
On the Synthesis of Choreographies

Thesis submitted for the degree of
Doctor of Philosophy at the
University of Leicester

by

Julien Lange

Department of Computer Science
University of Leicester

August 2013

On the Synthesis of Choreographies

Julien Lange

Abstract. The theories based on session types stand out as effective methodologies to specify and verify properties of distributed systems. A key result in the area shows the suitability of choreography languages and session types as a basis for a choreography-driven methodology for distributed software development. The methodology it advocates is as follows: a team of programmers designs a global view of the interactions to be implemented (i.e., a choreography), then the choreography is projected onto each role. Finally, each program implementing one or more roles in the choreography is validated against its corresponding projection(s).

This is an ideal methodology but it may not always be possible to design one set of choreographies that will drive the overall development of a distributed system. Indeed, software needs maintenance, specifications may evolve (sometimes also during the development), and issues may arise during the implementation phase. Therefore, there is a need for an alternative approach whereby it is possible to infer a choreography from local behavioural specifications (i.e., session types).

We tackle the problem of synthesising choreographies from local behavioural specifications by introducing a type system which assigns – if possible – a choreography to set of session types. We demonstrate the importance of obtaining a choreography from local specifications through two applications. Firstly, we give three algorithms and a methodology to help software designers refine a choreography into a global assertion, i.e., a choreography decorated with logical formulae specifying senders' obligations and receivers' requirements. Secondly, we introduce a formal model for distributed systems where each participant advertises its requirements and obligations as behavioural contracts (in the form of session types), and where multiparty sessions are started when a set of contracts allows to synthesise a choreography.

Acknowledgements

First and foremost I would like to thank Emilio Tuosto for being such a great supervisor, friend and colleague over these past four years. Thank you very much Emilio it has been great working with you (and still is)!

I am grateful to my collaborators Laura Bocchi and Alceste Scalas. I am very glad I had the chance to work with both of you.

I would like to thank Mariangiola Dezani-Ciancaglini and Nir Piterman for examining my thesis, also for great comments and discussions during my viva.

I would like to thank my colleagues from Leicester who have had – in a way or another – some influence on my work: Daniela Petrişan whose “puzzle” became somehow part of this thesis; Tadeusz Litak with whom I had many discussions on early ideas on a type system from local to global types; Roy Crole and Fer-Jan de Vries for discussions and advice as part of my thesis committee.

I would like to thank the Department of Computer Science of the University of Leicester for offering me the chance of working and studying here.

I am grateful to Nobuko Yoshida and Pierre-Malo Deniélou for insightful discussions and comments.

Thank you to my other friends and (ex-)colleagues without whom working in the department would not be what it was. Thank you to my “brother-in-PhD” Kyriakos; and thank you to Gabriela, Octavian, Alex, Oliver, Sokkar, Muz, and all those that I am forgetting here, for being great friends in and outside the department.

I am grateful to my parents, sisters, and grand-parents for their support and encouragements during these past four years. Merci Angèle for visiting me whenever you could!

Last but not least, I would like to thank Franz Preud’homme who, fifteen years ago, gave me a book on programming with Delphi. Had he not been around, I probably would not have studied computer science.

To Angèle.

Contents

1	Introduction	1
1.1	Motivations and Objectives	1
1.2	Synopsis and Contributions	4
1.3	Publications	5
2	Background	7
2.1	Foundations	7
2.1.1	Dyadic Session Types	8
2.1.2	Multiparty Session Types	14
2.1.3	Design-by-Contract for Distributed Multiparty Interactions	20
2.1.4	Session Types and Communicating Machines	22
2.2	Related Work	25
2.2.1	Global versus Local Specifications	25
2.2.2	On Synthesising Choreographies	32
2.2.3	Beyond Multiparty Session Types	36
2.2.4	Other Approaches to Multiparty Sessions	38
3	Synthesising Choreographies from Local Types	41

3.1	Introduction	41
3.2	Local Types	43
3.3	Global Types	47
3.3.1	Well-formed Global Types	49
3.3.2	Properties of Well-formed Global Types	57
3.4	Synthesising Global Types	62
3.4.1	Validation Rules	63
3.4.2	Applying the Rules	66
3.5	Properties of the Synthesis	71
3.5.1	Decidability	72
3.5.2	Uniqueness	79
3.5.3	Well-formedness and Projections	82
3.5.4	Subject Reduction	95
3.5.5	Equivalence with Original System	99
3.5.6	Completeness	100
3.6	Perspectives	106
3.7	Concluding Remarks	111
4	Amending Contracts for Choreographies	113
4.1	Introduction	113
4.2	Preliminaries	116
4.3	On Recovering History Sensitivity	122
4.3.1	Strengthening	125
4.3.2	Variable Propagation	128
4.3.3	Properties of Σ and Π	131
4.4	On Recovering Temporal Satisfiability	139
4.4.1	Lifting Algorithm	140

4.4.2	Applying Λ to Branching and Recursion	145
4.4.3	Properties of Λ	151
4.5	A Methodology for Amending Choreographies	158
4.5.1	Amendment Strategies	162
4.6	Applying the Methodology	163
4.6.1	Cash Withdrawal	164
4.6.2	Credit Request	166
4.7	Concluding Remarks	168
5	Choreography Synthesis as Contract Agreement	169
5.1	Introduction	169
5.2	A Motivating Example	171
5.3	A Choreography-Based Contract Model	174
5.4	Contract-Oriented Computing and Choreographies	176
5.4.1	A Choreography-Based CO ₂	177
5.4.2	On the Flexibility of Session Establishment	183
5.5	The Problem of Honesty	190
5.6	Properties of Honest Networks	197
5.7	Concluding Remarks	202
6	Conclusions and Future Directions	204
6.1	Summary of the Contributions	204
6.2	Future Directions	206
	Bibliography	210

List of Figures

1.1	Choreography driven development	3
2.1	The multiparty session types approach	15
2.2	Example of generalised global type	22
2.3	Example of communicating machines	23
3.1	Global view of S_{BS}	43
3.2	Rules for well-formedness	51
3.3	Validation rules for programs	65
3.4	Typing derivation of $S_{ex3.6}$	70
4.1	Seller-Buyer global type (extract)	115
4.2	ATM protocol	164
4.3	Credit request protocol	166
5.1	Semantics rules for CO_2	178
5.2	Congruence rules for CO_2	181

1.1 Motivations and Objectives

The design, implementation, and maintenance of distributed systems are not easy tasks. This is due, notably, to the fact that many parties – with different objectives and policies – may be involved in each one of these tasks. Also, the specification of such systems is generally done at different levels and possibly via different formalisms, e.g., each part of a system may be specified from internal and external perspectives, using different languages such as BPEL [61] and UML [62]. Addressing these tasks is made even more challenging when specifications change during the development cycle, when part of the system relies on legacy software, or when the system must evolve in a hostile environment, e.g., when a service relies on the dynamic discovery of other services, which might not be trusted.

In the last decade, message passing based distributed systems have been receiving much attention both from industry [64, 65, 67] and from the research community. From the research perspective, a major advantage is that such systems may be reasoned about via very well established theoretical tools such as the π -calculus [56]. The suitability of

these systems for theoretical work lead to a new approach to distributed software design and verification where types are not only used for checking, e.g., expressions and functions, but also to check that, e.g., the *behaviour* of a client matches the behaviour of a service offered by a server. Behavioural types, or *session types*, are used to assign an abstract behaviour to channels which are used by distributed programs to interact with one another.

Additionally, *choreography* languages such as WS-CDL [34] have been developed to model the interactions between different services from a global point of view. Notably, they allow software designers of different parts of a system to define – jointly – the rules of participation in the system, without delegating the control of their part of the system to another party. A groundbreaking work [44] relates choreography languages and session types to promote a choreography-driven methodology for distributed software development. The methodology it advocates, summarised in Figure 1.1, is as follows: a team of programmers designs a global view of the interactions to be implemented (i.e., a choreography), then the choreography is *projected* onto each role, or participant, finally each program implementing one or more roles in the choreography is *type-checked* against its corresponding projection(s). This approach permits, if the choreography satisfies some properties, to guarantee safety properties such as deadlock freedom and to ensure that an implementation abide by the specified protocol.

This is an ideal methodology, but it may not always be possible to design one set of choreographies that will drive the overall development of a distributed system. Indeed, software need maintenance, specifications may evolve (sometimes also during the development), and issues may arise during the implementation phase. Therefore, there is a need for an alternative approach whereby it is possible to infer a global point of view of a distributed system (i.e., a choreography) from local behavioural specifications (i.e., session types). Without such an approach the burden of (re-) constructing the choreographies of a distributed system is on the practitioners. Often, they may choose to omit this step,

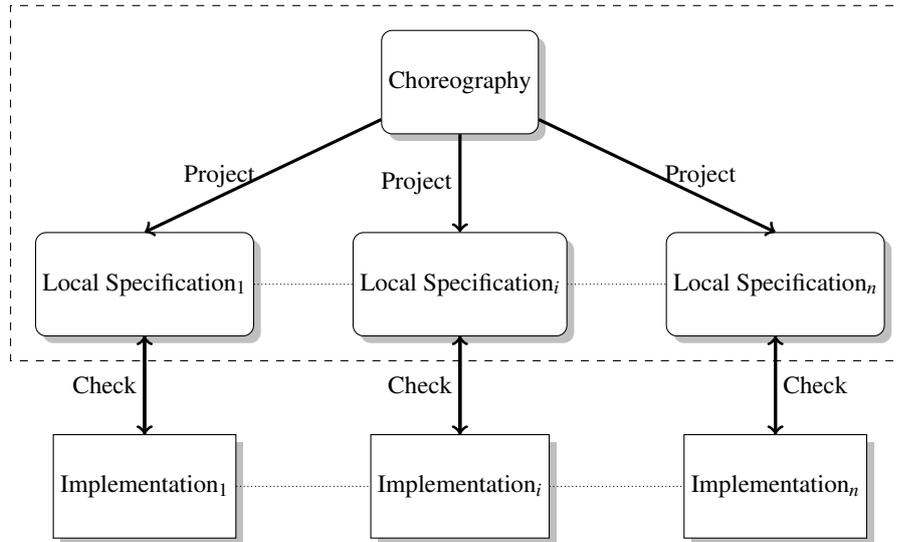


Figure 1.1: Choreography driven development

therefore creating a gap between an implementation and its specification. We propose to tackle the problem of synthesising choreographies from local behavioural specifications via a type system which assigns a choreography to a set of session types, if possible.

Obtaining a choreography from local specifications permits not only the methodology above to be applicable and thus to guarantee safety properties, but also allows to keep an implementation closer to its specification. Indeed, as the implementation progresses, one may infer choreographies from it and check that it still matches the specification. If there is a gap between implementation and specification, then it is possible to adapt both specification and implementation in a semi-automatic manner (using choreography projections and synthesis).

Keeping specifications and implementations close is not the only advantage of having a means to synthesise a choreography from local specifications. In fact, once a choreography has been obtained it may be refined into a more informative formalism, e.g, a *global assertion*, that is a choreography decorated with logical formulae specifying senders' obligations and receivers' requirements. We will see however that refining a choreography into a global assertion is not always a trivial task. We will give a few techniques to facilitate the task of the designer so that global assertions satisfy essential safety

properties.

In addition, synthesising choreographies may also be done at runtime to characterise a set of services, or contracts, which are fulfilling each other's requirements. We say that a set of contracts is compliant if it may be assigned a choreography. An advantage of having a choreography to represent the fact that a set of services is compliant is that it permits to reason about meta-level properties of the interactions, e.g., if the interactions always terminate or if the involved services will share some common properties once the interactions have terminated.

The main objectives of the thesis are as follows. Firstly, we tackle the problem of constructing a global view of a distributed system (i.e., a choreography) from local specifications: we focus on the top part of Figure 1.1 and reverse the “project” arrows. Secondly, we demonstrate the interest of obtaining a choreography from local specifications by (i) giving techniques to help refine a choreography into a global assertion and (ii) by introducing a model for distributed systems which use choreographies as a basis for combining services together.

1.2 Synopsis and Contributions

The thesis is divided in six chapters which we briefly summarise below.

- In the rest of this chapter, we list the published works on which the thesis is build.
- In Chapter 2, we introduce the context and foundations of our work and discuss related literature.
- In Chapter 3, we give a construction to synthesise a choreography from local specifications. This chapter represents the core contribution of this thesis: a type system which assigns a global type to a set of local types.

- In Chapter 4, we present three algorithms and a methodology to help designers decorate a choreography with logical predicate stating senders' obligations and receivers' requirements.
- In Chapter 5, we introduce a formal model for distributed systems where the synthesis of choreographies is used to define a compliance relation between contracts advertised by participants willing to cooperate.
- In Chapter 6, we conclude the thesis and discuss future research directions.

1.3 Publications

The thesis is based on the following publications (in chronological order). We describe their correspondance with the content of the thesis and the contributions of the author.

- **A Modular Toolkit for Distributed Interactions.** Julien Lange and Emilio Tuosto, *PLACES 2010*, [52].
 - *Author's contribution:* This is a paper of my own.
 - This work does not appear in the thesis, but the toolkit it describes was used (i) to identify the problem that Chapter 4 tackles (i.e., designing a *well-asserted* global assertion is not an easy task) and (ii) to test the algorithms presented in Chapter 4.
- **Amending Contracts for Choreographies.** Laura Bocchi, Julien Lange and Emilio Tuosto, *ICE 2011*, [13].
 - *Author's contribution:* In this work, my contributions include: part of the writing, part of the ideas and formalisations of the three algorithms; I had a major contribution in the proofs, the methodology, and the idea of the lifting algorithm.

- *Corresponding part:* Chapter 4.
- **Three Algorithms and a Methodology for Amending Contracts for Choreographies.** Laura Bocchi, Julien Lange and Emilio Tuosto, *Sci. Ann. Comp. Sci.*, 22, [14].
 - *Author's contribution:* This paper is an extended version of [13] above. This version includes all the proofs.
 - *Corresponding part:* Chapter 4 is essentially an extended and updated version of [14].
- **Synthesising Choreographies from Local Session Types.** Julien Lange and Emilio Tuosto, *CONCUR 2012*, [53].
 - *Author's contribution:* This is a paper of my own.
 - *Corresponding part:* Chapter 3 is major revision of [53]. The technicalities were updated in order to match recent developments on the relationship between communicating machines and session types. In addition, all the proofs are included in the chapter and many additional examples were added.
- **Choreography Synthesis as Contract Agreement.** Julien Lange and Alceste Scalas, *ICE 2013*, [51].
 - *Author's contribution:* In this work, my contributions include: part of the writing, part of the ideas and formalisations of our adaptation of CO₂, the variations of the fuse primitive, and the proofs.
 - *Corresponding part:* Chapter 5 is an extended version of [51], with all the proofs, additional examples and, notably, a larger discussion on variations of the fuse-primitive.

We review the key theories on which our work relies and discuss a few related approaches. We focus on *(i)* the basis on which our work is build – in Section 2.1 – and *(ii)* on works that study the relationship between global and local specifications – in Section 2.2.

2.1 Foundations

The theories based on session types stand out as effective methodologies to specify and verify properties of distributed systems. Amongst the research community, they have become a popular framework to reason about concurrent programs which communicate via (usually asynchronous) communication channels.

Since the first paper introducing session types as a means to design and verify properties of distributed systems in the mid-nineties, many extensions and applications have been proposed by many authors. These range from theoretical works establishing a Curry-Howard-like isomorphism between linear logic and session types [19, 20, 70], to works embedding session type theory into programming languages [45, 58–60], and industrial applications [64, 65].

In this section, we introduce the key elements of session type theory. We first present the theory of dyadic session types and its extension to multiparty session types. We then present two extensions which are essential for the next chapters: an extension which introduces design-by-contract for distributed interactions and an extension of the multiparty session types framework to communicating machines.

2.1.1 Dyadic Session Types

Session types were first introduced in [43, 68] by Honda et al. as a means to structure communications in concurrent programming. Analogously to data types in traditional programming language, where, e.g., type checking at compilation time ensures that the arguments of a function match its definition, dyadic session types guarantee that, e.g., the behaviour of a client-side application matches the behaviour of a service invoked on the server-side.

The theory relies on a few ingredients that we describe below. The key ingredient is the notion of *session*, that is a set of *structured* interactions which corresponds, for instance, to a communication protocol (from the communication initialisation to the termination). We give an example of a session in Example 2.3 below.

The framework introduced in [43] – and most of its extensions – is based on three main components: (i) a language used to write the processes which take part in the interactions (usually a variation of the π -calculus [56]); (ii) a syntax for *types* that describe the structure of the interactions (i.e., session types); and (iii) a type discipline which verifies that processes match their types and guarantees safety properties of the interactions. We describe informally each ingredient and illustrate how the theory works.

A session-oriented calculus. The first component is a calculus to model processes which interact via communication channels. In order to structure the interactions between parties, a few *interaction primitives* are made available to the programmers. They include

communication mechanisms à la π -calculus that deal with session channels as first-class values. We focus on the primitives and constructs of the language which directly relate to the rest of this thesis.

In the rest of this chapter, we let R range over process, n over public names, k over channels, and e over values.

$$\begin{array}{ll}
 R ::= \bar{n}(k) \text{ in } R & \text{[SESSION REQUEST]} \\
 | n(k) \text{ in } R & \text{[SESSION ACCEPTANCE]} \\
 | k!\langle e \rangle; R & \text{[SEND]} \\
 | k?(x); R & \text{[RECEIVE]} \\
 | k \triangleleft l; R & \text{[SELECT]} \\
 | k \triangleright \{l_1 : R_1, \dots, l_n : R_n\} & \text{[BRANCH]} \\
 | (\nu k) R & \text{[RESTRICTION]} \\
 | R | R' & \text{[PARALLEL]} \\
 | \dots &
 \end{array}$$

Primitives [SESSION REQUEST] and $\text{[SESSION ACCEPTANCE]}$ deal with session creation. The former indicates a request for a new session on public name n (i.e., a name known to other processes). Primitive $\text{[SESSION ACCEPTANCE]}$ is its dual: it indicates the acceptance of a new session. These two primitives bind k in their respective continuations.

Example 2.1. A process implementing an FTP server would start with a primitive $n(k) \text{ in } R$, indicating that it is waiting for a client request. Here n may be understood as *port 21*, k may be understood as the port number through which the data exchange will occur, and R is the implementation of the body of the server.

Similarly, a process implementing an FTP client would start with a primitive $\bar{n}(k) \text{ in } R$, indicating that the client sends a request to a server via port 21 (or n).¹ \diamond

¹ Note that n may also be understood as a pair of IP address and port number.

Primitives $[\text{SEND}]$ and $[\text{RECEIVE}]$ allow processes to exchange data. The former indicates the sending of a value e on a channel k ; while the latter indicates the reception of a datum from channel k , and binds x in the continuation R ; once a value is received, x is replaced by that value in R .

Primitives $[\text{SELECT}]$ and $[\text{BRANCH}]$ provide a *select* and *branch* mechanism which allows participants to make an internal choice or to let the environment make a decision (external choice). Notably, the mechanism allows to provide a simple model of (distributed) method invocations, where a label l models the name of a method. Primitive $[\text{SELECT}]$ indicates the sending of a label l on channel k ; and primitive $[\text{BRANCH}]$ indicates a choice made by another process, if label l_i is selected, then R_i is executed.

The construct $[\text{RESTRICTION}]$ allows to restrict the use of channel k to the process R . Note that $[\text{RESTRICTION}]$ is a runtime construct, i.e., it is not available to the programmer and appears only when sessions have been started. Concurrent processes are specified via construct $[\text{PARALLEL}]$ which indicates the concurrent execution of processes R and R' .

In addition to these primitives, the language features standard constructions such as if-then-else, process definitions and invocations, terminated process $\mathbf{0}$ (trailing occurrences are often omitted), and special primitives for sending and receiving session names. These special primitives allow *delegation*, namely the fact that session channels may be communicated so to allow a process to delegate its “role” to another process.

We illustrate the syntax of the language in Example 2.2 below.

Example 2.2. Consider the two processes R_s and R_c below.

$$\begin{aligned}
 R_s &= n(k) \text{ in} \\
 & k \triangleright \{ \quad i : k?(x); k!\langle x * 2 \rangle, \\
 & \quad \quad \quad b : k?(y); k!\langle -y \rangle \quad \}
 \end{aligned}$$

The process R_s may be understood as a server which waits for a request on a public name

n , i.e., a name known by other processes in the system. Once a request has been received, the server offers two operations, or branches. If the branch i is selected, then the server receives an integer, and returns the double of this integer. If the branch b is selected, the server receives a boolean and returns its negation.

$$R_c = \bar{n}(k) \text{ in} \\ \text{if}(\dots) \text{ then } k \triangleleft i; k! \langle 3 \rangle; k?(x) \\ \text{else } k \triangleleft b; k! \langle \text{true} \rangle; k?(y)$$

The process R_c is a client which, essentially, exhibits the dual behaviour. Once its session request has been accepted, it chooses either branch offered by the server – we abstract from the actual decision here. If the client chooses the i (resp. b) branch, it sends the value 3 (resp. true) via channel k , and expects an answer which will be stored in variable x (resp. y). \diamond

The semantics of the calculus is similar to the one of the π -calculus (but for the session creation primitives) and we illustrate it with Example 2.3.

Example 2.3. Composing the two processes of Example 2.2 in parallel, we obtain $R_s \mid R_c$, which reduces to

$$(\nu k') \left(\begin{array}{l} k' \triangleright \{ \quad i : k'?(x); k'! \langle x * 2 \rangle, \\ \quad b : k'?(y); k'! \langle \neg y \rangle \} \end{array} \left| \begin{array}{l} \text{if}(\dots) \text{ then } k' \triangleleft i; k'! \langle 3 \rangle; k'?(x) \\ \quad \text{else } k' \triangleleft b; k'! \langle \text{true} \rangle; k'?(y) \end{array} \right. \right)$$

where k' is a (fresh) channel restricted to these processes (i.e., no other process may use it).

At this point we can say that a *session* has started: two processes have started a new session by “meeting” via a public name n and they may now communicate “privately” via the fresh channel k' .

If the client chooses the i branch, the system reduces to

$$(\nu k') (k' \triangleright \{i : k'?(\mathbf{x}); k'!\langle \mathbf{x} * 2 \rangle, b : k'?(\mathbf{y}); k'!\langle \neg \mathbf{y} \rangle\} \quad | \quad k' \triangleleft i; k'!\langle 3 \rangle; k'?(\mathbf{x}))$$

and after a few steps, we obtain:

$$(\nu k') (k'!\langle 3 * 2 \rangle \quad | \quad k'?(\mathbf{x}))$$

and the process terminates once the client has received the value 6 via channel k' . \diamond

Session types. The next components of the theory are *session types*, also called local types, or end-point types. Essentially, a session type represents the abstract *behaviour* of a process on a given communication channel. We give their syntax below – adapted from [43] – and let P range over session types and a over message sorts.

$$\begin{array}{l}
P ::= !a;P \quad \text{[SEND]} \\
| ?a;P \quad \text{[RECEIVE]} \\
| \oplus\{l_1 : P_1, \dots, l_n : P_n\} \quad \text{[INTERNAL CHOICE]} \\
| +\{l_1 : P_1, \dots, l_n : P_n\} \quad \text{[EXTERNAL CHOICE]} \\
| \mu\chi.P \quad \text{[RECURSION]} \\
| \chi \quad \text{[RECURSIVE CALL]} \\
| \mathbf{0} \quad \text{[END]}
\end{array}$$

where a ranges over basic data types (or *sorts*) such as `bool`, `int`, etc.² There is a close correspondence with the primitives of the calculus presented above. Types `[SEND]` and `[RECEIVE]` correspond to the send and receive primitives. Types `[INTERNAL CHOICE]` and `[EXTERNAL CHOICE]` correspond to the select and branch primitives. Note that type `[INTERNAL CHOICE]` says that

² These may also include session types, but we abstract from such constructions in this thesis.

any of the labels l_i may be *sent*, then behaviour P_i takes places. Types $_{[RECURSION]}$ and $_{[RECURSIVE CALL]}$ are used to represent recursive behaviours. Finally, type $_{[END]}$ indicates the termination of the session.

An important notion regarding dyadic session types is the one of *co-type*. The co-type of a type P , written \bar{P} , is defined as follows:

$$\begin{aligned} \overline{!a;P} &= ?a;\bar{P} & \overline{\oplus\{l_1 : P_1, \dots, l_n : P_n\}} &= +\{l_1 : \bar{P}_1, \dots, l_n : \bar{P}_n\} & \bar{\mathbf{0}} &= \mathbf{0} \\ \overline{?a;P} &= !a;\bar{P} & \overline{+\{l_1 : P_1, \dots, l_n : P_n\}} &= \oplus\{l_1 : \bar{P}_1, \dots, l_n : \bar{P}_n\} \\ \overline{\mu\chi.P} &= \mu\chi.\bar{P} & \bar{\chi} &= \chi \end{aligned}$$

Typing discipline. The final component of the framework is a *typing discipline* which guarantees safety properties of the system under consideration. One of the main contributions in [43] is a type system that, given a process R , permits to assign session types to the channels used by R . The reader does not need to have a precise understanding of the type system in the scope of this thesis, thus we only highlight the basic ideas.

The main judgement of the type system is of the form:

$$\Theta; \Gamma \vdash R \triangleright \Delta$$

where Θ is an environment mapping process variables to processes (to deal with process definitions), environment Γ maps free names (e.g., public name n in Example 2.2) to pairs of local types, and expressions and variables to sorts. Environment Δ maps free channels (such as channel k in Example 2.2) to local types. We give a couple of typing rules to show the flavour of the type system.

$$\frac{\Theta; \Gamma \vdash R \triangleright \Delta \cdot k : P}{\Theta; \Gamma \cdot n : \langle P, \bar{P} \rangle \vdash n(k) \text{ in } R \triangleright \Delta} \qquad \frac{\Gamma \vdash e : a \quad \Theta; \Gamma \vdash R \triangleright \Delta \cdot k : P}{\Theta; \Gamma \vdash k!\langle e \rangle; R \triangleright \Delta \cdot k : !a; P}$$

The left-hand side rule validates session acceptance primitive; if the environment Γ maps a name n to a pair of dual local types and if the continuation is typable, with $k : P$ in environment Δ . Note that the rest of the derivation must ensure that the process uses the channel k as advertised by type P . The right-hand side rule validates sending actions; if the continuation is typable and the expression e is of type a .

Example 2.4. Continuing with Example 2.2, the type of channel k , from the point of view of R_s (resp. R_c) is given by P_s (resp. P_c) below.

$$P_s = +\{i : ?\text{int}; !\text{int}, b : ?\text{bool}; !\text{bool}\} \quad P_c = \oplus\{i : !\text{int}; ?\text{int}, b : !\text{bool}; ?\text{bool}\}$$

It is easy to see that we have $\overline{P_c} = P_s$ and the public name n has type $\langle P_s, P_c \rangle$. \diamond

The main results of [43] are that (i) if a process is typable then it will never reduce to an error (e.g., a process receives string but was expecting a bool) and (ii) typing is preserved by reduction (*subject reduction*).

2.1.2 Multiparty Session Types

The initial theory of session types allowed only to reason about bi-party communications. In particular, the notion of *co-type* only makes sense for sessions that involve two participants. Following preliminary results in [24, 25], the dyadic framework was extended in [44] to a methodology to design, implement, and verify distributed software. The methodology – summarised in Figure 2.1 – advocates that (i) a team of programmers specify a global description of the protocol, i.e., a *choreography* (or global type). Then, (ii) the choreography is *projected* onto each *participant* (or role) so to obtain a *local type*, i.e., a description of the role that the participant plays in the session. Finally, (iii) processes are type-checked to ensure that they validates their respective local types.

In the rest of this section, we present the main ideas introduced in [44], using the

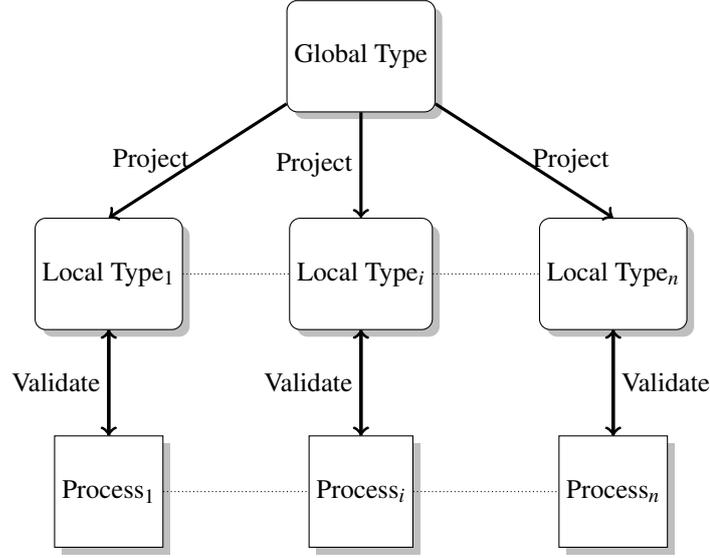


Figure 2.1: The multiparty session types approach

simplified syntax from [11].

Remark 2.1. *We will only briefly discuss the processes from which local types are extracted from. In this thesis, we focus on the relationship between global and local types, i.e., the top part of Figure 2.1. The extraction of session types from processes has been studied extensively in, e.g., [31, 43, 44].*

Choreographies. Choreographies are specified by global types. Intuitively, a global type establishes the interaction pattern for the harmonious coordination of distributed parties. The syntax of global types is given by the following productions:

$$\begin{array}{ll}
 \mathcal{G} ::= & \mathbf{s} \rightarrow \mathbf{r} : \mathbf{a}; \mathcal{G}' \quad \text{[VALUE PASSING]} \\
 | & \mathbf{s} \rightarrow \mathbf{r} : \{l_i : \mathcal{G}_i\}_{i \in I} \quad \text{[BRANCHING]} \\
 | & \mathcal{G}' \mid \mathcal{G}'' \quad \text{[PARALLEL]} \\
 | & \mu \chi. \mathcal{G}' \quad \text{[RECURSION]} \\
 | & \chi \quad \text{[RECURSIVE CALL]} \\
 | & \mathbf{0} \quad \text{[END]}
 \end{array}$$

Type $[\text{VALUE PASSING}]$ says that participant s sends a message (of sort a) to participant r , then the interactions in \mathcal{G}' take place. Type $[\text{BRANCHING}]$ says that participant s sends one of the labels l_i to participant r . If l_j is sent then the interactions in \mathcal{G}_j take place. Type $[\text{PARALLEL}]$ represents the concurrent run of the interactions described in \mathcal{G}' and \mathcal{G}'' . Type $[\text{RECURSION}]$ is a recursive type for recursive conversation structure, assuming that type variables ($[\text{RECURSIVE CALL}]$) are guarded in the standard way. Type $[\text{END}]$ represents the termination of the session.

We illustrate the theory via Examples 2.5, 2.6, and 2.7 adapted from [44].

Example 2.5. Below is a global type specifying a protocol with two buyers (b_1 and b_2) and a seller (s). The buyers b_1 and b_2 want to purchase a book from s by combining their funds.

$$\mathcal{G}_{\text{ex2.5}} = b_1 \rightarrow s : \text{string}; \quad (1)$$

$$s \rightarrow b_1 : \text{int}; \quad (2)$$

$$b_1 \rightarrow b_2 : \text{int}; \quad (3)$$

$$b_2 \rightarrow s : \{ \text{ok} : b_2 \rightarrow s : \text{string}; \quad (4)$$

$$s \rightarrow b_2 : \text{date}; \mathbf{0}, \quad (5)$$

$$\text{quit} : \mathbf{0} \} \quad (6)$$

In (1), b_1 and s interact and exchange a book title. In (2), s sends a quote to b_1 . In (3), b_1 tells b_2 its contribution to the purchase. In (4-6), b_2 may refuse (selecting label quit) or accept the deal (selecting label ok). In the former case, the protocol terminates there, otherwise it continues as in (4-5), namely b_2 and s exchange delivery address and date. \diamond

Remark 2.2. In [44], interactions specify a channel over which participants communicate. For simplicity, we use the syntax of [11] and omit channel identities. This approach is also reflected in recent works such as [33, 39, 40, 73].

When interactions specify channels, global types are required to satisfy a property

called linearity [44]. Linearity essentially ensures that no race may occur on a channel, i.e., each message put on a channel may be received by exactly one participant.

Local types & projections. Local types are similar to those presented in the previous section, except that they now specify to (resp. from) whom a message is sent (resp. expected). The syntax of (multiparty) *local types* is as follows

$$\begin{array}{ll}
P ::= r!a;P & \text{[SEND]} \\
| s?a;P & \text{[RECEIVE]} \\
| r\oplus\{l_1 : P_1, \dots, l_n : P_n\} & \text{[INTERNAL CHOICE]} \\
| s+\{l_1 : P_1, \dots, l_n : P_n\} & \text{[EXTERNAL CHOICE]} \\
| \mu\chi.P & \text{[RECURSION]} \\
| \chi & \text{[RECURSIVE CALL]} \\
| \mathbf{0} & \text{[END]}
\end{array}$$

For instance, the type [INTERNAL CHOICE] says that one of the labels l_i may be sent to r , then the interactions P_i take place.

If it satisfies certain conditions – which we will discuss in length in Chapter 3 – a global type may be *projected* onto each participant of the choreography. The projection is a simple syntactic operation which, given a global type \mathcal{G} and a participant s , returns a local type; we write $\mathcal{G} \downarrow_s$ for the projection of \mathcal{G} onto participant s .

Example 2.6. We give the local types obtained by projecting the global type $\mathcal{G}_{\text{ex2.5}}$ of Example 2.5:

$$\begin{aligned}
\mathcal{G}_{\text{ex2.5}} \downarrow_{b_1} &= s!\text{string}; s?\text{int}; b_2!\text{int}; \mathbf{0} \\
\mathcal{G}_{\text{ex2.5}} \downarrow_{b_2} &= b_1?\text{int}; s\oplus\{\text{ok} : s!\text{string}; s?\text{date}; \mathbf{0}, \text{quit} : \mathbf{0}\} \\
\mathcal{G}_{\text{ex2.5}} \downarrow_s &= b_1?\text{string}; b_1!\text{int}; b_2 + \{\text{ok} : b_2?\text{string}; b_2!\text{date}; \mathbf{0}, \text{quit} : \mathbf{0}\}
\end{aligned}$$

The meaning of each type is straightforward, knowing the global type of Example 2.5. For instance, the type $\mathcal{G}_{\text{ex2.5}} \downarrow_s$ says that the participant waits for a message of type string from participant b_1 and replies with a message of type int. Then it let b_2 choose either the quit branch, in which case its behaviour is finished; or the ok branch, in which case it receives a string from b_2 and sends back a date. \diamond

Session calculus. The main differences between the syntax of *processes* in the dyadic setting and the one in the multiparty setting are as follows. First, the semantics now adopts asynchronous communications, i.e., the sending of messages is a non-blocking operation. Second, the primitives for session request and acceptance now specify the role a process is playing in the session. We illustrate the syntax and semantics of the calculus in Example 2.7. below.

Example 2.7. Consider the three processes below, implementing the protocol specified by the global type $\mathcal{G}_{\text{ex2.5}}$ in Example 2.5.

$$\begin{aligned}
R_{b_1} &= \bar{n}_{[b_1]}(k) \text{ in } k!s \langle \text{”Loving Sabotage”} \rangle; k?s(x); k!b_2 \langle x/2 \rangle \\
R_{b_2} &= n_{[b_2]}(k) \text{ in } k?b_1(y); \text{if}(y > 50) \text{ then } k \triangleleft s : \text{quit}; \mathbf{0} \\
&\quad \text{else } k \triangleleft s : \text{ok}; k!s \langle \text{”Leicester”} \rangle; k?s(y') \\
R_s &= n_{[s]}(k) \text{ in } k?b_1(z); k!b_1 \langle 70 \rangle; k \triangleright b_2 \{ \text{quit} : \mathbf{0}, \\
&\quad \text{ok} : k?s(z'); k!s \langle \text{today} \rangle \}
\end{aligned}$$

The process R_{b_1} is the implementation of the first buyer and is the one requesting the creation of a new session. Once the session has been established, b_1 sends the title of a book to the seller s , along channel k , b_1 then expects a quote from the seller. Finally, b_1 sends the amount of money it is willing to share with b_2 . The process R_{b_2} is the implementation of the second buyer. When b_2 has accepted the session request, b_2 receives the price that b_1 is willing to pay then decides whether or not to proceed. If the price is greater than

50, b_2 aborts the purchase, otherwise, the buyer sends its address and expects and answer from the seller. The process R_s is the implementation of the seller and is essentially symmetric to the other two processes. Here, we abstract from the way s retrieves the price associated with a book and assume that all books cost 70, for simplicity.

The composition of these processes reduces as

$$R_{b_1} \mid R_{b_2} \mid R_s \quad \rightarrow \quad (\nu k') (R'_{b_1} \mid R'_{b_2} \mid R'_s)$$

where k' is a (fresh) channel private to these three processes and R'_{b_1} , R'_{b_2} , and R'_s are as follows.

$$R'_{b_1} = k'_{[b_1]}!s \langle \text{"Loving Sabotage"} \rangle; k'_{[b_1]}?s(x); k'_{[b_1]}!b_2 \langle x/2 \rangle$$

$$R'_{b_2} = k'_{[b_2]}?b_1(y); \text{if}(y > 50) \text{ then } k'_{[b_2]} \triangleleft s : \text{quit}; \mathbf{0}$$

$$\text{else } k'_{[b_2]} \triangleleft s : \text{ok}; k'_{[b_2]}!s \langle \text{"Leicester"} \rangle; k'_{[b_2]}?s(y')$$

$$R'_s = k'_{[s]}?b_1(z); k'_{[s]}!b_1 \langle 70 \rangle; k'_{[s]} \triangleright b_2 \{ \text{quit} : \mathbf{0},$$

$$\text{ok} : k'_{[s]}?s(z'); k'_{[s]}!s \langle \text{today} \rangle \}$$

Observe that each instance of channel k has been replaced by a channel of the form, e.g., $k'_{[b_2]}$. For instance $k'_{[b_2]}?b_1(y)$ says that participant b_2 expects a message from b_1 on the session channel k' .³ Similarly, the prefix $k'_{[b_2]} \triangleleft s : \text{quit}$ says that participant b_2 sends the label `quit` to participant s , via channel k' . \diamond

Typing & results. Typing processes in a multiparty setting is quite similar to the binary case. The main difference being that instead of checking whether a public name n is assigned a pair of dual types, one checks that a public name is typed by a global type and that each process taking part in the session may be typed by one of the projections of the

³ The semantics allows for message from different sender to be re-ordered within a channel such as k' .

global type. For instance, the rule for typing session acceptance is of the form:

$$\frac{\Gamma \vdash n : \langle \mathcal{G} \rangle \quad \Theta; \Gamma \vdash R \triangleright \Delta \cdot k : \mathcal{G} \downarrow_s}{\Theta; \Gamma \vdash n_{[s]}(k) \text{ in } R \triangleright \Delta}$$

which requires that n is of global type \mathcal{G} and that the way participant s uses channel k must match the projection of \mathcal{G} onto s .

The main result in [44] is that, for typable processes, *progress* is guaranteed within a session, i.e., given a set of processes involved in a session (which is not interleaved with any other session), these processes will either be able to reduce further or terminate. Another result shows that the behaviour of a typable process will always follow the interactions specified by the global type, this property is called *session fidelity*.

2.1.3 Design-by-Contract for Distributed Multiparty Interactions

The multiparty sessions types framework was extended in [12], applying the ideas of Design-by-Contract [55] to the specification and verification of distributed systems.

The main new notion is the one of *global assertions* which decorate global types with logical formulae (*predicates*) that constrain interactions, declaring senders' obligations and receivers' requirements on the values of the exchanged data and on the choice of the branches to follow. This adds fine-grained constraints to the specification of the interaction structure.

Example 2.8. Consider for instance the *global assertion* below, where the values of the messages are represented by the *interaction variables* x and y .

$$\begin{aligned} \mathcal{G}_{\text{ex2.8}} &= \text{Alice} \rightarrow \text{Bob} : \{x \mid x > 0\}; \\ &\quad \text{Bob} \rightarrow \text{Carol} : \{y \mid y > x\} \end{aligned}$$

$\mathcal{G}_{\text{ex2.8}}$ describes a protocol with three participants, Alice, Bob, and Carol, who agree on a “contract” constraining the interaction variables x and y . The contract stipulates that (i) Alice has to send Bob a positive value x in the first interaction, and that (ii) Bob must send Carol a value y strictly greater than x , fixed by Alice in the first interaction. Notice that Bob can fulfill his pledge (i.e., the predicate $y > x$ in the second interaction above) only after he has received the value x from Alice. \diamond

Similarly to a global type, a global assertion \mathcal{G} may be projected on *endpoint assertions* that are local types constrained according to the predicates in \mathcal{G} . For instance, the projection on Alice in $\mathcal{G}_{\text{ex2.8}}$ is an endpoint assertion prescribing that Alice has to send a positive value to Bob. Endpoint assertions can be used for static validation of the actual processes implementing one or more roles in a choreography represented by \mathcal{G} , and/or to synthesise monitor processes for runtime checking/enforcement [32].

The calculus of [44] is augmented with predicates for checking that the sent or received values validate the predicates so to obtain the following primitives:

$$k!\langle e \rangle(x)\{\psi\};R \quad [\text{SEND}] \qquad k?(x)\{\psi\};R \quad [\text{RECEIVE}]$$

where e is a value, x a variable and ψ a predicate. The idea is that, in the $[\text{SEND}]$ case, if $\psi\{e/x\}$ evaluates to true, then the continuation R is executed, otherwise the process reduces to an error. Similarly, in the $[\text{RECEIVE}]$ case, the value received by the process at runtime, say e , must be so that $\psi\{e/x\}$ evaluates to true, otherwise the process reduces to an error.

If a global assertion is *well-asserted*, namely when it obeys two precise design principles: *history-sensitivity* and *temporal satisfiability*, it may be projected, and processes may be type-checked against the projections. If type-checking succeeds then the system is guaranteed to be error free.

Informally, history-sensitivity demands that a party having an obligation on a predi-

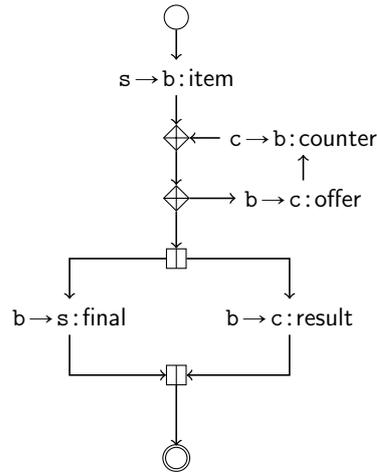


Figure 2.2: Example of generalised global type

cate has enough information for choosing a set of values that guarantees it. Instead, temporal satisfiability requires that the values sent in each interaction do not make predicates of future interactions unsatisfiable.

Remark 2.3. *We give a more precise account of global assertions, history-sensitivity, and temporal satisfiability in Chapter 4.*

2.1.4 Session Types and Communicating Machines

In [40], Deniélou and Yoshida introduced *generalised global types*, extending significantly the expressiveness of the global types of [44] (cf. Section 2.1.2) and characterising a relationship between communicating machines and multiparty session types. Below, we highlight the main novelties introduced in this work which we will refer to later in this thesis.

Overview. Generalised global types allow to describe protocols via general graphs, instead of a tree-like structure. Branching constructs, parallel operator and recursive definition are replaced by gates and edges which connect interactions of the form $s \rightarrow r : a$. A \diamond -gate indicates either a choice branch, merge, or a recursion; a \square -gate indicates either

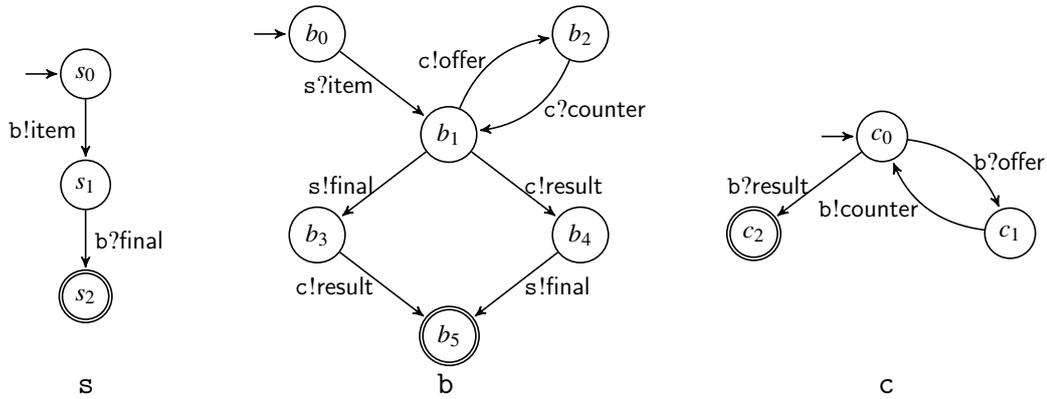


Figure 2.3: Example of communicating machines

fork or join of concurrent threads. We give an example of a generalised global type in Example 2.9 – borrowed from [40].

Example 2.9. The global type in Figure 2.2 specifies a protocol where “a seller (s) relies on a broker (b) to negotiate and sell an item to a client (c)” [40]. First, s sends a message `item` to b . The broker then chooses whether to (i) enter a loop where the broker and the client negotiate via `offer` and `counter-offer` messages; or (ii) the broker accepts the sale and terminates the protocol by sending concurrently messages `final` and `result` to the seller and the client, respectively. \diamond

Interestingly, if a generalised global type satisfies some conditions it may be projected onto local types that are, essentially, *communicating machines* – or communicating finite states machines (CFSM) [16]. This permits to identify a new class of CFSMs, called *Multiparty Session Automata*.

A system of communicating machines is a list of automata which communicate via unbounded FIFO channels, generally assuming that, given two machines, they communicate via a pair of distinguished uni-directional channels. We give an example of a system of CFSMs in Example 2.10 – again borrowed from [40].

Example 2.10. The machines in Figure 2.3 correspond to the projections of the global

type of Example 2.9. A transition $b!item$ in machine s says that s may put a message item on the queue from s to b . Once such a message is on the queue, machine b may fire its transition $-s?item$ – thus removing the message from the queue. \diamond

The authors also give a translation to/from traditional local types (such as the ones in Section 2.1.2) to communicating machines. Remarkably, these local types may now feature concurrent threads. For instance, the local type corresponding to machine b in Figure 2.3 explicitly shows that the actions $s!final$ and $c!result$ are executed concurrently.

Finally, a new session calculus is introduced in [40] so to allow these extensions to be fully supported at the programming language level.

Remark 2.4. *Global types, local types, and calculus in [40] are formally defined in terms of a set of mutually recursive variables which may be depicted as graphs (or automata) as shown in this section. Here, we only use the graphical notation for the sake of clarity.*

Safety properties. For a generalised global type to be projected onto “safe” local types, it must satisfy two essential conditions: *local choice* and *linearity*.

- The local choice condition requires the following for each choice in a global type.
 - There is exactly one participant who makes the decision, i.e., for each \diamond -gate with several outgoing edges, the first action of all the branches should have the same sender.
 - The choice must be propagated to the other participants involved in the choice, i.e., it cannot be the case that a participant behaves differently in two branches without knowing which branch was chosen by the sender.
- The linearity condition ensures that no races occur within the process of a participant. To achieve this, the authors “impose that no participant can be faced with two concurrent receptions where messages can have the same label” [40].

In addition, the authors require the generalised global types to be “sane”, e.g., there should be exactly one starting point in the choreography, there should be at most one finishing point, and all threads must be collected by a join.

When these conditions are satisfied, the global type may be projected onto its participant. The system of CFSMs thus obtained is then guaranteed to satisfy the usual safety and liveness properties. In particular, the projections are deadlock free and all messages sent are eventually received. We discuss these properties in more details in Section 3.3.2.

2.2 Related Work

We discuss recent works that tackle the design and verification of distributed systems and focus on those that are based either on session type theory or on choreographies. In Section 2.2.1, we present different approaches that formalise the relationship between global and local specifications. In Section 2.2.2, we discuss works that tackle the problem of inferring a global specification from local ones. In Section 2.2.3, we report a few major extensions of the theories based on session types. In Section 2.2.4, we discuss briefly a few approaches that tackle the verification of (multiparty) session-based distributed systems without choreographies.

2.2.1 Global versus Local Specifications

A main concern in choreography-driven development is that the projections of a choreography should respect the structure of the interactions specified by the choreography. We distinguish three main lines of work that tackle this problem. First, Castagna et al. [28, 29] give a presentation of global types as a formal language and relate global and local types via a trace semantics. Second, Lanese et al. [17, 47, 50] give different notions of *conformance* that specify whether the projections of a choreography indeed match the behaviour specified by the choreography. Third, Bultan et al. [9, 10] focus on the notion of realis-

ability. Essentially, a choreography is realisable if “it is possible to build a distributed system that communicates exactly as the choreography specifies” [9].

Global types and multiparty sessions. In [28, 29], the authors introduce a variation of global types, give them a trace semantics, and discuss their properties with respect to their projections. We survey the part of this work that is of specific interest in the scope of this thesis.

The authors allow general sequencing in global types, e.g., types of the form $(\mathcal{G}_1 \mid \mathcal{G}_2); (\mathcal{G}_3 \mid \mathcal{G}_4)$, but restrict recursive global types to iterations, i.e., global types of the form \mathcal{G}^* – the usual Kleene star. This allows them to define global types as a traditional formal language and have a quite natural definition of *traces* of global types (and local types).

They identify a few basic properties that an implementation of a global type should guarantee, we describe them below.

Sequentiality says that if a global type specifies an order between interactions, then all executions of its implementation must respect this ordering. For instance, in the global type:

$$s_1 \rightarrow r_1 : a; s_2 \rightarrow r_2 : b$$

if $\{s_1, r_1\} \cap \{s_2, r_2\} = \emptyset$ it is not possible for an implementation to guarantee that these two interactions will be executed in the prescribed order.

Alternativeness states that if a global type specifies an exclusive choice between two sets of interactions, then all the executions of its implementation must exhibit exactly one of these sets.

Shuffling states that all the executions of the implementation of a global type must exhibit some shuffling of the actions prescribed by the global type. For instance,

consider:

$$s_1 \rightarrow r_1 : a \quad | \quad s_2 \rightarrow r_2 : b$$

if an implementation of this global type exhibits the action on a but not the one on b , then it violates the specification.

Fitness says that all the sequences exhibited by an implementation of a global type must “fit” the interactions prescribed by the global type. Considering the global type

$$s \rightarrow r : a \quad + \quad s \rightarrow r : b \tag{2.1}$$

if an implementation exhibits an action unspecified by (2.1), say c , then the fitness condition is violated, since the action c is not prescribed by the global type.

Exhaustivity states that if a sequence of interactions is specified by a global type, then there must exist at least one execution of its implementation that exhibits these interactions (up-to possible re-ordering). For instance, an implementation of (2.1) such that s never sends b violates the condition.

Besides the global types violating the basic condition on sequentiality, the authors identify two kinds of “more serious” flaws, which may be illustrated as follows. The global type in (2.2) specifies a choreography where participant n has to behave differently in each branch of the choice, but is unaware of which branch was selected by s .

$$s \rightarrow r : a; r \rightarrow n : a; n \rightarrow s : a \quad + \quad s \rightarrow r : b; r \rightarrow n : a; n \rightarrow s : b \tag{2.2}$$

The global type in (2.3) specifies a choreography where the choice is “distributed” between participants s and n . This sort of choreography cannot be implemented without

covert channels to allow participants to agree on which branch to choose.

$$s \rightarrow r : a \quad + \quad r \rightarrow s : b \quad (2.3)$$

Remarkably, the authors are able to characterise these issues in terms of an ordering on the traces of global types and traces of local types.

Remark 2.5. *Observe that the local choice condition of [40] excludes problematic global types such as (2.2) and (2.3).*

Another important contributions of this work, is a novel (stronger) notion of progress which, essentially, states that, given an implementation of a global type, every participant must be able to (eventually) reach a terminated state. This is stronger than the progress of, e.g, [44] which only requires that the implementation (as a whole) is always able to reduce further or terminate.

Conformance. In [50], the authors identify two approaches to choreographies. The first approach, called *interaction-oriented*, is best represented by choreography languages such as WS-CDL [34] and the global types of Section 2.1.2. In our setting, we may consider interaction-oriented choreographies (IOC) as global types. For instance, adapting the first example of [50] to our syntax we have that

$$s_1 \rightarrow r_1 : a; s_2 \rightarrow r_2 : b \quad (2.4)$$

is a simple IOC describing that roles s_1 and r_1 interact on an operation a , then roles s_2 and r_2 interact on an operation b .

The second approach, called *process-oriented*, is best illustrated by BPEL4Chor [38] (extending WS-BPEL [61] to choreographies). In our setting, process-oriented chore-

ographies (POC) may be represented by systems of local types; for example

$$s_1[!a] \mid r_2[?a] \mid s_2[!b] \mid r_2[?b] \quad (2.5)$$

corresponds to the IOC of (2.4), i.e., system (2.5) may be seen as the projections of the choreography in (2.4). In fact, these two “choreographies” may exhibit different behaviours, e.g., there is no ordering between the actions on a and those on b in (2.5).

The authors define two languages, one for IOC and one for POC. Both languages consist of essentially the same constructs: sequential and parallel compositions, non-deterministic choice, empty process, and termination. The basic construct in the IOC language is the interaction, i.e., the construct $s \rightarrow r : a$ corresponds to an interaction between s and r on an operation a . The basic constructs in the POC language are output $!a$ and input $?a$. The semantics of the POC language is given in two versions: synchronous and asynchronous.

The relationship between an IOC and a POC differs depending on the communication models one considers. The authors identify four possible semantics that we describe below.

- The *synchronous semantics* demands that the POC behaves as specified by the IOC, following a synchronous semantics (i.e., corresponding send and receive occur at the same time). For instance, in a synchronous semantics the choreography (2.4) requires that both sending and receiving actions on a happen before the ones on b . Given a synchronous semantics, such an ordering may be enforced in (2.4) if we have

$$\{s_1, r_1\} \cap \{s_2, r_2\} \neq \emptyset$$

- The *sender semantics* demands that the sequentiality of sending actions are preserved. In this case, the choreography (2.4) requires that the sending on a happens

before the sending on b , while receiving actions (by r_1 and r_2) may happen in any order. This may be enforced by $s_2 = s_1$ or $s_2 = r_1$.

- The *receiver semantics* is the dual of the sender semantics, it requires that receiving actions are ordered. This may be enforced in (2.4) by $s_2 = r_1$ or $r_2 = r_1$.
- The *disjoint semantics* requires the executions of two sequential interactions to be totally disjoint. In the choreography (2.4), this means that the sending on b must happen *after* r_1 has received a . This may be enforced by $r_1 = s_2$.

For each of the semantics above, the authors formalise a notion of conformance, based on a bisimulation relation between POC and IOC. Then, for each notion of conformance, the authors give *connectedness* conditions which guarantee that a choreography and its projections are conformant.

Basically, the connectedness condition require that (i) there is a causality relation between sequential action in a IOC and (ii) that there is a unique point of choice.

- The causality relation requirement rules out IOC like

$$s_1 \rightarrow r_1 : a \quad | \quad s_2 \rightarrow r_2 : a$$

which specifies two un-ordered interactions on the *same* operation. The corresponding POC of this choreography are as follows:

$$s_1[!a] \quad | \quad r_2[?a] \quad | \quad s_2[!a] \quad | \quad r_2[?a]$$

which, according to the semantics of POC, would lead to a race condition, i.e., it may be the case that s_1 interacts with r_2 (instead of s_1).

- The unique point of choice condition requires that each choice is made by exactly one participant, as in other theories such as [29, 40].

For each variations of the semantics, the authors show that if a choreography satisfies the connectedness conditions, then it is conformant with its projections (with respect to the semantics being considered).

Based on this work, [47] introduces a few techniques to amend choreographies so that the connectedness conditions hold.

Realisability. Bultan et al. [9, 10] tackle the problem of determining whether a choreography is realisable. Essentially, a choreography is realisable if “it is possible to build a distributed system that communicates exactly as the choreography specifies”.

Choreographies in these works take the form of *conversation protocols*, that are finite state machines specifying the allowable sequence of interactions. A conversation protocol is akin to a generalised global type (cf. Section 2.1.4), except that parallel interactions cannot be explicitly represented, i.e., the \square -gate of the generalised global types has no counterpart in conversation protocols. Analogously, the counterpart of local types (or POC) is specified in the form of (asynchronous) communicating machines, similar to the ones in [40] (cf. Section 2.1.4). A subtle difference with the communicating machines we presented earlier is that each machine has now a unique buffer from which it can receive messages. This is different from the machines presented previously where there are two buffers for each pair of machines, one in each direction.

In [9], the authors give necessary and sufficient conditions for deciding choreography realisability. More precisely, a choreography is realisable if there exists a system of communicating machines which generates exactly the set of message sequences (i.e., traces) specified by the choreography. A choreography is realisable if and only if:

- the language accepted by the conversation protocol is equivalent to the 1-bounded system of its projections, i.e., the transition system obtained by projecting the choreography on each participant and executing the system with the additional constraint that there is at most one symbol in each buffer, and

- the 1-bounded system satisfies a temporal logic property stating that every message sent by any participant must eventually be consumed by a receiver.

A consequence of this result is that checking for realisability may be done efficiently, namely, using off-the-shelf model checking tools.

2.2.2 On Synthesising Choreographies

To be the best of our knowledge, little work has been done with respect to synthesising a global view of a distributed system from local specifications. In this section, we discuss works which attempt to obtain a *formal* global descriptions from local specifications.

Synthesis of global types

In the session type setting, the most advanced result concerns the synthesis of global types from communicating machines. We discuss briefly an older work which first mentioned the synthesis of global types from local types.

From local types. A *bottom-up* approach to build choreographies is briefly studied in [57]. This work relies on global and local types, but uses *local and global graphs*. A local graph is similar to a local type while a global graph is a disjoint union of family of local graphs. In fact, global graphs do not explicitly give the structure of the interactions that the local types specify.

From communicating machines. In [41], Deniérou and Yoshida propose a characterisation of communicating machines, from which it is possible to synthesise a global type, and a synthesis algorithm.

They use local types generated from the following grammar:

$$P ::= r?\{a_i; P_i\}_{i \in I} \quad | \quad s!\{a_i; P_i\}_{i \in I} \quad | \quad \mu\chi.P \quad | \quad \chi \quad | \quad \mathbf{0}$$

which may easily be translated to communicating machines (and vice versa). These communicating machines are called *basic* CFSMs and have the following properties

- they are deterministic, i.e., all the transitions leaving a same state have different labels,
- they have no mixed sates, i.e., all the transitions leaving a same state are either all receiving ($r?a$) or all sending ($r!a$), and
- they are directed, i.e., all the transitions leaving a same state are labelled by an action directed to the same machine.

The main contribution of this paper is the notion of *multiparty compatibility* which generalises the notion of duality between local types from [43] (cf. Section 2.1.1). Essentially, a system of CFSMs is multiparty compatible, if any action made by a machine can be “met” by the other machines, i.e., “the idea is to check the duality between each automaton and the rest, up to the internal communications (...) that the other machines will independently perform” [41]. Since CFSMs are generally undecidable [16], this check must be restricted to configurations of the system which are reachable by a 1-bounded execution, i.e., an execution of the system with the additional restriction that there is at most one symbol in each queue. Multiparty compatible systems enjoy the usual safety properties such as deadlock freedom, etc.

The algorithm to synthesise global types from multiparty compatible systems is quite simple and relies on a conversion from the transition system, corresponding to a synchronous execution of the CFSMs, to a global type. Remarkably, the global type so obtained is not unique and there is no special treatment of independent interactions, i.e., one may not synthesise a global type of the form $\mathcal{G} \mid \mathcal{G}'$. However, the projections of a synthesised global type are equivalent to the original system.

Other contributions of the authors in this paper include labelled transition system (LTS) semantics of both local types and global types, respectively.

Service composition

In [71, 72], Wang et al. give an algorithm to represent the composition of several services as a Petri net, such that an “optimal” representation is found. They consider three kinds of service models: input/output, precondition/effect and stateful models.

- In an *input/output model*, each operation of a service is represented by a pair of input and output data. For instance, a pair (I_a, O_a) says that, in order to execute operation a , the data I_a must have been previously generated (by other services or operations). After the execution of a , data O_a is generated.
- In a *precondition/effect model*, operations are defined as triples (p_a, E_a^+, E_a^-) . The element p_a is a set of propositions expressing the preconditions of the operation a : all the propositions must be true before executing operation a . The positive (resp. negative) effects E_a^+ (resp. E_a^-) specify the propositions which are true (resp. false) after the execution of a .
- In a *stateful model*, a service is represented as an automaton, whose transitions represent operations.

Remarkably, the authors show that both input/output and precondition/effect models may be easily represented as automata. In particular, they build an automaton per operation and per variable used in a service.

In order to compose services (i.e., automata representing services), the authors use the parallel product of automata. The parallel product of two automata is a new automaton which, essentially, covers all possible “private” transitions of each original automaton, and the common transition are synchronised. Once the services have been combined into a unique automaton, the authors use the theory of regions [3] to synthesise a Petri net, using slight variations of well-established algorithms from [37].

One of the results presented by the authors is that the Petri net obtained from such a transformation is optimal in the sense that it is “maximally concurrent”. However, in [71], no result shows that the Petri net obtained from the composition of services allows to recover the original services. In our terminology, it means that the authors do not guarantee that the global view of the services can be projected to a set of services equivalent to the original ones.

Inference of Message Sequence Charts.

In [1], Alur et al. focus on discovering whether Message Sequence Charts (MSC) [46, 54] imply unspecified scenarios. In fact, from a set of MSCs specifying what may seem to be correct scenarios, one may actually infer incorrect ones. This is due to the fact that, once implemented, each participant has only a local knowledge of the system. The authors assume that MSCs are implemented by concurrent automata, that are akin to the communicating machines we discussed earlier, but do not necessarily feature order-preserving communications.

The authors give precise conditions on MSCs for their implementation to be deadlock-free and realisable. Essentially, MSCs are realisable if no other MSC may be inferable from them, i.e., there is no other “implied” scenario. For the cases where MSCs are not realisable, the authors propose an algorithm to *synthesise* implied MSCs (i.e., unspecified scenarios which are in fact possible in the automata implementation).

Observe that the authors do not attempt to give an exhaustive global view of a distributed system, but instead focus on identifying “bad” executions. This is quite different from the approach we adopt in our work since we actually require that the projections of a synthesised choreography exhibits a behaviour that is equivalent to the original system.

2.2.3 Beyond Multiparty Session Types

In this section, we review a few remarkable extensions of the multiparty session types.

Dynamic session types. The original multiparty session types framework only caters for a fixed number of participants within a session. In [39], Deniérou and Yoshida introduced an extension of the framework to allow participants to join dynamically an already running session. To achieve this, the authors introduce a universal quantifier for *roles* in both global and local types; so that a (potentially unlimited) number of participants abiding by a common local type may join the session – provided that another participant exhibits a joining point.

Global programming. In [26], Carbone et al. introduce a novel methodology to distributed software design and implementation. They propose a language for global choreographies where instances of sessions are considered as first class element. In a sense this approach is a level of abstraction higher than global types – which describe only one session at a time. This permits to design and implement distributed systems from a global point of view. These “meta-choreographies” may then be projected onto end-point code. When some conditions holds for these global choreographies, the authors show that the implementation of the system, obtained by projection, is deadlock-free.

Global progress. One shortcoming of the original works on session types is that progress was only guaranteed within a single session. However, in realistic scenarios, a program may be involved in several sessions, possibly with different participants. In [11, 35, 36], the authors tackle the problem of *global progress*, i.e., ensuring that programs involved in several (interleaved) sessions do not deadlock. In Example 2.11 (adapted from [35]), we illustrate the paradigmatic example of a deadlock due to interleaved sessions.

Example 2.11. Consider the following processes:

$$\begin{aligned} R_s &= n_{1[s]}(k_1); n_{2[s]}(k_2); k_1 ?r(a); k_2 !r\langle b \rangle \\ R_r &= \bar{n}_{1[r]}(k_1); \bar{n}_{2[r]}(k_2); k_2 ?s(b); k_1 !s\langle a \rangle \end{aligned}$$

where each process initiates two sessions on n_1 and n_2 .

The parallel composition of these processes, $R_s \mid R_r$, deadlocks when it reduces to

$$(\forall k_1) (\forall k_2) (k_{1[s]} ?r(a); k_{2[s]} !r\langle b \rangle \mid k_{2[r]} ?s(b); k_{1[r]} !s\langle a \rangle)$$

where each process waits for a message from the other.

Notice that the program is well-typed in the sense of [44]. Indeed, assuming that the global type associated to n_1 is $r \rightarrow s : a$ and that the global type associated to n_2 is $s \rightarrow r : b$, both processes use each channel as prescribed by the projections of its corresponding global type. \diamond

Intuitively, a process has the global progress property if every receive action will eventually be matched by a corresponding message and every sent message is eventually consumed. However, an important aspect considered in these works is that a process should not be considered stuck if the only reason why it cannot reduce further is because some participants are missing in order to initiate a new session. Indeed, in this case, the process will be able to proceed as soon as other participants join the system. In other words, global progress requires that, once a session has started, all the interactions that are supposed to occur in this session will eventually happen.

In order to guarantee global progress, the authors rely on a slight variation on the type system of the multiparty session types [44], which guarantees progress within a session. In addition, they introduce an *interaction type system* which, roughly, ensures that there is no circular dependencies between channel (or service) names.

2.2.4 Other Approaches to Multiparty Sessions

Many different approaches tackle the problem of designing and verifying formal models of distributed systems, which, to some extent, distance themselves from the framework based on global types. In this section, we discuss three approaches that are of particular interest in the scope of this thesis.

Compliance. In [31], Castagna et al. propose a *contract* language for processes. The language is akin to the π -calculus, allowing channel names to be passed around, and the semantics of contracts is synchronous.

In order to characterise when a set of contracts has progress, they adopt a testing approach [15] and define the notion of *compliance*. Essentially, a contract S is compliant with a contract S' , if for any computation of their composition where S cannot progress by itself, either (i) the residual of S has terminated and S' will eventually terminate; or (ii) the two residuals (of S and S') can eventually synchronise. The notion of compliance is extended to the one of *well-formedness* of a system of contracts; which requires that each contract is compliant with rest of the system. Compliance is in fact similar to the notion of multiparty compatibility of [41].

Interestingly, the authors apply their approach to a calculus similar to the one of [11, 43]. They propose a type system for this calculus which extracts contracts from processes. Using the notion of well-formedness, they are able to give a progress result, stating that if a process is typable (i.e., by a system of contracts) and its type is “well-formed”, then the process is able to reduce further or has terminated.

In the same line of work, [63] proposes to “project” processes into session types (akin to the contracts of [31] but without name passing capabilities). The author gives an adapted notion of compliance, now called *completeness* which permits to characterise when a set of session types is able to reduce further. Similarly to [31], session types may be extracted from processes (based on a variation of the π -calculus). If the types assigned

to a process are “complete”, then progress and deadlock properties are guaranteed.

In addition, to ensure that channel delegation does not cause deadlock, the author defines a strict partial ordering on channels which is used to check whether a channel may be sent over another one. Interestingly, this allows channels to be used in a non-linear way, i.e, a process may send a channel to another process and still use this channel later.

Conversation types. Caires et al. [21, 23] propose a type-based framework to analyse and verify distributed systems. Their approach is somewhat similar to that of session types, in the sense that they propose a calculus (inspired from the π -calculus) and a type discipline for it. If a program is typable then it never violates the prescribed protocol and, under some condition, it is guaranteed to have progress. The framework allows for dynamical join and leave of participants, and *nested* conversations.

The framework is based on the conversation calculus [22, 69]. The main difference with respect to the approach of Honda et al. [44] is that interactions are passing through a *conversation*, i.e., a shared medium akin to a shared channel between several participants. A conversation may be dynamically created via a shared name – similarly to a session channels – and allows participants to dynamical join and leave a conversations by communicating the name of the conversation. A conversation may be typed by a *conversation type*. These behavioural types take into account both internal and external behaviours of participants in a conversation.

Contract oriented computing. In [5–7], Bartoletti et al. propose a novel approach, called contract-oriented computing, where the interactions between distributed entities are governed by contracts.

The life-cycle of contract-oriented computing consists of three phases. In the first phase, the entities’ contracts are used to find an agreement where each participant’s needs

and/or offers are matched by other contracts. In the second phase, the contracts become “binding”, i.e., each entity will be deemed *culpable* if it cannot meet its obligations. In the third phase, the entities execute their tasks while monitoring their contracts.

The authors defined a new calculus, CO_2 (for COntract-Oriented computing) which may be parametrised with respect to a contract model. The calculus features basic primitives that permit to instantiate the contract-oriented paradigm, such as:

- a *tell* primitive to advertise a contract to a broker,
- a *fuse* primitive allowing a broker to create a new session between a set of participants who advertised compliant contracts, and
- a *do* primitive, used by participants to fulfil an action specified in their contracts.

An important feature of CO_2 is that it allows to identify participants who do not meet their requirement – so that, e.g., they may face legal consequences. The authors also define the notion of *honesty*, i.e., the ability of a participant to always be able to exculpate itself. In other words, an honest participant will always fulfil the contracts it advertised.

A negative result in [7] is that it is not decidable whether or no a given process is honest. In [7], CO_2 is instantiated to a theory of bilateral contracts inspired by [30]. A decidable approximation of honesty is introduced in [4], relying on the product between a finite state system (approximating contracts) and a Basic Parallel Process (approximating a CO_2 process).

Remark 2.6. *In Chapter 5, we define formally a slight variation of CO_2 where contracts are local session types and they are compliant when it is possible to synthesise a choreography from them. We also revisit the notion of honesty to our needs.*

Synthesising Choreographies from Local Session Types

We introduce a type system which allows, under some conditions, to synthesise a choreography (i.e., a global type) from a set of local session types which describe end-point behaviours (i.e., local types). We show, notably, that the projections of a synthesised global type is equivalent to the original set of local types and that, if a set of local types is typable, then it enjoys safety and liveness properties.

3.1 Introduction

The multiparty session types framework proposes a methodology that (i) allows to design a global view of the interactions – aka *global type* –, (ii) provides an effective analysis of such a global view, (iii) automatically projects the global view to local end-points – aka *local types* –, and (iv) type checks end-points' code against local types. The theory guarantees that, when the global view enjoys suitable properties (phase (ii)), the end-points typable with local types enjoy, e.g., liveness properties such as progress.

A drawback of such an approach is that it cannot be applied when the local types describing the communication patterns of end-points are not obtained by an a priori designed global view. For instance, in service-oriented computing, one typically has independently

developed end-points that have to be combined to form larger services. Hence, deciding if the combined service respects its specification becomes non trivial. To illustrate this, we introduce a simple example used throughout this chapter.

Consider the system

$$S_{BS} = b_1[P_1] \mid s_1[S_1] \mid b_2[P_2] \mid s_2[S_2]$$

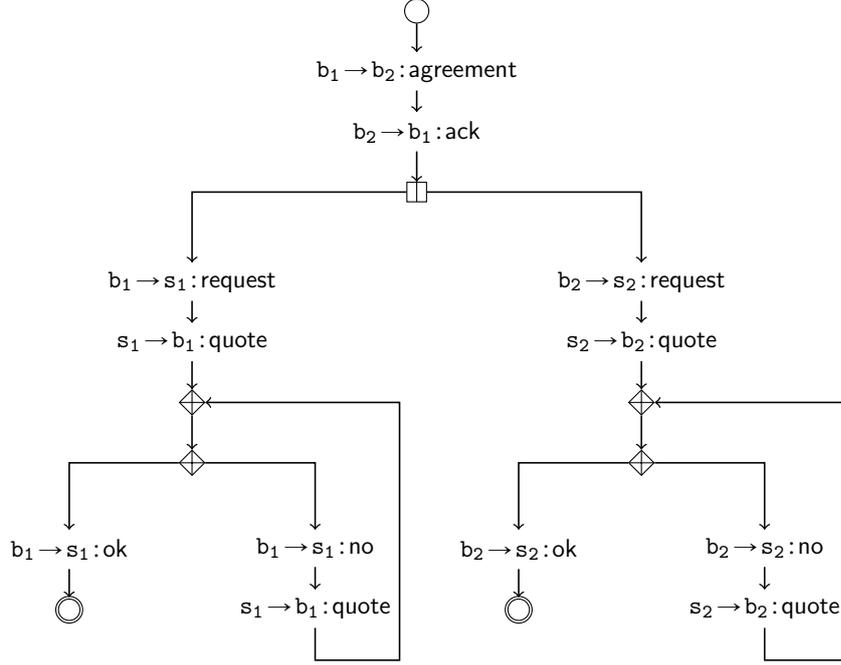
consisting of two buyers (b_1 and b_2) and two sellers (s_1 and s_2) running in parallel, and whose behaviours are specified as follows.

$$\begin{aligned} P_1 &= b_2!agreement; b_2?ack; Q_1 \\ P_2 &= b_1?agreement; b_1!ack; Q_2 \\ Q_i &= s_i!request; s_i?quote; \mu\chi.(s_i!ok \oplus s_i!no; s_i?quote; \chi) \quad i \in \{1, 2\} \\ S_i &= b_i?request; b_i!quote; \mu\chi'.(b_i?ok + b_i?no; b_i!quote; \chi') \quad i \in \{1, 2\} \end{aligned}$$

where \oplus represents an internal choice and $+$ a choice made by the environment. Intuitively, the overall behaviour of S_{BS} should be that two buyers (b_1 and b_2) agree on purchasing concurrently two items from their respective sellers and try to obtain the best price for these items. A natural question arises: is this intended behaviour actually realised by S_{BS} ? Arguably, it is not immediate to answer such a question by considering only S_{BS} , even more so if a system involves a large number of participants, with possibly complex behaviours.

We propose to construct a global view of distributed end-points like S_{BS} via a slight variation of the multiparty session types introduced in [44]. Such types formalise a global view of the behaviour which, for S_{BS} , resembles the diagram in Figure 3.1, where the choreography of the overall protocol becomes much clearer.

An advantage of our approach is that we can reuse the results of the theory of multiparty session types to prove properties of end-points, e.g., safety and progress. In fact,


 Figure 3.1: Global view of S_{BS}

we show that when the choreography can be constructed, its projections correspond to the initial end-points. Therefore, the well-formedness of the synthesised global choreography guarantees progress and safety properties of end-points.

We assume that session types are extracted from programs (relying on, e.g., [31, 43, 44]), and that they are readily available before addressing the construction of a global type.

3.2 Local Types

We use CCS-like processes (with guarded external and internal choices) to infer a global type from *local types* that correspond to the participants in the inferred choreography. Hereafter, \mathbb{P} is a denumerable set of *participant names* (ranged over by s, r, n, \dots) and $\mathbb{C} = (\mathbb{P} \times \mathbb{P}) \setminus \{(n, n) \mid n \in \mathbb{P}\}$ is a denumerable set of *channel identifiers*, we write sr when

$(s, r) \in \mathbb{C}$.

Syntax. The syntax of local types below is parametrised with respect to basic data types such as `bool`, `int`, ... (ranged over by `a`, `b`, `c`, ...):

$$\begin{array}{lcl}
 S, S' & ::= & n[P] \quad | \quad sr : \rho \quad | \quad S | S' \quad | \quad \mathbf{0} \\
 P, Q & ::= & \bigoplus_{i \in I} r_i ! a_i ; P_i \quad | \quad \sum_{i \in I} s ? a_i ; P_i \quad | \quad \mu \chi . P \quad | \quad \chi \\
 \rho & ::= & \varepsilon \quad | \quad a \cdot \rho
 \end{array}$$

A system S consists of the parallel composition of *processes* and *queues*. A *process* $n[P]$ is a behaviour P identified by $n \in \mathbb{P}$; we assume that, given a system S , the participant names are all different in S . A behaviour is either an internal choice, an external choice, or a recursive process. An internal choice $\bigoplus_{i \in I} r_i ! a_i ; P_i$ is guarded by output prefixes $r_i ! a_i$ representing the sending of a value of sort a_i on the queue sr_i , assuming the behaviour is the one of participant s . We assume that

$$\forall i, j \in I : i \neq j \implies (r_i, a_i) \neq (r_j, a_j)$$

An external choice $\sum_{i \in I} s ? a_i ; P_i$ is guarded by input prefixes $s ? a_i$ representing the reception of a value of sort a_i from the queue sr , assuming the behaviour is the one of participant r . We assume that

$$\forall i, j \in I : i \neq j \implies a_i \neq a_j$$

The asymmetry between internal and external choice, i.e., the fact that all branches of an external choice must be prefixed by an input from a same sender s allows us to guarantee that a choice is always made by exactly one participant (cf. Section 3.6).

We overload $\mathbf{0}$ to denote either an internal or external choice where $I = \emptyset$, i.e.,

$$\mathbf{0} \stackrel{def}{=} \bigoplus_{i \in \emptyset} r_i ! a_i ; P_i = \sum_{i \in \emptyset} s ? a_i ; P_i$$

In $\mu\chi.P$, $\mu\chi$ is a binder for the free occurrences of χ in P . Moreover, all such free occurrences are prefix-guarded in P . We consider closed behaviours only (that is, behaviours with no free occurrences of recursion variables). Also, for simplicity – and without loss of generality since bound variables are α -convertible – we assume that bound variables are pairwise distinct.

A *program* is a system with no queues, while a *runtime system* is a system S having exactly one queue $sr : \rho$ per pair of participants s and r in S . We write $sr : \varepsilon$ whenever the queue from s to r is empty. Given a program S , we write $Q(S)$ for the parallel composition of the empty queues connecting all the participants in S . In the following, S, S', \dots denote either a program or runtime system.

Semantics. The semantics of systems is reminiscent of CCS with asynchronous and order-preserving communications. The semantics of systems is a labelled transition system (LTS) with labels:

$$\begin{aligned} \alpha &::= r!a \mid s?a \\ \lambda &::= \alpha \mid sr \cdot a \mid a \cdot sr \mid n[\alpha] \mid s \rightarrow r : a \mid r \leftarrow s : a \end{aligned}$$

Label α indicates either sending or reception by a process. Labels $sr \cdot a$ and $a \cdot sr$ respectively indicate push and pop operations on queues. Label $n[\alpha]$ indicates a communication action done by participant n . The next two labels indicate synchronisations between a participant and a queue: label $s \rightarrow r : a$ indicates that s puts a datum on a queue sr , while label $r \leftarrow s : a$ indicates that r retrieves a datum from the queue sr .

The semantics is given up to the congruence rules below.

$$\begin{aligned} \mu\chi.P &\equiv P \{\mu\chi.P/\chi\} \\ S | S' &\equiv S' | S \quad S | \mathbf{0} \equiv S \quad S' | (S' | S'') \equiv (S' | S') | S'' \end{aligned}$$

Note that commutativity and associativity of internal and external choice follow by construction (i.e., by the set notation used in the sums). We often write $P \oplus P'$ (resp. $P + P$) for internal (resp. external) choice. We assume that these operators are commutative and associative, and that $_ ; _$ has a higher precedence than $_ \oplus _$ and $_ + _$.

The LTS $\xrightarrow{\lambda}$ is the smallest relation closed under the following rules:

$$\begin{array}{c} \text{[EXT]} \sum_{i \in I} s^?a_i ; P_i \xrightarrow{s^?a_j} P_j \quad j \in I \quad \text{[INT]} \bigoplus_{i \in I} r_i !a_i ; P_i \xrightarrow{r_j !a_j} P_j \quad j \in I \\ \\ \text{[BOX]} \frac{P \xrightarrow{\alpha} P'}{n[P] \xrightarrow{n[\alpha]} n[P']} \quad \text{[EQ-P]} \frac{P \equiv Q \xrightarrow{\alpha} Q' \equiv P'}{P \xrightarrow{\alpha} P'} \\ \\ \text{[POP]} sr : a \cdot \rho \xrightarrow{a \cdot sr} sr : \rho \quad \text{[PUSH]} sr : \rho \xrightarrow{sr \cdot a} sr : \rho \cdot a \\ \\ \text{[IN]} \frac{S \xrightarrow{r[s^?a]} S_1 \quad S' \xrightarrow{a \cdot sr} S_2}{S | S' \xrightarrow{r \leftarrow s : a} S_1 | S_2} \quad \text{[OUT]} \frac{S \xrightarrow{s[r!a]} S_1 \quad S' \xrightarrow{sr \cdot a} S_2}{S | S' \xrightarrow{s \rightarrow r : a} S_1 | S_2} \\ \\ \text{[EQ-S]} \frac{S \equiv S_1 \xrightarrow{\lambda} S_2 \equiv S'}{S \xrightarrow{\lambda} S'} \quad \text{[PAR]} \frac{S \xrightarrow{\lambda} S'}{S | S'' \xrightarrow{\lambda} S' | S''} \end{array}$$

Rules [EXT] and [INT] are trivial. By [POP] , a queue sends the first datum, if any. By [PUSH] , a queue receives a new datum. Processes can synchronise with queues according to rules [IN] and [OUT] . The remaining rules are standard. Let $S \rightarrow$ if and only if there are S' and λ such that $S \xrightarrow{\lambda} S'$ and $\xrightarrow{\lambda_1 \dots \lambda_n}$ (resp. \Longrightarrow) be the reflexive transitive closure of $\xrightarrow{\lambda}$ (resp. \rightarrow).

We adopt the following notation.¹

$$S(\mathfrak{n}) = \begin{cases} P & \text{if } S \equiv \mathfrak{n}[P] \mid S' \\ \perp & \text{otherwise} \end{cases} \quad S[\mathfrak{sr}] = \begin{cases} \rho & \text{if } S \equiv \mathfrak{sr} : \rho \mid S' \\ \perp & \text{otherwise} \end{cases}$$

Given a system S and a participant name \mathfrak{n} , $S(\mathfrak{n})$ returns the behaviour of \mathfrak{n} in S , if \mathfrak{n} is in S ; and is undefined otherwise. Similarly, $S[\mathfrak{sr}]$ returns the content of the queue \mathfrak{sr} , if the queue is part of S .

Example 3.1. Consider S_{BS} from the previous section, we have, e.g., $S_{\text{BS}}(\mathfrak{b}_1) = P_1$, while $S_{\text{BS}}[\mathfrak{b}_1\mathfrak{b}_2] = \perp$ since S_{BS} is a program, i.e., it does not contain queues. \diamond

3.3 Global Types

The global types we consider in this chapter are borrowed from [44] with minor variations. A global type is a term derivable from the following grammar:

$$\mathcal{G} ::= \sum_{i \in I} \mathfrak{s} \rightarrow \mathfrak{r}_i : \mathfrak{a}_i ; \mathcal{G}_i \quad | \quad \mathcal{G} \mid \mathcal{G}' \quad | \quad \mu \chi . \mathcal{G} \quad | \quad \chi \quad | \quad \mathbf{0}$$

An expression of the form $\mathfrak{s} \rightarrow \mathfrak{r} : \mathfrak{a}$ represents an interaction where $\mathfrak{s} \in \mathbb{P}$ sends a value of sort \mathfrak{a} to $\mathfrak{r} \in \mathbb{P}$ (we assume that $\mathfrak{s} \neq \mathfrak{r}$). The choice production $\sum_{i \in I} \mathfrak{s} \rightarrow \mathfrak{r}_i : \mathfrak{a}_i ; \mathcal{G}_i$ is guarded by prefixes $\mathfrak{s} \rightarrow \mathfrak{r}_i : \mathfrak{a}_i$. The production indicates a (exclusive) choice of interactions, if participant \mathfrak{s} sends a message \mathfrak{a}_i to \mathfrak{r}_i , then interactions \mathcal{G}_i take place; we assume that

$$\forall i, j \in I : i \neq j \implies (\mathfrak{r}_i, \mathfrak{a}_i) \neq (\mathfrak{r}_j, \mathfrak{a}_j)$$

We often write $\mathcal{G} + \mathcal{G}'$ for the binary version of the choice production, we assume that $_{+}$ is commutative and associative, and that the operator $_{;}$ has a higher precedence

¹We write $f(x) = \perp$ when the function f is undefined on x .

than $_{-}+_{-}$. Concurrent interactions are written $\mathcal{G} \mid \mathcal{G}'$ and indicate sets of independent interactions. In a recursive global type $\mu\chi.\mathcal{G}$, χ is bound and guarded in \mathcal{G} . We assume that global types are closed and often omit trailing occurrences of $\mathbf{0}$.

Remark 3.1. *The main difference with respect to the global types in [44] is that we merge the branching constructs with the interaction prefix, similarly to [28], and we do not require a choice construct to involve only two participants, e.g., we accept the following global type:*

$$s \rightarrow r_1 : a; s \rightarrow r_2 : a \quad + \quad s \rightarrow r_2 : b; s \rightarrow r_1 : b$$

Global types are taken up-to structural congruence, defined as the smallest equivalence relation satisfying the monoidal axioms for $_{-} \mid _{-}$ and $_{-} + _{-}$ (with $\mathbf{0}$ as the identity element); and the axiom

$$\mu\chi.\mathcal{G} \equiv \mathcal{G} \{ \mu\chi.\mathcal{G} / \chi \}$$

Example 3.2. The first two interactions between b_1 and b_2 in the example of Section 3.1 are modelled by the global type \mathcal{G}_1 :

$$\mathcal{G}_1 = b_1 \rightarrow b_2 : \text{agreement}; b_2 \rightarrow b_1 : \text{ack}$$

The type \mathcal{G}_1 says that participant b_1 sends a message of sort agreement to b_2 , then b_2 replies with a message of sort ack.

Each concurrent branch of the interactions in Figure 3.1 are specified by the following global type (where $i \in \{1, 2\}$)

$$\begin{aligned} \mathcal{G}_2^i &= b_i \rightarrow s_i : \text{request}; s_i \rightarrow b_i : \text{quote}; \\ &\mu\chi.(b_i \rightarrow s_i : \text{ok} + b_i \rightarrow s_i : \text{no}; s_i \rightarrow b_i : \text{quote}; \chi) \end{aligned}$$

The global type \mathcal{G}_2^i says that b_i sends a request to s_i , who replies with a message of type quote. Then, the two participants enter a recursion which is guarded by a choice made by

b_i . If b_i is satisfied with the quote, the interactions terminate here; otherwise s_i sends another quote and the interactions are repeated.

Finally, the global type representing all the interactions in Figure 3.1 is given below.

$$\mathcal{G}_{SBS} = b_1 \rightarrow b_2 : \text{agreement}; b_2 \rightarrow b_1 : \text{ack}; (\mathcal{G}_2^1 \mid \mathcal{G}_2^2)$$

◇

The syntax of global types may specify behaviours that are not implementable. The rest of this section borrows from [28] and [44] and adapts the requirements a global type must fulfil to ensure that the set of interactions it prescribes is indeed feasible.

3.3.1 Well-formed Global Types

We define the conditions for a global type to be well-formed. We write $\mathcal{P}(\mathcal{G})$ (resp. $\mathcal{C}(\mathcal{G})$) for the set of participant (resp. channel) names in \mathcal{G} , and $\text{fv}(\mathcal{G})$ (resp. $\text{bv}(\mathcal{G})$) for the set of free (resp. bound) variables in \mathcal{G} , similarly for a system S .

We give a few auxiliary functions. The *ready set* of a global type \mathcal{G} , written $\text{R}(\mathcal{G})$, is defined as follows.

$$\text{R}(\mathcal{G}) \stackrel{\text{def}}{=} \{s \rightarrow r : a \mid \mathcal{G} \equiv (s \rightarrow r : a; \mathcal{G}_1 + \mathcal{G}_2) \mid \mathcal{G}_3\}$$

Next we define a functions that computes the sets of participant which interact inde-

pendently “in the last part of a global type”. We define $\text{Indep}(\mathcal{G}) \stackrel{\text{def}}{=} \text{I}_0(\emptyset, \mathcal{G})$, where

$$\text{I}_0(\mathcal{P}, \mathcal{G}) \stackrel{\text{def}}{=} \begin{cases} \text{I}_0(\{\mathbf{s}, \mathbf{r}\} \cup \mathcal{P}, \mathcal{G}_1), & \text{if } \mathcal{G} = \mathbf{s} \rightarrow \mathbf{r} : \mathbf{a}; \mathcal{G}_1 \\ \text{I}_0(\emptyset, \mathcal{G}_1) \cup \text{I}_0(\emptyset, \mathcal{G}_2), & \text{if } \mathcal{G} = \mathcal{G}_1 \mid \mathcal{G}_2 \\ \text{I}_0(\mathcal{P}, \mathcal{G}_1), & \text{if } \mathcal{G} = \mathcal{G}_1 + \mathcal{G}_2 \text{ and } \text{I}_0(\mathcal{P}, \mathcal{G}_1) = \text{I}_0(\mathcal{P}, \mathcal{G}_2) \\ \text{I}_0(\mathcal{P}, \mathcal{G}_1), & \text{if } \mathcal{G} = \mu\chi.\mathcal{G}_1 \\ \{\mathcal{P}\}, & \text{if } \mathcal{G} = \mathbf{0} \text{ or } \mathcal{G} = \chi \\ \perp, & \text{otherwise} \end{cases}$$

$\text{I}_0(\mathcal{P}, \mathcal{G})$ is the family of sets of participants of \mathcal{G} , so that for all $N \neq M \in \text{I}_0(\mathcal{P}, \mathcal{G})$, the participants in N and those in M are in different concurrent branches “in the last part of \mathcal{G} ”. Note that $\text{I}_0(-, -)$ is a partial function.

Example 3.3. Continuing with the global types from Example 3.2, we have:

$$\text{R}(\mathcal{G}_{\text{SBS}}) = \text{R}(\mathcal{G}_1) = \{\mathbf{b}_1 \rightarrow \mathbf{b}_2 : \text{agreement}\} \quad \text{and} \quad \text{R}(\mathcal{G}_2^1) = \{\mathbf{b}_1 \rightarrow \mathbf{s}_1 : \text{request}\}$$

$$\text{Indep}(\mathcal{G}_{\text{SBS}}) = \{\{\mathbf{b}_1, \mathbf{s}_1\}, \{\mathbf{b}_2, \mathbf{s}_2\}\}, \quad \text{Indep}(\mathcal{G}_1) = \{\{\mathbf{b}_1, \mathbf{b}_2\}\}, \text{ and}$$

$$\text{Indep}(\mathcal{G}_2^1) = \{\{\mathbf{b}_1, \mathbf{s}_1\}\}$$

◇

We say that \mathcal{G} is *well-formed* if $\vdash \mathcal{G}$ can be derived from the rules given in Figure 3.2. We assume that each premise of the rules in Figure 3.2 does not hold whenever any of the auxiliary function used is undefined (e.g., in $[\text{WF-}\mu\chi]$, if $\text{Indep}(\mathcal{G}) = \perp$ then $\vdash \mu\chi.\mathcal{G}$ is not derivable). In the following, we discuss the rules of Figure 3.2, which are grouped according to three requirements: sequentiality, single threadness, and knowledge of choice.

Sequentiality [28]. Rule $[\text{WF-};]$ ensures that sequentiality is preserved. In $[\text{WF-};]$, the ordering dependency between a prefix and its continuation allows us to implement each

$$\begin{array}{c}
 \text{[WF-;]} \frac{\forall s' \rightarrow r' : _ \in R(\mathcal{G}) : \{s', r'\} \cap \{s, r\} \neq \emptyset \quad \vdash \mathcal{G}}{\vdash s \rightarrow r : a ; \mathcal{G}} \\
 \\
 \text{[WF-|]} \frac{\mathcal{P}(\mathcal{G}) \cap \mathcal{P}(\mathcal{G}') = \emptyset \quad \vdash \mathcal{G} \quad \vdash \mathcal{G}'}{\vdash \mathcal{G} \mid \mathcal{G}'} \\
 \\
 \text{[WF-}\mu\chi\text{]} \frac{\chi \in \text{fv}(\mathcal{G}) \Rightarrow |\text{Indep}(\mathcal{G})| = 1 \quad \vdash \mathcal{G}}{\vdash \mu\chi. \mathcal{G}} \\
 \\
 \text{[WF-}\chi\text{]} \frac{}{\vdash \chi} \qquad \text{[WF-}\mathbf{0}\text{]} \frac{}{\vdash \mathbf{0}} \\
 \\
 \text{[WF-+]} \frac{\forall s \rightarrow r : a \in R(\mathcal{G}) : \forall s' \rightarrow r' : a' \in R(\mathcal{G}') : s = s' \wedge (r, a) \neq (r', a') \quad \vdash \mathcal{G} \quad \vdash \mathcal{G}'}{\vdash \mathcal{G} + \mathcal{G}'}
 \end{array}$$

Figure 3.2: Rules for well-formedness

participant so that at least one action of the first prefix always happens before an action of the second prefix. More concretely, this rules out, e.g.,

$$s_1 \rightarrow r_1 : a ; s_2 \rightarrow r_2 : b \quad \mathbf{x}$$

where, evidently, it is not possible to guarantee that s_2 sends b after r_1 receives a from s_1 . Since we are working in an asynchronous setting, we do not want to force both send and receive actions of the first prefix to happen before both actions of the second one.

Single threadness [44]. A participant should not appear in different concurrent branches of a global type, so that each participant is single threaded. This is also reflected in the calculus of Section 3.2, where parallel composition is only allowed at the system level. Therefore, in $[\text{WF-}|]$, the participant names in concurrent branches must be disjoint. Rule $[\text{WF-}\mu\chi]$ requires that \mathcal{G} is single threaded, i.e., concurrent branches cannot appear under recursion. If that was the case, a participant could appear in different concurrent branches of the unfolding of a recursive global type. For instance, the one-time unfolding of \mathcal{G}

below is \mathcal{G}' .

$$\mathcal{G} = \mu\chi.(\mathfrak{s}_1 \rightarrow \mathfrak{r}_1 : \mathfrak{a}; \chi \mid \mathfrak{s}_2 \rightarrow \mathfrak{r}_2 : \mathfrak{b}; \chi) \quad \mathcal{G}' = (\mathfrak{s}_1 \rightarrow \mathfrak{r}_1 : \mathfrak{a}; \mathcal{G}) \mid (\mathfrak{s}_2 \rightarrow \mathfrak{r}_2 : \mathfrak{b}; \mathcal{G}) \quad \times$$

Observe that the sets of participants in the concurrent branches of \mathcal{G}' are not disjoint.

Knowledge of choice [28, 44]. Whenever a global type specifies a choice of two sets of interactions, the decision should be made by exactly one participant. For instance,

$$\mathfrak{s}_1 \rightarrow \mathfrak{r}_1 : \mathfrak{a}; \mathcal{G}_1 \quad + \quad \mathfrak{s}_2 \rightarrow \mathfrak{r}_2 : \mathfrak{b}; \mathcal{G}_2 \quad \times$$

specifies a choice made by \mathfrak{s}_1 in the first branch and by \mathfrak{s}_2 in the second one; this kind of choreographies cannot be implemented (without using hidden interactions). Also, we want to avoid global types where a participant n behaves differently in choice branches without being aware of the choice made by others. For instance, in

$$\mathfrak{s} \rightarrow \mathfrak{r} : \mathfrak{a}; \mathfrak{n} \rightarrow \mathfrak{r} : \mathfrak{c}; \mathcal{G}_1 \quad + \quad \mathfrak{s} \rightarrow \mathfrak{r} : \mathfrak{b}; \mathfrak{n} \rightarrow \mathfrak{r} : \mathfrak{d}; \mathcal{G}_2 \quad \times$$

where n ignores the choice of s and behaves differently in each branch. On the other hand, we want global types of the following form to be accepted.

$$\begin{array}{l} \mathfrak{s} \rightarrow \mathfrak{r} : \mathfrak{a}; \mathfrak{n} \rightarrow \mathfrak{s} : \mathfrak{b}; \mathfrak{s} \rightarrow \mathfrak{n} : \mathfrak{a}; \mathfrak{n} \rightarrow \mathfrak{r} : \mathfrak{a} \\ + \\ \mathfrak{s} \rightarrow \mathfrak{r} : \mathfrak{b}; \mathfrak{n} \rightarrow \mathfrak{s} : \mathfrak{b}; \mathfrak{s} \rightarrow \mathfrak{n} : \mathfrak{c}; \mathfrak{n} \rightarrow \mathfrak{r} : \mathfrak{d} \end{array} \quad \checkmark$$

Indeed, in this case n behaves differently in each branch, but only *after* “being informed” by s about the chosen branch, i.e., n always sends b , then if it receives a , it sends a again; if it receives c , then it sends d .

Together with the definitions of the projections of a global type (cf. Definition 3.1

below) rule $_{[WF+]}$ guarantees that “knowledge of choice” is respected. In particular, the rule requires that the participant who makes the decision is the same in every branch of a choice, while the prefixes guarding the choice must be distinct.

Definition 3.1 ($_{-}\downarrow_{-}$). *The projection of a global type \mathcal{G} with respect to a participant n – written $\mathcal{G} \downarrow_n$ – is defined by the (partial) function below.*

$$\mathcal{G} \downarrow_n \stackrel{\text{def}}{=} \begin{cases} s?a; \mathcal{G}' \downarrow_n, & \text{if } \mathcal{G} = s \rightarrow n : a; \mathcal{G}' \\ r!a; \mathcal{G}' \downarrow_n, & \text{if } \mathcal{G} = n \rightarrow r : a; \mathcal{G}' \\ \mathcal{G}' \downarrow_n, & \text{if } \mathcal{G} = s \rightarrow r : a; \mathcal{G}' \text{ and } s \neq n \neq r \\ \mathcal{G}_1 \downarrow_n \oplus \mathcal{G}_2 \downarrow_n, & \text{if } \mathcal{G} = \mathcal{G}_1 + \mathcal{G}_2 \text{ and } \exists n \rightarrow r : _ \in R(\mathcal{G}) \\ \mathcal{G}_1 \downarrow_n \uplus \mathcal{G}_2 \downarrow_n, & \text{if } \mathcal{G} = \mathcal{G}_1 + \mathcal{G}_2 \text{ and } \forall s \rightarrow r : _ \in R(\mathcal{G}) : n \neq s \\ \mathcal{G}_i \downarrow_n, & \text{if } \mathcal{G} = \mathcal{G}_1 | \mathcal{G}_2 \text{ and } n \notin \mathcal{P}(\mathcal{G}_j), i \neq j \in \{1, 2\} \\ \mu\chi. \mathcal{G}' \downarrow_n, & \text{if } \mathcal{G} = \mu\chi. \mathcal{G}' \text{ and } \mathcal{G}' \downarrow_n \neq \chi \\ \mathbf{0}, & \text{if } \mathcal{G} = \mu\chi. \mathcal{G}' \text{ and } \mathcal{G}' \downarrow_n = \chi \\ \mathcal{G}, & \text{if } \mathcal{G} = \chi \text{ or } \mathcal{G} = \mathbf{0} \\ \perp, & \text{otherwise} \end{cases}$$

We say that a global type is projectable if $\mathcal{G} \downarrow_n$ is defined for all $n \in \mathcal{P}(\mathcal{G})$.

The projection map is similar to the one given in [44], but for the use of $_{-}\uplus_{-}$ to project choice branches (see Definition 3.2 below) – a similar operation is also used in [39].

The first three cases deal with prefixes and are trivial. The fourth case projects choice branches when the considered participant is the sender. Note that if \mathcal{G} is well-formed then there is always a single sender participant in choice branches. The fifth case deals with branches when the considered participant is not the sender, in which cases the branches must be merged (see Definition 3.2 below). Concurrent branches may be projected only if

their sets of participants are disjoint. There are two cases to project recursive global types. The first case is used to project the behaviour of participants who are indeed involved in the recursion, while the second case deals with participants who do not appear after the recursion definition. The other cases are trivial.

Note that a global type may be projected even if it is not well-formed, but in that case none of the properties given below are guaranteed to hold.

Essentially, the function $_ \uplus _$ merges (if possible) the behaviour of a participant in different choice branches, cf. Example 3.4 below.

Definition 3.2 ($_ \uplus _$).

$$P \uplus Q = \begin{cases} P + Q, & \text{if } P = \sum_{i \in I} s?a_i; P_i \text{ and } Q = \sum_{j \in J} s?a_j; Q_j \\ & \text{and } \forall i \in I: \forall j \in J: a_i \neq a_j \text{ and } I, J \neq \emptyset \\ \sum_{i \in I} s?a_i; P_i \uplus P'_i, & \text{if } P = \sum_{i \in I} s?a_i; P_i \text{ and } Q = \sum_{i \in I} s?a_i; P'_i \\ \bigoplus_{i \in I} r_i!a_i; P_i \uplus P'_i, & \text{if } P = \bigoplus_{i \in I} r_i!a_i; P_i \text{ and } Q = \bigoplus_{i \in I} r_i!a_i; P'_i \\ P, & \text{if } P \equiv Q \\ \perp, & \text{otherwise} \end{cases}$$

In the first case, the function merges two guarded external choices, if the message sorts in their prefix are disjoint. This ensures that each participant knows which branch of the global type was chosen. In the second and third cases, the function merges the continuation of both processes, if both have the same prefix. This allows for participants to have a common behaviour in both branches until they have enough information to know which branch was chosen. The fourth case deals with process of the form $\mathbf{0}$ and χ .

Note that $_ \uplus _$ is a partial function, e.g., it might be the case that a participant behaves differently in two branches without being aware of which branch was chosen, in such a

case the projection of that participant is undefined, cf. the last case of the definition.

Example 3.4. We illustrate the use of $_ \uplus _$ in a projection. Consider the global type

$$\mathcal{G} = (s \rightarrow r : a; \mathcal{G}_3) + (s \rightarrow r : b; \mathcal{G}_3)$$

such that $\mathcal{G}_3 = s \rightarrow r' : a + s \rightarrow r' : b$, and $a \neq b$.

- For s , we have $\mathcal{G}_3 \downarrow_s = r' ! a \oplus r' ! b$ and $\mathcal{G} \downarrow_s = (r' ! a; \mathcal{G}_3 \downarrow_s) \oplus (r' ! b; \mathcal{G}_3 \downarrow_s)$

The projection onto the participant that makes the choice is simple since, by well-formedness, the sender is the same in each branch, and the guards are pairwise distinct.

- For r , we have $\mathcal{G}_3 \downarrow_r = \mathbf{0}$, since r does not appear in \mathcal{G}_3 , and

$$\mathcal{G} \downarrow_r = s ? a; \mathcal{G}_3 \downarrow_r \uplus s ? b; \mathcal{G}_3 \downarrow_r = s ? a + s ? b$$

i.e., $_ \uplus _$ is defined here since $a \neq b$.

- For r' , we have $\mathcal{G}_3 \downarrow_{r'} = s ? a \uplus s ? b$, which is defined since $a \neq b$; and we have $\mathcal{G} \downarrow_{r'} = \mathcal{G}_3 \downarrow_{r'} \uplus \mathcal{G}_3 \downarrow_{r'}$, which is defined since the arguments of \uplus are the same.

◇

For the sake of readability, we have defined the projection of a global type using the binary operator $_ + _$. In fact, the projection of a global type is sensitive to the way choice branches are associated, as we show in Example 3.5 below.

Example 3.5. Consider the global type below.

$$\begin{aligned}
 \mathcal{G}_{\text{ex3.5}} &= \mathbf{s} \rightarrow \mathbf{r} : \mathbf{a}; \mathbf{s} \rightarrow \mathbf{r}' : \mathbf{b} & (1) \\
 &+ \\
 &\mathbf{s} \rightarrow \mathbf{r}' : \mathbf{b}; \mathbf{s} \rightarrow \mathbf{r} : \mathbf{a} & (2) \\
 &+ \\
 &\mathbf{s} \rightarrow \mathbf{r} : \mathbf{c}; \mathbf{s} \rightarrow \mathbf{r}' : \mathbf{d} & (3)
 \end{aligned}$$

The success of the projection of $\mathcal{G}_{\text{ex3.5}}$ onto participant \mathbf{r} (and \mathbf{r}') depends on how the branches of the global type are partitioned. Let us project each branch onto \mathbf{r} separately:

$$(1) : \mathbf{s}?a \qquad (2) : \mathbf{s}?a \qquad (3) : \mathbf{s}?c$$

If $_ \uplus _$ is used to merge branches (1) and (2) first, we obtain $\mathbf{s}?a$ which may be merged with branch (3), i.e., $\mathbf{s}?a \uplus \mathbf{s}?c = \mathbf{s}?a + \mathbf{s}?c \neq \perp$.

However, if we first merge branches (2) and (3), we obtain $\mathbf{s}?a + \mathbf{s}?c$, which may not be merged with branch (1), i.e., $(\mathbf{s}?a + \mathbf{s}?c) \uplus \mathbf{s}?a = \perp$. \diamond

In order to cater for global types such as $\mathcal{G}_{\text{ex3.5}}$ above, we introduce a more precise version of the case of Definition 3.1 dealing with the projection of choice branches onto a *receiver* participant.

Let $\mathcal{G} = \sum_{i \in I} \mathbf{s} \rightarrow \mathbf{r}_i : \mathbf{a}_i; \mathcal{G}_i$, the projection of $\mathcal{G} \downarrow_{\mathbf{n}}$, with $\mathbf{n} \in \{\mathbf{r}_i \mid i \in I\}$, is defined as $\mathcal{G} \downarrow_{\mathbf{n}} \stackrel{\text{def}}{=} \mathcal{G}_1 \downarrow_{\mathbf{n}} \uplus \mathcal{G}_2 \downarrow_{\mathbf{n}}$, if the following holds:

$$\begin{aligned}
 \exists J, K : J \cup K = I \text{ and } J \cap K = \emptyset \text{ such that } & \mathcal{G}_1 = \sum_{j \in J} \mathbf{s} \rightarrow \mathbf{r}_j : \mathbf{a}_j; \mathcal{G}_j, \\
 & \mathcal{G}_2 = \sum_{k \in K} \mathbf{s} \rightarrow \mathbf{r}_k : \mathbf{a}_k; \mathcal{G}_k, \text{ and} \\
 & \mathcal{G}_1 \downarrow_{\mathbf{n}} \uplus \mathcal{G}_2 \downarrow_{\mathbf{n}} \neq \perp
 \end{aligned}$$

Note that because $_ \uplus _$ is not used to project choice branches onto the participant who

makes the choice, there is no need to cater for associativity in that case. In addition, if the projection of a global type onto a participant is defined, then it is unique, cf. Lemma 3.1 below.

3.3.2 Properties of Well-formed Global Types

We give the basic properties which hold for well-formed global types. Some of these properties are based on previous results from [40], thus we show that these results are applicable in our setting.

Lemma 3.1. *For all \mathcal{G} such that $\vdash \mathcal{G}$ and \mathcal{G} is projectable, let $S = \prod_{n \in \mathcal{P}(\mathcal{G})} n[\mathcal{G} \downarrow_n]$, the following holds:*

1. $\forall n \in \mathcal{P}(\mathcal{G}) : \mathcal{G} \downarrow_n$ is unique.
2. $\forall n \in \mathcal{P}(\mathcal{G}) : \mathcal{G} \downarrow_n$ is a process.
3. S is a program.
4. $S \mid Q(S)$ is a runtime system.

Proof. Item 1 follows from the condition that sets of participants are pairwise disjoint in $\mathcal{G} \mid \mathcal{G}'$, and commutativity and associativity of $+$. We show the latter by contradiction. Assume we have $\mathcal{G} = (\mathcal{G}_1 + \mathcal{G}_2) + \mathcal{G}_3$ and $\mathcal{G}' = \mathcal{G}_1 + (\mathcal{G}_2 + \mathcal{G}_3)$ with $\exists n \in \mathcal{P}(\mathcal{G})$ such that $\mathcal{G} \downarrow_n \neq \mathcal{G}' \downarrow_n$, $\mathcal{G} \downarrow_n \neq \perp$, and $\mathcal{G}' \downarrow_n \neq \perp$.

- If n is the sender, then, by Definition 3.1, we have

$$\mathcal{G} \downarrow_n = (\mathcal{G}_1 \downarrow_n \oplus \mathcal{G}_2 \downarrow_n) \oplus \mathcal{G}_3 \downarrow_n \quad \text{and} \quad \mathcal{G}' \downarrow_n = \mathcal{G}_1 \downarrow_n \oplus (\mathcal{G}_2 \downarrow_n \oplus \mathcal{G}_3 \downarrow_n)$$

and the result follows directly by definition of processes (i.e., by associativity of \oplus).

- If n is not the sender, then we have

$$\mathcal{G} \downarrow_n = (\mathcal{G}_1 \downarrow_n \uplus \mathcal{G}_2 \downarrow_n) \uplus \mathcal{G}_3 \downarrow_n \quad \text{and} \quad \mathcal{G} \downarrow_n = \mathcal{G}_1 \downarrow_n \uplus (\mathcal{G}_2 \downarrow_n \uplus \mathcal{G}_3 \downarrow_n) \quad (3.1)$$

By Definition 3.2, either all the projections must be an external choice, or all the projections must be an internal choice (otherwise merge would not be defined).

- If the projections are external choices, then by (3.1) and Definition 3.2, it must be the case that either all the prefixes are identical or all the prefixes are different. By contradiction, if we have

$$\mathcal{G}_1 \downarrow_n = s?a \quad \mathcal{G}_2 \downarrow_n = s?a \quad \mathcal{G}_3 \downarrow_n = s?b$$

then $\mathcal{G}_1 \downarrow_n \uplus (\mathcal{G}_2 \downarrow_n \uplus \mathcal{G}_3 \downarrow_n)$ is not defined. Similarly, if we have

$$\mathcal{G}_1 \downarrow_n = s?b \quad \mathcal{G}_2 \downarrow_n = s?a \quad \mathcal{G}_3 \downarrow_n = s?a$$

then $(\mathcal{G}_1 \downarrow_n \uplus \mathcal{G}_2 \downarrow_n) \uplus \mathcal{G}_3 \downarrow_n$ is not defined. Thus, associating the branches differently does not affect the projection.

- If the projections are internal choices then by (3.1) and Definition 3.2, it must be the case that all the prefixes are identical (as above). Thus, associating the branches differently does not affect the projection.

We show item 2 by induction on the structure of \mathcal{G} .

- If $\mathcal{G} \equiv \mu\chi.\mathcal{G}'$ the result follows by induction hypothesis and the fact $\mathcal{G} \downarrow_n = \mathbf{0}$ if $\mathcal{G}' \downarrow_n = \chi$.
- If $\mathcal{G} \equiv \mathbf{0}$ or $\mathcal{G} \equiv \chi$ the result holds trivially.

- If $\mathcal{G} \equiv s \rightarrow r : a; \mathcal{G}'$ the result follows directly by induction hypothesis and Definition 3.1.
- If $\mathcal{G} \equiv \mathcal{G}_0 + \mathcal{G}_1$ we have to distinguish two cases. Either (i) n is the sender on each branch, and therefore its projections are guarded by different prefixes (cf. syntax of global types and well-formedness rules); or (ii) n is not the sender.

In this case, we know by induction hypothesis that $\mathcal{G}_0 \downarrow_n$ and $\mathcal{G}_1 \downarrow_n$ are processes, we have to show that the result of $\mathcal{G}_0 \downarrow_n \uplus \mathcal{G}_1 \downarrow_n$ (Definition 3.2) is indeed a process. Note that since \mathcal{G} is projectable, $_ \uplus _$ must always be defined. We show that if P and Q are processes and $P \uplus Q \neq \perp$, then $P \uplus Q$ is a process. We proceed by induction on the structure of P and Q (abstracting from commutative invocations).

– If

$$P \equiv \sum_{i \in I} s ? a_i ; P'_i \quad \text{and} \quad Q \equiv \sum_{j \in J} s ? a_j ; Q'_j$$

then $P \uplus Q$ is defined if

$$\forall i \in I : \forall j \in J : a_i \neq a_j \text{ and } I, J \neq \emptyset$$

and thus we have a process since both P and Q are processes and they do not share any prefix, or if $I = J$ (i.e., the prefixes are the same in both processes), and we have the result since each $P'_i \uplus Q'_j$ is a process by induction hypothesis.

– If

$$P \equiv \bigoplus_{i \in I} r_i ! a_i ; P'_i \quad \text{and} \quad Q \equiv \bigoplus_{j \in J} r_j ! a_j ; Q'_j$$

then $P \uplus Q$ is defined if either $P \equiv Q$, in which case $P \uplus Q = P$, and the result follows directly; or if $I = J$ (i.e., the prefixes are the same in both processes), and we have the result since each $P'_i \uplus Q'_j$ is a process by induction hypothesis.

– If

$$P \equiv \sum_{i \in I} s_i ? a_i ; P'_i \quad \text{and} \quad Q \equiv \bigoplus_{j \in J} r_j ! a_j ; Q'_j$$

then $P \uplus Q = \perp$, i.e., a contradiction with the fact that \mathcal{G} is projectable.

– If

$$P \equiv \mu \chi . P' \quad \text{or} \quad P \equiv \mathbf{0} \quad \text{or} \quad P \equiv \chi$$

then $P \uplus Q = P$ only if $P \equiv Q$, and we are done by assumption that P is a process.

- If $\mathcal{G} \equiv \mathcal{G}_0 \mid \mathcal{G}_1$, the result follows directly from induction hypothesis.

Item 3 follows directly from items 1 and 2. Item 4 follows directly from item 3 and the definition of $Q(S)$. □

We state some results that were established in [40] and that are applicable in our setting for the following reasons. (i) Our local types can be easily translated to communicating machines, as shown in [40]. (ii) Our global types are essentially a subset of the *generalised global types* of [40]. Indeed, we allow only tree-shaped global types and do not allow participants to appear in different concurrent branches. (iii) Participants are single-threaded in our setting, thus the *linearity* condition of [40] holds trivially. (iv) The *local choice* condition is satisfied for all well-formed and projectable global types. This condition is satisfied for a global type \mathcal{G} if for each sub-term $\mathcal{G}_0 + \mathcal{G}_1$ of \mathcal{G} , and for each participant $n \in \mathcal{P}(\mathcal{G})$, n has enough information to “know” which branch of the choreography was chosen. Indeed, by Definition 3.2, if a participant has different behaviours in different branches of a global type, then these behaviours must be prefixed by different inputs.

Definition 3.3 (Properties of runtime systems [40]). *Let S be a runtime system, we say that*

- *S is stable if*

$$\forall \mathbf{sr} \in \mathcal{C}(S) : S[\mathbf{sr}] = \varepsilon$$

- *S is a deadlock if $\forall \mathbf{sr} \in \mathcal{C}(S) : S[\mathbf{sr}] = \varepsilon$, and*

$$\exists \mathbf{n} \in \mathcal{P}(S) : S(\mathbf{n}) \equiv \mathbf{s?a};P+P' \quad \text{and} \quad \forall \mathbf{n} \in \mathcal{P}(S) : S(\mathbf{n}) \equiv \mathbf{s?a};P+P' \vee S(\mathbf{n}) \equiv \mathbf{0}$$

- *S is an orphan message configuration if*

$$\forall \mathbf{n} \in \mathcal{P}(S) : S(\mathbf{n}) \equiv \mathbf{0} \quad \text{and} \quad \exists \mathbf{sr} \in \mathcal{C}(S) : S[\mathbf{sr}] \neq \varepsilon$$

- *S is an unspecified reception configuration if $\exists \mathbf{r} \in \mathcal{P}(S) :$*

$$S(\mathbf{r}) \equiv \mathbf{s?a};P+P' \implies S[\mathbf{sr}] \neq \varepsilon \wedge S[\mathbf{sr}] \neq \mathbf{a} \cdot \rho$$

A stable system is a system in which all the queues are empty. A system is a deadlock if all its queues are empty, all participants are either terminated or waiting for an input, and there is one participant expecting a message. An orphan message configuration is a system where all the participants have terminated their behaviours and there is one queue that is not empty. An unspecified reception configuration is a system for which there is a participant who is permanently unable to read a datum from one of its queues.

We can now recall the result established in [40], adapted to our setting.

Lemma 3.2 (Properties of projections [40]). *For all \mathcal{G} such that $\vdash \mathcal{G}$ and \mathcal{G} is projectable.*

Let

$$S \equiv \bigsqcup_{\mathbf{n} \in \mathcal{P}(\mathcal{G})} \mathbf{n}[\mathcal{G} \downarrow_{\mathbf{n}}] \quad \text{and} \quad \hat{S} \equiv S \mid \mathcal{Q}(S)$$

For all S' such that $\hat{S} \implies S'$,

- *S' is not a deadlock.*

- S' is not an orphan message configuration.
- S' is not an unspecified reception configuration.
- Either there is S'' such that $S' \rightarrow S''$, or

$$\forall \mathbf{sr} \in \mathcal{C}(S) : S[\mathbf{sr}] = \varepsilon \quad \text{and} \quad \forall \mathbf{n} \in \mathcal{P}(S) : S(\mathbf{n}) \equiv \mathbf{0}$$

Definition 3.4 (1-buffer execution). *We say that S' is reachable from S by a 1-buffer execution if there exists $\{S_i \mid 1 \leq i \leq n\}$ such that $S_1 = S$, $S_n = S'$, and*

$$\begin{aligned} &\forall 1 \leq i < n : S_i \rightarrow S_{i+1}, \text{ and} \\ &\exists \mathbf{sr} \in \mathcal{C}(S_i) : S_i[\mathbf{sr}] \neq \varepsilon \implies \forall \mathbf{s'r'} \in \mathcal{C}(S_i) \setminus \{\mathbf{sr}\} : S_i[\mathbf{s'r'}] = \varepsilon \end{aligned}$$

There is a 1-buffer execution from a system S to S' , if it is possible to reach S' by a series of intermediary configurations such that there is at most one non-empty buffer for each of them.

Lemma 3.3 (1-buffer execution [40]). *For all G such that $\vdash G$ and G is projectable. Let*

$$S \equiv \bigsqcup_{\mathbf{n} \in \mathcal{P}(G)} \mathbf{n}[G \downarrow_{\mathbf{n}}] \quad \text{and} \quad \hat{S} \equiv S \mid Q(S)$$

For all S' such that $\hat{S} \implies S'$, if S' is stable, then there is a 1-buffer execution from \hat{S} to S' .

3.4 Synthesising Global Types

We now introduce a type system to synthesise a global type G from a system S , so that S satisfies the same properties as the system consisting of the projections of a well-formed and projectable global type (cf. Lemma 3.2). Also, the system obtained by projecting a well-formed and projectable global type is always typable. One objective of the typing

system is to allow the diagram below to commute:

$$\begin{array}{ccc}
 S & \xrightarrow{\blacktriangleright} & \mathcal{G} \\
 \approx \Big\| & & \Big\| \equiv \\
 S' & \xleftarrow{\blacktriangleleft} & \mathcal{G}'
 \end{array}$$

Essentially, if a global type \mathcal{G} may be synthesised (\blacktriangleright) from an initial system S , then \mathcal{G} may be projected (\blacktriangleleft) so to obtain S' bisimilar to the original system. Analogously, the synthesis of the projections of a well-formed and projectable global type \mathcal{G}' yields a structurally congruent global type \mathcal{G} .

To synthesise \mathcal{G} from a system S , a careful analysis of what actions can occur at each possible state of S is necessary. We define the *ready set* of a system as follows:

$$\mathbf{R}(S) = \begin{cases} \{\mathbf{sr}\} \cup \mathbf{R}(S'), & \text{if } S \equiv \mathbf{r}[\sum_{i \in I} \mathbf{s} ? a_i ; P_i] \mid S' \\ \{\overline{\mathbf{sr}}_i \mid i \in I\} \cup \mathbf{R}(S'), & \text{if } S \equiv \mathbf{s}[\bigoplus_{i \in I} \mathbf{r}_i ! a_i ; P_i] \mid S' \\ \emptyset, & \text{if } S \equiv \mathbf{0} \end{cases}$$

We define

$$S \uparrow \iff \exists \mathbf{sr} \in \mathbb{C} : \mathbf{sr} \in \mathbf{R}(S) \wedge \overline{\mathbf{sr}} \in \mathbf{R}(S)$$

and write $S \not\uparrow$ if $S \uparrow$ does not hold.

3.4.1 Validation Rules

We use judgements of the form:

$$\Gamma \vdash S \blacktriangleright \mathcal{G}$$

saying that the system S forms a choreography defined by a global type \mathcal{G} , under the environment Γ . The environment Γ is a map from participant names and local recursion

variables to global recursion variables. Formally, we define Γ as follows

$$\Gamma ::= \circ \mid (\mathfrak{n}, \chi) : \chi' \cdot \Gamma'$$

Environment \circ is the empty context, and the environment $(\mathfrak{n}, \chi) : \chi'$ states that the *local* recursion variable χ of participant \mathfrak{n} is associated to the *global* recursion variable χ' . We write \cdot for the disjoint union of environments, i.e., we write $\Gamma \cdot \Gamma'$ if $\text{dom}(\Gamma) \cap \text{dom}(\Gamma') = \emptyset$. The domain and image of Γ are defined below.

$$\begin{aligned} \text{dom}(\Gamma) &= \begin{cases} \{(\mathfrak{n}, \chi)\} \cup \text{dom}(\Gamma'), & \text{if } \Gamma = (\mathfrak{n}, \chi) : \chi' \cdot \Gamma' \\ \emptyset, & \text{if } \Gamma = \circ \end{cases} \\ \text{img}(\Gamma) &= \begin{cases} \{\chi'\} \cup \text{img}(\Gamma'), & \text{if } \Gamma = (\mathfrak{n}, \chi) : \chi' \cdot \Gamma' \\ \emptyset, & \text{if } \Gamma = \circ \end{cases} \end{aligned}$$

A global type \mathcal{G} can be synthesised from the *program* S if the judgement

$$\circ \vdash S \blacktriangleright \mathcal{G}$$

is derivable from the rules in Figure 3.3. The rules are driven by the ready set of S and the structure of its processes, we detail each of them below.

Sequential interactions are dealt with by rule [;]. The rule validates prefixes provided that the continuation is typable and that no other interactions are possible in S . For instance, rule [;] does not apply to

$$s_1[r_1!a; P_1] \mid r_1[s_1?a; Q_1] \mid s_2[r_2!b; P_2] \mid r_2[s_2?b; Q_2] \quad \times$$

because there is no ordering relation between the actions of s_1 and r_1 on one hand, and the actions of s_2 and r_2 on the other hand. In such cases, the rule [||] should be used.

$$\begin{array}{c}
 \frac{\Gamma \vdash s[P] \mid r[Q] \mid S \triangleright \mathcal{G} \quad S \not\downarrow}{\Gamma \vdash s[r!a;P] \mid r[s?a;Q] \mid S \triangleright s \rightarrow r:a; \mathcal{G}} \quad [!:] \\
 \\
 \frac{\circ \vdash S \triangleright \mathcal{G} \quad \circ \vdash S' \triangleright \mathcal{G}'}{\Gamma \vdash S \mid S' \triangleright \mathcal{G} \mid \mathcal{G}'} \quad [||] \\
 \\
 \frac{\Gamma \vdash s[P] \mid S \triangleright \mathcal{G} \quad \Gamma \vdash s[Q] \mid S \triangleright \mathcal{G}' \quad S \not\downarrow}{\Gamma \vdash s[P \oplus Q] \mid S \triangleright \mathcal{G} + \mathcal{G}'} \quad [\oplus] \\
 \\
 \frac{\Gamma \vdash r[P] \mid S \triangleright \mathcal{G}}{\Gamma \vdash r[P+Q] \mid S \triangleright \mathcal{G}} \quad [+] \\
 \\
 \frac{\exists 1 \leq i, j \leq k : (n_i[P_i] \mid n_j[P_j]) \not\downarrow \quad \Gamma \cdot (n_1, \chi_1) : \chi, \dots, (n_k, \chi_k) : \chi \vdash n_1[P_1] \mid \dots \mid n_k[P_k] \triangleright \mathcal{G}}{\Gamma \vdash n_1[\mu\chi_1.P_1] \mid \dots \mid n_k[\mu\chi_k.P_k] \triangleright \mu\chi.\mathcal{G}} \quad [\mu] \\
 \\
 \frac{\forall 1 \leq i \leq k : \Gamma(n_i, \chi_i) = \chi}{\Gamma \vdash n_1[\chi_1] \mid \dots \mid n_k[\chi_k] \triangleright \chi} \quad [\chi] \\
 \\
 \frac{S \equiv S' \quad \Gamma \vdash S' \triangleright \mathcal{G}}{\Gamma \vdash S \triangleright \mathcal{G}} \quad [\equiv] \quad \frac{\forall n \in \mathcal{P}(S) : S(n) = \mathbf{0}}{\Gamma \vdash S \triangleright \mathbf{0}} \quad [0]
 \end{array}$$

Figure 3.3: Validation rules for programs

Concurrent branches are introduced by rule $[||]$. The rule validates concurrent branches when they can be validated using a partition of the system begin considered, recall that $\mathcal{P}(S) \cap \mathcal{P}(S') = \emptyset$.

Choice in a global type is dealt with by rules $[\oplus]$ and $[+]$. Rule $[\oplus]$ introduces the global type choice operator, it requires that both branches are typable and that no other interactions are possible in S – for the same reason as in rule $[!:]$. Rule $[+]$ allows to discharge a branch of an external choice. This allows to type systems of the form:

$$s[r!a \oplus r!b] \mid r[s?a + s?b + s?c] \quad \checkmark$$

Recursion is handled by rules $[\mu]$ and $[\chi]$. The former rule “guesses” the participants involved in a recursive behaviour. If two of them interact, $[\mu]$ validates the recursion provided that the system can be typed when such participants are associated to the global recursion variable χ – assuming that $\chi \notin \text{img}(\Gamma)$. Rule $[\chi]$ checks that all the participants in the recursion have reached a local recursion variable corresponding to the global recursion.

The termination $\mathbf{0}$ is introduced by rule $[\mathbf{0}]$, which only applies when all the participants in S are terminated.

Rule $[\equiv]$ validates a system up to structural congruence. This rule notably allows recursive behaviours to be unfolded, see Example 3.7 below.

3.4.2 Applying the Rules

We give a few examples of derivations. We give the full derivation of the system S_{BS} from Section 3.1, show that we support an example borrowed from [41], and we illustrate the use of rule $[\equiv]$ to unfold behaviours.

First, we show how to synthesise the global type corresponding to system S_{BS} of Section 3.1. Recall that this system is defined as follows:

$$S_{BS} = b_1[P_1] \mid s_1[S_1] \mid b_2[P_2] \mid s_2[S_2]$$

with

$$P_1 = b_2! \text{agreement}; b_2? \text{ack}; Q_1$$

$$P_2 = b_1? \text{agreement}; b_1! \text{ack}; Q_2$$

$$Q_i = s_i! \text{request}; s_i? \text{quote}; \mu\chi. (s_i! \text{ok} \oplus s_i! \text{no}; s_i? \text{quote}; \chi) \quad i \in \{1, 2\}$$

$$S_i = b_i? \text{request}; b_i! \text{quote}; \mu\chi'. (b_i? \text{ok} + b_i? \text{no}; b_i! \text{quote}; \chi') \quad i \in \{1, 2\}$$

We describe the derivation from the rule typing the overall system to the leaves of the type derivation. The derivation of the overall system S_{BS} is as follows.

$$\begin{array}{c}
 \frac{\frac{\frac{}{\nabla_l}}{\circ \vdash b_1[Q_1] \mid s_1[S_1] \blacktriangleright \mathcal{G}_2^1}{} \quad \frac{\frac{}{\nabla_r}}{\circ \vdash b_2[Q_2] \mid s_2[S_2] \blacktriangleright \mathcal{G}_2^2}{} }{[\]} \quad \frac{}{\circ \vdash b_1[Q_1] \mid s_1[S_1] \mid b_2[Q_2] \mid s_2[S_2] \blacktriangleright \mathcal{G}_2^1 \mid \mathcal{G}_2^2}}{[\]} \\
 \frac{[\]}{\circ \vdash b_1[b_2?ack; Q_1] \mid s_1[S_1] \mid b_2[b_1!ack; Q_2] \mid s_2[S_2] \blacktriangleright b_2 \rightarrow b_1 : ack; (\mathcal{G}_2^1 \mid \mathcal{G}_2^2)} \\
 \frac{[\]}{\circ \vdash b_1[P_1] \mid s_1[S_1] \mid b_2[P_2] \mid s_2[S_2] \blacktriangleright \mathcal{G}_{S_{BS}}}
 \end{array}$$

First, rule $[\]$ is applied twice, recording the two interactions between b_1 and b_2 . Then, rule $[\]$ is applied to separate the system in two concurrent sub-systems. The derivation of the left (∇_l) and right (∇_r) branches are similar and we give a parametrised version of them below. We pose

$$\begin{aligned}
 \hat{\mathcal{G}}_r &= b_i \rightarrow s_i : ok + b_i \rightarrow s_i : no; s_i \rightarrow b_i : quote; \chi \\
 \hat{\mathcal{Q}}_r &= s_i !ok \oplus s_i !no; s_i ?quote; \chi \\
 \hat{\mathcal{S}}_r &= b_i ?ok + b_i ?no; b_i !quote; \chi'
 \end{aligned}$$

in the following derivation:

$$\begin{array}{c}
 \frac{\frac{\frac{}{\nabla_{in-rec}}}{(b_i, \chi) : \chi \cdot (s_i, \chi') : \chi \vdash b_i[\hat{\mathcal{Q}}_r] \mid s_i[\hat{\mathcal{S}}_r] \blacktriangleright \hat{\mathcal{G}}_r}{} }{[\mu]} \quad \frac{}{\circ \vdash b_i[\mu\chi.\hat{\mathcal{Q}}_r] \mid s_i[\mu\chi'.\hat{\mathcal{S}}_r] \blacktriangleright \mu\chi.\hat{\mathcal{G}}_r}}{[\]} \\
 \frac{[\]}{\circ \vdash b_i[s_i?quote; \mu\chi.\hat{\mathcal{Q}}_r] \mid s_i[b_i!quote; \mu\chi'.\hat{\mathcal{S}}_r] \blacktriangleright s_i \rightarrow b_i : quote; \mu\chi.\hat{\mathcal{G}}_r} \\
 \frac{[\]}{\circ \vdash b_i[Q_i] \mid s_i[S_i] \blacktriangleright \mathcal{G}_2^1}
 \end{array}$$

In the sub-derivation above, rule $[\]$ is applied twice again to record the interactions be-

tween b_i and s_i , then rule $[\mu]$ is applied and the recursion environment is updated correspondingly. The next step in the sub-derivation $\nabla_{\text{in-rec}}$ is given below:

$$[\oplus] \frac{\nabla_{\text{end}} \quad \nabla_{\text{rec}}}{(b_i, \chi) : \chi \cdot (s_i, \chi') : \chi \vdash b_i[\hat{Q}_r] \mid s_i[\hat{S}_r] \blacktriangleright \hat{G}_r}$$

where $[\oplus]$ is used to split the derivation in two parts, one per branch of the internal choice in b_i , both sub-derivations are given below. First, we give ∇_{end} :

$$\begin{array}{c} [\emptyset] \frac{}{(b_i, \chi) : \chi \cdot (s_i, \chi') : \chi \vdash b_i[\mathbf{0}] \mid s_i[\mathbf{0}] \blacktriangleright \mathbf{0}} \\ [:] \frac{}{(b_i, \chi) : \chi \cdot (s_i, \chi') : \chi \vdash b_i[s_i!ok] \mid s_i[b_i?ok] \blacktriangleright b_i \rightarrow s_i : ok} \\ [+]\frac{}{(b_i, \chi) : \chi \cdot (s_i, \chi') : \chi \vdash b_i[s_i!ok] \mid s_i[\hat{S}_r] \blacktriangleright b_i \rightarrow s_i : ok} \end{array}$$

In the sub-derivation above, rule $[+]$ is used to discard the right branch of s_i , then rule $[:]$ is used for the ok interaction. The last rule applied is rule $[\emptyset]$ which terminates the global type.

We now give ∇_{rec} , where we pose $\hat{G}_l = b_i \rightarrow s_i : \text{no}; s_i \rightarrow b_i : \text{quote}; \chi$.

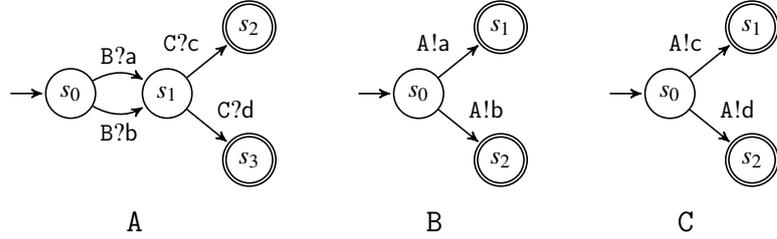
$$\begin{array}{c} [x] \frac{}{(b_i, \chi) : \chi \cdot (s_i, \chi') : \chi \vdash b_i[\chi] \mid s_i[\chi'] \blacktriangleright \chi} \\ [:] \frac{}{(b_i, \chi) : \chi \cdot (s_i, \chi') : \chi \vdash b_i[s_i?quote; \chi] \mid s_i[b_i!quote; \chi'] \blacktriangleright s_i \rightarrow b_i : \text{quote}; \chi} \\ [:] \frac{}{(b_i, \chi) : \chi \cdot (s_i, \chi') : \chi \vdash b_i[s_i!no; s_i?quote; \chi] \mid s_i[b_i?no; b_i!quote; \chi'] \blacktriangleright \hat{G}_l} \\ [+]\frac{}{(b_i, \chi) : \chi \cdot (s_i, \chi') : \chi \vdash b_i[s_i!no; s_i?quote; \chi] \mid s_i[\hat{S}_r] \blacktriangleright \hat{G}_l} \end{array}$$

In this, case we use rule $[+]$ to discard the left branch of s_i , then use rule $[:]$ twice, and finally we apply rule $[x]$ to introduce the recursion variable. Note that this rule is applicable

since the recursion environment maps each local recursion variable to the same global recursion variable.

We now give two additional examples of typing derivations. Example 3.6 shows the typing of local types obtained from communicating machines given in Remark 4.1 of [41]. Example 3.7 illustrates the use of rule $[\equiv]$, whose main role is to unfold behaviours when required.

Example 3.6. Consider the following communicating machines (borrowed from [41]):



These machines may be easily translated to our local types so to obtain the following system:

$$S_{\text{ex3.6}} = A[B?a;(C?c + C?d) + B?b;(C?c + C?d)] \mid B[A!a \oplus A!b] \mid C[A!c \oplus A!d]$$

In fact, we have $\circ \vdash S_{\text{ex3.6}} \triangleright \mathcal{G}_{\text{ex3.6}}$, with $\mathcal{G}_{\text{ex3.6}}$ defined as follows:

$$\mathcal{G}_{\text{ex3.6}} = B \rightarrow A : a; \mathcal{G}' + B \rightarrow A : b; \mathcal{G}' \quad \mathcal{G}' = C \rightarrow A : c + C \rightarrow A : d$$

The complete typing derivation is given in Figure 3.4. ◇

Example 3.7. Consider the system below.

$$S_{\text{ex3.7}} = s[r!a; r?b; \mu\chi.r!a; r?b; \chi] \mid r[\mu\chi.s?a; s!b; \chi]$$

3.5 Properties of the Synthesis

In this section, we show the properties of the type system, which, notably, allow us to prove that the diagram of Section 3.4 commutes. We summarise our results below.

- Theorem 3.1 gives safety and progress properties for typable systems.
- Theorem 3.2 states that typability is decidable.
- Theorem 3.3 states that the global type assigned to a system is unique (up to structural congruence).
- Lemma 3.8 characterises the relationship between the behaviours in the initial system and the projections of the synthesised global type. This lemma is crucial to show that a synthesised global type is always well-formed (Theorem 3.4).
- Theorem 3.5 gives a subject reduction result.
- Theorem 3.6 shows that the system consisting of the projections of a synthesised global type is bisimilar to the original system.
- Theorem 3.7 shows that any well-formed and projectable global type may be projected and synthesised again to an equivalent global type.

The following intermediary result follows directly from the rules of Figure 3.3.

Lemma 3.4. *If $\Gamma \vdash S \blacktriangleright G$ then $\forall n \in \mathcal{P}(S) : S(n) \neq \mathbf{0} \iff n \in \mathcal{P}(G)$*

Proof. Straightforward induction on the derivation. □

Theorem 3.1 below follows from the results in the rest of this section. Namely, we will show that a typable system and the system consisting of its projections are bisimilar, and that a synthesised global type is well-formed. This allows us to reuse the results of Lemma 3.2 to guarantee the same safety properties as with the usual top-down approach of the multiparty session types.

Theorem 3.1 (Properties of typable systems). *Let S be a stable runtime system, such that*

$\circ \vdash S \triangleright \mathcal{G}$. *For all S' such that $S \Longrightarrow S'$,*

- *S' is not a deadlock.*
- *S' is not an orphan message configuration.*
- *S' is not an unspecified reception configuration.*
- *Either there is S'' such that $S' \rightarrow S''$, or*

$$\forall \mathbf{sr} \in \mathcal{C}(S') : S'[\mathbf{sr}] = \varepsilon \quad \text{and} \quad \forall \mathbf{n} \in \mathcal{P}(S') : S'(\mathbf{n}) \equiv \mathbf{0}$$

- *If S and S' are stable, then there is a 1-buffer execution from S to S' .*

Proof. The result follows directly from the fact that \mathcal{G} is a well-formed global type (Theorem 3.4), the projections of a well-formed global types guarantee such properties (Lemmas 3.2 and 3.3), and the projections of a synthesised global type are \approx -equivalent to the original system (Theorem 3.6). □

3.5.1 Decidability

A crucial result is that it is decidable whether or not a system may be assigned a global type. We give a few examples of type derivations to show that the answer to this question may not be straightforward.

Example 3.8. Consider the typable system:

$$s[\mu\chi.r!a;\chi] \mid r[\mu\chi.s?a;\chi]$$

A possibly infinite derivation for this system would be of the form

$$\begin{array}{c}
 \vdots \\
 \text{[}\equiv\text{]} \frac{}{\circ \vdash s[\mu\chi.r!a;\chi] \mid r[\mu\chi.s?a;\chi] \blacktriangleright \mathcal{G}} \\
 \text{[!]} \frac{}{\circ \vdash s[r!a;\mu\chi.r!a;\chi] \mid r[s?a;\mu\chi.s?a;\chi] \blacktriangleright s \rightarrow r : a; \mathcal{G}} \\
 \text{[}\equiv\text{]} \frac{}{\circ \vdash s[\mu\chi.r!a;\chi] \mid r[\mu\chi.s?a;\chi] \blacktriangleright s \rightarrow r : a; \mathcal{G}}
 \end{array}$$

where rules $\text{[}\equiv\text{]}$ and [!] are applied successively, possibly endlessly. Of course, if we apply rule $\text{[}\mu\text{]}$ first we are done in a few steps:

$$\begin{array}{c}
 \text{[}\chi\text{]} \frac{}{(s, \chi) : \chi \cdot (r, \chi) : \chi \vdash s[\chi] \mid r[\chi] \blacktriangleright \chi} \\
 \text{[!]} \frac{}{(s, \chi) : \chi \cdot (r, \chi) : \chi \vdash s[r!a;\chi] \mid r[s?a;\chi] \blacktriangleright s \rightarrow r : a; \chi} \\
 \text{[}\mu\text{]} \frac{}{\circ \vdash s[\mu\chi.r!a;\chi] \mid r[\mu\chi.s?a;\chi] \blacktriangleright \mu\chi.s \rightarrow r : a; \chi}
 \end{array}$$

◇

Example 3.9. Consider the system

$$s[r!b;\mu\chi.r!b;r!a;\chi] \mid r[\mu\chi.s?b;s?a;\chi]$$

This system is not typable: consider the following partial derivation:

$$\begin{array}{c}
 \perp \\
 \text{[!]} \frac{}{\circ \vdash s[\mu\chi.r!b;r!a;\chi] \mid r[s?a;\mu\chi.s?b;s?a;\chi] \blacktriangleright \dots} \\
 \text{[}\equiv\text{]} \frac{}{\circ \vdash s[r!b;\mu\chi.r!b;r!a;\chi] \mid r[s?b;s?a;\mu\chi.s?b;s?a;\chi] \blacktriangleright \dots} \\
 \text{[}\equiv\text{]} \frac{}{\circ \vdash s[r!b;\mu\chi.r!b;r!a;\chi] \mid r[\mu\chi.s?b;s?a;\chi] \blacktriangleright \dots}
 \end{array}$$

Intuitively, this is due to the fact that participant s is ready to communicate with a participant (r) which exhibits directly a recursive behaviour. However, the sub-term $r!b;\mathbf{0}$ (i.e., the part of the behaviour outside the recursion, appended by $\mathbf{0}$) is not the same as behaviour within the recursion, i.e., $r!b;r!a;\mathbf{0}$. Thus, there is a mismatch between the behaviour of s outside the recursion and the one of r within the recursion. \diamond

We now introduce a few definitions that we will use in the proof of Theorem 3.2 below.

Definition 3.5 (Unfolding). *Let $\text{unfold}_i(P)$ be the i -time unfolding of P , defined as follows*

$$\text{unfold}_i(P) = \text{unfold}_1(\text{unfold}_{i-1}(P)) \quad i > 1$$

$$\text{unfold}_1(P) = \begin{cases} \bigoplus_{i \in I} r_i!a_i; \text{unfold}_1(P_i) & \text{if } P = \bigoplus_{i \in I} r_i!a_i; P_i \\ \sum_{i \in I} s?a_i; \text{unfold}_1(P_i) & \text{if } P = \sum_{i \in I} s?a_i; P_i \\ \chi & \text{if } P = \chi \\ P' \{ \mu\chi.P' / \chi \} & \text{if } P = \mu\chi.P' \end{cases}$$

We illustrate this definition with the example below.

Example 3.10. Consider the following behaviours

$$P = \mu\chi.P' \quad P' = r!a; \mu\chi'.(r!b;\chi \oplus r!c;\chi')$$

The one-time unfolding of P is

$$\text{unfold}_1(P) = r!a; \mu\chi'.(r!b;P \oplus r!c;\chi')$$

while the one-time unfolding of P' is

$$\text{unfold}_1(P') = r!a; (r!b;\chi \oplus r!c; \mu\chi'.(r!b;\chi \oplus r!c;\chi'))$$

◇

Definition 3.6 (Behaviour context). *Let a (non-recursive) behaviour context $C\langle-\rangle$ be defined as follows*

$$C\langle-\rangle ::= \bigoplus_{i \in I} r_i !a_i; C_i\langle-\rangle \mid \sum_{i \in I} s ?a_i; C_i\langle-\rangle \mid \langle-\rangle$$

and a (possibly recursive) behaviour context $C'\langle-\rangle$ be defined as follows

$$C'\langle-\rangle ::= \bigoplus_{i \in I} r_i !a_i; C'_i\langle-\rangle \mid \sum_{i \in I} s ?a_i; C'_i\langle-\rangle \mid \langle-\rangle \mid \mu\chi. C'\langle-\rangle \mid \chi$$

A non-recursive context $C\langle-\rangle$ is simply a context where no recursive definition may appear, while a possibly recursive context $C'\langle-\rangle$ may be any context defined by the syntax of behaviours. We illustrate this definition in Example 3.11 below.

Using Definitions 3.6 and 3.5, we define a “distance” between the non-recursive context of a behaviour and its unfolding.

Definition 3.7. *Given a process of the form:*

$$n[C\langle\mu\chi.C'\langle\chi\rangle\rangle] \quad \text{let } P_k = \text{unfold}_k(\mu\chi.C'\langle\chi\rangle)$$

We define $|n|$ to be the smallest k such that $C\langle\mathbf{0}\rangle$ is a sub-tree of P_k where all sub-terms of the form $\mu\chi.P'$ in P_k are replaced by $\mathbf{0}$; and $|n| = \perp$ if there is no such k .

Note that if $|n|$ is defined, it must be smaller or equal than the length of $C\langle\mathbf{0}\rangle$ (since recursion is guarded).

Example 3.11. We illustrate Definition 3.7.

- Assume we have n such that $S(n) = P$, and

$$P = C\langle\mu\chi.C'\langle\chi\rangle\rangle = r!a;r!a;\mu\chi.r!a;\chi$$

We have $|n| = 2$ since

$$\text{unfold}_2(\mu\chi.r!a;\chi) = r!a;r!a;\mu\chi.r!a;\chi$$

- Assume we have n such that $S(n) = P$, and

$$P = C\langle\mu\chi.C'\langle\chi\rangle\rangle = r!a;\mu\chi.(r!a;\chi\oplus r!b;\chi)$$

We have $|n| = 1$ since $\text{unfold}_1(\mu\chi.(r!a;\chi\oplus r!b;\chi))$ is

$$r!a;\mu\chi.(r!a;\chi\oplus r!b;\chi) \oplus r!b;\mu\chi.(r!a;\chi\oplus r!b;\chi)$$

Note that in this case, only one branch of the choice appears in $C\langle\mathbf{0}\rangle$, i.e., $r!a$.

- In Example 3.9, we had

$$P = C\langle\mu\chi.C'\langle\chi\rangle\rangle = r!b;\mu\chi.r!b;r!a;\chi$$

We have $|n| = \perp$ since

$$\begin{aligned} \text{unfold}_1(\mu\chi.r!b;r!a;\chi) &= r!b;r!a;\mu\chi.r!b;r!a;\chi \\ &\neq r!b;\mu\chi.r!b;r!a;\chi \end{aligned}$$

Note that $r!b;\mathbf{0}$ is not a sub-tree of $r!b;r!a;\mathbf{0}$, and it is clear that unfolding the behaviour again will not help.

◇

Theorem 3.2 (Decidability). *Typability is decidable.*

Proof. Decidability follows from the fact that the ready set of a system and the number of participants are finite. We show that the number of behaviour unfoldings needed to type

a system is also finite.

The need for unfolding occurs whenever a recursive participant interact with another participant, while not all the participants feature directly a recursive behaviour (cf. Examples 3.7 and 3.9). In this case, we need to unfold some participants (rule \equiv), then use rules \oplus , $+$, $;$, and/or 0 until rule μ is applicable. Note that rule $[\]$ empties the recursion environment, thus it cannot be used after a recursion definition if the axiom $[\chi]$ is to be used.

Consider the following system

$$S = S_0 | S_1$$

where

$$\begin{aligned} S_0 &= n_1[C_1\langle\mu\chi.C'_1\langle\chi\rangle\rangle] | \dots | n_j[C_j\langle\mu\chi.C'_j\langle\chi\rangle\rangle] \\ S_1 &= n_{j+1}[\mu\chi.C'_{j+1}\langle\chi\rangle] | \dots | n_k[\mu\chi.C'_{j+k}\langle\chi\rangle] \end{aligned}$$

such that $\forall 1 \leq i \leq j: C_i\langle-\rangle \neq \langle-\rangle$, $S \uparrow$, $S_0 \not\uparrow$ and $S_1 \not\uparrow$. The system S_0 consists of participants that do not exhibit a recursive behaviour (yet), while S_1 consists of participants exhibiting directly a recursive behaviour. There must be exactly one $n \in \mathcal{P}(S)$ such that

$$S \equiv n[S(n)] | S' \quad \text{and} \quad S' \not\uparrow$$

Since if more than one pair of participants may interact directly after the recursion, it means that rule $[\]$ will have to be used further in the derivation tree, and therefore the recursion environment will be emptied.

Let C_i for $j < i \leq j+k$ be the empty context, we can rewrite S such that

$$S \equiv \big|_{i \in I} n_i[C_i\langle\mu\chi.C'_i\langle\chi\rangle\rangle] \quad I = \{i \mid 1 \leq i \leq j+k\}$$

Given S as above, we will use the function $|n_i|$ on each participant. Note that if one $|n_i|$ is not defined, then S is not typable, cf. Example 3.9.

We define $M \stackrel{def}{=} \max\{|n_i| \mid i \in I\}$, and $K(i) \stackrel{def}{=} M - |n_i|$, in order to be able to unfold each behaviour so that all of them are unfolded to the same extent, let

$$S_* \equiv \prod_{i \in I} n_i [\text{unfold}_{K(i)}(C_i \langle \mu\chi. C'_i \langle \chi \rangle \rangle)]$$

We show that

$$\Gamma \vdash S_* \triangleright \mathcal{G}' \iff \Gamma \vdash S \triangleright \mathcal{G} \quad \text{with } \mathcal{G} \equiv \mathcal{G}'$$

By definition of $\text{unfold}(-)$ and since $C \langle - \rangle$ does not contain recursive definitions, we have

$$S_* \equiv \prod_{i \in I} n_i [C_i \langle \text{unfold}_{K(i)}(\mu\chi. C'_i \langle \chi \rangle) \rangle] \quad (3.4)$$

$$S_* \equiv \prod_{i \in I} n_i [C_i \langle C'_i \langle C'_i \langle \dots C'_i \langle \mu\chi. C'_i \langle \chi \rangle \rangle \dots \rangle \rangle] \quad (3.5)$$

Where $C'_i \langle - \rangle$ has been unfolded $K(i)$ times in (3.5). It is easy to see that S_* is typable if

$$\prod_{i \in I} n_i [C_i \langle C'_i \langle C'_i \langle \dots C'_i \langle \mathbf{0} \rangle \dots \rangle \rangle] \quad \text{and} \quad \prod_{i \in I} n_i [C'_i \langle \mathbf{0} \rangle]$$

are typable themselves, note that rule $[\equiv]$ does not need to be used to unfold the left-hand side system, since it is recursion free; and there is exactly one recursion less in the right hand side.

In fact, if we would unfold (3.4) once more, we would not get more chances to type S_* . Indeed, it would amount to add the right-hand sub-derivation to the left-hand sub-derivation. Thus, we have a bound on the number of required unfoldings and the type system is decidable. \square

3.5.2 Uniqueness

Theorem 3.3 below shows that, given a typable system S , the global type that is assigned to it is unique up to the congruence rules from Section 3.3.

First, we give two lemmas that we will use in the proof of Theorem 3.3.

Lemma 3.5. *Let S be a program such that $S \Downarrow$, if $\Gamma \vdash S \triangleright \mathcal{G}$ then*

- $\mathcal{G} = \mathbf{0}$ and $\forall n \in \mathcal{P}(S) : S(n) = \mathbf{0}$, or
- $\mathcal{G} = \chi$ and $\forall n \in \mathcal{P}(S) : S(n) = \chi'$ and $\Gamma(n, \chi') = \chi$

Proof. Since $S \Downarrow$, we must have that either (i) for each pair of participants s and r they do not have matching prefixes, or (ii) all the behaviours are of the form $\mathbf{0}$ or χ . In the former case, it will not be possible to reach one of the axioms (i.e., rules $[\mathbf{0}]$ or $[\chi]$). Thus, it contradicts the fact that the system is typable. In the latter case, if we have some participants with behaviour $\mathbf{0}$ and some with behaviours χ , then we would have to use rule $[\mid]$ to separate them in two subsystems, which implies that the environment Γ is emptied. Thus, we will not be able to reach the axioms for the subsystems consisting of processes with behaviour χ , and again we have a contradiction with the fact that S is typable. \square

Lemma 3.6. *Let $n[P + Q] \mid S$ be a program, if*

$$\Gamma \vdash r[P + Q] \mid S \triangleright \mathcal{G} \quad \text{and} \quad \Gamma \vdash r[P] \mid S \triangleright \mathcal{G}$$

are derivable, then $\Gamma \vdash r[Q] \mid S \triangleright \mathcal{G}$ is not derivable.

Proof. We show this by contradiction. Assume we have

$$\Gamma \vdash r[P + Q] \mid S \triangleright \mathcal{G} \tag{3.6}$$

such that $P \neq \mathbf{0}$ and $Q \neq \mathbf{0}$, note that we cannot have $P = \chi$ (or $Q = \chi$), by definition of processes. Take $S \Downarrow$, if $S \Downarrow$, we can “wait” before applying rule $[\mid]$ until the condition

holds. Indeed, the process of r may only be affected by the derivation once one of the branches is discarded (rule $[\dagger]$), then each of the branches P and Q may only be reduced if the rest of the system cannot “synchronise” without r (i.e., $S \not\Downarrow$ holds). Finally, recall that, by definition of processes, the process $P + Q$ synchronises with exactly one participant. Thus, we may “wait” until r ’s partner is ready to synchronise with r .

Assume (by contradiction) that we have

$$\Gamma \vdash r[P] \mid S \blacktriangleright \mathcal{G} \quad \text{and} \quad \Gamma \vdash r[Q] \mid S \blacktriangleright \mathcal{G}'$$

We must have

$$r[P] \mid S \Downarrow \quad \text{and} \quad r[Q] \mid S \Downarrow$$

otherwise the systems would not be typable by Lemma 3.5, since we have $S \not\Downarrow$.

Take $P + Q \equiv s?a;P_1 + s?b;Q_1$. We reason by case analysis on the behaviour of s (up-to commutativity of internal choice).

- If we have

$$S \equiv s[r!a;P_2 \oplus r!b;Q_2] \mid S'$$

From the judgement in (3.6) and rule $[\oplus]$, the two judgements below must be derivable

$$\Gamma \vdash r[P + Q] \mid s[r!a;P_2] \mid S' \blacktriangleright \mathcal{G} \quad \Gamma \vdash r[P + Q] \mid s[r!b;Q_2] \mid S' \blacktriangleright \mathcal{G}$$

and we should also have

$$\Gamma \vdash r[s?a;P_1] \mid s[r!a;P_2] \mid S' \blacktriangleright \mathcal{G} \quad \Gamma \vdash r[s?a;P_1] \mid s[r!b;Q_2] \mid S' \blacktriangleright \mathcal{G}$$

but the judgement on the right is not derivable since the prefix of the behaviours

of r and s do not match; and we have reached a contradiction. A dual reasoning shows that the Q branch is also not derivable.

- If we have

$$S \equiv s[r!b; Q_2] | S'$$

then it contradicts the fact that $\Gamma \vdash r[P] | S \triangleright \mathcal{G}$ is derivable since the prefix $r!b$ will never be eliminated. Recall that all the message sorts guarding an external choice must be distinct, thus no other branch than Q will permit to eliminate the prefix $r!b$.

- If we have

$$S \equiv s[r!a; P_2 \oplus \dots r!b; Q_2] | S'$$

where the ellipsis indicates a series of actions that are not towards r , i.e., $r!b$ is the first action towards r in the right-hand branch of the internal choice.

This also contradicts the fact that $\Gamma \vdash r[P] | S \triangleright \mathcal{G}$ is derivable since after all the prefixes in the right-hand branch of the internal choice have been eliminated, we return to the previous case. \square

Theorem 3.3 (Unique typing). *If $\Gamma \vdash S \triangleright \mathcal{G}$ and $\Gamma \vdash S \triangleright \mathcal{G}'$ then $\mathcal{G} \equiv \mathcal{G}'$.*

Proof. We show that each time a rule from Figure 3.3 is applicable, either no other rule is applicable, or the typing produces an equivalent global type.

- Due to their syntactic restrictions and the condition $S \not\equiv \chi$, the cases for rules $[\cdot]$, $[\mu]$, $[\chi]$ and $[\emptyset]$ are straightforward.
- The cases for rules $[\oplus]$ and $[\cdot]$ follow naturally from the structural congruence rules of both systems and global types.
- The case for rule $[\equiv]$ follows from the fact that associativity and commutativity in S do not affect \mathcal{G} . In addition, if one unfolds a behaviour once more than required, we have the result since $\mu\chi.\mathcal{G} \equiv \mathcal{G} \{\mu\chi.\mathcal{G}/\chi\}$.

- For the case of rule $[\text{+}]$, we have to show that we cannot type a system with either branch of an external choice, i.e., only one branch allows the system to be typable.

We show this case by contradiction, assume $\mathcal{G}_1 \neq \mathcal{G}_2$,

$$\frac{\frac{\vdots}{\Gamma \vdash r[P] \mid S \blacktriangleright \mathcal{G}_1}}{[\text{+}] \Gamma \vdash r[P+Q] \mid S \blacktriangleright \mathcal{G}_1} \quad \text{and} \quad \frac{\frac{\vdots}{\Gamma \vdash r[Q] \mid S \blacktriangleright \mathcal{G}_2}}{[\text{+}] \Gamma \vdash r[P+Q] \mid S \blacktriangleright \mathcal{G}_2}$$

By assumption, we have $\Gamma \vdash r[P+Q] \mid S \blacktriangleright \mathcal{G}_1$ or $\Gamma \vdash r[P+Q] \mid S \blacktriangleright \mathcal{G}_2$ and, by Lemma 3.6, we have a contradiction. \square

3.5.3 Well-formedness and Projections

We characterise the relationships between the behaviours in a typable system and the projections of the global types that is assigned to it. In addition, we show that a synthesised global type is well-formed and projectable.

The following lemma shows a relationship between the ready set of a system and the ready set of its global type. We will use this lemma in several occasions.

Lemma 3.7. *If $\circ \vdash S \blacktriangleright \mathcal{G}$ and \mathcal{G} is well-formed and projectable, then*

$$\mathcal{G} \equiv (\mathfrak{s} \rightarrow r : a; \mathcal{G}_1 + \mathcal{G}_2) \mid \mathcal{G}_3 \iff S \equiv \mathfrak{s}[r!a; P \oplus P'] \mid r[\mathfrak{s}?a; Q + Q'] \mid S'$$

Proof. (\Rightarrow) Assume that

$$\circ \vdash S \blacktriangleright \mathcal{G} \quad \text{such that} \quad \mathcal{G} \equiv (\mathfrak{s} \rightarrow r : a; \mathcal{G}_1 + \mathcal{G}_2) \mid \mathcal{G}_3$$

is derivable. We show that either a rule introducing the corresponding operator is applicable or that an equivalent \mathcal{G} can be inferred.

Assume $\mathcal{G} = \mathcal{G}' \mid \mathcal{G}_3$, if $\circ \vdash S \blacktriangleright \mathcal{G}$ is derivable then we must have either $S \equiv S_1 \mid S_2$ such that

$$\circ \vdash S_1 \blacktriangleright \mathcal{G}' \quad \text{and} \quad \circ \vdash S_2 \blacktriangleright \mathcal{G}_3$$

are derivable (by rule $[\mid]$), or $S_3 \equiv \mathbf{0}$ and $\mathcal{G}_3 \equiv \mathbf{0}$. Observe that we must have $\mathcal{P}(S_1) \cap \mathcal{P}(S_2) = \emptyset$, by Lemma 3.4 and the fact that \mathcal{G} is projectable.

Now, take $\mathcal{G}' = \mathcal{G}'' + \mathcal{G}_2$, we must have either

$$S_1 = s[P_1 \oplus P'_1] \mid S'_1$$

such that

$$\circ \vdash s[P_1] \mid S'_1 \blacktriangleright \mathcal{G}'' \quad \text{and} \quad \circ \vdash n[P'_1] \mid S'_1 \blacktriangleright \mathcal{G}_2$$

are derivable (by rules $[\oplus]$ and $S'_1 \not\downarrow$), or $P'_1 \equiv \mathbf{0}$ and $\mathcal{G}_2 \equiv \mathbf{0}$.

If we have $\mathcal{G}'' = s \rightarrow r : a; \mathcal{G}_1$, we must have

$$P_1 \equiv r!a; P \quad \text{and} \quad S'_1 \equiv r[s?a; Q + Q'] \mid S''_1$$

and

$$\circ \vdash s[r!a; P] \mid r[s?a; Q + Q'] \mid S''_1 \blacktriangleright s \rightarrow r : a; \mathcal{G}_1$$

derivable.

Finally, considering the systems S_1 and S_2 , and the processes of s and r , we have the required result:

$$S \equiv s[r!a; P \oplus P'] \mid r[s?a; Q + Q'] \mid S'$$

(\Leftarrow) Assume

$$\circ \vdash s[r!a; P \oplus P'] \mid r[s?a; Q + Q'] \mid S' \blacktriangleright \mathcal{G}$$

We show that either a rule introducing the corresponding operator is applicable or that an

equivalent \mathcal{G} can be inferred. Either there is S_1 and S_2 such that $S' \equiv S_1 \mid S_2$, and

$$\circ \vdash s[r!a; P \oplus P'] \mid r[s?a; Q + Q'] \mid S_1 \blacktriangleright \mathcal{G}' \quad \text{and} \quad \circ \vdash S_2 \blacktriangleright \mathcal{G}_3$$

or $\mathcal{G}_3 \equiv \mathbf{0}$. Note that we can assume that $S_1 \not\downarrow$ since one could apply rule $[\mid]$ as many times as necessary until it holds.

Now, either there is $\mathcal{G}' \equiv \mathcal{G}'' + \mathcal{G}_2$ such that

$$\circ \vdash s[r!a; P] \mid r[s?a; Q + Q'] \mid S_1 \blacktriangleright \mathcal{G}''$$

and

$$\circ \vdash s[P'] \mid r[s?a; Q + Q'] \mid S_1 \blacktriangleright \mathcal{G}_2$$

or $\mathcal{G}_2 = \mathbf{0}$.

Finally, for the judgement

$$\circ \vdash s[r!a; P] \mid r[s?a; Q + Q'] \mid S_1 \blacktriangleright \mathcal{G}''$$

to be derivable, we must have $\mathcal{G}'' \equiv s \rightarrow r : a ; \mathcal{G}_1$. Considering \mathcal{G}' and \mathcal{G}'' , we have the required result. \square

We show that there is a correspondence between the behaviours of the original system and the projections of its synthesised global type. The behaviour of a participant in S is a simulation of the projection of a synthesised global type from S onto the participant. Intuitively, the other direction is lost due to rule $[\mid]$, indeed external choice branches which are never chosen are not “recorded” in the synthesised global type.

Example 3.12. Consider the following system

$$S \equiv s[r!a] \mid r[\mu\chi.(s?b; \chi + s?a)]$$

We show that S has global type $s \rightarrow r : a$, via the derivation below.

$$\begin{array}{c}
 \text{[0]} \frac{}{\circ \vdash s[\mathbf{0}] \mid r[\mathbf{0}] \triangleright \mathbf{0}} \\
 \text{[!]} \frac{}{\circ \vdash s[r!a] \mid r[s?a] \triangleright s \rightarrow r : a} \\
 \text{[+]} \frac{}{\circ \vdash s[r!a] \mid r[s?b; (\mu\chi.(s?b; \chi + s?a)) + s?a] \triangleright s \rightarrow r : a} \\
 \text{[≡]} \frac{}{\circ \vdash S \triangleright s \rightarrow r : a}
 \end{array}$$

◇

Note that, given a participant with an external choice, it must be the case that at least one branch of the external choice is matched by a corresponding sending action. For instance, in Example 3.12, we can only eliminate the prefix $s?a$ via rule [!].

We now give our notion of simulation, which preserves sending actions in both directions; it requires all the receiving actions to be preserved in one direction and the *existence* of a matching receiving action in the other direction.

Definition 3.8 (\lesssim). $P \lesssim Q$ if and only if

- $P \xrightarrow{r!a} P'$ implies that there is Q' such that $Q \xrightarrow{r!a} Q'$ and $P' \lesssim Q'$,
- $Q \xrightarrow{r!a} Q'$ implies that there is P' such that $P \xrightarrow{r!a} P'$ and $P' \lesssim Q'$,
- $P \xrightarrow{s?a} P'$ implies that there is Q' such that $Q \xrightarrow{s?a} Q'$ and $P' \lesssim Q'$, and
- $Q \xrightarrow{s?a} Q'$ implies that there is b , Q'' , and P' such that $Q \xrightarrow{s?b} Q''$, $P \xrightarrow{s?b} P'$, and $P' \lesssim Q''$,

Definition 3.8 implies naturally the following proposition.

Proposition 3.1. *Let Q be a process. If $\mathbf{0} \lesssim Q$ then $Q \equiv \mathbf{0}$.*

Proof. By contradiction. If $Q = \bigoplus_{i \in I} r_i !a_i; Q_i$, then we reach a contradiction since $\mathbf{0}$ cannot match any of the sending actions. If $Q = \sum_{i \in I} s ?a_i; Q_i$, then there should exist at least one action that is matched by $\mathbf{0}$, which leads to a contradiction. \square

Lemma 3.8. *If $\circ \vdash n[P] \mid S \blacktriangleright \mathcal{G}$ then $\mathcal{G} \downarrow_n \lesssim P$.*

Proof. The result follows directly from following intermediary statement (since we consider only closed behaviours).

Let $S' \equiv n[P] \mid S$, \mathcal{G} , and Γ such that

- $\forall n' \in \mathcal{P}(S') : \forall \chi \in \text{fv}(S'(n')) : \exists \chi' : \Gamma(n', \chi) = \chi'$
- $\forall \chi \in \text{fv}(\mathcal{G}) : \forall n' \in \mathcal{P}(S') : \exists \chi' : \Gamma(n', \chi') = \chi$

if $\Gamma \vdash n[P] \mid S \blacktriangleright \mathcal{G}$ then $\mathcal{G} \downarrow_n \lesssim P$.

The proof is by case analysis on the actions enabled for P and $\mathcal{G} \downarrow_n$ (we treat the recursive cases separately), then by induction on the structure of the derivation.

P sends. If $P \xrightarrow{r!a} P'$, then we must have $P \equiv r!a; P' \oplus P''$. Assume we have $S \equiv r[n?a; Q + Q'] \mid S''$ and $S'' \not\downarrow$, if it is not the case, we may apply other rules until we obtain such a system. Indeed, the derivation may not affect the process of n until the system cannot synchronise without n .²

$$\begin{array}{c}
 \vdots \\
 \hline
 \Gamma \vdash n[P'] \mid r[Q] \mid S'' \blacktriangleright \mathcal{G}_1 \\
 \text{[!]} \hline
 \Gamma \vdash n[r!a; P'] \mid r[n?a; Q] \mid S'' \blacktriangleright n \rightarrow r : a; \mathcal{G}_1 \qquad \vdots \\
 \text{[+]} \hline
 \Gamma \vdash n[r!a; P'] \mid r[n?a; Q + Q'] \mid S'' \blacktriangleright n \rightarrow r : a; \mathcal{G}_1 \qquad \Gamma \vdash n[P''] \mid r[n?a; Q + Q'] \mid S'' \blacktriangleright \mathcal{G}_2 \\
 \text{[\oplus]} \hline
 \Gamma \vdash n[r!a; P' \oplus P''] \mid r[n?a; Q + Q'] \mid S'' \blacktriangleright n \rightarrow r : a; \mathcal{G}_1 + \mathcal{G}_2
 \end{array}$$

Thus, we have $\mathcal{G} \downarrow_n \xrightarrow{r!a} \mathcal{G}_1 \downarrow_n$, and, by induction hypothesis, we have

$$\mathcal{G}_1 \downarrow_n \lesssim P' \quad \text{and} \quad \mathcal{G}_2 \downarrow_n \lesssim P''$$

² Except for rule [!] , but it may only fold or unfold the process.

as required.

P receives. If $P \xrightarrow{s?a} P'$, then we must have $P \equiv s?a; P' + P''$. There are two cases: either $P'' \equiv \mathbf{0}$ or $P'' \not\equiv \mathbf{0}$.

- If $P'' \equiv \mathbf{0}$, then we must have a matching action for $s?a$ (we cannot apply rule $_{[+]}$). Assume we have $S \equiv s[n!a; Q] \mid S''$ and $S'' \not\downarrow$, if it is not the case, we may apply other rules until we obtain such a system (as above). We have the following derivation:

$$\begin{array}{c} \vdots \\ \hline \Gamma \vdash n[P'] \mid s[Q] \mid S'' \triangleright \mathcal{G}' \\ \text{[;]} \hline \Gamma \vdash n[s?a; P'] \mid s[n!a; Q] \mid S'' \triangleright s \rightarrow n : a; \mathcal{G}' \end{array}$$

Thus, we have $\mathcal{G} \downarrow_n \xrightarrow{s?a} \mathcal{G}' \downarrow_n$, and, by induction hypothesis, we have $\mathcal{G}' \downarrow_n \lesssim P'$ as required.

- If $P'' \not\equiv \mathbf{0}$ then there are two sub-cases:
 - Either there is matching action for $s?a$, in which case we obtain the result by discarding the P'' branch with rule $_{[+]}$, then applying the same reasoning as above to apply rule $_{[;]}$; or
 - there is not matching action for $s?a$, in which case we can discard the $s?a; P'$ branch with rule $_{[+]}$. Since $P'' \not\equiv \mathbf{0}$ and the system is typable, there must be one branch of n that will be recorded in \mathcal{G} , i.e., one cannot discard the “last” branch of an external choice (cf. the case where $P'' \equiv \mathbf{0}$). The result follows from the fact that *there exists* one branch of the external choice that will be “recorded” in \mathcal{G} .

$\mathcal{G} \downarrow_n$ sends. If $\mathcal{G} \downarrow_n \xrightarrow{r!a} \mathcal{G}_1 \downarrow_n$, then we must have $\mathcal{G} \downarrow_n \equiv r!a; \mathcal{G}_1 \downarrow_n \oplus \mathcal{G}_2 \downarrow_n$. We must have the following sub-derivation of $\Gamma \vdash n[P] \mid S \triangleright \mathcal{G}$ such that $P \equiv r!a; P' \oplus P''$ where

all the previous steps of the derivation did not affect $n[P]$ (otherwise we would not have $\mathcal{G} \downarrow_n \equiv r!a; \mathcal{G}_1 \downarrow_n \oplus \mathcal{G}_2 \downarrow_n$). Note that we apply rules $[\oplus]$ and $[+]$ in one step, for simplicity.

$$\begin{array}{c}
 \vdots \\
 \hline
 \Gamma \vdash n[P'] \mid r[Q] \mid S'' \triangleright \mathcal{G}_1 \qquad \qquad \qquad \vdots \\
 \text{[;]} \frac{\Gamma \vdash n[r!a; P'] \mid r[s?a; Q] \mid S'' \triangleright n \rightarrow r:a; \mathcal{G}_1 \quad \Gamma \vdash n[P''] \mid r[Q'] \mid S'' \triangleright \mathcal{G}_2}{\text{[\oplus][+]} \Gamma \vdash n[r!a; P' \oplus P''] \mid r[s?a; Q + Q'] \mid S'' \triangleright n \rightarrow r:a; \mathcal{G}_1 + \mathcal{G}_2}
 \end{array}$$

We have $P \xrightarrow{r!a} P'$, and, by induction hypothesis, we have

$$\mathcal{G}_1 \downarrow_n \lesssim P' \quad \text{and} \quad \mathcal{G}_2 \downarrow_n \lesssim P''$$

as required.

$\mathcal{G} \downarrow_n$ **receives**. If $\mathcal{G} \downarrow_n \xrightarrow{s?a} \mathcal{G}_1 \downarrow_n$, then we must have $\mathcal{G} \downarrow_n \equiv s?a; \mathcal{G}_1 \downarrow_n$.³ We must have the following sub-derivation of $\Gamma \vdash n[P] \mid S \triangleright \mathcal{G}$ such that $P \equiv s?a; P' + P''$ where all the previous steps of the derivation did not affect $n[P]$ (otherwise we would not have $\mathcal{G} \downarrow_n \equiv s?a; \mathcal{G}_1 \downarrow_n$).

$$\begin{array}{c}
 \vdots \\
 \hline
 \Gamma \vdash n[P'] \mid s[Q] \mid S'' \triangleright \mathcal{G}_1 \\
 \text{[;]} \frac{\Gamma \vdash n[s?a; P'] \mid s[n!a; Q] \mid S'' \triangleright s \rightarrow n:a; \mathcal{G}_1}{\text{[+]} \Gamma \vdash n[s?a; P' + P''] \mid s[n!a; Q] \mid S'' \triangleright s \rightarrow n:a; \mathcal{G}_1}
 \end{array}$$

We have $P \xrightarrow{s?a} P'$, and, by induction hypothesis, we have $\mathcal{G}_1 \downarrow_n \lesssim P'$ as required.

Recursion in projection. If $\mathcal{G} \downarrow_n \equiv \mu\chi.\mathcal{G}_1 \downarrow_n$, then we must have used rule $[\mu]$, thus we must have $P \equiv \mu\chi.P'$, and $\forall m \in \mathcal{P}(S) : S(m) \equiv \mu\chi.Q$ and the rest follows by induction

³ We abstract from possible external choice.

hypothesis and the fact that all participants must be added in Γ together with their local recursion variable.

Recursion in local type. If $P \equiv \mu\chi.P'$, then there are two cases either rule $[\mu]$ is the last rule applied, in which case, the result follows as above. Otherwise $[\equiv]$ is the last rule that has been applied, and we must have the following sub-derivation:

$$\frac{\Gamma \vdash_{\mathfrak{n}} [P' \{\mu\chi.P'/\chi\}] \mid S \triangleright \mathcal{G}}{[\equiv] \Gamma \vdash_{\mathfrak{n}} [\mu\chi.P'] \mid S \triangleright \mathcal{G}}$$

By induction hypothesis, we have $\mathcal{G} \downarrow_{\mathfrak{n}} \lesssim P' \{\mu\chi.P'/\chi\}$, and the result follows from the rule

$$P' \{\mu\chi.P'/\chi\} \equiv \mu\chi.P' \quad \square$$

Theorem 3.4 (Well-formedness). *If $\Gamma \vdash S \triangleright \mathcal{G}$ then $\vdash \mathcal{G}$ and \mathcal{G} is projectable.*

Proof. The proof is by induction on the derivation $\Gamma \vdash S \triangleright \mathcal{G}$. We make a case analysis on the last rule used.

Case $[\chi]$. Trivial by Definition 3.1.

Cases $[\circ]$: Trivial by Definition 3.1.

Cases $[\equiv]$: Trivial by induction hypothesis.

Case $[\cdot]$. We have $\Gamma \vdash s[P] \mid r[Q] \mid S' \triangleright \mathcal{G}'$,

$$\mathcal{G} = s \rightarrow r : a ; \mathcal{G}', \quad S = s[r!a; P] \mid r[s?a; Q] \mid S', \quad \text{and} \quad S' \not\Downarrow$$

- *WF.* We show that we have

$$\forall n_1 \rightarrow n_2 : _ \in R(\mathcal{G}') : \{s, r\} \cap \{n_1, n_2\} \neq \emptyset$$

by contradiction. By induction hypothesis, we know that

$$\Gamma \vdash s[P] \mid r[Q] \mid S' \triangleright \mathcal{G}'$$

If we had $\mathcal{G}' \equiv (n_1 \rightarrow n_2 : b; \mathcal{G}_0 + \mathcal{G}_1) \mid \mathcal{G}_2$ with $n_i \neq s$ and $n_i \neq r$ and $i \in \{1, 2\}$, by Lemma 3.7, we would have

$$S' \equiv n_1[n_2!b; P'_0 \oplus P'_1] \mid n_2[n_1?b; Q'_0 + Q'_1] \mid S''$$

which is in contradiction with the premise $S' \not\Downarrow$.

- *Projection.* By induction hypothesis, we know that $\mathcal{G}' \downarrow_n$ is defined for all $n \in \mathcal{P}(\mathcal{G}')$. By Definition 3.1, we have that $\mathcal{G} \downarrow_s = r!a; \mathcal{G}' \downarrow_s$, $\mathcal{G} \downarrow_r = s?a; \mathcal{G}' \downarrow_r$, and $\mathcal{G} \downarrow_n = \mathcal{G}' \downarrow_n$, for $s \neq n \neq r$.

Case $[\oplus]$. We have

$$\mathcal{G} = \mathcal{G}_0 + \mathcal{G}_1, \quad S = s[P \oplus Q] \mid S' \quad \text{and} \quad S' \not\Downarrow$$

- *WF.* We have to show that

$$\forall s \rightarrow r : a \in R(\mathcal{G}_0) : \forall s' \rightarrow r' : b \in R(\mathcal{G}_1) : s = s' \wedge (r, a) \neq (r', b)$$

$(r, a) \neq (r', b)$ follows directly from the syntax of processes, i.e., we have

$$P \equiv \bigoplus_{i \in I} r_i!a_i; P_i \quad \text{and} \quad Q \equiv \bigoplus_{j \in J} r_j!a_j; P_j \quad \text{and} \quad I \cap J = \emptyset$$

$$\forall i, j \in I \cup J : i \neq j \implies (r_i, a_i) \neq (r_j, a_j)$$

We have to show that for all prefixes in \mathcal{G}_0 and \mathcal{G}_1 , s is the sender. In other words,

the participant who makes the internal choice at the global type level must be the same in all the branches. By contradiction, assume we have

$$\mathcal{G}_0 = s \rightarrow r:a + s \rightarrow r:b \quad \text{and} \quad \mathcal{G}_1 = s' \rightarrow r'':b' \quad \text{and} \quad s \neq s'$$

then we must have a system of the form

$$S \equiv s[r!a \oplus r'!b] \mid s'[r''!b'] \mid r[s?a + s?b] \mid r''[s'?b'] \mid S''$$

which is in contradiction with the premise $S' \not\downarrow$.

By the result above and $\vdash \mathcal{G}_0$ and $\vdash \mathcal{G}_1$ by induction hypothesis, we have $\vdash \mathcal{G}$.

- *Projection.* If n is the sender, we have the result directly by induction hypothesis, i.e., both $\mathcal{G}_0 \downarrow_n$ and $\mathcal{G}_1 \downarrow_n$ are defined by induction hypothesis, thus so is $\mathcal{G}_0 \downarrow_n \oplus \mathcal{G}_1 \downarrow_n$.

For all n not the sender in \mathcal{G}_0 and \mathcal{G}_1 , we have to show that, knowing by induction hypothesis that both $\mathcal{G}_0 \downarrow_n$ and $\mathcal{G}_1 \downarrow_n$ are defined,

$$\mathcal{G}_0 \downarrow_n \uplus \mathcal{G}_1 \downarrow_n \neq \perp$$

We show this by contradiction, if $_ \uplus _$ is undefined, it may only originate from the following cases, by Definition 3.2 (we consider only the prefixes of projections, without loss of generality).

Send. If we have⁴

$$\mathcal{G}_0 \downarrow_n \equiv r!a \quad \text{and} \quad \mathcal{G}_1 \downarrow_n \equiv r!b$$

Take the system of the form below from which the projections above must originate,

⁴This case corresponds to the case where the behaviour of n is different in two branches, while being unaware of which branch was chosen: n has to send either a or b depending on which branch was chosen by s – without knowing which branch was selected.

by Lemma 3.8 and definition of projections.

$$S \equiv s[r!a \oplus r!b] \mid r[s?a; n!a + s?b; n!b] \mid n[r?a + r?b]$$

We must have the following derivation

$$\begin{array}{c}
\begin{array}{c}
\frac{[0] \text{-----}}{\Gamma \vdash s[0] \mid r[0] \mid n[0] \blacktriangleright \mathbf{0}} \\
\frac{[;] \text{-----}}{\Gamma \vdash s[0] \mid r[n!a] \mid n[r!a] \blacktriangleright \mathcal{G}'_0} \\
\frac{[+] \text{-----}}{\Gamma \vdash s[0] \mid r[n!a] \mid n[\dots] \blacktriangleright \mathcal{G}'_0} \\
\frac{[;] \text{-----}}{\Gamma \vdash s[r!a] \mid r[s?a; n!a] \mid n[\dots] \blacktriangleright \mathcal{G}_0} \\
\frac{[+] \text{-----}}{\Gamma \vdash s[r!a] \mid r[\dots] \mid n[\dots] \blacktriangleright \mathcal{G}_0} \\
\frac{[\oplus] \text{-----}}{\Gamma \vdash S \blacktriangleright \mathcal{G}}
\end{array}
\qquad
\begin{array}{c}
\frac{[0] \text{-----}}{\Gamma \vdash s[0] \mid r[0] \mid n[0] \blacktriangleright \mathbf{0}} \\
\frac{[;] \text{-----}}{\Gamma \vdash s[0] \mid r[n!b] \mid n[r!b] \blacktriangleright \mathcal{G}'_1} \\
\frac{[+] \text{-----}}{\Gamma \vdash s[0] \mid r[n!b] \mid n[\dots] \blacktriangleright \mathcal{G}'_1} \\
\frac{[;] \text{-----}}{\Gamma \vdash s[r!b] \mid r[s?b; n!b] \mid n[\dots] \blacktriangleright \mathcal{G}_1} \\
\frac{[+] \text{-----}}{\Gamma \vdash s[r!b] \mid r[\dots] \mid n[\dots] \blacktriangleright \mathcal{G}_1}
\end{array}
\end{array}$$

where

$$\mathcal{G}_0 = s \rightarrow r : a; \mathcal{G}'_0 \quad \text{and} \quad \mathcal{G}'_0 = r \rightarrow n : a; \mathbf{0}$$

and

$$\mathcal{G}_1 = s \rightarrow r : b; \mathcal{G}'_1 \quad \text{and} \quad \mathcal{G}'_1 = r \rightarrow n : b; \mathbf{0}$$

Clearly, we have that $\mathcal{G}_0 \downarrow_n \neq r?a + r?b$, which is a contradiction.

Send-Receive. The case where one projection is an internal choice, while the other is an external choice, is discarded by the syntax of processes.

Case [:].

- *WF.* By induction hypothesis, we have $\vdash \mathcal{G}$.
- *Projection.* By induction hypothesis.

Case $[|]$. We have

$$\mathcal{G} = \mathcal{G}_0 \mid \mathcal{G}_1$$

- *WF.* We have to show that

$$\mathcal{P}(\mathcal{G}_0) \cap \mathcal{P}(\mathcal{G}_1) = \emptyset$$

By definition of systems, we know that there cannot be two participants with the same name in a system, since $S \mid S'$ is a system we have that

$$\mathcal{P}(S) \cap \mathcal{P}(S') = \emptyset$$

and by Lemma 3.4, we have $\mathcal{P}(\mathcal{G}_0) \cap \mathcal{P}(\mathcal{G}_1) = \emptyset$.

- *Projection.* For all $n \in \mathcal{P}(\mathcal{G})$, $\mathcal{G} \downarrow_n$ is defined by induction hypothesis.

Case $[\mu]$. We have $\mathcal{G} = \mu\chi. \mathcal{G}'$,

$$S = n_1[\mu X_1.P_1] \mid \dots \mid n_k[\mu X_k.P_k] \quad \text{and} \quad \exists 1 \leq i, j \leq k : (n_i[P_i] \mid n_j[P_j])\Downarrow$$

- *WF.* We have to show that

$$\chi \in \text{fv}(\mathcal{G}') \Rightarrow |\text{Indep}(\mathcal{G}')| = 1$$

which follows from the fact that the context Γ is emptied each time the rule $[|]$ is used in the derivation (this rule is the only one introducing concurrent branches). In addition, for the axiom $[\chi]$ to be used in the derivation one must have $(-, -) : \chi \in \Gamma$. Therefore, the only way one could have $|\text{Indep}(\mathcal{G}')| > 1$ (i.e., at least two concurrent branches in \mathcal{G}) is if χ does not appear in \mathcal{G} .

Observe that the recursion is prefix guarded since we have $(n_i[P_i] \mid n_j[P_j])\Downarrow$.

- *Projection.* By induction hypothesis. □

3.5.4 Subject Reduction

In order to have a form of subject reduction for our type system, it is slightly extended to support (stable) runtime systems. Recall that a stable runtime system is such that all its queues are empty. First the rule

$$\text{[}\varepsilon\text{]} \frac{\Gamma \vdash S \triangleright \mathcal{G}}{\Gamma \vdash \text{sr} : \varepsilon \mid S \triangleright \mathcal{G}}$$

is added so to allow empty queues to be discarded. Second, the premises of rule [0] are updated to obtain

$$\text{[0]} \frac{\forall n \in \mathcal{P}(S) : S(n) = \mathbf{0} \quad C(S) = \emptyset}{\Gamma \vdash S \triangleright \mathbf{0}}$$

where $C(S) = \emptyset$ means that there is no queue left in S , i.e., S is a program where all the participants have an empty behaviour.

The lemma below follows directly from the extension above.

Lemma 3.9. *Let S be a program,*

$$\circ \vdash S \triangleright \mathcal{G} \iff \circ \vdash S \mid Q(S) \triangleright \mathcal{G}$$

Proof. (\Rightarrow) If $\circ \vdash S \triangleright \mathcal{G}$, then we can apply rule $\text{[}\varepsilon\text{]}$ on $S \mid Q(S)$ until all the queues have been removed and we obtain $\circ \vdash S \triangleright \mathcal{G}$ as a sub-derivation. (\Leftarrow) If $\circ \vdash S \mid Q(S) \triangleright \mathcal{G}$, then $\circ \vdash S \triangleright \mathcal{G}$ is a sub-derivation of the runtime system. \square

Theorem 3.5 (Subject reduction). *Let S be a stable runtime system, if $\circ \vdash S \triangleright \mathcal{G}$ and $S \xrightarrow{s \rightarrow r : a} \xrightarrow{r \leftarrow s : a} S'$, then $\circ \vdash S' \triangleright \mathcal{G}'$.*

In addition, if $S \equiv s[r!a; P \oplus P'] \mid r[s?a; Q + Q'] \mid S''$, then

$$\mathcal{G} \equiv (s \rightarrow r:a; \mathcal{G}_1 + \mathcal{G}_2) \mid \mathcal{G}_3 \quad S' \equiv s[P] \mid r[Q] \mid S'' \quad \mathcal{G}' \equiv \mathcal{G}_1 \mid \mathcal{G}_3$$

Proof. Assume S is a stable runtime system such that $S \xrightarrow{s \rightarrow r:a} \xrightarrow{r \leftarrow s:a} S'$, from the semantics of systems and processes, it is clear that we must have

$$S \equiv s[r!a; P \oplus P'] \mid r[s?a; Q + Q'] \mid S_1$$

Since $\circ \vdash S \blacktriangleright \mathcal{G}$, we have

$$\mathcal{G} \equiv (s \rightarrow r:a; \mathcal{G}_1 + \mathcal{G}_2) \mid \mathcal{G}_3$$

by Lemma 3.7 and the fact that by Theorem 3.4, \mathcal{G} is well-formed and projectable.

By a similar reasoning to the one of the proof of Lemma 3.7, we have the following derivation, with $S_1 \equiv S_2 \mid S_3$ and $S_2 \not\downarrow$.

$$\begin{array}{c} \vdots \\ \hline \nabla \quad \circ \vdash S_3 \blacktriangleright \mathcal{G}_3 \\ \hline \text{[!]} \quad \circ \vdash s[r!a; P \oplus P'] \mid r[s?a; Q + Q'] \mid S_2 \mid S_3 \blacktriangleright \mathcal{G} \end{array} \quad (3.7)$$

where ∇ is as follows

$$\begin{array}{c} \vdots \\ \hline \circ \vdash s[P] \mid r[Q] \mid S_2 \blacktriangleright \mathcal{G}_1 \\ \hline \text{[;]} \quad \circ \vdash s[r!a; P] \mid r[s?a; Q] \mid S_2 \blacktriangleright s \rightarrow r:a; \mathcal{G}_1 \quad \vdots \\ \hline \text{[+]} \quad \circ \vdash s[r!a; P] \mid r[s?a; Q + Q'] \mid S_2 \blacktriangleright s \rightarrow r:a; \mathcal{G}_1 \quad \circ \vdash s[P'] \mid r[s?a; Q + Q'] \mid S_2 \blacktriangleright \mathcal{G}_2 \\ \hline \text{[\oplus]} \quad \circ \vdash s[r!a; P \oplus P'] \mid r[s?a; Q + Q'] \mid S_2 \blacktriangleright s \rightarrow r:a; \mathcal{G}_1 + \mathcal{G}_2 \end{array}$$

From the semantics of systems, we must have

$$S' \equiv s[P] \mid r[Q] \mid S_1$$

and since $S_1 \equiv S_2 \mid S_3$ is not changed by reduction, we have the derivation:

$$\frac{\begin{array}{c} \vdots \\ \hline \circ \vdash s[P] \mid r[Q] \mid S_2 \triangleright \mathcal{G}_1 \end{array} \quad \begin{array}{c} \vdots \\ \hline \circ \vdash S_3 \triangleright \mathcal{G}_3 \end{array}}{\circ \vdash s[P] \mid r[Q] \mid S_2 \mid S_3 \triangleright \mathcal{G}'} \quad (3.8)$$

and the result follows from the fact that each sub-derivation of (3.8) is a sub-derivation of (3.7). \square

Corollary 3.1. *Let S be a stable runtime system such that $\circ \vdash S \triangleright \mathcal{G}$. If S' is a stable runtime system such that S' is reachable from S with a 1-buffer execution, then $\circ \vdash S' \triangleright \mathcal{G}'$.*

Proof. Direct from Theorem 3.5 and the fact that closeness of behaviours is preserved by reduction. \square

Lemma 3.10 below motivates the restriction of our subject reduction theorem to synchronisation of output/input. Indeed, this result shows that, for typable systems, there is always an execution for which each output is followed by its corresponding input.

Lemma 3.10. *If S is a stable runtime system such that $\circ \vdash S \triangleright \mathcal{G}$ and S' is a stable runtime system such that $S \Longrightarrow S'$, then S' is reachable by a 1-buffer execution from S .*

Proof. First we note that we have the following dependencies, by the semantics of sys-

tems. Let us define the following accessory function on labels:

$$\text{subj}(\lambda) \stackrel{\text{def}}{=} \begin{cases} s, & \text{if } \lambda = s \rightarrow r : a \\ r, & \text{if } \lambda = r \leftarrow s : a \end{cases}$$

A label λ' depends on λ iff

- $\lambda = s \rightarrow r : a$ and $\lambda' = r \leftarrow s : a$, or
- $\text{subj}(\lambda) = \text{subj}(\lambda')$

We write $\lambda < \lambda'$ if λ' depends on λ , and we call $<$ -causality chain a list $\lambda_1 \dots \lambda_n$ such that for all $1 \leq i < n$, $\lambda_i < \lambda_{i+1}$.

Assume, by contradiction, that we have S and S' such that $S \Longrightarrow S'$, and S' is *not* reachable by a 1-buffer execution from S . That is to say that at least two buffers are needed for S to reduce to S' . We have the following situation

$$S \xrightarrow{s \rightarrow r : a} \xrightarrow{\phi} \xrightarrow{r \leftarrow s : a} S'$$

where ϕ is a non-empty sequence of labels λ alternating between outputs and inputs (since S and S' are stable) and there is no 1-buffer execution (i.e., ϕ must include another label $s' \rightarrow r' : b$). There must be a $<$ -causality chain between the labels in ϕ and $s \rightarrow r : a$ (resp. $r \leftarrow s : a$), otherwise we could swap labels so that $r \leftarrow s : a$ follows directly $s \rightarrow r : a$. The smallest chain satisfying these requirements is $\phi = r \rightarrow s : b \cdot s \leftarrow r : b$. Such a sequence of transitions implies that we have the following system:

$$S \equiv s[r!a; r?a] \mid r[s!b; s?a] \mid S''$$

which is not typable, and thus leads to a contradiction. □

3.5.5 Equivalence with Original System

Since the branches of external choices that are discarded while typing are those that are never chosen, we can show that there is a bisimulation between the original system and the system consisting of the projections of the synthesised global type.

We first give our notion of bisimulation, then our main result (Theorem 3.6).

Definition 3.9 (\approx). $S \approx S'$ if and only if

- $S \xrightarrow{\lambda} S_1$ implies $S' \xrightarrow{\lambda} S_2$ for some $S_1 \approx S_2$, and
- $S' \xrightarrow{\lambda} S_2$ implies $S \xrightarrow{\lambda} S_1$ for some $S_1 \approx S_2$

where $\lambda \in \{s \rightarrow r : a, r \leftarrow s : a\}$.

Theorem 3.6. Let S be a program, if $\circ \vdash S \blacktriangleright \mathcal{G}$, then

$$\bigsqcup_{n \in \mathcal{P}(S)} n[\mathcal{G} \downarrow_n] \mid Q(S) \approx S \mid Q(S)$$

Proof. Let $\hat{S} \equiv \bigsqcup_{n \in \mathcal{P}(S)} n[\mathcal{G} \downarrow_n]$. From Lemmas 3.4 and 3.8, we have that

$$S \equiv n_1[P_1] \mid \dots \mid n_k[P_k]$$

such that for all $1 \leq i \leq k$: $\mathcal{G} \downarrow_i \lesssim P_i$. Thus, both directions of the bisimulation are straightforward if $\lambda = s \rightarrow r : a$. We also know that each receive action done by the projected system can be simulated by the original system from Lemma 3.8.

We show that each receiving action made by the original system can also be made by the projected system, by contradiction. Let S' be such that $S \mid Q(S) \Longrightarrow S'$ and \hat{S}' be such that $\hat{S} \mid Q(\hat{S}) \Longrightarrow \hat{S}'$; and assume that \hat{S}' and S' are the first configurations that diverge (i.e., the first reduced systems where S' can do more actions than \hat{S}'). Let S'' (resp. \hat{S}'') be the

last stable system before S' (resp. \hat{S}'). We have

$$S \mid Q(S) \Longrightarrow S'' \xrightarrow{\phi} S' \quad \text{and} \quad \hat{S} \mid Q(\hat{S}) \Longrightarrow \hat{S}'' \xrightarrow{\phi} \hat{S}'$$

where ϕ consists only of sending actions. Note that ϕ must be the same in both systems since sending actions are preserved by \lesssim . We must have systems of the form:

$$S' \equiv r[s?a; P + s?b; P'] \mid sr : b \cdot \rho \mid \dots \quad \text{and} \quad \hat{S}' \equiv r[s?a; P] \mid sr : b \cdot \rho \mid \dots$$

where clearly \hat{S}' cannot match the actions of S' . However, since S is typable and S'' is stable, we must have

$$\circ \vdash S'' \blacktriangleright \mathcal{G}''$$

by Lemma 3.10 and Theorem 3.5. Since ϕ consists only of sending actions, we must have

$$S'' \equiv r[s?a; P + s?b; P'] \mid s[r!b\dots] \mid \dots$$

and

$$\hat{S}'' \equiv r[s?a; P] \mid s[r!b\dots] \mid \dots$$

it is clear that we must have $\mathcal{G}'' \downarrow_r \equiv s?b; Q + Q'$, which is in contradiction with the behaviour of r in \hat{S}'' . \square

3.5.6 Completeness

Our completeness result shows that every well-formed and projectable global type is inhabited by the system consisting of the parallel composition of all its projections.

We give two lemmas that are necessary for the proof of Theorem 3.7 below. Recall that the simulation relation \lesssim is given in Definition 3.8 and the merge function $_ \uplus _$ is given in Definition 3.2.

Lemma 3.11. *Let P and Q be two processes, if $P \uplus Q \neq \perp$ then $P \lesssim P \uplus Q$.*

Proof. We show the result by induction on the structure of P and a case analysis following Definition 3.2.

- If $P \equiv Q$, the result follows trivially.
- If $P = \bigoplus_{i \in I} r_i !a_i; P_i$, then, by Definition 3.2, we must have $Q = \bigoplus_{i \in I} r_i !a_i; Q_i$, and $P_i \uplus Q_i$ for all $i \in I$. We have the result by definition of \lesssim (all sending actions are preserved in both directions since P and Q have the same prefixes) and induction hypothesis: $P_i \lesssim P_i \uplus Q_i$ for all $i \in I$.
- If $P = \sum_{i \in I} s?a_i; P_i$, then there are two cases, either
 - $Q = \sum_{i \in I} s?a_i; Q_i$, then the result follows by definition of \lesssim (the processes share the same prefixes) and induction hypothesis; or
 - $Q = \sum_{j \in J} s?a_j; Q_j$ and $\forall i \in I : \forall j \in J : a_i \neq a_j$ and $I, J \neq \emptyset$, then we have, by Definition 3.2,

$$P \uplus Q = \sum_{i \in I} s?a_i; P_i + \sum_{j \in J} s?a_j; Q_j$$

and by definition of \lesssim :

$$\sum_{i \in I} s?a_i; P_i \lesssim \sum_{i \in I} s?a_i; P_i + \sum_{j \in J} s?a_j; Q_j$$

Note that if $P \equiv \mathbf{0}$, then we must have $Q \equiv \mathbf{0}$, cf. the condition $I, J \neq \emptyset$ above. □

Lemma 3.12. *Let S and S' be two programs such that $\mathcal{P}(S) = \mathcal{P}(S')$ and $\forall \mathbf{n} \in \mathcal{P}(S) : S(\mathbf{n}) \lesssim S'(\mathbf{n})$. If $\Gamma \vdash S \blacktriangleright \mathcal{G}$ then $\Gamma \vdash S' \blacktriangleright \mathcal{G}$.*

Proof. The proof is by straightforward induction on the derivation of $\Gamma \vdash S \blacktriangleright \mathcal{G}$. We only give the main argument here.

First, observe that by Proposition 3.1, if $S(\mathbf{n}) \equiv \mathbf{0}$, then we must have $S'(\mathbf{n}) \equiv \mathbf{0}$.

By definition of \lesssim , we have that each process $S'(n)$ may only differ from $S(n)$ by having some additional external choice branches, i.e., if $S(n) = P$ (with P an external choice), we can have $S'(n) = P + P'$. It is easy to see that we can always apply rule $[\text{+}]$ to discard the additional P' branch and return to the original system.

Note that, by definition of \lesssim , if $S(n) = P \oplus P'$, then $S'(n) = Q \oplus Q'$, with $P \lesssim Q$ and $P' \lesssim Q'$; therefore the induction hypothesis may be applied easily for the case when the last rule is $[\oplus]$. \square

We show an intermediary result (Lemma 3.13) which allows us to take into account the set of participants of a recursive global type during its projection. We illustrate the utility of this lemma with Example 3.13 below.

Example 3.13. Consider the global type:

$$\mathcal{G} = \mu\chi. s \rightarrow r : a; s \rightarrow r' : b; \chi$$

where we have $\mathcal{P}(\mathcal{G}) = \{s, r, r'\}$. While projecting \mathcal{G} , we obtain the sub-term \mathcal{G}' :

$$\mathcal{G}' = s \rightarrow r' : b; \chi$$

where we have $\mathcal{P}(\mathcal{G}') = \{s, r'\}$, i.e., r does not appear in \mathcal{G}' anymore. However, we need to record the fact that r is part of the global type since r is involved in the recursion. \diamond

Lemma 3.13. Let $\text{Proj}(\mathcal{G}, Q) = \big|_{n \in Q} n[\mathcal{G} \downarrow_n]$ with $\mathcal{P}(\mathcal{G}) \subseteq Q$,

If $\Gamma \vdash \mathcal{G}$, \mathcal{G} is projectable, and $\forall \chi \in \text{fv}(\mathcal{G}) : \exists \chi' : \forall n \in Q : \Gamma(n, \chi') = \chi$ then

$$\Gamma \vdash \text{Proj}(\mathcal{G}, Q) \blacktriangleright \mathcal{G}' \quad \text{with } \mathcal{G} \equiv \mathcal{G}'$$

Proof. We show this by induction on the structure of \mathcal{G} . Throughout the proof, we let

$$S = \bigsqcup_{n \in Q} n[\mathcal{G} \downarrow_n] \quad \text{and} \quad S_{Q'} = \bigsqcup_{n \in Q \setminus Q'} n[\mathcal{G} \downarrow_n]$$

Case $\mathcal{G} = \chi$. Then, by Definition 3.1, we have

$$S \equiv \bigsqcup_{n \in Q} n[\chi]$$

By assumption, we have $\exists \chi : \forall n \in Q : \Gamma(n, \chi) : \chi$, hence we can apply rule $[\chi]$, and we obtain the required result.

Case $\mathcal{G} \equiv \mathbf{0}$. We have $S \equiv \bigsqcup_{n \in Q} n[\mathbf{0}]$ and the results holds by rule $[\mathbf{0}]$.

Case $\mathcal{G} \equiv s \rightarrow r : a; \mathcal{G}'$.

By definition of projection, we have

$$S \equiv s[r!a; \mathcal{G}' \downarrow_s] \mid r[s?a; \mathcal{G}' \downarrow_r] \mid S_{\{s,r\}}$$

with $Q = \mathcal{P}(S) \cup \{s, r\}$. We can apply rule $[\cdot]$ and the induction hypothesis in order to have the result, i.e.,

$$\frac{\text{by IH}}{\Gamma \vdash s[\mathcal{G}' \downarrow_s] \mid r[\mathcal{G}' \downarrow_r] \mid S_{\{s,r\}} \triangleright \mathcal{G}'} \quad [\cdot]}{\Gamma \vdash s[r!a; \mathcal{G}' \downarrow_s] \mid r[s?a; \mathcal{G}' \downarrow_r] \mid S_{\{s,r\}} \triangleright s \rightarrow r : a; \mathcal{G}'}$$

Observe that $S_{\{s,r\}} \not\downarrow$ otherwise it would mean that

$$\exists s' \rightarrow r' : a \in R(\mathcal{G}') \text{ such that } \{s, r\} \cap \{s', r'\} = \emptyset$$

which is in contradiction with $\vdash \mathcal{G}$.

Case $\mathcal{G} \equiv \mathcal{G}_0 + \mathcal{G}_1$.

Let us assume that for all $n \in \mathcal{P}(\mathcal{G})$ (that is not the sender on the prefixes in $\mathcal{R}(\mathcal{G})$) that $\mathcal{G}_0 \downarrow_n \uplus \mathcal{G}_1 \downarrow_n \neq \perp$; otherwise we can re-arrange \mathcal{G} such that this condition holds since \mathcal{G} is projectable.

By definition of projection, we have

$$S \equiv s[\mathcal{G}_0 \downarrow_s \oplus \mathcal{G}_1 \downarrow_s] \mid \prod_{n \in Q \setminus \{s\}} n[\mathcal{G}_0 \downarrow_n \uplus \mathcal{G}_1 \downarrow_n]$$

We can apply rule $[\oplus]$. Indeed, by well-formedness, rule $[\oplus]$ is applicable (s is “preventing” the other participants to interact, similarly to the case above) and we obtain:

$$\frac{\begin{array}{c} \vdots \\ \hline \Gamma \vdash s[\mathcal{G}_0 \downarrow_s] \mid \prod_{n \in Q \setminus \{s\}} n[\mathcal{G}_0 \downarrow_n \uplus \mathcal{G}_1 \downarrow_n] \triangleright \mathcal{G}_0 \end{array} \quad \begin{array}{c} \vdots \\ \hline \Gamma \vdash s[\mathcal{G}_1 \downarrow_s] \mid \prod_{n \in Q \setminus \{s\}} n[\mathcal{G}_0 \downarrow_n \uplus \mathcal{G}_1 \downarrow_n] \triangleright \mathcal{G}_1 \end{array}}{[\oplus] \quad \Gamma \vdash S \triangleright \mathcal{G}}$$

By induction hypothesis, we have the result for

$$\Gamma \vdash s[\mathcal{G}_0 \downarrow_s] \mid \prod_{n \in Q \setminus \{s\}} n[\mathcal{G}_0 \downarrow_n] \triangleright \mathcal{G}_0 \quad \Gamma \vdash s[\mathcal{G}_1 \downarrow_s] \mid \prod_{n \in Q \setminus \{s\}} n[\mathcal{G}_1 \downarrow_n] \triangleright \mathcal{G}_1$$

By Lemma 3.11 and the assumption that \mathcal{G} is projectable, we have

$$\forall n \in Q \setminus \{s\} : \mathcal{G}_i \downarrow_n \lesssim \mathcal{G}_0 \downarrow_n \uplus \mathcal{G}_1 \downarrow_n \quad i \in \{0, 1\}$$

Thus, by Lemma 3.12, we have

$$\Gamma \vdash s[\mathcal{G}_i \downarrow_s] \mid \prod_{n \in Q \setminus \{s\}} n[\mathcal{G}_i \downarrow_n] \triangleright \mathcal{G}_i \implies \Gamma \vdash s[\mathcal{G}_i \downarrow_s] \mid \prod_{n \in Q \setminus \{s\}} n[\mathcal{G}_0 \downarrow_n \uplus \mathcal{G}_1 \downarrow_n] \triangleright \mathcal{G}_i$$

and we have the required result.

Case $\mathcal{G} \equiv \mathcal{G}_0 \mid \mathcal{G}_1$.

We have S of the form below, by definition of projections (and well-formedness)

$$S \equiv \text{Proj}(\mathcal{G}_0, Q_0) \mid \text{Proj}(\mathcal{G}_1, Q_1) \quad \text{with } Q_0 \cap Q_1 = \emptyset$$

Note that since $|\text{Indep}(\mathcal{G})| > 1$, we have $\text{fv}(\mathcal{G}) = \emptyset$, thus, by induction hypothesis, we have

$$\circ \vdash \text{Proj}(\mathcal{G}_i, Q_i) \triangleright \mathcal{G}'_i \quad \text{and} \quad \mathcal{G}_i \equiv \mathcal{G}'_i$$

for $i \in \{0, 1\}$ and we have the result by applying rule $[\mid]$:

$$\frac{\frac{\text{by IH}}{\circ \vdash \text{Proj}(\mathcal{G}_0, Q_0) \triangleright \mathcal{G}_1} \quad \frac{\text{by IH}}{\circ \vdash \text{Proj}(\mathcal{G}_1, Q_1) \triangleright \mathcal{G}_1}}{[\mid] \quad \Gamma \vdash S \triangleright \mathcal{G}}$$

Case $\mathcal{G} \equiv \mu\chi.\mathcal{G}'$.

By definition of projections, we have

$$S \equiv \left|_{n \in Q_1} n[\mu\chi.\mathcal{G}' \downarrow_n] \mid \left|_{n \in Q_2} n[\mathbf{0}] \quad \text{with } Q = Q_1 \cup Q_2 \text{ and } Q_1 \cap Q_2 = \emptyset$$

Since \mathcal{G}' is prefix-guarded, there must be $s, r \in Q_1$ such that

$$s[\mu\chi.\mathcal{G}' \downarrow_s] \mid r[\mu\chi.\mathcal{G}' \downarrow_r] \Downarrow$$

Therefore, rule $[\mu]$ is applicable after having discarded the empty processes via rule $[\mid]$

(followed by the axiom $[\mathbf{0}]$). We have the following derivation:

$$\begin{array}{c}
 \text{by IH} \\
 \hline
 \frac{\Gamma' \vdash \prod_{n \in Q_1} n[\mathcal{G}' \downarrow_n] \blacktriangleright \mathcal{G}'}{[\mu]} \quad \frac{[\mathbf{0}]}{\Gamma \vdash \prod_{n \in Q_2} n[\mathbf{0}] \blacktriangleright \mathbf{0}} \\
 \frac{[\]}{\Gamma \vdash \prod_{n \in Q_1} n[\mu\chi.\mathcal{G}' \downarrow_n] \mid \prod_{n \in Q_2} n[\mathbf{0}] \blacktriangleright \mathcal{G} \mid \mathbf{0}}
 \end{array}$$

where $\Gamma' = \Gamma \cdot \Gamma_S$ and Γ_S such that $\forall n \in Q_1 : \Gamma(n, \chi) = \chi$

We have the required result by induction hypothesis. \square

Theorem 3.7. *If $\vdash \mathcal{G}$ and \mathcal{G} is projectable, then there is $\mathcal{G}' \equiv \mathcal{G}$ such that*

$$\circ \vdash \prod_{n \in \mathcal{P}(\mathcal{G})} n[\mathcal{G} \downarrow_n] \blacktriangleright \mathcal{G}'$$

Proof. By Lemma 3.13, with $Q = \mathcal{P}(\mathcal{G})$ and $\Gamma = \circ$ (since \mathcal{G} is closed by assumption). \square

3.6 Perspectives

In this section, we discuss possible generalisations of our framework to extend the set of typable systems and discuss a limitation of the synthesis as a type system.

Concurrent branches and recursion. In order to synthesise global types where concurrent branches appear under recursion (i.e., a construct which is not allowed in the current setting), we propose to introduce an additional recursion variable environment which is used linearly. We use judgements of the form

$$\hat{\Gamma}; \Gamma \vdash S \blacktriangleright \mathcal{G}$$

which may be read as “the system S forms a choreography defined by a global type \mathcal{G} , under the environments $\hat{\Gamma}$ and Γ ”. Environments $\hat{\Gamma}$ and Γ are as in Section 3.4, except from the fact that each of the “hypothesis” in $\hat{\Gamma}$ must be used exactly once. The rules from Section 3.4.1 remains essentially the same but for the rule introducing parallel branches and the axioms for termination and recursion variable. We give the new rules below.

$$[|] \frac{\hat{\Gamma} \cdot \Gamma; \circ \vdash S \triangleright \mathcal{G} \quad \hat{\Gamma} \cdot \Gamma'; \circ \vdash S' \triangleright \mathcal{G}'}{\hat{\Gamma}; \Gamma \cdot \Gamma' \vdash S | S' \triangleright \mathcal{G} | \mathcal{G}'}$$

$$[\chi] \frac{\forall 1 \leq i \leq k : \Gamma(n_i, \chi_i) = \chi}{\circ; \Gamma \vdash n_1[\chi_1] | \dots | n_k[\chi_k] \triangleright \chi} \quad [0] \frac{\forall n \in \mathcal{P}(S) : S(n) = \mathbf{0}}{\circ; \Gamma \vdash S \triangleright \mathbf{0}}$$

$$[\chi_L] \frac{\text{dom}(\hat{\Gamma}) = \{(n_i, \chi_i) \mid 1 \leq i \leq k\} \quad \text{img}(\hat{\Gamma}) = \{\chi\}}{\hat{\Gamma}; \Gamma \vdash n_1[\chi_1] | \dots | n_k[\chi_k] \triangleright \chi}$$

Rule $[|]$ allows to separate the system in two concurrent branches. In this case the normal environment Γ is split in two disjoint parts and moved to the linear environments. Rules $[\chi]$ and $[0]$ are updated so that they are applicable only when the linear environment is empty. The new rule $[\chi_L]$ allows to use the variables stored in $\hat{\Gamma}$ only if each behaviour in the system consists of recursion variables in $\hat{\Gamma}$ and each element in $\hat{\Gamma}$ is indeed used. Intuitively, these rules allows some participants to agree on a “meeting point” (i.e., a recursion definition), then they behave concurrently for a while, and return to the meeting point eventually.

Another change in rules concerns rule $[\mu]$ where we add the constraint that there must be exactly one pair of participants that interact after the recursion. This is to prevent rule

$[\mu]$ to be applicable on systems of the form

$$s_1[\mu\chi.r_1!a;\chi] \mid r_1[\mu\chi.s_1?a;\chi] \mid s_2[\mu\chi.r_2!b;\chi] \mid r_2[\mu\chi.s_2?b;\chi]$$

where it is clear that rule $[\mid]$ should be used before rule $[\mu]$ so that we obtain a global type of the form:

$$\mu\chi.s_1 \rightarrow r_1 : a; \chi \mid \mu\chi'.s_2 \rightarrow r_2 : b; \chi' \quad \checkmark$$

Instead of a global type of the form:

$$\mu\chi.(s_1 \rightarrow r_1 : a; \chi \mid s_2 \rightarrow r_2 : b; \chi) \quad \times$$

We show an example of a system that would be rejected by the original rules but is accepted by this extension.

Example 3.14. Consider the system

$$S = s[\mu\chi.r!a;r'!b;\chi] \mid r[\mu\chi.s?a;s'?c;\chi] \mid r'[\mu\chi.s?b;\chi] \mid s'[\mu\chi.r!c;\chi]$$

which has the following global type

$$\mathcal{G} = \mu\chi.s \rightarrow r : a; (s \rightarrow r' : b; \chi \mid s' \rightarrow r : c; \chi)$$

We give the type derivation below, where we pose

$$\Gamma_1 = (s, \chi) : \chi \cdot (r', \chi) : \chi \quad \text{and} \quad \Gamma_2 = (r, \chi) : \chi \cdot (s', \chi) : \chi$$

and we abstract from the intermediary global types in the derivation.

$$\begin{array}{c}
 \frac{[\chi_L] \text{-----}}{\Gamma_1; \circ \vdash s[\chi] \mid r'[\chi] \blacktriangleright \chi} \quad \frac{[\chi_L] \text{-----}}{\Gamma_2; \circ \vdash |r[\chi] \mid s'[\chi] \blacktriangleright \chi} \\
 \frac{[\cdot] \text{-----} \quad \frac{[\cdot] \text{-----}}{\Gamma_2; \circ \vdash |r[s'?c;\chi] \mid s'[r!c;\chi] \blacktriangleright \dots}}{\Gamma_1; \circ \vdash s[r'!b;\chi] \mid r'[s?b;\chi] \blacktriangleright \dots} \quad \frac{[\cdot] \text{-----}}{\Gamma_2; \circ \vdash |r[s'?c;\chi] \mid s'[r!c;\chi] \blacktriangleright \dots} \\
 \frac{[\cdot] \text{-----}}{\circ; \Gamma_1 \cdot \Gamma_2 \vdash s[r'!b;\chi] \mid r[s'?c;\chi] \mid r'[s?b;\chi] \mid s'[r!c;\chi] \blacktriangleright \dots} \\
 \frac{[\cdot] \text{-----}}{\circ; \Gamma_1 \cdot \Gamma_2 \vdash s[r!a; r'!b;\chi] \mid r[s?a; s'?c;\chi] \mid r'[s?b;\chi] \mid s'[r!c;\chi] \blacktriangleright \dots} \\
 \frac{[\mu] \text{-----}}{\circ; \circ \vdash S \blacktriangleright \mathcal{G}}
 \end{array}$$

Notice that, when rule $[\cdot]$ is used, the environment $\Gamma_1 \cdot \Gamma_2$ is split in two parts and moved to the linear environment. \diamond

Remarkably, if we relax the corresponding conditions on the well-formedness of global types and their projections, the properties presented in Section 3.5 also holds for this generalisation. In fact, the main change concerns *knowledge of choice*. This property holds in this extension since each participant involved in a “concurrent recursion” must *always* return to the recursion point.

Minimal local types. Some system of local types may fail to be assigned a global type because the behaviours of two participants do not match within a recursion, even if they do not deadlock – this is due to the syntactic nature of the type system. We may be given more chances to type a system of local types if each local type has been *minimised* up to bisimilarity beforehand. For instance, the system:

$$s[\mu\chi.r!a;\chi] \mid r[\mu\chi.s?a;s?a;\chi] \tag{3.9}$$

is not typable because the typing derivation will eventually reach a term of the form

$$(s, \chi) : \chi \cdot (r, \chi) : \chi \vdash s[\chi] \mid r[s?a; \chi] \blacktriangleright \dots$$

for which no rule from Figure 3.3 is applicable. In particular, observe that unfolding any of the behaviours in (3.9) will not help.

However, the system in (3.9) is typable if we minimise each behaviour. Indeed, it is possible to translate the behaviour of r into an automaton, which can be minimised and translated back to the process: $\mu\chi.s?a;\chi$. Note that, at the level of behavioural specifications, minimising a process (up to bisimilarity) is a sound transformation.

After minimisation, the system in (3.9) is assigned the global type $\mu\chi.s \rightarrow r : a ; \chi$. It is easy to see that the system consisting of the projections of this global type is bisimilar to the original one, i.e.,

$$s[\mu\chi.r!a;\chi] \mid r[\mu\chi.s?a;s?a;\chi] \approx s[\mu\chi.r!a;\chi] \mid r[\mu\chi.s?a;\chi]$$

The result follows directly from the fact that minimisation is done up to bisimilarity and the fact that a synthesised global type preserve the behaviour of the original system (cf. Theorem 3.6).

A limitation. A disadvantage of presenting a synthesis as a type system is that both checking that each input/output actions are matched and checking that there is a 1-buffer execution of the system (between two stable configurations) must be done at the same time. Instead, when working with communicating machines, one can first construct a transition system corresponding to all the possible (bounded) executions of the systems, so to check that, e.g., the system is deadlock free, there is no orphan message configuration etc. Then, one may check that every stable configuration is reachable by a 1-buffer execution and, therefore, can be captured by a global type.

This limitation is most noticeable if we were to allow a more liberal version of the external choice in the behaviours, e.g., a behaviour of the form:

$$P = s?a; s'?b + s'?b; s?c \quad \text{with } s \neq s'$$

Indeed, consider the following system

$$S = s[r!a; s'!d] \mid r[P] \mid s'[s?d; r!b]$$

In a 1-buffer execution, this system does not deadlock. Indeed, under such a restriction r must read a from the queue sr first, then b from the queue $s'r$. However, the runtime system $S \mid Q(S)$ may reduce to the following system:

$$S' = s[0] \mid sr : a \mid r[s?a; s'?b + s'?b; s?c] \mid s'r : b \mid s'[0]$$

in which case if r reads b first (i.e., the right branch of the external choice is chosen), it will deadlock because it will not be able to read c from the queue sr .

With the current type system, it is not possible to discard such systems as we must eventually discard one of the branches of participant r and thus the typing will not check the “faulty” branch. Indeed, this branch is “invisible” to the type system which only considers 1-buffer executions.

3.7 Concluding Remarks

We presented a type system to synthesise a choreography (i.e., a global type) from a set of local specifications (i.e., local types). Such a global type is unique, well-formed, and its projections are equivalent to the original local types. We have shown safety and progress properties for (typable) local session types and given a subject reduction theorem for our

type system.

Observe that our type system ensures that a synthesised global type meets the connectedness conditions of [50]. This follows from the fact that a synthesised global type is well-formed: local choice and sequentiality hold.

In addition, considering a set of typable local types as an implementation of its synthesised global type, the properties identified in [29], i.e., sequentiality, alternativeness, shuffling, fitness and exhaustivity, are guaranteed to hold.⁵ This follows from the fact that a global type is well-formed and its projections are equivalent to the original system, i.e., its “implementation”. In particular, *exhaustivity* – requiring that if a sequence of interactions is specified by a global type, then there must exist at least an execution of its implementation that exhibits these interactions – is guaranteed by construction.

In the rest of this thesis, we will present some applications of the synthesis. In Chapter 4, we show that once a choreography is available, it may be decorated with logic predicates so to constrain interactions declaring senders obligations and receivers requirements on the values of the exchanged data and on the choice of the branches to follow. In particular, we focus on algorithms and a methodology to correct these choreographies if they do not satisfy some requirements. In Chapter 5, we present a formal framework for distributed systems where the synthesis of choreographies is used as basis for finding sets of compliant contracts.

⁵See Section 2.2.1 for a description of these properties.

Amending Contracts for Choreographies

We show a few techniques that help software architects to amend *global assertions* during the design of distributed choreographies. Our results include: two algorithms to solve history sensitivity problems, one algorithm to solve temporal satisfiability problems, and a methodology for applying the algorithms to protocol design.

4.1 Introduction

Once a global type has been designed, or synthesised as in Chapter 3, it may be refined into a *global assertion* [12]. That is, a global type decorated with predicates that constrain interactions, declaring senders obligations and receivers requirements on exchanged data and on the choice of the branches to follow.

Once designed, a global assertion \mathcal{G} may be projected on endpoint assertions that are local types – modelling the behaviour of a specific participant – constrained according to the predicates of \mathcal{G} . However, a global assertion may be projected into “safe” local types only if it is *well-asserted* [12], namely when it obeys two precise design principles: *history-sensitivity* (HS for short) and *temporal satisfiability* (TS for short). Informally, HS demands that a participant having an obligation on a predicate has enough information for

choosing a set of values that guarantees it. Instead, TS requires that the values sent in each interaction do not make predicates of future interactions unsatisfiable.

The main motivation of our interest in HS and TS is that, in global assertions, they are the technical counterparts of the fundamental coordination issue that could be summarized in the slogan “who do what and when do they do it?”. In fact, HS pertains to *when* variables are constrained and *who* constrain them, while TS pertains to *which* values variables take. The contracts specified in global assertions are, on the one hand, “global” as they pertain to the whole choreography while, on the other hand, they are also “local” in (at least) two aspects. The first is that they assign responsibilities to participants (*who*) at definite moments of the computation (*when*). The second aspect is that the values assigned to variables are critical because one could either over-constrain variables fixed in the past or over-restrict the range of those assigned in the future (*which*). These conditions – especially TS – are rather crucial as global assertions that violate them may be infeasible or fallacious. Remarkably, a global assertion not satisfying TS may lead to conversations in which progress is not guaranteed unless one of the participants deliberately violates the contract.

Guaranteeing HS and TS is often non-trivial, and this burden is on the software architect; using tools like the ones described in [52], one only highlights the problems but does not help to fix them. HS and TS are global semantic properties that may be hard to achieve. Namely, TS requires to trace back for “under-constrained” interactions (i.e., which allow values causing future predicates to be unsatisfiable) and re-distribute there the unsatisfiable constraints.

A motivating example. Using the running example of Chapter 3, we illustrate how TS problems may easily occur at design time. In the global type depicted in Figure 4.1, a buyer b_i requests a quote from a seller s_i , then depending on whether the client is satisfied with the quote, b_i chooses either the *ok* branch and the interactions terminate, or

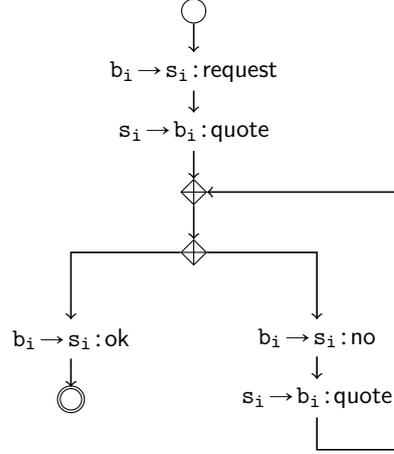


Figure 4.1: Seller-Buyer global type (extract)

b_i chooses the no branch and the seller has to offer a new quote.

A natural refinement of this global type is the following global assertion:¹

$$\hat{\mathcal{G}} = b_i \rightarrow s_i : \{request \mid \mathbf{true}\}; \quad (1)$$

$$s_i \rightarrow b_i : \{quote_f \mid \mathbf{true}\}; \quad (2)$$

$$\mu\chi\langle quote_f \, quote_f \rangle \{quote_c \, quote_p \mid quote_c \leq quote_p\}. \quad (3)$$

$$b_i \rightarrow s_i : \left\{ \begin{array}{l} \{\mathbf{true}\} \text{ ok} : \mathbf{0}, \\ \{\mathbf{true}\} \text{ no} : s_i \rightarrow b_i : \{quote_n \mid quote_n < quote_c\}; \end{array} \right. \quad (4)$$

$$\chi\langle quote_n \, quote_c \rangle \quad (5)$$

$$\left. \begin{array}{l} \\ \\ \end{array} \right\} \quad (6)$$

where $request$, $quote_f$, $quote_c$, $quote_p$ and $quote_n$ are *interaction variables*; we assume that the sort of $request$ is string and the sort of the other variables is int – for simplicity.

In lines (1-2), the buyer and the seller exchange variables $request$ and $quote_f$. The variables are unconstrained (the predicates are set to true).

Line (3) defines a recursion over the recursion parameters $quote_c$ and $quote_p$ – which

¹ Observe that the structure of $\hat{\mathcal{G}}$ is the same as the global type of Figure 4.1.

are both initialised to the value of $quote_f$. The invariant of the recursion, i.e., $quote_c \leq quote_p$ requires that the current quote ($quote_c$) is less than the previous one ($quote_p$). This invariant trivially holds in the first iteration.

At line (4), b_i selects either the ok or no branch (the choice of a branch may be constrained by a predicate, but in this case they are both set to **true**, so to model a non-deterministic choice).

In the recursive branch, at lines (5-6), the seller has to send a new quote ($quote_n$) to the buyer and this quote must be (strictly) less than the previous quote, i.e., $quote_n < quote_c$. The recursive call $\chi\langle quote_n quote_c \rangle$ says that the recursion arguments are now $quote_n$ and $quote_c$ (i.e., the current quote and the previous quote).

We will see that \hat{G} is *well-asserted*. However, the designer may soon realise that \hat{G} is under-constrained as it may imply that the seller has to make offers that are arbitrarily low. Assume that the designer tackles this flaw by updating the recursion invariant (line (3)) to the predicate

$$\text{MIN} < quote_c \leq quote_p \quad (4.1)$$

where MIN is a constant representing the minimum price that a seller is willing to sell an item. In fact, this change makes the global assertion violate temporal satisfiability since, e.g., if s_i sends a quote (at line (2) or (5)) that is less than MIN, then the invariant (4.1) does not hold.

In the rest of this chapter, we will see that solving this kind of problems is not trivial, but may be solved automatically by our algorithms and methodology.

4.2 Preliminaries

We now define global assertions formally. Let \mathcal{V} (ranged over by u, v, x, y, \dots) be a countably infinite set of *interaction variables*, such that $\mathbb{P} \cap \mathcal{V} = \emptyset$. Recall that \mathbb{P} is the set of

participants of a choreography (cf. Chapter 3).

Hereafter, $\vec{\cdot}$ represents a list of some elements (for instance, \vec{v} is a list of interaction variables). The concatenation of \vec{x} and \vec{y} is denoted by the juxtaposition $\vec{x} \vec{y}$; abusing notation, we identify a one-element list with its (unique) element and identify lists with the underlying sets of their elements (e.g., $a \in \vec{x}$ indicates that a occurs in the list \vec{x}).

As in [12], we parametrise our constructions by abstracting away from the logical language Ψ adopted. Here, it suffices to assume that Ψ is a decidable fragment of a first-order logic obtained by adding first-order quantification to a language of boolean expressions. In fact, we allow expressions (ranged over by e) that include constructors and operators/relations of common data types (e.g., strings, integers, booleans, etc.) and include variables drawn from \mathcal{V} . (For simplicity, our examples use basic numeric types or strings.) We write $\text{var}(e)$ to denote the set of variables occurring in e and use the symbol \implies to denote logic implication. Then a predicate $\psi \in \Psi$ is either a boolean expression e (understood to be a boolean expression in our language of expressions), or a quantified predicate $\forall \vec{v} : e$ or $\exists \vec{v} : e$. Given a predicate ψ in Ψ , $\text{var}(\psi)$ is the set of free interaction variables of ψ (we write $\psi(\vec{v})$ to emphasise that $\text{var}(\psi) \subseteq \vec{v}$).

The main ingredients of global assertions are *interactions*, abbreviated ι , like

$$\mathbf{s} \rightarrow \mathbf{r} : \{ \vec{v} \mid \psi \} \quad (4.2)$$

where $\mathbf{s}, \mathbf{r} \in \mathbb{P}$ are the *sender* and the *receiver*, $\vec{v} \subseteq \mathcal{V}$ is a pairwise distinct list of variables, and $\psi \in \Psi$. We say that the variables \vec{v} in (4.2) are *introduced* by \mathbf{s} . The interaction (4.2) reads as “ \mathbf{s} has to send to \mathbf{r} some values for \vec{v} that satisfy ψ ” or as “ \mathbf{r} relies on the fact that the values fixed by \mathbf{s} for \vec{v} satisfy ψ ”. For instance,²

$$\mathbf{s} \rightarrow \mathbf{r} : \{ v \ w \mid \exists u : v = u \times w \}$$

²For simplicity, we assume the typing of variables understood.

states that s has the obligation to send r two values such that the first is a multiple of the second. Given ι as in (4.2), we define

$$\text{snd}(\iota) \stackrel{\text{def}}{=} s, \quad \text{rcv}(\iota) \stackrel{\text{def}}{=} r, \quad \text{var}(\iota) \stackrel{\text{def}}{=} \vec{v}, \quad \text{and} \quad \text{cst}(\iota) \stackrel{\text{def}}{=} \Psi$$

Remark 4.1. *In [12], interactions specify a channel over which participants communicate. For consistency with the rest of this thesis, we omit channels since they are inconsequential to our results. In fact, the algorithms we present do not use identities of channels but only those of participants and variables.*

Global assertions are ranged over by \mathcal{G} and have the following syntax:

$$\begin{array}{ll} \mathcal{G} ::= \iota; \mathcal{G} & \text{[PREFIX]} \\ | \quad s \rightarrow r : \{ \{ \Psi_j \} |_{j : \mathcal{G}_j} \}_{j \in J} & \text{[BRANCHING]} \\ | \quad \mu \chi \langle \vec{e} \rangle \{ \vec{v} \mid \Psi \}. \mathcal{G} & \text{[RECURSIVE DEFINITION]} \\ | \quad \chi \langle \vec{e} \rangle & \text{[RECURSIVE CALL]} \\ | \quad \mathbf{0} & \text{[END SESSION]} \end{array}$$

where $\Psi, \Psi_j \in \Psi$ and l_j ranges over a set of labels.

The syntax above is essentially borrowed from [12] but for a slightly simplified notation. In [12], the semantics of global assertions is given in terms of *endpoint assertions* (by projecting global assertions to endpoint assertions and exploiting the operational semantics of the latter). In this chapter, only the syntactic aspects of global assertions given below are relevant; therefore, we give an informal account of the semantics of global assertions.

The prefix production $\iota; \mathcal{G}$ defines a global assertion where the interaction ι must precede the interactions in \mathcal{G} . The branching production allows the selector s to choose one of the labels l_j and send it to r , then the interactions in \mathcal{G}_j occur. Recursion is dealt with

as usual but for the presence of an initialisation vector \vec{e} (of the same length as \vec{v}) which specifies the initial values of each formal parameter in \vec{v} and onto which a recursion invariant ψ is specified. Finally, the last production represents a completed global assertion; trailing occurrences of $\mathbf{0}$ are often omitted.

In a recursive definition $\mu\chi\langle\vec{e}\rangle\{\vec{v} \mid \psi\}. \mathcal{G}$, occurrences of χ (i.e., recursive calls) in \mathcal{G} must be prefix-guarded and the length of \vec{e} is the same as \vec{v} . Also, we assume that variables χ are always in the scope of a recursive definition $\mu\chi\langle-\rangle\{- \mid -\}$, and that, in a recursive call $\chi\langle\vec{e}\rangle$, \vec{e} match the sorts of its corresponding recursive definition $\mu\chi\langle\vec{e}'\rangle\{\vec{v} \mid \psi\}. \mathcal{G}$ and the length of \vec{e} is the same as \vec{v} .

We denote with $\text{var}(\mathcal{G})$ the set of interaction variables and recursion parameters in \mathcal{G} . The interaction variables $\text{var}(\mathfrak{t})$ of global assertion $\mathfrak{t}; \mathcal{G}$ are bound in \mathcal{G} and in $\text{cst}(\mathfrak{t})$; similarly, the formal parameters \vec{v} in a recursive definition $\mu\chi\langle-\rangle\{\vec{v} \mid \psi\}. \mathcal{G}$ are bound in ψ and in the recursion body \mathcal{G} . We consider *closed* global assertions (i.e., for any occurrence of $v \in \mathcal{V}$ in \mathcal{G} either the occurrence is in a recursive definition having v as formal parameter or there is an interaction \mathfrak{t} in \mathcal{G} such that $v \in \text{var}(\mathfrak{t})$ which precedes that occurrence of v).

Remark 4.2. *For simplicity, we assume that bound interaction variables are pairwise distinct.*

Remark 4.3. *For consistency with [12], we use a restricted form of branching, where a choice must be made between exactly two participants (one sender and one receiver). However, we conjecture that our results are applicable to global assertions that features constructs of the form*

$$\sum_{j \in J} s \rightarrow r_j : \{l_j \mid \psi_j\}; \mathcal{G}_j$$

where we may have $r_j \neq r_i$ for $i \neq j \in J$.

Definition 4.1 (Knows). *Under the syntactic restrictions listed above, we say that a participant p knows a variable $v \in \text{var}(\mathcal{G})$ if one of the following conditions holds:*

- there is ι in \mathcal{G} such that $v \in \text{var}(\iota)$ and $p \in \{\text{snd}(\iota), \text{rcv}(\iota)\}$ or
- there is a recursive definition $\mu\chi\langle\vec{e}_1 e \vec{e}_2\rangle\{\vec{v}_1 v \vec{v}_2 \mid \Psi\}$. \mathcal{G}' in \mathcal{G} such that p knows all the variables in $\text{var}(e)$ and, for each recursive invocation $\chi\langle\vec{e}'_1 e' \vec{e}'_2\rangle$ in \mathcal{G}' , p knows all the variables in $\text{var}(e')$.

We denote with $\text{knows}_p(\mathcal{G})$ the set of variables in $\text{var}(\mathcal{G})$ that p knows.

Example 4.1. Consider the following global assertion

$$\begin{aligned} \mathcal{G}_{\text{ex4.1}} = & \text{I} \rightarrow \text{Server} : \{x \mid x \geq 3\}; \\ & \mu\chi\langle 3 \rangle \{r \mid \mathbf{true}\}. \text{Server} \rightarrow \text{Player} : \{ \\ & \quad \{r > x\} \text{less} : \text{Player} \rightarrow \text{Server} : \{y \mid \mathbf{true}\}; \chi\langle y \rangle, \\ & \quad \{r < x\} \text{greater} : \text{Player} \rightarrow \text{Server} : \{z \mid \mathbf{true}\}; \chi\langle z \rangle, \\ & \quad \{r = x\} \text{win} : \mathbf{0} \} \end{aligned}$$

where I initialises a value $x \geq 3$ for Server . Then, repeatedly, Server sends a label chosen in the set $\{\text{less}, \text{greater}, \text{win}\}$ to Player depending on r being greater, smaller, or equal to the value of x ; and Player replies with an integer in the first two cases while the interaction ends if win was sent by Server . In $\mathcal{G}_{\text{ex4.1}}$, both I and Server know x while Player does not know it; instead the recursion parameter r is known only to Server and Player . \diamond

It is convenient to treat global assertions as trees whose nodes are drawn from a set \mathcal{N} (ranged over by n, n', \dots) and labelled with information on the syntactic categories of the syntax of global assertions. Hereafter, we write $n \in T$ if n is a node of a tree T , \underline{n} to denote the label of n , and T^\bullet for the root of T .

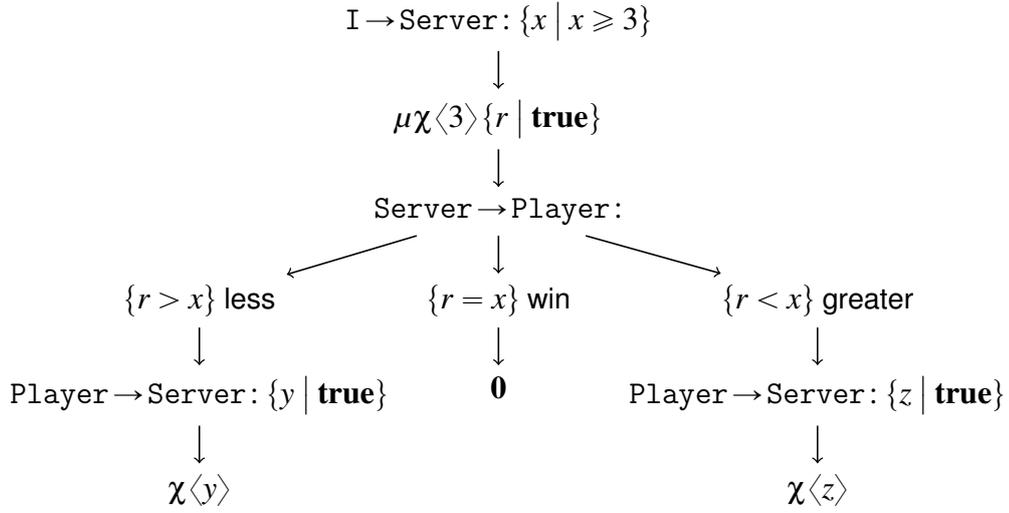
Definition 4.2 (Assertion Tree). *The assertion tree $\text{T}(\mathcal{G})$ of a global assertion \mathcal{G} is defined as follows:*

- If $\mathcal{G} = \iota$; \mathcal{G}' then $\text{T}(\mathcal{G})^\bullet$ has label ι and its unique child is $\text{T}(\mathcal{G}')$.

- If $\mathcal{G} = s \rightarrow r : \{\Psi_j\}_j : \mathcal{G}_j\}_{j \in J}$ then $\mathsf{T}(\mathcal{G})^\bullet$ has label $s \rightarrow r :$ and its children are $\{n_j\}_{j \in J} \subseteq \mathcal{N}$ such that, for each $j \in J$, $n_j = \{\Psi_j\}_j$ and $\mathsf{T}(\mathcal{G}_j)$ is the unique child of n_j .
- If $\mathcal{G} = \mu\chi\langle\vec{e}\rangle\{\vec{v} \mid \Psi\} . \mathcal{G}'$ then $\mathsf{T}(\mathcal{G})^\bullet$ has label $\mu\chi\langle\vec{e}\rangle\{\vec{v} \mid \Psi\}$ and its unique child is $\mathsf{T}(\mathcal{G}')$.
- If $\mathcal{G} = \chi\langle\vec{e}\rangle$ then $\mathsf{T}(\mathcal{G})$ consists of one node with label $\chi\langle\vec{e}\rangle$.
- If $\mathcal{G} = \mathbf{0}$ then $\mathsf{T}(\mathcal{G})$ consists of one node with label $\mathbf{0}$.

We denote the set of assertion trees as \mathcal{T} and let T, T', \dots range over \mathcal{T} .

Example 4.2. The assertion tree of $\mathcal{G}_{\text{ex4.1}}$ in Example 4.1 can be depicted as:



where identities of nodes are not shown and only their labels appear. ◇

For convenience, given $T \in \mathcal{T}$, we will use the partial functions

$$\text{var}_T : \mathcal{N} \rightarrow 2^{\mathcal{V}}, \quad \text{cst}_T : \mathcal{N} \rightarrow \Psi, \quad \text{and} \quad \text{snd}_T, \text{rcv}_T : \mathcal{N} \rightarrow \mathbb{P}$$

that are undefined on $\mathcal{N} \setminus \{n \mid n \in T\}$ and defined as follows otherwise:

$$\begin{aligned} \text{var}_T(n) &= \begin{cases} \text{var}(\mathfrak{t}), & \text{if } \underline{n} = \mathfrak{t} \\ \emptyset, & \text{otherwise} \end{cases} & \text{cst}_T(n) &= \begin{cases} \Psi, & \text{if } \underline{n} = \mathfrak{t} \text{ and } \text{cst}(\mathfrak{t}) = \Psi \\ \Psi, & \text{if } \underline{n} = \{\Psi\}! \\ \mathbf{true}, & \text{otherwise} \end{cases} \\ \text{snd}_T(n) &= \begin{cases} \text{snd}(\mathfrak{t}), & \text{if } \underline{n} = \mathfrak{t} \\ \mathfrak{s}, & \text{if } \underline{n} = \mathfrak{s} \rightarrow \mathfrak{r} \end{cases} & \text{rcv}_T(n) &= \begin{cases} \text{rcv}(\mathfrak{t}), & \text{if } \underline{n} = \mathfrak{t} \\ \mathfrak{r}, & \text{if } \underline{n} = \mathfrak{s} \rightarrow \mathfrak{r} \end{cases} \end{aligned}$$

Moreover, we shall use the following functions:

- $\text{parent}_T(n)$ returning ε if $n = T^\bullet$, the parent of n in T if $n \in T$, and \perp otherwise.
- $n \uparrow_T$ returning the path from T^\bullet to n (including n) if $n \in T$, and \perp otherwise.

Given $T \in \mathcal{T}$, let $A(T)$ be the global assertion obtained by appending the labels of the nodes in (depth-first) preorder traversal visit of T .

Proposition 4.1. $A(\mathbb{T}(\mathcal{G})) = \mathcal{G}$

Proposition 4.1 allows us to extend $\text{knows}_p(-)$ to \mathcal{T} by $\text{knows}_p(T) \stackrel{\text{def}}{=} \text{knows}_p(A(T))$.

Proposition 4.2. *If $T \in \mathcal{T}$ then $\mathbb{T}(A(T)) = T$*

Propositions 4.1 and 4.2 – whose proofs are trivial – basically induce an isomorphism between global assertions and their parsing trees.

4.3 On Recovering History Sensitivity

In a choreography, participants have to make local choices on the communicated values; such choices have an impact on the graceful coordination of the distributed parties. It is therefore crucial that the responsible party has “enough information” to commit to an

“appropriate” local choice, at each point of the choreography. For global assertions, this distills into *history sensitivity* (HS), a property defined in [12] as follows:

A predicate guaranteed by a participant s can only contain those interaction variables that s knows.

HS demands the sender/selector of each interaction in a given assertion to know all the variables involved in the predicate associated to that interaction. We illustrate HS with the following example.

Example 4.3. The global assertion $\mathcal{G}_{\text{ex4.3}}$ below violates HS.

$$\begin{aligned} \mathcal{G}_{\text{ex4.3}} &= \text{Alice} \rightarrow \text{Bob} : \{v_1 \mid v_1 > 0\}; \\ &\quad \text{Bob} \rightarrow \text{Carol} : \{v_2 \mid v_2 > 0\}; \\ &\quad \text{Carol} \rightarrow \text{Alice} : \{v_3 \mid v_3 > v_1\} \end{aligned}$$

Carol’s obligation $v_3 > v_1$ cannot be fulfilled because $v_1 \notin \text{knows}_{\text{Carol}}(\mathcal{G}_{\text{ex4.3}})$. \diamond

Given a global assertion \mathcal{G} , the function $\overline{\text{HS}}(\mathcal{G})$ below returns the nodes of $\text{T}(\mathcal{G})$ where HS is violated

$$\overline{\text{HS}}(\mathcal{G}) \stackrel{\text{def}}{=} \left\{ n \in \text{T}(\mathcal{G}) \mid \text{var}(\text{cst}_T(n)) \not\subseteq \text{knows}_s(n \uparrow_T) \text{ and } s = \text{resp}_T(\mathcal{G})(n) \right\}$$

where $\text{resp}_T(-) : \mathcal{N} \rightarrow \mathbb{P}$ yields the responsible party of a node and is defined as

$$\text{resp}_T(n) \stackrel{\text{def}}{=} \begin{cases} \text{snd}_T(n), & \text{if } \underline{n} = \mathbf{1} \\ \text{snd}_T(\text{parent}_T(n)), & \text{if } \underline{n} = \{\Psi\} \\ \perp, & \text{otherwise} \end{cases}$$

Intuitively, to determine whether a node $n \in \text{T}(\mathcal{G})$ violates HS, one checks if the responsible party of n knows all the variables involved in $\text{cst}_T(\mathcal{G})(n)$.

Given $T \in \mathcal{T}$, $\text{varHS}_T(-)$ is a function from the nodes of T to sets of variables and is defined as follows.

$$\text{varHS}_T(n) \stackrel{\text{def}}{=} \text{var}(\text{cst}_T(n)) \setminus \text{knows}_s(n \uparrow_T) \quad \text{where } s = \text{resp}_T(n)$$

Namely, $\text{varHS}_T(n)$ yields the variables of n not known to the responsible party of n . It is a simple observation that if HS is violated at a node n , then there exists a variable in the predicate of n which is not known to the responsible party of n (in fact, $n \in \overline{\text{HS}}(\mathcal{G})$ iff $\text{varHS}_T(n) \neq \emptyset$).

Example 4.4. Consider the following global assertion:

$$\begin{aligned} \mathcal{G}_{\text{ex4.4}} &= \mu\chi\langle 10 \rangle \{v \mid v > 0\}. \\ &\quad \text{Alice} \rightarrow \text{Bob}: \{v_1 \mid v \geq v_1\}; \\ &\quad \text{Bob} \rightarrow \text{Carol}: \{v_2 \mid v_2 > v_1\}; \\ &\quad \text{Carol} \rightarrow \text{Alice}: \{v_3 \mid v_3 > v_1\}; \\ &\quad \text{Carol} \rightarrow \text{Bob}: \{v_4 \mid v_4 > v\}; \\ &\quad \chi\langle v_1 \rangle \end{aligned}$$

$\overline{\text{HS}}(\mathcal{G}_{\text{ex4.4}}) = \{n_3, n_4\}$ where n_3 and n_4 are the nodes in $\text{T}(\mathcal{G}_{\text{ex4.4}})$ corresponding to the third and fourth interactions of $\mathcal{G}_{\text{ex4.4}}$, i.e., $n_3 = \text{Carol} \rightarrow \text{Alice}: \{v_3 \mid v_3 > v_1\}$ and $n_4 = \text{Carol} \rightarrow \text{Bob}: \{v_4 \mid v_4 > v\}$.

Carol is responsible for both violations (i.e., $\text{resp}_{\text{T}(\mathcal{G}_{\text{ex4.4}})}(n_3) = \text{resp}_{\text{T}(\mathcal{G}_{\text{ex4.4}})}(n_4) = \text{Carol}$). The violation in n_3 is on $\text{varHS}_{\text{T}(\mathcal{G}_{\text{ex4.4}})}(n_3) = \{v_1\}$ (i.e., Carol has to choose v_3 so that $v_3 > v_1$ without knowing v_1) and the violation in n_4 is on $\text{varHS}_{\text{T}(\mathcal{G}_{\text{ex4.4}})}(n_4) = \{v\}$ (i.e., Carol has to choose v_4 so that $v_4 > v$ without knowing v). Note that the violation of HS above does not imply that Carol will actually violate the condition $v_3 > v_1$. In fact, Carol could unknowingly choose either a violating or a non violating value for v_3 . \diamond

In Section 4.3.1 and Section 4.3.2, we present two algorithms that fix, when possible, HS violations in a global assertion. We discuss and compare their applicability, as well as the relationship between the amended global assertion and the original one. We shall use Example 4.4 as the running example of Section 4.3.1 and Section 4.3.2.

4.3.1 Strengthening

Throughout this section we fix a global assertion \mathcal{G} and its assertion tree $T = \mathsf{T}(\mathcal{G})$ and assume HS is violated at $n \in T$ with $\mathsf{cst}_T(n) = \psi$ and $\mathsf{resp}_T(n) = \mathfrak{s}$.

Violations occur when the responsible party \mathfrak{s} of n ignores at least one variable $v \in \mathsf{var}(\psi)$. The strengthening algorithm (cf. Definition 4.4) replaces ψ in \mathcal{G} with a predicate $\psi\{v'/v\}$ (if any) such that:

- (i) v' is a variable that \mathfrak{s} knows,
- (ii) if $\psi\{v'/v\}$ and the predicates occurring from T^\bullet to $\mathsf{parent}_T(n)$ are satisfied, then ψ is also satisfied.

Intuitively, the method above *strengthens* ψ by replacing it with $\psi\{v'/v\}$ so that: due to (i) the presence of variable v , which is unknown to the sender/selector, is removed, and due to (ii) ψ can still be guaranteed. In fact, relying on the information provided by all the predicates occurring before n , if the sender/selector guarantees $\psi\{v'/v\}$ then they also guarantee ψ . If there is no variable v' that ensures (i) and (ii) then we say that *strengthening is not applicable*.

Let $\mathsf{PRED}_T : \mathcal{N} \rightarrow \Psi$ yield the conjunction of the predicates on the path from T^\bullet to the

parent of a node:

$$\text{PRED}_T(n) \stackrel{\text{def}}{=} \begin{cases} \perp, & \text{if } \text{parent}_T(n) = \perp \\ \mathbf{true}, & \text{if } \text{parent}_T(n) = \varepsilon \\ \text{cst}_T(\text{parent}_T(n)) \wedge \text{PRED}_T(\text{parent}_T(n)), & \text{otherwise} \end{cases}$$

The function $\text{strengthen}(\mathcal{G})$ uses PRED_T to compute a global assertion \mathcal{G}' by replacing in \mathcal{G} , if possible, the predicate violating HS by a stronger one.

Definition 4.3 (*strengthen*). *The function strengthen is defined as follows.*

- If $\overline{\text{HS}}(\mathcal{G}) = \emptyset$ then $\text{strengthen}(\mathcal{G})$ returns \mathcal{G} .
- If $n \in \overline{\text{HS}}(\mathcal{G})$, $v \in \text{varHS}_T(n)$ and there exists $v' \in \text{knows}_s(n \uparrow_T)$ such that

$$\text{PRED}_T(n) \wedge \Psi\{v'/v\} \implies \Psi \quad \text{with} \quad \Psi = \text{cst}_T(n) \quad (4.3)$$

then $\text{strengthen}(\mathcal{G})$ returns $A(T')$ where T' is obtained from T by replacing Ψ with $\Psi\{v'/v\}$ in \underline{n} .

- When condition (4.3) does not hold for any $v \in \text{knows}_s(n \uparrow_T)$, $\text{strengthen}(\mathcal{G})$ returns $\mathcal{G}^{\perp n}$, indicating that \mathcal{G} violates HS at $n \in \overline{\text{HS}}(\mathcal{G})$.

Remark 4.4. *We assume that variables that are not under the scope of a quantifier are implicitly universally quantified.*

The algorithm Σ in Definition 4.4 below recursively applies $\text{strengthen}(_)$ until either the global assertion satisfies HS or Σ is not applicable anymore.

Definition 4.4 (Σ). *The algorithm Σ is defined as follows*

$$\Sigma(\mathcal{G}) \stackrel{\text{def}}{=} \begin{cases} \text{strengthen}(\mathcal{G}), & \text{if } \text{strengthen}(\mathcal{G}) \in \{\mathcal{G}, \mathcal{G}^{\perp n}\} \\ \Sigma(\text{strengthen}(\mathcal{G})), & \text{otherwise} \end{cases}$$

Example 4.5. Consider $\mathcal{G}_{\text{ex4.4}}$ from Example 4.4 and recall that $\overline{\text{HS}}(\mathcal{G}_{\text{ex4.4}}) = \{n_3, n_4\}$. Strengthening is applicable to n_3 since by substituting v_1 with v_2 in $v_3 > v_1$ (with $v_2 \in \text{knows}_{\text{Carol}}(n_3 \uparrow_{\text{T}}(\mathcal{G}_{\text{ex4.4}}))$) we have that condition (4.3) in Definition 4.3 holds:

$$(v > 0 \wedge v \geq v_1 \wedge v_2 > v_1) \wedge (v_3 > v_2) \implies (v_3 > v_1)$$

The invocation of $\text{strengthen}(\mathcal{G}_{\text{ex4.4}})$ returns (by substituting v_1 with v_2 in n_3):

$$\begin{aligned} \mathcal{G}_{\text{ex4.5}} &= \mu\chi \langle 10 \rangle \{v \mid v > 0\}. \\ &\quad \text{Alice} \rightarrow \text{Bob}: \{v_1 \mid v \geq v_1\}; \\ &\quad \text{Bob} \rightarrow \text{Carol}: \{v_2 \mid v_2 > v_1\}; \\ &\quad \text{Carol} \rightarrow \text{Alice}: \{v_3 \mid v_3 > v_2\}; \\ &\quad \text{Carol} \rightarrow \text{Bob}: \{v_4 \mid v_4 > v\}; \\ &\quad \chi \langle v_1 \rangle \end{aligned}$$

The invocation of $\text{strengthen}(\mathcal{G}_{\text{ex4.5}})$ returns $\mathcal{G}_{\text{ex4.5}}^{\perp n_4}$ since $\mathcal{G}_{\text{ex4.5}}$ has still one violating node n_4 for which strengthening is not applicable. In fact, $\text{knows}_{\text{Carol}}(n_4 \uparrow_{\text{T}}(\mathcal{G}_{\text{ex4.5}})) = \{v_2, v_3\}$ and:

- by substituting v with v_2 , condition (4.3) in Definition 4.3 does not hold since:

$$(v > 0 \wedge v \geq v_1 \wedge v_2 > v_1 \wedge v_3 > v_2) \wedge (v_4 > v_2) \not\implies (v_4 > v)$$

- by substituting v with v_3 , condition (4.3) in Definition 4.3 does not hold since:

$$(v > 0 \wedge v \geq v_1 \wedge v_2 > v_1 \wedge v_3 > v_2) \wedge (v_4 > v_3) \not\implies (v_4 > v)$$

◇

4.3.2 Variable Propagation

An alternative approach to solve HS problems is based on the modification of global assertions by letting responsible parties of the violating nodes know the variables causing the violation. The idea is that such variables are propagated within a “chain of interactions”.

Definition 4.5 ($<_T$ -chain and propagation). *Let $T \in \mathcal{T}$ and $n, n' \in T$, we write $n <_T n'$ iff n appears in $n' \uparrow_T$ and $\text{rcv}_T(n) = \text{snd}_T(n')$. A vector of nodes n_1, \dots, n_t is a $<_T$ -chain iff $n_i <_T n_{i+1}$ for all $1 \leq i < t$.*

Given a $<_T$ -chain $\vec{n} = n_1 \cdots n_t$ in T and $v_0 \in \text{var}_T(n_t)$, let the propagation of v_0 in \vec{n} – written $\text{P}_T(v_0, \vec{n})$ – be the tree T' obtained by updating the nodes in T as follows:

- *for $1 \leq i < t$, $\text{var}_{T'}(n_i) = \text{var}_T(n_i) \cup v_i$ and $\text{cst}_{T'}(n_i) = \text{cst}_T(n_i) \wedge (v_i = v_{i-1})$, with $v_1, \dots, v_{t-1} \in \mathcal{V}$ fresh and pairwise distinct,*
- *$\text{cst}_{T'}(n_t) = \text{cst}_T(n_t) \{v_{t-1}/v_0\}$,³*
- *all the other nodes of T remain unchanged.*

Example 4.6. In the global assertion $\mathcal{G}_{\text{ex4.6}}$ below assume Alice knows v from previous interactions (the ellipsis in $\mathcal{G}_{\text{ex4.6}}$).

$$\mathcal{G}_{\text{ex4.6}} = \dots \text{ Alice} \rightarrow \text{Bob}: \{u_1 \mid \Psi_1\}; \quad (1)$$

$$\text{Bob} \rightarrow \text{Carol}: \{u_2 \mid \Psi_2\}; \quad (2)$$

$$\text{Bob} \rightarrow \text{Dave}: \{u_3 \mid \Psi_3\}; \quad (3)$$

$$\text{Dave} \rightarrow \text{Alice}: \{u_4 \mid u_4 > v\} \quad (4)$$

For the $<_T$ -chain $\vec{n} = n_1 n_3 n_4$ in $\text{T}(\mathcal{G}_{\text{ex4.6}})$ – where n_i corresponds to line (i) in $\mathcal{G}_{\text{ex4.6}}$ –

³Note that $t > 1$.

$P_T(\mathcal{G}_{\text{ex4.6}})(v, \vec{n})$ returns T' such that $A(T')$ is

$$\mathcal{G}'_{\text{ex4.6}} = \dots \text{ Alice} \rightarrow \text{Bob}: \{u_1 \ v_1 \mid \Psi_1 \wedge v = v_1\}; \quad (1)$$

$$\text{Bob} \rightarrow \text{Carol}: \{u_2 \mid \Psi_2\}; \quad (2)$$

$$\text{Bob} \rightarrow \text{Dave}: \{u_3 \ v_2 \mid \Psi_3 \wedge v_1 = v_2\}; \quad (3)$$

$$\text{Dave} \rightarrow \text{Alice}: \{u_4 \mid u_4 > v_2\} \quad (4)$$

◇

Hereafter, we fix a global assertion \mathcal{G} . Let $T = T(\mathcal{G})$, $n \in \overline{\text{HS}}(\mathcal{G})$, $v \in \text{varHS}_T(n)$, and $s = \text{resp}_T(n)$. The *propagation* algorithm (cf. Definition 4.7) is *applicable* only if there exists a $<_T$ -chain in $n \uparrow_T$ through which v can be propagated from a node whose sender knows v to a node where s is a receiver.

We define a function *propagate* which takes a global assertion \mathcal{G} and returns \mathcal{G} itself if HS is satisfied, $\mathcal{G}^{\perp n}$ if HS is violated at $n \in T(\mathcal{G})$ and propagation is not applicable, and \mathcal{G}' otherwise, where \mathcal{G}' is obtained by propagating a violating variable v of node n . In the latter case, observe that v has necessarily been introduced in a node $n' \in n \uparrow_{T(\mathcal{G})}$ from which v can be propagated, since we assume \mathcal{G} closed.

Definition 4.6 (*propagate*). *The function $\text{propagate}(\mathcal{G})$ returns*

- \mathcal{G} , if $\overline{\text{HS}}(\mathcal{G}) = \emptyset$,
- $P_T(v, \vec{n})$, if $T = T(\mathcal{G})$ and there exists $n \in \overline{\text{HS}}(\mathcal{G})$, $v \in \text{varHS}_T(n)$, and $\vec{n} = n_0 \vec{n}_1 \dots n$ $<_T$ -chain in T such that $\text{snd}_T(n_0)$ knows v ,
- $\mathcal{G}^{\perp n}$ with $n \in \overline{\text{HS}}(\mathcal{G})$, otherwise.

Example 4.7. Consider again the global assertion $\mathcal{G}_{\text{ex4.5}}$ obtained after the invocation $\text{strengthen}(\mathcal{G}_{\text{ex4.4}})$ in Example 4.5. In this case $\overline{\text{HS}}(\mathcal{G}_{\text{ex4.5}}) = \{n_4\}$ with $\underline{n}_4 = \text{Carol} \rightarrow$

Bob: $\{v_4 \mid v_4 > v\}$. Propagation is applicable to n_4 and $\text{propagate}(\mathcal{G}_{\text{ex4.5}})$ returns

$$\begin{aligned} \mathcal{G}_{\text{ex4.7}} &= \mu\chi\langle 10 \rangle \{v \mid v > 0\}. \\ &\text{Alice} \rightarrow \text{Bob}: \{v_1 \mid v \geq v_1\}; \\ &\text{Bob} \rightarrow \text{Carol}: \{v_2 \ u_1 \mid v_2 > v_1 \wedge u_1 = v\}; \\ &\text{Carol} \rightarrow \text{Alice}: \{v_3 \mid v_3 > v_2\}; \\ &\text{Carol} \rightarrow \text{Bob}: \{v_4 \mid v_4 > u_1\}; \\ &\chi\langle v_1 \rangle \end{aligned}$$

by propagating v from the second interaction – where the sender Bob knows v – to Carol, $\mathcal{G}_{\text{ex4.7}}$ satisfies HS. The predicate of the last interaction originates from the substitution $(v_4 > v) \{u_1/v\}$. \diamond

The propagation algorithm is defined below and is based on a repeated application of $\text{propagate}(\cdot)$.

Definition 4.7 (Π). *Given a global assertion \mathcal{G} , the function Π is defined as follows:*

$$\Pi(\mathcal{G}) = \begin{cases} \text{propagate}(\mathcal{G}), & \text{if } \text{propagate}(\mathcal{G}) \in \{\mathcal{G}, \mathcal{G}^{\perp n}\} \\ \Pi(\text{propagate}(\mathcal{G})), & \text{otherwise} \end{cases}$$

Remark 4.5. *In distributed applications it is often necessary to guarantee that exchanged information is accessible only to intended participants. It is worth observing that Π discloses information about the propagated variable to the participants involved in the propagation chain. The architect should therefore evaluate whether it is appropriate to use Π .*

One could think of an extension of $\text{propagate}(\mathcal{G})$ which propagates variables only to participants entitled to know them. The existence of a propagation chain \vec{n} for a variable v may be parametrised by two sets of participants chosen by the architect: a set A containing the participants who are allowed to know the value of v and a set N of participants

not allowed to know it. Let $T = \mathbb{T}(\mathcal{G})$, and v_0 a variable causing an HS problem in \mathcal{G} , an acceptable chain \vec{n} is defined as in Section 4.3.2 and such that

- for all $n \in \vec{n}$, $\text{rcv}_T(n) \notin N$, and
- there is no other $<_T$ -chain \vec{n}' for which $\text{propagate}(_)$ is applicable such that $|P(\vec{n}')| < |P(\vec{n})|$, where

$$P(\vec{n}) = \{\mathfrak{r} \mid \text{there exist } n \in \vec{n} \text{ such that } \mathfrak{r} = \text{rcv}_T(n) \text{ and } \mathfrak{r} \notin A\}$$

Note that even though this additional condition provides a more fine-grained control on the way problems are solved, it also decreases the range of applicability of the algorithm since the existence of such a chain is not guaranteed.

4.3.3 Properties of Σ and Π

We now discuss the properties of the global assertions amended by each algorithm and we compare them. Hereafter, we say Σ (resp. Π) returns \mathcal{G} if it returns either \mathcal{G} or $\mathcal{G}^{\perp n}$ for some n .

The applicability of Σ depends on whether it is possible to find a variable known to the responsible party of the violating node such that condition (4.3) in Definition 4.3 is satisfied. The applicability of Π depends on whether there exists a chain through which the problematic variable can be propagated. Observe also that there are cases in which Σ is applicable and Π is not, and vice versa, e.g., Example 4.5. Moreover, $\Sigma(\mathcal{G}) \neq \Pi(\mathcal{G})$ in general, hence it may not always be clear which one should be preferred.

We first show that our algorithms preserve the structure of the initial global assertion.

Lemma 4.1. *Let \mathcal{G} be a global assertion; $\text{strengthen}(\mathcal{G})$ (resp. $\text{propagate}(\mathcal{G})$) always returns \mathcal{G}' such that $\mathbb{T}(\mathcal{G}')$ is isomorphic to $\mathbb{T}(\mathcal{G})$ up-to labels, namely it has the same tree structure (but possibly different node labels).*

Proof. By Definition 4.3 $\text{strengthen}(\mathcal{G})$ always returns either \mathcal{G} (which includes the case for $\mathcal{G}^{\perp n}$) or \mathcal{G}' . $T(\mathcal{G}')$ is isomorphic to $T(\mathcal{G})$ as the two trees either are the same or only differ in the label of one node.

By Definition 4.6, $\text{propagate}(\mathcal{G})$ always returns either \mathcal{G} (which includes the case for $\mathcal{G}^{\perp n}$) or $\mathcal{G}' = P_{T(\mathcal{G})}(v_0, \vec{n})$ for some \prec_T -chain \vec{n} and variable v_0 . By definition of $P_{T(\mathcal{G})}(v_0, \vec{n})$, $T(\mathcal{G}')$ only differs from $T(\mathcal{G})$ in the labels of the nodes in \vec{n} , hence $P_{T(\mathcal{G})}(v_0, \vec{n})$ is isomorphic to $T(\mathcal{G})$. \square

Proposition 4.3. *Let \mathcal{G} be a global assertion. If it terminates, $\Sigma(\mathcal{G})$ (resp. $\Pi(\mathcal{G})$) returns \mathcal{G}' such that $T(\mathcal{G}')$ is isomorphic to $T(\mathcal{G})$.*

Proof. By straightforward induction on the number of invocations of $\text{strengthen}(\mathcal{G})$ (resp. $\text{propagate}(\mathcal{G})$) in Σ (resp. Π), see [14] for more details. \square

Lemma 4.2 below shows that the number of HS problems in a global assertion decreases after an invocation of either of our algorithm (if applicable).

Lemma 4.2. *Let \mathcal{G} be a global assertion, $T = T(\mathcal{G})$, and k be the number of HS violations in T .⁴ Let k_1, k_2 be the number of HS violations in $T_1 = T(\text{strengthen}(\mathcal{G}))$ and $T_2 = T(\text{propagate}(\mathcal{G}))$, respectively; then either $T_i = T$ or $k_i = k - 1$ with $i \in \{1, 2\}$.*

Proof. By Proposition 4.3, T and T_1 are isomorphic. Let $n_1 \in T_1$ be the node corresponding to $n \in T$. By definition of $\overline{\text{HS}}$ the violation in a node $n' \in T_1$ is defined only in terms of the nodes in $n' \uparrow_{T_1}$. By definition of $\text{strengthen}(_)$ the only node from which T differs from T_1 is n_1 . Hence, if a violation is added in T_1 with respect to T it must be in the subtree of T_1 rooted at n_1 . However, a violation is not added in n_1 itself since the variable chosen to replace the problematic one is selected so that the responsible party *knows* it. No violation can be added in the subtree rooted at n' since $\text{strengthen}(_)$ does not modify the variables known by the participants (but only the predicates). Thus, either $T_1 = T$ or $k_1 = k - 1$.

⁴Note that more than one violation may occur in one node if the sender does not know several variables.

In the case $T_2 = T(\text{propagate}(\mathcal{G}))$, assume $\hat{n} \in \overline{\text{HS}}(T)$ with $\hat{n} = s \rightarrow r : \{\vec{x} \mid \phi\}$, $v \in \text{varHS}_T(\hat{n})$, and there exists an $<_T$ -chain $n_0 \vec{n} \hat{n}$ with $\underline{n_0} = s_0 \rightarrow r_0 : \{\vec{y}_0 \mid \Psi_0\}$ such that s_0 knows v . We proceed by case analysis showing that, in any node $n \in T$, no violation is introduced, and exactly one violation is removed from \hat{n} .

- if $n = n_0$ no violation is introduced since $\underline{n_0}$ becomes $s_0 \rightarrow r_0 : \{\vec{y} \ v_0 \mid \Psi_0 \wedge v_0 = v\}$ in T_2 where, by Definition 4.6, s_0 knows v .
- if $n_i \in \vec{n}$ by definition of propagation $\underline{n_i} = s_i \rightarrow r_i : \{\vec{y}_i \mid \Psi_i\}$ becomes $s_i \rightarrow r_i : \{\vec{y}_i \ v_i \mid \Psi_i \wedge v_{i-1} = v_i\}$ in T_2 where, by Definition 4.5, $r_{i-1} = s_i$. It follows that s_i has previously received v_{i-1} hence s_i knows it.
- if $n = \hat{n}$, following Definition 4.5, $\text{cst}_{T_2}(\hat{n}) = \phi \{v_0/v\}$ and the problem on v at \hat{n} has been solved since s knows v_0 , and v does not appear in ϕ anymore, hence $k_2 = k - 1$.
- if n does not belong to the $<_T$ -chain $n_0 \vec{n} \hat{n}$ then n remains unchanged and no violation is introduced. Note that if n is in the subtree rooted at \hat{n} , no violation is introduced since propagation does not decrease the knowledge of any participant. \square

Lemma 4.2 implies directly that both algorithms terminate.

Corollary 4.1. *Let \mathcal{G} be a global assertion; $\Sigma(\mathcal{G})$ and $\Pi(\mathcal{G})$ terminate.*

Whereas Σ does not change the global type underlying the global assertion, Π does. Indeed, in the resulting global assertion, more variables are exchanged in each interaction involved in the propagation. However, the structure of the tree remains the same.

Let $\text{erase}(\mathcal{G})$ be the function that returns the underlying global type [44] corresponding to \mathcal{G} (i.e., a global assertion without predicates).

Definition 4.8. Given a global assertion \mathcal{G} , $\text{erase}(\mathcal{G})$ is defined inductively as:

$$\text{erase}(\mathcal{G}) = \begin{cases} \mathbf{s} \rightarrow \mathbf{r} : \vec{\mathbf{a}}; \text{erase}(\mathcal{G}') & \text{if } \mathcal{G} = \mathbf{s} \rightarrow \mathbf{r} : \{\vec{\mathbf{v}} \mid \Psi\}; \mathcal{G}' \\ \mathbf{s} \rightarrow \mathbf{r} : \{l_j : \text{erase}(\mathcal{G}_j)\}_{j \in J} & \text{if } \mathcal{G} = \mathbf{s} \rightarrow \mathbf{r} : \{\{\Psi_j\}_{l_j} : \mathcal{G}_j\}_{j \in J} \\ \mu\chi.\text{erase}(\mathcal{G}') & \text{if } \mathcal{G} = \mu\chi \langle \vec{e} \rangle \{\vec{\mathbf{v}} \mid \Psi\}. \mathcal{G}' \\ \mathcal{G} & \text{if } \mathcal{G} = \mathbf{0} \text{ or } \mathcal{G} = \chi \end{cases}$$

where $\vec{\mathbf{a}}$ is a list of sorts, so that the length of $\vec{\mathbf{v}}$ is the same as $\vec{\mathbf{a}}$. If $\vec{\mathbf{a}} = a_1 \dots a_n$ and $\vec{\mathbf{v}} = v_1 \dots v_n$, then for $1 \leq i \leq n$, a_i is the sort of variable v_i .

Lemma 4.3. Let \mathcal{G} and \mathcal{G}' be two assertions that differ only in the predicates annotating interactions and branching constructs. Then $\text{erase}(\mathcal{G}) = \text{erase}(\mathcal{G}')$.

Proof. The proof is by straightforward structural induction on \mathcal{G} (cf. [14]). \square

Lemma 4.4. Let \mathcal{G} be a global assertion and $T = \mathbb{T}(\mathcal{G})$. Given a $<_T$ -chain $\vec{n} = n_1, \dots, n_t$ in T , if $\mathbb{P}_T(v_0, \vec{n}) = \mathcal{G}'$ for an interaction variable v_0 , then for all $n \in \mathbb{T}(\mathcal{G})$ and its corresponding node $n' \in \mathbb{T}(\mathcal{G}')$,

$$\text{var}_{\mathbb{T}(\mathcal{G})}(n) \subseteq \text{var}_{\mathbb{T}(\mathcal{G}')} (n')$$

Proof. Observe that, by Proposition 4.3, T is isomorphic to T' . For each $n \in T$, let n' denote the node in T' that corresponds to n ; and let $\vec{n} = n_1, \dots, n_t$. We proceed by case analysis on the (labels of the) nodes of T , which we divide in three groups according to the definition of $\mathbb{P}_T(v_0, \vec{n})$:

- for $i \in \{1, \dots, t-1\}$, $\text{var}_T(n_i) = \text{var}_{T'}(n'_i) v_i$.
- $\text{var}_T(n_t) = \text{var}_{T'}(n'_t)$.
- if $n \in T$ and $n \notin \vec{n}$, then n' has the same label as n since by definition of $\mathbb{P}_T(v_0, \vec{n})$ all the nodes that are not in \vec{n} are unchanged. \square

Proposition 4.4 (Underlying type structure). *Let \mathcal{G} be a global assertion,*

- *if $\Sigma(\mathcal{G})$ returns \mathcal{G}' then $\text{erase}(\mathcal{G}) = \text{erase}(\mathcal{G}')$*
- *if $\Pi(\mathcal{G})$ returns \mathcal{G}' then for all $n \in \mathsf{T}(\mathcal{G})$ and its corresponding node $n' \in \mathsf{T}(\mathcal{G}')$,*

$$\text{var}_{\mathsf{T}(\mathcal{G})}(n) \subseteq \text{var}_{\mathsf{T}(\mathcal{G}')} (n')$$

Proof. Consider Σ ; we observe that $\text{strengthen}(\mathcal{G})$ either returns \mathcal{G} or a \mathcal{G}' that differs from \mathcal{G} only in the predicate of one (interaction or branching) node. By Lemma 4.3:

$$\text{erase}(\mathcal{G}) = \text{erase}(\text{strengthen}(\mathcal{G}))$$

The rest follows by straightforward induction on the number of applications of Σ , see [14] for more details.

Consider Π ; we observe that, by Lemma 4.4, if $\text{propagate}(\mathcal{G}) = \mathcal{G}'$ then for all $n \in \mathsf{T}(\mathcal{G})$ and its corresponding node $n' \in \mathsf{T}(\mathcal{G}')$,

$$\text{var}_{\mathsf{T}(\mathcal{G})}(n) \subseteq \text{var}_{\mathsf{T}(\mathcal{G}')} (n') \tag{4.4}$$

In fact, $\text{propagate}(\mathcal{G})$ either returns $\mathcal{G}' = \mathcal{G}$ or $\mathcal{G}' = \mathsf{P}_{\mathsf{T}(\mathcal{G})}(v_0, \vec{n})$. In the former case (4.4) holds trivially, in the latter case it holds by Lemma 4.4. The rest follows by straightforward induction on the number of applications of Π (cf. [14]). \square

The application of Σ and Π affects the predicates of the original global assertion. In Σ , strengthening allows less values for the interaction variables of the amended interaction. Conversely, the predicates computed by Π are equivalent to the original ones (i.e., they allow sender and receiver to choose/expect the same set of values). Nevertheless, such predicates are syntactically different as Π adds the equality predicates on the propagated variables.

Proposition 4.5 (Assertion predicates). *Let \mathcal{G} be a global assertion,*

1. *if $\Sigma(\mathcal{G})$ returns \mathcal{G}' then for all $n \in \mathsf{T}(\mathcal{G})$ whose label is modified by Σ , and its corresponding node $n' \in \mathsf{T}(\mathcal{G}')$ (cf. Proposition 4.4), it holds that*

$$\text{PRED}_{\mathsf{T}(\mathcal{G}')} (n') \wedge \text{cst}_{\mathsf{T}(\mathcal{G}')} (n') \implies \text{cst}_{\mathsf{T}(\mathcal{G})} (n)$$

2. *if $\Pi(\mathcal{G})$ returns \mathcal{G}' then for all $n \in \mathsf{T}(\mathcal{G})$ whose label is modified by Π , and its corresponding node $n' \in \mathsf{T}(\mathcal{G}')$*

(a) *$\text{cst}_{\mathsf{T}(\mathcal{G}')} (n')$ is the predicate $\text{cst}_{\mathsf{T}(\mathcal{G})} (n)\sigma \wedge \Psi$*

(b) *$\text{PRED}_{\mathsf{T}(\mathcal{G})} (n) \wedge \text{cst}_{\mathsf{T}(\mathcal{G})} (n) \wedge \Psi \iff \text{PRED}_{\mathsf{T}(\mathcal{G}')} (n') \wedge \text{cst}_{\mathsf{T}(\mathcal{G}')} (n')$*

For some satisfiable $\Psi \in \Psi$ and variable substitution σ .

Proof. Item 1. The proof relies on the fact that Σ either does not change \mathcal{G} or replaces a problematic variable by a variable for which (4.3) holds. We show the result by showing that it holds for each invocation of $\text{strengthen}(_)$ by Σ . Indeed, for each invocation we have that, by Definition 4.3, if n is modified by Σ , then we have that $n \in \overline{\text{HS}}(\mathsf{T}(\mathcal{G}))$. In addition, there must be $v \in \text{varHS}_{\mathsf{T}(\mathcal{G})}(n)$ such that there exists $v' \in \text{knows}_{\mathsf{s}}(n \uparrow_{\mathsf{T}(\mathcal{G})})$ and (4.3) is satisfied. This gives us

$$\underbrace{\text{PRED}_{\mathsf{T}(\mathcal{G})} (n)}_{\text{PRED}_{\mathsf{T}(\mathcal{G}')} (n')} \wedge \underbrace{\Psi \{v'/v\}}_{\text{cst}_{\mathsf{T}(\mathcal{G}')} (n')} \implies \underbrace{\Psi}_{\text{cst}_{\mathsf{T}(\mathcal{G})} (n)}$$

Since only the predicate of node n is updated by substituting v by v' , by Definition 4.3.

Item 2. The proof relies on the definition of $\text{P}_{\mathsf{T}(\mathcal{G})}(_, _)$, i.e., a predicate of the form $v_1 = v_0$ or $v_i = v_{i-1}$ is added to each predicate of the nodes in the chain, and problematic variables are replaced by fresh ones. The additional predicates are satisfiable since

they constrain only fresh variables (i.e., v_i). We have these results by showing that each invocation of $\text{propagate}(_)$ by Π validates the result.

Item 2a. If n is modified by $\text{propagate}(_)$, then $n \in \vec{n}$, by Definition 4.6. Assume v_0 is the variable to be propagated.

- If n is not the last node of \vec{n} , by definition of $P_{T(\mathcal{G})}(v_0, \vec{n})$, we have that $\text{cst}_{T(\mathcal{G}')} (n') = \text{cst}_{T(\mathcal{G})} (n) \wedge (v_i = v_{i-1})$, which gives us the expected result if ψ is $v_i = v_{i-1}$ and σ is the empty substitution.
- If n is the last node of \vec{n} , by definition of $P_{T(\mathcal{G})}(v_0, \vec{n})$, we have that $\text{cst}_{T(\mathcal{G}')} (n') = \text{cst}_{T(\mathcal{G})} (n) \{v_{i-1}/v_0\}$, which gives the expected result with $\sigma = \{v_{i-1}/v_0\}$ and $\psi = \text{true}$.

Item 2b. If n is modified by $\text{propagate}(_)$, then $n \in \vec{n}$, by Definition 4.6. Assume v_0 is the variable to be propagated and the length of \vec{n} is k .

- If n is the first node of \vec{n} , then

$$\text{cst}_{T(\mathcal{G}')} (n') = \text{cst}_{T(\mathcal{G})} (n) \wedge (v_1 = v_0) \quad \text{and} \quad \text{PRED}_{T(\mathcal{G}')} (n') = \text{PRED}_{T(\mathcal{G})} (n)$$

since $n' \uparrow_{T(\mathcal{G}')}$ is unchanged. We have the expected result if ψ is $v_1 = v_0$. Note that by definition of $P_{T(\mathcal{G})}(v_0, \vec{n})$, v_1 is a fresh variable therefore $v_1 = v_0$ is satisfiable.

- If n is the i^{th} node in \vec{n} ($1 < i < k$) then $\text{cst}_{T(\mathcal{G}')} (n') = \text{cst}_{T(\mathcal{G})} (n) \wedge (v_i = v_{i-1})$, and

$$\text{PRED}_{T(\mathcal{G}')} (n') = \text{PRED}_{T(\mathcal{G})} (n) \wedge \bigwedge_{1 < j < i} v_j = v_{j-1}$$

where each $v_j = v_{j-1}$ is satisfiable since each variable is freshly introduced. We have the expected result with ψ as $v_i = v_{i-1}$.

- If n is the last node in \vec{n} , then $\text{cst}_{\mathcal{T}(\mathcal{G}')} (n') = \text{cst}_{\mathcal{T}(\mathcal{G})} (n) \{v_k/v_0\}$, and

$$\text{PRED}_{\mathcal{T}(\mathcal{G}')} (n') = \text{PRED}_{\mathcal{T}(\mathcal{G})} (n) \wedge \bigwedge_{1 < j < k} v_j = v_{j-1}$$

with each $v_j = v_{j-1}$ satisfiable, as before. We have the required result with ψ set to **true**. \square

The statement 2 (b) in Proposition 4.5 amounts to saying that $\text{cst}_{\mathcal{T}(\mathcal{G})} (n) \wedge \psi$ is equivalent to the predicate $\text{cst}_{\mathcal{T}(\mathcal{G}')} (n')$ when considered in its respective context.

Remarkably, Σ and Π do not add violations (of either HS or TS) to the amended global assertions. We postpone the discussion of this property to Section 4.4.3 (Proposition 4.7) after the formal introduction of TS.

Finally, we prove that if the value returned by Σ or Π is not of the type $\mathcal{G}^{\perp n}$ then the amended global assertion satisfies HS.

Theorem 4.1 (Correctness). *If there is \mathcal{G}' such that $\Sigma(\mathcal{G}) = \mathcal{G}'$ or $\Pi(\mathcal{G}) = \mathcal{G}'$ then $\overline{\text{HS}}(\mathcal{G}') = \emptyset$.*

Proof. **Case Σ .** By Definitions 4.3 and 4.4, Σ terminates successfully when $\overline{\text{HS}}(\mathcal{G}) = \emptyset$. We show that at each iteration of Σ , the number of HS violations decreases. Assume that there is k violations in \mathcal{G} , by Definition 4.4, we have either

- $\Sigma(\mathcal{G}) = \text{strengthen}(\mathcal{G}) = \mathcal{G}$ in which case, by Definition 4.3, $\overline{\text{HS}}(\mathcal{G}) = \emptyset$, i.e., $k = 0$, and the function terminates, or
- $\Sigma(\mathcal{G}) = \Sigma(\text{strengthen}(\mathcal{G})) = \mathcal{G}'$ with $\mathcal{G} \neq \mathcal{G}'$, and by Lemma 4.2 the number of HS violation in $\text{strengthen}(\mathcal{G})$ is strictly less than k .

Case Π . The case for Π is similar to the previous case, using Definition 4.6 (resp. 4.7) instead of Definition 4.3 (resp. 4.4). \square

4.4 On Recovering Temporal Satisfiability

In a choreography, the local choices made by some parties may restrict later choices of other parties to the point that no suitable value is available. This would lead to an abnormal termination since the choreography cannot continue. For global assertions, this distills into *temporal satisfiability* (TS) which requires that the values sent in each interaction do not compromise the satisfiability of future interactions. The formal definition of temporal satisfiability is adapted from [12].

Definition 4.9 (TS [12]). *A global assertion \mathcal{G} satisfies TS, written $\text{TS}(\mathcal{G})$, if and only if $\text{GSat}(\mathcal{G}, \mathbf{true})$ holds where*

$$\text{GSat}(\mathcal{G}, \Psi) \text{ iff } \left\{ \begin{array}{ll} \text{GSat}(\mathcal{G}', \Psi \wedge \text{cst}(\mathfrak{t})), & \text{if } \mathcal{G} = \mathfrak{t}; \mathcal{G}' \text{ and } \Psi \implies \exists \text{var}(\mathfrak{t}) : \text{cst}(\mathfrak{t}) \\ \bigwedge_{j \in J} \text{GSat}(\mathcal{G}_j, \Psi \wedge \Psi_j), & \text{if } \mathcal{G} = \mathfrak{s} \rightarrow \mathfrak{r} : \{\{\Psi_j\}_{j \in J} : \mathcal{G}_j\}_{j \in J} \text{ and } \Psi \implies \bigvee_{j \in J} (\Psi_j) \\ \text{GSat}(\mathcal{G}', \Psi \wedge \Psi'), & \text{if } \mathcal{G} = \mu \chi \langle \vec{e} \rangle \{ \vec{v} \mid \Psi' \}. \mathcal{G}' \text{ and } \Psi \implies \Psi' \{ \vec{e} / \vec{v} \} \\ \text{GSat}(\mathcal{G}', \Psi \wedge \Psi'), & \text{if } \mathcal{G} = \chi_{\Psi'(\vec{v})} \langle \vec{e} \rangle \text{ and } \Psi \implies \Psi' \{ \vec{e} / \vec{v} \} \\ \mathcal{G} = \mathbf{0}, & \text{otherwise} \end{array} \right.$$

For an assertion tree $T \in \mathcal{T}$, $\text{TS}(T)$ holds iff $\text{GSat}(\mathbf{A}(T), \mathbf{true})$.

Intuitively, the predicate Ψ in Definition 4.9 is the conjunction of all the predicates that precede an interaction in \mathcal{G} . In the first case, all the values satisfying Ψ allow to instantiate the interaction variables $\text{var}(\mathfrak{t})$ so to satisfy the constraint $\text{cst}(\mathfrak{t})$ of \mathfrak{t} . For branching, GSat requires that at least one branch can be chosen and that each possible path satisfies GSat . For recursive definition, we require that the initial parameters satisfy the invariant Ψ' . We assume that each recursive call is annotated with the invariant of its corresponding

recursive definition, i.e., in Definition 4.9, $\psi'(\vec{v})$ is the predicate corresponding to the invariant of the definition of χ .

Often, TS problems appear when one tries to restrict the domain of a variable after its introduction. To illustrate this, we introduce the following running example.

Example 4.8. Consider $\mathcal{G}_{\text{ex4.8}}$ below, where s constrains x and y :

$$\begin{aligned} \mathcal{G}_{\text{ex4.8}} &= s \rightarrow r: \{x \mid x < 10\}; \\ & \quad s \rightarrow r: \{y \mid y > 8\}; \\ & \quad r \rightarrow s: \{z \mid x > z \wedge z > 6 \wedge y \neq z\} \end{aligned}$$

When r introduces z , both x and y are further restricted. $\mathcal{G}_{\text{ex4.8}}$ violates TS because it does not hold that (cf. Definition 4.9)

$$\forall xy: \underbrace{(x < 10) \wedge (y > 8)}_{\psi} \implies \underbrace{\exists z:}_{\exists \text{var}(t):} \underbrace{(x > z \wedge z > 6 \wedge y \neq z)}_{\text{cst}(t)}$$

Noticeably, if s chooses, e.g., $x = 6$ then r cannot choose a value for z . ◇

Possibly, TS can be regained by rearranging some predicates. In particular, we can “lift” a predicate to a previous interaction node. For instance, in Example 4.8, one could lift the predicate $\exists z: x > z > 6$ (adapted from the last interaction) to the first interaction’s predicate.

We first consider TS violations occurring in interactions and recursive definitions. Amending violations arising in branching and recursive calls is similar but complicates the presentation; for the sake of clarity, such violations are considered in Section 4.4.2.

4.4.1 Lifting Algorithm

We formalise the lifting algorithm. First, we give a function telling us whether a *node* n violates TS.

Definition 4.10 (TSnode). *Given $T \in \mathcal{T}$, $\text{TSnode}_T(n)$ holds iff $n \in T$ and $\text{TS}(n \uparrow_T)$ holds. In addition, we assume that TSnode holds for nodes with label of the form $s \rightarrow r$: (since there is no predicate in these nodes, no TS problem can arise).*

We now define a function that returns a set of nodes violating TS such that all the previous nodes in the tree do not violate TS.

Definition 4.11 ($\overline{\text{TS}}$). *The function $\overline{\text{TS}} : \mathcal{T} \rightarrow \mathcal{N}$ is defined as follows:*

$$\overline{\text{TS}}(T) \stackrel{\text{def}}{=} \left\{ n \in T \mid \begin{array}{l} \text{TSnode}_T(n) \text{ is false and } \text{TSnode}_T(n') \\ \text{is true for all } n' \in \text{parent}_T(n) \uparrow_T \end{array} \right\}$$

For instance, in Example 4.8, we have that $\overline{\text{TS}}(T_{\text{ex4.8}})$ is the singleton $\{n_{\text{ex4.8}}\}$ where $T_{\text{ex4.8}} = T(\mathcal{G}_{\text{ex4.8}})$ and $n_{\text{ex4.8}}$ is the node corresponding to the last interaction of $\mathcal{G}_{\text{ex4.8}}$.

Once an *interaction* node $n \in \overline{\text{TS}}(T)$ is chosen, the next step is to identify which part of its predicate is the source of the problem. To this purpose, we define a relation among predicates ψ and ϕ in a context ψ' .

Definition 4.12 (Conflict). *The predicate $\psi \in \Psi$ is in conflict on $\vec{v} \subseteq \mathcal{V}$ with ϕ in ψ' iff*

$$\psi' \implies \exists \vec{v} : \phi \quad \text{and} \quad \psi' \not\Rightarrow \exists \vec{v} : (\phi \wedge \psi)$$

The notion of conflict is based on the definition of TS for interaction nodes (Definition 4.9). On the one hand, the part ϕ of the predicate does not spoil TS and, on the other hand, the part ψ , in conjunction with ϕ , invalidates TS.

Example 4.9. In Example 4.8, we have

$$\forall xy : x < 10 \wedge y > 8 \implies \exists z : y \neq z$$

and

$$\forall xy : x < 10 \wedge y > 8 \not\Rightarrow \exists z : y \neq z \wedge x > z \wedge z > 6$$

That is to say that $x > z \wedge z > 6$ is the “problematic” part of the predicate. \diamond

Using Definition 4.12 and $\text{PRED}_T(n)$ (cf. Section 4.3.1), we define

$$\text{split}_T(n, \psi) \stackrel{\text{def}}{=} \left\{ \psi' \mid \begin{array}{l} \psi \iff \psi' \wedge \phi \text{ and } \psi' \text{ is in conflict on } \text{var}(n) \text{ with } \phi \\ \text{in } \text{PRED}_T(n) \end{array} \right\}$$

which returns a set of problematic predicates. Note that \iff denotes a logical equivalence between ψ and $\psi' \wedge \phi$; we assume that ψ' and ϕ are sub-formulae of ψ .

Example 4.10. Considering Examples 4.8 and 4.9, the application of split yields

$$\text{split}_{T_{\text{ex4.8}}}(n_{\text{ex4.8}}, z > 6 \wedge x > z \wedge y \neq z) = \{z > 6 \wedge x > z\}$$

since $y \neq z$ allows to choose a suitable value for z . \diamond

Remark 4.6. For a tree $T \in \mathcal{T}$ and $n \in \overline{\text{TS}}(T)$ such that $\psi' \in \text{split}_T(n, \psi)$, we may have $\text{PRED}_T(n) \not\Rightarrow \exists \vec{v}: \psi'$. For instance, if the predicate ψ' is not satisfiable, e.g., $\psi' = v < 7 \wedge v > 7$. In this case the lifting algorithm is not applicable.

Remark 4.7. Note that at this level, it is not necessary to require ψ' to be minimal in the definition of split (in terms of, e.g., the size of the formula or the number of variables in ψ'). Indeed, as stated later in Definition 4.14, only the predicates which can be lifted successfully are used by the algorithm. However, an implementation of the algorithm could minimise the predicate in order to maximise the efficiency of the lifting algorithm.

The next definition formalises the construction of a new assertion tree which possibly regains TS, given a node and an assertion to be “lifted” (i.e., a “problematic” predicate).

Definition 4.13 (build). *The function $\text{build}_T(n, \psi)$ returns:*

- $\hat{T} \in \mathcal{T}$, if ψ is satisfiable and we can construct \hat{T} isomorphic to T except that, each

node

$$n' \in (\text{parent}_T(n)) \uparrow_T \text{ such that } \underline{n'} = \mathbf{s} \rightarrow \mathbf{r} : \{\vec{u} \mid \theta\} \text{ with } \vec{u} \cap \text{var}(\psi) \neq \emptyset$$

is replaced by a node \hat{n} with label

$$\mathbf{s} \rightarrow \mathbf{r} : \{\vec{u} \mid \theta \wedge \forall \vec{x} : \exists \vec{y} : \psi\} \text{ such that } \theta \wedge \forall \vec{x} : \exists \vec{y} : \psi \text{ is satisfiable}$$

where

- $\vec{x} \subseteq \text{var}(\psi) \setminus \text{knows}_s(T)$ are introduced in a node in $n' \uparrow_T$
- $\vec{y} \subseteq \text{var}(\psi)$ are introduced in a node in the subtree rooted at n'

and there is no $n'' \in (\text{parent}_T(n)) \uparrow_T$ such that $\underline{n''} = \mu \chi \langle \vec{e} \rangle \{\vec{v} \mid \psi'\}$ with $\vec{v} \cap \text{var}(\psi) \neq \emptyset$.

- \perp otherwise.

Essentially, Definition 4.13 duplicates a quantified version of the predicate ψ in the nodes which introduce variables in $\text{var}(\psi)$. For each updated node n' , the quantification of the variables in $\text{var}(\psi)$ operates in the following way. The variables which are introduced before n' in the tree and which are not known to \mathbf{s} are quantified *universally* (since \mathbf{s} has no control over them). The variables that are introduced later in the tree are quantified *existentially*, so that \mathbf{s} may choose values for the variables in \vec{u} which do not compromise the satisfiability of predicates down in the tree.

Remark 4.8. *In the definition of build , we assume that if either \vec{x} or \vec{y} is empty, the corresponding unnecessary quantifier is removed. Recall that global assertions are closed (cf. Section 4.2). Therefore all the variables in $\text{var}(\psi)$ are either quantified in the predicate of \hat{n} , or have been introduced before n' .*

Continuing from Example 4.8, we would invoke $\text{build}_{T_{\text{ex4.8}}}(n_{\text{ex4.8}}, z > 6 \wedge x > z)$ which returns a new assertion tree. The new tree can be transformed into a global assertion isomorphic to $\mathcal{G}_{\text{ex4.8}}$ with the first line updated to: $\mathfrak{s} \rightarrow \mathfrak{r}: \{x \mid x < 10 \wedge \exists z: x > z > 6\}$.

The function $\text{TSres}: \mathcal{T} \times \mathcal{N} \rightarrow \mathcal{T} \cup \perp$ brings the above definitions together in order to either fix a TS problem at n , or return \perp .

Definition 4.14 (TSres). *Given $T \in \mathcal{T}$ and $n \in \overline{\text{TS}}(T)$, we define*

$$\text{TSres}_T(n) = \begin{cases} \text{build}_T(n, \psi), & \text{if } \underline{n} = \mathfrak{t} \text{ and there is } \psi \in \text{split}_T(n, \text{cst}_T(n)) \\ & \text{s.t. } \text{build}_T(n, \psi) \neq \perp \\ \text{build}_T(n, \psi\{\vec{e}/\vec{v}\}), & \text{if } \underline{n} = \mu\chi\langle\vec{e}\rangle\{\vec{v} \mid \psi\} \\ \perp, & \text{otherwise} \end{cases}$$

The first case of Definition 4.14 handles TS problems in interaction nodes. If there is a predicate ψ in conflict such that it can be “lifted” by build successfully, then the function returns the result of build. The second case handles TS violations in recursive definitions. The problem is similar to the interaction case, but in this case, the values assigned to the recursion parameters are known (i.e., \vec{e}). It may be possible to lift the recursion invariant, where we replace the recursion parameters by the corresponding initialisation vector. Example 4.11 illustrates this case.

Example 4.11. For the global assertion $\mathcal{G}_{\text{ex4.11}}$ given below, $\text{TS}(\mathcal{G}_{\text{ex4.11}})$ does not hold because $\forall xy: \mathbf{true} \not\Rightarrow (x > y > 6)$.

$$\begin{aligned} \mathcal{G}_{\text{ex4.11}} &= \mathfrak{s} \rightarrow \mathfrak{r}: \{x \mid \mathbf{true}\}; \\ &\quad \mu\chi\langle 8 \rangle \{y \mid x > y > 6\}. \mathcal{G}' \end{aligned}$$

However, using the initialisation parameters, we can lift $x > 8 > 6$, i.e., the original predicate where we replaced y by 8, to the interaction preceding the recursion. TS now holds

in the new global assertion (assuming that $\text{TS}(\mathcal{G}')$ holds as well). \diamond

Remark 4.9. *In Example 4.11, if we had only lifted $x > y > 6$, as in the interaction case, it would not have solved the TS problem. Indeed, the predicate of the first interaction would have become $\exists y : x > y > 6$ which does not exclude values for x which are incompatible with the invariant (e.g., $x = 8$).*

The overall lifting procedure relies on a repeated application of TSres until either the assertion tree validates TS or the function fails to solve the problem. In the latter case, the function returns the most improved version of the tree and the node at which it failed.

Definition 4.15 (Λ). Λ is defined as follows, given a global assertion \mathcal{G} .

$$\Lambda(\mathcal{G}) = \begin{cases} \mathcal{G}, & \text{if } \text{TS}(\mathcal{G}) \\ \Lambda(\text{TSres}_{\text{T}(\mathcal{G})}(n)), & \text{if there is } n \in \overline{\text{TS}}(\text{T}(\mathcal{G})) \text{ s.t. } \text{TSres}_{\text{T}(\mathcal{G})}(n) \neq \perp \\ \mathcal{G}^{\perp n}, & \text{otherwise} \end{cases}$$

4.4.2 Applying Λ to Branching and Recursion

Branching. The underlying idea of branching is to enable the architect to design a choreography where a branch cannot be taken when some variables have a particular value. The architect should be involved in the resolution of the problem, because two options are possible; either the disjunction of all the predicates found in the branches is lifted, or one of the branches predicate is lifted. Arguably, the latter may also prohibit other branches to be chosen, as shown in Example 4.12 below.

Observe that according to Definition 4.9, TS fails on branching nodes only when there are values for which none of the predicates decorating the branches are satisfiable.

Example 4.12. As an illustration, we consider the following global assertion:

$$\begin{aligned} \mathcal{G}_{\text{ex4.12}} = & \mathbf{r} \rightarrow \mathbf{s} : \{v \mid \mathbf{true}\}; \\ & \mathbf{s} \rightarrow \mathbf{r} : \left\{ \begin{array}{l} \{v > 5\} l_1 : \mathcal{G}_1, \\ \{v < 5\} l_2 : \mathcal{G}_2 \end{array} \right\} \end{aligned}$$

Assuming that $\text{TS}(\mathcal{G}_1)$ and $\text{TS}(\mathcal{G}_2)$ hold, we have that $\text{TS}(\mathcal{G}_{\text{ex4.12}})$ does not hold because $\mathbf{true} \not\Rightarrow (v > 5 \vee v < 5)$. Clearly, if $v = 5$ then no branch may be selected.

Let us call \hat{n} the node corresponding to the branching in the second line of $\mathcal{G}_{\text{ex4.12}}$. Depending on the intention of the architect the problem could be fixed by one of these invocations to build (where, in both cases, superfluous quantifiers are removed).

- $\text{build}_{\text{T}(\mathcal{G}_{\text{ex4.12}})}(\hat{n}, v > 5 \vee v < 5)$ replaces the predicate in the first line by $\mathbf{true} \wedge (v > 5 \vee v < 5)$
- $\text{build}_{\text{T}(\mathcal{G}_{\text{ex4.12}})}(\hat{n}, v < 5)$ replaces the predicate in the first line by $\mathbf{true} \wedge (v < 5)$.

Both solutions solve the TS problem, however the second one prevents the first branch to be ever taken. \diamond

Given an assertion tree T and a branching node⁵ $n \in T$ such that TS does not hold. One can invoke $\text{build}_T(n, \psi)$ where ψ is either the disjunction of all the branching predicates or one of the branches predicate. If the function does not return \perp , then the TS problem is solved.

Recursion. The lifting algorithm can easily be extended to solve TS problems which occur in a recursive call, if we assume an annotation giving the invariant of its corresponding recursive definition (as in Definition 4.9). In fact, let a TS problem be located at a node $n \in T$ such that $\underline{n} = \chi \langle \vec{e} \rangle$ and let the invariant of the definition of χ be $\psi(\vec{v})$, then if the invocation of $\text{build}_T(n, \psi \{ \vec{e}/\vec{v} \})$ succeeds, the problem is solved.

⁵We also assume that TS is not violated in $\text{parent}_T(n) \uparrow_T$ as in Definition 4.11.

In Example 4.13, below we give a more complex example of the application of Λ , with TS problems in recursive calls.

Example 4.13. Consider the global assertion below.

$$\begin{aligned}
\mathcal{G}_{\text{ex4.13}} = & \text{Generator} \rightarrow \text{Server} : \{n \mid n > 0\}; \\
& \text{Player} \rightarrow \text{Server} : \{x \mid \mathbf{true}\}; \\
& \mu\chi\langle x \rangle \{r \mid r > 0\}. \\
& \text{Server} \rightarrow \text{Player} : \left\{ \begin{array}{l} \{r > n\} \text{ less} : \text{Player} \rightarrow \text{Server} : \{y \mid \mathbf{true}\}; \chi\langle y \rangle, \\ \{r < n\} \text{ greater} : \text{Player} \rightarrow \text{Server} : \{z \mid \mathbf{true}\}; \chi\langle z \rangle, \\ \{r = n\} \text{ win} : \mathbf{0} \end{array} \right\}
\end{aligned}$$

modelling a small game where a Player has to guess an integer n , following the hints given by a Server. The number is fixed by a Generator. Each time Player sends Server a number, Server says whether n is less or greater than that number. Let $T_{\text{ex4.13}}$ be the tree generated from $\mathsf{T}(\mathcal{G}_{\text{ex4.13}})$. There is a TS problem at the node corresponding to the recursive definition (let us call it n_3), indeed if $x \leq 0$, the invariant is not respected. After the first loop of $\Lambda(T_{\text{ex4.13}})$, the predicate $x > 0$ is added in the second interaction, i.e., $\text{TSres}_{T_{\text{ex4.13}}}(n_3)$ is invoked and returns a new tree, say $T'_{\text{ex4.13}}$, where the second interaction is updated to

$$\text{Player} \rightarrow \text{Server} : \{x \mid x > 0\}$$

Then, the algorithm loops two more times to solve the problems appearing before the recursive calls. Assuming n_4 (resp. n_5) is the node corresponding to the recursive call in the less (resp. greater) branch. $\text{TSres}_{T'_{\text{ex4.13}}}(n_4)$ is invoked, adding $y > 0$ in the interaction of the less branch, let us call this new tree $T''_{\text{ex4.13}}$. The updated interaction is now

$$\text{Player} \rightarrow \text{Server} : \{y \mid y > 0\}$$

Then, the algorithm invokes $\text{TSres}_{\text{ex4.13}}''(n_5)$, which adds $z > 0$ in the interaction of the greater branch, updating the interaction to

$$\text{Player} \rightarrow \text{Server} : \{z \mid z > 0\}$$

The global assertion now satisfies temporal satisfiability. \diamond

Recursion parameters. Even though lifting may be applied when a TS violation is detected in a recursion definition, lifting a predicate involving a recursion parameter v would require to strengthen the invariant where v is introduced. This is quite dangerous, therefore the lifting algorithm does not apply in this case (cf. the last line of the first part of Definition 4.13). In fact, for recursive definition and calls, Definition 4.9 requires $\Psi \implies \Psi' \{\bar{e}/\bar{v}\}$, where Ψ' is the recursion invariant and Ψ is the conjunction of the previous predicates. Hence, lifting a predicate involving a recursion parameter may strengthen the invariant and possibly create a new problem in a corresponding recursive call. Moreover, notice that, in recursive calls, *GSat* (Definition 4.9) requires that $\Psi \wedge \Psi' \implies \Psi' \{\bar{e}/\bar{v}\}$; namely, strengthening Ψ' would automatically strengthen $\Psi' \{\bar{e}/\bar{v}\}$ and therefore leave the TS problem unsolved.

We illustrate such a case with Example 4.14 below.

Example 4.14. Continuing with the example of Section 4.1, consider the following global

assertion which is supposed to have been “badly” modified by a designer.

$$\mathcal{G}_{\text{ex4.14}} = b_i \rightarrow s_i : \{request \mid \mathbf{true}\}; \quad (1)$$

$$s_i \rightarrow b_i : \{quote_f \mid \mathbf{true}\}; \quad (2)$$

$$\mu\chi \langle quote_f \ quote_f \rangle \{quote_c \ quote_p \mid \text{MIN} < quote_c \leq quote_p\}. \quad (3)$$

$$b_i \rightarrow s_i : \{ \mathbf{true} \} \text{ ok} : \mathbf{0}, \quad (4)$$

$$\{ \mathbf{true} \} \text{ no} : s_i \rightarrow b_i : \{quote_n \mid quote_n < quote_c\}; \quad (5)$$

$$\chi \langle quote_n \ quote_c \rangle \} \quad (6)$$

There are two TS problems in $\mathcal{G}_{\text{ex4.14}}$. The first problem is at the node corresponding to line (3) and is due to the fact that $quote_f$ is not required to be strictly greater than MIN – thus violate the invariant. The second problem is at the node corresponding to the recursive call (line (6)) and is due to the fact that $quote_n$ may not be strictly greater than MIN, and thus violate the invariant as well.

The first problem may be solved by lifting the predicate $\text{MIN} < quote_f \leq quote_f$ to line (2), however our algorithm fails to solve the second problem as it would amount to lifting the predicate

$$\text{MIN} < quote_n \leq quote_c$$

which contains a recursion parameter ($quote_c$).

We illustrate why lifting such a predicate is not supported by our algorithm. Assume that both problems were tackled by lifting, we would obtain the following global asser-

tion:

$$\mathcal{G}'_{\text{ex4.14}} = b_i \rightarrow s_i : \{request \mid \mathbf{true}\}; \quad (1)$$

$$s_i \rightarrow b_i : \{quote_f \mid \text{MIN} < quote_f\}; \quad (2)$$

$$\mu\chi \langle quote_f quote_f \rangle \{quote_c quote_p \mid \text{MIN} < quote_c \leq quote_p\}. \quad (3)$$

$$b_i \rightarrow s_i : \{ \mathbf{true} \} \text{ ok} : \mathbf{0}, \quad (4)$$

$$\{ \mathbf{true} \} \text{ no} : s_i \rightarrow b_i : \{quote_n \mid \text{MIN} < quote_n < quote_c\}; \quad (5)$$

$$\chi \langle quote_n quote_c \rangle \} \quad (6)$$

Observe that lifting $\text{MIN} < quote_n \leq quote_c$ to line (3) yields the predicate:

$$\text{MIN} < quote_c \leq quote_p \quad \wedge \quad \exists quote_n : \text{MIN} < quote_n \leq quote_c$$

which is equivalent to the initial invariant.

In fact, there is also a TS problem in $\mathcal{G}'_{\text{ex4.14}}$. At line (5), if, e.g., $quote_f = \text{MIN} + 1$, then s_i cannot choose a correct value for $quote_n$. Assuming again we may lift predicates that contain recursion parameters, we would lift the predicate $\text{MIN} < quote_n < quote_c$ to solve this issue. This would strengthen the invariant to⁶

$$\text{MIN} < quote_c \leq quote_p \quad \wedge \quad \exists quote_n : \text{MIN} < quote_n < quote_c$$

which is equivalent to $\text{MIN} + 1 < quote_c \leq quote_p$. At this point, we are essentially back to the global assertion $\mathcal{G}_{\text{ex4.14}}$, modulo the fact that MIN has increased by one. It is easy to see that, without the condition on recursion parameters in Definition 4.13, the algorithm would not terminate in such a case. \diamond

⁶ Observe the strict inequality between $quote_n$ and $quote_c$ in this case.

4.4.3 Properties of Λ

Similarly to the algorithms Σ and Π of Section 4.3, Λ does not modify the structure of the tree and preserves the properties of the initial assertion.

Proposition 4.6 (Underlying type structure - Λ). *Let \mathcal{G} be a global assertion. If $\Lambda(\mathcal{G})$ returns \mathcal{G}' then $\text{erase}(\mathcal{G}) = \text{erase}(\mathcal{G}')$ (cf. Definition 4.8).*

Proof. The proof is by induction on the structure of \mathcal{G} , similarly to the one of Proposition 4.4. □

Proposition 4.7 below guarantees that Λ does not introduce new HS or TS problems. Likewise, it gives a formal account of the remark in Section 4.3.3, showing that Σ and Π do not add violations (of either HS or TS) to the amended global assertions.

Proposition 4.7 (Properties Preservation). *Assume $F(\mathcal{G})$ returns \mathcal{G}' with $F \in \{\Sigma, \Pi, \Lambda\}$. If $\overline{\text{HS}}(\mathcal{G}) = \emptyset$ then $\overline{\text{HS}}(\mathcal{G}') = \emptyset$ and if $\overline{\text{TS}}(\mathcal{G}) = \emptyset$ then $\overline{\text{TS}}(\mathcal{G}') = \emptyset$.*

Proof. We first consider HS preservation and then TS preservation for $F(\mathcal{G})$ with $F \in \{\Sigma, \Pi, \Lambda\}$.

HS preservation. The proof of HS preservation by Σ and Π follows by the fact that Σ and Π return \mathcal{G} if $\overline{\text{HS}}(\mathcal{G}) = \emptyset$. For Λ , the preservation of HS follows from the fact that all the variables which are not known to a participant are quantified (either universally or existentially) in the modified predicates. We show that all the variables not known to the sender of an updated node are quantified. Let T be an assertion tree and ψ be the predicate lifted at a node $n \in T$ such that $\text{snd}_T(n) = s$. The predicate is quantified as in Definition 4.14 so to obtain $\forall \vec{x} : \exists \vec{y} : \psi$ such that

- $\vec{x} \subseteq \text{var}(\psi) \setminus \text{knows}_s(T)$ are introduced in a node in $n \uparrow_T$
- $\vec{y} \subseteq \text{var}(\psi)$ are introduced in a node in the subtree rooted at n

Let $z \in \text{var}(\psi)$, (i) if $z \notin \vec{x}$, by definition, either z is known to s (therefore z should not be quantified) or z is introduced after n (hence it would have been quantified in \vec{y}). (ii) If $z \notin \vec{y}$, by definition, either z is introduced at n (therefore known to s) or z is introduced before n . In that case, if it is known to s then it should not be quantified, and if it is not known to n , then $z \in \vec{x}$.

TS preservation. TS preservation in Σ follows from the fact that predicates may only be changed by a variable substitution. For $T = \mathbb{T}(\mathcal{G})$, such that $\overline{\text{TS}}(\mathcal{G}) = \emptyset$, we have that, for any $n \in T$

$$\text{PRED}_T(n) \implies \exists \text{var}_T(n) : \phi$$

by definition of TS (Definition 4.9), with ϕ being the predicate at node n . And, by (4.3) (cf. page 126), we have that

$$\text{PRED}_T(n) \implies \exists \text{var}_T(n) : \phi \{v'/v\}$$

thus, TS is preserved by Σ .

TS preservation in Π follows from the fact that the predicates of a global assertion are only modified by adding equalities between problematic variables and fresh variables (see statement 2b in Proposition 4.5). For $T = \mathbb{T}(\mathcal{G})$, such that $\overline{\text{TS}}(\mathcal{G}) = \emptyset$, we have that, for any $n \in T$

$$\text{PRED}_T(n) \implies \exists \text{var}_T(n) : \phi \tag{4.5}$$

by definition of TS, with ϕ being the predicate at node n . And, by construction of a \prec_T -chain, after each modification by Π , we obtain

$$\text{PRED}_T(n) \wedge v = v_0 \wedge \dots \wedge v_t = v_{t-1} \implies \exists \text{var}_T(n) : \phi \{v_i/v\}$$

with $v_0 \dots v_t$ fresh. This is equivalent to (4.5), i.e., TS is preserved by Π . The proof of TS preservation for Λ follows trivially from the first case of Definition 4.15. \square

Proposition 4.8 establishes an intermediate result for the correctness of Λ . It says that a successful invocation of TSres (cf. Definition 4.14) on a node removes the problem at that node.

Proposition 4.8 (Correctness - TSres). *Let T be an assertion tree. For each $n \in \overline{\text{TS}}(T)$, if $\text{TSres}_T(n) \neq \perp$, then $n \notin \overline{\text{TS}}(\text{TSres}_T(n))$.*

Proof. We start by giving the proof of the correctness for *interaction nodes*. Let T be an assertion tree with a node n such that $n \in \overline{\text{TS}}(T)$, and

$$\underline{n} = \mathbf{s} \rightarrow \mathbf{r} : \{\vec{v} \mid \psi\}$$

with $\psi \iff \beta \wedge \gamma$ such that β is in conflict on $\text{var}(n)$ with γ in $\text{PRED}_T(n)$. Then β is the predicate to be lifted. Assume $\hat{T} = \text{build}_T(n, \beta)$.

By Definition 4.13, we have that, for suitable $\vec{x}_1, \vec{y}_1 \dots \vec{x}_k, \vec{y}_k$,

$$\text{PRED}_{\hat{T}}(n) = \text{PRED}_T(n) \wedge \forall \vec{x}_1 : \exists \vec{y}_1 : \beta \wedge \dots \wedge \forall \vec{x}_k : \exists \vec{y}_k : \beta$$

We have that a quantified version of β is added k times in the assertion tree, above n . We show that

$$\bigwedge_{1 \leq i \leq k} \forall \vec{x}_i : \exists \vec{y}_i : \beta \implies \exists \vec{v} : \beta \quad (4.6)$$

Assume that each $\forall \vec{x}_i : \exists \vec{y}_i : \beta$ corresponds to the predicate added to the i^{th} node (n_i) modified by TSres (from the root to n). Then $\vec{y}_k = \vec{v} \cap \text{var}(\beta)$ since by Definition 4.13 \vec{y}_k is the set of variables introduced *after* n_k , and we assumed that the global assertion is closed (i.e., all the variables $\text{var}(\beta)$ have been introduced before they are used, in n). Since every $\forall \vec{x}_i : \exists \vec{y}_i : \beta$ is satisfiable by Definition 4.13, we have that the following holds

$$\exists \vec{z} : \forall \vec{x}_k : \exists \vec{y}_k : \beta \quad \text{with } \vec{z} = \text{var}(\forall \vec{x}_k : \exists \vec{y}_k : \beta)$$

this gives us (4.6) (note that \vec{z} , \vec{x}_k and \vec{y}_k are pairwise disjoint by definition). Since $\forall \vec{x}_k : \exists \vec{y}_k : \beta$ is one of the conjunct of $\text{PRED}_{\hat{T}}(n)$ we also have

$$\text{PRED}_{\hat{T}}(n) \implies \exists \vec{v} : \beta \quad (4.7)$$

By the definition of conflict (Definition 4.12), we have

$$\text{PRED}_T(n) \implies \exists \vec{v} : \gamma \quad \text{and} \quad \text{PRED}_T(n) \not\Rightarrow \exists \vec{v} : \beta \wedge \gamma$$

hence, $\text{PRED}_T(n)$ is satisfiable. Therefore, by weakening, we have that:

$$\text{PRED}_{\hat{T}}(n) \implies \exists \vec{v} : \gamma \quad (4.8)$$

TS must hold for n , which implies that $n \notin \overline{\text{TS}}(\hat{T})$ and $\text{TSnode}_{\hat{T}}(n)$ holds, i.e.,

$$\text{PRED}_{\hat{T}}(n) \implies \exists \vec{v} : \psi \quad (\text{with } \psi \iff \beta \wedge \gamma)$$

Otherwise, that would imply that

$$\text{PRED}_{\hat{T}}(n) \wedge \forall \vec{v} : (\neg\beta \vee \neg\gamma)$$

which is in contradiction with (4.7) (β) and (4.8) (γ).

Let us now show the result for *recursive nodes*, which is somewhat similar to the previous case. Assume we have

$$\underline{n} = \mu\chi \langle \vec{e} \rangle \{ \vec{v} \mid \beta \}$$

with $n \in \overline{\text{TS}}(T)$, thus we have that (by Definition 4.9)

$$\text{PRED}_T(n) \iff \beta\{\bar{e}/\bar{v}\} \quad (4.9)$$

Assuming $\hat{T} = \text{build}_T(n, \beta\{\bar{e}/\bar{v}\})$ (i.e., build succeeds), we have that a quantified version of $\beta\{\bar{e}/\bar{v}\}$ is added k times in the assertion tree, above n . Following a similar argument as before, this gives us

$$\text{PRED}_{\hat{T}}(n) = \text{PRED}_T(n) \wedge \beta\{\bar{e}/\bar{v}\}$$

By Definition 4.13, we also know that $\beta\{\bar{e}/\bar{v}\}$ is satisfiable. Moreover, $\text{PRED}_T(n)$ is satisfiable, by Definition 4.11 and (4.9).

Since

$$\text{PRED}_{\hat{T}}(n) \implies \beta\{\bar{e}/\bar{v}\} \iff \text{PRED}_T(n) \wedge \beta\{\bar{e}/\bar{v}\} \implies \beta\{\bar{e}/\bar{v}\}$$

we have that $n \notin \overline{\text{TS}}(\hat{T})$. □

Finally, we can say that, if a repeated application of lifting succeeds, the global assertion which is returned satisfies temporal satisfiability.

Theorem 4.2 (Correctness - Λ). *If $\Lambda(\mathcal{G}) = \mathcal{G}'$ then $\overline{\text{TS}}(\mathcal{G}') = \emptyset$.*

Proof. The proof is by induction on the number of problematic nodes and the minimum depth of these nodes in the tree. It relies on Proposition 4.8, i.e., the fact that $\text{TSres}_T(n)$ either solves the problem at n or fails.

Let $T = T(\mathcal{G})$ and N be the set of nodes in T which violates TS. We write $|n|$ for the depth of n in T (with $|T^\bullet| = 0$), and we denote by N' the number of problematic node after an invocation to TSres .

1. If $N = \emptyset$, then T is TS.

2. If $N \neq \emptyset$, let $n \in \overline{\text{TS}}(T) \subseteq N$, after an invocation to $\text{TSres}_T(n)$, we have
- (a) If $|n| > 1$ then either
 - i. $N' := N \setminus \{n\}$, i.e., the node is simply removed from the set of problematic nodes,
 - ii. $N' := N \cup N_* \setminus \{n\}$, where N_* is the set of problematic nodes *added* by TSres . We have that $\forall n'_i \in N_* : |n'_i| < |n|$, i.e., the problem at n is solved but other problematic nodes, *above* n in T , are added, or,
 - iii. the algorithm fails on n
 - (b) If $|n| \leq 1$ then either $N' := N \setminus \{n\}$, or the algorithm fails. In fact, once the algorithm reaches a problem located at a child of the root, then it either fails or solves the problem. Indeed, there cannot be a TS problem at the root node unless the predicate is unsatisfiable (see Definition 4.9), in which case, the algorithm fails.

Note that selecting $n \in \overline{\text{TS}}(T)$ implies that the depth of n is smaller or equal to the depth of the nodes in N .

Regarding step 2(a)ii, note that the algorithm cannot loop on a problematic node indefinitely. Indeed, the number of (sub)predicates available for lifting is finite and Λ invokes TSres only when a problematic node is found.

It is easily shown by induction that the algorithm terminates either with $\overline{\text{TS}}(T) = \emptyset$, or a failure. □

In addition, we have that Λ preserves the domain of possible values for each variable from the initial assertion.

Proposition 4.9 (Assertion predicates). *If $\Lambda(\mathcal{G}) = \mathcal{G}'$ then for all $n \in \text{T}(\mathcal{G})$ such that n is*

a leaf, and its corresponding node $n' \in T(\mathcal{G}')$

$$\text{PRED}_{T(\mathcal{G})}(n) \wedge \text{cst}_{T(\mathcal{G})}(n) \iff \text{PRED}_{T(\mathcal{G}')} (n') \wedge \text{cst}_{T(\mathcal{G}')} (n')$$

Proof. The proof follows from the observation that predicates are only duplicated in the tree, i.e., the lifting algorithm does not add any new constraints in the conjunction of the predicates found on the path from the root to a leaf. In addition, the algorithm modifies only predicates which appear above a problematic node in the tree (i.e., a predicate in a leaf will never be modified, by Definition 4.13).

We show the result for *interaction nodes*, the case for recursive nodes is similar. After each successful iteration of Λ on an assertion tree $T = T(\mathcal{G})$, where

$$\text{PRED}_T(n) = \psi_1 \wedge \dots \wedge \psi_i \wedge \dots \wedge \psi_k$$

and ψ_i is the predicate where the TS problem is located, and β is the lifted predicate (cf. proof of Proposition 4.8), such that we have, by Definition 4.12

$$\psi_i \iff \beta \wedge \gamma$$

We can rewrite $\text{PRED}_T(n)$ as

$$\text{PRED}_T(n) = \psi_1 \wedge \dots \wedge \gamma \wedge \beta \wedge \dots \wedge \psi_k$$

build returns a new tree T' such that the conjunction of predicates in the new tree is of the form

$$\text{PRED}_{T'}(n') = \text{PRED}_T(n) \wedge \bigwedge_{j \in J} \forall \vec{x}_j \exists \vec{y}_j : \beta$$

By Definition 4.13, we have $\vec{x}_j \vec{y}_j \subseteq \text{var}(\beta)$, for all $j \in J$. Therefore, we have

$$\beta \implies \forall \vec{x}_j \exists \vec{y}_j : \beta \quad \text{for all } j \in J$$

This means that the additional predicates do not constrain further the conjunction of predicates, and we have

$$\text{PRED}_T(n) \iff \text{PRED}_{T'}(n') \quad \square$$

4.5 A Methodology for Amending Choreographies

The algorithms Σ , Π , and Λ of Section 4.3 and Section 4.4 can be used to support a methodology for amending contracts in choreographies. The methodology consists of the following steps:

- (i) the architect designs a global assertion $\widehat{\mathcal{G}}$ (possibly based on a synthesised global type synthesised as in Chapter 3)
- (ii) the architect is notified if there are any HS or TS problems in $\widehat{\mathcal{G}}$
- (iii) using Σ and Π solutions may be offered for HS problems, while Λ can be used to offer solutions and/or hints on how to solve TS problems
- (iv) the architect selects one of the solutions offered in (iii)
- (v) steps (ii) to (iv) are repeated until all problems are addressed.

We illustrate our methodology using the following global assertion:

$$\widehat{\mathcal{G}} = \mu\chi\langle 10 \rangle \{v \mid v > 0\}. \quad (1)$$

$$\text{Alice} \rightarrow \text{Bob} : \{v_1 \mid v \geq v_1\}; \quad (2)$$

$$\text{Bob} \rightarrow \text{Carol} : \{v_2 \mid v_2 > v_1\}; \quad (3)$$

$$\text{Carol} \rightarrow \text{Alice} : \{v_3 \mid v_3 > v_1\}; \quad (4)$$

$$\text{Carol} \rightarrow \text{Bob} : \{v_4 \mid v_4 > v\}; \quad (5)$$

$$\text{Alice} \rightarrow \text{Bob} : \{ \{\mathbf{true}\} \text{ cont} : \chi\langle v_1 \rangle, \quad (6)$$

$$\{\mathbf{true}\} \text{ finish} : \text{Alice} \rightarrow \text{Bob} : \{v_5 \mid v_1 < v_5 < v_3 - 2\} \} \quad (7)$$

which extends the global assertion in Example 4.4 and is assumed to have been designed by the architect (step (i) of the methodology).

Firstly, $\widehat{\mathcal{G}}$ is inspected by history sensitivity and temporal satisfiability checkers, such as the ones described in [52]. If any HS problems are reported (step (ii) of the methodology), algorithms Σ and Π are used, while Λ is used for TS problems. This allows the architect to detect all the problems and to consider the ones for which (at least) one of the algorithms is applicable. In general, the architect can decide which problem to tackle first (step (iii) of our methodology). For $\widehat{\mathcal{G}}$, we focus on HS problems first. There are two HS problems in $\widehat{\mathcal{G}}$, both of which can be solved automatically, and the methodology will return the following suggestions.

1. At line (4), v_1 is not known by Carol; the problem is solvable by either
 - replacing $v_3 > v_1$ by $v_3 > v_2$ (algorithm Σ) at line (4), or
 - by revealing v_1 to Carol (algorithm Π); in this case, line (3) becomes

$$\text{Bob} \rightarrow \text{Carol} : \{v_2 \ u_1 \mid v_2 > v_1 \wedge u_1 = v_1\}$$

and the predicate at line (4) becomes $v_3 > u_1$.

2. At line (5), v is not known by Carol; the problem is solvable by revealing the value of v to Carol (algorithm Π) in which case line (3) becomes

$$\text{Bob} \rightarrow \text{Carol} : \{v_2 \ u_2 \mid v_2 > v_1 \wedge u_2 = v\}$$

and the assertion at line (5) becomes $v_4 > u_2$.

In the *propagation case* (i.e., Π), the methodology gives the architect information on which participants the value of a variable may be disclosed to. Indeed, as discussed in Remark 4.5, it may not be appropriate to use the suggested solution. Therefore, the actual adoption of the proposed solutions should be left to the architect. In addition, the order in which problems are tackled is also left to the architect (e.g., the same variable may be involved in several problems and solving one of them may automatically fix the others). Assuming that Σ is used to solve the first problem and Π to solve the second, HS is fixed, and the amended global assertion is as follows:

$$\widehat{\mathcal{G}} = \mu\chi\langle 10 \rangle \{v \mid v > 0\}. \quad (1)$$

$$\text{Alice} \rightarrow \text{Bob} : \{v_1 \mid v \geq v_1\}; \quad (2)$$

$$\text{Bob} \rightarrow \text{Carol} : \{v_2 \ u_1 \mid v_2 > v_1 \wedge u_1 = v\}; \quad (3)$$

$$\text{Carol} \rightarrow \text{Alice} : \{v_3 \mid v_3 > v_2\}; \quad (4)$$

$$\text{Carol} \rightarrow \text{Bob} : \{v_4 \mid v_4 > u_1\}; \quad (5)$$

$$\text{Alice} \rightarrow \text{Bob} : \{ \{\mathbf{true}\} \text{ cont} : \chi\langle v_1 \rangle, \quad (6)$$

$$\{\mathbf{true}\} \text{ finish} : \text{Alice} \rightarrow \text{Bob} : \{v_5 \mid v_1 < v_5 < v_3 - 2\} \} \quad (7)$$

Now HS is satisfied in $\widehat{\mathcal{G}}$, but TS problems are still present.

In case a TS problem cannot be solved automatically, additional information can be returned: (a) at which node the problem occurred, (b) which variables or recursion parameters are posing problems (i.e., using *split* and *build*), and (c) where liftings are not possible (i.e., when *build* fails to add a satisfiable predicate to a node). For $\widehat{\mathcal{G}}$ there

are two TS problems which are dealt with sequentially. The methodology would report that

1. At line (6), v_1 does not satisfy the invariant $v > 0$. This can be solved by lifting $v_1 > 0$ (i.e., the invariant where v is replaced by the actual parameter v_1) to the interaction at line (2), which would yield the new predicate $v \geq v_1 \wedge v_1 > 0$.
2. At line (7), there might be no value for v_5 such that $v_1 < v_5 < v_3 - 2$. The assertion is *in conflict* (cf. Definition 4.12) with the previous predicates; this problem cannot be solved since lifting would add the following predicates in lines (2) and (4), respectively.
 - $\exists v_3 v_5 : v_1 < v_5 < v_3 - 2$ which is indeed satisfiable, but remarkably does not constraint v_1 more than the initial predicate. Indeed, the updated predicate (i.e., $v \geq v_1 \wedge \exists v_3 v_5 : v_1 < v_5 < v_3 - 2$) does not constrain v_1 more than the original predicate, $v \geq v_1$.
 - $\forall v_1 : \exists v_5 : v_1 < v_5 < v_3 - 2$ which is not satisfiable, therefore the algorithm fails.

The failure of Λ is due to the fact that v_5 is constrained by v_1 and v_3 which are fixed by two different participants. They would have to somehow interact in order to guarantee that there exists a value for v_5 , this cannot be done using the proposed algorithms. Notice that in this case the methodology tells the architect that v_5 , fixed by Alice, is constrained by v_1 and v_3 which are fixed by Alice and Carol, respectively. Our approach can also suggest that the node introducing v_3 , or (the part of) the assertion over v_3 may be the source of the problem since v_3 is the only variable not known by Alice.

4.5.1 Amendment Strategies

The methodology above does not specify any particular order for tackling HS and TS problems. In fact, it is for the architect to assess the importance of each problem; furthermore, the architect should also proof-read the proposed solutions (e.g., in the case of propagation). One of our future work plans is to help the architect making choice regarding the order in which problems should be tackled by designing *amendment strategies*, which maximise the chances of having *all* problems solved using the proposed algorithms.

In fact, the application of an algorithm could be spoilt by the application of another one. For instance, the application of the strengthening algorithm (Σ) might compromise the applicability of the lifting algorithm (Λ), and vice versa. This happens when a variable v introduced by Σ in a node (say n) to solve an HS problem is also involved in a TS problem, in a descendant node (say n') of n ; indeed, both Σ and Λ will independently strengthen the predicate of n . This may compromise the application of the algorithm invoked last, as illustrated below.

Let T be an assertion tree, where there are $n, n' \in T$ such that $\text{varHS}_T(n) \neq \emptyset$, n' is a descendant of n ,

$$\underline{n} = \mathbf{s} \rightarrow \mathbf{r} : \{\vec{v}_1 \ v \ \vec{v}_2 \mid \psi\} \quad \text{and} \quad \underline{n'} = \mathbf{s}' \rightarrow \mathbf{r}' : \{\vec{u} \mid \gamma \wedge \beta\}$$

where $v \in \text{var}(\beta)$, and β is in conflict on \vec{u} with γ in $\text{PRED}_T(n')$.

- if Σ is used to solve the problem at n , ψ might be strengthened (by a variable substitution)
- if Λ is used to solve the problem at n' , β will be lifted to the node n since $v \in \text{var}(\beta)$.

Call ψ_1 the new predicate produced by Σ and ψ_2 the one produced by Λ . The application of Σ would give

$$\underline{n} = \mathbf{s} \rightarrow \mathbf{r} : \{\vec{v}_1 \ v \ \vec{v}_2 \mid \psi_1\}$$

which might prevent the application of Λ because, e.g., $\psi_1 \wedge \forall \vec{x} : \exists \vec{y} : \beta$ is not satisfiable, for suitable \vec{x} and \vec{y} . Likewise, the application of Λ first would give

$$\underline{n} = \mathbf{s} \rightarrow \mathbf{r} : \{\vec{v}_1 \vee \vec{v}_2 \mid \psi_2\}$$

for which Σ might not be applicable because, e.g., no substitution with a variable known to \mathbf{s} yields a satisfiable predicate for \underline{n} .

We conjecture that this is the *only* source of issues arising from the absence of prescribed order for addressing HS and TS problems in the methodology. Intuitively, the only way one algorithm could spoil the applicability of another is by modifying the satisfiability of a predicate of (at least) one common node. Propagation (Π) preserves the semantics of all the nodes it updates (by Proposition 4.5); instead, Σ and Λ may strengthen predicates. Note that Σ modifies only the node in which there is an HS problem, while Λ updates only the nodes *above* the TS-problematic one. Therefore, the only possible issue occurs when there is a node with an HS problem “above” another, with a TS problem. Since Λ modifies only the nodes that introduce variables which appear in a *problematic* predicate, we conjecture that this happens only in the case explained above.

Note that even though an occurrence of two inter-dependent TS and HS problems as above may compromise the applicability of an algorithm, thus preventing the amendment of existing problems, it will not introduce new violations, as stated in Proposition 4.7.

4.6 Applying the Methodology

To illustrate our methodology, we consider the design of a couple of services offered by an ATM to the customers of the bank where it is located. The first service offers cash withdrawal. The second service allows customers to request a small line of credit, provided that they are considered trusted by the bank. We propose two global assertions

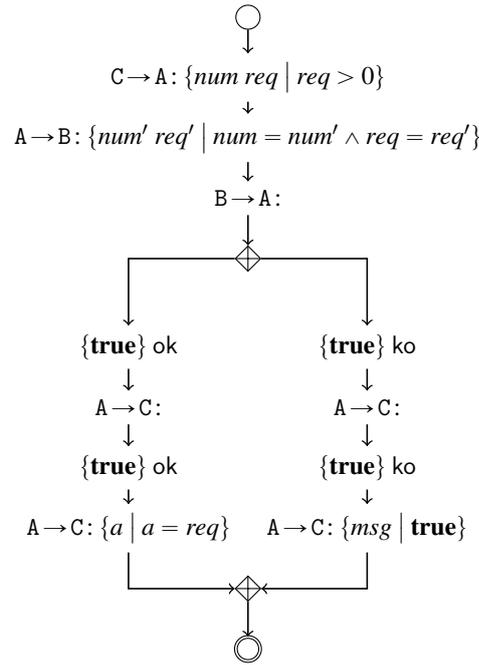


Figure 4.2: ATM protocol

for each of the two functionalities and discuss problems which may be encountered during their design.

4.6.1 Cash Withdrawal

Consider the following global assertion, also illustrated in Figure 4.2, where C is the customer, A is the ATM, and B is the bank.

$$C \rightarrow A: \{num\ req \mid req > 0\}; \quad (1)$$

$$A \rightarrow B: \{num' \ req' \mid num = num' \wedge req = req'\}; \quad (2)$$

$$B \rightarrow A: \left\{ \begin{array}{l} \{true\} \ ok : A \rightarrow C: \{true\} \ ok : A \rightarrow C: \{a \mid a = req\}, \\ \{true\} \ ko : A \rightarrow C: \{true\} \ ko : A \rightarrow C: \{msg \mid true\} \end{array} \right\} \quad (3)$$

$$\{true\} \ ko : A \rightarrow C: \{true\} \ ko : A \rightarrow C: \{msg \mid true\} \quad (4)$$

This global assertion models a simple cash withdrawal service under the assumption that the credentials of the customer have already been verified. The customer sends the ATM an account number num and the amount of money to withdraw req . The ATM

forwards the request to the bank. If the withdrawal is accepted, the bank selects branch ok; in this case the ATM gives the corresponding amount to the customer. Otherwise, the bank selects branch ko and ATM sends an error message to the customer.

This global assertion is well-asserted, but soon the architect realises that it contains a major flaw: the ATM is expected to give money to the customer even when there is not enough cash available in the machine. The architect corrects the problem by adding a predicate $a \leq \text{CASH}$ at line (3), where CASH is the money available at the beginning of the session:

$$C \rightarrow A: \{num\ req \mid req > 0\}; \quad (1)$$

$$A \rightarrow B: \{num' \ req' \mid num = num' \wedge req = req'\}; \quad (2)$$

$$B \rightarrow A: \left\{ \begin{array}{l} \{\mathbf{true}\} \text{ ok} : A \rightarrow C: \{\mathbf{true}\} \text{ ok} : A \rightarrow C: \{a \mid a = req \wedge a \leq \text{CASH}\}, \\ \{\mathbf{true}\} \text{ ko} : A \rightarrow C: \{\mathbf{true}\} \text{ ko} : A \rightarrow C: \{msg \mid \mathbf{true}\} \end{array} \right\}, \quad (3)$$

$$\left. \begin{array}{l} \{\mathbf{true}\} \text{ ok} : A \rightarrow C: \{\mathbf{true}\} \text{ ok} : A \rightarrow C: \{a \mid a = req \wedge a \leq \text{CASH}\}, \\ \{\mathbf{true}\} \text{ ko} : A \rightarrow C: \{\mathbf{true}\} \text{ ko} : A \rightarrow C: \{msg \mid \mathbf{true}\} \end{array} \right\} \quad (4)$$

Although this solves the flaw, a temporal satisfiability issue is introduced at line (3). In fact, A cannot guarantee its obligation if the amount requested req in the first interaction is greater than the cash available.

Fortunately, Λ is applicable and it can amend the global assertion automatically by returning the choreography below

$$C \rightarrow A: \{num\ req \mid req > 0 \wedge \exists a: a = req \wedge a \leq \text{CASH}\}; \quad (1)$$

$$A \rightarrow B: \{num' \ req' \mid num = num' \wedge req = req'\}; \quad (2)$$

$$B \rightarrow A: \left\{ \begin{array}{l} \{\mathbf{true}\} \text{ ok} : A \rightarrow C: \{\mathbf{true}\} \text{ ok} : A \rightarrow C: \{a \mid a = req \wedge a \leq \text{CASH}\}, \\ \{\mathbf{true}\} \text{ ko} : A \rightarrow C: \{\mathbf{true}\} \text{ ko} : A \rightarrow C: \{msg \mid \mathbf{true}\} \end{array} \right\}, \quad (3)$$

$$\left. \begin{array}{l} \{\mathbf{true}\} \text{ ok} : A \rightarrow C: \{\mathbf{true}\} \text{ ok} : A \rightarrow C: \{a \mid a = req \wedge a \leq \text{CASH}\}, \\ \{\mathbf{true}\} \text{ ko} : A \rightarrow C: \{\mathbf{true}\} \text{ ko} : A \rightarrow C: \{msg \mid \mathbf{true}\} \end{array} \right\} \quad (4)$$

which is well-asserted.

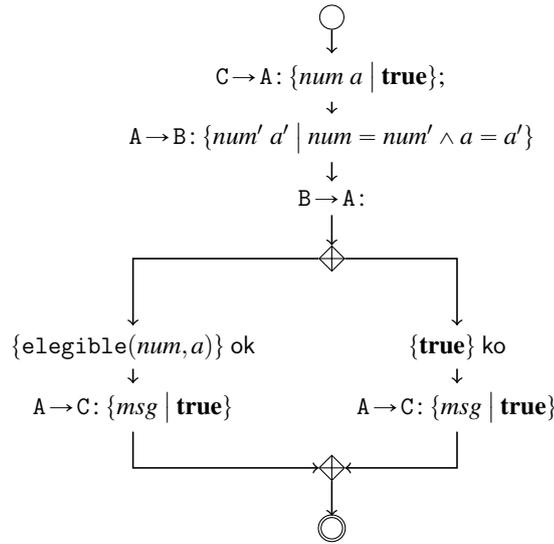


Figure 4.3: Credit request protocol

4.6.2 Credit Request

We now want to model a service through which a customer can request a small line of credit. The intuition of the protocol is illustrated in Figure 4.3. The customer C sends their account number num and the requested credit a to the ATM. The ATM forwards the request to the bank and, depending on whether or not C is eligible for the credit according to the bank's records (i.e., $eligible(num, a)$), the bank selects either branch ok or ko . Finally, the ATM sends a message to the customer notifying them of the decision.

Remark 4.10. *For simplicity, we use branch mergeability [73], a slight extension of multiparty session types. Otherwise, it would be necessary to add an extra branch in the inner branching between A and C to have the same behaviour of C in both branches of the outer branching. Note that this does not affect the applicability of our methodology.*

A naive global assertion modelling this service is as follows:

$$C \rightarrow A: \{num\ a \mid \mathbf{true}\}; \quad (1)$$

$$A \rightarrow B: \{num'\ a' \mid num = num' \wedge a = a'\}; \quad (2)$$

$$B \rightarrow A: \left\{ \begin{array}{l} \{eligible(num, a)\} \\ \mathbf{true} \end{array} \right. \quad \begin{array}{l} ok : A \rightarrow C: \{msg \mid \mathbf{true}\}, \\ ko : A \rightarrow C: \{msg \mid \mathbf{true}\} \end{array} \quad (3)$$

$$\quad \quad \quad \quad \quad \quad \quad ko : A \rightarrow C: \{msg \mid \mathbf{true}\} \quad (4)$$

The attentive reader will notice that there is an HS problem at line (3) of this global assertion. Indeed, B does not know num nor a and therefore could not guarantee that the customer is in fact eligible. Both Σ and Π algorithms are applicable here. The second algorithm would return the following global assertion:

$$C \rightarrow A: \{num\ a \mid \mathbf{true}\}; \quad (1)$$

$$A \rightarrow B: \{num'\ a' \ v_1 \ v_2 \mid num = num' \wedge a = a' \wedge v_1 = num \wedge v_2 = a\}; \quad (2)$$

$$B \rightarrow A: \left\{ \begin{array}{l} \{eligible(v_1, v_2)\} \\ \mathbf{true} \end{array} \right. \quad \begin{array}{l} ok : A \rightarrow C: \{msg \mid \mathbf{true}\}, \\ ko : A \rightarrow C: \{msg \mid \mathbf{true}\} \end{array} \quad (3)$$

$$\quad \quad \quad \quad \quad \quad \quad ko : A \rightarrow C: \{msg \mid \mathbf{true}\} \quad (4)$$

Although this solves the problem, we notice that this solution is not ideal. Indeed v_1 and v_2 are somewhat redundant with num' and a' , respectively.

Strengthening gives us a better solution in this case:

$$C \rightarrow A: \{num\ a \mid \mathbf{true}\}; \quad (1)$$

$$A \rightarrow B: \{num'\ a' \mid num = num' \wedge a = a'\}; \quad (2)$$

$$B \rightarrow A: \left\{ \begin{array}{l} \{eligible(num', a')\} \\ \mathbf{true} \end{array} \right. \quad \begin{array}{l} ok : A \rightarrow C: \{msg \mid \mathbf{true}\}, \\ ko : A \rightarrow C: \{msg \mid \mathbf{true}\} \end{array} \quad (3)$$

$$\quad \quad \quad \quad \quad \quad \quad ko : A \rightarrow C: \{msg \mid \mathbf{true}\} \quad (4)$$

Which is what one would expect. Note that the algorithm is applicable because

$$num = num' \wedge a = a' \wedge eligible(num', a') \implies eligible(num, a)$$

holds and B knows num' and a' .

4.7 Concluding Remarks

In this chapter, we investigated the problem of designing consistent assertions. We focused on two consistency criteria from [12]: history sensitivity and temporal satisfiability. We proposed and compared three algorithms (Σ , Π , and Λ) to amend global assertions. Since each algorithm is applicable only in certain circumstances, we proposed a methodology that supports the architect when violations are not automatically amendable.

On the theoretical side, the algorithms Σ , Π , and Λ address the general problem of guaranteeing the satisfiability of predicates when: (1) the parts of the system have a different perspective/knowledge of the available information (in the case of history sensitivity), and (2) the constraints are introduced progressively (in the case of temporal satisfiability). The proposed solutions can be adapted and used, for instance, to amend processes (rather than types), orchestrations (rather than choreographies, when we want to check for local constraints) expressed in formalisms such as CC-Pi [18], a language for distributed processes with constraints. Interestingly, temporal satisfiability is similar to the feasibility property in [2] requiring that any initial segment of a computation must be possibly extended to a full computation to prevent “a scheduler from ‘painting itself into a corner’ with no possible continuation”.

An interesting future development is to investigate more general accounts of satisfiability which is applicable to different scenarios. In scope of future work, we will design *amendment strategies* to so to refine our methodology and maximise the applicability of the proposed algorithms (cf. Section 4.5.1).

Choreography Synthesis as Contract Agreement

We introduce a formal model for distributed systems where each participant advertises its requirements and obligations as behavioural *contracts* (in the form of local types), and where multiparty sessions are started when a set of contracts allows to synthesise a *choreography* (i.e., a global type). The framework is based on the CO₂ calculus and preserves its main features, while allowing to establish sessions where the number of participants is not determined beforehand. We give progress and session fidelity results for CO₂ systems which are “honest”.

5.1 Introduction

Distributed applications are nowadays omnipresent but even for seemingly simple cases, there is still a pressing need to make sure they do work as their designers intended. Indeed, developing and maintaining such systems is rather difficult. This is not only due to intrinsic issues originating from concurrency and distribution; but also because such applications have to be designed within a strange paradox: they are made of components that, on the one hand, must collaborate with each other and, on the other hand, may compete for resources, or for achieving conflicting goals.

We propose a formal model for distributed systems based on *contracts-driven interac-*

tions: components advertise contracts (in the form of local types); such contracts are used at runtime to establish (multiparty) *agreements*, and such agreements steer the behaviour of components. Therefore, contracts are not just a specification or a design mechanism anymore, rather they become a pivotal element of the execution model.

We combine two approaches: *session types* [44] and *contract-oriented computing* [8]. From the former, we adopt concepts, syntax, and semantics – and in particular, the interplay between local behaviours and choreographies (i.e., between local types and global types) as a method for specifying and analysing the interactions of participants in a distributed system. However, in our framework we do not assume that a participant will necessarily always adhere to its specification, nor that a global description is available beforehand to validate the system. From the contract-oriented computing approach, we adopt CO₂ [7] (for CO_ntract-Oriented computing), a generic contract-oriented calculus where participants advertise their requirements and obligations through contracts, and interact with each other once *compliant* contracts have been found. In this work, we tailor CO₂ to a multiparty model where contracts have the syntax of local types. We say that contracts P_1, \dots, P_n are compliant when they can be used to synthesise a global type (cf. Chapter 3). Once a set of compliant contracts has been found, a CO₂ session may be established, wherein the participants who advertised the contracts can interact. However, in line with what may happen in real life cases, the runtime behaviour of these participants may then depart from the contracts: the calculus allows to model and identify these situations.

Our framework models multiparty contractual agreements as “tangible” objects, i.e., choreographies. This allows us to rely on results from Chapter 3; in particular, the fact that well-formed choreographies ensure that contractual agreements enjoy safety and liveness properties. Furthermore, it allows us to easily check that some meta-level properties are satisfied at runtime, e.g., on the number of involved participants, whether or not the session may terminate, etc.

Our adaptation of CO_2 to a choreography-based contract model preserves the properties of the original calculus. In particular, if a system gets stuck, it is possible to identify which participants violated their contracts.

We also introduce global progress and session fidelity in CO_2 , inspired by analogous concepts in theories based on session types (cf. Section 2.1.2). We show that they hold in systems where all participants are *honest* (i.e., they always respect their contracts in any context) – even when a participant takes part in multiple sessions.

5.2 A Motivating Example

We introduce a running example to illustrate our framework. We denote participant variables in bold font, i.e., \mathbf{s} , \mathbf{b}_1 , \mathbf{b}_2 , ... are participant variables; while s , b_1 , b_2 , ... are participant identifiers, as in previous chapters.

Consider the following distributed scenario: an online store s allows any two buyers \mathbf{b}_1 and \mathbf{b}_2 to make a joint purchase through a simplified protocol: once they both have requested the same item, a quote is sent to \mathbf{b}_1 , who is then expected to either place an order (*order*) or end the session (*bye*); the store also promises to let \mathbf{b}_2 know whether the order was placed (*ok*) or cancelled (*bye*). The behaviour of the store s is described by the following contract:

$$P_s = \mathbf{b}_1?req; \mathbf{b}_2?req; \mathbf{b}_1!quote; (\mathbf{b}_1?order; \mathbf{b}_2!ok + \mathbf{b}_1?bye; \mathbf{b}_2!bye)$$

We would like to know which contracts would be *compliant* with P_s . A possible pair of compliant contracts, advertised by buyers b_1 and b_2 , is as follows:

$$P_{b_1} = s!req; s?quote; (\mathbf{b}'_2!ok; s!order \oplus \mathbf{b}'_2!bye; s!bye)$$

$$P_{b_2} = s'!req; (\mathbf{b}'_1?ok; s'?ok + \mathbf{b}'_1?bye; s'?bye)$$

where b_1 promises to send the request to the store (s), wait for the quote, and then notify the other buyer (b'_2) before accepting or rejecting the store offer. Symmetrically, b_2 promises to send the request to the store (s'), and then expects to receive the same notification (either ok or bye) from both the other buyer (b'_1) and the store s' .

An agreement among s , b_1 and b_2 may be found by replacing the participant variables in each contract with actual identifiers, and composing them in a system such as:

$$S_{sb_1b_2} = s[P_s \{b_1/b_1, b_2/b_2\}] \mid b_1[P_{b_1} \{s/s, b_2/b'_2\}] \mid b_2[P_{b_2} \{s/s', b_1/b'_1\}]$$

from which we may obtain the following global type, as per the construction of Chapter 3,

$$\begin{aligned} \mathcal{G}_{sb_1b_2} = & b_1 \rightarrow s : \text{req}; b_2 \rightarrow s : \text{req}; s \rightarrow b_1 : \text{quote}; \\ & \left(\begin{array}{l} b_1 \rightarrow b_2 : \text{ok}; b_1 \rightarrow s : \text{order}; s \rightarrow b_2 : \text{ok} \\ + \\ b_1 \rightarrow b_2 : \text{bye}; b_1 \rightarrow s : \text{bye}; s \rightarrow b_2 : \text{bye} \end{array} \right) \end{aligned}$$

Recall that the synthesis of $\mathcal{G}_{sb_1b_2}$ guarantees that the global type is well-formed and projectable, and that the system consisting of its projections is equivalent to the original system of contracts.

In realistic scenarios, the existence of a contractual agreement among participants does not guarantee that progress and safety will also hold at runtime: in fact, a participant may advertise a contract promising some behaviour, and then fail to respect it – either maliciously or accidentally. Such failures may then cascade on other participants, e.g., if they remain stuck waiting for a promised message that is never sent.

This sort of situations can be modelled in the CO_2 calculus, formally defined in Section 5.4. A CO_2 system (or *network*) for the store-and-two-buyer example may be imple-

mented as follows:

$$T_{sb_1b_2} = (x, y, z) \left(s \llbracket \text{tell}_s \downarrow_x P_s \cdot \text{fuse} \cdot R_s \rrbracket \mid b_1 \llbracket \text{tell}_s \downarrow_y P_{b_1} \cdot R_{b_1} \rrbracket \mid b_2 \llbracket \text{tell}_s \downarrow_z P_{b_2} \cdot R_{b_2} \rrbracket \right)$$

Here, participant s advertises its contract P_s to itself via the primitive $\text{tell}_s \downarrow_x P_s$, where x is used as a session handle for interacting with other participants. Participants b_1 and b_2 advertise their respective contracts to s with a similar invocation.

In this example, s also plays the role of *contract broker*: once all the contracts have been advertised, the *fuse* primitive can establish a new session, based on the fact that the global agreement $G_{sb_1b_2}$ can be synthesised from P_s , P_{b_1} and P_{b_2} . This new session is shared among participants s , b_1 and b_2 .

At this point, the execution of the network (i.e., the continuation of processes R_s , R_{b_1} , and R_{b_2}) is not required to respect the contracts. In fact, we will see that when the contracts are violated, the calculus allows for *culpable* participants to be always identified. Furthermore, we will discuss *honesty*, i.e., the guarantee that a participant will always fulfil its advertised contracts – even in contexts where other participants fail to fulfil theirs. When such a guarantee holds, the contractual progress and safety are also reflected in the runtime behaviour of the network.

Other compliant contracts. Our contract model allows for other scenarios. For instance, a participant b_{12} may impersonate both buyers, and promise to always accept the store offer, by advertising the following contract:

$$P_{b_{12}} = s''! \text{req}; s''! \text{req}; s''? \text{quote}; s''! \text{order}; s''? \text{ok}$$

where the request to the store (s'') is sent twice (i.e., once for each impersonated buyer). In this case, if we combine P_s and $P_{b_{12}}$ with substitutions $\{s/s''\}$ and $\{b_{12}/b_1, b_2\}$, we can

find an agreement by synthesising the following global type:

$$\begin{aligned} \mathcal{G}_{sb_{12}} = & \text{b}_{12} \rightarrow \text{s} : \text{req}; \text{b}_{12} \rightarrow \text{s} : \text{req}; \\ & \text{s} \rightarrow \text{b}_{12} : \text{quote}; \text{b}_{12} \rightarrow \text{s} : \text{order}; \\ & \text{s} \rightarrow \text{b}_{12} : \text{ok} \end{aligned}$$

This scenario may be modelled with the following network:

$$T_{sb_{12}} = (x, w) \left(\text{s} \llbracket \text{tell}_{\text{s}} \downarrow_x P_{\text{s}} \cdot \text{fuse} \cdot R_{\text{s}} \rrbracket \mid \text{b}_{12} \llbracket \text{tell}_{\text{s}} \downarrow_w P_{\text{b}_{12}} \cdot R_{\text{b}_{12}} \rrbracket \right)$$

where the fuse prefix can now create a session involving s and b_{12} .

The participants in the networks $T_{sb_1b_2}$ and $T_{sb_{12}}$ may also be combined, so to obtain:

$$\begin{aligned} T_{sb_1b_2b_{12}} = & (x, y, z, w) \left(\text{s} \llbracket \text{tell}_{\text{s}} \downarrow_x P_{\text{s}} \cdot \text{fuse} \cdot R_{\text{s}} \rrbracket \right. \\ & \mid \text{b}_1 \llbracket \text{tell}_{\text{s}} \downarrow_y P_{\text{b}_1} \cdot R_{\text{b}_1} \rrbracket \mid \text{b}_2 \llbracket \text{tell}_{\text{s}} \downarrow_z P_{\text{b}_2} \cdot R_{\text{b}_2} \rrbracket \\ & \left. \mid \text{b}_{12} \llbracket \text{tell}_{\text{s}} \downarrow_w P_{\text{b}_{12}} \cdot R_{\text{b}_{12}} \rrbracket \right) \end{aligned}$$

In this case, after all contracts have been advertised to s , either a session corresponding to $\mathcal{G}_{sb_1b_2}$, or to $\mathcal{G}_{sb_{12}}$, may take place, thus involving a different number of participants depending on which contracts are fused.

5.3 A Choreography-Based Contract Model

We introduce a contract model based on the results of Chapter 3. Individual contracts are expressed using the syntax of local types; while contractual compliance is based on the synthesis of global types: a set of contracts is compliant if it is possible to synthesise a choreography from it.

Local types as contracts. Let \mathcal{P} be a set of *participant variables* ranged over by \mathbf{s} , \mathbf{r} , \mathbf{n} , etc. and let $\underline{\mathbf{s}}$, $\underline{\mathbf{r}}$, $\underline{\mathbf{n}}$, etc. range over $\mathbb{P} \cup \mathcal{P}$ (recall that \mathbb{P} is the set of *participant names*). The syntax of systems and contracts below is the same as in Section 3.2 but for the introduction of participant variables.

$$\begin{aligned} S, S' & ::= S \mid S' & | & \mathbf{n}[P] & | & \mathbf{sr} : \rho & | & \mathbf{0} \\ P, Q & ::= \bigoplus_{i \in I} \underline{\mathbf{r}}_i ! a_i ; P_i & | & \sum_{i \in I} \underline{\mathbf{s}}_i ? a_i ; P_i & | & \mu \chi . P & | & \chi \end{aligned}$$

As before, we assume that there is at most one queue per pair of participants, (i.e., one channel per direction), that participant names are pairwise distinct in a system S , and that local types are closed.

We write $S \xrightarrow{\mathbf{s} \leftrightarrow \mathbf{r} : a} S'$ when either $S \xrightarrow{\mathbf{s} \rightarrow \mathbf{r} : a} S'$ or $S \xrightarrow{\mathbf{s} \leftarrow \mathbf{r} : a} S'$, where $\mathbf{s} \rightarrow \mathbf{r} : a$ indicates that \mathbf{s} puts a datum on a queue \mathbf{sr} , and label $\mathbf{s} \leftarrow \mathbf{r} : a$ indicates that \mathbf{s} retrieves a datum from the queue \mathbf{rs} (cf. Section 3.2).

Choreography synthesis as compliance. Essentially, in the rest of this chapter, we say that a set of contracts is *compliant*, if it can be assigned a global type, possibly after applying a set of substitutions on the contracts.

Below, we give an example to illustrate the use of local types as contracts and choreography synthesis as a compliance relation between contracts.

Example 5.1. Building up on the example from the previous section, we combine the contract of store \mathbf{s} with those of customers \mathbf{b}_1 and \mathbf{b}_2 , and we obtain the system:

$$\begin{aligned} S_{\mathbf{s}\mathbf{b}_1\mathbf{b}_2} & = \mathbf{s}[P_{\mathbf{s}} \{ \mathbf{b}_1 / \mathbf{b}_1 \} \{ \mathbf{b}_2 / \mathbf{b}_2 \}] \mid \mathbf{b}_1[P_{\mathbf{b}_1} \{ \mathbf{s} / \mathbf{s} \} \{ \mathbf{b}_2 / \mathbf{b}'_2 \}] \mid \mathbf{b}_2[P_{\mathbf{b}_2} \{ \mathbf{s} / \mathbf{s}' \} \{ \mathbf{b}_1 / \mathbf{b}'_1 \}] \\ & = \mathbf{s}[\mathbf{b}_1 ? \text{req} ; \mathbf{b}_2 ? \text{req} ; \mathbf{b}_1 ! \text{quote} ; (\mathbf{b}_1 ? \text{order} ; \mathbf{b}_2 ! \text{ok} + \mathbf{b}_1 ? \text{bye} ; \mathbf{b}_2 ! \text{bye})] \\ & \quad \mid \mathbf{b}_1[\mathbf{s} ! \text{req} ; \mathbf{s} ? \text{quote} ; (\mathbf{b}_2 ! \text{ok} ; \mathbf{s} ! \text{order} \oplus \mathbf{b}_2 ! \text{bye} ; \mathbf{s} ! \text{bye})] \\ & \quad \mid \mathbf{b}_2[\mathbf{s} ! \text{req} ; (\mathbf{b}_1 ? \text{ok} ; \mathbf{s} ? \text{ok} + \mathbf{b}_1 ? \text{bye} ; \mathbf{s} ? \text{bye})] \end{aligned}$$

which is assigned the following global type:

$$\mathcal{G}_{sb_1b_2} = b_1 \rightarrow s : \text{req}; b_2 \rightarrow s : \text{req}; s \rightarrow b_1 : \text{quote};$$

$$\left(\begin{array}{c} b_1 \rightarrow b_2 : \text{ok}; b_1 \rightarrow s : \text{order}; s \rightarrow b_2 : \text{ok} \\ + \\ b_1 \rightarrow b_2 : \text{bye}; b_1 \rightarrow s : \text{bye}; s \rightarrow b_2 : \text{bye} \end{array} \right)$$

that is to say that $\circ \vdash S_{sb_1b_2} \blacktriangleright \mathcal{G}_{sb_1b_2}$ holds.

If we combine the store s with b_{12} we have

$$\begin{aligned} S_{sb_{12}} &= s[P_s \{b_{12}/b_1, b_2\}] \mid b_{12}[P_{b_{12}} \{s/s''\}] \\ &= s[b_{12}?req; b_{12}?req; b_{12}!quote; (b_{12}?order; b_{12}!ok + b_{12}?bye; b_{12}!bye)] \\ &\quad \mid b_{12}[s!req; s!req; s?quote; s!order; s?ok] \\ \mathcal{G}_{sb_{12}} &= b_{12} \rightarrow s : \text{req}; b_{12} \rightarrow s : \text{req}; \\ &\quad s \rightarrow b_{12} : \text{quote}; b_{12} \rightarrow s : \text{order}; \\ &\quad s \rightarrow b_{12} : \text{ok} \end{aligned}$$

and, again, the judgement $\circ \vdash S_{sb_{12}} \blacktriangleright \mathcal{G}_{sb_{12}}$ holds. \diamond

5.4 Contract-Oriented Computing and Choreographies

We first introduce a version of the CO₂ calculus [7] adapted to multiparty contracts and sessions. Secondly, we discuss several options for fine-grained control of session establishment, which are made possible thanks to the fact that we base contract agreements on the existence of a choreography.

5.4.1 A Choreography-Based CO₂

Let \mathbb{K} and \mathcal{K} be disjoint sets of, respectively, *session names* (ranged over by k, k', \dots) and *session variables* (ranged over by x, y, z, \dots). Let u, v, \dots range over $\mathbb{K} \cup \mathcal{K}$.

The syntax of CO₂ is as follows:

$$\begin{aligned}
 R & ::= \sum_{i \in I} p_i \cdot R_i \mid R \mid R \mid (\vec{u}, \vec{n})R \mid X(\vec{u}, \vec{n}) \mid \mathbf{0} && \text{[PROCESSES]} \\
 p & ::= \tau \mid \text{tell}_{\vec{n}} \downarrow_u P \mid \text{fuse} \mid \text{do}_{\vec{n}}^u a && \text{[PREFIXES]} \\
 K & ::= \downarrow_u \vec{n} \text{says } P \mid K \mid K && \text{[LATENT CONTRACTS]} \\
 T & ::= \vec{n} \llbracket R \rrbracket \mid \vec{n} \llbracket K \rrbracket \mid k \llbracket S \rrbracket \mid T \mid T \mid (\vec{u}, \vec{n})T \mid \mathbf{0} && \text{[NETWORKS]}
 \end{aligned}$$

CO₂ features CCS-style processes, equipped with branching \sum (not to be confused with the choice operator used in contracts), parallel composition \mid , restrictions of session and participant variables, and named process invocation. We often write $R + R'$ for the binary version of \sum , and we consider $+$ to be an associative and commutative operator. The prefixes are for internal action (τ), contract advertisement ($\text{tell} \downarrow$), session creation upon contractual agreement (fuse), and execution of contractual actions (do). A *latent contract* of the form $\downarrow_u \vec{n} \text{says } P$ represents the promise of participant \vec{n} to fulfil P by executing do -actions on a session variable u . A network T may be the parallel composition of processes of the form $\vec{n} \llbracket R \rrbracket$ (where \vec{n} is the participant executing R), *latent contracts* $\vec{n} \llbracket K \rrbracket$ (where \vec{n} is the participant to whom the contracts in K have been advertised), and established sessions $k \llbracket S \rrbracket$ (where k is a session name, and S is a system of *stipulated* contracts as in Section 5.3). Note that CO₂ process and network productions allow to delimit both session names/variables (\vec{u}) and participant variables (\vec{n}), but *not* participant names, which are considered public. Hereafter, we assume that bound participant variables are pairwise distinct.

Remark 5.1. *Syntactically, the only difference between our version of CO₂ and the calculus in [4] is that, in our adaptation, $\text{do}_{\vec{n}}^u a$ -prefixes mention the participant towards which*

$$\begin{array}{c}
 \text{[TELL]} \quad \mathfrak{s} \llbracket \text{tell}_r \downarrow_x P . R + R' \mid R'' \rrbracket \rightarrow \mathfrak{s} \llbracket R \mid R'' \rrbracket \mid r \llbracket \downarrow_x \mathfrak{s} \text{ says } P \rrbracket \\
 \\
 \text{[FUSE]} \quad \frac{K \triangleright_{\pi}^{\sigma} S \quad \vec{\mathfrak{s}} = \text{dom}(\pi) \quad \vec{u} = \text{dom}(\sigma) \quad \text{img}(\sigma) = \{k\} \quad k \text{ fresh}}{(\vec{u}, \vec{\mathfrak{s}}) (\mathfrak{s} \llbracket \text{fuse} . R + R' \mid R'' \rrbracket \mid \mathfrak{s} \llbracket K \rrbracket \mid T) \rightarrow (k) (\mathfrak{s} \llbracket R \mid R'' \rrbracket \sigma \pi \mid k \llbracket S \mid Q(S) \rrbracket \mid T \sigma \pi)} \\
 \\
 \text{[DO]} \quad \frac{S \xrightarrow{\mathfrak{s} \leftarrow r : a} S'}{k \llbracket S \rrbracket \mid \mathfrak{s} \llbracket \text{do}_r^k a . R + R' \mid R'' \rrbracket \rightarrow k \llbracket S' \rrbracket \mid \mathfrak{s} \llbracket R \mid R'' \rrbracket} \\
 \\
 \text{[TAU]} \quad \mathfrak{s} \llbracket \tau . R + R' \mid R'' \rrbracket \rightarrow \mathfrak{s} \llbracket R \mid R'' \rrbracket \\
 \\
 \text{[DEF]} \quad \frac{X(\vec{u}, \vec{\mathfrak{n}}) \stackrel{\text{def}}{=} R \quad (\vec{z}, \vec{\mathfrak{n}}') (\mathfrak{s} \llbracket R \{ \vec{v} / \vec{u}, \vec{x} / \vec{\mathfrak{n}} \} \mid R' \rrbracket \mid T) \rightarrow T'}{(\vec{z}, \vec{\mathfrak{n}}') (\mathfrak{s} \llbracket X(\vec{v}, \vec{x}) \mid R' \rrbracket \mid T) \rightarrow T'} \\
 \\
 \text{[PAR]} \quad \frac{T \rightarrow T'}{T \mid T'' \rightarrow T' \mid T''} \quad \text{[RES]} \quad \frac{T \rightarrow T'}{(\vec{u}, \vec{\mathfrak{s}}) T \rightarrow (\vec{u}, \vec{\mathfrak{s}}) T'}
 \end{array}$$

 Figure 5.1: Semantics rules for CO₂

the action a is directed.

The semantics of CO₂ is given by the rules in Figure 5.1. Rule [TELL] allows a participant \mathfrak{s} to advertise a contract P to r ; as a result, a new *latent* contract is created, recording the fact that it was promised by \mathfrak{s} . Rule [FUSE] establishes a new session: the latent contracts held in $\mathfrak{s} \llbracket K \rrbracket$ are combined, and their participant variables substituted, in order to find an *agreement*, i.e., a system of local types S that satisfies the relation $K \triangleright_{\pi}^{\sigma} S$ (see Definition 5.1 below). Recall that $Q(S)$ stands for the system of empty queues connecting all pairs of participants in S . Provided that an agreement is found, fresh session name k and participants names are shared among the parties, via substitutions σ and π ; within the session, the involved contracts become *stipulated* (as opposed to “latent”, before the agreement). Rule [DO] allows \mathfrak{s} to perform an input/output action a towards r on session k , provided that S permits it. Rule [TAU] allows a participant to make an internal move. Rule [DEF] deals with named process invocation. Observe that a process invocation may

contain session variables/names and participant variables/names. Rules $_{[PAR]}$ and $_{[RES]}$ are standard.

When needed, we label network transitions: $T \xrightarrow{s:p} T'$ means that T reduces to T' through a prefix p fired by participant s .

Remark 5.2. *In the following, we assume that sessions, i.e., sub-networks of the form $k[S]$, are only created at the runtime by rule $_{[FUSE]}$.*

Example 5.2. Consider the network:

$$T_{\text{ex5.2}} = s \left[\left[\text{do}_r^k \text{int} + \text{do}_r^k \text{bool} \right] \mid k[s[r!\text{int}] \mid r[s?\text{int}] \mid sr : \varepsilon \mid rs : \varepsilon] \mid r \left[\left[\text{do}_s^k \text{int} \right] \right]$$

Here, the CO_2 process of participant s can perform an action towards r on session k , with either a message of sort int or bool . However, s 's contract in k only specifies that s should send a message of sort int to r : therefore, according to rule $_{[DO]}$, only the first branch of s may be chosen, and the network reduces as follows:

$$\begin{array}{l} T_{\text{ex5.2}} \xrightarrow{s: \text{do}_r^k \text{int}} s[\mathbf{0}] \mid k[s[\mathbf{0}] \mid r[s?\text{int}] \mid sr : \text{int} \mid rs : \varepsilon] \mid r \left[\left[\text{do}_s^k \text{int} \right] \right] \\ \xrightarrow{r: \text{do}_s^k \text{int}} s[\mathbf{0}] \mid k[s[\mathbf{0}] \mid r[\mathbf{0}] \mid sr : \varepsilon \mid rs : \varepsilon] \mid r[\mathbf{0}] \end{array}$$

◇

A main difference between our adaptation of CO_2 and the original presentation comes from the way we specify session establishment. We adopt the session agreement relation defined below.

Definition 5.1 (Agreement relation $K \triangleright_{\pi}^{\sigma} S$). *Let*

$$K \equiv \prod_{i \in I} \downarrow_{x_i} s_i \text{ says } P_i \quad \text{such that } \forall i \neq j \in I : s_i \neq s_j$$

and let $\pi: \bigcup_{i \in I} \text{fv}(P_i) \rightarrow \mathbb{P}$ and $\sigma: \bigcup_{i \in I} \{x_i\} \rightarrow \mathbb{K}$ be two substitutions mapping participant variables to participant names, and session variables to session names, respectively.

Also, let S be the system of local types $\prod_{i \in I} s_i [P_i] \pi$.

We define:

$$K \triangleright_{\pi}^{\sigma} S \iff \forall i \in I: \forall \mathbf{n} \in \text{fv}(P_i): \pi(\mathbf{n}) \neq s_i \quad \wedge \quad \exists \mathcal{G} : \circ \vdash S \triangleright \mathcal{G}$$

Intuitively, a system of stipulated contracts S is constructed from a set of latent contracts K , using a substitution π that maps all the participant variables in K to the participant names in K . The definition requires that, within a contract P_i , belonging to s_i , no participant variable is substituted by s_i itself. If it is possible to synthesise a global type \mathcal{G} out of S , then the relation holds, and a contractual agreement exists. Note that due to the condition imposed on K , each participant may have at most one contract in a given session. Example 5.3 below illustrates Definition 5.1.

Example 5.3. Consider the following network, with s, b_1, b_2 from Section 5.2, and $S_{sb_1b_2}$ from Example 5.1:

$$\begin{aligned} T_{sb_1b_2} &= (x, y, z) (s \llbracket \text{tell}_s \downarrow_x P_s \cdot \text{fuse} \cdot R_s \rrbracket \mid b_1 \llbracket \text{tell}_s \downarrow_y P_{b_1} \cdot R_{b_1} \rrbracket \mid b_2 \llbracket \text{tell}_s \downarrow_z P_{b_2} \cdot R_{b_2} \rrbracket) \\ &\rightarrow \rightarrow \rightarrow \\ &(x, y, z) (s \llbracket \text{fuse} \cdot R_s \rrbracket \mid s \llbracket \downarrow_x s \text{ says } P_s \mid \downarrow_y b_1 \text{ says } P_{b_1} \mid \downarrow_z b_2 \text{ says } P_{b_2} \rrbracket \mid b_1 \llbracket R_{b_1} \rrbracket \mid b_2 \llbracket R_{b_2} \rrbracket) \\ &\xrightarrow{s: \text{fuse}} \\ (k) T'_{sb_1b_2} &= (k) (s \llbracket R_s \rrbracket \sigma \pi \mid k \llbracket S_{sb_1b_2} \mid Q(S_{sb_1b_2}) \rrbracket \mid b_1 \llbracket R_{b_1} \rrbracket \sigma \pi \mid b_2 \llbracket R_{b_2} \rrbracket \sigma \pi) \\ &\text{where } \sigma = \{k/x, y, z\} \quad \text{and} \quad \pi = \{s/s', b_1/b_1', b_2/b_2'\} \end{aligned} \quad (5.1)$$

The initial network $T_{sb_1b_2}$ is the one considered in Section 5.2, where all the participants are ready to advertise their respective contracts to the store s , by using a $\text{tell}_s \downarrow$ -primitive. This has the effect of creating corresponding latent contracts within s . Once all the latent contracts are in a same location, they may be fused.

$$\begin{aligned}
 (\vec{u}, \vec{r})\mathfrak{s}[(\vec{v}, \vec{n})P] &\equiv (\vec{u}, \vec{r})(\vec{v}, \vec{n})\mathfrak{s}[P] & \mathfrak{s}[\mathbf{0}] &\equiv \mathbf{0} & \mathfrak{s}[[K]] \mid \mathfrak{s}[[K']] &\equiv \mathfrak{s}[[K \mid K']] \\
 Z \mid \mathbf{0} &\equiv Z & Z \mid Z' &\equiv Z' \mid Z & (Z \mid Z') \mid Z'' &\equiv Z \mid (Z' \mid Z'') \\
 Z \mid (\vec{u}, \vec{n})Z' &\equiv (\vec{u}, \vec{n})(Z \mid Z') & \text{if } \vec{u} \cap \text{fnv}(Z) &= \vec{n} \cap \text{fnv}(Z) = \emptyset \\
 (\vec{u}, \vec{n})(\vec{v}, \vec{n})Z &\equiv (\vec{v}, \vec{n})(\vec{u}, \vec{n})Z & (\vec{u}, \vec{n})(\vec{v}, \vec{n})Z &\equiv (\vec{u}\vec{v}, \vec{n}\vec{n})Z \\
 (\vec{u}, \vec{n})Z &\equiv Z & \text{if } \vec{u} \cap \text{fnv}(Z) &= \vec{n} \cap \text{fnv}(Z) = \emptyset
 \end{aligned}$$

 Figure 5.2: Congruence rules for CO_2

In this case, given σ and π as in (5.1), Definition 5.1 is indeed applicable; we already saw that the system

$$\mathfrak{s}[P_s\pi] \mid \mathfrak{b}_1[P_{b_1}\pi] \mid \mathfrak{b}_2[P_{b_2}\pi] \quad (5.2)$$

may be assigned a global type. Hence, a new session k is created, based on the system of contracts in (5.2) and the queues connecting all pairs of participants.

The session variables of the latent contracts being fused (i.e., x for participant s , y for b_1 , and z for b_2) are all substituted with the fresh session name k in the processes R_s , R_{b_1} and R_{b_2} , via σ . Similarly for participant variables which are substituted with participant names, via π . \diamond

The semantics of CO_2 is to be considered up-to a standard structural congruence relation \equiv , given in Figure 5.2, where Z, Z', Z'' range over processes, networks, and latent contracts. We write $\text{fnv}(Z)$ for the free session and participant variables in Z . Remarkably, the rule

$$\mathfrak{s}[[K]] \mid \mathfrak{s}[[K']] \equiv \mathfrak{s}[[K \mid K']]$$

allows to select subsets of latent contracts before invoking a fuse primitive. This allows to (momentarily) disregard some contracts when searching for a global type to be synthesised. We illustrate this in Example 5.4 below.

Example 5.4. Consider the network:

$$\dots r \llbracket \text{fuse} . R \mid R' \rrbracket \mid r \llbracket \downarrow_x s_1 \text{ says } \mathbf{n}_1 ! \text{int} \mid \downarrow_y s_2 \text{ says } \mathbf{n}_2 ? \text{int} \mid \downarrow_z s_3 \text{ says } \mathbf{n}_3 ? \text{bool} \rrbracket \mid \dots$$

The fuse prefix cannot be fired: no contract matches s_3 's, and thus, together, the three latent contracts cannot be assigned a global type. However, by rearranging the network with congruence \equiv , we have:

$$\dots r \llbracket \text{fuse} . R \mid R' \rrbracket \mid r \llbracket \downarrow_x s_1 \text{ says } \mathbf{n}_1 ! \text{int} \mid \downarrow_y s_2 \text{ says } \mathbf{n}_2 ? \text{int} \rrbracket \mid r \llbracket \downarrow_z s_3 \text{ says } \mathbf{n}_3 ? \text{bool} \rrbracket \mid \dots$$

It is now possible to synthesise a global type $s_1 \rightarrow s_2 : \text{int}$, and a session may be created for s_1 and s_2 . The latent contract of s_3 may be fused later on. \diamond

In Example 5.5 below, we return to the running example of Chapter 3 and show that introducing participant variables adds flexibility to the synthesis of choreographies.

Example 5.5. We reuse the local types from the running example of Chapter 3 (cf. Section 3.1), and consider them as contracts advertised by their respective participants, these contracts may now feature participant variables.

Recall that the system consists of two buyers (b_1 and b_2) and two sellers (s_1 and s_2). The expected overall behaviour is that the two buyers should agree on purchasing concurrently two items from two sellers, and then try to obtain the best price for these items. The contracts are as follows:

$$\begin{aligned} P_{b_1} &= b_2 ! \text{agreement} ; b_2 ? \text{ack} ; Q_1 \\ P_{b_2} &= b_1 ? \text{agreement} ; b_1 ! \text{ack} ; Q_2 \\ Q_i &= s_i ! \text{request} ; s_i ? \text{quote} ; \mu \chi . (s_i ! \text{ok} \oplus s_i ! \text{no} ; s_i ? \text{quote} ; \chi) \quad i \in \{1, 2\} \\ S_{s_i} &= b_i ? \text{request} ; b_i ! \text{quote} ; \mu \chi' . (b_i ? \text{ok} + b_i ? \text{no} ; b_i ! \text{quote} ; \chi') \quad i \in \{1, 2\} \end{aligned} \tag{5.3}$$

Observe that the contracts now features four *participant variables*: b_i and s_i ($i \in \{1, 2\}$).

A network for these four participants may be implemented as follows:

$$T_{\text{ex5.5}} = (x, y, z, w) \left(\begin{array}{l} b_1 \llbracket \text{tell}_{b_1} \downarrow_x P_{b_1} \cdot \text{fuse} \cdot R_{b_1} \rrbracket \mid b_2 \llbracket \text{tell}_{b_1} \downarrow_y P_{b_2} \cdot R_{b_2} \rrbracket \\ \mid s_1 \llbracket \text{tell}_{b_1} \downarrow_z P_{s_1} \cdot R_{s_1} \rrbracket \mid s_2 \llbracket \text{tell}_{b_1} \downarrow_w P_{s_2} \cdot R_{s_2} \rrbracket \end{array} \right)$$

where each participant advertises its contract to b_1 . Interestingly, the pairs buyer/seller are not determined in advance, i.e., we may have that b_1 buys an item from s_2 , because two sets of substitutions are possible:

$$\pi_1 = \{s_1/s_1, s_2/s_2, b_1/b_1, b_2/b_2\} \quad \text{or} \quad \pi_2 = \{s_1/s_2, s_2/s_1, b_1/b_2, b_2/b_1\}$$

Now, assume that the contracts of b_1 and s_1 are slightly modified so that b_1 wants to buy a book, while the seller s_1 sells only books (i.e., it only accepts book requests). We have the following two new contracts:

$$\begin{aligned} P_{b_1} &= b_2! \text{agreement}; b_2? \text{ack}; Q_1 \\ Q_1 &= s_1! \text{book}; s_1? \text{quote}; \mu \chi. (s_1! \text{ok} \oplus s_1! \text{no}; s_1? \text{quote}; \chi) \\ S_{s_1} &= b_1? \text{book}; b_1! \text{quote}; \mu \chi'. (b_1? \text{ok} + b_1? \text{no}; b_1! \text{quote}; \chi') \end{aligned}$$

If the contracts of b_2 and s_2 are the same as in (5.3), the fusion of these contracts is deterministic: b_i will be instantiated to b_i and s_i will be instantiated to s_i (for $i \in \{1, 2\}$). Indeed, this is the only substitution that will allow a global type to be synthesised from the contracts. \diamond

5.4.2 On the Flexibility of Session Establishment

We highlight the flexibility of our definition of contract agreement, together with the semantics of CO_2 , by discussing three features: (i) contracts may use a mix of participant names and variables, (ii) different contracts may use common participant variables, and

(iii) the definition of agreement may be easily accommodated to different requirements.

Mixing participant variables and names

Mixing participant variables and names in contracts allows us to design fine-grained contracts and specify relationships between them.

- **Contracts with both participant names and variables.** A seller s may want to sell an item to a *specific* buyer b , via *any* shipping company that provides a package tracking system. The contract of s may be:

$$b!price; b?ack; n!request; n?tracking; b!tracking$$

saying that the seller s must send a price to the buyer b ; once b has acknowledged, s must send a shipping request to a shipper n ; who must send back a tracking number, which is then forwarded to b . This contract may be fused only if b takes part in the session, while the role of shipper n may be played by any participant (offering a compliant contract).

- **Contracts sharing participant variables.** Consider the network:

$$s \llbracket \text{tell}_{s \downarrow x}(\mathbf{n}!request; \mathbf{n}?ack) \dots X(\vec{z}, \mathbf{n}) \rrbracket \mid \dots$$

where $X(\vec{z}, \mathbf{n}) \stackrel{def}{=} (y, \mathbf{b}) \text{tell}_{s \downarrow y}(\mathbf{b}!quote; \dots \mathbf{n}!address) \dots X(\vec{z}, \mathbf{n})$

Here, s advertises two contracts: the first one ($\mathbf{n}!request; \mathbf{n}?ack$) is used by s to find a shipping company, and the second ($\mathbf{b}!quote; \dots \mathbf{n}!address$) to sell items. The two contracts are linked by the common variable \mathbf{n} : whenever the first one is fused, variable \mathbf{n} is instantiated to a participant name, say shipper, which is also substituted in the second. This means that whenever a new selling session starts, shipper will also be involved as the receiver of the address message.

- **Injective substitution.** Definition 5.1 is quite liberal with respect to the substitution mapping participant variables to participant names (π). The definition may be refined easily to, e.g., force π to be an injective substitution and therefore not allowing two participant variables to be mapped to the same participant identifier.

In the example of Section 5.2, such a variation would allow the seller to disregard the second set of contracts (where one participant impersonates both buyers).

Parametrised session establishment

The participants firing fuse primitives are playing the role of brokers in our framework. Depending on their implementation, brokers may also have some obligations in the contracts they fuse, or they may want to enforce some general policy – therefore they may have additional requirements before agreeing to start a session. For instance, a broker may not want to start a session with too many participants as it may be too resource demanding (too many connections etc.). Another broker may want to start sessions that terminates within a limited number of interactions, because it has a short life expectancy, e.g., due to an approaching scheduled maintenance. Another kind of broker may precisely want to start sessions which do not terminate, e.g., if the broker is interested in resilient services. Several variations of the fuse primitive are possible thanks to the fact that we base contract agreements on objects representing the overall choreography.

In this section, we discuss some of these variations, starting from simple ones to more complex ones. The overall idea of these variation is to parametrise the fuse primitive with a function, say F^{fuse} , that takes as input a global type \mathcal{G} and returns a boolean value.¹ Definition 5.1 stays essentially the same but for the slightly refined condition:

$$\exists \mathcal{G} : \circ \vdash S \blacktriangleright \mathcal{G} \quad \wedge \quad F^{\text{fuse}}(\mathcal{G})$$

¹ We consider F^{fuse} to be a computable function.

Remark 5.3. *This sort of variations does not affect the results that follow, since the original fuse primitive is also blocking. The variations only restrict the sets of compliant contracts.*

Some variations. We introduce $\text{fuse}[n]$, a version of fuse that only fuses sessions where there are at least n participants; fuse_T , which has the additional constraint that no recursive behaviour is allowed in the synthesised choreography (therefore ensuring that the session will eventually terminate), and fuse_R , which only creates sessions when the synthesised choreography never terminates (i.e., it only consists of recursive behaviours).

Such extensions may be defined directly via simple definitions of function F^{fuse} :

- for $\text{fuse}[n]$, we define $F^{\text{fuse}}(\mathcal{G}) \stackrel{\text{def}}{=} |\mathcal{P}(\mathcal{G})| \geq n$, recall that $\mathcal{P}(\mathcal{G})$ is the set of participants in \mathcal{G} ;
- for fuse_T , we define $F^{\text{fuse}}(\mathcal{G}) \stackrel{\text{def}}{=} \text{bv}(\mathcal{G}) = \emptyset$, i.e., the condition holds only if there is no recursion variable χ in \mathcal{G} (recall that synthesised global types are closed since we assumed the contracts to be closed);
- for fuse_R , we define $F^{\text{fuse}}(\mathcal{G})$ so that it holds only if $\mathbf{0}$ is not a sub-term of \mathcal{G} , i.e., we add the condition that $\mathbf{0}$ does not appear in \mathcal{G} .

Observe that this kind of properties must be checked for at the global level because it cannot always be decided by looking at the individual contracts. For instance, a participant might exhibit a recursive behaviour in one of the branches of an external choice, while the participant it interacts with may always choose a branch that is not recursive. Consider for instance, the following system of contracts:

$$s[r!a] \mid r[\mu\chi.(s?a + s?b;\chi)]$$

Looking only at the contract of r , it appears to be recursive. However, when composed

with the contract of s , their global type is $s \rightarrow r : a$, i.e., a non-recursive choreography.

Sharing knowledge. In order to illustrate the expressive power of our agreement relation – thanks to the fact that we base it on an object representing the choreography of the session – we give a more advanced version of F^{fuse} that checks semantic properties of the synthesised global type.

Consider a broker whose sole purpose is to fuse contracts only if their composition corresponds to a choreography such that all the participants share the same knowledge once the session has terminated (i.e, when all the participants have fulfilled their contracts in the session). Assume that k is a sort representing the knowledge of a participant; and that a sending operation $r!k$ made by a participant s means that s sends *all its current* knowledge to r , and that a receiving operation $s?k$ made by a participant r means that r acquires all the knowledge that s has (when s sent the message). We need to check that, given a set of contracts, all participants share the same knowledge at the end of the session. For instance, the contracts P_s and P_r below comply with the condition, but $\overline{P_s}$ and $\overline{P_r}$ do not.

$$P_s = r!k; r?k \text{ and } P_r = s?k; s!k \quad \checkmark \quad \overline{P_s} = r!k \text{ and } \overline{P_r} = s?k \quad \times$$

Indeed, on the right-hand side, s does not acquire r 's knowledge. Note that both pairs of contract may be assigned a global type.

We introduce a map $K : \mathbb{P} \rightarrow 2^{\mathbb{P}}$ that, given a participant name n , returns the set of names representing the participants from whom n acquired knowledge. Initially, we have that $K(n) = \{n\}$ for each participant n , i.e., a participant has only access to its own knowledge. Given a global type \mathcal{G} , we would like that

$$\forall n \in \mathcal{P}(\mathcal{G}) : K(n) = \mathcal{P}(\mathcal{G})$$

holds at the end of the session embodied by \mathcal{G} .

We define the boolean function $\hat{K}(\mathcal{G}, N, K)$ which returns true if all the participants in $N \supseteq \mathcal{P}(\mathcal{G})$ share the same knowledge (with respect to K) at the end of the session. We write $\{-/\cdot\}$ for the update operation on maps, and \cdot for the disjoint union of maps.

Definition 5.2. Let $\hat{K}(\mathcal{G}) \stackrel{def}{=} \hat{K}(\mathcal{G}, \mathcal{P}(\mathcal{G}), K_0)$, where K_0 is the initial knowledge map, i.e. $\forall n \in \mathcal{P}(\mathcal{G}) : K_0(n) = \{n\}$, and

$$\hat{K}(\mathcal{G}, N, K) \stackrel{def}{=} \begin{cases} \hat{K}(\mathcal{G}', N, K'), & \text{if } \mathcal{G} = s \rightarrow r : k; \mathcal{G}' \text{ and } K' = K \{K(r) \cup K(s) / r\} \\ \hat{K}(\mathcal{G}_1, N, K) \wedge \hat{K}(\mathcal{G}_2, N, K), & \text{if } \mathcal{G} = \mathcal{G}_1 + \mathcal{G}_2 \\ \hat{K}(\mathcal{G}_1, N, K_1) \wedge \hat{K}(\mathcal{G}_2, N, K_2), & \text{if } \mathcal{G} = \mathcal{G}_1 \mid \mathcal{G}_2 \text{ and } K = K_1 \cdot K_2 \\ \forall n \in \text{dom}(K) : K(n) = N, & \text{if } \mathcal{G} = \mathbf{0} \\ \mathbf{false}, & \text{otherwise} \end{cases}$$

The first case of Definition 5.2 handles interactions: r acquires all the current knowledge of s . The case of choice simply requires that the condition is satisfied in both branches. In the case of concurrent branches, we need to partition the knowledge map in two – since the sets of participants in the concurrent branches of \mathcal{G} are disjoint. In the case of $\mathcal{G} = \mathbf{0}$, we check that all participants share the same knowledge. For the sake of simplicity, the function is undefined for recursive global type.

We may use $\hat{K}(\cdot)$ in session establishment with $F^{\text{fuse}}(\mathcal{G}) \iff \hat{K}(\mathcal{G})$. We illustrate Definition 5.2 in Example 5.6 below.

Example 5.6. Consider the following contracts

$$\begin{aligned} P_s &= r_1 ?k; r_2 ?k; r_3 ?k; r_1 !k; r_2 !k; r_3 !k \\ P_{r_i} &= s !k; s ?k \quad i \in \{1, 2, 3\} \end{aligned}$$

The composition of these contracts, i.e.

$$s[P_s] \mid r_1[P_{r_1}] \mid r_2[P_{r_2}] \mid r_3[P_{r_3}]$$

gives us the following global type

$$\begin{aligned} \mathcal{G}_{\text{ex5.6}} = & r_1 \rightarrow s:k; r_2 \rightarrow s:k; r_3 \rightarrow s:k; \\ & s \rightarrow r_1:k; s \rightarrow r_2:k; s \rightarrow r_3:k; \mathbf{0} \end{aligned}$$

$\hat{K}(\mathcal{G}_{\text{ex5.6}})$ holds since at the end of the first line of $\mathcal{G}_{\text{ex5.6}}$, we have $K(s) = \{s, r_1, r_2, r_3\}$ and, in the second line, s shares its knowledge with all the other participants. \diamond

Interestingly, it is possible to refine session establishments so that only “efficient” contracts are fused. For instance, in Example 5.6, the sharing of knowledge is done in a very centralised way: s acquires the knowledge of all the other participants, then share back all s knows.

One may compute a bound on the number of consecutive interactions that are required for n participants to share their knowledge [48]. If the depth of the parsing tree of the synthesised global type is greater than this bound, then the contracts are not “efficient”. We give a more efficient set of contracts in Example 5.7 below.

Example 5.7. Consider the following contracts

$$\begin{aligned} P_s &= r_1?k; r_2?k; r_3?k; r_1!k; r_3!k & P_{r_2} &= s!k; r_1?k \\ P_{r_1} &= s!k; s?k; r_2!k & P_{r_3} &= s!k; s?k \end{aligned}$$

where the difference between these contracts and those in Example 5.6, is that r_1 is now the one who shares the overall knowledge to participant r_2 . The composition of these

contracts gives us the following global type:

$$\begin{aligned} \mathcal{G}_{\text{ex5.7}} = & \ r_1 \rightarrow s:k; r_2 \rightarrow s:k; r_3 \rightarrow s:k; \\ & \ s \rightarrow r_1:k; (r_1 \rightarrow r_2:k; \mathbf{0} \mid s \rightarrow r_3:k; \mathbf{0}) \end{aligned}$$

We have again that s has all the knowledge at the end of the first line of $\mathcal{G}_{\text{ex5.7}}$, but in this case s and r_1 share back their knowledge concurrently to r_3 and r_2 , respectively. Thus, $\mathcal{G}_{\text{ex5.7}}$ is (slightly) more efficient than $\mathcal{G}_{\text{ex5.6}}$. \diamond

5.5 The Problem of Honesty

In this section, we discuss and define the notion of *honesty* [7], i.e., the ability of a participant to always fulfil its contracts, in any context. In our contract-oriented setting, honesty is essentially the counterpart of well-typedness in a session type setting: the static proof that a participant always honours its contracts provides guarantees about its runtime behaviour.

As seen in Example 5.2, each do prefix within the process of a participant, say $s[[P]]$, is driven by the contract that s promised to abide by. In a sense, CO_2 is *culpability-driven*, according to Definition 5.3 below: when a participant is “culpable”, it has the duty of making the session progress according to its contract.

Definition 5.3 (Culpability). *Let T be a network with a session k , i.e.,*

$$T \equiv (\vec{u}, \vec{n}) (s[[R]] \mid k[[S]] \mid T_1) \mid T_2$$

We say that s is culpable in T when there exist r and a such that $S \xrightarrow{s \leftarrow r:a}$.

A culpable participant can overcome its status by firing its do prefixes, according to rule $[\text{DO}]$, until another participant becomes culpable or session k terminates. Hence, as long as a culpable participant s does not enable a do-prefix matching a contractual action,

s will remain culpable. Note that when a participant is involved in multiple sessions, it may end up being culpable in more than one of them.

When a participant s is always able to fulfil its contractual actions (i.e., overcome its culpability), no matter what other participants do, then it is said to be *honest* (cf. Definition 5.8). This is a desirable property in contract-oriented scenarios: a participant may be stuck in a culpable condition either due to “simple” bugs (cf. Example 5.9), or due to the unexpected (or malicious) behaviour of other participants (cf. Example 5.12). Therefore, before deploying a service, its developers may want to ensure that it will always be able to exculpate itself.

Formally, as in [4], we base the definition of honesty on the relationship between the ready set of a contract and the ready set of a CO₂ process. We call the former *contract ready set* and the latter *process ready set*. The concept of contract ready set is similar to [4, 7, 30], where only bilateral contracts are considered. Here, we adapt it to suit our multiparty contract model, the main difference being that we take into account towards which participant an action is directed.

Definition 5.4 (Contract ready set). *The ready set of a contract P , written $\text{Crs}(P)$, is:*

$$\text{Crs}(P) = \begin{cases} \text{Crs}(P') & \text{if } P = \mu\chi.P' \\ \{(r_i, a_i) \mid i \in I\} & \text{if } P = \bigoplus_{i \in I} r_i!a_i; P_i \text{ and } I \neq \emptyset \\ \{(s, a_i) \mid i \in I\} & \text{if } P = \sum_{i \in I} s?a_i; P_i \end{cases}$$

Intuitively, when a participant n is bound to a contract P , the ready set of P tells which interactions n must be able to perform towards other participants. Each interaction has the form of a pair, consisting of a participant name and a message sort. The interactions offered by an external choice are all available at once, while those offered by an internal choice are mutually exclusive.²

² Recall that for local types, $\mathbf{0}$ is defined as an external/internal choice where $I = \emptyset$ (cf. Section 3.2).

Example 5.8. Consider the system of contracts $S_{sb_1b_2}$ from Example 5.1 and, in particular, its stipulated contracts, with substitution $\pi = \{s/s', b_1/b_1, b'_1, b_2/b_2, b'_2\}$ from Example 5.3:

$$\hat{P}_s = P_s\pi = b_1?req; b_2?req; b_1!quote; (b_1?order; b_2!ok + b_1?bye; b_2!bye)$$

$$\hat{P}_{b_1} = P_{b_1}\pi = s!req; s?quote; (b_2!ok; s!order \oplus b_2!bye; s!bye)$$

$$\hat{P}_{b_2} = P_{b_2}\pi = s!req; (b_1?ok; s?ok + b_1?bye; s?bye)$$

We have $\text{CrS}(\hat{P}_s) = \{(b_1, req)\}$: in other words, at this point of the contract, an interaction is expected between s and b_1 (since s is waiting for req), while no interaction is expected between s and b_2 .

When \hat{P}_s reaches its external choice, we have:

$$\hat{P}'_s = b_1?order; b_2!ok + b_1?bye; b_2!bye$$

Now, the ready set becomes $\text{CrS}(\hat{P}'_s) = \{(b_1, order), (b_1, bye)\}$, i.e., s must handle both answers from b_1 . Instead, when \hat{P}_{b_1} reduces to its internal choice, we have:

$$\hat{P}'_{b_1} = b_2!ok; s!order \oplus b_2!bye; s!bye$$

Thus, its ready set becomes $\text{CrS}(\hat{P}'_{b_1}) = \{(b_2, ok), (b_2, bye)\}$: b_1 is free to choose either branch. \diamond

Example 5.8 shows that when a *contract* P of a participant evolves within a system S its ready set changes. Now we define the counterpart of contract ready set for CO₂ processes, i.e., the *process ready set*. Again, we adapt the definition from [4] to our multiparty contract model.

Definition 5.5 (Process ready set). *For all networks T , all participants s, r and sessions*

u , we define the set of pairs:

$$\text{Prs}_s^u(T) = \left\{ (r, a) \mid \begin{array}{l} \exists \vec{v}, \vec{n}, R, R', R'', T' : \\ T \equiv (\vec{v}, \vec{n}) (s \llbracket \text{do}_r^u a . R + R' \mid R'' \rrbracket \mid T_0) \mid T_1 \wedge u \notin \vec{v} \end{array} \right\}$$

Intuitively, Definition 5.5 says that the process ready set of s over a session u in a network T contains the interactions that s is immediately able to perform with other participants through its do-actions on u . As in a contract ready set, the interactions are represented by participant/sort pairs.

Next, we want to characterise a weaker notion of the process ready set, so it only takes into account the first actions *on a specific session* that a participant is ready to make.

Definition 5.6 (Weak Process ready set). We write $T \xrightarrow{\neq(s: \text{do}^u)} T'$ iff:

$$\exists s', p : T \xrightarrow{s': p} T' \implies (s \neq s' \vee \forall a : \forall r : p = \text{do}_r^v a \implies u \neq v)$$

We define the set of pairs $\text{WPrs}_s^u(T)$ as:

$$\text{WPrs}_s^u(T) = \left\{ (r, a) \mid \exists T' : T \xrightarrow{\neq(s: \text{do}^u)}^* T' \wedge (r, a) \in \text{Prs}_s^u(T') \right\}$$

In Definition 5.6, we are not interested in the actions that do not relate to the session u . Thus, we allow the network to evolve either by (i) letting any other participant other than s do an action, or (ii) letting s act on a different session than u , or (iii) do internal actions.

We now introduce the final ingredient for honesty, that is the notion of *readiness* of a participant.

Definition 5.7 (Readiness). We say that s is ready in T iff, whenever $T \equiv (\vec{u}, \vec{n}) T_1 \mid T_2$ for

some \vec{u}, \vec{n} and $T_1 = k[[s[P] \mid \dots]] \mid T_0$, the following holds:

$$\exists X \in \text{Crs}(P) : \left((r, a) \in X \implies (r, a) \in \text{WPrs}_s^k(T_1) \right)$$

Definition 5.7 says that a participant s is *ready* in a network T whenever its process ready set for a session k will eventually contain all the participant/sort pairs of a set in the contract ready set of s 's contract in k . When a participant s is “ready”, then, for any of its contracts P , the CO_2 process of s is (eventually) able to fulfil at least the interactions in P 's prefix.

Remark 5.4. The side condition “ $u \notin \vec{v}$ ” of Definition 5.5 deals with cases like

$$T_0 = (k) \left(s \left[\left[\text{do}_r^k \text{int} \right] \right] \right) \quad \text{and} \quad T = T_0 \mid k[[s[r!\text{int}] \mid \dots]] \mid \dots$$

without the side condition, $\text{Prs}_k^u(T_0) = \{ \{ (r, \text{int}) \} \}$; thus, by Definition 5.7, s would result to be ready in T .

Example 5.9. We have seen that, after fusion of the latent contracts of $T_{sb_1b_2}$ (in Example 5.3) we obtain:

$$(k)T'_{sb_1b_2} = (k) \left(s[[R_s \sigma \pi]] \mid k[[S_{sb_1b_2} \mid Q(S_{sb_1b_2})]] \mid b_1[[R_{b_1} \sigma \pi]] \mid b_2[[R_{b_2} \sigma \pi]] \right)$$

Let us define the processes (after substitutions):

$$R_s \sigma \pi = \text{do}_{b_1}^k \text{req} . \text{do}_{b_2}^k \text{req} . \text{do}_{b_1}^k \text{quote} . (\text{do}_{b_1}^k \text{order} . \text{do}_{b_2}^k \text{ok} + \text{do}_{b_1}^k \text{bye} . \text{do}_{b_2}^k \text{bye})$$

$$R_{b_1} \sigma \pi = \tau . \text{do}_s^k \text{req} . \text{do}_s^k \text{quote} . \text{do}_s^k \text{order}$$

$$R_{b_2} \sigma \pi = \text{do}_s^k \text{req} . (\text{do}_{b_1}^k \text{ok} . \text{do}_s^k \text{ok} + \text{do}_{b_1}^k \text{bye} . \text{do}_s^k \text{bye})$$

Thus, we have:

$$\begin{aligned}
 \text{Prs}_s^k(T'_{sb_1b_2}) &= \{(b_1, \text{req})\} = \text{WPrs}_s^k(T'_{sb_1b_2}) \\
 \text{Prs}_{b_1}^k(T'_{sb_1b_2}) &= \emptyset \neq \{(s, \text{req})\} = \text{WPrs}_{b_1}^k(T'_{sb_1b_2}) \\
 \text{Prs}_{b_2}^k(T'_{sb_1b_2}) &= \{(s, \text{req})\} = \text{WPrs}_{b_2}^k(T'_{sb_1b_2})
 \end{aligned}$$

Note that the τ prefix in P_{b_1} prevents b_1 from interacting immediately with s on session k , although it is “weakly ready” to do so. Hence, considering that each weak process ready set of each participant in $T'_{sb_1b_2}$ matches their respective contract ready set in $S_{sb_1b_2}$ (Example 5.8) according to Definition 5.7 we have that participants s , b_1 and b_2 are all ready in $(k)T'_{sb_1b_2}$. \diamond

Definition 5.8 (Honesty). *We say that $s[R]$ is honest iff, for all T such that neither latent/stipulated contracts of s nor $s[\dots]$ occur in T , and for all T' such that $s[R] \mid T \rightarrow^* T'$, s is ready in T' .*

A process $s[R]$ is said to be honest when, for all contexts and reductions that $s[R]$ may be engaged in, s is persistently ready. In other words, there is a continuous correspondence between the interactions exposed in the contract ready sets and the process ready sets of the possible reductions of any network involving $s[R]$. The definition rules out contexts with latent/stipulated contracts of s , otherwise s could be made trivially dishonest, e.g., by inserting a latent contract $\downarrow_u s \text{ says } P$ that s cannot fulfil.

Similarly to [4], the definition of honesty assumes a *fair* scheduler, eventually allowing participants to fire persistently (weakly) enabled do actions.

Example 5.10. Consider the process $b_1[\text{tell}_s \downarrow_y P_{b_1} \cdot R_{b_1}]$ of network $T_{sb_1b_2}$, as defined in Examples 5.3 and 5.9. We show that this process is *not* honest. In fact, $T_{sb_1b_2}$ reduces as follows:

$$T_{sb_1b_2} \rightarrow^* (k)T'_{sb_1b_2} \rightarrow^* (k)T''_{sb_1b_2}$$

where:

$$\begin{aligned}
 (k)T''_{sb_1b_2} = & (k) \left(s \left[\text{do}_{b_1}^k \text{ order} . \text{do}_{b_2}^k \text{ ok} + \text{do}_{b_1}^k \text{ bye} . \text{do}_{b_2}^k \text{ bye} \right] \right. \\
 & | k \left[s[b_1?order; b_2!ok + b_1?bye; b_2!bye] \right. \\
 & \quad | b_1[b_2!ok; s!order \oplus b_2!bye; s!bye] \\
 & \quad | b_2[b_1?ok; s?ok + b_1?bye; s?bye] \\
 & \quad | sb_1 : \varepsilon | b_1s : \varepsilon | sb_2 : \varepsilon | b_2s : \varepsilon | b_1b_2 : \varepsilon | b_2b_1 : \varepsilon \left. \right] \\
 & \left. | b_1 \left[\text{do}_s^k \text{ order} \right] \quad | \quad b_2 \left[\text{do}_{b_1}^k \text{ ok} . \text{do}_s^k \text{ ok} + \text{do}_{b_1}^k \text{ bye} . \text{do}_s^k \text{ bye} \right] \right)
 \end{aligned}$$

At this point, we notice a problem in the implementation of b_1 : it does not notify the other buyer before making an order.

In fact, b_1 's process is trying to perform $\text{do}_s^k \text{ order}$, but its contract requires that $\text{do}_{b_2}^k \text{ ok}$ is performed first (or $\text{do}_{b_2}^k \text{ bye}$, if the quote is rejected). This is reflected by the mismatch between b_1 's process ready set in $T''_{sb_1b_2}$ and its contract ready sets, in session k :

$$\begin{aligned}
 \text{Prs}_{b_1}^k(T''_{sb_1b_2}) &= \{(s, \text{order})\} \\
 \text{Crs}(b_2!ok; s!order \oplus b_2!bye; s!bye) &= \{(b_2, \text{ok})\}, \{(b_2, \text{bye})\}
 \end{aligned}$$

Using the vocabulary of Definitions 5.3, 5.7, and 5.8, we have:

- There exists a network $T_{sb_1b_2}$ which contains $b_1 \llbracket \text{tell}_s \downarrow_y P_{b_1} . R_{b_1} \rrbracket$
- and reduces to a network $(k)T''_{sb_1b_2}$, where b_1 is *not ready*.
- Thus, we have that $b_1 \llbracket \text{tell}_s \downarrow_y P_{b_1} . R_{b_1} \rrbracket$ is not honest, and
- b_1 is *culpable* in $(k)T''_{sb_1b_2}$.

◇

Remark 5.5. *In this section, we consider only fully instantiated contracts, i.e., contracts that do not use participant variables. This is due to two reasons. (i) Honesty is defined on*

all the possible executions of any networks, thus it encompasses any possible instantiation of participant variables. (ii) The results of Section 5.6 apply only whenever sessions have been started, which implies that participant variables have substituted by participant names.

Remark 5.6. *Honesty is not decidable in general [7], but for a bilateral contract model it has been approximated via an abstract semantics [7] and a type discipline [4] for CO_2 . We conjecture that these approximations may be adapted to our setting.*

5.6 Properties of Honest Networks

We give the properties that our framework guarantees. We ensure that two basic features of CO_2 hold in our multiparty adaptation: the state of a session always permits to establish who is responsible for making the network progress (Theorem 5.1) and honest participants can always exculpate themselves (Lemma 5.1). We then formalise a link between the honesty of participants and two key properties borrowed from the session types world: Theorem 5.2 introduces session fidelity in CO_2 ; and Theorem 5.3 introduces a notion of progress in CO_2 , based on the progress of the contractual agreement.

First, we define the kind of networks on which our results apply, i.e., *honest networks*. These networks consists only of honest participants and do not feature any runtime constructs (such as sessions or stipulated contracts).

Definition 5.9 (Initial & honest networks). *A network T is initial if*

- *for each participant s in T , there is at most one sub-network $s[[Z]]$ in T ,*
- *for each sub-network $s[[Z]]$ in T , there is no (latent or stipulated) contract in Z that does not belong to s , i.e., Z has no sub-term of the form $\text{tell}_{\underline{n}} \downarrow_u P$ or $\downarrow_u \underline{n} \text{says} P$, with $\underline{n} \neq s$, and*

- no session has been started in T , i.e., there is no sub-network of the form $k\llbracket S \rrbracket$ in T .

An honest network T is an initial network such that each process $s\llbracket R \rrbracket$ in T is honest.

Theorem 5.1 below says that, in an active session, there is always at least one participant $s\llbracket R \rrbracket$ who is responsible for the next interaction. Thus, if a corresponding $\text{do}_r^k a$ prefix is not in R , T may get stuck, and s is culpable.

Theorem 5.1 (Culpability). *Given a network T , if T contains a session $k\llbracket S \rrbracket$ such that $\exists n \in \mathcal{P}(S) : S(n) \neq \mathbf{0}$, then there exists at least one culpable participant.*

Proof. Since session k is not terminated, we must have:

$$S \equiv n[r!a; P \oplus P] \mid S' \quad \text{or} \quad S \equiv n[s?a; P + P] \mid sn : a \cdot \rho \mid S'$$

Note that, by Theorem 3.1, S cannot be a deadlock or reduce to an orphan message configuration.

Since k has been started (cf. Remark 5.2) and n is part of the session, we must have T of the form

$$T \equiv (\vec{u}, \vec{n})(n\llbracket R \rrbracket \mid k\llbracket S \rrbracket \mid T_1) \mid T_2$$

with $k \in \vec{u}$, \vec{n} possibly empty, and possibly $R \equiv \mathbf{0}$, note that we have $n\llbracket \mathbf{0} \rrbracket \equiv \mathbf{0}$ from the congruence rules of CO_2 . Thus, we have $S \xrightarrow{n \rightarrow r : a}$ or $S \xrightarrow{n \leftarrow s : a}$ as required by Definition 5.3. \square

Example 5.11. Consider network $T''_{sb_1b_2}$ of Example 5.10 and the system of contracts in its session k :

$$\begin{aligned} S_k &= s[b_1?order; b_2!ok + b_1?bye; b_2!bye] \\ &\quad \mid b_1[b_2!ok; s!order \oplus b_2!bye; s!bye] \mid b_2[b_1?ok; s?ok + b_1?bye; s?bye] \\ &\quad \mid sb_1 : \varepsilon \mid b_1s : \varepsilon \mid sb_2 : \varepsilon \mid b_2s : \varepsilon \mid b_1b_2 : \varepsilon \mid b_2b_1 : \varepsilon \end{aligned}$$

We have $S_k \xrightarrow{b_1 \Leftarrow b_2: \text{ok}}$ and $S_k \xrightarrow{b_1 \Leftarrow b_2: \text{bye}}$. Hence, b_1 is responsible for the next interaction, and culpable for $T''_{sb_1b_2}$ being stuck. \diamond

Lemma 5.1 follows from the definition of honesty. It states that honest participants can always overcome their culpability, by firing their contractual do actions (possibly after some internal actions).

Lemma 5.1 (Exculpation). *Given an initial network T_0 with an honest participant $s \llbracket R \rrbracket$, whenever $T_0 \rightarrow^* T \equiv (\vec{u}, \vec{n}) (k \llbracket S \rrbracket \mid T_1) \mid T_2$ and s is culpable in T , there exist r and a such that:*

$$T \xrightarrow{s: \tau}^* \xrightarrow{s: \text{do}_r^k a}$$

Proof. By Definition 5.8, we have that for any reduction of T_0 , s is *ready*. By assumption, s is culpable in T , so we must have $S \xrightarrow{s \Leftarrow r: a}$, and, by Definition 5.7, we must have $(r, a) \in \text{WPrs}_s^k(T_1)$. By Definition 5.6, we have

$$\exists T': T_1 \xrightarrow{\neq (s: \text{do}_r^k)}^* T' \wedge (r, a) \in \text{Prs}_s^k(T')$$

By Definition 5.5, we must have that $s \llbracket \text{do}_r^k a . R \rrbracket$ is a sub-network of T' .³ Thus, we just have to show that s may reduce to such a prefixed process via internal actions (τ) *only*. By contradiction, if it was the case that, in the process belonging to s , there was a blocking action (before $\text{do}_r^k a$), then one could easily find a network in which s is not ready (i.e., s would be waiting for this action to be dealt with by another participant); thus contradicting the assumption that s is honest. \square

Theorem 5.2 below says that each (honest) participant will strictly adhere to its contracts, once they have been fused in a session. It follows directly from the semantics of CO₂ (that forbid non-contractual do prefixes to be fired) and from the definition of honesty.

³We abstract from possible choice or variable restrictions, without loss of generality.

Theorem 5.2 (Fidelity). *Let T be an honest network, if*

$$T \rightarrow^* T' \equiv (\vec{u}, \vec{n}) (s[R] \mid k[S] \mid T_0) \mid T_1$$

then

$$T' \xrightarrow{-k}^* \xrightarrow{s: \text{do}_r^k a} \iff S \xrightarrow{s \Leftarrow r: a}$$

(where $\xrightarrow{-k}^*$ is any reduction not involving session k).

Proof. The direction (\implies) follows directly from the semantics of CO_2 (i.e., rule $[\text{DO}]$). The direction (\impliedby) follows from Lemma 5.1. Indeed, $S \xrightarrow{s \Leftarrow r: a}$ implies that s is culpable in session k , and by Lemma 5.1, s is able to exculpate itself after some internal actions (i.e., no actions on session k). \square

Theorem 5.3 below introduces a notion of global progress in CO_2 networks. Observe that progress in CO_2 is only meaningful *after* a session has been established and thus a culpable participant exists. A network without sessions may not progress because a set of compliant contracts cannot be found, or a fuse prefix is not enabled. In both cases, no participant may be deemed culpable, and thus responsible for the next move. However, the network may progress again if other (honest) participants join it, allowing a session to be established. This is analogous to the notion of progress of [35], cf. Section 2.2.3.

Theorem 5.3 (Global progress). *Given an honest network T_0 , whenever*

$$T_0 \rightarrow^* T \equiv (\vec{u}, \vec{n}) (k[S] \mid T_1) \mid T_2$$

and $\exists n \in \mathcal{P}(S) : S(n) \neq \mathbf{0}$, then $T \rightarrow$.

Proof. By contradiction. By Theorem 3.1, we know that S is deadlock free, thus if the network cannot make further reductions, it must be because one of the participant n cannot

meet its obligations. This is a contradiction with Lemma 5.1 since all participants are honest in T_0 . \square

Note that Theorem 5.3 holds for networks where a process takes part in multiple sessions: the honesty of all participants guarantees that all sessions will be completed. We illustrate such a situation in Example 5.12 below.

Example 5.12. We show how a seemingly honest process (r) could be deemed culpable due to the unexpected behaviour of other participants. Consider the following network:

$$\begin{aligned}
 T_{\text{ex5.12}} = (x, y, z, w) (& \text{ s } \llbracket \text{tell}_s \downarrow_x (r! \text{int}) . \text{fuse} . \text{fuse} \rrbracket \\
 & | \text{ r } \llbracket \text{tell}_s \downarrow_y (s? \text{int}) . \text{tell}_s \downarrow_z (n! \text{bool}) . \text{do}_s^y \text{int} . \text{do}_n^z \text{bool} \rrbracket \\
 & | \text{ n } \llbracket \text{tell}_s \downarrow_w (r? \text{bool}) . \text{do}_r^w \text{bool} \rrbracket)
 \end{aligned}$$

where participant s advertises a contract to itself (promising to sent a message int to r); participant r advertises two contracts to s , the first contract specifies an interaction with s , while the second specifies an interaction with n ; and participant n advertises a contract (to s) specifying an interaction with r .

After all four contracts have been advertised to s and fused, the network reduces to:

$$\begin{aligned}
 T'_{\text{ex5.12}} = (k_1, k_2) (& \text{ s } \llbracket \mathbf{0} \rrbracket \quad | \quad \text{ r } \llbracket \text{do}_s^{k_1} \text{int} . \text{do}_n^{k_2} \text{bool} \rrbracket \quad | \quad \text{ n } \llbracket \text{do}_r^{k_2} \text{bool} \rrbracket \\
 & | \quad k_1 \llbracket \text{s}[r? \text{int}] \quad | \quad \text{r}[s! \text{int}] \quad | \quad \text{sr} : \varepsilon \quad | \quad \text{rs} : \varepsilon \rrbracket \\
 & | \quad k_2 \llbracket \text{r}[n! \text{bool}] \quad | \quad \text{n}[r? \text{bool}] \quad | \quad \text{rn} : \varepsilon \quad | \quad \text{nr} : \varepsilon \rrbracket)
 \end{aligned}$$

Even if the systems of contracts in both sessions k_1 and k_2 are able to reduce further, $T'_{\text{ex5.12}}$ is stuck. Indeed, s does not perform the promised action, it is culpable in k_1 ; and r is stuck waiting in k_1 , thus remaining culpable⁴ in k_2 . Indeed, neither s nor r are *ready* in $T'_{\text{ex5.12}}$, and thus they are not honest in $T_{\text{ex5.12}}$. Hence, the network cannot progress

⁴In this case, r is deemed culpable in k_2 because its implementation did not expect s to misbehave.

further. ◇

Example 5.13. Let us now consider a variant of $T_{\text{ex5.12}}$ from Example 5.12, where all participants are honest:

$$\begin{aligned}
 (x, y, z, w) & \left(s \left[\left[(\text{tell}_s \downarrow_x (r! \text{int}) . \text{do}_r^x \text{int}) \mid \text{fuse} \mid \text{fuse} \right] \right. \right. \\
 & \quad \left. \left. \mid n \left[\left[\text{tell}_s \downarrow_w (r? \text{bool}) . \text{do}_r^w \text{bool} \right] \right] \right. \right. \\
 & \quad \left. \left. \mid r \left[\left[\text{tell}_s \downarrow_y (s? \text{int}) . \text{tell}_s \downarrow_z (n! \text{bool}) . (\text{do}_s^y \text{int} . \text{do}_n^z \text{bool} + \tau . (\text{do}_s^y \text{int} \mid \text{do}_n^z \text{bool})) \right] \right] \right] \right)
 \end{aligned}$$

In this case, s will respect its contractual duties, while r will be ready to fulfil its contracts on both sessions – even if one is not activated, or remains stuck (here, τ represents an internal action, e.g., a timeout: if the first $\text{do}_s^y \text{int}$ cannot reduce, r falls back to running the sessions in parallel). The honesty of all participants in the network above guarantees that, once a session is active, it will reach its completion. ◇

5.7 Concluding Remarks

In this chapter, we investigated the combination of the contract-oriented calculus CO_2 with a contract model that fulfils two basic design requirements: (i) it supports multiparty agreements, and (ii) it provides an explicit description of the choreography that embodies each agreement. To the best of our knowledge, no other contract model provides an explicit choreography synthesis.

We built our framework upon a simple version of session types and yet it turns out to be quite flexible, e.g., allowing for sessions where the number of participants is not known beforehand and permitting fine-grained characterisation of the contractual agreement.

We are considering two main directions of future research based on this work. The first direction concerns an extension of the fuse primitive so to allow that, whenever several sets of contracts are compliant, the “best” set of contracts is fused (according to

different criteria such as, e.g., the number of participants or a semantic characterisation of the choreography). The extensions we discussed in Section 5.4.2 relied on a non-deterministic selection of the set of contacts to be fused (via the \equiv congruence relation). One may integrate the selection of contracts within the agreement relation so that the best contracts are chosen and the ones which are unused stay in the network, to be fused later on.

Another research direction concerns the possibility of allowing a participant to be involved in a same session with multiple contracts – e.g., a bank advertising two services, and a customer publishing a contract which uses both contracts of the bank in a same session. This is quite challenging as it may require to somehow merge these contracts into one, before attempting to synthesise a global type. Indeed, one may not be able to consider that the contracts are simply two different participants in the choreography, e.g., a process advertising several contracts may not fulfil them in an independent manner.

Conclusions and Future Directions

6.1 Summary of the Contributions

We review the main contributions presented in Chapters 3, 4, and 5.

On synthesising choreographies. In Chapter 3, we introduced a theory whereby it is possible to infer a choreography from a set of local behavioural specifications. The main construction in this chapter is a type system which assigns a global type to a set of local types. We showed that, if it exists, such a global type is unique and well-formed; also, its projections are equivalent to the original local types. The type system enjoys a subject reduction property. Finally, we showed that for every well-formed global type, an equivalent global type can be assigned to its projections.

A major advantage of our bottom-up approach is that it allows the result of the top-down approach of [44] to be applicable even if no global view of the system is available. This allowed us to reuse well-established results to show that, if a system is typable, then it has progress and safety properties. As a consequence, our theory allows to lift the assumption that a global view of the system is available to type-check programs with

session types.

On amending global assertions. In Chapter 4, we gave a few techniques to help distributed software architects design global assertions. These techniques include two algorithms to solve history sensitivity problems, one algorithm to solve temporal satisfiability problems, and a methodology for applying the algorithms to protocol design.

We showed that our algorithms satisfy the following properties: *(i)* structure preservation, i.e., they do not modify the structure of the underlying global type, *(ii)* properties preservation, i.e., they do not introduce new violations, and *(iii)* correctness, i.e., if applicable they correct all the problems.

As the theory in Chapter 3 allows to lift the burden of (re-) designing global types, the methodology presented in Chapter 4 permits, to some extent, to lift the burden of making a global assertion well-asserted.

Runtime choreography synthesis. In Chapter 5, we introduced a formal model for distributed systems where participants advertise contracts and multiparty sessions are started when a set of contracts allows to synthesise a choreography.

Our model is based on an adaptation of the CO₂ calculus [6] and preserves its main properties. In particular, it is always possible to identify which participant is responsible for the session being stuck. Inspired by other results from session type theory, we also showed that progress and session fidelity properties hold for *honest* systems, i.e., systems consisting of honest participants only. In addition, we gave a few examples illustrating the utility of having a choreography as a tangible representation of an agreement between participants.

A remarkable feature of our model is that it adds flexibility to the synthesis of choreographies of Chapter 3 by allowing *participant variables* to appear in contracts.

6.2 Future Directions

We identify a few directions of future research based on the results established in this thesis. We focus on the synthesis of choreographies as the other topics where covered in the concluding remarks of their respective chapters.

From communicating machines to choreographies. A main direction of future work, which has been confirmed in preliminary results by the author and colleagues (and to some extent in [41]), is to study the synthesis of choreographies from communicating machines. Remarkably, such a synthesis allows to extend substantially the set of systems to which it is possible to assign a choreography. Indeed, contrarily to the type system of Chapter 3, there are fewer syntactic restrictions: this is most noticeable for recursive types. In fact, using the theory of regions [3, 37], it is possible to synthesise a choreography – or a generalised global types [40] (cf. Section 2.1.4) – from the *synchronous* transition system of a system of communicating machines. The synchronous transition system of a system of communicating machines amounts to, essentially, letting the machines run with the additional constraint that at most one buffer is non-empty, that the non-empty buffer contains at most one symbol, and that an output action is directly followed by its corresponding input action.

Since communicating machines are undecidable in general [16], one needs to restrict the set of “typable” systems. In fact, an important step is to check whether all the executions of the machines are “captured” by the synchronous transition system.

In addition, we conjecture that the results of Chapters 4 and 5 may be quite easily extended to generalised global types. On the one hand, the framework of Chapter 5 is in fact rather independent of the way choreographies are actually synthesised. It uses the type system of Chapter 3 as a “black box”. On the other hand, in order to adapt the results of Chapter 4, one would need to re-visit the notion of well-assertedness for general graphs and similarly for the three algorithms. In fact, the main difference is that there might be

several paths to a same node. Thus, our algorithms would have to update all the paths leading to a problematic node.

Beyond 1-buffer executions. One of the main restriction of our type system and, to the best of our knowledge, all the current works on the synthesis of choreographies [41], is that they do not capture systems that have executions which are not captured by a synchronous execution. Analogously, the conditions for the *realisability* of choreography (cf. Section 2.2.1 and [9, 10]) also requires a system to be *synchronisable*. See Section 3.6 for a discussion on the 1-buffer restriction in our type system.

Example 6.1 below shows that some systems cannot be assigned a choreography in any of the works we discussed in Chapter 2.

Example 6.1. Consider the following system of two local types:

$$s[r!a;r?b] \mid r[s!b;s?a] \quad (6.1)$$

it is easy to see that this system does not deadlock, however it is not typable by the type system of Chapter 3.¹ Essentially, this is due to the fact that the system “deadlocks” when each intermediary configuration may contain at most one non-empty buffer. Indeed, with such a constraint the system reduces, e.g, to

$$s[r?b] \mid r[s!b;s?a] \mid sr : a \mid rs : \varepsilon$$

where r cannot send b since one buffer is already non empty and a cannot be received since r must send b first.

In fact, the system (6.1) cannot be captured by any of the formalism we have seen so

¹ This system is also not *multiparty-compatible* in the sense of [41].

far. For instance, neither of these global types capture it:

$$s \rightarrow r:a; r \rightarrow s:b \quad \times \qquad r \rightarrow s:b; s \rightarrow r:a \quad \times$$

Even if we allow a same participant to appear in different concurrent branches and introduce parallel constructs at the local type level, projecting the global type

$$s \rightarrow r:a \mid r \rightarrow s:b$$

would yield the system:

$$s[r!a \mid r?b] \mid r[s!b \mid s?a] \tag{6.2}$$

which is not bisimilar to system (6.1). Observe that the system (6.2) is not directly representable in terms of communicating machines since CFSMs do not feature (explicit) concurrent activities. \diamond

Example 6.1 shows that taking asynchrony to its full potential when synthesising a choreography is not easy. In fact, choreography languages consider interactions of the form $s \rightarrow r:a$ to be *atomic*; and they generally offer three main constructs: sequence, parallel, and choice. None of these constructs allow to model system (6.1) from a global perspective. In fact, before attempting to tackle the problem of synthesising a choreography from systems that do not have a synchronous execution, one would need to consider a choreography language and a notion of projection that do cater for (explicit) asynchrony. For instance, one may consider a notion of projection that gives precedence to sending actions in concurrent branches so that the projection of

$$s \rightarrow r:a \mid r \rightarrow s:b$$

yields

$$s[r!a;r?b] \mid r[s!b;s?a]$$

Other directions. Other directions of research concerns a more precise comparison between our type system and other related approaches, such as the conversation types [23] and the work of Castagna et al. [31], so to characterise the systems that are supported by each theory.

Another interesting direction is to work on a methodology that combines our bottom-up approach with top-down approaches such as “global programming” [26]. The objective being that both directions, from global to local specifications, and vice versa, should be available to practitioners. Combining both approaches would not only help to keep specifications and implementations as close as possible to each other, but also facilitate the adoption of session type theory in more ad-hoc development cycles. Similarly to what is done in [1], our bottom-up approach may also be valuable to infer “bad scenarios” so that practitioners may more easily understand why an implementation does not match a given choreography.

Finally, we are currently working on a tool that constructs a choreography (in a syntax similar to the one of generalised global types) from a set of communicating machines, based on the preliminary results mentioned above. We are also considering implementing both the type system of Chapter 3 and the algorithms of Chapter 4, so to integrate them in tool described in [52].

Bibliography

- [1] Rajeev Alur, Kousha Etessami, and Mihalis Yannakakis. Inference of message sequence charts. *IEEE Trans. Software Eng.*, 29(7):623–633, 2003.
- [2] Krzysztof R. Apt, Nissim Francez, and Shmuel Katz. Appraising fairness in languages for distributed programming. *Distributed Computing*, 2(4):226–241, 1988.
- [3] Eric Badouel and Philippe Darondeau. Theory of regions. In Wolfgang Reisig and Grzegorz Rozenberg, editors, *Petri Nets*, volume 1491 of *Lecture Notes in Computer Science*, pages 529–586. Springer, 1996.
- [4] Massimo Bartoletti, Alceste Scalas, Emilio Tuosto, and Roberto Zunino. Honesty by typing. In Dirk Beyer and Michele Boreale, editors, *FMOODS/FORTE*, volume 7892 of *Lecture Notes in Computer Science*, pages 305–320. Springer, 2013.
- [5] Massimo Bartoletti, Emilio Tuosto, and Roberto Zunino. Contracts in distributed systems. In Silva et al. [66], pages 130–147.
- [6] Massimo Bartoletti, Emilio Tuosto, and Roberto Zunino. Contract-oriented computing in CO₂. *Scientific Annals in Comp. Sci.*, 22(1):5–60, 2012.
- [7] Massimo Bartoletti, Emilio Tuosto, and Roberto Zunino. On the realizability of

- contracts in dishonest systems. In Marjan Sirjani, editor, *COORDINATION*, volume 7274 of *Lecture Notes in Computer Science*. Springer, 2012.
- [8] Massimo Bartoletti and Roberto Zunino. A calculus of contracting processes. In *LICS*. IEEE Computer Society, 2010.
- [9] Samik Basu, Tevfik Bultan, and Meriem Ouederni. Deciding choreography realizability. In John Field and Michael Hicks, editors, *POPL*, pages 191–202. ACM, 2012.
- [10] Samik Basu, Tevfik Bultan, and Meriem Ouederni. Synchronizability for verification of asynchronously communicating systems. In Viktor Kuncak and Andrey Rybalchenko, editors, *VMCAI*, volume 7148 of *LNCS*. Springer, 2012.
- [11] Lorenzo Bettini, Mario Coppo, Loris D’Antoni, Marco De Luca, Mariangiola Dezani-Ciancaglini, and Nobuko Yoshida. Global progress in dynamically interleaved multiparty sessions. In Franck van Breugel and Marsha Chechik, editors, *CONCUR*, volume 5201 of *Lecture Notes in Computer Science*, pages 418–433. Springer, 2008.
- [12] Laura Bocchi, Kohei Honda, Emilio Tuosto, and Nobuko Yoshida. A theory of design-by-contract for distributed multiparty interactions. In Gastin and Laroussinie [42], pages 162–176.
- [13] Laura Bocchi, Julien Lange, and Emilio Tuosto. Amending contracts for choreographies. In Silva et al. [66], pages 111–129.
- [14] Laura Bocchi, Julien Lange, and Emilio Tuosto. Three algorithms and a methodology for amending contracts for choreographies. *Scientific Annals of Computer Science*, 22(1):61–104, 2012.

- [15] Michele Boreale and Rocco De Nicola. Testing equivalence for mobile processes. *Inf. Comput.*, 120(2):279–303, 1995.
- [16] Daniel Brand and Pitro Zafiropulo. On communicating finite-state machines. *JACM*, 30(2):323–342, 1983.
- [17] Mario Bravetti, Ivan Lanese, and Gianluigi Zavattaro. Contract-driven implementation of choreographies. In Christos Kaklamanis and Flemming Nielson, editors, *TGC*, volume 5474 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 2008.
- [18] Maria Grazia Buscemi and Ugo Montanari. Cc-pi: A constraint-based language for specifying service level agreements. In Rocco De Nicola, editor, *ESOP*, volume 4421 of *Lecture Notes in Computer Science*, pages 18–32. Springer, 2007.
- [19] Luís Caires and Frank Pfenning. Session types as intuitionistic linear propositions. In Gastin and Laroussinie [42], pages 222–236.
- [20] Luís Caires, Frank Pfenning, and Bernardo Toninho. Towards concurrent type theory. In Benjamin C. Pierce, editor, *TLDI*, pages 1–12. ACM, 2012.
- [21] Luís Caires and Hugo Torres Vieira. Conversation types. In Castagna [27], pages 285–300.
- [22] Luís Caires and Hugo Torres Vieira. Analysis of service oriented software systems with the conversation calculus. In Luís Soares Barbosa and Markus Lumpe, editors, *FACS*, volume 6921 of *Lecture Notes in Computer Science*, pages 6–33. Springer, 2010.
- [23] Luís Caires and Hugo Torres Vieira. Conversation types. *Theor. Comput. Sci.*, 411(51-52):4399–4440, 2010.

- [24] Marco Carbone, Kohei Honda, and Nobuko Yoshida. Structured communication-centred programming for web services. In Rocco De Nicola, editor, *ESOP*, volume 4421 of *Lecture Notes in Computer Science*, pages 2–17. Springer, 2007.
- [25] Marco Carbone, Kohei Honda, and Nobuko Yoshida. Structured communication-centered programming for web services. *ACM Trans. Program. Lang. Syst.*, 34(2):8, 2012.
- [26] Marco Carbone and Fabrizio Montesi. Deadlock-freedom-by-design: multiparty asynchronous global programming. In Roberto Giacobazzi and Radhia Cousot, editors, *POPL*, pages 263–274. ACM, 2013.
- [27] Giuseppe Castagna, editor. *Programming Languages and Systems, 18th European Symposium on Programming, ESOP 2009*, volume 5502 of *Lecture Notes in Computer Science*. Springer, 2009.
- [28] Giuseppe Castagna, Mariangiola Dezani-Ciancaglini, and Luca Padovani. On global types and multi-party sessions. In Roberto Bruni and Jürgen Dingel, editors, *FMOODS/FORTE*, volume 6722 of *Lecture Notes in Computer Science*, pages 1–28. Springer, 2011.
- [29] Giuseppe Castagna, Mariangiola Dezani-Ciancaglini, and Luca Padovani. On global types and multi-party session. *Logical Methods in Computer Science*, 8(1), 2012.
- [30] Giuseppe Castagna, Nils Gesbert, and Luca Padovani. A theory of contracts for web services. *ACM Trans. on Prog. Lang. and Sys.*, 31(5), 2009.
- [31] Giuseppe Castagna and Luca Padovani. Contracts for mobile processes. In Mario Bravetti and Gianluigi Zavattaro, editors, *CONCUR*, volume 5710 of *Lecture Notes in Computer Science*, pages 211–228. Springer, 2009.

- [32] Tzu-Chun Chen, Laura Bocchi, Pierre-Malo Deniérou, Kohei Honda, and Nobuko Yoshida. Asynchronous distributed monitoring for multiparty session enforcement. In Roberto Bruni and Vladimiro Sassone, editors, *TGC*, volume 7173 of *Lecture Notes in Computer Science*, pages 25–45. Springer, 2011.
- [33] Tzu-Chun Chen and Kohei Honda. Specifying stateful asynchronous properties for distributed programs. In Koutny and Ulidowski [49], pages 209–224.
- [34] World Wide Web Consortium. Web services choreography description language version 1.0. <http://www.w3.org/TR/ws-cdl-10/>, 11 2005.
- [35] Mario Coppo, Mariangiola Dezani-Ciancaglini, Luca Padovani, and Nobuko Yoshida. Inference of global progress properties for dynamically interleaved multiparty sessions. In Rocco De Nicola and Christine Julien, editors, *COORDINATION*, volume 7890 of *Lecture Notes in Computer Science*, pages 45–59. Springer, 2013.
- [36] Mario Coppo, Mariangiola Dezani-Ciancaglini, Nobuko Yoshida, and Luca Padovani. Global progress for dynamically interleaved multiparty sessions. Technical report, Dipartimento di Informatica, Università di Torino, 2013. Available at <http://www.di.unito.it/~padovani/Papers/CoppoDezaniYoshidaPadovani13.pdf>.
- [37] Jordi Cortadella, Michael Kishinevsky, Luciano Lavagno, and Alexandre Yakovlev. Deriving Petri Nets for Finite Transition Systems. *IEEE Trans. Computers*, 47(8):859–882, 1998.
- [38] Gero Decker, Oliver Kopp, Frank Leymann, and Mathias Weske. Bpel4chor: Extending bpel for modeling choreographies. In *ICWS*, pages 296–303. IEEE Computer Society, 2007.

- [39] Pierre-Malo Deniélou and Nobuko Yoshida. Dynamic multirole session types. In Thomas Ball and Mooly Sagiv, editors, *POPL*, pages 435–446. ACM, 2011.
- [40] Pierre-Malo Deniélou and Nobuko Yoshida. Multiparty session types meet communicating automata. In Helmut Seidl, editor, *ESOP*, volume 7211 of *Lecture Notes in Computer Science*, pages 194–213. Springer, 2012.
- [41] Pierre-Malo Deniélou and Nobuko Yoshida. Multiparty compatibility in communicating automata: Characterisation and synthesis of global session types. In Fedor V. Fomin, Rusins Freivalds, Marta Z. Kwiatkowska, and David Peleg, editors, *ICALP (2)*, volume 7966 of *Lecture Notes in Computer Science*, pages 174–186. Springer, 2013.
- [42] Paul Gastin and François Laroussinie, editors. *CONCUR 2010 - Concurrency Theory, 21th International Conference, CONCUR 2010, Paris, France, August 31-September 3, 2010. Proceedings*, volume 6269 of *Lecture Notes in Computer Science*. Springer, 2010.
- [43] Kohei Honda, Vasco Thudichum Vasconcelos, and Makoto Kubo. Language primitives and type discipline for structured communication-based programming. In Chris Hankin, editor, *ESOP*, volume 1381 of *Lecture Notes in Computer Science*, pages 122–138. Springer, 1998.
- [44] Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty asynchronous session types. In George C. Necula and Philip Wadler, editors, *POPL*, pages 273–284. ACM, 2008.
- [45] Raymond Hu, Nobuko Yoshida, and Kohei Honda. Session-based distributed programming in java. In Jan Vitek, editor, *ECOOP*, volume 5142 of *Lecture Notes in Computer Science*, pages 516–541. Springer, 2008.

- [46] ITU-T. Message sequence chart (MSC). Technical Report ITU-T Recommendation Z.120, International Telecommunication Union, 2011.
- [47] Fabrizio Montesi Ivan Lanese and Gianluigi Zavattaro. Amending choreographies. In Editors Antonio Ravara, Josep Silva, editor, *Proceedings of WWV 2013, 9th International Workshop on Automated Specification and Verification of Web Systems*, EPTCS, page 15 pages, 2013. to appear.
- [48] Walter Knodel. New gossips and telephones. *Discrete Mathematics*, 13(1):95 –, 1975.
- [49] Maciej Koutny and Irek Ulidowski, editors. *CONCUR 2012 - Concurrency Theory - 23rd International Conference, CONCUR 2012, Newcastle upon Tyne, UK, September 4-7, 2012. Proceedings*, volume 7454 of *Lecture Notes in Computer Science*. Springer, 2012.
- [50] Ivan Lanese, Claudio Guidi, Fabrizio Montesi, and Gianluigi Zavattaro. Bridging the gap between interaction- and process-oriented choreographies. In Antonio Cerone and Stefan Gruner, editors, *SEFM*, pages 323–332. IEEE Computer Society, 2008.
- [51] Julien Lange and Alceste Scalas. Choreography synthesis as contract agreement. <http://www.cs.le.ac.uk/people/jlange/papers/agreement.pdf>, March 2013. To appear in ICE'13.
- [52] Julien Lange and Emilio Tuosto. A modular toolkit for distributed interactions. In Kohei Honda and Alan Mycroft, editors, *PLACES*, volume 69 of *EPTCS*, pages 92–110, 2010.
- [53] Julien Lange and Emilio Tuosto. Synthesising Choreographies from Local Session Types. In Koutny and Ulidowski [49], pages 225–239.

- [54] Sjouke Mauw. The formalization of message sequence charts. *Computer Networks and ISDN Systems*, 28(12):1643–1657, 1996.
- [55] Bertrand Meyer. Applying "design by contract". *IEEE Computer*, 25(10):40–51, 1992.
- [56] Robin Milner. *Communicating and mobile systems - the Pi-calculus*. Cambridge University Press, 1999.
- [57] Dimitris Mostrous, Nobuko Yoshida, and Kohei Honda. Global principal typing in partially commutative asynchronous sessions. In Castagna [27], pages 316–332.
- [58] Matthias Neubauer and Peter Thiemann. An implementation of session types. In Bharat Jayaraman, editor, *PADL*, volume 3057 of *Lecture Notes in Computer Science*, pages 56–70. Springer, 2004.
- [59] Nicholas Ng, Nobuko Yoshida, and Kohei Honda. Multipart session c: Safe parallel programming with message optimisation. In Carlo A. Furia and Sebastian Nanz, editors, *TOOLS (50)*, volume 7304 of *Lecture Notes in Computer Science*, pages 202–218. Springer, 2012.
- [60] Nicholas Ng, Nobuko Yoshida, Olivier Pernet, Raymond Hu, and Yiannos Kryptis. Safe parallel programming with session java. In Wolfgang De Meuter and Grigore-Catalin Roman, editors, *COORDINATION*, volume 6721 of *Lecture Notes in Computer Science*, pages 110–126. Springer, 2011.
- [61] OASIS. Web services business process execution language (WS-BPEL) version 2.0. <https://www.oasis-open.org/committees/wsbpel>, April 2007.
- [62] OMG. Unified modeling language (UML). <http://www.omg.org/spec/UML/>, August 2011.

- [63] Luca Padovani. On projecting processes into session types. *Mathematical Structures in Computer Science*, 22(2):237–289, 2012.
- [64] SAVARA and testable architecture. <http://www.jboss.org/savara>.
- [65] Scribble. <http://www.jboss.org/scribble>.
- [66] Alexandra Silva, Simon Bliudze, Roberto Bruni, and Marco Carbone, editors. *Proceedings Fourth Interaction and Concurrency Experience*, volume 59 of *EPTCS*, 2011.
- [67] Singularity. <http://research.microsoft.com/projects/singularity/>.
- [68] Kaku Takeuchi, Kohei Honda, and Makoto Kubo. An interaction-based language and its typing system. In Constantine Halatsis, Dimitris G. Maritsas, George Philokyrou, and Sergios Theodoridis, editors, *PARLE*, volume 817 of *Lecture Notes in Computer Science*, pages 398–413. Springer, 1994.
- [69] Hugo Torres Vieira, Luís Caires, and João Costa Seco. The conversation calculus: A model of service-oriented computation. In Sophia Drossopoulou, editor, *ESOP*, volume 4960 of *Lecture Notes in Computer Science*, pages 269–283. Springer, 2008.
- [70] Philip Wadler. Propositions as sessions. In Peter Thiemann and Robby Bruce Findler, editors, *ICFP*, pages 273–286. ACM, 2012.
- [71] Yin Wang, Ahmed Nazeem, and Ram Swaminathan. Finding the optimal representation for service composition using the theory of regions. Technical Report HPL-2010-191, HP Laboratories, 2011.
- [72] Yin Wang, Ahmed Nazeem, and Ram Swaminathan. On the optimal petri net representation for service composition. In *ICWS*, pages 235–242. IEEE Computer Society, 2011.

- [73] Nobuko Yoshida, Pierre-Malo Deniérou, Andi Bejleri, and Raymond Hu. Parameterised multiparty session types. In C.-H. Luke Ong, editor, *FOSSACS*, volume 6014 of *Lecture Notes in Computer Science*, pages 128–145. Springer, 2010.