



HAL
open science

Dynamic Choreographies: Theory And Implementation

Mila Dalla Preda, Maurizio Gabbrielli, Saverio Giallorenzo, Ivan Lanese,
Jacopo Mauro

► **To cite this version:**

Mila Dalla Preda, Maurizio Gabbrielli, Saverio Giallorenzo, Ivan Lanese, Jacopo Mauro. Dynamic Choreographies: Theory And Implementation. Logical Methods in Computer Science, 2017, 13, pp.1 - 57. 10.23638/LMCS-13(2:1)2017. hal-01631394

HAL Id: hal-01631394

<https://inria.hal.science/hal-01631394>

Submitted on 9 Nov 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

DYNAMIC CHOREOGRAPHIES: THEORY AND IMPLEMENTATION

MILA DALLA PREDA^a, MAURIZIO GABBRIELLI^b, SAVERIO GIALLORENZO^c, IVAN LANESE^d,
AND JACOPO MAURO^e

^a Department of Computer Science, University of Verona
e-mail address: mila.dallapreda@univr.it

^{b,c,d} Department of Computer Science and Engineering, University of Bologna/INRIA
e-mail address: {gabbri, sgiallor, lanese}@cs.unibo.it

^e Department of Informatics, University of Oslo
e-mail address: jacpom@ifi.uio.no

ABSTRACT. Programming distributed applications free from communication deadlocks and race conditions is complex. Preserving these properties when applications are updated at runtime is even harder. We present a choreographic approach for programming updatable, distributed applications. We define a choreography language, called Dynamic Interaction-Oriented Choreography (DIOC), that allows the programmer to specify, from a global viewpoint, which parts of the application can be updated. At runtime, these parts may be replaced by new DIOC fragments from outside the application. DIOC programs are compiled, generating code for each participant in a process-level language called Dynamic Process-Oriented Choreographies (DPOC). We prove that DPOC distributed applications generated from DIOC specifications are deadlock free and race free and that these properties hold also after any runtime update. We instantiate the theoretical model above into a programming framework called Adaptable Interaction-Oriented Choreographies in Jolie (AIOCJ) that comprises an integrated development environment, a compiler from an extension of DIOCs to distributed Jolie programs, and a runtime environment to support their execution.

1. INTRODUCTION

Programming distributed applications is an error-prone activity. Participants send and receive messages but, if the application is badly programmed, they may get stuck waiting for messages that never arrive (communication deadlock), or they may receive messages in an unexpected order, depending on the speed of the other participants and of the network (races).

Key words and phrases: Choreographies, Adaptable Systems, Deadlock Freedom.

^c Supported by the EU EIT Digital project *SMAll*.

^d Supported by the GNCS group of INdAM via project *Logica, Automi e Giochi per Sistemi Auto-adattivi*.

^e Supported by the EU project FP7-644298 *HyVar: Scalable Hybrid Variability for Distributed, Evolving Software Systems*.

Recently, language-based approaches have been proposed to tackle the complexity of programming concurrent and distributed applications. These approaches provide high-level primitives that avoid by construction some of the risks of concurrent programming. Some renowned examples are the ownership/borrowing mechanisms of Rust [44] and the `separate` keyword of SCOOP [41]. In these settings most of the work needed to ensure a correct behaviour is done by the language compiler and runtime support. The use of these languages requires a conceptual shift from traditional ones, but reduces times and costs of development, testing, and maintenance by avoiding some of the most common programming errors.

We propose an approach based on *choreographic programming* [12, 13, 30, 47, 34] following a similar philosophy. A choreography is a global description of a distributed application expressed as the composition of the expected interactions between its components. Based on such an abstraction, we present high-level primitives that avoid some of the main errors linked to programming communication-centred concurrent and distributed applications that can be updated at run time.

For an informal introduction of our approach, let us consider a simple application where a **Buyer** asks the price of some product to a **Seller**.

```

1  priceReq : Buyer(prod) → Seller(order);
2  order_price@Seller = getPrice(order);
3  offer : Seller(order_price) → Buyer(prod_price)

```

A unique choreographic program describes the behaviour of multiple participants, here the **Buyer** and the **Seller**. The first line specifies that the **Buyer** sends, along channel `priceReq`, the name of the desired product `prod` to the **Seller**, which stores it in its local variable `order`. At Line 2, the **Seller** computes the price of the product with function `getPrice`. Finally, at Line 3 the **Seller** sends the computed price to the **Buyer** on channel `offer`. The **Buyer** stores the value in its local variable `prod_price`.

Choreographic languages focus on describing *message-based interactions* between distributed participants. As shown by the example, the distinctive feature of choreographies is that communications are atomic entities, i.e., not split into send and receive actions. This makes impossible to write programs with common errors of concurrent and distributed applications like deadlocks or race conditions. However, a choreographic description is not directly executable. Indeed, choreographies describe atomic interactions from a global point of view whilst executable programs are written in lower level languages (like Java or C) that support only local behaviours and communication/synchronisation primitives such as message send and message receive. Hence, to run a choreographic description, we have to compile it into a set of lower level programs. The correct compilation of choreographic descriptions is one of the main challenges of choreographic languages, yet the ease of development and the strong guarantees of deadlock and race freedom make it a worth endeavour.

In this work, we take this challenge one step further: we consider *updatable* distributed applications whose code can change dynamically, i.e., while the application is running. In particular, the applications we consider can integrate external code at runtime. Such a feature, tricky in a sequential setting and even more in a distributed one, has countless uses: to deal with emergency requirements, to cope with rules and requirements depending on contextual properties or to improve and specialise the application to user preferences. We propose a general mechanism to structure application updates. Inside applications, we delimit blocks of code, called *scopes*, that may be dynamically replaced by new blocks of

code, called *updates*. Updates are loosely related to scopes: it is not necessary to know the details of the behaviour of updates when writing scopes, and updates may be written while applications are running.

To illustrate our approach, let us consider the previous example and let us suppose that we would like to introduce new commercial offers at runtime, e.g., to provide a discount on the computed prices. Since we want to be able to change how prices are computed, we enclose Lines 2–3 of the example within a scope, as shown below.

```

1  priceReq : Buyer(prod) → Seller(order);
2  scope @Seller{
3    order_price@Seller = getPrice(order);
4    offer : Seller(order_price) → Buyer(prod_price)
5  }
```

In essence, a scope is a delimiter that defines which part of the application can be updated. Each scope identifies a *coordinator* of the update, i.e., the participant entitled to ensure that either none of the participants updates, or they all apply the same update. In the example above, the coordinator of the update is the **Seller** (Line 2).

Since now the code includes a scope, we can introduce runtime updates. Let us assume that the **Seller** issued a fidelity card to some frequent customers and (s)he wants to update the application to let **Buyers** insert their fidelity card to get a discount. The update in Figure 1 answers this business need. At runtime, if the **Seller** (which is the coordinator of the update) applies the update in Figure 1, the code of the update replaces the scope. When this new code executes, the **Buyer** sends his/her *card_id* to the **Seller**. If the *card_id* is valid, the **Seller** issues a 10% discount on the price of the selected good, otherwise it reports the standard price to the **Buyer**.

We remark that expressing the behaviour described above using lower level languages that lack a dedicated support for distributed runtime updates (e.g., Java, C, etc.) is extremely error prone. For instance, in the example above, let us consider the case in which the **Buyer** is updated first and it starts the execution whilst the **Seller** has not been updated yet. The **Buyer** waits for a message from the **Seller** on channel *cardReq*, whilst the **Seller** sends a message on channel *offer*. Thus, the application is deadlocked. Furthermore, in our setting the available updates may change at any time, posing an additional challenge. To avoid

```

1  cardReq : Seller(null) → Buyer(-);
2  card_id@Buyer = getInput();
3  cardRes : Buyer(card_id) → Seller(buyer_id);
4  if isValid(buyer_id)@Seller {
5    order_price@Seller = getPrice(order) * 0.9
6  } else {
7    order_price@Seller = getPrice(order)
8  };
9  offer : Seller(order_price) → Buyer(prod_price)
```

Figure 1: Fidelity Card Update.

errors, participants must select the same update and, since updates appear and disappear at runtime, they must also be able to retrieve the same selected update.

Since at choreographic level updates are applied atomically to all the involved participants, these problems cannot arise if both the original application and the updates are described as choreographies. However, to execute our choreographic specifications we need to provide a correct compilation from choreographies to lower level, executable programs. Such task is very challenging and, in particular, we need to make sure that our compilation generates correct behaviours that avoid inconsistencies on updates.

Contributions. In this paper, we present a correctness-by-construction approach to solve the problem of dynamic updates of distributed applications. We provide a general solution that comprises:

- the definition of a *Dynamic Interaction-Oriented Choreography* language, called DIOC, to program distributed applications and supporting runtime code update (§ 2);
- the definition of a *Dynamic Process-Oriented Choreography* language, called DPOC. DPOCs are based on standard send and receive primitives but they are tailored to program updatable systems. We introduce DPOCs to describe implementations corresponding to DIOCs. (§ 3);
- the definition of a behaviour-preserving *projection function* to compile DIOCs into DPOCs (§ 4);
- the *proof of the correctness* of the projection function (§ 7). Correctness is guaranteed even in a scenario where the set of available updates dynamically changes, at any moment and without notice;
- one instantiation of our theory into a *development framework* for adaptive distributed applications called AIOCJ (§ 8). In AIOCJ updates are embodied into adaptation rules, whose application is not purely non-deterministic (as in DIOCs), but depends on the state of the system and of its environment. AIOCJ comprises an Integrated Development Environment, a compiler from choreographies to executable programs, and a runtime environment to support their execution and update.

This paper integrates and extends material from [18], which outlines the theoretical aspects, and [19], which describes the AIOCJ framework. Main extensions include the full semantics of DIOC and DPOC, detailed proofs and technical definitions, and a thorough description of the example. Furthermore, both the presentation and the technical development have been deeply revised and refined.

2. DYNAMIC INTERACTION-ORIENTED CHOREOGRAPHIES

In this section we introduce the syntax of DIOCs, we illustrate the constructs of the DIOC language with a comprehensive example, and we finally present the semantics of DIOCs.

2.1. DIOC Syntax. DIOCs rely on a set of *Roles*, ranged over by $\mathbf{R}, \mathbf{S}, \dots$, to identify the participants in the choreography. We call them roles to highlight that they have a specific duty in the choreography. Each role has its local state.

Roles exchange messages over public channels, also called *operations*, ranged over by o . We denote with *Expr* the set of expressions, ranged over by e . We deliberately do not give a formal definition of expressions and of their typing, since our results do not depend on it. We only require that expressions include at least values, belonging to a set *Val* ranged over

by v , and variables, belonging to a set Var ranged over by x, y, \dots . We also assume a set of boolean expressions ranged over by b .

The syntax of *DIOC processes*, ranged over by $\mathcal{I}, \mathcal{I}', \dots$, is defined as follows:

$\mathcal{I} ::=$	$o : \mathbf{R}(e) \rightarrow \mathbf{S}(x)$	$(interaction)$		$\mathbf{0}$	(end)
	$\mathcal{I}; \mathcal{I}'$	$(sequence)$		$\mathbf{if} \ b@ \mathbf{R} \ \{\mathcal{I}\} \ \mathbf{else} \ \{\mathcal{I}'\}$	$(conditional)$
	$\mathcal{I} \mathcal{I}'$	$(parallel)$		$\mathbf{while} \ b@ \mathbf{R} \ \{\mathcal{I}\}$	$(while)$
	$x@ \mathbf{R} = e$	$(assignment)$		$\mathbf{scope} \ @ \mathbf{R} \ \{\mathcal{I}\}$	$(scope)$
	$\mathbf{1}$	$(inaction)$			

Interaction $o : \mathbf{R}(e) \rightarrow \mathbf{S}(x)$ means that role \mathbf{R} sends a message on operation o to role \mathbf{S} (we require $\mathbf{R} \neq \mathbf{S}$). The sent value is obtained by evaluating expression e in the local state of \mathbf{R} and it is then stored in the local variable x of \mathbf{S} . Processes $\mathcal{I}; \mathcal{I}'$ and $\mathcal{I} | \mathcal{I}'$ denote sequential and parallel composition, respectively. Assignment $x@ \mathbf{R} = e$ assigns the evaluation of expression e in the local state of \mathbf{R} to its local variable x . The empty process $\mathbf{1}$ defines a DIOC that can only terminate. $\mathbf{0}$ represents a terminated DIOC. It is needed for the definition of the operational semantics and it is not intended to be used by the programmer. We call *initial* a DIOC process where $\mathbf{0}$ never occurs. The conditional $\mathbf{if} \ b@ \mathbf{R} \ \{\mathcal{I}\} \ \mathbf{else} \ \{\mathcal{I}'\}$ and the iteration $\mathbf{while} \ b@ \mathbf{R} \ \{\mathcal{I}\}$ are guarded by the evaluation of the boolean expression b in the local state of \mathbf{R} . The construct $\mathbf{scope} \ @ \mathbf{R} \ \{\mathcal{I}\}$ delimits a subterm \mathcal{I} of the DIOC process that may be updated in the future. In $\mathbf{scope} \ @ \mathbf{R} \ \{\mathcal{I}\}$, role \mathbf{R} is the coordinator of the update: it ensures that either none of the participants update, or they all apply the same update.

A Running Example. We report in Figure 2 a running example of a DIOC process that extends the one presented in the Introduction: the example features a **Buyer** that orders a product from a **Seller**, and a **Bank** that supports the payment from the **Buyer** to the **Seller**. The DIOC process describes the behaviour of all of them. In this sense, the DIOC implements a protocol they agreed upon to integrate their business. The DIOC protocol also interacts with functionalities that are only available at the premises of some of the participants (for instance, the function `getPrice` is provided by the **Seller** IT system and may be only locally available). For this reason, it is important that the execution of the DIOC is distributed across the participants.

At Lines 1–2 the **Buyer** initialises its local variables `price_ok` and `continue`. These variables control the while loop used by the **Buyer** to ask the price of some product to the **Seller**. The loop execution is controlled by the **Buyer**, but it impacts also the behaviour of the **Seller**. We will see in § 4 that this is done using automatically generated auxiliary communications. A similar approach is used to manage conditionals. At Line 4, the **Buyer** takes the name of the *product* from the user with function `getInput`, which models interaction with the user, and proceeds to send it to the **Seller** on operation `priceReq` (Line 5). The **Seller** computes the price of the product calling the function `getPrice` (Line 7) and, via operation `offer`, it sends the price to the **Buyer** (Line 8), that stores it in a local variable `prod_price`. These last two operations are performed within a scope and therefore they can be updated at runtime to implement some new business policies (e.g., discounts). At Lines 10–12 the **Buyer** checks whether the user is willing to buy the product, and, if (s)he is not interested, whether (s)he wants to ask prices for other products. If the **Buyer** accepts the offer of the **Seller**, the **Seller** sends to the **Bank** the payment details (Line 16),

```

1  price_ok@Buyer = false;
2  continue@Buyer = true;
3  while( !price_ok and continue )@Buyer{
4    prod@Buyer = getInput();
5    priceReq : Buyer(prod) → Seller(order);
6    scope @Seller{
7      order_price@Seller = getPrice(order);
8      offer : Seller(order_price) → Buyer(prod_price)
9    };
10   price_ok@Buyer = getInput();
11   if( !price_ok )@Buyer{
12     continue@Buyer = getInput()
13   }
14 };
15 if( price_ok )@Buyer{
16   payReq : Seller( payDesc(order_price) ) → Bank(desc);
17   scope @Bank{
18     pay : Buyer( payAuth(prod_price) ) → Bank(auth)
19   };
20   payment_ok@Bank = makePayment(desc, auth);
21   if( payment_ok )@Bank{
22     confirm : Bank(null) → Seller(-)
23     |
24     confirm : Bank(null) → Buyer(-)
25   } else {
26     abort : Bank(null) → Buyer(-)
27   }
28 }

```

Figure 2: DIOC process for Purchase Scenario.

computed using function `payDesc`. Next, the **Buyer** authorises the payment via operation `pay`, computing the payment authorisation form using function `payAuth`. Since the payment may be critical for security reasons, the related communication is enclosed in a scope (Lines 17–19), thus allowing the introduction of more refined procedures later on. After the scope successfully terminates, at Line 20 the **Bank** locally executes the actual payment by calling function `makePayment`. The function `makePayment` abstracts an invocation to the **Bank** IT system. We show in § 8 that, using functions, one can integrate existing service-oriented software into a choreographic program.

Finally, the **Bank** acknowledges the payment to the **Seller** and the **Buyer** in parallel (Lines 22–24). If the payment is not successful, the failure is notified explicitly only to the **Buyer**. Note that at Lines 1–2 the annotation `@Buyer` means that the variables `price_ok` and `continue` belong to the **Buyer**. Similarly, at Line 3, the annotation `@Buyer` means that the guard of the while loop is evaluated by the **Buyer**. The term `@Seller` at Line 6 is part of the scope construct and indicates the **Seller** as coordinator of the update.

2.2. Annotated DIOCs and their Semantics. In the remainder of the paper, we define our results on an annotated version of the DIOC syntax. Annotations are numerical indexes $\mathbf{i} \in \mathbb{N}$ assigned to DIOC constructs. We only require indexes to be distinct. Any annotation that satisfies this requirement provides the same result. Indeed, programmers do not need to annotate DIOCs: the annotation with indexes is mechanisable and can be performed by the language compiler¹. Indexes are used both in the proof of our results and in the projection to avoid interferences between different constructs. From now on we consider only well-annotated DIOCs, defined as follows.

Definition 2.1 (Well-annotated DIOC). *Annotated DIOC processes* are obtained by indexing every interaction, assignment, conditional, while loop, and scope in a DIOC process with a positive natural number $\mathbf{i} \in \mathbb{N}$, resulting in the following grammar:

$$\begin{array}{l|l} \mathcal{I} ::= & \mathbf{i} : o : \mathbf{R}(e) \rightarrow \mathbf{S}(x) \quad | \quad \mathbf{0} \\ & \mathcal{I}; \mathcal{I}' \quad | \quad \mathbf{i} : \text{if } b @ \mathbf{R} \{ \mathcal{I} \} \text{ else } \{ \mathcal{I}' \} \\ & \mathcal{I} | \mathcal{I}' \quad | \quad \mathbf{i} : \text{while } b @ \mathbf{R} \{ \mathcal{I} \} \\ & \mathbf{i} : x @ \mathbf{R} = e \quad | \quad \mathbf{i} : \text{scope } @ \mathbf{R} \{ \mathcal{I} \} \\ & \mathbf{1} \end{array}$$

A DIOC process is *well annotated* if all its indexes are distinct.

DIOC processes do not execute in isolation: they are equipped with a *global state* Σ and a set of available updates \mathbf{I} , i.e., a set of DIOCs that may replace scopes. Set \mathbf{I} may change at runtime. A global state Σ is a map that defines the value v of each variable x in a given role \mathbf{R} , namely $\Sigma : \text{Roles} \times \text{Var} \rightarrow \text{Val}$. The local state of role \mathbf{R} is denoted as $\Sigma_{\mathbf{R}} : \text{Var} \rightarrow \text{Val}$ and it verifies that $\forall x \in \text{Var} : \Sigma(\mathbf{R}, x) = \Sigma_{\mathbf{R}}(x)$. Expressions are always evaluated by a given role \mathbf{R} : we denote the evaluation of expression e in local state $\Sigma_{\mathbf{R}}$ as $\llbracket e \rrbracket_{\Sigma_{\mathbf{R}}}$. We assume that $\llbracket e \rrbracket_{\Sigma_{\mathbf{R}}}$ is always defined (e.g., an error value is given as a result if evaluation is not possible) and that for each boolean expression b , $\llbracket b \rrbracket_{\Sigma_{\mathbf{R}}}$ is either **true** or **false**.

Remark 2.2. The above assumption on expressions is needed for our results. To satisfy it, when programming, one should prevent runtime errors and notify abnormal events to partners using normal constructs to guarantee error management or smooth termination (see e.g., Lines 21–26 in Figure 2). A more elegant way to deal with errors would be to include in the language well-known constructs, such as try-catch, which are however left as future work. This could be done by adapting the ideas presented in [9].

Definition 2.3 (DIOC systems). A DIOC system is a triple $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle$ denoting a DIOC process \mathcal{I} equipped with a global state Σ and a set of updates \mathbf{I} .

Definition 2.4 (DIOC systems semantics). The semantics of DIOC systems is defined as the smallest labelled transition system (LTS) closed under the rules in Figure 3, where symmetric rules for parallel composition have been omitted.

The rules in Figure 3 describe the behaviour of a DIOC system by induction on the structure of its DIOC process, with a case analysis on its topmost operator. We use μ to range over labels. The possible values for μ are described below.

$$\begin{array}{l|l} \mu ::= & o : \mathbf{R}(v) \rightarrow \mathbf{S}(x) \quad (\text{interaction}) \quad | \quad \tau \quad (\text{silent}) \\ & \mathcal{I} \quad (\text{update}) \quad | \quad \mathbf{no-up} \quad (\text{no update}) \\ & \mathbf{I} \quad (\text{change updates}) \quad | \quad \checkmark \quad (\text{termination}) \end{array}$$

¹In fact, the AIOCJ compiler implements such a feature.

$$\begin{array}{c}
\frac{\llbracket e \rrbracket_{\Sigma_{\mathbf{R}}} = v}{\langle \Sigma, \mathbf{I}, \mathbf{i}: o : \mathbf{R}(e) \rightarrow \mathbf{S}(x) \rangle \xrightarrow{o:\mathbf{R}(v) \rightarrow \mathbf{S}(x)} \langle \Sigma, \mathbf{I}, \mathbf{i}: x @ \mathbf{S} = v \rangle} \text{[DIOC | INTERACTION]} \\
\\
\frac{\llbracket e \rrbracket_{\Sigma_{\mathbf{R}}} = v}{\langle \Sigma, \mathbf{I}, \mathbf{i}: x @ \mathbf{R} = e \rangle \xrightarrow{\tau} \langle \Sigma[v/x, \mathbf{R}], \mathbf{I}, \mathbf{1} \rangle} \text{[DIOC | ASSIGN]} \\
\\
\frac{\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle \xrightarrow{\mu} \langle \Sigma, \mathbf{I}', \mathcal{I}' \rangle \quad \mu \neq \surd}{\langle \Sigma, \mathbf{I}, \mathcal{I}; \mathcal{J} \rangle \xrightarrow{\mu} \langle \Sigma, \mathbf{I}', \mathcal{I}'; \mathcal{J} \rangle} \text{[DIOC | SEQUENCE]} \\
\\
\frac{\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle \xrightarrow{\surd} \langle \Sigma, \mathbf{I}, \mathcal{I}' \rangle \quad \langle \Sigma, \mathbf{I}, \mathcal{J} \rangle \xrightarrow{\mu} \langle \Sigma, \mathbf{I}, \mathcal{J}' \rangle}{\langle \Sigma, \mathbf{I}, \mathcal{I}; \mathcal{J} \rangle \xrightarrow{\mu} \langle \Sigma, \mathbf{I}, \mathcal{J}' \rangle} \text{[DIOC | SEQ-END]} \\
\\
\frac{\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle \xrightarrow{\mu} \langle \Sigma, \mathbf{I}', \mathcal{I}' \rangle \quad \mu \neq \surd}{\langle \Sigma, \mathbf{I}, \mathcal{I} \parallel \mathcal{J} \rangle \xrightarrow{\mu} \langle \Sigma, \mathbf{I}', \mathcal{I}' \parallel \mathcal{J} \rangle} \text{[DIOC | PARALLEL]} \\
\\
\frac{\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle \xrightarrow{\surd} \langle \Sigma, \mathbf{I}, \mathcal{I}' \rangle \quad \langle \Sigma, \mathbf{I}, \mathcal{J} \rangle \xrightarrow{\surd} \langle \Sigma, \mathbf{I}, \mathcal{J}' \rangle}{\langle \Sigma, \mathbf{I}, \mathcal{I} \parallel \mathcal{J} \rangle \xrightarrow{\surd} \langle \Sigma, \mathbf{I}, \mathcal{I}' \parallel \mathcal{J}' \rangle} \text{[DIOC | PAR-END]} \\
\\
\frac{\llbracket b \rrbracket_{\Sigma_{\mathbf{R}}} = \mathbf{true}}{\langle \Sigma, \mathbf{I}, \mathbf{i}: \mathbf{if } b @ \mathbf{R} \{ \mathcal{I} \} \mathbf{ else } \{ \mathcal{I}' \} \rangle \xrightarrow{\tau} \langle \Sigma, \mathbf{I}, \mathcal{I} \rangle} \text{[DIOC | IF-THEN]} \\
\\
\frac{\llbracket b \rrbracket_{\Sigma_{\mathbf{R}}} = \mathbf{false}}{\langle \Sigma, \mathbf{I}, \mathbf{i}: \mathbf{if } b @ \mathbf{R} \{ \mathcal{I} \} \mathbf{ else } \{ \mathcal{I}' \} \rangle \xrightarrow{\tau} \langle \Sigma, \mathbf{I}, \mathcal{I}' \rangle} \text{[DIOC | IF-ELSE]} \\
\\
\frac{\llbracket b \rrbracket_{\Sigma_{\mathbf{R}}} = \mathbf{true}}{\langle \Sigma, \mathbf{I}, \mathbf{i}: \mathbf{while } b @ \mathbf{R} \{ \mathcal{I} \} \rangle \xrightarrow{\tau} \langle \Sigma, \mathbf{I}, \mathcal{I}; \mathbf{i}: \mathbf{while } b @ \mathbf{R} \{ \mathcal{I} \} \rangle} \text{[DIOC | WHILE-UNFOLD]} \\
\\
\frac{\llbracket b \rrbracket_{\Sigma_{\mathbf{R}}} = \mathbf{false}}{\langle \Sigma, \mathbf{I}, \mathbf{i}: \mathbf{while } b @ \mathbf{R} \{ \mathcal{I} \} \rangle \xrightarrow{\tau} \langle \Sigma, \mathbf{I}, \mathbf{1} \rangle} \text{[DIOC | WHILE-EXIT]} \\
\\
\frac{\text{roles}(\mathcal{I}') \subseteq \text{roles}(\mathcal{I}) \quad \mathcal{I}' \in \mathbf{I} \quad \text{connected}(\mathcal{I}') \quad \text{freshIndexes}(\mathcal{I}')}{\langle \Sigma, \mathbf{I}, \mathbf{i}: \mathbf{scope } @ \mathbf{R} \{ \mathcal{I} \} \rangle \xrightarrow{\mathcal{I}'} \langle \Sigma, \mathbf{I}, \mathcal{I}' \rangle} \text{[DIOC | UP]} \\
\\
\frac{}{\langle \Sigma, \mathbf{I}, \mathbf{i}: \mathbf{scope } @ \mathbf{R} \{ \mathcal{I} \} \rangle \xrightarrow{\text{no-up}} \langle \Sigma, \mathbf{I}, \mathcal{I} \rangle} \text{[DIOC | NoUP]} \\
\\
\frac{}{\langle \Sigma, \mathbf{I}, \mathbf{1} \rangle \xrightarrow{\surd} \langle \Sigma, \mathbf{I}, \mathbf{0} \rangle} \text{[DIOC | END]} \quad \frac{}{\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle \xrightarrow{\mathbf{I}'} \langle \Sigma, \mathbf{I}', \mathcal{I} \rangle} \text{[DIOC | CHANGE-UPDATES]}
\end{array}$$

Figure 3: Annotated DIOC system semantics.

Rule $\text{[DIOC | INTERACTION]}$ executes a communication from \mathbf{R} to \mathbf{S} on operation o , where \mathbf{R} sends to \mathbf{S} the value v of an expression e . The communication reduces to an assignment that inherits the index \mathbf{i} of the interaction. The assignment stores value v in variable x of role \mathbf{S} . Rule [DIOC | ASSIGN] evaluates the expression e in the local state $\Sigma_{\mathbf{R}}$ and stores the resulting value v in the local variable x in role \mathbf{R} ($[v/x, \mathbf{R}]$ represents the substitution).

$$\begin{aligned}
\text{roles}(\mathbf{i} : \mathbf{R}(e) \rightarrow \mathbf{S}(x)) &= \{\mathbf{R}, \mathbf{S}\} \\
\text{roles}(\mathbf{1}) &= \text{roles}(\mathbf{0}) = \emptyset \\
\text{roles}(\mathbf{i} : x @ \mathbf{R} = e) &= \{\mathbf{R}\} \\
\text{roles}(\mathcal{I}; \mathcal{I}') &= \text{roles}(\mathcal{I} | \mathcal{I}') = \text{roles}(\mathcal{I}) \cup \text{roles}(\mathcal{I}') \\
\text{roles}(\mathbf{i} : \text{if } b @ \mathbf{R} \{ \mathcal{I} \} \text{ else } \{ \mathcal{I}' \}) &= \{\mathbf{R}\} \cup \text{roles}(\mathcal{I}) \cup \text{roles}(\mathcal{I}') \\
\text{roles}(\mathbf{i} : \text{while } b @ \mathbf{R} \{ \mathcal{I} \}) &= \text{roles}(\mathbf{i} : \text{scope } @ \mathbf{R} \{ \mathcal{I} \}) = \{\mathbf{R}\} \cup \text{roles}(\mathcal{I})
\end{aligned}$$

Figure 4: Auxiliary function roles.

Rule $\llbracket \text{DIOC} \mid_{\text{SEQUENCE}} \rrbracket$ executes a step in the first process of a sequential composition, while Rule $\llbracket \text{DIOC} \mid_{\text{SEQ-END}} \rrbracket$ acknowledges the termination of the first process, starting the second one. Rule $\llbracket \text{DIOC} \mid_{\text{PARALLEL}} \rrbracket$ allows a process in a parallel composition to compute, while Rule $\llbracket \text{DIOC} \mid_{\text{PAR-END}} \rrbracket$ synchronises the termination of two parallel processes. Rules $\llbracket \text{DIOC} \mid_{\text{IF-THEN}} \rrbracket$ and $\llbracket \text{DIOC} \mid_{\text{IF-ELSE}} \rrbracket$ evaluate the boolean guard of a conditional, selecting the “then” and the “else” branch, respectively. Rules $\llbracket \text{DIOC} \mid_{\text{WHILE-UNFOLD}} \rrbracket$ and $\llbracket \text{DIOC} \mid_{\text{WHILE-EXIT}} \rrbracket$ correspond respectively to the unfolding of a while loop when its condition is satisfied and to its termination otherwise. Rule $\llbracket \text{DIOC} \mid_{\text{UP}} \rrbracket$ and Rule $\llbracket \text{DIOC} \mid_{\text{NOUP}} \rrbracket$ deal with updates: the former applies an update, while the latter allows the body of the scope to be executed without updating it. More precisely, Rule $\llbracket \text{DIOC} \mid_{\text{UP}} \rrbracket$ models the application of the update \mathcal{I}' to the scope $\text{scope } @ \mathbf{R} \{ \mathcal{I} \}$ which, as a result, is replaced by the DIOC process \mathcal{I}' . In the conditions of the rule, we use the function roles and the predicates connected and freshIndexes . Function $\text{roles}(\mathcal{I})$, defined in Figure 4, computes the roles of a DIOC process \mathcal{I} . The condition of the rule requires that the roles of the update are a subset of the roles of the body of the scope. Predicate $\text{connected}(\mathcal{I}')$ holds if \mathcal{I}' is connected. Connectedness is a well-formedness property of DIOCs and is detailed in § 6. Roughly, it ensures that roles involved in a sequential composition have enough information to enforce the correct sequence of actions. Predicate $\text{freshIndexes}(\mathcal{I}')$ holds if all indexes in \mathcal{I}' are fresh with respect to all indexes already present in the target DIOC². This is needed to avoid interferences between communications inside the update and communications in the original DIOC. This problem is discussed in more details in § 4, Example 4.3. Rule $\llbracket \text{DIOC} \mid_{\text{NOUP}} \rrbracket$, used when no update is applied, removes the scope boundaries and starts the execution of the body of the scope. Rule $\llbracket \text{DIOC} \mid_{\text{END}} \rrbracket$ terminates the execution of an empty process. Rule $\llbracket \text{DIOC} \mid_{\text{CHANGE-UPDATES}} \rrbracket$ allows the set \mathbf{I} of available updates to change. This rule is always enabled and models the fact that the set of available updates is not controlled by the system, but by the external world: the set of updates can change at any time, the system cannot forbid or delay these changes, and the system is not notified when they happen. Label \mathbf{I} is needed to make the changes to the set of available updates observable (cf. Definition 7.1).

Remark 2.5. Whether to update a scope or not, and which update to apply if many are available, is completely non-deterministic. We have adopted this view to maximise generality. However, for practical applications it is also possible to reduce the non-determinism using suitable policies to decide when and whether a given update applies. One of such policies is defined in AIOCJ (see § 8).

²We do not give a formal definition of $\text{freshIndexes}(\mathcal{I}')$ to keep the presentation simple. However, freshness of indexes can be formally ensured using restriction as in π -calculus [46].

We can finally provide the definition of *DIOC traces* and *weak DIOC traces*, which we use to express our results of behavioural correspondence between DIOCs and DPOCs. Intuitively, in DIOC traces all the performed actions are observed, whilst in weak DIOC traces silent actions τ are not visible.

Definition 2.6 (DIOC traces). A (*strong*) *trace* of a DIOC system $\langle \Sigma_1, \mathbf{I}_1, \mathcal{I}_1 \rangle$ is a sequence (finite or infinite) of labels μ_1, μ_2, \dots such that there is a sequence of DIOC system transitions $\langle \Sigma_1, \mathbf{I}_1, \mathcal{I}_1 \rangle \xrightarrow{\mu_1} \langle \Sigma_2, \mathbf{I}_2, \mathcal{I}_2 \rangle \xrightarrow{\mu_2} \dots$.

A *weak trace* of a DIOC system $\langle \Sigma_1, \mathbf{I}_1, \mathcal{I}_1 \rangle$ is a sequence of labels μ_1, μ_2, \dots obtained by removing all silent labels τ from a trace of $\langle \Sigma_1, \mathbf{I}_1, \mathcal{I}_1 \rangle$.

3. DYNAMIC PROCESS-ORIENTED CHOREOGRAPHIES

In this section we define the syntax and semantics of DPOCs, the target language of our projection from DIOCs. We remind that DIOCs are not directly executable since their basic primitives describe distributed interactions. On the contrary, mainstream languages like Java and C, used for implementation, describe distributed computations using local behaviours and communication/synchronisation primitives, such as message send and message receive. In order to describe implementations corresponding to DIOCs we introduce the DPOC language, a core language based on this kind of primitives, but tailored to program updatable systems. Indeed, differently from DIOC constructs, DPOC constructs are locally implementable in any mainstream language. In AIOcJ (see § 8) we implement the DPOC constructs in the Jolie [36] language.

3.1. DPOC syntax. DPOCs include *processes*, ranged over by P, P', \dots , describing the behaviour of participants. $(P, \Gamma)_{\mathbf{R}}$ denotes a *DPOC role* named \mathbf{R} , executing process P in a local state Γ . *Networks*, ranged over by $\mathcal{N}, \mathcal{N}', \dots$, are parallel compositions of DPOC roles with different names. DPOC systems, ranged over by \mathcal{S} , are DPOC networks equipped with a set of updates \mathbf{I} , namely pairs $\langle \mathbf{I}, \mathcal{N} \rangle$.

DPOCs, like DIOCs, feature operations o . Here we call them *programmer-specified operations* to mark the difference with respect to *auxiliary operations*, ranged over by o^* . We use $o^?$ to range over both programmer-specified and auxiliary operations. Differently from communications on programmer-specified operations, communications on auxiliary operations have no direct correspondent at the DIOC level. Indeed, we introduce auxiliary operations in DPOCs to implement the synchronisation mechanisms needed to realise the global constructs of DIOCs (conditionals, while loops, and scopes) at DPOC level.

Like DIOC constructs, also DPOC constructs are annotated using indexes. However, in DPOCs we use two kinds of indexes: *normal indexes* $\mathbf{i} \in \mathbb{N}$ and *auxiliary indexes* of the forms $\mathbf{i}_T, \mathbf{i}_F, \mathbf{i}_?$, and \mathbf{i}_C where $\mathbf{i} \in \mathbb{N}$. Auxiliary indexes are introduced by the projection, described in § 4, and are described in detail there. We let ι range over DPOC indexes.

In DPOCs, normal indexes are also used to prefix the operations in send and receive primitives³. Thus, a send and a receive can interact only if they are on the same operation and they are prefixed by the same normal index. This is needed to avoid interferences

³In principle, one may use just indexes and drop operations. However, explicit operations are the standard for communication-based systems, in particular in the area of Web Services, as they come handy to specify and to debug such systems.

between concurrent communications, in particular when one of them comes from an update. We describe in greater detail this issue in § 4, Example 4.3.

The syntax of DPOCs is the following.

$$\begin{array}{l|l|l}
P ::= & \iota: \mathbf{i.o}^? : x \text{ from } \mathbf{R} & \text{(receive)} & | & \mathbf{1} & \text{(inaction)} \\
& \iota: \mathbf{i.o}^? : e \text{ to } \mathbf{R} & \text{(send)} & | & \mathbf{0} & \text{(end)} \\
& \mathbf{i}: \mathbf{i.o}^* : X \text{ to } \mathbf{R} & \text{(send-update)} & | & \mathbf{i}: \text{if } b \{P\} \text{ else } \{P'\} & \text{(conditional)} \\
& P; P' & \text{(sequence)} & | & \mathbf{i}: \text{while } b \{P\} & \text{(while)} \\
& P | P' & \text{(parallel)} & | & \mathbf{i}: \text{scope } @\mathbf{R} \{P\} \text{ roles } \{S\} & \text{(scope-coord)} \\
& \iota: x = e & \text{(assignment)} & | & \mathbf{i}: \text{scope } @\mathbf{R} \{P\} & \text{(scope)} \\
\end{array}$$

$$X ::= \quad \mathbf{no} \quad | \quad P \quad \quad \mathcal{N} ::= \quad (P, \Gamma)_{\mathbf{R}} \quad | \quad \mathcal{N} \parallel \mathcal{N}'$$

DPOC processes include receive action $\iota: \mathbf{i.o}^? : x \text{ from } \mathbf{R}$ on a specific operation $\mathbf{i.o}^?$ (either programmer-specified or auxiliary) of a message from role \mathbf{R} to be stored in variable x , send action $\iota: \mathbf{i.o}^? : e \text{ to } \mathbf{R}$ of the value of an expression e to be sent to role \mathbf{R} , and higher-order send action $\mathbf{i}: \mathbf{i.o}^* : X \text{ to } \mathbf{R}$ of the higher-order argument X to be sent to role \mathbf{R} . Here X may be either a DPOC process P , which is the new code for a scope in \mathbf{R} , or a token \mathbf{no} , notifying that no update is needed. $P; P'$ and $P | P'$ denote the sequential and parallel composition of P and P' , respectively. Processes also feature assignment $\iota: x = e$ of the value of expression e to variable x , the process $\mathbf{1}$, that can only successfully terminate, and the terminated process $\mathbf{0}$. DPOC processes also include conditionals $\mathbf{i}: \text{if } b \{P\} \text{ else } \{P'\}$ and loops $\mathbf{i}: \text{while } b \{P\}$. Finally, there are two constructs for scopes. Construct $\mathbf{i}: \text{scope } @\mathbf{R} \{P\} \text{ roles } \{S\}$ defines a scope with body P and set of participants S , and may occur only inside role \mathbf{R} , which acts as coordinator of the update. The shorter version $\mathbf{i}: \text{scope } @\mathbf{R} \{P\}$ is used instead inside the code of some role \mathbf{R}_1 , which is not the coordinator \mathbf{R} of the update. In fact, only the coordinator \mathbf{R} needs to know the set S of involved roles to be able to send to them their updates.

3.2. DPOC semantics. DPOC semantics is defined in two steps: we define the semantics of DPOC roles first, and then we define how roles interact giving rise to the semantics of DPOC systems.

Definition 3.1 (DPOC roles semantics). The semantics of DPOC roles is defined as the smallest LTS closed under the rules in Figure 5, where we report the rules dealing with computation, and Figure 6, in which we define the rules related to updates. Symmetric rules for parallel composition have been omitted.

DPOC role semantics. We use δ to range over labels. The possible values for δ are as follows:

$$\begin{array}{l|l|l}
\delta ::= & \overline{\mathbf{i.o}^?} \langle v \rangle @\mathbf{S} : \mathbf{R} & \text{(send)} & | & \mathbf{i.o}^? (x \leftarrow v) @\mathbf{S} : \mathbf{R} & \text{(receive)} \\
& \overline{\mathbf{i.o}^*} \langle X \rangle @\mathbf{S} : \mathbf{R} & \text{(send-update)} & | & \mathbf{i.o}^* (\leftarrow X) @\mathbf{S} : \mathbf{R} & \text{(receive-update)} \\
& \mathcal{I} & \text{(update)} & | & \mathbf{no-up} & \text{(no-update)} \\
& \tau & \text{(silent)} & | & \surd & \text{(termination)}
\end{array}$$

$$\begin{array}{c}
\frac{}{(\mathbf{1}, \Gamma)_{\mathbf{R}} \xrightarrow{\checkmark} (\mathbf{0}, \Gamma)_{\mathbf{R}}} \quad [\text{DPOC} \mid_{\text{ONE}}] \quad \frac{\llbracket e \rrbracket_{\Gamma} = v}{(\iota: x = e, \Gamma)_{\mathbf{R}} \xrightarrow{\tau} (\mathbf{1}, \Gamma[v/x])_{\mathbf{R}}} \quad [\text{DPOC} \mid_{\text{ASSIGN}}] \\
\frac{\llbracket e \rrbracket_{\Gamma} = v}{(\iota: \mathbf{i.o}^? : e \text{ to } \mathbf{S}, \Gamma)_{\mathbf{R}} \xrightarrow{\overline{\mathbf{i.o}^?(v)}@S:\mathbf{R}}} (\mathbf{1}, \Gamma)_{\mathbf{R}}} \quad [\text{DPOC} \mid_{\text{SEND}}] \\
\frac{}{(\iota: \mathbf{i.o}^? : x \text{ from } \mathbf{S}, \Gamma)_{\mathbf{R}} \xrightarrow{\mathbf{i.o}^?(x \leftarrow v)@S:\mathbf{R}}} (\iota: x = v, \Gamma)_{\mathbf{R}}} \quad [\text{DPOC} \mid_{\text{RECV}}] \\
\frac{}{(\mathbf{i}: \mathbf{i.o}^? : X \text{ to } \mathbf{S}, \Gamma)_{\mathbf{R}} \xrightarrow{\overline{\mathbf{i.o}^?(X)}@S:\mathbf{R}}} (\mathbf{1}, \Gamma)_{\mathbf{R}}} \quad [\text{DPOC} \mid_{\text{SEND-UP}}] \\
\frac{(P, \Gamma)_{\mathbf{R}} \xrightarrow{\delta} (P', \Gamma')_{\mathbf{R}} \quad \delta \neq \checkmark}{(P; Q, \Gamma)_{\mathbf{R}} \xrightarrow{\delta} (P'; Q, \Gamma')_{\mathbf{R}}} \quad [\text{DPOC} \mid_{\text{SEQUENCE}}] \\
\frac{(P, \Gamma)_{\mathbf{R}} \xrightarrow{\checkmark} (P', \Gamma)_{\mathbf{R}} \quad (Q, \Gamma)_{\mathbf{R}} \xrightarrow{\delta} (Q', \Gamma')_{\mathbf{R}}}{(P; Q, \Gamma)_{\mathbf{R}} \xrightarrow{\delta} (Q', \Gamma')_{\mathbf{R}}} \quad [\text{DPOC} \mid_{\text{SEQ-END}}] \\
\frac{(P, \Gamma)_{\mathbf{R}} \xrightarrow{\delta} (P', \Gamma')_{\mathbf{R}} \quad \delta \neq \checkmark}{(P \mid Q, \Gamma)_{\mathbf{R}} \xrightarrow{\delta} (P' \mid Q, \Gamma')_{\mathbf{R}}} \quad [\text{DPOC} \mid_{\text{PARALLEL}}] \\
\frac{(P, \Gamma)_{\mathbf{R}} \xrightarrow{\checkmark} (P', \Gamma)_{\mathbf{R}} \quad (Q, \Gamma)_{\mathbf{R}} \xrightarrow{\checkmark} (Q', \Gamma)_{\mathbf{R}}}{(P \mid Q, \Gamma)_{\mathbf{R}} \xrightarrow{\checkmark} (P' \mid Q', \Gamma)_{\mathbf{R}}} \quad [\text{DPOC} \mid_{\text{PAR-END}}] \\
\frac{\llbracket b \rrbracket_{\Gamma} = \mathbf{true}}{(\mathbf{i}: \text{if } b \{P\} \text{ else } \{P'\}, \Gamma)_{\mathbf{R}} \xrightarrow{\tau} (P, \Gamma)_{\mathbf{R}}} \quad [\text{DPOC} \mid_{\text{IF-THEN}}] \\
\frac{\llbracket b \rrbracket_{\Gamma} = \mathbf{false}}{(\mathbf{i}: \text{if } b \{P\} \text{ else } \{P'\}, \Gamma)_{\mathbf{R}} \xrightarrow{\tau} (P', \Gamma)_{\mathbf{R}}} \quad [\text{DPOC} \mid_{\text{IF-ELSE}}] \\
\frac{\llbracket b \rrbracket_{\Gamma} = \mathbf{true}}{(\mathbf{i}: \text{while } b \{P\}, \Gamma)_{\mathbf{R}} \xrightarrow{\tau} (P; \mathbf{i}: \text{while } e \{P\}, \Gamma)_{\mathbf{R}}} \quad [\text{DPOC} \mid_{\text{WHILE-UNFOLD}}] \\
\frac{\llbracket b \rrbracket_{\Gamma} = \mathbf{false}}{(\mathbf{i}: \text{while } b \{P\}, \Gamma)_{\mathbf{R}} \xrightarrow{\tau} (\mathbf{1}, \Gamma)_{\mathbf{R}}} \quad [\text{DPOC} \mid_{\text{WHILE-EXIT}}]
\end{array}$$

Figure 5: DPOC role semantics. Computation rules. (Update rules in Figure 6)

The semantics is in the early style. Rule $[\text{DPOC} \mid_{\text{RECV}}]$ receives a value v from role \mathbf{S} and assigns it to local variable x of \mathbf{R} . Similarly to Rule $[\text{DIO} \mid_{\text{INTERACT}}]$ (see Figure 3), the reception reduces to an assignment that inherits the index \mathbf{i} from the receive primitive.

Rules $[\text{DPOC} \mid_{\text{SEND}}]$ and $[\text{DPOC} \mid_{\text{SEND-UP}}]$ execute send and higher-order send actions, respectively. Send actions evaluate expression e in the local state Γ . Rule $[\text{DPOC} \mid_{\text{ONE}}]$ terminates an empty process. Rule $[\text{DPOC} \mid_{\text{ASSIGN}}]$ executes an assignment ($[v/x]$ represents the substitution of value v for variable x). Rule $[\text{DPOC} \mid_{\text{SEQUENCE}}]$ executes a step in the first process of a sequential composition, while Rule $[\text{DPOC} \mid_{\text{SEQ-END}}]$ acknowledges the termination of the first process, starting the second one. Rule $[\text{DPOC} \mid_{\text{PARALLEL}}]$ allows a process in a

$$\begin{array}{c}
\frac{\text{roles}(\mathcal{I}) \subseteq S \quad \text{freshIndexes}(\mathcal{I}) \quad \text{connected}(\mathcal{I})}{(\mathbf{i}: \text{scope } @\mathbf{R} \{P\} \text{ roles } \{S\}, \Gamma)_{\mathbf{R}} \xrightarrow{\mathcal{I}} \left(\prod_{\mathbf{R}_j \in S \setminus \{\mathbf{R}\}} \mathbf{i}: \mathbf{i}.sb_{\mathbf{i}}^* : \pi(\mathcal{I}, \mathbf{R}_j) \text{ to } \mathbf{R}_j; \pi(\mathcal{I}, \mathbf{R}); \prod_{\mathbf{R}_j \in S \setminus \{\mathbf{R}\}} \mathbf{i}: \mathbf{i}.se_{\mathbf{i}}^* : - \text{from } \mathbf{R}_j, \Gamma \right)_{\mathbf{R}}} \text{[DPOC } |_{\text{LEAD-UP}}\text{]} \\
\frac{}{(\mathbf{i}: \text{scope } @\mathbf{R} \{P\} \text{ roles } \{S\}, \Gamma)_{\mathbf{R}} \xrightarrow{\text{no-up}} \left(\prod_{\mathbf{R}_j \in S \setminus \{\mathbf{R}\}} \mathbf{i}: \mathbf{i}.sb_{\mathbf{i}}^* : \text{no to } \mathbf{R}_j; P; \prod_{\mathbf{R}_j \in S \setminus \{\mathbf{R}\}} \mathbf{i}: \mathbf{i}.se_{\mathbf{i}}^* : - \text{from } \mathbf{R}_j, \Gamma \right)_{\mathbf{R}}} \text{[DPOC } |_{\text{LEAD-NOUP}}\text{]} \\
\frac{}{(\mathbf{i}: \text{scope } @\mathbf{S} \{P\}, \Gamma)_{\mathbf{R}} \xrightarrow{\mathbf{i}.sb_{\mathbf{i}}^* (\leftarrow P') @\mathbf{S}:\mathbf{R}} (P'; \mathbf{i}: \mathbf{i}.se_{\mathbf{i}}^* : \text{ok to } \mathbf{S}, \Gamma)_{\mathbf{R}}} \text{[DPOC } |_{\text{UP}}\text{]} \\
\frac{}{(\mathbf{i}: \text{scope } @\mathbf{S} \{P\}, \Gamma)_{\mathbf{R}} \xrightarrow{\mathbf{i}.sb_{\mathbf{i}}^* (\leftarrow \text{no}) @\mathbf{S}:\mathbf{R}} (P; \mathbf{i}: \mathbf{i}.se_{\mathbf{i}}^* : \text{ok to } \mathbf{S}, \Gamma)_{\mathbf{R}}} \text{[DPOC } |_{\text{NOUP}}\text{]}
\end{array}$$

Figure 6: DPOC role semantics. Update rules. (Computation rules in Figure 5)

parallel composition to compute, while Rule $[\text{DPOC } |_{\text{PAR-END}}]$ synchronises the termination of two parallel processes. Rules $[\text{DPOC } |_{\text{IF-THEN}}]$ and $[\text{DPOC } |_{\text{IF-ELSE}}]$ select the “then” and the “else” branch in a conditional, respectively. Rules $[\text{DPOC } |_{\text{WHILE-UNFOLD}}]$ and $[\text{DPOC } |_{\text{WHILE-EXIT}}]$ model respectively the unfolding and the termination of a while loop.

The rules reported in Figure 6 deal with code updates. Rules $[\text{DPOC } |_{\text{LEAD-UP}}]$ and $[\text{DPOC } |_{\text{LEAD-NOUP}}]$ specify the behaviour of the coordinator \mathbf{R} of the update, respectively when an update is performed and when no update is performed. In particular, \mathbf{R} non-deterministically selects whether to update or not and, in the first case, which update to apply. The coordinator communicates the selection to the other roles in the scope using operations $sb_{\mathbf{i}}^*$. The content of the message is either the new code that the other roles need to execute, if the update is performed, or a token **no**, if no update is applied. Communications on operations $sb_{\mathbf{i}}^*$ also ensure that no role starts executing the scope before the coordinator has selected whether to update or not. The communication is received by the other roles using Rule $[\text{DPOC } |_{\text{UP}}]$ if an update is selected and Rule $[\text{DPOC } |_{\text{NOUP}}]$ if no update is selected. Similarly, communications on $se_{\mathbf{i}}^*$ ensure that the updated code has been completely executed before any role can proceed to execute the code that follows the scope in its original *DPOC* process. Communications on $se_{\mathbf{i}}^*$ carry no relevant data: they are used for synchronisation purposes only.

As already discussed, Rule $[\text{DPOC } |_{\text{LEAD-UP}}]$ models the fact that the coordinator \mathbf{R} of the update non-deterministically selects an update \mathcal{I} . The premises of Rule $[\text{DPOC } |_{\text{LEAD-UP}}]$ are similar to those of Rule $[\text{DPOC } |_{\text{UP}}]$ (see Figure 3). Function `roles` is used to check that the roles in \mathcal{I} are included in the roles of the scope. Freshness of indexes is checked by predicate `freshIndexes`, and well formedness of \mathcal{I} by predicate `connected` (formally defined later on, in Definition 6.2 in § 6). In particular, the coordinator \mathbf{R} generates the processes to be executed by the roles in S using the process-projection function π (detailed in § 4). More precisely, $\pi(\mathcal{I}, \mathbf{R}_i)$ generates the code for role \mathbf{R}_i . The processes $\pi(\mathcal{I}, \mathbf{R}_i)$ are sent via auxiliary higher-order communications on $sb_{\mathbf{i}}^*$ to the roles that have to execute them. These

communications also notify the other roles that they can start executing the new code. Here, and in the remainder of the paper, we define $\prod_{R_i \in S} P_i$ as the parallel composition of DPOC processes P_i for $R_i \in S$. We assume that \prod binds more tightly than sequential composition. After the communication of the updated code to the other participants, \mathbf{R} starts its own updated code $\pi(\mathcal{I}, \mathbf{R})$. Finally, auxiliary communications se_i^* are used to synchronise the end of the execution of the update (here $_$ denotes a fresh variable to store the synchronisation message ok). Rule $[\text{DPOC}]_{\text{LEAD-NOUP}}$ defines the behaviour of the coordinator \mathbf{R} when no update is applied. In this case, \mathbf{R} sends a token \mathbf{no} to the other involved roles, notifying them that no update is applied and that they can start executing their original code. End of scope synchronisation is the same as that of Rule $[\text{DPOC}]_{\text{LEAD-UP}}$. Rules $[\text{DPOC}]_{\text{UP}}$ and $[\text{DPOC}]_{\text{NOUP}}$ define the behaviour of the other roles involved in the scope. The scope waits for a message from the coordinator. If the content of the message is \mathbf{no} , the body of the scope is executed. Otherwise, the content of the message is a process P' which is executed instead of the body of the scope.

We highlight the importance of the coordinator \mathbf{R} . Since the set of updates may change at any moment, we need to be careful to avoid that the participants in the scope get code projected from different updates. Given that only role \mathbf{R} obtains and delivers the new code, one is guaranteed that all the participants receive their projection of the same update.

DPOC system semantics.

Definition 3.2 (DPOC systems semantics). The semantics of DPOC systems is defined as the smallest LTS closed under the rules in Figure 7. Symmetric rules for parallel composition have been omitted.

We use η to range over DPOC systems labels. The possible values of η are as follows:

$$\begin{aligned} \eta ::= & \quad o^? : \mathbf{R}(v) \rightarrow \mathbf{S}(x) && \text{(interaction)} \\ & \quad | \quad o^* : \mathbf{R}(X) \rightarrow \mathbf{S}() && \text{(interaction-update)} \\ & \quad | \quad \delta && \text{(role label)} \end{aligned}$$

Rules $[\text{DPOC}]_{\text{LIFT}}$ and $[\text{DPOC}]_{\text{LIFT-UP}}$ lift role transitions to the system level. Rule $[\text{DPOC}]_{\text{LIFT-UP}}$ also checks that the update \mathcal{I} belongs to the set of currently available updates \mathbf{I} . Rule $[\text{DPOC}]_{\text{SYNCH}}$ synchronises a send with the corresponding receive, producing an interaction. Rule $[\text{DPOC}]_{\text{SYNCH-UP}}$ is similar, but it deals with higher-order interactions. Note that Rules $[\text{DPOC}]_{\text{SYNCH}}$ and $[\text{DPOC}]_{\text{SYNCH-UP}}$ remove the prefixes from DPOC operations in transition labels. The labels of these transitions store the information on the occurred communication: label $o^? : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x)$ denotes an interaction on operation $o^?$ from role \mathbf{R}_1 to role \mathbf{R}_2 where the value v is sent by \mathbf{R}_1 and then stored by \mathbf{R}_2 in variable x . Label $o^* : \mathbf{R}_1(X) \rightarrow \mathbf{R}_2()$ denotes a similar interaction, but concerning a higher-order value X , which can be either the code used in the update or a token \mathbf{no} if no update is performed. No receiver variable is specified, since the received value becomes part of the code of the receiving process. Rule $[\text{DPOC}]_{\text{EXT-PAR}}$ allows a network inside a parallel composition to compute. Rule $[\text{DPOC}]_{\text{EXT-PAR-END}}$ synchronises the termination of parallel networks. Finally, Rule $[\text{DPOC}]_{\text{CHANGE-UPDATES}}$ allows the set of updates to change arbitrarily.

We now define *DPOC traces* and *weak DPOC traces*, which we later use, along with *DIOC traces* and *weak DIOC traces*, to define our result of correctness.

$$\begin{array}{c}
\frac{\mathcal{N} \xrightarrow{\delta} \mathcal{N}' \quad \delta \neq \mathcal{I}}{\langle \mathbf{I}, \mathcal{N} \rangle \xrightarrow{\delta} \langle \mathbf{I}, \mathcal{N}' \rangle} \quad \left[\text{DPOC} \mid_{\text{LIFT}} \right] \quad \frac{\mathcal{N} \xrightarrow{\mathcal{I}} \mathcal{N}' \quad \mathcal{I} \in \mathbf{I}}{\langle \mathbf{I}, \mathcal{N} \rangle \xrightarrow{\mathcal{I}} \langle \mathbf{I}, \mathcal{N}' \rangle} \quad \left[\text{DPOC} \mid_{\text{LIFT-UP}} \right] \\
\frac{\langle \mathbf{I}, \mathcal{N} \rangle \xrightarrow{\overline{\mathbf{i.o}^?(v)@S:\mathbf{R}}} \langle \mathbf{I}, \mathcal{N}' \rangle \quad \langle \mathbf{I}, \mathcal{N}'' \rangle \xrightarrow{\mathbf{i.o}^?(x \leftarrow v)@R:\mathbf{S}} \langle \mathbf{I}, \mathcal{N}''' \rangle}{\langle \mathbf{I}, \mathcal{N} \parallel \mathcal{N}'' \rangle \xrightarrow{o^?:\mathbf{R}(v) \rightarrow \mathbf{S}(x)} \langle \mathbf{I}, \mathcal{N}' \parallel \mathcal{N}''' \rangle} \quad \left[\text{DPOC} \mid_{\text{SYNCH}} \right] \\
\frac{\langle \mathbf{I}, \mathcal{N} \rangle \xrightarrow{\overline{\mathbf{i.o}^*(X)@S:\mathbf{R}}} \langle \mathbf{I}, \mathcal{N}' \rangle \quad \langle \mathbf{I}, \mathcal{N}'' \rangle \xrightarrow{\mathbf{i.o}^*(\leftarrow X)@R:\mathbf{S}} \langle \mathbf{I}, \mathcal{N}''' \rangle}{\langle \mathbf{I}, \mathcal{N} \parallel \mathcal{N}'' \rangle \xrightarrow{o^*:\mathbf{R}(X) \rightarrow \mathbf{S}()} \langle \mathbf{I}, \mathcal{N}' \parallel \mathcal{N}''' \rangle} \quad \left[\text{DPOC} \mid_{\text{SYNCH-UP}} \right] \\
\frac{\langle \mathbf{I}, \mathcal{N} \rangle \xrightarrow{\eta} \langle \mathbf{I}, \mathcal{N}' \rangle \quad \eta \neq \surd}{\langle \mathbf{I}, \mathcal{N} \parallel \mathcal{N}'' \rangle \xrightarrow{\eta} \langle \mathbf{I}, \mathcal{N}' \parallel \mathcal{N}'' \rangle} \quad \left[\text{DPOC} \mid_{\text{EXT-PARALLEL}} \right] \\
\frac{\langle \mathbf{I}, \mathcal{N} \rangle \xrightarrow{\surd} \langle \mathbf{I}, \mathcal{N}' \rangle \quad \langle \mathbf{I}, \mathcal{N}'' \rangle \xrightarrow{\surd} \langle \mathbf{I}, \mathcal{N}''' \rangle}{\langle \mathbf{I}, \mathcal{N} \parallel \mathcal{N}'' \rangle \xrightarrow{\surd} \langle \mathbf{I}, \mathcal{N}' \parallel \mathcal{N}''' \rangle} \quad \left[\text{DPOC} \mid_{\text{EXT-PAR-END}} \right] \\
\frac{}{\langle \mathbf{I}, \mathcal{N} \rangle \xrightarrow{\mathbf{I}} \langle \mathbf{I}', \mathcal{N} \rangle} \quad \left[\text{DPOC} \mid_{\text{CHANGE-UPDATES}} \right]
\end{array}$$

Figure 7: DPOC system semantics.

Definition 3.3 (DPOC traces). A (*strong*) *trace* of a DPOC system $\langle \mathbf{I}_1, \mathcal{N}_1 \rangle$ is a sequence (finite or infinite) of labels η_1, η_2, \dots with

$$\eta_i \in \{ \tau, o^? : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x), o^* : \mathbf{R}_1(X) \rightarrow \mathbf{R}_2(), \surd, \mathcal{I}, \text{no-up}, \mathbf{I} \}$$

such that there is a sequence of transitions $\langle \mathbf{I}_1, \mathcal{N}_1 \rangle \xrightarrow{\eta_1} \langle \mathbf{I}_2, \mathcal{N}_2 \rangle \xrightarrow{\eta_2} \dots$

A *weak trace* of a DPOC system $\langle \mathbf{I}_1, \mathcal{N}_1 \rangle$ is a sequence of labels η_1, η_2, \dots obtained by removing all the labels corresponding to auxiliary communications, i.e., of the form $o^? : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x)$ or $o^* : \mathbf{R}_1(X) \rightarrow \mathbf{R}_2()$, and the silent labels τ , from a trace of $\langle \mathbf{I}_1, \mathcal{N}_1 \rangle$.

DPOC traces do not allow send and receive actions. Indeed these actions represent incomplete interactions, thus they are needed for compositionality reasons, but they do not represent relevant behaviours of complete systems. Note also that these actions have no correspondence at the DIOC level, where only whole interactions are allowed.

Remark 3.4. Contrarily to DIOCs, DPOCs in general can deadlock. For instance,

$$(\mathbf{i} : \mathbf{i.o} : x \text{ from } \mathbf{R}', \Gamma)_{\mathbf{R}}$$

is a deadlocked DPOC network: the process $\mathbf{i} : \mathbf{i.o} : x \text{ from } \mathbf{R}'$ is not terminated, and the only enabled actions are changes of the set of updates (i.e., transitions with label \mathbf{I}), which are not actual system activities, but are taken by the environment. Notably, the DPOC above cannot be obtained by projecting a DIOC. In fact, DPOCs generated from DIOCs are guaranteed to be deadlock free.

4. PROJECTION FUNCTION

We now introduce the projection function proj . Given a DIOC specification, proj returns a network of DPOC programs that enact the behaviour defined by the originating DIOC.

We write the projection of a DIOC \mathcal{I} as $\text{proj}(\mathcal{I}, \Sigma)$, where Σ is a global state. Informally, the projection of a DIOC is a parallel composition of terms, one for each role of the DIOC. The body of these roles is computed by the *process-projection* function π (defined below). Given a DIOC and a role name \mathbf{R} , the process-projection returns the process corresponding to the local behaviour of role \mathbf{R} . Since the roles executing the process-projections are composed in parallel, the projection of a DIOC program results into the DPOC network of the projected roles.

To give the formal definition of projection, we first define $\parallel_{i \in I} \mathcal{N}_i$ as the parallel composition of networks \mathcal{N}_i for $i \in I$.

Definition 4.1 (Projection). The projection of a DIOC process \mathcal{I} with global state Σ is the DPOC network defined by:

$$\text{proj}(\mathcal{I}, \Sigma) = \parallel_{\mathbf{S} \in \text{roles}(\mathcal{I})} (\pi(\mathcal{I}, \mathbf{S}), \Sigma_{\mathbf{S}})_{\mathbf{S}}$$

The process-projection function that derives DPOC processes from DIOC processes is defined as follows.

Definition 4.2 (Process-projection). Given an annotated DIOC process \mathcal{I} and a role \mathbf{R} the projected DPOC process $\pi(\mathcal{I}, \mathbf{R})$ is defined as in Figure 8.

With a little abuse of notation, we write $\text{roles}(\mathcal{I}, \mathcal{I}')$ for $\text{roles}(\mathcal{I}) \cup \text{roles}(\mathcal{I}')$. We assume that variables $x_{\mathbf{i}}$ are never used in the DIOC to be projected and we use them for auxiliary synchronisations.

The projection is homomorphic for sequential and parallel composition, $\mathbf{1}$ and $\mathbf{0}$. The projection of an assignment is the assignment on the role performing it and $\mathbf{1}$ on other roles. The projection of an interaction is a send on the sender role, a receive on the receiver, and $\mathbf{1}$ on any other role. The projection of a scope is a scope on all its participants. On its coordinator it also features a clause that records the roles of the involved participants. On the roles not involved in the scope the projection is $\mathbf{1}$. Projections of conditional and while loop are a bit more complex, since they need to coordinate a distributed computation. To this end they exploit communications on auxiliary operations. In particular, $\text{cnd}_{\mathbf{i}}^*$ coordinates the branching of conditionals, carrying information on whether the “then” or the “else” branch needs to be taken. Similarly, $\text{wb}_{\mathbf{i}}^*$ coordinates the beginning of a while loop, carrying information on whether to loop or to exit. Finally, $\text{we}_{\mathbf{i}}^*$ coordinates the end of the body of the while loop. This closing operation carries no relevant information and it is just used for synchronisation purposes. In order to execute a conditional \mathbf{i} : **if** $b@_{\mathbf{R}} \{\mathcal{I}\}$ **else** $\{\mathcal{I}'\}$, the coordinator \mathbf{R} of the conditional locally evaluates the guard and tells the other roles which branch to choose using auxiliary communications on $\text{cnd}_{\mathbf{i}}^*$. Finally, all the roles involved in the conditional execute their code corresponding to the chosen branch. Execution of a loop \mathbf{i} : **while** $b@_{\mathbf{R}} \{\mathcal{I}\}$ is similar, with two differences. First, end of loop synchronisation on operations $\text{we}_{\mathbf{i}}^*$ is used to notify the coordinator that an iteration is terminated, and a new one can start. Second, communication of whether to loop or to exit is more tricky than communication on the branch to choose in a conditional. Indeed, there are two points in the projected code where the coordinator \mathbf{R} sends the decision: the first is inside the body of the loop and it is used if the decision is to loop; the second is after the loop and it is

$$\begin{array}{l}
\boxed{\pi(\mathbf{1}, \mathbf{S})} = \mathbf{1} \quad \boxed{\pi(\mathbf{0}, \mathbf{S})} = \mathbf{0} \\
\boxed{\pi(\mathcal{I}; \mathcal{I}', \mathbf{S})} = \pi(\mathcal{I}, \mathbf{S}); \pi(\mathcal{I}', \mathbf{S}) \quad \boxed{\pi(\mathcal{I}|\mathcal{I}', \mathbf{S})} = \pi(\mathcal{I}, \mathbf{S})|\pi(\mathcal{I}', \mathbf{S}) \\
\boxed{\pi(\mathbf{i}: x @ \mathbf{R} = e, \mathbf{R})} = \mathbf{i}: x = e \\
\boxed{\pi(\mathbf{i}: x @ \mathbf{R} = e, \mathbf{S}) \text{ and } \mathbf{S} \neq \mathbf{R}} = \mathbf{1} \\
\boxed{\pi(\mathbf{i}: o : \mathbf{R}_1(e) \rightarrow \mathbf{R}_2(x), \mathbf{R}_1)} = \mathbf{i}: \mathbf{i}.o : e \text{ to } \mathbf{R}_2 \\
\boxed{\pi(\mathbf{i}: o : \mathbf{R}_1(e) \rightarrow \mathbf{R}_2(x), \mathbf{R}_2)} = \mathbf{i}: \mathbf{i}.o : x \text{ from } \mathbf{R}_1 \\
\boxed{\pi(\mathbf{i}: o : \mathbf{R}_1(e) \rightarrow \mathbf{R}_2(x), \mathbf{S})} \\
\text{and } \mathbf{S} \notin \{\mathbf{R}_1, \mathbf{R}_2\} = \mathbf{1} \\
\boxed{\pi(\mathbf{i}: \text{if } b @ \mathbf{R} \{ \mathcal{I} \} \text{ else } \{ \mathcal{I}' \}, \mathbf{R})} = \left\{ \begin{array}{l} \mathbf{i}: \text{if } b \left\{ \left(\prod_{\mathbf{R}' \in \text{roles}(\mathcal{I}, \mathcal{I}') \setminus \{ \mathbf{R} \}} \mathbf{i}_T: \mathbf{i}.cnd_i^* : \text{true to } \mathbf{R}' \right); \right. \\ \left. \pi(\mathcal{I}, \mathbf{R}) \right\} \\ \text{else} \left\{ \left(\prod_{\mathbf{R}' \in \text{roles}(\mathcal{I}, \mathcal{I}') \setminus \{ \mathbf{R} \}} \mathbf{i}_F: \mathbf{i}.cnd_i^* : \text{false to } \mathbf{R}' \right); \right. \\ \left. \pi(\mathcal{I}', \mathbf{R}) \right\} \end{array} \right\} \\
\boxed{\pi(\mathbf{i}: \text{if } b @ \mathbf{R} \{ \mathcal{I} \} \text{ else } \{ \mathcal{I}' \}, \mathbf{S})} \\
\text{and } \mathbf{S} \in \text{roles}(\mathcal{I}, \mathcal{I}') \setminus \{ \mathbf{R} \} = \mathbf{i}_T: \mathbf{i}.cnd_i^* : x_i \text{ from } \mathbf{R}; \mathbf{i}: \text{if } x_i \{ \pi(\mathcal{I}, \mathbf{S}) \} \text{ else } \{ \pi(\mathcal{I}', \mathbf{S}) \} \\
\boxed{\pi(\mathbf{i}: \text{if } b @ \mathbf{R} \{ \mathcal{I} \} \text{ else } \{ \mathcal{I}' \}, \mathbf{S})} \\
\text{and } \mathbf{S} \notin \text{roles}(\mathcal{I}, \mathcal{I}') \cup \{ \mathbf{R} \} = \mathbf{1} \\
\boxed{\pi(\mathbf{i}: \text{while } b @ \mathbf{R} \{ \mathcal{I} \}, \mathbf{R})} = \left\{ \begin{array}{l} \mathbf{i}: \text{while } b \left\{ \right. \\ \left. \left(\prod_{\mathbf{R}' \in \text{roles}(\mathcal{I}) \setminus \{ \mathbf{R} \}} \mathbf{i}_T: \mathbf{i}.wb_i^* : \text{true to } \mathbf{R}' \right); \pi(\mathcal{I}, \mathbf{R}); \right. \\ \left. \prod_{\mathbf{R}' \in \text{roles}(\mathcal{I}) \setminus \{ \mathbf{R} \}} \mathbf{i}_C: \mathbf{i}.we_i^* : _ \text{ from } \mathbf{R}' \right. \\ \left. \left. \right\}; \prod_{\mathbf{R}' \in \text{roles}(\mathcal{I}) \setminus \{ \mathbf{R} \}} \mathbf{i}_F: \mathbf{i}.wb_i^* : \text{false to } \mathbf{R}' \right. \\ \left. \mathbf{i}_T: \mathbf{i}.wb_i^* : x_i \text{ from } \mathbf{R}; \right. \\ \left. \mathbf{i}: \text{while } x_i \left\{ \begin{array}{l} \pi(\mathcal{I}, \mathbf{S}); \\ \mathbf{i}_C: \mathbf{i}.we_i^* : \text{ok to } \mathbf{R}; \\ \mathbf{i}_T: \mathbf{i}.wb_i^* : x_i \text{ from } \mathbf{R} \end{array} \right\} \right. \\ \left. \right\} \\
\boxed{\pi(\mathbf{i}: \text{while } b @ \mathbf{R} \{ \mathcal{I} \}, \mathbf{S})} \\
\text{and } \mathbf{S} \in \text{roles}(\mathcal{I}) \setminus \{ \mathbf{R} \} = \left\{ \begin{array}{l} \mathbf{i}_T: \mathbf{i}.wb_i^* : x_i \text{ from } \mathbf{R}; \\ \mathbf{i}: \text{while } x_i \left\{ \begin{array}{l} \pi(\mathcal{I}, \mathbf{S}); \\ \mathbf{i}_C: \mathbf{i}.we_i^* : \text{ok to } \mathbf{R}; \\ \mathbf{i}_T: \mathbf{i}.wb_i^* : x_i \text{ from } \mathbf{R} \end{array} \right\} \end{array} \right\} \\
\boxed{\pi(\mathbf{i}: \text{while } b @ \mathbf{R} \{ \mathcal{I} \}, \mathbf{S})} \\
\text{and } \mathbf{S} \notin \text{roles}(\mathcal{I}) \cup \{ \mathbf{R} \} = \mathbf{1} \\
\boxed{\pi(\mathbf{i}: \text{scope } @ \mathbf{R} \{ \mathcal{I} \}, \mathbf{R})} = \mathbf{i}: \text{scope } @ \mathbf{R} \{ \pi(\mathcal{I}, \mathbf{R}) \} \text{ roles } \{ \text{roles}(\mathcal{I}) \} \\
\boxed{\pi(\mathbf{i}: \text{scope } @ \mathbf{R} \{ \mathcal{I} \}, \mathbf{S})} \\
\text{and } \mathbf{S} \in \text{roles}(\mathcal{I}) \setminus \{ \mathbf{R} \} = \mathbf{i}: \text{scope } @ \mathbf{R} \{ \pi(\mathcal{I}, \mathbf{S}) \} \\
\boxed{\pi(\mathbf{i}: \text{scope } @ \mathbf{R} \{ \mathcal{I} \}, \mathbf{S})} \\
\text{and } \mathbf{S} \notin \text{roles}(\mathcal{I}) \cup \{ \mathbf{R} \} = \mathbf{1}
\end{array}$$

Figure 8: process-projection function π .

used if the decision is to exit. Also, there are two points where these communications are received by the other roles: before their loop at the first iteration, at the end of the previous iteration of the body of the loop in the others.

One has to keep attention since, by splitting an interaction into a send and a receive primitive, primitives corresponding to different interactions, but on the same operation, may interfere.

Example 4.3. We illustrate the issue of interferences using the two DPOC processes below, identified by their respective roles, \mathbf{R}_1 (right) and \mathbf{R}_2 (left), assuming that operations are not prefixed by indexes. We describe only \mathbf{R}_1 as \mathbf{R}_2 is its dual. At Line 1, \mathbf{R}_1 sends a message to \mathbf{R}_2 on operation o . In parallel with the send, \mathbf{R}_1 had a scope (Lines 3–5) that performed an update. The new code (Line 4) contains a send on operation o to role \mathbf{R}_2 . Since the two sends and the two receives share the same operation o and run in parallel, they can interfere with each other.

<p style="margin: 0;"><u>process \mathbf{R}_1</u></p> <ol style="list-style-type: none"> 1. 1: $o : e_1$ to \mathbf{R}_2 2. 3. // update auxiliary code 4. 2: $o : e_2$ to \mathbf{R}_2 5. // update auxiliary code 	<p style="margin: 0;"><u>process \mathbf{R}_2</u></p> <ol style="list-style-type: none"> 1. 1: $o : x_1$ from \mathbf{R}_2 2. 3. // update auxiliary code 4. 2: $o : x_2$ from \mathbf{R}_2 5. // update auxiliary code
---	---

Note that, since updates come from outside and one cannot know in advance which operations they use, this interference cannot be statically avoided.

For this reason, in § 3 we introduced indexes to prefix DPOC operations.

A similar problem may occur also for auxiliary communications. In particular, imagine to have two parallel conditionals executed by the same role. We need to avoid that, e.g., the decision to take the “else” branch on the first conditional is wrongly taken by some role as a decision concerning the second conditional. To avoid this problem, we prefix auxiliary operations using the index \mathbf{i} of the conditional. In this way, communications involving distinct conditionals cannot interact. Note that communications concerning the same conditional (or while loop) may share the same operation name and prefix. However, since all auxiliary communications are from the coordinator of the construct to the other roles involved in it, or vice versa, interferences are avoided.

We now describe how to generate indexes for statements in the projection. As a general rule, all the DPOC constructs obtained by projecting a DIOC construct with index \mathbf{i} have index \mathbf{i} . The only exceptions are the indexes of the auxiliary communications of the projection of conditionals and while loops.

Provided \mathbf{i} is the index of the conditional: *i*) in the projection of the coordinator we index the auxiliary communications for selecting the “then” branch with index \mathbf{i}_T , the ones for selecting the “else” branch with index \mathbf{i}_F ; *ii*) in the projection of the other roles involved in the conditional we assign the index $\mathbf{i}_?$ to the auxiliary receive communications. To communicate the evaluation of the guard of a while loop we use the same indexing scheme (\mathbf{i}_T , \mathbf{i}_F , and $\mathbf{i}_?$) used in the projection of conditional. Moreover, all the auxiliary communications for end of loop synchronisation are indexed with \mathbf{i}_C .

5. RUNNING EXAMPLE: PROJECTION AND EXECUTION

In this section we use our running example (see Figure 2) to illustrate the projection and execution of DIOC programs.

5.1. Projection. Given the code in Figure 2, we need to annotate it to be able to project it (we remind that in § 4 we defined our projection function on well-annotated DIOCs). Since we wrote one instruction per line in Figure 2, we annotate every instruction using its line number as index. This results in a well-annotated DIOC.

From the annotated DIOC, the projection generates three DPOC processes for the **Seller**, the **Buyer**, and the **Bank**, respectively reported in Figures 9, 11, and 10. To improve readability, we omit some **1** processes. In the projection of the program, we also omit to write the index that prefixes the operations since it is always equal to the numeric part of the index of their correspondent construct. Finally, we write auxiliary communications in grey.

5.2. Runtime Execution. We now focus on an excerpt of the code to exemplify how updates are performed at runtime. We consider the code of the scope at Lines 6–9 of Figure 2. In this execution scenario we assume to introduce in the set of available updates the update presented in Figure 1, which enables the use of a fidelity card to provide a price discount. Below we consider both the DIOC and the DPOC level, dropping some **1**s to improve readability.

Since we describe a runtime execution, we assume that the **Buyer** has just sent the name of the product (s)he is interested in to the **Seller** (Line 5 of Figure 2). The annotated DIOCs we execute is the following.

```

1  6: scope @Seller{
2    7: order_price@Seller = getPrice(order);
3    8: offer : Seller(order_price) → Buyer(prod_price)
4  }
```

At runtime we apply Rule $[\text{DIOC} \mid_{\text{UP}}]$ that substitutes the scope with the new code. The replacement is atomic. Below we assume that the instructions of the update are annotated with indexes corresponding to their line number plus 30.

```

1  31: cardReq : Seller(null) → Buyer(_);
2  32: card_id@Buyer = getInput();
3  33: card : Buyer(card_id) → Seller(buyer_id);
4  34: if isValid(buyer_id)@Seller{
5    35: order_price@Seller = getPrice(order) * 0.9
6  } else {
7    37: order_price@Seller = getPrice(order)
8  };
9  39: offer : Seller(order_price) → Buyer(prod_price)
```

Let us now focus on the execution at DPOC level, where the application of updates is not atomic. The scope is distributed between two participants. The first step of the

```

1  3?:  $wb_3^* : x_3$  from Buyer;
2  3: while( $x_3$ ){
3    5:  $priceReq : order$  from Buyer;
4    6: scope @Seller{
5      7:  $order\_price = getPrice(order)$ ;
6      8:  $offer : order\_price$  to Buyer
7    } roles { Seller, Buyer };
8    3C:  $we_3^* : ok$  to Buyer;
9    3?:  $wb_3^* : x_3$  from Buyer
10 };
11 15?:  $cmd_{15}^* : x_{15}$  from Buyer;
12 15: if( $x_{15}$ ){
13   16:  $payReq : payDesc(order\_price)$  to Bank;
14   21?:  $cmd_{21}^* : x_{21}$  from Bank;
15   21: if( $x_{21}$ ){
16     22:  $confirm : \_$  from Bank }
17 }

```

Figure 9: Seller DPOC Process.

```

1  15?:  $cmd_{15}^* : x_{15}$  from Buyer;
2  15: if( $x_{15}$ ){
3    16:  $payReq : desc$  from Seller;
4    17: scope @Bank{
5      18:  $pay : auth$  from Buyer
6    } roles { Buyer, Bank };
7    20:  $payment\_ok = makePayment(desc, auth)$ ;
8    21: if( $payment\_ok$ ){
9      {
10     21T:  $cmd_{21}^* : true$  to Seller
11     | 21T:  $cmd_{21}^* : true$  to Buyer
12     } ; {
13     22:  $confirm : null$  to Seller
14     | 24:  $confirm : null$  to Buyer
15     }
16   } else {
17     {
18     21F:  $cmd_{21}^* : false$  to Seller
19     | 21F:  $cmd_{21}^* : false$  to Buyer
20     };
21     26:  $abort : null$  to Buyer
22   }
23 }

```

Figure 10: Bank DPOC Process.

```

1  1:  $price\_ok = false$ ;
2  2:  $continue = true$ ;
3  3: while(! $price\_ok$  and  $continue$ ){
4    3T:  $wb_3^* : true$  to Seller;
5    4:  $prod = getInput()$ ;
6    5:  $priceReq : prod$  to Seller;
7    6: scope @Seller{
8      7:  $offer : prod\_price$  from Seller
9    };
10   10:  $price\_ok = getInput()$ ;
11   11: if(! $price\_ok$ ){
12     12:  $continue = getInput()$ 
13   };
14   3C:  $we_3^* : \_$  from Seller
15 };
16 3F:  $wb_3^* : false$  to Seller;
17 15: if( $price\_ok$ ){
18   {
19     15T:  $cmd_{15}^* : true$  to Seller
20     | 15T:  $cmd_{15}^* : true$  to Bank
21     };
22   17: scope @Bank{
23     19:  $pay : payAuth(prod\_price)$  to Bank
24   };
25   21?:  $cmd_{21}^* : x_{21}$  from Bank;
26   21: if( $x_{21}$ ){
27     24:  $confirm : \_$  from Bank
28   } else {
29     26:  $abort : \_$  from Bank}
30 }
31 }

```

Figure 11: Buyer DPOC Process.

update protocol is performed by the **Seller**, since (s)he is the coordinator of the update. The DPOC description of the **Seller** before the update is:

```

6: scope @Seller{
  7: order_price = getPrice(order);
  8: offer : order_price to Buyer
} roles {Seller, Buyer}

```

When the scope construct is enabled, the **Seller** non-deterministically selects whether to update or not and, in the first case, which update to apply. Here, we assume that the update using the code in Figure 1 is selected. Below we report on the left the reductum of the projected code of the **Seller** after the application of Rule $[\text{DPOC}_{\text{LEAD-UP}}]$. The **Seller** sends to the **Buyer** the code — denoted as $P_{\mathbf{B}}$ and reported below on the right — obtained projecting the update on role **Buyer**.

<pre> 1 6: sb₆[*] : P_B to Buyer; 2 31: cardReq : null to Buyer; 3 33: card : buyer_id from Buyer; 4 34: if isValid(buyer_id){ 5 35: order_price = getPrice(order) * 0.9 6 } else { 7 37: order_price = getPrice(order) 8 }; 9 39: offer : order_price to Buyer; 10 6: se₆[*] : _ from Buyer; </pre>	<pre> P_B := 31: cardReq : null from Seller; 32: card_id = getInput(); 33: card : card_id to Seller; 39: offer : prod_price from Seller </pre>
--	--

Above, at Line 1 the **Seller** requires the **Buyer** to update, sending to him the new DPOC fragment to execute. Then, the **Seller** starts to execute its own updated DPOC. At the end of the execution of the new DPOC code (Line 10) the **Seller** waits for the notification of termination of the DPOC fragment executed by the **Buyer**.

Let us now consider the process-projection of the **Buyer**, reported below.

```

6: scope @Seller{
  8: offer : order_price from Seller
}

```

At runtime, the scope waits for the arrival of a message from the coordinator of the update. In our case, since we assumed that the update is applied, the **Buyer** receives using Rule $[\text{DPOC}_{\text{UP}}]$ the DPOC fragment $P_{\mathbf{B}}$ sent by the coordinator. In the reductum, $P_{\mathbf{B}}$ replaces the scope, followed by the notification of termination to the **Seller**.

```

31: cardReq : null from Seller;
32: card_id = getInput();
33: card : card_id to Seller;
39: offer : prod_price from Seller
6: se6* : ok to Seller

```

Consider now what happens if no update is applied. At DIOC level the **Seller** applies Rule $[\text{DIOC} \mid_{\text{NoUP}}]$, which removes the scope and runs its body. At DPOC level, the update is not atomic. The code of the **Seller** is the following one.

```

1 6:  $sb_6^*$  : no to Buyer;
2 7:  $order\_price = \text{getPrice}(order)$ ;
3 8:  $offer : order\_price$  to Buyer;
4 6:  $se_6^*$  : _ from Buyer;

```

Before executing the code inside the scope, the **Seller** notifies the **Buyer** that (s)he can proceed with her execution (Line 1). Like in the case of update, the **Seller** also waits for the notification of the end of execution from the **Buyer** (Line 4).

Finally, we report the DPOC code of the **Buyer** after the reception of the message that no update is needed. Rule $[\text{DPOC} \mid_{\text{NoUP}}]$ removes the scope and adds the notification of termination (Line 2 below) to the coordinator at the end.

```

1 7:  $offer : prod\_price$  from Seller;
2 6:  $se_6^*$  : ok to Seller;

```

6. CONNECTED DIOCS

We now give a precise definition of the notion of connectedness that we mentioned in § 2.2 and § 3.2. In both DIOC and DPOC semantics we checked such a property of updates with predicate *connected*, respectively in Rule $[\text{DIOC} \mid_{\text{UP}}]$ (Figure 3) and Rule $[\text{DPOC} \mid_{\text{UP-LEAD}}]$ (Figure 6).

To give the intuition of why we need to restrict to connected updates, consider the scenario below of a DIOC (left side) and its projection (right side).

$$\begin{array}{l}
op1 : \mathbf{A}(e_1) \rightarrow \mathbf{B}(x); \\
op2 : \mathbf{C}(e_2) \rightarrow \mathbf{D}(y)
\end{array}
\quad \xrightarrow{\text{projection}} \quad
\begin{array}{c|c}
\frac{\text{process } \mathbf{A}}{op1 : e_1 \text{ to } \mathbf{B}} & \frac{\text{process } \mathbf{B}}{op1 : x \text{ from } \mathbf{A}} \\
\hline
\frac{\text{process } \mathbf{C}}{op2 : e_2 \text{ to } \mathbf{D}} & \frac{\text{process } \mathbf{D}}{op2 : y \text{ from } \mathbf{C}}
\end{array}$$

DIOCs can express interactions that, if projected as described in § 4, can behave differently with respect to the originating DIOC. Indeed, in our example we have a DIOC that composes in sequence two interactions: an interaction between **A** and **B** on operation op_1 followed by an interaction between **C** and **D** on operation op_2 . The projection of the DIOC produces four processes (identified by their role): **A** and **C** send a message to **B** and **D**, respectively. Dually, **B** and **D** receive a message from **A** and **C**, respectively. In the example, at the level of processes we lose the global order among the interactions: each projected process runs its code locally and it is not aware of the global sequence of interactions. Indeed, both sends and both receives are enabled at the same time. Hence, the semantics of DPOC lets the two interactions interleave in any order. It can happen that the interaction between **C** and **D** occurs before the one between **A** and **B**, violating the order of interactions prescribed by the originating DIOC.

Restricting to connected DIOCs avoids this kind of behaviours. We formalise *connectedness* as an efficient (see Theorem 6.3) syntactic check. We highlight that our definition of connectedness does not hamper programmability and it naturally holds in most real-world

scenarios (the interested reader can find in the website of the AIOCJ project [1] several such scenarios).

Remark 6.1. There exists a trade-off between efficiency and ease of programming with respect to the guarantee that all the roles are aware of the evolution of the global computation. This is a common element of choreographic approaches, which has been handled in different ways, e.g., *i*) by restricting the set of well-formed choreographies to only those on which the projection preserves the order of actions [10]; *ii*) by mimicking the non-deterministic behaviour of process-level networks at choreography level [13]; or *iii*) by enforcing the order of actions with additional auxiliary messages between roles [31].

Our choice of preserving the order of interactions defined at DIOC level follows the same philosophy of [10], whilst for scopes, conditionals, and while loops we enforce connectedness with auxiliary messages as done in [31]. We remind that we introduced auxiliary messages for coordination both in the semantics of scopes at DPOC level (§ 3.2) and in the projection (§ 4). We choose to add such auxiliary messages to avoid to impose strong constraints on the form of scopes, conditionals, and while loops, which in the end would pose strong limitations to the programmers of DIOCs. On the other hand, for sequential composition we choose to restrict the set of allowed *DIOCs* by requiring connectedness, which ensures that the order of interactions defined at DIOC level is preserved by projection.

As discussed above, the execution of conditionals and while loops rely on auxiliary communications to coordinate the different roles. Some of these communications may be redundant. For instance, in Figure 9, Line 8, the **Seller** notifies to the **Buyer** that (s)he has completed her part of the while loop. However, since her last contribution to the while loop is the auxiliary communication for end of scope synchronisation, the **Buyer** already has this information. Hence, Line 8 in Figure 9, where the notification is sent, and Line 14 in Figure 11, where the notification is received, can be safely dropped. Removing redundant auxiliary communications can be automatised using a suitable static analysis. We leave this topic for future work.

To formalise connectedness we introduce, in Figure 12, the auxiliary functions `transl` and `transF` that, given a DIOC process, compute sets of pairs representing senders and receivers of possible initial and final interactions in its execution. We represent one such pair as $\mathbf{R} \rightarrow \mathbf{S}$. Actions located at \mathbf{R} are represented as $\mathbf{R} \rightarrow \mathbf{R}$. For instance, given an interaction $\mathbf{i}: o : \mathbf{R}(e) \rightarrow \mathbf{S}(x)$ both its `transl` and `transF` are $\{\mathbf{R} \rightarrow \mathbf{S}\}$. For conditional, `transl(i: if $b@R$ { \mathcal{I} } else { \mathcal{I}' })` = $\{\mathbf{R} \rightarrow \mathbf{R}\}$ since the first action executed is the evaluation of the guard by role \mathbf{R} . The set `transF(i: if $b@R$ { \mathcal{I} } else { \mathcal{I}' })` is normally `transF(\mathcal{I})` \cup `transF(\mathcal{I}')`, since the execution terminates with an action from one of the branches. If instead the branches are both empty then `transF` is $\{\mathbf{R} \rightarrow \mathbf{R}\}$, representing guard evaluation.

Finally, we give the formal definition of connectedness.

Definition 6.2 (Connectedness). A DIOC process \mathcal{I} is connected if each subterm \mathcal{I}' ; \mathcal{I}'' of \mathcal{I} satisfies

$$\forall \mathbf{R}_1 \rightarrow \mathbf{R}_2 \in \text{transF}(\mathcal{I}'), \forall \mathbf{S}_1 \rightarrow \mathbf{S}_2 \in \text{transl}(\mathcal{I}'') . \{\mathbf{R}_1, \mathbf{R}_2\} \cap \{\mathbf{S}_1, \mathbf{S}_2\} \neq \emptyset$$

Connectedness can be checked efficiently.

Theorem 6.3 (Connectedness-check complexity).

The connectedness of a DIOC process \mathcal{I} can be checked in time $O(n^2 \log(n))$, where n is the number of nodes in the abstract syntax tree of \mathcal{I} .

$$\begin{aligned}
\text{transl}(\mathbf{i}: o : \mathbf{R}(e) \rightarrow \mathbf{S}(x)) &= \text{transF}(\mathbf{i}: o : \mathbf{R}(e) \rightarrow \mathbf{S}(x)) = \{\mathbf{R} \rightarrow \mathbf{S}\} \\
\text{transl}(\mathbf{i}: x @ \mathbf{R} = e) &= \text{transF}(\mathbf{i}: x @ \mathbf{R} = e) = \{\mathbf{R} \rightarrow \mathbf{R}\} \\
\text{transl}(\mathbf{1}) &= \text{transl}(\mathbf{0}) = \text{transF}(\mathbf{1}) = \text{transF}(\mathbf{0}) = \emptyset \\
\text{transl}(\mathcal{I} | \mathcal{I}') &= \text{transl}(\mathcal{I}) \cup \text{transl}(\mathcal{I}') \\
\text{transF}(\mathcal{I} | \mathcal{I}') &= \text{transF}(\mathcal{I}) \cup \text{transF}(\mathcal{I}') \\
\text{transl}(\mathcal{I}; \mathcal{I}') &= \begin{cases} \text{transl}(\mathcal{I}') & \text{if } \text{transl}(\mathcal{I}) = \emptyset \\ \text{transl}(\mathcal{I}) & \text{otherwise} \end{cases} \\
\text{transF}(\mathcal{I}; \mathcal{I}') &= \begin{cases} \text{transF}(\mathcal{I}) & \text{if } \text{transF}(\mathcal{I}') = \emptyset \\ \text{transF}(\mathcal{I}') & \text{otherwise} \end{cases} \\
\text{transl}(\mathbf{i}: \text{if } b @ \mathbf{R} \{ \mathcal{I} \} \text{ else } \{ \mathcal{I}' \}) &= \text{transl}(\mathbf{i}: \text{while } b @ \mathbf{R} \{ \mathcal{I} \}) = \{\mathbf{R} \rightarrow \mathbf{R}\} \\
\text{transF}(\mathbf{i}: \text{if } b @ \mathbf{R} \{ \mathcal{I} \} \text{ else } \{ \mathcal{I}' \}) &= \begin{cases} \{\mathbf{R} \rightarrow \mathbf{R}\} & \text{if } \text{transF}(\mathcal{I}) \cup \text{transF}(\mathcal{I}') = \emptyset \\ \text{transF}(\mathcal{I}) \cup \text{transF}(\mathcal{I}') & \text{otherwise} \end{cases} \\
\text{transF}(\mathbf{i}: \text{while } b @ \mathbf{R} \{ \mathcal{I} \}) &= \begin{cases} \{\mathbf{R} \rightarrow \mathbf{R}\} & \text{if } \text{transF}(\mathcal{I}) = \emptyset \\ \bigcup_{\mathbf{R}' \in \text{roles}(\mathcal{I}) \setminus \{\mathbf{R}\}} \{\mathbf{R} \rightarrow \mathbf{R}'\} & \text{otherwise} \end{cases} \\
\text{transl}(\mathbf{i}: \text{scope } @ \mathbf{R} \{ \mathcal{I} \}) &= \{\mathbf{R} \rightarrow \mathbf{R}\} \\
\text{transF}(\mathbf{i}: \text{scope } @ \mathbf{R} \{ \mathcal{I} \}) &= \begin{cases} \{\mathbf{R} \rightarrow \mathbf{R}\} & \text{if } \text{roles}(\mathcal{I}) \subseteq \{\mathbf{R}\} \\ \bigcup_{\mathbf{R}' \in \text{roles}(\mathcal{I}) \setminus \{\mathbf{R}\}} \{\mathbf{R}' \rightarrow \mathbf{R}\} & \text{otherwise} \end{cases}
\end{aligned}$$

Figure 12: Auxiliary functions `transl` and `transF`.

The proof of the theorem is reported in Appendix A.

We remind that we allow only connected updates. Indeed, replacing a scope with a connected update always results in a deadlock- and race-free DIOC. Thus, one just needs to statically check connectedness of the starting program and of the updates, and there is no need to perform expensive runtime checks on the whole application after updates have been performed.

7. CORRECTNESS

In the previous sections we have presented DIOCs, DPOCs, and described how to derive a DPOC from a given DIOC. This section presents the main technical result of the paper, namely the correctness of the projection. Moreover, as a consequence of the correctness, in Section 7.2 we prove that properties like deadlock freedom, termination, and race freedom are preserved by the projection.

Correctness here means that a connected DIOC and its projected DPOC are weak system bisimilar. Weak system bisimilarity is formally defined as follows.

Definition 7.1 (Weak System Bisimilarity). *A weak system bisimulation is a relation \mathcal{R} between DIOC systems and DPOC systems such that if $(\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle, \langle \mathbf{I}', \mathcal{N} \rangle) \in \mathcal{R}$ then:*

- if $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle \xrightarrow{\mu} \langle \Sigma'', \mathbf{I}'', \mathcal{I}'' \rangle$ then $\langle \mathbf{I}', \mathcal{N} \rangle \xrightarrow{\eta_1}, \dots, \xrightarrow{\eta_k} \xrightarrow{\mu} \langle \mathbf{I}''', \mathcal{N}''' \rangle$ with $\forall i \in [1 \dots k], \eta_i \in \{o^* : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x), o^* : \mathbf{R}_1(X) \rightarrow \mathbf{R}_2(), \tau\}$ and $(\langle \Sigma'', \mathbf{I}'', \mathcal{I}'' \rangle, \langle \mathbf{I}''', \mathcal{N}''' \rangle) \in \mathcal{R}$;
- if $\langle \mathbf{I}', \mathcal{N} \rangle \xrightarrow{\eta} \langle \mathbf{I}''', \mathcal{N}''' \rangle$ with $\eta \in \{o^? : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x), o^* : \mathbf{R}_1(X) \rightarrow \mathbf{R}_2(), \sqrt{\cdot}, \mathcal{I}, \mathbf{no-up}, \mathbf{I}''', \tau\}$ then one of the following two conditions holds:
 - $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle \xrightarrow{\eta} \langle \Sigma'', \mathbf{I}', \mathcal{I}'' \rangle$ and $(\langle \Sigma'', \mathbf{I}'', \mathcal{I}'' \rangle, \langle \mathbf{I}''', \mathcal{N}''' \rangle) \in \mathcal{R}$ or
 - $\eta \in \{o^* : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x), o^* : \mathbf{R}_1(X) \rightarrow \mathbf{R}_2(), \tau\}$ and $(\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle, \langle \mathbf{I}''', \mathcal{N}''' \rangle) \in \mathcal{R}$

A DIOC system $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle$ and a DPOC system $\langle \mathbf{I}', \mathcal{N} \rangle$ are *weak system bisimilar* iff there exists a weak system bisimulation \mathcal{R} such that $(\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle, \langle \mathbf{I}', \mathcal{N} \rangle) \in \mathcal{R}$.

In the proof, we provide a relation \mathcal{R} which relates each well-annotated connected DIOC system with its projection and show that it is a weak system bisimulation. Such a relation is not trivial since events that are atomic in the DIOC, e.g., the evaluation of the guard of a conditional, including the removal of the discarded branch, are not atomic at DPOC level. In the case of conditional, the DIOC transition is mimicked by a conditional performed by the role evaluating the guard, a set of auxiliary communications sending the value of the guard to the other roles, and local conditionals based on the received value. These mismatches are taken care by function \mathbf{upd} (Definition 7.18). This function needs also to remove the auxiliary communications used to synchronise the termination of scopes, which have no counterpart after the DIOC scope has been consumed. However, we have to record the impact of the mentioned auxiliary communications on the possible executions. Thus we define an event structure for DIOC (Definition 7.10) and one for DPOC (Definition 7.11) and we show that the two are related (Lemma 7.15).

Thanks to the existence of a bisimulation relating each well-annotated connected DIOC system with its projection we can prove that the projection is correct. Formally:

Theorem 7.2 (Correctness). For each initial, connected DIOC process \mathcal{I} , each state Σ , each set of updates \mathbf{I} , the DIOC system $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle$ and the DPOC system $\langle \mathbf{I}, \mathbf{proj}(\mathcal{I}, \Sigma) \rangle$ are weak system bisimilar.

As a corollary of the result above, a DIOC system and its projection are also trace equivalent. Trace equivalence is defined as follows.

Definition 7.3 (Trace equivalence). A DIOC system $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle$ and a DPOC system $\langle \mathbf{I}, \mathcal{N} \rangle$ are *(weak) trace equivalent* iff their sets of (weak) traces coincide.

The following lemma shows that, indeed, weak system bisimilarity implies weak trace equivalence.

Lemma 7.4. *Let $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle$ be a DIOC system and $\langle \mathbf{I}', \mathcal{N} \rangle$ a DPOC system. If $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle \sim \langle \mathbf{I}', \mathcal{N} \rangle$ then the DIOC system $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle$ and the DPOC system $\langle \mathbf{I}', \mathcal{N} \rangle$ are weak trace equivalent.*

Proof. The proof is by coinduction. Take a DIOC trace μ_1, μ_2, \dots of the DIOC system. From bisimilarity, the DPOC system has a sequence of transitions with labels $\eta_1, \dots, \eta_k, \mu_1$ where η_1, \dots, η_k are weak transitions. Hence, the first label in the weak trace is μ_1 . After the transition with label μ_1 , the DIOC system and the DPOC system are again bisimilar. By coinductive hypothesis, the DPOC system has a weak trace μ_2, \dots . By composition the DPOC system has a trace μ_1, μ_2, \dots as desired. The opposite direction is similar. \square

Hence the following corollary holds.

Corollary 7.5 (Trace Equivalence). *For each initial, connected DIOC process \mathcal{I} , each state Σ , each set of updates \mathbf{I} , the DIOC system $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle$ and the DPOC system $\langle \mathbf{I}, \text{proj}(\mathcal{I}, \Sigma) \rangle$ are weak trace equivalent.*

Proof. It follows from Theorem 7.2 and Lemma 7.4. \square

The following section presents the details of the proof of Theorem 7.2. The reader not interested in such details can safely skip § 7.1 and go to § 7.2.

7.1. Detailed Proof of Correctness. In our proof strategy we rely on the distinctness of indexes of DIOC constructs that, unfortunately, is not preserved by transitions due to while unfolding.

Example 7.6. Consider the DIOC $\mathbf{i}: \text{while } b@R \{ \mathbf{j}: x@R = e \}$. If the condition b evaluates to true, in one step the application of Rule $[\text{DPOC} \mid_{\text{WHILE-UNFOLD}}]$ produces the DIOC \mathcal{I} below

$$\mathbf{j}: x@R = e; \mathbf{i}: \text{while } b@R \{ \mathbf{j}: x@R = e \}$$

where the index \mathbf{j} occurs twice.

To solve this problem, instead of using indexes, we rely on *global indexes* built on top of indexes. Global indexes can be used both at the DIOC level and at the DPOC level and their distinctness is preserved by transitions.

Definition 7.7 (Global index). Given an annotated DIOC process \mathcal{I} , or an annotated DPOC network \mathcal{N} , for each annotated construct with index ι we define its global index ξ as follows:

- if the construct is not in the body of a while loop then $\xi = \iota$;
- if the innermost while construct that contains the considered construct has global index ξ' then the considered construct has global index $\xi = \xi' : \iota$.

Example 7.8. Consider the DIOC \mathcal{I} in Example 7.6. The first assignment with index \mathbf{j} also has global index \mathbf{j} , while the second assignment with index \mathbf{j} has global index $\mathbf{i}: \mathbf{j}$, since this last assignment is inside a while loop with global index \mathbf{i} .

Lemma 7.9 (Distinctness of Global Indexes). *Given a well-annotated DIOC \mathcal{I} , a global state Σ , and a set of updates \mathbf{I} , if $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle \xrightarrow{\eta_1} \dots \xrightarrow{\eta_n} \langle \Sigma', \mathbf{I}', \mathcal{I}' \rangle$ then all global indexes in \mathcal{I}' are distinct.*

Proof. The proof is by induction on the number n of transitions. Details are in Appendix B. \square

Using global indexes we can now define event structures corresponding to the execution of DIOCs and DPOCs. We start by defining DIOC events. Some events correspond to transitions of the DIOC, and we say that they are enabled when the corresponding transition is enabled, executed when the corresponding transition is executed.

Definition 7.10 (DIOC events). We use ε to range over events, and we write $[\varepsilon]_{\mathbf{R}}$ to highlight that event ε is performed by role \mathbf{R} . An annotated DIOC \mathcal{I} contains the following events:

Communication events: a sending event $\xi : \bar{o}@R_2$ in role R_1 and a receiving event $\xi : o@R_1$ in role R_2 for each interaction $i : o : R_1(e) \rightarrow R_2(x)$ with global index ξ ; we also denote the sending event as f_ξ or $[f_\xi]_{R_1}$ and the receiving event as t_ξ or $[t_\xi]_{R_2}$. Sending and receiving events correspond to the transition executing the interaction.

Assignment events: an assignment event ε_ξ in role R for each assignment $i : x@R = e$ with global index ξ ; the event corresponds to the transition executing the assignment.

Scope events: a scope initialisation event \uparrow_ξ and a scope termination event \downarrow_ξ for each scope $i : \text{scope } @R \{T\}$ with global index ξ . Both these events belong to all the roles in $\text{roles}(T)$. The scope initialisation event corresponds to the transition performing or not performing an update on the given scope. The scope termination event is just an auxiliary event (related to the auxiliary interactions implementing the scope termination).

If events: a guard if-event ε_ξ in role R for each construct $i : \text{if } b@R \{T\} \text{ else } \{T'\}$ with global index ξ ; the guard-if event corresponds to the transition evaluating the guard of the condition.

While events: a guard while-event ε_ξ in role R for each construct $i : \text{while } b@R \{T\}$ with global index ξ ; the guard-while event corresponds to the transition evaluating the guard of the while loop.

Function $\text{events}(T)$ denotes the set of events of the annotated DIOC T . A sending and a receiving event with the same global index ξ are called matching events. We denote with $\bar{\varepsilon}$ an event matching event ε . A communication event is either a sending event or a receiving event. A communication event is unmatched if there is no event matching it.

As a corollary of Lemma 7.9 events have distinct names. Note also that, for each while loop, there are events corresponding to the execution of just one iteration of the loop. If unfolding is performed, new events are created.

Similarly to what we have done for DIOC, we can define events for DPOC as follows.

Definition 7.11 (DPOC events). An annotated DPOC network \mathcal{N} contains the following events:

Communication events: a sending event $\xi : \bar{o}^?@R_2$ in role R_1 for each send $\iota : i.o^? : e \text{ to } R_2$ with global index ξ in role R_1 ; and a receiving event $\xi : o^?@R_1$ in role R_2 for each receive $\iota : i.o^? : x \text{ from } R_1$ with global index ξ in role R_2 ; we also denote the sending event as f_ξ or $[f_\xi]_{R_1}$; and the receiving event as t_ξ or $[t_\xi]_{R_2}$. Sending and receiving events correspond to the transitions executing the corresponding communication.

Assignment events: an assignment event ε_ξ in role R for each assignment $i : x = e$ with global index ξ ; the event corresponds to the transition executing the assignment.

Scope events: a scope initialisation event \uparrow_ξ and a scope termination event \downarrow_ξ for each $i : \text{scope } @R \{P\} \text{ roles } \{S\}$ or $i : \text{scope } @R \{P\}$ with global index ξ . Scope events with the same global index coincide, and thus the same event may belong to different roles; the scope initialisation event corresponds to the transition performing or not performing an update on the given scope for the role leading the update. The scope termination event is just an auxiliary event (related to the auxiliary interactions implementing the scope termination).

If events: a guard if-event ε_ξ in role R for each construct $i : \text{if } b \{P\} \text{ else } \{P'\}$ with global index ξ ; the guard-if event corresponds to the transition evaluating the guard of the condition.

While events: a guard while-event ε_ξ in role \mathbf{R} for each construct $\mathbf{i} : \mathbf{while} \ b \ \{P\}$ with global index ξ ; the guard-while event corresponds to the transition evaluating the guard of the while loop.

Let $\mathbf{events}(\mathcal{N})$ denote the set of events of the network \mathcal{N} . A sending and a receiving event with either the same global index ξ or with global indexes differing only for replacing index $\mathbf{i}_?$ with \mathbf{i}_T or \mathbf{i}_F are called matching events. We denote with $\bar{\varepsilon}$ an event matching event ε .

With a slight abuse of notation, we write $\mathbf{events}(P)$ to denote events originated by constructs in process P , assuming the network \mathcal{N} to be understood. We use the same syntax for events of DIOCs and of DPOCs. Indeed, the two kinds of events are strongly related (cf. Lemma 7.15).

We define below a causality relation \leq_{DIOC} among DIOC events based on the constraints given by the semantics on the execution of the corresponding transitions.

Definition 7.12 (DIOC causality relation). Let us consider an annotated DIOC \mathcal{I} . A causality relation $\leq_{DIOC} \subseteq \mathbf{events}(\mathcal{I}) \times \mathbf{events}(\mathcal{I})$ is the minimum reflexive and transitive relation satisfying:

Sequentiality: let $\mathcal{I}'; \mathcal{I}''$ be a subterm of DIOC \mathcal{I} . If ε' is an event in \mathcal{I}' and ε'' is an event in \mathcal{I}'' , then $\varepsilon' \leq_{DIOC} \varepsilon''$.

Scope: let $\mathbf{i} : \mathbf{scope} \ @\mathbf{R} \ \{\mathcal{I}'\}$ be a subterm of DIOC \mathcal{I} . If ε' is an event in \mathcal{I}' then $\uparrow_\xi \leq_{DIOC} \varepsilon' \leq_{DIOC} \downarrow_\xi$.

Synchronisation: for each interaction the sending event precedes the receiving event.

If: let $\mathbf{i} : \mathbf{if} \ b@\mathbf{R} \ \{\mathcal{I}'\} \ \mathbf{else} \ \{\mathcal{I}''\}$ be a subterm of DIOC \mathcal{I} , let ε_ξ be the guard if-event in role \mathbf{R} , then for every event ε' in \mathcal{I}' and for every event ε'' in \mathcal{I}'' we have $\varepsilon_\xi \leq_{DIOC} \varepsilon'$ and $\varepsilon_\xi \leq_{DIOC} \varepsilon''$.

While: let $\mathbf{i} : \mathbf{while} \ b@\mathbf{R} \ \{\mathcal{I}'\}$ be a subterm of DIOC \mathcal{I} , let ε_ξ be the guard while-event in role \mathbf{R} , then for every event ε' in \mathcal{I}' we have $\varepsilon_\xi \leq_{DIOC} \varepsilon'$.

As expected, the relation \leq_{DIOC} is a partial order.

Lemma 7.13. *Let us consider an annotated DIOC \mathcal{I} . The relation \leq_{DIOC} among events of \mathcal{I} is a partial order.*

Proof. A partial order is a relation which is reflexive, transitive, and antisymmetric. Reflexivity and transitivity follow by definition. We show antisymmetry by showing that \leq_{DIOC} does not contain any cycle. The proof is by structural induction on the DIOC \mathcal{I} . The base cases are interaction, assignment, $\mathbf{1}$, and $\mathbf{0}$, which are all trivial. In the case of sequence, $\mathcal{I}'; \mathcal{I}''$, by inductive hypothesis there are no cycles among the events of \mathcal{I}' nor among events of \mathcal{I}'' . Since all events of \mathcal{I}' precede all events of \mathcal{I}'' , there are no cycles among the events of $\mathcal{I}'; \mathcal{I}''$. In the case of parallel composition, $\mathcal{I}' | \mathcal{I}''$, there is no relation between events in \mathcal{I}' and in \mathcal{I}'' , hence the thesis follows. In the case of conditional $\mathbf{i} : \mathbf{if} \ b@\mathbf{R} \ \{\mathcal{I}'\} \ \mathbf{else} \ \{\mathcal{I}''\}$ there are no relations between events in \mathcal{I}' and in \mathcal{I}'' and all the events follow the guard-if event. Hence the thesis follows. The cases of while and scope are similar. \square

We can now define a causality relation \leq_{DPOC} among DPOC events.

Definition 7.14 (DPOC causality relation). Let us consider an annotated DPOC network \mathcal{N} . A causality relation $\leq_{DPOC} \subseteq \mathbf{events}(\mathcal{N}) \times \mathbf{events}(\mathcal{N})$ is the minimum reflexive and transitive relation satisfying:

Sequentiality: Let $P'; P''$ be a subterm of DPOC network \mathcal{N} . If ε' is an event in P' and ε'' is an event in P'' then $\varepsilon' \leq_{DPOC} \varepsilon''$.

Scope: Let $\mathbf{i: scope @R \{P\} roles \{S\}}$ or $\mathbf{i: scope @R \{P\}}$ be a subterm of DPOC \mathcal{N} with global index ξ . If ε' is an event in P then $\uparrow_{\xi} \leq_{DPOC} \varepsilon' \leq_{DPOC} \downarrow_{\xi}$.

Synchronisation: For each pair of events ε and ε' , $\varepsilon \leq \varepsilon'$ implies $\bar{\varepsilon} \leq_{DPOC} \varepsilon'$.

If: Let $\mathbf{i: if b \{P\} else \{P'\}}$ be a subterm of DPOC network \mathcal{N} with global index ξ , let ε_{ξ} be the guard if-event, then for every event ε in P and for every event ε' in P' we have $\varepsilon_{\xi} \leq_{DPOC} \varepsilon$ and $\varepsilon_{\xi} \leq_{DPOC} \varepsilon'$.

While: Let $\mathbf{i: while b \{P\}}$ be a subterm of DPOC network \mathcal{N} with global index ξ , let ε_{ξ} be the guard while-event, then for every event ε in P we have $\varepsilon_{\xi} \leq_{DPOC} \varepsilon$.

On *DPOC* networks obtained as projections of well-annotated *DIOC*s the relation \leq_{DPOC} is a partial order, as expected. However, since this result is not needed in the remainder of the paper, we do not present its proof.

There is a relation between *DIOC* events and causality relation and their counterparts at the *DPOC* level. Indeed, the events and causality relation are preserved by projection.

Lemma 7.15. *Given a well-annotated connected DIOC process \mathcal{I} and for each state Σ the DPOC network $\text{proj}(\mathcal{I}, \Sigma)$ is such that:*

- (1) $\text{events}(\mathcal{I}) \subseteq \text{events}(\text{proj}(\mathcal{I}, \Sigma))$;
- (2) $\forall \varepsilon_1, \varepsilon_2 \in \text{events}(\mathcal{I}). \varepsilon_1 \leq_{DIOC} \varepsilon_2 \Rightarrow \varepsilon_1 \leq_{DPOC} \varepsilon_2 \vee \varepsilon_1 \leq_{DPOC} \bar{\varepsilon}_2$

Proof. The events of a DPOC obtained by projecting a DIOC \mathcal{I} are included in the events of the DIOC \mathcal{I} by definition of projection. The preservation of the causality relation can be proven by a case analysis on the condition used to derive the dependency (i.e., sequentiality, scope, synchronisation, if and while). Details are in Appendix B. \square

To complete the definition of our event structure we now define a notion of conflict between (DIOC and DPOC) events, relating events which are in different branches of the same conditional.

Definition 7.16 (Conflicting events). Given a DIOC process \mathcal{I} , two events $\varepsilon, \varepsilon' \in \text{events}(\mathcal{I})$ are conflicting if they belong to different branches of the same conditional, i.e., there exists a subprocess $\mathbf{i: if b @R \{I'\} else \{I''\}}$ of \mathcal{I} such that $\varepsilon \in \text{events}(I') \wedge \varepsilon' \in \text{events}(I'')$ or $\varepsilon' \in \text{events}(I') \wedge \varepsilon \in \text{events}(I'')$.

Similarly, given a DPOC network \mathcal{N} , we say that two events $\varepsilon, \varepsilon' \in \text{events}(\mathcal{N})$ are conflicting if they belong to different branches of the same conditional, i.e., there exists a subprocess $\mathbf{i: if b \{P\} else \{P'\}}$ of \mathcal{N} such that $\varepsilon \in \text{events}(P) \wedge \varepsilon' \in \text{events}(P')$ or $\varepsilon' \in \text{events}(P) \wedge \varepsilon \in \text{events}(P')$.

Similarly to what we did for DIOCs, we define below well-annotated DPOCs. Well-annotated DPOCs include all DPOCs obtained by projecting well-annotated DIOCs. As stated in the definition below, and proved in Lemma 7.20, well-annotated DPOCs enjoy various properties useful for our proofs.

Definition 7.17 (Well-annotated DPOC). An annotated DPOC network \mathcal{N} is well annotated for its causality relation \leq_{DPOC} if the following conditions hold:

- C1: for each global index ξ there are at most two communication events on programmer-specified operations with global index ξ and, in this case, they are matching events;
- C2: only events which are minimal according to \leq_{DPOC} may correspond to enabled transitions;

- C3: for each pair of non-conflicting sending events $[f_\xi]_{\mathbf{R}_1}$ and $[f_{\xi'}]_{\mathbf{R}_1}$ on the same operation $\mathbf{i.o}^?$ with the same target \mathbf{R}_2 such that $\xi \neq \xi'$ we have $[f_\xi]_{\mathbf{R}_1} \leq_{DPOC} [f_{\xi'}]_{\mathbf{R}_1}$ or $[f_{\xi'}]_{\mathbf{R}_1} \leq_{DPOC} [f_\xi]_{\mathbf{R}_1}$;
- C4: for each pair of non-conflicting receiving events $[t_\xi]_{\mathbf{R}_2}$ and $[t_{\xi'}]_{\mathbf{R}_2}$ on the same operation $\mathbf{i.o}^?$ with the same sender \mathbf{R}_1 such that $\xi \neq \xi'$ we have $[t_\xi]_{\mathbf{R}_2} \leq [t_{\xi'}]_{\mathbf{R}_2}$ or $[t_{\xi'}]_{\mathbf{R}_2} \leq [t_\xi]_{\mathbf{R}_2}$;
- C5: if ε is an event inside a scope with global index ξ then its matching events $\bar{\varepsilon}$ (if they exist) are inside a scope with the same global index.
- C6: if two events have the same index but different global indexes then one of them, let us call it ε_1 , is inside the body of a while loop with global index ξ_1 and the other, ε_2 , is not. Furthermore, $\varepsilon_2 \leq_{DPOC} \varepsilon_{\xi_1}$ where ε_{ξ_1} is the guarding while-event of the while loop with global index ξ_1 .

Since scope update, conditional, and iteration at the DIOC level happen in one step, while they correspond to many steps of the projected DPOC, we introduce a function, denoted **upd**, that bridges this gap. More precisely, function **upd** is obtained as the composition of two functions, a function **compl** that completes the execution of DIOC actions which have already started, and a function **clean** that eliminates all the auxiliary closing communications of scopes (scope execution introduces in the DPOC auxiliary communications which have no correspondence in the DIOC).

Definition 7.18 (**upd function**). Let \mathcal{N} be an annotated DPOC (we drop indexes if not relevant). The **upd** function is defined as the composition of a function **compl** and a function **clean**. Thus, $\mathbf{upd}(\mathcal{N}) = \mathbf{clean}(\mathbf{compl}(\mathcal{N}))$. Network $\mathbf{compl}(\mathcal{N})$ is obtained from \mathcal{N} by repeating the following operations while possible.

- (1) Performing the reception of the positive evaluation of the guard of a while loop, by replacing for every $\mathbf{i.wb}_i^* : \mathbf{true to R}'$ enabled, all the terms

$$\mathbf{i.wb}_i^* : x_i \mathbf{from R}; \mathbf{while } x_i \{P; \mathbf{i.we}_i^* : \mathbf{ok to R}; \mathbf{i.wb}_i^* : x_i \mathbf{from R}\}$$

not inside another while construct, with

$$P; \mathbf{i.we}_i^* : \mathbf{ok to R}; \mathbf{i.wb}_i^* : x_i \mathbf{from R}; \mathbf{while } x_i \{P; \mathbf{i.we}_i^* : \mathbf{ok to R}; \mathbf{i.wb}_i^* : x_i \mathbf{from R}\}$$

and replace $\mathbf{i.wb}_i^* : \mathbf{true to R}'$ with **1**.

- (2) Performing the reception of the negative evaluation of the guard of a while loop by replacing, for every $\mathbf{i.wb}_i^* : \mathbf{false to R}'$ enabled, all the terms

$$\mathbf{i.wb}_i^* : x_i \mathbf{from R}; \mathbf{while } x_i \{P; \mathbf{i.we}_i^* : \mathbf{ok to R}; \mathbf{i.wb}_i^* : x_i \mathbf{from R}\}$$

not inside another while construct, with **1**, and replace $\mathbf{i.wb}_i^* : \mathbf{false to R}'$ with **1**.

- (3) Performing the unfolding of a while loop by replacing every

$$\mathbf{while } x_i \{P; \mathbf{i.we}_i^* : \mathbf{ok to R}; \mathbf{i.wb}_i^* : x_i \mathbf{from R}\}$$

enabled not inside another while construct, such that x_i evaluates to **true** in the local state, with

$$P; \mathbf{i.we}_i^* : \mathbf{ok to R}; \mathbf{i.wb}_i^* : x_i \mathbf{from R}; \mathbf{while } x_i \{P; \mathbf{i.we}_i^* : \mathbf{ok to R}; \mathbf{i.wb}_i^* : x_i \mathbf{from R}\}$$

- (4) Performing the termination of while loop by replacing every

$$\mathbf{while } x_i \{P; \mathbf{i.we}_i^* : \mathbf{ok to R}; \mathbf{i.wb}_i^* : x_i \mathbf{from R}\}$$

enabled not inside another while construct, such that x_i evaluates to **false** in the local state, with **1**.

- (5) Performing the reception of the positive evaluation of the guard of a conditional by replacing, for every $\mathbf{i.cnd}_i^* : \mathbf{true to R}'$ enabled, all the terms

$$\mathbf{i.cnd}_i^* : x_i \mathbf{from R}; \mathbf{if } x_i \{P'\} \mathbf{else } \{P''\}$$

not inside a while construct with P' , and replace $\mathbf{i.cnd}_i^* : \mathbf{true to R}'$ with $\mathbf{1}$.

- (6) Performing the reception of the negative evaluation of the guard of a conditional by replacing, for every $\mathbf{i.cnd}_i^* : \mathbf{false to R}'$ enabled, all the terms

$$\mathbf{i.cnd}_i^* : x_i \mathbf{from R}; \mathbf{if } x_i \{P'\} \mathbf{else } \{P''\}$$

not inside a while construct, with P'' , and replace $\mathbf{i.cnd}_i^* : \mathbf{false to R}'$ with $\mathbf{1}$.

- (7) Performing the selection of the “then” branch by replacing every

$$\mathbf{if } x_i \{P'\} \mathbf{else } \{P''\}$$

enabled, such that x_i evaluates to \mathbf{true} in the local state, with P' .

- (8) Performing the selection of the “else” branch by replacing every

$$\mathbf{if } x_i \{P'\} \mathbf{else } \{P''\}$$

enabled, such that x_i evaluates to \mathbf{false} in the local state, with P'' .

- (9) Performing the communication of the updated code by replacing, for every $\mathbf{i.sb}_i^* : P \mathbf{to S}$ enabled, all the terms

$$\mathbf{i: scope @R } \{P'\}$$

in role \mathbf{S} not inside a while construct with P , and replace $\mathbf{i.sb}_i^* : P \mathbf{to S}$ with $\mathbf{1}$.

- (10) Performing the communication that no update is needed by replacing, for each $\mathbf{i.sb}_i^* : \mathbf{no to S}$ enabled, all the terms

$$\mathbf{i: scope @R } \{P'\}$$

in role \mathbf{S} not inside a while construct with P' , and replace $\mathbf{i.sb}_i^* : P \mathbf{to S}$ with $\mathbf{1}$.

Network $\mathbf{clean}(\mathcal{N})$ is obtained from \mathcal{N} by repeating the following operations while possible:

- Removing the auxiliary communications for end of scope and end of while loop synchronisation by replacing each

$$\begin{array}{ll} \mathbf{i.se}_i^* : \mathbf{ok to R} & \mathbf{i_c: i.we}_i^* : \mathbf{ok to R} \\ \mathbf{i.se}_i^* : _ \mathbf{from R} & \mathbf{i_c: i.we}_i^* : _ \mathbf{from R} \end{array}$$

not inside a while construct with $\mathbf{1}$.

Furthermore \mathbf{clean} may apply 0 or more times the following operation:

- replace a subterm $\mathbf{1}; P$ by P or a subterm $\mathbf{1} \mid P$ by P .

Note that function \mathbf{compl} does not reduce terms inside a while construct. Assume, for instance, to have an auxiliary send targeting a receive inside the body of a while loop. These two communications should not interact since they have different global indexes. This explains why we exclude terms inside the body of while loops.

We proceed now to prove some of the proprieties of DIOC and DPOC. The first result states that in a well-annotated DPOC only transitions corresponding to events minimal with respect to the causality relation \leq_{DPOC} may be enabled.

Lemma 7.19. *If \mathcal{N} is a DPOC, \leq_{DPOC} its causality relation and ε is an event corresponding to a transition enabled in \mathcal{N} then ε is minimal with respect to \leq_{DPOC} .*

Proof. The proof is by contradiction. Suppose ε is enabled but not minimal, i.e., there exists ε' such that $\varepsilon' \leq_{DPOC} \varepsilon$. If there is more than one such ε' consider the one such that the length of the derivation of $\varepsilon' \leq_{DPOC} \varepsilon$ is minimal. This derivation should have length one, and following Definition 7.14 it may result from one of the following cases:

- **Sequentiality:** $\varepsilon' \leq_{DPOC} \varepsilon$ means that $\varepsilon' \in \text{events}(P')$, $\varepsilon \in \text{events}(P'')$, and $P'; P''$ is a subterm of \mathcal{N} . Because of the semantics of sequential composition ε cannot be enabled.
- **Scope:** let $\mathbf{i} : \text{scope } @\mathbf{R} \{P\} \text{ roles } \{S\}$ or $\mathbf{i} : \text{scope } @\mathbf{R} \{P\}$ be a subprocess of \mathcal{N} with global index ξ . We have the following cases:
 - $\varepsilon' = \uparrow_{\xi}$ and $\varepsilon \in \text{events}(P)$, and this implies that ε cannot be enabled since if ε' is enabled then the Rule $[\text{DPOC}]_{\text{UP}}$ or Rule $[\text{DPOC}]_{\text{NoUP}}$ for starting the execution of the scope have not been applied yet;
 - $\varepsilon' = \uparrow_{\xi}$ and $\varepsilon = \downarrow_{\xi}$, or $\varepsilon' \in \text{events}(P)$ and $\varepsilon = \downarrow_{\xi}$: this is trivial, since \downarrow_{ξ} is an auxiliary event and no transition corresponds to it;
- **If:** $\varepsilon' \leq_{DPOC} \varepsilon$ means that ε' is the evaluation of the guard of a subterm $\mathbf{i} : \text{if } x_i \{P'\} \text{ else } \{P''\}$ and $\varepsilon \in \text{events}(P') \cup \text{events}(P'')$. Event ε cannot be enabled because of the semantics of conditionals.
- **While:** $\varepsilon' \leq_{DPOC} \varepsilon$ means that ε' is the evaluation of the guard of a subterm $\mathbf{i} : \text{while } x_i \{P\}$ and $\varepsilon \in \text{events}(P)$. Event ε cannot be enabled because of the semantics of the while loop. \square

We now prove that all the DPOCs obtained as projection of well-annotated connected DIOCs are well-annotated.

Lemma 7.20. *Let \mathcal{I} be a well-annotated connected DIOC process and Σ a state. Then the projection $\mathcal{N} = \text{proj}(\mathcal{I}, \Sigma)$ is a well-annotated DPOC network with respect to \leq_{DPOC} .*

Proof. We have to prove that $\text{proj}(\mathcal{I}, \Sigma)$ satisfies the conditions of Definition 7.17. We have a case for each condition. Details in Appendix B. \square

The next lemma shows that for every starting set of updates \mathbf{I} the DPOC \mathcal{N} and $\text{upd}(\mathcal{N})$ have the same set of weak traces.

Lemma 7.21. *Let \mathcal{N} be a DPOC. The following properties hold:*

- (1) *if $\langle \mathbf{I}, \text{upd}(\mathcal{N}) \rangle \xrightarrow{\eta} \langle \mathbf{I}', \mathcal{N}' \rangle$ with $\eta \in \{o^? : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x), o^* : \mathbf{R}_1(X) \rightarrow \mathbf{R}_2(), \mathbf{I}', \surd, \mathcal{I}, \text{no-up}, \tau\}$ then there exists \mathcal{N}'' such that $\langle \mathbf{I}, \mathcal{N} \rangle \xrightarrow{\eta_1} \dots \xrightarrow{\eta_k} \xrightarrow{\eta} \langle \mathbf{I}', \mathcal{N}'' \rangle$ where $\eta_i \in \{o^* : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x), o^* : \mathbf{R}_1(X) \rightarrow \mathbf{R}_2(), \tau\}$ for each $i \in \{1, \dots, k\}$ and $\text{upd}(\mathcal{N}'') = \text{upd}(\mathcal{N}')$.*
- (2) *if $\langle \mathbf{I}, \mathcal{N} \rangle \xrightarrow{\eta} \langle \mathbf{I}', \mathcal{N}' \rangle$ for $\eta \in \{o^? : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x), o^* : \mathbf{R}_1(X) \rightarrow \mathbf{R}_2(), \mathbf{I}', \surd, \mathcal{I}, \text{no-up}, \tau\}$, then one of the following holds:*
 - (a) $\text{upd}(\mathcal{N}) = \text{upd}(\mathcal{N}')$ and $\eta \in \{o^* : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x), o^* : \mathbf{R}_1(X) \rightarrow \mathbf{R}_2(), \tau\}$, or
 - (b) $\langle \mathbf{I}, \text{upd}(\mathcal{N}) \rangle \xrightarrow{\eta} \langle \mathbf{I}', \mathcal{N}'' \rangle$ such that $\text{upd}(\mathcal{N}') = \text{upd}(\mathcal{N}'')$.

Proof.

- (1) Applying the **upd** function corresponds to perform weak transitions, namely transitions with labels in $\{o^* : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x), o^* : \mathbf{R}_1(X) \rightarrow \mathbf{R}_2(), \tau\}$. Some of such transitions may not be enabled yet. Hence, \mathcal{N} may perform the subset of the weak transitions above which are or become enabled, reducing to some \mathcal{N}''' . Then, η is enabled also in \mathcal{N}''' and we have $\langle \mathbf{I}, \mathcal{N}''' \rangle \xrightarrow{\eta} \langle \mathbf{I}', \mathcal{N}''' \rangle$. At this point we have that \mathcal{N}''' and \mathcal{N}' may differ only for the weak transitions that were never enabled, which can be executed by **upd**.
- (2) There are two cases. In the first case the transition with label η is one of the transitions executed by function **upd**. In this case the condition 2a holds. In the second case, the transition with label η is not one of the transitions executed by function **upd**. In this case the transition with label η is still enabled in $\mathbf{upd}(\mathcal{N})$ and can be executed. This leads to a network that differs from \mathcal{N}' only because of transitions executed by the **upd** function and case 2b holds. \square

We now prove a property of transitions with label \surd .

Lemma 7.22. *For each DIOC system $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle$ that reduces with a transition labelled \surd then, for each role $\mathbf{R} \in \text{roles}(\mathcal{I})$, the DPOC role $(\pi(\mathcal{I}, \mathbf{R}), \Sigma_{\mathbf{R}})_{\mathbf{R}}$ can reduce with a transition labelled \surd and vice versa.*

Proof. Note that a DIOC can perform a transition with label \surd only if it is a term obtained using sequential and/or parallel composition starting from $\mathbf{1}$ constructs. The projection has the same shape, hence it can perform the desired transition. The other direction is similar. \square

We can now prove our main theorem (Theorem 7.2, restated below) for which, given a connected well-annotated DIOC process \mathcal{I} and a state Σ , the DPOC network obtained as its projection has the same behaviours of \mathcal{I} .

Theorem 7.2 (Correctness). *For each initial, connected DIOC process \mathcal{I} , each state Σ , each set of updates \mathbf{I} , the DIOC system $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle$ and the DPOC system $\langle \mathbf{I}, \text{proj}(\mathcal{I}, \Sigma) \rangle$ are weak system bisimilar.*

Proof. We prove that the relation \mathcal{R} below is a weak system bisimulation.

$$\mathcal{R} = \left\{ \left(\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle, \langle \mathbf{I}, \mathcal{N} \rangle \right) \left| \begin{array}{l} \mathbf{upd}(\mathcal{N}) = \text{proj}(\mathcal{I}, \Sigma), \\ \text{events}(\mathcal{I}) \subseteq \text{events}(\text{compl}(\mathcal{N})), \\ \forall \varepsilon_1, \varepsilon_2 \in \text{events}(\mathcal{I}) . \\ \varepsilon_1 \leq_{DIOC} \varepsilon_2 \Rightarrow \varepsilon_1 \leq_{DPOC} \varepsilon_2 \vee \varepsilon_1 \leq_{DPOC} \bar{\varepsilon}_2 \end{array} \right. \right\}$$

where \mathcal{I} is obtained from a well-annotated connected DIOC via 0 or more transitions and $\mathbf{upd}(\mathcal{N})$ is a well-annotated DPOC.

To ensure that proving that the relation above is a weak system bisimulation implies our thesis, let us show that the pair $(\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle, \langle \mathbf{I}, \text{proj}(\mathcal{I}, \Sigma) \rangle)$ from the theorem statement belongs to \mathcal{R} . Note that here \mathcal{I} is well-annotated and connected, and for each such \mathcal{I} we have $\mathbf{upd}(\text{proj}(\mathcal{I}, \Sigma)) = \text{proj}(\mathcal{I}, \Sigma)$. From Lemma 7.20 $\text{proj}(\mathcal{I}, \Sigma)$ is well-annotated, thus $\mathbf{upd}(\text{proj}(\mathcal{I}, \Sigma))$ is well-annotated. Observe that compl is the identity on $\text{proj}(\mathcal{I}, \Sigma)$, thus from Lemma 7.15 we have that the conditions $\text{events}(\mathcal{I}) \subseteq \text{events}(\text{compl}(\mathcal{N}))$ and $\forall \varepsilon_1, \varepsilon_2 \in \text{events}(\mathcal{I}) . \varepsilon_1 \leq_{DIOC} \varepsilon_2 \Rightarrow \varepsilon_1 \leq_{DPOC} \varepsilon_2 \vee \varepsilon_1 \leq_{DPOC} \bar{\varepsilon}_2$ are satisfied.

We now prove that \mathcal{R} is a weak system bisimulation. To prove it, we show below that it is enough to consider only the case in which \mathcal{N} (and not $\text{upd}(\mathcal{N})$) is equal to $\text{proj}(\mathcal{I}, \Sigma)$. Furthermore, in this case the transition of $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle$ is matched by the first transition of $\langle \mathbf{I}, \text{proj}(\mathcal{I}, \Sigma) \rangle$.

Formally, for each $(\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle, \langle \mathbf{I}, \mathcal{N} \rangle)$ where $\mathcal{N} = \text{proj}(\mathcal{I}, \Sigma)$ we have to prove the following simplified bisimulation clauses.

- if $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle \xrightarrow{\mu} \langle \Sigma'', \mathbf{I}'', \mathcal{I}'' \rangle$ then $\langle \mathbf{I}, \mathcal{N} \rangle \xrightarrow{\mu} \langle \mathbf{I}'', \mathcal{N}'''' \rangle$ with $(\langle \Sigma'', \mathbf{I}'', \mathcal{I}'' \rangle, \langle \mathbf{I}'', \mathcal{N}'''' \rangle) \in \mathcal{R}$;
- if $\langle \mathbf{I}, \mathcal{N} \rangle \xrightarrow{\eta} \langle \mathbf{I}'', \mathcal{N}'''' \rangle$ with $\eta \in \{o : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x); \surd; \mathcal{I}; \mathbf{no-up}; \mathbf{I}''; \tau\}$ then $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle \xrightarrow{\eta} \langle \Sigma'', \mathbf{I}'', \mathcal{I}'' \rangle$ and $(\langle \Sigma'', \mathbf{I}'', \mathcal{I}'' \rangle, \langle \mathbf{I}'', \mathcal{N}'''' \rangle) \in \mathcal{R}$.

In fact, consider a general network \mathcal{N}_g with $\text{upd}(\mathcal{N}_g) = \text{proj}(\mathcal{I}, \Sigma)$. If $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle \xrightarrow{\mu} \langle \Sigma'', \mathbf{I}'', \mathcal{I}'' \rangle$, then by hypothesis $\langle \mathbf{I}, \text{upd}(\mathcal{N}_g) \rangle \xrightarrow{\mu} \langle \mathbf{I}'', \mathcal{N}'''' \rangle$. From Lemma 7.21 case 1 there exists \mathcal{N}'' such that $\langle \mathbf{I}, \mathcal{N}_g \rangle \xrightarrow{\eta_1} \dots \xrightarrow{\eta_k} \xrightarrow{\mu} \langle \mathbf{I}'', \mathcal{N}'''' \rangle$ where $\eta_i \in \{o^* : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x), o^* : \mathbf{R}_1(X) \rightarrow \mathbf{R}_2(), \tau\}$ for each $i \in \{1, \dots, k\}$ and $\text{upd}(\mathcal{N}''') = \text{upd}(\mathcal{N}'''')$. By hypothesis $(\langle \Sigma'', \mathbf{I}'', \mathcal{I}'' \rangle, \langle \mathbf{I}'', \mathcal{N}'''' \rangle) \in \mathcal{R}$, hence, by definition of \mathcal{R} , $\text{upd}(\mathcal{N}''') = \text{proj}(\mathcal{I}'', \Sigma'')$, and therefore also $\text{upd}(\mathcal{N}''') = \text{proj}(\mathcal{I}'', \Sigma'')$. The conditions on events hold by hypothesis since function upd has no effect on DPOC events corresponding to DIOC events. Furthermore, only enabled interactions have been executed, hence dependencies between DPOC events corresponding to DIOC events are untouched.

If instead $\langle \mathbf{I}, \mathcal{N}_g \rangle \xrightarrow{\eta} \langle \mathbf{I}'', \mathcal{N}'''' \rangle$ with $\eta \in \{o^? : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x), o^* : \mathbf{R}_1(X) \rightarrow \mathbf{R}_2(), \surd, \mathcal{I}, \mathbf{no-up}, \mathbf{I}'', \tau\}$ then thanks to Lemma 7.21 we have one of the following: (2a) $\text{upd}(\mathcal{N}_g) = \text{upd}(\mathcal{N}'''')$ and $\eta \in \{o^* : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x), o^* : \mathbf{R}_1(X) \rightarrow \mathbf{R}_2(), \tau\}$, or (2b) $\langle \mathbf{I}, \text{upd}(\mathcal{N}_g) \rangle \xrightarrow{\eta} \langle \mathbf{I}'', \mathcal{N}'''' \rangle$ such that $\text{upd}(\mathcal{N}''') = \text{upd}(\mathcal{N}'''')$. In case (2b) we have $\langle \mathbf{I}, \text{upd}(\mathcal{N}_g) \rangle \xrightarrow{\eta} \langle \mathbf{I}'', \mathcal{N}'''' \rangle$. Then, by hypothesis, we have $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle \xrightarrow{\eta} \langle \Sigma'', \mathbf{I}'', \mathcal{I}'' \rangle$ and $(\langle \Sigma'', \mathbf{I}'', \mathcal{I}'' \rangle, \langle \mathbf{I}'', \mathcal{N}'''' \rangle) \in \mathcal{R}$. To deduce that $(\langle \Sigma'', \mathbf{I}'', \mathcal{I}'' \rangle, \langle \mathbf{I}'', \mathcal{N}'''' \rangle) \in \mathcal{R}$, one can proceed using the same strategy as the case of the challenge from the DIOC above. In case (2a) the step is matched by the DIOC by staying idle, following the second option in the definition of weak system bisimulation. The proof is similar to the one above.

Thus, we have to prove the two simplified bisimulation clauses above. The proof is by structural induction on the DIOC \mathcal{I} . All the subterms of a well-annotated connected DIOC are well-annotated and connected, thus the induction can be performed. We consider both challenges from the DIOC (\rightarrow) and from the DPOC (\leftarrow). The case for label \surd follows from Lemma 7.22. The case for labels \mathbf{I} is trivial. Let us consider the other labels, namely $o : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x), \mathcal{I}, \mathbf{no-up}$, and τ .

Note that no transition (at the DIOC or at the DPOC level) with one of these labels can change the set of updates \mathbf{I} . Thus, in the following, we will not write it. Essentially, we will use DIOC processes and DPOC networks instead of DIOC systems and DPOC systems, respectively. Note that DPOC networks also include the state, while this is not the case for DIOC processes. For DIOC processes, we assume to associate to them the state Σ , and comment on its changes whenever needed.

Case 1, 0: trivial.

Case i: $x@R = e$: the assignment changes the global state in the DIOC, and its projection on the role \mathbf{R} changes the local state of the role in the DPOC in a corresponding way.

Case i: $o : \mathbf{R}_1(e) \rightarrow \mathbf{R}_2(x)$: trivial. Just note that at the DPOC level the interaction gives rise to one send and one receive with the same operation and prefixed by the same index.

Synchronisation between send and receive is performed by Rule $[^{\text{DPOC}} |_{\text{SYNCH}}]$ that also removes the index from the label.

Case $\mathcal{I}; \mathcal{I}'$: from the definition of the projection function we have that $\mathcal{N} = \parallel_{\mathbf{R} \in \text{roles}(\mathcal{I}; \mathcal{I}')} (\pi(\mathcal{I}, \mathbf{R}); \pi(\mathcal{I}', \mathbf{R}), \Sigma_{\mathbf{R}})_{\mathbf{R}}$.

\rightarrow : Assume that $\mathcal{I}; \mathcal{I}' \xrightarrow{\mu} \mathcal{I}''$ with $\mu \in \{o : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x); \mathcal{I}; \text{no-up}, \tau\}$. There are two possibilities: either (i) $\mathcal{I} \xrightarrow{\mu} \mathcal{I}'''$ and $\mathcal{I}'' = \mathcal{I}'''$; \mathcal{I}' or (ii) \mathcal{I} has a transition with label \checkmark and $\mathcal{I}' \xrightarrow{\mu} \mathcal{I}''$.

In case (i) by inductive hypothesis

$$\parallel_{\mathbf{R} \in \text{roles}(\mathcal{I})} (\pi(\mathcal{I}, \mathbf{R}), \Sigma_{\mathbf{R}})_{\mathbf{R}} \xrightarrow{\mu} \mathcal{N}''' \text{ and } \text{upd}(\mathcal{N}''') = \parallel_{\mathbf{R} \in \text{roles}(\mathcal{I})} (\pi(\mathcal{I}''', \mathbf{R}), \Sigma'_{\mathbf{R}})_{\mathbf{R}}$$

Thus

$$\begin{aligned} \parallel_{\mathbf{R} \in \text{roles}(\mathcal{I})} (\pi(\mathcal{I}, \mathbf{R}); \pi(\mathcal{I}', \mathbf{R}), \Sigma_{\mathbf{R}})_{\mathbf{R}} &\xrightarrow{\mu} \mathcal{N} \text{ and} \\ \text{upd}(\mathcal{N}) &= \parallel_{\mathbf{R} \in \text{roles}(\mathcal{I})} (\pi(\mathcal{I}''', \mathbf{R}); \pi(\mathcal{I}', \mathbf{R}), \Sigma'_{\mathbf{R}})_{\mathbf{R}} \end{aligned}$$

If $\text{roles}(\mathcal{I}') \subseteq \text{roles}(\mathcal{I})$ then the thesis follows. If $\text{roles}(\mathcal{I}') \not\subseteq \text{roles}(\mathcal{I})$ then at the DPOC level the processes in the roles in $\text{roles}(\mathcal{I}') \setminus \text{roles}(\mathcal{I})$ are not affected by the transition. Note however that the projection of \mathcal{I} on these roles is a term composed only by **1s**, and the ones corresponding to parts of \mathcal{I} that have been consumed can be removed by the clean part of function upd .

In case (ii), \mathcal{I} has a transition with label \checkmark and $\mathcal{I}' \xrightarrow{\mu} \mathcal{I}''$. By inductive hypothesis $\text{proj}(\mathcal{I}', \Sigma) \xrightarrow{\mu} \mathcal{N}''$ and $\text{upd}(\mathcal{N}'') = \text{proj}(\mathcal{I}'', \Sigma')$. The thesis follows since, thanks to Lemma 7.22, $\text{proj}(\mathcal{I}; \mathcal{I}', \Sigma) \xrightarrow{\mu} \mathcal{N}$ and $\text{upd}(\mathcal{N}) = \text{proj}(\mathcal{I}'', \Sigma')$, possibly using the clean part of function upd to remove the **1s** which are no more needed.

Note that, in both the cases, conditions on events follow by inductive hypothesis.

\leftarrow : Assume that

$$\mathcal{N} = \parallel_{\mathbf{R} \in \text{roles}(\mathcal{I}; \mathcal{I}')} (\pi(\mathcal{I}, \mathbf{R}); \pi(\mathcal{I}', \mathbf{R}), \Sigma_{\mathbf{R}})_{\mathbf{R}} \xrightarrow{\eta} \parallel_{\mathbf{R} \in \text{roles}(\mathcal{I}; \mathcal{I}')} (P_{\mathbf{R}}, \Sigma'_{\mathbf{R}})_{\mathbf{R}}$$

with $\eta \in \{o : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x), \mathcal{I}, \text{no-up}, \tau\}$. We have a case analysis on η .

If $\eta = o : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x)$ then

$$\begin{aligned} (\pi(\mathcal{I}; \mathcal{I}', \mathbf{R}_1), \Sigma_{\mathbf{R}_1})_{\mathbf{R}_1} &\xrightarrow{\overline{\mathbf{i}.o(v)@_{\mathbf{R}_2}:\mathbf{R}_1}} (P_{\mathbf{R}_1}, \Sigma_{\mathbf{R}_1})_{\mathbf{R}_1} \text{ and} \\ (\pi(\mathcal{I}; \mathcal{I}', \mathbf{R}_2), \Sigma_{\mathbf{R}_2})_{\mathbf{R}_2} &\xrightarrow{\mathbf{i}.o(x \leftarrow v)@_{\mathbf{R}_1}:\mathbf{R}_2} (P_{\mathbf{R}_2}, \Sigma_{\mathbf{R}_2})_{\mathbf{R}_2} \end{aligned}$$

The two events have the same global index since they have the same index \mathbf{i} (otherwise they could not synchronise) and they are both outside of any while loop (since they are enabled), hence the global index coincides with the index. Thus, they are either both from \mathcal{I} or both from \mathcal{I}' .

In the first case we have also

$$\parallel_{\mathbf{R} \in \text{roles}(\mathcal{I}; \mathcal{I}')} (\pi(\mathcal{I}, \mathbf{R}), \Sigma_{\mathbf{R}})_{\mathbf{R}} \xrightarrow{o:\mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x)} \parallel_{\mathbf{R} \in \text{roles}(\mathcal{I}; \mathcal{I}')} (P''_{\mathbf{R}}, \Sigma_{\mathbf{R}})_{\mathbf{R}}$$

with $P_{\mathbf{R}} = P''_{\mathbf{R}}; \pi(\mathcal{I}', \mathbf{R})$. Thus, by inductive hypothesis, $\mathcal{I} \xrightarrow{o:\mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x)} \mathcal{I}''$ and $\text{upd}(\parallel_{\mathbf{R} \in \text{roles}(\mathcal{I}; \mathcal{I}')} (P''_{\mathbf{R}}, \Sigma_{\mathbf{R}})_{\mathbf{R}}) = \text{proj}(\mathcal{I}'', \Sigma)$. Hence, we have that

$$\mathcal{I}; \mathcal{I}' \xrightarrow{o:\mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x)} \mathcal{I}''; \mathcal{I}'$$

and

$$\text{upd}(\|_{\mathbf{R} \in \text{roles}_{\mathcal{I}; \mathcal{I}'}} (P''_{\mathbf{R}}; \pi(\mathcal{I}', \mathbf{R}), \Sigma_{\mathbf{R}})_{\mathbf{R}}) = \text{proj}(\mathcal{I}''; \mathcal{I}', \Sigma)$$

The thesis follows.

In the second case, we need to show that the interaction $o : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x)$ is enabled. Assume that this is not the case. This means that there is a DIOC event ε corresponding to some construct in \mathcal{I} . Because of the definition of \mathcal{R} , ε is also a DPOC event and $\varepsilon \leq_{DPOC} \xi : \bar{o} @ \mathbf{R}_2 \vee \varepsilon \leq_{DPOC} \xi : o @ \mathbf{R}_1$. Hence, at least one of the two events is not minimal and the corresponding transition cannot be enabled, against our hypothesis. Therefore the interaction $o : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x)$ is enabled. Thus, \mathcal{I} has a transition with label \surd and $\mathcal{I}' \xrightarrow{o: \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x)} \mathcal{I}''$. Thanks to Lemma 7.22 then both $(\pi(\mathcal{I}, \mathbf{R}_1), \Sigma_{\mathbf{R}_1})_{\mathbf{R}_1}$ and $(\pi(\mathcal{I}, \mathbf{R}_2), \Sigma_{\mathbf{R}_2})_{\mathbf{R}_2}$ have a transition with label \surd . Thus, we have

$$\begin{aligned} & (\pi(\mathcal{I}', \mathbf{R}_1), \Sigma_{\mathbf{R}_1})_{\mathbf{R}_1} \xrightarrow{\bar{i}.o\langle v \rangle @ \mathbf{R}_2 : \mathbf{R}_1} (P_{\mathbf{R}_1}, \Sigma_{\mathbf{R}_1})_{\mathbf{R}_1} \\ & (\pi(\mathcal{I}', \mathbf{R}_2), \Sigma_{\mathbf{R}_2})_{\mathbf{R}_2} \xrightarrow{i.o(x \leftarrow v) @ \mathbf{R}_1 : \mathbf{R}_2} (P_{\mathbf{R}_2}, \Sigma_{\mathbf{R}_2})_{\mathbf{R}_2} \text{ and thus} \\ & \text{proj}(\mathcal{I}', \Sigma) \xrightarrow{o: \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x)} \|_{\mathbf{R} \in \text{roles}(\mathcal{I}')} (P_{\mathbf{R}}, \Sigma_{\mathbf{R}})_{\mathbf{R}} \end{aligned}$$

The thesis follows by inductive hypothesis.

For the other cases of η , all the roles but one are unchanged. The proof of these cases is similar to the one for interaction, but simpler.

Note that in all the above cases, conditions on events follow by inductive hypothesis.

Case $\mathcal{I}|\mathcal{I}'$: from the definition of the projection function we have

$$\mathcal{N} = \|_{\mathbf{R} \in \text{roles}(\mathcal{I}; \mathcal{I}')} (\pi(\mathcal{I}, \mathbf{R}) \mid \pi(\mathcal{I}', \mathbf{R}), \Sigma_{\mathbf{R}})_{\mathbf{R}}$$

- : We have a case analysis on the rule used to derive the transition. If the transition is derived using Rule $[^{\text{DIOC}} |_{\text{PARALLEL}}]$ and $\mathcal{I}|\mathcal{I}'$ can perform a transition with label μ then one of its two components can perform a transition with the same label μ and the thesis follows by inductive hypothesis. Additional roles not occurring in the term performing the transition are dealt with by the **clean** part of function **upd**. If instead the transition is derived using Rule $[^{\text{DIOC}} |_{\text{PAR-END}}]$ then the thesis follows from Lemma 7.22.
- ←: We have a case analysis on the label η of the transition. If $\eta = o^? : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x)$ then a send and a receive on the same operation are enabled. The two events have the same global index since they have the same index \mathbf{i} (otherwise they could not synchronise) and they are both outside of any while loop (since they are enabled), hence the global index coincides with the index. Thus, they are either both from \mathcal{I} or both from \mathcal{I}' . The thesis follows by inductive hypothesis. For the other cases of η , only the process of one role changes. The thesis follows by inductive hypothesis. In all the cases, roles not occurring in the term performing the transition are dealt with by function **upd**.

Case i: if $b@R \{I\}$ else $\{I'\}$: from the definition of projection

$$\mathcal{N} = \left(\parallel_{S \in \text{roles}(I, I') \setminus \{R\}} \left(\mathbf{i}_?: \mathbf{i}. \mathbf{cnd}_i^* : x_i \mathbf{from} \mathbf{R}; \mathbf{i} \mathbf{if} x_i \{ \pi(I, S) \} \mathbf{else} \{ \pi(I', S) \}, \Sigma_S \right)_S \parallel \right. \\ \left. \left(\mathbf{i} \mathbf{if} b \left\{ \left(\prod_{R' \in \text{roles}(I, I') \setminus \{R\}} \mathbf{i}_T: \mathbf{i}. \mathbf{cnd}_i^* : \mathbf{true} \mathbf{to} \mathbf{R}' \right); \pi(I, R) \right\} \right. \right. \\ \left. \left. \mathbf{else} \left\{ \left(\prod_{R' \in \text{roles}(I, I') \setminus \{R\}} \mathbf{i}_F: \mathbf{i}. \mathbf{cnd}_i^* : \mathbf{false} \mathbf{to} \mathbf{R}' \right); \pi(I', R) \right\}, \Sigma_R \right) \right)_R$$

Let us consider the case when the guard is true (the other one is analogous).

\rightarrow : The only possible transition from the DIOC is $\mathbf{i} \mathbf{if} b@R \{I\}$ else $\{I'\} \xrightarrow{\tau} I$. The DPOC can match this transition by reducing to

$$\mathcal{N}' = \left(\parallel_{S \in \text{roles}(I, I') \setminus \{R\}} \left(\begin{array}{l} \mathbf{i}_?: \mathbf{i}. \mathbf{cnd}_i^* : x_i \mathbf{from} \mathbf{R}; \\ \mathbf{i} \mathbf{if} x_i \{ \pi(I, S) \} \mathbf{else} \{ \pi(I', S) \} \end{array}, \Sigma_S \right)_S \parallel \right. \\ \left. \left(\left(\prod_{R' \in \text{roles}(I, I') \setminus \{R\}} \mathbf{i}_T: \mathbf{i}. \mathbf{cnd}_i^* : \mathbf{true} \mathbf{to} \mathbf{R}' \right); \pi(I, R), \Sigma_R \right) \right)_R$$

By applying function upd we get

$$\text{upd}(\mathcal{N}') = \left(\parallel_{S \in \text{roles}(I, I') \setminus \{R\}} (\pi(I, S), \Sigma_S)_S \parallel (\pi(I, R), \Sigma_R)_R \right)$$

Concerning events, at the DIOC level events corresponding to the guard and to the discarded branch are removed. The same holds at the DPOC level, thus conditions on the remaining events are inherited. This concludes the proof.

\leftarrow : The only possible transition from the DPOC is the evaluation of the guard from the coordinator. This reduces \mathcal{N} to \mathcal{N}' above and the thesis follows from the same reasoning.

Case i: while $b@R \{I\}$: from the definition of projection

$$\mathcal{N} = \left(\parallel_{S \in \text{roles}(I) \setminus \{R\}} \left(\begin{array}{l} \mathbf{i}_?: \mathbf{i}. \mathbf{wb}_i^* : x_i \mathbf{from} \mathbf{R}; \mathbf{i} \mathbf{while} x_i \{ \pi(I, S) \}; \\ \mathbf{i}_C: \mathbf{i}. \mathbf{we}_i^* : \mathbf{ok} \mathbf{to} \mathbf{R}; \mathbf{i}_?: \mathbf{i}. \mathbf{wb}_i^* : x_i \mathbf{from} \mathbf{R} \end{array}, \Sigma_S \right)_S \parallel \right. \\ \left. \left(\mathbf{i} \mathbf{while} b \left\{ \left(\prod_{R' \in \text{roles}(I) \setminus \{R\}} \mathbf{i}_T: \mathbf{i}. \mathbf{wb}_i^* : \mathbf{true} \mathbf{to} \mathbf{R}' \right); \pi(I, R); \right. \right. \\ \left. \left. \prod_{R' \in \text{roles}(I) \setminus \{R\}} \mathbf{i}_C: \mathbf{i}. \mathbf{we}_i^* : \mathbf{-} \mathbf{from} \mathbf{R}' \right\}; \right. \\ \left. \prod_{R' \in \text{roles}(I) \setminus \{R\}} \mathbf{i}_F: \mathbf{i}. \mathbf{wb}_i^* : \mathbf{false} \mathbf{to} \mathbf{R}', \Sigma_R \right) \right)_R$$

\rightarrow : Let us consider the case when the guard is true. The only possible transition from the DIOC is $\mathbf{i} \mathbf{while} b@R \{I\} \xrightarrow{\tau} I; \mathbf{i} \mathbf{while} b@R \{I\}$. The DPOC can match this

transition by reducing to

$$\mathcal{N}' = \left\|_{\mathbf{S} \in \text{roles}(\mathcal{I}) \setminus \{\mathbf{R}\}} \left(\begin{array}{l} \mathbf{i}_?: \mathbf{i}.wb_i^* : x_i \text{ from } \mathbf{R}; \mathbf{i}: \text{while } x_i \{ \pi(\mathcal{I}, \mathbf{S}); \\ \mathbf{i}_c: \mathbf{i}.we_i^* : \text{ok to } \mathbf{R}; \mathbf{i}_?: \mathbf{i}.wb_i^* : x_i \text{ from } \mathbf{R} \} , \Sigma_{\mathbf{S}} \end{array} \right)_{\mathbf{S}} \left\| \right. \left(\begin{array}{l} \left(\prod_{\mathbf{R}' \in \text{roles}(\mathcal{I}) \setminus \{\mathbf{R}\}} \mathbf{i}_T: \mathbf{i}.wb_i^* : \text{true to } \mathbf{R}' \right); \pi(\mathcal{I}, \mathbf{R}); \\ \left(\prod_{\mathbf{R}' \in \text{roles}(\mathcal{I}) \setminus \{\mathbf{R}\}} \mathbf{i}_c: \mathbf{i}.we_i^* : _ \text{ from } \mathbf{R}' \right); \\ \mathbf{i}: \text{while } b \left\{ \begin{array}{l} \left(\prod_{\mathbf{R}' \in \text{roles}(\mathcal{I}) \setminus \{\mathbf{R}\}} \mathbf{i}_T: \mathbf{i}.wb_i^* : \text{true to } \mathbf{R}' \right); \pi(\mathcal{I}, \mathbf{R}); \\ \prod_{\mathbf{R}' \in \text{roles}(\mathcal{I}) \setminus \{\mathbf{R}\}} \mathbf{i}_c: \mathbf{i}.we_i^* : _ \text{ from } \mathbf{R}' \end{array} \right\}; \\ \prod_{\mathbf{R}' \in \text{roles}(\mathcal{I}) \setminus \{\mathbf{R}\}} \mathbf{i}_F: \mathbf{i}.wb_i^* : \text{false to } \mathbf{R}', \Sigma_{\mathbf{R}} \end{array} \right)_{\mathbf{R}}$$

By applying function `upd` we get

$$\text{upd}(\mathcal{N}') = \left(\left\|_{\mathbf{S} \in \text{roles}(\mathcal{I}) \setminus \{\mathbf{R}\}} \left(\begin{array}{l} \pi(\mathcal{I}, \mathbf{S}); \mathbf{i}_?: \mathbf{i}.wb_i^* : x_i \text{ from } \mathbf{R}; \\ \mathbf{i}: \text{while } x_i \{ \pi(\mathcal{I}, \mathbf{S}); \\ \mathbf{i}_c: \mathbf{i}.we_i^* : \text{ok to } \mathbf{R}; \\ \mathbf{i}_?: \mathbf{i}.wb_i^* : x_i \text{ from } \mathbf{R} \} , \Sigma_{\mathbf{S}} \end{array} \right)_{\mathbf{S}} \right\| \left(\begin{array}{l} \pi(\mathcal{I}, \mathbf{R}); \\ \mathbf{i}: \text{while } b \left\{ \begin{array}{l} \left(\prod_{\mathbf{R}' \in \text{roles}(\mathcal{I}) \setminus \{\mathbf{R}\}} \mathbf{i}_T: \mathbf{i}.wb_i^* : \text{true to } \mathbf{R}' \right); \pi(\mathcal{I}, \mathbf{R}); \\ \prod_{\mathbf{R}' \in \text{roles}(\mathcal{I}) \setminus \{\mathbf{R}\}} \mathbf{i}_c: \mathbf{i}.we_i^* : _ \text{ from } \mathbf{R}' \end{array} \right\}; \\ \prod_{\mathbf{R}' \in \text{roles}(\mathcal{I}) \setminus \{\mathbf{R}\}} \mathbf{i}_F: \mathbf{i}.wb_i^* : \text{false to } \mathbf{R}' \end{array} \right)_{\mathbf{R}}$$

exactly the projection of $\mathcal{I}; \mathbf{i}: \text{while } b @ \mathbf{R} \{ \mathcal{I} \}$.

As far as events are concerned, in $\text{compl}(\mathcal{N}')$ we have all the needed events since, in particular, we have already done the unfolding of the while in all the roles. Concerning the ordering, at the DIOC level, we have two kinds of causal dependencies: (1) events in the unfolded process precede the guard event; (2) the guard event precedes the events in the body. The first kind of causal dependency is matched at the DPOC level thanks to the auxiliary synchronisations that close the unfolded body (which are not removed by `compl`) using synchronisation and sequentiality. The second kind of causal dependency is matched thanks to the auxiliary synchronisations that start the following iteration using synchronisation, sequentiality and while.

The case when the guard evaluates to **false** is simpler.

←: The only possible transition from the DPOC is the evaluation of the guard from the coordinator. This reduces \mathcal{N} to \mathcal{N}' above and the thesis follows from the same reasoning.

Case i: scope @R {I}: from the definition of the projection

$$\mathcal{N} = \left(\parallel_{\mathbf{R}' \in \text{roles}(\mathcal{I}) \setminus \{\mathbf{R}\}} (\mathbf{i}: \text{scope } @\mathbf{R} \{ \pi(\mathcal{I}, \mathbf{R}') \}, \Sigma_{\mathbf{R}'})_{\mathbf{R}'} \parallel \right. \\ \left. (\mathbf{i}: \text{scope } @\mathbf{R} \{ \pi(\mathcal{I}, \mathbf{R}) \} \text{ roles } \{ \text{roles}(\mathcal{I}) \}, \Sigma_{\mathbf{R}})_{\mathbf{R}} \right)$$

→: Let us consider the case when the scope is updated. At the DIOC level all the possible transitions have label of the form \mathcal{I}' and are obtained by applying Rule $[\text{DIOC}]_{\text{UP}}$. Correspondingly, at the DPOC level one applies Rule $[\text{DPOC}]_{\text{LEAD-UP}}$ to the coordinator of the update, obtaining

$$\mathcal{N}' = \left(\parallel_{\mathbf{R}' \in \text{roles}(\mathcal{I}) \setminus \{\mathbf{R}\}} (\mathbf{i}: \text{scope } @\mathbf{R} \{ \pi(\mathcal{I}, \mathbf{R}') \}, \Sigma_{\mathbf{R}'})_{\mathbf{R}'} \parallel \right. \\ \left. \left(\begin{array}{l} \left(\prod_{\mathbf{R}' \in \text{roles}(\mathcal{I}) \setminus \{\mathbf{R}\}} \mathbf{i}.sb_i^* : \pi(\mathcal{I}', \mathbf{R}') \text{ to } \mathbf{R}' \right); \\ \pi(\mathcal{I}', \mathbf{R}); \\ \prod_{\mathbf{R}' \in \text{roles}(\mathcal{I}) \setminus \{\mathbf{R}\}} \mathbf{i}.se_i^* : _ \text{ from } \mathbf{R}' \end{array} \right), \Sigma_{\mathbf{R}} \right)_{\mathbf{R}}$$

By applying the upd function we get:

$$\text{upd}(\mathcal{N}') = \left(\parallel_{\mathbf{R}' \in \text{roles}(\mathcal{I}) \setminus \{\mathbf{R}\}} (\pi(\mathcal{I}', \mathbf{R}'), \Sigma_{\mathbf{R}'})_{\mathbf{R}'} \parallel (\pi(\mathcal{I}', \mathbf{R}), \Sigma_{\mathbf{R}})_{\mathbf{R}} \right)$$

This is exactly the projection of the DIOC obtained after applying the rule $[\text{DIOC}]_{\text{UP}}$. The conditions on events are inherited. Observe that the closing event of the scope is replaced by events corresponding to the auxiliary interactions closing the scope. This allows us to preserve the causality dependencies also when the scope is inserted in a context.

The case of Rule $[\text{DIOC}]_{\text{NOUP}}$ is simpler.

←: The only possible transitions from the DPOC are the ones of the coordinator of the update checking whether to apply an update or not. This reduces \mathcal{N} to \mathcal{N}' above and the thesis follows from the same reasoning. \square

7.2. Deadlock freedom, termination, and race freedom. Due to the fact that the projection preserves weak traces, we have that trace-based properties of the DIOC are inherited by the DPOC. A first example of such properties is *deadlock freedom*.

Definition 7.23 (Deadlock freedom). An *internal* DIOC (resp. DPOC) *trace* is obtained by removing transitions labelled \mathbf{I} from a DIOC (resp. DPOC) trace. A DIOC (resp. DPOC) system is *deadlock free* if all its maximal finite internal traces have \checkmark as last label.

Intuitively, internal traces are needed since labels \mathbf{I} do not correspond to activities of the application and may be executed also after application termination. The fact that after a \checkmark only changes in the set of available updates are possible is captured by the following lemma.

Lemma 7.24. For each initial, connected DIOC \mathcal{I} , state Σ , and set of updates \mathbf{I} , if $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle \xrightarrow{\checkmark} \langle \Sigma', \mathbf{I}', \mathcal{I}' \rangle$ then each transition of $\langle \Sigma', \mathbf{I}', \mathcal{I}' \rangle$ has label \mathbf{I}'' for some \mathbf{I}'' .

Proof. The proof is by case analysis on the rules which can derive a transition with label \checkmark . All the cases are easy. \square

Since by construction initial DIOCs are deadlock free we have that also the DPOC obtained by projection is deadlock free.

Corollary 7.25 (Deadlock freedom). *For each initial, connected DIOC \mathcal{I} , state Σ , and set of updates \mathbf{I} the DPOC system $\langle \mathbf{I}, \text{proj}(\mathcal{I}, \Sigma) \rangle$ is deadlock free.*

Proof. Let us first prove that for each initial, connected DIOC \mathcal{I} , state Σ , and set of updates \mathbf{I} , the DIOC system $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle$ is deadlock free. This amounts to prove that all its maximal finite internal traces have \surd as last label. For each trace, the proof is by induction on its length, and for each length by structural induction on \mathcal{I} . The proof is based on the fact that \mathcal{I} is initial. The induction considers a reinforced hypothesis, saying also that:

- \surd may occur *only* as the last label of the internal trace;
- all the DIOC systems in the sequence of transitions generating the trace, but the last one, are initial.

We have a case analysis on the top-level operator in \mathcal{I} . Note that in all the cases, but $\mathbf{0}$, at least a transition is derivable.

Case $\mathbf{0}$: not allowed since we assumed an initial DIOC.

Case $\mathbf{1}$: trivial because by Rule $[\text{DIOC}_{\text{END}}]$ and Lemma 7.24 its only internal trace is \surd .

Case $x@\mathbf{R} = e$: the only applicable rule is $[\text{DIOC}_{\text{ASSIGN}}]$ that in one step leads to a $\mathbf{1}$ process. The thesis follows by inductive hypothesis on the length of the trace.

Case $o^? : \mathbf{R}_1(e) \rightarrow \mathbf{R}_2(x)$: the only applicable rule is $[\text{DIOC}_{\text{INTERACTION}}]$, which leads to an assignment. Then the thesis follows by inductive hypothesis on the length of the trace.

Case $\mathcal{I}; \mathcal{I}'$: the first transition can be derived either by Rule $[\text{DIOC}_{\text{SEQUENCE}}]$ or Rule $[\text{DIOC}_{\text{SEQ-END}}]$. In the first case the thesis follows by induction on the length of the trace. In the second case the trace coincides with a trace of \mathcal{I}' , and the thesis follows by structural induction.

Case $\mathcal{I}|\mathcal{I}'$: the first transition can be derived either by Rule $[\text{DIOC}_{\text{PARALLEL}}]$ or by Rule $[\text{PAR-END}]$. In the first case the thesis follows by induction on the length of the trace. In the second case the thesis follows by Lemma 7.24, since the label is \surd .

Case **if $b@\mathbf{R} \{ \mathcal{I} \} \text{ else } \{ \mathcal{I}' \}$:** the first transition can be derived using either Rule $[\text{DIOC}_{\text{IF-THEN}}]$ or Rule $[\text{DIOC}_{\text{IF-ELSE}}]$. In both the cases the thesis follows by induction on the length of the trace.

Case **while $b@\mathbf{R} \{ \mathcal{I} \}$:** the first transition can be derived using either Rule $[\text{DIOC}_{\text{WHILE-UNFOLD}}]$ or Rule $[\text{DIOC}_{\text{WHILE-EXIT}}]$. In both the cases the thesis follows by induction on the length of the trace.

Case **scope $@\mathbf{R} \{ \mathcal{I} \}$:** the first transition can be derived using either Rule $[\text{DIOC}_{\text{UP}}]$ or Rule $[\text{DIOC}_{\text{NOUP}}]$. In both the cases the thesis follows by induction on the length of the trace.

The weak internal traces of the DIOC coincide with the weak internal traces of the DPOC by Theorem 7.2, thus also the finite weak internal traces of the DPOC end with \surd . The same holds for the finite strong internal traces, since label \surd is preserved when moving between strong and weak traces, and no transition can be added after the \surd thanks to Lemma 7.24. □

DPOCs also inherit termination from terminating DIOCs.

Definition 7.26 (Termination). A DIOC (resp. DPOC) system terminates if all its internal traces are finite.

Note that if arbitrary sets of updates are allowed, then termination of DIOCs that contain at least a scope is never granted. Indeed, the scope can always be replaced by a non-terminating update or it can trigger an infinite chain of updates. Thus, to exploit this result, one should add constraints on the set of updates ensuring DIOC termination.

Corollary 7.27 (Termination). *If the DIOC system $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle$ terminates and \mathcal{I} is connected then the DPOC system $\langle \mathbf{I}, \text{proj}(\mathcal{I}, \Sigma) \rangle$ terminates.*

Proof. It follows from the fact that only a finite number of auxiliary actions are added when moving from DIOCs to DPOCs. \square

Other interesting properties derived from weak trace equivalence are freedom from races and orphan messages. A race occurs when the same receive (resp. send) may interact with different sends (resp. receives). In our setting, an orphan message is an enabled send that is never consumed by a receive. Orphan messages are more relevant in asynchronous systems, where a message may be sent, and stay forever in the network, since the corresponding receive operation may never become enabled. However, even in synchronous systems orphan messages should be avoided: the message is not communicated since the receive is not available, hence a desired behaviour of the application never takes place due to synchronisation problems. Trivially, DIOCs avoid races and orphan messages since send and receive are bound together in the same construct. Differently, at the DPOC level, since all receive of the form $\iota: \mathbf{i}.o^? : x$ **from** \mathbf{R}_1 in role \mathbf{R}_2 may interact with the sends of the form $\iota: \mathbf{i}.o^? : e$ **to** \mathbf{R}_2 in role \mathbf{R}_1 , races may happen. However, thanks to the correctness of the projection, race freedom holds also for the projected DPOCs.

Corollary 7.28 (Race freedom). *For each initial, connected DIOC \mathcal{I} , state Σ , and set of updates \mathbf{I} , if $\langle \mathbf{I}, \text{proj}(\mathcal{I}, \Sigma) \rangle \xrightarrow{\eta_1} \dots \xrightarrow{\eta_n} \langle \mathbf{I}', \mathcal{N} \rangle$, where $\eta_i \in \{\tau, o^? : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x), o^* : \mathbf{R}_1(X) \rightarrow \mathbf{R}_2(), \surd, \mathcal{I}, \text{no-up}, \mathbf{I}\}$ for each $i \in \{1, \dots, n\}$, then in \mathcal{N} there are no two sends (resp. receives) which can interact with the same receive (resp. send).*

Proof. We have two cases, corresponding respectively to programmer-specified and auxiliary operations.

For programmer-specified operations, thanks to Lemma 7.20, case C1, for each global index ξ there are at most two communication events with global index ξ . The corresponding DPOC terms can be enabled only if they are outside of the body of a while loop. Hence, their index coincides with their global index. Since the index prefixes the operation, then no interferences with other sends or receives are possible.

For auxiliary operations, the reasoning is similar. Note, in fact, that sends or receives with the same global index can be created only by a unique DIOC construct, but communications between the same pair of roles are never enabled together. This can be seen by looking at the definition of the projection. \square

As far as orphan messages are concerned, they may appear in infinite DPOC computations since a receive may not become enabled due to an infinite loop. However, as a corollary of trace equivalence, we have that terminating DPOCs are orphan-message free.

Corollary 7.29 (Orphan-message freedom). *For each initial, connected DIOC \mathcal{I} , state Σ , and set of updates \mathbf{I} , if $\langle \mathbf{I}, \text{proj}(\mathcal{I}, \Sigma) \rangle \xrightarrow{\eta_1} \dots \xrightarrow{\eta_n} \surd \langle \mathbf{I}', \mathcal{N} \rangle$, where $\eta_i \in \{\tau, o^? : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x), o^* : \mathbf{R}_1(X) \rightarrow \mathbf{R}_2(), \surd, \mathcal{I}, \text{no-up}, \mathbf{I}\}$, then \mathcal{N} contains no sends.*

Proof. The proof is by case analysis on the rules which can derive a transition with label \surd . All the cases are easy. \square

8. ADAPTABLE INTERACTION-ORIENTED CHOREOGRAPHIES IN JOLIE

In this section we present AIOCJ (Adaptable Interaction-Oriented Choreographies in Jolie), a development framework for adaptable distributed applications [28]. AIOCJ is one of the possible instantiations of the theoretical framework presented in the previous sections and it gives a tangible proof of the expressiveness and feasibility of our approach. We say that AIOCJ is an instance of our theory because it follows the theory, but it provides mechanisms to resolve the non-determinism related to the choice of whether to update or not and on which update to select. Indeed, in AIOCJ updates are chosen and applied according to the state of the application and of its running environment. To this end, updates are embodied into adaptation rules, which specify when and whether a given update can be applied, and to which scopes. AIOCJ also inherits all the correctness guarantees provided by our theory, in particular: *i*) applications are free from deadlocks and races by construction, *ii*) applications remain correct after any step of adaptation. As in the theory, adaptation rules can be added and removed while applications are running.

Remark 8.1. In order for the correctness guarantees to be provided, one needs to satisfy the required assumptions, in particular the fact that functions never block and always return a value (possibly an error notification). If this condition is not satisfied, at the theoretical level, the behaviour of the source DIOC and the behaviour of its projection may differ, as described in Remark 2.2. The behaviour of the corresponding AIOCJ application can be different from both of them. Usually, the violation of the assumption above leads the application to crash. This allows the programmer to realise that the assumption has been violated.

Below we give a brief overview of the AIOCJ framework by introducing its components: the Integrated Development Environment (IDE), the AIOCJ compiler, and the Runtime Environment.

Integrated Development Environment. AIOCJ supports the writing of programs and adaptation rules in the Adaptable Interaction-Oriented Choreography (AIOC) language, an extension of the DIOC language. We discuss the main novelties of the AIOC language in § 8.1. AIOCJ offers an integrated environment for developing programs and adaptation rules that supports syntax highlighting and on-the-fly syntax checking. Since checking for connectedness (see § 6) of programs and adaptation rules is polynomial (as proven by Theorem 6.3), the IDE also performs on-the-fly checks on connectedness of programs and rules.

Compiler. The AIOCJ IDE also embeds the AIOCJ compiler, which implements the procedure for projecting AIOCs and adaptation rules into distributed executable code. The implementation of the compiler is based on the rules for projecting DIOCs, described in § 4. The target language of the AIOCJ compiler is Jolie [26, 35], a Service-Oriented language with primitives similar to those of our DPOC language. Jolie programs are also called *services*. Given an AIOC program, the AIOCJ compiler produces one Jolie service for each role in the source AIOC. The compilation of an AIOC rule produces one Jolie service for each role and an additional service that describes the applicability condition of the rule. All these services are enclosed into an **Adaptation Server**, described below.

Runtime Environment. The AIOCJ runtime environment comprises a few Jolie services that support the execution and adaptation of compiled programs. The main services of the AIOCJ runtime environment are the **Adaptation Manager**, **Adaptation Servers**, and the **Environment**. The compiled services interact both among themselves and with an **Adaptation Manager**, which is in charge of managing the adaptation protocol. **Adaptation Servers** contain adaptation rules, and they can be added or removed dynamically, thus enabling dynamic changes in the set of rules, as specified by Rule $[^{\text{DIOC}} |_{\text{CHANGE-UPDATES}}]$. When started, an **Adaptation Server** registers itself at the **Adaptation Manager**. The **Adaptation Manager** invokes the registered **Adaptation Servers** to check whether their adaptation rules are applicable. In order to check whether an adaptation rule is applicable, the corresponding **Adaptation Server** evaluates its applicability condition. Applicability conditions may refer to the state of the role which coordinates the update, to properties of the scope, and to properties of the environment (e.g., time, temperature, etc.), stored in the **Environment** service.

In the remainder of this Section we detail the grammar of the AIOC language used by AIOCJ in § 8.1, we illustrate the use of AIOCJ on a simple example in § 8.2, we discuss some relevant implementation aspects of AIOCJ in § 8.3, and we present some guidelines on how to use scopes in AIOCJ in § 8.4.

8.1. DIOC Language Extensions in AIOCJ. AIOCs extend DIOCs with:

- the definition of adaptation rules, instead of updates, that include the information needed to evaluate their applicability condition;
- the definition of constructs to express the deployment information needed to implement real-world distributed applications.

Below we describe in detail, using the Extended Backus-Naur Form [3], the new or refined constructs introduced by the AIOC language.

Function inclusions. The AIOC language can exploit functionalities provided by external services via the `include` construct. The syntax is as follows.

```
Include ::= include FName [, FName]* from "URL" [with PROTOCOL]
```

This allows one to reuse existing legacy code and to interact with third-party external applications. As an example, the **Seller** of our running example can exploit an external database to implement the functionality for price retrieval `getPrice`, provided that such a functionality is exposed as a service. If the service is located at `"socket://myService:8000"` and accessible via the `"SOAP"` protocol we enable its use with the following inclusion:

```
include getPrice from "socket://myService:8000" with "SOAP"
```

Similarly, the **Bank** IT system is integrated in the example by invoking the function `makePayment`, also exposed as a service.

This feature enables a high degree of integration since AIOCJ supports all protocols provided by the underlying Jolie language, which include TCP/IP, RMI, SOAP, XML/RPC, and their encrypted correspondents over SSL.

External services perfectly fit the theory described in previous sections, since they are seen as functions, and thus introduced in expressions. Notice, in particular, that Remark 8.1 applies.

Adaptation rules. Adaptation rules extend DIOC updates, and are a key ingredient of the AIOCJ framework. The syntax of adaptation rules is as follows:

```

Rule ::= rule {
  [ Include ]*
  on { Condition }
  do { Choreography }
}

```

The applicability condition of the adaptation rule is specified using the keyword `on`, while the code to install in case adaptation is performed (which corresponds to the DIOC update) is specified using the keyword `do`. Optionally, adaptation rules can include functions they rely on.

The Condition of an adaptation rule is a propositional formula which specifies when the rule is applicable. To this end, it can exploit three sources of information: local variables of the coordinator of the update, environmental variables, and properties of the scope to which the adaptation rule is applied. Environmental variables are meant to capture contextual information that is not under the control of the application (e.g., temperature, time, available resources, ...). To avoid ambiguities, local variables of the coordinator are not prefixed, environment variables are prefixed by `E`, and properties of the scope are prefixed by `N`.

For example, if we want to apply an adaptation rule only to those scopes whose property name is equal to the string `"myScope"` we can use as applicability condition the formula `N.name == "myScope"`.

Scopes. As described above, scopes in AIOC also feature a set of properties (possibly empty). Scope properties describe the current implementation of the scope, including both functional and non-functional properties. Such properties are declared by the programmer, and the system only uses them to evaluate the applicability condition of adaptation rules, to decide whether a given adaptation rule can be applied to a given scope. Thus the syntax for scopes in AIOC is:

```

scope @Role { Choreography }
[ prop { Properties } ]?

```

where clause `prop` introduces a list of comma-separated assignments of the form `N.ID = Expression`.

For instance, the code

```

scope @R {
  //AIOC code
} prop { N.name = "myScope" }

```

specifies that the scope has a property name set to the string `"myScope"`, thus satisfying the applicability condition `N.name == "myScope"` discussed above.

Programs. AIOC programs have the following structure.

```

Program ::=
  [ Include ]*
  preamble {
    starter: Role
    [ Location ]*
  }
  aioc { Choreography }

```

where `Include` allows one to include external functionalities, as discussed above, and keyword `aioc` introduces the behaviour of the program, which is a DIOC apart for the fact that scopes may define properties, as specified above.

The keyword `preamble` introduces deployment information, i.e., the definition of the `starter` of the AIOC and the `Location` of participants.

The definition of a `starter` is mandatory and designs which role is in charge of waiting for all other roles to be up and running before starting the actual computation. Any role can be chosen as `starter`, but the chosen one needs to be launched first when running the distributed application.

Locations define where the participants of the AIOC will be deployed. They are specified using the keyword `location`:

```
Location ::= location@Role:"URL"
```

where `Role` is the name of a role (e.g., `Role1`) and `URL` specifies where the service can be found (e.g., `"socket://Role1:8001"`). When not explicitly defined, the projection automatically assigns a distinct local TCP/IP location to each role.

8.2. AIOCJ Workflow. Here we present a brief description of how a developer can write an adaptable distributed system in AIOCJ, execute it, and change its behaviour at runtime by means of adaptation rules. For simplicity, we reuse here the minimal example presented in the Introduction, featuring a scope that encloses a price offer from the **Seller** to the **Buyer** and an update — here an adaptation rule — that provides a discount for the **Buyer**. We report the AIOC program in the upper part and the adaptation rule in the lower part of Figure 13.

At Line 1 of the AIOC program we have the inclusion of function `getPrice`, which is provided by a service within the internal network of the Seller, reachable as a TCP/IP node at URL `"storage.seller.com"` on port `"80"`. At Line 2 of the AIOC program we have the `preamble`. The preamble specifies deployment information, and, in particular, defines the `starter`, i.e., the service that ensures that all the participants are up and running before starting the actual computation. No locations are specified, thus default ones are used. The actual code is at Lines 4–7, where we declare a scope. At Line 7 we define a property `scope_name` of this scope with value `"price_inquiry"`.

Figure 14 depicts the process of compilation ① and execution ② of the AIOC. From left to right, we write the AIOC and we compile it into a set of executable Jolie services (**Buyer service** and **Seller service**). To execute the projected system, we first launch the **Adaptation Manager** and then the two compiled services, starting from the **Seller**, which is the `starter`. Since there is no compiled adaptation rule, the result of the execution is the offering of the standard price to the Buyer.

Now, let us suppose that we want to adapt our system to offer discounts in the Fall season. To do that, we can write the adaptation rule shown in the lower part of Figure 13.

Since we want to replace the scope for the `price_inquiry`, we define, at Line 4, that this rule applies only to scopes with property `scope_name` set to `"price_inquiry"`. Furthermore, since we want the update to apply only in the Fall season, we also specify that the environment variable `E.season` should match the value `"Fall"`. The value of `E.season` is retrieved from the **Environment Service**.

The rule uses two functions. At Line 2, it includes function `getPrice`, which is the one used also in the body of the scope. The other function, `isValid`, is included at Line 3 and

```

1 include getPrice from "socket://storage.seller.com:80"
2 preamble{ starter: Seller }
3 aioc {
4   scope @Seller {
5     order_price@Seller = getPrice( order );
6     offer: Seller( order_price ) -> Buyer( prod_price )
7   } prop { N.scope_name = "price_inquiry" } }

```

```

1 rule {
2   include getPrice from "socket://storage.seller.com:80"
3   include isValid from "socket://discounts.seller.com:80"
4   on { N.scope_name == "price_inquiry" and E.season == "Fall" }
5   do {
6     cardReq: Seller( _ ) -> Buyer( _ );
7     card_id@Buyer = getInput( "Insert your customer card ID");
8     cardRes: Buyer( card_id ) -> Seller( buyer_id );
9     if isValid( buyer_id )@Seller {
10      order_price@Seller = getPrice( order )*0.9;
11    } else {
12      order_price@Seller = getPrice( order )
13    };
14    offer: Seller( order_price ) -> Buyer( prod_price )
15  }
16 }

```

Figure 13: An AIOC program (upper part) and an applicable adaptation rule (lower part).

it is provided by a different TCP/IP node, located at URL `"discounts.seller.com"` on port `"80"` within the internal network of the Seller.

The body of the adaptation rule (Lines 6–14) specifies the same behaviour described in Figure 1 in the Introduction: the Seller asks the Buyer to provide its `card_id`, which the Seller uses to provide a discount on the price of the ordered product. We depict the inclusion of the new adaptation rule (outlined with dashes) and the execution of the adaptation at point ③ of Figure 14. From right to left, we write the rule and we compile it. The compilation of a (set of) adaptation rule(s) in AIOCJ produces a service, called **Adaptation Server** (also outlined with dashes), that the **Adaptation Manager** can query to fetch adaptation rules at runtime. The compilation of the adaptation rule can be done while the application is running. After the compilation, the generated **Adaptation Server** is started and registers on the **Adaptation Manager**. Since the rule relies on the environment to check its applicability condition, we also need the **Environment service** to be running. In order for the adaptation rule to apply we need the environment variable `season` to have value `"Fall"`.

8.3. Implementation. AIOCJ is composed of two elements: the AIOCJ Integrated Development Environment (IDE), named `AIOCJ – ecl`, and the adaptation middleware that enables AIOC programs to adapt, called `AIOCJ – mid`.




Figure 14: Representation of the AIOcJ framework — Projection and execution of the example in Figure 13.

```

1 include getPrice, getPaymentMethods from "socket://localhost:800
2
3 preamble{
4   starter: Buyer
5 }
6
7 aioc {
8   price_ok@Buyer = false;
9   continue@Buyer = true;
10  while( !price_ok and continue )@Buyer {
11    prod@Buyer = getInput( "Insert product" );
12    //priceReq: Buyer( prod ) -> Seller( order );
13    scope @Seller {
14      order price@Seller = getPrice( order );
15      offer: Seller( order price ) -> Buyer( prod price )
16    } prop { N.scope_name = "price inquiry" };
17    getInput( "Do you accept the price?[y/n]" );
18    "n" )@Buyer {
19      price_ok@Buyer = false;
20      continue@Buyer = getInput( "Ask another product?[y/n]" );
  
```

The screenshot shows the Eclipse IDE with the Package Explorer on the left and the AIOc code editor on the right. The code editor displays the AIOc code for 'Purchase_Example.ioc'. A red error icon is visible next to line 13, with a tooltip that reads: "The sequence is not connected".

Figure 15: Check for connectedness.

AIOcJ – ecl is a plug-in for Eclipse [21] based on Xtext [49]. Starting from a grammar, Xtext generates the parser for programs written in the AIOc language. Result of the parsing is an abstract syntax tree (AST) we use to implement *i*) the checker of connectedness for AIOc programs and adaptation rules and *ii*) the generation of Jolie code for each role. Since the check for connectedness has polynomial computational complexity (cf. Theorem 6.3) it is efficient enough to be performed while editing the code. Figure 15 shows AIOcJ – ecl notifying the error on the first non-connected instruction (Line 13).

As already mentioned, we chose Jolie as target language of the compilation of AIOCJ because its semantics and language constructs naturally lend themselves to translate our theoretical results into practice. Indeed, Jolie supports architectural primitives like dynamic embedding, aggregation, and redirection, which ease the compilation of AIOCs.

Each scope at the AIOC level is projected into a specific sub-service for each role. The roles run the projected sub-services by embedding them and access them via redirection. In this way, we implement adaptation by disabling the default sub-service and by redirecting the execution to a new one, obtained from the **Adaptation server**.

When at runtime the coordinator of the update reaches the beginning of that scope, it queries the **Adaptation Manager** for adaptation rules to apply. The **Adaptation Manager** queries each **Adaptation Server** sequentially, based on their order of registration. On its turn, each **Adaptation Server** checks the applicability condition of each of its rules. The first rule whose applicability condition holds is applied. The adaptation manager sends to the coordinator the updates that are then distributed to the other involved roles. In each role, the new code replaces the old one.

To improve the performance, differently from the theory, in AIOCJ adaptation rules are compiled statically, and not by the coordinator of the update when the update is applied.

8.4. AIOCJ Practice. AIOCJ provides scopes as a way to specify which parts of the application can be updated. As a consequence, deciding which parts of the code to enclose in scopes is a relevant and non trivial decision. Roughly, one should enclose into scopes parts of the code that may need to be updated in the future. This includes for instance parts of the code which model business rules, hence may change according to business needs, parts of the code which are location dependent, hence may be updated to configure the software for running in different locations, parts of the code which are relevant for performance or security reasons, hence may be updated to improve the performance or the security level.

One may think that a simple shortcut would be to have either a big scope enclosing the whole code of the application or many small scopes covering all the code, but both these solutions have relevant drawbacks. Indeed, if one replaces a big scope in order to change a small piece of it, (s)he has to provide again all the code, while one would like to provide only the part that needs to change. Furthermore, update for the big scope is no more available as soon as its execution starts, which may happen quite earlier than when the activity to be updated starts. Using small scopes also is not a solution, since an update may need to cover more than one such scope, but there is no way to ensure that multiple scopes are updated in a coordinated way. Also, both the approaches have performance issues: in the first case large updates have to be managed, in the second case many checks for the updates are needed, causing a large overhead.

In order to write updates for a given scope, one needs to have some knowledge about the scope that the update is going to replace, since the update will run in the context where the scope is. This includes, for instance, knowing which variables are available, their types and their meaning, and in which variables results are expected. This information can be either tracked in the documentation of the application, or added to the scope properties. In this way, the adaptation rule embodying the update can check this information before applying the update (see for instance the approach of [29]). The fact that this information is needed for all scopes is another point against the approach of using many small scopes.

9. RELATED WORK AND DISCUSSION

Recently, languages such as Rust [44] or SCOOP [41] have been defined to provide high-level primitives to program concurrent applications avoiding by construction some of the risks of concurrent programming. For instance, Rust ensures that there is exactly one binding to any given resource using ownership and borrowing mechanisms to achieve memory safety. SCOOP instead provides the `separate` keyword to enable programming patterns that enforce data-race freedom. In industry, Rust has started to be used for the development of complex concurrent/distributed systems (e.g., the file storage system of Dropbox [33] or the parallel rendering engine of Mozilla [51]). However, industry has not yet widely adopted any language-based solution trading expressive power of the language for safety guarantees.

Inspired by the languages above, we have presented high-level primitives for the dynamic update of distributed applications. We guarantee the absence of communication deadlocks and races by construction for the running distributed application, even in presence of updates that were unknown when the application was started. More generally, the compilation of a DIOC specification produces a set of low-level DPOCs whose behaviour is compliant with the behaviour of the originating DIOC.

As already remarked, our theoretical model is very general. Indeed, whether to update a scope or not and which update to apply if many are available is completely non deterministic. Choosing a particular policy for decreasing the non-determinism of updates is orthogonal with respect to our results. Our properties are preserved by any such policy, provided that the same policy is applied both at the DIOC and at the DPOC level. AIOCJ presents a possible instantiation of our general approach with a concrete mechanism for choosing updates according to the state of the application and of its running environment.

Our work is on the research thread of choreographies and multiparty session types [12, 24, 13, 14, 5, 45]. One can see [25] for a description of the field. Like choreographies and multiparty session types, DIOCs target communication-centred concurrent and distributed applications, avoiding deadlocks and communication races. Thanks to the particular structure of the language, DIOCs provide the same guarantees without the need of types. Below we just discuss the approaches in this research thread closest to ours. The two most related approaches we are aware of are based on multiparty session types, and deal with dynamic software updates [2] and with monitoring of self-adaptive systems [16]. The main difference between [2] and our approach is that [2] targets concurrent applications which are not distributed. Indeed, it relies on a check on the global state of the application to ensure that the update is safe. Such a check cannot be done by a single role, thus it is impractical in a distributed setting. Furthermore, the language in [2] is much more constrained than ours, e.g., requiring each pair of participants to interact on a dedicated pair of channels, and assuming that all the roles that are not the sender or the receiver within a choice behave the same in the two branches. The approach in [16] is very different from ours, too. In particular, in [16], all the possible behaviours are available since the very beginning, both at the level of types and of processes, and a fixed adaptation function is used to switch between them. This difference derives from the distinction between self-adaptive applications, as they discuss, and applications updated from the outside, as in our case.

We also recall [20], which uses types to ensure safe adaptation. However, [20] allows updates only when no session is active, while we change the behaviour of running DIOCs.

We highlight that, contrarily to our approach, none of the approaches above has been implemented.

Various tools [17, 4, 8] support adaptation exploiting automatic planning techniques in order to elaborate, at runtime, the best sequence of activities to achieve a given goal. These techniques are more declarative than ours, but, to the best of our knowledge, they are not guaranteed to always find a plan to adapt the application.

[38] presents an approach to safe dynamic software updates for concurrent systems. According to [38], safe means that all the traces are version consistent, namely that for each trace there is a version of the application able to produce it. In our case, we want the effects of the update to be visible during the computation, hence we do not want this property to hold. Instead, for us safe means that the DIOC and the projected DPOC are weak system bisimilar, and that they are both free from deadlocks and races.

Among the non-adaptive languages, Chor [13] is the closest to ours. Indeed, like ours, Chor is a choreographic language that compiles to Jolie. Actually, AIOCJ shares part of the Chor code base. Our work shares with [37] the interest in choreographies composition. However, [37] uses multiparty session types and only allows static parallel composition, while we replace a term inside an arbitrary context at runtime.

Extensions of multiparty session types with error handling [11, 9] share with us the difficulties in coordinating the transition from the expected pattern to an alternative one, but in their case the error recovery pattern is known since the very beginning, thus considerably simplifying the analysis.

We briefly compare to some works that exploit choreographic descriptions for adaptation, but with very different aims. For instance, [27] defines rules for adapting the specification of the initial requirements for a choreography, thus keeping the requirements up-to-date in presence of run-time changes. Our approach is in the opposite direction: we are not interested in updating the system specification tracking system updates, but in programming and ensuring correctness of adaptation itself. Other formal approaches to adaptation represent choreographies as annotated finite state automata. In [43] choreographies are used to propagate protocol changes to the other peers, while [48] presents a test to check whether a set of peers obtained from a choreography can be reconfigured to match a second one. Differently from ours, these works only provide change recommendations for adding and removing message sequences.

An approach close to ours is multi-tier programming [39, 6, 15], where the programmer writes a single program that is later automatically distributed to different tiers. This is similar to what we do by projecting a DIOC on the different roles. In particular, [39] considers functional programs with role annotations and splits them into different sequential programs to be run in parallel. Multi-tier approaches such as HipHop [6] or Link [15] have been implemented and used to develop three tier web applications. Differently from ours, these works abstract away communication information as much as possible, while we emphasise this kind of information. Furthermore, they do not take into account code update. It would be interesting to look for cross-fertilisation results between our approach and multi-tier programming. For instance, one could export our techniques to deal with runtime updates into multi-tier programming. In the other direction, we could abstract away communications and let the compiler manage them according to the programmed flow of data, as done in multi-tier programming. For instance, in Figure 2 the communications at Lines 22, 24, and 26 could be replaced by local assignments, since the information on the chosen branch is carried by the auxiliary communications. In general, communications seem not strictly needed from a purely technical point of view, however they are relevant for the systems we currently consider. Furthermore, the automatic introduction of communications

in an optimised way is not always trivial since the problem is NP-hard even for languages that do not support threads [40].

On a broader perspective, our theory can be used to inject guarantees of freedom from deadlocks and races into many existing approaches to adaptation, e.g., the ones in the surveys [32, 22]. However, this task is cumbersome, due to the huge number and heterogeneity of those approaches. For each of them the integration with our techniques is far from trivial. Nevertheless, we already started it. Indeed, AIOCJ follows the approach to adaptation described in [29]. However, applications in [29] are not distributed and there are no guarantees on the correctness of the application after adaptation. Furthermore, in the website of the AIOCJ project [1], we give examples of how to integrate our approach with distributed [42] and dynamic [50] Aspect-Oriented Programming (AOP) and with Context-Oriented Programming (COP) [23]. In general, we can deal with cross-cutting concerns like logging and authentication, typical of AOP, viewing pointcuts as empty scopes and advices as updates. Layers, typical of COP, can instead be defined by updates which can fire according to contextual conditions. As future work, besides applying the approach to other adaptation mechanisms, we also plan to extend our techniques to deal with multiparty session types [12, 24, 13, 14]. The main challenge here is to deal with multiple interleaved sessions whilst each of our choreographies corresponds to a single session. An initial analysis of the problem is presented in [7].

Also the study of more refined policies for rule selection, e.g., based on priorities, is a topic for future work.

REFERENCES

- [1] AIOCJ website. <http://www.cs.unibo.it/projects/jolie/aioj.html>.
- [2] G. Anderson and J. Rathke. Dynamic software update for message passing programs. In *APLAS*, volume 7705 of *LNCS*, pages 207–222. Springer, 2012.
- [3] J. W. Backus. The syntax and semantics of the proposed international algebraic language of the zurich ACM-GAMM conference. In *IFIP Congress*, pages 125–131, 1959.
- [4] L. Baresi, A. Marconi, M. Pistore, and A. Sirbu. Corrective Evolution of Adaptable Process Models. In *BMMDS/EMMSAD*, volume 147 of *LNBP*, pages 214–229. Springer, 2013.
- [5] S. Basu, T. Bultan, and M. Ouederni. Deciding choreography realizability. In *POPL*, pages 191–202. ACM, 2012.
- [6] G. Berry and M. Serrano. Hop and HipHop: Multitier Web Orchestration. In *ICDCIT*, volume 8337 of *LNCS*, pages 1–13. Springer, 2014.
- [7] M. Bravetti et al. Towards global and local types for adaptation. In *SEFM Workshops*, volume 8368 of *LNCS*, pages 3–14. Springer, 2013.
- [8] A. Bucchiarone, A. Marconi, C. A. Mezzina, M. Pistore, and H. Raik. On-the-fly adaptation of dynamic service-based systems: Incrementality, reduction and reuse. In *ICSOC*, volume 8274 of *LNCS*, pages 146–161. Springer, 2013.
- [9] S. Capecchi, E. Giachino, and N. Yoshida. Global Escape in Multiparty Sessions. In *Proc. of FSTTCS 2010*, volume 8 of *LIPICs*, pages 338–351. Schloss Dagstuhl, 2010.
- [10] M. Carbone, K. Honda, and N. Yoshida. Structured Communication-Centred Programming for Web Services. In *ESOP*, volume 4421 of *LNCS*, pages 2–17. Springer, 2007.
- [11] M. Carbone, K. Honda, and N. Yoshida. Structured Interactional Exceptions in Session Types. In *Proc. of CONCUR’08*, volume 5201 of *LNCS*, pages 402–417. Springer, 2008.
- [12] M. Carbone, K. Honda, and N. Yoshida. Structured communication-centered programming for web services. *ACM Trans. Program. Lang. Syst.*, 34(2):8, 2012.
- [13] M. Carbone and F. Montesi. Deadlock-Freedom-by-Design: Multiparty Asynchronous Global Programming. In *POPL*, pages 263–274. ACM, 2013.

- [14] G. Castagna, M. Dezani-Ciancaglini, and L. Padovani. On global types and multi-party session. *Logical Methods in Computer Science*, 8(1), 2012.
- [15] E. Cooper, S. Lindley, P. Wadler, and J. Yallop. Links: Web Programming Without Tiers. In *FMCO*, volume 4709 of *LNCS*, pages 266–296. Springer, 2006.
- [16] M. Coppo, M. Dezani-Ciancaglini, and B. Venneri. Self-adaptive multiparty sessions. *Service Oriented Computing and Applications*, 9(3-4):249–268, 2015.
- [17] G. Cugola, C. Ghezzi, and L. S. Pinto. DSOL: a declarative approach to self-adaptive service orchestrations. *Computing*, 94(7):579–617, 2012.
- [18] M. Dalla Preda, M. Gabbrielli, S. Giallorenzo, I. Lanese, and J. Mauro. Dynamic choreographies - safe runtime updates of distributed applications. In *COORDINATION*, volume 9037 of *LNCS*, pages 67–82. Springer, 2015.
- [19] M. Dalla Preda, S. Giallorenzo, I. Lanese, J. Mauro, and M. Gabbrielli. AIOCJ: A choreographic framework for safe adaptive distributed applications. In *SLE*, volume 8706 of *LNCS*, pages 161–170. Springer, 2014.
- [20] C. Di Giusto and J. A. Pérez. Disciplined structured communications with consistent runtime adaptation. In *SAC*, pages 1913–1918. ACM, 2013.
- [21] Eclipse website. <http://www.eclipse.org/>.
- [22] C. Ghezzi, M. Pradella, and G. Salvaneschi. An evaluation of the adaptation capabilities in programming languages. In *SEAMS*, pages 50–59. ACM, 2011.
- [23] R. Hirschfeld, P. Costanza, and O. Nierstrasz. Context-oriented Programming. *Journal of Object Technology*, 7(3):125–151, 2008.
- [24] K. Honda, N. Yoshida, and M. Carbone. Multiparty Asynchronous Session Types. In *POPL*, pages 273–284. ACM, 2008.
- [25] H. Hüttl et al. Foundations of session types and behavioural contracts. *ACM Computing Surveys*, 2016.
- [26] Jolie website. <http://www.jolie-lang.org/>.
- [27] I. Jureta, S. Faulkner, and P. Thiran. Dynamic requirements specification for adaptable and open service-oriented systems. In *ICSOC*, volume 4749 of *LNCS*, pages 270–282. Springer, 2007.
- [28] J. O. Kephart and D. M. Chess. The vision of autonomic computing. *Computer*, 36(1):41–50, 2003.
- [29] I. Lanese, A. Bucchiarone, and F. Montesi. A Framework for Rule-Based Dynamic Adaptation. In *TGC*, volume 6084 of *LNCS*, pages 284–300. Springer, 2010.
- [30] I. Lanese, C. Guidi, F. Montesi, and G. Zavattaro. Bridging the Gap between Interaction- and Process-Oriented Choreographies. In *SEFM*, pages 323–332. IEEE, 2008.
- [31] I. Lanese, F. Montesi, and G. Zavattaro. Amending choreographies. In *WWV*, volume 123, pages 34–48. EPTCS, 2013.
- [32] L. A. F. Leite et al. A systematic literature review of service choreography adaptation. *Service Oriented Computing and Applications*, 7(3):199–216, 2013.
- [33] C. Metz. The epic story of dropbox’s exodus from the amazon cloud empire. Wired, <http://www.wired.com/2016/03/epic-story-dropboxs-exodus-amazon-cloud-empire/>.
- [34] F. Montesi. Kickstarting choreographic programming. In *WS-FM*, pages 3–10. Springer, 2015.
- [35] F. Montesi, C. Guidi, and G. Zavattaro. Composing services with JOLIE. In *ECOWS*, pages 13–22. IEEE, 2007.
- [36] F. Montesi, C. Guidi, and G. Zavattaro. Service-oriented programming with Jolie. In *Web Services Foundations*, pages 81–107. Springer, 2014.
- [37] F. Montesi and N. Yoshida. Compositional choreographies. In *CONCUR*, volume 8052 of *LNCS*, pages 425–439. Springer, 2013.
- [38] I. Neamtiu and M. W. Hicks. Safe and timely updates to multi-threaded programs. In *PLDI*, pages 13–24. ACM, 2009.
- [39] M. Neubauer and P. Thiemann. From sequential programs to multi-tier applications by program transformation. In *POPL*, pages 221–232. ACM, 2005.
- [40] M. Neubauer and P. Thiemann. Placement inference for a client-server calculus. In *ICALP*, volume 5126 of *LNCS*, pages 75–86. Springer, 2008.
- [41] P. Nienaltowski. *Practical framework for contract-based concurrent object-oriented programming*. PhD thesis, ETH Zurich, 2007.
- [42] R. Pawlak et al. JAC: an aspect-based distributed dynamic framework. *Software: Practice and Experience*, 34(12):1119–1148, 2004.

- [43] S. Rinderle, A. Wombacher, and M. Reichert. Evolution of Process Choreographies in DYCHOR. In *OTM Conferences (1)*, volume 4275 of *LNCS*, pages 273–290. Springer, 2006.
- [44] Rust website. <http://www.rust-lang.org/>.
- [45] G. Salaün, T. Bultan, and N. Roohi. Realizability of choreographies using process algebra encodings. *IEEE T. Services Computing*, 5(3):290–304, 2012.
- [46] D. Sangiorgi and D. Walker. *The pi-calculus: a Theory of Mobile Processes*. Cambridge university press, 2003.
- [47] Scribble website. <http://www.jboss.org/scribble>.
- [48] A. Wombacher. Alignment of choreography changes in BPEL processes. In *IEEE SCC*, pages 1–8. IEEE, 2009.
- [49] Xtext website. <http://www.eclipse.org/Xtext/>.
- [50] Z. Yang, B. H. C. Cheng, R. E. K. Stirewalt, J. Sowell, S. M. Sadjadi, and P. K. McKinley. An aspect-oriented approach to dynamic adaptation. In *WOSS*, pages 85–92. ACM, 2002.
- [51] S. Yegulalp. Mozilla’s Rust-based Servo browser engine inches forward. InfoWorld, <http://www.infoworld.com/article/2905688/applications/mozillas-rust-based-servo-browser-engine-inches-forward.html>.

APPENDIX A. PROOF OF THEOREM 6.3

In order to prove the bound on the complexity of the connectedness check we use the lemma below, showing that the checks to verify the connectedness for a single sequence operator can be performed in linear time on the size of the sets generated by `transl` and `transF`.

Lemma A.1. *Given S, S' sets of multisets of two elements, checking if $\forall s \in S . \forall s' \in S' . s \cap s' \neq \emptyset$ can be done in $O(n)$ steps, where n is the maximum of $|S|$ and $|S'|$.*

Proof. Without loss of generality, we can assume that $|S| \leq |S'|$. If $|S| \leq 9$ then the check can be performed in $O(n)$ by comparing all the elements in S with all the elements in S' . If $|S| > 9$ then at least 4 distinct elements appear in the multisets in S since the maximum number of multisets with cardinality 2 obtained by 3 distinct elements is 9. In this case the following cases cover all the possibilities:

- there exist distinct elements a, b, c, d s.t. $\{a, b\}$, $\{a, c\}$, and $\{a, d\}$ belong to S . In this case for the check to succeed all the multisets in S' must contain a , otherwise the intersection of the multiset not containing a with one among the multisets $\{a, b\}$, $\{a, c\}$, and $\{a, d\}$ is empty. Similarly, since $|S'| > 9$, for the check to succeed all the multisets in S must contain a . Hence, if $\{a, b\}$, $\{a, c\}$, and $\{a, d\}$ belong to S then the check succeeds iff a belongs to all the multisets in S and in S' .
- there exist distinct elements a, b, c, d s.t. $\{a, b\}$ and $\{c, d\}$ belong to S . In this case the check succeeds only if S' is a subset of $\{\{a, c\}, \{a, d\}, \{b, c\}, \{b, d\}\}$. Since $|S'| > 9$ the check can never succeed.
- there exist distinct elements a, b, c s.t. $\{a, a\}$ and $\{b, c\}$ belong to S . In this case the check succeeds only if S' is a subset of $\{\{a, b\}, \{a, c\}\}$. Since $|S'| > 9$ the check can never succeed.
- there exist distinct elements a, b s.t. $\{a, a\}$ and $\{b, b\}$ belong to S . In this case the check succeeds only if S' is a subset of $\{\{a, b\}\}$. Since $|S'| > 9$ the check can never succeed.

Summarising, if $|S| > 9$ the check can succeed iff all the multisets in S and in S' share a common element. The existence of such an element can be verified in time $O(n)$. \square

Theorem 6.3 (Connectedness-check complexity).

The connectedness of a DIOC process \mathcal{I} can be checked in time $O(n^2 \log(n))$, where n is the number of nodes in the abstract syntax tree of \mathcal{I} .

Proof. To check the connectedness of \mathcal{I} we first compute the values of the functions `transl` and `transF` for each node of the abstract syntax tree (AST). We then check for each sequence operator whether connectedness holds.

The functions `transl` and `transF` associate to each node a set of pairs of roles. Assuming an implementation of the data set structure based on balanced trees (with pointers), `transl` and `transF` can be computed in constant time for interactions, assignments, $\mathbf{1}$, $\mathbf{0}$, and sequence constructs. For while and scope constructs computing `transF`(\mathcal{I}') requires the creation of balanced trees having an element for every role of \mathcal{I}' . Since the roles are $O(n)$, `transF`(\mathcal{I}') can be computed in $O(n \log(n))$. For parallel and if constructs a union of sets is needed. The union costs $O(n \log(n))$ since each set generated by `transl` and `transF` contains at maximum n elements.

Since the AST contains n nodes, the computation of the sets generated by `transl` and `transF` can be performed in $O(n^2 \log(n))$.

To check connectedness we have to verify that for each node $\mathcal{I}'; \mathcal{I}''$ of the AST $\forall \mathbf{R}_1 \rightarrow \mathbf{R}_2 \in \text{transF}(\mathcal{I}'), \forall \mathbf{S}_1 \rightarrow \mathbf{S}_2 \in \text{transl}(\mathcal{I}'') . \{\mathbf{R}_1, \mathbf{R}_2\} \cap \{\mathbf{S}_1, \mathbf{S}_2\} \neq \emptyset$. Since `transF`(\mathcal{I}') and

$\text{transl}(\mathcal{I}'')$ have $O(n)$ elements, thanks to Lemma A.1, checking if $\mathcal{I}'; \mathcal{I}''$ is connected costs $O(n)$. Since in the AST there are less than n sequence operators, checking the connectedness on the whole AST costs $O(n^2)$.

The complexity of checking the connectedness of the entire AST is therefore limited by the cost of computing functions transl and transF and of checking the connectedness. All these activities have a complexity of $O(n^2 \log(n))$. \square

APPENDIX B. PROOFS OF SECTION 7.1

Lemma 7.9 (Distinctness of Global Indexes). *Given a well-annotated DIOC \mathcal{I} , a global state Σ , and a set of updates \mathbf{I} , if $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle \xrightarrow{\eta_1} \dots \xrightarrow{\eta_n} \langle \Sigma', \mathbf{I}', \mathcal{I}' \rangle$ then all global indexes in \mathcal{I}' are distinct.*

Proof. The proof is by induction on the number n of transitions, using a stronger inductive hypothesis: indexes are distinct but, possibly, inside DIOC subterms of the form $\mathcal{I}; \mathbf{i}; \mathbf{while} \ b@R \ \{\mathcal{I}'\}$. In this last case, the same index can occur both in \mathcal{I} and in \mathcal{I}' , attached to constructs with different global indexes. The statement of the Lemma follows directly: first, distinct indexes imply distinct global indexes; second, global indexes of \mathcal{I} and of \mathcal{I}' are distinct, since \mathcal{I}' is inside the while loop whilst \mathcal{I} is not.

In the base case ($n = 0$), thanks to well annotatedness, indexes are always distinct. The inductive case follows directly by induction for transitions with label \surd . Otherwise, we have a case analysis on the only axiom which derives a transition with label different from \surd .

The only difficult cases are $\llbracket^{\text{DIOC}} \text{WHILE-UNFOLD} \rrbracket$ and $\llbracket^{\text{DIOC}} \text{UP} \rrbracket$.

In the case of Rule $\llbracket^{\text{DIOC}} \text{WHILE-UNFOLD} \rrbracket$, note that the while is enabled, hence it cannot be part of a term of the form $\mathcal{I}; \mathbf{i}; \mathbf{while} \ b@R \ \{\mathcal{I}'\}$. Hence, indexes of the body of the while loop do not occur elsewhere. As a consequence, after the transition no clashes are possible with indexes in the context. Note also that indexes of the body of the loop are duplicated, but the resulting term has the form $\mathcal{I}; \mathbf{i}; \mathbf{while} \ b@R \ \{\mathcal{I}'\}$, thus global indexes are distinct by construction.

The case $\llbracket^{\text{DIOC}} \text{UP} \rrbracket$ follows thanks to the condition $\text{freshIndexes}(\mathcal{I}')$. \square

Lemma 7.15. *Given a well-annotated connected DIOC process \mathcal{I} and for each state Σ the DPOC network $\text{proj}(\mathcal{I}, \Sigma)$ is such that:*

- (1) $\text{events}(\mathcal{I}) \subseteq \text{events}(\text{proj}(\mathcal{I}, \Sigma))$;
- (2) $\forall \varepsilon_1, \varepsilon_2 \in \text{events}(\mathcal{I}). \varepsilon_1 \leq_{\text{DIOC}} \varepsilon_2 \Rightarrow \varepsilon_1 \leq_{\text{DPOC}} \varepsilon_2 \vee \varepsilon_1 \leq_{\text{DPOC}} \bar{\varepsilon}_2$

Proof.

- (1) By definition of projection.
- (2) Let $\varepsilon_1 \leq_{\text{DIOC}} \varepsilon_2$. We have a case analysis on the condition used to derive the dependency.

Sequentiality: Consider $\mathcal{I} = \mathcal{I}'; \mathcal{I}''$. If events are in the same role the implication follows from the sequentiality of the \leq_{DPOC} .

Let us show that there exists an event ε'' in an initial interaction of \mathcal{I}'' such that either $\varepsilon'' \leq_{\text{DPOC}} \varepsilon_2$ or $\varepsilon'' \leq_{\text{DPOC}} \bar{\varepsilon}_2$. The proof is by induction on the structure of \mathcal{I}'' . The only difficult case is sequential composition. Assume $\mathcal{I}'' = \mathcal{I}_1; \mathcal{I}_2$. If $\varepsilon_2 \in \text{events}(\mathcal{I}_1)$ the thesis follows from inductive hypothesis. If $\varepsilon_2 \in \text{events}(\mathcal{I}_2)$ then by induction there exists an event ε_3 in an initial interaction of \mathcal{I}_2 such that $\varepsilon_3 \leq_{\text{DPOC}} \varepsilon_2$ or $\varepsilon_3 \leq_{\text{DPOC}} \bar{\varepsilon}_2$. By synchronisation (Definition 7.14) we have that $\bar{\varepsilon}_3 \leq_{\text{DPOC}} \varepsilon_2$ or $\bar{\varepsilon}_3 \leq_{\text{DPOC}} \bar{\varepsilon}_2$. By connectedness we have that ε_3 or $\bar{\varepsilon}_3$ are in the same role of an

event ε_4 in \mathcal{I}' . By sequentiality (Definition 7.14) we have that $\varepsilon_4 \leq_{DPOC} \varepsilon_3$ or $\varepsilon_4 \leq_{DPOC} \bar{\varepsilon}_3$. By synchronisation we have that $\bar{\varepsilon}_4 \leq_{DPOC} \varepsilon_3$ or $\bar{\varepsilon}_4 \leq_{DPOC} \bar{\varepsilon}_3$. The thesis follows from the inductive hypothesis on ε_4 and by transitivity of \leq_{DPOC} .

Let us also show that there exists a final event $\varepsilon''' \in \text{events}(\mathcal{I}')$ such that $\varepsilon_1 \leq_{DPOC} \varepsilon'''$ or $\varepsilon_1 \leq_{DPOC} \bar{\varepsilon}'''$. The proof is by induction on the structure of \mathcal{I}' . The only difficult case is sequential composition. Assume $\mathcal{I}' = \mathcal{I}_1; \mathcal{I}_2$. If $\varepsilon_1 \in \text{events}(\mathcal{I}_2)$ the thesis follows from inductive hypothesis. If $\varepsilon_1 \in \text{events}(\mathcal{I}_1)$ then the proof is similar to the one above, finding a final event in \mathcal{I}_1 and applying sequentiality, synchronisation, and transitivity.

The thesis follows from the two results above again by sequentiality, synchronisation, and transitivity.

Scope: it means that either (a) $\varepsilon_1 = \uparrow_\xi$ and ε_2 is an event in the scope or (b) $\varepsilon_1 = \uparrow_\xi$ and $\varepsilon_2 = \downarrow_\xi$, or (c) ε_1 is an event in the scope and $\varepsilon_2 = \downarrow_\xi$. We consider case (a) since case (c) is analogous and case (b) follows by transitivity. If ε_2 is in the coordinator then the thesis follows easily. Otherwise it follows thanks to the auxiliary synchronisations with a reasoning similar to the one for sequentiality.

Synchronisation: it means that ε_1 is a sending event and ε_2 is the corresponding receiving event, namely $\varepsilon_1 = \bar{\varepsilon}_2$. Thus, since $\varepsilon_2 \leq_{DPOC} \varepsilon_2$ then $\bar{\varepsilon}_2 \leq_{DPOC} \varepsilon_2$.

If: it means that ε_1 is the evaluation of the guard and ε_2 is an event in one of the two branches. Thus, if ε_2 is in the coordinator then the thesis follows easily. Otherwise it follows thanks to the auxiliary synchronisations with a reasoning similar to the one for sequentiality.

While: it means that ε_1 is the evaluation of the guard and ε_2 is in the body of the while loop. Thus, if ε_2 is in the coordinator then the thesis follows easily. Otherwise it follows thanks to the auxiliary synchronisations with a reasoning similar to the one for sequentiality. \square

Lemma 7.20. *Let \mathcal{I} be a well-annotated connected DIOC process and Σ a state. Then the projection $\mathcal{N} = \text{proj}(\mathcal{I}, \Sigma)$ is a well-annotated DPOC network with respect to \leq_{DPOC} .*

Proof. We have to prove that $\text{proj}(\mathcal{I}, \Sigma)$ satisfies the conditions of Definition 7.17 of well-annotated DPOC:

- C1: For each global index ξ there are at most two communication events on programmer-specified operations with global index ξ and, in this case, they are matching events. The condition follows by the definition of the projection function, observing that in well-annotated DIOCs, each interaction has its own index, and different indexes are mapped to different global indexes.
- C2: Only events which are minimal according to \leq_{DPOC} may correspond to enabled transitions. This condition follows from Lemma 7.19.
- C3: For each pair of non-conflicting sending events $[f_\xi]_{\mathbf{R}}$ and $[f_{\xi'}]_{\mathbf{R}}$ on the same operation $o^?$ and with the same target \mathbf{R}' such that $\xi \neq \xi'$ we have $[f_\xi]_{\mathbf{R}} \leq_{DPOC} [f_{\xi'}]_{\mathbf{R}}$ or $[f_{\xi'}]_{\mathbf{R}} \leq_{DPOC} [f_\xi]_{\mathbf{R}}$. Note that the two events are in the same role \mathbf{R} , thus without loss of generality we can assume that there exist two processes P, P' such that $[f_\xi]_{\mathbf{R}} \in \text{events}(P)$ and $[f_{\xi'}]_{\mathbf{R}} \in \text{events}(P')$ and there is a subprocess of \mathcal{N} of one of the following forms:
 - $P; P'$: the thesis follows by sequentiality (Definition 7.14);
 - $P|P'$: this case can never happen for the reasons below. For events on programmer-specified operations this follows by the definition of projection, since the prefixes of the names of operations are different. For events on auxiliary operations originated

by the same construct this follows since all the targets are different. For events on auxiliary operations originated by different constructs this follows since the prefixes of the names of the operations are different.

- **if $b \{P\}$ else $\{P'\}$** : this case can never happen since the events are non-conflicting (Definition 7.16).
- C4: Similar to the previous case, with receiving events instead of sending events.
- C5: If ε is an event inside a scope with global index ξ then its matching events $\bar{\varepsilon}$ (if they exist) are inside a scope with the same global index. This case holds by definition of the projection function.
- C6: If two events have the same index but different global indexes then one of them, let us call it ε_1 , is inside the body of a while loop with global index ξ_1 and the other, ε_2 , is not. Furthermore, $\varepsilon_2 \leq_{DPOC} \varepsilon_{\xi_1}$ where ε_{ξ_1} is the guarding while-event of the while loop with global index ξ_1 . By definition of well-annotated DIOC and of projection the only case where there are two events with the same index but different global indexes is for the auxiliary communications in the projection of the while construct, where the conditions hold by construction. \square