

1 Asynchronous Multiparty Session Type 2 Implementability is Decidable – 3 Lessons Learned from Message Sequence Charts

4 Felix Stutz  

5 MPI-SWS, Kaiserslautern, Germany

6 — Abstract —

7 Multiparty session types (MSTs) provide efficient means to specify and verify asynchronous message-
8 passing systems. For a global type, which specifies all interactions between roles in a system,
9 the implementability problem asks whether there are local specifications for all roles such that
10 their composition is deadlock free and generates precisely the specified executions. Decidability of
11 the implementability problem is an open question. We answer it positively for global types with
12 generalised choice that allow a sender to send to different receivers and a receiver to receive from
13 different senders upon branching. To achieve this, we generalise results from the domain of high-level
14 message sequence charts (HMSCs). This connection also allows us to comprehensively investigate
15 how HMSC techniques can be adapted to the MST setting. This comprises techniques to make the
16 problem algorithmically more tractable as well as a variant of implementability which may open new
17 design space for MSTs. Inspired by potential performance benefits, we introduce a generalisation of
18 the implementability problem that we, unfortunately, prove to be undecidable.

19 **2012 ACM Subject Classification** Theory of computation → Concurrency

20 **Keywords and phrases** Multiparty session types, Verification, Message sequence charts

21 **Digital Object Identifier** 10.4230/LIPIcs.CVIT.2016.23

22 **1** Introduction

23 Distributed message-passing systems are omnipresent and, therefore, designing and imple-
24 menting them correctly is very important. However, this is a very difficult task at the same
25 time. In fact, it is well-known that the verification problem is algorithmically undecidable in
26 general due to the combination of asynchrony (messages are buffered) and concurrency [15].

27 Multiparty Session Type (MST) frameworks provide efficient means to specify and verify
28 such distributed message-passing systems. MSTs (and their binary counterpart) are not
29 only of theoretical interest but have been implemented for many mainstream programming
30 languages [6, 54, 62, 58, 74, 70, 25]. They have also been applied to various other domains
31 like operating systems [36], cyber-physical systems [65], timed systems [11], distributed
32 algorithms [57], web services [86], and smart contracts [33]. In MST frameworks, global types
33 are global specifications, which comprise all interactions between roles in a protocol. From a
34 design perspective, it makes sense to start with such a global protocol specification — instead
35 of a system with arbitrary communication between roles and a specification to satisfy.

36 Let us consider a variant of the well-known two buyer protocol from the MST literature,
37 e.g., [75, Fig. 4 (2)]. Two Buyers *a* and *b* purchase a sequence of items from Seller *s*. We
38 informally describe the protocol and *emphasise* the interactions. At the start and after
39 every purchase (attempt), Buyer *a* can decide whether to buy the next item or whether they
40 are *done*. For each item, Buyer *a* *queries* its price and the Seller *s* replies with the *price*.
41 Subsequently, Buyer *a* decides whether to *cancel* the purchase process for the current item
42 or proposes to *split* to Buyer *b* that can *accept* or *reject*. In both cases, Buyer *a* notifies the
43 Seller *s* whether they want to *buy* the item or *not*. This protocol can be specified with the
44 following global type:



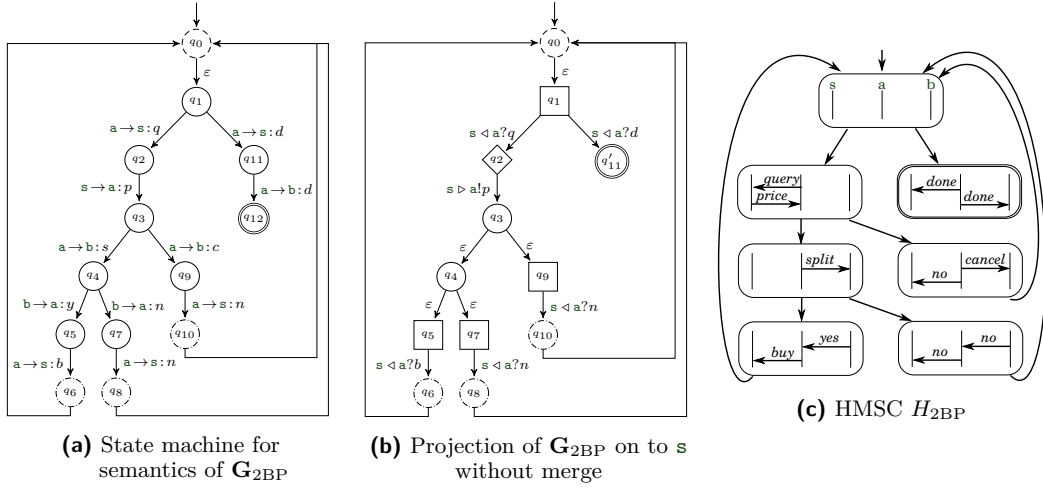
© Felix Stutz;
licensed under Creative Commons License CC-BY 4.0
42nd Conference on Very Important Topics (CVIT 2016).

Editors: John Q. Open and Joan R. Access; Article No. 23; pp. 23:1–23:43



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 1** Two Buyer Protocol: the finite state machine for the semantics of \mathbf{G}_{2BP} on the left, the first step of projection in the middle, and as HMSC on the right; a transition label $a \rightarrow s : q$ jointly specifies a send event $a \triangleright s ! q$ for Buyer a and a receive event $s \triangleleft a ? q$ for Seller s ; styles of states indicate their kind, e.g., recursion states (dashed lines) while final states have double lines

$$\mathbf{G}_{2BP} := \mu t. \left\{ \begin{array}{l} a \rightarrow s : \text{query}. s \rightarrow a : \text{price}. + \left\{ \begin{array}{l} a \rightarrow b : \text{split}. (b \rightarrow a : \text{yes}. a \rightarrow s : \text{buy}. t + b \rightarrow a : \text{no}. a \rightarrow s : \text{no}. t) \\ a \rightarrow b : \text{cancel}. a \rightarrow s : \text{no}. t \end{array} \right. \\ a \rightarrow s : \text{done}. a \rightarrow b : \text{done}. 0 \end{array} \right.$$

The first term μt binds the recursion variable t which is used at the end of the first two lines and allows the protocol to recurse back to this point. Subsequently, $+$ and the curly bracket indicate a choice that is taken by Buyer a as it is the sender for the next interaction, e.g., $a \rightarrow s : \text{query}$. For our asynchronous setting, this term jointly specifies the send event $a \triangleright s ! \text{query}$ for Buyer a and its corresponding receive event $s \triangleleft a ? \text{query}$ for Seller s , which may happen with arbitrary delay. The state machine in Figure 1a illustrates its semantics.

The Implementability Problem for Global Types and the MST Approach

A global type provides a global view of the intended protocol. However, when implementing a protocol in a distributed setting, one needs a local specification for each role. The *implementability problem* for a global type asks whether there are local specifications for all roles such that, when complying with their local specifications, their composition never gets stuck and exposes the same executions as specified by the global type. This is a challenging problem because roles can only partially observe the execution of a system: each role only knows the messages it sent and received and, in an asynchronous setting, a role does not know when one of its messages will be received by another role.

In general, one distinguishes between a role in a protocol and the process which implements the local specification of a role in a system. We use the local specifications directly as implementations so the difference is not essential and we use the term role instead of process.

Classical MST frameworks employ a partial *projection operator* with an in-built *merge operator* to solve the implementability problem. For each role, the projection operator takes the global type and removes all interactions the role is not involved in. Figure 1a illustrates the semantics of \mathbf{G}_{2BP} while Figure 1b gives the projection on to Seller s before the merge operator is applied — in both, messages are abbreviated with their first letter. It is easy to see that this introduces non-determinism, e.g., in q_3 and q_4 , which shall be resolved by the merge operator. Most merge operators can resolve the non-determinism in Figure 1b. A merge operator checks whether it is safe to merge the states and it might fail so it is

72 a partial operation. For instance, every kind of state, indicated by a state's style in Figure 1b,
 73 can only be merged with states of the same kind and states of circular shape. For a role, the
 74 result of the projection, if defined, is a local type. They act as local specifications and their
 75 syntax is similar to the one of global types.

76 Classical projection operators are a best-effort technique. This yields good (mostly
 77 linear) worst-case complexity but comes at the price of rejecting implementable global types.
 78 Intuitively, classical projection operators consider a limited search space for local types. They
 79 bail out early when encountering difficulties and do not unfold recursion. In addition, most
 80 MST frameworks do effectively not allow a role to send to different receivers or receive from
 81 different senders upon branching. This restriction is called *directed choice* — in contrast to
 82 *generalised choice* which allows such patterns. Among the classical projection operators, the
 83 one by Majumdar et al. [64] is the only to handle global types with *generalised choice* but
 84 suffers from the shortcomings of a classical projection approach. We define different merge
 85 operators from the literature and visually explain their supported features by example. We
 86 show that the presented classical projection/merge operators fail to project implementable
 87 variations of the two buyer protocol. This showcases the sources of incompleteness for the
 88 classical projection approach. For non-classical approaches, we refer to Section 7.

89 As a best-effort technique, it is natural to focus on efficiency rather than completeness. The
 90 work by Castagna et al. [19] is a notable exception even though their notion of completeness [19,
 91 Def. 4.1] is not as strict as the one considered in this work and only a restricted version of
 92 their characterisation is algorithmically checkable. In general, it is not known whether the
 93 implementability problem for global types, with directed or generalised choice, is decidable.
 94 We answer this open question positively for global types with generalised choice. To this end,
 95 we relate the implementability problem for global types with the safe realisability problem
 96 for high-level message sequence charts and generalise results for the latter.

97 **Lessons Learned from Message Sequence Charts**

98 The two buyer protocol \mathbf{G}_{2BP} can also be specified as high-level message sequence chart
 99 (HMSC). It is illustrated in Figure 1c. Each block is a basic message sequence chart (BMS
 100 which intuitively corresponds to straight-line code. In each of those, time flows from top to
 101 bottom and each role is represented by a vertical line. We only give the names in the initial
 102 block, which is marked by an incoming arrow at the top. An arrow between two role lines
 103 specifies sending and receiving a message with the corresponding label. The graph structure
 104 adds branching, which corresponds to choice in global types, and control flow. Top branches
 105 from the global type are on the left in the HMSC while bottom branches are on the right.

106 While research on MSTs and HMSCs has been pursued quite independently, the MST
 107 literature frequently uses HMSC-like visualisations for global types, e.g., [18, Fig. 1] and [49,
 108 Figs. 1 and 2]. The first formal connection was recently established by Stutz and Zufferey [77].

109 The HMSC approach to the implementability problem, studied as safe realisability, differs
 110 from the MST approach of checking conditions during the projection. For an HMSC, it is
 111 known that there is a candidate implementation [3], which implements the HMSC if it is
 112 implementable. Intuitively, one takes the HMSC and removes all interactions a role is not
 113 involved in and determinises the result. We generalise their result to infinite executions.¹

114 Hence, algorithms and conditions center around checking implementability of HMSCs. In
 115 general, this problem is undecidable [63]. For *globally-cooperative* HMSCs [39], Lohrey [63]
 116 proved it to be EXPSpace-complete. We show that any implementable global type naturally
 117 belongs to this class of HMSCs¹ which is far from trivial. These results give rise to the

¹ For this, we impose a mild assumption: all protocols can (but do not need to) terminate.

118 following algorithm to check implementability of a global type. One can check whether a
 119 global type is globally-cooperative (which is equivalent to checking its HMSC encoding). If
 120 it is not globally-cooperative, it cannot be implementable. If it is globally-cooperative, we
 121 apply the algorithm by Lohrey [63] to check whether its HMSC encoding is implementable. If
 122 it is, we use its candidate implementation and know that it generalises to infinite executions.

123 While this algorithm shows decidability, the complexity might not be tractable. Based
 124 on our results, we show how more tractable but still permissive approaches to check imple-
 125 mentability of HMSCs can be adapted to the MST setting. In addition, we consider *payload*
 126 *implementability*, which allows to add payload to messages of existing interactions and checks
 127 agreement when the additional payload is ignored. We present a sufficient condition for
 128 global types that implies payload implementability. These techniques can be used if the
 129 previous algorithms are not tractable or reject a global type.

130 Furthermore, we introduce a generalisation of the implementability problem. A network
 131 may reorder messages from different senders for the same receiver but the implementability
 132 problem still requires the receiver to receive them in the specified order. Our generalisation
 133 allows to consider such reorderings of arrival and can yield performance gains. In addition, it
 134 also renders global types implementable that are not implementable in the standard setting.
 135 Unfortunately, we prove it to be undecidable in general.

136 Contributions and Outline

137 We introduce our MST framework in Section 2 while Section 7 covers details on related work.
 138 In the other sections, we introduce the necessary concepts to establish our main *contributions*:

- 139 ■ We give a visual explanation of the classical projection operator with different merge
 140 operators and exemplify its shortcomings (Section 3).
- 141 ■ We prove decidability of the implementability problem for global types with generalised
 142 choice (Section 4) — provided that protocols can (but do not need to) terminate.
- 143 ■ We comprehensively investigate how MSC techniques can be applied to the MST setting,
 144 including algorithmics with better complexity for subclasses as well as an interesting
 145 variant of the implementability problem (Section 5).
- 146 ■ Lastly, we introduce a new variant of the implementability problem with a more relaxed
 147 role message ordering, which is closer to the network ordering, and prove it to be
 148 undecidable in general (Section 6).

149 2 Multiparty Session Types

150 In this section, we formally introduce our Multiparty Session Type (MST) framework. We
 151 define the syntax of global and local types and their semantics. Subsequently, we recall the
 152 implementability problem for global types which asks if there is a deadlock free communicating
 153 state machine that admits the same language (without additional synchronisation).

154 **Finite and Infinite Words.** Let Σ be an alphabet. We denote the set of finite words over Σ
 155 by Σ^* and the set of infinite words by Σ^ω . Their union is denoted by Σ^∞ . For two strings
 156 $u \in \Sigma^*$ and $v \in \Sigma^\infty$, we say that u is a *prefix* of v if there is some $w \in \Sigma^\infty$ such that $u \cdot w = v$
 157 and denote this with $u \leq v$. For a language $L \subseteq \Sigma^\infty$, we distinguish between the language of
 158 finite words $L_{\text{fin}} := L \cap \Sigma^*$ and the language of infinite words $L_{\text{inf}} := L \cap \Sigma^\omega$.

159 **Message Alphabet.** We fix a finite set of messages \mathcal{V} and a finite set of roles \mathcal{P} , ranged
 160 over with p, q, r , and s . With $\Sigma_{\text{sync}} = \{p \rightarrow q : m \mid p, q \in \mathcal{P} \text{ and } m \in \mathcal{V}\}$, we denote
 161 the set of interactions where sending and receiving a message is specified at the same
 162 time. For our asynchronous setting, we also define individual send and receive events:
 163 $\Sigma_p = \{p \triangleright q!m, p \triangleleft q?m \mid q \in \mathcal{P}, m \in \mathcal{V}\}$ for a role p . For both send events $p \triangleright q!m$ and receive

164 events $p \triangleleft q?m$, the first role is *active*, i.e., the sender in the first event and the receiver in
 165 the second one. The union for all roles yields all (asynchronous) events: $\Sigma = \bigcup_{p \in \mathcal{P}} \Sigma_p$. For
 166 the rest of this work, we fix the set of roles \mathcal{P} , the messages \mathcal{V} , and both sets Σ_{sync} and Σ .
 167 We may also use the term Σ_{async} for Σ . We define an operator that splits events from Σ_{sync} ,
 168 $\text{split}(p \rightarrow q:m) := p \triangleright q!m. q \triangleleft p?m$, which is lifted to sequences and languages as expected.
 169 Given a word, we might also project it to all letters of a certain shape. For instance, $w \Downarrow_{p \triangleright q! _}$
 170 is the subsequence of w with all of its send events where p sends any message to q . If we
 171 want to select all messages of w , we write $\mathcal{V}(w)$.

172 Global and Local Types – Syntax

173 We give the syntax of global and local types following work by Majumdar et al. [64], Honda
 174 et al. [48], Hu and Yoshida [50], as well as Scalas and Yoshida [75]. In this work, we
 175 consider global types as specifications for message-passing concurrency and omit features
 176 like delegation.

177 ► **Definition 2.1** (Syntax of global types). Global types for MSTs are defined by the grammar:

$$178 \quad G ::= 0 \mid \sum_{i \in I} p \rightarrow q_i : m_i . G_i \mid \mu t . G \mid t$$

180 The term 0 explicitly represents termination. A term $p \rightarrow q_i : m_i$ indicates an interaction
 181 where p sends message m_i to q_i . In our asynchronous semantics, it is split into a send event
 182 $p \triangleright q_i!m_i$ and a receive event $q_i \triangleleft p?m_i$. In a choice $\sum_{i \in I} p \rightarrow q_i : m_i . G_i$, the sender p chooses
 183 the branch. We require choices to be unique, i.e., $\forall i, j \in I. i \neq j \Rightarrow q_i \neq q_j \vee m_i \neq m_j$.
 184 If $|I| = 1$, which means there is no actual choice, we omit the sum operator. The operators
 185 μt and t allow to encode loops. We require them to be guarded, i.e., there must be at least
 186 one interaction between the binding μt and the use of the recursion variable t . Without loss
 187 of generality, all occurrences of recursion variables t are bound and distinct.

188 Our definition allows *generalised choice* as p can send to different receivers upon branching:
 189 $\sum_{i \in I} p \rightarrow q_i : m_i . G_i$. In contrast, *directed choice* requires a sender to send to a single receiver,
 190 i.e., $\forall i, j \in I. q_i = q_j$.

191 ► **Example 2.2** (Global types). The two buyer protocol \mathbf{G}_{2BP} from the introduction is a
 192 global type. Instead of \sum , we use $+$ with curly brackets for readability.

193 ► **Definition 2.3** (Syntax of local types). For a role p , the local types are defined as follows:

$$194 \quad L ::= 0 \mid \bigoplus_{i \in I} q_i!m_i . L_i \mid \&_{i \in I} q_i?m_i . L_i \mid \mu t . L \mid t$$

195 We call $\bigoplus_{i \in I} q_i!m_i$ an *internal choice* while $\&_{i \in I} q_i?m_i$ is an *external choice*. For both, we
 196 require the choice to be unique, i.e., $\forall i, j \in I. i \neq j \Rightarrow (q_i, m_i) \neq (q_j, m_j)$. Similarly to global
 197 types, we may omit \bigoplus or $\&$ if there is no actual choice and we require recursion to be guarded
 198 as well as recursion variables to be bound and distinct.

199 ► **Example 2.4** (Local type). For the global type \mathbf{G}_{2BP} , a local type for Seller s is

$$200 \quad \mu t . \& \begin{cases} a?query. a!price. (a?buy. t \& a?no. t) \\ a?done. 0 \end{cases}$$

201 Implementing in a Distributed Setting

202 Global types can be thought of as global protocol specifications. Thus, a natural question and
 203 a main concern in MST theory is whether a global type can be implemented in a distributed
 204 setting. We present communicating state machines, which are built from finite state machines,
 205 as the standard implementation model.

206 ► **Definition 2.5** (State machines [77]). A state machine $A = (Q, \Delta, \delta, q_0, F)$ is a 5-tuple with
 207 a finite set of states Q , an alphabet Δ , a transition relation $\delta \subseteq Q \times (\Delta \cup \{\varepsilon\}) \times Q$, an initial
 208 state $q_0 \in Q$ from the set of states, and a set of final states F with $F \subseteq Q$. If $(q, a, q') \in \delta$,
 209 we also write $q \xrightarrow{a} q'$. A sequence $q_0 \xrightarrow{w_0} q_1 \xrightarrow{w_1} \dots$, with $q_i \in Q$ and $w_i \in \Delta \cup \{\varepsilon\}$ for
 210 $i \geq 0$, such that q_0 is the initial state, and for each $i \geq 0$, it holds that $(q_i, w_i, q_{i+1}) \in \delta$, is
 211 called a run in A with its trace $w_0 w_1 \dots \in \Delta^\infty$. A run is maximal if it ends in a final state
 212 or is infinite. The language $\mathcal{L}(A)$ of A is the set of traces of all maximal runs. If Q is finite,
 213 we say A is a finite state machine (FSM).

214 ► **Definition 2.6** (Communicating state machines [77]). We call $\mathcal{A} = \{\{A_p\}_{p \in \mathcal{P}}\}$ a communi-
 215 cating state machine (CSM) over \mathcal{P} and \mathcal{V} if A_p is a finite state machine with alphabet Σ_p
 216 for every $p \in \mathcal{P}$. The state machine for p is denoted by $(Q_p, \Sigma_p, \delta_p, q_{0,p}, F_p)$. Intuitively, a
 217 CSM allows a set of state machines, one for each role in \mathcal{P} , to communicate by sending and
 218 receiving messages. For this, each pair of roles $p, q \in \mathcal{P}$, $p \neq q$, is connected by two directed
 219 message channels. A transition $q_p \xrightarrow{p \triangleright q!m} q'_p$ in the state machine of p denotes that p sends
 220 message m to q if p is in the state q_p and changes its local state to q'_p . The channel $\langle p, q \rangle$
 221 is appended by message m . For receptions, a transition $q_q \xrightarrow{q \triangleleft p?m} q'_q$ in the state machine
 222 of q corresponds to q retrieving the message m from the head of the channel when its local
 223 state is q_q which is updated to q'_q . The run of a CSM always starts with empty channels and
 224 each finite state machine is in its respective initial state. A deadlock of $\{\{A_p\}_{p \in \mathcal{P}}\}$ is the last
 225 configuration of a finite run for which cannot be extended with \rightarrow . The formalisation of this
 226 intuition is standard and can be found in Appendix A.1.

227 A global type always specifies send and receive events together. In a CSM execution, there
 228 may be independent events that can occur between a send and its respective receive event.

229 ► **Example 2.7** (Motivation for indistinguishability relation \sim). Let us consider the following
 230 global type which is a part of the two buyer protocol: $a \rightarrow b : \text{cancel}. a \rightarrow s : \text{no}. 0$. This is
 231 one of its traces: $a \triangleright b! \text{cancel}. b \triangleleft a? \text{cancel}. a \triangleright s! \text{no}. s \triangleleft a? \text{no}$. Because the active roles in
 232 $b \triangleleft a? \text{cancel}$ and $a \triangleright s! \text{no}$ are different and we do not reorder a receive event in front of its
 233 respective send event, any CSM that accepts the previous trace also accepts the following
 234 trace: $a \triangleright b! \text{cancel}. a \triangleright s! \text{no}. b \triangleleft a? \text{cancel}. s \triangleleft a? \text{no}$.

235 Majumdar et al. [64] introduced the following relation to capture this phenomenon.

236 ► **Definition 2.8** (Indistinguishability relation \sim [64]). We define a family of indistinguishability
 237 relations $\sim_i \subseteq \Sigma^* \times \Sigma^*$, for $i \geq 0$ as follows. For all $w \in \Sigma^*$, we have $w \sim_0 w$. For $i = 1$,
 238 we define:

- 239 1. If $p \neq r$, then $w.p \triangleright q!m.r \triangleright s!m'.u \sim_1 w.r \triangleright s!m'.p \triangleright q!m.u$.
- 240 2. If $q \neq s$, then $w.q \triangleleft p?m.s \triangleleft r?m'.u \sim_1 w.s \triangleleft r?m'.q \triangleleft p?m.u$.
- 241 3. If $p \neq s \wedge (p \neq r \vee q \neq s)$, then $w.p \triangleright q!m.s \triangleleft r?m'.u \sim_1 w.s \triangleleft r?m'.p \triangleright q!m.u$.
- 242 4. If $|w \downarrow_{p \triangleright q!} _ | > |w \downarrow_{q \triangleleft p?} _ |$, then $w.p \triangleright q!m.q \triangleleft p?m'.u \sim_1 w.q \triangleleft p?m'.p \triangleright q!m.u$.

243 Let w, w' , and w'' be words s.t. $w \sim_1 w'$ and $w' \sim_i w''$ for some i . Then, $w \sim_{i+1} w''$. We
 244 define $w \sim u$ if $w \sim_n u$ for some n . It is straightforward that \sim is an equivalence relation.
 245 Define $u \preceq_\sim v$ if there is $w \in \Sigma^*$ such that $u.w \sim v$. Observe that $u \sim v$ iff $u \preceq_\sim v$ and
 246 $v \preceq_\sim u$. For infinite words $u, v \in \Sigma^\omega$, we define $u \preceq_\sim^\omega v$ if for each finite prefix u' of u , there
 247 is a finite prefix v' of v such that $u' \preceq_\sim v'$. Define $u \sim v$ iff $u \preceq_\sim^\omega v$ and $v \preceq_\sim^\omega u$.

248 We lift the equivalence relation \sim on words to languages:

249 For a language L , we define $\mathcal{C}^\sim(L) = \left\{ w' \mid \bigvee \begin{array}{l} w' \in \Sigma^* \wedge \exists w \in \Sigma^*. w \in L \text{ and } w' \sim w \\ w' \in \Sigma^\omega \wedge \exists w \in \Sigma^\omega. w \in L \text{ and } w' \preceq_\sim^\omega w \end{array} \right\}$.

This relation characterises what can be achieved in a distributed setting using CSMs.

► **Lemma 2.9** (L. 21 [64]). *Let $\{\{A_p\}_{p \in \mathcal{P}}\}$ be a CSM. Then, $\mathcal{L}(\{\{A_p\}_{p \in \mathcal{P}}\}) = \mathcal{C}^\sim(\mathcal{L}(\{\{A_p\}_{p \in \mathcal{P}}\}))$.*

Global and Local Types – Semantics

Hence, we define the semantics of global types using the indistinguishability relation \sim .

► **Definition 2.10** (Semantics of global types). *We construct a state machine $\mathbf{GAut}(\mathbf{G})$ to obtain the semantics of a global type \mathbf{G} . We index every syntactic subterm of \mathbf{G} with a unique index to distinguish common syntactic subterms, denoted with $[G, k]$ for syntactic subterm G and index k . Without loss of generality, the index for \mathbf{G} is 0: $[\mathbf{G}, 0]$. For clarity, we do not quantify indices. We define $\mathbf{GAut}(\mathbf{G}) = (Q_{\mathbf{GAut}(\mathbf{G})}, \Sigma_{\text{sync}}, \delta_{\mathbf{GAut}(\mathbf{G})}, q_{0, \mathbf{GAut}(\mathbf{G})}, F_{\mathbf{GAut}(\mathbf{G})})$ where*

- $Q_{\mathbf{GAut}(\mathbf{G})}$ is the set of all indexed syntactic subterms $[G, k]$ of \mathbf{G}
- $\delta_{\mathbf{GAut}(\mathbf{G})}$ is the smallest set containing $([\sum_{i \in I} p \rightarrow q_i : m_i. [G_i, k_i], k], p \rightarrow q_i : m_i, [G_i, k_i])$ for each $i \in I$, and $([\mu t. [G', k^2], k^1], \varepsilon, [G', k^2])$ and $([t, k^3], \varepsilon, [\mu t. [G', k^2], k^1])$,
- $q_{0, \mathbf{GAut}(\mathbf{G})} = [\mathbf{G}, 0]$, and $F_{\mathbf{GAut}(\mathbf{G})} = \{[0, k] \mid k \text{ is an index for subterm } 0\}$

We consider asynchronous communication so each interaction is split into its send and receive event. In addition, we consider CSMs as implementation model for global types and, from Lemma 2.9, we know that CSM languages are always closed under the indistinguishability relation \sim . Thus, we also apply its closure to obtain the semantics of \mathbf{G} : $\mathcal{L}(\mathbf{G}) := \mathcal{C}^\sim(\text{split}(\mathcal{L}(\mathbf{GAut}(\mathbf{G})))$.

The closure $\mathcal{C}^\sim(-)$ corresponds to similar reordering rules in standard MST developments, e.g., [49, Def. 3.2 and 5.3].

► **Example 2.11.** In Figure 1a (p.2), we presented the FSM for the semantics of $\mathbf{GAut}(\mathbf{G}_{2\text{BP}})$. We give the semantics of a simple global type where p communicates a list of book titles to q : $\mu t. (p \rightarrow q : \text{title}. t + p \rightarrow q : \text{done}. 0)$. Its semantics is the union of two cases: if the list of book titles is finite, i.e., $\mathcal{C}^\sim((p \triangleright q ! \text{title}. q \triangleleft p ? \text{title})^* . p \triangleright q ! \text{done}. q \triangleleft p ? \text{done})$; and the one if the list is infinite, i.e., $\mathcal{C}^\sim((p \triangleright q ! \text{title}. q \triangleleft p ? \text{title})^\omega)$. Here, there are only two roles so $\mathcal{C}^\sim(-)$ can solely delay receive events (Rule 4 of \sim).

We distinguish states depending on which subterm they correspond to: *binder states* with their dashed line correspond to a recursion variable binder, while *recursion states* with their dash-dotted lines indicate the use of a recursion variable. We omit ε for transitions from recursion to binder states.

► **Definition 2.12** (Semantics for local types). *Given a local type L for role p , we index syntactic subterms as for the semantics of global types. We construct a state machine $\mathbf{LAut}(L) = (Q, \Sigma_p, \delta, q_0, F)$ where*

- Q is the set of all indexed syntactic subterms in L ,
- δ is the smallest set containing $([\oplus_{i \in I} q_i ! m_i. [L_i, k_i], k], p \triangleright q_i ! m_i, [L_i, k_i])$ and $([\&_{i \in I} q_i ? m_i. [L_i, k_i], k], p \triangleleft q_i ? m_i, [L_i, k_i])$ for each $i \in I$, as well as $([\mu t. [L', k^2], k^1], \varepsilon, [L', k^2])$ and $([t, k^3], \varepsilon, [\mu t. [L', k^2], k^1])$,
- $q_0 = [L, 0]$ and $F = \{[0, k] \mid k \text{ is an index for subterm } 0\}$

We define the semantics of L as language of this automaton: $\mathcal{L}(L) = \mathcal{L}(\mathbf{LAut}(L))$.

Compared to global types, we distinguish two more kinds of states for local types: a *send state* (internal choice) has a diamond shape while a *receive state* (external choice) has a rectangular shape. For states with ε as next action, we keep the circular shape and call them *neutral states*. Figure 1b (p.2) does not represent the state machine for any local type but illustrates the use of different styles for different kinds of states.

294 **The Implementability Problem for Global Types**

295 The implementability problem for global types asks whether a global type can be implemented
 296 in a distributed setting. The projection operator takes the intermediate representation of
 297 local types as local specifications for roles. We define implementability directly on the
 298 implementation model of CSMs. Intuitively, every set of local types constitutes a CSM
 299 through their semantics.

300 ► **Definition 2.13** (Implementability [64]). *A global type \mathbf{G} is said to be implementable if
 301 there exists a CSM $\{\{A_p\}_{p \in \mathcal{P}}\}$ such that*

- 302 ■ (deadlock freedom) $\{\{A_p\}_{p \in \mathcal{P}}\}$ is deadlock free, and
- 303 ■ (protocol fidelity) their languages are the same: $\mathcal{L}(\mathbf{G}) = \mathcal{L}(\{\{A_p\}_{p \in \mathcal{P}}\})$.

304 We say that $\{\{A_p\}_{p \in \mathcal{P}}\}$ implements \mathbf{G} .

305 ► **Remark 2.14** (Progress). Deadlock freedom is sometimes also studied as *progress* — in the
 306 sense that a system should never get stuck. However, for infinite executions, a role could
 307 starve in a non-final state by waiting for a message that is never sent [19, Sec. 3.2]. Castagna
 308 et al. [19] consider a stronger notion of progress (Def. 3.3: live session) which requires that
 309 every role could eventually reach a final state. Our results apply to this stronger notion of
 310 progress as we discuss in Section 4.2.

311 **3 Projection – From Global to Local Types**

312 In this section, we define and visually explain a typical approach to the implementability
 313 problem: the classical projection operator. It tries to translate global types to local types
 314 and, while doing so, checks if this is safe. Behind the scenes, these checks are conducted
 315 by a partial merge operator. We consider different variants of the merge operator from the
 316 literature and exemplify the features they support. We provide visual explanations of the
 317 classical projection operator with these merge operators on the state machines of global
 318 types by example. In Appendix B, we give general descriptions but they are not essential to
 319 explain our observations. Lastly, we summarise the shortcomings of the full merge operator
 320 and exemplify them with variants of the two buyer protocol from the introduction.

321 **Classical Projection Operator with Parametric Merge**

322 ► **Definition 3.1** (Projection operator). *For a merge operator \sqcap , the projection of a global
 323 type \mathbf{G} on to a role $r \in \mathcal{P}$ is a local type that is defined as follows:² $0|_r := 0$ $t|_r := t$*

$$324 \left(\sum_{i \in I} p \rightarrow q : m_i . G_i \right) |_r := \begin{cases} \oplus_{i \in I} q ! m_i . (G_i |_r) & \text{if } r = p \\ \&_{i \in I} p ? m_i . (G_i |_r) & \text{if } r = q \\ \sqcap_{i \in I} G_i |_r & \text{otherwise} \end{cases} \quad (\mu t . G) |_r := \begin{cases} \mu t . (G |_r) & \text{if } G |_r \neq t \\ 0 & \text{otherwise} \end{cases}$$

325 Intuitively, a projection operator takes the state machine $\mathbf{GAut}(\mathbf{G})$ for a global type \mathbf{G} and
 326 projects each transition label to the respective alphabet of the role, e.g., $p \rightarrow q : m$ becomes
 327 $q \triangleleft p ? m$ for role q . This can introduce non-determinism that ought to be resolved by a partial
 328 merge operator. Several merge operators have been proposed in the literature.

329 ► **Definition 3.2** (Merge Operators). *Let L_1 and L_2 be local types for a role r , and \sqcap be a
 330 merge operator. We define different cases for the result of $L_1 \sqcap L_2$:*

² The case split for the recursion binder changes slightly across different definitions. We chose a simple but also the least restrictive condition. We simply check whether the recursion is vacuous (as $\mu t . t$) and omit it in this case. We require to omit μt if t is never used in the result.

- 331 (1) L_1 if $L_1 = L_2$
 332 (2) $\left(\begin{array}{cc} \&_{i \in I \setminus J} q?m_i.L'_{1,i} & \& \\ \&_{i \in I \cap J} q?m_i.(L'_{1,i} \sqcap L'_{2,i}) & \& \\ \&_{i \in J \setminus I} q?m_i.L'_{2,i} & \end{array} \right)$ if $\begin{cases} L_1 = \&_{i \in I} q?m_i.L'_{1,i}, \\ L_2 = \&_{i \in J} q?m_i.L'_{2,i} \end{cases}$
 333 (3) $\mu t_1.(L'_1 \sqcap L'_2[t_2/t_1])$ if $L_1 = \mu t_1.L'_1$ and $L_2 = \mu t_2.L'_2$

334 Each merge operator is defined by a collection of cases it can apply. If none of the respective
 335 cases applies, the result of the merge is undefined. The plain merge \sqcap [28] can only apply
 336 Case (1). The semi-full merge \sqcap [85] can apply Cases (1) and (2). The full merge \sqcap [75] can
 337 apply all Cases (1), (2), and (3).

338 We will also consider the availability merge operator \sqcap by Majumdar et al. [64] which
 339 builds on the full merge operator but generalises Case (2) to allow generalised choice. We
 340 will explain the main differences in Remark 3.12.

341 ► **Remark 3.3 (Correctness of projection).** This would be the correctness criterion for projection:
 342 Let \mathbf{G} be some global type and let plain merge \sqcap , semi full merge \sqcap , full merge \sqcap , or
 343 availability merge \sqcap be the merge operator \sqcap . If $\mathbf{G} \upharpoonright_p$ is defined for each role p , then the
 344 CSM $\{\{\text{LAut}(\mathbf{G} \upharpoonright_p)\}\}_{p \in \mathcal{P}}$ implements \mathbf{G} .

345 We do not actually prove this so we do not state it as lemma. *But why does this hold?*
 346 The implementability condition is the combination of deadlock freedom and protocol fidelity.
 347 Coppo et al. [28] show that *subject reduction* entails protocol fidelity and progress while
 348 progress, in turn, entails deadlock freedom. Subject reduction has been proven for the plain
 349 merge operator [28, Thm. 1] and the semi-full operator [85, Thm. 1]. Scalas and Yoshida
 350 pointed out that several versions of classical projection with the full merge are flawed [75,
 351 Sec. 8.1]. Hence, we have chosen a full merge operator whose correctness follows from the
 352 correctness of the more general availability merge operator. For the latter, the correctness
 353 follows from work by Majumdar et al. [64, Thm. 16].

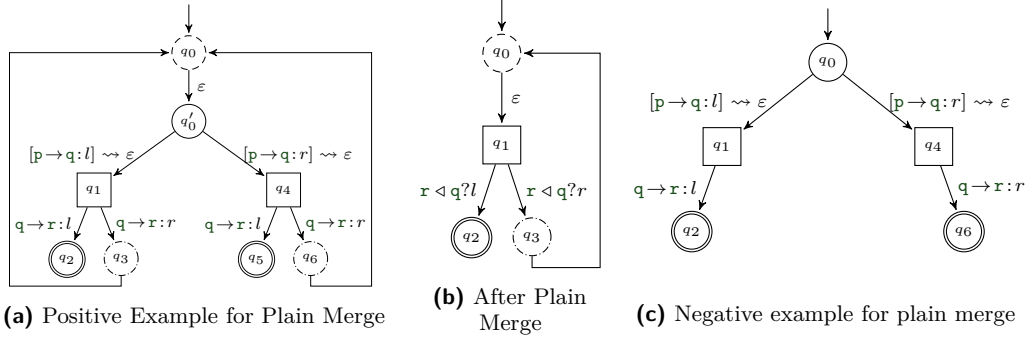
354 ► **Example 3.4 (Projection without merge / Collapsing erasure).** In the introduction, we
 355 considered $\mathbf{G}_{2\text{BP}}$ and the FSM for its semantics in Figure 1a. We projected (without merge)
 356 on to Seller s to obtain the FSM in Figure 1b. In general, we also collapse neutral states
 357 with a single ε -transition and their only successor. We call this *collapsing erasure*. We only
 358 need to actually collapse states for the protocol in Figure 4a. In all other illustrations, we
 359 indicate the interactions the role is not involved with the following notation: $[p \rightarrow q:l] \rightsquigarrow \varepsilon$.

360 On the Structure of $\text{GAut}(\mathbf{G})$

361 We now show that the state machine for every local and global type has a certain shape. This
 362 simplifies the visual explanations of the different merge operators. Intuitively, every such
 363 state machine has a tree-like structure where backward transitions only happen at leaves of
 364 the tree, are always labelled with ε , and only lead to ancestors. The FSM in Figure 1a (p.2)
 365 illustrates this shape where the root of the tree is at the top.

366 ► **Definition 3.5 (Ancestor-recursive, non-merging, intermediate recursion, etc.).** Let $A =$
 367 $(Q, \Delta, \delta, q_0, F)$ be a finite state machine. We say that A is ancestor-recursive if there is a
 368 function $\text{lvl}: Q \rightarrow \mathbb{N}$ such that, for every transition $q \xrightarrow{x} q' \in \delta$, one of the two holds:

- 369 (1) $\text{lvl}(q) > \text{lvl}(q')$, or
 370 (2) $x = \varepsilon$ and there is a run from the initial state q_0 (without going through q) to q' which
 371 can be completed to reach q : $q_0 \xrightarrow{\bar{x}} \dots \xrightarrow{\bar{x}} q_n$ is a run with $q_n = q'$ and $q \neq q_i$ for every
 372 $0 \leq i \leq n$, and the run can be extended to $q_0 \xrightarrow{\bar{x}} \dots \xrightarrow{\bar{x}} q_n \xrightarrow{\bar{x}} \dots \xrightarrow{\bar{x}} q_{n+m}$ with $q_{n+m} = q$.
 373 Then, the state q' is called ancestor of q .



■ **Figure 2** The FSM on the left represents an implementable global type that is accepted by plain merge. It implicitly shows the FSM after collapsing erasure: every interaction r is not involved in is given as $[p \rightarrow q: l] \rightsquigarrow \varepsilon$. The FSM in the middle is the result of the plain merge. The FSM on the right represents an implementable global type that is rejected by plain merge. It is obtained from the left one by removing one choice option in each branch of the initial choice.

374 We call the first (1) kind of transition forward transition while the second (2) kind is a
 375 backward transition. The state machine A is said to be free from intermediate recursion if
 376 every state q with more than one outgoing transition, i.e., $|\{q' \mid q \xrightarrow{\cdot} q' \in \delta\}| > 1$, has only
 377 forward transitions. We say that A is non-merging if every state only has one incoming edge
 378 with greater level, i.e., for every state q' , $\{q \mid q \xrightarrow{\cdot} q' \in \delta \wedge \text{lvl}(q) > \text{lvl}(q')\} \leq 1$. The state
 379 machine A is dense if, for every $q \xrightarrow{x} q' \in \delta$, the transition label x is ε implies that q has
 380 only one outgoing transition. Last, the cone of q are all states q' which are reachable from q
 381 and have a smaller level than q , i.e., $\text{lvl}(q) > \text{lvl}(q')$.

382 ► **Proposition 3.6** (Shape of $\text{GAut}(\mathbf{G})$ and $\text{LAut}(L)$). Let \mathbf{G} be some global type and L be some
 383 local type. Then, both $\text{GAut}(\mathbf{G})$ and $\text{LAut}(L)$ are ancestor-recursive, free from intermediate
 384 recursion, non-merging, and dense.

385 For both, the only forward ε -transitions occur precisely from binder states while backward
 386 transitions happen from variable states to binder states. The illustrations for our examples
 387 always have the initial state, which is the state with the greatest level, at the top. This is
 388 why we use greater and higher as well as smaller and lower interchangeably for levels.

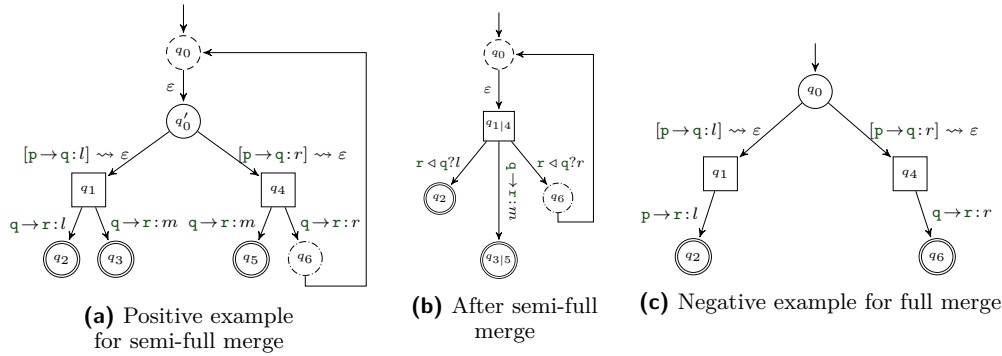
389 Features of Different Merge Operators by Example

390 In this section, we exemplify which features each of the merge operators does support. We
 391 present a sequence of implementable global types. Despite, some cannot be handled by
 392 some (or all) merge operators. If a global type is not projectable using some merge operator,
 393 we say it is *rejected* and a *negative* example for this merge operator. We focus on role r
 394 when projecting. Thus, rejected mostly means that there is (at least) no projection on to r .
 395 If a global type is projectable by some merge operator, we call it a *positive* example. All
 396 examples strive for minimality and follow the idea that roles decide whether to take a left (l)
 397 or right (r) branch of a choice.

398 ► **Example 3.7** (Positive example for plain merge). The following global type is implementable:

$$399 \quad \mu t. + \left\{ \begin{array}{l} p \rightarrow q: l. (q \rightarrow r: l. 0 + q \rightarrow r: r. t) \\ p \rightarrow q: r. (q \rightarrow r: l. 0 + q \rightarrow r: r. t) \end{array} \right.$$

400 The state machine for its semantics is given in Figure 2a. After collapsing erasure, there is
 401 a non-deterministic choice from q'_0 leading to q_1 and q_4 since r is not involved in the initial
 402 choice. The plain merge operator can resolve this non-determinism since both cones of q_1



■ **Figure 3** The FSM on the left represents an implementable global type (and implicitly the collapsing erasure on to \mathbf{r}) that is accepted by semi-merge. The FSM in the middle is the result of the semi-full merge. The FSM on the right is a negative example for the full merge operator.

403 and q_4 represent the same subterm. Technically, there is an isomorphism between the states
 404 in both cones which preserves the kind of states as well as the transition labels and the
 405 backward transitions from isomorphic recursion states lead to the same binder state. The
 406 result is illustrated in Figure 2b. It is also the FSM of a local type for \mathbf{r} which is the result
 407 of the (syntactic) plain merge: $\mu t. (\mathbf{r} \triangleleft q?l. 0 \ \& \ \mathbf{r} \triangleleft q?r. t)$.

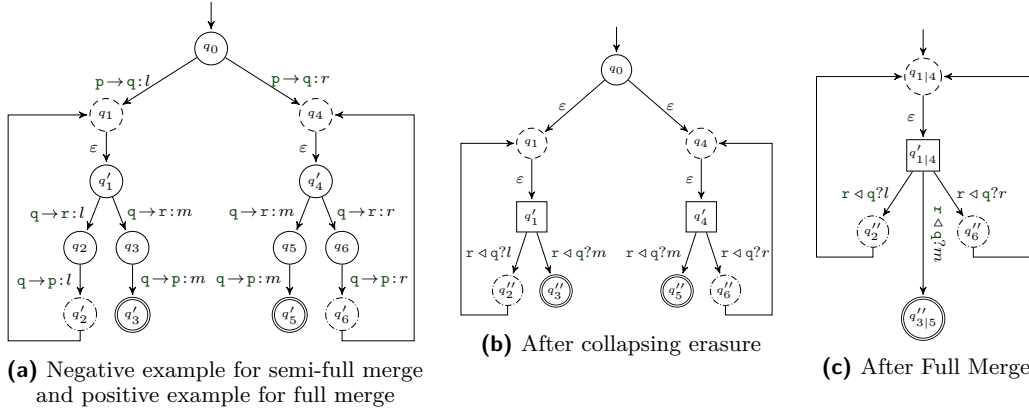
408 Our explanation on FSMs allows to check congruence of cones to merge while the definition
 409 requires syntactic equality. If we swap the order of branches $q \rightarrow \mathbf{r}: l$ and $q \rightarrow \mathbf{r}: r$ in Figure 2a
 410 on the right, the syntactic merge rejects. Still, because both are semantically the same
 411 protocol specification, we expect tools to check for such easy fixes.

412 ► **Example 3.8** (Negative example for plain merge). We consider the following simple imple-
 413 mentable global type where the choice by p is propagated to \mathbf{r} : $\left\{ \begin{array}{l} p \rightarrow q: l. q \rightarrow \mathbf{r}: l. 0 \\ p \rightarrow q: r. q \rightarrow \mathbf{r}: r. 0 \end{array} \right.$.
 414 The corresponding state machine is illustrated in Figure 2c. Here, q_0 exhibits non-determinism
 415 but the plain merge fails because q_1 and q_4 have different outgoing transition labels.

416 Intuitively, the plain merge operator forbids that any, but the two roles involved in a
 417 choice, can have different behaviour after the choice. It basically forbids propagating a choice.
 418 The semi-full merge overcomes this shortcoming and can handle the previous example. We
 419 present a slightly more complex one to showcase the features it supports.

420 ► **Example 3.9** (Positive example for semi-full merge). Let us consider the following imple-
 421 mentable global type: $\mu t. \left\{ \begin{array}{l} p \rightarrow q: l. (q \rightarrow \mathbf{r}: l. 0 + q \rightarrow \mathbf{r}: m. 0) \\ p \rightarrow q: r. (q \rightarrow \mathbf{r}: m. 0 + q \rightarrow \mathbf{r}: r. t) \end{array} \right.$. The state machine for its semantics
 422 is illustrated in Figure 3a. After applying collapsing erasure, there is a non-deterministic
 423 choice from q_0 leading to q_1 and q_4 since \mathbf{r} is not involved in the initial choice, We apply
 424 the semi-full merge for both states. Both are receive states so Case (2) applies. First, we
 425 observe that $\mathbf{r} \triangleleft q?l$ and $\mathbf{r} \triangleleft q?r$ are unique to one of the two states so both transitions,
 426 with the cones of the states they lead to, can be kept. Second, there is $\mathbf{r} \triangleleft q?m$ which
 427 is possible in both states. We recursively apply the semi-full merge and, with Case (1),
 428 observe that the result $q_{3,5}$ is simply a final state. Overall, we obtain the state machine in
 429 Figure 3b which is equivalent to the result of the syntactic projection with semi-full merge:
 430 $\mu t. (\mathbf{r} \triangleleft q?l. 0 + \mathbf{r} \triangleleft q?m. 0 + \mathbf{r} \triangleleft q?r. t)$.

431 ► **Example 3.10** (Negative example for semi-full merge and positive example for full merge).
 432 The semi-full merge operator rejects the following implementable global type:



■ **Figure 4** The FSM on the left represents an implementable global type that is rejected by the semi-full merge. It is accepted by the full merge: collapsing erasure yields the FSM in the middle and applying the full merge the FSM on the right.

$$+ \left\{ \begin{array}{l} p \rightarrow q:l. \mu t_1. (q \rightarrow r:l. q \rightarrow p:l. t_1 + q \rightarrow r:m. q \rightarrow p:m. 0) \\ p \rightarrow q:r. \mu t_2. (q \rightarrow r:m. q \rightarrow p:m. 0 + q \rightarrow r:r. q \rightarrow p:r. t_2) \end{array} \right.$$

Its FSM and the FSM after collapsing erasure is given in Figures 4a and 4b. Intuitively, it would need to recursively merge the parts after both recursion binders in order to merge the branches with receive event $r \triangleleft q?m$ but it cannot do so. The full merge can handle this global type. It can descend beyond q_1 and q_4 and is able to merge q'_1 and q'_4 . To obtain $q''_{3|5}$, it applies Case (1) while $q'_{1|4}$ is only feasible with Case (2). The result is embedded into the recursive structure to obtain the FSM in Figure 4c. It is equivalent to the (syntactic) result which renames the recursion variable for one branch: $\mu t_1. (r \triangleleft q?l. t_1 \ \& \ r \triangleleft q?m. 0 \ \& \ r \triangleleft q?r. t_1)$.

► **Example 3.11** (Negative example for full merge). We consider a simple implementable global type where p propagates its decision to r in the top branch while q propagates it in the bottom branch: $+ \left\{ \begin{array}{l} p \rightarrow q:l. p \rightarrow r:l. 0 \\ p \rightarrow q:r. q \rightarrow r:r. 0 \end{array} \right.$. It is illustrated in Figure 3c. This cannot be projected on to r by the full merge operator for which all receive events need to have the same sender.

► **Remark 3.12** (On generalised choice). Majumdar et al. [64] proposed a classical projection operator that allows to overcome this shortcoming. It can project the previous example. In general, allowing to receive from different senders has subtle consequences. Intuitively, messages from different senders could overtake each other in a distributed setting and one cannot rely on the FIFO order provided by the channel of a single sender. Thus, they employ a message availability analysis to ensure that there cannot be any confusion about which branch shall be taken. Except for the possibility to merge cases where a receiver receives from multiple senders, their merge operator suffers from the same shortcomings as any classical projection operator. We refrain from presenting their merge operator here but refer to their work for details on the availability merge operator $\overline{\mu}$.

Case (2) allows to descend for common receive events. One could also add a similar case for send events where one recursively applies the merge operator (but, in most cases, the set of send events ought to be the same). Such a case might render some global types projectable. However, it does not give any additional insights into the concept of the classical projection operator and its potential merge operators. Of course, one could consider the different cases in all combinations. Again, this does not really give insights which is why we deliberately chose this incremental style that concisely shows which cases support which features.

462 Shortcomings of Classical Projection/Merge Operators

463 In this section, we present slight variations of the two buyer protocol that are implementable
464 but rejected by all of the presented projection/merge operators.

465 ► **Example 3.13.** We obtain an implementable variant by omitting both message interactions
466 $a \rightarrow s : no$ with which Buyer a notifies Seller s that they will not buy the item:

$$467 \quad \mu t. + \begin{cases} a \rightarrow s : query. s \rightarrow a : price. (a \rightarrow b : split. (b \rightarrow a : yes. a \rightarrow s : buy. t + b \rightarrow a : no. t) + a \rightarrow b : cancel. t) \\ a \rightarrow s : done. a \rightarrow b : done. 0 \end{cases}$$

468 This global type cannot be projected on to Seller s . The merge operator would need to
469 merge a recursion variable with an external choice. Visually, the merge operator does not
470 allow to unfold the variable t and try to merge again. However, there is a local type:

$$471 \quad \mu t_1. \& \begin{cases} s \triangleleft a?query. \mu t_2. s \triangleright a!price. (s \triangleleft a?buy. t_1 \& s \triangleleft a?query. t_2 \& s \triangleleft a?done. 0) \\ s \triangleleft a?done. 0 \end{cases}$$

472 The local type has two recursion variable binders while the global type only has one. Our
473 explanations showed that classical projection operators can never yield such a structural
474 change: the merge operator can only merge states but not introduce new ones or introduce
475 new backward transitions.

476 ► **Example 3.14 (Two Buyer Protocol with Subscription).** In this variant, Buyer a first decides
477 whether to subscribe to a yearly discount offer or not — before purchasing the sequence of
478 items — and notifies Buyer b if it does so: $\mathbf{G}_{2BPWS} := + \begin{cases} a \rightarrow s : login. \mathbf{G}_{2BP} \\ a \rightarrow s : subscribe. a \rightarrow b : subscribed. \mathbf{G}_{2BP} \end{cases}$

479 The merge operator needs to merge a recursion variable binder μt with an external choice
480 $b \triangleleft a?subscribed$. Because Buyer a only sends *subscribed* at the beginning of the protocol,
481 it is safe to introduce one recursion variable earlier to obtain the following local type for
482 Buyer b . (In fact, we could also remove μt_2 and substitute t_2 by t_1 for the same reason.)

$$483 \quad \mu t_1. \& \begin{cases} b \triangleleft a?split. (b \triangleright a!yes. t_1 \oplus b \triangleright a!no. t_1) \\ b \triangleleft a?cancel. t_1 \\ b \triangleleft a?done. 0 \\ b \triangleleft a?subscribed. \mu t_2. (b \triangleleft a?split. (b \triangleright a!yes. t_2 \oplus b \triangleright a!no. t_2) \& b \triangleleft a?cancel. t_2 \& b \triangleleft a?done. 0) \end{cases}$$

484 Similarly, the classical projection operator cannot yield any local type which needs to
485 distinguish semantic properties to disambiguate a choice, e.g., counting modulo a constant.
486 Scalas and Yoshida [75] identified another shortcoming: most classical projection operators
487 require all branches of a loop to contain the same set of active roles. Thus, they cannot
488 project the following global type. It is implementable and if it was projectable, the result
489 would be equivalent to the local types given in their example [75, Fig. 4 (2)].

490 ► **Example 3.15 (Two Buyer Protocol with Inner Recursion).** This variant allows to recursively
491 negotiate how to split the price (and omits the outer recursion):

$$492 \quad \mathbf{G}_{2BPIR} := a \rightarrow s : query. s \rightarrow a : price. \mu t. + \begin{cases} a \rightarrow b : split. (b \rightarrow a : yes. a \rightarrow s : buy. 0 + b \rightarrow a : no. t) \\ a \rightarrow b : cancel. a \rightarrow s : no. 0 \end{cases}$$

493 These shortcomings have been addressed by some non-classical approaches. For example,
494 Scalas and Yoshida [75] employ model checking while Dagnino et al. [31] characterise
495 implementable global types with an undecidable well-formedness condition and give a sound
496 algorithmically checkable approximation. It is not known whether the implementability
497 problem for global types, neither with directed or generalised choice, is decidable. We answer
498 this question positively for the more general case of generalised choice.

499 **4** Implementability for Global Types from MSTs is Decidable

500 In this section, we show that the implementability problem for global types with generalised
501 choice is decidable. For this, we use results from the domain of message sequence charts.


502 We first introduce high-level message sequence charts (HMSCs) and recall an encoding of
 503 global types to HMSCs. In general, implementability of HMSCs is undecidable but we
 504 prove that global types belong, when encoded as HMSCs, to a class of HMSCs for which
 505 implementability is decidable.

506 4.1 High-level Message Sequence Charts

507 Our definitions of (high-level) message sequence charts follow work by Genest et al. [38] and
 508 Stutz and Zufferey [77]. If reasonable, we adapt terminology to the MST setting.

509 ► **Definition 4.1** (Message Sequence Charts). *A message sequence chart (MSC) is a 5-tuple*
 510 $M = (N, p, f, l, (\leq_p)_{p \in \mathcal{P}})$ where

- N is a set of send (S) and receive (R) event nodes such that $N = S \uplus R$ (where \uplus denotes disjoint union),
- $p: N \rightarrow \mathcal{P}$ maps each event node to the role acting on it,
- $f: S \rightarrow R$ is an injective function linking corresponding send and receive event nodes,
- $l: N \rightarrow \Sigma$ labels every event node with an event, and
- $(\leq_p)_{p \in \mathcal{P}}$ is a family of total orders for the event nodes of each role: $\leq_p \subseteq p^{-1}(p) \times p^{-1}(p)$.



511 ■ **Figure 5** Highlighting the elements of a MSC: $(N, p, f, l, (\leq_p)_{p \in \mathcal{P}})$

512 An MSC M induces a partial order \leq_M on N that is defined co-inductively:

$$\frac{e \leq_p e'}{e \leq_M e'} \text{ PROC} \quad \frac{s \in S}{s \leq_M f(s)} \text{ SND-RCV} \quad \frac{}{e \leq_M e} \text{ REFL} \quad \frac{e \leq_M e' \quad e' \leq_M e''}{e \leq_M e''} \text{ TRANS}$$

513 The labelling function l respects the function f : for every send event node e , we have that
 514 $l(e) = p(e) \triangleright p(f(e)) ! m$ and $l(f(e)) = p(f(e)) \triangleleft p(e) ? m$ for some $m \in \mathcal{V}$.

515 All MSCs in our work respect FIFO, i.e., there are no p and q such that there are
 516 $e_1, e_2 \in p^{-1}(p)$ with $e_1 \neq e_2$, $l(e_1) = l(e_2)$, $e_1 \leq_p e_2$ and $f(e_2) \leq_q f(e_1)$ (also called
 517 degenerate) and for every pair of roles p, q , and for every two event nodes $e_1 \leq_M e_2$ with
 518 $l(e_i) = p \triangleright q ! _$ for $i \in \{1, 2\}$, it holds that $\mathcal{V}(w_p) = \mathcal{V}(f(w_p))$ where w_p is the (unique)
 519 linearisation of $p^{-1}(p)$. A *basic MSC* (BMSC) has a finite number of nodes N and \mathcal{M} denotes
 520 the set of all BMSCs. When unambiguous, we omit the index M for \leq_M and write \leq . We
 521 define \leq as expected. The language $\mathcal{L}(M)$ of an MSC M collects all words $l(w)$ for which w
 522 is a linearisation of N that is compliant with \leq_M .

523 If one thinks of a BMSC as straight-line code, a high-level message sequence chart adds
 524 control flow. It embeds BMSCs into a graph structure which allows for choice and recursion.

525 ► **Definition 4.2** (High-level Message Sequence Charts [77]). *A high-level message sequence*
 526 *chart (HMSC) is a 5-tuple (V, E, v^I, V^T, μ) where V is a finite set of vertices, $E \subseteq V \times V$*
 527 *is a set of directed edges, $v^I \in V$ is an initial vertex, $V^T \subseteq V$ is a set of terminal vertices,*
 528 *and $\mu: V \rightarrow \mathcal{M}$ is a function mapping every vertex to a BMSC. A path in an HMSC is a*
 529 *sequence of vertices v_1, \dots from V that is connected by edges, i.e., $(v_i, v_{i+1}) \in E$ for every i .*
 530 *A path is maximal if it is infinite or ends in a vertex from V^T .*

531 Intuitively, the language of an HMSC is the union of all languages of the finite and infinite
 532 MSCs generated from maximal paths in the HMSC and is formally defined in Appendix C.1.

533 Like global types, an HMSC specifies a protocol. The implementability question was also
 534 posed for HMSCs and studied as *safe realisability*. If the CSM is not required to be deadlock
 535 free, it is called weak realisability.

536 ► **Definition 4.3** (Safe realisability of HMSCs [4]). *An HMSC H is said to be safely realisable*
 537 *if there exists a deadlock free CSM $\{\{A_p\}_{p \in \mathcal{P}}\}$ such that $\mathcal{L}(H) = \mathcal{L}(\{\{A_p\}_{p \in \mathcal{P}}\})$.*

538 Encoding Global Types from MSTs as HMSCs

539 Stutz and Zufferey [77] provide a formal connection from global types to HMSCs. We recall
540 their encoding and main correctness result.

541 ► **Definition 4.4** (Encoding global types as HMSCs [77]). *In the translation, the following
542 notation is used: M_\emptyset is the empty BMSC ($N = \emptyset$) and $M(\mathbf{p} \rightarrow \mathbf{q} : m)$ is the BMSC with two
543 event nodes: e_1, e_2 such that $f(e_1) = e_2$, $l(e_1) = \mathbf{p} \triangleright \mathbf{q} ! m$, and $l(e_2) = \mathbf{q} \triangleleft \mathbf{p} ? m$.*

544 *Let \mathbf{G} be a global type, we construct an HMSC $H(\mathbf{G}) = (V, E, v^I, V^T, \mu)$ with*

$$\begin{aligned}
 V &= \{G' \mid G' \text{ is a subterm of } \mathbf{G}\} \cup \\
 &\quad \{(\sum_{i \in I} \mathbf{p} \rightarrow \mathbf{q}_i : m_i.G_i, j) \mid \sum_{i \in I} \mathbf{p} \rightarrow \mathbf{q}_i : m_i.G_i \text{ occurs in } \mathbf{G} \wedge j \in I\} \\
 E &= \{(\mu t.G', G') \mid \mu t.G' \text{ occurs in } \mathbf{G}\} \cup \{(t, \mu t.G') \mid t, \mu t.G' \text{ occurs in } \mathbf{G}\} \\
 &\quad \cup \{(\sum_{i \in I} \mathbf{p} \rightarrow \mathbf{q}_i : m_i.G_i, (\sum_{i \in I} \mathbf{p} \rightarrow \mathbf{q}_i : m_i.G_i, j)) \mid (\sum_{i \in I} \mathbf{p} \rightarrow \mathbf{q}_i : m_i.G_i, j) \in V\} \\
 &\quad \cup \{((\sum_{i \in I} \mathbf{p} \rightarrow \mathbf{q}_i : m_i.G_i, j), G_j) \mid (\sum_{i \in I} \mathbf{p} \rightarrow \mathbf{q}_i : m_i.G_i, j) \in V\} \\
 v^I &= \mathbf{G} \quad V^T = \{0\} \quad \mu(v) = \begin{cases} M(\mathbf{p} \rightarrow \mathbf{q}_i : m_j) & \text{if } v = (\sum_{i \in I} \mathbf{p} \rightarrow \mathbf{q}_i : m_i.G_i, j) \\ M_\emptyset & \text{otherwise} \end{cases}
 \end{aligned}$$

546 We adapt the correctness result to our definitions. In particular, our semantics of \mathbf{G} use
547 the closure operator $\mathcal{C}^\sim(-)$ while they explicitly distinguish between a type and execution
548 language. We also omit the closure operator on the right-hand side because HMSCs are
549 closed with regard to this operator [77, Lm. 5].

550 ► **Theorem 4.5.** *Let \mathbf{G} be a global type. Then, the following holds: $\mathcal{L}(\mathbf{G}) = \mathcal{L}(H(\mathbf{G}))$.*

551 4.2 Implementability is Decidable

552 ► **Assumption (0-Reachable).** *We introduce a mild assumption for global types. We say a
553 global type \mathbf{G} is 0-reachable if every prefix of a word in its language can be completed to a
554 finite word. Equivalently, we require that the vertex 0 is reachable from any vertex in $H(\mathbf{G})$.³
555 Intuitively, this solely rules out global types that have loops without exit (cf. Example 4.19).*

556 The MSC approach to safe realisability for HMSCs is different from the classical projection
557 approach to implementability. Given an HMSC, there is a canonical candidate implementation
558 which always implements the HMSC if an implementation exists [3, Thm. 13]. Therefore,
559 approaches center around checking safe realisability of HMSC languages and establishing
560 conditions on HMSCs that entail safe realisability.

561 ► **Definition 4.6** (Canonical candidate implementation [3]). *Given an HMSC H and a role \mathbf{p} ,
562 let $A'_\mathbf{p} = (Q', \Sigma_\mathbf{p}, \delta', q'_0, F')$ be a state machine with $Q' = \{q_w \mid w \in \text{pref}(\mathcal{L}(H) \downarrow_{\Sigma_\mathbf{p}})\}$,
563 $F' = \{q_w \mid w \in \mathcal{L}_{\text{fin}}(H) \downarrow_{\Sigma_\mathbf{p}}\}$, and $\delta'(q_w, x, q_{wx})$ for $x \in \Sigma_{\text{async}}$. The resulting state
564 machine $A'_\mathbf{p}$ is not necessarily finite so $A'_\mathbf{p}$ is determinised and minimised which yields the
565 FSM $A_\mathbf{p}$. We call $\{\{A_\mathbf{p}\}_{\mathbf{p} \in \mathcal{P}}$ the canonical candidate implementation of H .*

566 Intuitively, the intermediate state machine $A'_\mathbf{p}$ constitutes a tree whose maximal finite
567 paths give $\mathcal{L}(H) \downarrow_{\Sigma_\mathbf{p}} \cap \Sigma_\mathbf{p}^*$. This set can be infinite and, thus, the construction might not be
568 effective. We give an effective construction of a deterministic FSM for the same language
569 which was very briefly hinted at by Alur et al. [4, Proof of Thm. 3].

³ An equivalent conditions is common in the HMSC domain [39, Sec. 2] [77, Sec. 4].

570 ► **Definition 4.7** (Projection by Erasure). Let p be a role and $M = (N, p, f, l, (\leq_p)_{p \in \mathcal{P}})$ be
 571 an MSC. We denote the set of nodes of p with $N_p := \{n \mid p(n) = p\}$ and define a two-ary
 572 next-relation on N_p : $\text{next}(n_1, n_2)$ iff $n_1 \leq n_2$ and there is no n' with $n_1 \leq n' \leq n_2$. We
 573 define the projection by erasure of M on to p : $M \Downarrow_p = (Q_M, \Sigma_p, \delta_M, q_{M,0}, \{q_{M,f}\})$ with

$$574 \quad Q_M = \{q_n \mid n \in N_p\} \uplus \{q_{M,0}\} \uplus \{q_{M,f}\} \text{ and}$$

$$575 \quad \delta_M = \{q_{n_1} \xrightarrow{l(n_1)} q_{n_2} \mid \text{next}(n_1, n_2)\} \uplus \{q_{n_2} \xrightarrow{l(n_2)} q_{M,f} \mid \forall n_1. n_1 \leq n_2\} \uplus \{q_{M,0} \xrightarrow{\varepsilon} q_{n_1} \mid \forall n_2. n_1 \leq n_2\}$$

576 where \uplus denotes disjoint union. Let $H = (V, E, v^I, V^T, \mu)$ be an HMSC. We construct the
 577 projection by erasure for every vertex and identify them with the vertex, e.g., Q_v instead
 578 of $Q_{\mu(v)}$. We construct an auxiliary FSM $(Q'_H, \Sigma_p, \delta'_H, q'_{H,0}, F'_H)$ with $Q'_H = \biguplus_{v \in V} Q_v$,
 579 $\delta'_H = \biguplus_{v \in V} \delta_v \uplus \{q_{v_1,f} \xrightarrow{\varepsilon} q_{v_2,0} \mid (v_1, v_2) \in E\}$, $q'_{H,0} = q_{v^I,0}$, and $F'_H = \biguplus_{v \in V^F} q_{v,f}$. We
 580 determinise and minimise $(Q'_H, \Sigma_p, \delta'_H, q'_{H,0}, F'_H)$ to obtain $H \Downarrow_p := (Q_H, \Sigma_p, \delta_H, q_{H,0}, F_H)$
 581 which we define to be the projection by erasure of H on to p . The CSM formed from the
 582 projections by erasure $\{\{H \Downarrow_p\}_{p \in \mathcal{P}}\}$ is called erasure candidate implementation.

584 ► **Lemma 4.8** (Correctness of Projection by Erasure). Let H be an HMSC, p be a role, and $H \Downarrow_p$
 585 be its projection by erasure. Then, the following language equality holds: $\mathcal{L}(H \Downarrow_p) = \mathcal{L}(H) \Downarrow_{\Sigma_p}$.

586 The proof is straightforward and can be found in Appendix C.2. From this result and
 587 the construction of the canonical candidate implementation, it follows that the projection by
 588 erasure admits the same finite language.

589 ► **Corollary 4.9**. Let H be an HMSC, p be a role, $H \Downarrow_p$ be its projection by erasure, and A_p
 590 be the canonical candidate implementation. Then, it holds that $\mathcal{L}_{\text{fin}}(H \Downarrow_p) = \mathcal{L}_{\text{fin}}(A_p)$.

591 The projection by erasure can be computed effectively and is also deterministic. Thus, we
 592 use it in place of the canonical candidate implementation. Given a global type, the erasure
 593 candidate implementation for its HMSC encoding implements it if it is implementable.

594 ► **Theorem 4.10**. Let \mathbf{G} be a global type and $\{\{H(\mathbf{G}) \Downarrow_p\}_{p \in \mathcal{P}}\}$ be its erasure candidate
 595 implementation. If $\mathcal{L}_{\text{fin}}(\mathbf{G})$ is implementable⁴, then $\{\{H(\mathbf{G}) \Downarrow_p\}_{p \in \mathcal{P}}\}$ is deadlock free and
 596 $\mathcal{L}_{\text{fin}}(\{\{H(\mathbf{G}) \Downarrow_p\}_{p \in \mathcal{P}}\}) = \mathcal{L}_{\text{fin}}(\mathbf{G})$.

597 The proof can be found in Appendix C.3. This result does only account for finite languages
 598 so we extend it for infinite sequences.


599 ► **Lemma 4.11** (Erasure candidate implementation generalises to infinite language if imple-
 600 mentable). Let \mathbf{G} be a 0-reachable global type and $\{\{H(\mathbf{G}) \Downarrow_p\}_{p \in \mathcal{P}}\}$ be its erasure candidate
 601 implementation. If \mathbf{G} is implementable, then $\mathcal{L}_{\text{inf}}(\{\{H(\mathbf{G}) \Downarrow_p\}_{p \in \mathcal{P}}\}) = \mathcal{L}_{\text{inf}}(\mathbf{G})$.

602 The proof can be found in Appendix C.4.

603 So far, we have shown that, if \mathbf{G} is implementable, its erasure candidate implementation
 604 implements it. For this, we actually took the detour and showed the same for $H(\mathbf{G})$, the
 605 HMSC encoding of \mathbf{G} . For HMSCs, this is undecidable in general [63]. We show that, because
 606 of their conditions on choice, global types fall into the class of globally-cooperative HMSCs
 607 for which implementability is decidable.

608 ► **Definition 4.12** (Communication graph [39]). Let $M = (N, p, f, l, (\leq_p)_{p \in \mathcal{P}})$ be an MSC.
 609 The communication graph of M is a directed graph with node p for every role p that sends
 610 or receives a message in M and edges $p \rightarrow q$ if M contains a message from p to q , i.e., there
 611 is $e \in N$ such that $p(e) = p$ and $p(f(e)) = q$.

⁴ Implementability is lifted to languages as expected.



■ **Figure 6** An implementable HMSC which is not globally-cooperative with its implementation

612 It is important that the communication graph of M does not have a node for every role
613 but only the active ones, i.e., that send or receive in M .

614 ► **Definition 4.13** (Globally-cooperative HMSCs [39]). *An HMSC $H = (V, E, v^I, V^T, \mu)$ is
615 called globally-cooperative if for every loop, i.e., v_1, \dots, v_n with $(v_i, v_{i+1}) \in E$ for every
616 $1 \leq i < n$ and $(v_n, v_1) \in E$, the communication graph of $\mu(v_1) \dots \mu(v_n)$ is weakly connected.⁵*

617 We can check this directly on for a global type \mathbf{G} . It is straightforward to define a communi-
618 cation graph for words from Σ_{sync}^* . We check it on $\mathbf{GAut}(\mathbf{G})$: for each binder state, we check
619 the communication graph for the shortest trace to every corresponding recursion state.

620 ► **Theorem 4.14** (Thm. 3.7 [63]). *Let H be a globally-cooperative HMSC. Restricted to its
621 finite language $\mathcal{L}_{\text{fin}}(H)$, safe realisability is EXPSPACE-complete.*

622 ► **Lemma 4.15.** *Let \mathbf{G} be an implementable 0-reachable global type. Then, its HMSC
623 encoding $H(\mathbf{G})$ is globally-cooperative.*

624 The proof can be found in Appendix C.5 and is far from trivial. We explain the main
625 intuition for the proof with the following example where we exemplify why the same result
626 does not hold for HMSCs in general.

627 ► **Example 4.16** (Implementable HMSC that is not globally cooperative). Let us consider
628 the HMSC H_{ing} in Figure 6a. It is implementable but not globally-cooperative and not
629 representable with a global type. This protocol consists of three loops. In the first one,
630 p sends a message m to q while r sends a message m to s . This is the loop for which the
631 communication graph is not weakly connected. In the second one, only the interaction
632 between p and q is specified, while, in the third one, it is only the one between r and s . For
633 a protocol, which consists of the first and third loop only, an implementation can always
634 expose an execution with more interactions between p and q than the ones between r
635 and s due to the lack of synchronisation. Here, the additional second loop can make up for
636 such executions so any execution has a path that specifies it. Thus, this protocol can be
637 implemented with the CSM built from the FSMs illustrated in Figure 6b. In Appendix C.6,
638 we explain in detail why there is a path in H_{ing} for any trace of the CSM and how to modify
639 the example not to have final states with outgoing transitions.

⁵ Weakly connected means that, when considering every edge not to be directed, every node is connected with every other node.

640 ► **Theorem 4.17.** *Let \mathbf{G} be a 0-reachable global type with generalised choice. Checking*
 641 *implementability of \mathbf{G} is in EXPSPACE.*

642 **Proof.** We construct $H(\mathbf{G})$ from \mathbf{G} and check if it is globally-cooperative. For this, we apply
 643 the coNP-algorithm by Genest et al. [39] which is based on guessing a subgraph and checking
 644 its communication graph. If $H(\mathbf{G})$ is not globally cooperative, we know from Lemma 4.15
 645 that \mathbf{G} is not implementable. If $H(\mathbf{G})$ is globally cooperative, we check safe realisability for
 646 $H(\mathbf{G})$. By Theorem 4.14, this is in EXPSPACE. If $H(\mathbf{G})$ is not safely realisable, it trivially
 647 follows that \mathbf{G} is not implementable. If $H(\mathbf{G})$ is safely realisable, \mathbf{G} is implementable by the
 648 erasure candidate implementation with Theorem 4.10 and Lemma 4.11. ◀

649 With this, the implementability problem for global types with generalised choice is decidable.

650 ► **Corollary 4.18.** *Let \mathbf{G} be a 0-reachable global type with generalised choice. It is decidable*
 651 *whether \mathbf{G} is implementable and there is an algorithm to obtain its implementation.*

652 Our results also apply to the stronger notion of progress (Remark 2.14). This also
 653 entails that any sent message can eventually be received in an implementation — a property
 654 sometimes called *eventual reception* [61, Def. 4]. This notion only asks for the possibility but
 655 we can ensure that no role starves in a non-final state during an infinite execution in two
 656 ways. First, we can impose a (strong) fairness assumption — as imposed by Castagna et
 657 al. [19]. Second, we can require that every loop branch contains at least all roles that occur
 658 in interactions of any path with which the protocol can finish.

659 **The Odd Case of Infinite Loops Without Exits** In practice, it is reasonable to assume a
 660 mechanism to terminate a protocol for maintenance for instance, justifying the 0-Reachable-
 661 Assumption (p.15). In theory, one can think of protocols for which it does not hold. They
 662 would simply recurse indefinitely and can never terminate. This allows interesting behaviour
 663 like two sets of roles that do not interact with each other as the following example shows.

664 ► **Example 4.19.** Consider the following global type: $\mathbf{G} = \mu t. p \rightarrow q : m. r \rightarrow s : m. t.$
 665 This global type is basically the protocol which consists only of the first loop of H_{ing} from
 666 Example 4.16. It describes an infinite execution with two pairs of roles that send and receive
 667 messages independently. While this can be implemented for an infinite setting, such a loop
 668 could never be exited since the set of roles would need to synchronise on the number of times
 669 the loop was taken to satisfy language equality.

670 **Expressiveness of Local Types** Local types, like their global counterparts, have a distinct
 671 expression for termination: 0. Thus, if one considers the FSM of a local type, every final
 672 state has no outgoing transitions. Our proposed algorithm might produce state machines for
 673 which this is not true. However, the language of such a state machine cannot be represented
 674 as local type. Both, our construction and local types are deterministic. Thus, if there is a
 675 final state with an outgoing transition, there cannot be any state machine that only has final
 676 states without outgoing transitions.

677 In addition, the syntax prescribes the structure of the state machines similarly as for
 678 global types: state machines for local types are also ancestor-recursive, free of intermediate
 679 recursion, non-merging and dense (Proposition 3.6). We believe that this is rather a result of
 680 the classical projection operator than a design choice. For our algorithm, this is not the case.
 681 This raises two obvious directions for future work. On the one hand, it might be feasible
 682 to find rewriting techniques that take arbitrary state machines without final states with

683 outgoing transitions and transform them in a way such that they correspond to a local type.
 684 A naive approach to establish ancestor-recursiveness will most likely involve copying parts of
 685 the state machine. Such a rewriting would allow to re-use existing work, e.g., on sub-typing,
 686 which intuitively attempts to give freedom to implementations while preserving the soundness
 687 properties. On the other hand, one could also waive the syntactic restrictions and study
 688 sub-typing for this potentially more general class of local specifications.

689 **On Lower Bounds for Implementability** For general globally-cooperative HMSCs, i.e., that
 690 are not necessary the encoding of a global type, safe realisability is EXPSPACE-hard [63].
 691 This hardness result does not carry over for $H(\mathbf{G})$ of a global type \mathbf{G} . The construction
 692 exploits that HMSCs do not impose any restrictions on choice. Global types, however, require
 693 every branch to be chosen by a single sender.

694 **5 MSC Techniques for MST Verification**

695 In the previous section, we generalised results from the MSC literature to show decidability of
 696 the implementability problem for global types from MSTs. However, the resulting algorithm
 697 suffers from high complexity. This is also true for the original problem of safe realisability of
 698 HMSCs. In fact, the problem is undecidable for general HMSCs. Besides globally-cooperative
 699 HMSCs, further restrictions of HMSCs have been studied to obtain algorithms with better
 700 complexity for global types. The results from the previous section, in particular Theorem 4.10
 701 and Lemma 4.11, make most of these results applicable to the MST setting. One solely needs
 702 to check that the global type (or its HMSC encoding) belongs to the respective class. First, we
 703 transfer the algorithms for \mathcal{I} -closed HMSCs, which requires an HMSC not to exhibit certain
 704 anti-patterns of communication, to global types. Second, we explain approaches for HMSCs
 705 that introduced the idea of choice to HMSCs and a characterisation of implementable MSC
 706 languages. These can be a reasonable starting point for the design of complete algorithms for
 707 the implementability problem with better worst-case complexity. Third, we present a variant
 708 of the implementability problem. It can make unimplementable global types implementable
 709 without changing a protocol's structure but also help if the complexity of the previous
 710 algorithms is intractable. From now on, we may use the term implementability for HMSCs
 711 instead of safe realisability.

712 \mathcal{I} -closed Global Types

713 For globally-cooperative HMSCs, the implementability problem is EXPSPACE-complete.
 714 The membership in EXPSPACE was shown by reducing the problem to implementability of
 715 \mathcal{I} -closed HMSCs [63, Thm. 3.7]. These require the language of an HMSC to be closed with
 716 regard to an independence relation \mathcal{I} , where, intuitively, two interactions are independent if
 717 there is no role which is involved in both. Implementability for \mathcal{I} -closed HMSCs is PSPACE-
 718 complete [63, Thm. 3.6]. As for the EXPSPACE-hardness for globally-cooperative HMSCs,
 719 the PSPACE-hardness exploits features that cannot be modelled with global types and there
 720 might be algorithms with better worst-case complexity.

721 We adapt the definitions [63] to the MST setting. These consider atomic BMSCs, which
 722 are BMSCs that cannot be split further. With the HMSC encoding for global types, it is
 723 straightforward that atomic BMSCs correspond to individual interactions for global types.
 724 Thus, we define the independence relation \mathcal{I} on the alphabet Σ_{sync} .

725 **► Definition 5.1** (Independence relation \mathcal{I}). *We define the independence relation \mathcal{I} on Σ_{sync} :*

726 $\mathcal{I} := \{(p \rightarrow q : m, r \rightarrow s : m') \mid \{p, q\} \cap \{r, s\} \neq \emptyset\}$

727 We lift this to an equivalence relation $\equiv_{\mathcal{I}}$ on words as its transitive and reflexive closure:

728 $\equiv_{\mathcal{I}} := \{(u.x_1.x_2.w, u.x_2.x_1.w) \mid u, w \in \Sigma_{sync}^* \text{ and } (x_1, x_2) \in \mathcal{I}\}$

729 We define its closure for language $L \subseteq \Sigma_{sync}^*$: $\mathcal{C}^{\equiv_{\mathcal{I}}}(L) := \{u \in \Sigma_{sync}^* \mid \exists w \in L \text{ with } u \equiv_{\mathcal{I}} w\}$.

730 **► Definition 5.2** (\mathcal{I} -closedness for global types). Let \mathbf{G} be a global type \mathbf{G} . We say \mathbf{G} is
 731 \mathcal{I} -closed if $\mathcal{L}_{fin}(\mathbf{GAut}(\mathbf{G})) = \mathcal{C}^{\equiv_{\mathcal{I}}}(\mathcal{L}_{fin}(\mathbf{GAut}(\mathbf{G})))$.

732 Note that \mathcal{I} -closedness is defined on the state machine $\mathbf{GAut}(\mathbf{G})$ of \mathbf{G} with alphabet Σ_{sync}
 733 and not on its semantics $\mathcal{L}(\mathbf{G})$ with alphabet Σ_{async} .

734 **► Example 5.3.** The global type \mathbf{G}_{2BP} is \mathcal{I} -closed. Buyer a is involved in every interaction.
 735 Thus, for every consecutive interactions, there is a role that is involved in both.

736 **► Algorithm 1** (Checking if \mathbf{G} is \mathcal{I} -closed). Let \mathbf{G} be a global type \mathbf{G} . We construct the state
 737 machine $\mathbf{GAut}(\mathbf{G})$. We need to check every consecutive occurrence of elements from Σ_{sync}
 738 for words from $\mathcal{L}(\mathbf{GAut}(\mathbf{G}))$. For binder states, incoming and outgoing transition labels are
 739 always ε . This is why we slightly modify the state machine but preserve its language. We
 740 remove all variable states and rebend their only incoming transition to the state their only
 741 outgoing transition leads to. In addition, we merge binder states with their only successor.
 742 For every state q of this modified state machine, we consider the labels $x, y \in \Sigma_{sync}$ of every
 743 combination of incoming and outgoing transition of q . We check if $x \equiv_{\mathcal{I}} y$. If this is true for
 744 all x and y , we return true. If not, we return false.

745 **► Lemma 5.4.** A global type \mathbf{G} is \mathcal{I} -closed iff Algorithm 1 returns true.

746 The proof can be found in Appendix D. This shows that the presented algorithm can
 747 be used to check \mathcal{I} -closedness. The algorithm considers every state and all combinations of
 748 transitions leading to and from it.

749 **► Proposition 5.5.** For global type \mathbf{G} , checking \mathcal{I} -closedness of \mathbf{G} is in $O(|\mathbf{G}|^2)$.

750 The tree-like shape of $\mathbf{GAut}(\mathbf{G})$ might suggest that this check can be done in linear time.
 751 However, the following example shows that recursion can lead to a quadratic number of checks.

752 **► Example 5.6.** Consider the following global type for some n :

$$753 \quad \mu t. + \begin{cases} p \rightarrow q_0 : m_0. q_0 \rightarrow r_0 : m_0. r_0 \rightarrow s_0 : m_0. 0 \\ p \rightarrow q_1 : m_1. q_1 \rightarrow r_1 : m_1. r_1 \rightarrow s_1 : m_1. t \\ \vdots \\ p \rightarrow q_n : m_n. q_n \rightarrow r_n : m_n. r_n \rightarrow s_n : m_n. t \end{cases}$$

754 It is obvious that $(p \rightarrow q_i : m_i, q_i \rightarrow r_i : m_i) \notin \mathcal{I}$ and $(q_i \rightarrow r_i : m_i, r_i \rightarrow s_i : m_i) \notin \mathcal{I}$ for every i .
 755 Because of the recursion, we need to check if $(r_i \rightarrow s_i : m_i, p \rightarrow q_j : m_j)$ is in \mathcal{I} for every
 756 $0 \neq i \neq j$. This might lead to a quadratic number of checks.

757 If a global type \mathbf{G} is \mathcal{I} -closed, we can apply the respective results for its HMSC $H(\mathbf{G})$
 758 and the resulting CSM is also an implementation for \mathbf{G} . If not, we need to consider other
 759 approaches — where the last resort are the algorithms for globally-cooperative HMSCs.
 760 There are global types that are not \mathcal{I} -closed but implementable.

761 **► Example 5.7.** The following implementable global type is not \mathcal{I} -closed: $p \rightarrow q : m. r \rightarrow s : m. 0$.

762 Detecting Non-local Choice in HMSCs

763 For HMSCs, there is no restriction on branching. Similar to choice for global types, the idea
 764 of imposing restrictions on choice was studied for HMSCs [10, 68, 66, 44, 39]. We refer to
 765 Section 7 for an overview. Here, we focus on results that seem most promising for developing
 766 algorithms to check implementability of global types with better worst-case complexity. The
 767 work by Dan et al. [32] centers around the idea of non-local choice. Intuitively, non-local choice
 768 yields scenarios which makes it impossible to implement the language. In fact, if a language is
 769 not implementable, there is some non-local choice. Thus, checking implementability amounts
 770 to checking non-local choice freedom. For this definition, they showed insufficiency of Baker's
 771 condition [7] and reformulated the closure conditions for safe realisability by Alur et al. [3].
 772 In particular, they provide a definition which is based on projected words of a language in
 773 contrast to explicit choice. While it is straightforward to check their definition for finite
 774 collections of k BMSCs with n events in $O(k^2 \cdot |\mathcal{P}| + n \cdot |\mathcal{P}|)$, it is unclear how to check their
 775 condition for languages with infinitely many elements. The design of such a check is far from
 776 trivial as their definition does not give any insight about local behaviour and their algorithm
 777 heavily relies on the finite nature of finite collections of BMSCs. Still, we believe that the
 778 observations based on the closure conditions by Alur et al. [3], which provide a sound and
 779 complete characterisation of implementable languages, can be key to more efficient complete
 780 algorithms for the implementability problem for global types from MSTs.

781 Payload Implementability

782 A deadlock free CSM implements a global type if their languages are precisely the same. In
 783 the HMSC domain, a variant of the implementability problem has been studied. Intuitively,
 784 it allows to add fresh data to the payload of an existing message and protocol fidelity allows
 785 to omit the additional payload data. This allows to add synchronisation messages to existing
 786 interactions and can make unimplementable global types implementable without changing
 787 the structure of the protocol. It can also be used if a global type is rejected by projection
 788 and the run time of the previous algorithms is not acceptable.

789 ► **Definition 5.8** (Payload implementability). *Let L be a language with message alphabet \mathcal{V}_1 .
 790 We say that L is payload implementable if there is a message alphabet \mathcal{V}_2 for a deadlock free
 791 CSM $\{\{A_p\}_{p \in \mathcal{P}}\}$ with A_p over $\{p \triangleright q!m, p \triangleleft q?m \mid q \in \mathcal{P}, m \in \mathcal{V}_1 \times \mathcal{V}_2\}$ such that its language
 792 is the same when projecting on to the message alphabet \mathcal{V}_1 , i.e., $\mathcal{C}^\sim(L) = \mathcal{L}(\{\{A_p\}_{p \in \mathcal{P}}\}) \downarrow_{\mathcal{V}_1}$,
 793 where $(p \triangleright q!(m_1, m_2)) \downarrow_{\mathcal{V}_1} := p \triangleright q!m_1$ and $(q \triangleleft p?(m_1, m_2)) \downarrow_{\mathcal{V}_1} := q \triangleleft p?m_1$ and is lifted to
 794 words and languages as expected.*

795 Genest et al. [39] identified a class of HMSCs which is always payload implementable
 796 with a deadlock free CSM of linear size.

797 ► **Definition 5.9** (Local HMSCs [39]). *Let $H = (V, E, v^I, V^T, \mu)$ be an HMSC. We say that
 798 H is local if $\mu(v^I)$ has a unique minimal event and there is a function $\text{root}: V \rightarrow \mathcal{P}$ such
 799 that for every $(v, u) \in E$, it holds that $\mu(u)$ has a unique minimal event e and e belongs
 800 to $\text{root}(v)$, i.e., for $\mu(u) = (N, p, f, l, (\leq_p)_{p \in \mathcal{P}})$, we have that $p(e) = \text{root}(v)$ and $e \leq e'$ for
 801 every $e' \in N$.*

802 ► **Proposition 5.10** (Prop. 21 [39]). *For any local HMSC H , $\mathcal{L}_{\text{fin}}(H)$ is payload implementable.*

803 The algorithm to construct a deadlock free CSM [39, Sec. 5.2] suggests that the BMSCs
 804 for such HMSCs need to be maximal – in the sense that any vertex with a single successor is
 805 collapsed with its successor. If this was not the case, the result would claim that the language

806 of the following global type is payload implementable: $\mu t. + \begin{cases} p \rightarrow q : m_1 . r \rightarrow s : m_2 . t \\ p \rightarrow q : m_3. \end{cases}$. However,
 807 is is easy to see that it is not payload implementable since there is no interaction between p ,
 808 which decides whether to stay in the loop or not, and r . Thus, we cannot simply check
 809 whether $H(\mathbf{G})$ is local. In fact, it would always be. Instead, we first need to minimise it and
 810 then check whether it is local. If we collapse the two consecutive vertices with independent
 811 pairs of roles in this example, the HMSC is not local. The representation of the HMSC
 812 matters which shows that local as property is rather a syntactic than a semantic notion.

813 ► **Algorithm 2** (Checking if \mathbf{G} is local). *Let \mathbf{G} be a global type \mathbf{G} . We consider the finite*
 814 *trace w' of every longest branch-free, loop-free and non-initial run in the state machine*
 815 *$\mathbf{GAut}(\mathbf{G})$. We split the (synchronous) interactions into asynchronous events: $w = \text{split}(w') =$*
 816 *$w_1 \dots w_n$. We need to check if there is $u \sim w$ with $u = u_1 \dots u_n$ such that $u_1 \neq w_1$. For this,*
 817 *we can construct an MSC for w' [38, Sec. 3.1] and check if there is a single minimal event,*
 818 *because MSCs are closed under \approx [77, Lm. 5]. If this is the case for any trace w' , we return*
 819 *false. If not, we return true.*

820 It is straightforward that this mimics the corresponding check for the HMSC $H(\mathbf{G})$ and,
 821 including similar modifications as for Algorithm 1, the check can be done in $O(|\mathbf{G}|)$.

822 ► **Proposition 5.11.** *For a global type \mathbf{G} , Algorithm 2 returns true iff $H(\mathbf{G})$ is local.*

823 Ben-Abdallah and Leue [10] introduced local-choice HMSCs which are as expressive as
 824 local HMSCs. Their condition also uses a root-function and minimal events but quantifies
 825 over paths. Every local HMSC is a local-choice HMSC and every local-choice HMSC can be
 826 translated to a local HMSC, which accepts the same language, with a quadratic blow-up [39].
 827 It is straightforward to adapt the Algorithm 2 to check if a global type is local-choice. If this
 828 is the case, we translate the protocol and use the implementation for the translated protocol.

829 6 Implementability with Intra-role Reordering

830 In this section, we introduce a generalisation of the implementability problem that relaxes
 831 the total event order for each role and allows to reorder receive events. We prove that this
 832 generalisation is undecidable in general.

833 A Case for More Reordering

834 From the perspective of a single role, each word in its language consists of a sequence of
 835 receive and send events. Choice in global types happens by sending (and not by receiving).
 836 Because of this, one can argue that a role should be able to receive messages from different
 837 senders in any order between sending two messages. In practice, receiving a message can
 838 induce a task with non-trivial computation, which is not reflected in our model. Thus, such
 839 a reordering for a sequence of receive events can have outsized performance benefits. In
 840 addition, there are global types that can be implemented with regard to this generalised
 841 relation even if no (standard) implementation exists.

842 ► **Example 6.1** (Example for intra-role reordering). Let us consider a global type where a
 843 central coordinator p distributes independent tasks to different roles in rounds:

$$844 \quad \mathbf{G}_{TC} := \mu t. \begin{cases} p \rightarrow q_1 : \text{task} \dots p \rightarrow q_n : \text{task} . q_1 \rightarrow p : \text{result} \dots q_n \rightarrow p : \text{result} . t \\ p \rightarrow q_1 : \text{done} \dots p \rightarrow q_n : \text{done} . 0 \end{cases}$$

845 Since all tasks in each round are independent, p can benefit from receiving the results in the
 846 order they arrive instead of busy-waiting.

847 We generalise the indistinguishability relation \sim accordingly.

848 ► **Definition 6.2** (Intra-role indistinguishability relation). We define a family of intra-role
 849 indistinguishability relations $\approx_i \subseteq \Sigma^* \times \Sigma^*$, for $i \geq 0$ as follows. For all $w, u \in \Sigma^*$, $w \sim_i u$
 850 entails $w \approx_i u$. For $i = 1$, we define: if $q \neq r$, then $w.p \triangleleft q?m.p \triangleleft r?m'.u \sim_1 w.p \triangleleft r?m'.p \triangleleft q?m.u$.
 851 Based on this, we define \approx analogously to \sim . Let w, w', w'' be words s.t. $w \approx_1 w'$ and
 852 $w' \approx_i w''$ for some i . Then $w \approx_{i+1} w''$. We define $w \approx u$ if $w \approx_n u$ for some n . It is
 853 straightforward that \approx is an equivalence relation. Define $u \preceq_{\approx} v$ if there is $w \in \Sigma^*$ such that
 854 $u.w \approx v$. Observe that $u \sim v$ iff $u \preceq_{\approx} v$ and $v \preceq_{\approx} u$. We extend \approx to infinite words and
 855 languages as for \sim .

856 ► **Definition 6.3** (Implementability with intra-role reordering \approx). A global type \mathbf{G} is said to
 857 be implementable with regard to \approx if there exists a deadlock free CSM $\{\{A_p\}_{p \in \mathcal{P}}\}$ such that
 858 (i) $\mathcal{L}(\mathbf{G}) \subseteq \mathcal{C}^{\approx}(\mathcal{L}(\{\{A_p\}_{p \in \mathcal{P}}\}))$ and (ii) $\mathcal{C}^{\approx}(\mathcal{L}(\mathbf{G})) = \mathcal{C}^{\approx}(\mathcal{L}(\{\{A_p\}_{p \in \mathcal{P}}\}))$. We say that $\{\{A_p\}_{p \in \mathcal{P}}\}$
 859 \approx -implements \mathbf{G} .

860 In this section, we emphasise the indistinguishability relation, e.g., \approx -implementable,
 861 which is considered. We could have also followed the definition of \sim -implementability and
 862 required $\mathcal{C}^{\approx}(\mathcal{L}(\mathbf{G})) = \mathcal{L}(\{\{A_p\}_{p \in \mathcal{P}}\})$. This, however, requires the CSM to be closed under \approx .
 863 In general, this might not be possible with a finite number of states. In particular, if there is
 864 a loop without send event for a role, the labels in the loop would introduce an infinite closure
 865 if we require that $\mathcal{C}^{\approx}(\mathcal{L}(\mathbf{G})) \downarrow_{\Sigma_r} = \mathcal{L}(A_r)$.

866 ► **Example 6.4.** We consider a variant of $\mathbf{G}_{\text{TCLog}}$ from Example 6.1 with $n = 2$ where q_1 and
 867 q_2 send a log message to r after receiving the task and before sending the result back:

$$868 \quad \mathbf{G}_{\text{TCLog}} := \mu t. \begin{cases} p \rightarrow q_1 : \text{task} . p \rightarrow q_2 : \text{task} . q_1 \rightarrow r : \text{log} . q_2 \rightarrow r : \text{log} . q_1 \rightarrow p : \text{result} . q_2 \rightarrow p : \text{result} . t \\ p \rightarrow q_1 : \text{done} . p \rightarrow q_2 : \text{done} . 0 \end{cases}$$

869 There is no FSM for r that precisely accepts $\mathcal{C}^{\approx}(\mathcal{L}(\mathbf{G}_{\text{TCLog}})) \downarrow_{\Sigma_r}$ as it would need keep count
 870 of the difference at any point in time which can be unbounded. If we rely on the fact that
 871 q_1 and q_2 send the same number of log-messages to r , we can use an FSM A_r with a single
 872 state (both initial and final) with two transitions: one for the log-message from q_1 and q_2
 873 each, that lead back to the same state. For this, it holds that $\mathcal{C}^{\approx}(\mathcal{L}(\mathbf{G}_{\text{TCLog}})) \downarrow_{\Sigma_r} \subseteq \mathcal{L}(A_r)$.


874 This is why we chose a more permissive definition which is required to cover at least as
 875 much as specified in the global type (i) and the \approx -closure of both are the same (ii).

876 It is trivial that any \sim -implementation for a global type does also \approx -implement it.

877 ► **Proposition 6.5.** Let \mathbf{G} be a global type that is \sim -implemented by the CSM $\{\{A_p\}_{p \in \mathcal{P}}\}$.
 878 Then, $\{\{A_p\}_{p \in \mathcal{P}}\}$ also \approx -implements \mathbf{G} .

879 For instance, the task coordination protocol from Example 6.1 can be \sim -implemented as
 880 well as \approx -implemented by an erasure candidate implementation. Still, \approx -implementability
 881 gives more freedom and allows to consider all possible combinations of arrivals of results.
 882 In addition, \approx -implementability renders some global types implementable which would not
 883 be otherwise. For instance, those with a role that would need to receive different sequences,
 884 which are related by \approx , in different branches it cannot distinguish (yet).

885 ► **Example 6.6** (\approx -implementable but not \sim -implementable). Let us consider the following
 886 global type: $(p \rightarrow q : l . p \rightarrow r : m . q \rightarrow r : m . 0) + (p \rightarrow q : r . q \rightarrow r : m . p \rightarrow r : m . 0)$. This cannot
 887 be \sim -implemented because r would need to know the branch to receive the messages from
 888 p and q in the correct order. However, it is \approx -implementable. The FSMs for p and q can
 889 be obtained with projection by erasure. For r , we can have an FSM that only accepts
 890 $r \triangleleft p?m . r \triangleleft q?m$ but also an FSM which accepts $r \triangleleft q?m . r \triangleleft p?m$ in addition. Note that r
 891 does not learn the choice in the second FSM even if it branches. Hence, it would not be



■ **Figure 7** HMSC encoding $H(\mathbf{G}_{\text{MPCP}})$ of the MPCP encoding

892 implementable if it sent different messages in both branches later on. However, it could still
 893 learn by receiving and, afterwards, send different messages.

894 Implementability with Intra-role Reordering is Undecidable

895 Unfortunately, checking implementability with regard to \approx for global types (with directed
 896 choice) is undecidable. Intuitively, the reordering allows roles to drift arbitrarily far apart as
 897 the execution progresses which makes it hard to learn which choices were made.

898 We reduce the *Post Correspondence Problem* (PCP) [73] to the problem of checking
 899 implementability with regard to \approx . An instance of PCP over an alphabet Δ , $|\Delta| > 1$, is given
 900 by two finite lists (u_1, u_2, \dots, u_n) and (v_1, v_2, \dots, v_n) of finite words over Δ . A solution to the
 901 instance is a sequence of indices $(i_j)_{1 \leq j \leq k}$ with $k \geq 1$ and $1 \leq i_j \leq n$ for all $1 \leq j \leq k$, such
 902 that $u_{i_1} \dots u_{i_k} = v_{i_1} \dots v_{i_k}$. To be precise, we present a reduction from the modified PCP
 903 (MPCP) [76, Sec. 5.2], which is also undecidable. It simply requires that a match starts
 904 with a specific pair — in our case we choose the pair with index 1. It is possible to directly
 905 reduce from PCP but the reduction of MPCP is more concise. Intuitively, we require that
 906 the solution starts with the first pair so there exists no trivial solution and choosing a single
 907 pair is more concise than all possible ones. Our encoding is the following global type where
 908 $x \in \{u, v\}$, $[x_i]$ denotes a sequence of message interactions with message $x_i[1], \dots, x_i[k]$ each
 909 for x_i of length k , message $c-x$ indicates choosing tile set x , and message $ack-x$ indicates
 910 acknowledging the tile set x :

$$\begin{aligned}
 911 \quad \mathbf{G}_{\text{MPCP}} &:= + \left\{ \begin{array}{l} G(u, r \rightarrow p: ack-u. 0) \\ G(v, r \rightarrow p: ack-v. 0) \end{array} \right. \quad \text{with} \quad \left\{ \begin{array}{l} p \rightarrow q: 1. p \rightarrow r: 1. q \rightarrow r: [x_1]. t_1 \\ \vdots \\ p \rightarrow q: n. p \rightarrow r: n. q \rightarrow r: [x_n]. t_1 \\ p \rightarrow q: d. p \rightarrow r: d. q \rightarrow r: d. X \end{array} \right. \\
 912 \quad G(x, X) &:= p \rightarrow q: c-x. p \rightarrow q: 1. p \rightarrow r: 1. q \rightarrow r: [x_1]. \mu t_1. + \left\{ \begin{array}{l} \vdots \\ p \rightarrow q: n. p \rightarrow r: n. q \rightarrow r: [x_n]. t_1 \\ p \rightarrow q: d. p \rightarrow r: d. q \rightarrow r: d. X \end{array} \right. \\
 913
 \end{aligned}$$

914 The HMSC encoding $H(\mathbf{G}_{\text{MPCP}})$ is illustrated in Figure 7. Intuitively, r eventually needs to
 915 know which branch was taken in order to match $ack-x$ with $c-x$ from the beginning. However,
 916 it can only know if there is no solution to the MPCP instance. In the full proof in Appendix E,
 917 we show that \mathbf{G}_{MPCP} is \approx -implementable iff the MPCP instance has no solution.

919 ► **Theorem 6.7.** *Checking implementability with regard to \approx for global types with directed
 920 choice is undecidable.*

921 This result carries over to HMSCs if we consider safe realisability with regard to \approx .

922 ► **Definition 6.8** (Safe realisability with regard to \approx). *An HMSC H is said to be safely
 923 realisable with regard to \approx if there exists a deadlock-free CSM $\{\{A_p\}\}_{p \in \mathcal{P}}$ such that the
 924 following holds: (i) $\mathcal{L}(H) \subseteq \mathcal{C}^\approx(\mathcal{L}(\{\{A_p\}\}_{p \in \mathcal{P}}))$ and (ii) $\mathcal{C}^\approx(\mathcal{L}(H)) = \mathcal{C}^\approx(\mathcal{L}(\{\{A_p\}\}_{p \in \mathcal{P}}))$.*

925 ► **Corollary 6.9.** *Checking safe realisability with regard to \approx for HMSCs is undecidable.*

926 In fact, the HMSC encoding for \mathbf{G}_{MPCP} satisfies a number of channel restrictions. The
 927 HMSC $H(\mathbf{G}_{\text{MPCP}})$ is existentially 1-bounded, 1-synchronisable and half-duplex [77]. For
 928 details on these channel restrictions, we refer to work by Stutz and Zufferey [77, Sec. 3.1].

929 ► Remark 6.10 (Sending to the rescue). The MPCP encoding only works since receive events
 930 can be reordered unboundedly in an execution. If we amended the definition of \approx to give
 931 each receive event a budget that depletes with every reordering, this encoding would not be
 932 possible. Alternatively, one could require every active role in a loop to send at least once.
 933 This also prevents such an unbounded reordering. For such restrictions on the considered
 934 indistinguishability relation, the corresponding implementability problem likely becomes
 935 decidable. We leave a detailed analysis for future work.

936 7 Related Work

937 In this section, we solely cover related work which we have not discussed before.
 938 Session types originate in process algebra and were first introduced by Honda et al. [46] for
 939 binary session. For systems with more than two roles, they have been extended to multiparty
 940 session types [48]. Their connection to linear logic [41] has been studied subsequently [37, 84,
 941 17]. In this work, we explained MST frameworks with classical projection operators. Other
 942 approaches do not focus on projection but, for instance, employ model checking [75] or only
 943 apply ideas from MST without the need for global types [61].

944 Our decidability result applies to global types with generalised choice. There are only
 945 few MST frameworks that effectively allow to generate local types from global types with
 946 generalised choice for an asynchronous setting. Castellani et al. [20] consider a synchronous
 947 setting. The same holds for the work by Jongmans and Yoshida [55] but their parallel
 948 operator allows to model some asynchrony with bag semantics. The setting in the work
 949 by Lange et al. [59] yields semantics similar to Petri nets. To the best of our knowledge,
 950 the work by Castagna et al. [19] is the only one to attempt completeness for global types
 951 with generalised choice. However, their notion of completeness allows to omit redundant
 952 executions for underspecified global types [19, Def. 4.1]. Their conditions, given as inference
 953 rules, are not effective and their algorithmically checkable conditions can only exploit local
 954 information to disambiguate choices. In contrast, Majumdar et al. [64] employ a global
 955 availability analysis but, as classical projection operator, it suffers from the shortcomings
 956 presented in this work. For a detailed overview of frameworks allowing generalised choice,
 957 we refer to their work [64]. They also present a counterexample to the implementability
 958 conditions formulated for Choreography Automata [8]. The global types by Dagnino et
 959 al. [31] specify send and receive events independently but each term requires to send to a
 960 single receiver and to receive from a single sender upon branching. They present a sound
 961 and complete type inference algorithm that infers all global types for a given system.

962 Here, we do not distinguish between local types and implementations but use the local
 963 types directly as implementations. Intuitively, subtyping studies possibilities to give freedom
 964 in the implementation while preserving the soundness properties. The intra-role indistin-
 965 guishability relation \approx , which allows to reorder receive events for a role, resembles subtyping
 966 to some extent, e.g., the work by Cutner et al. [30]. A detailed investigation of this relation
 967 is beyond the scope of this work. For details on subtyping, we refer to work by Chen et
 968 al. [27, 26], Lange and Yoshida [60], and Bravetti et al. [16].

969 Various extensions to make MST verification applicable to more scenarios were studied:
 970 for instance delegation [47, 48, 21], dependent session types [80, 35, 81], parametrised session
 971 types [24, 35], gradual session types [51], or dynamic self-adaption [43]. Context-free session

972 types [79, 56] provide a more expressive way to specify protocols in the MST domain.
 973 Recently, research on fault-tolerant MSTs [83, 9, 72] investigated ways to weaken the strict
 974 assumptions about reliable channels.

975 Choreographic programming [29, 40, 45] applies a similar approach as MSTs: they allow
 976 to specify a global protocol specification with joint send and receive events and project to
 977 end-point views. As for Pirouette by Hirsch and Garg [45], there are first mechanisations of
 978 MST frameworks [78, 22, 53, 52].

979 The connection of MSTs and CSMs was studied soon after MSTs had been proposed [34].
 980 CSMs are known to be Turing-powerful [15]. Decidable classes have been obtained for
 981 different semantics, e.g., half-duplex communication for two roles [23], input-bounded [12],
 982 and unreliable/lossy channels [2], as well for restricted communication topology [71, 82].
 983 Similar restrictions for CSMs are existential boundedness [38] and synchronisability [14, 42].
 984 It was shown that global types can only express existentially 1-bounded, 1-synchronisable
 985 and half-duplex communication [77] while Bollig et al. [13] established a connection between
 986 synchronisability and MSO logic.

987 Globally-cooperative HMSCs were independently introduced by Morin [67] as c-HMSCs.
 988 Their communication graph is weakly connected. The class of bounded HMSCs [5] requires
 989 it to be strongly connected. Historically, it was introduced before the class of globally-
 990 cooperative HMSCs and, after the latter has been introduced, safe realisability for bounded
 991 HMSCs was also shown to be EXPSPACE-complete [63]. This class was independently
 992 introduced as regular HMSCs by Muscholl and Peled [69]. Both terms are justified: the
 993 language generated by a *regular* HMSC is regular and every *bounded* HMSC can be imple-
 994 mented with universally bounded channels. In fact, a HMSC is bounded if and only if it is a
 995 globally-cooperative and it has universally bounded channels [39, Prop. 4].

996 We cover approaches which introduced the idea of choice to HMSCs that were not
 997 discussed in Section 5. Ben-Abdallah and Leue [10] approached the realisability problem by
 998 defining and detecting non-local choice, which are basically choices not made by a single role.
 999 Their semantics, however, incorporates queuing behaviour that renders their systems finite-
 1000 state. Another line of work [68, 66] identified that non-local choice and implied scenarios are
 1001 strongly coupled. An implied scenario is an execution, which is not specified in the HMSC,
 1002 but any candidate implementation necessarily exposes. Initial attempts by Muccini [68]
 1003 yielded contradictory observations as shown by Mooij et al. [66] so they proposed variants of
 1004 non-local choice but they accept the implied scenarios from such choices as given in the HMSC.
 1005 H elou et and Jard [44] pointed out that the absence of non-local choice does not guarantee
 1006 implementability but just less ambiguity. They proposed reconstructibility which shall entail
 1007 implementability. Majumdar et al. [64] showed that their notion of reconstructibility, with
 1008 the requirement of unique messages, is quite restrictive but also flawed.

1009 In addition to local HMSCs, Genest et al. [39] also introduced locally-cooperative HMSCs.
 1010 Intuitively, they require for every two successors that each of their communication graphs as
 1011 well as their concatenation’s communication graph is weakly connected but it is only known
 1012 that checking weak realisability (the one allowing deadlocks) has linear time complexity.
 1013 Non-FIFO channel semantics has also been considered for HMSCs for which the complexity
 1014 for safe realisability does not change while it has an influence on weak realisability [63].

1015 **8 Conclusion**

1016 We have proven decidability of the implementability problem for global types with generalised
 1017 choice from MSTs — under the mild assumption that protocols can (but do not need to)

1018 terminate. To point at the origin for incompleteness of classical projection operators, we gave
1019 a visual explanation of the projection with various merge operators on finite state machines,
1020 which define the semantics of global and local types. To prove decidability, we formally
1021 related the implementability problem for global types with the safe realisability problem for
1022 HMSCs. While safe realisability is undecidable, we showed that implementable global types
1023 do always belong to the class of globally-cooperative HMSCs. There are global types that
1024 are outside of this class but the syntax of global types allowed us to prove that those can not
1025 be implemented. Another key was the extension of the HMSC results to infinite executions.
1026 We gave a comprehensive overview of MSC techniques and adapted some to the MST setting.
1027 Furthermore, we introduced a performance-oriented generalisation of the implementability
1028 problem which, however, we proved to be undecidable in general.

- 1030 1 Martín Abadi and Leslie Lamport. The existence of refinement mappings. In *Proceedings of the*
1031 *Third Annual Symposium on Logic in Computer Science (LICS '88), Edinburgh, Scotland, UK,*
1032 *July 5-8, 1988*, pages 165–175. IEEE Computer Society, 1988. doi:10.1109/LICS.1988.5115.
- 1033 2 Parosh Aziz Abdulla, Ahmed Bouajjani, and Bengt Jonsson. On-the-fly analysis of systems
1034 with unbounded, lossy FIFO channels. In Alan J. Hu and Moshe Y. Vardi, editors, *Computer*
1035 *Aided Verification, 10th International Conference, CAV'98, Vancouver, BC, Canada, June 28*
1036 *- July 2, 1998, Proceedings*, volume 1427 of *Lecture Notes in Computer Science*, pages 305–318.
1037 Springer, 1998. doi:10.1007/BFb0028754.
- 1038 3 Rajeev Alur, Kousha Etessami, and Mihalis Yannakakis. Inference of message sequence charts.
1039 *IEEE Trans. Software Eng.*, 29(7):623–633, 2003. doi:10.1109/TSE.2003.1214326.
- 1040 4 Rajeev Alur, Kousha Etessami, and Mihalis Yannakakis. Realizability and verification of MSC
1041 graphs. *Theor. Comput. Sci.*, 331(1):97–114, 2005. doi:10.1016/j.tcs.2004.09.034.
- 1042 5 Rajeev Alur and Mihalis Yannakakis. Model checking of message sequence charts. In
1043 Jos C. M. Baeten and Sjouke Mauw, editors, *CONCUR '99: Concurrency Theory, 10th*
1044 *International Conference, Eindhoven, The Netherlands, August 24-27, 1999, Proceedings*,
1045 volume 1664 of *Lecture Notes in Computer Science*, pages 114–129. Springer, 1999. doi:
1046 10.1007/3-540-48320-9\10.
- 1047 6 Davide Ancona, Viviana Bono, Mario Bravetti, Joana Campos, Giuseppe Castagna, Pierre-
1048 Malo Deniérou, Simon J. Gay, Nils Gesbert, Elena Giachino, Raymond Hu, Einar Broch
1049 Johnsen, Francisco Martins, Viviana Mascardi, Fabrizio Montesi, Romyana Neykova,
1050 Nicholas Ng, Luca Padovani, Vasco T. Vasconcelos, and Nobuko Yoshida. Behavioral
1051 types in programming languages. *Found. Trends Program. Lang.*, 3(2-3):95–230, 2016.
1052 doi:10.1561/25000000031.
- 1053 7 Paul Baker, Paul Bristow, Clive Jervis, David J. King, Robert Thomson, Bill Mitchell, and
1054 Simon Burton. Detecting and resolving semantic pathologies in UML sequence diagrams. In
1055 Michel Wermelinger and Harald C. Gall, editors, *Proceedings of the 10th European Software*
1056 *Engineering Conference held jointly with 13th ACM SIGSOFT International Symposium on*
1057 *Foundations of Software Engineering, 2005, Lisbon, Portugal, September 5-9, 2005*, pages
1058 50–59. ACM, 2005. doi:10.1145/1081706.1081716.
- 1059 8 Franco Barbanera, Ivan Lanese, and Emilio Tuosto. Choreography automata. In Simon
1060 Bliudze and Laura Bocchi, editors, *Coordination Models and Languages - 22nd IFIP WG 6.1*
1061 *International Conference, COORDINATION 2020, Held as Part of the 15th International*
1062 *Federated Conference on Distributed Computing Techniques, DisCoTec 2020, Valletta, Malta,*
1063 *June 15-19, 2020, Proceedings*, volume 12134 of *Lecture Notes in Computer Science*, pages
1064 86–106. Springer, 2020. doi:10.1007/978-3-030-50029-0\6.
- 1065 9 Adam D. Barwell, Alceste Scalas, Nobuko Yoshida, and Fangyi Zhou. Generalised multiparty
1066 session types with crash-stop failures. In Bartek Klin, Slawomir Lasota, and Anca Muscholl,
1067 editors, *33rd International Conference on Concurrency Theory, CONCUR 2022, September*
1068 *12-16, 2022, Warsaw, Poland*, volume 243 of *LIPICs*, pages 35:1–35:25. Schloss Dagstuhl -
1069 Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICs.CONCUR.2022.35.
- 1070 10 Hanène Ben-Abdallah and Stefan Leue. Syntactic detection of process divergence and non-
1071 local choice in message sequence charts. In Ed Brinksma, editor, *Tools and Algorithms for*
1072 *Construction and Analysis of Systems, Third International Workshop, TACAS '97, Enschede,*
1073 *The Netherlands, April 2-4, 1997, Proceedings*, volume 1217 of *Lecture Notes in Computer*
1074 *Science*, pages 259–274. Springer, 1997. doi:10.1007/BFb0035393.
- 1075 11 Laura Bocchi, Maurizio Murgia, Vasco Thudichum Vasconcelos, and Nobuko Yoshida.
1076 Asynchronous timed session types - from duality to time-sensitive processes. In Luís
1077 Caires, editor, *Programming Languages and Systems - 28th European Symposium on Pro-*
1078 *gramming, ESOP 2019, Held as Part of the European Joint Conferences on Theory and*
1079 *Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6-11, 2019, Proceed-*

- ings, volume 11423 of *Lecture Notes in Computer Science*, pages 583–610. Springer, 2019. doi:10.1007/978-3-030-17184-1_21.
- 12 Benedikt Bollig, Alain Finkel, and Amrita Suresh. Bounded reachability problems are decidable in FIFO machines. In Igor Konnov and Laura Kovács, editors, *31st International Conference on Concurrency Theory, CONCUR 2020, September 1-4, 2020, Vienna, Austria (Virtual Conference)*, volume 171 of *LIPICs*, pages 49:1–49:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPICs.CONCUR.2020.49.
- 13 Benedikt Bollig, Cinzia Di Giusto, Alain Finkel, Laetitia Laversa, Étienne Lozes, and Amrita Suresh. A unifying framework for deciding synchronizability. In Serge Haddad and Daniele Varacca, editors, *32nd International Conference on Concurrency Theory, CONCUR 2021, August 24-27, 2021, Virtual Conference*, volume 203 of *LIPICs*, pages 14:1–14:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICs.CONCUR.2021.14.
- 14 Ahmed Bouajjani, Constantin Enea, Kailiang Ji, and Shaz Qadeer. On the completeness of verifying message passing programs under bounded asynchrony. In Hana Chockler and Georg Weissenbacher, editors, *Computer Aided Verification - 30th International Conference, CAV 2018, Held as Part of the Federated Logic Conference, FloC 2018, Oxford, UK, July 14-17, 2018, Proceedings, Part II*, volume 10982 of *Lecture Notes in Computer Science*, pages 372–391. Springer, 2018. doi:10.1007/978-3-319-96142-2_23.
- 15 Daniel Brand and Pitro Zafiropulo. On communicating finite-state machines. *J. ACM*, 30(2):323–342, 1983. doi:10.1145/322374.322380.
- 16 Mario Bravetti, Marco Carbone, and Gianluigi Zavattaro. On the boundary between decidability and undecidability of asynchronous session subtyping. *Theor. Comput. Sci.*, 722:19–51, 2018. doi:10.1016/j.tcs.2018.02.010.
- 17 Luís Caires, Frank Pfenning, and Bernardo Toninho. Linear logic propositions as session types. *Math. Struct. Comput. Sci.*, 26(3):367–423, 2016. doi:10.1017/S0960129514000218.
- 18 Marco Carbone, Kohei Honda, N. Yoshida, R. Milner, G. Brown, and Steve Ross-Talbot. A theoretical basis of communication-centred concurrent programming. 2005.
- 19 Giuseppe Castagna, Mariangiola Dezani-Ciancaglini, and Luca Padovani. On global types and multi-party session. *Log. Methods Comput. Sci.*, 8(1), 2012. doi:10.2168/LMCS-8(1:24)2012.
- 20 Ilaria Castellani, Mariangiola Dezani-Ciancaglini, and Paola Giannini. Reversible sessions with flexible choices. *Acta Informatica*, 56(7-8):553–583, 2019. doi:10.1007/s00236-019-00332-y.
- 21 Ilaria Castellani, Mariangiola Dezani-Ciancaglini, Paola Giannini, and Ross Horne. Global types with internal delegation. *Theor. Comput. Sci.*, 807:128–153, 2020. doi:10.1016/j.tcs.2019.09.027.
- 22 David Castro-Perez, Francisco Ferreira, Lorenzo Gheri, and Nobuko Yoshida. Zoid: a DSL for certified multiparty computation: from mechanised metatheory to certified multiparty processes. In Stephen N. Freund and Eran Yahav, editors, *PLDI '21: 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation, Virtual Event, Canada, June 20-25, 2021*, pages 237–251. ACM, 2021. doi:10.1145/3453483.3454041.
- 23 Gérard Cécé and Alain Finkel. Verification of programs with half-duplex communication. *Inf. Comput.*, 202(2):166–190, 2005. doi:10.1016/j.ic.2005.05.006.
- 24 Minas Charalambides, Peter Dinges, and Gul A. Agha. Parameterized, concurrent session types for asynchronous multi-actor interactions. *Sci. Comput. Program.*, 115-116:100–126, 2016. doi:10.1016/j.scico.2015.10.006.
- 25 Ruofei Chen, Stephanie Balzer, and Bernardo Toninho. Ferrite: A judgmental embedding of session types in rust. In Karim Ali and Jan Vitek, editors, *36th European Conference on Object-Oriented Programming, ECOOP 2022, June 6-10, 2022, Berlin, Germany*, volume 222 of *LIPICs*, pages 22:1–22:28. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICs.ECOOP.2022.22.
- 26 Tzu-Chun Chen, Mariangiola Dezani-Ciancaglini, Alceste Scalas, and Nobuko Yoshida. On the preciseness of subtyping in session types. *Log. Methods Comput. Sci.*, 13(2), 2017. doi:10.23638/LMCS-13(2:12)2017.

- 1132 27 Tzu-Chun Chen, Mariangiola Dezani-Ciancaglini, and Nobuko Yoshida. On the preciseness of
1133 subtyping in session types. In Olaf Chitil, Andy King, and Olivier Danvy, editors, *Proceedings*
1134 *of the 16th International Symposium on Principles and Practice of Declarative Programming,*
1135 *Kent, Canterbury, United Kingdom, September 8-10, 2014*, pages 135–146. ACM, 2014. doi:
1136 10.1145/2643135.2643138.
- 1137 28 Mario Coppo, Mariangiola Dezani-Ciancaglini, Luca Padovani, and Nobuko Yoshida. A gentle
1138 introduction to multiparty asynchronous session types. In Marco Bernardo and Einar Broch
1139 Johnsen, editors, *Formal Methods for Multicore Programming - 15th International School on*
1140 *Formal Methods for the Design of Computer, Communication, and Software Systems, SFM*
1141 *2015, Bertinoro, Italy, June 15-19, 2015, Advanced Lectures*, volume 9104 of *Lecture Notes in*
1142 *Computer Science*, pages 146–178. Springer, 2015. doi:10.1007/978-3-319-18941-3_4.
- 1143 29 Luís Cruz-Filipe and Fabrizio Montesi. A core model for choreographic programming. *Theor.*
1144 *Comput. Sci.*, 802:38–66, 2020. doi:10.1016/j.tcs.2019.07.005.
- 1145 30 Zak Cutner, Nobuko Yoshida, and Martin Vassor. Deadlock-free asynchronous message
1146 reordering in rust with multiparty session types. In Jaejin Lee, Kunal Agrawal, and Michael F.
1147 Spear, editors, *PPoPP '22: 27th ACM SIGPLAN Symposium on Principles and Practice of*
1148 *Parallel Programming, Seoul, Republic of Korea, April 2 - 6, 2022*, pages 246–261. ACM, 2022.
1149 doi:10.1145/3503221.3508404.
- 1150 31 Francesco Dagnino, Paola Giannini, and Mariangiola Dezani-Ciancaglini. Deconfined global
1151 types for asynchronous sessions. In Ferruccio Damiani and Ornela Dardha, editors, *Co-*
1152 *ordination Models and Languages - 23rd IFIP WG 6.1 International Conference, COOR-*
1153 *DINATION 2021, Held as Part of the 16th International Federated Conference on Dis-*
1154 *tributed Computing Techniques, DisCoTec 2021, Valletta, Malta, June 14-18, 2021, Pro-*
1155 *ceedings*, volume 12717 of *Lecture Notes in Computer Science*, pages 41–60. Springer, 2021.
1156 doi:10.1007/978-3-030-78142-2_3.
- 1157 32 Haitao Dan, Robert M. Hierons, and Steve Counsell. Non-local choice and implied scenarios.
1158 In José Luiz Fiadeiro, Stefania Gnesi, and Andrea Maggiolo-Schettini, editors, *8th IEEE*
1159 *International Conference on Software Engineering and Formal Methods, SEFM 2010, Pisa,*
1160 *Italy, 13-18 September 2010*, pages 53–62. IEEE Computer Society, 2010. doi:10.1109/SEFM.
1161 2010.14.
- 1162 33 Ankush Das, Stephanie Balzer, Jan Hoffmann, Frank Pfenning, and Ishani Santurkar. Resource-
1163 aware session types for digital contracts. In *34th IEEE Computer Security Foundations*
1164 *Symposium, CSF 2021, Dubrovnik, Croatia, June 21-25, 2021*, pages 1–16. IEEE, 2021.
1165 doi:10.1109/CSF51468.2021.00004.
- 1166 34 Pierre-Malo Deniérou and Nobuko Yoshida. Multiparty session types meet communicating
1167 automata. In Helmut Seidl, editor, *Programming Languages and Systems - 21st European*
1168 *Symposium on Programming, ESOP 2012, Held as Part of the European Joint Conferences on*
1169 *Theory and Practice of Software, ETAPS 2012, Tallinn, Estonia, March 24 - April 1, 2012.*
1170 *Proceedings*, volume 7211 of *Lecture Notes in Computer Science*, pages 194–213. Springer,
1171 2012. doi:10.1007/978-3-642-28869-2_10.
- 1172 35 Pierre-Malo Deniérou, Nobuko Yoshida, Andi Bejleri, and Raymond Hu. Parameterised
1173 multiparty session types. *Log. Methods Comput. Sci.*, 8(4), 2012. doi:10.2168/LMCS-8(4:
1174 6)2012.
- 1175 36 Manuel Fähndrich, Mark Aiken, Chris Hawblitzel, Orion Hodson, Galen C. Hunt, James R.
1176 Larus, and Steven Levi. Language support for fast and reliable message-based communication
1177 in singularity OS. In Yolande Berbers and Willy Zwaenepoel, editors, *Proceedings of the*
1178 *2006 EuroSys Conference, Leuven, Belgium, April 18-21, 2006*, pages 177–190. ACM, 2006.
1179 doi:10.1145/1217935.1217953.
- 1180 37 Simon J. Gay and Vasco Thudichum Vasconcelos. Linear type theory for asynchronous session
1181 types. *J. Funct. Program.*, 20(1):19–50, 2010. doi:10.1017/S0956796809990268.

- 1182 38 Blaise Genest, Dietrich Kuske, and Anca Muscholl. On communicating automata with bounded
1183 channels. *Fundam. Inform.*, 80(1-3):147–167, 2007. URL: [http://content.iospress.com/
1184 articles/fundamenta-informaticae/fi80-1-3-09](http://content.iospress.com/articles/fundamenta-informaticae/fi80-1-3-09).
- 1185 39 Blaise Genest, Anca Muscholl, Helmut Seidl, and Marc Zeitoun. Infinite-state high-level
1186 mscs: Model-checking and realizability. *J. Comput. Syst. Sci.*, 72(4):617–647, 2006. doi:
1187 10.1016/j.jcss.2005.09.007.
- 1188 40 Saverio Giallorenzo, Fabrizio Montesi, Marco Peressotti, David Richter, Guido Salvaneschi,
1189 and Pascal Weisenburger. Multiparty languages: The choreographic and multitier cases (pearl).
1190 In Anders Møller and Manu Sridharan, editors, *35th European Conference on Object-Oriented
1191 Programming, ECOOP 2021, July 11-17, 2021, Aarhus, Denmark (Virtual Conference)*, volume
1192 194 of *LIPICs*, pages 22:1–22:27. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
1193 doi:10.4230/LIPICs.ECOOP.2021.22.
- 1194 41 Jean-Yves Girard. Linear logic. *Theor. Comput. Sci.*, 50:1–102, 1987. doi:10.1016/
1195 0304-3975(87)90045-4.
- 1196 42 Cinzia Di Giusto, Laetitia Laversa, and Étienne Lozes. On the k-synchronizability of systems.
1197 In Jean Goubault-Larrecq and Barbara König, editors, *Foundations of Software Science and
1198 Computation Structures - 23rd International Conference, FOSSACS 2020, Held as Part of
1199 the European Joint Conferences on Theory and Practice of Software, ETAPS 2020, Dublin,
1200 Ireland, April 25-30, 2020, Proceedings*, volume 12077 of *Lecture Notes in Computer Science*,
1201 pages 157–176. Springer, 2020. doi:10.1007/978-3-030-45231-5_9.
- 1202 43 Paul Harvey, Simon Fowler, Ornela Dardha, and Simon J. Gay. Multiparty session types for
1203 safe runtime adaptation in an actor language. In Anders Møller and Manu Sridharan, editors,
1204 *35th European Conference on Object-Oriented Programming, ECOOP 2021, July 11-17, 2021,
1205 Aarhus, Denmark (Virtual Conference)*, volume 194 of *LIPICs*, pages 10:1–10:30. Schloss
1206 Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICs.ECOOP.2021.10.
- 1207 44 Loïc Hélouët. Some pathological message sequence charts, and how to detect them. In
1208 Rick Reed and Jeanne Reed, editors, *SDL 2001: Meeting UML, 10th International SDL
1209 Forum Copenhagen, Denmark, June 27-29, 2001, Proceedings*, volume 2078 of *Lecture Notes
1210 in Computer Science*, pages 348–364. Springer, 2001. doi:10.1007/3-540-48213-X_22.
- 1211 45 Andrew K. Hirsch and Deepak Garg. Pirouette: higher-order typed functional choreographies.
1212 *Proc. ACM Program. Lang.*, 6(POPL):1–27, 2022. doi:10.1145/3498684.
- 1213 46 Kohei Honda. Types for dyadic interaction. In Eike Best, editor, *CONCUR '93, 4th In-
1214 ternational Conference on Concurrency Theory, Hildesheim, Germany, August 23-26, 1993,
1215 Proceedings*, volume 715 of *Lecture Notes in Computer Science*, pages 509–523. Springer, 1993.
1216 doi:10.1007/3-540-57208-2_35.
- 1217 47 Kohei Honda, Vasco Thudichum Vasconcelos, and Makoto Kubo. Language primitives and
1218 type discipline for structured communication-based programming. In Chris Hankin, editor,
1219 *Programming Languages and Systems - ESOP'98, 7th European Symposium on Programming,
1220 Held as Part of the European Joint Conferences on the Theory and Practice of Software,
1221 ETAPS'98, Lisbon, Portugal, March 28 - April 4, 1998, Proceedings*, volume 1381 of *Lecture
1222 Notes in Computer Science*, pages 122–138. Springer, 1998. doi:10.1007/BFb0053567.
- 1223 48 Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty asynchronous session types. In
1224 George C. Necula and Philip Wadler, editors, *Proceedings of the 35th ACM SIGPLAN-SIGACT
1225 Symposium on Principles of Programming Languages, POPL 2008, San Francisco, California,
1226 USA, January 7-12, 2008*, pages 273–284. ACM, 2008. doi:10.1145/1328438.1328472.
- 1227 49 Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty asynchronous session types.
1228 *J. ACM*, 63(1):9:1–9:67, 2016. doi:10.1145/2827695.
- 1229 50 Raymond Hu and Nobuko Yoshida. Explicit connection actions in multiparty session types. In
1230 Marieke Huisman and Julia Rubin, editors, *Fundamental Approaches to Software Engineering
1231 - 20th International Conference, FASE 2017, Held as Part of the European Joint Conferences
1232 on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017*,

- 1233 *Proceedings*, volume 10202 of *Lecture Notes in Computer Science*, pages 116–133. Springer,
1234 2017. doi:10.1007/978-3-662-54494-5_7.
- 1235 **51** Atsushi Igarashi, Peter Thiemann, Yuya Tsuda, Vasco T. Vasconcelos, and Philip Wadler.
1236 Gradual session types. *J. Funct. Program.*, 29:e17, 2019. doi:10.1017/S0956796819000169.
- 1237 **52** Jules Jacobs, Stephanie Balzer, and Robbert Krebbers. Connectivity graphs: a method for
1238 proving deadlock freedom based on separation logic. *Proc. ACM Program. Lang.*, 6(POPL):1–
1239 33, 2022. doi:10.1145/3498662.
- 1240 **53** Jules Jacobs, Stephanie Balzer, and Robbert Krebbers. Multiparty GV: functional multiparty
1241 session types with certified deadlock freedom. *Proc. ACM Program. Lang.*, 6(ICFP):466–495,
1242 2022. doi:10.1145/3547638.
- 1243 **54** Thomas Bracht Laumann Jespersen, Philip Munksgaard, and Ken Friis Larsen. Session
1244 types for rust. In Patrick Bahr and Sebastian Erdweg, editors, *Proceedings of the 11th ACM*
1245 *SIGPLAN Workshop on Generic Programming, WGP@ICFP 2015, Vancouver, BC, Canada,*
1246 *August 30, 2015*, pages 13–22. ACM, 2015. doi:10.1145/2808098.2808100.
- 1247 **55** Sung-Shik Jongmans and Nobuko Yoshida. Exploring type-level bisimilarity towards more
1248 expressive multiparty session types. In Peter Müller, editor, *Programming Languages and*
1249 *Systems - 29th European Symposium on Programming, ESOP 2020, Held as Part of the*
1250 *European Joint Conferences on Theory and Practice of Software, ETAPS 2020, Dublin,*
1251 *Ireland, April 25-30, 2020, Proceedings*, volume 12075 of *Lecture Notes in Computer Science*,
1252 pages 251–279. Springer, 2020. doi:10.1007/978-3-030-44914-8_10.
- 1253 **56** Alex C. Keizer, Henning Basold, and Jorge A. Pérez. Session coalgebras: A coalgebraic view
1254 on regular and context-free session types. *ACM Trans. Program. Lang. Syst.*, 44(3):18:1–18:45,
1255 2022. doi:10.1145/3527633.
- 1256 **57** Dimitrios Kouzapas, Ramunas Gutkovas, A. Laura Voinea, and Simon J. Gay. A session type
1257 system for asynchronous unreliable broadcast communication. *CoRR*, abs/1902.01353, 2019.
1258 URL: <http://arxiv.org/abs/1902.01353>, arXiv:1902.01353.
- 1259 **58** Julien Lange, Nicholas Ng, Bernardo Toninho, and Nobuko Yoshida. A static verification
1260 framework for message passing in go using behavioural types. In Michel Chaudron, Ivica
1261 Crnkovic, Marsha Chechik, and Mark Harman, editors, *Proceedings of the 40th International*
1262 *Conference on Software Engineering, ICSE 2018, Gothenburg, Sweden, May 27 - June 03,*
1263 *2018*, pages 1137–1148. ACM, 2018. doi:10.1145/3180155.3180157.
- 1264 **59** Julien Lange, Emilio Tuosto, and Nobuko Yoshida. From communicating machines to graphical
1265 choreographies. In Sriram K. Rajamani and David Walker, editors, *Proceedings of the 42nd*
1266 *Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL*
1267 *2015, Mumbai, India, January 15-17, 2015*, pages 221–232. ACM, 2015. doi:10.1145/2676726.
1268 2676964.
- 1269 **60** Julien Lange and Nobuko Yoshida. On the undecidability of asynchronous session subtyping.
1270 In Javier Esparza and Andrzej S. Murawski, editors, *Foundations of Software Science and*
1271 *Computation Structures - 20th International Conference, FOSSACS 2017, Held as Part of*
1272 *the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala,*
1273 *Sweden, April 22-29, 2017, Proceedings*, volume 10203 of *Lecture Notes in Computer Science*,
1274 pages 441–457, 2017. doi:10.1007/978-3-662-54458-7_26.
- 1275 **61** Julien Lange and Nobuko Yoshida. Verifying asynchronous interactions via communicating
1276 session automata. In Isil Dillig and Serdar Tasiran, editors, *Computer Aided Verification - 31st*
1277 *International Conference, CAV 2019, New York City, NY, USA, July 15-18, 2019, Proceedings,*
1278 *Part I*, volume 11561 of *Lecture Notes in Computer Science*, pages 97–117. Springer, 2019.
1279 doi:10.1007/978-3-030-25540-4_6.
- 1280 **62** Sam Lindley and J. Garrett Morris. Embedding session types in haskell. In Geoffrey Mainland,
1281 editor, *Proceedings of the 9th International Symposium on Haskell, Haskell 2016, Nara, Japan,*
1282 *September 22-23, 2016*, pages 133–145. ACM, 2016. doi:10.1145/2976002.2976018.
- 1283 **63** Markus Lohrey. Realizability of high-level message sequence charts: closing the gaps. *Theor.*
1284 *Comput. Sci.*, 309(1-3):529–554, 2003. doi:10.1016/j.tcs.2003.08.002.

- 1285 **64** Rupak Majumdar, Madhavan Mukund, Felix Stutz, and Damien Zufferey. Generalising
1286 projection in asynchronous multiparty session types. In Serge Haddad and Daniele Varacca,
1287 editors, *32nd International Conference on Concurrency Theory, CONCUR 2021, August*
1288 *24-27, 2021, Virtual Conference*, volume 203 of *LIPICs*, pages 35:1–35:24. Schloss Dagstuhl -
1289 Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICs.CONCUR.2021.35.
- 1290 **65** Rupak Majumdar, Marcus Pirron, Nobuko Yoshida, and Damien Zufferey. Motion session types
1291 for robotic interactions (brave new idea paper). In Alastair F. Donaldson, editor, *33rd European*
1292 *Conference on Object-Oriented Programming, ECOOP 2019, July 15-19, 2019, London, United*
1293 *Kingdom*, volume 134 of *LIPICs*, pages 28:1–28:27. Schloss Dagstuhl - Leibniz-Zentrum für
1294 Informatik, 2019. doi:10.4230/LIPICs.ECOOP.2019.28.
- 1295 **66** Arjan J. Mooij, Nicolae Goga, and Judi Romijn. Non-local choice and beyond: Intricacies of
1296 MSC choice nodes. In Maura Cerioli, editor, *Fundamental Approaches to Software Engineering,*
1297 *8th International Conference, FASE 2005, Held as Part of the Joint European Conferences on*
1298 *Theory and Practice of Software, ETAPS 2005, Edinburgh, UK, April 4-8, 2005, Proceedings,*
1299 volume 3442 of *Lecture Notes in Computer Science*, pages 273–288. Springer, 2005. doi:
1300 10.1007/978-3-540-31984-9_21.
- 1301 **67** Rémi Morin. Recognizable sets of message sequence charts. In Helmut Alt and Afonso Ferreira,
1302 editors, *STACS 2002, 19th Annual Symposium on Theoretical Aspects of Computer Science,*
1303 *Antibes - Juan les Pins, France, March 14-16, 2002, Proceedings*, volume 2285 of *Lecture Notes*
1304 *in Computer Science*, pages 523–534. Springer, 2002. doi:10.1007/3-540-45841-7_43.
- 1305 **68** Henry Muccini. Detecting implied scenarios analyzing non-local branching choices. In Mauro
1306 Pezzè, editor, *Fundamental Approaches to Software Engineering, 6th International Conference,*
1307 *FASE 2003, Held as Part of the Joint European Conferences on Theory and Practice of Software,*
1308 *ETAPS 2003, Warsaw, Poland, April 7-11, 2003, Proceedings*, volume 2621 of *Lecture Notes*
1309 *in Computer Science*, pages 372–386. Springer, 2003. doi:10.1007/3-540-36578-8_26.
- 1310 **69** Anca Muscholl and Doron A. Peled. Message sequence graphs and decision problems on
1311 mazurkiewicz traces. In Miroslaw Kutylowski, Leszek Pacholski, and Tomasz Wierzbicki,
1312 editors, *Mathematical Foundations of Computer Science 1999, 24th International Symposium,*
1313 *MFCS'99, Szklarska Poreba, Poland, September 6-10, 1999, Proceedings*, volume 1672 of *Lecture*
1314 *Notes in Computer Science*, pages 81–91. Springer, 1999. doi:10.1007/3-540-48340-3_8.
- 1315 **70** Rumyana Neykova, Raymond Hu, Nobuko Yoshida, and Fahd Abdeljallal. A session type
1316 provider: compile-time API generation of distributed protocols with refinements in f#. In
1317 Christophe Dubach and Jingling Xue, editors, *Proceedings of the 27th International Conference*
1318 *on Compiler Construction, CC 2018, February 24-25, 2018, Vienna, Austria*, pages 128–138.
1319 ACM, 2018. doi:10.1145/3178372.3179495.
- 1320 **71** Wuxu Peng and S. Purushothaman. Analysis of a class of communicating finite state machines.
1321 *Acta Informatica*, 29(6/7):499–522, 1992. doi:10.1007/BF01185558.
- 1322 **72** Kirstin Peters, Uwe Nestmann, and Christoph Wagner. Fault-tolerant multiparty session
1323 types. In Mohammad Reza Mousavi and Anna Philippou, editors, *Formal Techniques for*
1324 *Distributed Objects, Components, and Systems - 42nd IFIP WG 6.1 International Conference,*
1325 *FORTE 2022, Held as Part of the 17th International Federated Conference on Distributed*
1326 *Computing Techniques, DisCoTec 2022, Lucca, Italy, June 13-17, 2022, Proceedings*, volume
1327 13273 of *Lecture Notes in Computer Science*, pages 93–113. Springer, 2022. doi:10.1007/
1328 978-3-031-08679-3_7.
- 1329 **73** Emil L. Post. A variant of a recursively unsolvable problem. *Bulletin of the American*
1330 *Mathematical Society*, 52:264–268, 1946.
- 1331 **74** Alceste Scalas and Nobuko Yoshida. Lightweight session programming in scala. In Shriram
1332 Krishnamurthi and Benjamin S. Lerner, editors, *30th European Conference on Object-Oriented*
1333 *Programming, ECOOP 2016, July 18-22, 2016, Rome, Italy*, volume 56 of *LIPICs*, pages
1334 21:1–21:28. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPICs.
1335 ECOOP.2016.21.

- 1336 **75** Alceste Scalas and Nobuko Yoshida. Less is more: multiparty session types revisited. *Proc.*
1337 *ACM Program. Lang.*, 3(POPL):30:1–30:29, 2019. doi:10.1145/3290343.
- 1338 **76** Michael Sipser. *Introduction to the theory of computation*. PWS Publishing Company, 1997.
- 1339 **77** Felix Stutz and Damien Zufferey. Comparing channel restrictions of communicating state
1340 machines, high-level message sequence charts, and multiparty session types. In Pierre Ganty
1341 and Dario Della Monica, editors, *Proceedings of the 13th International Symposium on Games,*
1342 *Automata, Logics and Formal Verification, GandALF 2022, Madrid, Spain, September 21-23,*
1343 *2022*, volume 370 of *EPTCS*, pages 194–212, 2022. doi:10.4204/EPTCS.370.13.
- 1344 **78** Peter Thiemann. Intrinsically-typed mechanized semantics for session types. In Ekaterina
1345 Komendantskaya, editor, *Proceedings of the 21st International Symposium on Principles and*
1346 *Practice of Programming Languages, PPDP 2019, Porto, Portugal, October 7-9, 2019*, pages
1347 19:1–19:15. ACM, 2019. doi:10.1145/3354166.3354184.
- 1348 **79** Peter Thiemann and Vasco T. Vasconcelos. Context-free session types. In Jacques Garrigue,
1349 Gabriele Keller, and Eijiro Sumii, editors, *Proceedings of the 21st ACM SIGPLAN International*
1350 *Conference on Functional Programming, ICFP 2016, Nara, Japan, September 18-22, 2016*,
1351 pages 462–475. ACM, 2016. doi:10.1145/2951913.2951926.
- 1352 **80** Bernardo Toninho, Luís Caires, and Frank Pfenning. Dependent session types via intuitionistic
1353 linear type theory. In Peter Schneider-Kamp and Michael Hanus, editors, *Proceedings of*
1354 *the 13th International ACM SIGPLAN Conference on Principles and Practice of Declarative*
1355 *Programming, July 20-22, 2011, Odense, Denmark*, pages 161–172. ACM, 2011. doi:10.1145/
1356 2003476.2003499.
- 1357 **81** Bernardo Toninho and Nobuko Yoshida. Depending on session-typed processes. In Christel
1358 Baier and Ugo Dal Lago, editors, *Foundations of Software Science and Computation Structures -*
1359 *21st International Conference, FOSSACS 2018, Held as Part of the European Joint Conferences*
1360 *on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14-20, 2018,*
1361 *Proceedings*, volume 10803 of *Lecture Notes in Computer Science*, pages 128–145. Springer,
1362 2018. doi:10.1007/978-3-319-89366-2_7.
- 1363 **82** Salvatore La Torre, P. Madhusudan, and Gennaro Parlato. Context-bounded analysis of
1364 concurrent queue systems. In C. R. Ramakrishnan and Jakob Rehof, editors, *Tools and*
1365 *Algorithms for the Construction and Analysis of Systems, 14th International Conference,*
1366 *TACAS 2008, Held as Part of the Joint European Conferences on Theory and Practice of*
1367 *Software, ETAPS 2008, Budapest, Hungary, March 29-April 6, 2008. Proceedings*, volume
1368 4963 of *Lecture Notes in Computer Science*, pages 299–314. Springer, 2008. doi:10.1007/
1369 978-3-540-78800-3_21.
- 1370 **83** Malte Viering, Raymond Hu, Patrick Eugster, and Lukasz Ziarek. A multiparty session typing
1371 discipline for fault-tolerant event-driven distributed programming. *Proc. ACM Program. Lang.*,
1372 5(OOPSLA):1–30, 2021. doi:10.1145/3485501.
- 1373 **84** Philip Wadler. Propositions as sessions. *J. Funct. Program.*, 24(2-3):384–418, 2014. doi:
1374 10.1017/S095679681400001X.
- 1375 **85** Nobuko Yoshida and Lorenzo Gheri. A very gentle introduction to multiparty session types.
1376 In Dang Van Hung and Meenakshi D’Souza, editors, *Distributed Computing and Internet*
1377 *Technology - 16th International Conference, ICDCIT 2020, Bhubaneswar, India, January 9-12,*
1378 *2020, Proceedings*, volume 11969 of *Lecture Notes in Computer Science*, pages 73–93. Springer,
1379 2020. doi:10.1007/978-3-030-36987-3_5.
- 1380 **86** Nobuko Yoshida, Raymond Hu, Rumyana Neykova, and Nicholas Ng. The scribble protocol
1381 language. In Martín Abadi and Alberto Lluch-Lafuente, editors, *Trustworthy Global Computing*
1382 *- 8th International Symposium, TGC 2013, Buenos Aires, Argentina, August 30-31, 2013,*
1383 *Revised Selected Papers*, volume 8358 of *Lecture Notes in Computer Science*, pages 22–41.
1384 Springer, 2013. doi:10.1007/978-3-319-05119-2_3.

A Definitions for Section 2: Multiparty Session Types

A.1 Semantics of Communicating State Machines [77, App. A.4]

With $\text{Chan} = \{\langle p, q \rangle \mid p, q \in \mathcal{P}, p \neq q\}$, we denote the set of channels. The set of global states of a CSM is given by $\prod_{p \in \mathcal{P}} Q_p$. Given a global state q , q_p denotes the state of p in q . A *configuration* of a CSM \mathcal{A} is a pair (q, ξ) , where q is a global state and $\xi : \text{Chan} \rightarrow \mathcal{V}^\infty$ is a mapping of each channel to its current content. The initial configuration (q_0, ξ_ε) consists of a global state q_0 where the state of each role is the initial state $q_{0,p}$ of A_p and a mapping ξ_ε , which maps each channel to the empty word ε . A configuration (q, ξ) is said to be *final* iff each individual local state q_p is final for every p and ξ is ξ_ε .

The global transition relation \rightarrow is defined as follows:

- $(q, \xi) \xrightarrow{p \triangleright q!m} (q', \xi')$ if $(q_p, p \triangleright q!m, q'_p) \in \delta_p$, $q_r = q'_r$ for every role $r \neq p$, $\xi'(\langle p, q \rangle) = \xi(\langle p, q \rangle) \cdot m$ and $\xi'(c) = \xi(c)$ for every other channel $c \in \text{Chan}$.
- $(q, \xi) \xrightarrow{q \triangleleft p?m} (q', \xi')$ if $(q_q, q \triangleleft p?m, q'_q) \in \delta_q$, $q_r = q'_r$ for every role $r \neq q$, $\xi(\langle p, q \rangle) = m \cdot \xi'(\langle p, q \rangle)$ and $\xi'(c) = \xi(c)$ for every other channel $c \in \text{Chan}$.
- $(q, \xi) \xrightarrow{\varepsilon} (q', \xi)$ if $(q_p, \varepsilon, q'_p) \in \delta_p$ for some role p , and $q_q = q'_q$ for every role $q \neq p$.

A run of the CSM always starts with an initial configuration (q_0, ξ_0) , and is a finite or infinite sequence $(q_0, \xi_0) \xrightarrow{w_0} (q_1, \xi_1) \xrightarrow{w_1} \dots$ for which $(q_i, \xi_i) \xrightarrow{w_i} (q_{i+1}, \xi_{i+1})$. The word $w_0 w_1 \dots \in \Sigma^\infty$ is said to be the *trace* of the run. A run is called maximal if it is either infinite or finite and ends in a final configuration. As before, the trace of a maximal run is maximal. The language $\mathcal{L}(\mathcal{A})$ of the CSM \mathcal{A} consists of its set of maximal traces.

B Additional Explanation for Different Merge Operators on FSMs from Section 3

Visual Explanation of the Parametric Projection Operator: Collapsing Erasure Here, we describe *collapsing erasure* more formally. Let \mathbf{G} be some global type and r be the role on to which we project. We apply the parametric projection operator to the state machine $\mathbf{GAut}(\mathbf{G})$. It projects each transition label on to the respective event for role r : every forward transition label $p \rightarrow q : m$ turns to $p \triangleright q!m$ if $r = p$, $p \triangleleft q?m$ if $r = q$, and ε otherwise. Then, it collapses neutral states with a single successor: $q_{1|2}$ replaces two states q_1 and q_2 if $q_1 \xrightarrow{\varepsilon} q_2$ is the only forward transition for q_1 and $q_1 \xrightarrow{x} q_2$ for $x \neq \varepsilon$ in $\delta_{\mathbf{GAut}(\mathbf{G})}$. In case there is only a backward transition from q_1 to q_2 , the state $q_{1|2}$ is also final. This accounts for loops a role is not part of.

We call this procedure collapsing erasure as it erases interactions that do not belong to a role and collapses some states. It is common to all the presented merge operators. This procedure yields a state machine over Σ_r . It is straightforward that it is still ancestor-recursive, free from intermediate recursion and non-merging. However, it might not be dense. In fact, it is not dense if r is not involved in some choice with more than one branch.

Parametric Merge in the Visual Explanation The parametric projection operator applies the merge operator for these cases. Visually, these correspond precisely to the remaining neutral states (since all neutral states with a single successor have been collapsed). For instance, we have a neutral state q_1 with $q_1 \xrightarrow{\varepsilon} q_2$ and $q_1 \xrightarrow{\varepsilon} q_3$ for $q_2 \neq q_3$. Through the parametric projection operator, the merge operator may be indirectly called recursively. Thus, we explain the merge operators for two states (and their cones) in general. No information is

1428 propagated when the merge operator recurses and recursion variables are never unfolded.
 1429 Thus, we can ignore backward transitions and consider the cones of q_2 and q_3 . Intuitively,
 1430 we iteratively apply the merge operator from lower to higher levels. However, we might need
 1431 descend again when merge is applied recursively. Similar to the syntactic version, we do only
 1432 explain the 2-ary case but the reasoning easily lifts to the n -ary case.

1433 **Visual Explanation of Plain Merge** The plain merge is not applied recursively. Thus, we
 1434 consider q_1 with $q_1 \xrightarrow{\varepsilon} q_2$ and $q_1 \xrightarrow{\varepsilon} q_3$ for $q_2 \neq q_3$ such that q_1 has the lowest level for which
 1435 this holds. Hence, we can assume that each cone of q_2 and q_3 does not contain neutral states.
 1436 Then, the plain merge is only defined if there is an isomorphism between the states of both
 1437 cones that satisfy the following conditions:

- 1438 ■ it preserves the transition labels and hence the kind of states, and
- 1439 ■ if a state has a backward transition to a state outside of the cone, its isomorphic state
 1440 has a transition to the some state

1441 If defined, the result is q_1 with its cone (and q_2 with its cone is removed).

1442 **Visual Explanation of Semi-full Merge** The semi-full merge applies itself recursively. Thus,
 1443 we consider two states $q_2 \neq q_3$ that shall be merged. As before, we can assume that each
 1444 cone of q_2 and q_3 does not contain neutral states. In addition to plain merge, the semi-full
 1445 merge allows to merge receive states. For these, we introduce a new receive state $q_{2|3}$ from
 1446 which all new transitions start. For all possible transitions from q_2 and q_3 , we check if there
 1447 is a transition with the same label from the other state. For the ones not in common, we
 1448 simply add the respective transition with the state it leads to and its respective cone. For
 1449 the ones in common, we recursively check if the two states, which both transitions lead to,
 1450 can be merged. If not, the semi-full merge is undefined. If so, we add the original transition
 1451 to the state of the respective merge and keep its cone.

1452 **Visual Explanation of Full Merge** Intuitively, the full merge simply applies the idea of
 1453 the semi-full merge to another case. For the semi-full merge, one can recursively apply the
 1454 merge operator when a reception was common between two states to merge. The full merge
 1455 operator allows to descend for recursion variable binders.

1456 **C Formalisation for Section 4:** 1457 **Implementability for Global Types from MSTs is Decidable**

1458 **C.1 Definitions for Section 4.1**

1459 ► **Definition C.1** (Concatenation of MSCs [77]). Let $M_i = (N_i, p_i, f_i, l_i, (\leq_p^i)_{p \in \mathcal{P}})$ for $i \in \{1, 2\}$
 1460 where M_1 is a BMSC and M_2 is an MSC with disjoint sets of events, i.e., $N_1 \cap N_2 = \emptyset$. We
 1461 define their concatenation $M_1 \cdot M_2$ as the MSC $M = (N, p, f, l, (\leq_p)_{p \in \mathcal{P}})$ where:

- 1462 ■ $N := N_1 \cup N_2,$
- 1463 ■ for $\zeta \in \{p, f, l\}$: $\zeta(e) := \begin{cases} \zeta(e) & \text{if } e \in N_1 \\ \zeta(e) & \text{if } e \in N_2 \end{cases},$ and
- 1464 ■ $\forall p \in \mathcal{P} : \leq_p := \leq_p^1 \cup \leq_p^2 \cup \{(e_1, e_2) \mid e_1 \in N_1 \wedge e_2 \in N_2 \wedge p(e_1) = p(e_2) = p\}.$

1465 ► **Definition C.2** (Language of an HMSC [77]). Let $H = (V, E, v^I, V^T, \mu)$ be an HMSC. The
 1466 language of H is defined as

$$\begin{aligned}
 1467 \quad \mathcal{L}(H) &:= \{w \mid w \in \mathcal{L}(\mu(v_1)\mu(v_2)\dots\mu(v_n)) \text{ with } v_1 = v^I \wedge \forall 0 \leq i < n : (v_i, v_{i+1}) \in E \wedge v_n \in V^T\} \\
 1468 \quad &\cup \{w \mid w \in \mathcal{L}(\mu(v_1)\mu(v_2)\dots) \text{ with } v_1 = v^I \wedge \forall i \geq 0 : (v_i, v_{i+1}) \in E\}.
 \end{aligned}$$

1470 C.2 Proof of Lemma 4.8: Projection by Erasure is Correct

1471 Let $H = (V, E, v^I, V^T, \mu)$ be an HMSC. For every $v \in V$, it is straightforward that the
 1472 construction of $\mu(v)\downarrow_p$ yields $\mathcal{L}(\mu(v))\downarrow_{\Sigma_p} = \mathcal{L}(\mu(v)\downarrow_p)$ (1). We recall that \sim does not reorder
 1473 events by the same role: $w \sim w'$ for $w \in \Sigma_p$ iff $w = w'$ (2).

1474 The following reasoning proves the claim where the first equivalence follows from the
 1475 construction of the transition relation of $H\downarrow_p$:

$$\begin{aligned}
 1476 \quad & w \in \mathcal{L}(H\downarrow_p) \\
 1477 \quad & \Leftrightarrow w = w_1 \dots, \text{ there is a path } v_1, \dots \text{ in } H \text{ and } w_i \in \mathcal{L}(\mu(v_i)\downarrow_p) \text{ for every } i \\
 1478 \quad & \stackrel{(1)}{\Leftrightarrow} w = w_1 \dots, \text{ there is a path } v_1, \dots \text{ in } H \text{ and } w_i \in \mathcal{L}(\mu(v_i))\downarrow_{\Sigma_p} \text{ for every } i \\
 1479 \quad & \stackrel{(2)}{\Leftrightarrow} w \in \mathcal{L}(H)\downarrow_{\Sigma_p}
 \end{aligned}$$

1480
1481

1482 C.3 Proof of Theorem 4.10: 1483 Erasure Candidate Implementation is Sufficient

1484 We first use the correctness of the global type encoding (Theorem 4.5) to observe that
 1485 $\mathcal{L}_{\text{fin}}(\mathbf{G}) = \mathcal{L}_{\text{fin}}(H(\mathbf{G}))$. Theorem 13 by Alur et al. [3] states that the canonical candidate
 1486 implementation implements $\mathcal{L}_{\text{fin}}(H(\mathbf{G}))$ if it is implementable. Corollary 4.9 and the fact
 1487 that the FSM for each role is deterministic by construction allows us to replace every A_p
 1488 from the canonical candidate implementation with the projection by erasure $H(\mathbf{G})\downarrow_p$ for
 1489 every role p which proves the claim. \blacktriangleleft

1490 C.4 Proof of Lemma 4.11: 1491 Erasure Candidate Implementation Generalises to Infinite Case

1492 Let us assume that \mathbf{G} is implementable. From Theorem 4.10, we know that $\{\{H(\mathbf{G})\downarrow_p\}_{p \in \mathcal{P}}\}_{p \in \mathcal{P}}$
 1493 is deadlock free and $\mathcal{L}_{\text{fin}}(\{\{H(\mathbf{G})\downarrow_p\}_{p \in \mathcal{P}}\}_{p \in \mathcal{P}}) = \mathcal{L}_{\text{fin}}(\mathbf{G})$. We prove the claim by showing both
 1494 inclusions.

1495 **First**, we show that $\mathcal{L}_{\text{inf}}(\{\{H(\mathbf{G})\downarrow_p\}_{p \in \mathcal{P}}\}_{p \in \mathcal{P}}) \subseteq \mathcal{L}_{\text{inf}}(\mathbf{G})$. For this direction, let w be a
 1496 word in $\mathcal{L}_{\text{inf}}(\{\{H(\mathbf{G})\downarrow_p\}_{p \in \mathcal{P}}\}_{p \in \mathcal{P}})$. We need to show that there is a run ρ in $\text{GAut}(\mathbf{G})$ such that
 1497 $w \preceq^{\omega} \text{split}(\text{trace}(\rho))$. From the 0-Reachable-Assumption (p.15), we know that for every
 1498 $u \in \text{pref}(w)$, it holds that $u \in \text{pref}(\mathcal{L}_{\text{fin}}(\mathbf{G}))$. Thus, there exists a finite run ρ (that does not
 1499 necessarily end in a final state) and u' such that $u.u' \sim \text{trace}(\rho)$. We call ρ a witness run.
 1500 Intuitively, we will need to argue that every such witness run for u can be extended when
 1501 appending the next event x from w to obtain ux . In general, this does not hold for every
 1502 choice of witness run. However, because of monotonicity, any run (or rather a prefix of it)
 1503 for an extension ux can also be used as witness run for u . Thus, we make use of the idea of
 1504 prophecy variables [1] and assume an oracle which picks the correct witness run for every
 1505 prefix u . This oracle does not restrict the next possible events in any way. From here, we
 1506 apply the same idea as Majumdar et al. for the proof of Lemma 41 [64]. We construct a
 1507 tree \mathcal{T} such that each node represents a run ρ of some finite prefix w' of w . The root's label is
 1508 the empty run. For every node labelled with ρ , the children's extend ρ by a single transition.
 1509 The tree \mathcal{T} is finitely branching by construction of $\text{GAut}(\mathbf{G})$ for every role p . With König's
 1510 Lemma, we obtain an infinite path in \mathcal{T} and thus an infinite run ρ in $\text{GAut}_{\text{async}}(\mathbf{G})$ with
 1511 $w \preceq^{\omega} \text{trace}(\rho)$. From this, it follows that $w \in \mathcal{L}_{\text{inf}}(\mathbf{G})$.

1512 **Second**, we show that $\mathcal{L}_{\text{inf}}(\mathbf{G}) \subseteq \mathcal{L}_{\text{inf}}(\{\{H(\mathbf{G})\downarrow_p\}_{p \in \mathcal{P}}\}_{p \in \mathcal{P}})$. Let w be a word in $\mathcal{L}_{\text{inf}}(\mathbf{G})$.
 1513 Eventually, we will apply the same reasoning with König's lemma to obtain an infinite run

1514 in $\{\{H(\mathbf{G})\downarrow_p\}_{p \in \mathcal{P}}\}$ for w . Inspired from the first statement of Lemma 25 by Majumdar et
 1515 al. [64], we show:

- 1516 (i) for every prefix $w' \in \text{pref}(w)$, there is a run ρ' in $\{\{H(\mathbf{G})\downarrow_p\}_{p \in \mathcal{P}}\}$ such that $w' \preceq \text{trace}(\rho')$,
 1517 and
 1518 (ii) for every extension $w'x$ where x is the next event in w , the run ρ' can be extended

1519 We prove Claim (i) first. We first observe that, with the 0-Reachable-Assumption (p.15),
 1520 there is an extension w'' of w' with $w'' \in \mathcal{L}(\mathbf{G})$. By construction, we know that there is a run
 1521 ρ'' in $\{\{H(\mathbf{G})\downarrow_p\}_{p \in \mathcal{P}}\}$ for w'' . For ρ' , we can simply take the prefix of ρ'' which matches w' .
 1522 This proves Claim (i).

1523 Now, let us prove Claim (ii). Similar to the first case, we will use prophecy variables and
 1524 an oracle to pick the correct witness run that we can extend. Again, because of monotonicity,
 1525 any run (or rather a prefix of it) for an extension $w'x$ can also be used as witness run for w' .
 1526 As before, we make use of the idea of prophecy variables [1], assume an oracle which picks
 1527 the correct witness run for every prefix w' , and this oracle does not restrict the roles in any
 1528 way. From this, Claim (ii) follows.

1529 From here, we (again) use the same reasoning as Majumdar et al. for the proof of
 1530 Lemma 41 [64]. We construct a tree \mathcal{T} such that each node represents a run ρ of some finite
 1531 prefix w' of w . The root's label is the empty run. For every node labelled with ρ , the children's
 1532 extend ρ by a single transition. The tree \mathcal{T} is finitely branching by construction of A_p for
 1533 every role p . With König's Lemma, we obtain an infinite path in \mathcal{T} and, thus, an infinite run ρ
 1534 in $\{\{H(\mathbf{G})\downarrow_p\}_{p \in \mathcal{P}}\}$ with $w \preceq_{\sim} \text{trace}(\rho)$. From this, it follows that $w \in \mathcal{L}(\{\{H(\mathbf{G})\downarrow_p\}_{p \in \mathcal{P}}\})$.
 1535 ◀

1536 C.5 Formalisation for Lemma 4.15: 1537 Implementability entails Globally Cooperative

1538 ► **Definition C.3** (Matching Sends and Receptions (Def. 2 [77])). *In a word $w = e_1 \dots \in \Sigma^\infty$,
 1539 a send event $e_i = p \triangleright q!m$ is matched by a receive event $e_j = q \triangleleft p?m$, denoted by $e_i \vdash e_j$, if
 1540 $i < j$ and $\mathcal{V}((e_1 \dots e_i)\downarrow_{p \triangleright q! _}) = \mathcal{V}((e_1 \dots e_j)\downarrow_{q \triangleleft p? _})$. A send event e_i is unmatched if there
 1541 is no such receive event e_j .*

1542 **Proof.** We prove our claim by contraposition: assume there is a loop v_1, \dots, v_n such that the
 1543 communication graph of $\mu(v_1) \dots \mu(v_n)$ is not weakly connected. By construction of $H(\mathbf{G})$,
 1544 we know that every vertex is reachable so there is a path $u_1 \dots u_m v_1 \dots v_n$ in $H(\mathbf{G})$ for
 1545 some m and vertices u_1 to u_n such that $u_1 = v^I$. By the 0-Reachable-Assumption (p.15), this
 1546 path can be completed to end in a terminal vertex to obtain $u_1 \dots u_m v_1 \dots v_n u_{m+1} \dots u_{m+k}$
 1547 for some k and vertices u_{m+1} to u_{m+k} such that $u_{k+m} \in V^T$. By the syntax of global types
 1548 and the construction of $H(\mathbf{G})$, there is a role p that is the (only) sender in v_1 and u_{m+1} .

1549 Without loss of generality, let \mathcal{S}_1 and \mathcal{S}_2 be the two sets of (active) roles whose communi-
 1550 cation graphs of $v_1 \dots v_n$ are weakly connected and their union consists of all active roles.
 1551 Similar reasoning applies if there are more than two sets.

1552 We want to consider the specific linearisations from the language of the BMSC of each
 1553 subpath. Intuitively, these simply follow the order prescribed by the global type and do
 1554 not exploit the partial order of BMSC or the closure of the semantics for global types. For
 1555 this, we say that w_1 is the *canonical word for path* u_1, \dots, u_m if $w_1 \in \{w'_1 \dots w'_m \mid w'_i \in$
 1556 $\mathcal{L}(\mu(u_i)) \text{ for } 1 \leq i \leq m\}$. Analogously, let w_2 be the canonical word for $v_1 \dots v_n$ and w_3 be
 1557 the canonical word for $u_{m+1} \dots u_{m+k}$. Without loss of generality, \mathcal{S}_1 contains the sender of
 1558 the first element in w_2 and w_3 — basically the role which decides when to exit the loop for
 1559 the considered loop branch. Let $\{\{H(\mathbf{G})\downarrow_p\}_{p \in \mathcal{P}}\}$ be the erasure candidate implementation.

1560 By its definition and the correctness of $H(\mathbf{G})$, it holds that: $\mathcal{L}(\mathbf{G}) = \mathcal{L}(H(\mathbf{G}))$. With the
 1561 equivalence of the canonical candidate implementation (Corollary 4.9), the reasoning for
 1562 Lemma 3.2 by Lohrey [63], and the fact that it generalises to infinite executions Lemma 4.11,
 1563 the erasure candidate implementation admits at least the language specified by $H(\mathbf{G})$:

$$1564 \quad \mathcal{L}(H(\mathbf{G})) \subseteq \mathcal{L}(\{\!\!\{H(\mathbf{G})\!\!\}_{\downarrow_p}\!\!\}_{p \in \mathcal{P}}).$$

1565 Thus, it holds that $\mathcal{L}(\mathbf{G}) = \mathcal{L}(\{\!\!\{H(\mathbf{G})\!\!\}_{\downarrow_p}\!\!\}_{p \in \mathcal{P}})$ if \mathbf{G} is implementable. Therefore, we know
 1566 that $w_1 . w_2 . w_3 \in \mathcal{L}(\{\!\!\{H(\mathbf{G})\!\!\}_{\downarrow_p}\!\!\}_{p \in \mathcal{P}})$.

1567 From the construction of $H(\mathbf{G})$ and the construction of w_i for $i \in \{1, 2, 3\}$, it also holds
 1568 that $w_1 . (w_2)^h . w_3 \in \mathcal{L}(H(\mathbf{G})) \subseteq \mathcal{L}(\{\!\!\{H(\mathbf{G})\!\!\}_{\downarrow_p}\!\!\}_{p \in \mathcal{P}})$ for any $h > 0$.

1569 By construction of \mathcal{S}_1 and \mathcal{S}_2 , no two roles from both sets communicate with each other
 1570 in w_2 : there are no $r \in \mathcal{S}_1$ and $s \in \mathcal{S}_2$ such that $r \triangleright s!m$ is in w_2 or $s \triangleright r!m$ is in w_2 (and
 1571 consequently $r \triangleleft s?m$ is in w_2 or $s \triangleleft r?m$ is in w_2) for any m .

1572 From the previous two observations, it follows that

$$1573 \quad w_1 . w_2 . (w_2 \downarrow_{\Sigma_{\mathcal{S}_1}})^h . w_3 \in \mathcal{L}(\{\!\!\{H(\mathbf{G})\!\!\}_{\downarrow_p}\!\!\}_{p \in \mathcal{P}})$$

1574 for any h where $\Sigma_{\mathcal{S}_1} = \bigcup_{r \in \mathcal{S}_1} \Sigma_r$. Intuitively, this means that the set of roles with the role
 1575 to decide when to exit the loop can continue longer in the loop than the roles in \mathcal{S}_2 .

1576 With $\mathcal{L}(\mathbf{G}) = \mathcal{L}(H(\mathbf{G}))$, it suffices to show the following to find a contradiction:
 1577 $w_1 . w_2 . (w_2 \downarrow_{\Sigma_{\mathcal{S}_1}})^h . w_3 \notin \mathcal{L}(H(\mathbf{G}))$.

1578 Towards a contradiction, we assume the membership holds. By determinacy of $H(\mathbf{G})$, we
 1579 need to find a path $v'_1 \dots v'_{m'}$, that starts at the beginning of the loop, i.e., $v'_1 = v_1$, with
 1580 canonical word w_4 such that $w_2 \downarrow_{\Sigma_{\mathcal{S}_1}} . w_3 \sim w_4$.

1581 We show such a path cannot exist and that we would need to diverge during the loop.

1582 For readability, we denote $w_2 . w_3$ with x and $w_2 \downarrow_{\Sigma_{\mathcal{S}_1}} . w_3$ with x' . We know that x' is
 1583 a subsequence of x , i.e., $x' = x'_1 \dots x'_l$ and $x = x_1 \dots x_l$. Let $x_1 \dots x_j = x'_1 \dots x'_j$ denote
 1584 the maximal prefix on which both agree. Since \mathcal{S}_2 is not empty, we know that j can be
 1585 at most $|w_2 \downarrow_{\Sigma_{\mathcal{S}_1}}|$. (Intuitively, j cannot be so big that it reaches w_3 because there will be
 1586 mismatches due to $w_2 \downarrow_{\Sigma_{\mathcal{S}_2}}$ before.) We also claim that the next event x_{j+1} cannot be a
 1587 receive event. If it was, there was a matching send event in $x_1 \dots x_j$ (which is equal to
 1588 $x'_1 \dots x'_j$ by construction). Such a matching send event exists by construction of x from a
 1589 path in $H(\mathbf{G})$. By definition of \downarrow , the matching receive event must be x'_{j+1} which would
 1590 contradict the maximality of j . Thus, x_{j+1} must be a send event.

1591 By determinacy of $H(\mathbf{G})$ and $j \leq |w_2 \downarrow_{\Sigma_{\mathcal{S}_1}}|$, we know that $x_1 \dots x_j = x'_1 \dots x'_j$ share a
 1592 path $v_1 \dots v_n$ which is a part of the loop, i.e., $x_1 \dots x_j \in \mathcal{L}(\mu(v_1) \dots \mu(v_n))$ with $n < n'$.
 1593 For $M(p \rightarrow q : m)$ — the BMSC with solely this interaction from Definition 4.4, we say that
 1594 p is its *sender*. The syntax of global types prescribes that choice is deterministic and the
 1595 sender in a choice is unique. This is preserved for $H(\mathbf{G})$: for every vertex, all its successors
 1596 have the same sender. Therefore, the path for x' can only diverge, but also needs to diverge,
 1597 from the loop $v_1 \dots v_n$ after the common prefix $v_1 \dots v_n$ with a different send event but with
 1598 the same sender. Let v_l be next vertex after $v_1 \dots v_n$ on the loop v_1, \dots, v_n for which $\mu(v_l)$
 1599 is not M_ε — the BMSC with an empty set of event nodes from Definition 4.4. Note that
 1600 x_{j+1} belongs to v_l : $x_{j+1} \in \text{pref}(\mathcal{L}(\mu(v_l)))$.

1601 We do another case analysis whether x_{j+1} belongs to \mathcal{S}_1 or not, i.e., if $x_{j+1} \in \Sigma_{\mathcal{S}_1}$.

1602 If $x_{j+1} \notin \Sigma_{\mathcal{S}_1}$, there cannot be a path that continues for x'_{j+1} as the sender for $\mu(v_l)$ is
 1603 not in \mathcal{S}_1 . If $x_{j+1} \in \Sigma_{\mathcal{S}_1}$, the choice of j was not maximal which yields a contradiction. ◀

1604 C.6 Further Explanation for Example 4.16

1605 Here, we show that any trace of the CSM is specified by the HMSC. Let us consider a finite
 1606 execution of the CSM for which we want to find a path in the HMSC. Let us assume there

1607 are i interactions between p and q and j interactions between r and s . In our asynchronous
 1608 setting, these interactions are split and can be interleaved. From the CSM, it is easy to see
 1609 that i is at least 2 and j is at least 1. The simplest path goes through the first loop once and
 1610 accounts for $i - 1$ iterations in the second loop and $j - 1$ iterations in the third one. A more
 1611 involved path could account for $\min(i, j) - 1$ iterations of the first loop, as many as possible,
 1612 and $i - \min(i, j) + 1$ iterations of the second loop as well as $j - \min(i, j) + 1$ iterations of
 1613 the third loop. The key that both paths are valid possibilities is that the interactions of p
 1614 and q in the first and second loop are indistinguishable, i.e., the executions can be reordered
 1615 with \sim such that both is possible. The syntactic restriction on choice does prevent this
 1616 for global types (and this protocol cannot be represented with a global type). Intuitively,
 1617 one cannot make up for a different number of loop iterations, that are the consequence of
 1618 missing synchronisation, in global types because the “loop exit”-message will be distinct
 1619 (compared to staying in the loop) and anything specified afterwards cannot be reordered
 1620 by \sim in front of it. It is straightforward to adapt the protocol so final states do not have
 1621 outgoing transitions. We add another vertex with a BMSC at the bottom, which has the
 1622 same structure as the top one but with another message l instead of m . We add an edge
 1623 from the previous terminal vertex to the new vertex and make the new one the only terminal
 1624 vertex. With this, p and r can eventually decide not to send m anymore and indicate their
 1625 choice with the distinct message l to the other two roles.

1626 **D** Proof for Lemma 5.4: 1627 Correctness of Algorithm 1 to check \mathcal{I} -closedness of Global Types

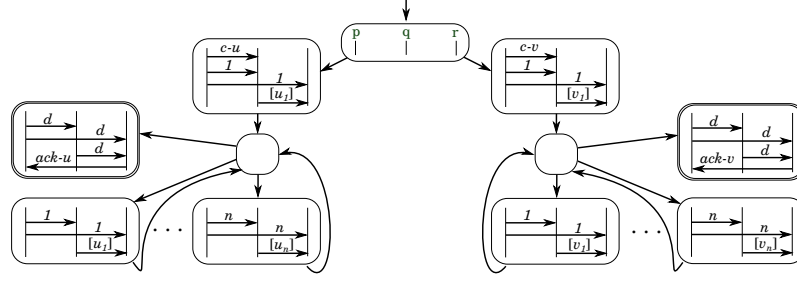
1628 It is obvious that the language is preserved by the changes to the state machine. (We basically
 1629 turned an unambiguous state machine into a deterministic one.)

1630 For soundness, we assume that Algorithm 1 returns *true* and let w be a word in
 1631 $\mathcal{C}^{\equiv_{\mathcal{I}}}(\mathcal{L}(\text{GAut}(\mathbf{G})))$. By definition, there is a run with trace w' in $\text{GAut}(\mathbf{G})$ such that $w' \equiv_{\mathcal{I}} w$.
 1632 The conditions in Algorithm 1 ensure that $w = w'$ because no two adjacent elements in w'
 1633 can be reordered with $\equiv_{\mathcal{I}}$. Therefore, $w \in \mathcal{L}(\text{GAut}(\mathbf{G}))$ which proves the claim.

1634 For completeness, we assume that the algorithm returns *false* and show that there is
 1635 $w \in \mathcal{C}^{\equiv_{\mathcal{I}}}(\mathcal{L}_{\text{fin}}(\mathbf{G}))$ such that $w \notin \mathcal{L}_{\text{fin}}(\text{GAut}(\mathbf{G}))$. Without loss of generality, let q_2 be the
 1636 state for which an incoming label x and outgoing label y can be reordered, i.e., $x \equiv_{\mathcal{I}} y$, and
 1637 let q_1 be the state from which the transition with label x originates: $q_1 \xrightarrow{x} q_2 \in \delta_{\text{GAut}(\mathbf{G})}$. We
 1638 consider a word w' which is the trace of a maximal run that passes q and the transitions
 1639 labelled with x and y . By construction, it holds that $w' \in \mathcal{L}_{\text{fin}}(\text{GAut}(\mathbf{G}))$. We swap x and y in
 1640 w' to obtain w . We denote x with $p \rightarrow q : m$ and y with $r \rightarrow s : m'$ such that $\{p, q\} \cap \{r, s\} \neq \emptyset$.
 1641 From the syntactic restrictions of global types, we know that any transition label from q_1 has
 1642 sender p while every transition label from q_2 has sender r . Because of this and determinacy
 1643 of the state machine, there is no run in $\text{GAut}(\mathbf{G})$ with trace w' . Thus, $w \notin \mathcal{L}_{\text{fin}}(\text{GAut}(\mathbf{G}))$
 1644 which concludes the proof. \blacktriangleleft

1645 **E** Proof for Theorem 6.7: Implementability with regard to 1646 Intra-role Reordering for Global Types from MSTs is Undecidable

1647 Let $\{(u_1, u_2, \dots, u_n), (v_1, v_2, \dots, v_n)\}$ be an instance of MPCP where 1 is the special index
 1648 with which each solution needs to start with. We construct a global type where, for a
 1649 word $w = a_1 a_2 \dots a_m \in \Delta^*$, a message labelled $[w]$ denotes a sequence of individual message
 1650 interactions with message a_1, a_2, \dots, a_m , each of size 1. We define a parametric global type



■ **Figure 8** HMSC encoding $H(\mathbf{G}_{\text{MPCP}})$ of the MPCP encoding (same as in main text)

1651 where $x \in \{u, v\}$:

$$1652 \quad G(x, X) := p \rightarrow q: c-x. p \rightarrow q: l. p \rightarrow r: l. q \rightarrow r: [x_1]. \mu t_1. + \begin{cases} p \rightarrow q: l. p \rightarrow r: l. q \rightarrow r: [x_1]. t_1 \\ \dots \\ p \rightarrow q: n. p \rightarrow r: n. q \rightarrow r: [x_n]. t_1 \\ p \rightarrow q: d. p \rightarrow r: d. q \rightarrow r: d. X \end{cases}$$

1653 where $c-x$ indicates *choosing* tile set x . Using this, we obtain our encoding:

$$1654 \quad \mathbf{G}_{\text{MPCP}} := + \begin{cases} G(u, r \rightarrow p: c-u. 0) \\ G(v, r \rightarrow p: c-v. 0) \end{cases}.$$

1656 Figure 8 illustrates its HMSC encoding $H(\mathbf{G}_{\text{MPCP}})$.

1657 It suffices to show the following equivalences:

$$1658 \quad \mathbf{G}_{\text{MPCP}} \text{ is } \approx\text{-implementable}$$

$$1659 \quad \Leftrightarrow_1 \mathcal{C}^{\approx}(\mathcal{L}(G(u, 0))) \downarrow_{\Sigma_r} \cap \mathcal{C}^{\approx}(\mathcal{L}(G(v, 0))) \downarrow_{\Sigma_r} = \emptyset$$

$$1660 \quad \Leftrightarrow_2 \text{MPCP instance has no solution}$$

1662 We prove \Rightarrow_1 by contraposition. Let $w \in \mathcal{C}^{\approx}(\mathcal{L}(G(u, 0))) \downarrow_{\Sigma_r} \cap \mathcal{C}^{\approx}(\mathcal{L}(G(v, 0))) \downarrow_{\Sigma_r}$. For
 1663 $x \in \{u, v\}$, let $w_x \in \mathcal{C}^{\approx}(\mathcal{L}(G(x, 0)))$ such that $w_x \downarrow_{\Sigma_r} = w$. By construction of \mathbf{G}_{MPCP} , we
 1664 know that $w_x . r \triangleright p!ack-x . p \triangleleft r?ack-x \in \mathcal{C}^{\approx}(\mathcal{L}(\mathbf{G}_{\text{MPCP}}))$.

1665 Suppose that CSM $\{\{A_p\}_{p \in \mathcal{P}}\} \approx$ -implements \mathbf{G}_{MPCP} . Then, it holds that
 1666 $w_x . r \triangleright p!ack-x . p \triangleleft r?ack-x \in \mathcal{C}^{\approx}(\mathcal{L}(\{\{A_p\}_{p \in \mathcal{P}}\}))$

1667 by (ii) from Definition 6.3. We also know that $w_x . r \triangleright p!ack-y . p \triangleleft r?ack-y \notin \mathcal{C}^{\approx}(\mathcal{L}(\mathbf{G}_{\text{MPCP}}))$
 1668 for $x \neq y$ where $x, y \in \{u, v\}$. By the choice of w_u and w_v , it holds that $w_u \downarrow_{\Sigma_r} = w = w_v \downarrow_{\Sigma_r}$.
 1669 Therefore, r needs to be in the same state of A_r after processing $w_u \downarrow_{\Sigma_r}$ or $w_v \downarrow_{\Sigma_r}$ and it can
 1670 either send both $ack-u$ and $ack-v$, only one of them or none of them to p . Thus, either one
 1671 of the following is true:

- 1672 a) (sending both) $w_x . r \triangleright p!ack-y \in \text{pref}(\mathcal{C}^{\approx}(\mathcal{L}(\{\{A_p\}_{p \in \mathcal{P}}\})))$ for $x \neq y$ where $x, y \in \{u, v\}$, or
- 1673 b) (sending u without loss of generality) $w_v . r \triangleright p!ack-u \notin \text{pref}(\mathcal{C}^{\approx}(\mathcal{L}(\{\{A_p\}_{p \in \mathcal{P}}\})))$, or
- 1674 c) (sending none) $w_x . r \triangleright p!ack-x \notin \text{pref}(\mathcal{C}^{\approx}(\mathcal{L}(\{\{A_p\}_{p \in \mathcal{P}}\})))$ for $x \in \{u, v\}$.

1675 All cases lead to deadlocks in $\{\{A_p\}_{p \in \mathcal{P}}\}$. For a) and for b) if $c-v$ was chosen in the beginning,
 1676 p cannot receive the sent message as it disagrees with its choice from the beginning $c-x$. In
 1677 all other cases, p waits for a message while no message will ever be sent. Having deadlocks
 1678 contradicts the assumption that $\{\{A_p\}_{p \in \mathcal{P}}\} \approx$ -implements \mathbf{G} (and there cannot be any CSM
 1679 that \approx -implements \mathbf{G}).

1680 We prove \Leftarrow_1 next. The language $\mathcal{C}^{\approx}(\mathcal{L}(\mathbf{G}_{\text{MPCP}}))$ is obviously non-empty. Therefore, let
 1681 $w' \in \mathcal{C}^{\approx}(\mathcal{L}(\mathbf{G}_{\text{MPCP}}))$. We split w' to obtain:

$$1682 \quad w' = w . r \triangleright p!ack-x . p \triangleleft r?ack-x \text{ for some } w \text{ and } x \in \{u, v\}.$$

1683 By construction of \mathbf{G}_{MPCP} , we know that

$$1684 \quad w \in \mathcal{C}^{\approx}(\mathcal{L}(G(u, 0))) \cup \mathcal{C}^{\approx}(\mathcal{L}(G(v, 0))).$$

1685 By assumption, it follows that exactly one of the following holds:

$$1686 \quad w \downarrow_{\Sigma_r} \in \mathcal{C}^{\approx}(\mathcal{L}(G(u, 0))) \downarrow_{\Sigma_r} \quad \text{or} \quad w \downarrow_{\Sigma_r} \in \mathcal{C}^{\approx}(\mathcal{L}(G(v, 0))) \downarrow_{\Sigma_r}.$$

1687 We give a \approx -implementation for \mathbf{G}_{MPCP} . It is straightforward to construct FSMs for both p
1688 and q . They are involved in the initial decision and \approx does not affect their projected languages.
1689 Thus, the projection by erasure can be applied to obtain FSMs A_p and A_q . We construct
1690 an FSM A_r for r with control state $i \in \{1, \dots, n\}$, $j \in \{1, \dots, \max(|u_i| \mid i \in \{1, \dots, n\})\}$,
1691 $d \in \{0, 1, 2\}$, and $x \in \{u, v\}$, where $|w|$ denotes the length of a word. The FSM is constructed
1692 in a way such that

$$1693 \quad w \downarrow_{\Sigma_r} \in \mathcal{C}^{\approx}(\mathcal{L}(G(u, 0))) \downarrow_{\Sigma_r} \quad \text{if and only if} \quad d \text{ is } 2 \text{ and } x \text{ is } u \quad \text{as well as}$$

$$1694 \quad w \downarrow_{\Sigma_r} \in \mathcal{C}^{\approx}(\mathcal{L}(G(v, 0))) \downarrow_{\Sigma_r} \quad \text{if and only if} \quad d \text{ is } 2 \text{ and } x \text{ is } v.$$

1696 We first explain that this characterisation suffices to show that $\{\{A_p\}_{p \in \mathcal{P}}\} \approx$ -implements \mathbf{G} .
1697 The control state d counts the number of received d -messages. Thus, there will be no more
1698 messages to r in any channel once d is 2 by construction of \mathbf{G}_{MPCP} . Once in a state
1699 for which d is 2, r sends message $ack-u$ to p if x is u and message $ack-v$ if x is v . With
1700 the characterisation, this message $ack-x$ matches the message $c-x$ sent from p to q in the
1701 beginning and, thus, p will be able to receive it and conclude the execution.

1702 Now, we will explain how to construct the FSM A_r . Intuitively, r keeps a tile number,
1703 which it tries to match against, and stores this in i . It is initially set to 0 to indicate no
1704 tile has been chosen yet. The index j denotes the position of the letter it needs to match in
1705 tile u_i next and, thus, is initialised to 1. The variable d indicates the number of d -messages
1706 received so far, so initially d is 0. With this, r knows when it needs to send $ack-x$. The FSM
1707 for r tries to match the received messages against the tiles of u , so x is initialised to u . If
1708 this matching fails at some point, x is set to v as it learned that v was chosen initially by p .
1709 In any of the following cases: if a received message is a d -message, d is solely increased by 1:

- 1710 ■ If x is u and i is 0, r receives a message z from p and sets i to z (technically the integer
1711 represented by z).
- 1712 ■ If x is u and i is not 0, r receives a message z from q .
 - 1713 ■ If z is the same as $u_i[j]$, we increment j by 1 and
 - 1714 check if $j > |u_i|$ and, if so, set i to 0 and j to 1
 - 1715 ■ If not, we set x to v
- 1716 ■ Once x is v , r can simply receive all remaining messages in any order.

1717 The described FSM can be used for r because it reliably checks whether a presented sequence
1718 of indices and words belongs to tile set u or v . It can do so because $\mathcal{C}^{\approx}(\mathcal{L}(G(u, 0))) \downarrow_{\Sigma_r} \cap$
1719 $\mathcal{C}^{\approx}(\mathcal{L}(G(v, 0))) \downarrow_{\Sigma_r} = \emptyset$ by assumption.

1720 We prove \Rightarrow_2 by contraposition. Suppose the MPCP instance has a solution. Let $i_1, \dots,$
1721 i_k be a non-empty sequence of indices such that $u_{i_1}u_{i_2} \cdots u_{i_k} = v_{i_1}v_{i_2} \cdots v_{i_k}$ and $i_1 = 1$. It is
1722 easy to see that

$$1723 \quad w_x := r \triangleleft p?i_1 r \triangleleft q?[x_{i_1}]. \cdots . r \triangleleft p?i_k . r \triangleleft q?[x_{i_k}]. r \triangleleft p?d . r \triangleleft q?d \in \mathcal{L}(G(x, 0)) \downarrow_{\Sigma_r} \text{ for } x \in \{u, v\}.$$

1724 By definition of \approx , we can re-arrange the previous sequences such that

$$1725 \quad r \triangleleft p?i_1 . \cdots . r \triangleleft p?i_k . r \triangleleft q?[x_{i_1}]. \cdots . r \triangleleft q?[x_{i_k}]. r \triangleleft p?d . r \triangleleft q?d \in \mathcal{C}^{\approx}(\mathcal{L}(G(x, 0))) \downarrow_{\Sigma_r} \text{ for } x \in \{u, v\}.$$

1726 Because i_1, \dots, i_k is a solution to the instance of MPCP, it holds that

$$1727 \quad r \triangleleft q?[u_{i_1}]. \cdots . r \triangleleft q?[u_{i_k}] = r \triangleleft q?[v_{i_1}]. \cdots . r \triangleleft q?[v_{i_k}]$$

1728 and, thus,

$$1729 \quad r \triangleleft p?i_1 . \cdots . r \triangleleft p?i_k . r \triangleleft q?[u_{i_1}]. \cdots . r \triangleleft q?[u_{i_k}]. r \triangleleft p?d . r \triangleleft q?d \text{ is in } \mathcal{C}^{\approx}(\mathcal{L}(G(v, 0))) \downarrow_{\Sigma_r}.$$

1730 This shows that $\mathcal{C}^{\approx}(\mathcal{L}(G(u, 0))) \downarrow_{\Sigma_r} \cap \mathcal{C}^{\approx}(\mathcal{L}(G(v, 0))) \downarrow_{\Sigma_r} \neq \emptyset$.

1731 Lastly, we prove \Leftarrow_2 . We know that the MPCP instance has no solution. Thus, there can-
 1732 not be a non-empty sequence of indices i_1, i_2, \dots, i_k such that $u_{i_1}u_{i_2} \cdots u_{i_k} = v_{i_1}v_{i_2} \cdots v_{i_k}$ and
 1733 $i_1 = 1$. For any possible word $w_u \in \mathcal{C}^\approx(\mathcal{L}(G(u, 0))) \downarrow_{\Sigma_x}$ and word $w_v \in \mathcal{C}^\approx(\mathcal{L}(G(v, 0))) \downarrow_{\Sigma_x}$.

1734 We consider the sequence of receive events $w_x \downarrow_{x \leftarrow p?}_-$ with sender p and the sequence
 1735 of messages $w_x \downarrow_{x \leftarrow q?}_-$ from q for $x \in \{u, v\}$. The intra-role indistinguishability relation \approx
 1736 allows to reorder events of both but for a non-empty intersection of both sets, we would still
 1737 need to find a word w_u and w_v such that

$$1738 \quad w_u \downarrow_{x \leftarrow p?}_- = w_v \downarrow_{x \leftarrow p?}_- \quad \text{and} \quad w_u \downarrow_{x \leftarrow q?}_- = w_v \downarrow_{x \leftarrow q?}_-.$$

1739 However, $G(x, 0)$ for $x \in \{u, v\}$ is constructed in a way that this is only possible if the MPCP
 1740 instance has a solution. Therefore, the intersection is empty which proves our claim. \blacktriangleleft