP.

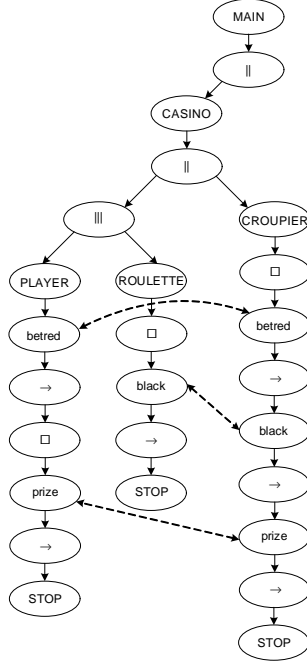


Fig. 1. Track of the program in Example 1

bet and the color of the roulette in order to give a prize to the player or just get the bet money.

This specification contains an error, because it allows the trace of events $t = \langle \mathbf{betred}, \mathbf{black}, \mathbf{prize} \rangle$ where the player bets for red and she wins a prize even though the roulette takes black.

Now assume that we execute the specification and discover the error after executing trace t . A track can be very useful to understand why the error was caused, and what part of the specification was involved in the wrong execution. For instance, if we look at the track of Fig. 1, we can easily see that the three processes run in parallel, and that the prize is given because there is a synchronization (dashed edges represent synchronizations) between **CROUPIER** and **PLAYER** that should never happen. Observe that the track is intuitive enough as to be a powerful program comprehension tool that provides much more information than the trace.

Moreover, observe that the track contains explicit information about the specification expressions that were involved in the execution. Therefore, it can be used for program slicing (see [17] for an explanation of the technique and [14] for an adaptation of program slicing to CSP). In particular, in this example, we can use the track to extract the part of the program that was involved in the execution—note that this is the only part that could cause the error—. This part has been underscored in the example. With a quick look, one can see that

the underscored part of process `CROUPIER` produced the wrong behavior. Event `prize` should be replaced by `getmoney`.

Another interesting application of tracks is related to component extraction and reuse. If we are interested in a particular trace, and we want to extract the part of the specification that models this trace to be used in another model, we can simply produce a slice, and slightly augment the code to make it syntactically correct (see [14] for an example and an explanation of this transformation). In our example, even though the system is very big due to the process `GAMBLING`, the track is able to extract the only information related to the trace.

We have implemented a tool [13] able to produce tracks and to automatically color parts of the code related to some point in the specification. This tool is integrated in the last version of ProB [11, 12] which is the most extended IDE for CSP.

In languages such as Haskell, the tracks (see, e.g., [3–5, 1]) are the basis of many analysis methods and tools. However, computing CSP tracks is a complex task due to the non-deterministic execution of processes, due to deadlocks, due to non-terminating processes and mainly due to synchronizations. This is probably the reason why no correctness result exists which formally relates the track of a specification to its execution. This semantics is needed because it would allow us to prove important properties (such as correctness and completeness) of the techniques and tools based on tracking.

To the best of our knowledge, there is only one attempt to define and build tracks for CSP [2]. Their notion of track is based on the standard *program dependence graph* [6]; therefore it is useful for program slicing but it is insufficient for other analyses that need a *context-sensitive graph* [9] (i.e., each different process call has a different representation). Moreover, their notion of track does not include synchronizations. Our tracks are able to represent synchronizations, and they are context-sensitive.

The main contributions of this work are the formal definition of tracks, the definition of the first tracking semantics for CSP and the proof that the trace of a computation can be extracted from the track of this computation. Concretely, we instrument the standard operational semantics of CSP in such a way that the execution of the semantics produces as a side-effect the track of the computation. It should be clear that the track of an infinite computation is also infinite. However, we design the semantics in such a way that the track is produced incrementally step by step. Therefore, if the execution is stopped (e.g., by the user because it is non-terminating or because a limit in the size of the track was specified), then the semantics produces the track of the computation performed so far. This semantics can serve as a theoretical foundation for tracking CSP computations because it formally relates the computations of the standard semantics with the tracks of these computations.

The rest of the paper has been organized as follows. Firstly, in Section 2 we recall the syntax and semantics of CSP. In Section 3 we define the concept of track for CSP. Then, in Section 4, we instrument the CSP semantics in such a way that its execution produces as a side-effect the track associated with the

performed computation. In Section 5, we present the main results of the paper proving that the instrumented semantics presented is a conservative extension of the standard semantics, its computed tracks are correct and the corresponding trace can be extracted from the track. Finally, Section 6 concludes.

2 The Syntax and Semantics of CSP

In order to make the paper self-contained, we recall in this section the syntax and semantics of CSP.

Figure 2 summarizes the syntax constructions used in CSP specifications. A *specification* is viewed as a finite set of process definitions. The left-hand side of each definition is the name of a process, which is defined in the right-hand side (abbrev. *rhs*) by means of an expression that can be a call to another process or a combination of the following operators:

- Prefixing It specifies that event a must happen before process P .
- Internal choice The system chooses non-deterministically to execute one of the two processes P or Q .
- External choice It is identical to internal choice but the choice comes from outside the system (e.g., the user).
- Sequential composition It specifies a sequence of two processes. When the first (successfully) finishes, the second starts.
- Synchronized parallelism Both processes are executed in parallel with a set X of synchronized events. In absence of synchronizations both processes can execute in any order. Whenever a synchronized event $a \in X$ happens in one of the processes it must also happen in the other at the same time. Whenever the set of synchronized events is not specified, it is assumed that processes are synchronized in all common events. A particular case of parallel execution is *interleaving* where no synchronizations exist (i.e., $X = \emptyset$).
- Skip It successfully finishes the current process. It allows us to continue the next sequential process.
- Stop Synonymous with deadlock: It finishes the current process and it does not allow the next sequential process to continue.

We now recall the standard operational semantics of CSP as defined by Roscoe [16]. It is presented in Fig. 3 as a logical inference system. A *state* of the semantics is a process to be evaluated called the *control*. The system starts with an initial state, and the rules of the semantics are used to infer how this state evolves. When no rules can be applied to the current state, the computation finishes. The rules of the semantics change the states of the computation due to the occurrence of events. The set of possible events is $\Sigma \cup \{\tau, \checkmark\}$. Events in $\Sigma = \{a, b, c, \dots\}$ are visible from the external environment, and can only happen with its co-operation (e.g., actions of the user). The special event τ cannot be observed from outside the system and it is an internal event that happens automatically as defined by the semantics. \checkmark is a special event representing the

$S ::= \{D_1, \dots, D_m\}$ (Entire specification)	<i>Domains</i> $M, N \dots \in \mathcal{N}$ (Process names) $P, Q \dots \in \mathcal{P}$ (Processes) $a, b \dots \in \Sigma$ (Events)
$D ::= N = P$ (Process definition)	
$P ::= M$ (Process call)	
$\quad a \rightarrow P$ (Prefixing)	
$\quad P \sqcap Q$ (Internal choice)	
$\quad P \square Q$ (External choice)	
$\quad P ; Q$ (Sequential composition)	
$\quad P \parallel_X Q$ (Synchronized parallelism)	where $X \subseteq \Sigma$
$\quad SKIP$ (Skip)	
$\quad STOP$ (Stop)	

Fig. 2. Syntax of CSP specifications

successful termination of a process. We use the special symbol Ω to denote any process that successfully terminated.

In order to perform computations, we construct an initial state (e.g., **MAIN**) and (non-deterministically) apply the rules of Fig. 3. The intuitive meaning of each rule is the following:

- (Process Call) The call is unfolded and the right-hand side of process named N is added to the control.
- (Prefixing) When event a occurs, process P is added to the control.
- (SKIP) After **SKIP**, the only possible event is \checkmark , which denotes the successful termination of the (sub)computation with the special symbol Ω . There is no rule for Ω (neither for **STOP**), hence, this (sub)computation has finished.
- (Internal Choice 1 and 2) The system, with the occurrence of the internal event τ , (non-deterministically) selects one of the two processes P or Q which is added to the control.
- (External Choice 1, 2, 3 and 4) The occurrence of τ develops one of the branches. The occurrence of an event $a \neq \tau$ is used to select one of the two processes P or Q and the control changes according to the event.
- (Sequential Composition 1) In $P;Q$, P can evolve to P' with any event except \checkmark . Hence, the control becomes $P';Q$.
- (Sequential Composition 2) When P successfully finishes (with event \checkmark), Q starts. Note that \checkmark is hidden from outside the whole process becoming τ .
- (Synchronized Parallelism 1 and 2) When event $a \notin X$ or events τ or \checkmark happen, one of the two processes P or Q evolves accordingly, but only a is visible from outside the parallelism operator.
- (Synchronized Parallelism 3) When event $a \in X$ happens, it is required that both processes synchronize, P and Q are executed at the same time and the control becomes $P' \parallel_X Q'$.

(Synchronized Parallelism 4) When both processes have successfully terminated the control becomes Ω , performing the event \checkmark .

(Process Call)	(Prefixing)	(SKIP)
$\frac{}{N \xrightarrow{\tau} rhs(N)}$	$\frac{}{(a \rightarrow P) \xrightarrow{a} P}$	$\frac{}{SKIP \xrightarrow{\checkmark} \Omega}$
(Internal Choice 1)	(Internal Choice 2)	
$\frac{}{(P \sqcap Q) \xrightarrow{\tau} P}$	$\frac{}{(P \sqcap Q) \xrightarrow{\tau} Q}$	
(External Choice 1)	(External Choice 2)	
$\frac{P \xrightarrow{\tau} P'}{(P \sqcap Q) \xrightarrow{\tau} (P' \sqcap Q)}$	$\frac{Q \xrightarrow{\tau} Q'}{(P \sqcap Q) \xrightarrow{\tau} (P \sqcap Q')}$	
(External Choice 3)	(External Choice 4)	
$\frac{P \xrightarrow{e} P'}{(P \sqcap Q) \xrightarrow{e} P'}$	$\frac{Q \xrightarrow{e} Q'}{(P \sqcap Q) \xrightarrow{e} Q'}$	$e \in \Sigma \cup \{\checkmark\}$
(Sequential Composition 1)	(Sequential Composition 2)	
$\frac{P \xrightarrow{e} P'}{(P; Q) \xrightarrow{e} (P'; Q)} \quad e \in \Sigma \cup \{\tau\}$	$\frac{P \xrightarrow{\checkmark} \Omega}{(P; Q) \xrightarrow{\tau} Q}$	
(Synchronized Parallelism 1)	(Synchronized Parallelism 2)	
$\frac{P \xrightarrow{e'} P'}{(P \parallel_X Q) \xrightarrow{e'} (P' \parallel_X Q)}$	$\frac{Q \xrightarrow{e'} Q'}{(P \parallel_X Q) \xrightarrow{e'} (P \parallel_X Q')}$	$(e = e' = a \wedge a \notin X)$ $\vee (e = \tau \wedge e' \in \{\tau, \checkmark\})$
(Synchronized Parallelism 3)	(Synchronized Parallelism 4)	
$\frac{P \xrightarrow{a} P' \quad Q \xrightarrow{a} Q'}{(P \parallel_X Q) \xrightarrow{a} (P' \parallel_X Q')} \quad a \in X$	$\frac{}{(\Omega \parallel_X \Omega) \xrightarrow{\checkmark} \Omega}$	

Fig. 3. CSP's operational semantics

We illustrate the semantics with the following example.

Example 2. Consider the next CSP specification:

$$\begin{aligned} \text{MAIN} &= (\mathbf{a} \rightarrow \text{STOP}) \parallel \{ \mathbf{a} \} (\mathbf{P} \square (\mathbf{a} \rightarrow \text{STOP})) \\ \mathbf{P} &= \mathbf{b} \rightarrow \text{SKIP} \end{aligned}$$

If we use **MAIN** as the initial state to execute the semantics, we get the computation shown in Fig. 4 where the final state is $((\mathbf{a} \rightarrow \text{STOP}) \parallel \Omega)$. This computation

corresponds to the execution of the left branch of the choice (i.e., **P**) and thus only event **b** occurs. Each rewriting step is labelled with the applied rule, and the example should be read from top to bottom.

$$\begin{array}{c} \text{(Process Call)} \\ \hline \text{MAIN} \xrightarrow{\tau} ((\mathbf{a} \rightarrow \text{STOP}) \parallel \{ \mathbf{a} \} (\mathbf{P} \square (\mathbf{a} \rightarrow \text{STOP}))) \\ \hline \text{(Synchronized Parallelism 2)} \quad \text{(External Choice 1)} \quad \frac{\text{(Process Call)} \quad \frac{\mathbf{P} \xrightarrow{\tau} (\mathbf{b} \rightarrow \text{SKIP})}{(\mathbf{P} \square (\mathbf{a} \rightarrow \text{STOP})) \xrightarrow{\tau} ((\mathbf{b} \rightarrow \text{SKIP}) \square (\mathbf{a} \rightarrow \text{STOP}))}}{((\mathbf{a} \rightarrow \text{STOP}) \parallel \{ \mathbf{a} \} (\mathbf{P} \square (\mathbf{a} \rightarrow \text{STOP}))) \xrightarrow{\tau} \text{State}_1}}{\text{where } \text{State}_1 = ((\mathbf{a} \rightarrow \text{STOP}) \parallel \{ \mathbf{a} \} ((\mathbf{b} \rightarrow \text{SKIP}) \square (\mathbf{a} \rightarrow \text{STOP})))} \\ \hline \text{(Synchronized Parallelism 2)} \quad \text{(External Choice 3)} \quad \frac{\text{(Prefixing)} \quad \frac{(\mathbf{b} \rightarrow \text{SKIP}) \xrightarrow{\mathbf{b}} \text{SKIP}}{((\mathbf{b} \rightarrow \text{SKIP}) \square (\mathbf{a} \rightarrow \text{STOP})) \xrightarrow{\mathbf{b}} \text{SKIP}}}{\text{State}_1 \xrightarrow{\mathbf{b}} ((\mathbf{a} \rightarrow \text{STOP}) \parallel \{ \mathbf{a} \} \text{SKIP})} \\ \hline \text{(Synchronized Parallelism 2)} \quad \text{(SKIP)} \quad \frac{\text{SKIP} \xrightarrow{\checkmark} \Omega}{((\mathbf{a} \rightarrow \text{STOP}) \parallel \{ \mathbf{a} \} \text{SKIP}) \xrightarrow{\tau} ((\mathbf{a} \rightarrow \text{STOP}) \parallel \{ \mathbf{a} \} \Omega)} \end{array}$$

Fig. 4. A computation with the operational semantics in Fig. 3

3 Tracking Computations

In this section we define the notion of track. Firstly, we introduce some notation that will be used throughout the paper.

A track is formed by the sequence of expressions that are evaluated during an execution. These expressions are conveniently connected to form a graph. However, several program analysis techniques such as program slicing make use of the locations of program expressions, and thus, this notion of track is insufficient for them. Therefore, we want our tracks to also store the location of each literal (i.e., events, operators and process names) in the specification so that the track can be used to know what portions of the source code have been executed and in what order. The inclusion of source positions in the track implies an additional level of complexity in the semantics, but the benefits of providing our tracks with this additional information are clear and, for some applications, essential. Therefore, we use labels (that we call *specification positions*) to uniquely identify each literal in a specification which roughly corresponds to nodes in the CSP specification's abstract syntax tree. We define a function \mathcal{Pos} to obtain the specification position of an element of a CSP specification and it is defined over nodes of an abstract syntax tree for a CSP specification. Formally,

Definition 1. (*Specification position*) A specification position is a pair (N, w) where $N \in \mathcal{N}$ and w is a sequence of natural numbers (we use Λ to denote the empty sequence). We let $\mathcal{Pos}(o)$ denote the specification position of an expression o . Each process definition $N = P$ of a CSP specification is labelled with specification positions. The specification position of its left-hand side is $\mathcal{Pos}(N) = (N, 0)$. The right-hand side is labelled with the call $\text{AddSpPos}(P, (N, \Lambda))$; where function AddSpPos is defined as follows:

$$\text{AddSpPos}(P, (N, w)) = \begin{cases} P_{(N,w)} & \text{if } P \in \mathcal{N} \\ \text{STOP}_{(N,w)} & \text{if } P = \text{STOP} \\ \text{SKIP}_{(N,w)} & \text{if } P = \text{SKIP} \\ a_{(N,w.1)} \rightarrow_{(N,w)} \text{AddSpPos}(Q, (N, w.2)) & \text{if } P = a \rightarrow Q \\ \text{AddSpPos}(Q, (N, w.1)) \text{ op}_{(N,w)} \text{AddSpPos}(R, (N, w.2)) & \text{if } P = Q \text{ op } R \quad \forall \text{op} \in \{\square, \square, ||, ;\} \end{cases}$$

Example 3. Consider again the CSP specification in Example 2 where literals are labelled with their associated specification positions (they are underlined) so that labels are unique:

$$\begin{aligned} \text{MAIN}_{(\text{MAIN},0)} &= (\underline{a}_{(\text{MAIN},1.1)} \rightarrow_{(\text{MAIN},1)} \text{STOP}_{(\text{MAIN},1.2)}) \parallel_{\{\underline{a}\}} (\underline{P}_{(\text{MAIN},2.1)} \square_{(\text{MAIN},2)} (\underline{a}_{(\text{MAIN},2.2.1)} \rightarrow_{(\text{MAIN},2.2)} \text{STOP}_{(\text{MAIN},2.2.2)})) \\ \underline{P}_{(\text{P},0)} &= \underline{b}_{(\text{P},1)} \rightarrow_{(\text{P},\Lambda)} \text{SKIP}_{(\text{P},2)} \end{aligned}$$

In the following, specification positions will be represented with greek letters (α, β, \dots) and we will often use indistinguishably an expression and its associated specification position when it is clear from the context (e.g., in Example 3 we will refer to $(P, 1)$ as \mathbf{b}).

In order to introduce the formal definition of track, we need first to define the concept of *control-flow*, which refers to the order in which the individual literals of a CSP specification are executed. Intuitively, the control can pass from a specification position α to a specification position β iff an execution exists where α is executed before β . This notion of control-flow is similar to the control-flow used in the *control-flow graphs* (CFG) [17] of imperative programming. We have adapted the same idea to CSP where choices and parallel composition appear; and in a similar way to the CFG, we use this definition to draw control arcs in our tracks. Formally,

Definition 2. (*Static control-flow*) Given a CSP specification \mathcal{S} and two specification positions α, β in \mathcal{S} , we say that the control can pass from α to β , denoted by $\alpha \Rightarrow \beta$, iff one of the following conditions holds:

- i) $\alpha = N \wedge \beta = \text{first}((N, A))$ with $N = \text{rhs}(N) \in \mathcal{S}$
- ii) $\alpha \in \{\square, \square, \|\}$ $\wedge \beta \in \{\text{first}(\alpha.1), \text{first}(\alpha.2)\}$
- iii) $\alpha \in \{\rightarrow, ;\}$ $\wedge \beta = \text{first}(\alpha.2)$
- iv) $\alpha = \beta.1 \wedge \beta = \rightarrow$
- v) $\alpha \in \text{last}(\beta.1) \wedge \beta = ;$

where $\text{first}(\alpha)$ is the specification position of the subprocess denoted by α which must be executed first:

$$\text{first}(\alpha) = \begin{cases} \alpha.1 & \text{if } \alpha = \rightarrow \\ \text{first}(\alpha.1) & \text{if } \alpha = ; \\ \alpha & \text{otherwise} \end{cases}$$

and $\text{last}(\alpha)$ is the set of all possible termination points of the subprocess denoted by α :

$$\text{last}(\alpha) = \begin{cases} \{\alpha\} & \text{if } \alpha = \text{SKIP} \\ \emptyset & \text{if } \alpha = \text{STOP} \vee \\ & (\alpha \in \{\|\}\ \wedge (\text{last}(\alpha.1) = \emptyset \vee \text{last}(\alpha.2) = \emptyset)) \\ \text{last}(\alpha.1) \cup \text{last}(\alpha.2) & \text{if } \alpha \in \{\square, \square\} \vee \\ & (\alpha \in \{\|\}\ \wedge \text{last}(\alpha.1) \neq \emptyset \wedge \text{last}(\alpha.2) \neq \emptyset) \\ \text{last}(\alpha.2) & \text{if } \alpha \in \{\rightarrow, ;\} \\ \text{last}((N, A)) & \text{if } \alpha = N \end{cases}$$

For instance, in Example 3, we can see how the control can pass from a specification position to another one, e.g., we have $(\text{MAIN}, 2) \Rightarrow (\text{MAIN}, 2.1)$ and $(\text{MAIN}, 2) \Rightarrow (\text{MAIN}, 2.2.1)$ due to rule ii). And $(\text{MAIN}, 2.2.1) \Rightarrow (\text{MAIN}, 2.2)$ due to rule iv); $(\text{MAIN}, 2.2) \Rightarrow (\text{MAIN}, 2.2.2)$ due to rule iii) and $(\text{MAIN}, 2.1) \Rightarrow (P, 1)$ due to rule i).

We also need to define the notions of *rewriting step* and *derivation*.

Definition 3. (*Rewriting Step, Derivation*) Given a CSP process P , a rewriting step for P , denoted by $P \xrightarrow{\Theta} P'$, is the transformation of P into P' by using a rule of the CSP semantics. Therefore, $P \xrightarrow{\Theta} P'$ iff a rule of the form $\frac{\Theta}{P \xrightarrow{e} P'}$ is applicable, where $e \in \Sigma \cup \{\tau, \checkmark\}$ and Θ is a (possibly empty) set of rewriting steps. Given a CSP process P_0 , we say that the sequence $P_0 \xrightarrow{\Theta_0} \dots \xrightarrow{\Theta_n} P_{n+1}$, $n \geq 0$, is a derivation of P_0 iff $\forall i, 0 \leq i \leq n, P_i \xrightarrow{\Theta_i} P_{i+1}$ is a rewriting step. We say that the derivation is complete iff there is no possible rewriting step for P_{n+1} . We say that the derivation has successfully finished iff P_{n+1} is Ω .

For instance, in Fig. 5(a), one (possible) complete derivation of Example 3 is shown (for the time being, the reader can ignore the underlined part). The rules applied in each rewriting step (ignoring subderivations) are (Process Call) and (Synchronized Parallelism 3) (abbrev. (PC) and (SP3), respectively).

Function *last* of Definition 2 can be used to determine the last specification position in a derivation. However, this function computes all possible final specification positions, and a derivation only reaches (non-deterministically) a set of them. Therefore, we will use in the following a modified version of *last* called *last'* whose behaviour is exactly the same as *last* except in the case of choices where only one of the branches is selected:

For each derivation $(P \sqcap P' \xrightarrow{\Theta} P)$ or $(P \sqcap P' \xrightarrow{\Theta_0} \dots \xrightarrow{\Theta_n} P'')$, $n \geq 0$ such that $P \xrightarrow{\Theta'_0} \dots \xrightarrow{\Theta'_m} P''$, $m \geq 0$, $last'(P \sqcap P') = last'(P \sqcap P') = last'(P)$.

Note that, while *last* is static, *last'* is dynamic; it is defined in the context of a particular derivation which implies one particular way of resolving any non-determinism. The same happens with the definition of control-flow. Control-flow is defined statically and says whether the control can pass from α to β in some derivation. However, the track is a dynamic structure produced for a particular derivation. Therefore, we produce a dynamic version of the definition of control-flow which is defined for a particular derivation.

Definition 4. (*Dynamic control-flow*) Let \mathcal{S} be a CSP specification and \mathcal{D} a derivation in \mathcal{S} . Given two specification positions α, β in \mathcal{S} , we say that the control can dynamically pass from α to β , denoted by $\alpha \Rightarrow \beta$, iff the control can pass from α to β ($\alpha \Rightarrow \beta$) in derivation \mathcal{D} . For each $P \xrightarrow{\Theta} P' \in \mathcal{D}$ and for all rewriting steps in Θ , we have that:

1. if P is a prefixing ($a \rightarrow Q$) or a sequential composition ($Q; R$), then $\mathcal{P}os(a) \Rightarrow \mathcal{P}os(\rightarrow)$ or $\forall p \in last'(Q), \mathcal{P}os(p) \Rightarrow \mathcal{P}os(;$ respectively,
2. if $P \Rightarrow first(P'')$ where $P'' \xrightarrow{\Theta'} P''' \in \Theta$, then $\mathcal{P}os(P) \Rightarrow \mathcal{P}os(first(P''))$,
3. if $P \Rightarrow first(P')$, then $\mathcal{P}os(P) \Rightarrow \mathcal{P}os(first(P'))$.

Clauses 1, 2 and 3 define the cases in which the control passes between two specification positions in a given derivation. In clause 1, if we have a prefixing in the control then Θ is empty and the rewriting step applied is of the form

$\frac{}{(a \rightarrow P) \xrightarrow{a} P}$. In this case, clause 1 guarantees that the control can dynamically pass from a to \rightarrow ; and clause 3 guarantees that the control can dynamically pass from \rightarrow to P . However, in general, Θ is not empty, and the rewriting step is of the form $\frac{P'' \rightarrow P'''}{P \rightarrow P'}$. Here, clause 2 ensures that the control can dynamically pass from P to P'' ; and clause 3 ensures that the control can dynamically pass from P to P' and from P'' to P''' . For instance, it is possible that we have a rewriting step to evaluate the process $P \square P'$. Clearly, the control can pass from \square to both P and P' ($\square \Rightarrow P$ and $\square \Rightarrow P'$), but in the rewriting step the control will only pass to one of them ($\square \Rightarrow P$ or $\square \Rightarrow P'$). In this case, clauses 2 and 3 are used.

We are now in a position to formally define the concept of *track* of a derivation.

Definition 5. (*Track*) Given a CSP specification \mathcal{S} , and a derivation \mathcal{D} in \mathcal{S} , the track of \mathcal{D} is a graph $\mathcal{G} = (N, E_c, E_s)$ where N is a set of nodes uniquely identified with a natural number and that are labelled with specification positions ($l(n)$ refers to the label of node n), and edges are divided into two groups:

- control-flow edges (E_c) are a set of one-way edges (denoted with \mapsto) representing the control-flow between two nodes, and
- synchronization edges (E_s) are a set of two-way edges (denoted with \leftrightarrow) representing the synchronization of two (event) nodes;

and

1. E_c contains a control-flow edge $a \mapsto a'$ iff $a \Rightarrow a'$ with respect to \mathcal{D} , and
2. E_s contains a synchronization edge $a \leftrightarrow a'$ for each synchronization occurring in \mathcal{D} where a and a' are the nodes of the synchronized events.

The only nodes in N are the nodes induced by E_c and E_s .

Example 4. Consider again the specification of Example 3. We show in Fig. 5(a) one possible derivation (ignoring subderivations) of this specification (for the time being, the underlined part should be ignored). Its associated track is shown in Fig. 5(b). In the example, we see that the track is a connected and directed graph. Apart from the control-flow edges, there is one synchronization edge between nodes (MAIN, 1.1) and (MAIN, 2.2.1) representing the synchronization of event **a**. To illustrate the inclusion of edges in Definition 5, we see that the edge between nodes 2 and 3 is introduced according to clause 1 of Definition 4; the edge between nodes 5 and 6 is introduced according to clause 2 of Definition 4 because, in the subderivations of (SP3), there is a rewriting

ing step $\frac{\text{(External Choice 4)} \frac{\text{(Prefixing)} \frac{}{(a \rightarrow \text{STOP}) \xrightarrow{a} (\text{STOP})}}{(P \square (a \rightarrow \text{STOP})) \xrightarrow{a} (\text{STOP})}}{}{} \quad \text{and } \text{first}(a \rightarrow \text{STOP}) = \mathbf{a};$

the edge between nodes 7 and 8 is introduced according to clause 3 of Definition 4 because there is also a rewriting step $\frac{\text{(Prefixing)} \frac{}{(a \rightarrow \text{STOP}) \xrightarrow{a} (\text{STOP})}}{}{} \quad \text{and}$

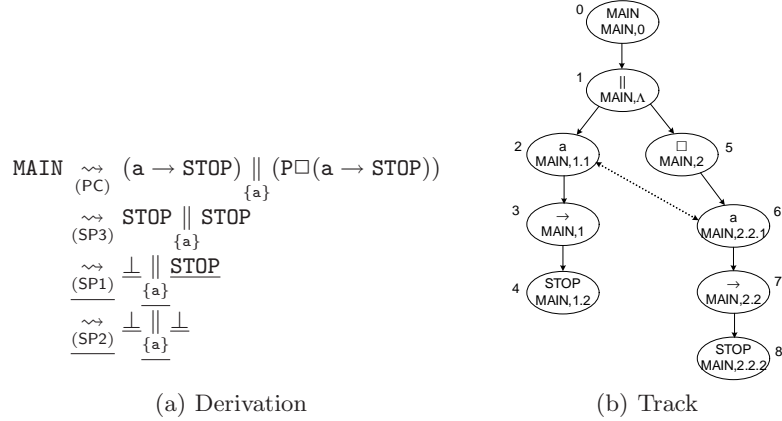


Fig. 5. Derivation and track associated with the specification of Example 3

$first(\text{STOP}) = \text{STOP}$; and the synchronization edge between nodes 2 and 6 is introduced according to clause 2 of Definition 5.

The trace associated with the derivation in Fig. 5(a) is $\langle \text{a} \rangle$. Therefore, note that the track is much more informative: it shows the exact processes that have been evaluated with an explicit causality relation; and, in addition, it shows the specification positions that have been evaluated and in what order.

4 Instrumenting the Semantics for Tracking

The generation of tracks in CSP introduces new challenges such as non-deterministic execution of processes, deadlocks, non-terminating processes and synchronizations. In this work, we design a solution that overcomes these difficulties. Firstly, we generate tracks with an augmented semantics which is conservative with respect to the standard operational semantics. Therefore, the execution order is the standard order, thus non-determinism and synchronizations are solved by the semantics. Moreover, the semantics generates the track incrementally, step by step. Therefore, infinite computations can be tracked until they are stopped. Hence, it is not needed to actually finish a computation to get the track of the subcomputations performed.

Example 5. In the following CSP specification two non-terminating processes run in parallel and synchronize infinitely.

$$\begin{aligned}
\text{MAIN}_{(\text{MAIN},0)} &= \text{P}_{(\text{MAIN},1)} \parallel \{ \text{a} \} (\text{MAIN},1) \text{P}_{(\text{MAIN},2)} \\
\text{P}_{(\text{P},0)} &= \text{a}_{(\text{P},1)} \rightarrow (\text{P},1) \text{P}_{(\text{P},2)}
\end{aligned}$$

Because the computation is infinite, the track (shown in Fig. 6) is also infinite.

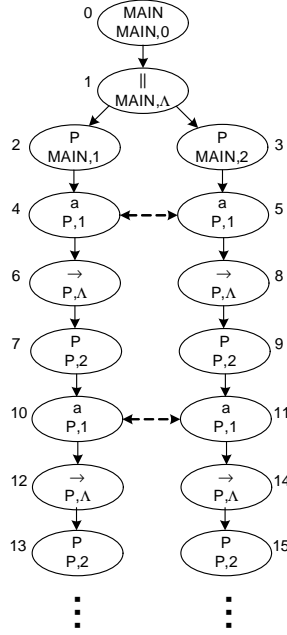


Fig. 6. Track of the program in Example 5

In order to solve the problem of deadlocks (that stop the computation), and have a representation for them in the tracks; when a deadlock happens, the semantics performs some additional steps to be able to generate a part of the track that represents the deadlock. These additional steps do not influence the other rules of the semantics, thus it remains conservative.

This section introduces an instrumented operational semantics of CSP which generates as a side-effect the tracks associated with the computations performed with the semantics. The tracking semantics is shown in Fig. 7, where we assume that every literal in the program has been labelled with its specification position (denoted by a subscript, e.g., P_α). In this semantics, a *state* is a tuple (P, G, m, Δ) , where P is the process to be evaluated (the *control*), G is a directed graph (i.e., the track built so far), m is a numeric reference to the current node in G , and Δ is a set of references to nodes that may be synchronized. Concretely, m references the node in G where the specification position of the control P must be stored. Reference m is a fresh² reference generated to add new nodes to G . The basic idea of the graph construction is to record the current control with the current reference in every step by connecting it to its parent. We use the notation $G[m \xrightarrow{n} \alpha]$ to introduce a node in G . For instance, if we are adding a node to G this new node has reference m , it is labelled with specification position α , and its successor is n (a fresh reference). Successor arrows are denoted

² We assume that fresh references are numeric and generated incrementally.

by $m \xrightarrow[n]{}$ which means that node n is the successor of node m . Every time an event in Σ happens during the computation, this event is stored in the set Δ of the current state. Therefore, when a synchronized parallelism is evaluated, all the events that must be synchronized are in Δ . We use the special symbol \perp to denote any process that is deadlocked. In order to perform computations, we construct an initial state (e.g., $(\text{MAIN}, \emptyset, 0, \emptyset)$) and (non-deterministically) apply the rules of Fig. 7. When the execution has finished or has been interrupted, the semantics has produced the track of the computation performed so far.

An explanation for each rule of the semantics follows:

- (Process Call) The called process N is unfolded, node m is added to the graph with specification position α and successor n (a fresh reference). The new process in the control is $rhs(N)$. The set Δ of events to be synchronized is put to \emptyset .
- (Prefixing) This rule adds nodes m (the prefix) and n (the prefixing operator) to the graph. In the new state, n becomes the parent reference and the fresh reference p represents the current reference. The new control is P . The set Δ is $\{m\}$ to indicate that event a has occurred and it must be synchronized when required by (Synchronized Parallelism 3).
- (SKIP and STOP) Whenever one of these rules is applied, the subcomputation finishes because Ω (for rule SKIP) and \perp (for rule STOP) are put in the control, and these special symbols have no associated rule. A node with the SKIP (respectively STOP) specification position is added to the graph.
- (Internal Choice 1 and 2) The choice operator is added to the graph, and the (non-deterministically) selected branch is put into the control with the fresh reference n as the successor of the choice operator.
- (External Choice 1, 2, 3 and 4) External choices can develop both branches while τ events happen (rules 1 and 2), until an event in $\Sigma \cup \{\checkmark\}$ occurs (rules 3 and 4). This means that the semantics can add nodes to both branches of the track alternatively, and thus, it needs to store the next reference to use in every branch of the choice. This is done by labelling choice operators with a tuple of the form (α, n_1, n_2) where α is the specification position of the choice operator; and n_1 and n_2 are respectively the references to be used in the left and right branches of the choice, and they are initialized to \bullet , a symbol used to express that the branch has not been evaluated yet. Therefore, the first time a branch is evaluated, we generate a new reference for this branch. For this purpose, function `FirstEval` is used:

$$\text{FirstEval}(G, n, m, \alpha) = \begin{cases} (G[m \xrightarrow[p]{\alpha}], p) & \text{if } n = \bullet \\ (G, n) & \text{otherwise} \end{cases}$$

This function checks whether this is the first time that the branch is evaluated (this only happens when the reference of this branch is empty, i.e., $n = \bullet$). In this case, the choice operator is added to G . For instance, consider the rewriting step (EC4) of Fig. 8. The choice operator in the rewriting step R is labelled with $((\text{MAIN}, A), \bullet, \bullet)$. Therefore, it is evaluated for the first

(Process Call)	$\frac{}{(N_\alpha, G, m, \Delta) \xrightarrow{\tau} (rhs(N), G[m \mapsto_n \alpha], n, \emptyset)}$
(Prefixing)	$\frac{}{(a_\alpha \rightarrow_\beta P, G, m, \Delta) \xrightarrow{a} (P, G[m \mapsto_n \alpha, n \mapsto_p \beta], p, \{m\})}$
(SKIP)	$\frac{}{(SKIP_\alpha, G, m, \Delta) \xrightarrow{\surd} (\Omega, G[m \mapsto_n \alpha], n, \emptyset)}$
(STOP)	$\frac{}{(STOP_\alpha, G, m, \Delta) \xrightarrow{\tau} (\perp, G[m \mapsto_n \alpha], n, \emptyset)}$
(Internal Choice 1)	$\frac{}{(P \sqcap_\alpha Q, G, m, \Delta) \xrightarrow{\tau} (P, G[m \mapsto_n \alpha], n, \emptyset)}$
(Internal Choice 2)	$\frac{}{(P \sqcap_\alpha Q, G, m, \Delta) \xrightarrow{\tau} (Q, G[m \mapsto_n \alpha], n, \emptyset)}$
(External Choice 1)	$\frac{(P_1, G', n', \Delta) \xrightarrow{\tau} (P', G'', n'', \emptyset)}{(P_1 \sqcap_{(\alpha, n_1, n_2)} P_2, G, m, \Delta) \xrightarrow{\tau} (P' \sqcap_{(\alpha, n'', n_2)} P_2, G'', m, \emptyset)}$ where $(G', n') = \text{FirstEval}(G, n_1, m, \alpha)$
(External Choice 2)	$\frac{(P_2, G', n', \Delta) \xrightarrow{\tau} (P', G'', n'', \emptyset)}{(P_1 \sqcap_{(\alpha, n_1, n_2)} P_2, G, m, \Delta) \xrightarrow{\tau} (P_1 \sqcap_{(\alpha, n_1, n'')} P', G'', m, \emptyset)}$ where $(G', n') = \text{FirstEval}(G, n_2, m, \alpha)$
(External Choice 3)	$\frac{(P_1, G', n', \Delta) \xrightarrow{e} (P', G'', n'', \Delta')}{(P_1 \sqcap_{(\alpha, n_1, n_2)} P_2, G, m, \Delta) \xrightarrow{e} (P', G'', n'', \Delta')} \quad e \in \Sigma \cup \{\surd\}$ where $(G', n') = \text{FirstEval}(G, n_1, m, \alpha)$
(External Choice 4)	$\frac{(P_2, G', n', \Delta) \xrightarrow{e} (P', G'', n'', \Delta')}{(P_1 \sqcap_{(\alpha, n_1, n_2)} P_2, G, m, \Delta) \xrightarrow{e} (P', G'', n'', \Delta')} \quad e \in \Sigma \cup \{\surd\}$ where $(G', n') = \text{FirstEval}(G, n_2, m, \alpha)$
(Sequential Composition 1)	$\frac{(P, G, m, \Delta) \xrightarrow{e} (P', G', m', \Delta')}{(P; Q, G, m, \Delta) \xrightarrow{e} (P'; Q, G', m', \Delta')} \quad e \in \Sigma \cup \{\tau\}$
(Sequential Composition 2)	$\frac{(P, G, m, \Delta) \xrightarrow{\surd} (\Omega, G', n, \emptyset)}{(P;_\alpha Q, G, m, \Delta) \xrightarrow{\tau} (Q, G'[n \mapsto_p \alpha], p, \emptyset)}$

Fig. 7. An instrumented operational semantics to generate CSP tracks

time, and thus, in the left-hand side state of the upper rewriting step, node $5 \mapsto_6 (\text{MAIN}, 2)$, which refers to the choice operator, is added to G .

(Sequential Composition 1 and 2) Sequential Composition 1 is used to evolve process P until it is finished. P is evolved to P' which is put into the control. When P successfully finishes (it becomes Ω), \surd happens. Then, Sequential Composition 2 is used and Q is put into the control. The sequential composition operator $;$ is added to the graph with successor p that is the reference to be used in the first node added in the subderivation associated with Q .

(Synchronized Parallelism 1)	$\frac{(P_1, G', n', \Delta) \xrightarrow{e'} (P', G'', n'', \Delta')}{(P_1 \parallel_X^{(\alpha, n_1, n_2)} P_2, G, m, \Delta) \xrightarrow{e} (P' \parallel_X^{(\alpha, n'', n_2)} P_2, G'', m, \Delta')}$ $\vee (e = e' = a \wedge a \notin X) \vee (e = \tau \wedge e' \in \{\tau, \checkmark\})$ <p style="margin-left: 20px;">where $(G', n') = \text{FirstEval}(G, n_1, m, \alpha)$</p>
(Synchronized Parallelism 2)	$\frac{(P_2, G', n', \Delta) \xrightarrow{e'} (P', G'', n'', \Delta')}{(P_1 \parallel_X^{(\alpha, n_1, n_2)} P_2, G, m, \Delta) \xrightarrow{e} (P_1 \parallel_X^{(\alpha, n_1, n'')} P', G'', m, \Delta')}$ $\vee (e = e' = a \wedge a \notin X) \vee (e = \tau \wedge e' \in \{\tau, \checkmark\})$ <p style="margin-left: 20px;">where $(G', n') = \text{FirstEval}(G, n_2, m, \alpha)$</p>
(Synchronized Parallelism 3)	$\frac{\text{RewritingStep}_1 \quad \text{RewritingStep}_2}{(P_1 \parallel_X^{(\alpha, n_1, n_2)} P_2, G, m, \Delta) \xrightarrow{a} (P_1 \parallel_X^{(\alpha, n'_1, n'_2)} P_2', G'', m, \Delta_1 \cup \Delta_2)} \quad a \in X$ <p style="margin-left: 20px;">where $G'' = G'_1 \cup G'_2 \cup \{s_1 \stackrel{a}{\leftrightarrow} s_2 \mid s_1 \in \Delta_1 \wedge s_2 \in \Delta_2\}$</p> <p style="margin-left: 40px;">$\wedge \text{RewritingStep}_1 = (P_1, G'_1, n'_1, \Delta) \xrightarrow{a} (P'_1, G''_1, n''_1, \Delta_1)$</p> <p style="margin-left: 40px;">$\wedge (G'_1, n'_1) = \text{FirstEval}(G, n_1, m, \alpha)$</p> <p style="margin-left: 40px;">$\wedge \text{RewritingStep}_2 = (P_2, G'_2, n'_2, \Delta) \xrightarrow{a} (P'_2, G''_2, n''_2, \Delta_2)$</p> <p style="margin-left: 40px;">$\wedge (G'_2, n'_2) = \text{FirstEval}(G, n_2, m, \alpha)$</p>
(Synchronized Parallelism 4)	$\frac{}{(\Omega \parallel_X^{(\alpha, n_1, n_2)} \Omega, G, m, \Delta) \xrightarrow{\checkmark} (\Omega, G', r, \emptyset)}$ <p style="margin-left: 20px;">where $G' = G[\{p \xrightarrow{r} \mid p \xrightarrow{q} \in G \text{ where } q \in \{n_1, n_2\}\}]$</p>

Fig. 7. An instrumented operational semantics to generate CSP tracks (cont.)

(Synchronized Parallelism 1 and 2) In a synchronized parallel composition, both parallel processes can be intertwiningly executed until a synchronized event is found. Therefore, nodes from both processes can be added interwoven to the graph. Hence, each parallelism operator is labelled with a tuple of the form (α, n_1, n_2) as it happens with external choices.

These rules develop the branches of the parallelism until they are finished or until they must synchronize. In order to introduce the parallelism operator into the graph, function `FirstEval` is used, as it happens in the external choice rules. For instance, consider the rewriting step (Synchronized Parallelism 3) of Fig. 8. The parallelism operator in the rewriting step *State 1* is labelled with $((\text{MAIN}, \lambda), \bullet, \bullet)$. Therefore, it is evaluated for the first time, and thus, in the left-hand side state of the rewriting step L , node $1 \xrightarrow{2} (\text{MAIN}, \lambda)$, which refers to the parallelism operator, is added to G .

(Synchronized Parallelism 3) This rule is used to synchronize the parallel processes. In this case, both branches must perform a rewriting step with the same visible (and synchronized) event. Each branch derivation has a non-empty set of events (Δ_1, Δ_2) to be synchronized (note that this is a set because many parallelisms could be nested). Then, all references in the sets Δ_1 and Δ_2 are mutually linked with synchronization edges. Both sets are joined to form the new set of synchronized events.

(Synchronized Parallelism 4) It is used when none of the parallel processes can proceed because they already successfully finished. In this case, the control becomes Ω indicating the successful termination of the synchronized parallelism. In the new state, the new (fresh) reference is r . This rule also adds to the graph the arcs from all the parents of the last references of each branch (n_1 and n_2) to r . Here, we use the notation $p \xrightarrow[r]$ to add an edge from p to r . Note that the fact of generating the next reference in each rule allows (Synchronized Parallelism 4) to connect the final node of both branches to the next node. This simplifies other rules such as (Sequential Composition) that already has the reference of the node ready.

We illustrate this semantics with a simple example.

Example 6. Consider again the specification in Example 3. Figure 5(a) shows one possible derivation (excluding subderivations) for this example. Note that the underlined part corresponds to the additional rewriting steps performed by the tracking semantics. This derivation corresponds to the execution of the instrumented semantics with the initial state $(\text{MAIN}, \emptyset, 0, \emptyset)$ shown in Fig. 8. Here, for clarity, each computation step is labelled with the applied rule; in each state, G denotes the current graph. This computation corresponds to the execution of the right branch of the choice (i.e., $\mathbf{a} \rightarrow \text{STOP}$). The final state is $(\perp \parallel_{\{\mathbf{a}\}} ((\text{MAIN}, \mathcal{A}), 9, 10) \perp, G', 1, \emptyset)$. The final track G' computed for this execution

is depicted in Fig. 5(b) where we can see that nodes are numbered with the references generated by the instrumented semantics. Note that nodes 9 and 10 were prepared by the semantics (edges to them were produced) but never used because the subcomputations were stopped in STOP . Note also that the track contains all the parts of the specification executed by the semantics. This means that if the left branch of the choice had been developed (i.e., unfolding the call to P , thus using rule (External Choice 3)), this branch would also belong to the track.

We show another example whose execution finishes with a deadlock and where a non-terminating process appears.

Example 7. Consider the following CSP specification where each (sub)process has been labelled with its associated specification position (underlining).

$$\begin{aligned} \text{MAIN}_{(\text{MAIN}, 0)} &= \underline{\mathbf{a}_{(\text{MAIN}, 1.1)} \rightarrow_{(\text{MAIN}, 1)} \mathbf{a}_{(\text{MAIN}, 1.2.1)} \rightarrow_{(\text{MAIN}, 1.2)} \text{STOP}_{(\text{MAIN}, 1.2.2)}} \\ &\quad \parallel_{\{\mathbf{a}\}} \underline{\text{P}_{(\text{MAIN}, 2)}} \\ \text{P}_{(\text{P}, 0)} &= \underline{\mathbf{a}_{(\text{P}, 1)} \rightarrow_{(\text{P}, \mathcal{A})} \text{P}_{(\text{P}, 2)}} \end{aligned}$$

We use the initial state $(\text{MAIN}, \emptyset, 0, \emptyset)$ to execute the semantics and get the computation of Fig. 10. The final state is $(\perp \parallel_{\{\mathbf{a}\}} ((\text{MAIN}, \mathcal{A}), 14, 15) (\mathbf{a} \rightarrow \text{P}), G', 1, \emptyset)$. The final track G' computed is the graph of Fig. 9.

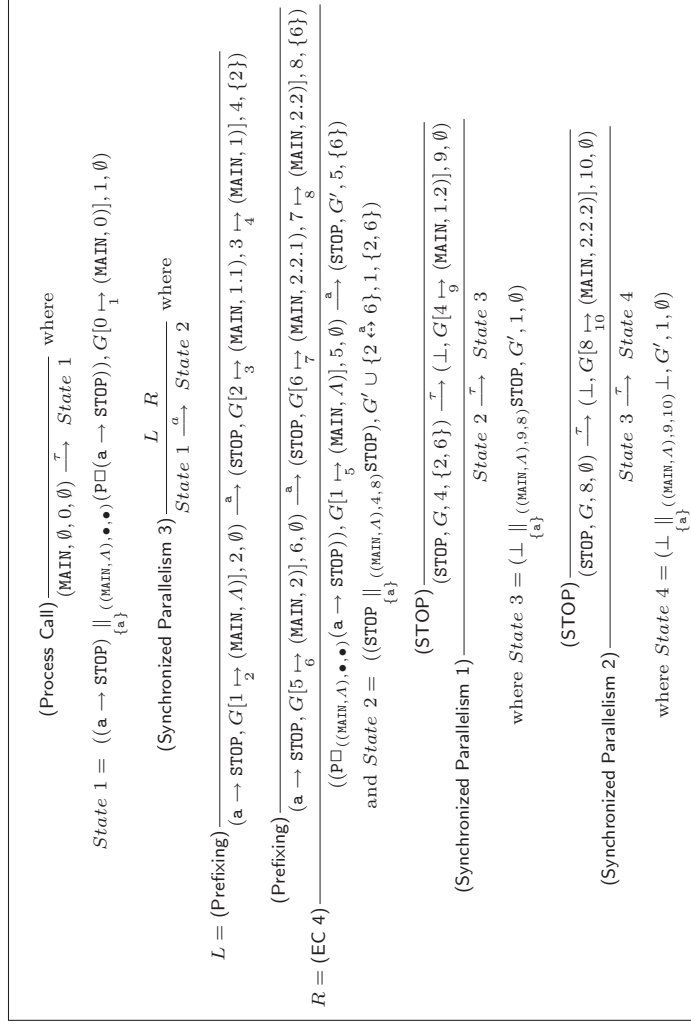


Fig. 8. An example of computation with the tracking semantics in Fig. 7

5 Correctness

In this section we prove the correctness of the tracking semantics (in Fig. 7) by showing that (i) the computations performed by the tracking semantics are equivalent to the computations performed by the standard semantics; and (ii) the graph produced by the tracking semantics is the track of the derivation. We also prove that the trace of a derivation can be automatically extracted from the track of this derivation.

The first theorem shows that the computations performed with the tracking semantics are all and only the computations performed with the standard

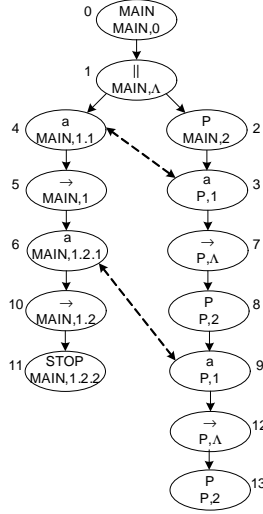


Fig. 9. Track of the program in Example 7

semantics. The only difference between them from an operational point of view is that the tracking semantics needs to perform one step when a **STOP** is evaluated (to add its specification position to the track) and then finishes, while the standard semantics finishes without performing any additional step.

Theorem 1 (Conservativeness). *Let \mathcal{S} be a CSP specification, P a process in \mathcal{S} , and \mathcal{D} and \mathcal{D}' the derivations of P performed with the standard semantics of CSP and with the tracking semantics, respectively. Then, the sequence of rules applied in \mathcal{D} and \mathcal{D}' is exactly the same except that \mathcal{D}' performs one rewriting step more than \mathcal{D} for each (sub)computation that finishes with **STOP**.*

Proof. Firstly, rule (**STOP**) of the tracking semantics is the only rule that is not present in the standard semantics. When a (**STOP**) is reached in a derivation, the standard semantics stops the (sub)computation because no rule is applicable. In the tracking semantics, when a **STOP** is reached in a derivation, the only rule applicable is (**STOP**) which performs τ and puts \perp in the control:

$$\overline{(\text{STOP}_\alpha, G, m, \Delta) \xrightarrow{\tau} (\perp, G[m \mapsto \alpha]_n, n, \emptyset)}$$

Then, the (sub)computation is stopped because no rule is applicable for \perp . Therefore, when the control in the derivation is **STOP**, the tracking semantics performs one additional rewriting step with rule (**STOP**).

The claim follows from the fact that both semantics have exactly the same number of rules except for rule (**STOP**), and these rules have the same control in all the states of the rules (thus the tracking semantics is a conservative extension of the standard semantics). Therefore, all derivations in both semantics have exactly the same number of steps and they are composed of the same sequences

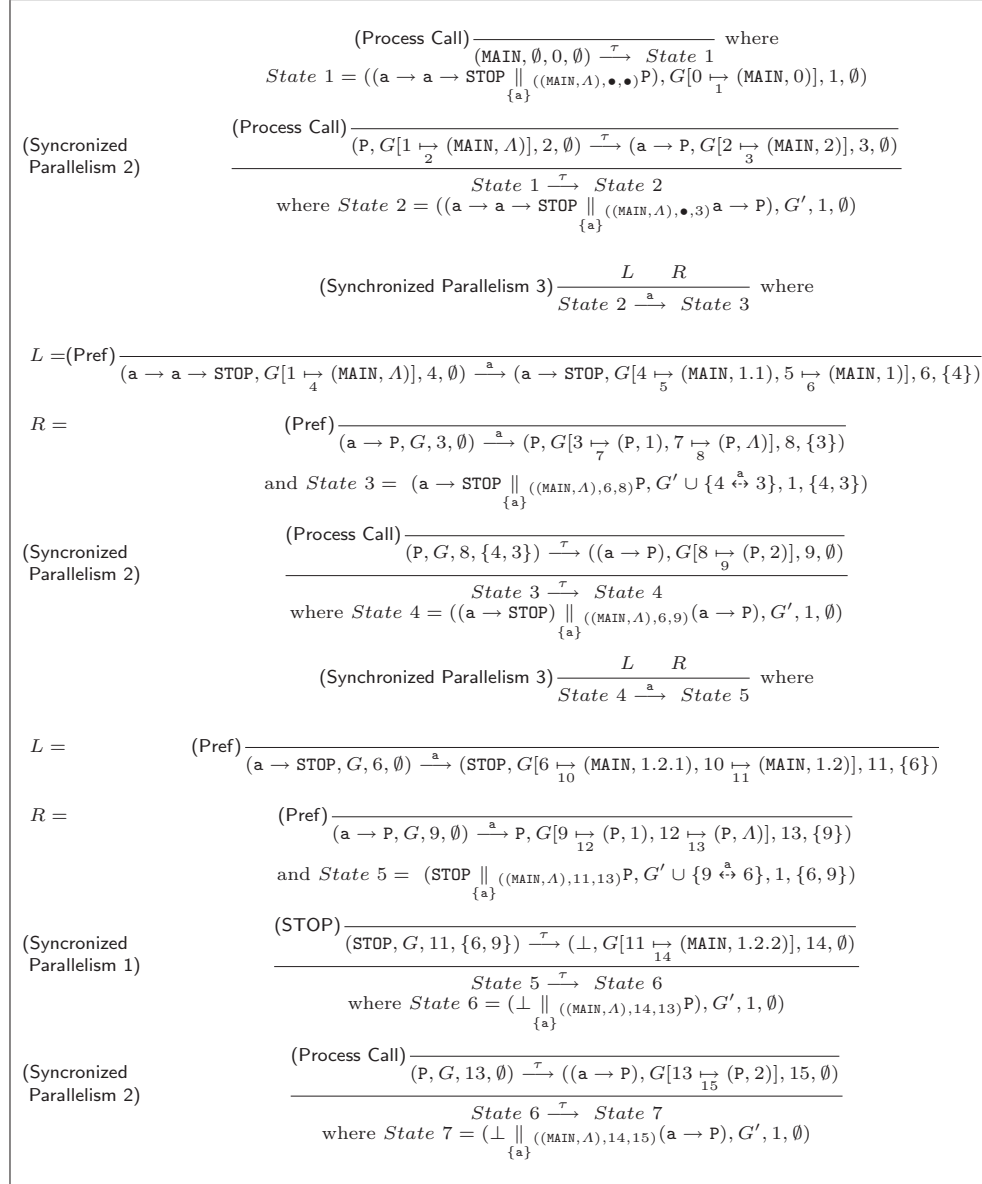


Fig. 10. Computation of Example 7 with the tracking semantics in Fig. 7

of rewriting steps except for (sub)derivations finishing with **STOP** that perform one rewriting step more (applying rule (**STOP**)).

The second theorem states the correctness of the tracking semantics by ensuring that the graph produced is the track of the computation. To prove this theorem, the following lemmas are used.

Lemma 1. *Let \mathcal{S} be a CSP specification, \mathcal{D} a complete derivation of \mathcal{S} performed with the tracking semantics, and \mathcal{G} the graph produced by \mathcal{D} . Then, for each prefixing $(a \rightarrow P)$ in the control of the left state of a rewriting step in \mathcal{D} , we have that $\mathcal{P}os(a)$ and $\mathcal{P}os(\rightarrow)$ are nodes of \mathcal{G} and $\mathcal{P}os(\rightarrow)$ is the successor of $\mathcal{P}os(a)$.*

Proof. If a prefixing $a \rightarrow P$ is in the control of the left state of a rewriting step, following the tracking semantics, only the rule (**Prefixing**) can be applied. By definition of this rule, m is the reference of the current node in \mathcal{G} . The rule (**Prefixing**) adds two new nodes to the graph: n and p . The node m is labelled with the specification position of event a and has successor n . The node n is labelled with the specification position of operator \rightarrow and has parent m and successor p (a fresh reference). Therefore, we have that $\mathcal{P}os(a)$ and $\mathcal{P}os(\rightarrow)$ are nodes of \mathcal{G} and $\mathcal{P}os(\rightarrow)$ is the successor of $\mathcal{P}os(a)$.

Auxiliary Lemma 1. *Given a derivation \mathcal{D} , for each rewriting step $R \xrightarrow{\Theta} R'$ in \mathcal{D} with $R' \neq \Omega$ and $R' \neq \perp$, $last'(R) = last'(R')$.*

Proof. We prove this lemma by induction on the length of Θ . In the base case, Θ is empty, and thus only the rules (**Process Call**), (**Prefixing**), (**Internal Choice 1 and 2**) and (**Synchronized Parallelism 4**) can be applied. In all cases, the lemma holds trivially by the definition of $last'$. We assume as the induction hypothesis that the lemma holds for a non-empty Θ_i with $i > 0$ rewriting steps; and prove that the lemma also holds for a Θ_{i+1} with $i + 1$ rewriting steps. We can assume that $\Theta_{i+1} = R \xrightarrow{\Theta'} R'$, thus, we have to prove that the lemma holds for any possible R and R' . The possible cases are the following:

- (**External Choice 1 and 2**) This case is trivial because the specification positions of R and R' are the same. Hence, $last'(R) = last'(R')$.
- (**External Choice 3 and 4**) Both cases are similar. Thus, we only discuss (**External Choice 3**). In the case of (**External Choice 3**), $last'(P_1 \square P_2) = last'(P_1)$. This rule puts P' in the control, and we know by the induction hypothesis that $last'(P_1) = last'(P')$ and, thus, the lemma holds.
- (**Sequential Composition 1**) This case is analogous to (**External Choice 1 and 2**).
- (**Sequential Composition 2**) $last'(P; Q) = last'(Q)$. Therefore the lemma holds trivially by the definition of $last'$.
- (**Synchronized Parallelism 1, 2 and 3**) It is the same case as (**External Choice 1 and 2**).

Lemma 2. *Let \mathcal{S} be a CSP specification, \mathcal{D} a complete derivation of \mathcal{S} performed with the tracking semantics, and \mathcal{G} the graph produced by \mathcal{D} . Then, for each sequential composition $(P; Q)$ in the control of the left state of a rewriting step in \mathcal{D} , we have that $last'(P)$ and $\mathcal{P}os(;;)$ are nodes of \mathcal{G} and $\mathcal{P}os(;;)$ is the successor of all the elements of the set $last'(P)$ whenever P has successfully finished.*

Proof. If a sequential composition $(P; Q)$ is in the left state of the control of a rewriting step, following the tracking semantics, only the rules (Sequential Composition 1) and (Sequential Composition 2) can be applied. (Sequential Composition 1) is only used to evolve process P until it is finished. The application of this rule is only possible with any event except \checkmark , remaining the sequential composition operator in the control. (Sequential Composition 2) can only be used when \checkmark happens and thus Ω is left in the control.

Therefore, when P has successfully finished evolving to Ω and with m' as the new reference, (Sequential Composition 2) is applied. This rule adds to the graph a new node n labelled with the specification position of $;;$ that has successor p (a fresh reference). Therefore, we have that $\mathcal{P}os(;;)$ is a node of \mathcal{G} and $\mathcal{P}os(;;)$ is the successor of the node m' . Then, we have to prove that $last'(P)$ is a set of nodes of the graph added before P successfully finished with reference m' and its successor is n .

We prove this claim by induction on the length of the derivation $P \xrightarrow{\Theta_0} \dots \xrightarrow{\Theta_n} \Omega$, $n \geq 0$. The base case happens when the last rewriting step of the derivation is done leaving Ω in the control. Only these rules can be used:

- (SKIP) In this case, $m \xrightarrow[n]{\text{SKIP}}$ is added to \mathcal{G} and n is the new reference. Because $last'(\text{SKIP}) = \{\text{SKIP}\}$, therefore, the claim follows.
- (External Choice 3) Here, $last'(P_1 \square P_2) = last'(P_1)$. This rule puts P' in the control which is Ω by the conditions of the lemma. Therefore, there must be a SKIP, which is $last'(P)$, at the top of Θ because we know that the derivation successfully finishes and thus Θ is finite.
- (External Choice 4) It is analogous to the previous case, but here $last'(P_1 \square P_2) = last'(P_2)$.
- (Synchronized Parallelism 4) $last'(P_1 \parallel P_2) = last'(P_1) \cup last'(P_2)$. The parents of the last nodes of P_1 and P_2 are connected to the new reference r . Therefore the claim follows.

The induction hypothesis states that for all rewriting step $R \xrightarrow{\Theta} R', R' \neq \Omega$ in the derivation $Q \xrightarrow{\Theta_0} \dots \xrightarrow{\Theta_n} \Omega, n \geq 0$ where $P \xrightarrow{\Theta'} Q \in \mathcal{D}$, $last'(P)$ is put in the control of a further rewriting step of the derivation together with its reference.

Then, we prove that this also holds for the previous rewriting step $R_0 \xrightarrow{\Theta''} R$. Only rules that do not perform \checkmark could be applied (because \checkmark puts Ω in the control of the right state and now, we are not considering the final rewriting step).

- (STOP) This rule could not be applied because it puts \perp in the control. There is no rule for \perp thus, if applied, P could not successfully finished.

- (Process Call), (Prefixing) and (Internal Choice 1 and 2) In these rules R is put in the control of the final state together with its reference. We know by Auxiliary Lemma 1 that $last'(R_0) = last'(R)$ thus, the claim follows by the induction hypothesis.
- (External Choice 1 and 2) Both rules keep the process in the control and the same references, thus the claim follows by the induction hypothesis.
- (External Choice 3 and 4) In this case, $last'(P_1 \square P_2) = last'(P_1)$. This rule puts P' in the control, and we know by Auxiliary Lemma 1 that $last'(P_1) = last'(P')$, thus the lemma holds by the induction hypothesis.
- (Sequential Composition 1 and 2) We know that P successfully finished, thus (Sequential Composition 1) is applied a number of times before (Sequential Composition 2), that puts Q in the control. We know that $last'(P; Q) = last'(Q)$ thus, the claim holds by the induction hypothesis.
- (Synchronized Parallelism 4) $last'(P_1 || P_2) = last'(P_1) \cup last'(P_2)$. The parents of the last nodes of P_1 and P_2 are connected to the new reference r . Therefore the claim follows.

In the following lemma, we need to extend the notion of rewriting step by including the graph reference to each expression. Therefore, we will use *extended* rewriting step denoted with $(R, m) \overset{\Theta}{\rightsquigarrow} (R', m')$.

Auxiliary Lemma 2. *Let \mathcal{S} be a CSP specification, \mathcal{D} a complete derivation of \mathcal{S} performed with the tracking semantics, and \mathcal{G} the graph produced by \mathcal{D} . Then, for each extended rewriting step in \mathcal{D} of the form $(R, m) \overset{\Theta}{\rightsquigarrow} (R', m')$ which is not associated with (Synchronized Parallelism 4) we have that a node for $first(R)$ is added to \mathcal{G} with reference m .*

Proof. We prove the lemma for each rule:

- (SKIP), (STOP), (Prefixing), (Process Call), and (Internal Choice 1 and 2) A node for $first(R)$ (in these rules α) is added to \mathcal{G} with reference m .
- (External Choice 1, 2, 3 and 4) and (Synchronized Parallelism 1, 2 and 3) In these rules, the node associated with $first(R)$ (it is α here) could be included or not, depending on whether it has been included by a previous rewriting step. If it is already included, it is due to the specification position of the previous expression in the control is the same as R , and its associated rewriting step or a previous one has added it. If other case, function `FirstEval` is called and it includes the node for R with reference m , since the corresponding n_1 and n_2 are equal to \bullet .
- (Sequential Composition 1 and 2) In both rules, the node for $first(R)$ (in this case $first(P)$) is included by a rewriting step in Θ . All possible rewriting steps must apply one of these previous rules, and thus, the claim recursively follows.

Lemma 3. *Let \mathcal{S} be a CSP specification, \mathcal{D} a complete derivation of \mathcal{S} performed with the tracking semantics, and \mathcal{G} the graph produced by \mathcal{D} . Then, for each rewriting step in \mathcal{D} of the form $R_i \overset{\Theta_i}{\rightsquigarrow} R_{i+1}$ we have that:*

1. E_c contains an edge $\mathcal{P}os(R_i) \mapsto \mathcal{P}os(\text{first}(R'))$ where $R' \stackrel{\Theta'}{\rightsquigarrow} R'' \in \Theta_i$ and $R_i \Rightarrow \text{first}(R')$, and
2. if $R_i \Rightarrow \text{first}(R_{i+1})$ then E_c contains an edge $\mathcal{P}os(R_i) \mapsto \mathcal{P}os(\text{first}(R_{i+1}))$.

Proof. We prove each claim separately:

1. Firstly, we know that Θ cannot be empty. Therefore, rules (SKIP), (STOP), (Process Call), (Prefixing), (Internal Choice 1 and 2) and (Synchronized Parallelism 4) could not be applied. Moreover, (Sequential Composition) could never be applied because if R_i is of the form $P; Q$, then the unique possible case is that $\mathcal{P}os(P; Q) \Rightarrow \mathcal{P}os(Q)$ (by Definition 2). And Q can only be in the control of the right state; hence, Q cannot appear in Θ . Then the only applicable rules are (External Choice) or (Synchronized Parallelism).

Let us consider extended rewriting steps $(R', m') \stackrel{\Theta'}{\rightsquigarrow} (R'', m'') \in \Theta_i$. First, we have to prove that a node with the specification position of R_i is included in the graph and the reference of its successor node is put in each m' of Θ_i . In rules (External Choice) and (Synchronized Parallelism) it is done using function `FirstEval`. The references associated with the selected branches of the operator must be \bullet , i.e., the branches have not been developed until now in the derivation. Otherwise, by Definition 2, there is no possible control flow between R_i and R' . In this case, if the corresponding reference is \bullet , then `FirstEval` adds to \mathcal{G} the specification position of R_i and the reference of the successor node is put in all possible m' .

2. In this case, R_i cannot be neither a SKIP nor a STOP, because $\mathcal{P}os(R_i) \neq \mathcal{P}os(R_{i+1})$ (Ω or \perp , respectively) by Definition 2. Process R_i cannot be a parallelism because $\mathcal{P}os(R_i) \neq \mathcal{P}os(R_{i+1})$ (itself or Ω).

If R_i is an external choice we have two possibilities. If we apply (External Choice 1 or 2) then R_i and R_{i+1} have the same specification position and thus, by Definition 2, no control flow is possible. If we apply (External Choice 3 or 4) the control cannot pass from R_i to R_{i+1} , because R_{i+1} is different to $\text{first}(R_i.1)$ or $\text{first}(R_i.2)$. This is due to the fact that the nodes associated with these positions have necessarily been added to \mathcal{G} by the rewriting step Θ_i or by a previous rewriting step on derivation \mathcal{D} . Therefore, process R_i must be a process call, a prefixing, an internal choice, or a sequential composition. If it is a sequential composition, rule (Sequential Composition 1) cannot be applied because in this case R_i and R_{i+1} have the same specification position. Therefore, only (Sequential Composition 2) can be applied.

We now prove that the application of any of remaining rules (Process Call), (Prefixing), (Internal Choice 1 and 2), and (Sequential Composition 2) satisfies the property.

Let $(R_i, n_i) \stackrel{\Theta_i}{\rightsquigarrow} (R_{i+1}, n_{i+1})$ be an extended rewriting step. In all the rules, a node labelled α is added to \mathcal{G} (except in (Prefixing) where is β) and the position of its successor is placed as n_{i+1} . Furthermore, we know by Auxiliary Lemma 2 that a node for $\mathcal{P}os(\text{first}(R_{i+1}))$ is included in the next rewriting step in the derivation $(R_{i+1}, n_{i+1}) \stackrel{\Theta_{i+1}}{\rightsquigarrow} (R_{i+2}, n_{i+2})$ having associated position n_{i+1} .

Note here again that Auxiliary Lemma 2 excludes rule (Synchronized Parallelism 4) but in this case both branches must be already in \mathcal{G} by a previous application of (Synchronized Parallelism 1, 2 or 3).

Lemma 4. *Let \mathcal{S} be a CSP specification, \mathcal{D} a derivation of \mathcal{S} performed with the tracking semantics, and \mathcal{G} the graph produced by \mathcal{D} . Then, there exists a synchronization edge $(a \leftrightarrow a')$ in \mathcal{G} for each synchronization in \mathcal{D} where a and a' are the nodes of the synchronized events.*

Proof. We prove this lemma by induction on the length of the derivation $\mathcal{D} = R_0 \xrightarrow{\Theta_0} R_1 \xrightarrow{\Theta_1} \dots \xrightarrow{\Theta_n} R_{n+1}$. We can assume that the derivation starts with the initial configuration $(\text{MAIN}, \emptyset, 0, \emptyset)$, thus in the base case, the only rule applicable is (Process Call) and hence no synchronization is possible. We assume as the induction hypothesis that there exists a synchronization edge $(a \leftrightarrow a') \in \mathcal{G}$ for each synchronization in $R_0 \xrightarrow{\Theta_0} \dots \xrightarrow{\Theta_{i-1}} R_i$ with $0 < i \leq n$ and prove that the lemma also holds for the next rewriting step $R_i \xrightarrow{\Theta_i} R_{i+1}$.

Firstly, only (Synchronized Parallelism 3) allows the synchronization of events. Therefore, only if R_i is a synchronizing parallelism, or if a (Synchronized Parallelism 3) is applied in Θ_i , $(a \leftrightarrow a') \in \mathcal{G}$. Then, let us consider the case where Θ_i is the application of rule (Synchronized Parallelism 3). This proof is also valid in the case where (Synchronized Parallelism 3) is applied in Θ_i . We have the following rewriting step:

$$\frac{\text{RewritingStep}_1}{(P_1 \parallel_{(\alpha, n_1, n_2)} P_2, G, m, \Delta) \xrightarrow{a} (P_1' \parallel_{(\alpha, n_1', n_2')} P_2', G'', m, \Delta_1 \cup \Delta_2)} \quad \frac{\text{RewritingStep}_2}{a \in X}$$

where $G'' = G_1'' \cup G_2'' \cup \{s_1 \overset{a}{\leftrightarrow} s_2 \mid s_1 \in \Delta_1 \wedge s_2 \in \Delta_2\}$

$$\wedge \text{RewritingStep}_1 = (P_1, G_1', n_1', \Delta) \xrightarrow{a} (P_1', G_1'', n_1'', \Delta_1)$$

$$\wedge (G_1', n_1') = \text{FirstEval}(G, n_1, m, \alpha)$$

$$\wedge \text{RewritingStep}_2 = (P_2, G_2', n_2', \Delta) \xrightarrow{a} (P_2', G_2'', n_2'', \Delta_2)$$

$$\wedge (G_2', n_2') = \text{FirstEval}(G, n_2, m, \alpha)$$

Because (Prefixing) is the only rule that performs an event a without further conditions, we know that P_1 must be a prefixing operator or a process containing a prefixing operator whose prefix is a , i.e., we know that the rule applied in RewritingStep_1 is fired with an event a ; and we know that all the rules of the semantics except (Prefixing) need to fire another rule with an event a as a condition. Therefore, at the top of the condition rules, there must be a (Prefixing). The same happens with P_2 . Hence, two prefixing rules (one for P_1 and one for P_2) have been fired as a condition of this rule.

In addition, the new graph G'' contains the synchronization set $\{s_1 \overset{a}{\leftrightarrow} s_2 \mid s_1 \in \Delta_1 \wedge s_2 \in \Delta_2\}$ where Δ_1 and Δ_2 are the sets of references to the

events that must synchronize in $RewritingStep_1$ and $RewritingStep_2$, respectively.

Hence, we have to prove that all and only the events (a) that must synchronize in $RewritingStep_1$ are in Δ_1 . We prove this by showing that all references to the synchronized events are propagated down by all rules from the (Prefixing) in the top to the (Synchronized Parallelism 3). And the proof is analogous for $RewritingStep_2$.

The possible rules applied in $(P_1, G', n', \Delta) \xrightarrow{a} (P'_1, G''_1, n''_1, \Delta_1)$ are: (Prefixing) In this case, the prefix a is added to Δ_1 . (External Choice 3), (External Choice 4), (Sequential Composition 1), (Synchronized Parallelism 1), (Synchronized Parallelism 2) In these cases, the set Δ is propagated down. (Synchronized Parallelism 3) In this case, the sets Δ_1 and Δ_2 are joined and propagated down.

Therefore, all the synchronized events are in the set Δ_1 and the claim follows.

Theorem 2 (Semantics correctness). *Let \mathcal{S} be a CSP specification, \mathcal{D} a derivation of \mathcal{S} performed with the tracking semantics, and \mathcal{G} the graph produced by \mathcal{D} . Then, \mathcal{G} is the track associated with \mathcal{D} .*

Proof. In order to prove that $\mathcal{G} = (N, E_c, E_s)$ is a track, we need to prove that it satisfies the properties of Definition 5. For each $R \overset{\Theta}{\rightsquigarrow} R' \in \mathcal{D}$ and for all rewriting steps in Θ we have

1. E_c contains a control-flow edge $a \mapsto a'$ iff $a \Rightarrow a'$ with respect to \mathcal{D} . This is ensured by the three clauses of Definition 4:
 - by Lemma 1, if R is a prefixing ($a \rightarrow P$), then E_c contains an edge $\mathcal{P}os(a) \mapsto \mathcal{P}os(\rightarrow)$;
 - by Lemma 2, if R is a sequential composition ($Q; P$), then E_c contains an edge $\forall p \in \mathit{last}'(Q), \mathcal{P}os(p) \mapsto \mathcal{P}os(;$);
 - by Lemma 3, if $R \Rightarrow \mathit{first}(R'')$ where $R'' \overset{\Theta'}{\rightsquigarrow} R''' \in \Theta$, then E_c contains an edge $\mathcal{P}os(R) \mapsto \mathcal{P}os(\mathit{first}(R''))$; and if $R \Rightarrow \mathit{first}(R')$ then E_c contains an edge $\mathcal{P}os(R) \mapsto \mathcal{P}os(\mathit{first}(R'))$; and
2. by Lemma 4, E_s contains a synchronization edge $a \leftrightarrow a'$ for each synchronization occurring in the rewriting step where a and a' are the synchronized events.

Moreover, we know that the only nodes in N are the nodes induced by E_c and E_s because all the nodes inserted in \mathcal{G} are inserted by connecting the new node to the last inserted node (i.e., if the current reference is m and the new fresh reference is n , then the new node is always inserted as $G[m \mapsto_n \alpha]$). Hence, all nodes are related by control or synchronization edges and thus the claim holds.

Our last result states that the trace of a derivation can be extracted from its associated track. To prove it, we define first an order on the event nodes of a track that corresponds to the order in which they were generated by the tracking semantics.

Definition 6. (*Event node order*) Given a track $\mathcal{G} = (N, E_c, E_s)$ and nodes $m, n \in N$ such that $l(m), l(n) \in \Sigma$, m is smaller than n , represented by $m \ll n$ iff $m' < n'$ where $(m, m'), (n, n') \in E_c$.

Intuitively, an event node m is smaller than an event node n if and only if the successor of m has a reference smaller than the reference of the successor of n . The following lemma is also necessary to prove that the order in which events occur in a derivation is directly related with the order of Definition 6. In the following we consider an augmented version of derivation \mathcal{D} which includes the event fired by the application of the rule. So, we can represent derivation \mathcal{D} as

$$P_1 \xrightarrow[e_1]{\Theta_1} \dots \xrightarrow[e_j]{\Theta_j} P_{j+1}.$$

Lemma 5. Given a derivation $\mathcal{D} = P_1 \xrightarrow[e_1]{\Theta_1} \dots \xrightarrow[e_j]{\Theta_j} P_{j+1}$ of the tracking semantics, and the track $\mathcal{G} = (N, E_c, E_s)$ produced by \mathcal{D} , then $\forall e_i \in \Sigma, 1 \leq i \leq j$,

- $\exists n \in N$ such that $l(n) = e_i$, and
- $\exists (n, n') \in E_c$ such that $n' = n + 1$.

Proof. In order to prove this lemma, we prove first that any rewriting step $P_i \xrightarrow[e_i]{\Theta_i} P_{i+1}$ in \mathcal{D} , $1 \leq i \leq j$, with $e_i \in \Sigma$ is a prefixing or it performs a prefixing in Θ_i . This can be easily proved by showing that the rewriting step is either a prefixing (thus $\Theta_i = \emptyset$), or Θ_i has a prefixing as the top rewriting step. We prove this by case analysis. The only possible rules applied are:

- (Prefixing) If this rule is applied, $\Theta_i = \emptyset$ and the claim follows trivially.
- (External Choice 3 and 4), (Synchronized Parallelism 1 and 2) In these rules Θ_i contains a single rewriting step whose event is also e_i ; hence, the claim follows by the induction hypothesis.
- (Sequential Composition 1) This case is completely analogous to the previous one.
- (Synchronized Parallelism 3) In this case, Θ_i is formed by two different rewriting steps, and both of them are similar to the previous case.

Now, both conditions hold trivially from the fact that the prefixing rule adds n to N with label $l(n) = e_i = \alpha$, and it also adds the prefixing operator (\rightarrow) to N as the successor of n .

Therefore, Lemma 5 ensures that the order of Definition 6 corresponds to the order in which the semantics generates the nodes, because each event is added to the graph together with a new fresh reference for the prefixing operator. Since references are generated incrementally, the occurrence of an event e will generate a reference which is less than the reference generated with a posterior event e' . With this order, we can easily define a transformation to extract a trace from a track based on the following proposition:

Proposition 1. Given a track $\mathcal{G} = (N, E_c, E_s)$, the trace induced by \mathcal{G} is the sequence of events $T = e_1, \dots, e_m$ that labels the associated sequence of nodes $T' = n_1, \dots, n_m$ (i.e., $\forall e_i \in T, n_i \in T', 1 \leq i \leq m, l(n_i) = e_i$ and $e_i \in \Sigma$) where:

1. $\forall n_i \in T', 0 < i < m, n_i \ll n_{i+1}$
2. $\forall n \in N$ such that $l(n) \in \Sigma$, if $(\nexists n' \in N \mid (n, n') \in E_s)$, then $n \in T'$
3. $\forall n \in N$ such that $l(n) \in \Sigma$, if $(\forall n' \in N \mid (n, n') \in E_s \wedge n' \ll n)$, then $n \in T'$

Proof. We consider a derivation $\mathcal{D} = P_1 \xrightarrow[e'_1]{\Theta_1} \dots \xrightarrow[e'_k]{\Theta_j} P_{j+1}$. Note that $e'_1, \dots, e'_k \neq e_1, \dots, e_m$ because the former contains events in $\{\tau, \checkmark\}$. Then, we have that the trace is the subsequence of e'_1, \dots, e'_k that only includes events of Σ . We will represent this subsequence with $\mathcal{E} = e'_{j_0}, \dots, e'_{j_k}$ with $0 \leq j_0 \leq j_1 \leq \dots \leq j_k \leq k$. Then, we have to show that $T = \mathcal{E}$. The proposition follows trivially from the fact that the sequence T follows the order of nodes imposed by Definition 6, and this order is the same order of the events that form the sequence \mathcal{E} as stated by Lemma 5.

Theorem 3 (Track correctness). *Let \mathcal{S} be a CSP specification, \mathcal{D} a derivation of \mathcal{S} produced by the sequence of events (i.e., the trace) $T = e_1, \dots, e_m$, and \mathcal{G} the track associated with \mathcal{D} . Then, there exists a function f that extracts the trace T from the track \mathcal{G} , i.e., $f(\mathcal{G}) = T$.*

Proof. Proposition 1 allows to trivially define a function f such that $f(\mathcal{G}) = T$ being \mathcal{G} the track of a derivation \mathcal{D} , and being T the trace of the same derivation. For a track $\mathcal{G} = (N, E_c, E_s)$ we have that

$$f((n : ns), E_c, E_s) = \begin{cases} \{f((ns), E_c, E_s)\} & \text{if } (\exists n' \in N \mid (n, n') \in E_s \wedge n \ll n') \\ \{l(n) : f((ns), E_c, E_s)\} & \text{otherwise} \end{cases}$$

where list $(n : ns)$ corresponds to the set $\{n \in N \mid l(n) \in \Sigma\}$ ordered with respect to order \ll of Definition 6.

6 Conclusions

This work introduces the first semantics of CSP instrumented for tracking. Therefore, it is an interesting result because it can serve as a reference mark to define and prove properties such as completeness of static analyses which are based on tracks [13–15]. The execution of the tracking semantics produces a graph as a side effect which is the track of the computation. This track is produced step by step by the semantics, and thus, it can be also used to produce a track of an infinite computation until it is stopped. The generated track can be useful not only for tracking computations but for debugging and program comprehension. This is due to the fact that our generated track also includes the specification positions associated with the expressions appearing in the track. Therefore, tracks could be used to analyse what parts of the program are executed (and in what order) in a particular computation. Also, this information allows a track viewer tool to highlight the parts of the code that are executed in each step. Notable analyses that use tracks are [3–5, 1, 13–15]. The introduction of this semantics allows us to adapt these analyses to CSP. On the practical side, we have implemented a tool called *SOC* [13] which is able to automatically generate tracks of a CSP specification. These tracks are later used for debugging.

SOC has been integrated into the most extended CSP animator and model-checker ProB [11, 12], that shows the maturity and usefulness of this tool and of tracks. The implementation, source code and several examples are publicly available at: <http://users.dsic.upv.es/~jsilva/soc/>

Acknowledgements

We want to thank the anonymous referees for many valuable comments and useful suggestions.

References

1. Brassel, B., Hanus, M., Huch, F., Vidal, G.: A Semantics for Tracing Declarative Multi-paradigm Programs. In: Moggi, E., Warren, D.S. (eds.) 6th ACM SIGPLAN Int'l Conf. on Principles and Practice of Declarative Programming (PPDP'04), pp. 179–190. ACM, New York, NY, USA (2004)
2. Brückner, I., Wehrheim, H.: Slicing an Integrated Formal Method for Verification. In: Lau, K.K., Banach, R. (eds.) Int'l Conf. on Formal Engineering Methods (ICFEM'05). LNCS, vol. 3785, pp. 360–374. Springer, Heidelberg (2005)
3. Chitil, O.: A Semantics for Tracing. In: Arts, T., Mohnen, M. (eds.) 13th Int'l Workshop on Implementation of Functional Languages (IFL'01), pp. 249–254. Ericsson CSL (2001)
4. Chitil, O., Runciman, C., Wallace, M.: Transforming Haskell for Tracing. In: Peña, R., Arts, T. (eds.) 14th Int'l Workshop on Implementation of Functional Languages (IFL'02), Revised Selected Papers. LNCS, vol. 2670, pp. 165–181. Springer, Heidelberg (2003)
5. Chitil, O., Lou, Y.: Structure and Properties of Traces for Functional Programs. *Electronic Notes in Theoretical Computer Science (ENTCS)*. 176(1), 39–63 (2007)
6. Ferrante, J., Ottenstein, K.J., Warren, J.D.: The Program Dependence Graph and its Use in Optimization. *ACM Transactions on Programming Languages and Systems*. 9(3), 319–349 (1987)
7. Hoare, C.A.R.: *Communicating Sequential Processes*. Prentice-Hall, Upper Saddle River, NJ, USA (1985)
8. Kavi, K.M., Sheldon, F.T., Shirazi, B., Hurson, A.R.: Reliability analysis of CSP specifications using Petri nets and Markov processes. In: 28th Annual Hawaii Int'l Conf. on System Sciences (HICSS'95), vol. 2 (Software Technology), pp. 516–524. IEEE Computer Society, Washington, DC, USA (1995)
9. Krinke, J.: Context-Sensitive Slicing of Concurrent Programs. *ACM SIGSOFT Software Engineering Notes*. 28(5) (2003)
10. Ladkin, P., Simons, B.: *Static Deadlock Analysis for CSP-Type Communications. Responsive Computer Systems (Chapter 5)*, Kluwer Academic Publishers (1995)
11. Leuschel, M., Butler, M.: ProB: an automated analysis toolset for the B method. *Journal of Software Tools for Technology Transfer*. 10(2), 185–203 (2008)
12. Leuschel, M., Fontaine, M.: Probing the depths of CSP-M: A new FDR-compliant validation tool. In: Liu, S., Maibaum, T., Araki, K. (eds.) 10th Int'l Conf. on Formal Engineering Methods (ICFEM'08). LNCS, vol. 5256, pp. 278–297. Springer, Heidelberg (2008)

13. Leuschel, M., Llorens, M., Oliver, J., Silva, J., Tamarit, S.: SOC: a Slicer for CSP Specifications. In: Puebla, G., Vidal, G. (eds.) 2009 ACM SIGPLAN Symposium on Partial Evaluation and Semantics-based Program Manipulation (PEPM'09), pp. 165–168. ACM, New York, NY, USA (2009)
14. Leuschel, M., Llorens, M., Oliver, J., Silva, J., Tamarit, S.: The MEB and CEB Static Analysis for CSP Specifications. In: Hanus, M. (ed.) 18th Int'l Symposium on Logic-Based Program Synthesis and Transformation (LOPSTR'08), Revised Selected Papers. LNCS, vol. 5438, pp. 103–118. Springer, Heildeberg (2009)
15. Llorens, M., Oliver, J., Silva, J., Tamarit, S.: An Algorithm to Generate the Context-sensitive Synchronized Control Flow Graph. In: 25th ACM Symposium on Applied Computing (SAC 2010) (to appear). ACM, New York, NY, USA (2010)
16. Roscoe, A.W.: The Theory and Practice of Concurrency. Prentice Hall, Upper Saddle River, NJ, USA (2005)
17. Tip, F.: A survey of program slicing techniques. *Journal of Programming Languages*. 3, 121–189 (1995)
18. Weiser, M.D.: Program Slicing. *IEEE Transactions on Software Engineering*. 10(4), 352–357 (1984)