Refinements for Multiparty Message-Passing Protocols

Specification-agnostic theory and implementation

Martin Vassor
University of Oxford, UK

Nobuko Yoshida
University of Oxford, UK

Abstract

Multiparty message-passing protocols are notoriously difficult to design, due to interaction mismatches that lead to errors such as deadlocks. Existing protocol specification formats have been developed to prevent such errors (e.g. multiparty session types (MPST)). In order to further constrain protocols, specifications can be extended with *refinements*, i.e. logical predicates to control the behaviour of the protocol based on previous values exchanged. Unfortunately, existing refinement theories and implementations are tightly coupled with specification formats.

This paper proposes a framework for multiparty message-passing protocols with refinements and its implementation in Rust. Our work *decouples* correctness of refinements from the underlying model of computation, which results in a *specification-agnostic* framework.

Our contributions are threefold. First, we introduce a trace system which characterises *valid refined traces*, i.e. a sequence of sending and receiving actions correct with respect to refinements. Second, we give a correct model of computation named *refined communicating system* (RCS), which is an extension of communicating automata systems with refinements. We prove that RCS only produce valid refined traces. We show how to generate RCS from mainstream protocol specification formats, such as *refined multiparty session types* (RMPST) or *refined choreography automata*. Third, we illustrate the flexibility of the framework by developing both a static analysis technique and an improved model of computation for dynamic refinement evaluation. Finally, we provide a Rust toolchain for decentralised RMPST, evaluate our implementation with a set of benchmarks from the literature, and observe that refinement overhead is negligible.

2012 ACM Subject Classification Software and its engineering → Specification languages; Theory of computation → Assertions; Theory of computation → Concurrency

Keywords and phrases Message-Passing Concurrency, Session Types, Specification

Digital Object Identifier 10.4230/LIPIcs.ECOOP.2024.35

Supplementary Material *Software (Artifact)*: `https://doi.org/10.4230/DARTS.3.2.13`

Funding Work supported by: EPSRC EP/T00006544/2, EP/K011715/1, EP/K034413/1, EP/L00058X/1, EP/N027833/2, EP/N028201/1, EP/T014709/2, EP/V000462, EP/X015955/1n NCSS/EPSRC VeTSS, and Horizon EU TaRDIS 101093006.

Acknowledgements We thank B. Ekici, M. Giunti, P. Hou, A. Suresh, and F. Zhou.

1 Introduction

Message passing programming is a notoriously difficult task with new bugs arising with respect to sequential programming, for instance deadlocks. To address this increased complexity, various specifications have been introduced (e.g., message sequence charts [24], multiparty session types [38, 19, 18], choreography automata [1]). In general, specifications are used to constrain messages, in order to prevent errors such as deadlocks (via message ordering) or

payload mismatch (by enforcing the sender and the receiver of a message to agree on the datatype exchanged).

In this paper, we tackle an important and advanced aspect of protocol specification, *logical constraints* (or *contracts*) on asynchronous message-passing communications. Contracts for heterogeneous systems are predominant for correctly designing, implementing, and composing software services, and have a long history in distributed software development as found in Design-by-Contracts [28], Service Level Agreements, and Component-Based Software Engineering. With contracts, software designers can define more precise (refined) and verifiable specifications for distributed software components. Contracts have been investigated from a variety of perspectives, using many different analysis techniques and formalisms. Our goal is to distill an essence of those models for protocol refinements by answering the following questions affirmatively: (i) what does it mean for an execution of contracts for message-passing systems to be correct; (ii) how do we integrate a theory to a variety of models; (iii) how do we analyse their correctness?; and (iv) how do we implement correct systems in a programming language?

To explain our framework, consider a *guessing game* (from [41]) with three participants where the first one (participant A) chooses a *secret* integer and sends it to the second participant (B). Then, the third participant (C) tries to *guess* this number. Depending on the guess, B replies with hints (*more* and *less*) until C succeeds in guessing the *correct* value.

The developer writing the specification for such protocol would like to ensure, *in the specification*, that hints from B are consistent with the previous values exchanged. For instance, if the *secret* is 5 and the guess is 10, the specification should constrain B to send *less*. Figure 1 shows a communication diagram of the protocol with constraints (which we call *refinements*) shown in light blue.

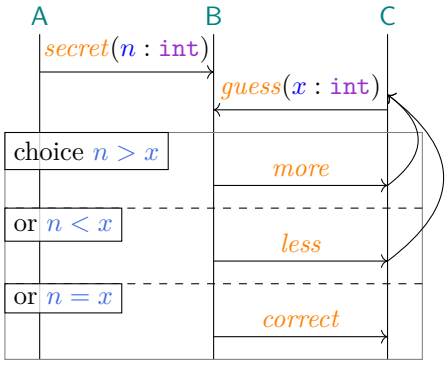

Figure 1 Communication diagram for the guessing game protocol with refinements.

In this paper, we develop a formal framework for refinements, agnostic to any particular specification formalism. Its core part is composed of a characterisation of refinement correctness: *Valid Refined Traces*, and a model of computation: *Refined Communicating Systems* (RCS), where communication is asynchronous and refinements are centrally and dynamically evaluated. For illustration, we use *Refined Multiparty Session Types* as the main specification format for multiparty protocols.

In addition, we demonstrate the versatility of our framework with multiple extensions. First, our framework can accommodate other protocol specification formats (e.g. choreography automata [1]). Second, it is used as a baseline for improved refinement evaluation: we present an optimised model of computation (decentralised refinement evaluation). Finally, it is also used as a baseline for implementing static analysis techniques: we present a simple strategy for statically removing redundant refinements.

Valid Refined Traces. The first building block is a common notion of *correct* executions with respect to added refinements. We introduce *valid refined traces* which are consistent traces with respect to refinements. This approach allows us to establish a general notion of refinements, which is applicable to different logics for constraints, type theories, models of computations, and programming languages. We consider asynchronous communications

(FIFOs), distinguishing *sending* and *receiving* actions in traces.

To illustrate our approach, consider the guessing game example shown above. Each execution of that protocol is recorded in a trace, i.e. a sequence of the individual events that take place during the execution (c.f. Section 2.2). For instance, a possible trace of the first four events of the protocol is the following:

$$\mathsf{A!B}\langle \mathit{secret}, \langle n, 5\rangle\rangle : \top \cdot \mathsf{A?B}\langle \mathit{secret}, \langle n, 5\rangle\rangle : \top \cdot \mathsf{C!B}\langle \mathit{guess}, \langle x, 5\rangle\rangle : \top \cdot \mathsf{C?B}\langle \mathit{guess}, \langle x, 5\rangle\rangle : \top$$

This trace contains four actions, and each action records an event, i.e. a message emission (denoted with a !) or reception (denoted with a ?). For instance, $\mathsf{A!B}\langle \mathit{secret}, \langle n, 5\rangle\rangle : \top$ records A sending a message to B, the payload of this message is a variable n, which has value 5. In the first four actions, we do not need any constraint, therefore actions are guarded by $\top$ which denotes a tautology predicate. The next action following this trace would be for B to send either *more*, *correct*, or *less*. Choosing *more* or *less* would be inconsistent with our protocol, since C guessed the correct number. For instance, choosing *more* would add the action $\mathsf{B!C}\langle \mathit{more}\rangle : n > x$ at the end of the queue: the refinement $n > x$ would be violated, since $x = n = 5$.

Valid Refined Traces characterise consistency based on the produced trace; and we aim to provide a model of computation constrained in a way that prevents such inconsistent choices.

Refined Communicating Systems. The second building block of our framework is a model of computation that only produces correct traces. *Communicating Systems* (CS) [5] are a model of concurrent computation, where *Communicating Finite State Machines* communicate asynchronously using unbounded FIFO queues. CS are often used to model and implement MPST [12, 13, 7]. We adapt CS to accommodate refinements, which we call *Refined Communicating Systems* (RCS). The semantics is modified in order to check refinements at every step. For this, we introduce a shared map in order to keep track of variables and their values that are exchanged in messages (e.g. the values of x and n in the guessing game example). This record of values is used to evaluate refinements, preventing undesired transitions. In this paper, we show that RCS only produce valid refined traces and we explain how to generate an RCS from a RMPST.

Refined MPST. Working with CS is cumbersome, and, in practice, we would prefer to adapt existing specification formats. We present in depth how to integrate refinements in *Multiparty Session Types* (MPST) [38, 19, 18], which are a family of type systems that aims to prevent communication bugs.

The following refined global type ($G_\pm$) is a specification of the guessing game protocol (Figure 1), with refinements: a participant A begins by sending a *secret* to B; the value of the *secret* is stored in the variable n. Then, C tries to guess the value (stored in variable x), and B replies with *more*, *less* (in which case the protocol loops and C can make another guess: $\mu\mathbf{T}.G$ denotes the recursion) or *correct*, at which point the protocol terminates ($\mathtt{end}$ denotes the termination). The refinements specify conditions upon which the *more*, *less*, and *correct* branches are possible. For instance, the protocol can take the *correct* branch only if the values in the *secret* and the *guess* messages are the same, i.e. if $x = n$.

$$G_\pm = \mathsf{A} \to \mathsf{B} \left\{ \mathit{secret}(n : \mathtt{int} \models \top).\mu\mathbf{T}.\mathsf{C} \to \mathsf{B} \left\{ \mathit{guess}(x : \mathtt{int} \models \top).\mathsf{B} \to \mathsf{C} \left\{ \begin{array}{l} \mathit{more}(\models x < n).\mathbf{T}, \\ \mathit{less}(\models x > n).\mathbf{T}, \\ \mathit{correct}(\models x = n).\mathtt{end} \end{array} \right. \right. \right.$$

Compared to standard MPST, *Refined MPST* (RMPST) contain variable names (n and x) and refinements (denoted with $\models r$ in the payloads, meaning that to send the message, r must hold). We present those extensions as well as the relation between RMPST and RCS.

Applications and Extensions. To show the versatility of our framework, we extend it:

Decentralised Refinement Evaluation: The canonical semantics for RCS we present uses a single shared map of variables to provide a simple way to reason about refinements. Having this global map would not be suited for a distributed implementation. We extend our framework with an alternative semantics where each participant of a protocol has a local map of variables. We show that if variables are not duplicated, then this alternative model also produces valid refined traces.

Static Elision of Redundant Refinements: At places where refinements are redundant (e.g. where it is entailed by previous refinements), we could benefit from removing those refinements. In order to show the versatility of our framework, we show how to develop a simple static analysis technique to remove such redundant refinements.

Refined Choreography Automata: While we mostly use RMPST as an example of protocol specification language, we sketch another specification by (informally) presenting how to integrate refinements in choreography automata (in Appendix G).

Rust Implementation. The last objective of our work is to implement RMPST into Rust. We choose Rust for several reasons: its affine type system makes it easy to avoid unwanted reuse of values, which helps to prevent a participant from duplicating actions; and thanks to its growing popularity, there are already a few existing toolchains for session types in Rust [27, 6, 26, 25]. Among them, we choose Rumpsteak [7] since it already uses CS to implement MPST participants inside its toolchain. We extend Rumpsteak with refinements using the decentralised refinement evaluation approach. We finally measure the refinement overhead in Rumpsteak.

Contributions and Outline. Our main contribution is to unify the different points presented above in a *single* framework as presented in Figure 2. We introduce a uniform framework which is agnostic to any particular specification formalism, model, semantics and language, defining the correctness of refinements as validity of traces. We then prove the safety of the framework (Theorem 18). We demonstrate the *versatility of our framework* by accommodating *multiple protocol specifications* such as (refined) multiparty session types [38, 19, 18, 42] and (refined) choreography automata [1, 16], *multiple semantics* such as (refined) communicating automata [5] with centralised and decentralised semantics, and *multiple analysis techniques* such as dynamic and static analyses. We provide an implementation of an instance of the framework in Rust. Our framework is the first, to the best of our knowledge, to achieve such versatility.

The framework is composed of the following parts (circled numbers refer to Figure 2):

① **Valid Refined Traces:** We introduce *valid refined traces* which characterise valid executions with respect to refinements.

② **Refined Communicating Systems (RCS):** We extend Communicating Systems to accommodate refinements. From a configuration of RCS, we induce a set of possible traces. One of our main results is Theorem 18 (④), which states that all traces produced by RCS are valid refined traces, which in turn proves the correctness of the RCS.

③ **Refined Multiparty Session Types (RMPST):** In Section 4, we adapt MPST (which consists of *global types* (which describe a multiparty protocol), *local types* (which describe the behaviour of a single participant), and a *projection* from global to local types which extracts the behaviour of a single participant) to accommodate for refinements. We show how to generate a RCS from a set of local types with refinements (⑤). In addition, in Appendix G, we sketch how to accommodate refinements in choreography automata, to illustrate the versatility of the framework (⑥).

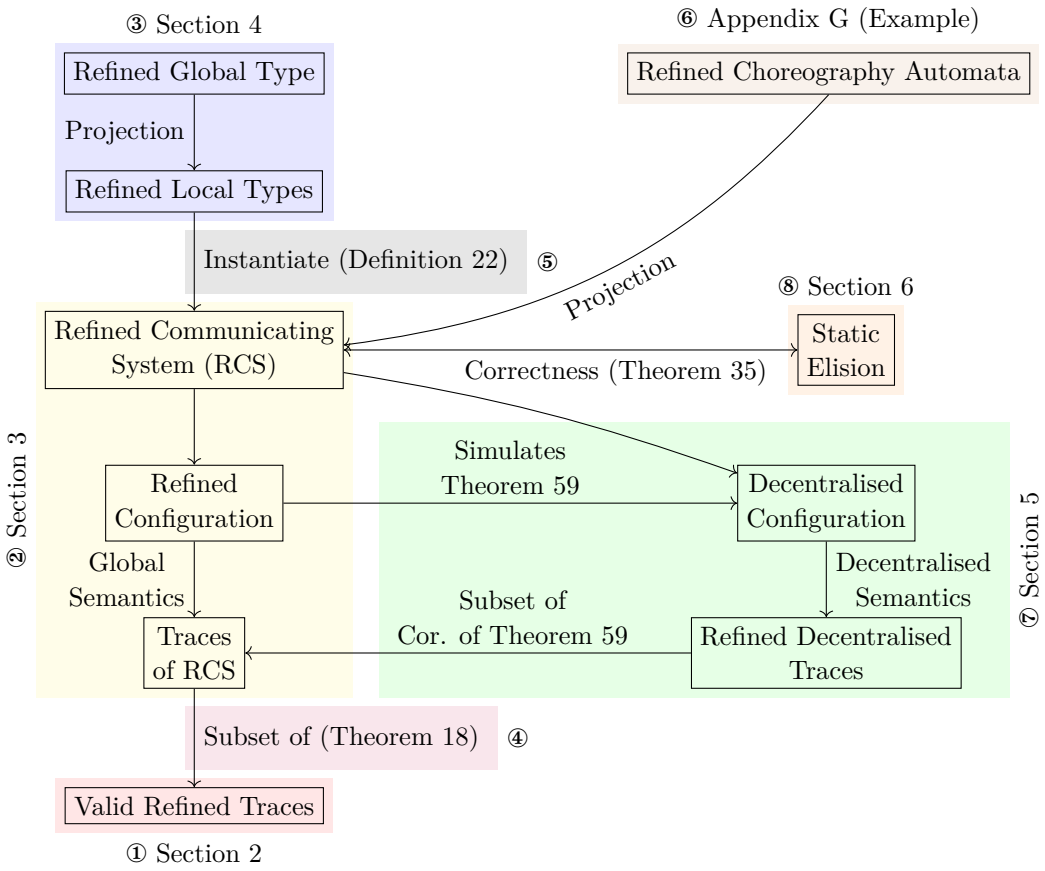

Figure 2 Overview of the framework for RMPST developed in this work. The coloured backgrounds show the main steps of this paper.

⑦ **and** ⑧ **Optimisations:** In Section 5 (⑦), we propose a *decentralised* model as an alternative for RCS. We show trace inclusion w.r.t. RCS, which ensures refinements are correctly checked. In Section 7, we implement this improved model in Rust. In addition, in Section 6, we demonstrate how to develop analysis techniques using the framework. We show how redundant refinements can, under some conditions, be statically removed (⑧).

2 Refined Traces and their Validity

This section introduces *refined traces* which are sequences of messages *actions*. We then define their *validity*, introducing two definitions on traces, *well-queued* and *well-predicated* traces. We precede this (in Section 2.1) with preliminary definitions used throughout this paper.

2.1 Preliminaries: Predicates Language and Semantics

This first subsection introduces the basic definitions we use in this paper.

Let $\mathbb{V}$ be a set of variables, ranged over by x, y, ...; and a finite set $\mathbb{C}$ of values (in this work, we take 32-bit integers: $\mathbb{Z}/2^{32}\mathbb{Z}$).

We use associative maps from variable names to values, noted M. $\mathtt{dom}(M)$ denotes the domain of a map, that is the set of variables that appear in the map. Maps are equipped with lookup ($M(x)$), update ($M[x \mapsto c]$) and removal ($M \backslash x$) operations. $M_1 \uplus M_2$ denotes the union of M_1 and M_2 if their domains are disjoint (see Appendix B.1 for the definition of all those operators). Finally, $M_\varnothing$ denotes an empty map.

In order to keep our work general, we do not strictly specify the language of predicates, nor their semantics rules. Instead, we suppose we are given a language to express refinements, whose terms are produced by a rule $\mathcal{R}$. In this paper, we intentionally leave the logic underspecified so that it can be fine tuned by the end user. In practice, in our implementation (Section 7), custom predicates can easily be added. In the following, we use a simple grammar with arithmetic and relational operators as predicates. Let $\mathbb{R}$ be the set of refinement expressions. We assume refinements can have free variables, and that there exist a function $\mathsf{fv} : \mathbb{R} \to \mathcal{P}(\mathbb{V})$ that gives the free variables of each refinement expression. We note $\mathbb{R}_\mathbb{W}$ be the set of refinements of $\mathbb{R}$ whose set of free variables is $\mathbb{W} \subseteq \mathbb{V}$. We assume a variable substitution function, $\mathcal{R}\{^{v_i}/_{x_i}\}$ that substitutes every free occurrence of each variable x_i for the value v_i. For any refinement expression r, $r\{^{\cdots}/_{\mathsf{fv}(r)}\}$ is a closed refinement. Since our predicates are abstract, we do not explicitly specify their semantics, nor their well-formedness. Instead, we assume each closed refinement formula evaluates to $\top$ or $\bot$. We assume there exists a function $\mathit{eval}(r)$ that evaluates the refinement r, provided that r is closed[1]. Finally, we assume the existence of a closed formula $\top$ that is a tautology, i.e. $\mathit{eval}(\top) = \top$.

Given a map M and a refinement r, we note $M \models r$ if and only if the refinement r is closed under the map M: $\mathsf{fv}(r) \subseteq \mathtt{dom}(M)$, and evaluates to $\top$ after substitution: $\mathit{eval}(r\{^{M(\mathsf{fv}(r))}/_{\mathsf{fv}(r)}\}) = \top$.

In a protocol with multiple participants, let $\mathbb{P}$ be a set of participants ranged over by $\mathsf{A}, \mathsf{B}, \ldots$ and $\mathsf{p}, \mathsf{q}, \ldots$ being meta-variables over participant names. In this work, messages contain a label, a variable, and a value. Let $\mathbb{L}$ be a set of labels; ℓ and its decorated variants range over labels in $\mathbb{L}$. We define $\mathbb{M} = \mathbb{L} \times (\mathbb{V} \times \mathbb{C})$ for the set of messages (as a reminder: $\mathbb{L}$ is the set of labels, $\mathbb{V}$ the set of variables, and $\mathbb{C}$ the set of values).

2.2 Traces

Let us denote $\vec{e} = e_1 :: \ldots :: e_n$ ($n \geq 0$) as a *FIFO*, i.e., a finite sequence of elements e_i (messages exchanged in this paper). We use ε for an empty FIFO ($n = 0$). We define: $\mathtt{enq}(\vec{e}, e) \stackrel{\text{def}}{=} e :: \vec{e}$; $\mathtt{deq}(\vec{e} :: e) \stackrel{\text{def}}{=} \vec{e}$ ($\mathtt{deq}(\varepsilon)$ is undefined); and $\mathtt{next}(\vec{e} :: e) \stackrel{\text{def}}{=} e$ ($\mathtt{next}(\varepsilon)$ is undefined). Notice that $\mathtt{deq}(\vec{e})$ is defined if and only if $\mathtt{next}(\vec{e})$ is defined. In this paper, we consider one FIFO channel per pair of participant. We call *queues* a map of all pairs of distinct participants to their communication FIFO of a system. We note $\mathtt{enq}_{(\mathsf{p},\mathsf{q})}(w, e)$, $\mathtt{deq}_{(\mathsf{p},\mathsf{q})}(w)$, $\mathtt{next}_{(\mathsf{p},\mathsf{q})}(w)$, where the indices indicates which FIFO of the set is affected (see Appendix B.1 for the formal definition). We write $w_\varnothing$ for the empty queue, which is the queue where $w_{(\mathsf{p},\mathsf{q})} = \varepsilon$ for all p and q.

Actions are tuples consisting of a sending participant p, a direction of communication $\dagger \in \{!, ?\}$ ($!$ stands for sending, and $?$ stands for receiving), a receiving participant q, a message m and a predicate r associated to the action (as a reminder: $\mathbb{R}$ is the set of refinements). We require participants to be distinct (i.e. $\mathsf{p} \neq \mathsf{q}$).

[1] We do not discuss the decidability of the actual chosen logic of refinements here. For undecidable logics, providing such function is, of course, not possible; however this is not in the scope of this work.

▶ **Definition 1** (Action and Trace). *An action is an element of* $\mathbb{A}$ *defined as follows:* $\mathbb{A} = \mathbb{P} \times \{!, ?\} \times \mathbb{P} \times \mathbb{M} \times \mathbb{R}$. *We write* $\alpha = \mathsf{p}\dagger\mathsf{q}\langle m \rangle : r$ $(\mathsf{p} \neq \mathsf{q})$ *when* $\langle \mathsf{p}, \dagger, \mathsf{q}, m, r \rangle \in \mathbb{A}$.

Traces (denoted by τ *and its decorated variants) are finite sequences of actions, defined inductively from the rule* $\mathcal{T} ::= \alpha \cdot \mathcal{T} \mid \epsilon$, *where* α *is an action. We write* $\mathbb{A}^\star$ *for the set of traces.* ◁

▶ **Example 2** (Trace). We presented a trace in Section 1.

We denote $\tau_1 \cdot \tau_2$ for the concatenation of two traces. We assume an intuitive notion of the size of trace, as well as lemmas that allow us to infer that, if the size is 0, then the trace is ϵ.

2.3 Properties of Refined Traces

In this subsection, we characterise the *correctness* of traces w.r.t. refinements.

There are two conditions valid traces should verify. First, the sending/reception of messages should be consistent (as with normal MPST). Second, for every action of the trace, predicates that guard the action should hold. We call traces that satisfy message consistency *well-queued traces*, and the traces that satisfy the predicates *well-predicated traces*. In the end, we consider traces that satisfy both conditions: we call those traces *valid refined traces*.

To start with well-queued traces, we first evaluate the impact of a trace on a queue, by looking at the effect of each action on that queue (Definition 3).

▶ **Definition 3** (Trace Ending Up with Queues, well-queued traces). *A trace* τ *ends up with the queue* w_f *w.r.t. a queue* w_i *if:*

1. *If* $\tau = \epsilon$, $w_i = w_f$; *and*
2. *If* $\tau = \mathsf{p}!\mathsf{q}\langle m \rangle : r \cdot \tau'$, *then* τ' *ends up with* w_f *w.r.t.* $\mathtt{enq}_{(\mathsf{p},\mathsf{q})}(w_i, m)$; *and*
3. *If* $\tau = \mathsf{p}?\mathsf{q}\langle m \rangle : r \cdot \tau'$, *then* τ' *ends up with* w_f *w.r.t.* $\mathtt{deq}_{(\mathsf{p},\mathsf{q})}(w_i)$ *and* $\mathtt{next}_{(\mathsf{p},\mathsf{q})}(w_i) = m$.

A trace τ *is* well-queued *with regards to the queue* w *if* τ *ends up with the empty queue* $w_\varnothing$ *with respect to an initial queue* w.

A trace τ *is valid if* τ *is well-queued with respect to the empty queue* $w_\varnothing$. ◁

▶ Remark 4. In Definition 3, we say w_i is the *initial* queue. ◁

Regarding well-predicated traces, the idea is to record the latest value of each variable in a map; and to use that map to evaluate refinements (Definition 5).

▶ **Definition 5** (Well-Predicated Traces). *A trace* τ *is* well-predicated *under a map* M, *if either* (i) $\tau = \epsilon$; *or* (ii) $\tau = \mathsf{p}\dagger\mathsf{q}\langle l, (x, c) \rangle : r \cdot \tau'$ *and* $M[x \mapsto c] \models r$ *and* τ' *is well-predicated under* $M[x \mapsto c]$. ◁

▶ **Example 6** (Well-Predicated Traces). In Section 1, we presented the trace τ:

$$\mathsf{A}!\mathsf{B}\langle \mathit{secret}, \langle n, 5 \rangle \rangle : \top \cdot \mathsf{A}?\mathsf{B}\langle \mathit{secret}, \langle n, 5 \rangle \rangle : \top \cdot \mathsf{C}!\mathsf{B}\langle \mathit{guess}, \langle x, 5 \rangle \rangle : \top \cdot \mathsf{C}?\mathsf{B}\langle \mathit{guess}, \langle x, 5 \rangle \rangle : \top$$

To illustrate Definition 5, we propose two actions after τ: (i) $\tau_1 = \mathsf{B}!\mathsf{C}\langle \mathit{more}, \langle \_, \_ \rangle \rangle : x > n$; and (ii) $\tau_2 = \mathsf{B}!\mathsf{C}\langle \mathit{correct}, \langle \_, \_ \rangle \rangle : x = n$. We can investigate whether $\tau \cdot \tau_1$ (resp. $\tau \cdot \tau_2$) is a well-predicated trace under $M_\varnothing$. According to Definition 5, we have to investigate whether τ_1 (resp. τ_2) is well predicated under $M = \{\langle n, 5 \rangle, \langle x, 5 \rangle\}$.

For τ_1, according to Item (ii) in Definition 5, then $x > n$ must hold under M, which is not the case, therefore $\tau \cdot \tau_1$ is not well-predicated.

Regarding τ_2, according to Item (ii) in Definition 5, then $x = n$ must hold under M, which is the case.

Finally, we consider traces that are both valid with respect to predicates and to messages. We call those *Valid Refined Traces*. Our overall goal is to show that our framework only produces such valid refined traces.

▶ **Definition 7** (Valid Refined Traces). *A refined trace τ is* valid *if* (i) *τ is well-queued with respect to the empty queue $w_\varnothing$; and* (ii) *τ is well-predicated under the empty map $M_\varnothing$.* ◁

3 Refined Communicating Automata

In this section, we model message-passing concurrent systems with refinements. We ensure that this model only generates valid refined traces (c.f. Definition 7). Our model of computation is an extension of *communicating systems* (CS) [5, 8], which are sets of Finite State Machines communicating using queues. We introduce *refined communicating systems* (RCS), a variant of CS which accounts for refinements and we show that all traces produced by RCS are valid refined traces (Theorem 18).

Refined Communicating Finite State Machines. *Communicating systems* [5] are a concurrent model of computation composed of a set of *communicating finite state machines* (CFSM) that interact with exchanges of messages. CFSM are standard finite state machines, where labels represent actions (i.e. sending or receiving messages). Individual FSM are then given a concurrent semantics, which performs messages exchanges. The state of the system is called a *configuration*, which records the state of the individual CFSMs as well as the content of the message queues. In this section, we adapt communicating systems for refinements.

First, we add refinements to the transitions of CFSM, which we call *refined* CFSM. This appears in the additional $\mathbb{R}$ in Definition 8 (we recall $\mathbb{R}$ is the set of refinements).

▶ **Definition 8** (Refined Communicating Finite State Machine (RCFSM)). *An RCFSM is a finite transition system given by $M = \langle Q, C, q_0, \mathbb{M}, \delta \rangle$, where Q is a set of states; $C = \{\mathsf{pq} \in \mathbb{P}^2 \mid \mathsf{p} \neq \mathsf{q}\}$ is a set of channels[2]; $q_0 \in Q$ is an initial state; $\mathbb{M}$ is a finite alphabet of messages; and $\delta \subseteq Q \times (C \times \{!, ?\} \times \mathbb{A} \times \mathbb{R}) \times Q$ is a finite set of transitions.* ◁

We write $s \xrightarrow{\mathsf{i}\dagger\mathsf{j}\langle m\rangle:r} s'$ for $\langle s, \langle \mathsf{ij}, \dagger, m, r\rangle, s'\rangle \in \delta$. *Refined communicating systems* (RCS) are analogous to their non-refined counterparts and simply consist of a tuple of RCFSM, with one RCFSM per participant. For *refined configurations*, as with (non-refined) configurations, we store the states of the individual CFSM and the content of queues. In addition, contrary to non-refined configurations, refined configurations also contain a map in order to keep track of the values of the variables in order to be able to evaluate refinements.

▶ **Definition 9** (Refined Communicating System (RCS)). *A refined communicating system is a tuple $R = \langle M_\mathsf{p} \rangle_{\mathsf{p} \in \mathbb{P}}$ of RCFSMs such that $M_\mathsf{p} = \langle Q_\mathsf{p}, C, q_{0\mathsf{p}}, \mathbb{M}, \delta_\mathsf{p} \rangle$.* ◁

An RCS uses one RCFSM per participant $\mathsf{i} \in \mathbb{P}$. A *configuration* represents the state of such RCS, where each participant i is in a local state s_i.

▶ **Definition 10** (Refined Configuration). *A* refined configuration *of an RCS R is a tuple S as follows: $S \stackrel{\text{def}}{=} \langle \langle s_1, \ldots, s_\mathsf{n} \rangle, w, M \rangle_R$ where each $s_\mathsf{i} \in Q_\mathsf{i}$, w is a queue of messages, and M is a map from variables names to values. Let $\mathbb{S}$ be the set of refined configurations.* ◁

[2] The original definition uses *channels*, which we do not use. We keep them for the sake of consistency.

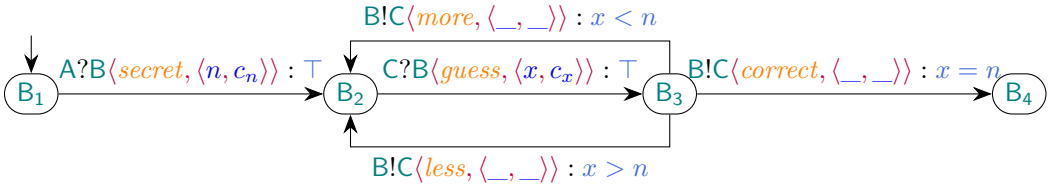

Figure 3 RCFSM of B in the $G_{\pm}$ protocol.

▶ Remark 11. Refined configurations are indexed by their RCS. This allows the configuration to store the automaton of the participant. The semantics developed below uses those (local) transitions to infer the global semantics. When the context is clear, we omit this index. ◁

From that, we characterise *initial* and *final* configurations. We call a configuration *initial* when it is a possible configuration where no actions have been taken yet. This means that there is no pending messages (which would imply a previous *send* action), nor known variables (which would imply a previous action initialised the variable). We say a configuration is *final* when there are no pending messages (otherwise, we would expect a *receive* action to take place). Notice that it does not mean the system cannot take action at all.

▶ **Definition 12** (Initial and Final Refined Configuration). *A refined configuration* $\langle\langle s_1, \ldots, s_n\rangle, w, M\rangle \in \mathbb{S}$ *is* initial *if and only if* (i) $w = w_\varnothing$; (ii) $M = M_\varnothing$; *and* (iii) *each* s_i *is initial in the RCFSM.*

A refined configuration $S = \langle\langle s_1, \ldots, s_n\rangle, w, M\rangle \in \mathbb{S}$ *is* final *if and only if* $w = w_\varnothing$. ◁

▶ **Example 13** (RCS). The RCFSM of participant B in the guessing game is shown in Figure 3. See Figure 9 for the RCFSM of A and C. Together, they form a RCS, which initial configuration is $\langle\langle \mathsf{A}_1, \mathsf{B}_1, \mathsf{C}_1\rangle, w_\varnothing, M_\varnothing\rangle$.

Refined Semantics. We now define the semantics of RCS in Definition 14 with two reduction rules GRRec and GRSnd (the initial GR stands for *global refined*, to distinguish the rules from variants in the following parts of this work), which are respectively used for receiving and sending messages. To avoid confusion with RCFSM reductions, we use a double arrow ($\Rightarrow$) to represent reductions at the refined communicating system level.

Rule GRSnd specifies that, if a participant i reduces from state s_i to state s_i' while sending a message m and if the refinement predicate r attached to the action holds, then the transition is taken at the global level. In the resulting refined configuration, the message is enqueued in the relevant queue and the map of known variables M is updated to take into account the new value of the carried variable c.

Rule GRRec is similar, with the additional requirement that the message received must be the next in the participant's queue (the third premise).

Notice that the verification of refinements is *dynamic*, as it is performed by the corresponding premise in each of the rules, i.e. at execution time.

▶ **Definition 14** (Refined Global Semantics). *Given a RCS* $R = \langle M_\mathsf{p}\rangle_{\mathsf{p}\in\mathbb{P}}$, *we define:*

$$GRRec\ \frac{t = s_i \xrightarrow{j?i\langle \ell, \langle x, c\rangle\rangle : r} s_i' \in \delta_i \qquad M[x \mapsto c] \models r \qquad \mathsf{next}_{(j,i)}(w) = \langle \ell, \langle x, c\rangle\rangle}{\langle\langle \ldots, s_i, \ldots\rangle, w, M\rangle_R \overset{t}{\Rightarrow} \langle\langle \ldots, s_i', \ldots\rangle, \mathsf{deq}_{(j,i)}(w), M[x \mapsto c]\rangle_R}$$

$$GRSnd\ \frac{t = s_i \xrightarrow{i!j\langle \ell, \langle x, c\rangle\rangle : r} s_i' \in \delta_i \qquad M[x \mapsto c] \models r}{\langle\langle \ldots, s_i, \ldots\rangle, w, M\rangle_R \overset{t}{\Rightarrow} \langle\langle \ldots, s_i', \ldots\rangle, \mathsf{enq}_{(i,j)}(w, \langle \ell, \langle x, c\rangle\rangle), M[x \mapsto c]\rangle_R}$$ ◁

▶ Remark 15. Global transitions are labelled with the underlying local transition. When the local transition is not relevant, we do not show it. ◁

▶ **Example 16** (Transitions of a RCS). Considering the RCS in Figure 9 in its initial configuration $C_i = \langle\langle \mathsf{A}_1, \mathsf{B}_1, \mathsf{C}_1\rangle, w_\varnothing, M_\varnothing\rangle$, we have that the automaton of A can fire a transition $\mathsf{A}_1 \xrightarrow{\mathsf{A!B}\langle secret, \langle n,5\rangle\rangle:\top} \mathsf{A}_2$, and $M_\varnothing[n \mapsto 5] \models \top$, by definition of $\top$. Therefore, C_i can take a GRSND transition and reduce to $\langle\langle \mathsf{A}_2, \mathsf{B}_1, \mathsf{C}_1\rangle, w, \{\langle n, 5\rangle\}\rangle$, where w contains a single message $\langle secret, \langle n, 5\rangle\rangle$ in $w_{(\mathsf{A},\mathsf{B})}$.

If the RCS is in the configuration $C = \langle\langle \mathsf{A}_2, \mathsf{B}_3, \mathsf{C}_2\rangle, w_\varnothing, M\rangle$ with $M = \{\langle x, 5\rangle, \langle n, 5\rangle\}$, the RCFSM of participant B offers three possible transitions: (i) $\mathsf{B}_3 \xrightarrow{\mathsf{B!C}\langle more, \langle \_,\_\rangle\rangle:x<n} \mathsf{B}_2$; (ii) $\mathsf{B}_3 \xrightarrow{\mathsf{B!C}\langle less, \langle \_,\_\rangle\rangle:x>n} \mathsf{B}_2$; and (iii) $\mathsf{B}_3 \xrightarrow{\mathsf{B!C}\langle correct, \langle \_,\_\rangle\rangle:x=n} \mathsf{B}_4$. The predicates carried in first two do not hold under M: $M \not\models x < n$ (resp. for $x > n$). Therefore, only $\mathsf{B}_3 \xrightarrow{\mathsf{B!C}\langle correct, \langle \_,\_\rangle\rangle:x=n} \mathsf{B}_4$ is feasible as a GRSND transition in the RCS. As we will see below (Theorem 18), this semantics prevents invalid traces.

Trace of Refined Communicating Systems. In order to show that the semantics of RCS captures the intuition of refinements, we study the traces formed by sequences of reductions (see Definition 47 for the formal definition of traces of RCS).

▶ **Example 17** (Trace of an RCS). The trace $\tau \cdot \tau_2$ (Example 6) is a trace of the RCS of $G_\pm$.

We conclude this section with our main result, which is that all traces produced by $\mathcal{S}(G)$ are valid refined traces. A trace is *initial* (resp. *final*) if it is obtained from a run whose first (resp. last) state is initial (resp. final).

▶ **Theorem 18** (Traces of Refined Communicating Systems are Valid Refined Traces). *For all RCS R, for all initial and final traces τ of R, τ is a valid refined trace.* ◁

The proof is in Appendix D.

4 Refined Multiparty Session Types (RMPST)

In the two previous sections, we introduced refinement validity and a variant of CS which is correct with respect to our validity criterion. However, working with RCS is cumbersome, in particular if we intend to prove additional properties (e.g. deadlock freedom). Fortunately, various models for message-passing concurrent computation have been developed in the literature, many of which can be encoded into CS. Multiparty session types (MPST) [38, 19, 18] is an example of such model. We focus on MPST as they have proved successful for many applications and the theory enjoy many useful properties (e.g. session fidelity, deadlock freedom, liveness etc). However, MPST is not the only possible choice, and we sketch different input models in Appendix G. In this section, we introduce *refined multiparty session types* (RMPST), which are an extension of MPST annotated with refinement predicates and we show how one can extend existing models to easily obtain refinements.

In Section 4.1, we first present the syntax of *global* and *local* refined multiparty session types, adapted for refinements. In Section 4.2, we present how to obtain RCS from local RMPST, extending a standard approach to implement MPST in CS [12] with refinements.

4.1 Syntax of RMPST

We define the syntax of RMPST. First we assume that messages carry different sorts of payload. As a reminder, for simplicity, in our examples, we only consider `int` payloads.

$$
\begin{array}{lcll}
\mathcal{G} & ::= & \mathsf{p} \to \mathsf{q}\{l_i(x_i : \mathcal{S} \models \mathcal{R}).\mathcal{G}\}_{i\in I} \mid \mu \mathbf{t}.\mathcal{G} & \textit{communication, recursive type} \\
& \mid & \mathbf{t} \mid \mathtt{end} & \textit{type variable, termination} \\
\mathcal{L} & ::= & \mathsf{p}\oplus\{l_i(x_i : \mathcal{S} \models \mathcal{R}).\mathcal{L}\}_{i\in I} \mid \mathbf{t} \mid \mathtt{end} & \textit{internal choice, type variable, termination} \\
& \mid & \mathsf{p}\&\{l_i(x_i : \mathcal{S} \models \mathcal{R}).\mathcal{L}\}_{i\in I} \mid \mu \mathbf{t}.\mathcal{L} & \textit{external choice, recursive type} \\
\mathcal{S} & ::= & \mathtt{int} \mid \ldots & \textit{sort (payload types)}
\end{array}
$$

Figure 4 Syntax of Global ($\mathcal{G}$) and Local ($\mathcal{L}$) Types and Sorts ($\mathcal{S}$).

Also, we recall the conventions from Section 2.1: $\mathbb{P}$ is the set of participants and $\mathbb{L}$ is the set of labels. For recursion, we introduce type variables that range over $\{\mathbf{T}, \mathbf{U}, \ldots\}$; $\mathbf{t}$ is a meta-variable taken over the set of type variables. We assume all type variables appearing in a type are distinct and we do not (syntactically) distinguish global and local type variables. Finally, x_i are meta-variables over payload variables taken from the set $\mathbb{V}$.

We first define *global refined multiparty session types*, which are inductive data types generated by the production $\mathcal{G}$ in Figure 4. The type $\mathsf{A} \to \mathsf{B}\{l_i(x_i : \mathsf{S}_i \models r_i).G_i\}_{i\in I}$ describes a protocol where A chooses a label l_i amongst possible I and sends a message to B. The message contains a payload of type $\mathsf{S_i}$, which is bound to x_i when sent. *Refinement predicates* we introduce guard the communication they are attached to, meaning the system can select a choice with predicate r_i only if r_i holds. In that case, the message is sent and the protocol continues with its continuation of type G_i. $\mu\mathbf{T}.G$ binds $\mathbf{T}$ in G, and a bound type variable $\mathbf{T}$ in a type denotes a protocol recursion. Let $\mathtt{frv}(G)$ denotes the free recursion variables occuring in G. Finally $\mathtt{end}$ describes a terminated protocol. Let $\mathtt{parts}(G)$ be the set of participants that appear in G (c.f. Definition 50 for the definition of $\mathtt{parts}(G)$). We write $\mathsf{p} \in G$ for $\mathsf{p} \in \mathtt{parts}(G)$.

▶ **Example 19** (Refined Global Multiparty Session Type)**.** The type $G_{\pm}$ presented in Section 1 is a refined global type; we have $\mathtt{parts}(G_{\pm}) = \{\mathsf{A}, \mathsf{B}, \mathsf{C}\}$.

To characterise the behaviour of individual participants, we define *refined local multiparty session types*, which are inductive datatypes generated by $\mathcal{L}$ in Figure 4. Recursion, type variables and termination are similar in local and global types. Only the communication specifications differs: in a local type $\mathsf{p}\oplus\{l_i(x_i : \mathsf{S}_i \models r_i).L_i\}_{i\in I}$ describes an *internal choice*, i.e. the participant chooses a label l_i and sends it to p. Conversely, $\mathsf{p}\&\{l_i(x_i : \mathsf{S}_i \models r_i).L_i\}_{i\in I}$ describes an *external choice*: p makes a choice amongst the possible l_i and the local participant *receives* this choice.

Global and local MPST are related: we can *project* a global type onto the local types of its participants. Below, we define a *projection* (partial) operator $G\upharpoonright_{\mathsf{p}}$, which returns the local type of p with respect to the global type G.

We define a projection with a *merge* (partial) operator, which merges multiple local types of a participant into a single local type. This is used to merge the (possibly different) types of the continuations present in the communication branches. The study of different variants of merge operators is an active field (e.g. [32, Section 3]). For the sake of simplicity, in this paper we use a naïve merge operator, which simply ensures that all types are the same.

▶ **Definition 20** (Projection)**.** *Given* p*,* q *and* r *three distinct participants:*

$$\mathsf{p} \to \mathsf{q}\{l_i(x_i : \mathtt{S}_i \models R_i).G_i\}_{i\in I}\upharpoonright_{\mathsf{p}} = \mathsf{q}\oplus\{l_i(x_i : \mathtt{S}_i \models R_i).G_i\upharpoonright_{\mathsf{p}}\}_{i\in I}$$
$$\mathsf{q} \to \mathsf{p}\{l_i(x_i : \mathtt{S}_i \models R_i).G_i\}_{i\in I}\upharpoonright_{\mathsf{p}} = \mathsf{q}\&\{l_i(x_i : \mathtt{S}_i \models \top).G_i\upharpoonright_{\mathsf{p}}\}_{i\in I}$$
$$\mathsf{q} \to \mathsf{r}\{l_i(x_i : \mathtt{S}_i \models R_i).G_i\}_{i\in I}\upharpoonright_{\mathsf{p}} = \sqcap_{i\in I}(G_i\upharpoonright_{\mathsf{p}})$$

$$\mu\mathbf{t}.G'\upharpoonright_{\mathsf{p}} = \begin{cases} \mu\mathbf{t}.(G'\upharpoonright_{\mathsf{p}}) & \text{if } \mathsf{p} \in G' \text{ or } \mathsf{frv}(G') \neq \varnothing \\ \mathit{end} & \text{otherwise} \end{cases} \qquad \mathbf{t}\upharpoonright_{\mathsf{p}} = \mathbf{t} \qquad \mathit{end}\upharpoonright_{\mathsf{p}} = \mathit{end}$$

where a merge *operator is defined as:* $\sqcap_{i\in I} L_i \stackrel{\text{def}}{=} L$ *if* $\forall i \in I \cdot L = L_i$, *undefined otherwise.* ◁

Notice that our local RMPST accept refinements on both receiving and sending, and the semantics developed in Section 3 accept any position for verification. When projecting a global type $G = \mathsf{A} \to \mathsf{B}\,\{\ell(x : \mathtt{int} \models r).\mathtt{end}\}$ onto local types, we therefore have a choice to project the refinement:

- on the send side: $G\upharpoonright_{\mathsf{A}} = \mathsf{B}\oplus\{\ell(x : \mathtt{int} \models r).\mathtt{end}\}$ and $G\upharpoonright_{\mathsf{B}} = \mathsf{A}\&\{\ell(x : \mathtt{int} \models \top).\mathtt{end}\}$
- on the receive side: $G\upharpoonright_{\mathsf{A}} = \mathsf{B}\oplus\{\ell(x : \mathtt{int} \models \top).\mathtt{end}\}$ and $G\upharpoonright_{\mathsf{B}} = \mathsf{A}\&\{\ell(x : \mathtt{int} \models r).\mathtt{end}\}$
- or a combination of both[3].

Our projection takes the first option, i.e. refinements are checked when the message is emitted, but with any of these choices, our developments would not substantially change.

▶ **Example 21** (Projection). We project $G_{\pm}$ (Section 1) onto participants A and B[4]:
$G_{\pm}\upharpoonright_{\mathsf{A}} = \mathsf{B}\oplus\{\mathit{secret}(n : \mathtt{int} \models \top).\mathtt{end}\}$
$G_{\pm}\upharpoonright_{\mathsf{B}} =$

$$\mathsf{A}\&\left\{\mathit{secret}(n : \mathtt{int} \models \top).\mu\mathbf{T}.\mathsf{C}\&\left\{\mathit{guess}(x : \mathtt{int} \models \top).\mathsf{C}\oplus\left\{\begin{array}{l}\mathit{more}(\models x < n).\mathbf{T},\\ \mathit{less}(\models x > n).\mathbf{T},\\ \mathit{correct}(\models x = n).\mathtt{end}\end{array}\right.\right\}\right\}$$

4.2 From Refined MPST to Refined Communicating System

In this subsection, we show how to generate an RCS from local RMPSTs. As shown in Definition 20, local types are projected from global multiparty session types. Therefore, this step allows us to complete the conversion from a global RMPST into an RCS. We adapt the translation from local type to CFSMs presented in [13] to accommodate refinements in types.

The intuition behind the translation is to decompose a local type into the individual steps it specifies. For this, we first need to retrieve all those steps. We define the set of types that occur nested in another type: a type T' *occurs* in a type T (noted $T' \in T$) if it appears in the continuations of T after one or multiple steps (see Definition 51).

Given this, we can proceed to the translation itself, in Definition 22. This definition says that the states of the RCFSM of a local type T_0 are composed of the (sub)types that appear in T_0, stripped of the leading $\mu\mathbf{t}.$ (the function `strip` removes the leading recursions variables; this formalises [13, Item (2) in Definition 3.4]) and of recursive variables $\mathbf{t}$. We define the set of transitions of this RCFSM by taking the action each type (i.e. each state) can take, and adding a transition with this action from the initial state to the continuation (stripped of leading $\mu\mathbf{t}.$). In the case that the continuation is a recursion variable $\mathbf{t}$, we have to search in

[3] For instance, if we want to implement a centralised server that communicates with (isolated) clients, we may want all refinements to be asserted by the server, independently of the direction.
[4] The projection onto C is similar to the recursive part of the projection onto B, with ! and ? swapped.

the original type the continuation. Compared to [13, Item (2) in Definition 3.4], we simply add the support for the refinements predicates, which appear both in the types (i.e. in the states) and in the actions (i.e. in the transitions).

▶ **Definition 22** (RCFSM of Refined Local Types (extends Definition 3.5 in [13])). *Given a global type G, the RCFSM of participant p (with local type $T_0 = G\upharpoonright_{\mathsf{p}}$) is the automaton $\mathcal{A}(T_0) = \langle Q, C, \mathsf{strip}(T_0), \mathbb{M}, \delta\rangle$ where:*

- $Q = \{T' \mid T' \in T_0 \wedge T' \neq \mathbf{t} \wedge T' \neq \mu\mathbf{t}.T_\mu\}$;
- $C = \{\mathsf{pq} \mid \mathsf{p}, \mathsf{q} \in G, \mathsf{p} \neq \mathsf{q}\}$; *and*
- δ *is the smallest set of transitions such that: for all* $T \in T_0$ *in* Q, *for all* $c \in \mathbb{C}$:
  - *if* T *is* $\mathsf{q}\oplus\{\ell_i(x_i : \mathsf{S}_i \models r_i).T_i\}_{i\in I}$, *for all* T_i:
    - *if* $T_i \neq \mathbf{t}$, *then* $\langle T, \mathsf{p}!\mathsf{q}\langle\ell_i, \langle x, c\rangle\rangle : r, \mathsf{strip}(T_i)\rangle \in \delta$
    - *if* $T_i = \mathbf{t}$ *with* $\mu\mathbf{t}.T' \in T_0$, *then* $\langle T, \mathsf{p}!\mathsf{q}\langle\ell_i, \langle x, c\rangle\rangle : r, \mathsf{strip}(T')\rangle \in \delta$
  - *if* T *is* $\mathsf{q}\&\{\ell_i(x_i : \mathsf{S}_i \models r_i).T_i\}_{i\in I}$, *for all* T_i:
    - *if* $T_i \neq \mathbf{t}$, *then* $\langle T, \mathsf{q}?\mathsf{p}\langle\ell_i, \langle x, c\rangle\rangle : r, \mathsf{strip}(T_i)\rangle \in \delta$
    - *if* $T_i = \mathbf{t}$ *with* $\mu\mathbf{t}.T' \in T_0$, *then* $\langle T, \mathsf{q}?\mathsf{p}\langle\ell_i, \langle x, c\rangle\rangle : r, \mathsf{strip}(T')\rangle \in \delta$

where $\mathsf{strip}(T) \stackrel{\text{def}}{=} \mathsf{strip}(T')$ *if* $T = \mu\mathbf{t}.T'$; *and* $\mathsf{strip}(T) \stackrel{\text{def}}{=} T$ *otherwise.* ◁

Finally, we define the *RCS of a type*.

▶ **Definition 23** (Refined Communicating System of a Type). *The RCS of a type G, noted $\mathcal{S}(G)$, is a tuple composed of the RCFSM of all participants $\mathcal{S}(G) \stackrel{\text{def}}{=} \langle\mathcal{A}(G\upharpoonright_{\mathsf{p}})\rangle_{\mathsf{p}\in\mathtt{parts}(G)}$. We note $\mathcal{C}(G)$ the initial configuration of $\mathcal{S}(G)$.* ◁

▶ **Example 24** (Refined Communicating System of $G_\pm$). The communicating system of $G_\pm$ is $\mathcal{S}(G_\pm) = \langle\mathcal{A}(G_\pm\upharpoonright_{\mathsf{A}}), \mathcal{A}(G_\pm\upharpoonright_{\mathsf{B}}), \mathcal{A}(G_\pm\upharpoonright_{\mathsf{C}})\rangle$, where the three automata are shown in Figure 9 .
The initial configuration $\mathcal{C}(G_\pm)$ of this RCS $\mathcal{S}(G_\pm)$ is $\langle\langle\mathsf{A}_1, \mathsf{B}_1, \mathsf{C}_1\rangle, w_\varnothing, M_\varnothing\rangle$.

Theorem 18 applies to RCS obtained from RMPST: RCS generated from Definition 23 only produce valid refined traces, with the refined global semantics presented in Definition 14. Notice also that, if refinements always hold, RMPST and their semantics coincide with the semantics presented in [12].

5 Decentralised Refined Multiparty Session Types

In the previous section, we presented RCS and we showed that every trace of an RCS is a valid refined trace. However, RCS are theoretical constructions and are not intended to be implemented directly, as they use a global shared map of variables. In practice, a user may want to develop more precise analysis techniques on specific classes of RCS to remove this global map, which allows a decentralised verification of refinements, while keeping the validity of refined traces.

The goal of this section is twofold: on the one hand, the decentralised semantics we develop serves as a theoretical background for our implementation (Section 7). On the other hand, it illustrates the modularity of our framework. We show that the decentralised approach produces valid refined traces by showing refined configurations we developed in Section 3 simulate decentralised systems. This approach is not specific to our variant: we expect other optimisations presented in the literature could be integrated similarly.

This section is divided in the following steps: first, we define what we mean by *decentralised verification of the refinements*, by adapting the semantics of RCS (Definitions 25 and 28). We

split the global map of variables' values into local maps (one per participant). Then, we show that despite being modified, the new variant still produces valid refined traces (Definition 7). We justify this claim by proving that under some conditions, the original RCS *simulates* (c.f. [30, Exercise 1.4.17, p. 26]) the decentralised variant (Theorem 59). Since trace equivalence is coarser than simulation, this is sufficient to prove that decentralised configurations that meet the said conditions produce valid refined traces.

The conditions we mentioned above are: (i) variables should not be duplicated; and (ii) when evaluating a predicate, the free variables of the predicate must be in the local map. Without the first condition, we can possibly have two distinct values assigned to the same variable without being able to distinguish which is the most recent. The second condition is required to verify the refinements locally (e.g. predicates that constraint an action of A should be checked by A itself). To illustrate the importance of the first condition, consider the type $\mathsf{A} \to \mathsf{B}\,\{\ell_1(x : \mathtt{int}).\mathsf{C} \to \mathsf{D}\,\{\ell_2(x : \mathtt{int}).\mathtt{end}\}\}$. In the centralised approach, x is aliased, while in the decentralised approach, the x exchanged between A and B is stored in a local map, and the x exchanged between C and D is stored in another local map; both are not aliased. To prevent different semantics, we need to prevent such difference, which is the goal of the first condition.

Decentralised Configurations and Decentralised Semantics. First, we define *decentralised* configurations in Definition 25. Compared to Definition 10, instead of a global map in the tuple, we associate a local map to each automata state. Those maps store the variables each participant has access to.

▶ **Definition 25** (Decentralised Configuration)**.** *A* Decentralised Configuration *of an RCS* $\mathcal{S}(G) = \langle\langle Q_{\mathsf{i}}, C_{\mathsf{i}}, q_{0,\mathsf{i}}, \mathbb{A}, \delta_{\mathsf{i}}\rangle\rangle_{\mathsf{i}\in\mathtt{parts}(G)}$ *is a tuple* $\langle\langle\langle s_1, M_1\rangle, \ldots, \langle s_{\mathsf{n}}, M_{\mathsf{n}}\rangle\rangle, w\rangle_{\mathcal{S}(G)}$ *where each* $s_{\mathsf{i}} \in Q_{\mathsf{i}}$*, each* M_{i} *is a local map of variables to values, and* w *is a queue of messages.*

Let $\mathbb{S}_D$ *be the set of decentralised configurations. We note* $\mathcal{D}(G)$ *the initial decentralised configuration of* $\mathcal{S}(G)$*.* ◁

Remark 11 also applies for decentralised configurations.

▶ **Example 26** (Initial decentralised configuration of $G_\pm$)**.** In Example 13, we presented the refined communicating system of $G_\pm$ and its associated refined configuration. The *initial decentralised configuration* of this system is $\langle\langle \mathsf{A}_1, M_\varnothing\rangle, \langle \mathsf{B}_1, M_\varnothing\rangle, \langle \mathsf{C}_1, M_\varnothing\rangle, w_\varnothing\rangle$. In particular, notice that it uses the same set of refined CFSM than the refined configuration.

The global reduction rules are adapted accordingly: in the rules DRec and DSnd ("D" stands for "decentralised"), when a message is sent or received, the corresponding local map is updated, instead of a global map as in GRRec and GRSnd.

▶ Remark 27. Contrary to Definition 14, when a variable is sent, it is removed from the local map of variables. Intuitively, when a participant sends a variable, it erases its knowledge of it, to prevent aliasing issues. A direct consequence of this is that, in the centralised implementation, the global map of variables is a *superset* of the local maps in the corresponding decentralised implementation. Indeed, while a variable is in transit, it appears neither in the sender's map, nor in the receiver's one. This observation will be proved together with the simulation proof (Theorem 59). ◁

▶ **Definition 28** (Decentralised Global Semantics). *Given an RCS* $R = \langle M_{\mathsf{p}} \rangle_{\mathsf{p} \in \mathbb{P}}$

$$DR_{EC}\ \frac{t = s_{\mathsf{i}} \xrightarrow{\mathsf{j}?\mathsf{i}\langle \ell, \langle x, c \rangle\rangle : r} s'_{\mathsf{i}} \in \delta_{\mathsf{i}} \qquad \mathtt{next}_{(\mathsf{j},\mathsf{i})}(w) = \langle \ell, \langle x, c \rangle\rangle \qquad M_{\mathsf{i}}[x \mapsto c] \models r}{\langle \langle \ldots, \langle s_{\mathsf{i}}, M_{\mathsf{i}} \rangle, \ldots \rangle, w \rangle_R \overset{t}{\Rightarrow} \langle \langle \ldots, \langle s_{\mathsf{i}}, M_{\mathsf{i}}[x \mapsto c] \rangle, \ldots \rangle, \mathtt{deq}_{(\mathsf{j},\mathsf{i})}(w) \rangle_R}$$

$$DS_{ND}\ \frac{t = s_{\mathsf{i}} \xrightarrow{\mathsf{i}!\mathsf{j}\langle \ell, \langle x, c \rangle\rangle : r} s'_{\mathsf{i}} \in \delta_{\mathsf{i}} \qquad M_{\mathsf{i}}[x \mapsto c] \models r}{\langle \langle \ldots, \langle s_{\mathsf{i}}, M_{\mathsf{i}} \rangle, \ldots \rangle, w \rangle_R \overset{t}{\Rightarrow} \langle \langle \ldots, \langle s_{\mathsf{i}}, M_{\mathsf{i}} \backslash x \rangle, \ldots \rangle, \mathtt{enq}_{(\mathsf{i},\mathsf{j})}(w, \langle \ell, \langle x, c \rangle\rangle) \rangle_R} \quad \triangleleft$$

Conditions for Decentralised Verification and Correctness Proofs. We now focus on proving that this decentralised semantics produces valid refined traces. As we mentioned above, this holds under two conditions, which we define first:

▶ **Definition 29** (Conditions for Decentralised Verification Simulation). *Given a decentralised configuration* $\langle \langle \langle s_{\mathsf{i}}, M_{\mathsf{i}} \rangle, \ldots \rangle, w \rangle$*, the* conditions for simulation *are:*

1. *No duplication:*
   - **a.** *if* $\exists M_{\mathsf{i}} \cdot x \in \mathtt{dom}(M_{\mathsf{i}})$*, then* $\forall \mathsf{i}, \mathsf{j} \cdot x \notin w_{(\mathsf{i},\mathsf{j})}$ *and* $\forall \mathsf{j} \neq \mathsf{i} \cdot x \notin \mathtt{dom}(M_{\mathsf{j}})$.
   - **b.** *if* $\exists \langle \mathsf{i}, \mathsf{j} \rangle \cdot x \in w_{(\mathsf{i},\mathsf{j})}$*, then* $\forall \mathsf{i} \cdot x \notin \mathtt{dom}(M_{\mathsf{i}})$ *and* $\forall \langle \mathsf{i}', \mathsf{j}' \rangle \neq \langle \mathsf{i}, \mathsf{j} \rangle \cdot x \notin w_{(\mathsf{i}',\mathsf{j}')}$.
2. *Free variables are in the map:* $\forall \mathsf{i} \cdot \forall s_{\mathsf{i}}' \cdot s_{\mathsf{i}} \xrightarrow{\mathsf{i}\dagger\mathsf{j}\langle \ell, \langle x, c \rangle\rangle : r} s'_{\mathsf{i}} \cdot \mathrm{fv}(r) \subseteq \mathtt{dom}(M_{\mathsf{i}}[x \mapsto c])$ $\triangleleft$

▶ **Definition 30** (Decentralisable Type). *A type* G *is* decentralisable *if the two conditions hold for all reachable decentralised configurations from* $\mathcal{D}(G)$. $\triangleleft$

Notice that the second condition is redundant, as the condition $M_{\mathsf{i}}[x \mapsto c] \models r$ (in the premises of the reduction rules) already requires that $\mathrm{fv}(r)$ is a subset of the variables in $M_{\mathsf{i}}[x \mapsto c]$. Even without making this condition explicit, the system would stall if a predicate cannot be verified. For the sake of clarity, we keep it explicit in the two required conditions.

We now observe a correspondence between the (centralised) refined configuration and the decentralised configuration of a global type G. To characterise the correspondence between centralised and decentralised configuration, we establish a *simulation* relation between the two (see Appendix E.2 and [30]). Intuitively, a simulation captures the fact that one system (the centralised configuration in our case) can mimic all transitions of another system (the decentralised one here).

We can now prove the main result of this section, which is that the decentralised semantics does not induce new (unwanted) behaviours, i.e. all decentralised transitions can be mimicked by centralised transitions, i.e. the centralised approach simulates the decentralised one.

▶ **Theorem 31** (Centralised simulates Decentralised). *For all decentralisable RMPST* G *(Definition 30),* $\mathcal{C}(G)$ *simulates* $\mathcal{D}(G)$. $\triangleleft$

Proof. The proof is available in Appendix E.3. ◀

This result shows that any type that verifies the conditions stated in Definition 29 can be verified in a decentralised way. The difficulty is that the conditions are about the execution: we do not know whether a predicate will have a missing variable during the execution. With a knowledge flow algorithm, we can infer (from the communication specifications in the global type) which participant has access to which variables at any point in the execution of the protocol, i.e. we can *localise* each variable throughout the execution of the protocol. This algorithm (which we present in Appendix E.4) does not present major challenges.

Notice that the reverse simulation does not hold: $\mathcal{D}(G)$ does not simulate $\mathcal{C}(G)$. Indeed, $\mathcal{C}(G)$ can verify a predicate whose variables are spread over different participants, i.e. where variables would be spread across multiple M_{i} in the decentralised variant.

6 Static Elision of Redundant Refinements

In this section, we present a second optimisation, which is complimentary from the first one. The main idea is to statically analyse a given protocol to find and remove redundant refinements. Our approach is to consider a *target* transition, which we aim to remove the refinement, if possible. Our optimisation can then be applied successively to different target transitions one after each other. For instance, consider the following protocol G_s. We target the second refinement, $x < 10$, which necessarily holds if the first one does (since x does not change). Therefore it is redundant and can be removed.

$$G_s = \mathsf{A} \to \mathsf{B}\,\{\ell_1(x : \mathtt{int} \models x < 0).\mathsf{A} \to \mathsf{B}\,\{\ell_2(y : \mathtt{int} \models x < 10).\mathtt{end}\}\}$$

However, removing refinements is not always trivial, since the communication semantics is asynchronous. Consider for instance the following type :

$$\mathsf{A} \to \mathsf{C}\,\{\ell_1(x : \mathtt{int} \models x > 20).\mathsf{A} \to \mathsf{B}\,\{\ell_2(x : \mathtt{int} \models x < 0).\mathsf{C} \to \mathsf{B}\,\{\ell_3(y : \mathtt{int} \models x < 10).\mathtt{end}\}\}\}$$

A naïve approach would be to remove the refinement of the last communication ($x < 10$), since the previous communication has a stronger guarantee ($x < 0$). However, due to the asynchrony of communications, the second and third communications could be swapped at runtime, but the refinement ($x < 10$) does not hold before the action $\mathsf{A} \to \mathsf{B}\,\{\ell_2(x : \mathtt{int} \models x < 0)\ldots\}$ occurs. Therefore, in this case, removing the last refinement is incorrect. The optimisation we present takes into account those cases, by keeping track of causal relations between actions.

This section is independent of the previous one, although this second optimisation can help to make some protocols localisable: for instance, G_s above is not localisable. Since the second step $\mathsf{A} \to \mathsf{B}\,\{\ell_2(y : \mathtt{int} \models x < 10).\mathtt{end}\}$ requires A to access x, which is at B. However, once removed, the protocol becomes localisable, and can therefore be decentralised, helping the first optimisation introduced in Section 5.

As with the previous section, the optimisation we present could easily be further improved. Here, we focus on a simple case, as our goal is not to discuss the optimisation itself, but rather to show the versatility of the framework.

We present this section in two steps: first, in Section 6.1, we focus on RCS, which form the core of our framework; then, in Section 6.2, we apply the above result to RMPST.

6.1 Static Elision of Refinements in RCS

In a first step, we develop and prove the correctness of our analysis in RCS. The question is therefore whether, given a RCS R with one CFSM containing a transition with refinement r, this RCS R is equivalent (bisimilar) to an RCS where r is replaced with $\top$.

For the sake of simplicity, in this subsection, we'll explain the static elision of refinements in RCS using examples. Formal definitions, lemmas and their proofs are available in Appendix F.1. We use the RCS of G_s shown in Figure 5.

If we aim to i.e. transitions which payload modify variables that do not appear free in the refinement of the considered transition.

▶ **Example 32** (Independent transitions). In $\mathcal{S}(G_s)$, $\mathsf{A}_2 \xrightarrow{\mathsf{A!B}\langle \ell_2, \langle y, \_\rangle\rangle : x<10} \mathsf{A}_3$ depends on the variable $x \in \mathrm{fv}(x < 10)$. This transition is self-independent. Since the payload of $\mathsf{A}_1 \xrightarrow{\mathsf{A!B}\langle \ell_1, \langle x, \_\rangle\rangle : x<0} \mathsf{A}_2$ is x, the former transition depends on the later.

▶ Remark 33. We note $\mathbb{T}_x$ the set of transitions $\sigma \xrightarrow{\_\dagger\_\langle \_, \langle x, \_\rangle\rangle : \_} \sigma'$. Given a transition t with refinement r, if $x \in \mathrm{fv}(r)$, then t depends on all transitions of $\mathbb{T}_x$. ◁

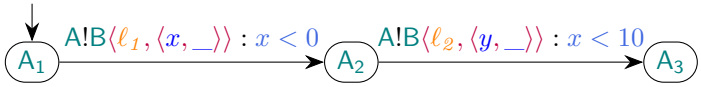

(a) RCFSM of A in the protocol G_s ($\mathcal{A}(G_s\upharpoonright_{\mathsf{A}})$).

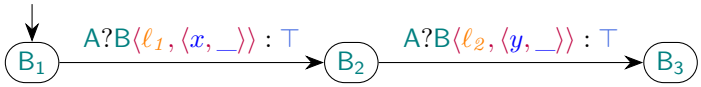

(b) RCFSM of B in the protocol G_s ($\mathcal{A}(G_s\upharpoonright_{\mathsf{B}})$).

Figure 5 RCFSM of the RCS of G_s, the running example of Section 6

Essentially, when attempting to remove a refinement from a target transition t, we can disregard all transitions t is independent of.

The second definition we will need is about transitions being *well-defined*. So far, nothing prevents us to use refinements with undefined free variables, we simply consider the refinement does not hold (c.f. Definition 14). In this section, we specifically focus on systems where free variables of refinements are in the map when the refinement is evaluated. When it is the case, we call transitions with such refinements *well-defined*.

▶ **Example 34** (Well-defined transition)**.** Considering the RCS in Figure 5. In the RCFSM of A, the (local) state A_2 is only accessible with a transition $\mathsf{A}_1 \xrightarrow{\mathsf{A!B}\langle \ell_1, \langle x, \_\rangle\rangle : x<0} \mathsf{A}_2$. Therefore, any global state $\langle\langle \mathsf{A}_2, \mathsf{B}_{\{1,2,3\}}\rangle, \_, M\rangle$ necessarily contains a preceding transition $\mathsf{A}_1 \xrightarrow{\mathsf{A!B}\langle \ell_1, \langle x, \_\rangle\rangle : x<0} \mathsf{A}_2$. Therefore, x is always in the map M of that state.

Therefore, the transition $\mathsf{A}_2 \xrightarrow{\mathsf{A!B}\langle \ell_2, \langle y, \_\rangle\rangle : x<10} \mathsf{A}_3$ is well-defined.

We can now conclude our analysis technique: consider a target transition t with refinement r that is self-independent (it does not modify the variables of its refinement) and well-define. If all transitions that modify the free variables of r can guarantee (via their refinement) that the modification they do is correct with respect to r, then we can safely remove r.

▶ **Theorem 35** (Correctness of refinement elision)**.** *Given an RCS R containing an RCFSM $M = \langle Q, C, q_0, \mathbb{A}, \delta\rangle$, and $t = s_{\mathsf{i}} \xrightarrow{\mathsf{p}\dagger\mathsf{q}\langle m\rangle : r} s_{\mathsf{i}}' \in \delta$, a well-defined self-independent transition. Let $t' = s_{\mathsf{i}} \xrightarrow{\mathsf{p}\dagger\mathsf{q}\langle m\rangle : \top} s_{\mathsf{i}}'$; $\delta' = \delta \setminus \{t\} \cup t'$; $M' = \langle Q, C, q_0, \mathbb{A}, \delta'\rangle$; and R' be R where M is replaced with M'. If, for each transition $t_w = \_ \xrightarrow{\_!\_\langle\_\rangle : r_w} \_$ in $\bigcup_{x\in \mathrm{fv}(r)} \mathbb{T}_x$, for all map M, $M \models r_w$ entails $M \models r$, then there exists a bisimulation relating the states of R' and R.* ◁

Proof. Proving each direction of the bisimulation is direct (see the proof in Appendix F.1). ◀

▶ **Example 36** (Application of Theorem 35)**.** The following RCFSM, where $x < 10$ is removed, is a valid replacement for $\mathcal{A}(G_s\upharpoonright_{\mathsf{A}})$ in $\mathcal{S}(G_s)$.

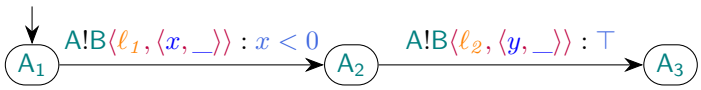

6.2 Application to RMPST Protocols

The above subsection explains how to remove some redundant refinements in RCS. In this subsection, we intend to do the same, focusing on RMPST instead of RCS.

Our goal is the following: we are given an RMPST G, and we would like to remove one of its refinement (which we call the *target* refinement r). For the sake of simplicity, in this

section, we assume all labels are uniquely used (which we use to prove Lemma 74). For the general case, we can simply uniquely rename redundant labels. Overall, the roadmap for this subsection is to show that given the type G', which is G where r is *replaced* by $\top$, G and G' behave similarly, i.e. the RCS the generate are bisimilar. To achieve this, we show that Theorem 35 applies to $\mathcal{S}(G)$ and $\mathcal{S}(G')$. Therefore, the main point is finding conditions on RMPST that ensures hypothesis of Theorem 35 holds; we have to verify the following items:

1. all transitions our refinement depends on should entail the refinement itself;
2. the transition that carries the refinement must be well-defined (Definition 63). Since variables cannot be removed from the map, the first occurrence of the target transition must respect the domain condition. Therefore, for this step, we can ignore recursion.

The main difference with automata is that, in types, we have *communications*, which possibly contains choices with multiple branches; and we our goal is to remove the refinement of one of those branches. Therefore, we first introduce *steps* of a communication, i.e. given a choice, what are the possible choices it can take. We then extend this to types. We show that steps in a type correspond to transitions in the automata of that type.

▶ **Example 37** (Step). The type $G_y = \mathsf{A} \to \mathsf{B}\,\{\ell_2(y : \mathtt{int} \models x < 10).\mathtt{end}\}$ has the step $\mathsf{A} \to \mathsf{B}\langle \ell_2, y\rangle \models x < 10$. Since G_y occurs in one of the branches of G_s (from the introduction of this section), this step *occurs* in G_s.

Given this notion of steps occurring in a type that is analogous to transitions in the RCFSM of that type, we can now focus on the conditions of Theorem 35. Therefore, we have to characterise what corresponds to *well-defined transitions* in a type. Since transitions (in RCS) and steps (in types) are analogous, we introduce *well-define steps* in a type. We recall that, in a RCS, a transition is well-defined if the free variables of the refinement it carries are always known when the transition is fired. Since variables are never removed from the map, we can focus on the first occurrence of the transition. So far, we do not have a notion of *run* for a type. Therefore, we first define an *happens-before* relation in RMPST, and we use this relation to define *well-defined steps* as steps that contain a refinement which free variable are all exchanged in a communication that *happens-before* the step we consider. With those two definitions, we can finally prove that a well-defined step in a type corresponds to a well-defined transition in the corresponding RCS.

▶ **Example 38** (Well-define step in a type). Consider G_s and G_y as in Example 37. The step $\mathsf{A} \to \mathsf{B}\langle \ell_2, y\rangle \models x < 10$ is well-defined. Indeed, $\mathsf{fv}(x < 10) = \{x\}$, $G_s < G_y$, and G_s contains a branch that sends x and which continuation contains G_y.

We can finally proceed to the overall goal of this section: showing that the type with and without the target refinement behave similarly. Thanks to the above lemmata, we simply have to target a refinement with the appropriate conditions and apply Theorem 35.

▶ **Theorem 39** (Static elision of redundant refinements in types). *Given two a global types G and $G_s = \mathsf{p} \to \mathsf{q}\{\ell_i(x_i : \mathsf{S}_i \models r_i).G_i\}_{i\in I} \in G$, such that, for one $t \in I$, $\mathsf{p} \to \mathsf{q}\langle \ell_t, x_t\rangle \models r_t$ is a well-defined step with $x_t \notin \mathsf{fv}(r_t)$. Let $\ell_{t'} = \ell_t$, $x_{t'} = x_t$, $\mathsf{S}_{t'} = \mathsf{S}_t$, $r'_t = \top$, $G_{t'} = G_t$, $G_{s'} = \mathsf{p} \to \mathsf{q}\{\ell_i(x_i : \mathsf{S}_i \models r_i).G_i\}_{i\in I\setminus\{t\}\cup\{t'\}}$; and G' be G where G_s is replaced with $G_{s'}$. If, for all steps, $\mathsf{r} \to \mathsf{s}\langle \_, x_w\rangle \models r_w$ occurring in G (for each $x \in \mathsf{fv}(r)$), $M \models r_w$ entails $M \models r$ (for all M), there exists a bisimulation between the states of $\mathcal{S}(G)$ and those of $\mathcal{S}(G')$.* ◁

Proof. We prove this by showing that Theorem 35 applies to $\mathcal{S}(G)$ and $\mathcal{S}(G')$. The proof is provided Appendix F.2. ◀

```
(*# RefinementTypes #*)

global protocol PlusMinus
(role A, role B, role C)
{
Secret(n: int) from A to B;
rec Loop {
  Guess(x: int) from C to B;
  choice at B {
    More(x: int {x < n}) from B to C;
    continue Loop;
  } or {
    Less(x: int {x > n}) from B to C;
    continue Loop;
  } or {
    Correct(x: int {x = n}) from B to C;
  }}}
```

(a) νScr description of the guessing game protocol.

```
type PlusMinusA =
Send<B, 'n',
Secret,
Tautology::<Name, Value, Secret>,
Constant<Name, Value>, End>;
// ...
async fn a(role: &mut A)
-> Result<(), Box<dyn Error>> {
  try_session(role,
  HashMap::new(),
  |s: PlusMinusA<'_, _>| async {
    let s =
    s.send(Secret(10)).await?;
    return Ok(((), s))
  })
  .await
}
```

(b) Rust type and implementation of participant A of the guessing game protocol. The handwritten code (Line 7 to Line 17) is the same than with Vanilla Rumpsteak.

Figure 6 Implementation of the guessing game using Rumpsteak.

▶ **Example 40** (Application of Theorem 39)**.** Given G_s as in Example 37 and G'_s as follows (notice the second refinement is replaced by $\top$), G_s and the following G'_s have the same behaviour:

$$G'_s = \mathsf{A} \to \mathsf{B}\,\{\ell_1(x : \mathtt{int} \models x < 0).\mathsf{A} \to \mathsf{B}\,\{\ell_2(y : \mathtt{int} \models \top).\mathtt{end}\}\}$$

7 Implementation

In the previous section, we introduced an instance of our framework: a system that accommodates refinements using a decentralised verification mechanism. In this section, we follow up on this example with an implementation, based on Rumpsteak, of this system.

Rumpsteak [7] is a framework to write Rust programs according to an MPST specification. The framework is divided into two parts: (i) a runtime library that provides primitives to write asynchronous programs in Rust; and (ii) a tool (`rumpsteak-generate`) to generate skeleton Rust files from specification files (i.e. from global types), in two steps.

Working with Rumpsteak takes two manual steps. The user specifies (step 1) the protocol in a global type(written as Scribble files [39], see Figure 6a). This global type is automatically projected using νScr [15] and the projected types are used to generate skeleton Rust files (see Figure 6b). The generated Rust code contains Rust types that encode local types (e.g. the type for A is shown in Line 1 in Figure 6b). The user then manually implements (step 2) the process of each participant, following their type (Line 7), using provided communication primitives (Line 13). Rumpsteak relies on Rust's typechecker to ensure the consistency of the implementation. For the sake of clarification where needed, we call *Vanilla Rumpsteak* the framework without refinements (i.e. as presented in [7]), and *Refined Rumpsteak* the framework modified to accommodate refinements.

In this section, we explain the main differences between Vanilla and Refined Rumpsteak: we introduce refinements in the types used in the runtime library, we modify the program generation step accordingly, and we introduce tools that ensure the localisation conditions

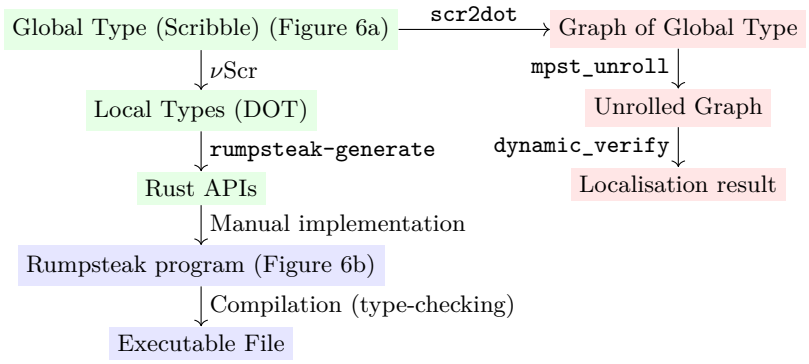

Figure 7 Workflow of Rumpsteak. Green nodes represent steps that already existed in Vanilla Rumpsteak and that have been adapted to accommodate for refinements, red nodes represent new steps, and blue nodes represent unmodified steps. The three new steps (`scr2dot`, `mpst_unroll`, and `dynamic_verify`) verify the conditions mentioned in Definition 29.

are met (Definition 29 in Section 5). The overall workflow is presented in Figure 7. We conclude this section by measuring the overhead induced by the refinement w.r.t. Vanilla Rumpsteak and the time needed for asserting the localisation conditions.

7.1 Refinement Implementation

Modifications to the Rumpsteak Library. In order to accommodate for refinements, we have to introduce new elements in to the Rumpsteak's encoding of local types. Consider the local type of participant A introduced in Example 21 $\mathsf{B}\oplus\{\mathit{secret}(n : \texttt{int} \models \top).\texttt{end}\}$: Rumpsteak now has to take into account the name of the variable sent (n), and the refinement attached to the transition ($\top$). Consider the type declaration in Line 1 to Line 5, Figure 6b. Compared to Vanilla Rumpsteak, we introduce `'n'`, a const generic[5], that carries the name of the variable sent (Line 4). Regarding the refinement, we introduce `Tautology::<Name, Value, Secret>`, which represent the refinement $\top$. The generic parameters are used to specify the type of variable names (`char`s in our case) and values (`i32`) as well as the label of the message (`Secret`). We modified νScr and `rumpsteak-generate` to generate skeleton files (the content of the file up to Line 5). Rumpsteak provides a set of available refinements, and additional ones can be written ad-hoc (for specific needs). To add an ad-hoc refinement, the user simply implements the trait `Predicate` (which extends `Default`), which requires a method `check` that asserts whether the predicate holds. For instance, the `check` function of `Tautology` simply returns `true`.

Verification of the Conditions for Decentralised Refinement Assertion. As we explained in Section 5, to make sure that refinements can be verified in a decentralised way, we require to check that variables needed for the refinements are located correctly (Definition 29). To perform this verification, we implemented new tools for the Rumpsteak framework (in red in Figure 7).

Our tools: (i) convert the global type into a graph (`scr2dot`); (ii) unroll the loops once to precisely capture variables initialisations (`unroll_mpst`); and (iii) localise variables on the unrolled graph (`dynamic_verify`).

[5] https://github.com/rust-lang/rfcs/blob/master/text/2000-const-generics.md

The core part of this verification, `dynamic_verify`, finds variables locations with simple inference rules written in Datalog. We use the *crepe* library [40] which provides a Datalog DSL for Rust. We provide more details on the algorithm in Appendix E.4.

Limitations. The current implementation makes extensive use of the Rust feature *const generics*[9] which introduces a limited form of dependent types in Rust. It allows to use constant values in types. As of today, only some *basic types* can be used as const generics, in particular `char`s and the various integer types. We use such const generics to encode informations about the variables into the types: for instance, the predicate $x < 5$ would have the type `LTnConst<L, 'x', 5>`, where the `'x'` and the `5` are const generics.

For readability, we choose to set variables to `char`s, meaning that in the current implementation, we can only accommodate a limited number of distinct variables. Should more be needed, one could easily modify our implementation to replace them with `u64`, which allows 2^{64} variables names. Similarly, we only consider `i32` as message payloads. Should different types of messages be needed, they could be encoded in an `enum`.

Finally, the static elision optimisation (Section 6) is not implemented.

7.2 Runtime and Localisation Benchmarks

We evaluate how Rumpsteak with refinements performs with respect to Rumpsteak without refinements. First, we measure the runtime of our analysis tool which verifies the two conditions in Definition 29 (`scr2dot`, `unroll_mpst` and `dynamic_verify`). Although not a runtime cost, and while we expect this analysis to be possibly expensive, we would like to ensure that it is still practical for test cases from the literature. Secondly, we evaluate the runtime overhead of adding refinements with respect to Rumpsteak without refinements.

Setup and Benchmark Programs. We evaluate the performance of Rumpsteak with refinement with benchmarks. Most of them are taken from the literature (Table 1). This set of program contains various micro-benchmarks with a variety of combination of properties (whether the protocol is binary or multiparty, contains recursivity or choice). Notice that protocols that contain recursivity with no choice (e.g. *simple auth* are infinite). Therefore, such protocols are only measured in the variable localisation paragraph. Also, where it applies, protocols were modified in order to add relevant refinements; such modifications are listed below. By default, we add `Tautology` predicates (Section 7.1). The tests were performed on a machine running Ubuntu 22.04.1 LTS x86_64 (kernel 5.15.0-60) with an Intel i7-6700 processor (4 cores, 8 threads running at 4GHz maximum) and 16GB of memory[6]. We compare Rumpsteak with refinement vs. Vanilla Rumpsteak. For a comparison between Vanilla Rumpsteak and other libraries, see [7, Figure 6].

Name	MP	Rec	Choice
① simple adder [21]	no	no	no
② travel_agency [23]	no	no	yes
③ ping pong [42]	no	yes	no
④ simple auth.	no	yes	yes
⑤ ring max	yes	no	no
⑥ three_buyers [19]	yes	no	yes
⑦ plus or minus	yes	yes	yes

Table 1 The set of micro benchmarks together with their characteristics. "MP" denotes a multiparty protocol, "Rec" the presence of recursion, and "Choice" the presence of choice.

Added Refinements & Protocol Modifications. Some benchmarks from the literature were adapted in order to accommodate refinements. In addition, we introduce three

[6] The micro-benchmarks are not memory intensive. The memory size is not a limiting factor. However, the benchmarks seem to be dominated by the startup time, which includes memory access time.

benchmarks. Those benchmarks are close to examples from the literature, adapted to better highlight refinements.

simple adder: This example is adjusted from the *Adder* ([21]) protocol, but we remove the choice of operation in order to increase the benchmark diversity;

ping pong: In [42], some of the loops were statically unrolled, and the protocol contained a choice to exit. Ours is equivalent to an infinite $PingPong_1$ in [42].

simple authentication: This example is a binary example of an authentication protocol (e.g. OAuth [31]). The added refinements enforce that access is granted if and only if the given password is correct.

ring max: A multiparty protocol where participants receive a value from their predecessor (except for the initial participant), and forward an other value to their successor (the final participant forwards it to the initial one). Refinements ensure that the value forwarded is greater than or equal to the value received.

plus or minus: An implementation of our running example.

	$\|S\|$	$\|U\|$	$\|V\|$	et $(\mu \pm \sigma)$
①	4	4	3	5.5 ± 0.2
②	7	7	6	5.5 ± 0.2
③	2	4	1	5.5 ± 0.2
④	6	11	3	5.6 ± 0.2
⑤	8	8	7	5.7 ± 0.2
⑥	10	10	7	5.6 ± 0.2
⑦	4	19	2	5.6 ± 0.2

Table 2 Benchmark of the localisation analysis (Red branch in Figure 7). $|S|$ denotes the number of states of the graph of the protocol; $|U|$ denotes the number of states after unrolling the recursion loops once; and $|V|$ denotes the number of variables in the protocol. $|S|$, $|U|$ and $|V|$ are computed manually to give an insight on how protocols compare. et is the execution time, measured by the benchmark (in ms).

	p	m	m_r
①	0.00	0.7	0.8
②	0.11	0.8	N/A
④	0.29	0.8	N/A
⑤	0.17	0.8	N/A
⑥	0.68	0.7	N/A
⑦	0.04	0.7	0.8

Table 3 Evaluation of the runtime overhead due to the addition of refinements in Rumpsteak. p is the MWU p-value, m is the baseline median runtime and m_r is the median runtime with refinements when applicable ($p < 0.05$). All times are in ms.

the following:

Static Analysis of Variable Locations. Table 2 shows the decentralised verification time cost for each refined global label. As shown in Figure 7, this static analysis is performed with three tools. The results shown account for the whole pipeline, and were measured over 50 samples, with 10 warmup runs (excluded from the measurements). Overall, the runtime for variable localisation is stable (around 5.6ms). We suspect that, for graphs with a low number of states, the runtime is dominated by the accesses to the file.

Runtime Overhead of Refinement Feature. Our second set of benchmarks aims to measure the overhead of runtime refinement verification with respect to the original Rumpsteak framework. We are expecting Rumpsteak with refinements to be slower than the original Rumpsteak, due to the additional cost of evaluating refinements. This benchmark has two objectives: first, to find out whether there is an actual, statistically significant, overhead; and second, if so, estimate this overhead. To measure this overhead, we only consider the protocols that terminate from the benchmark set.

To fulfil the first objective, we use a Mann-Whitney U test (MWU). We used MWU as it is a non-parametric test, and our runtime distributions do not follow a normal distribution, which prevents us to do simpler analysis. As MWU is sensitive to the number of samples, we run each benchmark 30 times, on both the original Rumpsteak and Rumpsteak with refinements. We perform the MWU test on the collected 30 samples, preceded by 10 iterations to warm the system up. Our hypothesis for the MWU test are

H_0: The distributions of runtimes with and without refinements are identical.
H_1: The distributions of runtimes with and without refinements are distincts.

The p-values obtained from the MWU test are reported in the first column of Table 3. We also report the baseline (Rumpsteak without refinements) median run time (over the 30 runs) in the second column of the table. Most often, the overhead is not significant ($p \geq 0.05$) and H_0 can not be rejected. When the overhead is statistically significant, we also report the median runtime (over the 30 runs) of Rumpsteak with refinements in the third column. With our set of microbenchmarks, in most cases we cannot distinguish Rumpsteak with refinement from Rumpsteak without refinements. We suspect Rumpsteak runtime is dominated by communications and context switching. However, as our refinements can be arbitrarily complex, specific instances could show real slowdown due to refinement evaluation.

8 Related Work and Conclusion

Design-by-Contract for (Multiparty) Session Types. In binary session types, [37] introduces contracts for binary sessions, and provides an analysis tool which verifies whether a given program comply with its associated contract. The verification is done with symbolic execution. Compared to this paper, we address multiparty sessions. Besides, our framework is more generic (specific instances could be based on symbolic execution, but we can also accommodate other verification methods). Bocchi et al. [4] present a variant of MPST that allows predicates on exchanges, that must hold for a typed process to take transitions. The main difference with our work is that their approach focuses on *correctness by construction*, i.e. they accept only correct protocols, while we can accept protocols that fail, and we simply prevent them to generate incorrect traces. More precisely, the authors statically ensure that there is a satisfiable path, which prevents some valid runs to be accepted. For instance, consider the following type:

$$\mathsf{A} \to \mathsf{B}\,\{\ell_1(x : \mathtt{int} \models x < 10).\mathsf{B} \to \mathsf{A}\,\{\ell_2(y : \mathtt{int} \models x > y \wedge y > 6).\mathtt{end}\}\}$$

This type would be rejected in [4] since if A sends $x = 5$ (which is allowed by $x < 10$), then there is no y that satisfies $5 > y \wedge y > 6$. By rejecting this, they also reject all possibly valid runs (e.g. if A sends $x = 9$ and B replies with $y = 7$). A follow-up on this work is [3] which introduces local states, i.e. the authors allow participants to have local variables, which can be updated during process execution. The session types reflect those elements and contain predicates on exchanged variables and local variables.

With respect to these two papers, our criteria for the validity of refinements (expressed as a property of the generated trace) is decoupled from the semantics of the model. This approach allows us to be more flexible than enforcing statically the refinements, and to lower the cost of adopting refinements, in particular to retrofit refinements into existing systems. For instance, using our framework, one can simply use the centralised semantics at first, which is very expressive, without having to prove the correctness of the implementation. In a second step, users can then develop different verification or analysis techniques which can be plugged-in transparently. For instance, switching from Vanilla Rumpsteak to Refined Rumpsteak does not involve changes in the implementation, as the modifications do not happen in the programming interface. Also, compared to these papers, our framework is not bound to MPST only, and provide an actual implementation of our framework.

Design-by-Contract in Choreography Automata. Choreography Automata (CA) are graphs that represent the global behaviour of a concurrent system. The behaviour of individual participants is obtained by *projecting* well-formed CA, i.e. erasing all actions that

do not concern a given participant. The result is a FSM which, after determinising and minimising, is used as a CFSM. The projection of all participants leads to a CS. Notice that CA accept some protocols that would be rejected by MPST, and vice-versa.

Gheri et al. [16] study the verification of CA with assertions. Their work and ours are distinct with respect to the following aspects: (i) the communication semantics; (ii) the choices; (iii) the logic for predicates; and (iv) the implementation presented in [16] is limited to CA without assertions (i.e., the design-by-contract approach was not implemented and left as their future work).

Regarding Item (i), Gheri et al. [16] defines choreography automata with *synchronous* communication semantics, while the one we developed in this work is asynchronous. Gheri et al. [16, Section 7] discusses asynchronous semantics but it remains future works.

Regarding Item (ii), we are constrained by the syntax of RMPST, in which choices can only happen between two selected participants, while choreography automata accept protocols with choices where a (single) participant A sends to multiple receivers (B and C) [16, Definition 4.15]. Explicit connections [22] is an extension of MPST that accommodates with choices with multiple receivers.

Regarding Item (iii), we kept our refinement logic abstract, while it is fixed in choreography automata, with a form of first order logic. Besides, predicates are handled differently in both frameworks as well: Gheri et al. [16] require choreography automata to be *history-sensitive* [4], a definition which serves a similar purpose to our definition of *variable localisation* (Section 5 and Appendix E.4.2), which constrains our decentralised semantics. Our centralised semantics (Definition 10) is not constrained by variable localisation. For instance, the RMPST $\mathsf{A} \to \mathsf{B}\,\{\ell_1(x : \mathtt{int} \models \top).\mathsf{C} \to \mathsf{D}\,\{\ell_2(y : \mathtt{int} \models x = y).\mathtt{end}\}\}$ produces valid traces with our centralised semantics, while the corresponding choreography automata would be rejected.

Besides, our work introduces a general framework that can accommodate refined CA in addition to RMPST. We show Appendix G a possible way to do so.

Implementations of Refinements in MPST. Neykova et al. [29] develop an F# library for static verification of MPST with refinements. They present a compiler plugin which uses an SMT solver (Z3) to statically verify some refinements. They use a notion of similar to our variable localisation criterion (which they call *variable knowledge*), and a variant of CFSM with refinements that is similar to ours. In their work, refinements that are statically asserted by the SMT solver are pruned in the CFSM, while the rest of refinements are kept in the CFSM and are dynamically checked. Similarly, [41, 42] develop a framework for multiparty session types with refinements in F*. They delegate the management of refinements to F* type system (which internally uses an SMT solver). They define refinements on global types, which are then projected onto local types. They show that a global type and its projection are trace equivalent. Those two works focus on the *implementation* of MPST with refinements. [29] does not focus on the theory of refinements and the theory developed in [42] is tightly coupled to F*. For instance, they do not present a *correctness* criterion such as *valid refined traces* we present. Contrary to both works, our correctness criteria (based on valid refined traces) is *decoupled from* (i.e. independent of) any target type theory, programming language or model of computation: we only require an LTS labelled with actions. Besides, the logic used for refinements is also a parameter of our framework, and users could use alternatives, leading to a greater expressivity of our framework.

The main syntactical difference between our RMPST and those developed in [42] is that we attach refinements to the messages of the protocol, while [42] attach refinements to the payload value. This is due to a different approach: correctness in [42] is related to payload types being inhabited while our criteria of correctness (developed in Definition 7) relies on

actions being allowed. In binary linear logic-based session types, [9] study the metatheory of binary session types with arithmetic refinements. In particular, they focus on the type equality, showing that added refinements make the type equality undecidable (they provide a sound but incomplete algorithm for type equality). [10] also implement a library for session types with refinements, although it only accounts for arithmetic refinements.

Other Related Works. There are various papers on the dynamic verification of MPST. For instance [2] present a framework that allows for both static and dynamic verification of MPST. This paper introduces a theory for (dynamically) monitoring assertions on messages (i.e. the equivalent of our refinements). Furthermore, the authors introduce theoretical tools (bisimulations) to relate monitored processes with correct unmonitored processes. This paper, however, suffers a few limitations. First, it focuses on *monitorable* types (which intuitively correspond to types satisfying our conditions for decentralised verification Definition 29). Second, it focuses on dynamic verification of assertions. The paper is compatible with statically verified processes (which allows turning off the dynamic monitoring), but it does not present techniques for static verification in itself.

On the other hand, our paper takes a different approach, by decoupling the correctness criterion from the verification technique. This allows us to have a more general framework (our framework accept types that are not localisable/monitorable, although not all semantics can accommodate those), as well as to develop static verification techniques.

In Rust, the `refinement` crate [11] provides refinement data types. Their approach of refinements is similar to ours, with a `Predicate` trait that provides a method to perform the predicate verification (at runtime). Refinement data types have also been implemented in multiple languages (e.g. F*, Haskell [36], etc.). On the practical side, we can note the similarities between typestates and session types [20]. [14] implements typestates in Rust with a DSL to verify protocol conformance. While Rumpsteak does not use their library, it internally uses similar constructs.

Regarding implementations of session types in Rust, there are several frameworks beside Rumpsteak. [25] first integrate binary session types in Rust, but their implementation suffers a few drawbacks (see [26, Section 3] for a detailed explanation). Sesh [26] and Ferrite [6] are two Rust libraries for *binary* session types, and they implement synchronous and asynchronous ones, respectively. MultiCrusty [27] implements synchronous MPST on top of Sesh, with a mesh of binary sessions. Compared to MultiCrusty, Rumpsteak implements directly MPST instead of wrapping them into binary sessions, and focuses on asynchronous MPST. None of the aforementioned tools develops refinements. It would be an interesting future work to apply our criteria to extend their tools with refinements.

Finally, we note the proximity between (MP)ST with refinements and dependent (MP)ST. For instance, [33] introduce a session type calculus with label-dependency (their approach does not explicitly account for payload value refinement). Other approaches exist, for instance, an intuitionistic linear logic-based type theory for building value-dependent session types [34], and separation logic-based work for reasoning about session types [17].

Future Work While, in our work, we consider MPST with payloads (some variants only consider messages with labels), we restrict our MPST with a single payload (i.e. *monadic* MPST, where each message carries a single value). The extension to polyadic MPST, where a message can carry multiple values, is straightforward, by adapting the RCS rules (GRSnd and GRRec, Definition 14).

We presented two optimisations, in order to illustrate the flexibility of our theoretical framework. Regarding the decentralised verification (Section 5), there is room for an extension, e.g. with specific domains (i.e. some class of protocols with specific refinements). Regarding

the static elision of redundant refinements, we envision improving the technique with use of SMT solvers could be promising. The main difficulty lies in asynchronous communications: one would need to consider all possible message orderings before solving constraints.

References

1 Franco Barbanera, Ivan Lanese, and Emilio Tuosto. Choreography automata. In Simon Bliudze and Laura Bocchi, editors, *Coordination Models and Languages - 22nd IFIP WG 6.1 International Conference, COORDINATION 2020, Held as Part of the 15th International Federated Conference on Distributed Computing Techniques, DisCoTec 2020, Valletta, Malta, June 15-19, 2020, Proceedings*, volume 12134 of *Lecture Notes in Computer Science*, pages 86–106. Springer, 2020. doi:10.1007/978-3-030-50029-0_6.

2 Laura Bocchi, Tzu-Chun Chen, Romain Demangeon, Kohei Honda, and Nobuko Yoshida. Monitoring networks through multiparty session types. *Theoretical Computer Science*, 669:33–58, 2017. URL: https://www.sciencedirect.com/science/article/pii/S0304397517301263, doi:10.1016/j.tcs.2017.02.009.

3 Laura Bocchi, Romain Demangeon, and Nobuko Yoshida. A Multiparty Multi-Session Logic. In *7th International Symposium on Trustworthy Global Computing*, volume 8191 of *LNCS*, pages 111–97. Springer, 2012.

4 Laura Bocchi, Kohei Honda, Emilio Tuosto, and Nobuko Yoshida. A Theory of Design-by-Contract for Distributed Multiparty Interactions. In Paul Gastin and François Laroussinie, editors, *CONCUR 2010 - Concurrency Theory*, Lecture Notes in Computer Science, pages 162–176, Berlin, Heidelberg, 2010. Springer. doi:10.1007/978-3-642-15375-4_12.

5 Daniel Brand and Pitro Zafiropulo. On Communicating Finite-State Machines. *Journal of the ACM*, 30(2):323–342, April 1983. doi:10.1145/322374.322380.

6 Ruofei Chen and Stephanie Balzer. Ferrite: A Judgmental Embedding of Session Types in Rust, 2021. (repository is found at https://github.com/ferrite-rs/ferrite). arXiv:2009.13619.

7 Zak Cutner, Nobuko Yoshida, and Martin Vassor. Deadlock-free asynchronous message reordering in rust with multiparty session types. In *Proceedings of the 27th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, PPoPP '22, pages 246–261, New York, NY, USA, April 2022. Association for Computing Machinery. doi:10.1145/3503221.3508404.

8 Gérard Cécé and Alain Finkel. Verification of programs with half-duplex communication. *Information and Computation*, 202(2):166–190, November 2005. URL: https://www.sciencedirect.com/science/article/pii/S0890540105001082, doi:10.1016/j.ic.2005.05.006.

9 Ankush Das and Frank Pfenning. Session Types with Arithmetic Refinements. In Igor Konnov and Laura Kovács, editors, *31st International Conference on Concurrency Theory (CONCUR 2020)*, volume 171 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 13:1–13:18, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.CONCUR.2020.13.

10 Ankush Das and Frank Pfenning. Rast: A Language for Resource-Aware Session Types. *Logical Methods in Computer Science*, Volume 18, Issue 1, January 2022. URL: https://lmcs.episciences.org/8954, doi:10.46298/lmcs-18(1:9)2022.

11 Brady Dean and Joey Ezechiëls. refinement crate, 2021. (repository is found at https://github.com/2bdkid/refinement).

12 Pierre-Malo Deniélou and Nobuko Yoshida. Multiparty Session Types Meet Communicating Automata. In *21st European Symposium on Programming*, volume 7211 of *LNCS*, pages 194–213. Springer, 2012.

13 Pierre-Malo Deniélou and Nobuko Yoshida. Multiparty Compatibility in Communicating Automata: Characterisation and Synthesis of Global Session Types. In *40th International*

Colloquium on Automata, Languages and Programming, volume 7966 of *LNCS*, pages 174–186, Berlin, Heidelberg, 2013. Springer. doi:10.1007/978-3-642-39212-2_18.

14 José Duarte and António Ravara. Retrofitting Typestates into Rust. In *25th Brazilian Symposium on Programming Languages*, pages 83–91, Joinville Brazil, September 2021. ACM. URL: https://dl.acm.org/doi/10.1145/3475061.3475082, doi:10.1145/3475061.3475082.

15 Francisco Ferreira, Fangyi Zhou, Simon Castellan, and Benito Echarren. NuScr, 2019. URL: https://github.com/nuscr/nuscr.

16 Lorenzo Gheri, Ivan Lanese, Neil Sayers, Emilio Tuosto, and Nobuko Yoshida. Design-By-Contract for Flexible Multiparty Session Protocols. In Karim Ali and Jan Vitek, editors, *36th European Conference on Object-Oriented Programming (ECOOP 2022)*, volume 222 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 8:1–8:28, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. URL: https://drops.dagstuhl.de/opus/volltexte/2022/16236, doi:10.4230/LIPIcs.ECOOP.2022.8.

17 Jonas Kastberg Hinrichsen, Jesper Bengtson, and Robbert Krebbers. Actris: Session-type based reasoning in separation logic. *Proceedings of the ACM on Programming Languages*, 4(POPL):1–30, January 2020. doi:10.1145/3371074.

18 Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty asynchronous session types. *ACM SIGPLAN Notices*, 43(1):273–284, January 2008. URL: https://dl.acm.org/doi/10.1145/1328897.1328472, doi:10.1145/1328897.1328472.

19 Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty Asynchronous Session Types. *Journal of the ACM*, 63(1):9:1–9:67, March 2016. doi:10.1145/2827695.

20 Raymond Hu, Dimitrios Kouzapas, Olivier Pernet, Nobuko Yoshida, and Kohei Honda. Type-safe eventful sessions in java. In *Proceedings of the 24th European conference on Object-oriented programming*, ECOOP'10, pages 329–353, Berlin, Heidelberg, June 2010. Springer-Verlag.

21 Raymond Hu and Nobuko Yoshida. Hybrid Session Verification Through Endpoint API Generation. In Perdita Stevens and Andrzej Wąsowski, editors, *Fundamental Approaches to Software Engineering*, Lecture Notes in Computer Science, pages 401–418, Berlin, Heidelberg, 2016. Springer. doi:10.1007/978-3-662-49665-7_24.

22 Raymond Hu and Nobuko Yoshida. Explicit Connection Actions in Multiparty Session Types. In Marieke Huisman and Julia Rubin, editors, *Fundamental Approaches to Software Engineering*, Lecture Notes in Computer Science, pages 116–133, Berlin, Heidelberg, 2017. Springer. doi:10.1007/978-3-662-54494-5_7.

23 Raymond Hu, Nobuko Yoshida, and Kohei Honda. Session-Based Distributed Programming in Java. In Jan Vitek, editor, *ECOOP 2008 – Object-Oriented Programming*, Lecture Notes in Computer Science, pages 516–541, Berlin, Heidelberg, 2008. Springer. doi:10.1007/978-3-540-70592-5_22.

24 International Telecommunication Union. Z.120 : Message Sequence Chart (MSC), February 2011.

25 Thomas Bracht Laumann Jespersen, Philip Munksgaard, and Ken Friis Larsen. Session types for Rust. In *Proceedings of the 11th ACM SIGPLAN Workshop on Generic Programming*, pages 13–22, Vancouver BC Canada, August 2015. ACM. doi:10.1145/2808098.2808100.

26 Wen Kokke. Rusty Variation: Deadlock-free Sessions with Failure in Rust. *Electronic Proceedings in Theoretical Computer Science*, 304:48–60, 2019. (repository is found at https://github.com/wenkokke/sesh). doi:10.4204/eptcs.304.4.

27 Nicolas Lagaillardie, Rumyana Neykova, and Nobuko Yoshida. Stay Safe Under Panic: Affine Rust Programming with Multiparty Session Types. In Karim Ali and Jan Vitek, editors, *36th European Conference on Object-Oriented Programming (ECOOP 2022)*, volume 222 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 4:1–4:29, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. ISSN: 1868-8969. URL: https://drops.dagstuhl.de/opus/volltexte/2022/16232, doi:10.4230/LIPIcs.ECOOP.2022.4.

28 Bertrand Meyer. Design by Contract. *Advances in Object-Oriented Software Engineering*, pages 1–35, 1991.

29 Rumyana Neykova, Raymond Hu, Nobuko Yoshida, and Fahd Abdeljallal. A session type provider: compile-time API generation of distributed protocols with refinements in F#. In *Proceedings of the 27th International Conference on Compiler Construction*, CC 2018, pages 128–138, New York, NY, USA, February 2018. Association for Computing Machinery. URL: https://dl.acm.org/doi/10.1145/3178372.3179495, doi:10.1145/3178372.3179495.

30 Davide Sangiorgi. *An Introduction to Bisimulation and Coinduction*. Cambridge University Press, Cambridge ; New York, 2012.

31 Alceste Scalas and Nobuko Yoshida. Less is more: multiparty session types revisited. *Proceedings of the ACM on Programming Languages*, 3(POPL):30:1–30:29, January 2019. doi:10.1145/3290343.

32 Felix Stutz. Asynchronous Multiparty Session Type Implementability is Decidable - Lessons Learned from Message Sequence Charts. In *DROPS-IDN/v2/Document/10.4230/LIPIcs.ECOOP.2023.32*. Schloss-Dagstuhl - Leibniz Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.ECOOP.2023.32.

33 Peter Thiemann and Vasco T. Vasconcelos. Label-dependent session types. *Proceedings of the ACM on Programming Languages*, 4(POPL):1–29, January 2020. doi:10.1145/3371135.

34 Bernardo Toninho, Luís Caires, and Frank Pfenning. Dependent session types via intuitionistic linear type theory. In *Proceedings of the 13th International ACM SIGPLAN Symposium on Principles and Practices of Declarative Programming*, pages 161–172, Odense Denmark, July 2011. ACM. doi:10.1145/2003476.2003499.

35 Martin Vassor and Nobuko Yoshida. Refinements for multiparty message-passing protocols: Specification-agnostic theory and implementation, 2024. Full version on Arxiv.

36 Niki Vazou. *Liquid Haskell: Haskell as a Theorem Prover*. PhD thesis, University of California, San Diego, USA, 2016. URL: http://www.escholarship.org/uc/item/8dm057ws.

37 Jules Villard. *Heaps and Hops*. PhD thesis, Laboratoire Spécification et Vérification, École Normale Supérieure de Cachan, France, February 2011.

38 Nobuko Yoshida and Lorenzo Gheri. A Very Gentle Introduction to Multiparty Session Types. In Dang Van Hung and Meenakshi D´Souza, editors, *Distributed Computing and Internet Technology*, Lecture Notes in Computer Science, pages 73–93, Cham, 2020. Springer International Publishing. doi:10.1007/978-3-030-36987-3_5.

39 Nobuko Yoshida, Raymond Hu, Rumyana Neykova, and Nicholas Ng. The scribble protocol language. In Martín Abadi and Alberto Lluch Lafuente, editors, *Trustworthy Global Computing*, pages 22–41, Cham, 2014. Springer International Publishing.

40 Erik Zhang. Crepe, 2022. URL: https://crates.io/crates/crepe.

41 Fangyi Zhou, Francisco Ferreira, Raymond Hu, Rumyana Neykova, and Nobuko Yoshida. Statically Verified Refinements for Multiparty Protocols. *Proc. ACM Program. Lang.*, 4(OOPSLA), November 2020. doi:10.1145/3428216.

42 Fangyi Zhou, Francisco Ferreira, Raymond Hu, Rumyana Neykova, and Nobuko Yoshida. Statically Verified Refinements for Multiparty Protocols. *arXiv:2009.06541 [cs]*, September 2020. arXiv: 2009.06541. URL: http://arxiv.org/abs/2009.06541.

A Extra examples illustrating interesting aspects of RMPST.

A.1 Usefulness of multiparty

Our guessing game running example is somewhat simple and could be implemented using two successive binary sessions. Here, we present a slightly more complex example that uses RMPST moore intensively. This example is based on the guessing game, with an extra communication (a final official guess) y from C to A at the end[7].

$$\mathsf{A} \to \mathsf{B} \left\{ \begin{array}{l} \mathit{secret}(n : \mathtt{int}).\mu \mathbf{T}.\mathsf{C} \to \mathsf{B} \left\{ \begin{array}{l} \mathit{guess}(x : \mathtt{int}).\mathsf{B} \to \mathsf{C} \left\{ \begin{array}{l} \mathit{more}(\models x < n).\mathbf{T}, \\ \mathit{less}(\models x > n).\mathbf{T}, \\ \mathit{correct}(\models x = n).G_{\text{cont}} \end{array} \right\} \end{array} \right\} \end{array} \right\}$$

with

$$G_{\text{cont}} = \mathsf{C} \to \mathsf{A} \{\mathit{validate}(y : \mathtt{int} \models y = x).\mathtt{end}\}$$

A.2 List Adder

This example is more involved: A sends a list of numbers to B, and eventually B returns the sum of all numbers. We want to have refinements to make sure the returned value is correct. To achieve that, upon each number received, B computes the partial sum. This example requires extra participants in order to update and store the partial sum.

$$\mu \mathbf{T}.\mathsf{A} \to \mathsf{B} \left\{ \begin{array}{l} \mathit{add}(n : \mathtt{int}).G_{\text{update}}, \\ \mathit{done}().\mathsf{B} \to \mathsf{A} \{\mathit{total}(\mathit{tot} : \mathtt{int} \models \mathit{tot} = \mathit{partial}).\mathtt{end}\} \end{array} \right\}$$

with

$$G_{\text{update}} = \mathsf{B} \to \mathsf{C} \{\mathit{partial}(\mathit{tmp} : \mathtt{int} \models \mathit{tmp} = \mathit{partial} + n).G'_{\text{update}}\}$$
$$G'_{\text{update}} = \mathsf{C} \to \mathsf{B} \{\mathit{update}(\mathit{partial} : \mathtt{int} \models \mathit{partial} = \mathit{tmp}).\mathbf{T}\}$$

A.3 Diffie-Hellman protocol

The Diffie-Hellman protocol is a protocol that allows two participants to securely establish a shared secret, relying on the difficulty of the discrete logarithm [**?**].

In a nutshell, two participants A and B each have a private key (resp. a and b). In addition, public values p and g are known and not secret[8]. During the protocol A (resp. B) sends $A = g^a \bmod p$ to B (resp. $B = g^b \bmod p$ to A). The shared secret is then obtain on each side with $A^b = B^a$. In our implementation, we have an extra participant G used for setting-up the values. In practice, g and p can be shared beforehand or agreed-on on the fly. The private keys can be generated from a random number generator local to each participant. The protocol we implement is shown in Figure 8.

The RMPST of this protocol is:

$$G_1 = \mathsf{G} \to \mathsf{A} \{\mathit{generatorA}(g : \mathtt{int}).\mathsf{A} \to \mathsf{B} \{\mathit{generatorB}(h : \mathtt{int} \models g = h).G_2\}\}$$

[7] Based on an idea from ECOOP reviewer C.
[8] In addition, p and g must be co-prime. In our example, we only focus on the secret keys not being disclosed, and we therefore ignore this condition.

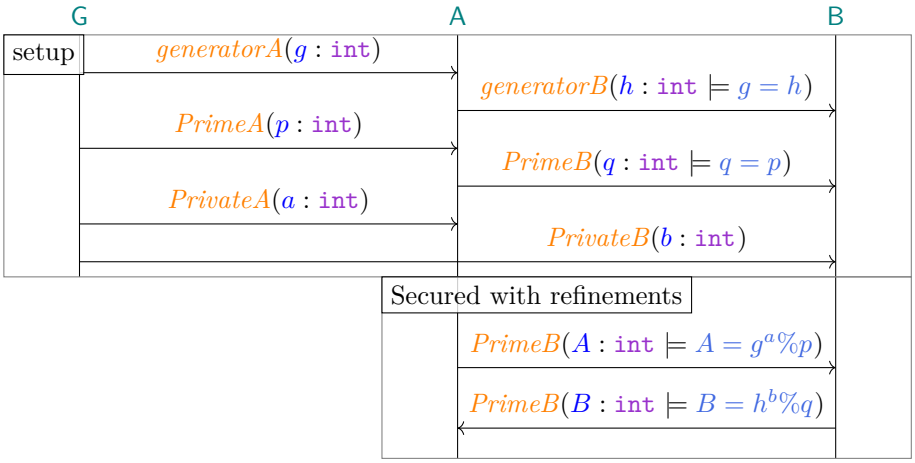

Figure 8 Communication diagram for the Diffie-Hellman key exchange protocol with refinements.

$$G_2 = \mathsf{G} \to \mathsf{A}\,\{PrimeA(p : \mathtt{int}).\mathsf{A} \to \mathsf{B}\,\{PrimeB(q : \mathtt{int} \models p = q).G_3\}\}$$

$$G_3 = \mathsf{G} \to \mathsf{A}\,\{PrivateA(a : \mathtt{int}).\mathsf{G} \to \mathsf{B}\,\{PrivateB(b : \mathtt{int}).G_4\}\}$$

$$G_4 = \mathsf{A} \to \mathsf{B}\,\{SharedA(A : \mathtt{int} \models A = g^a\%p).\mathsf{B} \to \mathsf{A}\,\{SharedB(B : \mathtt{int} \models B = g^b\%q).\mathtt{end}\}\}$$

While this example is theoretically simple (no recursion), we implemented this example in Rumpsteak, in order to show how to accommodate arithmetic refinements.

B Refined MPST

B.1 Preliminary definitions

▶ **Definition 41** (Map). *A* map M *is a set of pairs* $\langle t, v\rangle$*, where* t *is a variable and* v *is a value, such that there are no two pairs with the same variable. Maps are equipped with the following operations: lookup* $M(x)$*, update* $M[x \mapsto v]$*, domain* $\mathrm{dom}(M)$*, and removal* $M \backslash x$*.*

$$M(x) \stackrel{\text{def}}{=} \begin{cases} v & \text{if } \langle x, v\rangle \in M \\ \text{undefined} & \text{otherwise} \end{cases} \qquad M[x \mapsto v] \stackrel{\text{def}}{=} (M \setminus \{\langle x, v'\rangle \mid \forall v'\}) \cup \{\langle x, v\rangle\}$$

$$\mathrm{dom}(M) \stackrel{\text{def}}{=} \{x \mid \exists v \cdot \langle x, v\rangle \in M\} \qquad M \backslash x \stackrel{\text{def}}{=} M \setminus \{\langle x, v\rangle \mid \forall v\}$$

$$M_1 \biguplus M_2 \stackrel{\text{def}}{=} \begin{cases} M_1 \cup M_2 & \text{if } \mathrm{dom}(M_1) \cap \mathrm{dom}(M_2) = \varnothing \\ \text{undefined} & \text{otherwise} \end{cases}$$

We write $M_\varnothing$ *for the empty map.* ◁

Given a map M and a refinement r, we note $M \models r$ if and only if the refinement r is closed under the map M: $\mathrm{fv}(r) \subseteq \mathrm{dom}(M)$, and evaluates to truth after substitution: $\mathtt{eval}(r\{^{M(\mathrm{fv}(r))}/_{\mathrm{fv}(r)}\}) = \top$.

▶ **Definition 42** (Queues). *A* queue w *is a set of FIFOs for each pair of distinct participants in* $\mathbb{P} \times \mathbb{P}$*.* $w_{(\mathsf{p},\mathsf{q})}$ *denotes a FIFO from* p *to* q *in* w*. We define:*

1. $\mathtt{enq}_{(\mathsf{p},\mathsf{q})}(w, e) \stackrel{\text{def}}{=} \{w_{(\mathsf{p}',\mathsf{q}')} \mid \mathsf{p}' \neq \mathsf{p} \vee \mathsf{q}' \neq \mathsf{q}\} \cup \{\mathtt{enq}(w_{(\mathsf{p},\mathsf{q})}, e)\}$
2. $\mathtt{deq}_{(\mathsf{p},\mathsf{q})}(w) \stackrel{\text{def}}{=} \{w_{(\mathsf{p}',\mathsf{q}')} \mid \mathsf{p}' \neq \mathsf{p} \vee \mathsf{q}' \neq \mathsf{q}\} \cup \{\mathtt{deq}(w_{(\mathsf{p},\mathsf{q})})\}$ *if* $\mathtt{deq}(w_{(\mathsf{p},\mathsf{q})})$ *is defined*
3. $\mathtt{next}_{(\mathsf{p},\mathsf{q})}(w) \stackrel{\text{def}}{=} \mathtt{next}(w_{(\mathsf{p},\mathsf{q})})$

We write $w_\varnothing$ *for the empty queue, which is the queue where* $w_{(\mathsf{p},\mathsf{q})} = \varepsilon$ *for all* p *and* q. ◁

▶ **Definition 43** (Trace Ending-Up with Map). *A trace* τ ends up *with* M_τ *w.r.t. an initial map* M_I *if and only if:*

1. *if* τ *is* ϵ*, then* $M_\mathrm{I} = M_\tau$ *; and*
2. *if* τ *is* $\mathsf{p}\dagger\mathsf{q}\langle l, (x, c)\rangle : r \cdot \tau'$*, then* τ' *ends up with* M_τ *w.r.t.* $M_\mathrm{I}[x \mapsto c]$. ◁

B.2 Run and trace of an RCS

In order to define traces of RCSs, we first define *runs* (sequences of states, where the order is consistent with the reduction rules) and explain how to obtain a trace (as defined in Definition 1) from a run: the function `trace_step` extracts the action that happens between two consecutive configurations; by running this function on all successive states of a run, we retrieve the sequence of actions that took place, i.e. the trace of the run.

▶ **Definition 44** (Run of an RCS). *A* run *of an RCS is a sequence* $\sigma_0; \dots$ *of refined configurations such that* (i) *for all* $i \in \{1, \dots\}$, $\sigma_{i-1} \Rightarrow \sigma_i$*;* (ii) σ_0 *is initial ; and* (iii) *if the sequence is finite, then the last configuration* σ_n *is final.* ◁

▶ Remark 45 (Reachable State). We say a state σ of an RCS is *reachable* if there is a run of R that contains σ. This implies the run begins from an initial state (Item (ii) in Definition 44). ◁

▶ **Definition 46** (Trace of a Reduction Step).

$$\mathtt{trace\_step}(\langle \sigma_1, \sigma_2\rangle) \stackrel{\text{def}}{=} \begin{cases} \mathsf{j}?\mathsf{i}\langle m\rangle : r & \text{if } \sigma_1 \xRightarrow{s_\mathsf{i} \xrightarrow{\mathsf{j}?\mathsf{i}\langle m\rangle : r} s'_\mathsf{i}} \sigma_2 \\ \mathsf{i}!\mathsf{j}\langle m\rangle : r & \text{if } \sigma_1 \xRightarrow{s_\mathsf{i} \xrightarrow{\mathsf{i}!\mathsf{j}\langle m\rangle : r} s'_\mathsf{i}} \sigma_2 \\ \textit{undefined} & \textit{otherwise} \end{cases}$$

◁

▶ **Definition 47** (Trace of a Refined Communicating System). *Given a run* $\sigma_0; \dots$ *of an RCS, the trace of those reductions is given with the following function:*

$$\mathtt{trace}(\sigma_0; \sigma_1; \dots) = \begin{cases} \mathtt{trace\_step}(\langle \sigma_0, \sigma_1\rangle) \cdot \mathtt{trace}(\sigma_1; \dots) & \textit{if } \mathtt{trace\_step}(\langle \sigma_0, \sigma_1\rangle) \textit{ is defined} \\ \textit{undefined} & \textit{otherwise} \end{cases}$$

$$\mathtt{trace}(\epsilon) = \epsilon \qquad \mathtt{trace}(\sigma_0) = \epsilon$$

◁

▶ Remark 48 (Trace of step is invertible). `trace_step` is injective and therefore invertible. For the sake of simplicity, we implicitly convert traces into runs and vice-versa. ◁

▶ **Example 49** (Run and Trace of an RCS). A possible run of the RCS of $G_{\pm}$ is:

$\langle\langle \mathsf{A}_1, \mathsf{B}_1, \mathsf{C}_1\rangle, w_\varnothing, M_\varnothing\rangle; \langle\langle \mathsf{A}_2, \mathsf{B}_1, \mathsf{C}_1\rangle, w_1, \{\langle n, 5\rangle\}\rangle; \langle\langle \mathsf{A}_2, \mathsf{B}_2, \mathsf{C}_1\rangle, w_\varnothing, \{\langle n, 5\rangle\}\rangle;$
$\langle\langle \mathsf{A}_2, \mathsf{B}_2, \mathsf{C}_2\rangle, w_2, M\rangle; \langle\langle \mathsf{A}_2, \mathsf{B}_3, \mathsf{C}_2\rangle, w_\varnothing, M\rangle; \langle\langle \mathsf{A}_2, \mathsf{B}_4, \mathsf{C}_2\rangle, w_3, M\rangle; \langle\langle \mathsf{A}_2, \mathsf{B}_4, \mathsf{C}_3\rangle, w_\varnothing, M\rangle$

with the queues $w_1 = \mathtt{enq}_{(\mathsf{A},\mathsf{B})}(w_\varnothing, \langle secret, \langle n, 5\rangle\rangle)$; $w_2 = \mathtt{enq}_{(\mathsf{C},\mathsf{B})}(w_\varnothing, \langle guess, \langle x, 5\rangle\rangle)$; $w_3 = \mathtt{enq}_{(\mathsf{B},\mathsf{C})}(w_\varnothing, \langle correct, \langle \_, \_\rangle\rangle)$; and the map $M = \{\langle n, 5\rangle, \langle x, 5\rangle\}$.

This run produces the trace $\tau \cdot \tau_2$ presented in Example 6. Notice that, from the configuration with $\langle \mathsf{A}_2, \mathsf{B}_3, \mathsf{C}_2\rangle$, the system cannot take any of the GRSND transition with $\langle more, \langle \_, \_\rangle\rangle$ and $\langle less, \langle \_, \_\rangle\rangle$, as the refinement in the corresponding transition in the RCFSM of B does not hold with M: $M \not\models x < n$ (resp. $M \not\models x > n$) (c.f. Example 16).

B.3 Syntax

B.3.1 Roles in a global type

▶ **Definition 50** (Set of Roles in a Global Type).

$$\mathtt{parts}(G) = \begin{cases} \{\mathsf{p}, \mathsf{q}\} \cup \bigcup_{i\in I} \mathtt{parts}(G_i) & \text{if } G = \mathsf{p} \to \mathsf{q}\{l_i(x_i : \mathsf{S}_i \models R_i).G_i\}_{i\in I} \\ \mathtt{parts}(G') & \text{if } G = \mu \mathbf{t}.G' \\ \varnothing & \text{otherwise} \end{cases}$$

We note $\mathsf{p} \in G$ *for* $\mathsf{p} \in \mathtt{parts}(G)$. ◁

B.3.2 Type occurring in a type

Cases (i) — (iii) of Definition 51 recursively delve into the continuations, until we match on the exact type with case (iv).

▶ **Definition 51** (Type Occurring in a Type). *We say a type* T' *occurs in* T *(noted* $T' \in T$*) if and only if at least one of the following conditions holds:* (i) *if* T *is* $\mathsf{p} \oplus \{\ell_i(x_i : \mathsf{S}_i \models R_i).T_i\}_{i\in I}$*, there exist* $i \in I$ *such that* $T' \in T_i$*;* (ii) *if* T *is* $\mathsf{p} \& \{\ell_i(x_i : \mathsf{S}_i \models R_i).T_i\}_{i\in I}$*, there exist* $i \in I$ *such that* $T' \in T_i$*;* (iii) *if* T *is* $\mu \mathbf{t}.T_\mu$*,* $T' \in T_\mu$*; or* (iv) $T' = T$. ◁

▶ **Definition 52** (Global Type Occurring in a Global Type). *We say a type* G' *occurs in* G *(noted* $G' \in G$*) if and only if at least one of the following conditions holds:*

- *if* T *is* $\mathsf{p} \to \mathsf{q}\{\ell_i(x_i : \mathsf{S}_i \models r_i).G_i\}_{i\in I}$*, there exist* $i \in I$ *such that* $G' \in G_i$
- *if* G *is* $\mu \mathbf{t}.G_\mu$*,* $G' \in G_\mu$
- $G' = T$.

◁

C Section 2

C.1 Proofs of Lemmas

▶ **Lemma 53** (Concatenating Well-Queued Traces). *For any traces* τ_1 *and* τ_2*, for any queues* w_i*,* w_t *and* w_f*, if* τ_1 *ends up with queue* w_t *with respect to* w_i*, and* τ_2 *ends up with queue* w_f *with respect to* w_t*, then* $\tau_1 \cdot \tau_2$ *ends up with queue* w_f *with respect to* w_i. ◁

Proof. By induction on the size of τ_1. Notice that we leave w_i, w_t, w_f and τ_2 quantified; hence the property we want to prove by induction is $\forall \tau_2, w_i, w_t, w_f$, if τ_1 ends up with queue w_t w.r.t. w_i and τ_2 ends up with queue w_f w.r.t. w_t, then $\tau_1 \cdot \tau_2$ ends up with queue w_f w.r.t. w_i.

Base case, size of τ_1 is 0: in this case, $\tau_1 = \epsilon$, which trivially holds.

Base case, size of τ_1 is 1: in this case, $\tau_1 = \alpha$. We have to show that, for all τ_2, if α ends up with queue w_t w.r.t. w_i, and if τ_2 ends up with queue w_f w.r.t. w_t, then $\alpha \cdot \tau_2$ ends up with w_f w.r.t. w_i.
By case analysis on α:

- **If α is $\mathsf{p?q}\langle m \rangle : r$,** since α (i.e. $\alpha \cdot \epsilon$) ends up with w_t w.r.t. w_i, then, from Item 2 in Definition 3, ϵ ends up with queue w_t w.r.t. $\mathtt{deq}_{(\mathsf{p},\mathsf{q})}(w_i)$, and $\mathtt{next}_{(\mathsf{p},\mathsf{q})}(w_i) = m$.
  From Item 1 in Definition 3, $w_t = \mathtt{deq}_{(\mathsf{p},\mathsf{q})}(w_i)$.
  Therefore, we have that:
  - τ_2 ends up with w_f w.r.t. $w_t = \mathtt{deq}_{(\mathsf{p},\mathsf{q})}(w_i)$ (by hypothesis and above); and
  - $\mathtt{next}_{(\mathsf{p},\mathsf{q})}(w_i) = m$ (as shown above).

  Therefore, from Item 2 in Definition 3, $\alpha \cdot \tau_2$ ends up with w_f w.r.t. w_i.
- **If α is $\mathsf{p!q}\langle m \rangle : r$,** since α (i.e. $\alpha \cdot \epsilon$) ends up with w_t w.r.t. w_i, then, from Item 3 in Definition 3, ϵ ends up with queue w_t w.r.t. $\mathtt{enq}_{(\mathsf{p},\mathsf{q})}(w_i, m)$.
  From Item 1 in Definition 3, $w_t = \mathtt{enq}_{(\mathsf{p},\mathsf{q})}(w_i, m)$.
  Therefore, we have that τ_2 ends up with w_f w.r.t. $w_t = \mathtt{enq}_{(\mathsf{p},\mathsf{q})}(w_i, m)$ (by hypothesis and above). Therefore, from Item 3 in Definition 3, $\alpha \cdot \tau_2$ ends up with w_f w.r.t. w_i.

Inductive case, size of τ_1 is $n+1$ ($n \geq 1$): The induction hypothesis (IH) is: for all ${\tau_1}^{IH}$ with length less or equal to n, for all ${\tau_2}^{IH}, {w_i}^{IH}, {w_t}^{IH}, {w_f}^{IH}$, if ${\tau_1}^{IH}$ ends up with queue ${w_t}^{IH}$ w.r.t ${w_i}^{IH}$, and ${\tau_2}^{IH}$ ends up with queue ${w_f}^{IH}$ w.r.t. ${w_t}^{IH}$, then ${\tau_1}^{IH} \cdot {\tau_2}^{IH}$ ends up with queue ${w_f}^{IH}$ w.r.t. ${w_i}^{IH}$.
Since the length of τ_1 is $n+1 \geq 2$, $\tau_1 = \alpha \cdot \tau_1'$. Notice that the length of τ_1' is n.
Since τ_1 ends up with w_t w.r.t. w_i, then τ_1' ends up with queue w_t w.r.t. either $\mathtt{enq}_{(\mathsf{p},\mathsf{q})}(w_i, m)$ or $\mathtt{deq}_{(\mathsf{p},\mathsf{q})}(w_i)$ (depending on α), let call it w_α.
Therefore, by applying the induction hypothesis with ${\tau_1}^{IH} = \tau_1'$, ${\tau_2}^{IH} = \tau_2$, ${w_i}^{IH} = w_\alpha$, ${w_t}^{IH} = w_t$, ${w_f}^{IH} = w_f$, we have that $\tau_1' \cdot \tau_2$ ends up with queue w_f w.r.t. w_α.
By applying the induction hypothesis a second time, with ${\tau_1}^{IH} = \alpha$, ${\tau_2}^{IH} = \tau_1' \cdot \tau_2$, ${w_i}^{IH} = w_i$, ${w_t}^{IH} = w_{t_i}$, ${w_f}^{IH} = w_f$, we have that $\alpha \cdot \tau_1' \cdot \tau_2 = \tau_1 \cdot \tau_2$ ends up with queue w_f w.r.t. w_i, which concludes the inductive step. ◀

▶ **Lemma 54** (Concatenating Well-Predicated Traces). *For any map M, for any traces τ_1 and τ_2, if τ_1 is well-predicated under M and ends up with M_{τ_1} with respect to M, and if τ_2 is well-predicated under M_{τ_1}, then $\tau_1 \cdot \tau_2$ is well-predicated under M.* ◁

Proof. By induction on the size of τ_1:

Case size of τ_1 is 0: In that case, $\tau_1 = \epsilon$, therefore $\tau_1 \cdot \tau_2 = \tau_2$ and $M_{\tau_1} = M$ (from Item 1 in Definition 43). The result then trivially holds.

Case size of τ_1 is $n+1$: the induction hypothesis is: "for all M', τ_1' and τ_2', (H1) if the size of τ_1' is less than or equal to n, then (H2) if τ_1' is well-predicated with respect to M' and ends up with $M'_{\tau_1'}$ w.r.t. to M', and (H3) if τ_2' is well-predicated with respect to $M'_{\tau_1'}$, then $\tau_1' \cdot \tau_2'$ is well-predicated with respect to M'".
Let $\tau_1 = \alpha_0 \cdot \alpha_1 \cdot \ldots \cdot \alpha_n$, with $\alpha_0 = \mathsf{p}\dagger\mathsf{q}\langle l, (x, c)\rangle : r$.
Given that τ_1 is well-predicated with respect to M, then (i) r holds under $M[x \mapsto c]$; and (ii) $\alpha_1 \cdot \ldots \cdot \alpha_n$ is well-predicated with respect to $M[x \mapsto c]$.
Let $M_{\alpha_1 \cdot \ldots \cdot \alpha_n}$ be the map $\alpha_1 \cdot \ldots \cdot \alpha_n$ ends up to with respect to $M[x \mapsto c]$.

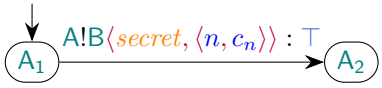

(a) RCFSM of A in the $G_{\pm}$ protocol.

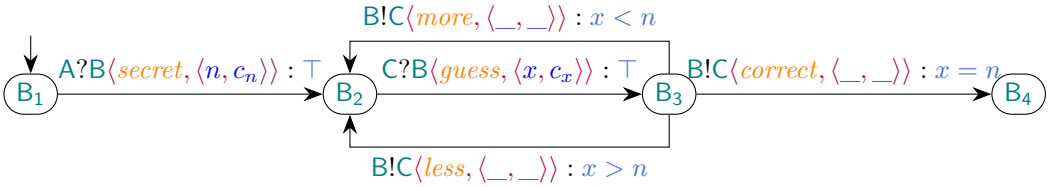

(b) RCFSM of B in the $G_{\pm}$ protocol.

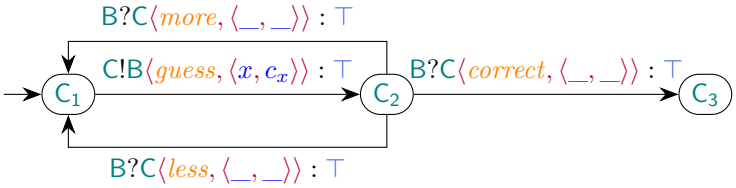

(c) RCFSM of C in the $G_{\pm}$ protocol.

Figure 9 RCFSM of the participants of the $G_{\pm}$ protocol.

From the induction hypothesis (with M' being $M[x \mapsto c]$, τ_1' being $\alpha_1 \cdot \ldots \cdot \alpha_n$, τ_2' being τ_2), and given that (i) (H1) holds trivially; (ii) (H2) holds from Item (ii); and (iii) (H3) holds from Item 2 in Definition 43 (which directly shows that $M_{\tau_1} = M_{\alpha_1 \cdot \ldots \cdot \alpha_n}$) then $\alpha_1 \cdot \ldots \cdot \alpha_n \cdot \tau_2$ is well-predicated under $M[x \mapsto c]$.
From Item (i) above and Item (ii) in Definition 5, we have that $\alpha_0 \cdot \alpha_1 \cdot \ldots \cdot \alpha_n \cdot \tau_2 = \tau_1 \cdot \tau_2$ is well-predicated under M, which concludes the inductive step. ◀

D Refined Automata

▶ **Lemma 55.** *For all runs $\sigma_0; \ldots; \sigma_{i-1}; \sigma_i; \ldots$, if $\sigma_{i-1} \overset{t}{\Rightarrow} \sigma_i$ with $t = s_{\mathsf{i}} \xrightarrow{\mathsf{j?i}\langle \ell, \langle x, c \rangle\rangle : r} s_{\mathsf{i}}'$, then exist $j \leq i-1$ and r' such that $\sigma_{j-1} \overset{t'}{\Rightarrow} \sigma_j$ with $t' = s_{\mathsf{i}} \xrightarrow{\mathsf{j!i}\langle \ell, \langle x, c \rangle\rangle : r'} s_{\mathsf{i}}'$.* ◁

Proof. Let w' be the queues of σ'. From the premises of GRRec, $m = \langle \ell, \langle x, c \rangle\rangle \in w_{(\mathsf{j},\mathsf{i})}$. The only rule that enqueues m in $w_{(\mathsf{j},\mathsf{i})}$ is GRSnd, with $s_{\mathsf{i}} \xrightarrow{\mathsf{j!i}\langle \ell, \langle x, c \rangle\rangle : r'} s_{\mathsf{i}}'$ (for some r') ◀

▶ **Theorem 18** (Traces of Refined Communicating Systems are Valid Refined Traces). *For all RCS R, for all initial and final traces τ of R, τ is a valid refined trace.* ◁

Proof. From Definition 7, we have to show that (i) τ is well-queued with regards to the empty queue $w_\varnothing$; and (ii) τ is well-predicated by the empty variable context $V_\varnothing$. We show the two points separately.

Well-Queued

By case analysis on the length of τ.

Case length of τ is 0 ($\tau = \epsilon$): since the trace is empty, initial and final, the corresponding run of the automaton is composed of a single state $\sigma = \langle \langle s_{\mathsf{i}} \rangle_{\mathsf{i} \in I}, w, m \rangle$, which is therefore both initial and final.
From Definition 12 (Initial Refined Global State), the queues w of the state are empty ($w = w_\varnothing$).
Therefore, from Definition 3, τ is well-queued with regards to $w_\varnothing$.

Case length of τ is 1 ($\tau = \alpha$): We prove this case leads to a contradiction, and therefore cannot happen.
From the definition of δ, the label of the transition is either $\mathsf{p}?\mathsf{q}\langle \ell, \langle x, c \rangle \rangle : r$ or $\mathsf{p}!\mathsf{q}\langle \ell, \langle x, c \rangle \rangle : r$.
Therefore, for any σ_1, σ_2 such that $\mathtt{trace}(\sigma_1; \sigma_2) = \alpha$, we show that it is not simulatenously possible for σ_1 to be initial and for σ_2 to be final. Let w_{σ_1} (resp. w_{σ_2}) be the queue of σ_1 (resp. σ_2).
If the label is $\mathsf{p}?\mathsf{q}\langle \ell, \langle x, c \rangle \rangle : r$, then from Definitions 46 and 47, $\sigma_1 \Rightarrow \sigma_2$ with a GRRec transition. From the premise of GRRec, $\mathtt{next}_{(\mathsf{p},\mathsf{q})}(w_{\sigma_1}) = \langle \ell, \langle x, c \rangle \rangle$, i.e. $w_{\sigma_1 (\mathsf{p},\mathsf{q})}$ is not ε, i.e. $w_{\sigma_1} \neq w_\varnothing$, therefore σ_1 is not initial.
Similarly, if the label is $\mathsf{p}!\mathsf{q}\langle \ell, \langle x, c \rangle \rangle : r$, then from Definitions 46 and 47, $\sigma_1 \Rightarrow \sigma_2$ with a GRSnd transition. From the conclusion of GRSnd, $w_{\sigma_2} = \mathtt{enq}_{(\mathsf{p},\mathsf{q})}(w_{\sigma_1}, \langle \ell, \langle x, c \rangle \rangle)$.
Therefore, from Definition 42, $w_{\sigma_2} \neq w_\varnothing$, therefore σ_2 is not final.
Therefore, if τ contains a single element, then τ is not initial and final, which contradicts the hypothesis.

Case length of τ is greater than 1 ($\tau = \tau_1 \cdot \tau_2$ for non-empty τ_1 and τ_2): first, let's notice that τ is initial and final; therefore the initial state σ_i and the final state σ_f of τ both have $w_\varnothing$ as their queues (from Definition 12). Therefore, we only have to show that τ ends up with w_f (the queue of σ_f) with respect to w_i (the queue of σ_i).
By contradiction, suppose τ does not end up in w_f w.r.t w_i.
Let σ_t be the final state of τ_1, starting from σ_i. Let w_t be the queue of σ_t.
From the contraposition of Lemma 53, either (i) τ_1 does not end up with queue w_t w.r.t. w_i; or (ii) τ_2 does not end up with queue w_f w.r.t. w_t.
In either case, we can recursively apply the contraposition of Lemma 53 until we have a trace composed of a single transition which trace is α_c that does not end up in ${w_c}'$ w.r.t. w_c.
By case analysis of α_c:

- **Case $\alpha_c = \mathsf{p}!\mathsf{q}\langle m \rangle : r$:** From Remark 48, we deduce that α_c is a GRSnd transition $s_{\mathsf{p}} \xrightarrow{\mathsf{p}!\mathsf{q}\langle m \rangle : r} {s_{\mathsf{p}}}'$.
  From the definition of GRSnd, we have that ${w_c}' = \mathtt{enq}_{(\mathsf{p},\mathsf{q})}(w_c, m)$.
- **Case $\alpha_c = \mathsf{p}?\mathsf{q}\langle m \rangle : r$:** From Remark 48, we deduce that α_c is a GRRec transition with $s_{\mathsf{q}} \xrightarrow{\mathsf{p}?\mathsf{q}\langle m \rangle : r} {s_{\mathsf{q}}}'$.
  From the definition of GRSnd, we have that ${w_c}' = \mathtt{deq}_{(\mathsf{p},\mathsf{q})}(w_c)$ and $\mathtt{next}_{(\mathsf{p},\mathsf{q})}(w_c) = m$.

In both cases, α_c ends up in ${w_c}'$ w.r.t. w_c. Contradiction.

Well-Predicated

By induction on the length of τ.

Case τ is ϵ: the result trivially hold from Item (i) in Definition 5.

Case τ is $\tau' \cdot \alpha_n$: the induction hypothesis is that τ' is well-predicated by the empty map $M_\varnothing$.
Suppose that τ' ends up with map $M_{\tau'}$. By case analysis of the action α_n:

Case α_n is $\mathsf{p?q}\langle \ell, \langle x, c\rangle\rangle : r$: from Remark 48, this corresponds to a GRRec transition, which corresponds to a reduction:

$$\langle\langle s_1, \ldots, s_\mathsf{p}, \ldots, s_\mathsf{n}\rangle, w, M_{\tau'}\rangle \Rightarrow \langle\langle s_1, \ldots, s_\mathsf{p}', \ldots, s_\mathsf{n}\rangle, \mathtt{deq}_{(\mathsf{q},\mathsf{p})}(w), M_{\tau'}[x \mapsto c]\rangle$$

with $s_\mathsf{p} \xrightarrow{\mathsf{p?q}\langle \ell, \langle x, c\rangle\rangle : r} s_\mathsf{p}'$ and $M_{\tau'}[x \mapsto c] \models r$ (from the premises of GRRec).
Since τ' is a well-predicated trace, that ends up with $M_{\tau'}$, from Lemma 54 and the induction hypothesis, we simply have to show that α_n is well-predicated under $M_{\tau'}$.
This holds directly from that $M_{\tau'}[x \mapsto c] \models r$ (from above) and Item (ii) in Definition 5.

Case α_n is $\mathsf{p!q}\langle \ell, \langle x, c\rangle\rangle : r$: from Remark 48, this corresponds to a GRSnd transition, which corresponds to a reduction:

$$\langle\langle s_1, \ldots, s_\mathsf{p}, \ldots, s_\mathsf{n}\rangle, w, M_{\tau'}\rangle \Rightarrow \langle\langle s_1, \ldots, s_\mathsf{p}', \ldots, s_\mathsf{n}\rangle, \mathtt{enq}_{(\mathsf{p},\mathsf{q})}(w, \langle \ell, \langle x, c\rangle\rangle), M_{\tau'}[x \mapsto c]\rangle$$

with $s_\mathsf{p} \xrightarrow{\mathsf{p!q}\langle \ell, \langle x, c\rangle\rangle : r} s_\mathsf{p}'$ and $M_{\tau'}[x \mapsto c] \models r$ (from the premises of GRSnd).
Since τ' is a well-predicated trace, that ends up with $M_{\tau'}$, from Lemma 54 and the induction hypothesis, we simply have to show that α_n is well-predicated under $M_{\tau'}$.
This holds directly from that $M_{\tau'}[x \mapsto c] \models r$ (from above) and Item (ii) in Definition 5.
◀

E Decentralised Verification

E.1 Decentralised run of an RCS

▶ **Example 56** (Run of a Decentralised Configuration)**.** We present a run of a decentralised configuration, which corresponds the run of the (centralised) configuration in Example 49. Compared to that example, notice that there is no single global map, but instead each local state is associated with a local map (with w_1, w_2, w_3 and M as in Example 49).

$\langle\langle \mathsf{A}_1, M_\varnothing\rangle, \langle \mathsf{B}_1, M_\varnothing\rangle, \langle \mathsf{C}_1, M_\varnothing\rangle, w_\varnothing\rangle; \langle\langle \mathsf{A}_2, M_\varnothing\rangle, \langle \mathsf{B}_1, M_\varnothing\rangle, \langle \mathsf{C}_1, M_\varnothing\rangle, w_1\rangle;$
$\langle\langle \mathsf{A}_2, M_\varnothing\rangle, \langle \mathsf{B}_2, \{\langle n, 5\rangle\}\rangle, \langle \mathsf{C}_1, M_\varnothing\rangle, w_\varnothing\rangle; \langle\langle \mathsf{A}_2, M_\varnothing\rangle, \langle \mathsf{B}_2, \{\langle n, 5\rangle\}\rangle, \langle \mathsf{C}_2, M_\varnothing\rangle, w_2\rangle;$
$\langle\langle \mathsf{A}_2, M_\varnothing\rangle, \langle \mathsf{B}_3, M\rangle, \langle \mathsf{C}_2, M_\varnothing\rangle, w_\varnothing\rangle; \langle\langle \mathsf{A}_2, M_\varnothing\rangle, \langle \mathsf{B}_4, M\rangle, \langle \mathsf{C}_2, M_\varnothing\rangle, w_3\rangle;$
$\langle\langle \mathsf{A}_2, M_\varnothing\rangle, \langle \mathsf{B}_4, M\rangle, \langle \mathsf{C}_3, M_\varnothing\rangle, w_\varnothing\rangle;$

E.2 Simulation

▶ **Definition 57** (Simulation [30])**.** *Given two labelled transition systems with states taken from S_1 and S_2, a relation $\mathcal{R} \subseteq S_1 \times S_2$ is a simulation if: $\forall s_1, s_1' \in S_1$, if $s_1 \to s_1'$, then $\forall s_2 \in S_2$ such that $\langle s_1, s_2\rangle \in \mathcal{R}$, there exists s_2' such that $s_2 \to s_2'$ and $\langle s_1', s_2'\rangle \in \mathcal{R}$.* ◁

▶ Remark 58 (Bisimulation). A relation $\mathcal{R}$ is a *bisimulation* if $\mathcal{R}$ and $\mathcal{R}^{-1}$ are simulations. ◁

E.3 Main theorem

▶ **Theorem 59** (Centralised simulates Decentralised)**.** *For all decentralisable RMPST G (Definition 30), $\mathcal{C}(G)$ simulates $\mathcal{D}(G)$.* ◁

Proof. The proof is quite straightforward, we match each DRec by GRRec and each DSnd by GRSnd. We can trivially note that the queue of messages is the same in both $\mathcal{D}(G)$ and $\mathcal{C}(G)$. Similarly, each s_i match in both systems, and remain matched after taking transitions.
Finally, the map of $\mathcal{C}(G)$ is always a superset of the union of all M_i in $\mathcal{D}(G)$.
We show the proof for DSnd and GRSnd. The proof for DRec and GRRec is similar.

Suppose that we have $D = \langle \langle \ldots, \langle s_i, M_i \rangle \ldots \rangle, w \rangle$ reducing to $D' = \langle \langle \ldots, \langle s_i', M_i' \rangle \ldots \rangle, w' \rangle$ with DSND and let j be the destination of the message. Suppose $S = \langle \langle \ldots, s_i, \ldots \rangle, w, M \rangle$, such that $M \subseteq \biguplus_{i \in \mathtt{parts}(G)} M_i$.

Notice that the premises of DSND entails the premises of GRSND. Therefore, GRSND can be triggered: let $S' = \langle \langle \ldots, s_i', \ldots \rangle, w', M' \rangle$ be the resulting configuration. To prove the simulation, we just have to show that the resulting s_i' of both configurations are the same, which is direct from the rule; that the queues w' of both configurations are the same, which is also direct from the rule; and that M' is indeed a subset of $\biguplus_{i \in \mathtt{parts}(G)} M_i'$.

Let x be the variable sent and c its associated value. By definition $M_i' = M_i \backslash x$, and all other local maps are unchanged. Also, by hypothesis (no duplication hypothesis), no other map contains x. Therefore, $M' = M[x \mapsto c] \supsetneq M \backslash x \supseteq (\biguplus_{i \in \mathtt{parts}(G)} M_i) \backslash x = \biguplus_{i \in \mathtt{parts}(G)} M_i'$. Therefore, $\mathcal{C}(G)$ simulates $\mathcal{D}(G)$. ◀

E.4 Static Verification of the Two Conditions

In the previous section, we introduced two conditions for the refined configurations to simulate decentralised configurations (Definition 29). We now aim to statically verify whether those two conditions hold for a given type. The algorithm we present keeps track of variable moves and tries to find which participant has a copy of which variable at any point in the execution of the protocol, which we call *variable localisation*.

Intuitively, from a global type G, it is possible to infer constraints on variable localisation: for instance, given the global type $G = \mathsf{A} \to \mathsf{B}\,\{\ell(x : \mathtt{int} \models x \geq 0).\mathtt{end}\}$, we can infer that participant A has x before the communication, and that x is transfered to B during the communication, therefore x is at B after the communication. The algorithm we propose generates such constraints to find variables localisations at any time. We represent RMPST as graphs (Appendix E.4.1), and the algorithm finds an assignation of variables that is consistent with the actions of the type. Finding such an assignation is trivially solved with logic programming (Appendix E.4.2). The only non-trivial part is to carefully pinpoint where variables are first used in loops. We illustrate this problem and solve it by unrolling loops in the graph of the type, which we present in Appendix E.4.3.

E.4.1 Type Graphs

First, we need to create the *graph of a global type*. This is analogous to the CFSM of a local type, but for global types instead.

▶ **Definition 60** (Graph of a Global Type). *Given a global type T_0, the* graph $\mathsf{GG}(T_0)$ *of that type is the labeled graph $\langle V, E \rangle$, where $V = \{T' | T' \in T_0 \wedge T' \neq \mathbf{t} \wedge T' \neq \mu \mathbf{t}.T_\mu\}$, $E = \delta_g$, and where δ_g is defined as the smallest relation such that*[9]*, for all $T \in V$, with $T = \mathsf{p} \to \mathsf{q}\{\ell_i(x_i : \mathsf{S}_i \models r_i).T_i\}_{i \in I}$:* (i) *if $\mathsf{strip}(T_i)$ is a communication or a termination, then, for all $i \in I$, $\langle T, \langle \mathsf{p}, \mathsf{q}, x_i, \mathsf{fv}(r_i) \rangle, \mathsf{strip}(T_i) \rangle \in \delta_g$; and* (ii) *if $\mathsf{strip}(T_i) = \mathbf{t}$ with $\mu \mathbf{t}.T' \in T_0$, then $\langle T, \langle \mathsf{p}, \mathsf{q}, x_i, \mathsf{fv}(r_i) \rangle, \mathsf{strip}(T') \rangle \in \delta_g$* ◁

Where `strip` is defined on global types as on local types.

▶ **Example 61** (Graph of a global type). The type graph of $G_\pm$ is shown below. Notice that the two transitions with *more* and *less* result in a single edge in the graph, as we erase the label, and both lead to the same state.

[9] Notice that δ_g is almost like δ, but we erase the value in the messages' payloads

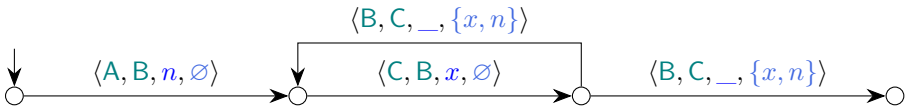

Among the vertices of the graph type $\mathsf{GG}(T_0)$ of a global type T_0, we distinguish the *initial* node, which is the vertice labeled T_0. In the graphical representation, this state is shown with an arrow.

E.4.2 Localising Algorithm

We now present the *localisation algorithm*. The goal of this algorithm is to find whether there exists a state, in the execution of the protocol specified by a type G, where a variable is duplicated, and whether there exists an action which refinement requires a variable that is not locally available.

Our algorithm infers the location of each variable (i.e. which participant has access to which variables), at each step of the protocol. We infer such location information from the actions of the protocol. This algorithm is expressed as a set of inference rules, which can be directly encoded in a logic programming language as DataLog.

The input of our algorithm is given by two provided atoms that are extracted from the global type graph (we show provided atoms in green and computed atoms in orange):

- $\mathit{Send}(s_1, \mathsf{p}, x, \mathsf{q}, s_2)$ holds when the graph contains an edge from s_1 to s_2 with label $\langle \mathsf{p}, \mathsf{q}, x, \_ \rangle$;
- $\mathit{FVRefinement}(s_1, x, s_2)$ holds when there is an edge from s_1 to s_2 with a label $\langle \_, \_, \_, \mathsf{fv}(r) \rangle$ and $\mathsf{fv}(r)$ contains x.

The core part of our algorithm is the atom $\mathit{In}(s, \mathsf{p}, x)$, which holds if p has access to variable x in state s (or, equivalently, if x is at p in s).

The first rule simply capture the fact that, if a participant p sends a variable x to q, then x is located at q after the exchange (notice that x is *not* necessarily available at p before the exchange, as it might be the first time it appears): $\mathit{In}(s_2, \mathsf{q}, x) \to \mathit{Send}(\_, \_, x, \mathsf{q}, s_2)$.

We then define two rules that infer which variables are available at which location after a communication. When p sends x, all $y \neq x$ are kept at p. For the other participants $\mathsf{r} \neq \mathsf{p}$, all variables are preserved.

$$\mathit{In}(s_2, \mathsf{p}, x) \leftarrow \mathit{In}(s_1, \mathsf{p}, x), \mathit{Send}(s_1, \mathsf{p}, y, \_, s_2), (x \neq y);$$

$$\mathit{In}(s_2, \mathsf{r}, x) \leftarrow \mathit{In}(s_1, \mathsf{r}, x), \mathit{Send}(s_1, \mathsf{p}, \_, \_, s_2), (\mathsf{p} \neq \mathsf{r});$$

Once variables are localised, we can check whether the two conditions hold. First, the atom $\mathit{NotVerifFV}$ checks whether there is a state s in which p sends a message with a refinement that contains a free variable x which p cannot access (that is a variable x that is neither located at p nor in the message sent). If that happens, the system is not verifiable, as p will not be able to check its refinement when taking the transition. Second, the atom $\mathit{NotVerifDupl}$ checks whether a variable is duplicated. It simply records whether there is a state s in which two distinct participants p and q can access the same variable x.

$$\mathit{NotVerifFV}(s, s_2, x, \mathsf{p}) \leftarrow \mathit{FVRefinement}(s, x, s_2), \mathit{Send}(s, \mathsf{p}, y, \_, \_), !\mathit{In}(s, \mathsf{p}, x), (x \neq y);$$

$$\mathit{NotVerifDupl}(s, x, \mathsf{p}, \mathsf{q}) \leftarrow \mathit{In}(s, \mathsf{p}, x), \mathit{In}(s, \mathsf{q}, x), (\mathsf{p} \neq \mathsf{q});$$

Finally, the two conditions hold if there is no $\mathit{NotVerifFV}$ nor $\mathit{NotVerifDupl}$.

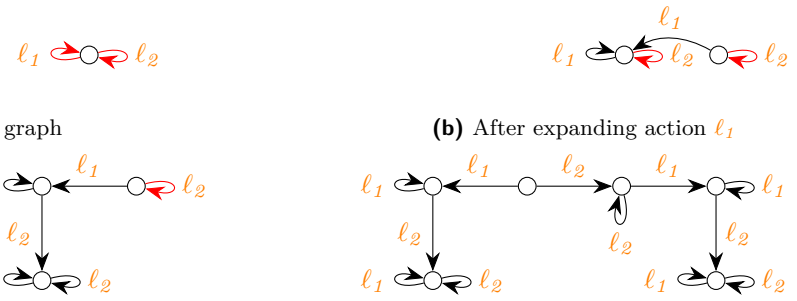

(a) The initial graph

(b) After expanding action ℓ_1

(c) After expanding the lower ℓ_2 action

(d) The graph type fully unrolled

Figure 10 The graph of G', with unvisited edges in red. For the sake of simplicity, we only use ℓ_1 and ℓ_2 as the two communication labels. Notice that the unrolled graph is not the smallest (in that some loops might be unrolled more than needed), but still unrolls every loop.

E.4.3 Recursion Unrolling

The localisation algorithm does not consider the case where a variable is first sent in a loop. Consider the type:

$$\mu \mathbf{T}.\mathsf{A} \to \mathsf{B} \left\{ \begin{array}{l} \ell_1(x : \mathtt{int} \models x = y).\mathsf{B} \to \mathsf{A}\,\{\ell_3(x : \mathtt{int}).\mathbf{T}\} \\ \ell_2(y : \mathtt{int} \models x \neq y).\mathsf{B} \to \mathsf{A}\,\{\ell_4(y : \mathtt{int}).\mathbf{T}\} \end{array} \right\}$$

This type cannot be verified, as if A chooses the branch ℓ_1 on the first iteration of the loop, then y is not defined, and similarly for the branch ℓ_2. However, our localisation algorithm first computes the location of variables (In), which results in the following graph and variable locations:

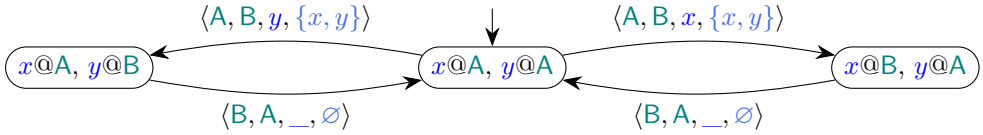

Therefore, we have to distinguish the first time a branch is taken from the following iterations. We therefore introduce our *unrolling* algorithm, which unrolls each loop once to distinguish the first iteration from the following ones. This task is non trivial, as it is not a simple syntactical replacement of a recursion variable by its definition since we have to take into account all branches. Specifically, in the example above, A can possibly take the branch ℓ_1 multiple times before choosing ℓ_2: we have to take all declarations orders into account.

For the sake of simplicity, consider the type G' (which is not localisable, but greatly reduces the size of the graph and is sufficient to explain our algorithm): $\mu \mathbf{T}.\mathsf{A} \to \mathsf{B}\,\{\ell_1(x : \mathtt{int}).\mathbf{T}; \ell_2(y : \mathtt{int}).\mathbf{T}\}$. The different steps of the execution of our algorithm are shown in Figure 10.

We first distinguish forward and backward edges in the type graph: a backward arrow is an arrow that contains a recursion, i.e. it ends in a previous state or in the same state. We initially mark all those backward edges as unvisited. In Figure 10, those unvisited backward arrows are shown in red. We then unroll those unvisited backward edges one at a time. For each unvisited backward edge E from N to N', we copy the graph, we mark E as visited in the copied graph and we replace E by a forward visited edge from N in the original graph to N' in the copied graph.

Data: G – A type graph
while $\exists e = \langle v_1, v_2\rangle \in \mathbb{B}_G \cap \mathbb{U}_G$ **do**
  Mark e as visited in G;
  Let G' be a copy of the continuation of v_1 using e in G;
  Let v_2' be the copy of v_2 in G';
  Remove e from G;
  Let G be the union of G and G';
  Add an edge from v_1 to v_2' in G;
end
Remove unreachable vertices;

Depth of a node: The depth of a node N is the length of the shortest path from the initial node to N.

Backward edge: An edge from N_1 to N_2 is said *backward* if the depth of N_1 is greater than or equal to the depth of N_2. We note $\mathbb{B}_G$ the set of backward edges of graph G.

State of an edge: We assign each edge a state taken from $\{\text{visited}, \text{unvisited}\}$. We note $\mathbb{U}_G$ the set of unvisited edges of graph G.

Continuation: The continuation of a node N following an edge E is the (sub)graph reachable from the destination of E, provided that the origin of E is N.

F Static Elision of Redundant Refinements

F.1 Static Elision of Refinements in RCS

F.1.1 Definitions.

▶ **Definition 62** (Independent transitions). *A transition* $t = \sigma \xrightarrow{\mathsf{p}\dagger\mathsf{q}\langle \ell, \langle x, \_\rangle\rangle : r} \sigma'$ depends *on* $\mathrm{fv}(r)$*, the free variables of its refinement.*

We say that t depends *on another transition* $t' = \rho \xrightarrow{\mathsf{r}\dagger\mathsf{s}\langle \ell', \langle y, \_\rangle\rangle : r'} \rho'$ *if* t *depends on* y*, the payload of* t'*. Otherwise, we say* t *is independent of* t'*. When* t *depends (resp. does not depend) on itself, we say it is* self-dependent *(resp.* self-independent*).* ◁

▶ **Definition 63** (Well-defined transitions). *Given a RCS with a CFSM containing a transition* $t = s_\mathsf{i} \xrightarrow{\_\dagger\_\langle \_, \_\rangle : r} s_\mathsf{i}'$*, we say* t *is* well-defined *if and only if, in all reachable states* $\langle \langle \ldots, s_\mathsf{i}, \ldots \rangle, \_, M\rangle$*,* $\mathrm{fv}(r) \subseteq \mathrm{dom}(M)$*.* ◁

F.1.2 Lemmas and proofs.

▶ **Lemma 64.** *Let* $t_1 = \sigma \xrightarrow{\mathsf{p}\dagger\mathsf{q}\langle \ell, \langle x, c\rangle\rangle : r} \sigma'$ *and* $t_2 = \rho \xrightarrow{\mathsf{r}\dagger\mathsf{s}\langle \ell', \langle y, c'\rangle\rangle : r'} \rho'$ *such that* t_1 *is independent from* t_2*.*

For all M*,* M_2 *such that* $\langle \_, \_, M\rangle \Rightarrow \langle \_, \_, M_2\rangle$ *with* t_2*;* $M \models r$ *if and only if* $M_2 \models r$*.* ◁

Proof of Lemma 64. First, t_1 is independent of t_2, i.e. (Definition 62), t_1 does not depend on y; that is $y \notin \mathrm{fv}(r)$. Also, since $\langle \_, \_, M\rangle \Rightarrow \langle \_, \_, M_2\rangle$, then $M_2 = M[y \mapsto c']$.

In addition, by definition (Section 2.1), $M \models r$ if and only if: (i) $\mathrm{fv}(r) \subseteq \mathrm{dom}(M)$; and (ii) $\mathrm{eval}(r\{^{M(\mathrm{fv}(r))}/_{\mathrm{fv}(r)}\}) = \top$.

Therefore, $r\{^{M(\mathrm{fv}(r))}/_{\mathrm{fv}(r)}\} = r\{^{M_2(\mathrm{fv}(r))}/_{\mathrm{fv}(r)}\}$ and $\mathrm{fv}(r) \subseteq \mathrm{dom}(M)$ if and only if $\mathrm{fv}(r) \subseteq \mathrm{dom}(M) \cup \{y\} = \mathrm{dom}(M_2)$.

Therefore, $M \models r$ if and only if $M_2 \models r$. ◀

▶ **Lemma 65.** *Given an RCS* R *containing an RCFSM with a transition* $s_\mathsf{i} \xrightarrow{\_\dagger\_\langle \_\rangle : r} s_\mathsf{i}'$*, and such that for each transition* $\_ \xrightarrow{\_!\_\langle \_\rangle : r_w} \_ \in \bigcup_{x \in \mathrm{fv}(r)} \mathbb{T}_x$*, for all map* M*,* $M \models r_w$ *entails* $M \models r$*.*

For all reachable states $\sigma = \langle \langle \vec{s_\mathsf{i}}\rangle, w, M\rangle$ *of* R*, if* $\mathrm{fv}(r) \subseteq \mathrm{dom}(M)$*, then* $M \models r$*.* ◁

Proof. Since σ is reachable, there is a run $\sigma_0; \ldots; \sigma; \ldots$ that contains σ. By induction on the run.

Base case (σ is initial): Since σ is an initial state, then $M = M_\varnothing$, our claim vacuously holds.

Inductive case ($\sigma' \xRightarrow{t} \sigma$ with $t = \_ \xrightarrow{\_\dagger\_\langle\_,\langle x,c_x\rangle\rangle:r'} \_$) Let $\sigma' = \langle\langle \vec{s_i}'\rangle, w', M'\rangle$. The induction hypothesis is that if $\mathrm{fv}(r) \subseteq M'$, then $M' \models r$. Also, $M \models r'$, from the premises of GRRec or GRSnd. By case analysis on the transition t:

- **If $x \notin \mathrm{fv}(r)$:** In that case, $\mathrm{fv}(r) \subseteq M$ if and only if $\mathrm{fv}(r) \subseteq M'$. The conclusion holds directly from the induction hypothesis.
- **If $x \in \mathrm{fv}(r)$:** In that case, $t \in \bigcup_{x \in \mathrm{fv}(r)} \mathbb{T}_x$. We distinguish the case of send and receive:
  - **Case receive** In that case, from Lemma 55, there are $\sigma_{j-1}; \sigma_j$ in $\sigma_0; \ldots; \sigma_{i-1}$ such that $\sigma_{j-1} \xRightarrow{\_!\_\langle\_,\langle x,c_x\rangle\rangle:r'} \sigma_j$, therefore $x \in \mathrm{dom}(M')$, therefore $\mathrm{fv}(r) \subseteq \mathrm{dom}(M)$ if and only if $\mathrm{fv}(r) \subseteq \mathrm{dom}(M')$. Let M_j be the map of σ_j. From the induction hypothesis, if $\mathrm{fv}(r) \subseteq \mathrm{dom}(M')$, then $M' \models r$, therefore if $\mathrm{fv}(r) \subseteq \mathrm{dom}(M)$, then $M' \models r$.
    - **If $\forall y \in \mathrm{fv}(r) \cdot M(y) = M'(y)$:** From the rule GRRec, we therefore have that $M = M'$, therefore $\mathrm{fv}(r) \subseteq \mathrm{dom}(M)$ implies $M \models r$.
    - **If $\exists y \in \mathrm{fv}(r) \cdot M(y) \neq M'(y)$:** Without loss of generality, suppose there is a single such y. In that case, there is at least one σ_k ($k > j$, $k < i$) such that $\sigma_{k-1} \xRightarrow{\_\dagger\_\langle\_,\langle y,c_y\rangle\rangle:r_y} \sigma_k$. Without loss of generality, consider there is a single such σ_k. Let M_k and M_{k-1} be the map of σ_k and σ_{k-1} respectively.
      - **If $x = y$:** in that case: (i) $M(x) = M_j(x) = c_x$; (ii) for all $v \in \mathrm{fv}(r)$ such that $v \neq x$, $M(v) = M'(v) = M_k(v) = M_{k-1}(v) = M_j(v)$; and (iii) $\mathrm{fv}(r) \subseteq M$ if and only if $\mathrm{fv}(r) \subseteq M_j$. Therefore, $M \models r'$ if and only if $M_j \models r'$. Therefore, if $\mathrm{fv}(r) \subseteq M$, then $M \models r'$, which itself entails $M \models r$.
      - **If $x \neq y$:** in that case, $M = M'[x \mapsto c_x] = M'$. The conclusion holds directly from the induction hypothesis.
  - **Case send** In that case, $M \models r'$ entails $M \models r$.

◀

▶ **Theorem 35** (Correctness of refinement elision). *Given an RCS R containing an RCFSM $M = \langle Q, C, q_0, \mathbb{A}, \delta\rangle$, and $t = s_i \xrightarrow{\mathsf{p}\dagger\mathsf{q}\langle m\rangle:r} s_i' \in \delta$, a well-defined self-independent transition. Let $t' = s_i \xrightarrow{\mathsf{p}\dagger\mathsf{q}\langle m\rangle:\top} s_i'$; $\delta' = \delta \setminus \{t\} \cup t'$; $M' = \langle Q, C, q_0, \mathbb{A}, \delta'\rangle$; and R' be R where M is replaced with M'. If, for each transition $t_w = \_ \xrightarrow{\_!\_\langle\_\rangle:r_w} \_$ in $\bigcup_{x \in \mathrm{fv}(r)} \mathbb{T}_x$, for all map M, $M \models r_w$ entails $M \models r$, then there exists a bisimulation relating the states of R' and R.* ◁

Proof. Let $\mathcal{R}$ being the identity of reachable states of R and R'. We show that $\mathcal{R}$ is a bisimulation.

R' simulating R is trivial, since the only difference is the lack of one refinement, on transition t' w.r.t. t.

We now show that R can simulate R'. Since our candidate simulation relation $\mathcal{R}$ is the reachable state identity, we have to prove that for each R' transition $\sigma \xRightarrow{t''} \sigma'$, there exists a transition from σ to σ' in R. We distinguish two cases:

If t'' is t': In that case, we have to prove that R can take the transition $\sigma \overset{t}{\Rightarrow} \sigma'$, i.e. we have to show that $M_{\sigma'} \models r$.

First, since $\sigma \overset{t'}{\Rightarrow} \sigma'$, from the definition of GRRec or GRSnd: $\sigma' = \langle \langle \dots, s_i', \dots \rangle, \_, M_{\sigma'} \rangle$. Since σ' is reachable, from Lemma 65, if $\mathsf{fv}(r) \subseteq \mathsf{dom}(M_{\sigma'})$, then $M_{\sigma'} \models r$. Since t is well-defined, from Definition 63, $\mathsf{fv}(r) \subseteq \mathsf{dom}(M_{\sigma'})$. Therefore $M_{\sigma'} \models r$, i.e. R can take the transition $\sigma \overset{t}{\Rightarrow} \sigma'$.

If t'' is not t': In that case, $t'' \in \delta$, so R can also take the transition $\sigma \overset{t''}{\Rightarrow} \sigma'$.

◀

F.2 Application to RMPST Protocols

F.2.1 Definitions

▶ **Definition 66** (Step of a communication). *A type* $\mathsf{p} \to \mathsf{q}\{\ell_i(x_i : \mathsf{S}_i \models r_i).G_i\}_{i \in I}$ *has step* z *(noted* $\mathsf{p} \to \mathsf{q}\langle \ell, x \rangle \models r$*) if there is an* $i \in I$ *such that* $x = x_i$, $\ell = \ell_i$ *and* $r = r_i$. ◁

▶ **Definition 67** (Step in a Type). *A step* z occurs *in a type* G *if:*

- *if* $G = \mathsf{r} \to \mathsf{s}\{\ell_i(x_i : \mathsf{S}_i \models r_i).G_i\}_{i \in I}$, *either:*
  - z *is a step of* G*; or*
  - *there exists* $i \in I$ *such that* z *occurs in* G_i*;*
- *if* $G = \mu \mathbf{t}.G'$, z *occurs in* G'.

◁

▶ **Definition 68** (Happens-before in type). *Given a global type* $G = \mathsf{p} \to \mathsf{q}\{\ell_i(x_i : \mathsf{S}_i \models r_i).G_i\}_{i \in I}$, G happens-before *all* $G' = \mathsf{r} \to \mathsf{s}\{\ell_j(x_j : \mathsf{S}_j \models r_j).G_j\}_{j \in J} \in G_i$ *where* $\mathsf{r} = \mathsf{p}$ *or* $\mathsf{r} = \mathsf{q}$, *noted* $G <' G'$.

The happens-before relation (noted $<$*) is the transitive closure of* $<'$. ◁

▶ Remark 69. The *happens-before* characterises the order of the *first occurrence* of each step, in particular in recursive types, where a step can occur multiple times. ◁

▶ **Example 70** (Happens-before in a type). Considering the same G_s and G_y than in Example 37, since they both have the same sender A, G_s happens-before G_y.

▶ **Definition 71** (Well-defined step in a type). *Given two global types* G *and* G_s *such that* $G_s \in G$, *and* $z = \mathsf{p} \to \mathsf{q}\langle \ell, x \rangle \models r$ *a step of* G_s, *we say* z *is* well-defined *if for all* $x \in \mathsf{fv}(r)$, *there exists* $G_x = \mathsf{r} \to \mathsf{s}\{\_(x_i : \_ \models \_).G_i\}_{i \in I}$ *such that,* $G_x < G_s$ *and for one* $i \in I$, $G_s \in G_i$ *and* $x = x_i$. ◁

F.2.2 Lemmata

▶ **Lemma 72.** *Given a projectable type* G, *if a step* $z = \mathsf{p} \to \mathsf{q}\langle \ell, x \rangle \models r$ *occurs in* G, *then there exists a transition* $s \xrightarrow{\mathsf{p}!\mathsf{q}\langle \ell, \langle x, \_ \rangle \rangle : r} s'$ *in* $\mathcal{A}(G\upharpoonright_{\mathsf{p}})$, *and* $s \xrightarrow{\mathsf{p}?\mathsf{q}\langle \ell, \langle x, \_ \rangle \rangle : \top} s'$ *in* $\mathcal{A}(G\upharpoonright_{\mathsf{q}})$. ◁

Proof. We prove the case for $s \xrightarrow{\mathsf{p}!\mathsf{q}\langle \ell, \langle x, \_ \rangle \rangle : r} s'$. The case for $s \xrightarrow{\mathsf{p}?\mathsf{q}\langle \ell, \langle x, \_ \rangle \rangle : \top} s'$ is similar.

By structural induction on G:

- if $G = \mathsf{r} \to \mathsf{s}\{\ell_i(x_i : \mathsf{S}_i \models r_i).G_i\}_{i \in I}$

- $\mathsf{r} = \mathsf{p}$, $\mathsf{s} = \mathsf{q}$, and there exist $i \in I$ such that $\ell = \ell_i$, $x = x_i$, and $r = r_i$: then, by definition of projection (Definition 20), $G\upharpoonright_{\mathsf{p}} = \mathsf{q}\&\{\ell_i(x_i : \mathsf{S}_i \models r_i).G_i\}_{i\in I}$. From Definition 22, we have that $G\upharpoonright_{\mathsf{p}} \xrightarrow{\mathsf{p}!\mathsf{q}\langle \ell, \langle x, \_\rangle\rangle : r} \_$. (Notice that the final state is not necessarily $G_i\upharpoonright_{\mathsf{p}}$, as there are possibly recursions.)
- if there exist $i \in I$ such that z occurs in G_i: direct, from the induction hypothesis.

- if $G = \mu \mathbf{t}.G'$ and s occurs in G': direct, from the induction hypothesis.

◀

▶ **Lemma 73** (Converse of Lemma 72). *For all types G and all participants p of G, if there exists a transition $s \xrightarrow{\mathsf{p}!\mathsf{q}\langle \ell, \langle x, \_\rangle\rangle : r} s'$ in $\mathcal{A}(G\upharpoonright_{\mathsf{p}})$, then a step $\mathsf{p} \to \mathsf{q}\langle \ell, x\rangle \models r$ occurs in G.* ◁

Proof. From Definition 22, if $s \xrightarrow{\mathsf{p}!\mathsf{q}\langle \ell, \langle x, \_\rangle\rangle : r} s'$ is in the set of transitions of $\mathcal{A}(G\upharpoonright_{\mathsf{p}})$, there exist some T_1 and T_2 in $\mathcal{A}(G\upharpoonright_{\mathsf{p}})$ such that $T_1 = \mathsf{q}\oplus\{\ell_i(x_i : \mathsf{S}_i \models r_i).T_i\}_{i\in I}$ and $T_2 = T_i$ for some $i \in I$. From Definition 20, $T_1 \in G\upharpoonright_{\mathsf{p}}$ only if there exist $G_1 = \mathsf{p} \to \mathsf{q}\{\ell_i(x_i : \mathsf{S}_i \models r_i).G_i\}_{i\in I}$ in G, where there is one $i \in I$ such that $\ell = \ell_i$, $x = x_i$, and $r = r_i$. Therefore, from Definition 67, $\mathsf{p} \to \mathsf{q}\langle \ell, x\rangle \models r$ occurs in G. ◀

▶ **Lemma 74.** *Let $G_1 = \mathsf{p} \to \mathsf{q}\{\ell_i(x_i : \mathsf{S}_i \models r_i).G_i\}_{i\in I}$. For any global type $G_2 \in G_i$, such that $G_1 < G_2$ and such that r is the sender of G_2. All paths p to a state $\langle\langle \ldots, s_{\mathsf{r}}, \ldots\rangle, w, M\rangle$, where $s_{\mathsf{r}} = G_2\upharpoonright_{\mathsf{r}}$ contain a transition $t = \sigma \xRightarrow{\mathsf{p}!\mathsf{q}\langle \ell_i, \langle x_i, \_\rangle\rangle : r_i} \sigma'$.* ◁

Proof. Since $<$ is the transitive closure of $<'$, we prove the result for $<'$, the general case is then direct by induction.

In that case, $G_1 = \mathsf{p} \to \mathsf{q}\{\ell_i(x_i : \mathsf{S}_i \models r_i).G_i\}_{i\in I}$ and $G_2 = \mathsf{r} \to \mathsf{s}\{\ell_j(x_j : \mathsf{S}_j \models r_j).G_j\}_{j\in J}$ ($\mathsf{r} = \mathsf{p}$ or $\mathsf{r} = \mathsf{q}$) and $G_2 \in G_i$ for some i.

Case $\mathsf{r} = \mathsf{p}$: $G_1\upharpoonright_{\mathsf{p}} = \mathsf{q}\oplus\{\ell_i(x_i : \mathsf{S}_i \models r_i).L_i\}_{i\in I}$, where each $L_i = G_i\upharpoonright_{\mathsf{p}}$. Since $G_2\upharpoonright_{\mathsf{p}}$ only appears in one of L_i, then all paths to $G_2\upharpoonright_{\mathsf{p}}$ in $\mathcal{A}(G\upharpoonright_{\mathsf{p}})$ contain a transition to the state $L_i\upharpoonright_{\mathsf{p}}$, which is only reachable with a transition $G_1\upharpoonright_{\mathsf{p}} \xrightarrow{\mathsf{p}?\mathsf{q}\langle \ell_i, \langle x_i, \_\rangle\rangle : \top} L_i$. The result then follows directly from the global reduction rules.

Case $\mathsf{r} = \mathsf{q}$: Therefore, $G_1\upharpoonright_{\mathsf{q}} = \mathsf{p}\&\{\ell_i(x_i : \mathsf{S}_i \models r_i).L_i\}_{i\in I}$ (with $L_i = G_i\upharpoonright_{\mathsf{q}}$), and $G_2\upharpoonright_{\mathsf{q}} \in L_i$. Since labels are uniquely used, $G_2\upharpoonright_{\mathsf{q}}$ only appears in L_i. Therefore, all paths to $G_2\upharpoonright_{\mathsf{q}}$ in $\mathcal{A}(G\upharpoonright_{\mathsf{q}})$ contain a transition $G_1\upharpoonright_{\mathsf{q}} \xrightarrow{\mathsf{p}?\mathsf{q}\langle \ell_i, \langle x_i, \_\rangle\rangle : \top} L_i$. The result then follows directly from the global reduction rules.

◀

▶ **Lemma 75** (Well-defined steps imply well-defined transitions). *For all projectable types G, for all well-defined steps $z = \mathsf{p} \to \mathsf{q}\langle \ell, x\rangle \models r$ that occur in G, the transition $s_{\mathsf{p}} \xrightarrow{\mathsf{p}!\mathsf{q}\langle \ell, \langle x, \_\rangle\rangle : r} s_{\mathsf{p}}'$ in $\mathcal{A}(G\upharpoonright_{\mathsf{p}})$ is well-defined in $\mathcal{S}(G)$.* ◁

Proof. Given a well-defined step $z = \mathsf{p} \to \mathsf{q}\langle \ell, x\rangle \models r$, without loss of generality, consider r has a single free variable y.

Since z is a well-defined step occurring in G, let $G_s \in G$ the type z is the step of. From Lemma 72, there exist $t = s_{\mathsf{p}} \xrightarrow{\mathsf{p}!\mathsf{q}\langle \ell, \langle x, \_\rangle\rangle : r} s_{\mathsf{p}}'$ in $\mathcal{A}(G\upharpoonright_{\mathsf{p}})$. From Definition 71, there exists $G_y = \mathsf{r} \to \mathsf{s}\{\_(y_i : \_ \models \_).G_i\}_{i\in I} \in G$ such that, $G_y < G_s$ and, for one $i \in I$, $G_s \in G_i$ and $y = y_i$.

By contradiction, suppose transition t is not well-defined in $\mathcal{S}(G)$, i.e. there is a reachable state $\sigma_{\mathsf{p}} = \langle\langle \ldots, s_{\mathsf{p}}, \ldots\rangle, \_, M_{\mathsf{p}}\rangle$ such that y is not in the map M_{p}.

Consider a path p to s_p. From Lemma 74, there is a transition $\sigma \xRightarrow{\mathsf{r}!\mathsf{s}\langle \_, \langle y, \_\rangle\rangle : \_} \sigma'$ with $\sigma' = \langle \_, \_, M_{\sigma'} \rangle$ in the path p. From the premises of rule GRSND, $y \in \mathsf{dom}(M_{\sigma'})$. Since variables cannot be removed from the map, M_p contains y. Contradiction. ◀

▶ **Theorem 39** (Static elision of redundant refinements in types). *Given two a global types G and $G_s = \mathsf{p} \to \mathsf{q}\{\ell_i(x_i : \mathsf{S}_i \models r_i).G_i\}_{i \in I} \in G$, such that, for one $t \in I$, $\mathsf{p} \to \mathsf{q}\langle \ell_t, x_t \rangle \models r_t$ is a well-defined step with $x_t \notin \mathsf{fv}(r_t)$. Let $\ell_{t'} = \ell_t$, $x_{t'} = x_t$, $\mathsf{S}_{t'} = \mathsf{S}_t$, $r'_t = \top$, $G_{t'} = G_t$, $G_{s'} = \mathsf{p} \to \mathsf{q}\{\ell_i(x_i : \mathsf{S}_i \models r_i).G_i\}_{i \in I \setminus \{t\} \cup \{t'\}}$; and G' be G where G_s is replaced with $G_{s'}$. If, for all steps, $\mathsf{r} \to \mathsf{s}\langle \_, x_w \rangle \models r_w$ occurring in G (for each $x \in \mathsf{fv}(r)$), $M \models r_w$ entails $M \models r$ (for all M), there exists a bisimulation between the states of $\mathcal{S}(G)$ and those of $\mathcal{S}(G')$.* ◁

Proof. We prove this by showing that Theorem 35 applies to $\mathcal{S}(G)$ and $\mathcal{S}(G')$.

First, since G' is G where G_s is replaced with $G_{s'}$, i.e. the only difference is that r_i is replaced with $\top$, then the RCFSM of p in $\mathcal{S}(G)$ contains $s_\mathsf{p} \xrightarrow{\mathsf{p}\dagger\mathsf{q}\langle m \rangle : r} s_\mathsf{p}'$, and $\mathcal{S}(G')$ contains $s_\mathsf{p} \xrightarrow{\mathsf{p}\dagger\mathsf{q}\langle m \rangle : \top} s_\mathsf{p}'$. Also, from Lemma 75, $s_\mathsf{p} \xrightarrow{\mathsf{p}\dagger\mathsf{q}\langle m \rangle : r} s_\mathsf{p}'$ is well-defined in $\mathcal{S}(G)$, and since $x_i \notin \mathsf{fv}(r_i)$, it is also self-independent. This is the only change and the RCFSM of all others participants are unchanged.

Second, we have to show that for each transition $t_w = r \xrightarrow{\_!\_\langle \_, x_w \rangle : r_w} s$ in $\bigcup_{x \in \mathsf{fv}(r)} \mathbb{T}_x$, for all map M, $M \models r_w$ entails $M \models r$. By contradiction, suppose there exists such a transition such that for all map M, $M \models r_w$ does not entails $M \models r$. From Lemma 73, if such transition exists, a step $\mathsf{r} \to \mathsf{s}\langle \_, x_w \rangle \models r_w$ occurs in G. By hypothesis, $M \models r_w$ entails $M \models r$. Contradiction.

Therefore, from Theorem 35, there is a bisimulation between the states of $\mathcal{S}(G)$ and those of $\mathcal{S}(G')$. ◀

G RCS from Choreography Automata

Despite those differences, with our framework, we can adapt CA to accommodate refinements, which we call *Refined Choreography Automata* (RCA). In this paragraph, in order to show the versatility of our framework and to show that the source formalism is not bound to RMPST only, we informally present how to obtain RCS from a *refined* variant of asynchronous *choreography automata* [1], obtaining correctness result w.r.t. valid refined traces (Definition 7).

In order to introduce refinements into CA, we can simply adapt the transition labels in the CA, to add payloads and predicates. Adapting RCA projection would then return RCFSMs, and we would obtain RCS with our framework, as with RMPST. To illustrate this, we expand on (1) p. 87 in [1]. In this example, a client C requests an entry from a server S. A logger L is used to log the information. After receiving the result, C can ask S to refine the result (*ref*), to restart the protocol (*ok*), or to quit (*bye*). In the original protocol, C performs the choice. Our goal is to let S decide whether C is allowed to continue, or if C must terminate the protocol (in state 3). For that, S can embed a boolean variable q in its communications, which is then tested in the outgoing transitions of state 3 (notice that C can always decide to quit, independently of q, so we only test it for the transitions *ok* and *ref*). For the sake of clarity, we only show attached variables and predicates when they are needed.

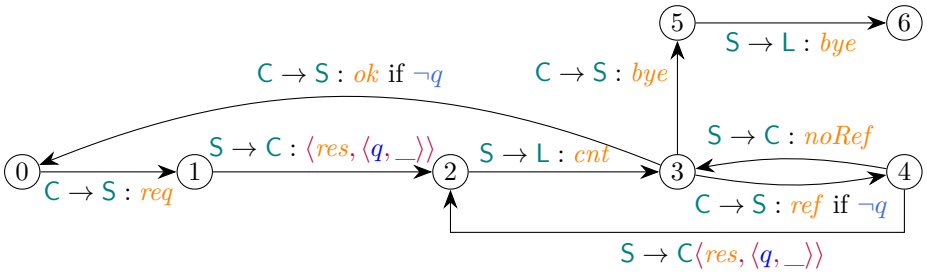

By preserving predicates, the projection onto CFSM is adapted to accommodate newly added refinements. We therefore obtain RCFSMs, which we compose into RCS. As an illustration, the RCFSM of participant C would be the following:

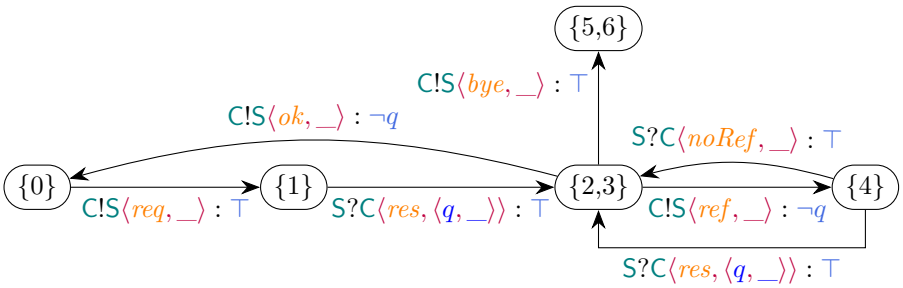

Similarly, we can obtain the RCFSM of S and L. Thanks to Theorem 18, we know that traces produced are valid refined traces.

H Artifact

The submitted artifact contains scripts to reproduce the benchmark results. In order to run the benchmarks, the following is required (version numbers indicate the version we tested our artifact on, we expect it to work on newer version, although not tested):

- a linux system (tested on Debian)
- an internet connection
- an Ocaml compiler (4.11.2) release with Dune (3.6.1) and Opam (2.0.8).
- Regarding Opam, one needs to install:
  - the ocamlgraph (2.0.0) library:
    `opam install ocamlgraph`
  - a specific branch of νScr:
    ```
    git clone -b https://github.com/Bromind/nuscr/tree/develop-refinements
    cd nuscr
    opam pin add nuscr.dev-refinements-local .
    ```
- a Rust compiler (rustc 1.67.0-nightly (e9493d63c 2022-11-16)) with Cargo (rustc 1.67.0-nightly (e9493d63c 2022-11-16)).
- Python3 (3.9.2) with the libraries: argparse, json, unittest, numpy, statistics, matplotlib, and scipy.

The artifact is composed of 2 parts: one for Table 2 and one for Table 3. Each part is covered by a shell script, which manages the *download*, the *preparation* and the *cleanup* of the system.

H.1 Refined Rumpsteak v. Vanilla Rumpsteak

The script `compare.sh` performs the comparision between Rumpsteak with refinements and without refinements. The scripts prints detailled results on `stdout`. Each microbenchmark begins with:

```
********* Analysis NAME_OF_BENCHMARK **********
```

In Table 3, we report the p-value printed on the line `MannwhitneyuResult(statistic=..., pvalue=...)` In addition, the scripts produces a folder in `/tmp/Rumpsteak_benchmarks.XXX` (where `XXX` is a random sequence of three characters) which contains records of each run of each microbenchmark. The runtimes reported in Table 3 are the median runtime for both the `vanilla` and `refinements` measurements, which are shown, for each benchmark after the lines:

***** 1st quartile, median, 3rd quartile (vanilla) *******

and

***** 1st quartile, median, 3rd quartile (refinements) *******

H.2 Evaluation of the localisation algorithm

The script `dynamic_verify.sh` measures the runtime of the localisation tools. After setting up the tools, the scripts runs the benchmarks and prints results on `stdout`. Each microbenchmark begins with:

```
********** NAME_OF_BENCHMARK ***********
```

For each benchmark, we output a result in the form:

```
Time (mean ± σ):  22.4 ms ± 0.8 ms [User:  10.5 ms, System:  11.4 ms]
```

Results reported in Table 2 are the printed mean and standard deviation as output by the script.