



HAL
open science

Real-World Choreographies

Saverio Giallorenzo

► **To cite this version:**

Saverio Giallorenzo. Real-World Choreographies. Distributed, Parallel, and Cluster Computing [cs.DC]. Università degli studi di Bologna, 2016. English. NNT : . tel-01336757

HAL Id: tel-01336757

<https://inria.hal.science/tel-01336757>

Submitted on 27 Jun 2016

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.

Alma Mater Studiorum — Università di Bologna

DOTTORATO DI RICERCA IN INFORMATICA
Ciclo: XXVIII

Settore Concorsuale di Afferenza: 01/B1
Settore Scientifico Disciplinare: INF/01

Real-World Choreographies

Presentata da: Saverio Giallorenzo

Coordinatore Dottorato:
Paolo Ciaccia

Relatore:
Maurizio Gabbrielli

Correlatore:
Fabrizio Montesi

Esame finale anno 2016

Abstract

Since the early days of the Internet, distributed software applications have become one of the leading forces behind the development and economic growth of our society. Nonetheless, the practice of programming distributed systems is one of the most error-prone. Developers strive to correctly implement separate components that, put together, enact an agreed protocol. If one component fails to follow such protocol, it could lead to system blocks or misbehaviours. Ensuring that all components correctly follow the intended protocol is very difficult due to the inherent non-determinism of distributed programs.

This led practitioners and theoretical researchers to explore new tools to assist the development of distributed systems. Choreographies are one of these tools. They have been introduced to describe from a global viewpoint the exchange of messages among the components of a distributed system. Moreover, since they describe atomic communications (not split into I/Os), they are free from deadlocks and race conditions by design. Recent theoretical results proved that it is possible to define proper Endpoint Projection (EPP) functions to compile choreographic specifications into their single components. Since EPPs are behaviour preserving, projected systems also enjoy freedom from deadlocks and races by construction. Some of these results have been implemented, however much work has to be done to make choreographies a suitable tool for real-world programming.

Aim of this PhD is to formalise non-trivial features of distributed systems with choreographies and to translate our theoretical results into the practice of implemented systems. To this purpose, we provide two main contributions.

The first contribution tackles one of the most challenging features of distributed development: *programming correct and consistent runtime updates of distributed systems*. There is no affirmed technology for structuring runtime updates of distributed applications. Moreover, the non-determinism of distributed systems easily leads to partial applications of updates and to inconsistent systems. Our solution is a theoretical model of dynamic choreographies, called DIOC. DIOC provides a clear definition of which components and behaviours can be updated. We prove that systems compiled from a DIOC definition are correct and consistent after any update. Finally, we refine our theoretical model with constructs for a finer control over updates. On this refinement, we develop a framework for programming adaptable distributed systems, called AIOCJ.

The second contribution covers one of the main issues of implementing theoretical results on choreographies: *formalising the compilation from choreographies to executable programs*. There is a sensible departure between choreographic frameworks like Chor (the first on this paradigm) and AIOCJ and their theoretical models: their theories abstract communications with synchronisation on names (à la CCS/ π -calculus) yet they compile to Jolie programs, an executable language that uses *correlation* — a renowned technology of Service-Oriented Computing — for message routing. This discontinuity breaks the chain of proven correctness from choreographies to implemented systems. Our solution is a theory of Applied Choreographies (AC) that models correlation-based message passing. With AC, we formalise the key theoretical problems and the guiding principles that developers should follow to obtain correct implementations. Finally, we prove our approach by defining a correct compiler from AC to the calculus behind the Jolie language.

Acknowledgements

If I have seen further, it is by standing on the
shoulders of giants.

Sir Isaac Newton

A PhD is a formidable endeavour, not only academically speaking but also at the personal level. It gave to me the opportunity to face my limits, to experiment failure, and, in many senses, to grow up. Thankfully, I was very lucky to have the best people around me. These are my acknowledgements to them all.

I thank my parents, Patrizia and Vito, for supporting all my (mostly eccentric) projects and for giving me the freedom to be wrong. I owe to them my creativity and love for Science. They taught me patience and perseverance, principles that “saved” me countless times and that I like to think constitute some of my best qualities.

In these years, I also contracted an infinite debt of gratitude with my girlfriend Daniela. She has been my most passionate supporter and best friend. Times and times she encouraged me while I lingered in self-doubt and comforted me when I complained about broken formalisations, buggy programs, and rejected papers. I owe to her all the sacrifices she made for me: the days and nights passed alone while I was working for a deadline, the missed vacations, and the long trips to Denmark. I deeply thank her for all the love she gave to me and for accepting a simple “Ti amo” (“I love you”) in exchange for all of that.

I dedicate this work to Daniela, Ma, and Pa, for their unbreakable belief in me.

Thanks also to Roberta and Andrea for having accepted this chaotic guy in their family and to Giorgia for all the joyful games and talks about dinosaurs (!).

In my PhD I had the good fortune to have many teachers, whom I thank.

Thanks to my two co-advisors: Maurizio Gabbrielli and Fabrizio Montesi.

I thank Maurizio for having believed in me and that a poor “Scenziato di Internet” (graduate in IT and Management) could pull off some basic research. I owe to Maurizio many thanks: for having set up a nice research team at the beginning of my PhD, for the freedom he has always given to me, for having pushed me beyond my limits, and for cheering me up during my saddest moments.

I owe a lot also to Fabrizio. He is the crazy guy who coded the most part of the Jolie language, which brought me in 2012 to work with him in Copenhagen (again, thanks to Maurizio) while he was working on his PhD on “Choreographic Programming” — of which this thesis can be considered a continuation. Fabrizio dedicated a lot of his time to teach me academic rigour and to seek clarity in the definition of problems before venturing in their resolution. I really enjoyed all the time we spent together trying to bring elegance and balance to our formalisations, Jolie, and choreographies¹.

Thanks to Ivan Lanese for the many times he proved me (and, above all, calmly explained why I was) wrong. I hope I have been able to steal a little of Ivan’s mathematical genius and cleverness during our restless sessions on our theoretical model and its proofs.

¹and, by extension, to the Force.

IV

Thanks to Jacopo — O Captain, My Captain! — Mauro. I had the good fortune to have a big brother like him in the Underground Laboratory in Bologna. Thanks for having given to me stability when I desperately needed some, for all the advices and chats, and for begin a very good friend to me.

Thanks to Claudio Guidi for having introduced me into the academic world, “Galeotto fu il linguaggio (Jolie) e chi lo scrisse²” and for all the inspiring conversations we had on languages, abstractions, and “visions”.

Thanks to Marco Carbone, my (now) academic grandfather who (back then) hosted me for my first stay-abroad in Copenhagen. It was a pleasure to know him as a person and I thank him for contributing in making all this possible.

Thanks to Davide Sangiorgi for all the inestimable support he gave to me before and during my PhD, both academically and personally. I am also really grateful to Davide for having let me teach Jolie in his course on Operating Systems (yay!).

Thanks to Gianluigi Zavattaro for the interesting discussions and alternative visions on choreographies and for having given to me the opportunity to present Jolie at the Bertinoro International Spring School, that was really great!

I also want to thank my reviewers, Ilaria Castellani and Thomas Hildebrandt, who not only took the duty of refereeing this thesis but provided encouraging and stimulating comments to improve its contents.

I thank all the members of the INRIA Focus team in Bologna for all the inspiring talks. Thanks to the people in the research group of Programming, Logic and, Semantics at the IT University of Copenhagen and in the department of Mathematics and Computer Science (IMADA) at the Southern Denmark University. I also thank Maja, Borna, Marco, and Luís for the joyful time in Odense.

Finally I owe a great deal of gratitude to all my fellow colleagues and PhD students for having shared with me the joys, the grumbles, the aperitivi, and the awesome discussions in and out of the (Hardcore) Underground Laboratory.

Many thanks to Valeria Vignudelli for the endless hours spent to teach me proof techniques, for all the laughs, and the cheerful pats on the shoulder that I needed.

Thanks to Jean-Marie Madiot for sharing with me many serious (and definitely not serious) talks, many pat-pats on the shoulder, and for that “You don’t know Hacker News?!”.

Thanks to Francesco Gavazzo, for the desperate mission to explain to me Category Theory and for all the fruitful chats on the philosophy of dark knights and the super men.

Thanks also to Mila, Ornela, Paolo, Gustavo, Giulio, Roberto, Abel and Yiyi (¡gracias por el ron!), Stefano, Vincenzo, Giacomo, (the other) Vincenzo, Federico, and Luca for all the nice time spent together.

Thank you all, my giants.

²Accurséd was the language (Jolie) and those who wrote it.

List of Publications

- Dalla Preda, Mila, Maurizio Gabbrielli, Saverio Giallorenzo, Ivan Lanese, and Jacopo Mauro. **Dynamic Choreographies: From Theory to Practice**. Submitted.
- Maurizio Gabbrielli, Saverio Giallorenzo, Claudio Guidi, Fabrizio Montesi, Jacopo Mauro. *Self-configuring Microservices*. In Theory and Practice of Formal Methods, pp. 194–210. Springer International Publishing, 2016.
- Maurizio Gabbrielli, Saverio Giallorenzo, and Fabrizio Montesi. **Applied Choreographies**. Submitted.
- Dalla Preda, Mila, Maurizio Gabbrielli, Saverio Giallorenzo, Ivan Lanese, and Jacopo Mauro. **Dynamic Choreographies — Safe Runtime Update of Distributed Applications**. In Coordination Models and Languages, pp. 67–82. Springer International Publishing, 2015.
- Dalla Preda, Mila, Maurizio Gabbrielli, Saverio Giallorenzo, Ivan Lanese, and Jacopo Mauro. *Developing correct, distributed, adaptive software*. Science of Computer Programming 97, pp. 1–46. Elsevier, 2015.
- Dalla Preda, Mila, Saverio Giallorenzo, Ivan Lanese, Jacopo Mauro, and Maurizio Gabbrielli. **AIOCJ: A choreographic framework for safe adaptive distributed applications**. In Software Language Engineering, pp. 161–170. Springer International Publishing, 2014.
- Maurizio Gabbrielli, Saverio Giallorenzo, and Fabrizio Montesi. *Service-Oriented Architectures: from design to production exploiting workflow patterns*. In Distributed Computing and Artificial Intelligence, 11th International Conference, pp. 131–139. Springer International Publishing, 2014.
- Claudio Guidi, Saverio Giallorenzo, and Maurizio Gabbrielli. *Towards a Composition-based APIaaS Layer*, In CLOSER 2014 - Proceedings of the 4th International Conference on Cloud Computing and Services Science, pp. 425–432. SciTePress, 2014.

Items with titles in **bold** are part of this dissertation.

Contents

Abstract	I
Acknowledgements	III
List of Publications	V
Table of Contents	VI
List of Figures	XI
1 Introduction to the dissertation	1
1.1 Problem Description	1
1.2 Aim and Thesis Statement	1
1.3 Contributions	3
1.4 Structure of the Dissertation	3
I Introduction and Background	7
2 Introduction to Choreographies	9
2.1 Distributed Programming, in brief	9
2.1.1 Problems of Distributed Programming	10
2.2 The Service-Oriented Approach	12
2.2.1 Breaking down complexity	14
2.3 Orchestration	15
2.4 Choreography	16

2.5	Programming with Choreographies	18
2.6	Contributions	20
II	Adaptable Choreographies	23
3	Runtime Software Update and Adaptation	25
3.1	Linguistic Constructs for Adaptation	25
3.1.1	Reflection and Metaprogramming	27
3.1.2	Aspect- and Context-Oriented Programming	28
3.1.3	Adaptation in Process-Aware Information Systems	29
3.2	Safe Update of Distributed Systems	32
4	Dynamic Choreographies	33
4.1	Introduction	33
4.2	Dynamic Interaction-Oriented Choreographies	36
4.2.1	DIOC Syntax	36
4.2.2	Annotated DIOCs and their Semantics	37
4.3	Dynamic Process-Oriented Choreographies	42
4.3.1	DPOC syntax	42
4.3.2	DPOC semantics	44
4.4	Projection Function	49
4.5	Running Example: Projection and Execution	53
4.5.1	Projection	53
4.5.2	Runtime Execution	53
4.6	Connected DIOCs	57
4.7	Correctness	60
5	Adaptable Interaction-Oriented Choreographies in Jolie	65
5.1	Introduction	65
5.1.1	DIOC Language Extensions in AIOCJ	66
5.2	AIOCJ Practice	69
5.3	Implementation	70
5.4	Validation	72
III	Applied Choreographies	77
6	Applied Choreographies	79
6.1	Introduction	79
6.2	Applied Choreography Language	82

6.2.1	Syntax	82
6.2.2	Semantics	84
6.2.3	An example	90
6.3	Typing	92
6.3.1	Types and Type Projection	92
6.3.2	Type checking	93
6.3.3	Runtime Typing	96
6.3.4	Properties	101
6.4	Endpoint Projection	102
6.4.1	Projection Example	104
6.4.2	Properties	105
6.5	Dynamic Correlation Calculus Language	106
6.6	Compiler from AC to DCC and Properties	107
6.6.1	Compiler	107
6.6.2	Example of Compilation	111
6.6.3	Properties	113
IV	Conclusion	115
7	Conclusion	117
7.1	Adaptable Choreographies	117
7.1.1	Adaptable Choreographies and their degree of flexibility	119
7.2	Applied Choreographies	119
V	Appendix	135
A	Adaptable Choreographies: Proofs	137
A.1	Proof of Theorem 1	137
A.2	Proof of Theorem 2	138
B	Adaptable Choreographies: Test Code	161
B.1	Code used for validation	161
B.1.1	Pipe and fork-join code	161
B.1.2	AIO CJ programs used for benchmarking primitives	164
C	Adaptable Choreographies: Models of Adaptation	167
C.1	A distributed adaptive document system	167
C.1.1	Pointcuts	170
C.1.2	Dynamic wrappers	171

C.2	ContextChat	172
C.2.1	Online/Offline switch	177
D	Applied Choreographies: Additional Material	179
D.1	Applied Choreographies	179
D.2	Typing	180
D.3	Endpoint Projection	182
D.4	Dynamic Correlation Calculus	185
E	Applied Choreographies: Proofs	187
E.1	Proofs of Subject Reduction and Session Fidelity	187
E.1.1	Local and Typing Environment Subtyping	189
E.2	Proof of Deadlock Freedom	203
E.3	Proof of Endpoint Projection	204
E.3.1	Minimal Typing	205
E.3.2	Typing Projection	205
E.3.3	Proof of Theorem 6	207
E.3.4	EPP Theorem	214
E.3.5	Proof of Theorem 7	217
E.4	Proof of Compilation	228
E.4.1	Proof of Theorem 8	228

List of Figures

2.1	Sequence chart of the ATM withdrawal example.	11
2.2	Example of orchestration.	15
2.3	Upper part: example of interaction-oriented choreography. Lower part: from left to right, the process-oriented choreographies of the ATM, the Bank, and the Card Issuer.	18
4.1	Fidelity Card Update.	34
4.2	DIOC process for Purchase Scenario.	38
4.3	Annotated DIOC system semantics.	40
4.4	Auxiliary function roles.	41
4.5	DPOC role semantics. Computation rules. (Update rules in Figure 4.6)	45
4.6	DPOC role semantics. Update rules. (Computation rules in Figure 4.5)	46
4.7	DPOC system semantics.	48
4.8	process-projection function π	51
4.9	Seller DPOC Process.	54
4.10	Bank DPOC Process.	54
4.11	Buyer DPOC Process.	54
4.12	Auxiliary functions transl and transF	59
5.1	An AIOC program (upper part) and an applicable adaptation rule (lower part).	70
5.2	Representation of the AIOCJ framework — Projection and execution of the example in Figure 5.1.	71
5.3	Check for connectedness: sequence.	72

5.4	Times of execution of the pipe (left) and the fork-join (right) scenarios	74
6.1	Overall methodology of Applied Choreographies.	81
6.2	Choreography Calculus - Syntax	84
6.3	Choreography calculus, semantics.	87
6.4	Semantics of effects on D	88
6.5	Choreography Example	90
6.6	Example of structure of Session Descriptor	92
6.7	Choreography Calculus - Global Type Projection	94
6.8	Choreography Calculus - Typing Rules (selected)	95
6.9	Protocols example	100
6.10	EPPs of $C_c \mid C_s$. Projection on a, dm (excepts), and l.	104
6.11	Dynamic Correlation Calculus, syntax	106
6.12	Correlation Calculus, semantics	108
6.13	Compiler from AC to DCC.	110
6.14	Compilation of $C = C_c \mid C_s$, behaviours of c (excepts), a, and l.	112
D.1	Choreography Calculus, swap relation \simeq_C	179
D.2	Choreography Calculus - Typing Rules	180
D.3	Global types, Swap Relation \simeq_G	181
D.4	Choreography Calculus - Buffer Type Projection	181
D.5	Choreography Calculus — Annotation operator	182
D.6	Choreography Calculus, process projection	183
D.7	Merging Function	184
D.8	Service Grouping	184
D.9	Correlation Calculus, structural congruence	185
E.1	Choreography calculus, annotated semantics.	188
E.2	Choreography Calculus — Minimal typing rules	206
E.3	Correlation Calculus, annotated semantics	243

Introduction to the dissertation

Here we present the scientific context of our research and the problems that we want to address in this dissertation. The purpose of this chapter is to provide an overview of the structure of this dissertation. We give a brief introduction of the concepts that are extensively presented in the next chapters.

1.1 Problem Description

Since the early days of the Internet [1], distributed software applications have become one of the leading forces behind the development and economic growth of our society. A distributed system is a software application whose behaviour emerges from the interaction of several components (programs) that run in parallel and rely on message passing to communicate and coordinate their actions [2]. News systems, messaging, governance, healthcare, and transportation are just some of the contexts recently revolutionised by distributed applications.

Nonetheless, the practice of programming distributed systems is one of the most error-prone. Developers strive to correctly implement separate components that, put together, enact an agreed global behaviour (protocol). If one or more of said components fail to follow their common protocol, the distributed system can block or misbehave [3, 4]. Unfortunately, ensuring that all components correctly *play* their *part* in the interaction — i.e., that they follow their intended protocol — is very difficult due to the inherent non-determinism of several distributed programs running in parallel.

1.2 Aim and Thesis Statement

Here, we consider “safe” a distributed system in which its components correctly enact the intended global behaviour of the system, i.e., a safe distributed system is one free from defective behaviours due to faulty communications between its

components. On this definition, we state that the aim of this dissertation is

to advance the scientific knowledge on programming safe distributed systems.

Having proper language abstractions to express and analyse concurrent and distributed software has proven to be one of the most effective techniques to harness their complexity and to provide guarantees of their safety [5, 6]. Hence, we pursue our aim investigating the linguistic abstractions that can support the development of safe distributed applications.

Choreographies [7] are one of such linguistic abstractions. Introduced for the definition of protocols of interactions [8, 9], they describe the communications between the components of a distributed system from a global point of view. Distinguishing feature of choreographic languages is that they describe communications as atomic entities, i.e., not split into inputs and outputs. Since it is impossible to write mismatched I/Os, distributed systems described in a choreographic language are free from deadlocks and race conditions by design. However, a choreographic description of a distributed system is not directly executable. Recent theoretical results proved [10, 11, 12, 13, 14] that, given a proper choreographic description of a distributed system, it is possible to correctly compile it into its single components, which are also deadlock- and race-free by construction. This makes choreographies a very interesting linguistic abstraction for the development of real-world safe distributed systems. Although promising, the present choreography theories and their implementation prototypes still cover a very limited set of properties of distributed systems. Making choreographies a suitable tool for real-world programming [15] still requires a lot of work. Extensions of the theoretical models are needed to capture desirable requirements of distributed applications, like safety-preserving integrations with existing executable languages (e.g., the possibility to write Java or Scala code directly in the choreography), exception handling [16], and runtime updates [17]. Finally, also the correct implementation of the theoretical models is extremely important to the aim of bringing into the real-world the desirable properties of choreographies. Indeed, providing a formal treatment of such compilation is of paramount importance to guarantee the preservation of the properties of correctness and safety of choreographies down to compiled executable programs.

The thesis of this dissertation affirms that choreographies are a suitable linguistic abstraction that can be successfully extended to support the development of real-world safe distributed systems.

Thesis Statement

Choreographies can be used to program real-world safe distributed systems.

1.3 Contributions

We prove our thesis presenting two main contributions.

The first contribution tackles one of the most challenging features of distributed programming: guaranteeing the correct update of the behaviour of distributed systems at runtime. Indeed, beside the lack of proper abstractions to structure distributed updates, the non-determinism of distributed systems easily lead to partial applications of updates, leaving systems in an inconsistent state. Our solution is a theory of Dynamic Choreographies, called DIOCs. DIOCs can update at runtime in a coherent way and provide a clear definition of which components and behaviours can be changed. We prove that systems compiled from a DIOC definition are correct and consistent after any update. Our theory of dynamic choreographies is deliberately general and can be refined to model specific policies of updates. We extend the DIOC language with constructs for providing a fine control on the updates and we implement AIOCI, a framework for the development of correct, adaptable distributed systems.

The second contribution covers one of the main issues of implementing theoretical results on choreographies: providing a formal compilation from a choreographic language to an executable language. Our solution is to develop a theory of Applied Choreographies (AC) that models correlation-based message passing. We pinpoint the key theoretical problems and formalise the guiding principles that developers should follow to obtain correct implementations. Finally, we prove our approach by defining a correct compiler from AC to the foundational calculus of the executable language Jolie.

1.4 Structure of the Dissertation

The remainder of this dissertation is structured in three main parts in which we introduce the reasons and issues that motivate our investigations and report our results. All presented results have been produced during the course of the PhD studies and are here presented in an extended form.

- **Part I — Introduction to the Dissertation and Background**

It comprises this Chapter about the aim, thesis, and structure of this dissertation. In the next Chapter of this Part we present the features and issues of concurrent, distributed programming, the scientific context of our research, and the state of the art, on which we build our results. This introduction is intended to provide the reader with the basic knowledge to interpret the results presented in the next Parts.

- **Part II — Adaptable Choreographies**

We tackle the problem of programming distributed applications that update their behaviour at runtime, adopting some new, external behaviours that were unforeseen at the time of their development. The Part is divided into three Chapters.

In the first Chapter we provide a general overview on runtime software update and adaptation. In particular we focus on the renowned technologies and abstractions introduced to enable and structure runtime software updates.

In the second Chapter we show how choreographies can be modelled into a language, called DIOC, that provides suitable abstractions for programming distributed systems that update at runtime. We prove that DIOCs can be compiled (projected) to sets of processes that correctly enact the indented behaviour of their specification. We also prove that correctness is preserved after any step of update.

DIOCs provide a general support for software update and do not specify a particular policy for the application of updates (indeed such choice is orthogonal wrt the guarantees of correctness). In the third Chapter of this Part we present a refinement of our theoretical model of DIOCs aimed at modelling adaptation of distributed systems. We report how we implemented such model into a framework, called AIO CJ, for programming safe and adaptable distributed programs.

- **Part III — Applied Choreographies**

We address the issue of formalising the implementation of choreography models presenting our choreography theory of Applied Choreographies (AC). In particular, we model message passing based on correlation in the semantics of AC. We provide AC with a typing discipline that ensures the correct use of the low-level mechanism of message correlation, thus avoiding communication errors. We also define a two-step compilation from AC to a low-level Correlation Calculus, which is the basis of the executable language Jolie. Finally, we prove an operational correspondence theorem, ensuring that compiled programs behave as the original choreography.

- **Part IV — Conclusion**

We summarise the contributions of this dissertation and relate them to similar existing work. We also discuss some interesting directions of future investigation.

1.4. Structure of the Dissertation

- **Part V — Appendices**

- Chapter A — Adaptable Choreographies: Proofs.
- Chapter B — Adaptable Choreographies: Test Code.
- Chapter C — Adaptable Choreographies: Models of Adaptation.
- Chapter D — Applied Choreographies: Additional Material.
- Chapter E — Applied Choreographies: Proofs.

Part I

Introduction and Background

Introduction to Choreographies

Clara pacta, boni amici.

Latin proverb

2.1 Distributed Programming, in brief

A distributed system is a software application whose behaviour emerges from the interaction of several components (programs) that run in parallel and rely on message passing to communicate and coordinate their actions [2]. Although simple, this definition covers all kinds of distributed systems, ranging from networks of computers to communicating threads of the same program. The practice of writing distributed applications is called *distributed programming*.

Distributed programming born with the introduction of the first wide [18] and local [19] area networks of computers between the late 1970s and the early 1980s. Since that time, distributed programs evolved along the networks that supported their communications. During the 1990s the Internet [1] widely spread to constitute an heterogeneous collection of networks — a network of networks — including wired and wireless sub-networks connected by different technologies (e.g., Ethernet, Bluetooth, WiFi, etc.). Lately, Pervasive Computing [20] and Cloud Computing [21] are major threads of research, ranging from room-scale to global-scale networks of interacting component.

Distributed software constitutes the backbone of our society: it enables people to carry out tasks otherwise unmanageable, automating and ensuring the correct execution of complex social systems. Healthcare, transportation, commerce, and manufacturing are some of the industries that are currently experiencing strong benefits from the integration of software.

The reasons for writing distributed programs can be either compulsory or deliberate. In some cases the nature of the application requires to be distributed, e.g.,

when production and consumption of data happens in different places. In other cases, distributed programming fits a particular choice of the designer of the system, e.g., a distributed system that uses replicated and (geographically) distributed components to ensure reliability.

More in general, distributed programs benefit from two main features:

- *modularity* as message passing imposes no bonds on the components of a system. Any resource that behaves according to the specification of a component of the system can be employed as such. If said resource fails or becomes busy, the system might replace it with a new, compliant one by simply redirecting its messages.
- *scalability*, both in functions and size. Once deployed, a distributed program can become itself a component employed by another program, enhancing its functionalities with minimal effort. In the same way, since components are modular, if the workload on a component can be shared, it is possible to add a new instance of the same component in the system and separate the workload between the two.

Although distributed programming contributed to shape the present, the practice of correctly developing and debugging distributed systems remains one of the most challenging.

2.1.1 Problems of Distributed Programming

Distributed systems are developed to implement some protocol of interaction.

Let us take as an example the sequence of messages exchanged between an ATM terminal, the bank accounting system, and the issuer of the credit card to approve a withdrawal. When a client requires a withdrawal, the ATM forwards to the bank the *card ID* and the inserted *PIN*. Next the bank asks the card issuer to validate the request. Finally, the card issuer approves or denies the withdrawal and the bank forwards the outcome to the ATM. We depict in Figure 2.1 the sequence chart diagram of the example.

2.1. Distributed Programming, in brief

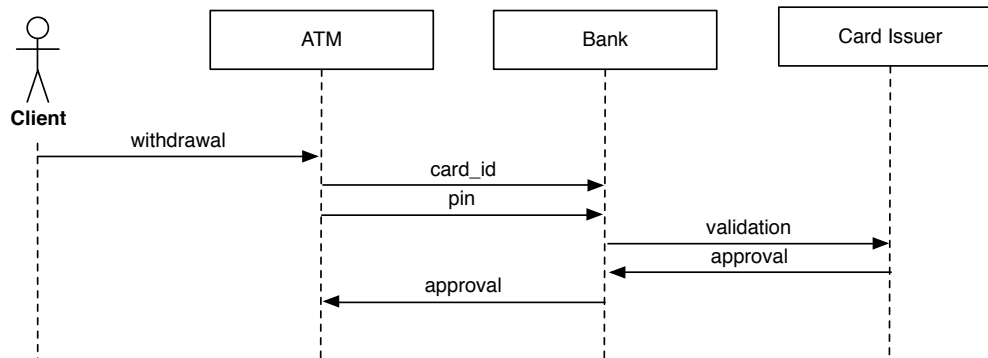


Figure 2.1: Sequence chart of the ATM withdrawal example.

Since the early days of distributed computing, developers introduced and used several tools to precisely describe the order of interactions between the components of a system, like message sequence charts [22] and sequence diagrams [23]. Baseline of all these tools is to avoid or at least minimise the ambiguity of the description of the sequence of messages in the system. However, developers struggle to make sure that the final interplay between the components of a distributed system correctly implement the global description of that system. The problem here is that they have to infer the logic of each component from a global protocol where interactions are considered as a whole and not broken down to sequences of send and receive actions.

Failing to correctly implement the global protocol of interaction of a system is the main source of system freezes and misbehaviours. Such errors are classified in literature as *deadlocks* [3] and *race conditions* [4].

Deadlocks. Deadlocks are one of the main causes of blocks of distributed systems. A system is in a *deadlock* state when one or more components hang, waiting for another component to release a resource (e.g., a message). The components cannot proceed with their computation and the whole system blocks.

Deadlocks usually occur when one of the components of a distributed application does not implement the global protocol of the system.

As an example, assume that the developer of the ATM mistakenly reversed the first two interactions with the bank. As a result, the system ends up in a deadlock state because *i*) the ATM hangs waiting for the bank to receive the message carrying the *PIN* whilst *ii*) the bank waits indefinitely for the message carrying the *card ID*.

Race Conditions. Race conditions are even subtler errors than deadlocks. When a system has a race condition, the runtime behaviour of an application depends on the timing or the sequencing of communications among its components. Due to their non-deterministic nature, race conditions are one of the most difficult bugs to spot and may potentially lead to worst scenarios than the block of the system; among these are data inconsistency and information leakage.

On the example in Figure 2.1, let us assume that the ATM sends the two messages for the *card ID* and *PIN* in parallel to the bank. Contrarily, the implementation of the bank is sequential: it waits to first receive the *card ID* and then the *PIN*. In this case we have two scenarios, depending on the order in which the bank receives the messages sent from the ATM. If the bank receives the *card ID* first then the interaction proceeds as intended. Otherwise the system ends up in a deadlock state where the bank waits to receive a message carrying the *card ID* whilst the first consumable message is the one carrying the *PIN*.

To avoid these faulty executions, programmers must make sure that all components of a distributed system correctly follow its designed protocol, however such strong guarantee is very difficult to assess.

On one hand, the static analysis of the non-deterministic interaction of many programs, executing in parallel, easily leads to the computation of exponentially many traces of execution, which makes such analysis computationally impractical.

On the other hand, non-determinism makes also the practice of testing and debugging the system unreliable. Plus, corner cases may be hard to predict and reproduce in all their possible combinations.

In the following sections we provide a brief introduction to the theoretical and practical solutions that have recently been proposed and adopted to ease the development of *safe* distributed systems. We introduce Service Oriented-Computing (SOC) and how SOC can be implemented following two opposite approaches, namely the Orchestration or the Choreographic approach. Finally, we present some seminal theoretical works that paved the way for the employment of choreographies as language abstraction for the development of safe distributed systems.

2.2 The Service-Oriented Approach

The World Wide Web [24] (WWW or Web for short) is an established technology that, although invented to provide human-readable content to users, has gone acquiring importance also for distributed programming.

Technologies like Uniform Resource Locators (URLs) [25], HyperText Transfer Protocol [26] (HTTP), and the eXtensible Markup Language [27] (XML) provided a common platform to make distributed programs communicate. In the context of the Web, distributed systems are composed by *Web services* that offer a

2.2. The Service-Oriented Approach

specific interface of interaction. The World Wide Web Consortium (W3C) defines a Web Service [28] as:

[...] a software system designed to support interoperable machine-to-machine interaction over a network. It has an interface described in a machine-processable format (specifically WSDL [29]). Other systems interact with the Web service in a manner prescribed by its description using SOAP [30] messages, typically conveyed using HTTP with an XML serialization in conjunction with other Web-related standards.

Paraphrasing the above definition, Web services pose no restriction on the internal technologies that implement their functionalities. Provided that a service describes its interface of interaction in a machine-processable format (WSDL) and interacts through SOAP messages, it will be able to interoperate with other Web services by receiving and sending requests through messages.

Notably, the computing paradigm described by Web services is not bound to Web standards but can be further abstracted to the so called *Service-Oriented* [31] paradigm.

Service orientation is platform-agnostic, it abstracts from the underlying technologies that enable components to interact in a network, and dictates that any resource (i.e., any program) of the network shall self-describe its functionalities through a standard interface. Services offer their functionalities following the client-server principle. In particular service orientation hinges on two principles:

- *Neutrality*: clients must be able to require functions of services through open, standardised, and wide-adopted technologies. Therefore all the layers that support the invocation of a request should comply with widely accepted standards.
- *Loose Coupling*: clients and services must agree only on the terms of the invocation of a function. Their internal implementations remain unknown to each other and the interaction does not rely on any knowledge of the internal structures that enact the execution of that functionality.

The service-oriented paradigm applies to any context that *i*) provides a common channel of communication for programs which *ii*) comply with some shared specification regarding protocols of communication and serialisation of data. Other than the Web stack, some renowned technologies that apply the service-oriented approach are *D-Bus* [32] and *DCOM* [33], which are respectively an open-source and a proprietary message bus that provide a context for one-to-one message passing for applications of the same system. Enterprise Service Bus [34] (ESB) tech-

nologies offer an analogous support for communication between distributed programs. In particular ESBs usually offer built-in transparent protocol conversion (e.g., HTTP, SOAP, DCOM, SAP RFC [35], etc.).

2.2.1 Breaking down complexity

The main novelty introduced by service orientation regards a new architectural perspective for organising distributed systems. Since the construction blocks of SOC are services, a service-oriented distributed system results from the composition of interacting services that form a Service-Oriented Architecture (SOA). The purpose of developing a distributed system as an SOA is to break down its complexity into simple services and build up the system in a hierarchical way.

This has several benefits, among which:

- *Parallel and compartmentalised development.* Since services communicate only through their standardised interfaces, an SOA can be quickly build up by different teams of developers working in parallel on a distinct subset of such services.
- *Reuse.* Usually projects share a common set of functionalities like logging, monitoring or security. Since services are compositional, it is easy to include an existing service that provides one of the required functionalities into a new project. Beside fostering re-usability of code, it also naturally lends itself to a market of specialised third-party services.
- *Simpler debugging.* Services that accomplish one simple task are easier to test and check for compliance (wrt specifics) than more complex ones.
- *Systems of systems.* SOAs can be built on top of other SOAs. The guiding principle remains that of ease of testability. Given a set of “simple” services, a first level of “simple” (to test) SOAs can accomplish more elaborate tasks as composition of such services. Iteratively, the most complex behaviours of a distributed system result from high-level SOAs that compose other architectures.

These general ideas gave birth to two opposite ways of composing services [7]:

Orchestration where the logic and order of the interactions in an SOA are defined and controlled by a single party, called the *orchestrator*.

Choreography which defines the sequence of messages each participant exchanges. Each party implements its own part and the intended interaction emerges from their interplay.

2.3. Orchestration

In the next section we present a brief introduction to the concept of orchestrated services and discuss the relationship between the former and choreography.

2.3 Orchestration

Nowadays, orchestration is the most adopted solution for service composition and the main reasons behind its wide adoption are *i*) that exposing and composing legacy programs (already developed and owned by businesses) requires little to no adjustments and *ii*) orchestrators become themselves new services that other orchestrators can use in their compositions.

Whilst the industrial standard for orchestration is WS-BPEL [36] (BPEL for short) all orchestration languages share a common base of functionalities to compose services at message level:

- *asynchronous* or *synchronous* delivery and reception of messages;
- sequential composition — an interaction can occur only after the one that precedes it;
- parallel composition — two or more interactions can occur in parallel;
- guarded composition — only one of a set of interactions can occur while the others are discarded.

As an example, we report in Figure 2.2 the orchestrated implementation of the protocol presented in Figure 2.1.

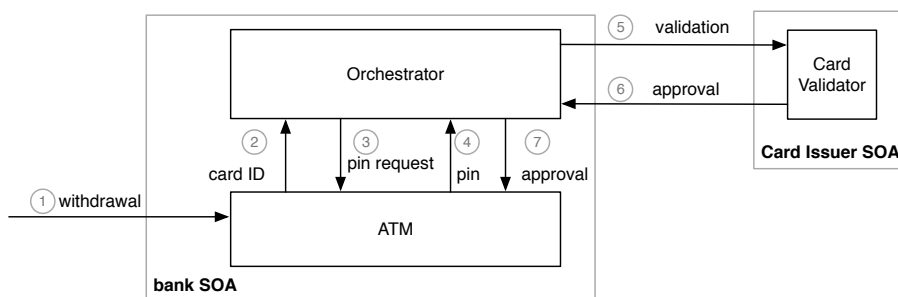


Figure 2.2: Example of orchestration.

In the example, since the bank is the owner of the SOA to which also the ATM belongs, we do not have a “bank” service but rather we introduce the Orchestrator service that implements the withdrawal protocol.

Following the orchestrative approach, the Orchestrator coordinates all interactions among the services in the systems. When the ATM receives a withdrawal request ①, it sends the *card ID* to the orchestrator. Then ② the Orchestrator asks the ATM for the *PIN*. After having received the *PIN* ③, the Orchestrator asks to validate the request to the Card Validator service (offered and owned by the card issuer). Finally, the Orchestrator receives the outcome of the validation (approval) ⑥ and ⑦ forwards it to the ATM.

In orchestration the logic of a distributed system is mostly implemented and enclosed in the orchestrator and therefore it becomes easier to check and test compliance between protocols and implementations.

However, orchestration poses a strong limitation in the development of distributed systems because it requires a central entity that manages the whole interaction whilst in many distributed contexts that assumption might not hold. For example, OpenID [37] is a widely adopted protocol that provides authentication in a decentralised way. In OpenID a service (e.g., a web site) delegates a relying party to authenticate users. The benefits are twofold: services do not have to manage users and users avoid to give their credentials to possibly untrusted services. In these cases, it is not possible to orchestrate the interaction and the only available solution is to develop each component separately.

2.4 Choreography

Unlike orchestration, choreography has no central point of coordination and describes SOAs in terms of interactions between the participant services. In a choreography all services are peers and the logic of the composition is spread among them. Choreographic composition is an unexplored topic wrt orchestration. In fact, the paradigm has been mostly used for protocol specification [8, 9].

W3C proposed a candidate recommendation, called WS-CDL [8]. In WS-CDL the developer specifies the protocol of observable interactions between services from a global point of view. As already remarked, the main reason behind the need for a choreographic composition is that the orchestration paradigm requires a full control on outputs and inputs toward and from the invoked services. On the contrary, there exist various systems that are composed by several services that communicate to each other but cannot be coordinated from a centralised position.

However, WS-CDL is not a language for programming composition of services, it is rather a formal language for specifying protocols, i.e., the correct sequences of interactions in the composition. Quoting the definition of the purpose of WS-CDL given by the W3C:

2.4. Choreography

The WS-CDL specification is aimed at being able to precisely describe collaborations between any type of participant regardless of the supporting platform or programming model used by the implementation of the hosting environment.

Using the WS-CDL specification, a contract containing a “global” definition of the common ordering conditions and constraints under which messages are exchanged, is produced that describes, from a global viewpoint, the common and complementary observable behaviour of all the participants involved. Each participant can then use the global definition to build and test solutions that conform to it. The global specification is in turn realized by combination of the resulting local systems, on the basis of appropriate infrastructure support.

Following to the above declaration of purposes, a choreography describes the interactions between several participants but it does not provide an implementation of such interaction. It is up to each participant to correctly implement its individual part in the choreography.

Beside the interpretation of choreographies given by the W3C, several works investigated the possibility to use choreographies as abstractions to implement distributed systems.

For example, OASIS proposed a different perspective to choreographic composition with BPEL4Chor [38]. Whilst WS-CDL sees choreographies as a top-down tool that can be used to specify protocols, OASIS considers choreographies from bottom-up, as an abstraction to define interoperability.

In particular, BPEL4Chor introduces a layer over BPEL where the developer defines only the activities of sending and reception of messages. Then, given an abstract specification in BPEL4Chor, the developer can implement the components that interact in the choreography as BPEL services. Once a BPEL4Chor choreography has been “grounded” into several BPEL orchestrators (one for each participant) the choreography is enacted by the parallel execution of the implemented orchestrators.

The two approaches have interests at odds (analysed in detail in [12]):

- top-down choreographies, à la WS-CDL, benefit from the maximum degree of abstraction, focussing only on the protocol of interaction. Communications are atomic and therefore they inherently exclude deadlocks and races (as seen in § 2.1.1). However, these kind of interaction-oriented choreographies are meant as design tools and cannot be used to implement distributed systems.
- contrarily, bottom-up choreographies written in BPEL4Chor are implemen-

table, allowing developers to follow a common schema of interaction. However, since these process-oriented choreographies maintain such close relation to the processes that implement the distributed system, they model communications as sequences of send and receive actions, possibly leading to deadlocks and races.

To exemplify the difference between the interaction-oriented and the process-oriented approach, we report in Figure 2.3 how the two approaches would model the scenario presented in Figure 2.1.

In the example we use the “Alice and Bob” notation [39] for interaction-oriented choreographies, where $A \rightarrow B : action$ represents a communication between two endpoints (components) of the system, specifically A sending a message through an *action* to B, and $;$ is the sequential composition of interactions. For process-oriented choreographies we use the notations *from* A : *action* and *to* A : *action* to respectively represent the receiving from and the sending to the endpoint A through *action*, $;$ still means sequential composition.

```
Client → ATM : withdrawal;
ATM → Bank : card_id;
ATM → Bank : pin;
Bank → Card Issuer : validation;
Card Issuer → Bank : approval;
Bank → ATM : approval
```

<p>ATM process</p> <pre>from Client : withdrawal; to Bank : card_id; to Bank : pin; from Bank : approval</pre>	<p>Bank process</p> <pre>from ATM : card_id; from ATM : pin; to Card Issuer : validation; from Card Issuer : approval; to ATM : approval</pre>	<p>Card Issuer process</p> <pre>from Card Issuer : validation; to Bank : approval</pre>
---	---	--

Figure 2.3: Upper part: example of interaction-oriented choreography. Lower part: from left to right, the process-oriented choreographies of the ATM, the Bank, and the Card Issuer.

2.5 Programming with Choreographies

Recent theoretical investigations explored how the interaction- and process-oriented aspects of choreographies could be merged into one language, namely one that enjoys the minimality and safety (deadlock- and race-freedom) of top-down chore-

2.5. Programming with Choreographies

ographies but able to express process-level interactions, needed to implement distributed systems.

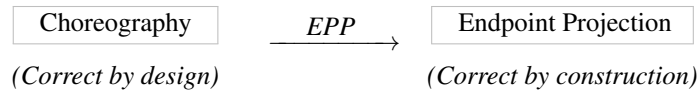
Seminal work in such endeavour are [10, 40, 11].

In [10] Qiu et al. give a first theoretical account of the fundamental issues between interaction-oriented choreographies and implementations. In that work, the authors introduce the concept of *projection* to map the global behaviour of a distributed system, described as a choreography, into a set of processes that implement it.

In [40], Bravetti and Zavattaro build on the concept of *service contracts* [41] — i.e., descriptions of the behaviour of services in terms of their inputs and outputs — and define a projection that, given an interaction-oriented choreography, generates a set of contracts, one for each of its participants. Then, it is possible to implement the original choreography by executing those services whose contracts match the generated ones.

Finally, in [11] Carbone et al. formalise a core WS-CDL language and relate interaction-oriented choreographies and implementations by means of an Endpoint Projection (EPP) function. The EPP is a mapping from a choreography specification to a set of processes which, run in parallel, enact the behaviour described by the choreography. The target of the EPP is an applied π -calculus [5] and not an actual executable language. However in [11] the authors prove that interaction-oriented choreographies can be made expressive enough to define the implementation of safe distributed systems. Essentially, the projected processes enact all and only the behaviours described in the choreography (protocol) and since choreographies cannot express deadlocks and races also the projected system is deadlock- and race-free.

The most notable results of [11] are that *i*) it paved the way for the concept of choreographic programming [42] — i.e., to use choreographies as abstraction to implement distributed systems —, *ii*) it inspired the introduction of *multiparty session types* [43] as evolution of *binary session types* [44], and *iii*) it provided a methodology for the development of distributed software, based on a correctness-by-construction approach, which we can depict as:



Subsequent theoretical works [12, 13, 14] followed a similar approach, extending the choreography model to support multiparty sessions, channel mobility, and modularity.

On the other side of the spectrum, some early works [45, 46, 47] explored how choreographies could be used to support the implementation of distributed programs, however none of these proposals uses choreographies as a programming

abstraction and rather employs them to check endpoint programs.

Chor [48] is the first work that brought the theoretical results on choreographies into the world of implementation languages. Based on the theoretical framework presented in [13], Chor supports the definition of global descriptions (protocols), the programming of compliant choreographies, and the safe projection of said choreographies into a collection of distributable and executable orchestrators.

Chor gives a tangible proof that choreographic programming is a suitable paradigm for the implementation of real-world distributed systems because: *i*) it lets developers focus on the description of the interactions between the components in the system and *ii*) it generates deadlock- and race-free distributed applications that are guaranteed to follow the designed protocol.

Nonetheless, Chor is just the first attempt at bringing into the real world the promising theoretical results on choreographies. The language lacks some standard features of programming languages like *modularity*¹, *dynamic updates* or *error handling*. Beside the expressiveness of the language, the major limitation of Chor is that its EPP is not formally proven. The theoretical treatment on which Chor is based uses as target language a process calculus that models communications through synchronisations on names (as in CCS [49] and the π -calculus [5]). Contrarily, the EPP equipped with Chor targets Jolie [50], which is an orchestration language that uses correlation (à la BPEL) to route messages. Hence, to ensure that programs projected from Chor correctly implement their originating choreography, it is necessary to formalise and prove an Endpoint Projection procedure for the targeted executable language.

2.6 Contributions

In Chapters 4–5 we give a positive answer to the question:

Can we model safe and dynamic updates of distributed systems with choreographies?

In Chapter 4 we present our theoretical model of Dynamic Choreographies and our results of correctness of projection and runtime update. In Chapter 5 we refine our theoretical model of Dynamic Choreographies to provide a fine control over updates. As a result we develop a programming framework for developing *adaptive* [51] (cf. Chapter 3) distributed systems. Like Chor, AIOCJ proves that choreographic programming can be used to deal with different and real-world application domains, preserving the desirable properties of correctness and safety of interaction-oriented choreographies.

¹The possibility to compose different choreographies — or better said, different systems projected from different choreographies — to implement applications in a compositional way.

2.6. Contributions

In Chapter 6 we progress our venture by tackling one of the most important problems of bringing choreographies to the real world. As discussed in § 2.5, Chor — but also AIOCJ — is based on a theoretical model that abstracts the mechanics of communication with synchronisation on names. Although such abstraction allows for an elegant theoretical treatment, it sensibly departs from the reality of distributed systems, where some underlying technology must be set and used to support the routing of messages between processes.

As a result, the implementations of Chor and AIOCJ loosely follow the EPP defined in their foundational models. This breaks the chain of guarantee of correctness from choreographies to the implemented systems. For instance, to establish the communication channels between the processes in the system, the programs projected by Chor and AIOCJ make some preliminary communications, unexpected by their theoretical models. Whether those communications are deadlock or race-free and that they do not interfere with the behaviour defined at the choreography level is not investigated nor theoretically proven. Moreover, these additional interactions weaken one of the main advantages of choreographic programming: the clear specification of the communications carried out during execution. The leading research question that we pose in Chapter 6 is therefore:

Can we formalise the implementation of communications in choreographies?

Interestingly, this question easily overcomes the boundaries of choreographies. The problem of bringing the results on name-based theoretical models for distributed systems into the real world of implemented languages has been faced in various works, following different custom practices [52, 53, 54]. Therefore, solving the problem of formalising the implementation of communications in choreographies would be a useful reference to formalise the implementation of name-based languages in general.

Our answer, presented in Chapter 6, is to develop a theoretical model for Applied Choreographies (AC). ACs support the most advanced features of recent choreography models like modularity [14], asynchrony, parallelism, and the dynamic creation of multiparty sessions and processes, yet it models message-passing based on *correlation*, one of the main technologies used in Service-Oriented Computing (BPEL, Java/JMS/ C#.NET, etc.) to support message routing.

On the definition of AC, we pinpoint the key theoretical problems and formalise the guiding principles that developers should follow to obtain correct implementations. Finally, to prove our approach, we define a correct compiler from AC choreographies to the theoretical model of the Jolie language.

Part II

Adaptable Choreographies

Runtime Software Update and Adaptation

In Chapter 2 we discussed how distributed software has become a commodity used on a daily basis like phone lines, electricity, and water.

As our society becomes more and more dependent on software, it also demands for solutions that quickly and continuously update and adapt to the changes in objectives and the surrounding environment. Software systems that respond to change by self-adapting while running — and providing services — reflect the vision of *autonomic* computing [55]. Along with this requirement of flexibility, software must remain reliable and predictable, in particular when adapting business- or safety-critical applications.

This requirements gave birth to an important branch of research interested in providing solutions like design patterns, language abstractions, and frameworks for the reliable development of systems that are able to adapt, preserving some desirable properties of correctness and coherence after the adaptation.

In the next sections we present some of the renowned technologies that enable runtime software update and to structure adaptation, focussing on the constructs introduced by these technologies to:

- harness the inherent complexity of updating different parts of a software system, from a software engineering viewpoint;
- guarantee coherence between all the updated components and the correctness of the system after any update.

3.1 Linguistic Constructs for Adaptation

The practice of adapting software at runtime spans various levels of abstraction, ranging over parameters, methods, aspects, components, applications, architectures, systems, and data-centres [56].

In [57] the authors clarify that there exist two main approaches to software adaptation:

This research has been addressing both the theoretical foundations of adaptation and, more pragmatically, how adaptation techniques can be applied to solve the problems at hand. [...] To achieve these goals, solutions to engineer self-adaptive behaviours are often sought at the software architecture level, including middleware and component-based design. Accordingly, architectural approaches to dynamic adaptation have been extensively studied by researchers [58, 59, 60]. As a complementary approach, researchers also adopted specialized programming paradigms to implement adaptive systems, such as metaprogramming and Aspect-Oriented Programming.

Essentially, the first approach is an architectural one and structures adaptation at the level of software components. The other approach is at language-level, which aims at providing language constructs for structuring which parts of a programs shall change at runtime. In their analysis, the authors of [57] remark that:

Supporting behavioural variations at the language level is appealing with respect to dealing with adaptation at a higher architectural level. First, language-level approaches provide means to adapt at a very fine-grain detail. Second, many language concepts (e.g., polymorphism and late binding) are well known to programmers and can be easily specialized to support adaptation [61].

Hence, language-level solutions provide a finer control over the adaptation actions wrt architectural ones. In particular, constructs for adaptation embedded in programming languages provide a fine-grained granularity on adaptation whilst architectures for software adaptation usually rely on a component-based design that allows adaptation only at a coarser level.

Moreover, adaptation at language level brings several benefits:

- research on constructs for adaptation at language level tends to elicit general concepts that can be integrated, mixed, and evolved within new languages;
- languages take advantage of various techniques like syntactic and type checkers, compilers, etc. that enforce correctness of software;
- adaptation at architectural level introduces strong dependencies on the frameworks that enable adaptation.

Focussing on the practice of programming adaptive systems, we highlight two aspects that a language has to address for programming adaptation.

3.1. Linguistic Constructs for Adaptation

The first is an *enabling* aspect, i.e., the possibility of a language to control and change its own programming at runtime. The second is a *structuring* aspect, i.e., the abstractions offered by the language to program updates in a structured way.

In the following sections we briefly overview some of the main linguistic constructs that enable software updates and provide structure to their application.

3.1.1 Reflection and Metaprogramming

In [62], Pattie defines computation reflection as:

[...] the behaviour exhibited by a reflective system, where a reflective system is a computational system which is about itself in a causally connected way.

In the context of the definition, a software system is said to be *causally connected* if there is a bond between its internal structures and the domain they represent. In particular, a reflective system is a computational system endowed with structures that represent aspects of itself. Having access to its self-representation, the system can respond to queries about its structural state and change it. Since a self-representing system is causally connected to the domain of the aspects it represents, *i*) the system has always an accurate representation of itself and *ii*) its computational status is always compliant with its representation, i.e., it can modify itself by means of its own computation.

Metaprogramming is the practice of using *computational reflection*, i.e., the use of a reflective interface on the program to change its behaviour. As exposed above, meta (domain) and base levels are causally connected, therefore changes at meta-level “reflect” in behavioural variations at base level.

Reflection is the most general and powerful mechanism to obtain runtime behavioural variation in programs. Examples of its usage span from higher-order languages like Lisp [63], which uses reflection to handle its own programming structures as *first-class* data types, to language APIs, which enable for runtime structure definitions (e.g., classes).

In general, the main two features of any metaprogramming facility are *inspection* and *modification*. Inspection supports queries on the state of execution of a program, whilst modification supports the change of the behaviour of a program, e.g., intercepting calls, modifying dispatching policies, augmenting data with new structures, and injecting new functions or modifying the body of existing ones. [64, 65, 66] are some established examples of usage of metaprogramming. In [64] Ledoux explains how it is possible to modify dynamically the behaviour of objects through dynamic classes. In [65] Dowling et al. employ reflection to enable a program to change the binding between its methods and their implemen-

tations. In [66] Xu et al. propose a method that exploits call interception and code injection for implementing fault tolerance.

All cited works prove that metaprogramming is the staple linguistic element that enables developers to pursue any kind of runtime adaptation on programs. However, due to its unstructured nature, the direct use of metaprogramming tends to generate code that is difficult to read and debug because programmers have to figure out which variations affected the behaviour of the program at runtime. More in general, metaprogramming does not guarantee any property on the state of the system after adaptation.

This led to the creation of a new set of language constructs *i*) to precisely define the boundaries of the effects of metaprogramming and to *ii*) to clearly specify the conditions that fire the application of each variation.

3.1.2 Aspect- and Context-Oriented Programming

Aspect-Oriented programming [67] (AOP) is one of the most renowned paradigms for managing runtime software update. The main reason behind the development of the aspect-oriented paradigm is to be able to separate the essential functionalities of a program from orthogonal modules like logging, persistence, synchronisation, and fault handling. The practice of uncoupling orthogonal modules from the main functionalities of a software is called *separation of concerns* and it is achieved by two specific constructs of AOP, namely *point-cuts* and *join-points*.

A point-cut corresponds to a set of join-points. When the execution of the main code of a program reaches one of the join-points of a point-cut, the control is transferred from the main code to the code that implements a separate concern of that point-cut. The code corresponding to a join-point of a point-cut is called *advice*. The *aspect weaver* is the metaprogramming utility that “weaves” (merges) together the main code of an application, its point-cuts, and a set of advices to generate the final implementation of an aspect-oriented application.

Aspect-oriented programming has been successfully employed in many works on dynamic adaptation. Some examples are [68], [69], whilst [70] employed it to adapt distributed systems.

Although AOP provides a powerful yet easy abstraction to manage separation of concerns, its scope of application remains that of single programs and lacks mechanisms that ensure coherence among the adapted parts of a distributed system. Moreover, it is possible that various conditions coexist in the environment of execution and that several variations apply on the same system. Hence, beside making sure that adaptation is coherent over the system, it should be possible to express *conflicts* between variations. Adding such conflicting conditions in plain AOP quickly clutters the code, tangles the present point-cuts, and makes the insertion of new ones harder.

3.1. Linguistic Constructs for Adaptation

This motivations paved the way for the development of an evolution of the aspect-oriented paradigm for context-aware software: Context-Oriented programming [71] (COP).

The abstractions introduced by COP regard the modularisation of software updates to support the coherent activation of behavioural variations, based on *contexts* [71], that Hirschfeld et al. define as:

Any information which is computationally accessible may form part of the context upon which behavioural variations depend.

Given its recent introduction, one major direction of research on COP is involved in defining the essential features a context-oriented language should support. The context-oriented paradigm relies on *i*) proper abstractions for representing variations and *ii*) explicit mechanisms for triggering variations. Usually, COP extensions for mainstream languages [72], like ContextJ [71], provide abstractions for representing variations as first-class structures, which can be referenced and returned. The language provides ad-hoc mechanisms to dynamically activate and deactivate variations. When a variation is enabled, the set of behaviours it models activate and modify the runtime behaviour of the application. Vice versa, if the variation is deactivated the control returns to the main code.

Variations can also be activated simultaneously. In that case, the final behaviour of the application derives from the combination of the active variations. Since the applicability of variations is specified wrt both environmental conditions and the state of other variations, context-oriented languages ensure two properties: *i*) all the advices related to a variation are coherently active or inactive *ii*) conflicting advices cannot occur.

ContextErlang [73] is an interesting example of context-orientation brought in the (possibly) distributed world, by enabling the coherent application of variations on a per-thread basis. Variations can be both expressed internally and be included at runtime from an external source via message passing.

Although AOP and COP give a sensible contribution to ease the programming of runtime software updates, their scope remains that of single programs. Indeed, to the best of our knowledge, there are no formal results on the employment of such programming abstractions to preserve the coherency of updated distributed programs nor on the correctness of the programs after the application of updates.

3.1.3 Adaptation in Process-Aware Information Systems

In § 3.1, we underlined that in this work we do not consider architectural-level solutions wrt to language-level ones because the latter give a finer control over the adaptation actions.

Nonetheless, architectural-level approaches like that of *business process management* [74], and in particular that of Process-Aware Information Systems [75] (PAIS), investigated in depth how flexibility [76] — to support dynamic process adaptations — impacts on business processes and their realisation.

PAISs provide a high-level perspective over the whole composition of processes and, since the implementation is detached from the logic of composition, it is possible to plan for various degrees of flexibility by selecting which and when parts of the composition can change.

The community of BPM and PAIS did a formidable work in breaking down the nuances of such flexible systems into their attributes. In the context of this work, it is interesting to summarise such attributes and patterns, using them to clarify the possibilities offered by our language-level solution presented in Chapter 4.

In particular, we focus on the notion of *loosely specified processes* as given in [76] and defined as: *process models that are not fully pre-specified and that keep some parts unspecified at built-time by deferring decisions to runtime*. We report below the same breakdown of the attributes of such processes, drawn from [76]:

freedom given to the user to change the process. The lowest degree of freedom is *none*, which entails fixed processes that cannot change at runtime. Above it, we have the runtime *selection* of parts from a set of predefined alternatives. The highest level is the *modelling or composition* of parts of fragments of processes out of a set of activity templates;

planning approach used to take decisions when changing some parts. From most constrained to loosest: the *plan-driven* approach defines a detailed model of change up-front, the *iterative and continuous* approach relies on a coarse plan which is refined at runtime, and the *ad-hoc modelling* has no plan and resorts to the interaction with some user-manned or automatic system;

scope of change, i.e., which “regions” of the process model can be changed. The scope ranges from *predefined regions* to the *entire process*;

process perspective define which properties of a process can change. Changes can regard: the *behaviour* of a process, the *actors* executing a particular activity, the *data* used in the computation, the *functions* available in the process and their concrete implementations (called *operational decisions*), and finally the concrete *time* of start or end of activities;

making and support of decisions to concretise loosely specified processes at runtime. Decision systems are classified by the policy on which they base their decisions. These can be *goal-based*, which take into account the overall goals of the process, *rule-based*, which rely on a set of rules, *experience-based* which suggest solutions based on past experiences made in similar contexts (e.g., execution logs), *user decision* where decision making is left to the human expert;

degree of automation given to the system to compose instances of processes.

3.1. Linguistic Constructs for Adaptation

The division goes from full *automation*, *system-supported* interaction with the user (e.g., by providing context-specific templates which can be adopted), and *manual*, without system support.

Focussing in particular on the first two attributes of flexibility, namely the degree of *freedom* and the *planning approach*, the authors of [76] elicit some patterns of decision deferral, listed below in order of degree of flexibility:

Pattern	Description	Degree of freedom	Planning approach
Fully pre-specified processes	Processes that offer little flexibility. The model is fully pre-specified. They use a strict plan-driven approach separating modelling and execution entirely.	none	plan-driven
Late selection	Processes that allow placeholder activities (whose <i>scope</i> regards predefined regions). The placeholder activities are refined during runtime with concrete implementations from the available set of alternatives, which must be specified at build-time following a plan-driven approach.	selection	plan-driven
Late modelling and composition	Processes that allow also the automatic composition (not only selection) of activities. The <i>scope</i> can encompass regions or the whole process. The runtime refinement is still driven by a pre-specified plan.	modelling and composition	plan-driven
Iterative refinement	Processes that allow users to iteratively compose the content of a placeholder activity by selecting any process fragment from a given pool. The runtime refinement is not bound to a pre-specified plan. The <i>scope</i> spans the whole process or regions of it.	modelling and composition	iterative
Ad-hoc composition	Processes are composed in an ad-hoc manner, without any planning, with the <i>scope</i> spanning the whole process or regions of it.	modelling and composition	ad-hoc

Table 3.1: Decision deferral patterns.

As stated at the beginning of this section, in § 7.1.1 we use the fine-grained

definitions of flexibility reported above to clarify the possibilities offered by our language-level choreography-based solution presented in Chapter 4.

3.2 Safe Update of Distributed Systems

Updating a single (non-distributed) program and making sure that its behaviour remains correct and coherent is a difficult task. However, also distributed systems and, by extension, the programs composing them, need to be updated and adapted. Guaranteeing that a distributed system remains correct and coherent after a step of adaptation is even more challenging.

Some interesting works on dynamic adaptation of distributed systems regard application of Aspect-Oriented programming to Service-Oriented Applications (introduced in § 2.2). For instance, in [77] Charfi et al. proved that adaptive service composition can be effectively addressed combining the constructs of AOP in BPEL [36]. However, the current technology of development is not well suited for writing and verifying adaptive distributed systems. This is due to the lack of high-level structuring abstractions needed for handling complex communication behaviours and context-aware adaptation.

Recently, several works introduced interesting abstractions for providing high-level specifications of the expected behaviour of a distributed systems. In particular we cite session types [11, 78, 43], choreography languages [12, 13], behavioural contracts [40], and ad-hoc scripting languages [79].

These concepts usually consider static techniques, which alone are not sufficient to model dynamically adaptable software. Indeed, the assumption on either the ability to type-check the component source code, or the availability of its complete behavioural interface, may not be realistic in presence of adaptation.

On the practice of adapting a distributed system, [80, 81, 61, 82, 83] are some interesting proposals of middlewares and architectures to enable runtime adaptation (an interesting survey can be found in [84]). Other works applied concepts of adaptive languages in the context of distributed systems like [85] with traits, [86] with role-based modelling, and [87] with aspect-oriented programming.

All the mentioned approaches provide tools for programming distributed adaptive systems, although they do not provide any guarantee on the expected behaviour of the system after the execution of some steps of adaptation. In the case of dynamically adaptive distributed programs, techniques like static analysis are impracticable due to the impossibility to know in advance the structure of the adapted system. For this reason, most of the approaches in the literature offer no guarantee on the behaviour of the adaptive system after adaptation.

In the next Chapter we report our choreographic solution to the problem of the safe, runtime (i.e., dynamic) update of distributed system.

Dynamic Choreographies

Intelligence is the ability to adapt to change.

Stephen W. Hawking

4.1 Introduction

As illustrated in Chapter 2, programming safe distributed systems is a difficult activity, which can be made even more difficult if those systems shall update their behaviour at runtime. In Chapter 3 we presented the difficulties of such task within a single process and explained why updating a distributed system at runtime in a correct and coherent way is even more challenging. However, enabling distributed software to safely update at runtime has countless uses, spanning from dealing with emergencies, to coping with unexpected requirements or to improving and specialising applications to user preferences.

Here, we argue that choreographies (presented in § 2.4) are a suitable abstraction to express the dynamic update of distributed systems and to guarantee coherency and correctness of the system after each step of update.

In our treatment, we consider applications whose code can change at runtime by integrating other code provided by some external party. We propose a general mechanism for the structuring of updates that consists in delimiting inside the application blocks of code, called *scopes*, that may be dynamically replaced with new code, called an *update*. Remarkably, scopes and updates are independent and the details of the behaviour of updates do not need to be foreseen when writing the scopes. Indeed, updates may be written while the application is running.

Moreover, choreographies naturally lend themselves to the clear definition of the boundaries of the update and to enforce the coherent application of updates on all (and only) involved participants. To exemplify this concept, let us consider a simple interaction of a buyer that asks to a seller the price of some product.


```

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

```

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*. The Seller computes the price of the product calling the function *getPrice* (Line 2) and, via the channel *offer*, it sends the price to the Buyer (Line 3), that stores it in a local variable *prod_price*.

On the above example, let us suppose that the Seller would like to introduce new commercial offers in the future, e.g., to provide a discount on the computed prices. Since we want to be able to update the computation of prices, we enclose Lines 2–3 of the example within a scope, as shown below.

```

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

```

In essence, a scope is a delimiter that defines which part of the application can be updated. Each scope identifies a participant as the *coordinator* of the update, responsible to decide whether to update and which (among possibly many) update to apply. In our example, the coordinator of the update is the Seller (Line 1).

Now that we have a scope, we can introduce runtime updates. Let us assume that the Seller issued a fidelity card to some returning customers and now the Seller wants to update the system to let Buyers insert their fidelity card to get a discount. The update in Figure 4.1 answers this business need.

```

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 4.1: Fidelity Card Update.

At runtime, if the Seller (which is the coordinator of the update) selects the

4.1. Introduction

update in Figure 4.1, the code of the update replaces the one in the body of 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 separately writing the code for the Buyer and the Seller with some process-level language equipped with update primitives (cf. 3) would be extremely error prone. For instance, let us consider the case in which the Buyer is updated first and it starts the new interaction whilst the Seller cannot find the same update. The Buyer proceeds to send its card ID to the Seller. Contrarily, the Seller is executing the original interaction, it does not expect to receive the message carrying the card ID from the Buyer and the system is stuck (deadlocked): the Buyer waits for the Seller to receive its message carrying the card ID whilst the Seller waits for the Buyer to receive the message carrying the price of the product. In our setting, the available updates may change at any time, posing an additional challenge. Extra precautions are needed to ensure that all the participants agree on which code is used for a given update.

If both the original application and the updates are programmed using a choreographic language, these problems cannot arise. In fact, at the choreographic level, the update is applied atomically to all the involved participants.

However, providing the correct compilation from choreographic specifications to process-level distributed code is a challenging task. In particular, at the level of processes, the different participants have to coordinate their updates avoiding inconsistencies.

In this Chapter we propose our theory of dynamic choreographies. We present our choreographic language, called DIOC, for the programming of distributed applications supporting code update (§ 4.2). Next, we introduce a process-level language, called DPOC, based on standard send and receive primitives (§ 4.3) and we define our Endpoint Projection to compile DIOCs into DPOCs (§ 4.4). The EPP is proven correct (behaviour-preserving) (§ 4.7). We also prove that the property of correctness is preserved after any step of update.

In Chapter 5 we present one instantiation of our theory into a development framework for adaptive distributed applications called AIOCJ. 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.

4.2 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.

4.2.1 DIOC Syntax

DIOCs rely on a set of *Roles*, ranged over by R, 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:

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

Interaction $o : R(e) \rightarrow S(x)$ means that role R sends a message on operation o to role S (we require $R \neq S$). The sent value is obtained by evaluating expression e in the local state of R and it is then stored in the local variable x of S . Processes $\mathcal{I}; \mathcal{I}'$ and $\mathcal{I} | \mathcal{I}'$ denote sequential and parallel composition, respectively. Assignment $x@R = e$ assigns the evaluation of expression e in the local state of 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 $\text{if } b@R \{ \mathcal{I} \} \text{ else } \{ \mathcal{I}' \}$ and the iteration $\text{while } b@R \{ \mathcal{I} \}$ are guarded by the evaluation of the boolean expression b in the local state of R . The construct $\text{scope } @R \{ \mathcal{I} \}$ delimits a subterm \mathcal{I} of the DIOC process that may be updated in the future. In $\text{scope } @R \{ \mathcal{I} \}$, role R is the coordinator of the update: it decides whether to update or not and which update to apply.

A Running Example. We report in Figure 4.2 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.

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 Buyer takes the name of the *product* from the user with function *getInput*, which models interaction with the user (Line 4), 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 Seller offer, the Seller sends to the Bank the payment details (Line 16). Next, the Buyer authorises the payment via operation *pay*. We omit the details of the local execution of the payment at the Bank. Since the payment may be critical for security reasons, the related communications are enclosed in a scope (Lines 17–21), thus allowing the introduction of more refined procedures later on. After the scope successfully terminates, the application ends with the Bank acknowledging the payment to the Seller and the Buyer in parallel (Lines 23–25). If the payment is not successful, the failure is notified to the Buyer only. 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.

4.2.2 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 $i \in \mathbb{N}$ 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.

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

```

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     payment_ok@Bank = true;
19     pay : Buyer( payAuth( prod_price ) ) → Bank(auth);
20     // code for the payment
21   };
22   if(payment_ok)@Bank{
23     confirm : Bank(null) → Seller(_)
24     |
25     confirm : Bank(null) → Buyer(_)
26   }else{
27     abort : Bank(null) → Buyer(_)
28   }
29 }

```

Figure 4.2: DIOC process for Purchase Scenario.

4.2. Dynamic Interaction-Oriented Choreographies

Definition 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 $i \in \mathbb{N}$, resulting in the following grammar:

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

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 R , namely $\Sigma : Roles \times Var \rightarrow Val$. The local state of role R is denoted as $\Sigma_R : Var \rightarrow Val$ and it verifies that $\forall x \in Var : \Sigma(R, x) = \Sigma_R(x)$. Expressions are always evaluated by a given role R : we denote the evaluation of expression e in local state Σ_R as $\llbracket e \rrbracket_{\Sigma_R}$. We assume that $\llbracket e \rrbracket_{\Sigma_R}$ 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 , $\llbracket b \rrbracket_{\Sigma_R}$ is either `true` or `false`.

Definition 2 (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 3 (DIOC systems semantics). The semantics of DIOC systems is defined as the smallest labelled transition system (LTS) closed under the rules in Figure 4.3, where symmetric rules for parallel composition have been omitted.

The rules in Figure 4.3 describe the behaviour of a DIOC system by induction on the structure of its DIOC process. We use μ to range over labels. The possible values for μ are described below.

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

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

$$\begin{array}{c}
 \frac{\llbracket e \rrbracket_{\Sigma_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} \quad [\text{DIOC} \mid \text{INTERACTION}] \\
 \\
 \frac{\llbracket e \rrbracket_{\Sigma_R} = v}{\langle \Sigma, \mathbf{I}, \mathbf{i}: x @ R = e \rangle \xrightarrow{\tau} \langle \Sigma[v/x, \mathbf{R}], \mathbf{I}, \mathbf{1} \rangle} \quad [\text{DIOC} \mid \text{ASSIGN}] \\
 \\
 \frac{\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle \xrightarrow{\mu} \langle \Sigma, \mathbf{I}', \mathcal{I}' \rangle \quad \mu \neq \checkmark}{\langle \Sigma, \mathbf{I}, \mathcal{I}; \mathcal{J} \rangle \xrightarrow{\mu} \langle \Sigma, \mathbf{I}', \mathcal{I}'; \mathcal{J} \rangle} \quad [\text{DIOC} \mid \text{SEQUENCE}] \\
 \\
 \frac{\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle \xrightarrow{\checkmark} \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} \quad [\text{DIOC} \mid \text{SEQ-END}] \\
 \\
 \frac{\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle \xrightarrow{\mu} \langle \Sigma, \mathbf{I}', \mathcal{I}' \rangle \quad \mu \neq \checkmark}{\langle \Sigma, \mathbf{I}, \mathcal{I} \parallel \mathcal{J} \rangle \xrightarrow{\mu} \langle \Sigma, \mathbf{I}', \mathcal{I}' \parallel \mathcal{J} \rangle} \quad [\text{DIOC} \mid \text{PARALLEL}] \\
 \\
 \frac{\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle \xrightarrow{\checkmark} \langle \Sigma, \mathbf{I}, \mathcal{I}' \rangle \quad \langle \Sigma, \mathbf{I}, \mathcal{J} \rangle \xrightarrow{\checkmark} \langle \Sigma, \mathbf{I}, \mathcal{J}' \rangle}{\langle \Sigma, \mathbf{I}, \mathcal{I} \parallel \mathcal{J} \rangle \xrightarrow{\checkmark} \langle \Sigma, \mathbf{I}, \mathcal{I}' \parallel \mathcal{J}' \rangle} \quad [\text{DIOC} \mid \text{PAR-END}] \\
 \\
 \frac{\llbracket b \rrbracket_{\Sigma_R} = \text{true}}{\langle \Sigma, \mathbf{I}, \mathbf{i}: \text{if } b @ \mathbf{R} \{ \mathcal{I} \} \text{ else } \{ \mathcal{I}' \} \rangle \xrightarrow{\tau} \langle \Sigma, \mathbf{I}, \mathcal{I} \rangle} \quad [\text{DIOC} \mid \text{IF-THEN}] \\
 \\
 \frac{\llbracket b \rrbracket_{\Sigma_R} = \text{false}}{\langle \Sigma, \mathbf{I}, \mathbf{i}: \text{if } b @ \mathbf{R} \{ \mathcal{I} \} \text{ else } \{ \mathcal{I}' \} \rangle \xrightarrow{\tau} \langle \Sigma, \mathbf{I}, \mathcal{I}' \rangle} \quad [\text{DIOC} \mid \text{IF-ELSE}] \\
 \\
 \frac{\llbracket b \rrbracket_{\Sigma_R} = \text{true}}{\langle \Sigma, \mathbf{I}, \mathbf{i}: \text{while } b @ \mathbf{R} \{ \mathcal{I} \} \rangle \xrightarrow{\tau} \langle \Sigma, \mathbf{I}, \mathcal{I}; \text{while } b @ \mathbf{R} \{ \mathcal{I} \} \rangle} \quad [\text{DIOC} \mid \text{WHILE-UNFOLD}] \\
 \\
 \frac{\llbracket b \rrbracket_{\Sigma_R} = \text{false}}{\langle \Sigma, \mathbf{I}, \mathbf{i}: \text{while } b @ \mathbf{R} \{ \mathcal{I} \} \rangle \xrightarrow{\tau} \langle \Sigma, \mathbf{I}, \mathbf{1} \rangle} \quad [\text{DIOC} \mid \text{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}: \text{scope } @r \{ \mathcal{I} \} \rangle \xrightarrow{\mathcal{I}'} \langle \Sigma, \mathbf{I}, \mathcal{I}' \rangle} \quad [\text{DIOC} \mid \text{UP}] \\
 \\
 \frac{}{\langle \Sigma, \mathbf{I}, \mathbf{i}: \text{scope } @r \{ \mathcal{I} \} \rangle \xrightarrow{\text{no-up}} \langle \Sigma, \mathbf{I}, \mathcal{I} \rangle} \quad [\text{DIOC} \mid \text{NOUP}] \\
 \\
 \frac{}{\langle \Sigma, \mathbf{I}, \mathbf{1} \rangle \xrightarrow{\checkmark} \langle \Sigma, \mathbf{I}, \mathbf{0} \rangle} \quad [\text{DIOC} \mid \text{END}] \quad \frac{}{\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle \xrightarrow{\mathbf{I}'} \langle \Sigma, \mathbf{I}', \mathcal{I} \rangle} \quad [\text{DIOC} \mid \text{CHANGE-UPDATES}]
 \end{array}$$

Figure 4.3: Annotated DIOC system semantics.

4.2. Dynamic Interaction-Oriented Choreographies

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

Figure 4.4: Auxiliary function roles.

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. The rules $\llbracket \text{DIOC} \mid_{\text{UP}} \rrbracket$ and $\llbracket \text{DIOC} \mid_{\text{NOUP}} \rrbracket$ 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 roles(\mathcal{I}), defined in Figure 4.4, 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 connected(\mathcal{I}') holds if \mathcal{I}' is connected. Connectedness is a well-formedness property of DIOCs and is detailed in § 4.6. Predicate freshIndexes(\mathcal{I}') holds if all indexes in \mathcal{I}' are fresh with respect to all indexes already present in the target DIOC². 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.

Remark 1. *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*

²We do not give a formal definition of freshIndexes(\mathcal{I}') to keep the presentation simple. However, freshness of indexes can be formally ensured using restriction as in π -calculus [5].

the non-determinism using suitable policies to decide when and whether a given update applies. One of such policies is defined in AIOCJ (see § 5.1).

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 4 (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 $\langle \Sigma_1, \mathbf{I}_1, \mathcal{I}_1 \rangle$.

4.3 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 § 5.1) we implement the DPOC constructs in the Jolie [88] language.

4.3.1 DPOC syntax

DPOCs include *processes*, ranged over by P, P', \dots , describing the behaviour of participants. $(P, \Gamma)_R$ denotes a DPOC role named 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 *public operations* to mark the difference with respect to *private operations*, also called auxiliary operations, ranged over by o^* . We use $o^?$ to range over both public and private operations. Differently from communications on public operations, communications on private operations have no direct correspondent at the DIOC level. Indeed, we

4.3. Dynamic Process-Oriented Choreographies

introduce private 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 § 4.4, and are described in detail there. We range over DPOC indexes with ι .

In DPOCs, normal indexes are also used to prefix the operations of sends and receives. 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 between different communications, in particular when one of them comes from an update. We will describe in greater detail this issue later on.

The syntax of DPOCs is the following.

$$\begin{array}{l}
 P ::= \iota: \mathbf{i}.o^? : x \text{ from } R \text{ (receive)} \quad | \quad \mathbf{1} \quad \text{(inaction)} \\
 \quad | \iota: \mathbf{i}.o^? : e \text{ to } R \quad \text{(send)} \quad | \quad \mathbf{0} \quad \text{(end)} \\
 \quad | \mathbf{i}: \mathbf{i}.o^* : X \text{ to } R \quad \text{(send-update)} \quad | \quad \mathbf{i}: \text{if } b \{P\} \text{ else } \{P'\} \quad \text{(conditional)} \\
 \quad | P; P' \quad \text{(sequence)} \quad | \quad \mathbf{i}: \text{while } b \{P\} \quad \text{(while)} \\
 \quad | P \mid P' \quad \text{(parallel)} \quad | \quad \mathbf{i}: \text{scope } @R \{P\} \text{ roles } \{S\} \text{ (scope-coord)} \\
 \quad | \iota: x = e \quad \text{(assignment)} \quad | \quad \mathbf{i}: \text{scope } @R \{P\} \quad \text{(scope)} \\
 \\
 X ::= \text{no} \quad | \quad P \quad \quad \mathcal{N} ::= (P, \Gamma)_R \quad | \quad \mathcal{N} \parallel \mathcal{N}'
 \end{array}$$

DPOC processes include receive action $\iota: \mathbf{i}.o^? : x \text{ from } R$ on a specific operation $\mathbf{i}.o^?$ (either public or private) of a message from role R to be stored in variable x , send action $\iota: \mathbf{i}.o^? : e \text{ to } R$ of the value of an expression e to be sent to role R , and higher-order send action $\mathbf{i}: \mathbf{i}.o^* : X \text{ to } R$ of the higher-order argument X to be sent to role R . Here X may be either a DPOC process P , which is the new code for a scope in R , or a token no , notifying that no update is needed. $P; P'$ and $P \mid P'$ 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 $\mathbf{i}: \text{scope } @R \{P\} \text{ roles } \{S\}$ defines a scope with body P and set of participants S , and may occur only inside role R , which acts as coordinator of the update. The shorter version $\mathbf{i}: \text{scope } @R \{P\}$ is used instead when the role R is not the coordinator of the update. In fact, only the coordinator needs to know the set S of involved roles to be able to send to them their updates.

4.3.2 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 5 (DPOC roles semantics). *The semantics of DPOC roles is defined as the smallest LTS closed under the rules in Figure 4.5, where we report the rules dealing with computation, and Figure 4.6, 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. Possible values for δ are as follows:

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

The semantics is in the early style. Rule $\llbracket \text{DPOC} \mid_{\text{RECV}} \rrbracket$ receives a value v from role S and assigns it to local variable x of R . Similarly to Rule $\llbracket \text{DPOC} \mid_{\text{INTERACT}} \rrbracket$ (see Figure 4.3), the reception reduces to an assignment that inherits the index \mathbf{i} from the receive primitive.

Rules $\llbracket \text{DPOC} \mid_{\text{SEND}} \rrbracket$ and $\llbracket \text{DPOC} \mid_{\text{SEND-UP}} \rrbracket$ execute send and higher-order send actions, respectively. Send actions evaluate expression e in the local state Γ . Rule $\llbracket \text{DPOC} \mid_{\text{ONE}} \rrbracket$ terminates an empty process. Rule $\llbracket \text{DPOC} \mid_{\text{ASSIGN}} \rrbracket$ executes an assignment ($[v/x]$ represents the substitution of value v for variable x). Rule $\llbracket \text{DPOC} \mid_{\text{SEQUENCE}} \rrbracket$ executes a step in the first process of a sequential composition, while Rule $\llbracket \text{DPOC} \mid_{\text{SEQ-END}} \rrbracket$ acknowledges the termination of the first process, starting the second one. Rule $\llbracket \text{DPOC} \mid_{\text{PARALLEL}} \rrbracket$ allows a process in a parallel composition to compute, while Rule $\llbracket \text{DPOC} \mid_{\text{PAR-END}} \rrbracket$ synchronises the termination of two of parallel processes. Rules $\llbracket \text{DPOC} \mid_{\text{IF-THEN}} \rrbracket$ and $\llbracket \text{DPOC} \mid_{\text{IF-ELSE}} \rrbracket$ select the “then” and the “else” branch in a conditional, respectively. Rules $\llbracket \text{DPOC} \mid_{\text{WHILE-UNFOLD}} \rrbracket$ and $\llbracket \text{DPOC} \mid_{\text{WHILE-EXIT}} \rrbracket$ model respectively the unfolding and the termination of a loop.

The rules reported in Figure 4.6 deal with code updates. Rules $\llbracket \text{DPOC} \mid_{\text{LEAD-UP}} \rrbracket$ and $\llbracket \text{DPOC} \mid_{\text{LEAD-NOUP}} \rrbracket$ specify the behaviour of the coordinator of the update, respectively when an update is performed and when no update is performed. Rules $\llbracket \text{DPOC} \mid_{\text{UP}} \rrbracket$ and $\llbracket \text{DPOC} \mid_{\text{NOUP}} \rrbracket$ are the corresponding rules for other roles. The rules exploit private operations sb^* and se^* to coordinate the beginning and the end of the update, respectively. Communications on sb^* also carry the new code if the

4.3. Dynamic Process-Oriented Choreographies

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

Figure 4.5: DPOC role semantics. Computation rules. (Update rules in Figure 4.6)

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

Figure 4.6: DPOC role semantics. Update rules. (Computation rules in Figure 4.5)

update is performed, and a token no otherwise. Communications on se^* carry no relevant data: they are used for synchronisation purposes only.

Rule $\text{[DPOC |}_{\text{LEAD-UP}}\text{]}$ models the fact that the coordinator R of the update applies an update \mathcal{I} . The premises of Rule $\text{[DPOC |}_{\text{LEAD-UP}}\text{]}$ are similar to those of Rule $\text{[DIOC |}_{\text{UP}}\text{]}$ (see Figure 4.3). 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 10 in § 4.6).

It is important that the decision on whether to update or not is taken by the unique coordinator R for two reasons. First, R ensures that all involved roles agree on whether to update or not. Second, since the set of updates may change at any time, the choice of the update inside \mathbf{I} needs to be atomic, and this is guaranteed using a unique coordinator. The coordinator R also generates the processes to be executed by the roles in S using the process-projection function π (detailed in § 4.4). More precisely, $\pi(\mathcal{I}, R_i)$ generates the code for role R_i . The processes $\pi(\mathcal{I}, R_i)$ are sent via auxiliary higher-order communications 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 $R_i \in S$.

After the communication of the updated code to the other participants, R starts

4.3. Dynamic Process-Oriented Choreographies

its own updated code $\pi(\mathcal{I}, R)$. 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). Summarising, during scope execution auxiliary communications ensure that the update is performed in a coordinated way, i.e., the roles agree on when the scope starts and terminates and on whether the update is performed or not.

Rule $\llbracket \text{DPOC} \mid_{\text{LEAD-NOUP}} \rrbracket$ defines the behaviour of the coordinator R when no update is applied. In this case, R sends a token no to the other involved roles, notifying them that no update is applied and that they can start executing their code. End of scope synchronisation follows that of Rule $\llbracket \text{DPOC} \mid_{\text{LEAD-UP}} \rrbracket$.

Rules $\llbracket \text{DPOC} \mid_{\text{UP}} \rrbracket$ and $\llbracket \text{DPOC} \mid_{\text{NOUP}} \rrbracket$ 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 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.

DPOC system semantics.

Definition 6 (DPOC systems semantics). *The semantics of DPOC systems is defined as the smallest LTS closed under the rules in Figure 4.7. Symmetric rules for parallel composition have been omitted.*

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

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

Rules $\llbracket \text{DPOC} \mid_{\text{LIFT}} \rrbracket$ and $\llbracket \text{DPOC} \mid_{\text{LIFT-UP}} \rrbracket$ lift role transitions to the system level. Rule $\llbracket \text{DPOC} \mid_{\text{LIFT-UP}} \rrbracket$ also checks that the update \mathcal{I} belongs to the set of currently available updates \mathbf{I} . Rule $\llbracket \text{DPOC} \mid_{\text{SYNCH}} \rrbracket$ synchronises a send with the corresponding receive, producing an interaction. Rule $\llbracket \text{DPOC} \mid_{\text{SYNCH-UP}} \rrbracket$ is similar, but it deals with higher-order interactions. Note that Rules $\llbracket \text{DPOC} \mid_{\text{SYNCH}} \rrbracket$ and $\llbracket \text{DPOC} \mid_{\text{SYNCH-UP}} \rrbracket$ remove the prefixes from DPOC operations in transition labels. The labels of these transitions store the information on the occurred communication: label $o^? : R_1(v) \rightarrow R_2(x)$ denotes an interaction on operation $o^?$ from role R_1 to role R_2 where the value v is sent by R_1 and then stored by R_2 in variable x . Label $o^* : R_1(X) \rightarrow R_2()$ denotes a similar interaction, but concerning a higher-order

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

Figure 4.7: DPOC system semantics.

4.4. Projection Function

value X , which can be either the code used in the update or a token 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 $\lfloor^{\text{DPOC}} \rfloor_{\text{EXT-PAR}}$ allows a network inside a parallel composition to compute. Rule $\lfloor^{\text{DPOC}} \rfloor_{\text{EXT-PAR-END}}$ synchronises the termination of parallel networks. Finally, Rule $\lfloor^{\text{DPOC}} \rfloor_{\text{CHANGE-UPDATES}}$ 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.

Definition 7 (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 private 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 2. *Contrarily to DIOCs, DPOCs can deadlock. For instance,*

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

is a deadlocked DPOC network where processes are 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.

4.4 Projection Function

We now introduce the projection function proj . Given a DIOC specification, proj returns a network of DPOC programs that, by interacting, 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 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 8 (Projection). *The projection of a DIOC process \mathcal{I} with global state Σ is the DPOC network defined by:*

$$\text{proj}(\mathcal{I}, \Sigma) = \parallel_{s \in \text{roles}(\mathcal{I})} (\pi(\mathcal{I}, \mathcal{S}), \Sigma_s)_s$$

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

Definition 9 (Process-projection). *Given an annotated DIOC process \mathcal{I} and a role R the projected DPOC process $\pi(\mathcal{I}, R)$ is defined as in Figure 4.8.*

With 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_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, **1** and **0**. The projection of an assignment is the assignment on the role performing it and **1** on other roles. The projection of an interaction is a send on the sender role, a receive on the receiver, and **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 **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 private operations. In particular, cmd_i^* coordinates the branching of conditionals, carrying information on whether the “then” or the “else” branch needs to be taken. Similarly, wb_i^* coordinates the beginning of a while loop, carrying information on whether to loop or to exit. Finally, we_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 i : `if $b@R$ { \mathcal{I} } else { \mathcal{I}' }`, the coordinator R of the conditional locally evaluates the guard and tells the other roles which branch to choose using auxiliary communications on cmd_i^* . Finally, all the roles involved in the conditional execute their code corresponding to the chosen branch. Execution of a loop i : `while $b@R$ { \mathcal{I} }` is similar, with two differences. First, end of loop synchronisation on operations we_i^* is used to agree on when an iteration is terminated, and a new one can be started. Second, communication of whether to loop or to

4.4. Projection Function

$$\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}_{\mathbf{T}}: \mathbf{i}.cnd_{\mathbf{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}_{\mathbf{F}}: \mathbf{i}.cnd_{\mathbf{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}_{?}: \mathbf{i}.cnd_{\mathbf{i}}^* : x_{\mathbf{i}} \text{ from } \mathbf{R}; \mathbf{i}: \text{if } x_{\mathbf{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}_{\mathbf{T}}: \mathbf{i}.wb_{\mathbf{i}}^* : \text{true to } \mathbf{R}' \right); \pi(\mathcal{I}, \mathbf{R}); \right. \\ \prod_{\mathbf{R}' \in \text{roles}(\mathcal{I}) \setminus \{ \mathbf{R} \}} \mathbf{i}_{\mathbf{C}}: \mathbf{i}.we_{\mathbf{i}}^* : _ \text{ from } \mathbf{R}' \\ \left. \right\}; \prod_{\mathbf{R}' \in \text{roles}(\mathcal{I}) \setminus \{ \mathbf{R} \}} \mathbf{i}_{\mathbf{F}}: \mathbf{i}.wb_{\mathbf{i}}^* : \text{false to } \mathbf{R}' \\ \mathbf{i}_{?}: \mathbf{i}.wb_{\mathbf{i}}^* : x_{\mathbf{i}} \text{ from } \mathbf{R}; \\ \mathbf{i}: \text{while } x_{\mathbf{i}} \left\{ \begin{array}{l} \pi(\mathcal{I}, \mathbf{S}); \\ \mathbf{i}_{\mathbf{C}}: \mathbf{i}.we_{\mathbf{i}}^* : \text{ok to } \mathbf{R}; \\ \mathbf{i}_{?}: \mathbf{i}.wb_{\mathbf{i}}^* : x_{\mathbf{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} \in \text{roles}(\mathcal{I}) \setminus \{ \mathbf{R} \} = \left\{ \begin{array}{l} \mathbf{i}_{?}: \mathbf{i}.wb_{\mathbf{i}}^* : x_{\mathbf{i}} \text{ from } \mathbf{R}; \\ \mathbf{i}: \text{while } x_{\mathbf{i}} \left\{ \begin{array}{l} \pi(\mathcal{I}, \mathbf{S}); \\ \mathbf{i}_{\mathbf{C}}: \mathbf{i}.we_{\mathbf{i}}^* : \text{ok to } \mathbf{R}; \\ \mathbf{i}_{?}: \mathbf{i}.wb_{\mathbf{i}}^* : x_{\mathbf{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 4.8: process-projection function π .

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 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 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, inside the body of the loop during previous iteration 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 1. *We illustrate the issue of interferences using the two DPOC processes below, identified by their respective roles, R_1 (right) and R_2 (left), assuming that operations are not prefixed by indexes. We describe only R_1 as R_2 is its dual. At Line 1, R_1 sends a message to R_2 on operation o . In parallel with the send, R_1 had a scope (Lines 3–5) that performed an update. The new code (Line 4) contains a send on operation o for role R_2 . Since the two sends and the two receives share the same channel (o) and run in parallel, they can interfere with each other.*

<p><i>process</i> R_1</p> <ol style="list-style-type: none"> 1. 1: $o : e_1$ to R_2 2. 3. // update auxiliary code 4. 2: $o : e_2$ to R_2 5. // update auxiliary code 	<p><i>process</i> R_2</p> <ol style="list-style-type: none"> 1. 1: $o : x_1$ from R_2 2. 3. // update auxiliary code 4. 2: $o : x_2$ from R_2 5. // update auxiliary code
---	---

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

For this reason, in § 4.3 we introduced prefixes for 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 private operations using the index 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.

4.5. Running Example: Projection and Execution

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

Provided 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 i_T , the ones for selecting the “else” branch with index i_F ; *ii*) in the projection of the other roles involved in the conditional we assign the index $i_?$ to the auxiliary receive communications. To communicate the evaluation of the guard of a while loop we use the same indexing scheme (i_T , i_F , and $i_?$) used in the projection of conditional. Moreover, all the auxiliary communications for end of loop synchronisation are indexed with i_C .

4.5 Running Example: Projection and Execution

In this section we use our running example (see Figure 4.2) to illustrate the projection and execution of DIOC programs.

4.5.1 Projection

Given the code in Figure 4.2, we need to annotate it to be able to project it (we remind that in § 4.4 we defined our projection function on well-annotated DIOCs). Since we wrote one instruction per Line in Figure 4.2, 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 4.9, 4.11, and 4.10. To improve readability, we omit some 1 processes. In the projection of the program, we also omit to write the prefix operations since it is always equal to the numeric part of the index of their correspondent construct. Finally, we write auxiliary communications in grey.

4.5.2 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 4.2. In this execution scenario we assume to introduce in the set of available updates the update presented in Figure 4.1, 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 1s to improve readability.

```

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?:  $cnd_{15}^*$ :  $x_{15}$  from Buyer;
12 15: if ( $x_{15}$ ){
13   16:  $payReq$ :  $payDesc(order\_price)$  to Bank;
14   22?:  $cnd_{22}^*$ :  $x_{22}$  from Bank;
15   22: if ( $x_{22}$ ){
16     23:  $confirm$ : _ from Bank}
17 }

```

Figure 4.9: Seller DPOC Process.

```

1  15?:  $cnd_{15}^*$ :  $x_{15}$  from Buyer;
2  15: if ( $x_{15}$ ){
3    16:  $payReq$ :  $desc$  from Seller;
4    17: scope @Bank{
5      18:  $payment\_ok$  = true;
6      19:  $pay$ :  $auth$  from Buyer;
7      // code for the payment
8    } roles { Buyer, Bank };
9    22: if ( $payment\_ok$ ){
10     {22T:  $cnd_{22}^*$ : true to Seller
11      | 22T:  $cnd_{22}^*$ : true to Buyer};
12     { 23:  $confirm$ :  $null$  to Seller
13      | 25:  $confirm$ :  $null$  to Buyer
14     }
15   } else {
16     {22F:  $cnd_{22}^*$ : false to Seller
17      | 22F:  $cnd_{22}^*$ : false to Buyer};
18     27:  $abort$ :  $null$  to Buyer
19   }
20 }

```

Figure 4.10: 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:  $cnd_{15}^*$ : true to Seller
20     | 15T:  $cnd_{15}^*$ : true to Bank};
21   17: scope  $payment$ @Bank{
22     19:  $pay$ :  $payAuth(prod\_price)$  to Bank
23     // code for the payment
24   }
25 };
26 22?:  $cnd_{22}^*$ :  $x_{22}$  from Bank;
27 22: if ( $x_{22}$ ){
28   25:  $confirm$ : _ from Bank
29 } else {
30   27:  $abort$ : _ from Bank}
31 }

```

Figure 4.11: Buyer DPOC Process.

4.5. Running Example: Projection and Execution

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 4.2). 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 $[^{\text{DIOC}} |_{\text{UP}}]$ 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 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 case, which update to apply. Here, we assume that the update using the code in Figure 4.1 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_B and reported below on the right — obtained projecting the update on role Buyer.

```

1 6:  $sb_6^*$  :  $P_B$  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 =
6       getPrice(order) * 0.9
7 } else {
8   37: order_price = getPrice(order)
9 };
10 39: offer : order_price to Buyer;
11 6:  $se_6^*$  : _ from Buyer;

```

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

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 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 $[^{DPOC} |_{UP}]$ the DPOC fragment P_B sent by the coordinator. In the reductum, P_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:  $se_6^*$  : ok to Seller

```

Consider now what happens if no update is applied. At DIOC level the Seller applies Rule $[^{DIOC} |_{NoUP}]$, 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 = getPrice(order);
3 8: offer : order_price to Buyer;
4 6:  $se_6^*$  : _ from Buyer;

```

4.6. Connected DIOCs

Before executing the code inside the scope, the Seller has to notify the other roles that they can proceed with their 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 $\lfloor^{\text{DPOC}} \lfloor_{\text{NoUP}} \rfloor$ removes the scope and adds the notification of termination (Line 6) to the coordinator at the end.

```

1 7: offer : prod_price from Seller;
2 6: se_6^* : ok to Seller;

```

4.6 Connected DIOCs

We now give a precise definition of the notion of connectedness that we mentioned in § 4.2.2 and § 4.3.2. In both DIOC and DPOC semantics we checked such a property of updates with predicate connected, respectively in Rules $\lfloor^{\text{DIOC}} \lfloor_{\text{UP}} \rfloor$ (Figure 4.3) and $\lfloor^{\text{DPOC}} \lfloor_{\text{UP-LEAD}} \rfloor$ (Figure 4.6).

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).

$$\begin{array}{l}
 op_1 : A(a) \rightarrow B(b); \\
 op_2 : C(c) \rightarrow D(d)
 \end{array}
 \xrightarrow{\text{projection}}
 \begin{array}{c|c}
 \text{process A} & \text{process B} \\
 \hline
 op_1 : a \text{ to B} & op_1 : b \text{ from A} \\
 \hline
 \text{process C} & \text{process D} \\
 \hline
 op_2 : c \text{ to D} & op_2 : d \text{ from C}
 \end{array}$$

DIOCs can express interactions that, if projected as described in § 4.4, 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 originating DIOC.

Restricting to connected DIOCs avoids this kind of behaviours. We formalise connectedness as an efficient (see Theorem 1) syntactic check. More precisely, we check for *connectedness (for sequence)*, because, intuitively, we make sure that the DPOC network obtained by projecting a sequence $\mathcal{I}; \mathcal{I}'$ executes first the actions in \mathcal{I} and then those in \mathcal{I}' , thus respecting the intended semantics of sequential composition and avoiding undesired behaviours. We highlight that our definition of connectedness does not hamper programmability and it naturally holds in most of real-world scenarios (the interested reader can find at the website [89] several of such scenarios).

Remark 3. *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 that do not produce non-deterministic projections [11]; ii) by mimicking the non-deterministic behaviour of process-level networks at choreography level [90]; or iii) by enforcing the order of interaction with additional auxiliary messages between roles [91].*

Our choice of preserving the order of interactions defined at DIOC level follows the same philosophy of [11], whilst for scopes, conditionals, and while loops we enforce connectedness with auxiliary messages as done in [91]. We remind that we introduced auxiliary messages for coordination in the semantics of scopes at DPOC level (§ 4.3.2) and in the projection (§ 4.4). We choose to add such auxiliary message 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 interactions we choose to follow a stricter approach by preserving the order of interactions defined at DIOC level.

To formalise connectedness we introduce, in Figure 4.12, the auxiliary functions transI 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 $R \rightarrow S$. Actions located at R are represented as $R \rightarrow R$. For instance, given an interaction $i: o: R(e) \rightarrow S(x)$ both its transI and transF are $\{R \rightarrow S\}$. For conditional, $\text{transI}(i: \text{if } b@R \{ \mathcal{I} \} \text{ else } \{ \mathcal{I}' \}) = \{R \rightarrow R\}$ since the first action executed is the evaluation of the guard by role R . The set $\text{transF}(i: \text{if } b@R \{ \mathcal{I} \} \text{ else } \{ \mathcal{I}' \})$ is normally $\text{transF}(\mathcal{I}) \cup \text{transF}(\mathcal{I}')$, since the execution terminates with an action from one of the branches. If instead the branches are both empty then transF is $\{R \rightarrow R\}$, representing guard evaluation.

Finally, we give the formal definition of connectedness.

4.6. Connected DIOCs

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

Figure 4.12: Auxiliary functions transl and transF.

Definition 10 (Connectedness). *A DIOC process \mathcal{I} is connected if each subterm \mathcal{I}' ; \mathcal{I}'' of \mathcal{I} satisfies*

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

Connectedness can be checked efficiently.

Theorem 1 (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} .

The proof of the theorem is reported in Appendix A.1.

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.

4.7 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 this work, namely the correctness of the projection. Moreover, as a consequence of the correctness, we prove that properties like deadlock-freedom, termination, and race-freedom are preserved by the projection.

Correctness here means that the weak traces of a connected DIOC coincide with the weak traces of the projected DPOC.

Definition 11 (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 proof strategy to prove our main result relies on *i*) the definition of a notion of bisimilarity which implies weak trace equivalence and *ii*) the definition of a suitable bisimulation relating each well-annotated connected DIOC system with its projection.

Definition 12 (Weak System Bisimulation). *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 $\eta \in \{o^? : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x); o^* : \mathbf{R}_1(X) \rightarrow \mathbf{R}_2(); \sqrt{}; \mathcal{I}; \text{no-up}; \mathbf{I}''', \tau\}$ 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 $(\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle, \langle \mathbf{I}''', \mathcal{N}''' \rangle) \in \mathcal{R}$

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

Theorem 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 trace equivalent.*

We report the full proof of Theorem 2 in Appendix A.2.

Properties. Since the projection preserves weak traces, we have that trace-based properties of the DIOC are inherited by the DPOC. A first examples of such properties is *deadlock freedom*.

Definition 13 (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 \surd 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 \surd only rule updates are possible is captured by the following lemma.

Lemma 1. *For each initial, connected DIOC \mathcal{I} , state Σ , and set of updates \mathbf{I} , if $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle \xrightarrow{\surd} \langle \Sigma', \mathbf{I}', \mathcal{I}' \rangle$ then the only transitions of $\langle \Sigma', \mathbf{I}', \mathcal{I}' \rangle$ have label \mathbf{I}'' for some \mathbf{I}'' .*

Proof. The proof is by case analysis on the rules which can derive a transition with label \surd . 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 1 (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 at least a transition is derivable.

Case 0 not allowed since we assumed an initial DIOC.

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

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

Case $o' : R_1(e) \rightarrow R_2(x)$ 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 $[\text{DIOC}_{\text{SEQ-END}}]$. 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 1, since the label is \surd .

Case $\text{if } b@R \{ \mathcal{I} \} \text{ else } \{ \mathcal{I}' \}$ 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 $\text{while } b@R \{ \mathcal{I} \}$ 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 $\text{scope } @R \{ \mathcal{I} \}$ 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 2, thus the finite weak internal traces 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 1. \square

DPOCs also inherit termination from terminating DIOCs.

Definition 14 (Termination). *A DIOC (resp. DPOC) system terminates if all its internal traces are finite.*

Note that with arbitrary sets of updates DIOC termination is never granted if they contain at least a scope, since it 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 2 (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 proprieties 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 R_1 in role R_2 may interact with the sends of the form $\iota: \mathbf{i}.o^? : e$ to R_2 in role R_1 , races may happen. However, thanks to the correctness of the projection, race freedom holds also for the projected DPOCs.

Corollary 3 (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 $\eta_i \in \{\tau, o^? : R_1(v) \rightarrow R_2(x), o^* : R_1(X) \rightarrow R_2(), \sqrt{\cdot}, \mathcal{I}, \text{no-up}, \mathbf{I}\}$ 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, respectively for public and private operations.

For public operations, thanks to Lemma 7, 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 private 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 they 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 4 (Orphan-message freedom). *For each initial, connected DIOC \mathcal{I} , state Σ , and set of updates \mathbf{I} , if $\langle \mathbf{I}, \text{proj}(\mathcal{I}, \Sigma) \rangle \xrightarrow{\eta_1} \dots \xrightarrow{\eta_m} \xrightarrow{\surd} \langle \mathbf{I}', \mathcal{N} \rangle$, where $\eta_i \in \{\tau, o^? : R_1(v) \rightarrow R_2(x), o^* : R_1(X) \rightarrow R_2(), \surd, \mathcal{I}, \text{no-up}, \mathbf{I}\}$, 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

Adaptable Interaction-Oriented Choreographies in Jolie

5.1 Introduction

In this section, we present AIOCJ (Adaptable Interaction-Oriented Choreographies in Jolie), a development framework for adaptable distributed applications [55]. AIOCJ is one of the possible instantiations of our theory of Dynamic Choreographies presented in Chapter 4 and it gives a tangible proof of the expressiveness and feasibility of our approach, enabling adaptation of distributed programs (cf. Chapter 3).

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 system 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, and *iii*) it is possible to add and remove adaptation rules at runtime.

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 § 5.1.1. AIOCJ offers an integrated environment for developing programs and adaptation rules that supports syntax highlighting and

real-time¹ syntax checking. Since checking for *connectedness* (see § 4.6) of programs and adaptation rules is polynomial (as proven by Theorem 1), the IDE also performs real-time 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 § 4.4. The target language of the AIOCJ compiler is the Jolie [92] language, which sports primitives similar to those of our DPOC language. Given an AIOC program, the AIOCJ compiler produces one Jolie program, also called 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, stored in the Environment service.

In the remainder of this Section we detail the grammar of the AIOC language used by AIOCJ in § 5.1.1, we illustrate the use of AIOCJ on a simple example in § 5.2, we discuss some relevant implementation aspects of AIOCJ in § 5.3, and we provide a preliminary validation in § 5.4.

5.1.1 DIOC Language Extensions in AIOCJ

AIOCs extend DIOCs with:

¹In this context, real-time means “while the developer is writing AIOC programs and rules”.

5.1. Introduction

- 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 Extended Backus-Naur Form [93], 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 ID [, ID]* 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 `"HTTP"` protocol we enable its use with the following inclusion:

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

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. In order for the theory to apply they need to satisfy the condition of never blocking and always returning a value (possibly an error notification), exactly as other expressions.

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 use 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 @ID { Choreography }
  [ prop { Properties } ]?
```

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

For instance, the code

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

specifies that the scope has a property name set to the string "myScope", thus satisfying the condition of the adaptation rule 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 compu-

tation. Any role can be chosen as `starter`, but the chosen one needs to be started 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@ID:"URL"
```

where `ID` 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.

5.2 AIOCJ Practice

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 5.1.

At Line 1 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 3–6, where we declare a scope. At Line 6 we define a property `scope_name` of this scope with value `"price_inquiry"`.

Figure 5.2 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 5.1.

Since we want to replace the scope for the `price_inquiry`, we define 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 body of the adaptation rule specifies the same behaviour described in Fig-

```
1 preamble{ starter: Seller }
2 aioc {
3   scope @Seller {
4     order_price@Seller = getPrice( order );
5     offer: Seller( order_price ) -> Buyer( prod_price )
6   } prop { N.scope_name = "price_inquiry" }

```

```
1 rule {
2   on { N.scope_name == "price_inquiry" and E.season == "Fall" }
3   do {
4     cardReq: Seller( _ ) -> Buyer( _ );
5     card_id@Buyer = getInput( "Insert your customer card ID");
6     cardRes: Buyer( card_id ) -> Seller( buyer_id );
7     if isValid( buyer_id )@Seller {
8       order_price@Seller = getDiscountedPrice( order, buyer_id );
9     } else {
10      order_price@Seller = getPrice( order )
11    };
12    offer: Seller( order_price ) -> Buyer( prod_price )
13  }
14 }
```

Figure 5.1: An AIOC program (upper part) and an applicable adaptation rule (lower part).

ure 4.1 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 and the execution of the adaptation at point ③ of Figure 5.2 (outlined with dashes). From right to left, we write the rule and we compile it. The compilation of a (set of) adaptation rule(s) in AIOCI produces a service, called *Adaptation Server*, 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, in order 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"`.

5.3 Implementation

AIOCI is composed of two elements: the AIOCI Integrated Development Environment (IDE), named `AIOCI – ecl`, and the adaptation middleware that enables

5.3. Implementation

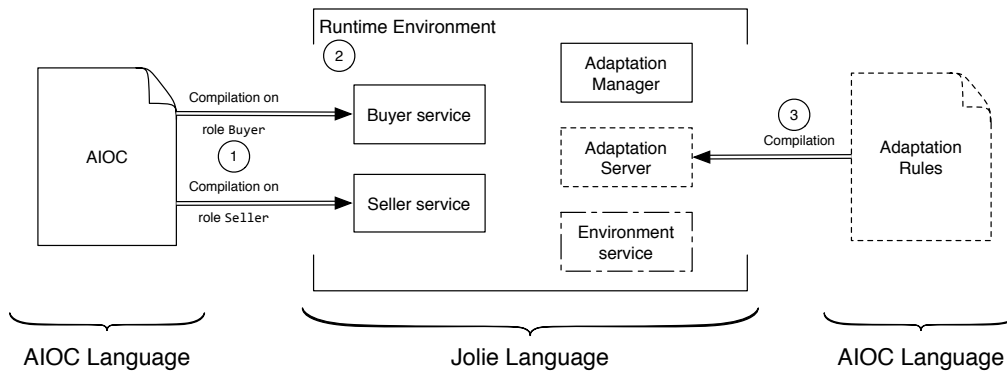


Figure 5.2: Representation of the AIOCJ framework — Projection and execution of the example in Figure 5.1.

AIOC programs to adapt, called AIOCJ – mid.

AIOCJ – ecl is a plug-in for Eclipse [94] based on Xtext [95]. 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. § 4.6) it is efficient enough to be performed on-the-fly while editing the code. Figure 5.3 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

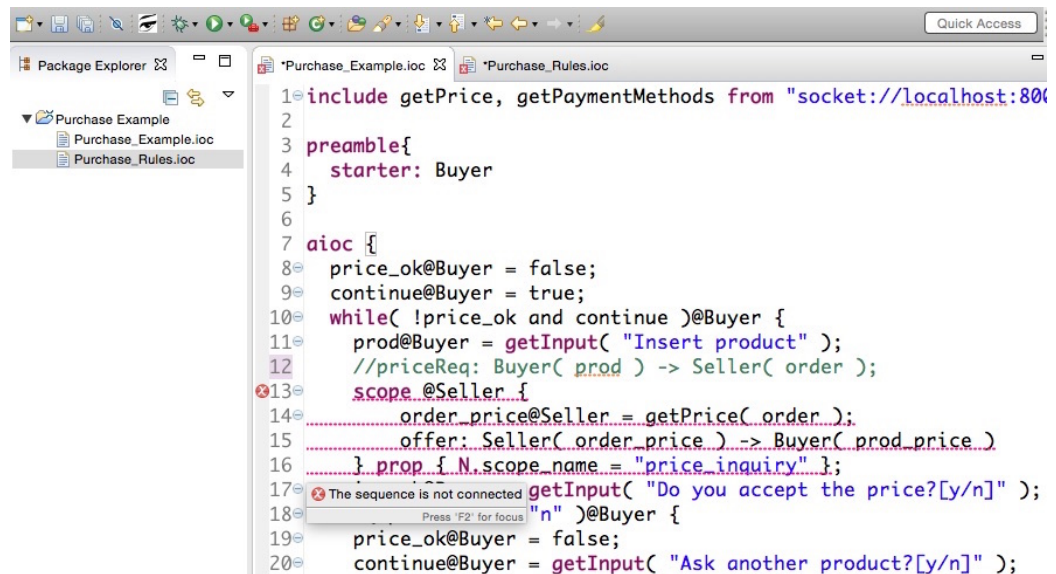


Figure 5.3: Check for connectedness: sequence.

the old one.

For reason of performances, the implementation differs from the theory in two respects. First, adaptation rules are compiled statically and not by the coordinator of the update when the update is applied. Second, to ensure that operations coming from distinct constructs do not interfere — what is guaranteed in the theory by prefixing operations with fresh indexes — we statically check that each program and each adaptation rule never features two interactions using the same operation in parallel (we call this property *connectedness for parallel*), and we run AIOCs coming from different programs or updates in different namespaces, thus forbidding interferences between them.

5.4 Validation

In this section, we give a preliminary empirical validation of our implementation. The main aim is to test how our mechanisms for adaptation impact on performances.

In the literature, to the best of our knowledge, there is no approach to adaptation based on choreography programming. Thus, it is difficult to directly compare our results with other existing approaches. Moreover, we are not aware of any established benchmark to evaluate adaptive applications. For this reason, we tested AIOCIJ performances by applying it to two typical programming patterns: *pipes*

and *fork-joins*.

Since we are interested in studying the cost of adaptation, our scenarios contain minimal computation and are particularly affected by the overhead of the adaptation process. Clearly, the percentage of the overhead due to adaptation will be far lower in real scenarios, which are usually more computationally intensive. In the first scenario, we program a pipe executing n tasks (in a pipe, the output of task t_i is given as input to task t_{i+1} , for $i \in \{1, \dots, n - 1\}$). To keep computation to a minimum, each task simply computes the increment function.

In the fork-join scenario, n tasks are computed in parallel. Each task processes one character of a message of length n , shifting it by one position. The message is stored in an external service.

The code of both scenarios is in Appendix B.1.

To enable adaptation, each task is enclosed in a scope. We test both scenarios with an increasing number of tasks $n \in \{10, 20, \dots, 100\}$ to study how performances scale as the number of adaptation scopes increases. We evaluate performances in different contexts, thus allowing us to understand the impact of different adaptation features, such as scopes, adaptation servers, and adaptation rules.

Context 1 no scopes, no adaptation servers, no rules;

Context 2 each task is enclosed in a scope, no adaptation servers, no rules;

Context 3 each task is enclosed in a scope, one adaptation server, no rules;

Context 4 as Context 3, but now the adaptation server contains 50 rules. Each rule is applicable to a unique scope i , and no rule is applicable to scopes with $i > 50$. The rules are stored in random order.

Context 5 as Context 4, but with 100 rules, one for each scope.

Each rule in Contexts 4 and 5 is applicable to one specific scope only (through a unique property of the scope), hence when testing for 50 rules, only the first 50 scopes adapt.

In order to reduce variability of performances, we repeated every test 5 times. We performed our tests on a machine equipped with a 2.6GHz quad-core Intel Core i7 processor and 16GB RAM. The machine runs Mavericks 10.9.3, Java 1.7.55, and Jolie r.2728. Figure 5.4 shows the tests for the pipe (left) and the fork-join (right). Both charts display on the x-axis the number of tasks/scopes and on the y-axis the execution time in milliseconds.

As expected, in both scenarios there is a significant gap between Contexts 1 and 2. In words, the introduction of scopes has a strong effect on performances. The ratio is 1:13 for the pipe scenario and 1:5.5 for the fork-join scenario. This is due to the auxiliary communications needed to correctly execute a scope. The

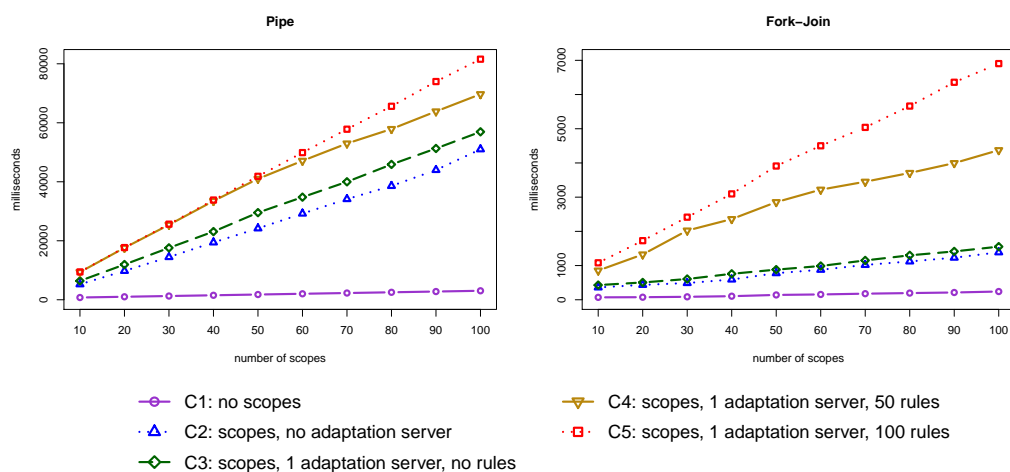


Figure 5.4: Times of execution of the pipe (left) and the fork-join (right) scenarios

observed overhead is higher in the pipe scenario, since different scopes check for adaptation in sequence, while this is done in parallel for the fork-join scenario.

Adding an adaptation server (from Context 2 to Context 3) has little impact on performances: 19% of decay for pipe, and 17% for fork-join. That is reasonable, considered that Context 3 just adds one communication wrt Context 2.

On the contrary, there is a notable difference when adding rules to the adaptation server (Context 4 is 1.4 times slower than Context 3 for the pipe scenario, 2.9 for the fork-join scenario). In Contexts 4 and 5, performances are really close up to 50 scopes (in the pipe scenario they almost overlap) although Context 5 has twice the rules of Context 4. This illustrates that the time to test for applicability of rules is negligible. Hence, the highest toll on performances is related to actual adaptation, since it requires to transfer and embed the new code. This is particularly evident in the fork-join scenario where multiple adaptations are executed in parallel and the adaptation server becomes a bottleneck. This problem can be mitigated using multiple distributed adaptation servers.

The fact that the most expensive operations are scope execution and actual adaptation is highlighted also by the results below. The table shows the cost of different primitives, including scopes in different contexts. Times are referred to 5 executions of the sample code in Appendix B.1.

As future work we will exploit these results to increase the performances of our framework, concentrating on the bottlenecks highlighted above. For instance, scope execution (as well as conditionals and cycles) currently requires many auxiliary communications ensuring that all the processes agree on the chosen path. In many cases, some of these communications are not needed, since a process will eventually discover the chosen path from the protocol communications. Static

5.4. Validation

Test	Time (ms)	Test	Time (ms)
assignment	2.2	scope, 1 adaptation server, 1 matching rule	280.6
interaction	4.2	scope, 1 adaptation server, 50 rules, none matching	254.2
if statement	16.6	scope, 1 adaptation server, 50 rules, 1 matching	338.6
scope, no adaptation server	129.4	scope, 1 adaptation server, 100 rules, none matching	310.2
scope, 1 adaptation server, no rule	203.8	scope, 1 adaptation server, 100 rules, 1 matching	385

analysis can discover redundant communications and remove them. Another improvement is letting the adaptation server send the new code directly to the involved roles, skipping the current forward chain.

Part III

Applied Choreographies

Applied Choreographies

There is nothing as practical as a good theory.

Kurt Z. Lewin

6.1 Introduction

Our work on Dynamic Choreographies in Chapter 4 and the implementation described in Chapter 5 proved that choreographies are a suitable model to tackle challenging and unsolved problems of distributed systems.

However, we highlight that in the development of AIOCJ, we took a sensible departure from the theoretical model of DIOCs. Indeed, in our theory we defined an EPP that targets the DPOC language, a process-level language that bases message passing on *names*, like DIOCs do. Synchronisation on names (as in CCS and the π -calculus [49, 5]) simplifies the theoretical treatment and it is used in many other works on choreographies as programming abstractions [10, 12, 43, 96, 13, 97]. Nevertheless, the compiler provided by AIOCJ targets the Jolie language [92, 50], which handles communications with *message correlation*, a standard technique in Service-Oriented Computing and Web Services that routes messages inspecting the data they carry (e.g., headers) [36, 98]. This makes the EPP provided by AIOCJ much more technically involved than its formal specification, including the management of underlying data structures (e.g., message queues) and unexpected — not present in the theoretical treatment — communications in the resulting executable code. We remark that a similar occurrence exists between the choreographic language Chor [99] and its theoretical model [13] and, more in general, that it is a typical characteristic of implementations of process calculi, see, e.g., [100, 52].

This key difference between formal models and implementations can compromise the benefits of choreographic programming: the correctness-by-construction

approach and the clear specification of the communications carried out during execution. Thus we ask:

How can we formalise the implementation of communications in choreographies?

A satisfactory answer should preserve the correctness-by-construction guarantees of choreography models down to the level of how communications are concretely implemented.

A challenging task that requires the definition of a model with *i*) the typical clarity of choreography languages and with *ii*) all the necessary details to formally reason on how to support communications at the lower level.

Our answer is to develop a theory of *Applied Choreographies* (AC). AC provides the simple and intuitive syntax of choreography languages along with the most advanced features of recent choreography models like modularity, asynchrony, parallelism, and the dynamic creation of multiparty sessions and processes. The key contribution of AC lies in its semantics: equipped with a novel notion of *deployment*, it lets us reason on how messages are routed between participants. Deployments separate the logic of choreographies (message passing, creation of sessions and processes) from their implementation. Notably, by just changing the definition of deployments and their effects we can easily formalise different implementation models, e.g., distributed objects [101], correlation, etc..

In this work, we chose to model correlation because clarifying the relationship between this communication model and choreographies is a relevant and challenging task, whose solution immediately yields a practical impact. Indeed, most practitioners use choreographies to describe correct distributed systems, but they have to deal with how channel-based interactions (e.g., as in WS-CDL [8]) can be correctly implemented with message correlation (used in most major frameworks, including WS-BPEL [36], Java/JMS, C#.NET, etc.).

Here, we model correlation on notions from Service-Oriented Computing, the setting where choreographies are used the most as design tool (cf. § 2.4). We pinpoint the key theoretical problems and formalise the guiding principles developers should follow to obtain correct implementations. On this result, we define a correct compiler from AC choreographies to the theoretical model behind the Jolie language.

Contributions. We list below our main contributions:

Applied Choreographies. AC is a choreography model for the modular development of choreographies based on asynchronous message passing. AC captures how communications among processes are concretely supported via message correlation (§ 6.2). The main novel aspect of AC is a notion of *deployment* for chore-

6.1. Introduction

ographies which, for each process in the choreography, keeps track of *i*) the location of the service in which it executes (multiple processes can run in the same service), *ii*) its input message queues (for asynchronous communications), and *iii*) its state (the value of its variables). This yields a close representation of how real-world service-oriented scenarios implement communications, which is the basis for all of our results.

Type system. We modelled real-world message passing in AC to formalise how programs may encounter execution errors. For example, it may be impossible to route a message from one endpoint to another, due to missing information in the deployment of a choreography (e.g., a suitable input queue at the receiver). This is a significant departure from previous choreography models, where it is assumed that communications between two processes can always be performed [10, 43, 12, 102, 96, 13] (cf. Chapter 4). In § 6.3 we present a typing discipline for AC that prevents communication errors. Our type language is that of standard multiparty session types [43]. This is a remarkable feature of AC: despite the additional complexity of its semantics, the complexity of the types a programmer needs to specify to check a choreography remain unchanged.

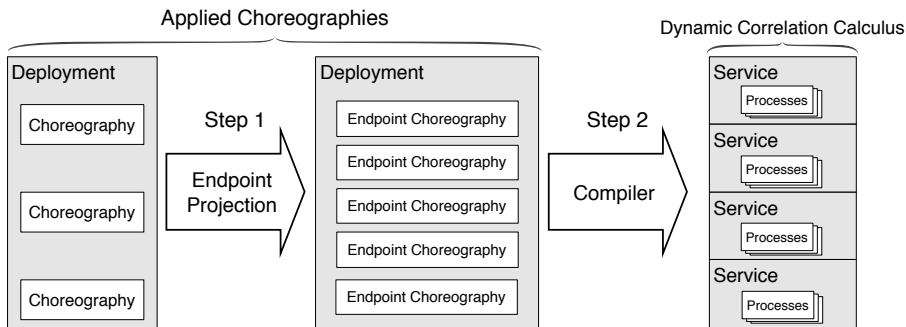


Figure 6.1: Overall methodology of Applied Choreographies.

Compilation of AC programs. We define a two-step methodology to transform a choreography in AC into programs of the *Dynamic Correlation Calculus* (DCC), a model for executable code based on message correlation. We depict our methodology in Figure 6.1. The first step is a source-to-source Endpoint Projection (as seen in [14]) which projects a choreography describing the behaviour of many participants to a series of choreography modules, called *Endpoint Choreographies*, each describing the behaviour of a single participant. The second step is a compilation from endpoint choreographies to DCC programs, which define the behaviour of *services* following the Service-Oriented model. Specifically, DCC formalises a syntax and a semantics for a fragment of the Jolie language. We chose Jolie because its reference implementation is based on a formal specification [98] called

Correlation Calculus (CC). DCC improves CC by adding message queues that can be created at runtime, a necessary feature to support our choreography model. We prove that the compiled processes implement the behaviour of the EPP and therefore, that of the originating AC program.

6.2 Applied Choreography Language

We introduce Applied Choreographies (AC) that, on key elements of choreography calculi like *processes*, *sessions*, and *roles*, introduces the notion of *locations*.

6.2.1 Syntax

In the syntax of AC (Figure 6.2) C denotes a choreography, p, q processes, A, B roles, k sessions, o operations, l locations, X procedures, and x variables¹. We consider all sets of identifiers disjoint. Processes are independent execution units that proceed in parallel. They can communicate with each other through sessions. Roles track which role each process plays in a session. As in standard multiparty session types [43] Roles are the basis for our typing discipline in § 6.3. Locations represent publicly reachable addresses, where we assume that an always-available service supports the creation and the execution of new processes. We introduce the notion of location to model the deployment of processes into services. Following the SOC model, a service is a container, reachable at a defined address (the location), where processes execute in parallel. In § 6.6, when compiling ACs to the lower language DCC, we map each location to a concrete service implementation. Using locations makes us depart from previous choreography languages, which do not consider the deployment of processes in their models. AC includes complete and partial actions to support modularity, as in previous choreography models [14]. A complete action specifies the behaviour of all participants involved in the action. A partial action describes the behaviour of only some participants. Partial actions enable compositionality: choreographies with compatible partial actions can be composed in parallel.

Complete Actions. Term (*start*) denotes session initiation: process p starts a new session k together with processes \tilde{q} . p , called *active process*, is already running, whereas each process q in $\tilde{l}.q$, called *service process*, is dynamically created at its respective location l in $\tilde{l}.q$. Service locations can be used repeatedly to

¹Variables are *paths* to traverse structured data, e.g., x can be a path $\underline{x.y.z}$ where “.” is the path nesting operator. § 6.2.2 formalises variables and paths.

6.2. Applied Choreography Language

spawn multiple processes, even inside recursions. We assume $\widetilde{l}.q$ always non-empty. Term (com) models a communication: on session k , process p sends to process q a message for operation o ; the message carries the evaluation of expression e on the local state of p whilst x is the variable where q will store the content of the message.

Partial Actions. In term (req) , process p requests the creation of some external processes at their respective locations \widetilde{l} to start a new session k . $\widetilde{l}.B$ means that each location l in $\widetilde{l}.B$ is expected to spawn a process that behaves as specified by the related role B . The dual of (req) is term (acc) , which provides the implementation of service processes. Specifically, term (acc) defines a reusable choreography module that accepts, at locations \widetilde{l} , the creation of processes \widetilde{q} playing their respective roles \widetilde{B} . Following the design idea that services should always be available, shared by other models [96, 13], we assume that all (acc) terms in a choreography are at top level (not guarded by other actions). Term $(send)$ models the sending of a message, for operation o , from process p to an external process playing role B in session k . Dually, in term $(recv)$, process q receives a message from an external process playing role A in session k ; q proceeds with the continuation corresponding to the operation where the message was received, e.g., receiving the message on o_j , q proceeds with continuation C_j — like in the standard branching found in session-oriented calculi [44].

Other Terms. In a conditional $(cond)$ process p evaluates a condition e in its local state to choose between the continuations C_1 and C_2 . Term (par) is the standard parallel composition of choreographies, as in [14, 97]. Terms (def) , $(call)$, and $(inact)$ denote respectively the standard definition of recursive procedures, procedure calls, and inaction. Some terms bind identifiers in continuations. In terms $(start)$ and (acc) , the session identifier k and the process identifiers \widetilde{q} are bound (as they are freshly created). All other identifiers are free. In terms (com) and $(recv)$, the variables used by the receiver to store the message are bound (x and all the x_i , respectively). In term (req) , the session identifier k is bound. Finally, in term (def) , the procedure identifier X is bound. In the remainder, we omit $\mathbf{0}$ or irrelevant variables (e.g., in communications with empty messages).

Remark 4. As in [13], AC terms, except for $(start)$, (req) , and (acc) , do not need to be annotated with roles as we can infer them from session identifiers. However, we include roles in all AC terms to ease the treatment of our typing discipline in § 6.3.

$$\begin{array}{l}
 C ::= \eta; C \quad (seq) \quad | \quad \text{if } p.e \{C_1\} \text{ else } \{C_2\} \quad (cond) \\
 \quad | \quad C_1 \mid C_2 \quad (par) \quad | \quad \text{def } X = C' \text{ in } C \quad (def) \\
 \quad | \quad X \quad (call) \quad | \quad \mathbf{acc} \ k : \widetilde{l.q[B]}; C \quad (acc) \\
 \quad | \quad \mathbf{0} \quad (inact) \quad | \quad k : A \rightarrow q[B].\{o_i(x_i); C_i\}_{i \in I} \quad (recv) \\
 \\
 \eta ::= \mathbf{start} \ k : p[A] \Leftrightarrow \widetilde{l.q[B]} \quad (start) \quad | \quad \mathbf{req} \ k : p[A] \Leftrightarrow \widetilde{l.B} \quad (req) \\
 \quad | \quad k : p[A].e \rightarrow q[B].o(x) \quad (com) \quad | \quad k : p[A].e \rightarrow B.o \quad (send)
 \end{array}$$

Figure 6.2: Choreography Calculus - Syntax

6.2.2 Semantics

The semantics of AC is one of our major contributions: it formally captures, on the level of choreographies, the real-world communication mechanism found in SOC, called message correlation [36]. We first give an informal overview of this communication mechanism. Communications in SOC are asynchronous: each process has a set of FIFO input queues that act as buffers, managed by its enclosing service. Each queue is equipped with some data, here called *correlation key*, that can be used to distinguish (identify) the queue among the many present inside a service. When a service receives a message from the network, it inspects the content of the message for a portion of data that matches the correlation key of one of its queues. If a queue can be found, the message is inserted at the end of it. The process owning the queue will be able to consume the message later on in its execution. Thus, when a sender process p sends a message to a receiver process q , it needs to know *i*) the location of the service where q is running and *ii*) the correlation key of one of the queues owned by q . In practice, the correlation key in the message can be a part of the message payload itself or be in some separate headers (in this work we abstract from such details). Below, we formalise the semantics of AC by equipping choreographies with *deployments*, ranged over by D . We use deployments to formalise the elements of SOC that we have just informally described, in particular the *state* and *message queues* of processes located at services. We define each element separately.

Data and Process state. Data in SOC is typically structured following a tree-like format, e.g., XML [27] and JSON [103]. In this work, we use trees to represent both messages and the state of running processes (as in [98]). Formally, we consider a set \mathcal{T} of rooted trees, ranged over by t , where edges are labelled by names, ranged over by $\underline{x}, \underline{y}, \dots$. We assume that all outgoing edges of a node have distinct labels and that only leafs contain values, which can be either a location l or some basic data (integer, string, etc.). Variables in AC, ranged over by x , are formalised as paths to traverse a tree: $x, y, z ::= \underline{x}.x \mid \varepsilon$. ε is the empty

6.2. Applied Choreography Language

path; we often omit the tailing ε in paths. Given a path x and a nonempty tree t , we denote by $x(t)$ the node reached following the path x in t . Observe that $x(t)$ is partially defined since the path x may not be valid in t in general. By a slight abuse of notation, when $x(t)$ is a leaf we denote by $x(t)$ also the value of the node. In our semantics, we will also use the replacement operator $t \triangleleft (x, t')$. If $x(t)$ is defined, $t \triangleleft (x, t')$ returns the tree obtained replacing in t the subtree rooted in $x(t)$ by t' . If $x(t)$ is undefined, $t \triangleleft (x, t')$ adds the smallest chain of empty nodes to t such that $x(t)$ is defined and inserts t' .

Deployment. A deployment D is an overloaded partial function on locations and processes. Intuitively, we map locations to the set of processes running at that location, and we map processes to their local state and input queues. Therefore, given a location l , we read $D(l) = \{\tilde{p}\}$ as “the processes \tilde{p} are running at the location l ” (for any process p and deployment D , we assume that each process p is always located at only one location). For processes, instead, given a process p then $D(p)$ returns a pair (t, M) , where t is a tree representing the local state of p and M is a *queue map*. A queue map defines the input queues that p can use to receive messages from other processes. Formally, a queue map M is a partial function of type $M : \mathcal{T} \rightarrow \text{Seq}(\mathcal{O} \times \mathcal{T})^2$. A map $M(t_c) = \widetilde{(o, t)}$ means that the process owning M has an input queue containing (o, t) (a sequence of messages³) that *correlates* with the data t_c , called the correlation key of the queue. In our semantics, we use correlation keys to formalise our mechanism of message correlation. When receiving a message with some correlation data, our semantics places the message in the queue pointed by the correlation key that matches the correlation data of the message. In a message (o, t) , o is the operation used by the sender and t is the payload of the message. In the remainder, for $D(p) = (t, M)$, we use the shortcuts $D(p).\text{st}$ to refer to t (the state of p) and $D(p).\text{que}$ to refer to M (the queues of p), respectively.

Observe that, in our model, programmers do not need to specify the deployment of a choreography⁴. Concretely, for any choreography program C with no free session names (all sessions are under a start term), we can define a *default deployment* where all active processes are assigned to some location and given an empty state and queue map. Below, we denote the set of free process names in a choreography C with $\text{fp}(C)$, and we write t_\perp for the empty tree.

² $\text{Seq}(_)$ is the type constructor of sequences.

³We denote the empty list ε and a list of messages $(o_1, t_1) :: \dots :: (o_n, t_n)$. We use $::$ rather than commas for a clearer definition of the semantics of AC.

⁴As in other languages states of programs are not specified by programmers.

Definition 15 (Default Deployment). *Let C be a choreography with no free session names. Then, D is a default deployment for C if for all $p \in \text{fp}(C) : D(p) = (t_\perp, \emptyset)$ and $p \in D(l)$ for some l .*

Reductions. We present the semantics for AC in terms of reductions of the form $D, C \rightarrow D', C'$, where D, C is a *running choreography*. Formally, the reduction relation \rightarrow is the smallest closed under the rules reported in Figure 6.3. In the rules for session creation and communications, we make use of the auxiliary relation $D, \delta \blacktriangleright D'$ to model how effect δ changes a deployment D into D' . δ ranges over

$$\delta ::= \mathbf{start} \ k : \widetilde{l.p[A]} \mid k : p[A].e \rightarrow B.o \mid k : A \rightarrow q[B].o(\underline{x})$$

denoting, from left to right, the start of a session and sending and reception of a message. Below we separately discuss our main rules for $D, C \rightarrow D', C'$, along with the definition of $D, \delta \blacktriangleright D'$ for each δ . We report the semantics of effects in Figure 6.4.

Rule $[^C]_{\text{START}}$. Rule $[^C]_{\text{START}}$ starts a fresh (denoted by $\#^5$) session k' with a complete action, together with some fresh processes \tilde{r} (k' and \tilde{r} replace respectively k and \tilde{q} in the continuation). Intuitively, in the premise $D, \delta \blacktriangleright D'$ for $\delta = \mathbf{start} \ k : \widetilde{l.p[A]}$, we obtain D' from D adding the needed information to support the newly created session in the continuation C' . In particular, we need to add information about the location of each new process and to set up the state and queue maps of each process involved in the new session to enable communications via message correlation. We formalise this requirement in the predicate of *session support*, sup :

Definition 16 (Session Support (sup)). *We say that $(t, \{M_A\}_{A \in \tilde{A}})$ is a session support for a deployment D and some located roles $\widetilde{l.A}$, denoted by the predicate $\text{sup}(t, \{M_A\}_{A \in \tilde{A}}, D, \widetilde{l.A})$, if and only if:*

- (Locations) For all $l.A \in \widetilde{l.A}$, $\underline{A.l}(t) = l$.
- (Correlation Keys) For all pairwise-distinct $l.A, l'.B$ in $\widetilde{l.A}$, $\underline{A.B}(t) = t_c$ and $M_B(t_c) = \varepsilon$ for some t_c such that $t_c \notin \text{dom}(D(s).\text{que})$ for all $s \in D(l')$.

The predicate $\text{sup}(t, \{M_A\}_{A \in \tilde{A}}, D, \widetilde{l.A})$ holds if the tree t , called *session descriptor*, contains the locations of the processes playing the roles in the session and the correlation keys that they use. Note that in t we use roles (which are static

⁵We let $\#$ informal for a simpler presentation, yet it is easy to formalise freshness of names on D .

6.2. Applied Choreography Language

$$\begin{array}{c}
\frac{\eta = k : p[A].e \rightarrow B.o \quad D, \eta \blacktriangleright D'}{D, \eta; C \rightarrow D', C} \quad [C]_{\text{SEND}} \\
\frac{j \in I \quad D, k : A \rightarrow q[B].o_j(x_j) \blacktriangleright D'}{D, k : A \rightarrow q[B].\{o_i(x_i); C_i\}_{i \in I} \rightarrow D', C_j} \quad [C]_{\text{RECV}} \\
\frac{\eta = k : p[A].e \rightarrow q[B].o(x) \quad D, k : p[A].e \rightarrow B.o \blacktriangleright D'}{D, \eta; C \rightarrow D', k : A \rightarrow q[B].o(x); C} \quad [C]_{\text{COM}} \\
\frac{i = 1 \text{ if } \text{eval}(e, D(p).\text{st}) = \text{true}, i = 2 \text{ otherwise}}{D, \text{if } p.e \{C_1\} \text{ else } \{C_2\} \rightarrow D, C_i} \quad [C]_{\text{COND}} \\
\frac{D, C_1 \rightarrow D', C'_1}{D, \text{def } X = C_2 \text{ in } C_1 \rightarrow D', \text{def } X = C_2 \text{ in } C'_1} \quad [C]_{\text{CTX}} \\
\frac{\mathcal{R} \in \{\equiv, \simeq_C\} \quad C_1 \mathcal{R} C'_1 \quad D, C'_1 \rightarrow D', C'_2 \quad C'_2 \mathcal{R} C_2}{D, C_1 \rightarrow D', C_2} \quad [C]_{\text{EQ}} \\
\frac{D, C_1 \rightarrow D', C'_1}{D, C_1 | C_2 \rightarrow D', C'_1 | C_2} \quad [C]_{\text{PAR}} \\
\frac{\#\tilde{r} \quad \#k' \quad p \in D(l) \quad \delta = \text{start } k' : l.p[A], \widetilde{l.r[B]} \quad D, \delta \blacktriangleright D'}{D, \text{start } k : p[A] \Leftrightarrow \widetilde{l.q[B]}; C \rightarrow D', C[k'/k][\tilde{r}/\tilde{q}]} \quad [C]_{\text{START}} \\
\frac{\begin{array}{l} i \in \{1, \dots, n\} \quad \#k' \\ \{\widetilde{l.B}\} = \biguplus_i \{\widetilde{l_i.B_i}\} \quad \#\tilde{r} \quad \{\tilde{r}\} = \bigcup_i \{\tilde{r}_i\} \quad p \in D(l) \\ \delta = \text{start } k' : l.p[A], \widetilde{l_1.r_1[B_1]}, \dots, \widetilde{l_n.r_n[B_n]} \quad D, \delta \blacktriangleright D' \end{array}}{D, \text{req } k : p[A] \Leftrightarrow \widetilde{l.B}; C | \prod_i (\text{acc } k : \widetilde{l_i.q_i[B_i]}; C_i) \rightarrow D', C[k'/k] | \prod_i (C_i[k'/k][\tilde{r}_i/\tilde{q}_i]) | \prod_i (\text{acc } k : l_i.q_i[B_i]; C_i)} \quad [C]_{\text{PSTART}}
\end{array}$$

Figure 6.3: Choreography calculus, semantics.

$$\begin{array}{c}
 \frac{l \in \underline{\mathbf{k.B.l}}(D(\mathbf{p}).\text{st}) \quad \mathbf{q} \in D(l) \quad t_c = \underline{\mathbf{k.A.B}}(D(\mathbf{p}).\text{st}) \quad D(\mathbf{q}).\text{que}(t_c) = \tilde{m} \quad t_m = \text{eval}(e, D(\mathbf{p}).\text{st})}{D, k : \mathbf{p[A].e} \rightarrow \mathbf{B.o} \blacktriangleright D[\mathbf{q} \mapsto (D(\mathbf{q}).\text{st}, D(\mathbf{q}).\text{que}[t_c \mapsto \tilde{m} :: (o, t_m)])]} \quad [\mathbf{D}|_{\text{SEND}}] \\
 \\
 \frac{t_c = \underline{\mathbf{k.A.B}}(D(\mathbf{q}).\text{st}) \quad D(\mathbf{q}).\text{que}(t_c) = (o, t_m) :: \tilde{m}}{D, k : \mathbf{A} \rightarrow \mathbf{q[B].o}(x) \blacktriangleright D[\mathbf{q} \mapsto (D(\mathbf{q}).\text{st} \triangleleft (x, t_m), D(\mathbf{q}).\text{que}[t_c \mapsto \tilde{m})]} \quad [\mathbf{D}|_{\text{RECV}}] \\
 \\
 \frac{\text{sup}(t, \{M_C\}_{C \in \{\mathbf{A}, \mathbf{B}\}}, D, (l.\mathbf{A}, \widetilde{l.\mathbf{B}})) \quad l'.\mathbf{q[B]} \in \widetilde{l.\mathbf{q[B]}} \quad t_p = D(\mathbf{p}).\text{st} \quad M_p = D(\mathbf{p}).\text{que} \quad D' = D \quad [l' \mapsto D(l') \cup \{\mathbf{q}\}] \quad [\mathbf{p} \mapsto (t_p \triangleleft (\underline{\mathbf{k}}, t), M_p \cup M_A)] \quad [\mathbf{q} \mapsto (t_\perp \triangleleft (\underline{\mathbf{k}}, t), M_B)]}{D, \text{start } k : l.\mathbf{p[A]}, \widetilde{l.\mathbf{q[B]}} \blacktriangleright D'} \quad [\mathbf{D}|_{\text{START}}]
 \end{array}$$

 Figure 6.4: Semantics of effects on D

names) instead of process identifiers because in the message correlation mechanism found in SOC, a sender does not know the process identifier of the intended receiver (this will be reflected in the development of our target language, DCC, in § 6.5). The set $\{M_A\}_{A \in \bar{\mathbf{A}}}$ contains the respective queue map M_A for each role A . M_A must contain an empty queue correlating with the corresponding key in t for each other role in the session, e.g., if role A receives from B , M_A contains an empty queue correlating with the key at $\underline{\mathbf{B.A}}$ in t . Thus, each role has a queue to receive messages from each other role, as in standard multiparty session types with session-role channels [104]. Correlation keys must be fresh within the location of the receiver⁶.

Rule $[\mathbf{C}|_{\text{START}}]$ applies if Rule $[\mathbf{D}|_{\text{START}}]$ applies. In $[\mathbf{D}|_{\text{START}}]$, we choose in a non-deterministic way the session descriptor t and the queue maps $\{M_C\}_{C \in \{\mathbf{A}, \mathbf{B}\}}$ under the only requirement that they respect predicate sup . This requirement makes our semantics very general as it supports potentially many session descriptors and therefore different implementations (e.g., based on cookies, random sets of data, API keys, etc.), as long as they comply with our definition of session support (see § 7.2). In $[\mathbf{D}|_{\text{START}}]$, according to the previous definition, we copy the session descriptor under $\underline{\mathbf{k}}$, the path named after the session. This eases the access to the structure, since all in-session interactions happen under a certain session name k , which is globally fresh and therefore cannot clash with other pre-existing structures.

Remark 5. In $[\mathbf{D}|_{\text{START}}]$ the session descriptor assigned to each process is not minimal, indeed it contains unused data like correlation keys used by other processes.

⁶In AC, like in SOC, the sender delivers its message in the queue correlating with the key of the message and owned by a process located at the location of the receiver. To make correlation deterministic we forbid duplicate keys within a location.

6.2. Applied Choreography Language

However, since it does not alter our result, we favoured this simpler definition.

Rule $[^C|_{\text{SEND}}]$. The rule describes a partial sending and, as shown in Figs. 6.3 and 6.4, it updates the deployment with Rule $[^D|_{\text{SEND}}]$ to account for asynchronous message passing. We comment the conditions of $[^D|_{\text{SEND}}]$, referring paths as applied on the state of the sender p . The first four conditions (from top left to down right) guess the receiving process q , which is that process located at l — as retrieved under k.B.1 — and that owns the queue correlating with the key at k.A.B. This models real-world message correlation, where the sender guesses the receiver from its location and the correlation key of the message. $t_m = \text{eval}(e, D(p).\text{st})$ is the content of the message, result of the evaluation of expression e on the state of p . $[^D|_{\text{SEND}}]$ modifies D such that q stores in the queue correlating with the key at k.A.B the message (o, t_m) .

Rule $[^C|_{\text{RECV}}]$. Rule $[^C|_{\text{RECV}}]$ implements a partial reception, updating the deployment with Rule $[^D|_{\text{RECV}}]$. In particular, $[^C|_{\text{RECV}}]$ records on which of the available operations $(o_i, i \in I)$ q received a message from the process playing role A. In $[^D|_{\text{RECV}}]$, the first condition finds the queue that q , playing B, uses to receive from A. The second one provides a new deployment D' if the head of the queue has a message on operation o . If it does, D' removes the message from the head of the queue and copies in the state of q , under path x , the content of the message t_m .

Other rules. Rule $[^C|_{\text{COM}}]$ describes a complete communication. Intuitively, the reduction of a complete communication i has the same effect on D as a (send) term and ii) it reduces the (com) to a (recv) on the same operation. Thanks to deployments and effects, we could model complete asynchronous communications in a fairly simple way wrt previous approaches (see § 7.2). Formally, the effect of Rule $[^C|_{\text{COM}}]$ on D is $\delta = k : p[A].e \rightarrow B.o$ which applies Rule $[^D|_{\text{SEND}}]$. In the reductum, the (com) term reduces to a partial reception on operation o . Rules $[^C|_{\text{COND}}]$, $[^C|_{\text{CTX}}]$, and $[^C|_{\text{PAR}}]$ are standard, while rule $[^C|_{\text{EQ}}]$ accounts for the structural congruence \equiv and the swapping relation \simeq_C . The structural congruence, defined as the smallest congruence supporting α -conversion and satisfying the rules below

$$\begin{aligned} C \mid C' &\equiv C' \mid C & (C_1 \mid C_2) \mid C_3 &\equiv C_1 \mid (C_2 \mid C_3) \\ \text{def } X = C' \text{ in } \mathbf{0} &\equiv \mathbf{0} & \text{def } X = C' \text{ in } C[X] &\equiv \text{def } X = C' \text{ in } C[C'] \end{aligned}$$

abstracts from purely syntactic differences in processes and treats recursion in a standard way. As in [13], the swap relation \simeq_C exchanges the order of some actions. This enables more interleaving among processes. We report the rules of

$$\begin{array}{l}
 C_c = \\
 \left\{ \begin{array}{l}
 1. \text{ req } k_d : c[C] \Leftrightarrow l_A.A, l_{DM}.DM, l_L.L; \\
 2. k_d : c[C].mkReq() \rightarrow A.get; \\
 3. \text{ def STORE} = \\
 4. k_d : DM \rightarrow c[C].\{ \\
 5. \text{ pkt}(\underline{bb}); \\
 6. k_s : c[U].\underline{bb} \rightarrow f[F].append(\underline{data}); \\
 7. \text{ STORE}, \\
 8. \text{ checksum}(\underline{cs}); \\
 9. k_s : c[U].\underline{cs} \rightarrow f[F].check(\underline{cs}); \\
 10. \text{ if } f.check(\underline{fname}, \underline{cs}) \{ \\
 11. k_s : f[F] \rightarrow c[U].saved \\
 12. \} \text{ else } \{ \\
 13. k_s : f[F] \rightarrow c[U].discarded \\
 14. \} \\
 15. \text{ in} \\
 16. k_d : A \rightarrow c[C].\{ \\
 17. ok; \\
 18. \text{ start } k_s : c[U] \Leftrightarrow l_C.f[F]; \\
 19. k_s : c[U].name() \rightarrow f[F].create(\underline{fname}); \\
 20. \text{ STORE}, \\
 21. ko \}
 \end{array} \right.
 \end{array}
 \qquad
 \begin{array}{l}
 C_s = \\
 \left\{ \begin{array}{l}
 1. \text{ acc } k_d : l_A.a[A], l_{DM}.dm[DM], l_L.l[L]; \\
 2. k_d : C \rightarrow a[A].get(\underline{req}); \\
 3. \text{ if } a.isValid(\underline{req}) \{ \\
 4. k_d : a[A].\underline{req.rsc} \rightarrow dm[DM].ok(\underline{rsc}); \\
 5. k_d : a[A] \rightarrow C.ok; \\
 6. k_d : dm[DM].logok(\underline{rsc}) \rightarrow l[L].log(\underline{log}); \\
 7. \text{ def TRANSFER} = \\
 8. \text{ if } dm.more(\underline{rsc}) \{ \\
 9. k_d : dm[DM].next(\underline{rsc}) \rightarrow C.pkt; \\
 10. \text{ TRANSFER} \\
 11. \} \text{ else } \{ \\
 12. k_d : dm[DM].checksum(\underline{rsc}) \rightarrow C.checksum \\
 13. \} \\
 14. \text{ in} \\
 15. \text{ TRANSFER} \\
 16. \} \text{ else } \{ \\
 17. k_d : a[A].\underline{req.rsc} \rightarrow dm[DM].ko(\underline{rsc}); \\
 18. k_d : a[A] \rightarrow C.ko; \\
 19. k_d : dm[DM].logko(\underline{rsc}) \rightarrow l[L].log(\underline{log}) \\
 20. \}
 \end{array} \right.
 \end{array}$$

Figure 6.5: Choreography Example

the swap relation in Appendix D.1. As an example, consider the Rule below that swaps η and η' if they share no processes (returned by $\text{pn}(\eta)$).

$$\text{pn}(\eta) \cap \text{pn}(\eta') = \emptyset \quad \Rightarrow \quad \eta; \eta' \simeq_C \eta'; \eta \quad [CS|_{\text{ETAETA}}]$$

Swapping terms means that, although a choreography defines the global order in which its processes send and receive messages, this order can change at run-time. Despite the change, we guarantee to preserve the order of messages between each couple of processes in a session. Rule $[C|_{\text{PSTART}}]$ starts a new session synchronising a partial choreography that requests to start a session with other choreographies that can accept the request. The premise of the rule $\{\widetilde{l}.B\} = \biguplus_i \{\widetilde{l}_i.B_i\}$ (\biguplus being the union of the disjoint lists of located roles) requires that, in the accepting choreographies, the list of locations and their supported roles match the corresponding list of the request. The accepting choreographies remain available afterwards, for reuse. The rest of the rule is similar to $[C|_{\text{START}}]$.

6.2.3 An example

We show an example of a system for verified file transfer written in AC. We use the example in § 6.6 to illustrate our compilation.

Verified file transfer. The program *i*) validates the file request of a Client on a Server, *ii*) transfers a file in multiple parts, and *iii*) verifies the transfer with a checksum. The Server logs all requests. We report in Figure 6.5 the code of program C , parallel composition of the two partial choreographies $C = C_c \mid C_s$. C_c and C_s respectively define the client- and the server-side code.

In C , process c — the only active process — plays the role of the Client C . In C_s , a plays the Access Manager (A), dm the Download Manager (DM) and l the Logger (L). In C_c (Line 18 of C_c) process f plays F , accessing the file system. We locate A , DM , L , and F at respectively l_A , l_{DM} , l_L , and l_C .

Lines 1 of C_c and C_s start session k_d between c , a , dm , and l . At Lines 2 of C_c and C_s , c makes a request for a file to a . Lines 4-15 of C_s define the outcome of a valid request (Line 3). Following the Lines, a forwards the resource request to dm and confirms (on *ok*) to c the request. dm asks l to *log* the request. Lines 7-14 define the recursive procedure TRANSFER, called at Line 15, for the multi-part file download. If dm has more packets of the resource, it sends the next to c on operation *pkt* and TRANSFER repeats. Else dm ends k_d sending to c the checksum of the file. Lines 17-19 of C_s specify the outcome of an invalid request: a notifies dm and c of the failed attempt (*ko*) and dm asks to *log* the event to l .

Observe that C_c defines at Lines 3-15 a procedure to STORE the file. At Line 18 of C_c , after the approval, c (playing user U) starts a new session k_s with f , asking f to create a file to store the incoming packets. Then, in STORE, if c receives a new packet of bytes from dm (Line 5), it asks f to append them to the local file and STORE repeats. Else c receives the checksum from dm (Line 8) and c forwards it to f . Finally, f notifies c whether it *saved* or *discarded* the file wrt the checksum.

Remark 6. *To illustrate compositionality of ACs, we (logically) divided the example into two choreographies, (resp. the client- and server-side described above). However, we could specify the system as single choreography, e.g.,*

1. *start* $k_d : c[C] \Leftrightarrow l_A.a[A], l_{DM}.dm[DM], l_L.l[L];$
2. $k_d : c[C].mkReq() \rightarrow a[A].get(\underline{req}); \dots$

We comment some features of AC programs in the example.

Session Descriptors. At Lines 1 of C_c and C_s c requests to start a new session k_d with some service processes a , dm , and l . We describe the structure of the session descriptor, stored in the state of the mentioned processes under path $\underline{k_d}$.

We report in Figure 6.6 such structure. Following the path \underline{A} , we find a subtree with *i*) a node \underline{l} that stores the location l_A of a and *ii*) a set of nodes, named after all the other roles in k_d , leading to subtrees containing the correlation keys related to the roles in the paths. For example, the path $\underline{A.C}$ points the tree t_{AC} , which contains the correlation key used by a to send messages to c .

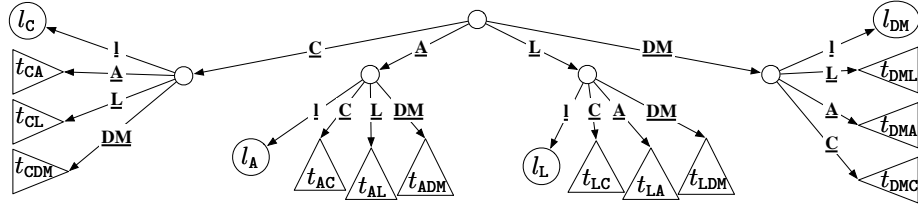


Figure 6.6: Example of structure of Session Descriptor

Parallelism. Lines 5-6 of C_s show how Rule $[^C|_{EQ}]$ can exchange the order of execution of actions. The actions at Lines 5 and 6 regard different processes (resp. a, s, and l) and the two instructions can swap along Rule $[^{CS}|_{ETAETA}]$ of the swap relation. A possible reduction of C , starting from Line 5 of C_s , can apply Rule $[^C|_{EQ}]$ to swap Lines 5 and 6, reduce Line 6 with Rule $[^C|_{COM}]$ to a partial reception for process l, and swap back the two Lines. Next, either a delivers its message with $[^C|_{SEND}]$ or $[^C|_{EQ}]$ applies, letting l consume its message. Observe that the swap is non-deterministic, allowing other possible executions.

Asynchrony. AC supports asynchronous communication with queues, however, we need the swap relation to let a process send or receive a message asynchronously. Consider Lines 4-5 of C_s . We can apply $[^C|_{COM}]$ on Line 4 and let a send its message to dm. This reduces Line 4 to a partial reception on dm. Then, we can apply $[^C|_{EQ}]$, swap the redex of Line 4 with Line 5, and let Line 5 execute with $[^C|_{SEND}]$. Finally, we let dm consume its message with $[^C|_{RECV}]$ on Line 4.

6.3 Typing

In this section we define our typing discipline for Applied Choreographies. Our typing checks the behaviour of sessions against protocols, given as multiparty session types [43, 105]. The main novelty of our type system is checking that the evolution of a deployment (states for message correlation and queues for asynchronous messaging) correctly implements the sessions described in a program, ensuring absence of errors such as deadlocks. We detail this part in § 6.3.3.

6.3.1 Types and Type Projection

Global and Local types. As in standard multiparty session types, we use *global types* to represent protocols from a global viewpoint and *local types* to describe the behaviour of each participant. Our type system checks that the local types, that

6.3. Typing

abstract the behaviour of processes in a choreography, coherently follow a global type. We report below the syntax of global types G and local types T .

$$\begin{array}{l}
 G ::= A \rightarrow B.\{o_i(U_i); G_i\}_i \\
 \quad | \text{rec } \mathbf{t}; G \\
 \quad | \mathbf{t} \\
 \quad | \text{end}
 \end{array}
 \qquad
 \begin{array}{l}
 T ::= \alpha.\{o_i(U_i); T_i\}_i \\
 \quad | \text{rec } \mathbf{t}; T \\
 \quad | \mathbf{t} \\
 \quad | \text{end}
 \end{array}$$

$$\alpha ::= !A \mid ?B \qquad U ::= S\{\underline{x}_i : U_i\}_i \qquad S ::= \text{int} \mid \text{bool} \mid \text{str} \mid \dots$$

A global type $A \rightarrow B.\{o_i(U_i); G_i\}_i$ abstracts a communication, where A can send to B a message on any of the operations o_i and continue with the respective continuation G_i . A carried type U types the tree value exchanged in the message. A tree type $S\{\underline{x}_i : U_i\}_i$ abstracts a tree with root value S (a basic type) and subtrees reachable from the root node by following \underline{x}_i with respective types U_i (our notation recalls that for record types in [106]). In local types, $!A.\{o_i(U_i); T_i\}_i$ abstracts the sending of a message of type U_i to role A on one of the operations o_i , with continuation T_i . Dually, $?A.\{o_i(U_i); T_i\}_i$ abstracts the offering of an input choice for all the operations o_i , with continuation T_i . Other terms for recursion and termination (end) are standard.

Type Projection. To relate global types to the behaviour of endpoints, we project a global type G onto the local type of a single role. We report in Figure 6.7 the projection of global types, defined following [14]. $\llbracket G \rrbracket_A$ denotes the projection of G onto the role A . Intuitively, $\llbracket G \rrbracket_A$ gives an encoding of the local actions expected by role A in the global type G . When projecting a communication we require the local behaviour of all roles not involved in it to be merged with the merging operator \sqcup . Like in [14] $T \sqcup T'$ is isomorphic to T and T' up to branching, where all branches of T or T' with distinct operations are also included.

6.3.2 Type checking

We present our system that guarantees that sessions follow their types.

Environments. We define our typing environments Γ, Γ', \dots as:

$$\begin{array}{l}
 \Gamma ::= \Gamma, \tilde{l} : G\langle A | \tilde{B} | \tilde{C} \rangle \quad (\text{service typing}) \quad | \quad \Gamma, \mathbf{p}.x : U \quad (\text{variable typing}) \\
 \quad | \quad \Gamma, k[A] : T \quad (\text{local typing}) \quad | \quad \Gamma, X : \Gamma \quad (\text{procedure typing}) \\
 \quad | \quad \Gamma, \mathbf{p} : k[A] \quad (\text{ownership}) \quad | \quad \emptyset \quad (\text{empty env.})
 \end{array}$$

A service typing $\tilde{l} : G\langle A | \tilde{B} | \tilde{C} \rangle$ types with G all sessions created by contacting the services at the locations \tilde{l} . We explain the role annotations: A is the role that the

$$\begin{aligned}
 \llbracket \mathbf{t} \rrbracket_{\mathbf{A}} &= \mathbf{t} & \llbracket \text{end} \rrbracket_{\mathbf{A}} &= \text{end} & \llbracket \text{rec } \mathbf{t}; G \rrbracket_{\mathbf{A}} &= \begin{cases} \text{rec } \mathbf{t}; \llbracket G \rrbracket_{\mathbf{A}} & \text{if } \mathbf{A} \in G \\ \text{end} & \text{otherwise} \end{cases} \\
 \llbracket \mathbf{A} \rightarrow \mathbf{B}. \{o_i(U_i); G_i\}_i \rrbracket_{\mathbf{C}} &= \begin{cases} !\mathbf{B}. \{o_i(U_i); \llbracket G_i \rrbracket_{\mathbf{C}}\}_i & \text{if } \mathbf{C} = \mathbf{A} \\ ?\mathbf{A}. \{o_i(U_i); \llbracket G_i \rrbracket_{\mathbf{C}}\}_i & \text{if } \mathbf{C} = \mathbf{B} \\ \bigsqcup_i \llbracket G_i \rrbracket_{\mathbf{A}} & \text{otherwise} \end{cases}
 \end{aligned}$$

Figure 6.7: Choreography Calculus - Global Type Projection

active process (the starter) should play; roles $\tilde{\mathbf{B}}$ are the roles respectively played by each l in \tilde{l} (each l plays one role, so the length of $\tilde{\mathbf{B}}$ is the same as the length of \tilde{l}); finally, $\tilde{\mathbf{C}}$ (where $\tilde{\mathbf{C}} \subseteq \tilde{\mathbf{B}}$) are the roles implemented by the choreography that we are typing. The annotation $\tilde{\mathbf{C}}$ enables composition of choreographies yet ensuring that only one choreography implements a specific role, as in [14]. When we write $\Gamma, \tilde{l} : G \langle \mathbf{A} | \tilde{\mathbf{B}} | \tilde{\mathbf{C}} \rangle$, we always assume that: *i*) $\{\mathbf{A}, \tilde{\mathbf{B}}\} = \text{roles}(G)$, where roles returns the set of roles in G ; *ii*) the locations \tilde{l} are ordered lexicographically; and, *iii*) the locations in l do not appear in any other service typing in Γ . A local typing $k[\mathbf{A}] : T$ states that role \mathbf{A} in session k follows the local type T . We assume that roles in a session are typed by a single local typing, as in standard multiparty session types [43]. An ownership typing $p : k[\mathbf{A}]$ states that process p owns the role \mathbf{A} in session k : each process can participate in multiple sessions, but can play only one role in each of such sessions. Hence, a process p may appear in more than one ownership typings in a Γ , but never more than once per session. The other typings for variables and recursive procedures are standard.

Typing Judgements and Rules. A judgement $\Gamma \vdash C$ states that the choreography C follows the specifications given in Γ . We report in Figure 6.8 some select typing rules to derive valid typing judgements and comment the rules below. We report the full typing rules in Appendix D.2.

Rule $\llbracket^T \text{START} \rrbracket$ types a session start. In the first premise, the service typing $\tilde{l} : G \langle \mathbf{A} | \tilde{\mathbf{B}} | \tilde{\mathbf{B}} \rangle$ checks that the continuation implements all the roles in protocol G . The auxiliary function init^7 intuitively returns an environment containing all the ownerships and local typings to correctly type a freshly-started session. The type of each process is the local type projection of the global type G on the role

⁷Formally $\text{init}(\widetilde{p[\mathbf{A}]}, k, G) = \{q : k[\mathbf{B}], k[\mathbf{B}] : \llbracket G \rrbracket_{\mathbf{B}} \mid q[\mathbf{B}] \in \widetilde{p[\mathbf{A}]}\}$.

6.3. Typing

$$\begin{array}{c}
\frac{j \in I \quad \Gamma \vdash p : k[A], q : k[B] \quad \Gamma \vdash p.e : U_j \quad \Gamma, q.x : U_j, k[A] : T_j, k[B] : T'_j \vdash C}{\Gamma, k[A] : !B.\{o_i(U_i); T_i\}_{i \in I}, k[B] : ?A.\{o_i(U_i); T'_i\}_{i \in I} \vdash k : p[A].e \rightarrow q[B].o_j(x); C} \quad [T]_{\text{COM}} \\
\\
\frac{j \in I \quad \Gamma \vdash p : k[A] \quad \Gamma \vdash p.e : U_j \quad \Gamma, k[A] : T_j \vdash C}{\Gamma, k[A] : !B.\{o_i(U_i); T_i\}_{i \in I} \vdash k : p[A].e \rightarrow B.o_j; C} \quad [T]_{\text{SEND}} \\
\\
\frac{\Gamma \vdash q : k[B] \quad \forall j \in I. \Gamma, q.x_j : U_j, k[B] : T_j \vdash C_j}{\Gamma, k[B] : ?A.\{o_i(U_i); T_i\}_{i \in I} \vdash k : A \rightarrow q[B].\{o_j(x_j); C_j\}_{j \in I \cup J}} \quad [T]_{\text{RECV}} \\
\\
\frac{\Gamma, \tilde{l} : G\langle A|\tilde{B}|\tilde{B} \rangle, \text{init}(r[\tilde{C}], k, G) \vdash C \quad r[\tilde{C}] = p[A], q[\tilde{B}] \quad \tilde{q} \notin \Gamma}{\Gamma, \tilde{l} : G\langle A|\tilde{B}|\tilde{B} \rangle \vdash \mathbf{start} \ k : p[A] \Leftrightarrow \tilde{l}.q[\tilde{B}]; C} \quad [T]_{\text{START}} \\
\\
\frac{\frac{\Gamma_i \vdash C_i \quad \Gamma_2 \vdash C_2}{\Gamma_1 \circ \Gamma_2 \vdash C_1 | C_2} \quad [T]_{\text{PAR}} \quad \frac{\Gamma, p : k[A], k[A] : \llbracket G \rrbracket_A \vdash C \quad \Gamma \vdash \tilde{l} : G\langle A|\tilde{B}|\emptyset \rangle}{\Gamma \vdash \mathbf{req} \ k : p[A] \Leftrightarrow \tilde{l}.B; C} \quad [T]_{\text{REQ}}}{\frac{\text{end}(\Gamma) \quad \Gamma'' \subseteq \Gamma'}{\Gamma, \Gamma'', X\langle \tilde{p} \rangle : \Gamma' \vdash X\langle \tilde{p} \rangle} \quad [T]_{\text{CALL}} \quad \frac{\tilde{l} \subseteq \tilde{l}' \quad \Gamma, \tilde{l}' : G\langle A|\tilde{B}|\emptyset \rangle, \text{init}(q[\tilde{C}], k, G) \vdash C}{\Gamma, \tilde{l}' : G\langle A|\tilde{B}|\tilde{C} \rangle \vdash \mathbf{acc} \ k : \tilde{l}.q[\tilde{C}]; C} \quad [T]_{\text{ACC}}}
\end{array}$$

Figure 6.8: Choreography Calculus - Typing Rules (selected)

owned by the process in the session. Similar to $[T]_{\text{START}}$, Rule $[T]_{\text{REQ}}$ performs the checks only for the process requesting the creation of a new session. Dually, $[T]_{\text{ACC}}$ checks that the processes created after a request correctly implement their expected behaviour.

Rule $[T]_{\text{COM}}$ types a complete communication, checking that: *i*) the chosen operation o_j is among the ones that the sender can select according to its local type; *ii*) similarly, o_j is among the ones offered by the receiver according to its local type; *iii*) the sender and the receiver processes own their respective roles in the session; *iv*) the expression of the sender (e) has the type⁸ U_j expected by the protocol; *v*) the receiver uses the reception variable accordingly in the continuation C ; and *vi*) processes p and q proceed according to their respective types in Γ . Similar to rule $[T]_{\text{COM}}$, $[T]_{\text{SEND}}$ and $[T]_{\text{RECV}}$ respectively check (*send*) and (*recv*) actions.

Rule $[T]_{\text{PAR}}$ uses the *role distribution operator* $\Gamma_1 \circ \Gamma_2$, from [14], to merge service typings and to check that choreographies executing in parallel do not implement overlapping roles at locations. Formally, $\Gamma_1 \circ \Gamma_2$ is defined if Γ_1 and Γ_2 type different processes⁹ and for $\Gamma_i = \Gamma'_i, \Gamma_i^l, i \in \{1, 2\}$ where Γ_i^l contains only

⁸The judgement $\vdash t : U$ reads as “tree t has type U ”.

⁹Formally, we can write $\Gamma_1 \circ \Gamma_2$ if $\text{pn}(\Gamma_1) \cap \text{pn}(\Gamma_2) = \emptyset$ for $\text{pn}(\Gamma) = \{p \mid p : k[A] \in \Gamma\}$.

service typings

$$\Gamma_1 \circ \Gamma_2 = \Gamma'_1, \Gamma'_2, \Gamma'_1 \circ \Gamma'_2$$

$$\Gamma'_1 \circ \Gamma'_2 = \left\{ \tilde{l}: G\langle A|\tilde{B}|\tilde{C} \rangle \mid \tilde{l}: G\langle A|\tilde{B}|\tilde{D} \rangle \in \Gamma'_1 \wedge \tilde{l}: G\langle A|\tilde{B}|\tilde{E} \rangle \in \Gamma'_2 \wedge \tilde{D} \uplus \tilde{E} = \tilde{C} \right\}$$

All the other typing rules are standard.

6.3.3 Runtime Typing

To prove that well-typed AC programs never go wrong, we need to pay attention to how their deployments evolve at runtime. For example, in Rule $[^C|_{\text{COM}}]$ the sender needs the necessary information in its state to “find” the receiver through correlation: a remarkable difference wrt previous works on choreographies, where such conditions do not exist and choreographies can always continue execution (see, e.g., [10, 107, 96, 13]). To address this issue, we extend our typing discipline to check runtime states.

Wrong Deployments. We need to prevent deployments from “going wrong” during execution. Intuitively, we say that a deployment is wrong wrt a choreography if any of these conditions holds:

- (*uninitialised variables*) processes have undefined variables;
- (*incompatible session descriptors*) processes in a session store different locations or keys for the same session descriptor;
- (*correlation race*) a correlation key is used for more than one queue at a location;
- (*protocol violations*) a message queue does not contain messages as expected by the protocol of the session in which it is used.

Wrong deployments may cause unpredictable executions or undesired behaviours, e.g., deadlocks. We illustrate the consequences of having wrong deployments with this simple running choreography: $D, k : p[A].\underline{y} \rightarrow q[B].\underline{x}; 0$.

The only way this choreography can reduce is to apply Rule $[^C|_{\text{COM}}]$ in a reduction derivation. The second premise of Rule $[^C|_{\text{COM}}]$ requires a deployment reduction of the form $D, k : p[A].\underline{y} \rightarrow B \blacktriangleright D'$ for some D' . Hence, the deployment D must respect the conditions given in Rule $[^D|_{\text{SEND}}]$. Let us see how a wrong deployment D may cause problems, following the above list:

- (*uninitialised variables*) Assume that D is such that $D(p).\text{st}$ is a tree with no node under path \underline{y} ; then the condition $\text{eval}(\underline{y}, D(p).\text{st})$ given in $[^D|_{\text{SEND}}]$ is undefined and $[^C|_{\text{COM}}]$ cannot be applied, causing the choreography to get stuck.

6.3. Typing

- (*incompatible session descriptors*) Assume $q \in D(l')$ for $l \neq l'$ and $\underline{k.B.I}(D(p).st) = l$. Again, $[^D]_{SEND}$ cannot apply and we have a deadlock, caused by the sender “pointing” at the wrong receiving service. A similar case happens if $t \notin \text{dom}(D(q).que)$ and $\underline{k.A.B}(D(p).st) = t$ because we cannot find the queue of the addressee.
- (*correlation race*) Assume D is such that $D(l) = \{q, r\}$ (r and q are at the same location). Assume also $D(q).que = D(r).que = M$ for some M such that $\underline{k.A.B}(D(p).st) \in M$, i.e., both q and r have a queue correlating with the key that p uses to send its message. Since $[^D]_{SEND}$ guesses the receiver from its location and correlation key, p non-deterministically delivers its message to either q or r . In the second case, we get:

$$D, k : p[A].\underline{y} \rightarrow q[B].o(\underline{x}); \mathbf{0} \rightarrow D', k : A \rightarrow q[B].\underline{x}; \mathbf{0}$$

where in D' the queue of r contains the message received from p . The choreography is now deadlocked because q cannot consume the expected message from p .

- (*protocol violations*) Assume that $D(q).que = M$ for some $M(\underline{k.A.B}(D(p).st)) = (o', t')$ where $o \neq o'$, i.e., q has a message in the queue used by p . If we let the choreography reduce like at the previous point, it ends up deadlocked. After the reduction, the queue used by p contains in its head the message (o', t') and Rule $[^C]_{RECV}$ cannot apply as it expects to find a message for o at that position.

Below we extend our type system to prove that, given a well-typed choreography, our semantics never produces wrong deployments (provided that we do not start from a wrong deployment). Observe that this development is transparent to programmers, since default deployments are never wrong.

Runtime Global Types. We extend the syntax of global types to capture partial runtime states, following the idea presented in [108]:

$$G ::= \dots \mid A \rightsquigarrow B.o(U); G$$

where the new term $A \rightsquigarrow B.o(U)$ means that the sender A has sent the message but the receiver B has still to consume it.

Semantics of Global Types. To express the (abstract) execution of protocols, we give a semantics for global types. Formally, $G \rightarrow G'$ is the smallest relation

on the recursion-unfolding of global types satisfying the rules below

$$\begin{aligned}
 [{}^G|_{\text{SEND}}] \quad j \in I &\Rightarrow A \rightarrow B.\{o_i(U_i); G_i\} \rightarrow A \rightsquigarrow B.o_j(U_j); G_j \\
 [{}^G|_{\text{RECV}}] \quad A \rightsquigarrow B.o(U); G &\rightarrow G' \\
 [{}^G|_{\text{REC}}] \quad G[\text{rec } \mathbf{t}; G/\mathbf{t}] &\rightarrow G' \Rightarrow \text{rec } \mathbf{t}; G \rightarrow G' \\
 [{}^G|_{\text{SWAP}}] \quad G_1 \simeq_G G_2 \wedge G_2 &\rightarrow G'_2 \wedge G'_2 \simeq_G G'_1 \Rightarrow G_1 \rightarrow G'_1
 \end{aligned}$$

The rules are similar to those for AC. Rule $[{}^G|_{\text{SEND}}]$ allows a sender role to proceed before the corresponding receiver has actually received the message. Based on the selected operation, e.g., o_j , the Rule reduces the type to a reception followed by the respective continuation (G_j). The reception is executed in Rule $[{}^G|_{\text{RECV}}]$. In $[{}^G|_{\text{SWAP}}]$, we model parallelism with the relation for global types \simeq_G , which follows the same intuition of \simeq_C . Swapping allows us to model asynchrony in global types as done in AC. We report the rules of \simeq_G in Appendix D.3. As an example, consider the swap between a reception and a complete communication:

$$[{}^{\text{GS}}|_{\text{COMRECV}}] \{A, B\} \cap \{D\} = \emptyset \Rightarrow \begin{aligned} &A \rightarrow B.\{o_i(U_i); C \rightsquigarrow D.o(U); G_i\} \\ &\simeq_G C \rightsquigarrow D.o(U); A \rightarrow B.\{o_i(U_i); G_i\} \end{aligned}$$

Runtime Type checking and Typing Rules. We extend the typing rules given in the previous section to check runtime terms. The extension consists in *i*) new terms for Γ , and *ii*) the introduction of rule $[{}^T|_{\text{DC}}]$ to type runtime choreographies. We extend the grammar of typing environments Γ as follows:

$$\begin{aligned}
 \Gamma ::= & \dots \\
 & | \Gamma, \mathbf{p}@l \quad (\text{location}) \\
 & | \Gamma, b[k]_B^A : T \quad (\text{buffer typing})
 \end{aligned}$$

$\Gamma, \mathbf{p}@l$ states that process \mathbf{p} runs at location l . A buffer typing $b[k]_B^A : T$ types the messages in the queue where the process implementing role B in session k receives messages from role A . To relate the typings of queues (see Definition 18) to the buffer types expected by the protocol of sessions, we define the *buffer type projection* $\llbracket G \rrbracket_B^A$, which returns the expected buffer type of role B from A in G . $\llbracket G \rrbracket_B^A$ extracts from G the partial receptions of the form $A \rightsquigarrow B.o(U)$, translating it to a local type $?A.o(U)$. Below we report the definition of $\llbracket G \rrbracket_B^A$ for the only interesting case of receptions. The full definition is reported in Appendix D.4.

$$\llbracket C \rightsquigarrow D.o(U); G \rrbracket_B^A = \begin{cases} ?A.o(U); \llbracket G \rrbracket_B^A & \text{if } C = A \wedge D = B \\ \llbracket G \rrbracket_B^A & \text{otherwise} \end{cases}$$

We also extend type projection to handle receptions:

$$\llbracket A \rightsquigarrow B.o(U); G \rrbracket_C = \begin{cases} ?A.o(U); \llbracket G \rrbracket_C & \text{if } C = B \\ \llbracket G \rrbracket_C & \text{otherwise} \end{cases}$$

6.3. Typing

We now proceed defining the *partial coherence* predicate $\text{pco}(\Gamma)$, which holds if and only if for all sessions k , the local and buffer typings of k follow (are projection of) the same global type G . The idea is that, since D is a global deployment, we can check for partial coherence of all sessions in Rule $[\text{T}|_{\text{DC}}]$.

Definition 17 (Partial Coherence). *We write $\text{pco}(\Gamma)$ when, for all sessions k in Γ , there exists a global type G such that,*

$$\forall k[\text{B}]: T \in \Gamma \Rightarrow T = \llbracket G \rrbracket_{\text{B}} \wedge \forall \text{A} \in \text{roles}(G)/\{\text{B}\} \Rightarrow \Gamma \vdash b[k]_{\text{B}}^{\text{A}}: \llbracket G \rrbracket_{\text{B}}^{\text{A}}$$

Finally, we define the Rule to type a running choreography:

$$\frac{\text{pco}(\Gamma) \quad \Gamma \vdash D \quad \Gamma \vdash C}{\Gamma \vdash D, C} [\text{T}|_{\text{DC}}]$$

A judgement $\Gamma \vdash D, C$ states that C and D are coherent according to Γ and all sessions in Γ are coherent. Γ is an abstraction between D and C and guarantees D to not go wrong. $\Gamma \vdash D$ checks D to be well-typed (not wrong) wrt to Γ . Formally:

Definition 18 (Deployment Judgements).

$$\Gamma \vdash D \iff \begin{cases} \forall l \in D, \forall \{p, q\} \subseteq D(l), \text{dom}(D(p).\text{que}) \cap \text{dom}(D(q).\text{que}) = \emptyset \\ \forall p.x: U \in \Gamma, \vdash x(D(p).\text{st}): U \\ \forall \{p: k[\text{A}], q: k[\text{B}]\} \subseteq \Gamma, \underline{\mathbf{k}}(D(p).\text{st}) = \underline{\mathbf{k}}(D(q).\text{st}) \\ \forall p: k[\text{A}] \in \Gamma, \Gamma \vdash p@l \wedge p \in D(l) \wedge \underline{\mathbf{k.A.l}}(D(p).\text{st}) = l \\ \forall p: k[\text{A}] \in \Gamma \wedge \forall b[k]_{\text{A}}^{\text{B}}: T \in \Gamma, \text{bte}(\text{B}, D(p).\text{que}(\underline{\mathbf{k.B.A}}(D(p).\text{st}))) = T \end{cases}$$

We comment, from top to bottom, the checks performed by $\Gamma \vdash D$:

- locations must have unique keys: for all locations in D and for all pairs of processes in that location, the queue maps of the two processes have no common correlation key (i.e., the domain of their queue maps are distinct).
- Γ and D must agree on the type of variables: for each typing $p.x: U$ in Γ , D must associate x , in the state of process p , to a value of type U ;
- session descriptors must match: for all couples of processes p and q , playing the respective roles A and B in a session k in Γ , the session descriptors for k stored by p and q are the same;
- Γ and D must agree on the location of all processes: for each process within a session k : *i*) its location according to Γ , *ii*) its location according to D , and *iii*) its location in the session descriptor of k must coincide;
- Γ and D must agree on the state of all queues: for each process p playing role A in a session k in Γ and for each role B such that the buffer type

$$\begin{aligned}
 G_d &= C \rightarrow A.get \left(\left\{ \begin{array}{l} \underline{\text{creds}} : \text{str}, \\ \underline{\text{rsc}} : \text{str} \end{array} \right\} \right); \\
 &A \rightarrow DM. \left\{ \begin{array}{l} ko(\text{str}); A \rightarrow C.ko(); DM \rightarrow L.log(\text{str}), \\ ok(\text{str}); A \rightarrow C.ok(); DM \rightarrow L.log(\text{str}); \text{rec } t; DM \rightarrow C. \left\{ \begin{array}{l} pkt(\text{bytes}); t, \\ chksum(\text{str}) \end{array} \right\} \end{array} \right\} \\
 G_s &= U \rightarrow F.create(\text{str}); \text{rec } s; U \rightarrow F. \left\{ \begin{array}{l} append(\text{bytes}); s, \\ check(\text{str}); F \rightarrow U. \left\{ \begin{array}{l} saved(), \\ discarded() \end{array} \right\} \end{array} \right\}
 \end{aligned}$$

Figure 6.9: Protocols example

$b[k]_A^B : T \in \Gamma$, the extracted buffer type of $\underline{k.A.B}(D(p).st)$ must be equal T . To check this last requirement, we define the *buffer typing extractor* $bte(B, \tilde{m})$ which, given a queue \tilde{m} and a sender role B , returns the buffer type of \tilde{m} ¹⁰.

Typing Example

We show, in Figure 6.9, how we can write the protocols that types the example choreographies in § 6.5. Notably, we define two independent protocols that are implemented, composed, and interleaved at runtime by the choreographies they type. The first protocol, called G_d , specifies the interaction of the system that allows a Client to authenticate on a Server, download a file from it, and verify its integrity, also logging all requests. The roles in the protocol are the same used in § 6.5, i.e., C for the Client, DM for the download manager, A for the access manager and L for the logger.

Following G_d , C asks to A to *get* a file with a message that must contain two values of type str : $\underline{\text{creds}}$ for the credentials of C and $\underline{\text{rsc}}$ for the resource identifier of the file (e.g., a URL). A notifies DM whether the credentials and the resource identifier are valid or not:

- If they are not valid, A forwards to DM the outcome of the validation (str) on operation ko . A also sends an empty message¹¹ to C on operation ko as response to the failed attempt. Finally, DM ends the branch forwarding to L , on log , the error received from A .
- If they are valid, A sends to DM the identifier of the file on operation ok . A confirms the request to C with an empty message on operation ok . DM asks

¹⁰Formally, let $\vdash t_1 : U_1, \dots, \vdash t_n : U_n$, and $\tilde{m} = (o_1, t_1) :: \dots :: (o_n, t_n)$ then $bte(A, \tilde{m}) = ?A.o_1(U_1); \dots; ?A.o_n(U_n)$.

¹¹empty parentheses are a short-cut for type **unit**

6.3. Typing

to *log* the request of the file (*str*) to *L*. The protocol continues with the definition of procedure *rec t* where *DM* sends to *C* either: *i*) a chunk (packet) of bytes of the file, on *pkt*, and repeats *t*, or *ii*) the checksum of the file, on *checksum*, if there are no more packets, ending the protocol.

The second protocol, G_s , defines the interaction between the User *U* of an operative system and a process *F* that accesses the filesystem. In G_s , *U* asks to *F* to *create* a new file (*str*). Then the protocol defines the procedure *rec s*. *U* can either ask *F* to *i*) *append* a new packet of bytes to the file, repeating *s* or *ii*) *check* the file against its checksum (*str*). Based on the outcome of the checksum, *F* either notifies *U* that it *saved* the file, or *discarded* it.

6.3.4 Properties

We close this section with the main guarantees of our type system.

First, our semantics preserves well-typedness:

Theorem 3 (Subject Reduction). $\Gamma \vdash D, C$ and $D, C \rightarrow D', C'$ imply $\Gamma' \vdash D', C'$ for some Γ' .

We now relate the behaviour of sessions in a well-typed choreography to their respective global types. We denote $\llbracket G \rrbracket_k$ the projection of a global type G for a session k and let $\llbracket G \rrbracket_k$ be the set of local and buffer typings as obtained by the projection of G on each of its roles:

Definition 19 (Global Type Projection).

$$\llbracket G \rrbracket_k = \left\{ \begin{array}{l} \{k[A]: \llbracket G \rrbracket_A \mid A \in \text{roles}(G)\}, \\ \{b[k]_B^A: \llbracket G \rrbracket_B^A \mid A \in \text{roles}(G), B \in \text{roles}(G) \setminus \{A\}\} \end{array} \right\}$$

We say that a reduction is “at session k ” if it is obtained by consuming a communication term for session k (as in [43]), and we write $k \notin \Gamma$ when k does not appear in any local typing in Γ . Then we have:

Theorem 4 (Session Fidelity). *Let $\Gamma, \Gamma_k \vdash D, C$, $k \notin \Gamma$. Then, $D, C \rightarrow D', C'$ with a redex at session k implies that, for some G and Γ' , $k \notin \Gamma'$, (i) $\Gamma_k \subseteq \llbracket G \rrbracket_k$, (ii) $G \rightarrow G'$, (iii) $\Gamma'_k \subseteq \llbracket G' \rrbracket_k$, and (iv) $\Gamma', \Gamma'_k \vdash D', C'$.*

Theorem 4 states that all communications on sessions follow the expected protocols (Γ' may differ from Γ for the instantiation of a new variable).

We report the full proof of Theorem 3 and that of Theorem 4 (below) in Appendix E.1.

We can now present one of our major results: well-typed applied choreographies never deadlock when all the necessary participants are defined. Let the coherence predicate *co* be defined as follows:

Definition 20 (Coherence). $\text{co}(\Gamma)$ holds iff $\forall k \in \Gamma, \exists G$ s.t.

- $\tilde{l} : G\langle A|\tilde{B}|\tilde{C}\rangle \in \Gamma \wedge \tilde{C} = \tilde{B}$ and
- $\forall A \in \text{roles}(G), k[A] : T \in \Gamma \wedge T = \llbracket G \rrbracket_A \wedge \forall B \in \text{roles}(G)/\{A\}, b[k]_A^B = \llbracket G \rrbracket_A^B$

Coherence extends partial coherence to check that *i*) all needed services to start new sessions are present and *ii*) all the roles in every open session are correctly implemented by some processes. Coherent and well-typed systems are deadlock-free:

Theorem 5 (Deadlock-freedom). $\Gamma \vdash D, C$ and $\text{co}(\Gamma)$ imply that either (i) $C \equiv \mathbf{0}$ or (ii) there exist D' and C' such that $D, C \rightarrow D', C'$.

We report the full proof of Theorem 5 in Appendix E.2.

6.4 Endpoint Projection

We now present the Endpoint Projection and its properties. The EPP returns a correct composition¹² of endpoint choreographies that implements the behaviour of a given choreography. Intuitively, an endpoint (applied) choreography is an applied choreography that does not contain complete actions. Formally, let $\text{fp}(C)$ return the set of free processes in a choreography C :

Definition 21 (Endpoint Choreography [14]). We say that a choreography C is an endpoint choreography if C does not contain complete actions and one of the two following conditions holds:

$$i) C = \mathbf{acc} \ k : l.q[B]; C' \wedge \{q\} = \text{fp}(C'); \quad ii) \text{fp}(C) = \{p\}$$

An endpoint choreography defines the identity of only one process: either *i*) a service process or *ii*) an active process. Notably, our definition of EPP captures the description of the behaviour of single endpoints up to complete choreographies.

To give the definition of the EPP of a choreography, we first define the notion of *process projection*, which defines the behaviour of a single process p in a choreography C , written $\llbracket C \rrbracket_p$. We report the formalisation of $\llbracket C \rrbracket_p$ in Appendix D.6. Intuitively, the process projection follows the structure of the originating choreography. A (*start*) projects to a (*req*) on the active process and a set of (*acc*) terms on the service processes. Likewise, a (*com*) projects to a partial (*send*) for

¹²Here and in the following “composition” means “parallel composition”.

6.4. Endpoint Projection

the sender and a partial (*recv*) for the receiver. Any partial term is projected as it is for its respective process, following the structure of the choreography. We report below the projections of (*recv*) and (*cond*) terms which use the *merging* partial operator \sqcup [96] (reported in Appendix D.7) to merge the behaviour of all processes in their branches.

$$\begin{aligned} \llbracket k : p[A].e \rightarrow q[B].o(x); C \rrbracket_r &= \begin{cases} k : p[A].e \rightarrow B.o; \llbracket C \rrbracket_r & \text{if } r = p \\ k : A \rightarrow q[B].o(x); \llbracket C \rrbracket_r & \text{if } r = q \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \\ \llbracket \text{if } p.e \{C_1\} \text{ else } \{C_2\} \rrbracket_r &= \begin{cases} \text{if } p.e \{ \llbracket C_1 \rrbracket_r \} \text{ else } \{ \llbracket C_2 \rrbracket_r \} & \text{if } r = p \\ \llbracket C_1 \rrbracket_r \sqcup \llbracket C_2 \rrbracket_r & \text{otherwise} \end{cases} \end{aligned}$$

$C \sqcup C'$ is defined only for endpoint choreographies and returns an endpoint choreography isomorphic to C and C' up to branching, i.e., with all branches on distinct operations. We show the only special rule of \sqcup for merging (*recv*) terms.

$$\begin{aligned} k : A \rightarrow p[B]. \{o_i(x_i); C_i\}_{i \in I} \sqcup k : A \rightarrow q[B]. \{o_j(x_j); C'_j\}_{j \in J} = \\ k : A \rightarrow p[B]. \left\{ \begin{array}{l} \{o_i(x_i); C_i\}_{i \in I/J} \cup \{o_i(x_i); C'_i\}_{i \in J/I} \\ \cup \{o_i(x_i); C_i \sqcup C'_i\}_{i \in I \cap J} \end{array} \right\} \end{aligned}$$

On the definition of process projection, we define the EPP of a whole system. Here, and in the following, $\lfloor C \rfloor_l$ is the grouping operator that returns the set of grouped (bound) service processes at the same location. We show only the rules (reported in full in Appendix D.8) of $\lfloor C \rfloor_l$ for (*start*) and (*acc*) as the other rules are trivially recursive applications.

$$\begin{aligned} \lfloor \mathbf{start} \ k : p[D] \Leftrightarrow \widetilde{l.q[B]}; C \rfloor_l &= \lfloor \mathbf{acc} \ k : \widetilde{l.q[B]}; C \rfloor_l \\ \lfloor \mathbf{acc} \ k : \widetilde{l.q[B]}; C \rfloor_l &= \begin{cases} \{r\} \cup \lfloor C \rfloor_l & \text{if } l.r[A] \in \widetilde{l.q[B]} \\ \lfloor C \rfloor_l & \text{otherwise} \end{cases} \end{aligned}$$

Definition 22 (Endpoint Projection). *Let C be a term of AC. The endpoint projection of C , denoted by $\llbracket C \rrbracket$, is defined as:*

$$\llbracket C \rrbracket = \prod_{p \in \text{fp}(C)} \llbracket C \rrbracket_p \mid \prod_l \left(\bigsqcup_{p \in \lfloor C \rfloor_l} \llbracket C \rrbracket_p \right)$$

The EPP of a choreography C is the composition of the projections of all active processes and the merged projections of the service processes in C .

6.4.1 Projection Example

$$\begin{aligned}
 \llbracket C \rrbracket_a &= \begin{cases} 1. \mathbf{acc} \ k_d : l_A.a[A]; \\ 2. \ k_d : C \rightarrow a[A].get(\underline{\mathbf{req}}); \\ 3. \ \text{if } a.isValid(\underline{\mathbf{req}}) \{ \\ 4. \ \ k_d : a[A].\underline{\mathbf{req.rsc}} \rightarrow S.ok; \\ 5. \ \ k_d : a[A] \rightarrow C.ok \\ 6. \ } \text{ else } \{ \\ 7. \ \ k_d : a[A].\underline{\mathbf{req.rsc}} \rightarrow S.ko; \\ 8. \ \ k_d : a[A] \rightarrow C.ko \\ 9. \ } \end{cases} & \llbracket C \rrbracket_{dm} = \begin{cases} 1. \ \mathbf{acc} \ k_d : l_S.dm[DM]; \\ 2. \ \ k_d : A \rightarrow dm[DM].\{ \\ 3. \ \ ok(\underline{\mathbf{rsc}}); \\ 4. \ \ k_d : dm[DM].\logok(\underline{\mathbf{req}}) \rightarrow L.log; \\ 5. \ \ \dots, \\ 6. \ \ ko(\underline{\mathbf{rsc}}); \\ 7. \ \ k_d : dm[DM].\logko(\underline{\mathbf{req}}) \rightarrow L.log \\ 8. \ } \end{cases} \\
 \llbracket C \rrbracket_l &= \begin{cases} 1. \ \mathbf{acc} \ k_d : l_L.l[L]; \\ 2. \ \ k_d : S \rightarrow l[L].log(\underline{\mathbf{log}}); \end{cases}
 \end{aligned}$$

Figure 6.10: EPPs of $C_c \mid C_s$. Projection on a , dm (excepts), and l .

As an example, we project the choreography $C = C_c \mid C_s$ presented in § 6.2.3. Applying the definition of EPP to C we obtain

$$\llbracket C \rrbracket = \llbracket C \rrbracket_c \mid \llbracket C \rrbracket_f \mid \llbracket C \rrbracket_a \mid \llbracket C \rrbracket_{dm} \mid \llbracket C \rrbracket_l$$

Figure 6.10 reports some excerpts of $\llbracket C \rrbracket$, i.e., $\llbracket C \rrbracket_a$, $\llbracket C \rrbracket_{dm}$, and $\llbracket C \rrbracket_l$. We illustrate how $\llbracket C \rrbracket$ implements the same behaviour of C .

Start. The start of a new session, like the one at Lines 1 of C_c and C_s , applies the same effect in the projection. Let $\llbracket C \rrbracket'_p$ be the continuation of the process projection of C on p after Line 1. The execution of Lines 1 of the process projections composing $\llbracket C \rrbracket$ is:

$$\begin{aligned}
 D, \quad & \mathbf{req} \ k_d : c[C] \Leftrightarrow l_A.a[A], l_{DM}.dm[DM], l_L.l[L]; \llbracket C \rrbracket'_c \\
 & \mid \mathbf{acc} \ k_d : l_A.a[A]; \llbracket C \rrbracket'_a \mid \mathbf{acc} \ k_d : l_{DM}.dm[DM]; \llbracket C \rrbracket'_s \rightarrow D', \llbracket C \rrbracket'_c \mid \llbracket C \rrbracket'_a \mid \llbracket C \rrbracket'_{dm} \mid \llbracket C \rrbracket'_l \mid \llbracket C \rrbracket'_f \\
 & \mid \mathbf{acc} \ k_d : l_L.l[L]; \llbracket C \rrbracket'_l \mid \llbracket C \rrbracket'_f
 \end{aligned}$$

Let the location of process c be l_C in D . The effect of the start is $\delta = \mathbf{start} \ k_d : l_C.c[C], l_A.a[A], l_{DM}.dm[DM], l_L.l[L]$ which can generate the same D' as the start at Lines 1 of C_c and C_s .

Communications. The EPP projects a complete communication to a partial send for the sender and a partial receive for the receiver. Line 6 of C_s is $k_d : dm[DM].\logok(\underline{\mathbf{rsc}}) \rightarrow l[L].log(\underline{\mathbf{log}})$ and in $\llbracket C \rrbracket$ it is projected into a partial send for dm (at Line 4 of $\llbracket C \rrbracket_{dm}$) and a partial receive for l (at Line 2 of $\llbracket C \rrbracket_l$). The semantics of Line 6 of C_s and that of Line 4 of $\llbracket C \rrbracket_{dm}$ and Line 2 of $\llbracket C \rrbracket_l$ is equal

6.4. Endpoint Projection

as the complete action breaks into a partial send, equal to the effect of Line 4 of $\llbracket C \rrbracket_{dm}$, and continues as a partial receive, equal to $\llbracket C \rrbracket_l$.

Conditionals. The EPP of a conditional merges the behaviour of all processes in its branches as branched receptions. For example, at Line 3 of C_s , process a evaluates a condition that either branches in the behaviour described at Lines 4–15 or at Lines 17–19. At Line 3 of $\llbracket C \rrbracket_a$ the conditional is preserved verbatim. For the other processes present in the branches of the conditional, the EPP merges their behaviours on a branched reception: $\llbracket C \rrbracket_{dm}$ projects the conditional on the branches guarded by operations ok or ko (resp. at Lines 3 and 6). $\llbracket C \rrbracket_l$ merges the branches of the conditional into a single branch on operation log , at Line 2.

6.4.2 Properties

We present the properties of EPP. We prove that the EPP implements the same behaviour of its originating choreography, we build on the foundation that the EPP of a choreography is still typable and then we establish a bisimilarity relation between the semantics of the EPP and the originating choreography.

First we state our type preservation result using the *minimal typing* of choreographies \vdash_{\min} , which types the branches in rules $\llbracket T \rrbracket_{SEND}$ and $\llbracket T \rrbracket_{RECV}$ using the respective minimal branch types. As in [13], \vdash_{\min} (defined in Appendix E.3.1) takes into account that, due to merging, the EPP of a complete choreography may still offer branches that the originating choreography discarded with a conditional. Thus we have:

Theorem 6 (EPP Typing Preservation). $\Gamma \vdash_{\min} D, C$ implies $\llbracket \Gamma \rrbracket \vdash_{\min} D, \llbracket C \rrbracket$

Where $\llbracket \Gamma \rrbracket$, defined in Appendix E.3.2, replaces the typings of recursive procedures in Γ with the typing of each procedure for each endpoint process taking part in it.

Combining the above Theorem with Theorem 4 we can prove:

Theorem 7 (EPP Operational Correspondence). *Let D, C be well-typed. Then,*

1. (Completeness) $D, C \rightarrow D', C'$ implies $D, \llbracket C \rrbracket \rightarrow D', C''$ and $\llbracket C' \rrbracket \prec C''$.
2. (Soundness) $D, \llbracket C \rrbracket \rightarrow D', C'$ implies $D, C \rightarrow D', C''$ and $\llbracket C'' \rrbracket \prec C'$.

In the theorem above $C \prec C'$ is the *pruning relation* (see Appendix E.3.4), a strong typed bisimilarity [96] such that C has some unused branches and always-available accepts.

We report the full proofs of Theorem 6 and Theorem 7 in Appendix E.3.

6.5 Dynamic Correlation Calculus Language

We introduce the Dynamic Correlation Calculus (DCC), the target language of our compilation, modelled on the Correlation Calculus [98] (CC) and extended to support the dynamic creation of queues. The reason to target CC is that it is the formal model of the executable language Jolie [92]. In this way, our results are immediately applicable to an actual implemented language, yet preserving the formality and simplicity of a theoretical approach. Moreover, the semantics of CC formalises message correlation à-la SOC, as in the standard service-oriented language BPEL [36], by following similar concepts to those that we used in our semantics for AC. However, it is complex to implement multiparty sessions in CC because processes are statically linked to a single queue and, through it, to a single correlation key. For these reasons, we extend CC to DCC by enabling processes to create queues at runtime and selectively read messages from them.

$$\begin{array}{l}
 S ::= \langle B_s, P \rangle_l \quad (\textit{service}) \\
 \quad | S \mid S' \quad (\textit{network}) \\
 \\
 B ::= ?@_{e_1}(e_2); B \quad (\textit{request}) \quad | \sum_i [o_i(x_i) \textit{ from } e] \{B_i\} \quad (\textit{choice}) \\
 \quad | o(x) \textit{ from } e; B \quad (\textit{input}) \quad | o@_{e_1}(e_2) \textit{ to } e_3; B \quad (\textit{output}) \\
 \quad | \textit{ if } e \{B_1\} \textit{ else } \{B_2\} \quad (\textit{cond}) \quad | X \quad (\textit{call}) \\
 \quad | \textit{ def } X = B' \textit{ in } B \quad (\textit{def}) \quad | x = e; B \quad (\textit{assign}) \\
 \quad | \textit{ cq}(x); B \quad (\textit{cqueue}) \quad | \mathbf{0} \quad (\textit{inact})
 \end{array}$$

Figure 6.11: Dynamic Correlation Calculus, syntax

Syntax. The syntax of DCC (in Figure 6.11) is divided in two layers: *services*, ranged over by S , and *processes*, ranged over by P . The term (*service*) describes a service, located at l , as a container of a *start behaviour* B_s and a (system of) processes P . A start behaviour creates new processes on request: following the syntax of B_s , a service can spawn a new process that stores the message of the request under the bound variable x and implements the behaviour B (see below). The term (*network*) supports interactions among services. The process layer P defines the structure of processes, which run inside services, and their composition. In the syntax, a process is the association of a behaviour B , a state t , and a queue map M . The state t and the map M are defined exactly like the states and queue maps found in AC (§ 6.2). Still from AC, we also make use of names for operations (o), procedures (X), and variables (x , which we recall are paths). Expressions, ranged over by e , are evaluated at runtime on the state of the re-

lated process. Terms (*input*) and (*output*) model communications. In (*input*) the process stores in x a message from the head of the queue correlating with e and expectedly received on operation o . Dually, (*output*) models the delivery of a message on operation o with content e_2 ; e_1 defines the location of the service where the addressee (process) is running, whilst e_3 is the key that correlates with the receiving queue of the addressee. The term (*choice*) models a branching input-choice. The argument of the choice is modelled after (*input*). If one of the inputs can execute — receiving a message from the queue correlating with e on operation o_i — it discards all other receptions and executes the related behaviour B_i . Term (*request*) is the dual of (*accept*) and asks the service located at e_1 to spawn a new process, passing to it the message in e_2 . Finally (*cqueue*) models the creation of a new queue. After the creation, variable x contains the key that correlates with the new queue. All other terms of the syntax are standard.

Semantics. In Figure 6.12 we report the semantics of DCC, given as rules for a reduction relation \rightarrow . Relation \rightarrow is closed under a structural congruence \equiv defined in the standard way (reported in Appendix D.9); in particular, it supports commutativity and associativity for parallel composition. We comment the rules. Rule $\llbracket^{\text{DCC}}\rrbracket_{\text{RECV}}$ models the reception of messages: if the queue correlating with e has a message on operation o then the message is removed from the queue and its content is assigned to the variable x in the state of the process. Rule $\llbracket^{\text{DCC}}\rrbracket_{\text{CQ}}$ adds to M an empty queue (ε) correlating with a fresh key, stored in x . The key is unique within the service of the process, to avoid ambiguity, but we impose no requirements on its structure (it can be any tree as long as it is unique in the enclosing service). Rule $\llbracket^{\text{DCC}}\rrbracket_{\text{SEND}}$ models the delivery of a message between processes in different services. In the Rule, the message from the sender is added to (the end of) the correlating queue of the receiver. Similarly, Rule $\llbracket^{\text{DCC}}\rrbracket_{\text{START}}$ models the matching of a request to create a new process with the service that accepts it. The newly created process has B as its behaviour, a state initialised with the content of the request message, and an empty queue map.

6.6 Compiler from AC to DCC and Properties

We present our main result: a formally-correct compiler from AC to DCC that models how applied choreographies can be implemented in real-world languages.

6.6.1 Compiler

$\llbracket D, C \rrbracket^{\Gamma}$ denotes the compilation of a well-typed running choreography D, C into an operationally-equivalent network of DCC services. Before giving the formal

$$\begin{array}{c}
 \frac{t_c = \text{eval}(e, t) \quad M(t_c) = (o, t') :: \tilde{m}}{o(x) \text{ from } e; B \cdot t \cdot M \rightarrow B \cdot t \triangleleft (x, t') \cdot M[t_c \mapsto \tilde{m}]} \quad [\text{DCC}|_{\text{RECV}}] \\
 \\
 \frac{B \cdot t \cdot M \rightarrow B' \cdot t' \cdot M'}{\text{def } X = B_1 \text{ in } B \cdot t \cdot M \rightarrow \text{def } X = B_1 \text{ in } B' \cdot t' \cdot M'} \quad [\text{DCC}|_{\text{CTX}}] \\
 \\
 \frac{i = 1 \text{ if } \text{eval}(e, t) = \text{true}, i = 2 \text{ otherwise}}{\text{if } e \{B_1\} \text{ else } \{B_2\} \cdot t \cdot M \rightarrow B_i \cdot t \cdot M} \quad [\text{DCC}|_{\text{COND}}] \\
 \\
 \frac{j \in I \quad t_c = \text{eval}(e, t) \quad M(t_c) = (o_j, t') :: \tilde{m}}{\sum_{i \in I} [o_i(x_i) \text{ from } e] \{B_i\} \cdot t \cdot M \rightarrow B_j \cdot t \triangleleft (x_j, t') \cdot M[t_c \mapsto \tilde{m}]} \quad [\text{DCC}|_{\text{CHOICE}}] \\
 \\
 \frac{t' = \text{eval}(e, t)}{x = e; B \cdot t \cdot M \rightarrow B \cdot t \triangleleft (x, t') \cdot M} \quad [\text{DCC}|_{\text{ASSIGN}}] \\
 \\
 \frac{P \rightarrow P'}{\langle B_s, P \mid P_1 \rangle_l \rightarrow \langle B_s, P' \mid P_1 \rangle_l} \quad [\text{DCC}|_{\text{PPAR}}] \\
 \\
 \frac{P = \text{cq}(x); B \cdot t \cdot M \quad t_c \notin \bigcup_i \text{dom}(M_i) \cup \text{dom}(M) \quad t' = t \triangleleft (x, t_c)}{\langle B_s, P \mid \prod_i B_i \cdot t_i \cdot M_i \rangle_l \rightarrow \langle B_s, B \cdot t' \cdot M[t_c \mapsto \varepsilon] \mid \prod_i B_i \cdot t_i \cdot M_i \rangle_l} \quad [\text{DCC}|_{\text{CQ}}] \\
 \\
 \frac{P = o @_{e_1}(e_2) \text{ to } e_3; B \cdot t \cdot M \quad \text{eval}(e_1, t) = l \quad \text{eval}(e_3, t) = t_c}{\text{eval}(e_2, t) = t_m \quad M'' = M'[t_c \mapsto M'(t_c) :: (o, t_m)]} \quad [\text{DCC}|_{\text{INSEND}}] \\
 \\
 \frac{P = o @_{e_1}(e_2) \text{ to } e_3; B \cdot t \cdot M \quad \text{eval}(e_1, t) = l' \quad \text{eval}(e_3, t) = t_c}{\text{eval}(e_2, t) = t_m \quad M'' = M'[t_c \mapsto M'(t_c) :: (o, t_m)]} \quad [\text{DCC}|_{\text{SEND}}] \\
 \\
 \frac{\langle B_s, P \mid P_1 \rangle_l \mid \langle B'_s, B' \cdot t' \cdot M' \mid P_2 \rangle_{l'}}{\rightarrow \langle B_s, B \cdot t \cdot M \mid P_1 \rangle_l \mid \langle B'_s, B' \cdot t' \cdot M'' \mid P_2 \rangle_{l'}} \\
 \\
 \frac{P_1 = ? @_{e_1}(e_2); B_1 \cdot t_1 \cdot M_1 \quad \text{eval}(e_1, t_1) = l \quad Q = B \cdot t_\perp \triangleleft (x, \text{eval}(e_2, t_1)) \cdot \emptyset}{\langle ! (x); B, P \rangle_l \mid \langle B'_s, P_1 \mid P_2 \rangle_{l'} \rightarrow \langle ! (x); B, Q \mid P \rangle_l \mid \langle B'_s, B_1 \cdot t_1 \cdot M_1 \mid P_2 \rangle_{l'}} \quad [\text{DCC}|_{\text{START}}] \\
 \\
 \frac{P = ? @_{e_1}(e_2); B \cdot t \cdot M \quad \text{eval}(e_1, t) = l \quad Q = B \cdot t_\perp \triangleleft (x, \text{eval}(e_2, t)) \cdot \emptyset}{\langle ! (x); B', P \mid P_1 \rangle_l \rightarrow \langle ! (x); B', Q \mid B \cdot t \cdot M \mid P_1 \rangle_l} \quad [\text{DCC}|_{\text{INSTART}}] \\
 \\
 \frac{S \equiv S_1 \quad S_1 \rightarrow S'_1 \quad S'_1 \equiv S'}{S \rightarrow S'} \quad [\text{DCC}|_{\text{EQ}}] \quad \frac{S_1 \rightarrow S'_1}{S_1 \mid S \rightarrow S'_1 \mid S} \quad [\text{DCC}|_{\text{SPAR}}]
 \end{array}$$

Figure 6.12: Correlation Calculus, semantics

6.6. Compiler from AC to DCC and Properties

definition of $\boxed{D, C}^\Gamma$, we need to introduce some additional notation. Let $\mathcal{C}[\cdot]$ be an AC context defined as usual: $\mathcal{C}[\cdot] ::= \cdot \mid \mathcal{C}[\cdot] \mid C \mid C \mid \mathcal{C}[\cdot]$. We then formalise the filtering operators $C|_l$ and $C|_p$ as follows:

$$\begin{aligned} C|_l &= C' \text{ if } C = \mathcal{C}[C'] \text{ and } C' = \mathbf{acc} \ k : l.p[A]; C'', \mathbf{0} \text{ otherwise} \\ C|_p &= C' \text{ if } C = \mathcal{C}[C'] \text{ and } \{p\} = \text{fp}(C'), \mathbf{0} \text{ otherwise} \end{aligned}$$

Intuitively, $C|_l$ returns the accept term at location l in C and $C|_p$ returns the endpoint choreography of process p in C . Next, we denote with \boxed{C}^Γ the compilation of a process projection defined on the rules reported in Figure 6.13 (and commented below). Below, we define our compiler abusing the notation $l \in \Gamma$ to say that $p@l \in \Gamma$ for some p or that there is a service typing in Γ containing l .

$$\begin{array}{l}
 \boxed{\text{req } k : p[A] \Leftrightarrow \widetilde{l.B}; C}^\Gamma = \\
 \quad \text{start}(k, l'.A, \widetilde{l.B}); \boxed{C}^\Gamma, p @ l' \in \Gamma \\
 \boxed{k : p[A].e \rightarrow B.o; C}^\Gamma = \\
 \quad o @ \mathbf{k.A.B}.l(e) \text{ to } \mathbf{k.A.B}; \boxed{C}^\Gamma \\
 \boxed{\text{def } X = C' \text{ in } C}^\Gamma = \\
 \quad \text{def } X = \boxed{C'}^\Gamma \text{ in } \boxed{C}^\Gamma \\
 \boxed{X}^\Gamma = X \\
 \boxed{\text{acc } k : l.q[B]; C}^\Gamma = \\
 \quad \text{accept}(k, B, \Gamma(\widetilde{l})); \boxed{C}^\Gamma, l \in \widetilde{l}, \widetilde{l} \in \Gamma \\
 \boxed{k : A \rightarrow q[B].\{o_i(x_i); C_i\}_{i \in I}}^\Gamma = \\
 \quad \sum_{i \in I} [o_i(x_i) \text{ from } \mathbf{k.A.B}] \{\boxed{C_i}^\Gamma\} \\
 \boxed{\text{if } p.e \{C_1\} \text{ else } \{C_2\}}^\Gamma = \\
 \quad \text{if } e \{\boxed{C_1}^\Gamma\} \text{ else } \{\boxed{C_2}^\Gamma\} \\
 \boxed{0}^\Gamma = 0 \\
 \\
 \text{start}(k, l_A.A, \widetilde{l_B.B}) = \\
 \underbrace{\bigcirc_{\mathbf{I} \in \{A, \widetilde{B}\}} \mathbf{k.I.l} = l_I}_{(s_1) \text{ locations}} ; \underbrace{\bigcirc_{\mathbf{I} \in \{\widetilde{B}\}} \left(\overbrace{\text{cq}(\mathbf{k.I.A})}^{(s_{2.1})} ; \overbrace{?@ \mathbf{k.I.l}(\mathbf{k})}^{(s_{2.2})} ; \overbrace{\text{sync}(\mathbf{k}) \text{ from } \mathbf{k.I.A}}^{(s_{2.3})} \right)}_{(s_2) \text{ creation of correlation keys and service processes}} ; \underbrace{\bigcirc_{\mathbf{I} \in \{\widetilde{B}\}} \text{start}@ \mathbf{k.I.l}(\mathbf{k}) \text{ to } \mathbf{k.A.I}}_{(s_3) \text{ session start}} \\
 \\
 \text{accept}(k, B, G(A|\widetilde{C}|\widetilde{D})) = \\
 \underbrace{!(\mathbf{k})}_{(a_1)} ; \underbrace{\bigcirc_{\mathbf{I} \in \{A, \widetilde{C}\}/\{B\}} \left(\text{cq}(\mathbf{k.I.B}) \right)}_{(a_2)} ; \underbrace{\text{sync}@ \mathbf{k.A.l}(\mathbf{k}) \text{ to } \mathbf{k.B.A}}_{(a_3)} ; \underbrace{\text{start}(\mathbf{k}) \text{ from } \mathbf{k.A.B}}_{(a_4)}
 \end{array}$$

Figure 6.13: Compiler from AC to DCC.

Definition 23 (Compilation). *Let C be a parallel composition of endpoint choreographies and let $\Gamma \vdash D, C$ for some Γ and D . The compilation $\boxed{D, C}^\Gamma$ is:*

$$\boxed{D, C}^\Gamma = \prod_{l \in \Gamma} \left\langle \boxed{C|_l}^\Gamma, \prod_{p \in D(l)} \boxed{C|_p}^\Gamma \cdot D(p).\text{st} \cdot D(p).\text{que} \right\rangle_l$$

Intuitively, for each service $\langle B_s, P \rangle_l$ in the compiled network: *i*) the start behaviour B_s is the compilation of the endpoint choreography in C accepting the creation of processes at location l ; *ii*) P is the parallel composition of the compilation of all active processes located at l , equipped with their respective states and queue maps according to D . Let us now comment the rules in Figure 6.13. We use the auxiliary notation \bigcirc for DCC behaviours, defined as $\bigcirc_{i \in [1, n]} (B_i) = B_1; \dots; B_n$.

Requests. Function `start` defines the compilation of (*req*) terms. `start` compiles (*req*) terms to create the queues and a part of the session descriptor of a valid session support (see Definition 16) for the starter. Given a session identifier k , the located role of the starter ($l_A.A$), and the other located roles in the session

$(\widetilde{l_B.B})$, start returns DCC code that: (s_1) includes in the session descriptor all the locations of the processes involved in the session and $(s_{2.1})$ all the keys correlating with the queues of the starter for the session. Then, $(s_{2.2})$ it requests the creation of all the service processes for the session, $(s_{2.3})$ waits for them to be ready using the reserved operation *sync* and, finally, (s_3) sends them the complete session descriptor obtained after receiving from all processes their correlation keys (in the *sync* step).

Accepts. (*acc*) terms define the start behaviour of the service accepting the creation of processes at a location. Given a session identifier k , the role B of the service process, and the service typing $G\langle A|\widetilde{C}|\widetilde{D}\rangle$ of the location, function *accept* defines the compilation of (*acc*) terms. *accept* complements function *start* by compiling the code that (a_1) accepts the spawn of a new process which, in turn, (a_2) creates its queues (including their keys in the session descriptor passed by the starter), (a_3) returns the session descriptor to the starter, and (a_4) waits for the signal to start the session.

Other terms. We compile (*send*) terms to (*output*) terms. Notably, the compiled code contains the same elements used by the semantics of AC to implement correlation, i.e., the location of the receiver ($\mathbf{k.B.I}$) and the key that correlates with its queue ($\mathbf{k.A.B}$). (*recv*) compiles to (*choice*), which defines the path ($\mathbf{k.A.B}$) of the key correlating with the receiving queue, used also in the semantics of AC. Terms (*def*), (*cond*), (*call*), and (*inact*) compile to the relative terms in DCC.

6.6.2 Example of Compilation

We report in Figure 6.14 the compilation of the example in § 6.2.3.

We associate a default deployment D to $\llbracket C \rrbracket$ such that for some Γ and D *i)* $\Gamma \vdash D, \llbracket C \rrbracket$ holds; *ii)* the location of process c is l_c . l_c is also the location of the service process f , which lets us to show how the compilation of active and service processes work. Figure 6.14 reports the result of the compilation $\llbracket D, \llbracket C \rrbracket \rrbracket^\Gamma$ where, on top it shows the composition of compiled services and below it, from left to right, the compiled behaviours of processes c , a , and l .

Compiled Services. The network of compiled services, on top of Figure 6.14, shows that there is a compiled service for all locations $(l_c, l_A, l_{DM}, \text{ and } l_L)$ in C . In particular, S_c , the service located at l_c , contains *i)* as its start behaviour the compilation of the process projection of f (which accepts requests on l_c in C) and *ii)* a running process whose behaviour is the compilation of the process projection of c , and whose state and queue map are those of c in D .

$$\begin{array}{c}
 \overbrace{\langle \llbracket C \rrbracket_f^\Gamma, \llbracket C \rrbracket_c^\Gamma \cdot D(c).st \cdot D(c).que \rangle_{l_c}}^{S_c} \mid \overbrace{\langle \llbracket C \rrbracket_a^\Gamma, \mathbf{0} \rangle_{l_A}}^{S_a} \mid \overbrace{\langle \llbracket C \rrbracket_{dm}^\Gamma, \mathbf{0} \rangle_{l_{DM}}}^{S_{dm}} \mid \overbrace{\langle \llbracket C \rrbracket_l^\Gamma, \mathbf{0} \rangle_{l_L}}^{S_l} \\
 \\
 \llbracket C \rrbracket_c^\Gamma = \begin{cases} 1. \underline{\mathbf{kd.C.l}} = l_c; \\ 2. \underline{\mathbf{kd.A.l}} = l_A; \\ 3. \underline{\mathbf{kd.DM.l}} = l_{DM}; \\ 4. \underline{\mathbf{kd.L.l}} = l_L; \\ 5. \text{cq}(\underline{\mathbf{kd.A.C}}); \\ 6. ?@\underline{\mathbf{kd.A.l}}(\underline{\mathbf{kd}}); \\ 7. \text{sync}(\underline{\mathbf{kd}}) \text{ from } \underline{\mathbf{kd.A.C}}; \\ 8. \dots \\ 9. \dots \\ 10. \text{start}@\underline{\mathbf{kd.A.l}}(\underline{\mathbf{kd}}) \text{ to } \underline{\mathbf{kd.C.A}}; \\ 11. \text{start}@\underline{\mathbf{kd.S.l}}(\underline{\mathbf{kd}}) \text{ to } \underline{\mathbf{kd.C.S}}; \\ 12. \text{start}@\underline{\mathbf{kd.C.l}}(\underline{\mathbf{kd}}) \text{ to } \underline{\mathbf{kd.C.L}}; \\ 13. \text{get}@\underline{\mathbf{kd.A.l}}(\text{mkReq}()) \text{ to } \underline{\mathbf{kd.C.A}}; \\ 14. \dots \end{cases} & \llbracket C \rrbracket_a^\Gamma = \begin{cases} 1. !(\underline{\mathbf{kd}}); \\ 2. \text{cq}(\underline{\mathbf{kd.C.A}}); \\ 3. \text{cq}(\underline{\mathbf{kd.S.A}}); \\ 4. \text{cq}(\underline{\mathbf{kd.L.A}}); \\ 5. \text{sync}@\underline{\mathbf{kd.C.l}}(\underline{\mathbf{kd}}) \text{ to } \underline{\mathbf{kd.A.C}}; \\ 6. \text{start}(\underline{\mathbf{kd}}) \text{ from } \underline{\mathbf{kd.C.A}}; \\ 7. \text{get}(\underline{\mathbf{req}}) \text{ from } \underline{\mathbf{kd.C.A}}; \\ 8. \text{if } \text{isValid}(\underline{\mathbf{req}}) \{ \\ 9. \quad \text{ok}@\underline{\mathbf{kd.S.l}}(\underline{\mathbf{req.rsc}}) \text{ to } \underline{\mathbf{kd.A.S}}; \\ 10. \quad \text{ok}@\underline{\mathbf{kd.C.l}}() \text{ to } \underline{\mathbf{kd.A.C}} \\ 11. \} \text{ else } \{ \\ 12. \quad \text{ko}@\underline{\mathbf{kd.S.l}}(\underline{\mathbf{req.rsc}}) \text{ to } \underline{\mathbf{kd.A.S}}; \\ 13. \quad \text{ko}@\underline{\mathbf{kd.C.l}}() \text{ to } \underline{\mathbf{kd.A.C}} \\ 14. \} \end{cases} \\
 \\
 \llbracket C \rrbracket_l^\Gamma = \begin{cases} 1. !(\underline{\mathbf{kd}}); 2. \text{cq}(\underline{\mathbf{kd.C.L}}); 3. \text{cq}(\underline{\mathbf{kd.A.L}}); 4. \text{cq}(\underline{\mathbf{kd.S.L}}); 5. \text{sync}@\underline{\mathbf{kd.C.l}}(\underline{\mathbf{kd}}) \text{ to } \underline{\mathbf{kd.L.C}}; \\ 6. \text{start}(\underline{\mathbf{kd}}) \text{ from } \underline{\mathbf{kd.C.L}}; 7. \text{log}(\underline{\mathbf{log}}) \text{ from } \underline{\mathbf{kd.S.L}} \end{cases}
 \end{array}$$

 Figure 6.14: Compilation of $C = C_c \mid C_s$, behaviours of c (excepts), a , and l .

Starting a new session. Lines 1-12 of $\llbracket C \rrbracket_c^\Gamma$ show an excerpt of the compiled code of the request to start session k_d for process c : $\text{req } k_d : c[C] \Leftrightarrow l_A.A, l_{DM}.DM, l_L.L$. As defined by function `start`, at Lines 1-4 the process sets the locations of all involved roles in the Session Descriptor, rooted in $\underline{\mathbf{kd}}$. Lines 5-7 define the interaction to start the process playing role A: line-wise, the process creates the queue to received from role A, it requests the creation of a new process at the relative location $\underline{\mathbf{kd.A.l}}$, passing to it the Session Descriptor, and, it waits for the process to return the updated Session Descriptor. Lines 10-12 signal the start of the session to all the started processes. Lines 1-6 of $\llbracket C \rrbracket_a^\Gamma$ and $\llbracket C \rrbracket_l^\Gamma$ show the code that complements the start. At the end of the procedure, D , $\llbracket C \rrbracket$ and $\llbracket D, \llbracket C \rrbracket \rrbracket^\Gamma$ properly set a Session Support for session k_d .

In-session message passing. We show that the originating choreography and its compilation handle communications similarly. We compare the execution of Lines 2 of $\llbracket C \rrbracket_c$ ¹³ and $\llbracket C \rrbracket_a$ with that of Line 13 of $\llbracket C \rrbracket_c^\Gamma$ and Line 7 of $\llbracket C \rrbracket_a^\Gamma$. Let D be the deployment associated with $\llbracket C \rrbracket'$, reduction of $\llbracket C \rrbracket$ to the considered Lines. $D, \llbracket C \rrbracket' \rightarrow D', \llbracket C \rrbracket''$ applying Rule $\llbracket C \rrbracket_{\text{SEND}}$ on Line 2 of $\llbracket C \rrbracket_c$. In D' the queue in which a receives from c contains the sent message: $D'(a).que(t_c) = (\text{get}, t_m)$, t_c correlation key rooted at $\underline{\mathbf{kd.C.A}}$ in the state of c and t_m evaluation

¹³Although not shown here, Line 2 of $\llbracket C \rrbracket_c$ corresponds to Line 2 of C_c as partial actions are projected verbatim on their process by the EPP.

of function $\text{mkReq}()$. Similarly, Line 13 of $\llbracket C \rrbracket_c^\Gamma$ reduces with Rule $[\text{DCC}]_{\text{SEND}}$. The Rule delivers the message on operation get , carrying the evaluation of e , to the service located at $\underline{\text{kd.A.I}}$, which contains all service processes compilation of a , and in the queue correlating with $\underline{\text{kd.A.C}}$. By the compilation, the compiled process relative to c has the state of c , hence, the value rooted at $\underline{\text{kd.A.C}}$ is equal to t_c , $\underline{\text{kd.A.I}}$ is equal to the location of A , and the evaluation of e is equal to t_m . Moreover, since the network has just started (at Line 13 of $\llbracket C \rrbracket_c^\Gamma$ and correspondent), we know that the receiving process had no message in the queue correlating with $\underline{\text{kd.A.C}}$, thus, after the reduction, it only has message (get, t_m) . The state of D' , $\llbracket C \rrbracket''$ corresponds to the state of $\llbracket D', \llbracket C \rrbracket'' \rrbracket^\Gamma$. On this result, if $D', \llbracket C \rrbracket''$ reduces with Rule $[\text{C}]_{\text{RECV}}$, $\llbracket D', \llbracket C \rrbracket'' \rrbracket^\Gamma$ can apply Rule $[\text{DCC}]_{\text{RECV}}$ and we are sure that *i*) the consumed message by the receiver (a and its compiled correspondent) is (get, t_m) , *ii*) the relative queues are now empty, and *iii*) the receiver included in its state t_m under path $\underline{\text{req}}$.

6.6.3 Properties

We present our main result: the operational correspondence between the semantics of a well-typed running Applied Choreography and the semantics of its compilation.

Theorem 8 (Compilation Operational Correspondence). *Let C be a composition of endpoint choreographies such that $\Gamma \vdash D, C$. Then we have that:*

1. (Completeness) $D, C \rightarrow D', C'$ implies $\llbracket D, C \rrbracket^\Gamma \rightarrow^+ \llbracket D', C' \rrbracket^{\Gamma'}$ for some Γ' such that $\Gamma' \vdash D', C'$.
2. (Soundness) $\llbracket D, C \rrbracket^\Gamma \rightarrow^* S$ implies (i) $D, C \rightarrow^* D', C'$ and (ii) $S \rightarrow^* \llbracket D', C' \rrbracket^{\Gamma'}$ for some D', C' and Γ' such that $\Gamma' \vdash D', C'$.

We report the full proof of Theorem 8 in Appendix E.4.

Definition 23 is defined only for endpoint choreographies, however, using EPP, we can always reconduce a choreography to this case.

Lemma 2. *Let $\Gamma \vdash C$ then $\llbracket C \rrbracket$ is a composition of endpoint choreographies.*

Proof. Trivial from Definition 21 and Definition 23. □

We conclude by answering our research question. The compilation of the EPP of an Applied Choreography is always a correct implementation.

Corollary 5 (Applied Choreographies). *Let $\Gamma \vdash D, C$. Then we have that:*

1. (Completeness) $D, C \rightarrow D', C'$ implies $\llbracket D, C \rrbracket^\Gamma \rightarrow^+ \llbracket D', C' \rrbracket^{\Gamma'}$ for some Γ' such that $\Gamma' \vdash D', C'$.
2. (Soundness) $\llbracket D, C \rrbracket^\Gamma \rightarrow^* S$ implies (i) $D, C \rightarrow^* D', C'$ and (ii) $S \rightarrow^* \llbracket D', C' \rrbracket^{\Gamma'}$ for some D', C' and Γ' such that $\Gamma' \vdash D', C'$.

Proof. Trivial from the application of Theorem 8 and Lemma 2. □

Part IV
Conclusion

Conclusion

In this final chapter we summarise the presented contributions, we discuss how they compare with related work, and define some direction of future investigation.

7.1 Adaptable Choreographies

In Chapters 4 and 5 we presented our theory for the safe dynamic update of distributed applications. Our model guarantees 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. We build on the strong theoretical results of our model and develop AIOCJ, a programming framework that offers the minimality and conciseness of the choreographic approach yet compiling choreographic specifications to executable distributed programs that can adapt in a consistent and coherent way at runtime.

The two approaches closest to ours we are aware of are in the area of multiparty session types [43, 13, 78], and deal with dynamic software updates [17] and with monitoring of self-adaptive systems [109]. The main difference between [17] and our approach is that [17] targets concurrent applications which are not distributed and 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 is impractical in a distributed setting. Furthermore, the language in [17] 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 not involved in a choice behave the same in the two branches. The approach in [109] is very different from ours, too. In particular, in [109] 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 [110], which uses types to ensure safe adaptation. However,

[110] allows updates only when no session is active, while we change the behaviour of running DIOCs. Our work shares with [14] the interest in choreographies composition. However, [14] 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 [16, 111] share with us the difficulties in coordinating the transition from the expected pattern to an alternative pattern, but in their case the error recovery pattern is known since the very beginning, thus considerably simplifying the analysis.

We briefly compare some works that exploit choreographic descriptions for adaptation, but with very different aims. For instance, [112] 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 [113] choreographies are used to propagate protocol changes to the other peers, while [114] 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.

In principle, our update mechanism can be used to inject guarantees of freedom from deadlocks and races into existing approaches to adaptation (cf. Chapter 3). 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, we apply our technique to the adaptation mechanism described in [82]. While applications in [82] are not distributed and there are no guarantees on the correctness of the application after adaptation, applications in AIOCJ, based on the same adaptation mechanisms, are distributed and free from deadlocks and races by construction.

Furthermore, on the website [89], we give examples of how to integrate our approach with distributed [70] and dynamic [87] Aspect-Oriented Programming (AOP) and with Context-Oriented Programming (COP) (cf. § 3.1.2). We report in Appendix C two examples which respectively model Aspect- and Context-Oriented Programming. 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. We are also planning to apply our techniques to multiparty session types [43, 13, 78]. The main challenge here is to deal with multiple interleaved sessions. An initial analysis of the problem is

presented in [115].

7.1.1 Adaptable Choreographies and their degree of flexibility

In Table 7.1, we report in § 3.1.3 how our solution of Adaptable Choreographies, as instantiated in AIOCJ, fits against the attributed of PAIS.

In summation, being a language-level solution, AIOCJ can implement the most part of decision deferral patterns, as reported in Table 3.1, from *Fully pre-specified processes* to *Ad-hoc composition*.

7.2 Applied Choreographies

In Chapter 6 we tackled the issue of formalising the implementation of choreography models. Our solution comprises *i*) a model of Applied Choreographies, *ii*) a type system to check AC against multiparty protocol specifications, and *iii*) a formally-correct compiler to obtain executable programs from choreographies.

The main novelty of AC regards its original semantics that abstracts the features of choreographies (message passing, creation of new sessions and processes) from their implementation. To this end we *i*) equip choreographies with a global deployment and *ii*) define a separate semantics of effects on deployments. Interestingly, deployments and effects hide the complexity of the implementation away from the semantics of choreographies. This means that we can preserve the semantics of choreographies and compose it with other definitions of deployment and semantics of effects to model different communication semantics (e.g., synchronous, asynchronous with buffers) and implementations (e.g., distributed objects [101]).

Deployments let us formalise how choreographies can go wrong (see §6.3.3) and prove that the theory of session types is useful not only to type communications on choreographies (like in [13, 14]) but also to check the correctness of deployments. We highlight that, except for the declaration of locations, AC has the same types and syntax from previous works [13, 14], i.e., developers only specify protocols and choreographies and do not deal with deployment information or correlation data.

Implementation Model. To the best of our knowledge, this is the first formal work on a compiler and an execution model for choreographies. All previous works on formal choreography languages only specify an EPP procedure towards a calculus based on name synchronisation, leaving the design of its concrete support to implementors.

Chor [99] and AIOCJ (cf. Chapter 5) are the only projects (that we are aware of) that aim at providing choreography languages with strong safety guarantees (e.g.,

Chapter 7. Conclusion

Attribute	Degree/type	Description
Freedom	<i>modelling and composition</i>	Updates can be introduced at runtime and they are not bound to a pre-specified set of behaviours. Moreover, updates can compose other updates, e.g., using scope properties.
Planning approach	<i>plan-driven, iterative and continuous, and ad-hoc</i>	It is possible to implement all planning approaches by controlling which updates are available during the execution. The <i>plan-driven</i> approach provides only those updates that belong to the pre-specified model, the <i>iterative and continuous</i> approach changes the updates at runtime to adapt a coarse plan (implemented in the choreography) to upcoming requirements, and the <i>ad-hoc</i> approach do not constrain the presence of updates wrt a plan.
Scope	<i>regions and entire processes</i>	The basic usage of scopes is to enclose a region of a choreography, yet it is possible to update the whole process by enclosing the whole choreography into a scope.
Process perspective	<i>behaviours, data, functions, operational decisions, time, and, actors (limited)</i>	Updates can bring into the original choreography new <i>behaviours</i> , use different <i>data</i> , introduce new <i>functions</i> (which are linked to their <i>implementations</i>) and thanks to the presence of the Environment, it is also possible to encode <i>time-triggered</i> events. It is possible to change which <i>actors</i> perform tasks, yet they are limited to those that belonged to the original choreography.
Making and support of decision	<i>goal-based, rule-based, experience-based, and user decision</i>	AIO CJ implements the application of updates on rules. Updates can encode <i>goal-based</i> decisions by sharing a property that is common to the same goal. It is also possible to implement <i>experience-based</i> and <i>user decision</i> systems by e.g., letting experience systems and users change the values of environmental variables.
Degree of automation	<i>automation, system-supported, and manual</i>	It is possible to either let the framework apply updates in an automatic way or require (different degrees of) interaction with the user.

Table 7.1: Attributes of flexibility as implemented in AIO CJ.

deadlock-freedom). The languages respectively implement the formal models presented in [13] and Chapter 4. However, both implementations depart significantly from their respective formal models as they generate code based on message correlation whilst their formalisations of EPP use synchronisations on names. This gap between theoretical models and their implementations has two consequences: *i)* it breaks the correctness-by-construction guarantee of choreographies — there is no proof that the implementation correctly supports synchronisation on names — and *ii)* users must look into the complexity of the compilers to understand how the generated code implements the originating choreography. Implementations of other frameworks based on sessions share similar issues and follow different custom practices to implement the semantics of name synchronisation [52, 53, 54]. Thus, our work is a useful reference to formalise the implementation of session-based languages in general.

Choreography Language and Deployment. The fragment of endpoint choreographies in AC is remarkably similar to standard process languages based on sessions, e.g., those used in [44, 116, 43, 96]. Moreover, the syntax of AC resembles that of Compositional Choreographies [14], which introduced a notion of compositionality in choreographies. This similarity is intentional: our aim is to show that it is possible to provide a suitable implementation model for this kind of languages.

Notably, our semantics easily captures asynchronous message passing through the interplay of the swapping relation \simeq_C , partial choreographies, and message queues. By contrast, previous works [13, 14] simulate asynchrony with encumbering ad-hoc rules that check if actions guarded by communications in continuations are allowed.

Several works have already analysed the behavioural expressiveness of choreographic descriptions as protocols (e.g., multiparty session types [108]) but, to the best of our knowledge, there are none on the expressiveness of choreographies meant as implementation languages. We leave this investigation as an interesting future work.

Integrating Dynamic Choreographies in Applied Choreographies. By integrating our theory of Dynamic Choreographies into that of Applied Choreographies we can proceed in the realisation of a formally proven programming framework for adaptable distributed systems. However, merging the two theoretical approaches constitutes a formidable challenge because we need to put together the typical static approach of typed sessions of ACs with dynamic distributed behaviours of Dynamic Choreographies. Some already-cited works in this direction are [43, 13, 78, 16, 111] whilst [117] discusses the main issues of integrating typed sessions that can update at runtime.

Delegation. Delegation is a standard feature of session-typed choreography models [13, 14] and enables to transfer the responsibility to continue a session

from a process to another. We can support delegation in AC by adding a rule that atomically updates the session descriptors of all processes involved in a session when there is a delegation. However, we leave delegation to future work, as its introduction would bring more complexity in our compiler; specifically, we would have to compile appropriate communications to update the local session descriptors of the involved DCC processes. This is a renowned issue of implementing delegation that we plan to address by formalising the techniques proposed in [52, 13].

Correlation keys. In the semantics of AC, we abstract how correlation keys are generated. With this loose definition we can capture several implementations, provided they satisfy the requirement of unicity of keys (wrt to locations). As future work, we plan to implement a language, based on AC, able to support custom procedures for the generation of correlation keys (e.g., from database queries, cookies, etc.).

Bibliography

- [1] B. M. Leiner, V. G. Cerf, D. D. Clark, R. E. Kahn, L. Kleinrock, D. C. Lynch, J. B. Postel, L. G. Roberts, and S. S. Wolff, “A brief history of the internet,” *CoRR*, vol. cs.NI/9901011, 1999. (Cited on pages 1 and 9.)
- [2] G. F. Coulouris and J. Dollimore, *Distributed Systems: Concepts and Design*. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc., 1988. (Cited on pages 1 and 9.)
- [3] E. G. Coffman, M. Elphick, and A. Shoshani, “System deadlocks,” *ACM Comput. Surv.*, vol. 3, pp. 67–78, June 1971. (Cited on pages 1 and 11.)
- [4] R. H. B. Netzer and B. P. Miller, “What are race conditions?: Some issues and formalizations,” *ACM Lett. Program. Lang. Syst.*, vol. 1, pp. 74–88, Mar. 1992. (Cited on pages 1 and 11.)
- [5] D. Sangiorgi and D. Walker, *The π -calculus: a Theory of Mobile Processes*. Cambridge University Press, 2001. (Cited on pages 2, 19, 20, 41, and 79.)
- [6] S. Lu, S. Park, E. Seo, and Y. Zhou, “Learning from mistakes: a comprehensive study on real world concurrency bug characteristics,” in *ASPLOS*, pp. 329–339, ACM, 2008. (Cited on page 2.)
- [7] C. Peltz, “Web services orchestration and choreography,” *Computer*, vol. 36, pp. 46–52, Oct 2003. (Cited on pages 2 and 14.)
- [8] W3C WS-CDL Working Group, “Web services choreography description language version 1.0.” <http://www.w3.org/TR/ws-cdl-10/>, 2004. (Cited on pages 2, 16, and 80.)

Bibliography

- [9] “Business Process Model and Notation.” <http://www.omg.org/spec/BPMN/2.0/>. (Cited on pages 2 and 16.)
- [10] Z. Qiu, X. Zhao, C. Cai, and H. Yang, “Towards the theoretical foundation of choreography,” in *WWW*, (United States), pp. 973–982, IEEE Computer Society Press, 2007. (Cited on pages 2, 19, 79, 81, and 96.)
- [11] M. Carbone, K. Honda, and N. Yoshida, “Structured communication-centred programming for web services,” in *Proc. of ESOP*, vol. 4421 of *LNCS*, pp. 2–17, Springer-Verlag, 2007. (Cited on pages 2, 19, 32, and 58.)
- [12] I. Lanese, C. Guidi, F. Montesi, and G. Zavattaro, “Bridging the gap between interaction-and process-oriented choreographies,” in *SEFM’08*, pp. 323–332, IEEE, 2008. (Cited on pages 2, 17, 19, 32, 79, and 81.)
- [13] M. Carbone and F. Montesi, “Deadlock-freedom-by-design: multiparty asynchronous global programming,” in *POPL*, pp. 263–274, 2013. (Cited on pages 2, 19, 20, 32, 79, 81, 83, 89, 96, 105, 117, 118, 119, 121, 122, and 203.)
- [14] F. Montesi and N. Yoshida, “Compositional choreographies,” in *CONCUR*, pp. 425–439, 2013. (Cited on pages 2, 19, 21, 81, 82, 83, 93, 94, 95, 102, 118, 119, 121, 189, 203, 205, and 216.)
- [15] F. Montesi, “Kickstarting choreographic programming,” *CoRR*, vol. abs/1502.02519, 2015. (Cited on page 2.)
- [16] M. Carbone, K. Honda, and N. Yoshida, “Structured Interactional Exceptions in Session Types,” in *Proc. of CONCUR’08*, vol. 5201 of *LNCS*, pp. 402–417, Springer, 2008. (Cited on pages 2, 118, and 121.)
- [17] G. Anderson and J. Rathke, “Dynamic software update for message passing programs,” in *APLAS*, vol. 7705 of *LNCS*, pp. 207–222, Springer, 2012. (Cited on pages 2 and 117.)
- [18] F. E. Heart, R. E. Kahn, S. M. Ornstein, W. R. Crowther, and D. C. Walden, “The interface message processor for the arpa computer network,” in *Proceedings of the May 5-7, 1970, Spring Joint Computer Conference, AFIPS ’70* (Spring), (New York, NY, USA), pp. 551–567, ACM, 1970. (Cited on page 9.)
- [19] R. M. Metcalfe and D. R. Boggs, “Ethernet: Distributed packet switching for local computer networks,” *Commun. ACM*, vol. 19, pp. 395–404, July 1976. (Cited on page 9.)

Bibliography

- [20] M. Weiser, “The computer for the 21st century,” *SIGMOBILE Mob. Comput. Commun. Rev.*, vol. 3, pp. 3–11, July 1999. (Cited on page 9.)
- [21] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, and M. Zaharia, “A view of cloud computing,” *Commun. ACM*, vol. 53, pp. 50–58, Apr. 2010. (Cited on page 9.)
- [22] ITU-T, Geneva, “120: Message sequence chart (MSC),” 1996. (Cited on page 11.)
- [23] G. Booch, *The unified modeling language user guide*. Pearson Education India, 2005. (Cited on page 11.)
- [24] T. Berners-Lee, R. Cailliau, A. Luotonen, H. F. Nielsen, and A. Secret, “The world-wide web,” *Commun. ACM*, vol. 37, pp. 76–82, Aug. 1994. (Cited on page 12.)
- [25] T. Berners-Lee, L. Masinter, M. McCahill, *et al.*, “Uniform resource locators (url),” 1994. (Cited on page 12.)
- [26] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, and T. Berners-Lee, “Hypertext transfer protocol–http/1.1,” 1999. (Cited on page 12.)
- [27] T. Bray, J. Paoli, C. M. Sperberg-McQueen, E. Maler, and F. Yergeau, “Extensible markup language (xml),” *World Wide Web Consortium Recommendation REC-xml-19980210*. <http://www.w3.org/TR/1998/REC-xml-19980210>, 1998. (Cited on pages 12 and 84.)
- [28] D. Booth, H. Haas, F. McCabe, E. Newcomer, M. Champion, C. Ferris, and D. Orchard, “Web services architecture, w3c working group note, 11 february 2004,” *World Wide Web Consortium, article available from: http://www.w3.org/TR/ws-arch*, 2004. (Cited on page 13.)
- [29] E. Christensen, F. Curbera, and G. Meredith, “Web services description language (WSDL) 1.1. w3c,” tech. rep., Note 15, 2001, www.w3.org/TR/wsdl, 2001. (Cited on page 13.)
- [30] M. Gudgin, M. Hadley, N. Mendelsohn, J.-J. Moreau, and H. F. Nielsen, “SOAP version 1.2 part 1: Messaging framework. w3c recommendation 24 june 2003,” *Retrieved from http://www.w*, vol. 3, 2003. (Cited on page 13.)
- [31] M. P. Papazoglou, “Service-oriented computing: Concepts, characteristics and directions,” in *WISE 2003*, pp. 3–12, IEEE, 2003. (Cited on page 13.)

Bibliography

- [32] Red Hat, “D-bus website.” <http://www.freedesktop.org/wiki/Software/dbus/>, 2014. (Cited on page 13.)
- [33] E. Frank and I. Redmond, “Dcom: Microsoft distributed component object model,” *Redmond III November*, 1997. (Cited on page 13.)
- [34] D. A. Chappell, *Enterprise service bus*. " O'Reilly Media, Inc.", 2004. (Cited on page 13.)
- [35] SAP, “Sap remote function call.” http://help.sap.com/saphelp_nw04/helpdata/en/6f/1bd5b6a85b11d6b28500508b5d5211/content.htm, 2014. (Cited on page 14.)
- [36] OASIS, “Web Services Business Process Execution Language.” <http://docs.oasis-open.org/wsbpel/2.0/wsbpel-v2.0.html>. (Cited on pages 15, 32, 79, 80, 84, and 106.)
- [37] D. Recordon and D. Reed, “Openid 2.0: a platform for user-centric identity management,” in *Proceedings of the second ACM workshop on Digital identity management*, pp. 11–16, ACM, 2006. (Cited on page 16.)
- [38] G. Decker, O. Kopp, F. Leymann, and M. Weske, “Bpel4chor: Extending bpel for modeling choreographies,” in *Web Services, 2007. ICWS 2007. IEEE International Conference on*, pp. 296–303, IEEE, 2007. (Cited on page 17.)
- [39] R. M. Needham and M. D. Schroeder, “Using encryption for authentication in large networks of computers,” *Commun. ACM*, vol. 21, pp. 993–999, Dec. 1978. (Cited on page 18.)
- [40] M. Bravetti and G. Zavattaro, “Towards a unifying theory for choreography conformance and contract compliance,” in *Proc. of SC'07*, vol. 4829, pp. 34–50, 2007. (Cited on pages 19 and 32.)
- [41] S. Carpineti and C. Laneve, “A basic contract language for web services.” in *ESOP'06*, LNCS, pp. 197–213, 2006. (Cited on page 19.)
- [42] F. Montesi, *Choreographic Programming*. Ph.D. thesis, IT University of Copenhagen, 2013. http://www.fabriziomontesi.com/files/choreographic_programming.pdf. (Cited on page 19.)
- [43] K. Honda, N. Yoshida, and M. Carbone, “Multiparty asynchronous session types,” in *Proc. of POPL*, vol. 43(1), pp. 273–284, ACM, 2008. (Cited on pages 19, 32, 79, 81, 82, 92, 94, 101, 117, 118, and 121.)

Bibliography

- [44] K. Honda, V. Vasconcelos, and M. Kubo, “Language primitives and type disciplines for structured communication-based programming,” in *ESOP’98*, vol. 1381 of *LNCS*, (Heidelberg, Germany), pp. 22–138, Springer-Verlag, 1998. (Cited on pages 19, 83, and 121.)
- [45] K. Honda, A. Mukhamedov, G. Brown, T.-C. Chen, and N. Yoshida, “Scribbling interactions with a formal foundation,” in *Proc. of ICDCIT*, vol. 6536 of *LNCS*, pp. 55–75, Springer, 2011. (Cited on page 19.)
- [46] PI4SOA, “<http://www.pi4soa.org>,” 2008. (Cited on page 19.)
- [47] Savara, “JBoss Community.” <http://www.jboss.org/savara/>. (Cited on page 19.)
- [48] F. Montesi, “Chor language website.” <http://www.chor-lang.org>, 2014. (Cited on page 20.)
- [49] R. Milner, *A Calculus of Communicating Systems*, vol. 92 of *Lecture Notes in Computer Science*. Springer, 1980. (Cited on pages 20 and 79.)
- [50] F. Montesi, C. Guidi, and G. Zavattaro, “Service-oriented programming with jolie,” in *Web Services Foundations*, pp. 81–107, Springer, 2014. (Cited on pages 20 and 79.)
- [51] M. A. Hiltunen and R. D. Schlichting, “Adaptive distributed and fault-tolerant systems,” *International Journal of Computer Systems Science and Engineering*, vol. 11, pp. 125–133, 1995. (Cited on page 20.)
- [52] R. Hu, N. Yoshida, and K. Honda, “Session-based distributed programming in java,” in *ECOOP*, pp. 516–541, 2008. (Cited on pages 21, 79, 121, and 122.)
- [53] R. Hu, R. Neykova, N. Yoshida, R. Demangeon, and K. Honda, “Practical interruptible conversations - distributed dynamic verification with session types and python,” in *RV*, pp. 130–148, 2013. (Cited on pages 21 and 121.)
- [54] R. Neykova and N. Yoshida, “Multiparty session actors,” in *COORDINATION 2014, Proceedings*, pp. 131–146, 2014. (Cited on pages 21 and 121.)
- [55] J. O. Kephart and D. M. Chess, “The vision of autonomic computing,” *Computer*, vol. 36, no. 1, pp. 41–50, 2003. (Cited on pages 25 and 65.)
- [56] M. Salehie and L. Tahvildari, “Self-adaptive software: Landscape and research challenges,” *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*, vol. 4, no. 2, p. 14, 2009. (Cited on page 25.)

Bibliography

- [57] G. Salvaneschi, C. Ghezzi, and M. Pradella, “An analysis of language-level support for self-adaptive software,” *ACM Trans. Auton. Adapt. Syst.*, vol. 8, pp. 7:1–7:29, July 2013. (Cited on pages 25 and 26.)
- [58] P. Oreizy, N. Medvidovic, and R. N. Taylor, “Architecture-based runtime software evolution,” in *Proceedings of the 20th international conference on Software engineering*, pp. 177–186, IEEE Computer Society, 1998. (Cited on page 26.)
- [59] J. Kramer and J. Magee, “Self-managed systems: an architectural challenge,” in *Future of Software Engineering, 2007. FOSE’07*, pp. 259–268, IEEE, 2007. (Cited on page 26.)
- [60] S. R. White, D. M. Chess, J. O. Kephart, J. E. Hanson, and I. Whalley, “An architectural approach to autonomic computing,” in *Autonomic Computing, International Conference on*, pp. 2–9, IEEE Computer Society, 2004. (Cited on page 26.)
- [61] C. Ghezzi, M. Pradella, and G. Salvaneschi, “An Evaluation of the Adaptation Capabilities in Programming Languages,” in *Proc. of SEAMS 2011*, pp. 50–59, ACM, 2011. (Cited on pages 26 and 32.)
- [62] P. Maes, “Concepts and experiments in computational reflection,” in *ACM Sigplan Notices*, vol. 22, pp. 147–155, ACM, 1987. (Cited on page 27.)
- [63] J. McCarthy, *LISP 1.5 programmer’s manual*. MIT press, 1965. (Cited on page 27.)
- [64] T. Ledoux, “Implementing proxy objects in a reflective orb,” in *Proc. ECOOP*, vol. 97, Citeseer, 1997. (Cited on page 27.)
- [65] J. Dowling, T. Schäfer, V. Cahill, P. Haraszti, and B. Redmond, “Using reflection to support dynamic adaptation of system software: A case study driven evaluation,” in *Reflection and Software Engineering*, pp. 169–188, Springer, 2000. (Cited on page 27.)
- [66] J. Xu, B. Randell, and A. F. Zorzo, “Implementing software-fault tolerance in c++ and open c++: An object-oriented and reflective approach,” *Proc. CADTED*, vol. 96, pp. 224–229, 1996. (Cited on pages 27 and 28.)
- [67] G. Kiczales, J. Lamping, A. Mendhekar, C. Maeda, C. Lopes, J.-M. Longtier, and J. Irwin, *Aspect-oriented programming*. Springer, 1997. (Cited on page 28.)

Bibliography

- [68] A. Popovici, T. Gross, and G. Alonso, “Dynamic weaving for aspect-oriented programming,” in *Proceedings of the 1st international conference on Aspect-oriented software development*, pp. 141–147, ACM, 2002. (Cited on page 28.)
- [69] R. Pawlak, L. Duchien, G. Florin, and L. Seinturier, “JAC: A flexible solution for aspect-oriented programming in java,” in *Metalevel architectures and separation of crosscutting concerns*, pp. 1–24, Springer, 2001. (Cited on page 28.)
- [70] R. Pawlak, L. Seinturier, L. Duchien, G. Florin, F. Legond-Aubry, and L. Martelli, “JAC: an aspect-based distributed dynamic framework,” *Software: Practice and Experience*, vol. 34, no. 12, pp. 1119–1148, 2004. (Cited on pages 28 and 118.)
- [71] R. Hirschfeld, P. Costanza, and O. Nierstrasz, “Context-oriented programming,” *Journal of Object Technology*, vol. 7, no. 3, 2008. (Cited on page 29.)
- [72] M. Appeltauer, R. Hirschfeld, M. Haupt, J. Lincke, and M. Perscheid, “A comparison of context-oriented programming languages,” in *International Workshop on Context-Oriented Programming*, p. 6, ACM, 2009. (Cited on page 29.)
- [73] C. Ghezzi, M. Pradella, and G. Salvaneschi, “Programming language support to context-aware adaptation: a case-study with erlang,” in *ICSE*, pp. 59–68, ACM, 2010. (Cited on pages 29 and 172.)
- [74] M. Weske, *Business process management: concepts, languages, architectures*. Springer Science & Business Media, 2012. (Cited on page 30.)
- [75] M. Dumas, W. M. Van der Aalst, and A. H. Ter Hofstede, *Process-aware information systems: bridging people and software through process technology*. John Wiley & Sons, 2005. (Cited on page 30.)
- [76] M. Reichert and B. Weber, *Enabling flexibility in process-aware information systems: challenges, methods, technologies*. Springer Science & Business Media, 2012. (Cited on pages 30 and 31.)
- [77] A. Charfi and M. Mezini, “Aspect-oriented web service composition with ao4bpel,” in *Web Services*, pp. 168–182, Springer, 2004. (Cited on page 32.)

Bibliography

- [78] G. Castagna, M. Dezani-Ciancaglini, and L. Padovani, “On global types and multi-party session,” *Logical Methods in Computer Science*, vol. 8, no. 1, 2012. (Cited on pages 32, 117, 118, and 121.)
- [79] J. A. Bergstra and P. Klint, “The discrete time TOOLBUS - A software co-ordination architecture,” *Sci. Comput. Program.*, vol. 31, no. 2-3, pp. 205–229, 1998. (Cited on page 32.)
- [80] A. Bucchiarone, A. Marconi, M. Pistore, and H. Raik, “Dynamic Adaptation of Fragment-Based and Context-Aware Business Processes,” in *Proc. of ICWS 2012*, pp. 33–41, 2012. (Cited on page 32.)
- [81] W.-K. Chen, M. A. Hiltunen, and R. D. Schlichting, “Constructing Adaptive Software in Distributed Systems,” in *Proc. of ICDCS’01*, vol. 6084, pp. 635–643, 2001. (Cited on page 32.)
- [82] I. Lanese, A. Bucchiarone, and F. Montesi, “A Framework for Rule-Based Dynamic Adaptation,” in *Proc. of TGC 2010*, vol. 6084, pp. 284–300, 2010. (Cited on pages 32 and 118.)
- [83] J. Zhang, H. Goldsby, and B. H. C. Cheng, “Modular Verification of Dynamically Adaptive Systems,” in *Proc. of AOSD’09*, pp. 161–172, ACM, 2009. (Cited on page 32.)
- [84] L. Leite, G. Ansal di Oliva, G. Nogueira, M. Gerosa, F. Kon, and D. Milojevic, “A systematic literature review of service choreography adaptation,” *Service Oriented Computing and Applications*, pp. 1–18, 2012. (Cited on page 32.)
- [85] S. González, K. Mens, M. Colacioiu, and W. Cazzola, “Context traits: dynamic behaviour adaptation through run-time trait recomposition,” in *AOSD*, pp. 209–220, 2013. (Cited on page 32.)
- [86] S. Goetz and I. Savga, “Exploring Role Based Adaptation,” in *RAM-SE*, pp. 21–26, Fakultät für Informatik, Universität Magdeburg, 2008. (Cited on page 32.)
- [87] 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*, pp. 85–92, ACM, 2002. (Cited on pages 32 and 118.)
- [88] F. Montesi, C. Guidi, and G. Zavattaro, “Service-oriented programming with jolie,” in *Web Services Foundations*, pp. 81–107, 2014. (Cited on page 42.)

Bibliography

- [89] “AIOCI website.” <http://www.cs.unibo.it/projects/jolie/aiocj.html>. (Cited on pages 58 and 118.)
- [90] M. Carbone and F. Montesi, “Deadlock-Freedom-by-Design: Multiparty Asynchronous Global Programming,” in *POPL*, pp. 263–274, ACM, 2013. (Cited on page 58.)
- [91] I. Lanese, F. Montesi, and G. Zavattaro, “Amending choreographies,” in *WWV*, vol. 123, pp. 34–48, EPTCS, 2013. (Cited on page 58.)
- [92] Jolie, “Programming Language.” <http://www.jolie-lang.org/>. (Cited on pages 66, 79, and 106.)
- [93] J. W. Backus, “The syntax and semantics of the proposed international algebraic language of the zurich ACM-GAMM conference,” in *IFIP Congress*, pp. 125–131, 1959. (Cited on page 67.)
- [94] I. Eclipse, “The eclipse foundation,” 2007. (Cited on page 71.)
- [95] “Xtext website.” <http://www.eclipse.org/Xtext/>. (Cited on page 71.)
- [96] M. Carbone, K. Honda, and N. Yoshida, “Structured communication-centered programming for web services,” *ACM Trans. Program. Lang. Syst.*, vol. 34, no. 2, p. 8, 2012. (Cited on pages 79, 81, 83, 96, 103, 105, 121, 189, and 216.)
- [97] M. Carbone, F. Montesi, and C. Schürmann, “Choreographies, logically,” in *CONCUR*, pp. 47–62, 2014. (Cited on pages 79 and 83.)
- [98] F. Montesi and M. Carbone, “Programming services with correlation sets,” in *ICSOC*, pp. 125–141, 2011. (Cited on pages 79, 81, 84, and 106.)
- [99] Chor, “Programming Language.” <http://www.chor-lang.org/>. (Cited on pages 79 and 119.)
- [100] S. Carpineti, C. Laneve, and P. Milazzo, “Bopi - A distributed machine for experimenting web services technologies,” in *ACSD 2005, 6-9 June 2005, St. Malo, France*, pp. 202–211, 2005. (Cited on page 79.)
- [101] R. S. Chin and S. T. Chanson, “Distributed object-based programming systems,” *ACM Comput. Surv.*, vol. 23, no. 1, pp. 91–124, 1991. (Cited on pages 80 and 119.)
- [102] S. Basu, T. Bultan, and M. Ouederni, “Deciding choreography realizability,” in *POPL*, pp. 191–202, 2012. (Cited on page 81.)

- [103] T. Bray, “The javascript object notation (json) data interchange format,” 2014. (Cited on page 84.)
- [104] L. Bettini, M. Coppo, L. D’Antoni, M. D. Luca, M. Dezani-Ciancaglini, and N. Yoshida, “Global progress in dynamically interleaved multiparty sessions,” in *CONCUR*, vol. 5201 of *LNCS*, pp. 418–433, Springer, 2008. (Cited on page 88.)
- [105] M. Coppo, M. Dezani-Ciancaglini, N. Yoshida, and L. Padovani, “Global progress for dynamically interleaved multiparty sessions,” *MSCS*, vol. 760, pp. 1–65, 2015. (Cited on page 92.)
- [106] B. C. Pierce, *Types and Programming Languages*. MA, USA: MIT Press, 2002. (Cited on page 93.)
- [107] N. Busi, R. Gorrieri, C. Guidi, R. Lucchi, and G. Zavattaro, “Choreography and orchestration conformance for system design,” in *COORDINATION’06*, vol. 4038 of *LNCS*, (Heidelberg, Germany), pp. 63–81, Springer-Verlag, 2006. (Cited on page 96.)
- [108] P. Deniérou and N. Yoshida, “Multiparty compatibility in communicating automata: Characterisation and synthesis of global session types,” in *ICALP 2013, Proceedings, Part II*, pp. 174–186, Springer, 2013. (Cited on pages 97 and 121.)
- [109] M. Coppo, M. Dezani-Ciancaglini, and B. Venneri, “Self-adaptive multiparty sessions,” *Service Oriented Computing and Applications*, pp. 1–20, 2014. (Cited on page 117.)
- [110] C. Di Giusto and J. A. Pérez, “Disciplined structured communications with consistent runtime adaptation,” in *SAC*, pp. 1913–1918, ACM, 2013. (Cited on pages 117 and 118.)
- [111] S. Capecchi, E. Giachino, and N. Yoshida, “Global Escape in Multiparty Sessions,” in *Proc. of FSTTCS 2010*, vol. 8 of *LIPICs*, pp. 338–351, Schloss Dagstuhl, 2010. (Cited on pages 118 and 121.)
- [112] I. Jureta, S. Faulkner, and P. Thiran, “Dynamic requirements specification for adaptable and open service-oriented systems,” in *ICSOC*, vol. 4749 of *LNCS*, pp. 270–282, Springer, 2007. (Cited on page 118.)
- [113] S. Rinderle, A. Wombacher, and M. Reichert, “Evolution of process choreographies in dychor,” in *OTM Conferences (1)*, vol. 4275 of *LNCS*, pp. 273–290, Springer, 2006. (Cited on page 118.)

Bibliography

- [114] A. Wombacher, “Alignment of choreography changes in BPEL processes,” in *IEEE SCC*, pp. 1–8, IEEE, 2009. (Cited on page 118.)
- [115] M. Bravetti *et al.*, “Towards global and local types for adaptation,” in *SEFM Workshops*, vol. 8368 of *LNCS*, pp. 3–14, Springer, 2013. (Cited on page 119.)
- [116] S. Gay and M. Hole, “Subtyping for session types in the pi calculus,” *Acta Informatica*, vol. 42, pp. 191–225, Nov. 2005. (Cited on pages 121 and 189.)
- [117] M. Bravetti, M. Carbone, T. T. Hildebrandt, I. Lanese, J. Mauro, J. A. Pérez, and G. Zavattaro, “Towards global and local types for adaptation,” in *Software Engineering and Formal Methods - SEFM 2013 Collocated Workshops: BEAT2, WS-FMDS, FM-RAIL-Bok, MoKMaSD, and OpenCert, Madrid, Spain, September 23-24, 2013, Revised Selected Papers*, pp. 3–14, Springer, 2013. (Cited on page 121.)
- [118] R. Pawlak *et al.*, “JAC: an aspect-based distributed dynamic framework,” *Softw., Pract. Exper.*, vol. 34, no. 12, pp. 1119–1148, 2004. (Cited on page 167.)

Bibliography

Part V
Appendix

Adaptable Choreographies: Proofs

A.1 Proof of Theorem 1

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 3. *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 1 (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 $\text{transF}(\mathcal{I}')$ requires the creation of balanced trees having an element for every role of \mathcal{I}' . Since the roles are $O(n)$, $\text{transF}(\mathcal{I}')$ 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 $\forall R_1 \rightarrow R_2 \in \text{transF}(\mathcal{I}'), \forall S_1 \rightarrow S_2 \in \text{transl}(\mathcal{I}'') . \{R_1, R_2\} \cap \{S_1, S_2\} \neq \emptyset$. Since $\text{transF}(\mathcal{I}')$ and $\text{transl}(\mathcal{I}'')$ have $O(n)$ elements, thanks to Lemma 3, 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

A.2 Proof of Theorem 2

Before entering into the details of the proof, we recall the proof strategy that we described in § 4.7, which consists in defining a notion of bisimulation which implies weak trace equivalence. We report below the definition of bisimulation.

Definition 12 (Weak System Bisimulation). 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 $\eta \in \{o^? : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x); o^* : \mathbf{R}_1(X) \rightarrow \mathbf{R}_2(); \sqrt{}; \mathcal{I}; \text{no-up}; \mathbf{I}''', \tau\}$ 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 $(\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle, \langle \mathbf{I}''', \mathcal{N}''' \rangle) \in \mathcal{R}$

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 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 `upd` (Definition 31). 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 25) and one for DPOC (Definition 26) and we show that the two are related (Lemma 5).

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.

In our proof strategy we rely on the uniqueness of indexes that unfortunately it is not preserved by transitions due to while unfolding. Indeed, as an example, consider the DIOC $i: \text{while } b@R \{j: x@R = e\}$. If the condition b evaluates to true, in one step the application of Rule $\lfloor^{\text{DPOC}} \rfloor_{\text{WHILE-UNFOLD}}$ produces the DIOC $j: x@R = e; i: \text{while } b@R \{j: x@R = e\}$ where the index j is used 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 uniqueness is preserved by transitions.

Definition 24 (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 then $\xi = \iota$;*
- *if the innermost while construct that contains the considered construct has global index ξ' then the considered construct has global index $\xi = \xi' : \iota$.*

Lemma 4 (Distinctness of Global Indexes). *Given a well-annotated DIOC \mathcal{I} , a global state Σ , and a set of updates \mathbf{I} , if $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle \xrightarrow{\eta_1} \dots \xrightarrow{\eta_m} \langle \Sigma', \mathbf{I}', \mathcal{I}' \rangle$ then all global indexes 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 $\mathcal{I}; \mathbf{i} : \text{while } b @ \mathbf{R} \{ \mathcal{I}' \}$. In this last case, the same index can occur both in \mathcal{I} and in \mathcal{I}' , in constructs with different global indexes. The statement of the Lemma follows directly since distinct indexes imply distinct global indexes.

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 $\llbracket^{\text{DIOC}} \mid_{\text{WHILE-UNFOLD}} \rrbracket$ and $\llbracket^{\text{DIOC}} \mid_{\text{UP}} \rrbracket$.

In the case of Rule $\llbracket^{\text{DIOC}} \mid_{\text{WHILE-UNFOLD}} \rrbracket$, note that the while is enabled, hence it cannot be part of a term of the form $\mathcal{I}; \mathbf{i} : \text{while } b @ \mathbf{R} \{ \mathcal{I}' \}$. 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 $\mathcal{I}; \mathbf{i} : \text{while } b @ \mathbf{R} \{ \mathcal{I}' \}$, thus global indexes are distinct by construction.

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

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 25 (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 $\xi : \bar{o} @ \mathbf{R}_2$ in role \mathbf{R}_1 and a receiving event $\xi : o @ \mathbf{R}_1$ in role \mathbf{R}_2 for each interaction $\mathbf{i} : o : \mathbf{R}_1(e) \rightarrow \mathbf{R}_2(x)$ with global*

A.2. Proof of Theorem 2

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 $i: \text{scope } @ R \{ \mathcal{I} \}$ with global index ξ . Both these events belong to all the roles in $\text{roles}(\mathcal{I})$. 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 $i: \text{if } b @ R \{ \mathcal{I} \} \text{ else } \{ \mathcal{I}' \}$ 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 $i: \text{while } b @ R \{ \mathcal{I} \}$ with global index ξ ; the guard-while event corresponds to the transition evaluating the guard of the while loop.

Function events(\mathcal{I}) denotes the set of events of the annotated DIOC \mathcal{I} . 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 4 events have distinct names. Note also that, for each while loop, there are events corresponding to the execution of just one loop. If unfolding is performed, new events are created.

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

Definition 26 (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 $\mathbf{i} : \text{scope } @R \{P\}$ roles $\{S\}$ or $\mathbf{i} : \text{scope } @R \{P\}$ 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 $\mathbf{i} : \text{if } b \{P\} \text{ else } \{P'\}$ 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 R for each construct $\mathbf{i} : \text{while } b \{P\}$ with global index ξ ; the guard-while event corresponds to the transition evaluating the guard of the while loop.

Let $\text{events}(\mathcal{N})$ 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 $\text{events}(P)$ 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 5).

The relation below defines a causality relation DIOC among events based on the constraints given by the semantics on the execution of the corresponding transitions.

Definition 27 (DIOC causality relation). *Let us consider an annotated DIOC \mathcal{I} . A causality relation $\leq_{DIOC} \subseteq \text{events}(\mathcal{I}) \times \text{events}(\mathcal{I})$ is a partial order among events in \mathcal{I} . We define \leq_{DIOC} as the minimum partial order 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 $\mathbf{i} : \text{scope } @R \{\mathcal{I}'\}$ 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 $\mathbf{i} : \text{if } b @R \{\mathcal{I}'\} \text{ else } \{\mathcal{I}''\}$ be a subterm of DIOC \mathcal{I} , let ε_ξ be the guard if-event in role 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 $\mathbf{i} : \text{while } b @R \{\mathcal{I}'\}$ be a subterm of DIOC \mathcal{I} , let ε_ξ be the guard while-event in role R , then for every event ε in \mathcal{I}' we have $\varepsilon_\xi \leq_{DIOC} \varepsilon$.*

Similarly for DPOC:

Definition 28 (DPOC causality relation). *Let us consider an annotated DPOC network \mathcal{N} . A causality relation $\leq_{DPOC} \subseteq \text{events}(\mathcal{N}) \times \text{events}(\mathcal{N})$ is a partial order among events in \mathcal{N} . We define \leq_{DPOC} as the minimum partial order 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 $\mathbf{i}: \text{scope } @R \{P\}$ roles $\{S\}$ or $\mathbf{i}: \text{scope } @R \{P\}$ 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 $\mathbf{i}: \text{if } b \{P\}$ 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 $\mathbf{i}: \text{while } b \{P\}$ 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 $\varepsilon_{\xi} \leq_{DPOC} \varepsilon$.*

Thanks to the definition above, we have that the events and causality relation are preserved by projection.

Lemma 5. *Given a 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. 1. By definition of projection.

2. Let $\varepsilon_1 \leq_{DIOC} \varepsilon_2$. 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 $\varepsilon'' \leq_{DPOC} \varepsilon_2$ or $\varepsilon'' \leq_{DPOC} \bar{\varepsilon}_2$. 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 $\varepsilon_3 \leq_{DPOC} \varepsilon_2$ or $\varepsilon_3 \leq_{DPOC} \bar{\varepsilon}_2$. By synchronisation (Definition 28) we have that $\bar{\varepsilon}_3 \leq_{DPOC} \varepsilon_2$ or $\bar{\varepsilon}_3 \leq_{DPOC} \bar{\varepsilon}_2$. 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 28) 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.

□

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 29 (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 $i: \text{if } b @ R \{ \mathcal{I}' \} \text{ else } \{ \mathcal{I}'' \}$ of \mathcal{I} such that $\varepsilon \in \text{events}(\mathcal{I}') \wedge \varepsilon' \in \text{events}(\mathcal{I}'')$ or $\varepsilon' \in \text{events}(\mathcal{I}') \wedge \varepsilon \in \text{events}(\mathcal{I}'')$.*

A.2. Proof of Theorem 2

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 $i: \text{if } b \{P\} \text{ else } \{P'\}$ 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 a well-annotated DIOC and executing the resulting network. They enjoy various properties useful for our proofs.

Definition 30 (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 public 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 $[f_\xi]_{R_1}$ and $[f_{\xi'}]_{R_1}$ on the same operation $i.o^?$ with the same target R_2 such that $\xi \neq \xi'$ we have $[f_\xi]_{R_1} \leq_{DPOC} [f_{\xi'}]_{R_1}$ or $[f_{\xi'}]_{R_1} \leq_{DPOC} [f_\xi]_{R_1}$;*
- C4 *for each pair of non-conflicting receiving events $[t_\xi]_{R_2}$ and $[t_{\xi'}]_{R_2}$ on the same operation $i.o^?$ with the same sender R_1 such that $\xi \neq \xi'$ we have $[t_\xi]_{R_2} \leq [t_{\xi'}]_{R_2}$ or $[t_{\xi'}]_{R_2} \leq [t_\xi]_{R_2}$;*
- 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 31 (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, $\text{upd}(\mathcal{N}) = \text{clean}(\text{compl}(\mathcal{N}))$. Network $\text{compl}(\mathcal{N})$ 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 $\mathbf{i.wb}_i^* : \text{true to } R'$ enabled, all the terms*

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

not inside another while construct, with

$$P; \mathbf{i.we}_i^* : \text{ok to } R; \mathbf{i.wb}_i^* : x_i \text{ from } R; \text{ while } x_i \left\{ \begin{array}{l} P; \mathbf{i.we}_i^* : \text{ok to } R; \\ \mathbf{i.wb}_i^* : x_i \text{ from } R \end{array} \right\}$$

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

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

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

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

3. *Performing the unfolding of a while loop by replacing every*

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

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

$$P; \mathbf{i.we}_i^* : \text{ok to } R; \mathbf{i.wb}_i^* : x_i \text{ from } R; \text{ while } x_i \left\{ \begin{array}{l} P; \mathbf{i.we}_i^* : \text{ok to } R; \\ \mathbf{i.we}_i^* : x_i \text{ from } R \end{array} \right\}$$

4. *Performing the termination of while loop by replacing every*

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

*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 $\mathbf{i.cnd}_i^* : \text{true to } R'$ enabled, all the terms*

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

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

A.2. Proof of Theorem 2

6. Performing the reception of the negative evaluation of the guard of a conditional by replacing, for every $i.cnd_i^* : \text{false}$ to R' enabled, all the terms

$$i.cnd_i^* : x_i \text{ from } R; \text{if } x_i \{P'\} \text{ else } \{P''\}$$

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

7. Performing the selection of the “then” branch by replacing every

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

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

8. Performing the selection of the “else” branch by replacing every

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

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

9. Performing the communication of the updated code by replacing, for every $i.sb_i^* : P$ to S enabled, all the terms

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

in role S not inside a while construct with P , and replace $i.sb_i^* : P$ to S with **1**.

10. Performing the communication that no update is needed by replacing, for each

$i.sb_i^* : \text{no}$ to S enabled, all the terms

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

in role S not inside a while construct with P' , and replace $i.sb_i^* : P$ to S with **1**.

Network $\text{clean}(\mathcal{N})$ 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

$$\begin{array}{ll} i.se_i^* : \text{ok to } R & i_c: i.we_i^* : \text{ok to } R \\ i.se_i^* : _ \text{ from } R & i_c: i.we_i^* : _ \text{ from } R \end{array}$$

not inside a while construct with **1**.

Furthermore clean may apply 0 or more times the following operation:

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

Note that function `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 6. *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 28 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 $i: \text{scope } @R \{P\} \text{ roles } \{S\}$ or $i: \text{scope } @R \{P\}$ 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 Rules $[^{DPOC}_{UP}]$ or $[^{DPOC}_{NOUP}]$ 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 $i: \text{if } x_i \{P'\} \text{ else } \{P''\}$ 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 $i: \text{while } x_i \{P\}$ and $\varepsilon \in \text{events}(P)$. Event ε cannot be enabled because of the semantics of the while loop.

□

A.2. Proof of Theorem 2

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

Lemma 7. *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 30 of well-annotated DPOC:

C1 For each global index ξ there are at most two communication events on public 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 6.

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 28);
- $P|P'$: this case can never happen for the reasons below. For events on public operations this follow by the definition of projection, since the prefixes of the names of operations are different. For events on private operations originated by the same construct this follows since all the targets are different. For events on private operations originated by different constructs this follows since the prefixes of the names of the operations are different.
- $\text{if } b \{P\} \text{ else } \{P'\}$: this case can never happen since the events are non-conflicting (Definition 29).

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 and different global indexes is for the auxiliary communications in the projection of the while construct, where the conditions hold by construction. \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 8. *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 $\eta \in \{o^? : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x), o^* : \mathbf{R}_1(X) \rightarrow \mathbf{R}_2(), \mathbf{I}', \sqrt{\cdot}, \mathcal{I}, \text{no-up}, \tau\}$ then there exists \mathcal{N}'' s.t. $\langle \mathbf{I}, \mathcal{N} \rangle \xrightarrow{\eta_1} \dots \xrightarrow{\eta_k} \xrightarrow{\eta} \langle \mathbf{I}', \mathcal{N}'' \rangle$ 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 $\eta \in \left\{ \begin{array}{l} o^? : \mathbf{R}_1(v) \rightarrow \mathbf{R}_2(x), o^* : \mathbf{R}_1(X) \rightarrow \mathbf{R}_2(), \\ \mathbf{I}', \sqrt{\cdot}, \mathcal{I}, \text{no-up}, \tau \end{array} \right\}$, 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 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 not 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

A.2. Proof of Theorem 2

in $\text{upd}(\mathcal{N})$ 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. □

We now prove a few properties of transitions with label \surd .

Lemma 9. *For each DIOC system $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle$ that reduces with a transition labelled \surd then, for each role $R \in \text{roles}(\mathcal{I})$, the DPOC role $(\pi(\mathcal{I}, R), \Sigma_R)_R$ 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. □

The next lemma shows that if two matching events are enabled in the projection of a DIOC, then the corresponding interaction is enabled in the DIOC.

Lemma 10. *Let \mathcal{I} be a DIOC obtained from a well-annotated connected DIOC via 0 or more transitions and $\mathbf{i}: o : R_1(e) \rightarrow R_2(x)$ be an interaction in \mathcal{I} . If $\mathbf{i}: \mathbf{i}.o : e$ to R_2 and $\mathbf{i}: \mathbf{i}.o : x$ from R_1 are matching events and are both enabled in $\text{proj}(\mathcal{I}, \Sigma)$ then $\mathbf{i}: o : R_1(e) \rightarrow R_2(x)$ is enabled.*

Proof. Note that well-annotatedness is preserved under reduction, hence \mathcal{I} is well-annotated.

If \mathcal{I} is connected for sequence, then the proof is by structural induction on \mathcal{I} . The cases for $\mathbf{1}$, $\mathbf{0}$, scopes, conditionals, and while loops are trivial. For parallel composition consider that, since the two events are matching, then they have the same global index. As a consequence, they are from the same component, and the thesis follows by inductive hypothesis. Let us consider sequential composition. Suppose $\mathcal{I} = \mathcal{I}'; \mathcal{I}''$. If $\mathbf{i}: o : R_1(e) \rightarrow R_2(x) \in \mathcal{I}'$ then the thesis follows by inductive hypothesis. Otherwise, by inductive hypothesis $\mathbf{i}: o : R_1(e) \rightarrow R_2(x)$ is enabled in \mathcal{I}'' . Thus, $R_1 \rightarrow R_2 \in \text{transl}(\mathcal{I}'')$. From connectedness, for each $S_1 \rightarrow S_2 \in \text{transF}(\mathcal{I}')$ we have $\{R_1, R_2\} \cap \{S_1, S_2\} \neq \emptyset$. This is not possible since otherwise at least one of the events $\mathbf{i}: \mathbf{i}.o : e$ to R_2 and $\mathbf{i}: \mathbf{i}.o : x$ from R_1 would not be enabled. Thus, the only possibility is $\text{transF}(\mathcal{I}') = \emptyset$. This implies that \mathcal{I}' has a transition with label \surd . Thus, $\mathbf{i}: o : R_1(e) \rightarrow R_2(x)$ is enabled in \mathcal{I} .

Let us now consider the case in which \mathcal{I} is not connected for sequence. This can happen only if $\mathcal{I} = \mathcal{I}'; \mathcal{I}''$ and \mathcal{I}' contained a scope which has been updated. Assume the scope has coordinator R and involves roles $\{R_i\}_{i \in I}$. The only tricky case is when $\mathbf{i}: o : R_1(e) \rightarrow R_2(x)$ is an initial interaction in \mathcal{I}'' . Since the term

was connected for sequence before the update was performed, we know that for each $R_i, i \in I$, $\{R_i, R\} \cap \{R_1, R_2\} \neq \emptyset$. We have two cases: either $R_1 = R \vee R_2 = R$ or $\{R_i\}_{i \in I} \subseteq \{R_1, R_2\}$.

In the first case, let us assume that $R_1 = R$ (the other case is analogous). Since inside the process of R_1 the send primitive $i : i.o : e$ to R_2 is enabled, all the auxiliary communications closing the scope have been performed, and so also the execution of the body of the scope has been completed.

The second case is analogous and follows from the fact that since the send primitive for R_1 and the receive primitive for R_2 are enabled, then it is not possible that either R_1 or R_2 (or both) is sending an auxiliary communication to R for closing the scope. \square

The following result states that weak system bisimilarity implies weak trace equivalence.

Lemma 11. *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

We can now prove our main theorem (Theorem 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 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 trace equivalent.*

Proof. We prove that the relation \mathcal{R} below is a weak system bisimulation. This implies weak trace equivalence by Lemma 11.

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

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

To ensure that proving that the relation above is a 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 $\text{upd}(\text{proj}(\mathcal{I}, \Sigma)) = \text{proj}(\mathcal{I}, \Sigma)$. From Lemma 7 $\text{proj}(\mathcal{I}, \Sigma)$ is well-annotated, thus $\text{upd}(\text{proj}(\mathcal{I}, \Sigma))$ is well-annotated. Observe that compl is the identity on $\text{proj}(\mathcal{I}, \Sigma)$, thus from Lemma 5 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 $\eta \in \{o : R_1(v) \rightarrow R_2(x); \sqrt{}; \mathcal{I}; \text{no-up}; \mathbf{I}''; \tau\}$ 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 8 case 1 there exists \mathcal{N}'' s.t. $\langle \mathbf{I}, \mathcal{N}_g \rangle \xrightarrow{\eta_1} \dots \xrightarrow{\eta_k} \xrightarrow{\mu} \langle \mathbf{I}'', \mathcal{N}'' \rangle$ where $\eta_i \in \{o^* : R_1(v) \rightarrow R_2(x), o^* : R_1(X) \rightarrow R_2(), \tau\}$ 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 $\eta \in \{o^? : R_1(v) \rightarrow R_2(x), o^* : R_1(X) \rightarrow R_2(), \sqrt{}; \mathcal{I}; \text{no-up}; \mathbf{I}'', \tau\}$ then thanks to Lemma 8 we have one of the following: (2a) $\text{upd}(\mathcal{N}_g) = \text{upd}(\mathcal{N}''')$ and $\eta \in \{o^* : R_1(v) \rightarrow R_2(x), o^* : R_1(X) \rightarrow R_2(), \tau\}$, 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 $\langle \Sigma, \mathbf{I}, \mathcal{I} \rangle \xrightarrow{\eta} \langle \Sigma'', \mathbf{I}'', \mathcal{I}'' \rangle$

$\langle \Sigma'', \mathbf{I}'', \mathcal{I}'' \rangle$ 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 bisimilarity. 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 9. The case for labels \mathbf{I} is trivial. Let us consider the other labels $o : R_1(v) \rightarrow R_2(x), \mathcal{I}$, 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 R changes the local state of the role in the DPOC in a corresponding way.

Case i: $o : R_1(e) \rightarrow R_2(x)$ trivial. Just note that at the DPOC level the interaction gives rise to one send and one receive on the same prefix. Synchronisation between send and receive is performed by Rule $[\text{DPOC}]_{\text{SYNCH}}$ that also removes the prefix from the label.

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

$$\mathcal{N} = \parallel_{R \in \text{roles}(\mathcal{I}; \mathcal{I}')} (\pi(\mathcal{I}, R); \pi(\mathcal{I}', R), \Sigma_R)_R.$$

\rightarrow Assume that $\mathcal{I}; \mathcal{I}' \xrightarrow{\mu} \mathcal{I}''$ with $\mu \in \{o : R_1(v) \rightarrow R_2(x); \mathcal{I}; \text{no-up}, \tau\}$. 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 \surd and $\mathcal{I}' \xrightarrow{\mu} \mathcal{I}''$.

In case (i) by inductive hypothesis

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

Thus

$$\parallel_{R \in \text{roles}(\mathcal{I})} (\pi(\mathcal{I}, R); \pi(\mathcal{I}', R), \Sigma_R)_R \xrightarrow{\mu} \mathcal{N} \quad \text{and}$$

A.2. Proof of Theorem 2

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

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 \surd 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 9, $\text{proj}(\mathcal{I}; \mathcal{I}', \Sigma) \xrightarrow{\mu} \mathcal{N}$ 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.

← 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

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

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 the

$\text{upd}(\parallel_{\mathbf{R} \in \text{roles}(\mathcal{I}; \mathcal{I}')} (P''_{\mathbf{R}}; \pi(\mathcal{I}', \mathbf{R}), \Sigma_{\mathbf{R}})_{\mathbf{R}}) = \text{proj}(\mathcal{I}''; \mathcal{I}', \Sigma)$. 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 $o : R_1(v) \rightarrow R_2(x)$ is enabled. Thus, \mathcal{I} has a transition with label \surd and $\mathcal{I}' \xrightarrow{o:R_1(v) \rightarrow R_2(x)} \mathcal{I}''$. Thanks to Lemma 9 then both $(\pi(\mathcal{I}, R_1), \Sigma_{R_1})_{R_1}$ and $(\pi(\mathcal{I}, R_2), \Sigma_{R_2})_{R_2}$ have a transition with label \surd . Thus, we have

$$\begin{aligned} & (\pi(\mathcal{I}', R_1), \Sigma_{R_1})_{R_1} \xrightarrow{\bar{i}.o\langle v \rangle @_{R_2:R_1}} (P_{R_1}, \Sigma_{R_1})_{R_1} \\ & (\pi(\mathcal{I}', R_2), \Sigma_{R_2})_{R_2} \xrightarrow{i.o\langle x \leftarrow v \rangle @_{R_1:R_2}} (P_{R_2}, \Sigma_{R_2})_{R_2} \text{ and thus} \\ & \text{proj}(\mathcal{I}', \Sigma) \xrightarrow{o:R_1(v) \rightarrow R_2(x)} \parallel_{R \in \text{roles}(\mathcal{I}')} (P_R, \Sigma_R)_R \end{aligned}$$

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

$$\mathcal{N} = \parallel_{R \in \text{roles}(\mathcal{I}, \mathcal{I}')} (\pi(\mathcal{I}, R) \mid \pi(\mathcal{I}', R), \Sigma_R)_R$$

→ We have a case analysis on the rule used to derive the transition. If the transition is derived using Rule $\lfloor \text{DIOC} \mid_{\text{PARALLEL}} \rfloor$ 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 $\lfloor \text{DIOC} \mid_{\text{PAR-END}} \rfloor$ then the thesis follows from Lemma 9.

← We have a case analysis on the label η of the transition. If $\eta = o^? : R_1(v) \rightarrow R_2(x)$ 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 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: $\text{if } b @ \mathbb{R} \{ \mathcal{I} \} \text{ else } \{ \mathcal{I}' \}$ from the definition of projection

$$\begin{aligned} \mathcal{N} = & \left(\parallel_{s \in \text{roles}(\mathcal{I}, \mathcal{I}') \setminus \{ \mathbb{R} \}} \left(\begin{array}{l} \mathbf{i}_? : \mathbf{i}. \mathbf{cnd}_i^* : x_i \text{ from } \mathbb{R}; \\ \mathbf{i} : \text{if } x_i \{ \pi(\mathcal{I}, \mathbb{S}) \} \text{ else } \{ \pi(\mathcal{I}', \mathbb{S}) \}, \Sigma_{\mathbb{S}} \end{array} \right)_s \right) \parallel \\ & \left(\mathbf{i} : \text{if } b \left\{ \left(\prod_{\mathbb{R}' \in \text{roles}(\mathcal{I}, \mathcal{I}') \setminus \{ \mathbb{R} \}} \mathbf{i}_{\mathbb{T}} : \mathbf{i}. \mathbf{cnd}_i^* : \text{true to } \mathbb{R}' \right); \pi(\mathcal{I}, \mathbb{R}) \right\} \right. \\ & \left. \text{else } \left\{ \left(\prod_{\mathbb{R}' \in \text{roles}(\mathcal{I}, \mathcal{I}') \setminus \{ \mathbb{R} \}} \mathbf{i}_{\mathbb{F}} : \mathbf{i}. \mathbf{cnd}_i^* : \text{false to } \mathbb{R}' \right); \pi(\mathcal{I}', \mathbb{R}) \right\}, \Sigma_{\mathbb{R}} \right)_{\mathbb{R}} \end{aligned}$$

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

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

$$\begin{aligned} \mathcal{N}' = & \left(\parallel_{s \in \text{roles}(\mathcal{I}, \mathcal{I}') \setminus \{ \mathbb{R} \}} \left(\begin{array}{l} \mathbf{i}_? : \mathbf{i}. \mathbf{cnd}_i^* : x_i \text{ from } \mathbb{R}; \\ \mathbf{i} : \text{if } x_i \{ \pi(\mathcal{I}, \mathbb{S}) \} \text{ else } \{ \pi(\mathcal{I}', \mathbb{S}) \}, \Sigma_{\mathbb{S}} \end{array} \right)_s \right) \parallel \\ & \left(\left(\prod_{\mathbb{R}' \in \text{roles}(\mathcal{I}, \mathcal{I}') \setminus \{ \mathbb{R} \}} \mathbf{i}_{\mathbb{T}} : \mathbf{i}. \mathbf{cnd}_i^* : \text{true to } \mathbb{R}' \right); \pi(\mathcal{I}, \mathbb{R}), \Sigma_{\mathbb{R}} \right)_{\mathbb{R}} \end{aligned}$$

By applying function upd we get

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

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: $\text{while } b @ R \{ \mathcal{I} \}$ from the definition of projection

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

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

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

By applying function `upd` we get

$$\text{upd}(\mathcal{N}') = \left(\parallel_{s \in \text{roles}(\mathcal{I}) \setminus \{R\}} \left(\begin{array}{l} \pi(\mathcal{I}, S); \mathbf{i}_?: \mathbf{i}.wb_i^* : x_i \text{ from } R; \\ \mathbf{i}: \text{while } x_i \{ \pi(\mathcal{I}, S); \\ \mathbf{i}_c: \mathbf{i}.we_i^* : \text{ok to } R; \\ \mathbf{i}_?: \mathbf{i}.wb_i^* : x_i \text{ from } R \} \end{array} , \Sigma_S \right) \parallel \right)_S$$

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

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

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

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

$$(\mathbf{i}: \text{scope } @R \{ \pi(\mathcal{I}, R) \} \text{ roles } \{ \text{roles}(\mathcal{I}) \}, \Sigma_R)_R$$

→ 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 $[\text{DIOC} \mid_{\text{UP}}]$. Correspondingly, at the DPOC level one applies Rule $[\text{DPOC} \mid_{\text{LEAD-UP}}]$

to the coordinator of the update, obtaining

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

By applying the upd function we get:

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

This is exactly the projection of the DIOC obtained after applying the rule $\llbracket^{\text{DIOC}} \mid_{\text{UP}} \rrbracket$. 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 $\llbracket^{\text{DIOC}} \mid_{\text{NoUP}} \rrbracket$ 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.

□

Adaptable Choreographies: Test Code

B.1 Code used for validation

B.1.1 Pipe and fork-join code

In this section, we provide the code of the AIOCJ programs used in Section 5.4. For simplicity, we show the programs fixing the number of tasks $n = 5$, and showing just two adaptation rules per scenario.

Pipe scenario. In the pipe scenario each task computes the increment function, passing the output of its computation as input to the next task.

```
1 include startTimer, endTimer from "socket://localhost:8000"
2 preamble{ starter: a }
3
4 aioc {
5   { x@a = 0 | x@b = 0 };
6   _r@a = startTimer( n );
7   x@a = 1 + x; pass: a( x ) -> b( x );
8   x@b = 1 + x; pass: b( x ) -> a( x );
9   x@a = 1 + x; pass: a( x ) -> b( x );
10  x@b = 1 + x; pass: b( x ) -> a( x );
11  x@a = 1 + x; pass: a( x ) -> b( x );
12  _r@a = stopTimer( n )
13 }
```

Listing B.1: Pipe scenario without scopes.

Listing B.1 shows the 5 tasks that increment variable x . After each increment the variable is sent from role a to role b or vice versa. Lines 6 and 12 are calls to the Timer service (Line 1), used to log the times of execution.

Appendix B. Adaptable Choreographies: Test Code

```
1 include startTimer, endTimer from "socket://localhost:8000"
2 preamble{ starter: a }
3
4 aioc {
5   { x@a = 0 | x@b = 0 };
6   _r@a = startTimer( n );
7   scope @a {x@a = 1 + x; pass: a( x ) -> b( x )} prop {N.x = 1};
8   scope @b {x@b = 1 + x; pass: b( x ) -> a( x )} prop {N.x = 2};
9   scope @a {x@a = 1 + x; pass: a( x ) -> b( x )} prop {N.x = 3};
10  scope @b {x@b = 1 + x; pass: b( x ) -> a( x )} prop {N.x = 4};
11  scope @a {x@a = 1 + x; pass: a( x ) -> b( x )} prop {N.x = 5};
12  _r@a = stopTimer( n )
13 }
```

Listing B.2: Pipe scenario with scopes.

Listing B.2 shows the 5 scopes enclosing the tasks of Listing B.1. Notably, each scope has a unique property `N.x` assigned incrementally. This property is used to identify uniquely each scope.

Rules of the pipe scenario. Listing B.3 shows two rules used in the pipe scenario. The adapted behaviour increments the variable `x` by 2. In particular, the rule at Lines 1-4 applies to the scope at Line 7 in Listing B.2, while the subsequent rule (Lines 5-8) applies to the scope at Line 8.

```
1 rule {
2   on { N.x == 1 }
3   do { x@a = 2 + x; pass: a( x ) -> b( x ) }
4 }
5 rule {
6   on { N.x == 2 }
7   do { x@b = 2 + x; pass: b( x ) -> a( x ) }
8 }
```

Listing B.3: Rules for pipe scenario.

Fork-join scenario. In the fork-join scenario tasks run in parallel. Each task takes one character of a message and replaces it with the next character in the alphabet. The program makes use of an external service to retrieve the character of the message, get the next letter, and store the new character in the message.

B.1. Code used for validation

```
1 include start, end from "socket://localhost:8000"
2 include getNthChar, getNext, setNthChar from "socket://localhost:8001"
3
4 preamble{ starter: a }
5
6 aioc {
7   _r@a = start( n );
8   {
9     { 10@a = getNthChar( 0 ); 10@a = getNext( 10 ); _r@a = setNthChar( 0, 10 ) }
10    | { 11@b = getNthChar( 1 ); 11@b = getNext( 11 ); _r@b = setNthChar( 1, 11 ) }
11    | { 12@a = getNthChar( 2 ); 12@a = getNext( 12 ); _r@a = setNthChar( 2, 12 ) }
12    | { 13@b = getNthChar( 3 ); 13@b = getNext( 13 ); _r@b = setNthChar( 3, 13 ) }
13    | { 14@a = getNthChar( 4 ); 14@a = getNext( 14 ); _r@a = setNthChar( 4, 14 ) }
14    };
15   _r@a = end( n )
16 }
```

Listing B.4: Fork-join scenario without scopes.

Listing B.4 shows the 5 tasks of the fork-join scenario. Each task uses a dedicated variable (10, 11, ..., 1n) to ensure thread-safety of the computation.

```
1 include start, end from "socket://localhost:8000"
2 include getNthChar, getNext, setNthChar from "socket://localhost:8001"
3
4 preamble{ starter: a }
5
6 aioc {
7   _r@a = startTimer( n ); {
8     scope @a {
9       10@a = getNthChar( 0 ); 10@a = getNext( 10 ); _r@a = setNthChar( 0, 10 ) }
10     prop { N.char = 0 } |
11     scope @b {
12       11@b = getNthChar( 1 ); 11@b = getNext( 11 ); _r@b = setNthChar( 1, 11 ) }
13     prop { N.char = 1 } |
14     scope @a {
15       12@a = getNthChar( 2 ); 12@a = getNext( 12 ); _r@a = setNthChar( 2, 12 ) }
16     prop { N.char = 2 } |
17     scope @b {
18       13@b = getNthChar( 3 ); 13@b = getNext( 13 ); _r@b = setNthChar( 3, 13 ) }
19     prop { N.char = 3 } |
20     scope @a {
21       14@a = getNthChar( 4 ); 14@a = getNext( 14 ); _r@a = setNthChar( 4, 14 ) }
22     prop { N.char = 4 }
23   };
24   _r@a = stopTimer( n )
25 }
```

Listing B.5: Fork-join scenario with scopes.

Likewise the pipe scenario, the scopes in Listing B.5 enclose the tasks of List-

ing B.4. Also in this scenario each scope has a unique property `N.char` that identifies it uniquely.

Rules of the fork-join scenario. The adapted behaviour introduces the new function `GetDoubleNext` that gets a character and returns the character two positions after it in the alphabet. Listing B.6 shows two rules used in the fork scenario. The rule at Lines 1-6 applies to the scope at Lines 8-10 of Listing B.5, while the rule at Lines 7-12 applies to the scope at Lines 11-13.

```

1 rule {
2   include getNthChar, getDoubleNext, setNthChar from "socket://localhost:8001"
3   on { N.char == 0 }
4   do { l0@b = getNthChar( 0 ); l0@b = getDoubleNext( l0 );
5       _r@b = setNthChar( 0, l0 ) }
6 }
7 rule {
8   include getNthChar, getDoubleNext, setNthChar from "socket://localhost:8001"
9   on { N.char == 1 }
10  do { l1@b = getNthChar( 1 ); l1@b = getDoubleNext( l1 );
11     _r@b = setNthChar( 1, l1 ) }
12 }

```

Listing B.6: Rules of pipe scenario.

B.1.2 AIO CJ programs used for benchmarking primitives

```

1 include start, end from "socket://localhost:8000"
2 preamble { starter: a }
3 aioc {
4   _r@a = start( "assignment" );
5   x@a = 1;
6   _r@a = end( "assignment" );
7
8   _r@a = start( "interaction" );
9   pass: a( x ) -> b( x );
10  _r@a = end( "interaction" );
11
12  _r@a = start( "if statement" );
13  if ( x == 1 )@a { skip };
14  _r@a = end( "if statement" )
15 }

```

Listing B.7: Code for benchmarking assignment, interaction, and the if statement.

```

1 include start, end
2   from "socket://localhost:8000"
3

```

B.1. Code used for validation

```
4 preamble { starter: a }
5
6 aioc {
7   _r@a = start( "scope" );
8   scope @a{ skip }
9   prop { N.scope_name = "applicable"};
10  _r@a = end( "scope" )
11 }
```

Listing B.8: Code for benchmarking the scope primitive.

Listings B.7 and B.8 show the code used to measure the performances of primitives of AIOCI. We tested the performances of the scope (Listing B.8) under 7 contexts:

1. without adaptation servers;
2. with 1 adaptation server, but without rules;
3. with 1 adaptation server and 1 matching rule;
4. with 1 adaptation server, 50 rules but none matching;
5. with 1 adaptation server, 50 rules, and 1 matching;
6. with 1 adaptation server, 100 rules but none matching;
7. with 1 adaptation server, 100 rules, and 1 matching;

```
1 rule {
2   on { N.scope_name == "applicable" }
3   do { skip }
4 }
5 rule {
6   on { N.scope_name == "non-applicable" }
7   do { skip }
8 }
```

Listing B.9: Rules used for benchmarking scopes.

Listing B.9 shows the two kinds of rules used for testing performances of Listing B.8. The first one (Lines 1-4) is the matching rule used in contexts 3, 5, and 7. The second one (Lines 5-8) is a non matching rule present, in different copies, in contexts 4, 5, 6, and 7.

Adaptable Choreographies: Models of Adaptation

To model in AIOCJ the advanced examples below, we introduce the keyword `roles` for AIOC `scopes`. Consider the code below:

```
scope @Role1 {  
    /* code */  
} roles { Role2 }
```

the keyword `roles` attached to the `scope` means that `Role2` passively participates to the `scope`, i.e., it has no interactions. At runtime, this feature enables the application of updates that comprise `Role2` among its roles. In this way, scopes comprise passive roles that can play an important part in the updated code, e.g., a `Logger` that is activated by an update.

This feature does not detach the implementation of AIOCJ from the theoretical model of Dynamic Choreographies, yet it allows to model advanced behaviours like the ones below.

C.1 A distributed adaptive document system

Here we provide an account of how we can model aspect orientation in AIOCs.

The example proposed here is taken from [118] and gives a pragmatic assessment on the expressiveness of the AIOC language with respect to distributed dynamic aspect-oriented systems.

In the scenario proposed in [118], the authors describe the classic server-client interaction in which the server serves documents to Clients. The server implements three “methods” to interact with documents, namely:

- Search, which allows a Client to search a certain string on a document;
- Read, which returns the content of a document;
- Write, which modifies the content of a document.

Appendix C. Adaptable Choreographies: Models of Adaptation

```
include search, read, write from "socket://localhost:8002"

preamble {
  starter: Client
}

aioc {
  continue@Client = true;

  while( continue )@Client{
    r@Client = getInput( "Select operation: (S)earch, " +
      "(R)ead, (W)rite, or (E)xit" );

    if( r == "S" )@Client{
      s@Client = getInput( "Insert search expression" );
      // for logging purposes
      scope @Client{ skip } prop { N.scopename = "log" }
      roles { Logger };
      // adapt the search protocol
      scope @Logger{
        search: Client( s ) -> Server( s );
        res@Server = search( s );
        response: Server( res ) -> Client( res )
      } prop{ N.scopename = "search" }
      roles { Balancer, Rserver };
      _r@Client = show( res )
    };

    if( r == "R" )@Client{
      s@Client = getInput( "Insert the page to read" );
      // for logging purposes
      scope @Client{ skip } prop { N.scopename = "log" }
      roles { Logger };
      // adapt the read protocol
      scope @Logger{
        read: Client( s ) -> Server( s );
        res@Server = read( s );
        response: Server( res ) -> Client( res )
      } prop { N.scopename = "read" }
      roles { Balancer, Rserver };
      _r@Client = show( res )
    };

    if( r == "W" )@Client{
      s@Client = getInput( "Insert the page and the " +
        "content of the modification (page, content)" );
      scope @Client{ skip } prop { N.scopename = "log" }
      roles { Logger };
      scope @Logger{
```

C.1. A distributed adaptive document system

```
    write: Client( s ) -> Server( s );
    res@Server = write( s );
    response: Server( res ) -> Client( res )
} prop{ N.scopename = "write" }
roles { Balancer, Rserver };
_r@Client = show( res )
};

if( r == "E" )@Client{
    continue@Client = false
};

if( r != "S" and r != "E" and
    r != "R" and r != "W" )@Client{
    r@Client = show( "Insert command not valid" )
}
}
}

rule {
    include log, getLoadBalancing from "socket://localhost:8000"
    on { N.scopename == "log" }
    do {
        log: Client( r ) -> Logger( log );
        _r@Logger = log( log );
        if( log == "S" )@Logger {
            search_balance@Logger = getLoadBalancing( log )
        };
        if( log == "R" )@Logger {
            read_balance@Logger = getLoadBalancing( log )
        }
    }
}

rule {
    include getServer from "socket://localhost:8001"
    include search from "socket://localhost:8002"
    include rSearch from "socket://localhost:8003"
    on { N.scopename == "search" and search_balance == true }
    do {
        ser@Balancer = getServer( "S" );
        if( ser )@Balancer{
            search: Client( s ) -> Server( s );
            res@Server = search( s );
            response: Server( res ) -> Client( res )
        } else {
            search: Client( s ) -> Rserver( s );
            res@Rserver = rSearch( s );
            response: Rserver( res ) -> Client( res )
        }
    }
}
```



```

    }
  }
}

rule {
  include getServer from "socket://localhost:8001"
  include read from "socket://localhost:8002"
  include rRead from "socket://localhost:8003"
  on { N.scopename == "read" and read_balance == true }
  do {
    ser@Balancer = getServer( "R" );
    if( ser )@Balancer{
      read: Client( s ) -> Server( s );
      res@Server = read( s );
      response: Server( res ) -> Client( res )
    } else {
      read: Client( s ) -> Rserver( s );
      res@Rserver = rRead( s );
      response: Rserver( res ) -> Client( res )
    }
  }
}

rule {
  include getServer from "socket://localhost:8001"
  include write from "socket://localhost:8002"
  include rWrite from "socket://localhost:8003"
  on {
    N.scopename == "write" and
    ( read_balance == true or search_balance == true )
  }
  do {
    write: Client( s ) -> Server( s );
    res@Server = write( s );
    // propagates changes to Rserver
    write: Server( s ) -> Rserver( s );
    res@Rserver = rWrite( s );
    ok: Rserver() -> Server();
    res@Server = "[Synchronised] " + res;
    response: Server( res ) -> Client( res )
  }
}

```

C.1.1 Pointcuts

First, we model one of the basic concepts of AOP which is pointcut. A pointcut describes a set of join points. Each time the execution of a program reaches one of

C.1. A distributed adaptive document system

the join points relative to a pointcut, a piece of code associated with the pointcut is executed. We implement pointcuts as empty scopes defined before each operation. The empty scopes make use of the feature `roles{...}` that lets us to specify additional roles in the scope. Additional roles do not take an active part in the scope they are added in, but their behaviour relative to that scope can be changed by an adaptation rule. In the example, if the adaptation rule `N.scope_name = log` applies, we allow a `Logger` to count the number of interactions relative to search, read, and write operations.

```
scope @Client{ skip }
prop { N.scopename = "log" }
roles { Logger };

rule {
  include log, getLoadBalancing
  from "socket://localhost:8000"
  on { N.scopename == "log" }
  do {
    log: Client( r ) -> Logger( log );
    _r@Logger = log( log );
    if( log == "S" )@Logger {
      search_balance@Logger =
        getLoadBalancing( log )
    };
    if( log == "R" )@Logger {
      read_balance@Logger =
        getLoadBalancing( log )
    }
  }
}
```

C.1.2 Dynamic wrappers

Dynamic wrappers, as defined in JAC, are software entities that allow the modification or the enhancement of the semantics of some base objects. They provide the ability to add code before, after, or around existing methods.

In the proposed example, the authors suggest to apply distributed aspects to avoid the overload of a single server. Namely, they introduce aspects that manage distribution of documents, load-balancing, and coherence among copies of the same document.

As for pointcuts, we make use of the feature `roles{...}` to include in the interaction of each operation a `Balancer` and an additional server (`Rserver`). Then, in the adaptation rules the `Balancer` routes the operations between the two servers to balance the load. Furthermore, in order to preserve coherence among documents, each time a write operation is executed, the server that applied the modifications notifies the other to synchronise the state of the copy of the document.

C.2 ContextChat

In this section we model context orientation in the AIOC language.

The scenario proposed here is taken from [73] in which the authors introduce context-oriented programming within the actor model of the Erlang language. Below we provide a pragmatic assessment on the expressiveness of the AIOC language with respect to distributed dynamic context-oriented systems.

Context-oriented evolve from aspect-oriented programming by introducing mechanisms to enforce the coherent application of aspects.

In [73] the authors propose the example of ContextChat to highlight the features of context-oriented programming applied to the Erlang actor model. In particular the authors introduce a fixed stack of adaptation procedures. The slots that compose said stack describe the behaviour of adaptable procedures (also called variations):

- activatable slots contain one variation which can be independently activated;
- switch slots contain several mutually exclusive variations;
- free slots contain an undefined single variation which can be acquired and assigned from other programs.

In ContextChat connected clients exchange messages in real time. The system implements advanced features as context variants:

- offline reception: if an user goes offline, other users can still send messages to him. The system saves offline messages and shows them to the recipient when s/he returns online. Offline and online statuses are described in a switch slot, therefore mutually exclusive;
- backup feature: the user can activate (activatable slot) a backup mode that saves all messages (sent and received) on a remote server;
- tracing feature: the system can activate (activatable slot) a tracing mode that collects information on client communication to handle network communication in a more efficient way (e.g., using internal messages in place of network ones);
- text effects: users can submit (free slot) customisable filters to the messages (e.g., for text emphasising, emoticons, etc.).

Modelling contexts with scopes

Let us analyse how the different behaviours described by activatable, switch, and free slots map to AIOCI scopes.

Activatable and free slots coalesce in the same type of scopes since in AIOCI the source of variations are rules that can be added and removed at runtime (i.e., the set of running Adaptation Servers, cf. § 5.1). Activation of such variations is based on variables belonging to the environment (prefixed with E.), to the scope, which are non-functional (prefixed with N.), and from the status of the coordinator of the adaptation.

Switch slots require a choice in the design of the adaptation with respect to the basic and the altered mutually exclusive behaviours.

Referring to the offline/online statuses of ContextChat, we assume that "online" is the basic behaviour which alternates with the "offline" status following the preferences of the User. Hence, the resulting AIOC contains the basic behaviour (online) and, depending on the availability of Adaptation Servers and the User's status, mutually exclusive behaviours will overwrite the basic one. Since after adaptation scope boundaries are removed, each time the basic behaviour adapts only the matching rule overwrites its behaviour, excluding any other applicable variant.

To keep the example simple and since activatable slots directly map to scopes, we implement the example of ContextChat with the online/offline variants and the free slot for text emphasising.

We report below the full code of the AIOC and the adaptation rules.

Appendix C. Adaptable Choreographies: Models of Adaptation

```
preamble { starter: User1 }

aioc {
  { name@User2 = "User2"    | name@User1 = "User1"
    | continue@User1 = true  | continue@User2 = true
    | status@User1 = "online" | status@User2 = "online"
  };
  {
    while( continue == true )@User1{
      scope @User1{
        msg@User1 = getInput( name + ": insert a message" )
      } prop { N.scope_name = "get_msg" };
      scope @User1{
        send: User1( msg ) -> User2( msg );
        scope @User2{
          _r@User2 = show( name + ": " + msg )
        } prop { N.scope_name = "display_msg" }
      }
      |
      c@User1 = getInput( name + ": do you want to continue [y/n]? ");
      if( c == "n" )@User1{
        continue@User1 = false | continue@User2 = false
      }
    } prop { N.scope_name = "status_switch_user1" }
  }
  |
  while( continue == true )@User2{
    scope @User2 {
      msg@User2 = getInput( name + ": insert a message" )
    } prop { N.scope_name = "get_msg" };
    scope @User2{
      send: User2( msg ) -> User1( msg );
      scope @User1{
        _r@User1 = show( name + ": " + msg )
      } prop { N.scope_name = "display_msg" }
    }
    |
    c@User2 = getInput( name + ": do you want to continue [y/n]? ");
    if( c == "n" )@User2{
      continue@User1 = false | continue@User2 = false
    }
  } prop { N.scope_name = "status_switch_user2" }
}
}
```

C.2. ContextChat

```
// the rule applies on any scope with
// property scope_name equal to "get_msg"
rule {
  on { N.scope_name == "get_msg" }
  do {
    msg@User1 = getInput( name + ": insert a message or " +
      "write 'exit' to go offline" );
    if( msg == "exit" )@User1{ status@User1 = "offline" }
  }
}

// the rule applies when the User is offline. It acts as a switch
// and, when the User returns online, displays the stored messages
rule {
  on { status == "offline" and N.scope_name == "status_switch_user1" }
  do {
    _r@User1 = show( name + ": click ok to go online" );
    status@User1 = "online";
    if( stored_msg_count > 0 )@User1 {
      _r@User1 = show( name + " received: " + stored_msg );
      stored_msg@User1 = ""; stored_msg_count@User1 = 0
    }
  }
}

// the rule applies when the User is offline.
// It stores the message for future display
rule {
  on { status == "offline" and N.scope_name == "display_msg" }
  do {
    if( stored_msg_count > 0 )@User1{
      stored_msg@User1 = stored_msg + "; " + msg
    } else { stored_msg@User1 = msg };
    stored_msg_count@User1 = stored_msg_count + 1
  }
}

// the rule applies when the User is online. It adds text effects to
// the displayed message we apply text effect only on online display
rule {
  on { status == "online" and N.scope_name == "display_msg" and
    N.emph_effect != "applied" }
  do {
    msg@User1 = "*" + msg + "*";
    scope @User1{
      _r@User1 = show( msg )
    } prop { N.scope_name = "display_msg", N.emph_effect = "applied" }
  }
}
```

Appendix C. Adaptable Choreographies: Models of Adaptation

```
// Rules below mirror the ones above but apply only on User2
rule {
  on { N.scope_name == "get_msg" }
  do {
    msg@User2 = getInput( name + ": insert a message or " +
      "write 'exit' to go offline" );
    if( msg == "exit" )@User2{
      status@User2 = "offline"
    }
  }
}

rule {
  on { status == "offline" and
    N.scope_name == "status_switch_user2" }
  do {
    _r@User2 = show( name + ": click ok to go online" );
    status@User2 = "online";
    if( stored_msg_count > 0 )@User2 {
      _r@User2 = show( name + " received: " + stored_msg );
      stored_msg@User2 = "";
      stored_msg_count@User2 = 0
    }
  }
}

rule {
  on { status == "offline" and N.scope_name == "display_msg" }
  do {
    if( stored_msg_count > 0 )@User2{
      stored_msg@User2 = stored_msg + "; " + msg
    } else {
      stored_msg@User2 = msg
    };
    stored_msg_count@User2 = stored_msg_count + 1
  }
}

rule {
  on { status == "online" and N.scope_name == "display_msg"
    and N.emph_effect != "applied" }
  do {
    msg@User2 = "*" + msg + "*";
    scope @User2{
      _r@User2 = show( msg )
    } prop { N.scope_name = "display_msg",
      N.emph_effect = "applied" }
  }
}
```

C.2.1 Online/Offline switch

Since Online status is the basic behaviour, the choreography implements it as its default. However, the programmer can specify that the procedure for acquire a User's message can be changed (below). In this case the rule always applies and augments the functionality of the interaction with the User which can either send a message or go offline by writing "exit".

```
scope @User1{
  msg@User1 = getInput( name + ": insert a message" )
} prop { N.scope_name = "get_msg" }
```

```
// rule applies on any scope with property
// scope_name equal to "get_msg"
rule {
  on { N.scope_name == "get_msg" }
  do { msg@User1 = getInput( name +
    ": write a message or 'exit' to go offline" );
    if( msg == "exit" )@User1{
      status@User1 = "offline"
    }
  }
}
```

Note that the scope, led and participated only by User1, indicates as `N.scope_name = "get_msg"` its only non-functional property. User2 shows an equivalent scope (with the non-functional property `N.scope_name = "get_msg"`), yet the adaptation mechanisms of AIOCJ apply the corresponding rule to the right scope as it checks applicability of rules on the participants to the scope. Since in this case the leaders are also the only participants to the scope, no further properties are needed to ensure coherent application of adaptation.

Subsequently, the scope with `N.scope_name = "status_switch_user1"` overwrites the sending procedure. The rules procedure waits for the user to return online and displays the messages received in the meanwhile. Notably the rule applies only if the status variable of the leaders is set to "offline".

```
scope @User1{
  send: User1( msg ) -> User2( msg );
{
  scope @User2{
    _r@User2 = show( name + ": " + msg )
  } prop { N.scope_name = "display_msg" }
  |
  c@User1 = getInput( name + ": do you want to continue [y/n]? ");
  if( c == "n" )@User1{
    continue@User1 = false | continue@User2 = false
  }
}
}
prop { N.scope_name = "status_switch_user1" }
```



```

rule {
  on { status == "offline" and
        N.scope_name == "status_switch_user1" }
  do { _r@User1 = show( name + ": click ok to go online" );
        status@User1 = "online";
        if( stored_msg_count > 0 )@User1 {
          _r@User1 = show( name + " received: " + stored_msg );
          stored_msg@User1 = "";
          stored_msg_count@User1 = 0
        }
      }
}

```

Finally, if the user is online and sends a message, the addressee (User2, in case of User1) enters the scope with `N.scope_name = "display_msg"`. The scope serves a twofold purpose: in case the addressee is offline, the overwriting rule (which is led and participated by User2) stores the message into a storage variable for future visualisation, if otherwise User2 is online the text-effect rule applies. Notably text-effects can stack as the adaptation code of the rule bears a nested scope with the same properties of the adapted scope, beside the one `N.emph_effect = "applied"` which prevents the infinite application of the same rule.

```

rule {
  on { status == "offline" and N.scope_name == "display_msg" }
  do {
    if( stored_msg_count > 0 )@User1{
      stored_msg@User1 = stored_msg + "; " + msg
    } else {
      stored_msg@User1 = msg
    };
    stored_msg_count@User1 = stored_msg_count + 1
  }
}

rule {
  on { status == "online" and N.scope_name == "display_msg" and
        N.emph_effect != "applied" }
  do {

    msg@User1 = "*" + msg + "*";
    scope @User1{
      _r@User1 = show( msg )
    } prop { N.scope_name = "display_msg",
             N.emph_effect = "applied" }
  }
}

```

Applied Choreographies: Additional Material

D.1 Applied Choreographies

$$\begin{array}{c}
\frac{\text{pn}(\eta) \cap \text{pn}(\eta') = \emptyset}{\eta; \eta' \simeq_C \eta'; \eta} \quad [\text{CS}|_{\text{ETAETA}}] \quad \frac{\text{p} \notin \text{pn}(\eta)}{\text{if p.e } \{\eta; C_1\} \text{ else } \{\eta; C_2\} \simeq_C \eta; \text{if p.e } \{C_1\} \text{ else } \{C_2\}} \quad [\text{CS}|_{\text{ETACND}}] \\
\frac{\text{q} \notin \text{pn}(\eta)}{k : \mathbf{A} \rightarrow \text{q}[\mathbf{B}].\{o_i(x_i); \eta; C_i\}_{i \in I} \simeq_C \eta; k : \mathbf{A} \rightarrow \text{q}[\mathbf{B}].\{o_i(x_i); C_i\}_{i \in I}} \quad [\text{CS}|_{\text{ETARCV}}] \\
\frac{\text{p} \neq \text{q}}{k : \mathbf{A} \rightarrow \text{p}[\mathbf{B}].\{o_i(x_i); k' : \mathbf{C} \rightarrow \text{q}[\mathbf{D}].\{o'_{ij}(x'_{ij}); C_{ij}\}_{j \in J}\}_{i \in I} \simeq_C k' : \mathbf{C} \rightarrow \text{q}[\mathbf{D}].\{o'_{ij}(x'_{ij}); k : \mathbf{A} \rightarrow \text{p}[\mathbf{B}].\{o_{ij}(x_{ij}); C_{ij}\}_{i \in I}\}_{j \in J}} \quad [\text{CS}|_{\text{RCVRCV}}] \\
\frac{\text{p} \neq \text{q}}{\text{if p.e } \{\text{if q.e}' \{C_1\} \text{ else } \{C_2\}\} \text{ else } \{\text{if q.e}' \{C'_1\} \text{ else } \{C'_2\}\} \simeq_C \text{if q.e}' \{\text{if p.e } \{C_1\} \text{ else } \{C'_1\}\} \text{ else } \{\text{if p.e } \{C_2\} \text{ else } \{C'_2\}\}} \quad [\text{CS}|_{\text{CNDCND}}] \\
\frac{\text{p} \neq \text{q}}{\text{if q.e } \{k : \mathbf{A} \rightarrow \text{p}[\mathbf{B}].\{o_i(x_i); C_{i1}\}_{i \in I}\} \text{ else } \{k : \mathbf{A} \rightarrow \text{p}[\mathbf{B}].\{o_i(x_i); C_{i2}\}_{i \in I}\} \simeq_C k : \mathbf{A} \rightarrow \text{p}[\mathbf{B}].\{o_i(x_i); \text{if q.e } \{C_{i1}\} \text{ else } \{C_{i2}\}\}_{i \in I}} \quad [\text{CS}|_{\text{RCVCND}}]
\end{array}$$

Figure D.1: Choreography Calculus, swap relation \simeq_C

$$\begin{array}{c}
 \frac{\Gamma, \tilde{l} : G\langle A|\tilde{B}|\tilde{B}\rangle, \text{init}(r[\tilde{C}], k, G) \vdash C \quad r[\tilde{C}] = p[A], q[\tilde{B}] \quad \tilde{q} \notin \Gamma}{\Gamma, \tilde{l} : G\langle A|\tilde{B}|\tilde{B}\rangle \vdash \mathbf{start} \ k : p[A] \Leftrightarrow \tilde{l}.q[\tilde{B}]; C} [T]_{\text{START}} \\
 \frac{\Gamma, p : k[A], k[A] : \llbracket G \rrbracket_A \vdash C \quad \Gamma \vdash \tilde{l} : G\langle A|\tilde{B}|\emptyset\rangle}{\Gamma \vdash \mathbf{req} \ k : p[A] \Leftrightarrow \tilde{l}.B; C} [T]_{\text{REQ}} \\
 \frac{\tilde{l} \subseteq \tilde{l}' \quad \Gamma, \tilde{l}' : G\langle A|\tilde{B}|\emptyset\rangle, \text{init}(q[\tilde{C}], k, G) \vdash C}{\Gamma, \tilde{l}' : G\langle A|\tilde{B}|\tilde{C}\rangle \vdash \mathbf{acc} \ k : \tilde{l}.q[\tilde{C}]; C} [T]_{\text{ACC}} \\
 \frac{\Gamma \vdash p.e : \mathbf{bool} \quad \Gamma \vdash C_1 \quad \Gamma \vdash C_2}{\Gamma \vdash \mathbf{if} \ p.e \ \{C_1\} \ \mathbf{else} \ \{C_2\}} [T]_{\text{COND}} \\
 \frac{j \in I \quad \Gamma \vdash p : k[A], q : k[B] \quad \Gamma \vdash p.e : U_j \quad \Gamma, q.x : U_j, k[A] : T_j, k[B] : T'_j \vdash C}{\Gamma, k[A] : !B.\{o_i(U_i); T_i\}_{i \in I}, k[B] : ?A.\{o_i(U_i); T'_i\}_{i \in I} \vdash k : p[A].e \rightarrow q[B].o_j(x); C} [T]_{\text{COM}} \\
 \frac{j \in I \quad \Gamma \vdash p : k[A] \quad \Gamma \vdash p.e : U_j \quad \Gamma, k[A] : T_j \vdash C}{\Gamma, k[A] : !B.\{o_i(U_i); T_i\}_{i \in I} \vdash k : p[A].e \rightarrow B.o_j; C} [T]_{\text{SEND}} \\
 \frac{\Gamma \vdash q : k[B] \quad \forall j \in I. \Gamma, q.x_j : U_j, k[B] : T_j \vdash C_j}{\Gamma, k[B] : ?A.\{o_i(U_i); T_i\}_{i \in I} \vdash k : A \rightarrow q[B].\{o_j(x_j); C_j\}_{j \in I \cup J}} [T]_{\text{RECV}} \\
 \frac{\Gamma, X : \Gamma' \vdash C \quad \Gamma', X : \Gamma' \vdash C' \quad \Gamma'_{\text{locs}} \subseteq \Gamma \quad \tilde{p} = \text{pn}(C')}{\Gamma \vdash \mathbf{def} \ X\langle \tilde{p} \rangle = C' \ \mathbf{in} \ C} [T]_{\text{DEF}} \\
 \frac{\Gamma_i \vdash C_i \quad \Gamma_2 \vdash C_2}{\Gamma_1 \circ \Gamma_2 \vdash C_1 \mid C_2} [T]_{\text{PAR}} \quad \frac{\text{end}(\Gamma)}{\Gamma \vdash \mathbf{0}} [T]_{\text{END}} \quad \frac{\text{end}(\Gamma) \quad \Gamma'' \subseteq \Gamma'}{\Gamma, \Gamma'', X\langle \tilde{p} \rangle : \Gamma' \vdash X\langle \tilde{p} \rangle} [T]_{\text{CALL}}
 \end{array}$$

Figure D.2: Choreography Calculus - Typing Rules

D.2 Typing

In addition to the description in § 6.3 we comment the additional rules in Figure D.2.

In Rule $[T]_{\text{END}}$ $\text{end}(\Gamma)$ holds if the protocols for all sessions have terminated (i.e., all local typings have type end). In Rule $[T]_{\text{DEF}}$ the condition $\Gamma'_{\text{locs}} \subseteq \Gamma$ checks that the body of the recursive procedure does not introduce unexpected services (Γ'_{locs} returns all service typings in Γ').

D.2. Typing

$$\begin{array}{c}
\frac{\{A, B\} \cap \{C, D\} = \emptyset}{A \rightarrow B.\{o_i(U_i); C \rightarrow D.\{o'_j(U'_j); G_{ij}\}_{j \in J}\}_{i \in I} \simeq_G C \rightarrow D.\{o'_j(U'_j); A \rightarrow B.\{o_i(U_i); G_{ij}\}_{i \in I}\}_{j \in J}} \quad [GS]_{\text{COMCOM}} \\
\frac{\{A, B\} \cap \{D\} = \emptyset}{A \rightarrow B.\{o_i(U_i); C \rightsquigarrow D.o(U); G_i\} \simeq_G C \rightsquigarrow D.o(U); A \rightarrow B.\{o_i(U_i); G_i\}} \quad [GS]_{\text{COMRECV}} \\
\frac{B \neq D}{A \rightsquigarrow B.o(U); C \rightsquigarrow D.o'(U); G \simeq_G C \rightsquigarrow D.o'(U); A \rightsquigarrow B.o(U); G} \quad [GS]_{\text{RECVRECV}}
\end{array}$$

Figure D.3: Global types, Swap Relation \simeq_G .

$$\begin{array}{l}
\llbracket C \rightarrow D.\{o_i(U_i); G_i\} \rrbracket_B^A = \begin{cases} \text{end} & \text{if } C = A \wedge D = B \\ \bigsqcup_i \llbracket G_i \rrbracket_B^A & \text{otherwise} \end{cases} \\
\llbracket C \rightsquigarrow D.o(U); G \rrbracket_B^A = \begin{cases} ?A.o(U); \llbracket G \rrbracket_B^A & \text{if } C = A \wedge D = B \\ \llbracket G \rrbracket_B^A & \text{otherwise} \end{cases} \\
\llbracket t \rrbracket_B^A = \text{end} \quad \llbracket \text{end} \rrbracket_B^A = \text{end} \quad \llbracket \text{rec } t; G \rrbracket_B^A = \text{end}
\end{array}$$

Figure D.4: Choreography Calculus - Buffer Type Projection

D.3 Endpoint Projection

We report the complete definition of choreography annotation, process projection, merging function, and grouping function, respectively in Figure D.5, Figure D.6, Figure D.7, and Figure D.8.

$$\begin{aligned}
 @C &= \emptyset @C \\
 d @ C &= \begin{cases} \eta; d @ C' & \text{if } C = \eta; C' \\ d @ C_1 \mid d @ C_2 & \text{if } C = C_1 \mid C_2 \\ k : A \rightarrow \mathbf{q}[\mathbf{B}].\{o_i(x_i); d @ C_i\}_{i \in I} & \text{if } C = k : A \rightarrow \mathbf{q}[\mathbf{B}].\{o_i(x_i); C_i\}_{i \in I} \\ \text{def } X \langle \tilde{p} \rangle = d' @ C'' \text{ in } d' @ C' & \text{if } \begin{cases} C = \text{def } X = C'' \text{ in } C' \wedge \\ \tilde{p} = \text{fp}(C'') \wedge d' = d \cup (X, X \langle \tilde{p} \rangle) \end{cases} \\ \mathbf{acc } k : l.\widetilde{\mathbf{q}[\mathbf{B}]}; d @ C' & \text{if } C = \mathbf{acc } k : l.\widetilde{\mathbf{q}[\mathbf{B}]}; C' \\ \text{if } p.e \{d @ C_1\} \text{ else } \{d @ C_2\} & \text{if } C = \text{if } p.e \{C_1\} \text{ else } \{C_2\} \\ X \langle \tilde{p} \rangle & \text{if } C = X \wedge (X, X \langle \tilde{p} \rangle) \in d \\ \mathbf{0} & \text{if } C = \mathbf{0} \end{cases}
 \end{aligned}$$

Figure D.5: Choreography Calculus — Annotation operator

D.3. Endpoint Projection

$$\begin{aligned}
\llbracket \mathbf{start} \ k : p[A] \Leftrightarrow l.q[B] \rrbracket_r &= \begin{cases} \mathbf{req} \ k : p[A] \Leftrightarrow l.B; \llbracket C \rrbracket_r & \text{if } r = p \\ \mathbf{acc} \ k : l.r[C]; \llbracket C \rrbracket_r & \text{if } l.r[C] \in \{l.q[B]\} \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \\
\llbracket \mathbf{acc} \ k : l.q[B] \rrbracket_r &= \begin{cases} \mathbf{acc} \ k : l.r[C]; \llbracket C \rrbracket_r & \text{if } l.r[C] \in \{l.q[B]\} \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \\
\llbracket \eta; C \rrbracket_r, \eta \in \{\mathit{req}, \mathit{send}\} &= \begin{cases} \eta; \llbracket C \rrbracket_r & \text{if } r \in \text{fn}(\eta) \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \\
\llbracket k : p[A].e \rightarrow q[B].o(x); C \rrbracket_r &= \begin{cases} k : p[A].e \rightarrow B.o; \llbracket C \rrbracket_r & \text{if } r = p \\ k : A \rightarrow q[B].o(x); \llbracket C \rrbracket_r & \text{if } r = q \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \\
\llbracket k : A \rightarrow q[B].\{o_i(x_i); C_i\}_{i \in I} \rrbracket_r &= \begin{cases} k : A \rightarrow q[B].\{o_i(x_i); \llbracket C_i \rrbracket_r\}_{i \in I} & \text{if } r = q \\ \bigsqcup_{i \in I} \llbracket C_i \rrbracket_r & \text{otherwise} \end{cases} \\
\llbracket \text{if } p.e \ \{C_1\} \ \text{else} \ \{C_2\} \rrbracket_r &= \begin{cases} \text{if } p.e \ \{\llbracket C_1 \rrbracket_r\} \ \text{else} \ \{\llbracket C_2 \rrbracket_r\} & \text{if } r = p \\ \llbracket C_1 \rrbracket_r \sqcup \llbracket C_2 \rrbracket_r & \text{otherwise} \end{cases} \\
\llbracket \mathbf{def} \ X \langle \tilde{p} \rangle = C' \ \text{in} \ C \rrbracket_r &= \begin{cases} \mathbf{def} \ X_r = \llbracket C' \rrbracket_r \ \text{in} \ \llbracket C \rrbracket_r & \text{if } r \in \{\tilde{p}\} \\ \llbracket C \rrbracket_r & \text{otherwise} \end{cases} \\
\llbracket X \langle \tilde{p} \rangle \rrbracket_r &= \begin{cases} X_r & \text{if } r \in \{\tilde{p}\} \\ \mathbf{0} & \text{otherwise} \end{cases} \\
\llbracket C_1 \mid C_2 \rrbracket_r &= \llbracket C_1 \rrbracket_r \mid \llbracket C_2 \rrbracket_r \\
\llbracket \mathbf{0} \rrbracket_r &= \mathbf{0}
\end{aligned}$$

Figure D.6: Choreography Calculus, process projection

$$\begin{aligned}
 \mathbf{req} \ k : p[A] \Leftrightarrow \widetilde{l.B}; C_1 \sqcup \mathbf{req} \ k : q[A] \Leftrightarrow \widetilde{l.B}; C_2 &= \mathbf{req} \ k : p[A] \Leftrightarrow \widetilde{l.B}; (C_1 \sqcup C_2) \\
 \mathbf{acc} \ k : l.p[A]; C_1 \sqcup \mathbf{acc} \ k : l.q[A]; C_2 &= \mathbf{acc} \ k : l.p[A]; (C_1 \sqcup C_2) \\
 k : p[A].e \rightarrow B.o; C_1 \sqcup k : q[A].e \rightarrow B.o; C_2 &= k : p[A].e \rightarrow B.o; (C_1 \sqcup C_2) \\
 k : A \rightarrow p[B]. \{o_i(x_i); C_i\}_{i \in I} \sqcup k : A \rightarrow q[B]. \{o_j(x_j); C'_j\}_{j \in J} &= \\
 k : A \rightarrow p[B]. \left\{ \begin{array}{l} \{o_i(x_i); C_i\}_{i \in I/J} \\ \cup \{o_i(x_i); C'_i\}_{i \in J/I} \\ \cup \{o_i(x_i); C_i \sqcup C'_i\}_{i \in I \cap J} \end{array} \right\} & \\
 \text{if } p.e \{C_1\} \text{ else } \{C'_1\} \sqcup \text{if } q.e \{C_2\} \text{ else } \{C'_2\} &= \text{if } p.e \{C_1 \sqcup C_2\} \text{ else } \{C'_1 \sqcup C'_2\} \\
 \text{def } X = C'_1 \text{ in } C_1 \sqcup \text{def } X = C'_2 \text{ in } C_2 &= \text{def } X = C'_1 \sqcup C'_2 \text{ in } C_1 \sqcup C_2 \\
 X \sqcup X &= X \qquad \mathbf{0} \sqcup \mathbf{0} = \mathbf{0}
 \end{aligned}$$

Figure D.7: Merging Function

$$\begin{aligned}
 [\mathbf{start} \ k : p[D] \Leftrightarrow \widetilde{l.q[B]}; C]_l &= [\mathbf{acc} \ k : l.q[B]; C]_l \\
 [\mathbf{acc} \ k : l.q[B]; C]_l &= \begin{cases} \{r\} \cup [C]_l & \text{if } l.r[A] \in \widetilde{l.q[B]} \\ [C]_l & \text{otherwise} \end{cases} \\
 [\eta; C]_l &= [C]_l \quad \text{if } \eta \neq (\text{start}) \\
 [\text{if } p.e \{C_1\} \text{ else } \{C_2\}]_l &= [C_1]_l \cup [C_2]_l \\
 [\text{def } X = C' \text{ in } C]_l &= [C']_l \cup [C]_l \\
 [X]_l &= \emptyset \\
 [\mathbf{0}]_l &= \emptyset \\
 [C_1 \mid C_2]_l &= [C_1]_l \cup [C_2]_l
 \end{aligned}$$

Figure D.8: Service Grouping

D.4 Dynamic Correlation Calculus

$$\begin{aligned} \text{def } X = B' \text{ in } B[X] \cdot t \cdot M &\equiv \text{def } X = B' \text{ in } B[B'] \cdot t \cdot M \\ P &\equiv P \mid \mathbf{0} \cdot t \cdot \emptyset \quad (P_1 \mid P_2) \mid P_3 \equiv P_1 \mid (P_2 \mid P_3) \\ P \mid P' &\equiv P' \mid P \quad S \mid S' \equiv S' \mid S \\ (S_1 \mid S_2) \mid S_3 &\equiv S_1 \mid (S_2 \mid S_3) \end{aligned}$$

Figure D.9: Correlation Calculus, structural congruence

Applied Choreographies: Proofs

E.1 Proofs of Subject Reduction and Session Fidelity

We define the semantics of annotated ACs by marking transitions with the name of the session whose term has reduced. We annotate other reductions as τ . We range over annotated labels with

$$\beta ::= k : A \rightarrow B.o \mid k : A \rightsquigarrow B.o(x) \mid \tau$$

We report the annotated semantics of AC in Figure E.1.

$$\begin{array}{c}
 \frac{\eta = k : \mathfrak{p}[A].e \rightarrow B.o \quad D, \eta \blacktriangleright D'}{D, \eta; C \xrightarrow{k: A \rightarrow B.o} D', C} [C]_{\text{SEND}} \\
 \\
 \frac{j \in I \quad D, k : A \rightarrow \mathfrak{q}[B].o_j(x_j) \blacktriangleright D'}{D, k : A \rightarrow \mathfrak{q}[B].\{o_i(x_i); C_i\}_{i \in I} \xrightarrow{k: A \rightsquigarrow B.o_j(x_j)} D', C_j} [C]_{\text{RECV}} \\
 \\
 \frac{\eta = k : \mathfrak{p}[A].e \rightarrow \mathfrak{q}[B].o(x) \quad D, k : \mathfrak{p}[A].e \rightarrow B.o \blacktriangleright D'}{D, \eta; C \xrightarrow{k: A \rightarrow B.o} D', k : A \rightarrow \mathfrak{q}[B].o(x); C} [C]_{\text{COM}} \\
 \\
 \frac{\#\tilde{r} \quad \#k' \quad \mathfrak{p} \in D(l) \quad \delta = \mathbf{start} \ k' : l.\mathfrak{p}[A], \widetilde{l.r[B]} \quad D, \delta \blacktriangleright D'}{D, \mathbf{start} \ k : \mathfrak{p}[A] \Leftrightarrow \widetilde{l.q[B]}; C \xrightarrow{\tau} D', C[k'/k][\tilde{r}/\tilde{q}]} [C]_{\text{START}} \\
 \\
 \frac{i \in \{1, \dots, n\} \quad \#k' \quad \{\widetilde{l.B}\} = \bigsqcup_i \{\widetilde{l_i.B_i}\} \quad \#\tilde{r} \quad \{\tilde{r}\} = \bigcup_i \{\tilde{r}_i\} \quad \mathfrak{p} \in D(l) \quad \delta = \mathbf{start} \ k' : l.\mathfrak{p}[A], \widetilde{l_1.r_1[B_1]}, \dots, \widetilde{l_n.r_n[B_n]} \quad D, \delta \blacktriangleright D'}{D, \mathbf{req} \ k : \mathfrak{p}[A] \Leftrightarrow \widetilde{l.B}; C \mid \prod_i (\mathbf{acc} \ k : \widetilde{l_i.q_i[B_i]}; C_i) \xrightarrow{\tau} D', C[k'/k] \mid \prod_i (C_i[k'/k][\tilde{r}_i/\tilde{q}_i]) \mid \prod_i (\mathbf{acc} \ k : \widetilde{l_i.q_i[B_i]}; C_i)} [C]_{\text{PSTART}} \\
 \\
 \frac{i = 1 \text{ if } \text{eval}(e, D(\mathfrak{p}).\text{st}) = \text{true}, i = 2 \text{ otherwise}}{D, \text{if } \mathfrak{p}.e \{C_1\} \text{ else } \{C_2\} \xrightarrow{\tau} D, C_i} [C]_{\text{COND}} \\
 \\
 \frac{D, C_1 \xrightarrow{\beta} D', C'_1}{D, \text{def } X = C_2 \text{ in } C_1 \xrightarrow{\beta} D', \text{def } X = C_2 \text{ in } C'_1} [C]_{\text{CTX}} \\
 \\
 \frac{D, C_1 \xrightarrow{\beta} D', C'_1}{D, C_1 \mid C_2 \xrightarrow{\beta} D', C'_1 \mid C_2} [C]_{\text{PAR}} \\
 \\
 \frac{\mathcal{R} \in \{\equiv, \simeq_c\} \quad C_1 \mathcal{R} C'_1 \quad D, C'_1 \xrightarrow{\beta} D', C'_2 \quad C'_2 \mathcal{R} C_2}{D, C_1 \xrightarrow{\beta} D', C_2} [C]_{\text{EQ}}
 \end{array}$$

Figure E.1: Choreography calculus, annotated semantics.

E.1. Proofs of Subject Reduction and Session Fidelity

We give a formal definition to $k \notin \Gamma$, which means that Γ has no local typing and buffer types for session k , formally

$$k \notin \Gamma \iff \nexists A, B \text{ s.t. } k[A]: T \in \Gamma \vee b[k]_A^B: T' \in \Gamma$$

E.1.1 Local and Typing Environment Subtyping

We define a subtyping relation on local types following [116, 96, 14]. We write the subtyping relation as $T \prec T'$, which intuitively indicates that T is more constrained than T' in its behaviour. Note that, like in [96, 14], the input type is covariant and the output type is contravariant for this relation.

Definition 32 (Local Subtyping). *We define the subtyping relation between local types as $T' \prec T$, which is the smallest relation over closed and unfolded local types, satisfying the rules below*

$$\begin{array}{c} \frac{T'' \prec T' \quad (T \approx T'' \vee T \simeq_{\top} T'')}{T \prec T'} \quad [SubT|_{EQ}] \\ \\ \frac{J \subseteq I \quad \forall i \in J \mid T_i \prec T'_i \wedge U_i \prec U'_i}{!A.\{o_i(U_i); T_i\}_{i \in I} \prec !A.\{o_i(U'_i); T'_i\}_{i \in J}} \quad [SubT|_{SEND}] \\ \\ \frac{I \subseteq J \quad \forall i \in I \mid T_i \prec T'_i \wedge U_i \prec U'_i}{?A.\{o_i(U_i); T_i\}_{i \in I} \prec ?A.\{o_i(U'_i); T'_i\}_{i \in J}} \quad [SubT|_{RECV}] \\ \\ \frac{S \prec S' \quad \forall i, U_i \prec U'_i}{S\{\underline{x}_i; U_i\}_i \prec S'\{\underline{x}_i; U'_i\}_i} \quad [SubT|_{PATH}] \quad \frac{\text{end} \approx T}{\text{end} \prec T} \quad [SubT|_{END}] \end{array}$$

Above, $T \approx T'$ is the standard tree isomorphism on recursive types whilst $T \simeq_{\top} T'$ is the swap relation for local types. The swap relation for local types is the smallest relation between local types that satisfies the Rule below. Let $\text{role}(\alpha) = A$ if either $\alpha = !A$ or $\alpha = ?A$.

$$\frac{\text{role}(\alpha) \neq \text{role}(\alpha')}{\alpha.\{o_i(U_i); \alpha'.\{o'_j(U'_j); T_{ij}\}_j\}_i \simeq_{\top} \alpha'.\{o'_j(U'_j); \alpha.\{o_i(U_i); T_{ij}\}_i\}_j} \quad [T|_{SWAP}]$$

We also define a subtyping relation between Typing Environments. Intuitively $\Gamma \prec \Gamma'$ means that Γ and Γ' are identical Typing Environments up to some local types that are more constrained in Γ — i.e., subtypes of a correspondent local type — than in Γ' .

Definition 33 (Typing Environment Subtyping). *Let Γ and Γ' be two typing environments then*

$$\Gamma' \prec \Gamma \iff \begin{cases} \Gamma = \Gamma_e, \Gamma_t \text{ s.t. } \nexists k[\mathbf{A}]: T \in \Gamma_e \\ \Gamma' = \Gamma_e, \Gamma'_t \\ \text{dom}(\Gamma_t) \cap \text{dom}(\Gamma'_t) = \emptyset \\ \forall k[\mathbf{A}]: T \in \Gamma, T \prec \Gamma'_t(k[\mathbf{A}]) \end{cases}$$

In Lemma 12 we prove that if $\Gamma' \prec \Gamma$ and Γ types a running choreography D, C also Γ' types that choreography.

Lemma 12 (Subsumption). *Let $\Gamma' \prec \Gamma$ and $\Gamma \vdash D, C$ for some D, C then $\Gamma' \vdash D, C$.*

Proof. Immediate by Definition 32 and Rules $[T]_{\text{RECV}}$ and $[T]_{\text{COM}}$. \square

We annotate the reductions of global types with labels

$$\gamma ::= A \rightarrow B.o \mid A \rightsquigarrow B.o$$

and report below the correspondent annotated semantics.

$$\begin{array}{c} \frac{j \in I}{A \rightarrow B.\{o_i(U_i); G_i\} \xrightarrow{A \rightarrow B.o_j} A \rightsquigarrow B.o_j(U_j); G_j} [G]_{\text{SEND}} \\ \frac{}{A \rightsquigarrow B.o(U); G \xrightarrow{A \rightsquigarrow B.o} G} [G]_{\text{RECV}} \quad \frac{G[\text{rec } \mathbf{t}; G/\mathbf{t}] \xrightarrow{\gamma} G'}{\text{rec } \mathbf{t}; G \xrightarrow{\gamma} G'} [G]_{\text{REC}} \\ \frac{G_1 \simeq_G G_2 \quad G_2 \xrightarrow{\gamma} G'_2 \quad G'_2 \simeq_G G'_1}{G_1 \xrightarrow{\gamma} G'_1} [G]_{\text{SWAP}} \end{array}$$

In Lemma 13 we account for the fact that any output reduction at the level of global types can constrain the projected local types of the roles not involved in the reduction. Indeed, referring to Rule $[G]_{\text{SEND}}$, the output operation chooses one of the available continuations G_j , $j \in I$ and discards all the others. Therefore the local types of the other roles not involved in the reduction can be constrained by the removal of the other branches.

Lemma 13 (Projection Subtyping). *Let $T = \llbracket G \rrbracket_{\mathbf{C}}$, $T' = \llbracket G' \rrbracket_{\mathbf{C}}$, and $\{A, B, C\} \in \text{roles}(G)$, $C \notin \{A, B\}$, then $G \xrightarrow{A \rightarrow B.o} G'$ implies $T' \prec T$.*

Proof. Easy by induction on the derivation of $G \xrightarrow{\gamma} G'$. \square

E.1. Proofs of Subject Reduction and Session Fidelity

In Lemma 14 we prove that the typing of choreographies is invariant wrt buffer types.

Lemma 14 (Buffer types invariance). *Let $\Gamma = \Gamma', \Gamma_b$ where Γ_b contains only buffer typings. If $\Gamma' \vdash C$ then $\Gamma \vdash C$.*

Proof. Trivial as buffer typings do not affect the typing of choreographies. \square

Below we restate the definition of *Deployment Judgements* enriched with pointers of the kind (DX.Y) for a clearer referencing in the proofs.

Definition 18 (Deployment Judgements) $\Gamma \vdash D \iff$

$$(D18.1) \quad \forall l \in D, \forall \{p, q\} \subseteq D(l), \quad \text{dom}(D(p).\text{que}) \cap \text{dom}(D(q).\text{que}) = \emptyset$$

$$(D18.2) \quad \forall p.x:U \in \Gamma, \vdash x(D(p).\text{st}):U$$

$$(D18.3) \quad \forall \{p: k[A], q: k[B]\} \subseteq \Gamma, \quad \underline{k}(D(p).\text{st}) = \underline{k}(D(q).\text{st})$$

$$(D18.4) \quad \forall p: k[A] \in \Gamma, \Gamma \vdash p@l \wedge p \in D(l) \wedge \underline{k.A.l}(D(p).\text{st}) = l$$

$$(D18.5) \quad \forall p: k[A] \in \Gamma \wedge \forall b[k]_A^B: T \in \Gamma, \quad \text{bte}(B, D(p).\text{que}(\underline{k.B.A}(D(p).\text{st}))) = T$$

We define a reduction relation for typing environments of the form $\Gamma \rightarrow \Gamma'$ where \rightarrow is the smallest closed under the rules below. Note that the annotation labels are a subset of the labels used to annotate the semantics of AC, ranged over by β .

$$\frac{k \notin \Gamma \quad \Gamma_k \subseteq \llbracket G \rrbracket_k \quad \{k[A]: T, k[B]: T'\} \in \Gamma_k \quad j \in I \quad G \xrightarrow{A \rightarrow B.o_j} G'}{\Gamma, \Gamma_k \xrightarrow{k: A \rightarrow B.o_j} \Gamma, \{k[C]: \llbracket G' \rrbracket_C \mid k[C] \in \Gamma_k\}, \{b[k]_C^D: \llbracket G' \rrbracket_C^D \mid b[k]_C^D \in \Gamma_k\}} \quad [\Gamma]_{\text{SEND}}$$

$$\frac{k \notin \Gamma \quad \Gamma_k \subseteq \llbracket G \rrbracket_k \quad \{k[A]: T, k[B]: T'\} \in \Gamma_k \quad \Gamma \vdash q: k[B] \quad G \xrightarrow{A \rightsquigarrow B.o_j} G'}{\Gamma, \Gamma_k \xrightarrow{k: A \rightsquigarrow B.o_j(x)} \Gamma, \{k[C]: \llbracket G' \rrbracket_C \mid k[C] \in \Gamma_k\}, \{b[k]_C^D: \llbracket G' \rrbracket_C^D \mid b[k]_C^D \in \Gamma_k\}, q.x: U_j} \quad [\Gamma]_{\text{RECV}}$$

We define the correspondence operator on label $G_{\text{act}}(\beta)$ between β and γ

$$G_{\text{act}}(\beta) = \begin{cases} A \rightarrow B.o & \text{if } \beta = k: A \rightarrow B.o \\ A \rightsquigarrow B.o & \text{if } \beta = k: A \rightsquigarrow B.o(x). \end{cases}$$

In Lemma 15 we prove that if a typing environment Γ includes local types that are projection of a global type G , then if the global type can reduce also the typing environment can reduce. The reduction preserves the correspondence between the reduced global type and the reduced local types in Γ .

Lemma 15 (Type-Environment Fidelity). *Let $\Gamma = \Gamma_*, \llbracket G \rrbracket_k$ for some $\Gamma_*, k \notin \Gamma_*$, and $G \xrightarrow{G_{\text{act}}(\beta)} G'$ then $\Gamma \xrightarrow{\beta} \Gamma'$ and for some $\Gamma'_*, k \notin \Gamma'_*, \Gamma' = \Gamma'_*, \llbracket G' \rrbracket_k$.*

Proof. Direct by cases on the derivation of Γ . □

We also report Lemmas 16 and 17 that prove that typing is invariant wrt, respectively, structural equivalence and swapping.

Lemma 16 (Subject Congruence). $\Gamma \vdash D, C$ and $C \equiv C'$ imply $\Gamma \vdash D, C'$ (up to α -renaming)

Proof. By induction on the rules that define \equiv . □

Lemma 17 (Subject Swap). $\Gamma \vdash D, C$ and $C \simeq_C C'$ imply $\Gamma \vdash D, C'$

Proof. By induction on the derivation of $C \simeq_C C'$. □

Finally, we prove Theorem 3 by proving the stronger result below.

Theorem 2 (Typing Soundness). *Let D, C be an annotated AC and (T2.1) $\Gamma \vdash D, C$ for some Γ :*

if (T2.2) $\beta \neq \tau$ and $D, C \xrightarrow{\beta} D', C'$ then (T2.3) $\Gamma \xrightarrow{\beta} \Gamma'$ and (T2.4) $\Gamma' \vdash D', C'$;

if (T2.5) $D, C \xrightarrow{\tau} D', C'$ then, for some Γ' , (T2.6) $\Gamma' \vdash D', C'$.

Proof. Proof by induction on the derivation of $D, C \xrightarrow{\beta} D', C'$.

Case $[^C]_{\text{SEND}}$

The case is:

$$\frac{\eta = k : p[A].e \rightarrow B.o_j \quad D, \eta \blacktriangleright D'}{D, \eta; C \xrightarrow{k: A \rightarrow B.o_j} D', C} [^C]_{\text{SEND}}$$

Where (T2.2) has $C' = C$ and $D' = D[q \mapsto (D(q).\text{st}, D(q).\text{que}[t_c \mapsto \tilde{m} :: (o_j, t_m)])]$ by Rule $[^D]_{\text{SEND}}$.

To prove (T2.3) we must prove Rule $[^\Gamma]_{\text{SEND}}$ to be applicable.

From (T2.1) we know that there exists a global type G for session k such that $\text{pco}(\Gamma)$ holds. We can partition $\Gamma = \Gamma_*, \Gamma_k$ such that $\Gamma_* = \Gamma / \llbracket G \rrbracket_k$ and $\Gamma_k = \Gamma / \Gamma_*$.

From (T2.1) we can write the derivation (with $\Gamma = \Gamma_1, k[A] : !B. \{o_i(U_i); \llbracket G_i \rrbracket_A\}_{i \in I}, j \in I$)

E.1. Proofs of Subject Reduction and Session Fidelity

$$\text{pco}(\Gamma) \quad \frac{\Gamma \vdash D \quad \frac{j \in I \quad \Gamma_1 \vdash \mathbf{p}: k[\mathbf{A}] \quad \Gamma_1 \vdash \mathbf{p}.e: U_j \quad \Gamma_1, k[\mathbf{A}]: \llbracket G_j \rrbracket_{\mathbf{A}} \vdash C}{\Gamma_1, k[\mathbf{A}]: !\mathbf{B}.\{o_i(U_i); \llbracket G_i \rrbracket_{\mathbf{A}}\}_{i \in I} \vdash k: \mathbf{p}[\mathbf{A}].e \rightarrow \mathbf{B}.o_j; C} [\mathbf{T}|_{\text{SEND}}]}{\Gamma \vdash D, k: \mathbf{p}[\mathbf{A}].e \rightarrow \mathbf{B}.o_j; C} [\mathbf{T}|_{\text{DC}}]}$$

Since $\Gamma \vdash k[\mathbf{A}]: !\mathbf{B}.\{o_i(U_i); T_i\}_{i \in I}$, G must be swap-equivalent to type $G_* = \mathbf{A} \rightarrow \mathbf{B}.\{o_i(U_i); G_i\}_{i \in I}$, where $\forall i \in I$, $\llbracket G_i \rrbracket_{\mathbf{B}} = T_i$. G_* reduces with Rule $[\mathbf{G}|_{\text{SEND}}]$

$$\frac{j \in I}{\mathbf{A} \rightarrow \mathbf{B}.\{o_i(U_i); G_i\} \xrightarrow{\mathbf{A} \rightarrow \mathbf{B}.o_j} \mathbf{A} \rightsquigarrow \mathbf{B}.o_j(U_j); G_j} [\mathbf{G}|_{\text{SEND}}]$$

and we can apply rule $[\mathbf{T}|_{\text{SEND}}]$ where $G'_* = \mathbf{A} \rightsquigarrow \mathbf{B}.o_j(U_j); G_j$ and $G' \simeq_{\mathbf{G}} G'_*$

$$\frac{\frac{G \simeq_{\mathbf{G}} G_* \quad G_* \xrightarrow{\mathbf{A} \rightarrow \mathbf{B}.o_j} G'_* \quad G'_* \simeq_{\mathbf{G}} G'}{G \xrightarrow{\mathbf{A} \rightarrow \mathbf{B}.o_j} G'} [\mathbf{G}|_{\text{SWAP}}]}{\frac{k \notin \Gamma_* \quad \Gamma_k \subseteq \llbracket G \rrbracket_k \quad \{k[\mathbf{A}]: T, k[\mathbf{B}]: T'\} \in \Gamma_k \quad j \in I}{\Gamma_*, \Gamma_k \xrightarrow{k: \mathbf{A} \rightarrow \mathbf{B}.o_j} \Gamma_*, \{k[\mathbf{C}]: \llbracket G' \rrbracket_{\mathbf{C}} \mid k[\mathbf{C}] \in \Gamma_k\}, \{b[k]_{\mathbf{D}}^{\mathbf{C}}: \llbracket G' \rrbracket_{\mathbf{C}}^{\mathbf{D}} \mid b[k]_{\mathbf{C}}^{\mathbf{D}} \in \Gamma_k\}} [\mathbf{T}|_{\text{SEND}}]}$$

Hence (T2.3) holds and $\Gamma' = \Gamma_*, \{k[\mathbf{C}]: \llbracket G' \rrbracket_{\mathbf{C}} \mid k[\mathbf{C}] \in \Gamma_k\}, \{b[k]_{\mathbf{C}}^{\mathbf{D}}: \llbracket G' \rrbracket_{\mathbf{C}}^{\mathbf{D}} \mid b[k]_{\mathbf{C}}^{\mathbf{D}} \in \Gamma_k\}$.

We now prove (T2.4) by proving that Rule $[\mathbf{T}|_{\text{DC}}]$ applies to $\Gamma' \vdash D', C'$.

$$\frac{\text{pco}(\Gamma') \quad \Gamma' \vdash C' \quad \Gamma' \vdash D'}{\Gamma' \vdash D', C'} [\mathbf{T}|_{\text{DC}}]$$

Hence we need to prove ① $\text{pco}(\Gamma')$, ② $\Gamma' \vdash C'$. and ③ $\Gamma' \vdash D'$

Proof of ①. For all sessions $k' \in \Gamma_*$, $\text{pco}(\Gamma')$ holds as $\text{pco}(\Gamma)$ holds by (T2.1). For session k , $\text{pco}(\Gamma')$ holds by construction. □

Proof of ②. From the derivation on $\Gamma \vdash D, k: \mathbf{p}[\mathbf{A}].e \rightarrow \mathbf{B}.o_j; C$ we know that $\Gamma_1, k[\mathbf{A}]: \llbracket G_j \rrbracket_{\mathbf{A}} \vdash C$. Let $\Gamma'' = \Gamma_1, k[\mathbf{A}]: \llbracket G_j \rrbracket_{\mathbf{A}}$ and $\Gamma'_k = \Gamma_1/\Gamma_* = \Gamma_k/\{k[\mathbf{A}]: \llbracket G \rrbracket_{\mathbf{A}}\}$. We can write $\Gamma'' = \Gamma_*, \Gamma'_k, k[\mathbf{A}]: \llbracket G_j \rrbracket_{\mathbf{A}}$. Notably the application of Rule $[\mathbf{T}|_{\text{SEND}}]$

does not modify the buffer types in Γ and therefore $\Gamma''(b[k]_{\mathbb{B}}^{\mathbb{A}}) \neq \Gamma'(b[k]_{\mathbb{B}}^{\mathbb{A}})$, however from Lemma 14 we know that we can omit to consider buffer types as they are irrelevant for the typing of choreographies. For all sessions $k' \neq k$ in Γ'' their local typings are the same in Γ' . For session k , the typing $\Gamma''(k[\mathbb{A}]) = \Gamma'(k[\mathbb{A}]) = \llbracket G_j \rrbracket_{\mathbb{A}}$. From Lemma 13, for all other $k[\mathbb{C}] \in \Gamma''$, $\mathbb{C} \neq \mathbb{A}$ it holds that $\Gamma''(k[\mathbb{C}]) = \llbracket G \rrbracket_{\mathbb{C}}$, $\Gamma'(k[\mathbb{C}]) = \llbracket G' \rrbracket_{\mathbb{C}}$, and $\llbracket G' \rrbracket_{\mathbb{C}} \prec \llbracket G \rrbracket_{\mathbb{C}}$. Therefore $\Gamma' \prec \Gamma''$ and ② holds by Lemma 12. \square

Proof of ③. To prove $\Gamma' \vdash D'$ we need to prove that the conditions of Definition 18 hold. (D18.1–D18.4) hold by the application of Rule $\llbracket \text{D} \rrbracket_{\text{SEND}}$, by construction of Γ' , and by (T2.1). (D18.5) holds for all sessions $k' \neq k$ by application of Rule $\llbracket \text{D} \rrbracket_{\text{SEND}}$ and the construction of Γ' . The same holds true for session k and any process $q: k[\mathbb{C}] \in \Gamma' \mid \mathbb{C} \neq \mathbb{B}$.

Finally, let $q: k[\mathbb{B}] \in \Gamma$, from $\llbracket \text{C} \rrbracket_{\text{SEND}}$ we know that

$$D'(q).\text{que}(\underline{\mathbf{k.A.B}}(D'(q).\text{st})) = \tilde{m} :: (o_j, t_m)$$

G does not contain any partial reception from A to B, hence \tilde{m} must be empty

$$D'(q).\text{que}(\underline{\mathbf{k.A.B}}(D'(q).\text{st})) = (o_j, t_m)$$

From the definition of type projection we have that $\Gamma'(b[k]_{\mathbb{B}}^{\mathbb{A}}) = ?\mathbb{A}.o_j(U_j)$.

From $\llbracket \text{T} \rrbracket_{\text{SEND}}$ we know that $p.e \vdash U_j$ and from $\llbracket \text{C} \rrbracket_{\text{SEND}}$ that $t_m = \text{eval}(e, D(p).\text{st})$, thus t_m has type U_j and $\text{bte}(\mathbb{A}, (o_j, U_j)) = ?\mathbb{A}.o_j(U_j)$. \square

Case $\llbracket \text{C} \rrbracket_{\text{RECV}}$

The case is:

$$\frac{j \in I \quad D, k : \mathbb{A} \rightarrow q[\mathbb{B}].o_j(x_j) \blacktriangleright D'}{D, k : \mathbb{A} \rightarrow q[\mathbb{B}].\{o_i(x_i); C_i\}_{i \in I} \xrightarrow{k:\mathbb{A} \rightsquigarrow \mathbb{B}.o_j(x_j)} D', C_j} \llbracket \text{C} \rrbracket_{\text{RECV}}$$

(T2.2) has $C' = C_j$ and $D' = D[q \mapsto (D(q).\text{st} \triangleleft (x, t_m), D(q).\text{que}[t_c \mapsto \tilde{m}])]$ by application of Rule $\llbracket \text{D} \rrbracket_{\text{RECV}}$.

To prove (T2.3) we must prove that Rule $\llbracket \text{T} \rrbracket_{\text{SEND}}$ is applicable.

Since D, C reduces with $\llbracket \text{C} \rrbracket_{\text{RECV}}$, the conditions of Rule $\llbracket \text{D} \rrbracket_{\text{RECV}}$ must hold, i.e., $t_c = \underline{\mathbf{k.A.B}}(D(q).\text{st}) \wedge D(q).\text{que}(t_c) = (o_j, t_m) :: \tilde{m}$. Since (T2.1) holds $\text{pco}(\Gamma)$ and $\Gamma \vdash D$ hold and therefore we know that, by (D18.5), $\Gamma(b[k]_{\mathbb{B}}^{\mathbb{A}}) = \text{bte}(\mathbb{A}, (o_j, t_m) :: \tilde{m})$.

E.1. Proofs of Subject Reduction and Session Fidelity

Let $\vdash t_m : U_j$, then $\text{bte}(\mathbf{A}, (o_j, t_m) :: \tilde{m}) = ?\mathbf{A}.o_j(U_j); T$ where $T = \text{bte}(\mathbf{A}, \tilde{m})$ by definition and $\Gamma(b[k]_{\mathbf{B}}^{\mathbf{A}}) = ?\mathbf{A}.o_j(U_j); T$. Since $\text{pco}(\Gamma)$ holds, there exists a global type G for session k such that G is swap-equivalent to a type $G_* = \mathbf{A} \rightsquigarrow \mathbf{B}.o_j(U_j); G_j$ where $\llbracket G_j \rrbracket_{\mathbf{B}}^{\mathbf{A}} = T$.

G_* reduces with Rule $[\mathbf{G}|_{\text{RECV}}]$

$$\frac{}{\mathbf{A} \rightsquigarrow \mathbf{B}.o(U); G \xrightarrow{\mathbf{A} \rightsquigarrow \mathbf{B}.o(U)} G} [\mathbf{G}|_{\text{RECV}}]$$

and we can apply Rule $[\Gamma|_{\text{RECV}}]$ with $G'_* = G_j$ and $G' \simeq_{\mathbf{G}} G'_*$

We partition $\Gamma = \Gamma_*, \Gamma_k$ such that $\Gamma_* = \Gamma / \llbracket G \rrbracket_k$ and $\Gamma_k = \Gamma / \Gamma_*$.

$$\frac{\frac{G \simeq_{\mathbf{G}} G_* \quad G_* \xrightarrow{\mathbf{A} \rightsquigarrow \mathbf{B}.o_j} G'_* \quad G'_* \simeq_{\mathbf{G}} G'}{G \xrightarrow{\mathbf{A} \rightsquigarrow \mathbf{B}.o_j} G'} [\mathbf{G}|_{\text{SWAP}}]}{\frac{k \notin \Gamma_* \quad \Gamma_k \subseteq \llbracket G \rrbracket_k \quad \{k[\mathbf{A}]: T, k[\mathbf{B}]: T'\} \in \Gamma_k \quad \Gamma_* \vdash \mathbf{q}: k[\mathbf{B}]}{\Gamma_*, \Gamma_k \xrightarrow{k:\mathbf{A} \rightsquigarrow \mathbf{B}.o_j(x)} \Gamma_*, \{k[\mathbf{C}]: \llbracket G' \rrbracket_{\mathbf{C}} \mid k[\mathbf{C}] \in \Gamma_k\}, \{b[k]_{\mathbf{C}}^{\mathbf{D}}: \llbracket G' \rrbracket_{\mathbf{C}}^{\mathbf{D}} \mid b[k]_{\mathbf{C}}^{\mathbf{D}} \in \Gamma_k\}, \mathbf{q}.x: U_j} [\Gamma|_{\text{RECV}}]}} [\Gamma|_{\text{RECV}}]$$

Hence (T2.3) holds and $\Gamma' = \Gamma_*, \{\llbracket G' \rrbracket_{\mathbf{C}} \mid k[\mathbf{C}] \in \Gamma_k\}, \mathbf{q}.x: U_j$.

(T2.4) holds if we can apply Rule $[\mathbf{T}|_{\text{DC}}]$ on $\Gamma' \vdash D', C'$

$$\frac{\text{pco}(\Gamma') \quad \Gamma' \vdash C' \quad \Gamma' \vdash D'}{\Gamma' \vdash D', C'} [\mathbf{T}|_{\text{DC}}]$$

and we need to prove ① $\text{pco}(\Gamma')$, ② $\Gamma' \vdash C'$. and ③ $\Gamma' \vdash D'$

The proof of ① for this case is similar to that of ① for case $[\mathbf{C}|_{\text{SEND}}]$.

Proof of ②. From (T2.1), partitioning $\Gamma = \Gamma_1, k[\mathbf{B}]: ?\mathbf{A}.o_j(U_j); \llbracket G_j \rrbracket_{\mathbf{B}}$ and since $j \in I$ from Rule $[\mathbf{C}|_{\text{RECV}}]$, we can write the derivation

$$\frac{\text{pco}(\Gamma) \quad \Gamma \vdash D \quad \frac{j \in I \quad \Gamma_1 \vdash \mathbf{q}: k[\mathbf{B}] \quad \Gamma_1, \mathbf{q}.x_j: U_j, k[\mathbf{B}]: \llbracket G_j \rrbracket_{\mathbf{B}} \vdash C_j}{\Gamma_1, k[\mathbf{B}]: ?\mathbf{A}.o_j(U_j); \llbracket G_j \rrbracket_{\mathbf{B}} \vdash k: \mathbf{A} \rightarrow \mathbf{q}[\mathbf{B}].\{o_i(x_i); C_i\}_{i \in I}} [\mathbf{T}|_{\text{RECV}}]}{\Gamma \vdash D, k: \mathbf{A} \rightarrow \mathbf{q}[\mathbf{B}].\{o_i(x_i); C_i\}_{i \in I}} [\mathbf{T}|_{\text{DC}}]$$

hence we know that $\Gamma_1, \mathbf{q}.x_j: U_j, k[\mathbf{B}]: \llbracket G_j \rrbracket_{\mathbf{B}} \vdash C_j$.

Let $\Gamma'' = \Gamma_1, \mathbf{q}.x_j: U_j, k[\mathbf{B}]: \llbracket G_j \rrbracket_{\mathbf{B}} \vdash C_j$ and $\Gamma'_k = \Gamma_1 / \Gamma_* = \Gamma_k / \{k[\mathbf{B}]: \llbracket G \rrbracket_{\mathbf{B}}\}$.

We can write $\Gamma'' = \Gamma_*, \Gamma'_k, k[\mathbf{B}]: \llbracket G_j \rrbracket_{\mathbf{B}}$. Similarly to the proof of ② for case

$\llbracket \text{C} \rrbracket_{\text{SEND}} \Gamma''(b[k]_{\text{B}}^{\text{A}}) \neq \Gamma'(b[k]_{\text{B}}^{\text{A}})$, but we omit to consider buffer types as they are irrelevant for the typing of choreographies by Lemma 14. For all sessions in Γ'' , their local typings are the same as in Γ' . We consider in particular k on which we applied the reduction for this case for which it holds

$$\forall k[\text{C}] \in \Gamma'', \Gamma''(k[\text{C}]) = \Gamma'(k[\text{C}]) = \llbracket G' \rrbracket_{\text{C}}$$

□

Proof of ③. To prove $\Gamma' \vdash D'$ we prove the conditions in Definition 18. (D18.1) holds from the application of Rule $\llbracket \text{D} \rrbracket_{\text{RECV}}$, (T2.1), and the construction of Γ' . (D18.2) holds for all $p.x$ from the application of Rule $\llbracket \text{D} \rrbracket_{\text{RECV}}$, (T2.1), and the construction of Γ' , except for $q.x_j$ which is not defined in Γ . However the condition holds by construction of $\Gamma' = \Gamma_1, q.x_j : U_j, k[\text{B}] : \llbracket G_j \rrbracket_{\text{B}}$. (D18.3–D18.4) hold by application of Rule $\llbracket \text{D} \rrbracket_{\text{RECV}}$ and (T2.1). (D18.5) holds for all sessions $k' \neq k$ by the application of Rule $\llbracket \text{D} \rrbracket_{\text{RECV}}$ and the construction of Γ' . The same holds true for session k and any process $p : k[\text{C}] \in \Gamma \mid \text{C} \neq \text{B}$.

For $q : k[\text{B}]$ and role A we know from the application of $\llbracket \text{C} \rrbracket_{\text{SEND}}$ that

$$D'(q).\text{que}(\mathbf{k.A.B}(D'(q).\text{st})) = \tilde{m}$$

Since we took G such that $\llbracket G \rrbracket_{\text{B}}^{\text{A}} = ?\text{A}.o_j(U_j); T$, where $T = \text{bte}(\text{A}, \tilde{m})$ then $\llbracket G' \rrbracket_{\text{B}}^{\text{A}} = T$.

□

Case $\llbracket \text{C} \rrbracket_{\text{COM}}$

The case is:

$$\frac{\eta = k : p[\text{A}].e \rightarrow q[\text{B}].o_j(x) \quad D, k : p[\text{A}].e \rightarrow \text{B}.o_j \blacktriangleright D'}{D, \eta; C \xrightarrow{k : \text{A} \rightarrow \text{B}.o_j} D', k : \text{A} \rightarrow q[\text{B}].o_j(x); C} \llbracket \text{C} \rrbracket_{\text{COM}}$$

Where (T2.2) has $C' = k : \text{A} \rightarrow q[\text{B}].o(x); C$ and

$D' = D[q \mapsto (D(q).\text{st}, D(q).\text{que}[t_c \mapsto \tilde{m} :: (o_j, t_m)])]$ by application of Rule $\llbracket \text{D} \rrbracket_{\text{SEND}}$.

To prove (T2.3) we must prove Rule $\llbracket \Gamma \rrbracket_{\text{SEND}}$ to be applicable.

Similarly to the proof of case $\llbracket \text{C} \rrbracket_{\text{SEND}}$ we know there exists a global type G for session k such that $\text{pco}(\Gamma)$ holds. Likewise, we partition $\Gamma = \Gamma_*, \Gamma_k$ wrt $\llbracket G \rrbracket_k$.

Partitioning $\Gamma = \Gamma_1, k[\text{A}] : !\text{B}.\{o_i(U_i); \llbracket G_i \rrbracket_{\text{A}}\}_{i \in I}, k[\text{B}] : ?\text{A}.\{o_i(U_i); \llbracket G_i \rrbracket_{\text{B}}\}_{i \in I}$ we can write the derivation

E.1. Proofs of Subject Reduction and Session Fidelity

$$\text{pco}(\Gamma) \quad \Gamma \vdash D \quad \frac{\begin{array}{c} j \in I \quad \Gamma_1 \vdash \mathbf{p}: k[\mathbf{A}], \mathbf{q}: k[\mathbf{B}] \quad \Gamma_1 \vdash \mathbf{p}.e: U_j \\ \Gamma_1, \mathbf{q}.x: U_j, k[\mathbf{A}]: \llbracket G_j \rrbracket_{\mathbf{A}}, k[\mathbf{B}]: \llbracket G_j \rrbracket_{\mathbf{B}} \vdash C \end{array}}{\Gamma_1, k[\mathbf{A}]: !\mathbf{B}.\{o_i(U_i); \llbracket G_i \rrbracket_{\mathbf{A}}\}_i, k[\mathbf{B}]: ?\mathbf{A}.\{o_i(U_i); \llbracket G_i \rrbracket_{\mathbf{B}}\}_i \vdash k: \mathbf{p}[\mathbf{A}].e \rightarrow \mathbf{q}[\mathbf{B}].o_j(x); C} \begin{array}{l} [\mathbf{T}]_{\text{COM}} \\ [\mathbf{T}]_{\text{DC}} \end{array}$$

Since $\Gamma \vdash k[\mathbf{A}]: !\mathbf{B}.\{o_i(U_i); \llbracket G_i \rrbracket_{\mathbf{A}}\}_i$ and $\Gamma \vdash k[\mathbf{B}]: ?\mathbf{A}.\{o_i(U_i); \llbracket G_i \rrbracket_{\mathbf{B}}\}_i$, G must be a swap-equivalent to type $G_* = \mathbf{A} \rightarrow \mathbf{B}.\{o_i(U_i); G_i\}_i$, which reduces along Rule $[\mathbf{G}]_{\text{SEND}}$

$$\frac{j \in I}{\mathbf{A} \rightarrow \mathbf{B}.\{o_i(U_i); G_i\} \rightarrow \mathbf{A} \rightsquigarrow \mathbf{B}.o_j(U_j); G_j} [\mathbf{G}]_{\text{SEND}}$$

where $G'_* = \mathbf{A} \rightsquigarrow \mathbf{B}.o_j(U_j); G_j$ and $G'_* \simeq_{\mathbf{G}} G'$. Therefore we can apply Rule $[\mathbf{T}]_{\text{SEND}}$

$$\frac{\frac{G \simeq_{\mathbf{G}} G_* \quad G_* \xrightarrow{\mathbf{A} \rightarrow \mathbf{B}.o_j} G'_* \quad G'_* \simeq_{\mathbf{G}} G'}{G \xrightarrow{\mathbf{A} \rightarrow \mathbf{B}.o_j} G'} [\mathbf{G}]_{\text{SWAP}}}{\frac{k \notin \Gamma_* \quad \Gamma_k \subseteq \llbracket G \rrbracket_k \quad \{k[\mathbf{A}]: T, k[\mathbf{B}]: T'\} \in \Gamma_k \quad j \in I}{\Gamma_*, \Gamma_k \xrightarrow{k: \mathbf{A} \rightarrow \mathbf{B}.o_j} \Gamma_*, \{k[\mathbf{C}]: \llbracket G' \rrbracket_{\mathbf{C}} \mid k[\mathbf{C}] \in \Gamma_k\}, \{b[k]_{\mathbf{C}}^{\mathbf{D}}: \llbracket G' \rrbracket_{\mathbf{C}}^{\mathbf{D}} \mid b[k]_{\mathbf{C}}^{\mathbf{D}} \in \Gamma_k\}} [\mathbf{T}]_{\text{SEND}}}$$

We prove that for $\Gamma' = \Gamma_*, \{\llbracket G' \rrbracket_{\mathbf{C}} \mid k[\mathbf{C}] \in \Gamma_k\}$ (T2.4) holds, i.e., that Rule $[\mathbf{T}]_{\text{DC}}$ is applicable on $\Gamma' \vdash D', C'$ and therefore that ① $\text{pco}(\Gamma')$, ② $\Gamma' \vdash C'$ and ③ $\Gamma' \vdash D'$ hold.

The proof of ① and ③ for this case is similar to those of, respectively, ① and ③ for case $[\mathbf{C}]_{\text{SEND}}$.

Proof of ②. Let $\Gamma'' = \Gamma_1, \mathbf{q}.x: U_j, k[\mathbf{A}]: \llbracket G_j \rrbracket_{\mathbf{A}}, k[\mathbf{B}]: \llbracket G_j \rrbracket_{\mathbf{B}}$, we know from the derivation on $\Gamma \vdash D, k: \mathbf{p}[\mathbf{A}].e \rightarrow \mathbf{q}[\mathbf{B}].o_j(x); C$ that $\Gamma'' \vdash C$.

We need to prove that $\Gamma' \vdash C'$ and therefore that Rule $[\mathbf{T}]_{\text{RECV}}$ applies. $\Gamma'(k[\mathbf{B}]) = \llbracket G' \rrbracket_{\mathbf{B}} = ?\mathbf{A}.o_j(U_j); \llbracket G_j \rrbracket_{\mathbf{B}}$ by construction and thus we can write $\Gamma' = \Gamma'_1, k[\mathbf{B}]: ?\mathbf{A}.o_j(U_j); \llbracket G_j \rrbracket_{\mathbf{B}}$ and apply Rule $[\mathbf{T}]_{\text{RECV}}$

$$\frac{\Gamma'_1 \vdash \mathbf{q}: k[\mathbf{B}] \quad \Gamma'_1, \mathbf{q}.x: U_j, k[\mathbf{B}]: \llbracket G_j \rrbracket_{\mathbf{B}} \vdash C}{\Gamma'_1, k[\mathbf{B}]: ?\mathbf{A}.o_j(U_j); \llbracket G_j \rrbracket_{\mathbf{B}} \vdash k: \mathbf{p}[\mathbf{A}].e \rightarrow \mathbf{q}[\mathbf{B}].o_j(x); C} [\mathbf{T}]_{\text{RECV}}$$

Again from the derivation of $\Gamma \vdash D, k : p[A].e \rightarrow q[B].o_j(x); C$ we know that

$$\Gamma_1, q.x : U_j, k[A] : \llbracket G_j \rrbracket_A, k[B] : \llbracket G_j \rrbracket_B \vdash C$$

Similarly to the proof of ② of case $[^C|_{\text{SEND}}]$, $\Gamma' \vdash C'$ since 1) we can omit to consider buffer types as of Lemma 14, 2) for all session $k' \neq k$ its local typings remain unchanged from Γ to Γ' , and 3) for session k , $\Gamma''(k[A]) = \Gamma'(k[A]) = \llbracket G_j \rrbracket_A$ and for all other local typings for $C \in \tilde{G}/\{A\}$, $\Gamma''(k[C]) = \llbracket G \rrbracket_C$, $\Gamma'(k[C]) = \llbracket G' \rrbracket_C$, and $\llbracket G' \rrbracket_C \prec \llbracket G \rrbracket_C$. Hence, $\Gamma' \prec \Gamma''$ and ② holds by Lemma 12. \square

Case $[^C|_{\text{START}}]$

The case is:

$$\frac{\# \tilde{r} \quad \# k' \quad p \in D(l) \quad \delta = \mathbf{start} \ k' : l.p[A], l.r[B] \quad D, \delta \blacktriangleright D'}{D, \mathbf{start} \ k : p[A] \Leftrightarrow l.q[B]; C \quad \xrightarrow{\tau} \quad D', C[k'/k][\tilde{r}/\tilde{q}]} [^C|_{\text{START}}]$$

Where (T2.5) has $C' = C[k'/k][\tilde{r}/\tilde{q}]$. D' is defined non-deterministically but abides the requirements defined in Rule $[^D|_{\text{START}}]$. Let $s[\tilde{C}] = p[A], r[B]$. Since (T2.1) holds, we can apply Rule $[^T|_{\text{START}}]$. We partition $\Gamma = \Gamma_1, \tilde{l} : G\langle A|\tilde{B}|\tilde{B} \rangle$

$$\frac{\Gamma_1, \tilde{l} : G\langle A|\tilde{B}|\tilde{B} \rangle, \text{init}(s[\tilde{C}], k, G) \vdash C \quad s[\tilde{C}] = p[A], q[B] \quad \tilde{q} \notin \Gamma_1}{\Gamma_1, \tilde{l} : G\langle A|\tilde{B}|\tilde{B} \rangle \vdash \mathbf{start} \ k : p[A] \Leftrightarrow l.q[B]; C} [^T|_{\text{START}}]$$

We take $\Gamma' = \Gamma, \text{init}(s[\tilde{C}], k', G)$ and we prove the case by proving that we can apply Rule $[^T|_{\text{DC}}]$ on $\Gamma' \vdash D', C'$, i.e, that the following hold: ① $\text{pco}(\Gamma')$, ② $\Gamma' \vdash C'$ and ③ $\Gamma' \vdash D'$.

Proof of ①. ① holds for all session $k'' \in \Gamma', k'' \neq k'$ by (T2.1). For session k' ① holds by construction. \square

Proof of ②. By (T2.1) we could apply $[^T|_{\text{START}}]$ where $\Gamma, \text{init}(s[\tilde{C}], k, G) \vdash C$. We apply α -renaming on k and \tilde{q} along what done in Rule $[^C|_{\text{START}}]$ and we obtain that $\Gamma, \text{init}(s[\tilde{C}], k, G)[k'/k][\tilde{r}/\tilde{q}] \vdash C[k'/k][\tilde{r}/\tilde{q}]$.

Since $\Gamma, \text{init}(s[\tilde{C}], k, G)[k'/k][\tilde{r}/\tilde{q}] = \Gamma, \text{init}(s[\tilde{C}], k', G)$ then ② holds. \square

E.1. Proofs of Subject Reduction and Session Fidelity

Proof of ③. To prove ③ we prove the conditions in Definition 18. (D18.1–D18.5) hold by the application of Rule $[^D]_{\text{START}}$ and the construction of Γ' . \square

Case $[^C]_{\text{PSTART}}$

The case is:

$$\frac{\begin{array}{l} i \in \{1, \dots, n\} \quad \#k' \quad \{\widetilde{l.B}\} = \biguplus_i \{\widetilde{l_i.B_i}\} \quad \#\tilde{r} \quad \{\tilde{r}\} = \bigcup_i \{\tilde{r}_i\} \\ p \in D(l) \quad \delta = \mathbf{start} \ k' : l.p[A], \widetilde{l_1.r_1[B_1]}, \dots, \widetilde{l_n.r_n[B_n]} \quad D, \delta \blacktriangleright D' \end{array}}{D, \mathbf{req} \ k : p[A] \Leftrightarrow \widetilde{l.B}; C \mid \prod_i (\mathbf{acc} \ k : \widetilde{l_i.q_i[B_i]}; C_i) \xrightarrow{\tau} D', C[k'/k] \mid \prod_i (C_i[k'/k][\tilde{r}_i/\tilde{q}_i]) \mid \prod_i (\mathbf{acc} \ k : \widetilde{l_i.q_i[B_i]}; C_i)} \quad [^C]_{\text{PSTART}}$$

Where (T2.5) has $C' = C[k'/k] \mid \prod_i (C_i[k'/k][\tilde{r}_i/\tilde{q}_i]) \mid \prod_i (\mathbf{acc} \ k : \widetilde{l_i.q_i[B_i]}; C_i)$. D' is defined non-deterministically but abides the requirements defined in Rule $[^D]_{\text{START}}$.

We partition Γ such that:

- $\Gamma = \Gamma_r \circ \Gamma_a$
- $\Gamma_r \vdash \tilde{l} : G\langle A | \tilde{B} | \emptyset \rangle$
- $\Gamma_a \vdash \tilde{l} : G\langle A | \tilde{B} | \tilde{B} \rangle$
- $\Gamma_a = \Gamma_1, \tilde{l} : G\langle A | \tilde{B} | \tilde{B}_1 \rangle \circ \dots \circ \Gamma_n, \tilde{l} : G\langle A | \tilde{B} | \tilde{B}_n \rangle$
- $\Gamma_a^i = \Gamma_i, \tilde{l} : G\langle A | \tilde{B} | \tilde{B}_i \rangle \circ \dots \circ \Gamma_n, \tilde{l} : G\langle A | \tilde{B} | \tilde{B}_n \rangle$

and we can write the derivation

$$\frac{\frac{\frac{\Gamma_r, p : k[A], k[A] : \llbracket G \rrbracket_A \vdash C \quad \Gamma_r \vdash \tilde{l} : G\langle A | \tilde{B} | \emptyset \rangle}{\Gamma_r \vdash \mathbf{req} \ k : p[A] \Leftrightarrow \widetilde{l.B}; C} \quad [^T]_{\text{REQ}} \quad \Delta_1}{\text{pco}(\Gamma) \quad \Gamma \vdash D \quad \Gamma \vdash \mathbf{req} \ k : p[A] \Leftrightarrow \widetilde{l.B}; C \mid \prod_{i \in I} (\mathbf{acc} \ k : \widetilde{l_i.q_i[B_i]}; C_i)} \quad [^T]_{\text{PAR}}}{\Gamma \vdash D, \mathbf{req} \ k : p[A] \Leftrightarrow \widetilde{l.B}; C \mid \prod_{i \in I} (\mathbf{acc} \ k : \widetilde{l_i.q_i[B_i]}; C_i)} \quad [^T]_{\text{DC}}$$

$$\Delta_i = \left\{ \begin{array}{l} \frac{\tilde{l}_i \subseteq \tilde{l} \quad \Gamma_i, \tilde{l} : G\langle A | \tilde{B} | \emptyset \rangle, \text{init}(\widetilde{q_i[B_i]}, k, G) \vdash C_i}{\Gamma_i, \tilde{l} : G\langle A | \tilde{B} | \tilde{B}_i \rangle \vdash \mathbf{acc} \ k : \widetilde{l_i.q_i[B_i]}; C_i} \quad [^T]_{\text{ACC}} \\ \frac{\Gamma_a^i \vdash \mathbf{acc} \ k : \widetilde{l_i.q_i[B_i]}; C_i \mid \prod_{j \in I/\{1, \dots, i\}} (\mathbf{acc} \ k : \widetilde{l_j.q_j[B_j]}; C_j)}{\Gamma_a^i \vdash \mathbf{acc} \ k : \widetilde{l_i.q_i[B_i]}; C_i} \quad \Delta_{i+1} \quad [^T]_{\text{PAR}} \end{array} \right.$$

Appendix E. Applied Choreographies: Proofs

Let $\widetilde{s[\mathbb{C}]} = \rho[\mathbb{A}], \widetilde{r_1[\mathbb{B}_1]}, \dots, \widetilde{r_n[\mathbb{B}_n]}$.

To prove (T2.6) we take

$$\Gamma' = \Gamma, \text{init}(\widetilde{s[\mathbb{C}]}, k', G) = \Gamma_r \circ \Gamma_a \circ \text{init}(\widetilde{s[\mathbb{C}]}, k', G)$$

By the definition of the merging operator \circ we can partition $\text{init}(\widetilde{s[\mathbb{C}]}, k', G)$ such that

$$\Gamma' = \Gamma'_r \circ \Gamma'_a \circ \Gamma_a$$

Where

- $\Gamma'_r = \Gamma_r, \text{init}(\rho[\mathbb{A}], k', G)$
- $\Gamma'_a = \Gamma'_1 \circ \dots \circ \Gamma'_n$
- $\Gamma'_i = \Gamma_i, \tilde{l}: G\langle \mathbb{A} | \tilde{\mathbb{B}} | \emptyset \rangle, \text{init}(\widetilde{r_i[\mathbb{B}_i]}, k', G)$

To prove (T2.6) we must prove we can apply Rule $[T]_{\text{DC}}$ on $\Gamma' \vdash D', C'$.

$$\frac{\begin{array}{l} \textcircled{1} \text{ pco}(\Gamma') \\ \textcircled{3} \Gamma' \vdash D' \end{array} \quad \textcircled{2} \left\{ \begin{array}{l} \textcircled{2a} \Gamma'_r \vdash C[k'/k] \\ \textcircled{2b} \Gamma'_a \vdash \prod_i (C_i[k'/k][\tilde{r}_i/\tilde{q}_i]) \\ \textcircled{2c} \Gamma_a \vdash \prod_i (\mathbf{acc} \ k : l_i \cdot \mathbf{q}_i[\mathbb{B}_i]; C_i) \end{array} \right.}{\Gamma' \vdash C[k'/k] \mid \prod_i (C_i[k'/k][\tilde{r}_i/\tilde{q}_i]) \mid \prod_i (\mathbf{acc} \ k : l_i \cdot \mathbf{q}_i[\mathbb{B}_i]; C_i)} \quad \begin{array}{l} \Gamma'_a \circ \Gamma_a \\ \vdash \prod_i (C_i[k'/k][\tilde{r}_i/\tilde{q}_i]) \mid \prod_i (\mathbf{acc} \ k : l_i \cdot \mathbf{q}_i[\mathbb{B}_i]; C_i) \\ \hline \Gamma' \vdash D', C' \end{array} \quad \begin{array}{l} [T]_{\text{PAR}} \\ [T]_{\text{PAR}} \\ [T]_{\text{DC}} \end{array}$$

$\textcircled{1}$ holds by by construction.

Proof of $\textcircled{2}$. $\textcircled{2}$ holds as

- $\textcircled{2a}$ holds by α -renaming $(\Gamma_r, \rho: k[\mathbb{A}], k[\mathbb{A}]: \llbracket G \rrbracket_{\mathbb{A}})[k'/k] \vdash C[k'/k]$ and by omitting to consider buffer types as of Lemma 14;
- similarly to $\textcircled{2a}$, $\textcircled{2b}$ holds by α -renaming on the derivation of

$$(\Gamma_i, \tilde{l}: G\langle \mathbb{A} | \tilde{\mathbb{B}} | \emptyset \rangle, \text{init}(\widetilde{q_i[\mathbb{B}_i]}, k, G))[k'/k][r_i/q_i] \vdash C_i[k'/k][r_i/q_i]$$

and by Lemma 14;

- $\textcircled{2c}$ holds by (T2.1).

E.1. Proofs of Subject Reduction and Session Fidelity

□

Proof of ③. The proof of ③ of this case is similar to the proof of ③ for Case $[^C|_{\text{START}}]$. □

Case $[^C|_{\text{COND}}]$

The case is:

$$\frac{i = 1 \text{ if } \text{eval}(e, D(\mathbf{p}).\text{st}) = \text{true}, i = 2 \text{ otherwise}}{D, \text{ if } \mathbf{p}.e \{C_1\} \text{ else } \{C_2\} \quad \xrightarrow{\tau} \quad D, C_i} \quad [^C|_{\text{COND}}]$$

In (T2.5) $D' = D$ and we have two cases for $C' = C_1$ or $C' = C_2$.

From (T2.1) we can write

$$\frac{\Gamma \vdash \mathbf{p}.e : \mathbf{bool} \quad \Gamma \vdash C_1 \quad \Gamma \vdash C_2}{\Gamma \vdash \text{ if } \mathbf{p}.e \{C_1\} \text{ else } \{C_2\}} \quad [^T|_{\text{COND}}]$$

The proof of (T2.6) follows directly from the premises of the typing derivation as $\Gamma \vdash D = D'$ and in both cases that $C' = C_1$ or $C' = C_2$ it holds that $\Gamma \vdash C'$ from the premises of $[^T|_{\text{COND}}]$.

Case $[^C|_{\text{CTX}}]$

The case is:

$$\frac{D, C_1 \xrightarrow{\beta} D', C'_1}{D, \text{ def } X = C_2 \text{ in } C_1 \quad \xrightarrow{\beta} \quad D', \text{ def } X = C_2 \text{ in } C'_1} \quad [^C|_{\text{CTX}}]$$

From (T2.1) we know that, $\Gamma = \Gamma_1, X : \Gamma_x$

$$\frac{\text{pco}(\Gamma) \quad \frac{\Gamma_1, X : \Gamma_x \vdash C_1 \quad \Gamma_x, X : \Gamma_x \vdash C_2 \quad \Gamma_x |_{\text{locs}} \subseteq \Gamma}{\Gamma \vdash \text{ def } X = C_2 \text{ in } C_1} \quad [^T|_{\text{DEF}}] \quad \Gamma \vdash D}{\Gamma \vdash D, \text{ def } X = C_2 \text{ in } C_1} \quad [^T|_{\text{DC}}]$$

The proof is divided in two cases on the type of β .

Case $\beta \neq \tau$

D, C_1 reduces on some session k . By the induction hypothesis since $\Gamma \vdash D, C_1$

we can find Γ' such that (T2.3) holds. We prove (T2.4) by proving that we can apply $[T]_{\text{DC}}$ on $\Gamma' \vdash D', \text{def } X = C_2$ in C'_1 and therefore that ① $\text{pco}(\Gamma')$ holds, ② $\Gamma' \vdash \text{def } X = C_2$ in C_1 and ③ $\Gamma' \vdash D'$.

① holds by the construction of Γ' and ③ holds by the induction hypothesis.

To prove ② we have to prove that $\Gamma' \vdash X : C_2$ and $\Gamma_x|_{\text{locs}} \subseteq \Gamma'$.

From the induction hypothesis we have that $\Gamma \xrightarrow{\beta} \Gamma'$ and $\Gamma' \vdash D', C'_1$. By construction of Γ' it holds that $\Gamma' = \Gamma'_*, \Gamma'_k$ where $\Gamma' \cap \Gamma = \Gamma_*$ such that $k \notin \Gamma_*$ and $\Gamma = \Gamma_*, \Gamma_k$ where $\Gamma_k \subseteq \llbracket G \rrbracket_k$ for some G . Therefore it holds that $\Gamma_* \vdash X : \Gamma_x$ and thus that $\Gamma' \vdash X : \Gamma_x$. The same applies to $\Gamma_x|_{\text{locs}} \subseteq \Gamma_*$ which proves $\Gamma_x|_{\text{locs}} \subseteq \Gamma'$.

Case $\beta = \tau$

from the induction hypothesis for any considered derivation we have $\Gamma \subseteq \Gamma'$.

We prove (T2.6) by proving that we can apply $[T]_{\text{DC}}$ on $\Gamma' \vdash D', \text{def } X = C_2$ in C'_1 .

①, ②, and ③ hold by construction of Γ' .

Case $[C]_{\text{PAR}}$

The case is:

$$\frac{D, C_1 \xrightarrow{\beta} D', C'_1}{D, C_1 | C_2 \xrightarrow{\beta} D', C'_1 | C_2} [C]_{\text{PAR}}$$

From (T2.1) we have the derivation below, with Γ partitioned as $\Gamma = \Gamma_1 \circ \Gamma_2$

$$\frac{\text{pco}(\Gamma) \quad \frac{\Gamma_1 \vdash C_1 \quad \Gamma_2 \vdash C_2}{\Gamma \vdash C_1 | C_2} [T]_{\text{PAR}} \quad \Gamma \vdash D}{\Gamma \vdash D, C_1 | C_2} [T]_{\text{DC}}$$

The proof is divided in two cases on the type of β .

Case $\beta \neq \tau$

D, C_1 reduces on some session k . By the induction hypothesis and since $\Gamma_1 \vdash D, C_1$ we can find Γ'_1 such that $\Gamma_1 \xrightarrow{\beta} \Gamma'_1$ and $\Gamma'_1 \vdash D', C'_1$. Then we take $\Gamma' = \Gamma'_1 \circ \Gamma_2$ which proves (T2.3) to hold. We prove (T2.4) by proving that we can apply $[T]_{\text{DC}}$ on $\Gamma' \vdash D', C'_1 | C_2$ and therefore that ① $\text{pco}(\Gamma')$, ② $\Gamma' \vdash C'_1 | C_2$ and ③ $\Gamma' \vdash D'$ hold. ①, ②, and ③ hold by construction and the induction hypothesis.

Case $\beta = \tau$

from the induction hypothesis, for any derivation we have that $\Gamma'_1 \vdash D', C'_1$ and $\Gamma_1 \subseteq \Gamma'_1$. Also in this case we take $\Gamma' = \Gamma'_1 \circ \Gamma_2$ and prove (T2.6) by proving that we can apply $[T]_{\text{DC}}$ on $\Gamma' \vdash D', C'_1 | C_2$. ①, ②, and ③ hold by construction of Γ' and the induction hypothesis.

Case $\llbracket^C|_{\text{EQ}}\rrbracket$

The case is:

$$\frac{\mathcal{R} \in \{\equiv, \simeq_C\} \quad C_1 \mathcal{R} C'_1 \quad D, C'_1 \xrightarrow{\beta} D', C'_2 \quad C'_2 \mathcal{R} C_2}{D, C_1 \xrightarrow{\beta} D', C_2} \llbracket^C|_{\text{EQ}}\rrbracket$$

The proof is divided into two subcases on the type of \mathcal{R} .

Case $\mathcal{R} = \equiv$

The case is proven by induction hypothesis and Lemma 16.

Case $\mathcal{R} = \simeq_C$

The case is proven by induction hypothesis and Lemma 17.

□

The proof of Theorem 4 follows directly from the proof of Theorem 2 and Lemma 15.

E.2 Proof of Deadlock Freedom

We report below the statement of Theorem 5 enriched with pointers for clearer referencing in the proof.

Theorem 5 (Deadlock-freedom) (D5.1) $\Gamma \vdash D, C$ and (D5.2) $\text{co}(\Gamma)$ imply that either (D5.3) $C \equiv \mathbf{0}$ or (D5.4) there exist D' and C' such that $D, C \rightarrow D', C'$.

Like in [13, 14], our choreographies enjoy deadlock-freedom provided that they *i*) do not contain free variable names and *ii*) are *well-sorted*, i.e., have no undefined procedure calls. Notably, well-sortedness is guaranteed by the type system.

Proof. Proof by induction on the structure of C .

Case $C \equiv \mathbf{0}$

trivial.

Case $C = k : p[A].e \rightarrow B.o; C_1$

from (D5.1) and (D5.2) we know that the requirements of $\llbracket^D|_{\text{SEND}}\rrbracket$ hold and we can find D' such that $D, k : p[A].e \rightarrow B.o \blacktriangleright D'$. We can apply Rule $\llbracket^C|_{\text{SEND}}\rrbracket$ for which $C' = C_1$.

Case $C = k : p[A].e \rightarrow q[B].o(x); C_1$

since (D5.1) holds both receiver and sender are typed by Γ . Similarly to the previous case, we apply Rule $\llbracket^C|_{\text{COM}}\rrbracket$ for which $C' = k : A \rightarrow q[B].o(x)$.

Case $C = k : A \rightarrow q[B].\{o_i(x_i); C_i\}_{i \in I}$

from (D5.1) and (D5.2) we know that the requirements of Rule $[^D|_{\text{RECV}}]$ hold and $D(q).\text{que}(\underline{k.A.B}(D(q).\text{st})) = (o_j, t_m) :: \tilde{m}$ for some $j \in I$. We can find D' such that $D, k : A \rightarrow q[B].o_j(x_j) \blacktriangleright D'$ and apply Rule $[^C|_{\text{RECV}}]$ for which $C' = C_j$.

Case $C = \text{start } k : p[A] \Leftrightarrow \widetilde{l.q[B]}; C_1$

from (D5.1) and (D5.2) $[^D|_{\text{START}}]$ applies and we can find D' such that $D, \text{start } k' : l.p[A], \widetilde{l.r[B]} \blacktriangleright D'$ for some k', \tilde{r} fresh. We can apply Rule $[^C|_{\text{START}}]$ for which $C' = C_1[k'/k][\tilde{r}/\tilde{q}]$.

Case $C = \text{req } k : p[A] \Leftrightarrow \widetilde{l.B}; C \mid \prod_{i=1}^n (\text{acc } k : \widetilde{l_i.q_i[B_i]}; C_i)$

similarly to the previous case the requirements of $[^D|_{\text{START}}]$ hold and we can find D' such that $D, \text{start } k' : l.p[A], \widetilde{l_1.r_1[B_1]}, \dots, \widetilde{l_n.r_n[B_n]} \blacktriangleright D'$ for some k' and $\tilde{r}_1, \dots, \tilde{r}_n$ fresh. We can apply Rule $[^C|_{\text{PSTART}}]$ for which $C' = C[k'/k] \mid \prod_{i=1}^n C_i[k'/k][\tilde{r}_i/\tilde{q}_i] \mid \prod_{i=1}^n (\text{acc } k : \widetilde{l_i.q_i[B_i]}; C_i)$.

Case $C = C_1 \mid C_2$

we can apply the induction hypothesis and Rule $[^C|_{\text{PAR}}]$ such that $D, C_1 \rightarrow D_1, C'_1$ and in (D5.4) $D' = D_1$ and $C' = C'_1 \mid C_2$.

Case $C = \text{def } X = C_2 \text{ in } C_1$

applies the induction hypothesis and Rule $[^C|_{\text{CTX}}]$ for which $D, C_1 \rightarrow D_1, C_1$, $D' = D_1$ and $C' = \text{def } X = C_2 \text{ in } C'_1$.

Case $\text{def } X = C_2 \text{ in } X; C_1$

applies Rule $[^C|_{\text{EQ}}]$ for $\text{def } X = C_2 \text{ in } X; C_1 \equiv \text{def } X = C_2 \text{ in } C_2; C_1$ and by the induction hypothesis $D, C_2 \rightarrow D_2, C'_2$ and therefore $D' = D_2$ and $C' = \text{def } X = C_2 \text{ in } C'_2; C_1$.

Case $C = \text{if } p.e \{C_1\} \text{ else } \{C_2\}$

from (D5.1) we know that $\Gamma \vdash p.e : \text{bool}$ and therefore we can apply Rule $[^C|_{\text{COND}}]$ and, according to the evaluation of e we have $C' = C_1$ or $C' = C_2$.

□

E.3 Proof of Endpoint Projection

To prove our result on the Endpoint Projection we first define the minimal typing system \vdash_{min} for AC.

E.3.1 Minimal Typing

We take the definition of Local and Typing Environment Subtyping from Section E.1.1 and we define the subtyping for global types $G \prec G'$.

Definition 34 (Global and Local Subtyping). *$G \prec G'$ is the smallest relation over closed and unfolded global types satisfying the rules below*

$$\frac{I \subseteq J \quad \forall i \in I, G_i \prec G'_i \wedge U_i \prec U'_i}{A \rightarrow B.\{o_i(U_i); G_i\}_{i \in I} \prec A \rightarrow B.\{o_j(U'_j); G'_j\}_{j \in J}} \text{[SubG]}_{\text{COM}}$$

$$\frac{U \prec U' \quad G \prec G'}{A \rightsquigarrow B.o(U); G \prec A \rightsquigarrow B.o(U'); G'} \text{[SubG]}_{\text{RECV}}$$

$$\frac{G'' \prec G' \quad (G'' \approx G \vee G'' \simeq_G G')}{G \prec G'} \text{[SubG]}_{\text{EQ}} \quad \frac{\text{end} \approx G}{\text{end} \prec G} \text{[SubG]}_{\text{END}}$$

We extend the subtyping relations for local and global types to set inclusion and point-wise to *i*) the typing of services (i.e., of kind $\tilde{l} : G \langle A | \tilde{B} | \tilde{C} \rangle$) and *ii*) the typing of sessions, respectively. Given two types G and G' , we denote their least upper bound wrt \prec with $G \nabla G'$ (the same for local types and typing environments). Our definition of subtyping for global types follows [14].

We define the minimal typing system \vdash_{min} on this notion of subtyping. The minimal typing uses the minimal global and local types for typing sessions and services such that the projection of the choreography is still typable. We report the rules for minimal typing in Figure E.2.

E.3.2 Typing Projection

Like in [14] our EPP projects recursive definitions of the same procedure from the point of view of different processes, therefore the same definition has different types according to the process on which it has been projected. Similarly to [14], to deal with this discrepancy we define the projection of environments. Note that we give the definition of $\llbracket \Gamma \rrbracket$ over annotated ACs, i.e., ACs where any procedure X is annotated, e.g., $X \langle \tilde{p} \rangle$, where \tilde{p} are the process identifiers of the active processes in the body of X . Formally, given a choreography C it is always possible to have its annotated version with operator $@C$, defined following the rules in Figure D.5.

Definition 35 (Typing Projection). *Let $\Gamma = \Gamma', \Gamma_{\text{def}}$ where $\nexists X : \Gamma^* \in \Gamma'$ and $\Gamma_{\text{def}} = \Gamma / \Gamma'$ (Γ_{def} contains only typings of recursive procedures). The projection of Γ , written $\llbracket \Gamma \rrbracket$, is defined as:*

$$\llbracket \Gamma \rrbracket = \Gamma', \left\{ X \langle \tilde{p} \rangle : \llbracket \Gamma_x \rrbracket_{\tilde{p}} \mid X \langle \tilde{r} \rangle : \Gamma_x \in \Gamma \wedge \tilde{p} \in \{\tilde{r}\} \right\}$$

$$\begin{array}{c}
 \frac{\Gamma, \text{init}(r[\widetilde{C}], k, G) \vdash_{\min} C \quad r[\widetilde{C}] = p[A], q[\widetilde{B}] \quad \tilde{q} \notin \Gamma}{\Gamma, \tilde{l}: G\langle A|\widetilde{B}|\widetilde{B} \rangle \vdash_{\min} \mathbf{start} \ k : p[A] \Leftrightarrow \tilde{l}.q[\widetilde{B}]; C} [\text{Min}|_{\text{START1}}] \\
 \\
 \frac{\Gamma, \tilde{l}: G\langle A|\widetilde{B}|\widetilde{B} \rangle, \text{init}(r[\widetilde{C}], k, G') \vdash_{\min} C \quad r[\widetilde{C}] = p[A], q[\widetilde{B}] \quad \tilde{q} \notin \Gamma}{\Gamma, \tilde{l}: G\nabla G'\langle A|\widetilde{B}|\widetilde{B} \rangle \vdash_{\min} \mathbf{start} \ k : p[A] \Leftrightarrow \tilde{l}.q[\widetilde{B}]; C} [\text{Min}|_{\text{START2}}] \\
 \\
 \frac{\Gamma, p : k[A], k[A] : \llbracket G \rrbracket_A \vdash_{\min} C \quad \Gamma \vdash \tilde{l} : G\langle A|\widetilde{B}|\emptyset \rangle}{\Gamma \vdash_{\min} \mathbf{req} \ k : p[A] \Leftrightarrow \tilde{l}.B; C} [\text{Min}|_{\text{REQ}}] \\
 \\
 \frac{\tilde{l} \subseteq \tilde{l}' \quad \Gamma, \tilde{l}': G\langle A|\widetilde{B}|\emptyset \rangle, \Gamma' \vdash_{\min} C \quad \Gamma' \prec \text{init}(q[\widetilde{C}], k, G)}{\Gamma, \tilde{l}': G\langle A|\widetilde{B}|\widetilde{C} \rangle \vdash_{\min} \mathbf{acc} \ k : \tilde{l}.q[\widetilde{C}]; C} [\text{Min}|_{\text{ACC}}] \\
 \\
 \frac{\Gamma_1 \nabla \Gamma_2 \vdash p.e : \mathbf{bool} \quad \Gamma_1 \vdash_{\min} C_1 \quad \Gamma_2 \vdash_{\min} C_2}{\Gamma_1 \nabla \Gamma_2 \vdash_{\min} \mathbf{if} \ p.e \ \{C_1\} \ \mathbf{else} \ \{C_2\}} [\text{Min}|_{\text{COND}}] \\
 \\
 \frac{\Gamma \vdash p : k[A], q : k[B] \quad \Gamma \vdash p.e : U \quad \Gamma, q.x : U, k[A] : T, k[B] : T' \vdash_{\min} C}{\Gamma, k[A] : !B.\{o(U); T\}, k[B] : ?A.\{o(U); T'\} \vdash_{\min} k : p[A].e \rightarrow q[B].o(x); C} [\text{Min}|_{\text{COM}}] \\
 \\
 \frac{j \in I \quad \Gamma \vdash p : k[A] \quad q : k[B] \notin \Gamma \quad \Gamma \vdash p.e : U_j \quad \Gamma, k[A] : T_j \vdash_{\min} C}{\Gamma, k[A] : !B.\{o_i(U_i); T_i\}_{i \in I} \vdash_{\min} k : p[A].e \rightarrow B.o_j; C} [\text{Min}|_{\text{SEND}}] \\
 \\
 \frac{\Gamma \vdash q : k[B] \quad p : k[A] \notin \Gamma \quad \forall i \in I. \Gamma, q.x_i : U_i, k[B] : T_i \vdash C_i}{\Gamma, k[B] : ?A.\{o_i(U_i); T_i\}_{i \in I} \vdash_{\min} k : A \rightarrow q[B].\{o_i(x_i); C_i\}_{i \in I}} [\text{Min}|_{\text{RECV}}] \\
 \\
 \frac{\Gamma_1 \vdash_{\min} C_1 \quad \Gamma_2 \vdash_{\min} C_2}{\Gamma_1 \circ \Gamma_2 \vdash_{\min} C_1 \mid C_2} [\text{T}|_{\text{PAR}}] \quad \frac{\text{end}(\Gamma)}{\Gamma \vdash_{\min} \mathbf{0}} [\text{T}|_{\text{END}}] \\
 \\
 \frac{\Gamma, X : \Gamma' \vdash_{\min} C \quad \Gamma', X : \Gamma' \vdash_{\min} C' \quad \Gamma' |_{\text{locs}} \subseteq \Gamma}{\Gamma \nabla \Gamma' \vdash_{\min} \mathbf{def} \ X = C' \ \mathbf{in} \ C} [\text{T}|_{\text{DEF}}] \\
 \\
 \frac{\text{end}(\Gamma) \quad \Gamma'' \subseteq \Gamma'}{\Gamma, \Gamma'', X : \Gamma' \vdash_{\min} X} [\text{T}|_{\text{CALL}}] \quad \frac{\text{pco}(\Gamma) \quad \Gamma \vdash_{\min} D \quad \Gamma \vdash_{\min} C}{\Gamma \vdash_{\min} D, C} [\text{Min}|_{\text{DC}}]
 \end{array}$$

Figure E.2: Choreography Calculus — Minimal typing rules

E.3. Proof of Endpoint Projection

where

$$\begin{aligned} \llbracket \Gamma \rrbracket_{\mathfrak{p}} = & \{ \mathfrak{p}.x : U \mid \mathfrak{p}.x : U \in \Gamma \} \cup \\ & \{ \mathfrak{p} : k[A], k[A] : T \mid \{ \mathfrak{p} : k[A], k[A] : T \} \subseteq \Gamma \} \cup \\ & \{ X \langle \mathfrak{p} \rangle : \llbracket \Gamma' \rrbracket_{\mathfrak{p}} \mid X \langle \tilde{r} \rangle : \Gamma' \in \Gamma \wedge \mathfrak{p} \in \{ \tilde{r} \} \} \cup \\ & \{ l : G \langle A \mid \tilde{B} \mid \tilde{C} \rangle \mid l : G \langle A \mid \tilde{B} \mid \tilde{C} \rangle \in \Gamma \} \end{aligned}$$

Notably, in the definition of $\llbracket \Gamma \rrbracket_{\mathfrak{p}}$ we omit to include in the typing of procedures *buffer types* ($b[k]_{\mathfrak{B}}^{\mathfrak{A}}$), *process locations* ($\mathfrak{p}@l$), since they are not used in the typing of C . Contrarily, we include all service typings present in the typings of procedures, although they might not be minimal wrt the process that is object of the projection.

E.3.3 Proof of Theorem 6

We define some auxiliary lemmas used in the proof of Theorem 6.

Lemma 18 (Weakening). *Let $\Gamma \vdash_{\min} C$ and \tilde{l} be a service name such that $\tilde{l} \notin \Gamma$ then $\Gamma, \tilde{l} : G \langle A \mid \tilde{B} \mid \tilde{C} \rangle \vdash C$ for any $G, A, \tilde{B}, \tilde{C}$.*

Proof. Immediate since $\tilde{l} : G \langle A \mid \tilde{B} \mid \tilde{C} \rangle$ is never used in C . □

Lemma 19 (Composability of Typing Projections). *Let $\Gamma \circ \Gamma' = \Gamma''$ then $\llbracket \Gamma \rrbracket \circ \llbracket \Gamma' \rrbracket = \llbracket \Gamma'' \rrbracket$.*

Proof. Directly from the fact that since the projection $\llbracket \Gamma \rrbracket$ returns exactly Γ except for the projection of the typings of the procedures. Since $\Gamma \circ \Gamma'$ is applicable, all procedure typings in common are equal and the same holds for their projections. □

We proceed defining well-annotated choreographies. Well annotated choreographies are ACs in which annotated procedures report the names of the active processes in their bodies.

Definition 36 (Well-annotated Choreographies). *Let C be an annotated choreography, C is well-annotated iff*

$$\forall X \langle \tilde{\mathfrak{p}} \rangle \text{ s.t } C \equiv \text{def } X \langle \tilde{\mathfrak{p}} \rangle = C'' \text{ in } C', \tilde{\mathfrak{p}} = \text{fp}(C'')$$

We prove Lemma 20 that states that given a well-annotated choreography C and a typing environment Γ for which $\Gamma \vdash_{\min} C$ then the projection of Γ , $\llbracket \Gamma \rrbracket$ types minimally the projection of C , $\llbracket C \rrbracket$.

Lemma 20 (Choreography EPP Typing Preservation). *Let C be a well-annotated choreography and let $\Gamma \vdash_{\min} C$ then $\llbracket \Gamma \rrbracket \vdash_{\min} \llbracket C \rrbracket$.*

Proof. Like for the proof of Theorem 5, we assume our choreographies to be well-sorted. The proof is by induction on the typing derivation of $\Gamma \vdash_{\min} C$.

Case $[\text{Min}|_{\text{START1}}]$

From the premises we have $C = \mathbf{start} \ k : p[A] \Leftrightarrow l.\widetilde{q[B]}; C'$. We can partition $\Gamma = \tilde{l} : G\langle A|\tilde{B}|\tilde{B}\rangle, \Gamma'$ and we can write the derivation

$$\frac{\Gamma', \text{init}(r[\widetilde{C}], k, G) \vdash_{\min} C' \quad r[\widetilde{C}] = p[A], q[\widetilde{B}] \quad \tilde{q} \notin \Gamma'}{\Gamma', \tilde{l} : G\langle A|\tilde{B}|\tilde{B}\rangle \vdash_{\min} \mathbf{start} \ k : p[A] \Leftrightarrow l.\widetilde{q[B]}; C'} \quad [\text{Min}|_{\text{START1}}]$$

Let $\widetilde{l.q[B]} = l_1.q_1.[B_1], \dots, l_n.q_n.[B_n]$.

Let $\Gamma_c = \Gamma', \text{init}(r[\widetilde{C}], k, G)$, from the induction hypothesis we have that $\Gamma_c \vdash_{\min} C'$ and therefore $\llbracket \Gamma_c \rrbracket \vdash_{\min} \llbracket C' \rrbracket$.

By its definition $\llbracket C' \rrbracket \equiv C'_s \mid C''$ where

$$C'_s = \llbracket C' \rrbracket_p \mid \llbracket C' \rrbracket_{q_1} \mid \dots \mid \llbracket C' \rrbracket_{q_n}$$

and

$$C'' = \prod_{r \in \text{fp}(C')/\{p, \tilde{q}\}} \llbracket C' \rrbracket_r \mid \prod_l \left(\bigsqcup_{s \in \llbracket C' \rrbracket_l} \llbracket C' \rrbracket_s \right)$$

We partition $\llbracket \Gamma_c \rrbracket$ as

$$\llbracket \Gamma_c \rrbracket = \Gamma'_p \circ \Gamma_{\tilde{q}} \circ \Gamma''$$

where

$$\Gamma'_p = \Gamma_p, p : k[A], k[A] : \llbracket G \rrbracket_p$$

and

$$\Gamma_{\tilde{q}} = \Gamma'_{q_1} \circ \dots \circ \Gamma'_{q_n}$$

where

$$\Gamma'_{q_i} = \Gamma_{q_i}, q_i : k[A], k[A] : \llbracket G \rrbracket_{q_i}$$

. We such that we can write the derivation

$$\frac{\Gamma'' \vdash_{\min} C'' \quad \frac{\Gamma'_{q_1} \vdash_{\min} \llbracket C' \rrbracket_{q_1} \quad \frac{\Gamma'_{q_2} \circ \dots \circ \Gamma'_{q_n} \vdash_{\min} \llbracket C' \rrbracket_{q_2} \mid \dots \mid \llbracket C' \rrbracket_{q_n}}{\Gamma'_{q_1} \circ \dots \circ \Gamma'_{q_n} \vdash_{\min} \llbracket C' \rrbracket_{q_1} \mid \dots \mid \llbracket C' \rrbracket_{q_n}} \quad [\text{Min}|_{\text{PAR}}]}{\Gamma'_p \circ \Gamma'_{q_1} \circ \dots \circ \Gamma'_{q_n} \vdash_{\min} \llbracket C' \rrbracket_p \mid \llbracket C' \rrbracket_{q_1} \mid \dots \mid \llbracket C' \rrbracket_{q_n}} \quad [\text{Min}|_{\text{PAR}}]}{\Gamma'' \circ \Gamma'_p \circ \Gamma_{\tilde{q}} \vdash_{\min} C'' \mid C'_s} \quad [\text{Min}|_{\text{PAR}}]$$

E.3. Proof of Endpoint Projection

Since the ownership and session typings for k in Γ_c belong to $\text{init}(r[\widetilde{C}], k, G)^1$ we have that $\Gamma'' \circ \Gamma_p \circ \Gamma_{q_1} \circ \dots \circ \Gamma_{q_n} = \llbracket \Gamma \rrbracket$ by Lemma 19.

Therefore we can partition $\llbracket \Gamma \rrbracket$ as (using Lemma 18 for Γ_p)

$$\llbracket \Gamma \rrbracket = \Gamma'' \circ \Gamma_p, \tilde{l}: G\langle A|\tilde{B}|\emptyset \rangle \circ \Gamma_{q_1}, \tilde{l}: G\langle A|\tilde{B}|B_1 \rangle \circ \dots \circ \Gamma_{q_n} \circ \tilde{l}: G\langle A|\tilde{B}|B_n \rangle$$

Let $\widetilde{l.q[B]}_i = \{l_i.q_i[B_i], \dots, l_n.q_n[B_n]\}$.

We prove the case by proving the typing derivation for $\llbracket \Gamma \rrbracket \vdash_{\min} \llbracket C \rrbracket$

From the definition of EPP (Definition 22) we can write

$$\llbracket C \rrbracket \equiv C_s \mid C''$$

where

$$C_s = \mathbf{req} k : p[A] \Leftrightarrow \widetilde{l.B}; \llbracket C' \rrbracket_p \quad | \quad \prod_{l.r[C] \in \{\widetilde{l.q[B]}\}} \mathbf{acc} k : l.r[C]; \llbracket C' \rrbracket_r$$

and

$$C'' = \prod_{r \in \text{fp}(C'')/\{p\}} \llbracket C' \rrbracket_r \quad | \quad \prod_{l' \neq l} \left(\bigsqcup_{s \in \llbracket C' \rrbracket_{l'}} \llbracket C' \rrbracket_s \right)$$

We now prove we can derive the typing of $\llbracket \Gamma \rrbracket \vdash_{\min} \llbracket C \rrbracket$

$$\frac{\frac{\Gamma_p, p: k[A], k[A]: \llbracket G \rrbracket_A \vdash_{\min} \llbracket C' \rrbracket_p}{\Gamma_p, \tilde{l}: G\langle A|\tilde{B}|\emptyset \rangle \vdash \tilde{l}: G\langle A|\tilde{B}|\emptyset \rangle} \quad [\text{Min}|_{\text{REQ}}] \quad \Delta_1}{\Gamma_p, \tilde{l}: G\langle A|\tilde{B}|\emptyset \rangle \vdash_{\min} \mathbf{req} k : p[A] \Leftrightarrow \widetilde{l.B}; \llbracket C' \rrbracket_p} \quad [\text{Min}|_{\text{PAR}}]}{\Gamma'' \vdash_{\min} C'' \quad \Gamma_p, \tilde{l}: G\langle A|\tilde{B}|\emptyset \rangle \circ \Gamma_{q_1}, \tilde{l}: G\langle A|\tilde{B}|B_1 \rangle \circ \dots \circ \Gamma_{q_n} \circ \tilde{l}: G\langle A|\tilde{B}|B_n \rangle \vdash_{\min} C_s} \quad [\text{Min}|_{\text{PAR}}]}{\Gamma'' \circ \Gamma_p, \tilde{l}: G\langle A|\tilde{B}|\emptyset \rangle \circ \Gamma_{q_1}, \tilde{l}: G\langle A|\tilde{B}|B_1 \rangle \circ \dots \circ \Gamma_{q_n} \circ \tilde{l}: G\langle A|\tilde{B}|B_n \rangle \vdash_{\min} C_s \mid C''} \quad [\text{Min}|_{\text{PAR}}]}$$

where

$$\Delta_i = \frac{\frac{l_i \subseteq \tilde{l} \quad \Gamma_{q_i}, \tilde{l}: G\langle A|\tilde{B}|\emptyset \rangle \circ \Gamma''_{q_i} \vdash_{\min} \llbracket C' \rrbracket_{q_i}}{\Gamma''_{q_i} \prec \text{init}(q_i[B_i], k, G) \quad q_i \notin \Gamma_{q_i}} \quad [\text{Min}|_{\text{ACC}}]}{\Delta_{i+1} \quad \Gamma_{q_i}, \tilde{l}: G\langle A|\tilde{B}|B_i \rangle \vdash_{\min} \mathbf{acc} k : l_i.q_i[B_i]; \llbracket C' \rrbracket_{q_i}} \quad [\text{Min}|_{\text{PAR}}]}{\Gamma_{q_i}, \tilde{l}: G\langle A|\tilde{B}|B_i \rangle \circ \dots \circ \Gamma_{q_n} \circ \tilde{l}: G\langle A|\tilde{B}|B_n \rangle \vdash_{\min} \mathbf{acc} k : l_i.q_i[B_i]; \llbracket C' \rrbracket_{q_i} \quad | \quad \prod_{l.r[C] \in \widetilde{l.q[B]}_{i+1}} \mathbf{acc} k : l.r[C]; \llbracket C' \rrbracket_r} \quad [\text{Min}|_{\text{PAR}}]}$$

In particular we prove

¹we omit buffer types as of Lemma 14.

Appendix E. Applied Choreographies: Proofs

- ① $\Gamma'' \vdash_{\min} C''$ and ② $\Gamma_p, p: k[A], k[A]: \llbracket G \rrbracket_A \vdash_{\min} \llbracket C'' \rrbracket_p$ which hold by the induction hypothesis;
- ③ $\Gamma_{q_i}, \tilde{l}: G\langle A | \tilde{B} | \emptyset \rangle \circ \Gamma''_{q_i} \vdash_{\min} \llbracket C'' \rrbracket_{q_i}$ which holds by the induction hypothesis and Lemma 18;
- ④ $\Gamma''_{q_i} \prec \text{init}(q_i[B_i], k, G)$ holds as of Definition 33 and because, for $\Gamma'_{q_i}, i \in \{1, \dots, n\}$, $\text{init}(q_i[B_i], k, G) \subseteq \Gamma'_{q_i}$.

Case $[\text{Min}|_{\text{START2}}]$
 Similar to case $[\text{Min}|_{\text{START1}}]$.

Case $[\text{Min}|_{\text{REQ}}]$
 Follows the proof of ② of case $[\text{Min}|_{\text{START1}}]$.

Case $[\text{Min}|_{\text{ACC}}]$
 Follows the proof of ③ of case $[\text{Min}|_{\text{START1}}]$.

Case $[\text{Min}|_{\text{COND}}]$
 By induction hypothesis on C_1 or C_2 .

Case $[\text{Min}|_{\text{COM}}]$
 From the premises we have $C = k: p[A].e \rightarrow q[B].o(x); C'$ on which we can apply the typing derivation

$$\frac{\Gamma' \vdash p: k[A], q: k[B] \quad \Gamma' \vdash p.e: U \quad \Gamma', q.x: U, k[A]: T, k[B]: T' \vdash_{\min} C'}{\Gamma', k[A]: !B.o(U); T, k[B]: ?A.o(U); T' \vdash_{\min} k: p[A].e \rightarrow q[B].o(x); C'} \quad [\text{Min}|_{\text{COM}}]$$

From the definition of EPP (Definition 22) we have $\llbracket C \rrbracket \equiv C_c \mid C''$ where

$$C_c = k: p[A].e \rightarrow B.o; \llbracket C'' \rrbracket_p \mid k: A \rightarrow q[B].o(x); \llbracket C'' \rrbracket_q$$

$$C'' = \prod_{r \in \{\text{fp}(C')/\{p,q\}\}} \llbracket C'' \rrbracket_r \mid \prod_l \left(\bigsqcup_{s \in \llbracket C'' \rrbracket_l} \llbracket C'' \rrbracket_s \right)$$

By the definition of $\llbracket \Gamma \rrbracket$ we can write

$$\llbracket \Gamma \rrbracket = \llbracket \Gamma' \rrbracket, k[A]: !B.o(U); T, k[A]: ?A.o(U); T'$$

from the induction hypothesis we have that, let $\Gamma_c = \Gamma', q.x: U, k[A]: T, k[B]: T'$, $\Gamma_c \vdash_{\min} C'$ and therefore $\llbracket \Gamma_c \rrbracket \vdash_{\min} \llbracket C' \rrbracket$. We can partition $\llbracket \Gamma_c \rrbracket$ as

E.3. Proof of Endpoint Projection

$$\llbracket \Gamma_c \rrbracket = \Gamma_p, k[A] : T \circ \Gamma_q, q.x : U, k[B] : T' \circ \Gamma''$$

such that

$$\frac{\Gamma'' \vdash_{\min} C'' \quad \frac{\Gamma_p, k[A] : T \vdash_{\min} \llbracket C' \rrbracket_p \quad \Gamma_q, q.x : U, k[B] : T' \vdash_{\min} \llbracket C' \rrbracket_q}{\Gamma_p, k[A] : T \circ \Gamma_q, q.x : U, k[B] : T' \vdash_{\min} \llbracket C' \rrbracket_p \mid \llbracket C' \rrbracket_q} \llbracket \text{Min} \rrbracket_{\text{PAR}}}{\Gamma'' \circ \Gamma_p, k[A] : T \circ \Gamma_q, q.x : U, k[B] : T' \vdash_{\min} \llbracket C' \rrbracket_p \mid \llbracket C' \rrbracket_q \mid C''} \llbracket \text{Min} \rrbracket_{\text{PAR}}$$

From the derivation on Rule $\llbracket \text{Min} \rrbracket_{\text{COM}}$ we know that

$$\llbracket \Gamma' \rrbracket = \Gamma'' \circ \Gamma_p \circ \Gamma_q$$

and therefore that

$$\llbracket \Gamma \rrbracket = \Gamma'' \circ \Gamma_p, k[A] : !B.o(U); T \circ \Gamma_q, k[B] : !A.o(U); T'$$

To prove $\llbracket \Gamma \rrbracket \vdash_{\min} \llbracket C \rrbracket$ we prove that we can apply Rule $\llbracket \text{Min} \rrbracket_{\text{PAR}}$ on it.

$$\frac{\frac{\frac{\Gamma_p \vdash p : k[A] \quad q : k[B] \notin \Gamma_q \quad \Gamma_p \vdash p.e : U}{\Gamma_p, k[A] : T \vdash_{\min} \llbracket C' \rrbracket_p} \llbracket \text{Min} \rrbracket_{\text{SEND}} \quad \frac{\Gamma_q \vdash p : k[B] \quad p : k[A] \notin \Gamma_q}{\Gamma_q, q.x : U, k[A] : T' \vdash_{\min} \llbracket C' \rrbracket_q} \llbracket \text{Min} \rrbracket_{\text{RECV}}}{\frac{\Gamma_p, k[A] : !B.o(U); T \quad \Gamma_q, q.x : U, k[A] : T' \vdash_{\min} \llbracket C' \rrbracket_q}{\Gamma_p, k[B] : ?A.o(U); T' \quad \vdash_{\min} k : A \rightarrow q[B].o(x); \llbracket C' \rrbracket_p} \llbracket \text{Min} \rrbracket_{\text{PAR}}}}{\frac{\Gamma'' \vdash_{\min} C'' \quad \frac{\Gamma_p, k[A] : !B.o(U); T \circ \Gamma_q, k[B] : !A.o(U); T' \quad \Gamma_q, q.x : U, k[A] : T' \vdash_{\min} \llbracket C' \rrbracket_q}{\Gamma_p, k[A] : !B.o(U); T \circ \Gamma_q, k[B] : !A.o(U); T' \quad \vdash_{\min} k : p[A].e \rightarrow B.o; \llbracket C' \rrbracket_p} \llbracket \text{Min} \rrbracket_{\text{PAR}}}}{\Gamma'' \circ \Gamma_p, k[A] : !B.o(U); T \circ \Gamma_q, k[B] : !A.o(U); T' \quad \vdash_{\min} k : p[A].e \rightarrow B.o; \llbracket C' \rrbracket_p \mid k : A \rightarrow q[B].o(x); \llbracket C' \rrbracket_q \mid C''} \llbracket \text{Min} \rrbracket_{\text{PAR}}$$

Case $\llbracket \text{Min} \rrbracket_{\text{SEND}}$

From the premises we have $C = k : p[A].e \rightarrow B.o; C'$ on which we can apply the typing derivation

$$\frac{j \in I \quad \Gamma' \vdash q : k[A] \quad q : k[B] \notin \Gamma' \quad \Gamma' \vdash p.e : U_j \quad \Gamma', k[A] : T_j \vdash_{\min} C''}{\Gamma', k[A] : !B.\{o_i(U_i); T_i\}_{i \in I} \vdash_{\min} k : p[A].e \rightarrow B.o; C''} \llbracket \text{Min} \rrbracket_{\text{SEND}}$$

the case is proven by applying the induction hypothesis similarly to case $\llbracket \text{Min} \rrbracket_{\text{COM}}$.

Case $[\text{Min}]_{\text{RECV}}$
 Analogous to case $[\text{Min}]_{\text{SEND}}$.

Case $[\text{Min}]_{\text{PAR}}$
 From the premises we know that $C = C_1 \mid C_2$ on which we can apply the typing derivation

$$\frac{\Gamma_1 \vdash_{\text{min}} C_1 \quad \Gamma_2 \vdash_{\text{min}} C_2}{\Gamma_1 \circ \Gamma_2 \vdash_{\text{min}} C_1 \mid C_2} [\text{Min}]_{\text{PAR}}$$

the case is proven by Lemma 18 and the induction hypothesis.

Case $[\text{Min}]_{\text{END}}$
 Trivial.

Case $[\text{Min}]_{\text{DEF}}$
 From the premises we know that $C = \text{def } X \langle \tilde{r} \rangle = C''$ in C' on which we can apply the typing derivation, with $\Gamma = \Gamma_d \nabla \Gamma'$

$$\frac{\Gamma', X \langle \tilde{r} \rangle : \Gamma_d \vdash_{\text{min}} C' \quad \Gamma_d, X \langle \tilde{r} \rangle : \Gamma_d \vdash_{\text{min}} C'' \quad \Gamma_d \upharpoonright_{\text{locs}} \subseteq \Gamma'}{\Gamma_d \nabla \Gamma' \vdash_{\text{min}} \text{def } X \langle \tilde{r} \rangle = C'' \text{ in } C'} [\text{Min}]_{\text{DEF}}$$

By its definition

$$\llbracket \text{def } X \langle \tilde{r} \rangle = C'' \text{ in } C' \rrbracket = C_d \mid C'''$$

Where

$$C_d = \prod_{p \in \tilde{r}} \text{def } X_p = \llbracket C'' \rrbracket_p \text{ in } \llbracket C' \rrbracket_p$$

and

$$C''' = \prod_{q \in \text{fp}(C')/\{\tilde{r}\}} \llbracket C' \rrbracket_q \mid \prod_l \left(\bigsqcup_{s \in \llbracket C' \rrbracket_l} \llbracket C' \rrbracket_s \right)$$

From the induction hypothesis we know that $\llbracket \Gamma', X \langle \tilde{r} \rangle : \Gamma_d \rrbracket \vdash_{\text{min}} \llbracket C' \rrbracket$.

Let $\llbracket \Gamma', X \langle \tilde{r} \rangle : \Gamma_d \rrbracket = \llbracket \Gamma' \rrbracket \circ \llbracket X \langle \tilde{r} \rangle : \Gamma_d \rrbracket = \llbracket \Gamma' \rrbracket \circ X_{p_1} : \llbracket \Gamma_d \rrbracket_{p_1} \circ \dots \circ X_{p_n} : \llbracket \Gamma_d \rrbracket_{p_n}$ for $\{p_1, \dots, p_n\} = \{\tilde{r}\}$ where $\llbracket \Gamma' \rrbracket$ can be partitioned as $\llbracket \Gamma' \rrbracket = \Gamma'' \circ \Gamma_d^{p_1} \circ \dots \circ \Gamma_d^{p_n}$ such that $\Gamma_{ds} = \Gamma_d^{p_1}, X_{p_1} : \llbracket \Gamma_d \rrbracket_{p_1} \circ \dots \circ \Gamma_d^{p_n}, X_{p_n} : \llbracket \Gamma_d \rrbracket_{p_n}$ and $\llbracket \Gamma', X \langle \tilde{r} \rangle : \Gamma_d \rrbracket = \Gamma'' \circ \Gamma_{ds}$.

Likewise, we apply the induction hypothesis on $\llbracket \Gamma_d, X \langle \tilde{r} \rangle : \Gamma_d \rrbracket \vdash_{\text{min}} C''$ and we can partition $\llbracket \Gamma_d, X \langle \tilde{r} \rangle : \Gamma_d \rrbracket$ in a similar way wrt $\llbracket \Gamma', X \langle \tilde{r} \rangle : \Gamma_d \rrbracket$.

E.3. Proof of Endpoint Projection

Finally we can write the derivation that proves the case

$$\frac{\Gamma'' \vdash_{\min} C''' \quad \frac{\Delta \quad \frac{\Gamma_d / \Gamma_d^{\mathfrak{p}_1} \vdash_{\min} \prod_{i=2}^n \text{def } X_{\mathfrak{p}_i} = \llbracket C'' \rrbracket_{\mathfrak{p}_i} \text{ in } \llbracket C' \rrbracket_{\mathfrak{p}_i}}{\Gamma_{ds} \vdash_{\min} \text{def } X_{\mathfrak{p}_1} = \llbracket C'' \rrbracket_{\mathfrak{p}_1} \text{ in } \llbracket C' \rrbracket_{\mathfrak{p}_1} \mid \prod_{i=2}^n \text{def } X_{\mathfrak{p}_i} = \llbracket C'' \rrbracket_{\mathfrak{p}_i} \text{ in } \llbracket C' \rrbracket_{\mathfrak{p}_i}} \left[\text{Min} \mid_{\text{PAR}} \right]}{\llbracket \Gamma \rrbracket \vdash_{\min} C_d \mid C'''} \left[\text{Min} \mid_{\text{PAR}} \right]} \left[\text{Min} \mid_{\text{DEF}} \right]$$

where

$$\Delta = \left\{ \frac{\Gamma_d^{\mathfrak{p}_1}, X_{\mathfrak{p}_1} : \llbracket \Gamma_d \rrbracket_{\mathfrak{p}_1} \vdash_{\min} \llbracket C' \rrbracket_{\mathfrak{p}_1} \quad \llbracket \Gamma_d \rrbracket_{\mathfrak{p}_1}, X_{\mathfrak{p}_1} : \llbracket \Gamma_d \rrbracket_{\mathfrak{p}_1} \vdash_{\min} \llbracket C'' \rrbracket_{\mathfrak{p}_1} \quad \llbracket \Gamma_d \rrbracket_{\mathfrak{p}_1} \Big|_{\text{locs}} \subseteq \Gamma_d^{\mathfrak{p}_1}}{\Gamma_d^{\mathfrak{p}_1}, X_{\mathfrak{p}_1} : \llbracket \Gamma_d \rrbracket_{\mathfrak{p}_1} \vdash_{\min} \text{def } X_{\mathfrak{p}_1} = \llbracket C'' \rrbracket_{\mathfrak{p}_1} \text{ in } \llbracket C' \rrbracket_{\mathfrak{p}_1}} \right\} \left[\text{Min} \mid_{\text{DEF}} \right]$$

Case $\left[\text{Min} \mid_{\text{CALL}} \right]$

From the premises we know that $C = X \langle \tilde{r} \rangle$, on which we can apply the typing derivation

$$\frac{\text{end}(\Gamma) \quad \Gamma' \subseteq \Gamma_d}{\Gamma, \Gamma', X \langle \tilde{r} \rangle : \Gamma_d \vdash_{\min} X \langle \tilde{r} \rangle} \left[\text{Min} \mid_{\text{CALL}} \right]$$

Since $\Gamma, \Gamma', X \langle \tilde{r} \rangle : \Gamma_d \vdash_{\min} X \langle \tilde{r} \rangle$ we know that

$$\llbracket X \langle \tilde{r} \rangle \rrbracket = \prod_{\mathfrak{p} \in \tilde{r}} X_{\mathfrak{p}}$$

Let $\Gamma_c = \Gamma', \Gamma'', X \langle \tilde{r} \rangle : \Gamma_d$, following Lemma 19 we can partition $\llbracket \Gamma_c \rrbracket$ such that, let $\{\mathfrak{p}_1, \dots, \mathfrak{p}_n\} = \tilde{r}$

$$\llbracket \Gamma_c \rrbracket = \llbracket \Gamma' \rrbracket, \llbracket \Gamma'' \rrbracket, \llbracket X \langle \tilde{r} \rangle : \Gamma_d \rrbracket$$

where

$$\llbracket X \langle \tilde{r} \rangle : \Gamma_d \rrbracket = X_{\mathfrak{p}_1} : \llbracket \Gamma_d \rrbracket_{\mathfrak{p}_1} \circ \dots \circ X_{\mathfrak{p}_n} : \llbracket \Gamma_d \rrbracket_{\mathfrak{p}_n}$$

$$\llbracket \Gamma' \rrbracket = \Gamma'_{\mathfrak{p}_1} \circ \dots \circ \Gamma'_{\mathfrak{p}_n}$$

$$\llbracket \Gamma'' \rrbracket = \Gamma''_{\mathfrak{p}_1} \circ \dots \circ \Gamma''_{\mathfrak{p}_n}$$

We can write the derivation that proves the case

$$\frac{\frac{\text{end}(\Gamma'_{p_1}) \quad \Gamma''_{p_1} \subseteq \llbracket \Gamma_d \rrbracket_{p_1}}{\Gamma'_{p_1}, \Gamma''_{p_1}, X_{p_1} : \llbracket \Gamma_d \rrbracket_{p_1} \vdash_{\min} X_{p_1}} \quad \llbracket \text{Min} \rrbracket_{\text{CALL}} \quad \Delta}{\llbracket \Gamma' \rrbracket, \llbracket \Gamma'' \rrbracket \llbracket X \langle \tilde{r} \rangle : \Gamma_d \rrbracket \vdash_{\min} \llbracket X \langle \tilde{r} \rangle \rrbracket} \quad \llbracket \text{Min} \rrbracket_{\text{PAR}}}{\Delta = \frac{\vdots}{\Gamma'_{p_2}, \Gamma''_{p_2}, X_{p_2} : \llbracket \Gamma_d \rrbracket_{p_2} \circ \dots \circ \Gamma'_{p_n}, \Gamma''_{p_n}, X_{p_n} : \llbracket \Gamma_d \rrbracket_{p_n} \vdash_{\min} \prod_{p \in \{\tilde{p}\}/p_1} X_p} \quad \llbracket \text{Min} \rrbracket_{\text{PAR}}}$$

as, for all $i \in \{1, \dots, n\}$, $\text{end}(\Gamma'_{p_i})$ and $\Gamma''_{p_i} \subseteq \llbracket \Gamma_d \rrbracket_{p_i}$ hold by the premises and by construction. □

We finally prove Theorem 6.

Theorem 6 (EPP Type Preservation) Let $\Gamma \vdash_{\min} D, C$ then $\llbracket \Gamma \rrbracket \vdash_{\min} D, \llbracket C \rrbracket$.

Proof. To prove Theorem 6 we need to prove that we can apply Rule $\llbracket \text{Min} \rrbracket_{\text{DC}}$ on $\llbracket \Gamma \rrbracket \vdash_{\min} D, \llbracket C \rrbracket$

$$\frac{\text{pco}(\llbracket \Gamma \rrbracket) \quad \llbracket \Gamma \rrbracket \vdash_{\min} D \quad \llbracket \Gamma \rrbracket \vdash_{\min} \llbracket C \rrbracket}{\llbracket \Gamma \rrbracket \vdash_{\min} D, \llbracket C \rrbracket} \quad \llbracket \text{Min} \rrbracket_{\text{DC}}$$

where

- $\text{pco}(\llbracket \Gamma \rrbracket)$ and $\llbracket \Gamma \rrbracket \vdash_{\min} D$ hold because $\llbracket \Gamma \rrbracket$ alters only procedure definitions in Γ ;
- $\llbracket \Gamma \rrbracket \vdash_{\min} \llbracket C \rrbracket$ holds from Lemma 20 and the assumption of well-sortedness on C (if C is well-sorted also $\llbracket C \rrbracket$ is well-sorted and typable by Γ). □

E.3.4 EPP Theorem

Before proving Theorem 7 we define some auxiliary concepts to establish a correspondence between a choreography and its projection.

Lemma 21 (EPP Swap Invariance). *Let $C \simeq_C C'$ then $\llbracket C \rrbracket \simeq_C \llbracket C' \rrbracket$.*

E.3. Proof of Endpoint Projection

Proof Sketch. In the proof we show that the projection is invariant under the rules for the swapping relation \simeq_C defined in Figure D.1. $\llbracket \text{CS} \rrbracket_{\text{ETAETA}}$ is trivial. For Rule $\llbracket \text{CS} \rrbracket_{\text{ETACND}}$ we need to check that the projections of the processes in the swapped interaction η do not change, which holds by the definition of EPP for (*cond*) terms and the merging operator (merging the same η returns η). The same reasoning on the EPP and the merging operator applies to all other cases. \square

Lemma 22 (EPP under \equiv). *Let $C \equiv C'$ then $\llbracket C \rrbracket \equiv \llbracket C' \rrbracket$.*

Proof. Easy by cases on the rules of \equiv for C . \square

Lemma 23 (Compositional EPP). *Let C be well-typed and $C = C_1 \mid C_2$ then $\llbracket C \rrbracket \equiv \llbracket C_1 \rrbracket \mid \llbracket C_2 \rrbracket$.*

Proof. By definition of EPP

$$\llbracket C \rrbracket = \prod_{p \in \text{fp}(C)} \llbracket C \rrbracket_p \mid \prod_l \left(\bigsqcup_{s \in [C]_l} \llbracket C \rrbracket_s \right)$$

Since C is well-typed and $C = C_1 \mid C_2$, Rule $\llbracket \text{T} \rrbracket_{\text{PAR}}$ applies and by definition of $\Gamma_1 \circ \Gamma_2$ there cannot be a process p such that $p \in \text{fp}(C_1) \cap \text{fp}(C_2)$. Therefore we can write

$$\llbracket C \rrbracket \equiv \prod_{p \in \text{fp}(C_1)} \llbracket C_1 \rrbracket_p \mid \prod_{q \in \text{fp}(C_2)} \llbracket C_2 \rrbracket_q \mid \prod_l \left(\bigsqcup_{s \in [C]_l} \llbracket C \rrbracket_s \right)$$

By the definition of service typing we know that *i*) locations can implement only one role in a choreography and *ii*) a location can appear only in one service typing. Therefore there cannot be two service processes at the same location in C_1 and C_2 . Thus we can write

$$\llbracket C \rrbracket \equiv \underbrace{\prod_{p \in \text{fp}(C_1)} \llbracket C_1 \rrbracket_p}_{C_1^a} \mid \underbrace{\prod_{q \in \text{fp}(C_2)} \llbracket C_2 \rrbracket_q}_{C_2^a} \mid \underbrace{\prod_l \left(\bigsqcup_{r \in [C_1]_l} \llbracket C_1 \rrbracket_r \right)}_{C_1^s} \mid \underbrace{\prod_{l'} \left(\bigsqcup_{s \in [C_2]_{l'}} \llbracket C_2 \rrbracket_s \right)}_{C_2^s}$$

where $\llbracket C_1 \rrbracket = C_1^a \mid C_1^s$ and $\llbracket C_2 \rrbracket = C_2^a \mid C_2^s$ by definition of EPP. \square

Pruning

Following our definition of EPP, the projection of (*start*) terms on service processes yield a parallel composition of (*acc*) terms on the locations subject of the

(*start*). However, the reduction of a (*start*) term might remove the availability to start new processes on the locations subject of the (*start*) (i.e., if the reductum does not contain another (*start*) term on the same locations). Contrarily, (*acc*) terms remain always available.

A similar observation can be drawn between conditional branches that contain (*com*) terms whose projection merges all possible communications into (*recv*) and (*send*) terms. Also in this case, reducing the condition and projecting the result yields a subset of all possible branches for the considered communication.

Similarly to [14] and [96], we deal with these asymmetries by introducing the *pruning relation*, which allows us to ignore unused *i*) endpoint services and *ii*) input branches.

Definition 37 (Pruning). Let $\Gamma \vdash_{\min} \llbracket C \rrbracket$, $\Gamma = \Gamma_p, \Gamma_{acc}$ where $\Gamma_{acc}|_{\text{locs}} = \Gamma_{acc}$ and $\Gamma_p|_{\text{locs}} = \emptyset$ and $\llbracket C \rrbracket = C_p \mid C_{acc}$, $\Gamma_p \vdash_{\min} C_p$, $\Gamma_{acc} \vdash_{\min} C_{acc}$.

Let $C' = C'_p \mid C'_{acc}$ such that $\Gamma_p, \Gamma_{acc}, \Gamma'_{acc} \vdash_{\min} C'$ where

- $\Gamma_{acc}, \Gamma'_{acc} \vdash_{\min} C'_{acc}$, $\Gamma'_{acc}|_{\text{locs}} = \Gamma'_{acc}$ and $\Gamma'_{acc} \cap \Gamma_{acc} = \emptyset$,
- $\Gamma_p \vdash_{\min} C'_p$ and $C'_p \sqcup C_p = C'_p$

then $\llbracket C \rrbracket$ prunes C' under Γ , written $\Gamma \vdash_{\min} \llbracket C \rrbracket \prec C'$.

Like in [96] we can write $\llbracket C \rrbracket \prec C'$ as it does not lose any precision since it is always possible to reconstruct appropriate typings.

The \prec relation is a typed strong bisimulation because if $\llbracket C \rrbracket \prec C'$ then $\llbracket C \rrbracket$ and C' have the same observable behaviours except for the visible input actions at pruned inputs (either input branches or services).

Lemma 24 (Pruning Lemma). \prec is a strong reduction bisimulation over partial choreographies if $\llbracket C \rrbracket \prec C'$ and for some D, D'

- if $D, \llbracket C \rrbracket \rightarrow D', C''$ then $D, C' \rightarrow D', C'''$ and $C'' \prec C'''$ and
- if $D, C' \rightarrow D, C'''$ then $D, \llbracket C \rrbracket \rightarrow D, C''$ and $C'' \prec C'''$

Proof Sketch. Since $\llbracket C \rrbracket \prec C'$ we know that C' does not use any of the locations not typed by the minimal typing of $\llbracket C \rrbracket$. Similarly, the input branches pruned in $\llbracket C \rrbracket$ are only those that are never used in C' and cannot contribute in the reduction. \square

Before continuing with the last auxiliary results and the proof of Theorem 7 we need to augment the labels of the semantics of annotated ACs with the identifiers of the processes involved in a reduction

$$\beta ::= k : p[A] \rightarrow B.o \mid A \rightsquigarrow q[B].o(x) \mid \tau @ p \mid \tau$$

E.3. Proof of Endpoint Projection

and the annotation of the reduction with Rule $[C]_{\text{COND}}$ as

$$\frac{i = 1 \text{ if } \text{eval}(e, D(\mathbf{p}).\text{st}) = \text{true}, i = 2 \text{ otherwise}}{D, \text{ if } \mathbf{p}.e \{C_1\} \text{ else } \{C_2\} \xrightarrow{\tau@p} D, C_i} [C]_{\text{COND}}$$

Let also $\text{pn}(k : \mathbf{p}[A] \rightarrow B.o) = \{\mathbf{p}\}$, $\text{pn}(A \rightsquigarrow \mathbf{q}[B].o(x)) = \{\mathbf{q}\}$, $\text{pn}(\tau@p) = \{\mathbf{p}\}$, and $\text{pn}(\tau) = \emptyset$

Lemma 25 (Passive Processes Pruning Invariance). $D, C \xrightarrow{\beta} D', C'$ implies that for all $\mathbf{p} \in \text{fp}(C)/\text{pn}(\beta)$, $\llbracket C' \rrbracket_{\mathbf{p}} \prec \llbracket C \rrbracket_{\mathbf{p}}$.

Sketch Proof. By cases on the derivation of C . The only interesting case is $[C]_{\text{COND}}$ in which the the projection of the processes receiving selections are merged. The thesis follows directly from Definition 37 and Lemmas 21 and 22. \square

Lemma 26 (Deployment Invariance). Let D and D' be two deployments such that

1. $\forall \mathbf{p}.x : U \in \Gamma, x(D(\mathbf{p}).\text{st}) = x(D'(\mathbf{p}).\text{st})$
2. $\forall \mathbf{p} : k[A], \mathbf{q} : k[B] \in \Gamma, D(\mathbf{p}).\text{que}(\underline{\mathbf{k.A.B}}(D(\mathbf{p}).\text{st})) = D'(\mathbf{p}).\text{que}(\underline{\mathbf{k.A.B}}(D'(\mathbf{p}).\text{st}))$
3. $\Gamma \vdash_{\min} D, C$ and $\Gamma \vdash_{\min} D', C$ for some C .

then $D, C \sim D', C$.

Sketch Proof. By cases on the derivation of C . Reductions on C are invariant wrt to a specific deployment as long as it is well-typed (3.) and for each process: typed variables have the same value (1.) and queues hold the same messages (2.). \square

E.3.5 Proof of Theorem 7

We restate Theorem 7 to include annotated reductions.

Theorem 7 (EPP Operational Correspondence)

Let D, C be well-typed and well-annotated. Then,

1. (Completeness) $D, C \xrightarrow{\beta} D', C'$ implies $D, \llbracket C \rrbracket \xrightarrow{\beta} D', C''$ and $\llbracket C' \rrbracket \prec C''$.
2. (Soundness) $D, \llbracket C \rrbracket \xrightarrow{\beta} D', C''$ implies $D, C \xrightarrow{\beta} D', C'$ and $\llbracket C' \rrbracket \prec C''$.

We report below the respective proofs of (Completeness) and (Soundness) separately.

Proof (Completeness).

Proof by induction on the derivation of $D, C \xrightarrow{\beta} D', C'$.

Case $[^C|_{\text{COM}}]$

we know that $C = k : p[A].e \rightarrow q[B].o(x); C_c$ and we can write the derivation

$$\frac{\eta = k : p[A].e \rightarrow q[B].o(x) \quad D, k : p[A].e \rightarrow B.o \blacktriangleright D'}{D, \eta; C \xrightarrow{k: p[A] \rightarrow B.o} D', k : A \rightarrow q[B].o(x); C_c} [^C|_{\text{COM}}]$$

and $C' = k : A \rightarrow q[B].o(x); C_c$.

From the definition of EPP we have that $\llbracket C \rrbracket = C_{act} \mid C_s$ such that

$$C_{act} = k : p[A].e \rightarrow B.o; \llbracket C_c \rrbracket_p \mid k : A \rightarrow q[B].o(x); \llbracket C_c \rrbracket_q \mid \prod_{r \in \text{fp}(C)/\{p,q\}} \llbracket C_c \rrbracket_r$$

and

$$C_s = \prod_l \left(\bigsqcup_{s \in [C]_l} \llbracket C_c \rrbracket_s \right)$$

Whilst $\llbracket C' \rrbracket \equiv C'_{act} \mid C_s$

$$C'_{act} = \llbracket C_c \rrbracket_p \mid k : A \rightarrow q[B].o(x); \llbracket C_c \rrbracket_q \mid \prod_{r \in \text{fp}(C')/\{p,q\}} \llbracket C_c \rrbracket_r$$

We can apply Rules $[^C|_{\text{PAR}}]$, $[^C|_{\text{EQ}}]$, and $[^C|_{\text{SEND}}]$ on $D, \llbracket C \rrbracket$ such that

$$\frac{\eta = k : p[A].e \rightarrow B.o \quad D, \eta \blacktriangleright D''}{D, \llbracket C \rrbracket \xrightarrow{k: p[A] \rightarrow B.o} D'', C''} [^C|_{\text{SEND}}]$$

for which it holds that $D' = D''$ by Rule $[^D|_{\text{SEND}}]$.

$$C'' = \llbracket C_c \rrbracket_p \mid k : A \rightarrow q[B].o(x); \llbracket C_c \rrbracket_q \mid \prod_{r \in \text{fp}(C')/\{p,q\}} \llbracket C_c \rrbracket_r \mid C_s$$

for which it holds that $\llbracket C' \rrbracket \prec C''$.

E.3. Proof of Endpoint Projection

Case $\llbracket^C|_{\text{SEND}} \rrbracket$

Similar to case $\llbracket^C|_{\text{COM}} \rrbracket$.

Case $\llbracket^C|_{\text{RECV}} \rrbracket$

we know that $C = D$, $k : A \rightarrow \mathbf{q}[\mathbf{B}].\{o_i(x_i); C_i\}_{i \in I}$ and we can write the derivation

$$\frac{j \in I \quad D, k : A \rightarrow \mathbf{q}[\mathbf{B}].o_j(x_j) \blacktriangleright D'}{D, k : A \rightarrow \mathbf{q}[\mathbf{B}].\{o_i(x_i); C_i\}_{i \in I} \xrightarrow{k:A \rightsquigarrow \mathbf{B}.o_j(x_j)} D', C_j} \llbracket^C|_{\text{RECV}} \rrbracket$$

for $\beta = k : A \rightsquigarrow \mathbf{B}.o_j(x_j)$ and $C' = C_j$.

By the definition of EPP we have

$$\llbracket C \rrbracket \equiv k : A \rightarrow \mathbf{q}[\mathbf{B}].\{o_i(x_i); \llbracket C_i \rrbracket_{\mathbf{q}}\}_{i \in I} \mid \prod_{\mathbf{p} \in \text{fp}(C)/\{\mathbf{q}\}} \left(\bigsqcup_{i \in I} \llbracket C_i \rrbracket_{\mathbf{p}} \right) \mid \prod_l \left(\bigsqcup_{r \in \llbracket C \rrbracket_l} \llbracket C \rrbracket_r \right)$$

Then we can apply Rules $\llbracket^C|_{\text{PAR}} \rrbracket$, $\llbracket^C|_{\text{EQ}} \rrbracket$, and $\llbracket^C|_{\text{SEND}} \rrbracket$ such that

$$\frac{j \in I \quad D, k : A \rightarrow \mathbf{q}[\mathbf{B}].o_j(x_j) \blacktriangleright D'}{D, \llbracket C \rrbracket \xrightarrow{k:A \rightsquigarrow \mathbf{B}.o_j(x_j)} D'', \llbracket C_j \rrbracket_{\mathbf{q}} \mid \prod_{\mathbf{p} \in \text{fp}(C)/\{\mathbf{q}\}} \left(\bigsqcup_{i \in I} \llbracket C_i \rrbracket_{\mathbf{p}} \right) \mid \prod_l \left(\bigsqcup_{r \in \llbracket C \rrbracket_l} \llbracket C \rrbracket_r \right)} \llbracket^C|_{\text{RECV}} \rrbracket$$

and

$$C'' = \llbracket C_j \rrbracket_{\mathbf{q}} \mid \prod_{\mathbf{p} \in \text{fp}(C)/\{\mathbf{q}\}} \left(\bigsqcup_{i \in I} \llbracket C_i \rrbracket_{\mathbf{p}} \right) \mid \prod_l \left(\bigsqcup_{r \in \llbracket C \rrbracket_l} \llbracket C \rrbracket_r \right)$$

From Rule $\llbracket^D|_{\text{RECV}} \rrbracket$ we know that $D'' = D'$.

$\llbracket C' \rrbracket \prec C''$ by Definition 37 and Lemma 25.

Case $\llbracket^C|_{\text{START}} \rrbracket$

we know that $C = \mathbf{start} \ k : \mathbf{p}[\mathbf{A}] \Leftrightarrow \widetilde{l.q[\mathbf{B}]}; C_c$ and we can write the derivation

$$\frac{\#\tilde{r} \quad \#k' \quad \mathbf{p} \in D(l) \quad \delta = \mathbf{start} \ k' : l.\mathbf{p}[\mathbf{A}], \widetilde{l.r[\mathbf{B}]} \quad D, \delta \blacktriangleright D'}{D, \mathbf{start} \ k : \mathbf{p}[\mathbf{A}] \Leftrightarrow \widetilde{l.q[\mathbf{B}]}; C_c \xrightarrow{\tau} D', C_c[k'/k][\tilde{r}/\tilde{q}]} \llbracket^C|_{\text{START}} \rrbracket$$

and $C' = C_c[k'/k][\tilde{r}/\tilde{q}]$.

From the definition of EPP we have

$$\llbracket C' \rrbracket = \prod_{q \in \text{fp}(C')} \llbracket C' \rrbracket_q \mid \prod_l \left(\bigsqcup_{s \in \llbracket C' \rrbracket_l} \llbracket C' \rrbracket_s \right)$$

and

$$\llbracket C \rrbracket \equiv \begin{array}{l} \mathbf{req} \ k : p[A] \Leftrightarrow \widetilde{l.B}; \llbracket C_c \rrbracket_p \mid \prod_{r \in \text{fp}(C)/\{p\}} \llbracket C \rrbracket_r \\ \mid \prod_{l.q[B] \in \widetilde{l.q[B]}} \mathbf{acc} \ k : l.q[B]; \llbracket C_c \rrbracket_q \mid \prod_{l' \notin \tilde{l}} \left(\prod_{s \in \llbracket C \rrbracket_{l'}} \llbracket C \rrbracket_s \right) \end{array}$$

we can apply Rules $\llbracket C \rrbracket_{\text{PAR}}$, $\llbracket C \rrbracket_{\text{EQ}}$, $\llbracket C \rrbracket_{\text{PSTART}}$ such that

$$\begin{array}{c} i \in \{1, \dots, n\} \quad \#k'' \quad \{\widetilde{l.B}\} = \bigsqcup_i \{\widetilde{l_i.B_i}\} \\ \# \tilde{r}' \quad \{\tilde{r}'\} = \bigcup_i \{\tilde{r}'_i\} \quad p \in D(l) \\ \delta = \mathbf{start} \ k'' : l.p[A], l_1.r'_1[B_1], \dots, l_n.r'_n[B_n] \quad D, \delta \blacktriangleright D'' \\ \hline D, \llbracket C \rrbracket \xrightarrow{\tau} D'', C'' \quad \llbracket C \rrbracket_{\text{PSTART}} \end{array}$$

where

$$C'' \equiv \left\{ \begin{array}{l} \llbracket C_c \rrbracket_p[k''/k] \mid \prod_{(q,r') \in \{(q_1,r'_1), \dots, (q_n,r'_n)\}} \llbracket C_c \rrbracket_q[k''/k][q/r'] \\ \mid \prod_{r \in \text{fp}(C_c)/\{p,q\}} \llbracket C_c \rrbracket_r \mid \prod_{l.q[B] \in \widetilde{l.q[B]}} \mathbf{acc} \ k : l.q[B]; \llbracket C_c \rrbracket_q \\ \mid \prod_{l' \notin \tilde{l}} \left(\prod_{s \in \llbracket C_c \rrbracket_{l'}} \llbracket C_c \rrbracket_s \right) \end{array} \right.$$

Observe that we can α -rename k'' to k' and \tilde{r}' to \tilde{r} as k'' , k' , \tilde{r}' , and \tilde{r} are all fresh wrt D, C .

From the application of Rule $\llbracket D \rrbracket_{\text{START}}$ we can find Γ such that

$$\Gamma \vdash_{\min} (D'', C'')[k'/k''][\tilde{r}/\tilde{r}']$$

and

$$\Gamma \vdash_{\min} D', C''[k'/k''][\tilde{r}/\tilde{r}']$$

E.3. Proof of Endpoint Projection

and by Lemma 26 we have that

$$D, \llbracket C \rrbracket \xrightarrow{\tau} D', C''[k'/k''][\tilde{r}/\tilde{r}']$$

Finally $\llbracket C' \rrbracket \prec C''[k'/k''][\tilde{r}/\tilde{r}']$ by Lemma 25.

Case $\llbracket C \rrbracket_{\text{PSTART}}$

Similar to (in particular the second part of) the proof of case $\llbracket C \rrbracket_{\text{START}}$.

Case $\llbracket C \rrbracket_{\text{COND}}$

we know that $C \equiv \text{if } p.e \{C_1\} \text{ else } \{C_2\}$ and we can write the derivation

$$\frac{i = 1 \text{ if } \text{eval}(e, D(p).\text{st}) = \text{true}, i = 2 \text{ otherwise}}{D, \text{ if } p.e \{C_1\} \text{ else } \{C_2\} \xrightarrow{\tau @ p} D, C_i} \llbracket C \rrbracket_{\text{COND}}$$

We only consider the case for $\text{eval}(e, D(p).\text{st}) = \text{true}$ as $\text{eval}(e, D(p).\text{st}) = \text{false}$ follows alike.

$C' = C_1$ and by the definition of EPP

$$\llbracket C \rrbracket \equiv \text{if } p.e \{ \llbracket C_1 \rrbracket_p \} \text{ else } \{ \llbracket C_2 \rrbracket_p \} \mid \prod_{q \in \text{fp}(C')/\{p\}} \llbracket C_1 \rrbracket_q \sqcup \llbracket C_2 \rrbracket_q \mid \prod_l \left(\bigsqcup_{r \in \llbracket C \rrbracket_l} \llbracket C \rrbracket_r \right)$$

and

$$\llbracket C' \rrbracket \equiv \llbracket C_1 \rrbracket_p \mid \prod_{q \in \text{fp}(C')/\{p\}} \llbracket C_1 \rrbracket_q \mid \prod_l \left(\bigsqcup_{r \in \llbracket C_1 \rrbracket_l} \llbracket C_1 \rrbracket_r \right)$$

We can apply rules $\llbracket C \rrbracket_{\text{PAR}}$, $\llbracket C \rrbracket_{\text{EQ}}$, and $\llbracket C \rrbracket_{\text{COND}}$ such that $D, \llbracket C \rrbracket \xrightarrow{\tau @ p} D, C''$ where

$$C'' = \llbracket C_1 \rrbracket_p \mid \prod_{q \in \text{fp}(C')/\{p\}} \llbracket C_1 \rrbracket_q \sqcup \llbracket C_2 \rrbracket_q \mid \prod_l \left(\bigsqcup_{r \in \llbracket C \rrbracket_l} \llbracket C \rrbracket_r \right)$$

and $\llbracket C' \rrbracket \prec C''$ by Lemma 25.

Case $[^C|_{\text{CTX}}]$ and **Case** $[^C|_{\text{PAR}}]$

Proven by the definition of EPP and the induction hypothesis.

Case $[^C|_{\text{EQ}}]$

We can write the derivation

$$\frac{\mathcal{R} \in \{\equiv, \simeq_C\} \quad C_1 \mathcal{R} C'_1 \quad D, C'_1 \xrightarrow{\beta} D', C'_2 \quad C'_2 \mathcal{R} C_2}{D, C_1 \xrightarrow{\beta} D', C_2} [^C|_{\text{EQ}}]$$

For $\mathcal{R} = \equiv$, proven by the definition of EPP and the induction hypothesis.

For $\mathcal{R} = \simeq_C$, proven by the definition of EPP, Lemma 21, and the induction hypothesis.

□

Proof (Soundness). Proof by induction on the structure of C .

Case $C = k : p[A].e \rightarrow q[B].o(x); C_c$

From the definition of EPP we have

$$\llbracket C \rrbracket \equiv k : p[A].e \rightarrow B.o; \llbracket C_c \rrbracket_p \mid k : A \rightarrow q[B].o(x); \llbracket C_c \rrbracket_q \mid \prod_{r \in \text{fp}(C)} \llbracket C_c \rrbracket_r \mid \prod_l \left(\bigsqcup_{s \in [C]_l} \llbracket C \rrbracket_s \right)$$

we proceed by subcases on the last applied Rule in the derivation of $D, \llbracket C \rrbracket \xrightarrow{\beta} D', C''$.

Case $[^C|_{\text{SEND}}]$

Divided into subcases whether $\beta = k : p[A] \rightarrow B.o$ holds or not.

Case $\beta = k : p[A] \rightarrow B.o$

$D, \llbracket C \rrbracket$ reduces to D', C'' with Rules $[^C|_{\text{PAR}}]$, $[^C|_{\text{EQ}}]$, ending with Rule $[^C|_{\text{SEND}}]$ such that

$$C'' = \llbracket C_c \rrbracket_p \mid k : A \rightarrow q[B].o(x); \llbracket C_c \rrbracket_q \mid \prod_{r \in \text{fp}(C)} \llbracket C_c \rrbracket_r \mid \prod_l \left(\bigsqcup_{s \in [C]_l} \llbracket C \rrbracket_s \right)$$

E.3. Proof of Endpoint Projection

D, C mimics $D, \llbracket C \rrbracket$ with Rule $[^C|_{\text{COM}}]$ for which $D, C \xrightarrow{\beta} D'', C', D' = D''$ by Rule $[^D|_{\text{SEND}}]$,

$$\llbracket C' \rrbracket \equiv \llbracket C_c \rrbracket_p \mid k : A \rightarrow q[B].o(x); \llbracket C_c \rrbracket_q \mid \prod_{r \in \text{fp}(C)} \llbracket C_c \rrbracket_r \mid \prod_l \left(\bigsqcup_{s \in [C]_l} \llbracket C \rrbracket_s \right)$$

and $\llbracket C' \rrbracket \prec C''$ by Lemma 25.

Case $\beta \neq k : p[A] \rightarrow B.o$

In this case D, C can mimic $D, \llbracket C \rrbracket$ with the application of Rules $[^C|_{\text{EQ}}]$ and $[^C|_{\text{PAR}}]$, ending the derivation with either Rule $[^C|_{\text{COM}}]$ or $[^C|_{\text{SEND}}]$. The thesis follows by the induction hypothesis.

Case $[^C|_{\text{RECV}}]$, $[^C|_{\text{PSTART}}]$, or $[^C|_{\text{COND}}]$

In this case $D, \llbracket C \rrbracket$ reduces with Rules $[^C|_{\text{EQ}}]$, $[^C|_{\text{PAR}}]$, and respectively ends the derivation with either $[^C|_{\text{RECV}}]$, $[^C|_{\text{PSTART}}]$, or $[^C|_{\text{COND}}]$, i.e., some process $r \in \text{fp}(C)$ (p and q included) either receives a message, starts a new session with some service processes, or reduces to some branch. D, C can mimic $D, \llbracket C \rrbracket$ applying Rules $[^C|_{\text{EQ}}]$, $[^C|_{\text{PAR}}]$ and terminates the derivation with either Rules $[^C|_{\text{RECV}}]$, $[^C|_{\text{PSTART}}]$, (or $[^C|_{\text{START}}]$, depending on the form of C) or $[^C|_{\text{COND}}]$. The thesis follows by the induction hypothesis.

Case $C = k : p[A].e \rightarrow B.o; C_c$

Similar to case $C = k : p[A].e \rightarrow q[B].o(x); C_c$.

Case $C = k : A \rightarrow q[B].\{o_i(x_i); C_i\}_{i \in I}$

From the definition of EPP we have

$$\llbracket C \rrbracket \equiv k : A \rightarrow q[B].\{o_i(x_i); \llbracket C_i \rrbracket_q\}_{i \in I} \mid \prod_{i \in I} \left(\bigsqcup_{p \in \text{fp}(C_i)} \llbracket C_i \rrbracket_p \right) \mid \prod_k \left(\bigsqcup_{r \in [C]_l} \llbracket C \rrbracket_r \right)$$

we proceed by subcases on last applied Rule in the derivation of $D, \llbracket C \rrbracket \xrightarrow{\beta} D', C''$.

Case $[^C|_{\text{RECV}}]$

Divided into subcases whether $\beta = k : A \rightsquigarrow q[B].o_j, j \in I$ or not.

Case $\beta = k: A \rightsquigarrow q[B].o_j, j \in I$
 $D, \llbracket C \rrbracket$ reduces to D', C'' with Rules $\llbracket C \rrbracket_{\text{PAR}}$, $\llbracket C \rrbracket_{\text{EQ}}$, and terminates with Rule $\llbracket C \rrbracket_{\text{RECV}}$ such that

$$C'' = \llbracket C_j \rrbracket_q \mid \prod_{i \in I} \left(\bigsqcup_{p \in \text{fp}(C_i)} \llbracket C_i \rrbracket_p \right) \mid \prod_k \left(\bigsqcup_{r \in \llbracket C \rrbracket_l} \llbracket C \rrbracket_r \right)$$

D, C mimics $D, \llbracket C \rrbracket$ with Rule $\llbracket C \rrbracket_{\text{RECV}}$ for which $D, C \xrightarrow{\beta} D'', C'$ where $D'' = D'$ by Rule $\llbracket D \rrbracket_{\text{RECV}}$ and

$$C'' = \llbracket C_j \rrbracket_q \mid \prod_{p \in \text{fp}(C_j)} \llbracket C_j \rrbracket_p \mid \prod_k \left(\bigsqcup_{r \in \llbracket C_j \rrbracket_l} \llbracket C_j \rrbracket_r \right)$$

and $\llbracket C' \rrbracket \prec C''$ by Lemma 25.

Case $\beta \neq k: A \rightsquigarrow q[B].o_j$

For any β of this case D, C can mimic $D, \llbracket C \rrbracket$ with the application of Rules $\llbracket C \rrbracket_{\text{EQ}}$ and $\llbracket C \rrbracket_{\text{PAR}}$, terminating with Rule $\llbracket C \rrbracket_{\text{RECV}}$ and the thesis follows by the induction hypothesis.

Case $\llbracket C \rrbracket_{\text{SEND}}$, $\llbracket C \rrbracket_{\text{PSTART}}$, or $\llbracket C \rrbracket_{\text{COND}}$

is similar to subcase **Case** $\llbracket C \rrbracket_{\text{RECV}}$, $\llbracket C \rrbracket_{\text{PSTART}}$, or $\llbracket C \rrbracket_{\text{COND}}$ of **Case**

$C = k: p[A].e \rightarrow q[B].o(x); C_c$.

Case $C = \text{start } k: p[A] \Leftrightarrow \widetilde{l}.q[B]; C_c$

$$\llbracket C \rrbracket \equiv \mathbf{req } k: p[A] \Leftrightarrow \widetilde{l}.B; C_c \mid \prod_{r \in \text{fp}(C_c)} \llbracket C_c \rrbracket_r \mid \prod_l \left(\bigsqcup_{s \in \llbracket C \rrbracket_l} \llbracket C \rrbracket_s \right)$$

we proceed by subcases on last applied Rule in the derivation of $D, \llbracket C \rrbracket \xrightarrow{\beta} D, C''$.

Case $\llbracket C \rrbracket_{\text{PSTART}}$

$D, \llbracket C \rrbracket$ can reduce to D', C'' with a process r (including p) that starts a new session with some service processes. D, C can reduce to D'', C' mimicking $D, \llbracket C \rrbracket$ by applying Rules $\llbracket C \rrbracket_{\text{EQ}}$, $\llbracket C \rrbracket_{\text{PAR}}$, terminating with either Rule $\llbracket C \rrbracket_{\text{PSTART}}$ or $\llbracket C \rrbracket_{\text{START}}$.

The proof follows from Lemma 26 as D'' and D' differ only on the correlation values used in for session k .

E.3. Proof of Endpoint Projection

Case $\llbracket C \rrbracket_{\text{SEND}}$, $\llbracket C \rrbracket_{\text{RECV}}$, and $\llbracket C \rrbracket_{\text{COND}}$
are similar to the corresponding proof for the previous cases.

Case $C = \text{if } p.e \{C_1\} \text{ else } \{C_2\}$

From the definition of EPP we have

$$\llbracket C \rrbracket \equiv \text{if } p.e \{ \llbracket C_1 \rrbracket_p \} \text{ else } \{ \llbracket C_2 \rrbracket_p \} \mid \prod_{q \in \text{fp}(C_1) \cup \text{fp}(C_2) / \{p\}} \llbracket C_1 \rrbracket_q \sqcup \llbracket C_2 \rrbracket_q \mid \prod_l \left(\bigsqcup_{r \in \llbracket C \rrbracket_l} \llbracket C \rrbracket_r \right)$$

we proceed by subcases on the derivation of $D, \llbracket C \rrbracket \xrightarrow{\beta} D', C''$.

Case $\llbracket C \rrbracket_{\text{COND}}$
 $D, \llbracket C \rrbracket$ can reduce to D', C'' with:

Case $\beta = \tau@p$

that reduces to a branch. D, C can mimic $D, \llbracket C \rrbracket$ applying Rules $\llbracket C \rrbracket_{\text{EQ}}$, $\llbracket C \rrbracket_{\text{PAR}}$, and terminating the derivation with Rule $\llbracket C \rrbracket_{\text{COND}}$. The case is proven by Lemma 25.

Case $\beta = \tau@r$

for $r \neq p$. A process r reduced to a branch. The case follows the proof of the previous case where the thesis follows by the induction hypothesis.

Case $\llbracket C \rrbracket_{\text{RECV}}$, $\llbracket C \rrbracket_{\text{SEND}}$, $\llbracket C \rrbracket_{\text{PSTART}}$
are similar to the corresponding proof for the previous cases.

Case $C = \mathbf{req} \ k : p[A] \Leftrightarrow \widetilde{l.B}; C_c$

Case not allowed by the hypothesis that $D, \llbracket C \rrbracket \xrightarrow{\beta} D, C''$.

Case $C = \mathbf{acc} \ k : \widetilde{l.q[B]}; C_c$

Case not allowed by the hypothesis that $D, \llbracket C \rrbracket \xrightarrow{\beta} D, C''$.

Case $C = \mathbf{def} \ X \langle \tilde{p} \rangle = C_d \text{ in } C_c$

Proven by Lemma 22 and the induction hypothesis.

Case $C = X \langle \tilde{p} \rangle$

Case not allowed by the hypothesis that C is well-sorted.

Case $C = C_1 \mid C_2$

$\llbracket C \rrbracket \equiv \llbracket C_1 \rrbracket \mid \llbracket C_2 \rrbracket$ by Lemma 23.

we proceed by subcases for n equal to the length of the derivation of D , $\llbracket C \rrbracket \xrightarrow{\beta} D', C''$

Case $n = 1$

In this case the only applicable Rule is $\llbracket C \rrbracket_{\text{PSTART}}$ where, by the definition of EPP, we can infer, let

$$\widetilde{l.q[B]} = l_1.q_1[B_1], \dots, l_i.q_i[B_i], l_{i+1}.q_{i+1}[B_{i+1}] \dots, l_n.q_n[B_n]$$

that

$$\llbracket C_1 \rrbracket \equiv \mathbf{req} \ k : p[A] \Leftrightarrow \widetilde{l.B}; \llbracket C_1^r \rrbracket_p \mid \prod_{j=1}^i \mathbf{acc} \ k : l_j.q_j[B_j]; \llbracket C_1^j \rrbracket_{q_j}$$

$$\llbracket C_2 \rrbracket \equiv \prod_{j=i+1}^n \mathbf{acc} \ k : l_j.q_j[B_j]; \llbracket C_2^j \rrbracket_{q_j}$$

Observe that we can proceed without loss of generality as the symmetric case (with $p \in \text{fp}(C_2)$) follows the same structure.

$$\begin{array}{c} i \in \{1, \dots, n\} \quad \#k' \quad \{\widetilde{l.B}\} = \uplus_i \{\widetilde{l_i.B_i}\} \\ \#\tilde{r} \quad \{\tilde{r}\} = \bigcup_i \{\tilde{r}_i\} \quad p \in D(l) \\ \delta = \mathbf{start} \ k' : l.p[A], \widetilde{l_1.r_1[B_1]}, \dots, \widetilde{l_n.r_n[B_n]} \quad D, \delta \blacktriangleright D'' \\ \hline D, \llbracket C_1 \rrbracket \mid \llbracket C_2 \rrbracket \quad \xrightarrow{\tau} \quad D', C'' \quad \llbracket C \rrbracket_{\text{PSTART}} \end{array}$$

where

$$C'' \equiv \llbracket C_1^r \rrbracket_p [k'/k] \mid \left(\begin{array}{c} \prod_{j=1}^i \llbracket C_1^j \rrbracket_{q_j} \\ \mid \prod_{j=i+1}^n \llbracket C_2^j \rrbracket_{q_j} \end{array} \right) [k'/k][\tilde{r}/\tilde{q}] \mid \left(\begin{array}{c} \prod_{j=1}^i \mathbf{acc} \ k : l_j.q_j[B_j]; \llbracket C_1^j \rrbracket_{q_j} \\ \mid \prod_{j=i+1}^n \mathbf{acc} \ k : l_j.q_j[B_j]; \llbracket C_2^j \rrbracket_{q_j} \end{array} \right)$$

Since both $\llbracket C_1 \rrbracket$ and $\llbracket C_2 \rrbracket$ reduce, we can infer from the definition of EPP that,

$$C_1 \equiv \mathbf{req} \ k : p[A] \Leftrightarrow \widetilde{l.B}; C_1^r \mid \prod_{j=1}^i \mathbf{acc} \ k : l_j.q_j[B_j]; C_1^j$$

E.3. Proof of Endpoint Projection

$$C_2 \equiv \prod_{j=i+1}^n \mathbf{acc} \ k : l_j \cdot \mathbf{q}_j[B_j]; C_2^j$$

Then D, C can mimic $D, \llbracket C \rrbracket$ applying Rule $[C]_{\text{PSTART}}$ with reduction

$$\frac{\begin{array}{l} i \in \{1, \dots, n\} \quad \#k'' \quad \{\widetilde{l.B}\} = \biguplus_i \{\widetilde{l_i.B_i}\} \quad \#\widetilde{r}' \quad \{\widetilde{r}'\} = \bigcup_i \{\widetilde{r}'_i\} \\ \mathbf{p} \in D(l) \quad \delta = \mathbf{start} \ k'' : l \cdot \mathbf{p}[A], l_1 \cdot r'_1[B_1], \dots, l_n \cdot r'_n[B_n] \quad D, \delta \blacktriangleright D'' \end{array}}{D, \llbracket C_1 \rrbracket \mid \llbracket C_2 \rrbracket \xrightarrow{\tau} D'', C'} \quad [C]_{\text{PSTART}}$$

where

$$C' \equiv C_1^r[k''/k] \mid \left(\frac{\prod_{j=1}^i C_1^j \mid}{\prod_{j=i+1}^n C_2^j} \right) [k''/k][\widetilde{r}'/\widetilde{q}] \mid \left(\frac{\prod_{j=1}^i \mathbf{acc} \ k : l_j \cdot \mathbf{q}_j[B_j]; C_1^j}{\mid \prod_{j=i+1}^n \mathbf{acc} \ k : l_j \cdot \mathbf{q}_j[B_j]; C_2^j} \right)$$

Following the structure of the second part of the proof of **Case** $[C]_{\text{START}}$ for the proof of *Completeness* of Theorem 7, we have, by Lemma 26, $D'', C' \sim D', C'[k'/k''][\widetilde{r}/\widetilde{r}']$ and $\llbracket C' \rrbracket \prec C''$.

Case $n > 1$

For $n > 1$ we have a derivation similar to

$$\frac{\begin{array}{l} R \\ \vdots \quad n-1 \text{ times, each either} \\ \quad [C]_{\text{PAR}} \text{ or } [C]_{\text{EQ}} \end{array}}{D, \llbracket C_1 \rrbracket \mid \llbracket C_2 \rrbracket \xrightarrow{\beta} D', C''_1 \mid \llbracket C_2 \rrbracket} \quad [C]_{\text{PAR}}$$

where R is the last applied Rule, $R \in \{[C]_{\text{SEND}}, [C]_{\text{COM}}, [C]_{\text{SEND}}, [C]_{\text{PSTART}}, [C]_{\text{COND}}\}$. The thesis follows from the induction hypothesis.

The proof for the mirror case $D, \llbracket C_1 \rrbracket \mid \llbracket C_2 \rrbracket \xrightarrow{\beta} D', \llbracket C_1 \rrbracket \mid C''_2$ follows the same structure.

Case $C = 0$
trivial.

□

E.4 Proof of Compilation

E.4.1 Proof of Theorem 8

Lemma 27 (Compilation Invariance). *Let $\Gamma \vdash D, C, D, C \rightarrow D', C'$, and $\Gamma' \vdash D', C'$ for some $\Gamma', \Gamma|_{\text{locs}} = \Gamma'|_{\text{locs}}$ and $\{p@l \mid p@l \in \Gamma\} = \{p@l \mid p@l \in \Gamma'\}$, then $\boxed{D', C'}^\Gamma = \boxed{D', C'}^{\Gamma'}$.*

Proof Sketch. By cases on the derivation of $D, C \rightarrow D', C'$. From Definition 23, the only information used by the compiler present in Γ (respectively Γ') are *i*) the location of running processes ($p@l$) and *ii*) the service typings ($\tilde{l}: G\langle A|\tilde{B}|\tilde{C}\rangle$). Therefore, provided Γ and Γ' have the same service typings and locations the compilation of the given choreography under Γ produces the same result as the compilation of such choreography under Γ' . \square

We provide some auxiliary results on variable substitution. We remind that the only bound names in *DCC* are the variables in (*accept*) terms (e.g., x in $!(x); B$). However, the following lemmas prove that renaming free variables with fresh names in processes (and, by extension, in services) preserves bisimilarity.

In the following, we abuse the notation for α -renaming to denote variable renaming in running processes. We define the variable renaming operator for *DCC* processes $P[x'/x]$.

Definition 38 (DCC Variable Renaming Operator). *Let $B \cdot t \cdot M$ be a DCC process, then $(B \cdot t \cdot M)[x'/x] = B[x'/x] \cdot t \triangleleft (x', x(t)) \triangleleft (x, t_\perp) \cdot M$ where $B[x'/x]$ substitutes every occurrence of x with x' .*

Lemma 28 (DCC Process Variable Renaming). *Let P and P' be two DCC processes such that $P = B \cdot t \cdot M$ and $P' = P[x'/x]$ and x' is fresh in B then $P \xrightarrow{\lambda} P'' \iff P' \xrightarrow{\lambda[x'/x]} P''[x'/x]$.*

Proof. The proof is by induction on the form of P . We report the most interesting cases. Below we consider $t' = t \triangleleft (x', x(t)) \triangleleft (x, t_\perp)$.

Case $P = o(y)$ from $e; B \cdot t \cdot M$

On both P and P' the only applicable Rule is $[\text{DCC}]_{\text{RECV}}$, hence we consider the interesting case in which M contains a message for the queue defined by e . In the other case the Lemma holds as both P and P' cannot reduce. The case unfolds on the combination of whether *i*) $y \neq x$ and *ii*) expression e contains x . Below we consider the comprehensive case for $y = x$ and e that contains x . The proof of the other cases is either trivial or a slight modification of the reported one.

E.4. Proof of Compilation

Since we assume we can apply Rule $[\text{DCC}|_{\text{RECV}}]$ we take $t_c = \text{eval}(e, t)$ and $M(t_c) = (o, t') :: \tilde{m}$. From Definition 38 we have that $t_c = \text{eval}(e[x'/x], t')$.

Therefore we have the reductions $P \xrightarrow{o \text{ from } e} B \cdot t \triangleleft (x, t_m) \cdot M[t_c \mapsto \tilde{m}]$ and $P' \xrightarrow{o \text{ from } e[x'/x]} B[x'/x] \cdot t' \triangleleft (x', t_m) \cdot M[t_c \mapsto \tilde{m}]$.

Case $P = \sum_{i \in I} [o_i(x_i) \text{ from } e] \{B_i\} \cdot t \cdot M$

The only applicable Rule on both P and P' is $[\text{DCC}|_{\text{CHOICE}}]$. The most comprehensive case is for M that contains a message for operation o_j , $j \in I$ where $x_j = x$ and expression e contains x . The remainder of the proof follows that of the previous case.

Case $P = \text{if } e \{B_1\} \text{ else } \{B_2\} \cdot t \cdot M$

Trivial by Definition 38 for which $\text{eval}(e, t) = \text{eval}(e[x'/x], t')$.

Case $P = y = e; B \cdot t \cdot M$

The only applicable Rule on both P and P' is $[\text{DCC}|_{\text{ASSIGN}}]$. The most comprehensive case is for $y = x$ and expression e that contains x . The case is proven considering that, by Definition 38, it holds that $\text{eval}(e, t) = \text{eval}(e[x'/x], t')$.

Case $P = \text{def } X = B_1 \text{ in } B \cdot t \cdot M$

The thesis follows from the application of Rule $[\text{DCC}|_{\text{CTX}}]$ and the induction hypothesis.

□

Lemma 29 (DCC Service Variable Renaming). *Let S and S' be two DCC services such that $S = \langle B_s, P \mid Q \rangle_l$ and $S' = \langle B_s, P[x'/x] \mid Q \rangle_l$ then*

$$S \xrightarrow{\lambda} \langle B_s, P'' \mid Q' \rangle_l \iff S' \xrightarrow{\lambda[x'/x]} \langle B_s, P''[x'/x] \mid Q' \rangle_l.$$

Proof. The proof is by induction on the derivation of S . Below we consider $t' = t \triangleleft (x', x(t)) \triangleleft (x, t_\perp)$. We report the most interesting cases.

Case $[\text{DCC}|_{\text{EQ}}]$

The thesis follows from the application of the induction hypothesis on $S \equiv S_1$, $S_1 \rightarrow S'_1$, and $S'_1 \equiv S'$.

Case $[\text{DCC}|_{\text{PAR}}]$

From Lemma 28 we have that, for any form of P of the kind $P = B \cdot t \cdot M$, $P[x'/x]$ can mimic the reduction of P such that $P \rightarrow P''$ and $P[x'/x] \rightarrow P''[x'/x]$ and vice versa. In this case Q remains the same in both reductions.

Case $[\text{DCC}|_{\text{CQ}}]$

We consider the case where $P = \text{cq}(x); B \cdot t \cdot M$. Let $Q = \prod_i B_i \cdot t_i \cdot M_i$ and $t_c \notin \bigcup_i \text{dom}(M_i) \cup \text{dom}(M)$. We have the reduction

$S \xrightarrow{\text{cq}(x)} \langle B_s, B \cdot t \triangleleft (x, t_c, \cdot) M[t_c \mapsto \varepsilon] \mid Q \rangle_l$. S' can mimic such behaviour by taking the fresh value $t'_c = t_c$, obtaining the reduction

$S' \xrightarrow{\text{cq}(x')} \langle B_s, B \cdot t \triangleleft (x', t'_c, \cdot) M[t'_c \mapsto \varepsilon] \mid Q \rangle_l$. The same holds if we let S' reduce and prove that S can mimic it.

Case $[\text{DCC}|_{\text{INSEND}}]$

We consider the case where $P = o@e_1(e_2) \text{ to } e_3; B \cdot t \cdot M$ and expressions e_1, e_2 and e_3 contain x . From Definition 38 we know that $\text{eval}(e_1, t) = \text{eval}(e_1[x'/x], t')$. Similarly the couples e_2 and $e_2[x'/x]$ and e_3 and $e_3[x'/x]$ enjoy the same property when evaluated respectively on t and t' .

We analyse the case in which P moves and $P[x'/x]$ mimics it. The other case, for $P[x'/x]$ that reduces and P that mimics it, follows the same structure. Since $[\text{DCC}|_{\text{INSEND}}]$ applies we can assume that $Q = B \cdot t \cdot M \mid B_r \cdot t_r \cdot M_r \mid P_1$.

$$\frac{P = o@e_1(e_2) \text{ to } e_3; B \cdot t \cdot M \quad \text{eval}(e_1, t) = l \quad \text{eval}(e_3, t) = t_c \quad \text{eval}(e_2, t) = t_m \quad M'_r = M_r[t_c \mapsto M_r(t_c) :: (o, t_m)]}{\langle B_s, P \mid B_r \cdot t_r \cdot M_r \mid P_1 \rangle_l \xrightarrow{o@e_3} \langle B_s, B \cdot t \cdot M \mid B_r \cdot t_r \cdot M'_r \mid P_1 \rangle_l} [\text{DCC}|_{\text{INSEND}}]$$

and

$$\frac{P[x'/x] = o@e_1[x'/x](e_2[x'/x]) \text{ to } e_3[x'/x]; B[x'/x] \cdot t' \cdot M \quad \text{eval}(e_1[x'/x], t) = l \quad \text{eval}(e_3[x'/x], t) = t_c \quad \text{eval}(e_2[x'/x], t) = t_m \quad M'_r = M_r[t_c \mapsto M_r(t_c) :: (o, t_m)]}{\langle B_s, P[x'/x] \mid B_r \cdot t_r \cdot M_r \mid P_1 \rangle_l \xrightarrow{o@e_3[x'/x]} \langle B_s, B[x'/x] \cdot t' \cdot M \mid B_r \cdot t_r \cdot M'_r \mid P_1 \rangle_l} [\text{DCC}|_{\text{INSEND}}]$$

Case $[\text{DCC}|_{\text{INSTART}}]$

We consider the case where $P = ?@e_1(e_2); B \cdot t \cdot M$ and expressions e_1 and e_2 contain x . From Definition 38 we know that $\text{eval}(e_1, t) = \text{eval}(e_1[x'/x], t')$. Similarly e_2 and $e_2[x'/x]$ and enjoy the same property when evaluated respectively on t and t' .

Below we describe the case in which P moves and $P[x'/x]$ mimics it. The other case, for $P[x'/x]$ that reduces and P that mimics it, follows the same structure. We assume the start behaviour $B_s = ?(y); B_c$.

E.4. Proof of Compilation

$$\frac{P = ?@e_1(e_2); B \cdot t \cdot M \quad \text{eval}(e_1, t) = l \quad P_1 = B_c \cdot t_\perp \triangleleft (y, \text{eval}(e_2, t)) \cdot \emptyset}{\langle !(y); B_c, P \mid Q \rangle_l \xrightarrow{?(e_2)} \langle !(y); B_c, P_1 \mid B \cdot t \cdot M \mid Q \rangle_l} \quad \text{[DCC]}_{\text{INSTART}}$$

and

$$\frac{P' = ?@e_1[x'/x](e_2[x'/x]); B[x'/x] \cdot t' \cdot M \quad \text{eval}(e_1[x'/x], t') = l \quad P_1 = B_c \cdot t_\perp \triangleleft (y, \text{eval}(e_2[x'/x], t')) \cdot \emptyset}{\langle !(y); B, P' \mid Q \rangle_l \xrightarrow{?(e_2[x'/x])} \langle !(y); B', P_1 \mid B[x'/x] \cdot t' \cdot M \mid Q \rangle_l} \quad \text{[DCC]}_{\text{INSTART}}$$

□

Lemma 30 (DCC Network Variable Renaming). *Let S and S' be two DCC network such that $S = \langle B_s, P \mid Q \rangle_l \mid S_*$ and $S' = \langle B_s, P[x'/x] \mid Q \rangle_l \mid S_*$ then $S \xrightarrow{\lambda} \langle B_s, P'' \mid Q' \rangle_l \mid S'_* \iff S' \xrightarrow{\lambda[x'/x]} \langle B_s, P''[x'/x] \mid Q' \rangle_l \mid S'_*$.*

Proof. The proof is by induction on the derivation of S . Below we consider $t' = t \triangleleft (x', x(t)) \triangleleft (x, t_\perp)$. We report the most interesting cases.

Case $\text{[DCC]}_{\text{SEND}}$

We know that $P = o@e_1(e_2) \text{ to } e_3; B \cdot t \cdot M$. We consider the case in which all expressions e_1 , e_2 , and e_3 contain x . Let suppose that $\text{eval}(e_1, t) = l'$ and $\text{eval}(e_3, t) = t_c$ then, since $\text{[DCC]}_{\text{SEND}}$ applies, we know that $S_* \equiv \langle B'_s, B_r \cdot t_r \cdot M_r \mid P_2 \rangle_{l'}$ where $t_c \in \text{dom}(M_r)$.

From Definition 38 we know that $\text{eval}(e_1, t) = \text{eval}(e_1[x'/x], t')$. Similarly the couples e_2 and $e_2[x'/x]$ and e_3 and $e_3[x'/x]$ enjoy the same property when evaluated respectively on t and t' .

We proceed proving that S reduces and S' can mimic it. The other case, for S' that reduces and S that mimics it, follows the same structure.

$$\frac{P = o@e_1(e_2) \text{ to } e_3; B \cdot t \cdot M \quad \text{eval}(e_1, t) = l' \quad \text{eval}(e_3, t) = t_c \quad \text{eval}(e_2, t) = t_m \quad M'_r = M_r[t_c \mapsto M_r(t_c) :: (o, t_m)]}{\langle B_s, P \mid P_1 \rangle_l \mid \langle B'_s, B_r \cdot t_r \cdot M_r \mid P_2 \rangle_{l'} \xrightarrow{o@e_3} \langle B_s, B \cdot t \cdot M \mid P_1 \rangle_l \mid \langle B'_s, B_r \cdot t_r \cdot M'_r \mid P_2 \rangle_{l'}} \quad \text{[DCC]}_{\text{SEND}}$$

and

$$\frac{
 \begin{array}{l}
 P[x'/x] = o@e_1[x'/x](e_2[x'/x]) \text{ to } e_3[x'/x]; B[x'/x] \cdot t \cdot M \\
 \text{eval}(e_1[x'/x], t) = l' \quad \text{eval}(e_3[x'/x], t) = t_c \\
 \text{eval}(e_2[x'/x], t) = t_m \quad M'_r = M_r[t_c \mapsto M_r(t_c) :: (o, t_m)]
 \end{array}
 }{
 \langle B_s, P[x'/x] \mid P_1 \rangle_l \mid \langle B'_s, B_r \cdot t_r \cdot M_r \mid P_2 \rangle_{l'} \xrightarrow{o@e_3[x'/x]} \langle B_s, B[x'/x] \cdot t' \cdot M \mid P_1 \rangle_l \mid \langle B'_s, B_r \cdot t_r \cdot M'_r \mid P_2 \rangle_{l'}
 } \text{ [DCC|SEND] }$$

Case $[\text{DCC}|_{\text{START}}]$

We know that $P = ?@e_1(e_2); B \cdot t \cdot M$. We consider the case in which both expressions e_1 and e_2 contain x . Let suppose that $\text{eval}(e_1, t) = l'$. Since $[\text{DCC}|_{\text{SEND}}]$ applies, we know that $S_* \equiv \langle !(y); B_c, P_1 \rangle_{l'}$.

From Definition 38 we know that $\text{eval}(e_1, t) = \text{eval}(e_1[x'/x], t')$. Similarly also e_2 and $e_2[x'/x]$ enjoy the same property when evaluated respectively on t and t' .

We proceed proving that S reduces and S' can mimic it. The other case, for S' that reduces and S that mimics it, follows the same structure.

$$\frac{
 \begin{array}{l}
 P = ?@e_1(e_2); B \cdot t \cdot M \quad \text{eval}(e_1, t_1) = l \\
 P_2 = B \cdot t_\perp \triangleleft (y, \text{eval}(e_2, t)) \cdot \emptyset
 \end{array}
 }{
 \langle !(y); B_c, P_1 \rangle_{l'} \mid \langle B'_s, P \mid Q \rangle_l \xrightarrow{?(e_2)} \langle !(y); B_c, P_2 \mid P_1 \rangle_{l'} \mid \langle B'_s, B \cdot t \cdot M \mid Q \rangle_l
 } \text{ [DCC|START] }$$

$$\frac{
 \begin{array}{l}
 P[x'/x] = ?@e_1[x'/x](e_2[x'/x]); B[x'/x] \cdot t' \cdot M \quad \text{eval}(e_1[x'/x], t_1) = l' \\
 P_2 = B \cdot t_\perp \triangleleft (y, \text{eval}(e_2[x'/x], t')) \cdot \emptyset
 \end{array}
 }{
 \langle !(y); B_c, P_1 \rangle_{l'} \mid \langle B'_s, P[x'/x] \mid Q \rangle_l \xrightarrow{?(e_2[x'/x])} \langle !(y); B_c, P_2 \mid P_1 \rangle_{l'} \mid \langle B'_s, B[x'/x] \cdot t' \cdot M \mid Q \rangle_l
 } \text{ [DCC|START] }$$

Case $[\text{DCC}|_{\text{EQ}}]$ and $[\text{DCC}|_{\text{SPAR}}]$

In both cases the thesis follows from the application of the induction hypothesis and Lemma 29.

□

We report below the statement of Theorem 8, enriched with annotation on the transitions of D, C .

Theorem 8 (*Compilation Operational Correspondence*)

Let C be a composition of endpoint choreographies such that $\Gamma \vdash D, C$. Then we have that:

E.4. Proof of Compilation

1. (Completeness) $D, C \xrightarrow{\beta} D', C'$ implies i) $\boxed{D, C}^\Gamma \rightarrow^+ \boxed{D', C'}^{\Gamma'}$ for some Γ' such that ii) $\Gamma' \vdash D', C'$.
2. (Soundness) $\boxed{D, C}^\Gamma \rightarrow^* S$ implies i) $D, C \rightarrow^* D', C'$ and ii) $S \rightarrow^* \boxed{D', C'}^{\Gamma'}$ for some D', C' , and Γ' such that iii) $\Gamma' \vdash D', C'$.

As a convention in the following we use the shortcuts $t_p = D(p).st$ and $M_p = D(p).que$ for p process in D .

Proof (Completeness). We proceed by induction on the derivation of $D, C \rightarrow D', C'$.

Case $[^C|_{SEND}]$

We know that

- $C \equiv C_p \mid C_c$ with $C_p = k : p[A].e \rightarrow B.o; C'_p$;
- $D, C \xrightarrow{\beta} D', C'$ with $[^C|_{SEND}]$ being the last applied Rule, where $\beta = k : p[A].e \rightarrow B.o$. $C' = C'_p \mid C_c$; $D' = D[q \mapsto (t_q, M_q[t_c \mapsto \tilde{m} :: (o_j, t_m)])]$ by Rule $[^D|_{SEND}]$, $t_c = \mathbf{k.A.B}(t_p)$ and $t_m = \text{eval}(e, t_p)$

Since $[^C|_{SEND}]$ applies and there exists $\Gamma \vdash D, C$, then there exists a queue associated with a process q in D where the message is delivered.

We have two cases, whether the receiving process q is in the same location of the sender p or not. Let $p \in D(l)$

Case $q \in D(l)$

From Definition 23 we have that $\boxed{D, C}^\Gamma \equiv S \mid S_c$ where

- $S = \left\langle \boxed{C_c|_l}^\Gamma, P \mid Q \mid R \right\rangle_l$
- $P = o@k.B.l(e) \text{ to } k.A.B; \boxed{C'_p}^\Gamma \cdot t_p \cdot M_p$
- $Q = \boxed{C_c|_q}^\Gamma \cdot t_q \cdot M_q$
- $R = \prod_{r \in D(l)/\{p,q\}} \boxed{C_c|_r}^\Gamma \cdot t_r \cdot M_r$
- $S_c = \prod_{l' \in \Gamma/\{l\}} \left\langle \boxed{C_c|_{l'}}^\Gamma, \prod_{s \in D(l')} \boxed{C_c|_s}^\Gamma \cdot t_s \cdot M_s \right\rangle_{l'}$

In this case $\boxed{D, C}^\Gamma$ can mimic D, C applying Rules $[^{DCC}|_{EQ}]$, $[^{DCC}|_{SPAR}]$, and $[^{DCC}|_{INSEND}]$ where $S \mid S_c \rightarrow S' \mid S_c$ with $[^{DCC}|_{SPAR}]$ and $S \rightarrow S'$ with

$$\frac{o@k.\mathbf{B.l}(e) \text{ to } k.\mathbf{A.B}; \boxed{C'_p}^\Gamma \cdot t_p \cdot M_p \quad \text{eval}(k.\mathbf{B.l}, t_p) = l \quad \text{eval}(k.\mathbf{A.B}, t_p) = t_c \quad \text{eval}(e, t_p) = t_m \quad M'_q = M_q[t_c \mapsto M_q(t_c) :: (o, t_m)]}{\left\langle \boxed{C_c|_l}^\Gamma, P \mid Q \mid R \right\rangle_l \xrightarrow{k: p[\mathbf{A}].e \rightarrow \mathbf{B.o}} \left\langle \boxed{C_c|_l}^\Gamma, P' \mid Q' \mid R \right\rangle_l} [\text{DCC}|_{\text{INSEND}}]$$

where $P' = \boxed{C'_p}^\Gamma \cdot t_p \cdot M_p$ and $Q' = \boxed{C_c|_q}^\Gamma \cdot t_q \cdot M'_q$.

From the hypothesis we know that $\Gamma = \Gamma_1, b[k]_{\mathbf{B}}^{\mathbf{A}}: T, k[\mathbf{A}]: !\mathbf{B.o}(U); T'$ (as $\Gamma \vdash D, C$).

By Rules $[\text{T}|_{\text{DC}}]$ and $[\text{T}|_{\text{SEND}}]$ we can find $\Gamma' = \Gamma_1, b[k]_{\mathbf{B}}^{\mathbf{A}}: T; ?\mathbf{A.o}(U), k[\mathbf{A}]: T'$ such that $\Gamma' \vdash D', C'$.

From Rule $[\text{D}|_{\text{SEND}}]$ and 23, we have that

$$\boxed{D', C'}^\Gamma \equiv \overbrace{\left\langle \boxed{C_c|_l}^\Gamma, \overbrace{\boxed{C'_p}^\Gamma \cdot t_p \cdot M_p}^{P'} \mid \overbrace{\boxed{C_c|_q}^\Gamma \cdot t_q \cdot M'_q}^{Q'} \mid \overbrace{\prod_{r \in D'(l)/\{p,q\}} \boxed{C_c|_r}^\Gamma \cdot t_r \cdot M_r}^R \right\rangle_l}^{S'}} \underbrace{\left| \prod_{l' \in \Gamma/\{l\}} \left\langle \boxed{C_c|_{l'}}^\Gamma, \prod_{s \in D(l')} \boxed{C_c|_s}^\Gamma \cdot t_s \cdot M_s \right\rangle_{l'} \right|}_{S_c}$$

Therefore $\boxed{D', C'}^\Gamma = S' \mid S_c$, by construction of Γ' we can apply Lemma 27 for which $\boxed{D', C'}^\Gamma = \boxed{D', C'}^{\Gamma'}$ and therefore $S' \mid S_c = \boxed{D', C'}^{\Gamma'}$.

Case $q \notin D(l)$

Similar to **Case** $q \in D(l)$ except the last applied Rule for $\boxed{D, C}^\Gamma \rightarrow \boxed{D', C'}^{\Gamma'}$ is $[\text{DCC}|_{\text{SEND}}]$.

Case $[\text{C}|_{\text{COM}}]$

Does not apply since C is a composition of endpoint choreographies by hypothesis.

Case $[\text{C}|_{\text{RECV}}]$

We know that

- $C \equiv C_q \mid C_c$ with $C_q = k : \mathbf{A} \rightarrow \mathbf{q}[\mathbf{B}].\{o_i(x_i); C_i\}_{i \in I}$

E.4. Proof of Compilation

- $D, C \xrightarrow{\beta} D, C'$ with Rule $[C]_{\text{RECV}}$ where $\beta = k: A \rightsquigarrow \mathbf{q}[B].o_j(x_j)$, $C' \equiv C_j \mid C_c$, and $D' = D[\mathbf{q} \mapsto (t_q \triangleleft (x_j, t_m), M_q[t_c \mapsto \tilde{m}])]$ where $t_c = \mathbf{k.A.B}(t_q)$ and $M_q(t_c) = (o_j, t_m) :: \tilde{m}$;

Let $\mathbf{q}@l \in \Gamma$, from the Definition 23 we have

$\boxed{D, C}^\Gamma \equiv S \mid S_c$ where

- $S = \left\langle \boxed{C_c|l}^\Gamma, Q \mid R \right\rangle_l$
- $Q = \sum_{i \in I} [o_i(x_i) \text{ from } \mathbf{k.A.B}] \{ \boxed{C_i}^\Gamma \} \cdot t_q \cdot M_q$
- $\prod_{r \in D(l)/\{q\}} \boxed{C_c|r}^\Gamma \cdot t_r \cdot M_r$
- $S_c = \prod_{\nu \in \Gamma/\{l\}} \left\langle \boxed{C_c|\nu}^\Gamma, \prod_{s \in D(\nu)} \boxed{C_c|s}^\Gamma \cdot t_s \cdot M_s \right\rangle_\nu$

In this case $\boxed{D, C}^\Gamma$ can mimic D, C applying Rules $[\text{DCC}|_{\text{EQ}}]$, $[\text{DCC}|_{\text{SPAR}}]$, $[\text{DCC}|_{\text{PPAR}}]$, and $[\text{DCC}|_{\text{CHOICE}}]$ where

$$\frac{\frac{j \in I \quad t_c = \text{eval}(\mathbf{k.A.B}, t_q) \quad M_q = (o_j, t_m) :: \tilde{m}}{\sum_{i \in I} [o_i(x_i) \text{ from } \mathbf{k.A.B}] \{ \boxed{C_i}^\Gamma \} \cdot t_q \cdot M_q \rightarrow \boxed{C_j}^\Gamma \cdot t_q \triangleleft (x_j, t_m) \cdot M_q[t_c \mapsto \tilde{m}]} \begin{array}{l} [\text{DCC}|_{\text{CHOICE}}] \\ [\text{DCC}|_{\text{PPAR}}] \end{array}}{\frac{\langle C_c|l, Q \mid R \rangle_l \rightarrow \langle C_c|l, \boxed{C_j}^\Gamma \cdot t_q \triangleleft (x_j, t_m) \cdot M_q[t_c \mapsto \tilde{m}] \mid R \rangle_l}{S \mid S_c \rightarrow S' \mid S_c} [\text{DCC}|_{\text{SPAR}}]}$$

Let $Q' = \boxed{C_j}^\Gamma \cdot t_q \triangleleft (x_j, t_m) \cdot M_q[t_c \mapsto \tilde{m}]$.

From the hypothesis we know that $\Gamma = \Gamma_1, b[k]_B^A: ?A.o_j(U_j); T, k[B]: ?A.\{o_i(U_i); T_i\}_{i \in I}$ and we can find $\Gamma' = \Gamma_1, b[k]_B^A: T, k[B]: T_j$ such that $\Gamma' \vdash D', C'$.

From Rule $[D]_{\text{RECV}}$ and Definition 23, we have that, let $t'_q = t_q \triangleleft (x_j, t_m)$ and $M'_q = M_q[t_c \mapsto \tilde{m}]$

$$\boxed{D', C'}^\Gamma \equiv \left\langle \boxed{C_c|l}^\Gamma, \overbrace{\left\langle \boxed{C_j|q}^\Gamma \cdot t'_q \cdot M'_q \mid \prod_{r \in D'(l)/\{q\}} \boxed{C_c|r}^\Gamma \cdot t_r \cdot M_r \right\rangle_l}^{S'} \right\rangle_l \underbrace{\prod_{\nu \in \Gamma/\{l\}} \left\langle \boxed{C_c|\nu}^\Gamma, \prod_{s \in D(\nu)} \boxed{C_c|s}^\Gamma \cdot t_s \cdot M_s \right\rangle_\nu}_{S_c}$$

Therefore $\boxed{D', C'}^\Gamma = S' \mid S_c$, by construction of Γ' we can apply Lemma 27 for which $\boxed{D', C'}^\Gamma = \boxed{D', C'}^{\Gamma'}$ and therefore $S' \mid S_c = \boxed{D', C'}^{\Gamma'}$.

Case $\lfloor^C_{\text{START}}\rfloor$

Does not apply since C is a composition of endpoint choreographies by hypothesis.

Case $\lfloor^C_{\text{PSTART}}\rfloor$

We know that

- $C \equiv C_r \mid C_a \mid C_c$ where, let $\tilde{l}: G\langle A \mid \tilde{B} \mid \tilde{B} \rangle \in \Gamma$
- $C_r = \mathbf{req} \ k : p[A] \Leftrightarrow \tilde{l}.B; C'_r$
- let $l_1.B_1, \dots, l_n.B_n = \tilde{l}.B$, $C_a = \prod_{i=1}^n \mathbf{acc} \ k : l_i.q_i[B_i]; C_{q_i}$

We can apply Rules $\lfloor^C_{\text{PAR}}\rfloor$ and $\lfloor^C_{\text{EQ}}\rfloor$ and lastly Rule $\lfloor^C_{\text{PSTART}}\rfloor$ such that

$$\frac{\begin{array}{c} i \in \{1, \dots, n\} \quad \#k' \quad \{\tilde{l}.B\} = \uplus_i \{l_i.B_i\} \quad \#\tilde{r} \quad \{\tilde{r}\} = \bigcup_i \{\tilde{r}_i\} \\ p \in D(l) \quad \delta = \mathbf{start} \ k' : l.p[A], l_1.r_1[B_1], \dots, l_n.r_n[B_n] \quad D, \delta \blacktriangleright D' \end{array}}{D, C_r \mid C_a \xrightarrow{\tau} D', C'_r[k'/k] \mid \prod_{i=1}^n (C_{q_i}[k'/k][r_i/q_i]) \mid C_a} \lfloor^C_{\text{PSTART}}\rfloor$$

and

$$D, C_r \mid C_a \mid C_c \rightarrow D', C'_r[k'/k] \mid \prod_i (C_{q_i}[k'/k][r_i/q_i]) \mid C_a \mid C_c$$

thus $C' = C'_r[k'/k] \mid \prod_i (C_{q_i}[k'/k][r_i/q_i]) \mid C_a \mid C_c$

From the hypothesis we know that $\Gamma \vdash D, C$ and therefore that $\Gamma = \Gamma_1, \tilde{l}: G\langle A \mid \tilde{B} \mid \tilde{B} \rangle$.

We can find $\Gamma' = \Gamma, \text{init}(k', (p[A], \tilde{q}[B]), G)$ and $\Gamma' \vdash D', C'$.

Remark 7. We have two cases for, let $p@l \in \Gamma$, whether $l \in \{\tilde{l}\}$ or not. For a clearer treatment of the case we proceed considering that $l \notin \{\tilde{l}\}$ (i.e., no service process is created in the same location — service — of the requester p). The other case follows the same structure of $l \notin \{\tilde{l}\}$ although the service located at l has $\boxed{C_a|l}^\Gamma$ as start behaviour and $\boxed{D, C'}^\Gamma$ applies Rule $\lfloor^{DCC}_{\text{INSTART}}\rfloor$ in place of the $\lfloor^{DCC}_{\text{START}}\rfloor$ for starting the DCC process located at l .

Henceforth we proceed analysing the case for $l \notin \{\tilde{l}\}$.

From Definition 23 we have

E.4. Proof of Compilation

$$\boxed{D', C'}^{\Gamma'} = \left\langle \boxed{C_c|l}^{\Gamma'}, P'' \mid R' \right\rangle_l \mid \prod_{i=1}^n \langle Q''_i, Q_i^* \mid R'_{l_i} \rangle_{l_i} \mid S'_c$$

In the following, we use the shortcuts $t_s^* = D'(s).st$ and $M_s^* = D'(s).que$ for s process in D' .

- $P'' = \boxed{C'_r[k'/k]}^{\Gamma'} \cdot t_p^* \cdot M_p^*$
- $R' = \prod_{p' \in D(l)/\{p\}} \boxed{C_c|p'}^{\Gamma'} \cdot t_{p'}^* \cdot M_{p'}^*$
- $Q''_i = \text{accept}(k, B_i, G\langle A|\tilde{B}|\tilde{B} \rangle); \boxed{C_{q_i}}^{\Gamma'}$
- $Q_i^* = \boxed{C_{q_i}[k'/k][r_i/q_i]}^{\Gamma'} \cdot t_{q_i}^* \cdot M_{q_i}^*$
- $R'_{l_i} = \prod_{s \in D(l_i)} \boxed{C_c|s}^{\Gamma'} \cdot t_s^* \cdot M_s^*$
- $S'_c = \prod_{l' \in \Gamma/\{l, \tilde{l}\}} \left\langle \boxed{C_c|l'}^{\Gamma'}, \prod_{s' \in D(l')} \boxed{C_c|s'}^{\Gamma'} \cdot t_{s'}^* \cdot M_{s'}^* \right\rangle_{l'}$

From Rule $[D]_{\text{START}}$ we know that

$$\underline{k}'(t_p^*) = \underline{k}'(t_{q_1}^*) = \dots = \underline{k}'(t_{q_n}^*) = t_{k'}$$

for some $t_{k'}$ session descriptor of session k' .

We proceed by proving that we can reduce $\boxed{D, C}^{\Gamma} \rightarrow^+ S$.

From Definition 23 we have

$$\boxed{D, C}^{\Gamma} \equiv \left\langle \boxed{C_c|l}^{\Gamma}, P \mid R \right\rangle_l \mid \prod_{i=1}^n \langle Q_i, R_{l_i} \rangle_{l_i} \mid S_c$$

where

$$P = \text{start}(k, (l.A, \tilde{l}.B)); \boxed{C'_r}^{\Gamma} \cdot t_p \cdot M_p =$$

$$\bullet \left(\begin{array}{l} \odot \underline{\mathbf{k.I.I}} = l_{\Gamma}; \\ \odot_{\mathbf{I} \in \{A, \tilde{B}\}} \left(\text{cq}(\underline{\mathbf{k.I.A}}); ?@ \underline{\mathbf{k.I.I}}(\mathbf{k}); \text{sync}(\underline{\mathbf{k}}) \text{ from } \underline{\mathbf{k.I.A}} \right); \\ \odot_{\mathbf{I} \in \{\tilde{B}\}} \text{start}@ \underline{\mathbf{k.I.I}}(\mathbf{k}) \text{ to } \underline{\mathbf{k.A.I}}; \boxed{C'_r}^{\Gamma} \end{array} \right) \cdot t_p \cdot M_p$$

$$\begin{aligned}
 & \text{!}(\underline{\mathbf{k}}); \bigodot_{\substack{\mathbf{I} \in \{\underline{\mathbf{A}}, \tilde{\mathbf{B}}\} / \{\mathbf{B}_i\} \\ \text{sync@}\underline{\mathbf{k}}.\underline{\mathbf{A}}.\mathbf{I}(\underline{\mathbf{k}}) \text{ to } \underline{\mathbf{k}}.\mathbf{B}_i.\underline{\mathbf{A}}; \\ \text{start}(\underline{\mathbf{k}}) \text{ from } \underline{\mathbf{k}}.\underline{\mathbf{A}}.\mathbf{B}_i; \overline{C_{q_i}}^\Gamma}} \text{cq}(\underline{\mathbf{k}}.\mathbf{I}.\mathbf{B}_i); \\
 \bullet \quad Q_i &= \text{accept}(k, \mathbf{B}_i, G\langle \underline{\mathbf{A}} | \tilde{\mathbf{B}} | \tilde{\mathbf{B}} \rangle); \overline{C_{q_i}}^\Gamma = \\
 \bullet \quad R &= \prod_{p' \in D(l)/\{p\}} \overline{C_c|_{p'}}^\Gamma \cdot t_{p'} \cdot M_{p'} \\
 \bullet \quad R_{l_i} &= \prod_{s \in D(l_i)} \overline{C_c|_s}^\Gamma \cdot t_s \cdot M_s \\
 \bullet \quad S_c &= \prod_{l' \in \Gamma/\{l, \bar{l}\}} \left\langle \overline{C_c|_{l'}}^\Gamma, \prod_{s' \in D(l')} \overline{C_c|_{s'}}^\Gamma \cdot t_{s'} \cdot M_{s'} \right\rangle_{l'}
 \end{aligned}$$

$\overline{D, C}^\Gamma$ can mimic D, C with the following sequence of reductions. Note that we make use of renaming on (*accept*) terms in Q_1, \dots, Q_n and variable renaming on P (as of Definition 38) to align the evolution of $\overline{D, C}^\Gamma$ with the evolution of D, C , in which k has been replaced with the fresh name k' .

Therefore we take $S_0^* \sim \overline{D, C}^\Gamma$

$$S_0^* = \left\langle \overline{C_c|_l}^\Gamma, P[\underline{\mathbf{k}}'/\underline{\mathbf{k}}] \mid R \right\rangle_l \mid \prod_{i=1}^n \langle Q_i[\underline{\mathbf{k}}'/\underline{\mathbf{k}}], R_{l_i} \rangle_{l_i} \mid S_c$$

$$\begin{array}{c}
 \begin{array}{c}
 \overline{C_c|_l}^\Gamma \xrightarrow{[DCC|_{EQ}] \quad [DCC|_{SPAR}] \quad [DCC|_{PPAR}] \quad [DCC|_{ASSIGN}]} \\
 \textcircled{2.1} \xrightarrow{[DCC|_{EQ}] \quad [DCC|_{SPAR}] \quad [DCC|_{PPAR}] \quad [DCC|_{CQ}]} \\
 \textcircled{2.2} \xrightarrow{[DCC|_{EQ}] \quad [DCC|_{SPAR}] \quad [DCC|_{START}]} \\
 \textcircled{2.3} \xrightarrow{[DCC|_{EQ}] \quad [DCC|_{SPAR}] \quad [DCC|_{PPAR}] \quad [DCC|_{CQ}]} \\
 \textcircled{2.4} \xrightarrow{[DCC|_{EQ}] \quad [DCC|_{SPAR}] \quad [DCC|_{SEND}]} \\
 \textcircled{2.5} \xrightarrow{[DCC|_{EQ}] \quad [DCC|_{SPAR}] \quad [DCC|_{RECV}]} \\
 \textcircled{3.1} \xrightarrow{[DCC|_{EQ}] \quad [DCC|_{SPAR}] \quad [DCC|_{SEND}]} \\
 \textcircled{3.2} \xrightarrow{[DCC|_{EQ}] \quad [DCC|_{SPAR}] \quad [DCC|_{RECV}]}
 \end{array}
 \end{array}
 \begin{array}{c}
 \left. \begin{array}{l}
 \textcircled{1} \\
 n+1 \text{ times} \\
 \textcircled{2} \\
 n \text{ times} \\
 n \text{ times} \\
 \textcircled{3} \\
 n \text{ times}
 \end{array} \right\} \rightarrow S_1^*
 \end{array}$$

We briefly comment the numbered transitions.

- In $\textcircled{1}$ $P[k'/k]$ proceeds to store (for $n+1$ times, l plus $l_i, i \in \{1, \dots, n\}$) the locations of all roles under $\underline{\mathbf{k}}'$.
- In $\textcircled{2}$, for each location $l_i, i \in \{1, \dots, n\}$ (for each service process):

E.4. Proof of Compilation

- P creates its receiving queue for the service process (2.1);
 - in (2.2) we apply Rule $[{}^{\text{DCC}}|_{\text{EQ}}]$ such that we rename
 - P synchronises with the service at location l_i starting ($[{}^{\text{DCC}}|_{\text{START}}]$) a new service process;
 - in (2.3) the service process creates its own queues for all other roles in the session (hence n times);
 - in (2.4) the service process sends the correlation values to P ;
 - finally P receives the message in (2.5).
- In (3) for each service process (n times) (3.1) the starter sends a message to the service process to start the session and (3.2) the addressee receives it.

Finally we have

$$S_1^* = \left\langle \overline{[C_c]_l}^\Gamma | P' | R \right\rangle_l \mid \prod_{i=1}^n \langle Q_i[\underline{\mathbf{k}}'/\underline{\mathbf{k}}], Q'_i | R_{l_i} \rangle_{l_i} \mid S_c$$

where

- $P' = \overline{[C'_r]^\Gamma}[\underline{\mathbf{k}}'/\underline{\mathbf{k}}] \cdot t'_p \cdot M'_p$, and
- $Q'_i = \overline{[C'_{q_i}]^\Gamma}[\underline{\mathbf{k}}'/\underline{\mathbf{k}}] \cdot t_{k'} \cdot M_{q_i}$

From the transitions presented above we know that there exists $t'_{k'}$ such that $t'_p = t_p \triangleleft (\underline{\mathbf{k}}', t'_{k'})$, where $t'_{k'}$ is a session descriptor for session k' (i.e., it contains all the locations and correlations keys used by the processes in session k'). In this case, we take $t'_k = t_{k'}$ obtained from the derivation $D, C \xrightarrow{\tau} D', C'$.

Similarly, M'_p and M_{q_1}, \dots, M_{q_n} contain the necessary (empty) queues to support communication in session k' .

$$M'_p = M_p[\underline{\mathbf{k}}'.\mathbf{B}_1.\mathbf{A}(t_{k'}) \mapsto \varepsilon] \dots [\underline{\mathbf{k}}'.\mathbf{B}_n.\mathbf{A}(t_{k'}) \mapsto \varepsilon]$$

and

$$M_{q_i} = \emptyset \quad \overline{[\underline{\mathbf{k}}'.\mathbf{A}.\mathbf{B}_i(t_k) \mapsto \varepsilon]} \overline{[\underline{\mathbf{k}}'.\mathbf{B}_1.\mathbf{B}_i(t_k) \mapsto \varepsilon]} \dots \overline{[\underline{\mathbf{k}}'.\mathbf{B}_{i-1}.\mathbf{B}_i(t_k) \mapsto \varepsilon]} \dots \\ \dots \overline{[\underline{\mathbf{k}}'.\mathbf{B}_{i+1}.\mathbf{B}_i(t_k) \mapsto \varepsilon]} \dots \overline{[\underline{\mathbf{k}}'.\mathbf{B}_n.\mathbf{B}_i(t_k) \mapsto \varepsilon]}$$

We proceed with the proof taking $S \sim S_1^*$ as S is simply the renaming of $\underline{\mathbf{k}}'$ to $\underline{\mathbf{k}}$ on start behaviours $Q_i, i \in \{1, \dots, n\}$ (trivially $Q_i[\underline{\mathbf{k}}'/\underline{\mathbf{k}}][\underline{\mathbf{k}}/\underline{\mathbf{k}}'] = Q_i$)

$$S = \left\langle \boxed{C_{cl}}^\Gamma \mid P' \mid R \right\rangle_l \mid \prod_{i=1}^n \langle Q_i, Q'_i \mid R_{l_i} \rangle_{l_i} \mid S_c$$

We now proceed to prove that $\boxed{D, C}^\Gamma \rightarrow^+ \boxed{D', C'}^{\Gamma'}$, i.e. that $\boxed{D', C'}^{\Gamma'} = S$ with $\Gamma' \vdash D', C'$.

We prove that

$$\overbrace{\left\langle \boxed{C_{cl}}^{\Gamma'}, P'' \mid R' \right\rangle_l \mid \prod_{i=1}^n \langle Q''_i, Q_i^* \mid R'_{l_i} \rangle_{l_i} \mid S'_c}^{\boxed{D', C'}^{\Gamma'}} = \overbrace{\left\langle \boxed{C_{cl}}^\Gamma \mid P' \mid R \right\rangle_l \mid \prod_{i=1}^n \langle Q_i, Q'_i \mid R_{l_i} \rangle_{l_i} \mid S_c}^S$$

- $\boxed{C_{cl}}^\Gamma = \boxed{C_{cl}}^{\Gamma'}$ as $\Gamma|_{\text{locs}} = \Gamma'|_{\text{locs}}$ by construction;
- $P'' = P'$ is proven by

$$\boxed{C'_r[k'/k]}^{\Gamma'} \cdot t_p^* \cdot M_p^* = \boxed{C'_r[\underline{k}'/\underline{k}]}^\Gamma \cdot t'_p \cdot M'_p$$

which holds as

- $\boxed{C'_r[k'/k]}^{\Gamma'} = \boxed{C'_r[\underline{k}'/\underline{k}]}^\Gamma$ since
 - Γ' does not contain any new process used in C'_r ;
 - by renaming, and
 - Definition 38.
 - $t_p^* = t'_p$ by construction and Rule $[D]_{\text{START}}$;
 - $M_p^* = M'_p$ by construction and Rule $[D]_{\text{START}}$.
- $Q_i^* = Q'_i$ proven by

$$\boxed{C_{q_i}[k'/k][r_i/q_i]}^{\Gamma'} \cdot t_{q_i}^* \cdot M_{q_i}^* = \boxed{C_{q_i}[\underline{k}'/\underline{k}]}^\Gamma \cdot t_{q_i}' \cdot M_{q_i}$$

whose proof of equivalence follows that of $P'' = P'$, except that Γ' contains the location of the process (r_i) used in $C_{q_i}[k'/k][r_i/q_i]$.

- $Q''_i = Q_i$ proven by

$$\text{accept}(k, B_i, G\langle A|\tilde{B}|\tilde{B} \rangle); \boxed{C_{q_i}}^{\Gamma'} = \text{accept}(k, B_i, G\langle A|\tilde{B}|\tilde{B} \rangle); \boxed{C_{q_i}}^\Gamma$$

which holds as $\boxed{C_{q_i}}^{\Gamma'} = \boxed{C_{q_i}}^\Gamma$ because Γ and Γ' contain the same service typings.

E.4. Proof of Compilation

- $R' = R$ is proven by

$$\prod_{p' \in D(l)/\{p\}} \boxed{C_c|_{p'}}^{\Gamma'} \cdot t_{p'}^* \cdot M_{p'}^* = \prod_{p' \in D(l)/\{p\}} \boxed{C_c|_{p'}}^{\Gamma} \cdot t_{p'} \cdot M_{p'}$$

for which

- i) $\boxed{C_c|_{p'}}^{\Gamma'} = \boxed{C_c|_{p'}}^{\Gamma}$ as Γ' does not contain any new process used in C_c .
 - ii) $t_{p'}^* = t_{p'}$ unchanged by the reduction of D, C and $\boxed{D, C}^{\Gamma}$;
 - iii) $M_{p'}^* = M_{p'}$ unchanged by the reduction of D, C and $\boxed{D, C}^{\Gamma}$.
- $R'_{l_i} = R_{l_i}$ whose proof follows that of $R' = R$.
 - $S'_c = S_c$ following the proof of $\boxed{C_c|_l}^{\Gamma} = \boxed{C_c|_l}^{\Gamma'}$ and $R'_{l_i} = R_{l_i}$.

Case \boxed{C}_{COND}

We have $C = C_p \mid C_c$ where $C_p = \text{if } p.e \{C_1\} \text{ else } \{C_2\}$. Let $p@l \in \Gamma$ and

- $P = \text{if } e \{ \boxed{C_1}^{\Gamma} \} \text{ else } \{ \boxed{C_2}^{\Gamma} \} \cdot t_p \cdot M_p$;
- $R = \prod_{r \in D(l)/\{p\}} \boxed{C_c|_r}^{\Gamma} \cdot t_r \cdot M_r$
- $S_c = \prod_{l' \in \Gamma/\{l\}} \left\langle \boxed{C_c|_{l'}}^{\Gamma}, \prod_{r \in D(l')} \boxed{C_c|_r}^{\Gamma} \cdot t_r \cdot M_r \right\rangle_{l'}$

From Definition 23 we have,

$$\boxed{D, C}^{\Gamma} \equiv \left\langle \boxed{C_c|_l}^{\Gamma}, P \mid R \right\rangle_l \mid S_c$$

we reduce D, C applying Rules \boxed{C}_{PAR} , \boxed{C}_{EQ} and lastly Rule \boxed{C}_{COND} . We analyse only the case for $\text{eval}(e, t_p) = \text{true}$ as the other case for $\text{eval}(e, t_p) = \text{false}$ follows the same structure.

$$D, C \xrightarrow{\tau@p} D, C_1 \mid C_c$$

and $C' = C_1 \mid C_c$ and $D' = D$. We can choose $\Gamma = \Gamma'$, for which it holds that $\Gamma \vdash \boxed{D, C}'$.

From Definition 23 we have

$$\boxed{D', C}'^{\Gamma'} = \boxed{D, C}'^{\Gamma} = \left\langle \boxed{C_c|_l}^{\Gamma}, \boxed{C_1}^{\Gamma} \cdot t_p \cdot M_p \mid R \right\rangle_l \mid S_c$$

$\boxed{D, C}'^{\Gamma}$ can mimic D, C applying Rules $\boxed{\text{DCC}}_{\text{EQ}}$, $\boxed{\text{DCC}}_{\text{SPAR}}$, $\boxed{\text{DCC}}_{\text{PPAR}}$, and lastly $\boxed{\text{DCC}}_{\text{COND}}$ for which

$$\boxed{D, C}^\Gamma \rightarrow \left\langle \boxed{C_c|l}^\Gamma, \boxed{C_1}^\Gamma \cdot t_p \cdot M_p \mid R \right\rangle_l \mid S_c$$

Case $\llbracket^C\rrbracket_{\text{CTX}}$

The thesis follows from the induction hypothesis as D, C applies Rule $\llbracket^C\rrbracket_{\text{CTX}}$ and $\boxed{D, C}^\Gamma$ can mimic it with Rule $\llbracket^{\text{DCC}}\rrbracket_{\text{CTX}}$.

Case $\llbracket^C\rrbracket_{\text{PAR}}$

The thesis follows from the induction hypothesis.

Case $\llbracket^C\rrbracket_{\text{EQ}}$

The thesis follows from the induction hypothesis. Starting from any configuration of D, C , $\boxed{D, C}^\Gamma$ can always mimic the evolution of D, C when it applies Rule $\llbracket^C\rrbracket_{\text{EQ}}$: in both cases that $\mathcal{R} = \equiv$ or $\mathcal{R} = \simeq_C$, $\boxed{D, C}^\Gamma$ can apply $\llbracket^{\text{DCC}}\rrbracket_{\text{EQ}}$, $\llbracket^{\text{DCC}}\rrbracket_{\text{SPAR}}$, and $\llbracket^{\text{DCC}}\rrbracket_{\text{PPAR}}$ to mimic D, C .

□

Before proceeding with the proof of (Soundness) of Theorem 8 we extend the semantics of DCC with annotations on transitions to bear the paths on which operations are executed. We range over transition labels with λ .

$$\lambda ::= x \mid \text{cq}(x) \mid ?(x) \mid o \text{ from } x \mid o@x \mid \tau$$

We report in Figure E.3 the annotated semantics of DCC.

E.4. Proof of Compilation

$$\begin{array}{c}
\frac{t_c = \text{eval}(e, t) \quad M(t_c) = (o, t') :: \tilde{m}}{o(x) \text{ from } e; B \cdot t \cdot M \xrightarrow{o \text{ from } e} B \cdot t \triangleleft (x, t') \cdot M[t_c \mapsto \tilde{m}]} \text{ [DCC|RECV]} \\
\\
\frac{B \cdot t \cdot M \xrightarrow{\lambda} B' \cdot t' \cdot M'}{\text{def } X = B_1 \text{ in } B \cdot t \cdot M \xrightarrow{\lambda} \text{def } X = B_1 \text{ in } B' \cdot t' \cdot M'} \text{ [DCC|CTX]} \\
\\
\frac{i = 1 \text{ if } \text{eval}(e, t) = \text{true}, i = 2 \text{ otherwise}}{\text{if } e \{B_1\} \text{ else } \{B_2\} \cdot t \cdot M \xrightarrow{\tau} B_i \cdot t \cdot M} \text{ [DCC|COND]} \\
\\
\frac{j \in I \quad t_c = \text{eval}(e, t) \quad M(t_c) = (o_j, t') :: \tilde{m}}{\sum_{i \in I} [o_i(x_i) \text{ from } e] \{B_i\} \cdot t \cdot M \xrightarrow{o_j \text{ from } e} B_j \cdot t \triangleleft (x_j, t') \cdot M[t_c \mapsto \tilde{m}]} \text{ [DCC|CHOICE]} \\
\\
\frac{t' = \text{eval}(e, t)}{x = e; B \cdot t \cdot M \xrightarrow{x} B \cdot t \triangleleft (x, t') \cdot M} \text{ [DCC|ASSIGN]} \\
\\
\frac{P \xrightarrow{\lambda} P'}{\langle B_s, P \mid P_1 \rangle_l \xrightarrow{\lambda} \langle B_s, P' \mid P_1 \rangle_l} \text{ [DCC|PPAR]} \\
\\
\frac{P = \text{cq}(x); B \cdot t \cdot M \quad t_c \notin \bigcup_i \text{dom}(M_i) \cup \text{dom}(M) \quad t' = t \triangleleft (x, t_c)}{\langle B_s, P \mid \prod_i B_i \cdot t_i \cdot M_i \rangle_l \xrightarrow{\text{cq}(x)} \langle B_s, B \cdot t' \cdot M[t_c \mapsto \varepsilon] \mid \prod_i B_i \cdot t_i \cdot M_i \rangle_l} \text{ [DCC|CQ]} \\
\\
\frac{P = o@e_1(e_2) \text{ to } e_3; B \cdot t \cdot M \quad \text{eval}(e_1, t) = l \quad \text{eval}(e_3, t) = t_c}{\text{eval}(e_2, t) = t_m \quad M'' = M'[t_c \mapsto M'(t_c) :: (o, t_m)]} \text{ [DCC|INSEND]} \\
\langle B_s, P \mid B' \cdot t' \cdot M' \mid P_1 \rangle_l \xrightarrow{o@e_3} \langle B_s, B \cdot t \cdot M \mid B' \cdot t' \cdot M'' \mid P_1 \rangle_l \\
\\
\frac{P = o@e_1(e_2) \text{ to } e_3; B \cdot t \cdot M \quad \text{eval}(e_1, t) = l' \quad \text{eval}(e_3, t) = t_c}{\text{eval}(e_2, t) = t_m \quad M'' = M'[t_c \mapsto M'(t_c) :: (o, t_m)]} \text{ [DCC|SEND]} \\
\langle B_s, P \mid P_1 \rangle_l \mid \langle B'_s, B' \cdot t' \cdot M' \mid P_2 \rangle_{l'} \xrightarrow{o@e_3} \\
\langle B_s, B \cdot t \cdot M \mid P_1 \rangle_l \mid \langle B'_s, B' \cdot t' \cdot M'' \mid P_2 \rangle_{l'} \\
P_1 = ?@e_1(e_2); B_1 \cdot t_1 \cdot M_1 \\
\text{eval}(e_1, t_1) = l \quad Q = B \cdot t_{\perp} \triangleleft (x, \text{eval}(e_2, t_1)) \cdot \emptyset \\
\text{ [DCC|START]} \\
\langle !(x); B, P \rangle_l \mid \langle B'_s, P_1 \mid P_2 \rangle_{l'} \xrightarrow{?(e_2)} \\
\langle !(x); B, Q \mid P \rangle_l \mid \langle B'_s, B_1 \cdot t_1 \cdot M_1 \mid P_2 \rangle_{l'} \\
P = ?@e_1(e_2); B \cdot t \cdot M \quad \text{eval}(e_1, t) = l \quad Q = B \cdot t_{\perp} \triangleleft (x, \text{eval}(e_2, t)) \cdot \emptyset \\
\text{ [DCC|INSTART]} \\
\langle !(x); B', P \mid P_1 \rangle_l \xrightarrow{?(e_2)} \langle !(x); B', Q \mid B \cdot t \cdot M \mid P_1 \rangle_l \\
\\
\frac{S \equiv S_1 \quad S_1 \xrightarrow{\lambda} S'_1 \quad S'_1 \equiv S'}{S \xrightarrow{\lambda} S'} \text{ [DCC|EQ]} \quad \frac{S_1 \xrightarrow{\lambda} S'_1}{S_1 \mid S \xrightarrow{\lambda} S'_1 \mid S} \text{ [DCC|SPAR]}
\end{array}$$

Figure E.3: Correlation Calculus, annotated semantics

We also introduce the filtering and complement operators on sequences of DCC transition labels. Let $\lambda, \tilde{\lambda}$ be a sequence of DCC labels, the filtering of $\lambda, \tilde{\lambda}$ on k , written $(\lambda, \tilde{\lambda})|_k$ is defined as

$$(\lambda, \tilde{\lambda})|_k = \begin{cases} \lambda, (\tilde{\lambda})|_k & \text{if } \lambda \in \left\{ \begin{array}{l} \underline{\mathbf{k.x.y}}, \text{cq}(\underline{\mathbf{k.x.y}}), ?(\underline{\mathbf{k}}), \\ \text{sync from } \underline{\mathbf{k.x.y}}, \text{sync@}\underline{\mathbf{k.x.y}}, \\ \text{start from } \underline{\mathbf{k.x.y}}, \text{start@}\underline{\mathbf{k.x.y}} \end{array} \right\} \\ \tilde{\lambda}|_k & \text{otherwise} \end{cases}$$

Let $\lambda_1, \tilde{\lambda}_1$ and $\lambda_2, \tilde{\lambda}_2$ be two sequences of DCC labels, the complement of $\lambda_1, \tilde{\lambda}_1$ on $\lambda_2, \tilde{\lambda}_2$, written $(\lambda_1, \tilde{\lambda}_1)/(\lambda_2, \tilde{\lambda}_2)$ is defined as

$$(\lambda_1, \tilde{\lambda}_1)/(\lambda_2, \tilde{\lambda}_2) = \begin{cases} \tilde{\lambda}_1/\tilde{\lambda}_2 & \text{if } \lambda_1 = \lambda_2 \wedge \tilde{\lambda}_2 \subseteq \tilde{\lambda}_1 \\ \lambda_1, (\tilde{\lambda}_1)/(\lambda_2, \tilde{\lambda}_2) & \text{if } \lambda_1 \neq \lambda_2 \wedge (\lambda_2, \tilde{\lambda}_2) \subseteq \tilde{\lambda}_1 \end{cases}$$

Below we state Lemma 32 that proves that we can permute the order of transitions in DCC systems and obtain the same final configurations. In particular prove such result on transitions that regard the start of new sessions. The proof of Lemma 32 is based on the simpler Lemma 31 that proves that we can change the order of one start transition and eventually obtain the same DCC system.

Lemma 31 (Single DCC Start Permutation). *Let S be a composition of DCC services such that $S \xrightarrow{\tilde{\lambda}_1} S_1$ and $S_1 \xrightarrow{\lambda_2} S_2$ where $\tilde{\lambda}_1|_k = \varepsilon$ and $\lambda_2|_k = \lambda_2$ then $S \xrightarrow{\lambda_2} S'_1$ and $S'_1 \xrightarrow{\tilde{\lambda}_1} S_2$.*

Proof Sketch. The proof is by induction on the length of $\tilde{\lambda}_1$. Note that we can exclude to have transitions on session k in $\tilde{\lambda}_1$ since by the premises they happen before λ_2 which, by the premises, at most may correspond to the last transition of the start of session k . \square

Lemma 32 (DCC Start Permutations). *Let S be a composition of DCC services such that $S \xrightarrow{\tilde{\lambda}} S'$ and $\tilde{\lambda} = (\underline{\mathbf{k.C.I}}, \tilde{\lambda}')$ then $S \xrightarrow{\tilde{\lambda}|_k} S''$ and $S'' \xrightarrow{\tilde{\lambda}/\tilde{\lambda}|_k} S'$.*

Proof Sketch. Let $\tilde{\lambda}|_k = \lambda_1^k, \dots, \lambda_m^k$ and

$$\tilde{\lambda} = \lambda_1^k, \lambda_{1.1}, \dots, \lambda_{1.n_1}, \dots, \lambda_i^k, \lambda_{i.1}, \dots, \lambda_{i.n_i}, \dots, \lambda_m^k, \lambda_{m.1}, \dots, \lambda_{m.n_m}$$

we can apply Lemma 32 inductively on each couple of the kind λ_i^k and $\lambda_{i.1}, \dots, \lambda_{i.n_i}$, $i \in \{1, \dots, n\}$. Note that the proof holds also for configurations of $\tilde{\lambda}$ in which the elements of $\tilde{\lambda}|_k$ appear next to each other (e.g., $\tilde{\lambda} = \lambda_1^k, \dots, \lambda_i^k, \lambda_i^k+1, \dots, \lambda_{m.n_m}$) since it is a base case of Lemma 31 in which there are no transitions in $\tilde{\lambda}_1$. \square

E.4. Proof of Compilation

Next we state Lemma 33 that proves that, *i*) given a well-typed AC choreography, *ii*) its DCC projection, and *iii*) the DCC system that results from a sequence of reductions belonging to the start of a session, we can complete the remaining sequences of interactions and obtain the originating AC choreography after the reduction to start the session.

Lemma 33 (DCC Start Completion). *Let $\Gamma \vdash D, C$,*

$$C = \mathbf{req} \ k : p[A] \Leftrightarrow l_1.[B_1], \dots, l_n.[B_n]; C_r \mid \prod_{i=1}^n \mathbf{acc} \ k : l_i.q_i[B_i]; C_{q_i}$$

and $\boxed{D, C}^\Gamma = S$ such that $S \xrightarrow{\tilde{\lambda}} S'$ where $\tilde{\lambda} \Big|_k = \tilde{\lambda}$ then *i*) $S' \xrightarrow{\tilde{\lambda}'} S''$, *ii*) $D, C \rightarrow D', C'$, and *iii*) there exists some Γ' s.t. $\Gamma' \vdash D', C'$ and $\boxed{D', C'}^{\Gamma'} = S''$.

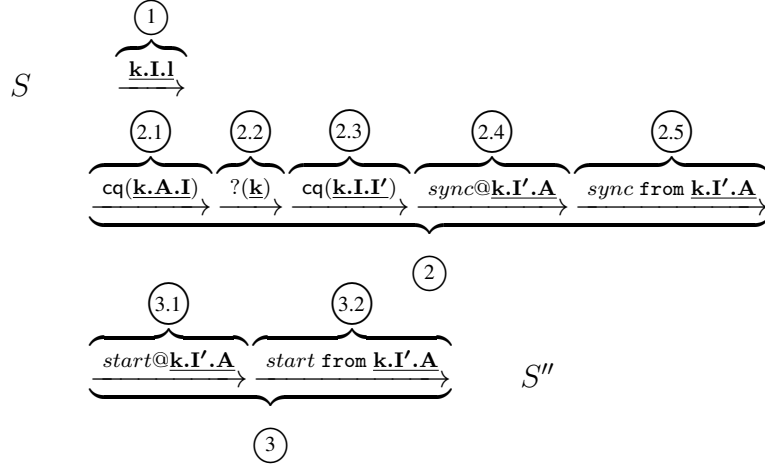
Proof. Proof by case analysis on the length of $\tilde{\lambda}$.

Let $p@l \in \Gamma$. To proceed, we have two subcases whether $l \in \{l_1, \dots, l_n\}$, i.e., whether one of the service processes is at the same location of p . Since the subcases follow the same structure, we detail only the proof for $l \notin \{l_1, \dots, l_n\}$ which allows for a uniform treatment. In the other case, *i*) we should account for transitions on the same service of p with Rules $[\text{DCC}|_{\text{INSTART}}]$ and $[\text{DCC}|_{\text{INSEND}}]$ and *ii*) we would have a newly created process in parallel with p in S'' and $\boxed{D, C}^\Gamma$.

Provided n is the number of service processes involved in the start of the session k , from Definition 23 we can count the number of transitions needed to complete the start of a session. Indeed, given a D, C with

$$C = \mathbf{req} \ k : p[A] \Leftrightarrow l_1.[B_1], \dots, l_n.[B_n]; C_r \mid \prod_{i=1}^n \mathbf{acc} \ k : l_i.q_i[B_i]; C_{q_i}$$

and $\boxed{D, C}^\Gamma = S$ then we can write the sequence of transitions of the projected DCC system



and count the number of all the transitions to complete the start, let it be m , as the sum of:

- ① $n+1$ times, for $\underline{I} \in \{A, \tilde{B}\}$, with last Rule $[^{\text{DCC}}|_{\text{ASSIGN}}]$;
- ② n times, for $\underline{I} \in \tilde{B}$:
 - ②.1 reduces with last applied Rule $[^{\text{DCC}}|_{\text{CQ}}]$;
 - ②.2 reduces with last applied Rule $[^{\text{DCC}}|_{\text{START}}]$;
 - ②.3 n times for $\underline{I}' \in \{A, \tilde{B}\}/\{\underline{I}\}$, reduces with last applied Rule $[^{\text{DCC}}|_{\text{CQ}}]$;
 - ②.4 reduces with last applied Rule $[^{\text{DCC}}|_{\text{SEND}}]$;
 - ②.5 reduces with last applied Rule $[^{\text{DCC}}|_{\text{RECV}}]$;
- ③ n times, for $\underline{I} \in \tilde{B}$:
 - ③.1 reduces with last applied Rule $[^{\text{DCC}}|_{\text{SEND}}]$;
 - ③.2 reduces with last applied Rule $[^{\text{DCC}}|_{\text{RECV}}]$;

and $m = n^2 + 7n + 1$. We proceed unfolding the proof on the length of $\tilde{\lambda}$.

Case $|\{\tilde{\lambda}\}| = 1$

Since the cardinality of $\tilde{\lambda}$ is one and that from the premises we know that $\tilde{\lambda}$ contains only transitions belonging to the start of session k , we can infer that $\tilde{\lambda} = \underline{k.C.I}$ where $C \in \{A, \tilde{B}\}$.

To prove the thesis we let S' do all the remaining transitions to start the session and show that D, C can mimic it. Let $\tilde{l.B} = l_1.B_1, \dots, l_n.B_n$ and $\tilde{l}: G\langle A|\tilde{B}|\tilde{B}\rangle \in \Gamma$.

E.4. Proof of Compilation

From Definition 23 we have

$$\boxed{D, C}^\Gamma \equiv \left\langle \boxed{C_c|_l}^\Gamma, P \mid R \right\rangle_l \mid \prod_{i=1}^n \langle Q_i, R_{l_i} \rangle_{l_i} \mid S_c$$

where

$$\begin{aligned}
 P &= \text{start}(k, (l.A, \widetilde{l.B}); \boxed{C_r}^\Gamma \cdot t_p \cdot M_p = \\
 &\bullet \left(\begin{array}{l} \odot_{\mathbf{I} \in \{\mathbf{A}, \widetilde{\mathbf{B}}\}} \underline{\mathbf{k.I.l}} = l_{\mathbf{I}}; \\ \odot_{\mathbf{I} \in \{\widetilde{\mathbf{B}}\}} \left(\text{cq}(\underline{\mathbf{k.I.A}}); ?@ \underline{\mathbf{k.I.l}}(\underline{\mathbf{k}}); \text{sync}(\underline{\mathbf{k}}) \text{ from } \underline{\mathbf{k.I.A}} \right); \\ \odot_{\mathbf{I} \in \{\widetilde{\mathbf{B}}\}} \text{start}@ \underline{\mathbf{k.I.l}}(\underline{\mathbf{k}}) \text{ to } \underline{\mathbf{k.A.I}}; \end{array} \right); \boxed{C_r}^\Gamma \cdot t_p \cdot M_p \\
 \bullet Q_i &= \text{accept}(k, B_i, G\langle \mathbf{A} | \widetilde{\mathbf{B}} | \widetilde{\mathbf{B}} \rangle); \boxed{C_{q_i}}^\Gamma = \begin{array}{l} !(\underline{\mathbf{k}}); \odot_{\mathbf{I} \in \{\mathbf{A}, \widetilde{\mathbf{B}}\} / \{B_i\}} \text{cq}(\underline{\mathbf{k.I.B}_i}); \\ \text{sync}@ \underline{\mathbf{k.A.l}}(\underline{\mathbf{k}}) \text{ to } \underline{\mathbf{k.B}_i.A}; \\ \text{start}(\underline{\mathbf{k}}) \text{ from } \underline{\mathbf{k.A.B}_i}; \boxed{C_{q_i}}^\Gamma \end{array} \\
 \bullet R &= \prod_{p' \in D(l)/\{p\}} \boxed{C_c|_{p'}}^\Gamma \cdot t_{p'} \cdot M_{p'} \\
 \bullet R_{l_i} &= \prod_{s \in D(l_i)} \boxed{C_c|_s}^\Gamma \cdot t_s \cdot M_s \\
 \bullet S_c &= \prod_{l' \in \Gamma/\{l, \widetilde{l}\}} \left\langle \boxed{C_c|_{l'}}^\Gamma, \prod_{s' \in D(l')} \boxed{C_c|_{s'}}^\Gamma \cdot t_{s'} \cdot M_{s'} \right\rangle_{l'}
 \end{aligned}$$

The first transition, $\lambda = \underline{\mathbf{k.C.l}}$ consumed the first assignment of location and assigned the location of role C to $\underline{\mathbf{k.C.l}}$ in the state of the starter t_p .

Let us suppose, without loss of generality, that $C = A$, then we have

$$P' = \left(\begin{array}{l} \odot_{\mathbf{I} \in \{\widetilde{\mathbf{B}}\}} \underline{\mathbf{k.I.l}} = l_{\mathbf{I}}; \\ \odot_{\mathbf{I} \in \{\widetilde{\mathbf{B}}\}} \left(\text{cq}(\underline{\mathbf{k.I.A}}); ?@ \underline{\mathbf{k.I.l}}(\underline{\mathbf{k}}); \text{sync}(\underline{\mathbf{k}}) \text{ from } \underline{\mathbf{k.I.A}} \right); \\ \odot_{\mathbf{I} \in \{\widetilde{\mathbf{B}}\}} \text{start}@ \underline{\mathbf{k.I.l}}(\underline{\mathbf{k}}) \text{ to } \underline{\mathbf{k.A.I}}; \end{array} \right); \boxed{C_r}^\Gamma \cdot t_p \triangleleft (\underline{\mathbf{k.A.l}}, l) \cdot M_p$$

and $\boxed{D, C}^\Gamma \xrightarrow{\underline{\mathbf{k.A.l}}} S'$ where

$$S' = \left\langle \boxed{C_c|_l}^\Gamma, P' \mid R \right\rangle_l \mid \prod_{i=1}^n \langle Q_i, R_{l_i} \rangle_{l_i} \mid S_c$$

Since in its reduction D, C renames the new session with a fresh name, we first rename session k in P and the service processes Q_i to k' , which is fresh. We take

$$P'' = P'[k'/k] = \left(\begin{array}{c} \odot \frac{\bullet}{\mathbf{I} \in \{\bar{\mathbf{B}}\}} \underline{\mathbf{k}' . \mathbf{I} . \mathbf{l}} = l_{\mathbf{I}} ; \\ \odot \frac{\bullet}{\mathbf{I} \in \{\bar{\mathbf{B}}\}} \left(\text{cq}(\underline{\mathbf{k}' . \mathbf{I} . \mathbf{A}}); ?@ \underline{\mathbf{k}' . \mathbf{I} . \mathbf{l}}(\underline{\mathbf{k}'}); ?@ \underline{\mathbf{k}' . \mathbf{I} . \mathbf{l}}(\underline{\mathbf{k}'}); \right); \\ \odot \frac{\bullet}{\mathbf{I} \in \{\bar{\mathbf{B}}\}} \text{sync}(\underline{\mathbf{k}'}) \text{ from } \underline{\mathbf{k}' . \mathbf{I} . \mathbf{A}} \end{array} \right) ; \overline{C_r}^\Gamma[\underline{\mathbf{k}'}/\underline{\mathbf{k}}] \cdot t'_p \cdot M_p$$

where, let $t'_p = t_p \triangleleft (\underline{\mathbf{k}} . \mathbf{A} . \mathbf{l}, l)$, $t''_p = t'_p \triangleleft (\underline{\mathbf{k}'}, \underline{\mathbf{k}}(t'_p)) \triangleleft (\underline{\mathbf{k}}, t_\perp)$.

By Lemma 28 we have $P' \sim P''$. We take

$$S_0^* = \left\langle \overline{C_c}^\Gamma, P'' \mid R \right\rangle_l \mid \prod_{i=1}^n \langle Q_i[\underline{\mathbf{k}'}/\underline{\mathbf{k}}], R_{l_i} \rangle_{l_i} \mid S_c$$

and by Corollary 30 we have $S_0^* \sim S'$.

Now we can proceed with the rest of the transitions of the start procedure, as defined at the beginning of the proof. Finally we have

$$S'' = \left\langle \overline{C_c}^\Gamma \mid P''' \mid R \right\rangle_l \mid \prod_{i=1}^n \langle Q_i[\underline{\mathbf{k}'}/\underline{\mathbf{k}}], Q'_i \mid R_{l_i} \rangle_{l_i} \mid S_c$$

where $P''' = \overline{C_r}^\Gamma[\underline{\mathbf{k}'}/\underline{\mathbf{k}}] \cdot t'_p \cdot M'_p$ and $Q'_i = \overline{C_{q_i}}^\Gamma[\underline{\mathbf{k}'}/\underline{\mathbf{k}}] \cdot t_{k'} \cdot M_{q_i}$

From the transitions presented above we know that there exists $t'_{k'}$ such that $t'_p = t_p \triangleleft (\underline{\mathbf{k}'}, t'_{k'})$, where $t'_{k'}$ is a session descriptor for session k' (i.e., it contains all the locations and correlations keys used by the processes in session k').

We proceed by proving that D, C can mimic $\overline{D, C}^\Gamma$.

We can apply Rules $[C]_{\text{PAR}}$ and $[C]_{\text{EQ}}$ and lastly Rule $[C]_{\text{PSTART}}$ such that

$$\frac{\begin{array}{l} i \in \{1, \dots, n\} \quad \#k' \quad \{\widetilde{l.B}\} = \uplus_i \{l_i.B_i\} \quad \#\tilde{r} \quad \{\tilde{r}\} = \bigcup_i \{\tilde{r}_i\} \\ p \in D(l) \quad \delta = \mathbf{start} \ k' : l.p[A], l_1.r_1[B_1], \dots, l_n.r_n[B_n] \quad D, \delta \blacktriangleright D' \end{array}}{D, C \rightarrow D', C_r[k'/k] \mid \prod_{i=1}^n (C_{q_i}[k'/k][r_i/q_i]) \mid C_a} \quad [C]_{\text{PSTART}}$$

and

$$\begin{array}{l} D, C \mid C_c \rightarrow \\ D', C_r[k'/k] \mid \prod_i (C_{q_i}[k'/k][r_i/q_i]) \mid \prod_{i=1}^n \mathbf{acc} \ k : l_i.q_i[B_i]; C_{q_i} \mid C_c \end{array}$$

E.4. Proof of Compilation

thus $C' = C_r[k'/k] \mid \prod_{i=1}^n (C_{q_i}[k'/k][r_i/q_i]) \mid \prod_{i=1}^n \mathbf{acc} \ k : l_i.q_i[B_i]; C_{q_i} \mid C_c$.

From the hypothesis we know that $\Gamma \vdash D, C$ and therefore that $\Gamma = \Gamma_1, \tilde{l} : G\langle A|\tilde{B}|\tilde{B} \rangle$. We can find $\Gamma' = \Gamma, \text{init}(k', (p[A], q[\tilde{B}]), G)$ and $\Gamma' \vdash D', C'$.

Finally, we need to prove that $S_1^* = \overline{D', C'}^{\Gamma'}$.

From Definition 23 we have

$$\overline{D', C'}^{\Gamma'} = \left\langle \overline{C_c|l}^{\Gamma'}, P'' \mid R' \right\rangle_l \mid \prod_{i=1}^n \left\langle Q_i'', Q_i^* \mid R_{l_i} \right\rangle_{l_i} \mid S'_c$$

In the following, we use the shortcuts $t_s^* = D'(s).\text{st}$ and $M_s^* = D'(s).\text{que}$ for s process in D' .

- $P'' = \overline{C_r[k'/k]}^{\Gamma'} \cdot t_p^* \cdot M_p^*$
- $R' = \prod_{p' \in D(l)/\{p\}} \overline{C_c|p'}^{\Gamma'} \cdot t_{p'}^* \cdot M_{p'}^*$
- $Q_i'' = \text{accept}(k, B_i, G\langle A|\tilde{B}|\tilde{B} \rangle); \overline{C_{q_i}}^{\Gamma'}$
- $Q_i^* = \overline{C_{q_i}[k'/k][r_i/q_i]}^{\Gamma'} \cdot t_{q_i}^* \cdot M_{q_i}^*$
- $R_{l_i} = \prod_{s \in D(l_i)} \overline{C_c|s}^{\Gamma'} \cdot t_s^* \cdot M_s^*$
- $S'_c = \prod_{l' \in \Gamma/\{l, \tilde{l}\}} \left\langle \overline{C_c|l'}^{\Gamma'}, \prod_{s' \in D(l')} \overline{C_c|s'}^{\Gamma'} \cdot t_{s'}^* \cdot M_{s'}^* \right\rangle_{l'}$

From Rule $[D]_{\text{START}}$ we know that

$$\underline{k'}(t_p^*) = \underline{k'}(t_{q_1}^*) = \dots = \underline{k'}(t_{q_n}^*) = t_{k'}$$

for some $t_{k'}$ session descriptor of session k' .

We prove the case by taking $t_{k'} = t'_{k'}, t'_{k'}$ obtained from the derivation of $\overline{D, C}^{\Gamma}$ and $M_p^* = M'_p$ and $M_{q_i}^* = M_{q_i}, i \in \{1, \dots, n\}$.

Case $1 < |\{\tilde{\lambda}\}| < m - 1$

The case follows the same structure of the previous case. We rename k to k' on p and all the newly created service processes. Then we let the system complete all the transitions and prove that the reductum corresponds to the compilation of D', C' .

Case $|\{\tilde{\lambda}\}| = m$

Since $|\{\tilde{\lambda}\}| = m$ then $S = S'$ where S' has terminated all the transitions to start the session. Here we only have to rename k to k' for all the involved processes and prove that $S' = \boxed{D', C'}^{\Gamma'}$.

□

We now proceed to prove the (*Soundness*) of Theorem 8. In the following we use the shortcut

$$C_{start} = \mathbf{req} \ k : p[A] \Leftrightarrow l_1.[B_1], \dots, l_n.[B_n]; C_r \mid \prod_{i=1}^n \mathbf{acc} \ k : l_i.q_i[B_i]; C_{q_i}$$

Proof (Soundness). We proceed by induction on the cardinality of $\tilde{\lambda}$.

Case $|\{\tilde{\lambda}\}| = 0$

Trivial, $\boxed{D, C}^{\Gamma} = S = \boxed{D', C'}^{\Gamma'}$, $D, C = D', C'$, and $\Gamma \vdash D', C'$.

Case $|\{\tilde{\lambda}\}| = 1$

We unfold the case on the structure of D, C . We do not consider the inapplicable cases for C containing (*com*) or (*start*) terms, which cannot be present in endpoint choreographies.

Case $C = k : p[A].e \rightarrow B.o; C_p \mid C_c$

We proceed analysing the last Rule applied in the derivation of $\boxed{D, C}^{\Gamma}$

Case \boxed{DCC}_{INSEND}

We have $\tilde{\lambda} = \lambda = o@k.A.B$ where $o \notin \{start, sync\}$. Indeed operations *start* and *sync* are reserved for the starting of new sessions and cannot appear as first action of the compilation of $\boxed{D, C}^{\Gamma}$. Therefore we can exclude the case that $\boxed{D, C}^{\Gamma}$ is starting a new session.

Since \boxed{DCC}_{INSEND} applies we can infer that, let $p@l \in \Gamma$, then $q@l \in \Gamma$ such that

$$\boxed{D, C}^{\Gamma} \equiv \left\langle \boxed{C_c|l}^{\Gamma}, P \mid Q \mid R \right\rangle_l \mid S_c$$

where

E.4. Proof of Compilation

- $P = o@k.B.l \text{ to } k.A.B; [C_p]^\Gamma \cdot t_p \cdot M_p$
- $Q = [C_c|_q]^\Gamma \cdot t_q \cdot M_q$
- $R = \prod_{r \in D(l)/\{p,q\}} [C_c|_r]^\Gamma \cdot t_r \cdot M_r$
- $S_c = \prod_{l' \in \Gamma/\{l\}} \left\langle [C_c|_{l'}]^\Gamma, \prod_{s \in D(l')} [C_c|_s]^\Gamma \cdot t_s \cdot M_s \right\rangle_{l'}$

$[D, C]^\Gamma$ reduces, applying as last Rule $[^{\text{DCC}}|_{\text{INSEND}}]$, let $t_c = \text{eval}(k.A.B, t_p)$, $t_m = \text{eval}(e, t_p)$, and $M_q(t_c) = \tilde{m}$

$$[D, C]^\Gamma \xrightarrow{o@k.A.B} S$$

where

$$S = \left\langle [C_c|_l]^\Gamma, [C_p]^\Gamma \cdot t_p \cdot M_p \mid [C_c|_q]^\Gamma \cdot t_q \cdot M_q[t_c \mapsto \tilde{m} :: (o, t_m)] \mid R \right\rangle_l \mid S_c$$

D, C can mimic $[D, C]^\Gamma$ with Rules $[^{\text{C}}|_{\text{EQ}}]$, $[^{\text{C}}|_{\text{PAR}}]$, and $[^{\text{C}}|_{\text{SEND}}]$ for which

$$D, C \rightarrow D', C_p \mid C_c$$

where $D' = D[q \mapsto (t_q, M_q[t_c \mapsto \tilde{m} :: (o, t_m)])]$.

Since from the premises $\Gamma \vdash D, C$ then $\Gamma = \Gamma_1, k[A] : !B.o(U); T, b[k]_B^A : T'$ and we can find $\Gamma' = \Gamma_1, k[A] : T, b[k]_B^A : T'; ?A.o(U)$ such that $\Gamma' \vdash D, C$.

From Lemma 27 we have that $[D', C']^{\Gamma'} = [D', C']^\Gamma$. Applying Definition 23 to $[D', C']^\Gamma$, we have $[D', C']^\Gamma = S$ and therefore $S = [D', C']^{\Gamma'}$.

Case $[^{\text{DCC}}|_{\text{SEND}}]$

Similar to the proof of case $[^{\text{DCC}}|_{\text{INSEND}}]$.

Case $[^{\text{DCC}}|_{\text{CHOICE}}]$

The case unfolds following the proof of case $C = k : A \rightarrow q[B].\{o_i(x_i); C_i\}_i; C_q \mid C_c$, subcase $[^{\text{DCC}}|_{\text{CHOICE}}]$, and the thesis follows by applying the induction hypothesis.

Case $[^{\text{DCC}}|_{\text{COND}}]$

The case unfolds following the proof of case $C = \text{if } p.e \{C_1\} \text{ else } \{C_2\}; C_q \mid C_c$, subcase $[^{\text{DCC}}|_{\text{COND}}]$, and the thesis follows by applying the induction hypothesis.

Case $[^{\text{DCC}}|_{\text{ASSIGN}}]$

The case unfolds following the proof of case $C = C_{start} \mid C_c$, subcase $[^{\text{DCC}}|_{\text{ASSIGN}}]$, and the thesis follows by applying the induction hypothesis.

Case $[^{\text{DCC}}|_{\text{RECV}}]$ or $[^{\text{DCC}}|_{\text{INSTART}}]$ or $[^{\text{DCC}}|_{\text{START}}]$ or $[^{\text{DCC}}|_{\text{CQ}}]$

None of the Rules $[^{\text{DCC}}|_{\text{RECV}}]$, $[^{\text{DCC}}|_{\text{INSTART}}]$, $[^{\text{DCC}}|_{\text{START}}]$, $[^{\text{DCC}}|_{\text{CQ}}]$ apply as

$\boxed{D, C}^\Gamma$ makes only one transition whilst all the listed Rules can appear only in a sequence of interactions greater than one.

Case $C = k : A \rightarrow q[B].\{o_i(x_i); C_i\}_{i \in I} \mid C_c$

We proceed analysing the last Rule applied in the derivation of $\boxed{D, C}^\Gamma$

Case $[\text{DCC}]_{\text{INSEND}}$ or $[\text{DCC}]_{\text{SEND}}$

The case unfolds following the proof of case $C = k : p[A].e \rightarrow B.o; C_q \mid C_c$, respectively subcases $[\text{DCC}]_{\text{INSEND}}$ and $[\text{DCC}]_{\text{SEND}}$. The thesis follows by applying the induction hypothesis.

Case $[\text{DCC}]_{\text{CHOICE}}$

We have $\tilde{\lambda} = \lambda = o_j$ from $\mathbf{k.A.B}$ where $o_j \notin \{\text{start}, \text{sync}\}$, following the same reasoning of case $C = k : p[A].e \rightarrow B.o; C_p \mid C_c$ subcase $[\text{DCC}]_{\text{SEND}}$.

Let $q@l \in \Gamma$, from Definition 23 we have

$$\boxed{D, C}^\Gamma \equiv \left\langle \boxed{C_c|l}^\Gamma, Q \mid R \right\rangle_l \mid S_c$$

where

- $Q = \sum_{i \in I} [o_i(x_i) \text{ from } \mathbf{k.A.B}] \{ \boxed{C_i}^\Gamma \} \cdot t_q \cdot M_q$
- $R = \prod_{r \in D(l)/\{q\}} \boxed{C_c|r}^\Gamma \cdot t_r \cdot M_r$
- $S_c = \prod_{l' \in \Gamma/\{l\}} \left\langle \boxed{C_c|l'}^\Gamma, \prod_{s \in D(l')} \boxed{C_c|s}^\Gamma \cdot t_s \cdot M_s \right\rangle_{l'}$

and can apply Rules $[\text{DCC}]_{\text{EQ}}$, $[\text{DCC}]_{\text{SPAR}}$, $[\text{DCC}]_{\text{PPAR}}$ and lastly $[\text{DCC}]_{\text{CHOICE}}$ such that, let $t_c = \text{eval}(\mathbf{k.A.B}, t_p)$, $t_m = \text{eval}(e, t_p)$, and $M_q(t_c) = (o_j, t_m) :: \tilde{m}$

$$\boxed{D, C}^\Gamma \xrightarrow{o_j \text{ from } \mathbf{k.A.B}} S$$

where

$$S = \left\langle \boxed{C_c|l}^\Gamma, \boxed{C_j}^\Gamma \cdot t_q \triangleleft (x_j, t_m) \cdot M_p[t_c \mapsto \tilde{m}] \mid R \right\rangle_l \mid S_c$$

D, C can mimic $\boxed{D, C}^\Gamma$ with Rules $[\text{C}]_{\text{EQ}}$, $[\text{C}]_{\text{PAR}}$, and $[\text{C}]_{\text{RECV}}$ for which

$$D, C \rightarrow D', C_p \mid C_c$$

where $D' = D[q \mapsto (t_q \triangleleft (x_j, t_m), M_q[t_c \mapsto \tilde{m}])]$.

Since from the premises $\Gamma \vdash D, C$ then

$\Gamma = \Gamma_1, k[A] : ?A.\{o_i(U_i); T_i\}_{i \in I}, b[k]_B^A : ?A.o_j(U_j); T'$ and we can find

$\Gamma' = \Gamma_1, k[A] : T_j, b[k]_B^A : T'$ such that $\Gamma' \vdash D, C$.

From Lemma 27 we have that $\boxed{D', C'}^{\Gamma'} = \boxed{D', C'}^\Gamma$. Applying Definition 23 to $\boxed{D', C'}^\Gamma$, we have $\boxed{D', C'}^\Gamma = S$ and therefore $S = \boxed{D', C'}^{\Gamma'}$.

E.4. Proof of Compilation

Case $\lfloor \text{DCC} \rfloor_{\text{COND}}$

The case unfolds following the proof of case $C = \text{if } p.e \{C_1\} \text{ else } \{C_2\}; C_q \mid C_c$, subcase $\lfloor \text{DCC} \rfloor_{\text{COND}}$, and the thesis follows by applying the induction hypothesis.

Case $\lfloor \text{DCC} \rfloor_{\text{ASSIGN}}$

The case unfolds following the proof of case $C = C_{start} \mid C_c$, subcase $\lfloor \text{DCC} \rfloor_{\text{ASSIGN}}$, and the thesis follows by applying the induction hypothesis.

Case $\lfloor \text{DCC} \rfloor_{\text{RECV}}$ or $\lfloor \text{DCC} \rfloor_{\text{INSTART}}$ or $\lfloor \text{DCC} \rfloor_{\text{START}}$ or $\lfloor \text{DCC} \rfloor_{\text{CQ}}$

None of the Rules $\lfloor \text{DCC} \rfloor_{\text{RECV}}$, $\lfloor \text{DCC} \rfloor_{\text{INSTART}}$, $\lfloor \text{DCC} \rfloor_{\text{START}}$, $\lfloor \text{DCC} \rfloor_{\text{CQ}}$ apply as $\lfloor \overline{D}, C \rfloor^\Gamma$ makes only one transition whilst all the listed Rules can appear only in a sequence of interactions greater than one.

Case $C = C_{start} \mid C_c$

We proceed analysing the last Rule applied in the derivation of $\lfloor \overline{D}, C \rfloor^\Gamma$

Case $\lfloor \text{DCC} \rfloor_{\text{INSEND}}$ or $\lfloor \text{DCC} \rfloor_{\text{SEND}}$

The case unfolds following the proof of case $C = k : p[A].e \rightarrow B.o; C_q \mid C_c$, respectively subcases $\lfloor \text{DCC} \rfloor_{\text{INSEND}}$ and $\lfloor \text{DCC} \rfloor_{\text{SEND}}$. The thesis follows by applying the induction hypothesis.

Case $\lfloor \text{DCC} \rfloor_{\text{CHOICE}}$

The case unfolds following the proof of case

$C = k : A \rightarrow q[B].\{o_i(x_i); C_i\}_{i \in I}; C_q \mid C_c$, respectively subcases $\lfloor \text{DCC} \rfloor_{\text{INSEND}}$ and $\lfloor \text{DCC} \rfloor_{\text{SEND}}$. The thesis follows by applying the induction hypothesis.

Case $\lfloor \text{DCC} \rfloor_{\text{COND}}$

The case unfolds following the proof of case $C = \text{if } p.e \{C_1\} \text{ else } \{C_2\}; C_q \mid C_c$, subcase $\lfloor \text{DCC} \rfloor_{\text{COND}}$, and the thesis follows by applying the induction hypothesis.

Case $\lfloor \text{DCC} \rfloor_{\text{ASSIGN}}$

Let $p@l \in \Gamma$, we have two subcases whether $l \in \{l_1, \dots, l_n\}$, i.e., whether one of the service processes is at the same location of p .

Case $l \notin \{l_1, \dots, l_n\}$

Let $C \in \{A, \tilde{B}\}$, we have two subcases whether $\tilde{\lambda} = \lambda = \underline{k.C.1}$ or $\tilde{\lambda} = \lambda = \underline{k''.C.1}$, i.e., whether we are stating the session on k or we are starting another session k'' .

Case $\lambda = \underline{k.C.1}$

In this case $\lfloor \overline{D}, C \rfloor^\Gamma$ is starting a new session on k . The case is proven applying Lemma 33.

Case $\lambda \neq \underline{k'.C.1}$

In this case we are starting a session on $k'' \neq k$. The case unfolds following the proof of case $C = C_{start} \mid C_c$, subcase $\lfloor C \rfloor_{\text{ASSIGN}}$, subcase $\lambda = \underline{k.C.1}$. The thesis follows by applying the induction hypothesis.

Case $l \in \{l_1, \dots, l_n\}$

The case follows the same structure of the case $l \notin \{\tilde{l}\}$ although the service located at l has $\boxed{C_{a|l}}^\Gamma$ as start behaviour and $\boxed{D, C}^\Gamma$ applies Rule $[\text{DCC}|_{\text{INSTART}}]$ in place of the $[\text{DCC}|_{\text{START}}]$ for starting the DCC process located at l .

Case $[\text{DCC}|_{\text{RECV}}]$ or $[\text{DCC}|_{\text{INSTART}}]$ or $[\text{DCC}|_{\text{START}}]$ or $[\text{DCC}|_{\text{CQ}}]$

None of the Rules $[\text{DCC}|_{\text{RECV}}]$, $[\text{DCC}|_{\text{INSTART}}]$, $[\text{DCC}|_{\text{START}}]$, $[\text{DCC}|_{\text{CQ}}]$ apply as $\boxed{D, C}^\Gamma$ makes only one transition whilst all the listed Rules can appear only in a sequence of interactions greater than one.

Case $C = \text{if } p.e \{C_1\} \text{ else } \{C_2\} | C_c$

We proceed analysing the last Rule applied in the derivation of $\boxed{D, C}^\Gamma$.

Case $[\text{DCC}|_{\text{INSEND}}]$ or $[\text{DCC}|_{\text{SEND}}]$

The case unfolds following the proof of case $C = k : p[A].e \rightarrow B.o; C_q | C_c$, respectively subcases $[\text{DCC}|_{\text{INSEND}}]$ and $[\text{DCC}|_{\text{SEND}}]$. The thesis follows by applying the induction hypothesis.

Case $[\text{DCC}|_{\text{CHOICE}}]$

The case unfolds following the proof of case

$C = k : A \rightarrow q[B].\{o_i(x_i); C_i\}_i; C_q | C_c$, subcase $[\text{DCC}|_{\text{CHOICE}}]$, and the thesis follows by applying the induction hypothesis.

Case $[\text{DCC}|_{\text{COND}}]$

We have $\lambda = \lambda = \tau$. Let $p@l \in \Gamma$. From Definition 23 we have

$$\boxed{D, C}^\Gamma \equiv \left\langle \boxed{C_{c|l}}^\Gamma, P | R \right\rangle_l | S_c$$

where

- $P = \text{if } p.e \{C_1\}^\Gamma \text{ else } \{C_2\}^\Gamma \cdot t_p \cdot M_p$
- $R = \prod_{r \in D(l)/\{p\}} \boxed{C_{c|r}}^\Gamma \cdot t_r \cdot M_r$
- $S_c = \prod_{l' \in \Gamma/\{l\}} \left\langle \boxed{C_{c|l'}}^\Gamma, \prod_{s \in D(l')} \boxed{C_{c|s}}^\Gamma \cdot t_s \cdot M_s \right\rangle_{l'}$

The case unfolds into two cases, on whether $\text{eval}(e, D.\text{st}(p)) = \text{true}$. Here we proceed with the positive case. The other case follows the same structure.

We proceed considering that $\text{eval}(e, D.\text{st}(p)) = \text{true}$. $\boxed{D, C}^\Gamma$ reduces with Rules $[\text{DCC}|_{\text{EQ}}]$, $[\text{DCC}|_{\text{SPAR}}]$, $[\text{DCC}|_{\text{PAR}}]$, and $[\text{DCC}|_{\text{COND}}]$ such that

$$\boxed{D, C}^\Gamma \rightarrow \left\langle \boxed{C_{c|l}}^\Gamma, \boxed{C_{c|p}}^\Gamma \cdot t_p \cdot M_p | R \right\rangle_l | S_c$$

where $S = \left\langle \boxed{C_{c|l}}^\Gamma, \boxed{C_{c|l}}^\Gamma \cdot t_p \cdot M_p | R \right\rangle_l | S_c$. D, C can mimic $\boxed{D, C}^\Gamma$ with Rules $[\text{C}|_{\text{EQ}}]$, $[\text{C}|_{\text{PAR}}]$, and $[\text{C}|_{\text{COND}}]$ such that

E.4. Proof of Compilation

$$D, C \rightarrow D, C_1 \mid C_c$$

We choose $\Gamma' = \Gamma$ for which it holds that $\Gamma \vdash D, C_1 \mid C_c$.

Finally, $\boxed{D, C_1 \mid C_c}^\Gamma = S$ by Definition 23.

Case $\llbracket \text{DCC} \rrbracket_{\text{ASSIGN}}$

The case unfolds following the proof of case $C = C_{start} \mid C_c$, subcase $\llbracket \text{DCC} \rrbracket_{\text{ASSIGN}}$, and the thesis follows by applying the induction hypothesis.

Case $\llbracket \text{DCC} \rrbracket_{\text{RECV}}$ or $\llbracket \text{DCC} \rrbracket_{\text{INSTART}}$ or $\llbracket \text{DCC} \rrbracket_{\text{START}}$ or $\llbracket \text{DCC} \rrbracket_{\text{CQ}}$

None of the Rules $\llbracket \text{DCC} \rrbracket_{\text{RECV}}$, $\llbracket \text{DCC} \rrbracket_{\text{INSTART}}$, $\llbracket \text{DCC} \rrbracket_{\text{START}}$, $\llbracket \text{DCC} \rrbracket_{\text{CQ}}$ apply as $\boxed{D, C}^\Gamma$ makes only one transition whilst all the listed Rules can appear only in a sequence of interactions greater than one.

Case $C = \text{def } X = C' \text{ in } C_p \mid C_c$ or $C = C_1 \mid C_2$

In either cases the proof unfolds with the same structure of the previous cases and the thesis follows by applying the induction hypothesis.

Case $C = 0$

Impossible since $|\{\tilde{\lambda}\}| = 1$.

Case $|\{\tilde{\lambda}\}| > 1$

We unfold the case on the structure of D, C . Like in the previous case, we do not consider the inapplicable cases for C containing (*com*) or (*start*) terms, which cannot be present in endpoint choreographies.

For each AC configuration, we analyse the subcases on the structure of $\tilde{\lambda} = \lambda', \tilde{\lambda}'$. Note that $\tilde{\lambda}$ contains all the annotation labels of execution since the compilation of $\boxed{D, C}^\Gamma$. Hence, we can exclude that the labels of the kind $\text{cq}(x)$ and $?(x)$ as they can only appear later in $\tilde{\lambda}'$ after a sequence of assignments.

Case $C = k : p[A].e \rightarrow B.o; C_p \mid C_c$

Case $\lambda' = o @ x$

We can follow the same structure of the proof for case $|\{\lambda\}| = 1$, $C = k : p[A].e \rightarrow B.o; C_p \mid C_c$, Rule $\llbracket \text{DCC} \rrbracket_{\text{INSEND}}$ (or $\llbracket \text{DCC} \rrbracket_{\text{SEND}}$, depending on whether $\{p@l, q[B] : k, q@l\} \subseteq \Gamma$ or not), to prove that $D, C \rightarrow D', C'$, $\boxed{D, C}^\Gamma \xrightarrow{\lambda'} S$ and that there exists a Γ' such that $\Gamma' \vdash D', C'$ and $\boxed{D, C}^\Gamma = S$. Then we can prove the thesis applying the induction hypothesis on the remaining transitions in $\tilde{\lambda}'$ and D', C' .

Case $\lambda' = o$ from x

The case unfolds following the proof of case $|\{\tilde{\lambda}\}| > 1$,

$C = k : A \rightarrow q[B].\{o_i(x_i); C_i\}_{i \in I} \mid C_c$, $\lambda' = o$ from x and the thesis follows by applying the induction hypothesis on the remaining transitions in $\tilde{\lambda}'$.

Case $\lambda' = x$

The case unfolds following the proof of case $|\{\tilde{\lambda}\}| > 1$, $C = C_{start} \mid C_c$, $\lambda' = x$. Let $x = \mathbf{k.C.1}$, the thesis follows by applying Lemma 32, Lemma 33 and the induction hypothesis on the remaining transitions in $\tilde{\lambda} / \tilde{\lambda} \Big|_k$.

Case $\lambda' = \tau$

The case unfolds following the proof of case $|\{\tilde{\lambda}\}| > 1$, $C = \text{if } p.e \{C_1\} \text{ else } \{C_2\} \mid C_c$, $\lambda' = \tau$ and the thesis follows by applying the induction hypothesis on the remaining transitions in $\tilde{\lambda}'$.

Case $C = k : A \rightarrow q[B].\{o_i(x_i); C_i\}_{i \in I} \mid C_c$

Case $\lambda' = o@x$

The case unfolds following the proof of case $|\{\tilde{\lambda}\}| > 1$, $C = k : p[A].e \rightarrow B.o; C_p \mid C_c$, $\lambda' = o@x$ and the thesis follows by applying the induction hypothesis on the remaining transitions in $\tilde{\lambda}'$.

Case $\lambda' = o \text{ from } x$

We follow the same structure of the proof for case $|\{\lambda\}| = 1$, $C = k : A \rightarrow q[B].\{o_i(x_i); C_i\}_{i \in I} \mid C_c$, Rule $[\text{DCC}|_{\text{CHOICE}}]$ to prove that $D, C \rightarrow D', C'$, $\boxed{D, C}^\Gamma \xrightarrow{\lambda'} S$ and that there exists a Γ' such that $\Gamma' \vdash D', C'$ and $\boxed{D, C}^{\Gamma'} = S$. Then we can prove the thesis applying the induction hypothesis on the remaining transitions in $\tilde{\lambda}'$. Note that we can exclude that λ' has been generated applying Rule $[\text{DCC}|_{\text{RECV}}]$ since such occurrence cannot happen right after the compilation of D, C .

Case $\lambda' = x$

The case unfolds following the proof of case $|\{\tilde{\lambda}\}| > 1$, $C = C_{start} \mid C_c$, $\lambda' = x$. Let $x = \mathbf{k.C.1}$, the thesis follows by applying Lemma 32, Lemma 33 and the induction hypothesis on the remaining transitions in $\tilde{\lambda} / \tilde{\lambda} \Big|_k$.

Case $\lambda' = \tau$

The case unfolds following the proof of case $|\{\tilde{\lambda}\}| > 1$, $C = \text{if } p.e \{C_1\} \text{ else } \{C_2\} \mid C_c$, $\lambda' = \tau$ and the thesis follows by applying the induction hypothesis on the remaining transitions in $\tilde{\lambda}'$.

Case $C = C_{start} \mid C_c$

Case $\lambda' = o@x$

The case unfolds following following the proof of case $|\{\tilde{\lambda}\}| > 1$, $C = k : p[A].e \rightarrow B.o; C_p \mid C_c$, $\lambda' = o@x$ and the thesis follows by applying the induction hypothesis on the remaining transitions in $\tilde{\lambda}'$.

Case $\lambda' = o \text{ from } x$

We follow the proof of case $|\{\tilde{\lambda}\}| > 1$, $C = k : A \rightarrow q[B].\{o_i(x_i); C_i\}_{i \in I} \mid C_c$,

E.4. Proof of Compilation

$\lambda' = o$ from x and the thesis follows by applying the induction hypothesis on the remaining transitions in $\tilde{\lambda}'$.

Case $\lambda' = x$

Since $\lambda' = x$, we know that we are at the beginning of the start of a new session. We have two subcases on whether $x = \mathbf{k.C.I}$ or $x = \mathbf{k''.C.I}$. The latter case is proven following the same structure of this case for C that contains the start of a session on k'' . Below we develop the first case.

First we extract all the transitions in $\tilde{\lambda}$ regarding the start of session k . Let $\tilde{\lambda}_k = \tilde{\lambda} \Big|_k$. From Lemma 33 we know that, given the set of transitions $\tilde{\lambda}_k$,

we can complete the start at DCC level such that $\boxed{D, C}^\Gamma \xrightarrow{\tilde{\lambda}_k, \tilde{\lambda}_{m-k}} S'$ where $\tilde{\lambda}_{m-k}$ are the remaining transitions to complete the start. We remind that, since at AC level we replace k with a fresh session name k' , we follow the same principle at DCC level applying the renaming $[k'/k]$ — k' fresh — on all the involved processes. Then we can find $D, C \rightarrow D', C'$ with Rule $[C]_{\text{PSTART}}$ such that, for some Γ' , $\Gamma' \vdash D', C'$ and $\boxed{D', C'}^{\Gamma'} = S'$. Finally, we prove the thesis applying Lemma 32 and the induction hypothesis on the remaining and properly renamed transitions $(\tilde{\lambda} / \tilde{\lambda} \Big|_k) [k'/k]$.

Case $\lambda' = \tau$

The case follows the proof of case $|\{\tilde{\lambda}\}| > 1$, $C = \text{if p.e } \{C_1\} \text{ else } \{C_2\} \mid C_c$, $\lambda' = \tau$ and the thesis follows by applying the induction hypothesis on the remaining transitions in $\tilde{\lambda}'$.

Case $C = \text{if p.e } \{C_1\} \text{ else } \{C_2\} \mid C_c$

Case $\lambda' = o@x$

The case follows the proof of case $|\{\tilde{\lambda}\}| > 1$, $C = k : \text{p}[A].e \rightarrow \text{B.o}; C_p \mid C_c$, $\lambda' = o@x$ and the thesis follows by applying the induction hypothesis on the remaining transitions in $\tilde{\lambda}'$.

Case $\lambda' = o$ from x

The case unfolds following the proof of case $|\{\tilde{\lambda}\}| > 1$, $C = k : \text{A} \rightarrow \text{q}[\text{B}].\{o_i(x_i); C_i\}_{i \in I} \mid C_c$, $\lambda' = o$ from x and the thesis follows by applying the induction hypothesis on the remaining transitions in $\tilde{\lambda}'$.

Case $\lambda' = x$

The case unfolds following the proof of case $|\{\tilde{\lambda}\}| > 1$, $C = C_{\text{start}} \mid C_c$, $\lambda' = x$. Let $x = \mathbf{k.C.I}$, the thesis follows by applying Lemma 32, Lemma 33, and the induction hypothesis on the remaining transitions in $\tilde{\lambda} / \tilde{\lambda} \Big|_k$.

Case $\lambda' = \tau$

The case follows the same proof of case $|\{\tilde{\lambda}\}| = 1$, $C = \text{if p.e } \{C_1\} \text{ else } \{C_2\} \mid C_c$,

$\lambda' = \tau$ and the thesis follows by applying the induction hypothesis on the remaining transitions in $\tilde{\lambda}'$.

Case $C = \text{def } X = C' \text{ in } C_p \mid C_c \text{ or } C = C_1 \mid C_2$

In either cases the proof unfolds with the same structure of the previous cases and the thesis follows by applying the induction hypothesis.

Case $C = 0$

Impossible since $|\{\tilde{\lambda}\}| > 1$.

□