



Realisability of Global Models of Interaction

Maurice H. ter Beek¹(✉) , Rolf Hennicker², and José Proença³(✉) 

¹ ISTI-CNR, Pisa, Italy

`maurice.terbeek@isti.cnr.it`

² Ludwig-Maximilians-Universität München, Munich, Germany

`hennicker@ifi.lmu.de`

³ CISTER & University of Porto, Porto, Portugal

`jose.proenca@fc.up.pt`

Abstract. We consider *global models* of communicating agents specified as transition systems labelled by *interactions* in which multiple senders and receivers can participate. A *realisation* of such a model is a set of local transition systems—one per agent—which are executed concurrently using synchronous communication. Our core challenge is how to check whether a global model is realisable and, if it is, how to synthesise a realisation. We identify and compare two variants to realise global interaction models, both relying on bisimulation equivalence. Then we investigate, for both variants, *realisability conditions* to be checked on global models. We propose a synthesis method for the construction of realisations by grouping locally indistinguishable states. The paper is accompanied by a tool that implements realisability checks and synthesises realisations.

1 Introduction

We deal with the development of systems of collaborating computing entities which interact by message exchange, like communicating component systems, multi-agent systems (MAS), collective adaptive systems (CAS), groupware systems, multi-party sessions, etc. Such systems are often presented by a set of components whose local behaviour is formally described by labelled transition systems (LTS) or process expressions. Their interaction behaviour is then captured by (synchronous or asynchronous) parallel composition of the local models.

Before designing such local models it is, however, safer to first model the interaction behaviour of the components from a *global* perspective. This led to the investigation of various forms of global models, like global (session) types [8, 13, 23, 24], global choreographies [35] and global languages [2]; also message sequence charts [20] and UML interaction diagrams [15, 30] serve this purpose.

An important question is, of course, whether a global model \mathcal{M} is indeed realisable by a system $\mathcal{S} = (\mathcal{M}_i)_{i \in \mathcal{I}}$ of local component models \mathcal{M}_i (where \mathcal{I} ranges over a set of component names). Possible solutions are investigated for global languages in [2] and, for global session types, in various papers (cf., e.g.,

[8, 13, 23, 24]) by imposing syntactic restrictions on global types. These approaches use projections to generate local models from global ones.

A different idea is to provide, instead of a global model, a requirements specification Sp describing properties of the desired global interaction behaviour by means of some logical formalism like in [9, 21, 22]. Then local models \mathcal{M}_i are constructed from scratch and their (synchronous) composition $\otimes(\mathcal{M}_i)_{i \in \mathcal{I}}$ must be proven to satisfy the requirements of Sp .

From Requirements to Realisations. We combine the advantages of logical specifications and global models for interaction-based systems by using both in a stepwise manner. Our development method is summarised in Fig. 1.




Fig. 1. Workflow for the development of interaction-based systems

We start by providing a (*system*) signature $\Sigma = (I, M)$, which determines finite sets I of component names and M of message names. Σ induces the set $\Gamma(\Sigma)$ of (*global*) Σ -interactions of the form $\text{out} \rightarrow \text{in} : m$ where out and in are disjoint sets of component names (such that $\text{out} \cup \text{in} \neq \emptyset$) and $m \in M$. The multi-interaction $\text{out} \rightarrow \text{in} : m$ expresses that all components of out send message m to all components of in such that all send and receive events occur simultaneously. Since usually not all interactions in $\Gamma(\Sigma)$ are desired in an application context (one may wish, e.g., to consider only binary or multicast communication) we consider pairs (Σ, Γ) where $\Gamma \subseteq \Gamma(\Sigma)$ is a user-defined *interaction set* which restricts the set of all Σ -interactions to admissible ones. For logical specifications of global interaction behaviour, we propose an action-based logic following a dynamic-logic style which has been successfully applied for specifying ensembles (cf., e.g., [22]). The logic uses the usual diamond and box modalities of dynamic logic ($\langle \alpha \rangle \varphi$ and $[\alpha] \varphi$, resp.) which range over (structured) interactions α built over Γ by sequential composition, choice and iteration. A *global interaction-behaviour specification* is then a triple $Sp = (\Sigma, \Gamma, Ax)$, where Ax is a set of formulas, called *axioms*, expressing requirements for the global interaction behaviour (e.g., safety and liveness properties and/or desired and forbidden interaction scenarios).

Given a (global) requirements specification $Sp = (\Sigma, \Gamma, Ax)$, we construct a global model \mathcal{M} for the system's intended interaction behaviour. To formalise such models we use *global LTS* whose transitions are labelled by interactions according to Γ . If only binary interactions are admitted a global LTS is a chore-

ography automaton as in [1]. Of course, we must check that the constructed global LTS \mathcal{M} satisfies the requirements of the specification Sp , i.e. $\mathcal{M} \models Ax$.

The central part of our work concerns the realisation (decomposition) of a global LTS \mathcal{M} in terms of a (possibly distributed) system of interacting components whose individual behaviour is modelled by *local LTS*. First we must determine, for each component name $i \in I$, which local actions component i should provide. To do so, any interaction in which i participates must be mapped to an appropriate local action for component i . We study two variants. The first follows approaches to multi-party session types and choreography languages where the names of the communication partners are kept in local actions. For instance, a binary interaction $i \rightarrow j : \mathbf{m}$ leads to a local output action $ij! \mathbf{m}$ for i and a local input action $ij? \mathbf{m}$ for j . In approaches to component-based development, however, transitions describing local behaviour are often labelled just by message names accompanied by information whether it is an output or an input of a component. This makes components better reusable and supports interface-based design [3, 4, 17, 26, 28]. In this case, a binary interaction $i \rightarrow j : \mathbf{m}$ leads to a local output action $! \mathbf{m}$ for i and a local input action $? \mathbf{m}$ for j . In this paper, we generalise both localisation styles to deal with multi-interactions and call the former “rich local actions” and the latter “poor local actions”. From a software designer’s point of view the poor localisation style better supports the principle of loose coupling, whereas the rich style better avoids undesired synchronisations.

Once a localisation style $x \in \{\mathbf{r}, \mathbf{p}\}$ is chosen (\mathbf{r} for “rich” and \mathbf{p} for “poor”) one can proceed with the actual construction of a realisation of \mathcal{M} in terms of a system presented by a family $\mathcal{S} = (\mathcal{M}_i^x)_{i \in I}$ of local LTS. We say that \mathcal{M} is *realisable* (with localisation style x) if such a system exists such that \mathcal{M} is bisimilar (denoted by \sim) to the synchronous composition $\otimes_I^x (\mathcal{M}_i^x)_{i \in I}$ of all \mathcal{M}_i^x taking into account the interactions Γ and the localisation style x . Hence, our realisation notion is generic w.r.t. Γ and parametrised by the chosen localisation style. We show that realisability with poor local actions implies realisability with rich local actions (Theorem 1) but the converse does not hold (Example 3). Since our realisability notion is based on bisimilarity we can deal with non-deterministic behaviour, differently from language-based approaches like [2, 13].

Race Example. We illustrate our methodology outlined so far by developing a (small) system, called *Race*, which is meant to model the competition of two runner components **R1** and **R2** under the control of a third component **Ctrl**. To start, we provide a signature $\Sigma_{\text{Race}} = (I_{\text{Race}}, M_{\text{Race}})$ with component names $I_{\text{Race}} = \{\mathbf{R1}, \mathbf{R2}, \mathbf{Ctrl}\}$ and message names $M_{\text{Race}} = \{\mathbf{start}, \mathbf{finish}\}$. The idea is that the controller starts the two runners simultaneously, while each runner signals individually to the controller when it has finished its run. Therefore, we use the interaction set on the left of Fig. 2. We do not model the actual running of a runner component, which would be an internal action (cf. [7]).




Fig. 2. Interaction set Γ_{Race} (left) and global LTS $\mathcal{M}_{\text{Race}}$ (right); we write **Ctrl** for $\{\text{Ctrl}\}$ and similarly for **R1**, **R2**.

We require that no runner should finish before starting and that any started runner should be able to finish running. This will be expressed by dynamic logic formulas to be detailed in the requirements specification Sp_{Race} in Example 1.

Next we construct the global LTS $\mathcal{M}_{\text{Race}}$ shown on the right of Fig. 2, which models the required interaction behaviour of the system so that the requirements of the specification are satisfied. The system starts in the initial (global) state 0, where the controller **starts** both runners at once. Each runner separately sends a **finish** signal to the controller (in arbitrary order). After that a new run can **start**.

Finally, we want to realise the system by three local LTS such that their composition is bisimilar to the global LTS $\mathcal{M}_{\text{Race}}$. We distinguish the two variants.

Rich Local Actions. From Γ_{Race} we derive the following sets of rich local actions:

$$\begin{aligned} \mathcal{A}_{\text{Ctrl}}^r &= \{ \text{Ctrl} \{\text{R1}, \text{R2}\} !\text{start}, \text{R1 Ctrl} ?\text{finish}, \text{R2 Ctrl} ?\text{finish} \}, \\ \mathcal{A}_{\text{R1}}^r &= \{ \text{Ctrl} \{\text{R1}, \text{R2}\} ?\text{start}, \text{R1 Ctrl} !\text{finish} \}, \text{ and} \\ \mathcal{A}_{\text{R2}}^r &= \{ \text{Ctrl} \{\text{R1}, \text{R2}\} ?\text{start}, \text{R2 Ctrl} !\text{finish} \}. \end{aligned}$$

For each $i \in \{\text{Ctrl}, \text{R1}, \text{R2}\}$, we use the local LTS \mathcal{M}_i^r in the upper row of Table 1 to build the system $\mathcal{S}_{\text{Race}}^r = \{\mathcal{M}_{\text{Ctrl}}^r, \mathcal{M}_{\text{R1}}^r, \mathcal{M}_{\text{R2}}^r\}$ with rich local actions. One can prove that the “rich” composition (Definition 4) of the three LTS by synchronisation w.r.t. Γ_{Race} is bisimilar (even isomorphic) to $\mathcal{M}_{\text{Race}}$; i.e., we have found a realisation with rich local actions.

Table 1. Local LTS for each localisation style and for each component

Localisation	Local Ctrl	Local R1	Local R2
Rich			
Poor			

Poor Local Actions. In this case, we derive from Γ_{Race} the following sets of poor local actions, where information on communication partners is omitted:

$$A_{\text{Ctrl}}^{\text{P}} = \{\text{!start}, \text{?finish}\} \text{ and } A_{\text{R1}}^{\text{P}} = A_{\text{R2}}^{\text{P}} = \{\text{?start}, \text{!finish}\}.$$

For each $i \in \{\text{Ctrl}, \text{R1}, \text{R2}\}$, we use the local LTS \mathcal{M}_i^p in the lower row of Table 1 to build the system $\mathcal{S}_{\text{Race}}^{\text{P}} = \{\mathcal{M}_{\text{Ctrl}}^{\text{P}}, \mathcal{M}_{\text{R1}}^{\text{P}}, \mathcal{M}_{\text{R2}}^{\text{P}}\}$ with poor local actions. Also the “poor” composition (Definition 7) of the three LTS by synchronisation w.r.t. Γ_{Race} is bisimilar (even isomorphic) to $\mathcal{M}_{\text{Race}}$.

Checking Realisability, Local Quotients, and System Synthesis. So far, we considered the case in which the realisation of a global interaction model \mathcal{M} is “invented”. However, there might be no realisation of \mathcal{M} and it would be better to know this as soon as possible to align the global model. Next, we consider the following two important issues and proceed as shown in Fig. 3.

1. How to check whether a given global LTS \mathcal{M} is realisable (rich/poor case)?
2. If it is, how can we build/synthesise a concrete realisation \mathcal{S} (rich/poor case)?




Fig. 3. Approach to check realisability and system synthesis

To tackle the first question we propose, similarly to [14], to find a family $\equiv = (\equiv_i)_{i \in I}$ of equivalence relations on the global state space Q of \mathcal{M} such that, for each component name $i \in I$ and states $q, q' \in Q$, $q \equiv_i q'$ expresses that q and q' are not distinguishable from the viewpoint of i . This suggests that q and q' , though globally different, can be locally interpreted as the same states. In particular, it is required that any two states q and q' which are related by a global transition $q \xrightarrow{\text{out} \rightarrow \text{in} : m} \mathcal{M} q'$ should be indistinguishable for any $i \in I$ which does not participate in the interaction, i.e. $i \notin \text{out} \cup \text{in}$. On the basis of a given I -equivalence \equiv , we formulate realisability conditions $\text{RC}(\mathcal{M}, \equiv)^{\text{r}}$ and $\text{RC}(\mathcal{M}, \equiv)^{\text{p}}$ for both localisation styles. We show that in the rich and in the poor case our condition is sufficient for realisability (cf. Theorems 2 and 3).

In both cases, the principle idea how to synthesise a realisation is the same. Given a family $(\equiv_i)_{i \in I}$ of I -equivalences for which the realisability condition holds, we construct, for each $i \in I$, a local quotient $(\mathcal{M}/\equiv_i)^{\text{r/p}}$ by identifying global states (in \mathcal{M}) which are i -equivalent. Thus we get the desired system (which might still benefit from minimisations w.r.t. bisimilarity).

Note that the I -equivalences found for satisfying the realisability condition in the rich case may not be the same as in the poor case and thus also the local quotients may show different behaviour. Moreover, the technique of building local quotients differs from projections used in the field of multi-party session

types, since projections are partial operations depending on syntactic conditions (cf., e.g., [8]). A less syntactic and more expressive approach is proposed in [25].

As an example, recall the global LTS $\mathcal{M}_{\text{Race}}$ shown in Fig. 2 (right). The three local LTS with *rich* local actions shown in the upper row of Table 1 are, up to renaming of states, local quotients of $\mathcal{M}_{\text{Race}}$. To construct the local quotient for **R1**, global states 0 and 2 are identified, as well as states 1 and 3 (and symmetrically for the local quotient for **R2**). For **Ctrl**, no proper identification is applied (cf. Example 4 for details). Also the three local LTS with *poor* local actions in the lower row of Table 1 are, up to renaming of states, local quotients of $\mathcal{M}_{\text{Race}}$. In this case, however, to construct the local quotient for **Ctrl**, two global states of $\mathcal{M}_{\text{Race}}$ are identified, namely states 2 and 3 (cf. Example 6 for details).

Contributions and Related Work

1. We propose a rigorous discipline for developing interaction-based systems following a step-wise development method from dynamic-logic requirements specifications over global models of interaction down to systems of (possibly distributed) components. Thus our approach supplements approaches to realisations of global behaviour descriptions (in the form of global languages, e.g. [2], or global session types, e.g. [23]), by an abstract logical layer.
2. Our approach is driven by specified sets of multi-interactions supporting any kind of synchronous communication between multiple senders and multiple receivers. To the best of our knowledge, realisations of global models with arbitrary multi-interactions have not yet been studied in the literature.
3. Our correctness notion for realisation of global models by systems of communicating local components is based on bisimulation, thus letting us deal with non-determinism and going beyond language-based approaches like [2, 13]. Bisimulation also fits well with global requirements specifications since dynamic logic formulas are invariant under bisimulation and therefore hold in any realisation of a global model of a global specification.
4. For constructing realisations we consider two localisation styles (rich and poor local actions) and analyse their relationship. This is a novel result.
5. A global interaction model may, in general, not be realisable. We provide conditions for realisability with respect to both localisation styles. Our conditions are related to the work in [14] which, however, does not deal with multi-interactions and uses a stronger condition ensuring realisation up to isomorphism of LTS; cf. our discussion in Sect. 5.1.
6. For realisable global models, we construct realisations in terms of systems of local quotients. Similar quotient constructions have been used in the proofs of [14], but not for multi-interactions and for different localisation styles. The technique of building local quotients differs from projections used in the field of multi-party session types, since projections are partial operations depending on syntactic conditions (cf., e.g., [8]). In our approach, no restrictions on the form of global models are assumed. However, it must be said that the syntactic restrictions used for global types guarantee some kind of communication properties of a resulting system which we do not consider.
7. We developed a prototypical tool Ceta which checks realisability conditions and, if they are satisfied, generates local quotients and hence realisations.

Outline. After some formal preliminaries in Sect. 2, we show how to specify requirements for global models of interaction in Sect. 3 and how to realise the latter in Sect. 4. The conditions that guarantee realisability are studied in Sect. 5. In Sect. 6, we present a tool that implements our analyses. It is available at <https://lmf.di.uminho.pt/ceta> and all examples of the paper are predefined in the tool, like $\mathcal{M}_{\text{Race}}$, including a hyperlink to open the tool with the specific example. Section 7 wraps up the paper. A companion report [6] includes all proofs of our results, more details of the tool, and a few additional examples.

2 Formal Preliminaries

LTS and Bisimulation. Let A be a finite set of actions. A *labelled transition system* (LTS) over A is a tuple $\mathcal{L} = (Q, q_0, A, T)$ such that Q is a finite set of states, $q_0 \in Q$ is the initial state, and $T \subseteq Q \times A \times Q$ is a transition relation. Note that we consider finite-state LTS, which makes the realisability conditions presented later *decidable*. We write $q \xrightarrow{a}_{\mathcal{L}} q'$ to denote $(q, a, q') \in T$. A state $q \in Q$ is *reachable* if there exists a finite sequence of transitions from initial state q_0 to q .

Let $\mathcal{L}_i = (Q_i, q_{i,0}, A, T_i)$ be two LTS (for $i = 1, 2$) over the same action set A . A *bisimulation relation* between \mathcal{L}_1 and \mathcal{L}_2 is a relation $B \subseteq Q_1 \times Q_2$ such that for all $(q_1, q_2) \in B$ and for all $a \in A$ the following holds:

1. if $q_1 \xrightarrow{a}_{\mathcal{L}_1} q'_1$ then there exist $q'_2 \in Q_2$ and $q_2 \xrightarrow{a}_{\mathcal{L}_2} q'_2$ such that $(q'_1, q'_2) \in B$;
2. if $q_2 \xrightarrow{a}_{\mathcal{L}_2} q'_2$ then there exist $q'_1 \in Q_1$ and $q_1 \xrightarrow{a}_{\mathcal{L}_1} q'_1$ such that $(q'_1, q'_2) \in B$.

\mathcal{L}_1 and \mathcal{L}_2 are *bisimilar*, denoted by $\mathcal{L}_1 \sim \mathcal{L}_2$, if there exists a bisimulation relation B between \mathcal{L}_1 and \mathcal{L}_2 such that $(q_{1,0}, q_{2,0}) \in B$.

Dynamic Logic. We use (test-free) propositional dynamic logic (PDL) [19] to formulate behavioural properties. Let A be a finite set of (*atomic*) actions. Let the grammar $\alpha := a \mid \alpha; \alpha \mid \alpha + \alpha \mid \alpha^*$, with $a \in A$, sequential composition $;$, non-deterministic choice $+$, and iteration $*$, define the set $Act(A)$ of *structured actions* over A . If $A = \{a_1, \dots, a_n\}$, we write *some* for structured action $a_1 + \dots + a_n$. We may also refer to all actions of A but one, say a , and express this by $-a$.

The set $Frm(A)$ of *A-formulas* is defined by the grammar

$$\varphi := true \mid \neg\varphi \mid \varphi \vee \varphi \mid \langle \alpha \rangle \varphi \tag{formulas}$$

where $\alpha \in Act(A)$. Formula $\langle \alpha \rangle \varphi$ expresses that at the current state it is possible to execute α such that φ holds in the next state.

Abbreviations. We use the usual abbreviations like *false*, $\varphi \wedge \varphi'$, $\varphi \rightarrow \varphi'$, and the modal box operator $[\alpha] \varphi$ which stands for $\neg \langle \alpha \rangle \neg \varphi$ and expresses that whenever in the current state α is executed, then φ holds afterwards.

For the interpretation of formulas we use LTS. Let $\mathcal{L} = (Q, q_0, A, T)$ be an LTS over A . First we extend the transition relation of \mathcal{L} to structured actions:

$q \xrightarrow{\alpha_1; \alpha_2}_{\mathcal{L}} q'$ if there exists $\hat{q} \in Q$ such that $q \xrightarrow{\alpha_1}_{\mathcal{L}} \hat{q}$ and $\hat{q} \xrightarrow{\alpha_2}_{\mathcal{L}} q'$;
 $q \xrightarrow{\alpha_1 + \alpha_2}_{\mathcal{L}} q'$ if $q \xrightarrow{\alpha_1}_{\mathcal{L}} q'$ or $q \xrightarrow{\alpha_2}_{\mathcal{L}} q'$; and
 $q \xrightarrow{\alpha^*}_{\mathcal{L}} q'$ if $q = q'$ or there exists $\hat{q} \in Q$ such that $q \xrightarrow{\alpha}_{\mathcal{L}} \hat{q}$ and $\hat{q} \xrightarrow{\alpha^*}_{\mathcal{L}} q'$.

The *satisfaction* of a formula $\varphi \in \text{Frm}(A)$ by \mathcal{L} at a state $q \in Q$, denoted by $\mathcal{L}, q \models \varphi$, is inductively defined as follows:

$\mathcal{L}, q \models \text{true}$;
 $\mathcal{L}, q \models \neg\varphi$ if not $\mathcal{L}, q \models \varphi$;
 $\mathcal{L}, q \models \varphi_1 \vee \varphi_2$ if $\mathcal{L}, q \models \varphi_1$ or $\mathcal{L}, q \models \varphi_2$; and
 $\mathcal{L}, q \models \langle \alpha \rangle \varphi$ if there exists $q' \in Q$ such that $q \xrightarrow{\alpha}_{\mathcal{L}} q'$ and $\mathcal{L}, q' \models \varphi$.

\mathcal{L} *satisfies* a formula $\varphi \in \text{Frm}(A)$, denoted by $\mathcal{L} \models \varphi$, if $\mathcal{L}, q_0 \models \varphi$. Hence, for the satisfaction of a formula by an LTS the non-reachable states are irrelevant (deviating from the classical semantics of PDL [19]). We can express safety properties, like $[\text{some}^*] \varphi$, and some kinds of liveness properties like, e.g., $[\text{some}^*] \langle \text{some}^* \rangle a \varphi$.

Satisfaction of formulas in PDL is invariant under bisimulation [10]: Let $\mathcal{L}_1, \mathcal{L}_2$ be two LTS over A . If $\mathcal{L} \sim \mathcal{L}'$ then, for any $\varphi \in \text{Frm}(A)$, $\mathcal{L} \models \varphi$ iff $\mathcal{L}' \models \varphi$.

3 Specifying Requirements for Global Models of Interaction

We focus on the stepwise development of systems whose components interact by synchronous message exchange. We support “multi-interactions”, in which several senders and receivers may participate in a communication. Our starting point are *signatures* $\Sigma = (I, M)$, where I is a finite set of component names (also called participants) and M is a finite set of message names. Any signature Σ induces a set $\Gamma(\Sigma)$ of (global) Σ -*interactions* defined by

$$\Gamma(\Sigma) = \{\text{out} \rightarrow \text{in} : \mathbf{m} \mid \text{out}, \text{in} \subseteq I, \text{out} \cup \text{in} \neq \emptyset, \mathbf{m} \in M\}.$$

An interaction $\text{out} \rightarrow \text{in} : \mathbf{m}$ expresses that all components whose name occurs in **out** send a message named **m** to all components whose name occurs in **in**. Such interactions involving arbitrarily many senders and receivers are also called *multi-interactions*. They will be interpreted by **synchronous** (handshake) communication. As a shorthand notation we write i for $\{i\}$. Special cases are binary interactions between two components i, j , denoted by $i \rightarrow j : \mathbf{m}$, or multicast communication with one sender i and a group **in** of receivers, denoted by $i \rightarrow \text{in} : \mathbf{m}$.

Usually not all Σ -interactions are meaningful for a certain application. Therefore our approach will be driven by user-definable *interaction sets* $\Gamma \subseteq \Gamma(\Sigma)$.

General Assumption. In the sequel, we assume that (Σ, Γ) denotes a system signature $\Sigma = (I, M)$ together with an interaction set Γ . When we talk about a signature we always mean a system signature.

We propose to use interactions as atomic actions in dynamic logic formulas for specifying desired and forbidden interaction properties from a global perspective.

Definition 1 (global Sp). A global interaction behaviour specification is a triple $Sp = (\Sigma, \Gamma, Ax)$ where $Ax \subseteq \text{Frm}(\Gamma)$ is a set of Γ -formulas, called axioms.

Example 1. A requirements specification for the interaction behaviour of the Race system is given by $Sp_{\text{Race}} = (\Sigma_{\text{Race}}, \Gamma_{\text{Race}}, Ax_{\text{Race}})$ where Σ_{Race} and Γ_{Race} are defined in Sect. 1 and Ax_{Race} consists of the following two dynamic logic formulas expressing the two informal requirements described in Sect. 1.

1. “No runner should finish before it has been started by the controller”.

$$\left[\left(- (\text{Ctrl} \rightarrow \{\text{R1}, \text{R2}\} : \text{start}) \right)^* ; \left(\begin{array}{l} \text{R1} \rightarrow \text{Ctrl} : \text{finish} + \\ \text{R2} \rightarrow \text{Ctrl} : \text{finish} \end{array} \right) \right] \text{false}$$
2. “For any started runner it should be possible to finish its run”.

$$\left[\text{some}^*; \text{Ctrl} \rightarrow \{\text{R1}, \text{R2}\} : \text{start} \right] \left(\langle \text{some}^*; \text{R1} \rightarrow \text{Ctrl} : \text{finish} \rangle \text{true} \wedge \langle \text{some}^*; \text{R2} \rightarrow \text{Ctrl} : \text{finish} \rangle \text{true} \right) \triangleright$$

Given a specification Sp , the goal of our next step is to model the global interaction behaviour of the intended system in accordance with Sp . For this purpose we use LTS with interactions from Γ on the transitions.

Definition 2 (global LTS). A global LTS over (Σ, Γ) is defined as an LTS $\mathcal{M} = (Q, q_0, \Gamma, T)$ over Γ .

To check that a global LTS satisfies the axioms of a specification, we may use the mCRL2 toolset [12] and, as explained in [7], the translation of LTS to process expressions as well as the translation of our dynamic logic formulas to the syntax used by mCRL2. For instance, the global LTS $\mathcal{M}_{\text{Race}}$ provided for the race example in Sect. 1 satisfies the axioms of the specification Sp_{Race} above.

4 Realisations of Global Models of Interaction

A crucial step in our development method concerns the realisation of a global interaction model in terms of a system of (possibly distributed) components modelled by *local* LTS (cf. Fig. 1). In this section, we formally define what we mean by a realisation. For modelling local components we must first determine, for each component name $i \in I$, which are the local actions that component i is supposed to support. We study two variants obeying different localisation styles and leading to different instantiations of our realisability notion.

4.1 Realisations Using Rich Local Actions

It is common in approaches to global (session) types and choreography languages to preserve the names of communication partners when moving from global interactions to local actions. In [2], e.g., a binary interaction $i \rightarrow j : \mathbf{m}$ leads to a local output action $ij! \mathbf{m}$ for i and a local input action $ij? \mathbf{m}$ for j . We generalise this approach to multi-interactions and call the resulting local actions *rich*.

Definition 3 (rich local actions and local LTS). For each $i \in I$, the set of rich local i -actions derived from Γ is $A_i^r(\Gamma) = A_{i,out}^r(\Gamma) \cup A_{i,in}^r(\Gamma)$ where

$$\begin{aligned} A_{i,out}^r(\Gamma) &= \{\text{out in!m} \mid \exists (\text{out} \rightarrow \text{in} : \text{m}) \in \Gamma \text{ such that } i \in \text{out}\} \text{ and} \\ A_{i,in}^r(\Gamma) &= \{\text{out in?m} \mid \exists (\text{out} \rightarrow \text{in} : \text{m}) \in \Gamma \text{ such that } i \in \text{in}\}. \end{aligned}$$

A local LTS for i with rich local actions is an LTS $\mathcal{M}_i^r = (Q_i, q_{i,0}, A_i^r(\Gamma), T_i)$.

A system over (Σ, Γ) with rich local actions is a family $\mathcal{S}^r = (\mathcal{M}_i^r)_{i \in I}$ of local LTS \mathcal{M}_i^r over $A_i^r(\Gamma)$ for $i \in I$. The behaviour of such a system is modelled by the synchronous Γ -composition of its components \mathcal{M}_i^r ($i \in I$) where for all interactions $(\text{out} \rightarrow \text{in} : \text{m}) \in \Gamma$ a global transition exists (in a composed state) if for all $i \in \text{out}$ ($i \in \text{in}$, resp.) there is a transition in \mathcal{M}_i^r with the local action out in!m (out in?m , resp.) leaving the current local state of \mathcal{M}_i^r .

Definition 4 (synchronous Γ -composition with rich local actions). Let $(\mathcal{M}_i^r)_{i \in I}$ be a family of local LTS $\mathcal{M}_i^r = (Q_i, q_{i,0}, A_i^r(\Gamma), T_i)$ with rich local actions. The synchronous Γ -composition of $(\mathcal{M}_i^r)_{i \in I}$ with rich local actions is the global LTS, denoted by $\otimes_{\Gamma}^r (\mathcal{M}_i^r)_{i \in I}$, over (Σ, Γ) with initial state $(q_{i,0})_{i \in I}$ and with (product) states $(q_i)_{i \in I}$ (with $q_i \in Q_i$ for all $i \in I$) and transitions generated from the initial state by the following rule:


$$\frac{(\text{out} \rightarrow \text{in} : \text{m}) \in \Gamma \quad \forall i \in \text{out} : q_i \xrightarrow{\text{out in!m}}_{\mathcal{M}_i^r} q'_i \quad \forall i \in \text{in} : q_i \xrightarrow{\text{out in?m}}_{\mathcal{M}_i^r} q'_i}{(q_i)_{i \in I} \xrightarrow{\text{out} \rightarrow \text{in} : \text{m}}_{\otimes_{\Gamma}^r (\mathcal{M}_i^r)_{i \in I}} (q'_i)_{i \in I} \text{ where } q'_i = q_i \text{ for all } i \in I \setminus (\text{out} \cup \text{in})}$$

Definition 5 (realisability with rich local actions). Let \mathcal{M} be a global LTS over (Σ, Γ) . A system $\mathcal{S} = (\mathcal{M}_i^r)_{i \in I}$ over (Σ, Γ) with rich local actions is a (rich) realisation of \mathcal{M} , if $\mathcal{M} \sim \otimes_{\Gamma}^r (\mathcal{M}_i^r)_{i \in I}$ are bisimilar. \mathcal{M} is realisable with rich local actions if such a realisation exists.

Our realisability notion relies on bisimulation. Thus we are able to deal with non-determinism. In particular, according to the invariance of dynamic logic under bisimulation (cf. Sect. 2), we know that global models and their realisations satisfy the same formulas. Hence, once a global model of a global specification is provided, any realisation will be correct with respect to the global specification.

Example 2. Consider a non-deterministic example with two participants, a **Person** and a **Coin**, and **tossing** the **Coin** by the **Person** is modelled as a non-deterministic action that leads to either **head** or **tail** (cf. [29]). Formally, $\Sigma_{\text{Toss}} = (\{\text{Coin, Person}\}, \{\text{toss, head, tail}\})$ and $\Gamma_{\text{Toss}} = \{\text{Person} \rightarrow \text{Coin} : \text{toss}, \text{Coin} \rightarrow \text{Person} : \text{head}, \text{Coin} \rightarrow \text{Person} : \text{tail}\}$. The global LTS $\mathcal{M}_{\text{Toss}}^{\square}$ and one of its realisations by the two LTS $\mathcal{M}_{\text{Person}}^r$ and $\mathcal{M}_{\text{Coin}}^r$ (with rich local actions) are shown in Table 2. Although $\mathcal{M}_{\text{Toss}} \sim \otimes_{\Gamma_{\text{Toss}}}^r \{\mathcal{M}_{\text{Person}}^r, \mathcal{M}_{\text{Coin}}^r\}$ there would be no bisimulation when considering a deterministic version for both **Person** and **Coin**. \triangleright

Table 2. Non-deterministic toss of a **Coin** by a **Person**



4.2 Realisations Using Poor Local Actions

We now consider a variant where we omit the communication partners when we move from a global interaction $(\text{out} \rightarrow \text{in} : m) \in \Gamma$ to local actions. In this case only the message name m is kept together with output information $!m$ for $i \in \text{out}$ and input information $?m$ for $i \in \text{in}$. This complies with the idea of component automata used in teams [4, 5, 7] and many other approaches to component-based design (e.g., I/O automata [28] and interface automata [17]). We call the resulting local actions “poor” since they do not specify communication partners.

Definition 6 (poor local actions and local LTS). For each $i \in I$, the set of poor local i -actions derived from Γ is given by $\Lambda_i^p(\Gamma) = \Lambda_{i,\text{out}}^p(\Gamma) \cup \Lambda_{i,\text{in}}^p(\Gamma)$ where

$$\Lambda_{i,\text{out}}^p(\Gamma) = \{!m \mid \exists (\text{out} \rightarrow \text{in} : m) \in \Gamma \text{ such that } i \in \text{out}\} \text{ and}$$

$$\Lambda_{i,\text{in}}^p(\Gamma) = \{?m \mid \exists (\text{out} \rightarrow \text{in} : m) \in \Gamma \text{ such that } i \in \text{in}\}.$$

A local LTS for i with poor local actions is an LTS $\mathcal{M}_i^p = (Q_i, q_{i,0}, \Lambda_i^p(\Gamma), T_i)$ over $\Lambda_i^p(\Gamma)$.

The notion of a **system with poor local actions** is defined completely analogously to the rich case in Sect. 4.1. Γ -composition with poor local actions needs, however, special care since for matching local actions only the message name and input/output information is relevant.

Definition 7 (synchronous Γ -composition with poor local actions). Let $(\mathcal{M}_i^p)_{i \in I}$ be a family of local LTS $\mathcal{M}_i^p = (Q_i, q_{i,0}, \Lambda_i^p(\Gamma), T_i)$ with poor local actions. The synchronous Γ -composition of $(\mathcal{M}_i^p)_{i \in I}$ with poor local actions is the global LTS, denoted by $\otimes_\Gamma^p (\mathcal{M}_i^p)_{i \in I}$, over (Σ, Γ) with initial state $(q_{i,0})_{i \in I}$ and with (product) states $(q_i)_{i \in I}$ (such that $q_i \in Q_i$ for all $i \in I$) and transitions generated from the initial state by the following rule:

$$\frac{(\text{out} \rightarrow \text{in} : m) \in \Gamma \quad (\forall i \in \text{out} : q_i \xrightarrow{!m}_{\mathcal{M}_i^p} q'_i) \quad (\forall i \in \text{in} : q_i \xrightarrow{?m}_{\mathcal{M}_i^p} q'_i)}{(q_i)_{i \in I} \xrightarrow{\text{out} \rightarrow \text{in} : m}_{\otimes_\Gamma^p (\mathcal{M}_i^p)_{i \in I}} (q'_i)_{i \in I} \text{ where } q'_i = q_i \text{ for all } i \in I \setminus (\text{out} \cup \text{in})}$$

The notion of **realisability with poor local actions** is defined completely analogously to the rich case (cf. Definition 5) replacing “rich (r)” by “poor (p)”.

An obvious question is whether realisability with respect to the two different localisation styles can be formally compared. This is indeed the case.

Theorem 1 (poor realisation implies rich realisation). *Let \mathcal{M} be a global LTS over (Σ, Γ) which is realisable by a system $\mathcal{S}^p = (\mathcal{M}_i^p)_{i \in I}$ with poor local actions. Then there exists a system $\mathcal{S}^r = (\mathcal{M}_i^r)_{i \in I}$ with rich local actions which is a realisation of \mathcal{M} .*

The converse of Theorem 1 is not true, as demonstrated by the next example.

Example 3 We consider a variant of the global LTS $\mathcal{M}_{\text{Race}}^{\square}$ (Fig. 2) where the transitions $1 \xrightarrow{\text{R2} \rightarrow \text{Ctrl} : \text{finish}} 3 \xrightarrow{\text{R1} \rightarrow \text{Ctrl} : \text{finish}} 0$ are removed, enforcing **R1** to finish before **R2**. Let us call the resulting LTS $\mathcal{M}'_{\text{Race}}^{\square}$. Moreover, consider the variant of the local controller $\mathcal{M}'_{\text{Ctrl}}$ (upper row of Table 1, left) where the local transitions $1 \xrightarrow{\text{R2Ctrl?finish}} 3 \xrightarrow{\text{R1Ctrl?finish}} 0$ are removed and call it $\mathcal{M}^r_{\text{Ctrl}}$. Now let $\mathcal{S}^r = \{\mathcal{M}^r_{\text{Ctrl}}, \mathcal{M}^r_{\text{R1}}, \mathcal{M}^r_{\text{R2}}\}$ be the system with rich local actions (where $\mathcal{M}^r_{\text{R1}}$ and $\mathcal{M}^r_{\text{R2}}$ are shown in the upper row of Table 1, middle and right). It is easy to check that \mathcal{S}^r is a realisation of $\mathcal{M}'_{\text{Race}}$ with rich local actions, since $\mathcal{M}'_{\text{Race}}$ is even isomorphic to the (rich) Γ_{Race} -composition of $\{\mathcal{M}^r_{\text{Ctrl}}, \mathcal{M}^r_{\text{R1}}, \mathcal{M}^r_{\text{R2}}\}$.

The situation is different if we consider the poor case with controller $\mathcal{M}^p_{\text{Ctrl}}$ (lower row of Table 1, left) which accepts two times in a row a “finish” signal but, due to the poor local actions, cannot fix an acceptance order. The only candidate for a realisation with poor local actions is then the system $\mathcal{S}^p = \{\mathcal{M}^p_{\text{Ctrl}}, \mathcal{M}^p_{\text{R1}}, \mathcal{M}^p_{\text{R2}}\}$ consisting of the local LTS with poor local actions shown in the lower row of Table 1. Obviously, the Γ_{Race} -composition of these local LTS with poor actions does allow a sequence of transitions $1 \xrightarrow{\text{R2} \rightarrow \text{Ctrl} : \text{finish}} 3 \xrightarrow{\text{R1} \rightarrow \text{Ctrl} : \text{finish}} 0$ and therefore cannot be bisimilar to $\mathcal{M}'_{\text{Race}}$. \triangleright

5 Realisability Conditions

In general a global LTS may not be realisable. Therefore we are interested in (i) conditions that guarantee realisability and (ii) techniques to synthesise realisations from a global LTS \mathcal{M} . The notion of I -equivalence provides a helpful tool. The basic idea is to consider the source and target states of a global transition $q \xrightarrow{\text{out} \rightarrow \text{in} : m} \mathcal{M} q'$ to be indistinguishable for a component $i \in I$ if i does not participate in the interaction, i.e. $i \notin \text{out} \cup \text{in}$ (cf. also [14] and the discussion below).

Definition 8 (I -equivalence). *Let $\mathcal{M} = (Q, q_0, \Gamma, T)$ be a global LTS over (Σ, Γ) . An I -equivalence over \mathcal{M} is a family $\equiv = (\equiv_i)_{i \in I}$ of equivalence relations $\equiv_i \subseteq Q \times Q$ (reflexive, symmetric, and transitive) such that $q \equiv_i q'$ holds whenever there exists a transition $q \xrightarrow{\text{out} \rightarrow \text{in} : m} \mathcal{M} q'$ with $i \notin \text{out} \cup \text{in}$. The equivalence class of a state $q \in Q$ w.r.t. \equiv_i is the set $[q]_{\equiv_i} = \{q' \in Q \mid q' \equiv_i q\}$.*

5.1 Condition for Realisability Using Rich Local Actions

First, we will formulate our realisability condition for the case of rich local actions. We consider a global LTS \mathcal{M} over (Σ, Γ) . Our goal is to find an I -equivalence $(\equiv_i)_{i \in I}$ over \mathcal{M} such that for each interaction $(\text{out} \rightarrow \text{in} : m) \in \Gamma$ the

following holds. Assume, for simplicity, that $\text{out} \cup \text{in} = \{1, \dots, n\}$. Whenever there is a combination q_1, \dots, q_n of n (not necessarily different) global states together with a global “glue” state g , i.e. for each $j \in \text{out} \cup \text{in}$, $q_j \equiv_j g$, then we expect: if $\text{out} \rightarrow \text{in} : m$ is enabled in each global state q_1, \dots, q_n then each $j \in \text{out} \cup \text{in}$ should also be able to participate in $\text{out} \rightarrow \text{in} : m$ when the global state is g , since j cannot distinguish g from q_j . Thus the global interaction $\text{out} \rightarrow \text{in} : m$ should be enabled in g and preserve I -equivalences.

Definition 9 (realisability condition (rich case)). *Let \mathcal{M} be a global LTS over (Σ, Γ) . The realisability condition $RC(\mathcal{M})^r$ for \mathcal{M} with respect to rich local actions says that there exists an I -equivalence $\equiv = (\equiv_i)_{i \in I}$ over \mathcal{M} such that the following property $RC(\mathcal{M}, \equiv)^r$ holds.*

For all $\gamma = (\text{out} \rightarrow \text{in} : m) \in \Gamma$ with $\text{out} \cup \text{in} = \{k_1, \dots, k_n\}$ we have:

$$\forall \left(\begin{array}{c} q_1 \xrightarrow{\gamma}_{\mathcal{M}} q'_1 \quad \dots \quad q_n \xrightarrow{\gamma}_{\mathcal{M}} q'_n \\ g \in \bigcap_{j=1}^n [q_j]_{\equiv_{k_j}} \end{array} \right) \quad \exists g' : \left(\begin{array}{c} g \xrightarrow{\gamma}_{\mathcal{M}} g' \\ g' \in \bigcap_{j=1}^n [q'_j]_{\equiv_{k_j}} \end{array} \right)$$

Theorem 2 will provide a constructive argument why condition $RC(\mathcal{M})^r$ ensures realisability with rich local actions. Local quotients are crucial for this.

Definition 10 (local quotients with rich local actions). *Let $\mathcal{M} = (Q, q_0, \Gamma, T)$ be a global LTS over (Σ, Γ) and $\equiv = (\equiv_i)_{i \in I}$ an I -equivalence over \mathcal{M} . For each $i \in I$ the local i -quotient of \mathcal{M} with rich local actions is the LTS $(\mathcal{M}/\equiv_i)^r = (Q/\equiv_i, [q_0]_{\equiv_i}, A_i^r(\Gamma), (T/\equiv_i)^r)$ where*


- $Q/\equiv_i = \{[q]_{\equiv_i} \mid q \in Q\}$,
- $(T/\equiv_i)^r$ is the least set of transitions generated by the following rules:

$$\frac{q \xrightarrow{\text{out} \rightarrow \text{in} : m}_{\mathcal{M}} q' \quad i \in \text{out}}{[q]_{\equiv_i} \xrightarrow{\text{out in} ! m}_{(\mathcal{M}/\equiv_i)^r} [q']_{\equiv_i}} \qquad \frac{q \xrightarrow{\text{out} \rightarrow \text{in} : m}_{\mathcal{M}} q' \quad i \in \text{in}}{[q]_{\equiv_i} \xrightarrow{\text{out in} ? m}_{(\mathcal{M}/\equiv_i)^r} [q']_{\equiv_i}}$$


Note that $[q]_{\equiv_i} \xrightarrow{\text{out in} ! m}_{(\mathcal{M}/\equiv_i)^r} [q']_{\equiv_i}$ implies that there exist $\hat{q} \in [q]_{\equiv_i}$, $\hat{q}' \in [q']_{\equiv_i}$ and a transition $\hat{q} \xrightarrow{\text{out} \rightarrow \text{in} : m}_{\mathcal{M}} \hat{q}'$ with $i \in \text{out}$ (and similarly for $\text{out in} ? m$).

Theorem 2. *Let \mathcal{M} be a global LTS over (Σ, Γ) and let $\equiv = (\equiv_i)_{i \in I}$ be an I -equivalence over \mathcal{M} . If $RC(\mathcal{M}, \equiv)^r$ holds, then $\mathcal{M} \sim \otimes_I^r ((\mathcal{M}/\equiv_i)^r)_{i \in I}$.*

Example 4.

Consider the global LTS $\mathcal{M}_{\text{Race}}$  (Fig. 2). We show that $RC(\mathcal{M}_{\text{Race}})^r$ holds and how to construct, following Theorem 2, a realisation of $\mathcal{M}_{\text{Race}}$. More concretely, we take the family of equivalences $\equiv = (\equiv_i)_{i \in \{\text{Ctrl}, \text{R1}, \text{R2}\}}$ that obeys $RC(\mathcal{M}_{\text{Race}}, \equiv)^r$ (see below) and partitions the state space Q as follows: $Q/\equiv_{\text{Ctrl}} = \{\{0\}, \{1\}, \{2\}, \{3\}\}$, $Q/\equiv_{\text{R1}} = \{\{0, 2\}, \{1, 3\}\}$, and $Q/\equiv_{\text{R2}} = \{\{0, 3\}, \{1, 2\}\}$.

Using these equivalences, the local quotients for **R1** and **R2** are as follows:



The local quotient for **Ctrl** is isomorphic to $\mathcal{M}_{\text{Race}}$ but with local labels. Thus we have obtained a system which is a realisation with rich local actions of $\mathcal{M}_{\text{Race}}$. The local quotients coincide, up to renaming of states, with the local LTS used in the system $\mathcal{S}_{\text{Race}}^r$ considered in Sect. 1.

Now, we illustrate how to verify $RC(\mathcal{M}_{\text{Race}}, \equiv)^r$ using, as an example, the interaction $\gamma = \mathbf{R1} \rightarrow \mathbf{Ctrl} : \mathbf{finish}$ which appears twice in $\mathcal{M}_{\text{Race}}$: $t_{12} = (1 \xrightarrow{\gamma} 2)$ and $t_{30} = (3 \xrightarrow{\gamma} 0)$. There are two participants involved: **R1** and **Ctrl**. Hence we need to consider four combinations: (t_{12}, t_{12}) , (t_{12}, t_{30}) , (t_{30}, t_{12}) , and (t_{30}, t_{30}) . For example, using the combination (t_{30}, t_{12}) , we compute the glue $[3]_{\equiv_{\mathbf{R1}}} \cap [1]_{\equiv_{\mathbf{Ctrl}}} = \{1\}$. Then, trivially, there exists a transition $1 \xrightarrow{\gamma} 2$, and $2 \in [0]_{\equiv_{\mathbf{R1}}} \cap [2]_{\equiv_{\mathbf{Ctrl}}}$. The same can be shown for all the glues found for the other three combinations. ▷

In general it may happen that a global LTS \mathcal{M} does not satisfy $RC(\mathcal{M})^r$ but nevertheless is realisable. We can prove that $RC(\mathcal{M})^r$ is a necessary condition to obtain a realisation which is related to a global model by a *functional* bisimulation. More interesting, however, would be to weaken $RC(\mathcal{M})^r$ such that it becomes necessary for realisability with respect to arbitrary bisimulations. This is an open and challenging question.

Example 5. Consider the signature $\Sigma = (I, M)$ with $I = \{\mathbf{a}, \mathbf{b}, \mathbf{c}\}$, $M = \{\mathbf{m}\}$ and the set $\Gamma = \{\mathbf{a} \rightarrow \mathbf{b} : \mathbf{m}, \mathbf{c} \rightarrow \mathbf{b} : \mathbf{m}, \mathbf{c} \rightarrow \mathbf{a} : \mathbf{m}\}$ of Σ -interactions. The global LTS \mathcal{M} in Table 3 (left) is realisable by the system $\mathcal{S}^r = \{\mathcal{M}_a^r, \mathcal{M}_b^r, \mathcal{M}_c^r\}$. To see this, we compute the Γ -composition $\mathcal{M}' = \otimes_{\Gamma}^r \{\mathcal{M}_a^r, \mathcal{M}_b^r, \mathcal{M}_c^r\}$ shown in Table 3 (right). Obviously \mathcal{M} is bisimilar to \mathcal{M}' and hence \mathcal{M} is realisable. However, \mathcal{M} does not satisfy the realisability condition $RC(\mathcal{M})^r$.

We prove this by contradiction. Assume that $\equiv = \{\equiv_a, \equiv_b, \equiv_c\}$ is an I -equivalence such that $RC(\mathcal{M}, \equiv)^r$ holds. Now consider the interaction $\mathbf{a} \rightarrow \mathbf{b} : \mathbf{m}$, the global state 0 of \mathcal{M} and the transition $0 \xrightarrow{\mathbf{a} \rightarrow \mathbf{b} : \mathbf{m}}_{\mathcal{M}} 1$. Obviously, $0 \equiv_a 1$

Table 3. Global LTS \mathcal{M} not satisfying $RC(\mathcal{M})^r$ but with realisation $\mathcal{S}^r = \{\mathcal{M}_a^r, \mathcal{M}_b^r, \mathcal{M}_c^r\}$ where $\mathcal{M}' = \otimes^r \{\mathcal{M}_a^r, \mathcal{M}_b^r, \mathcal{M}_c^r\}$

Global $\mathcal{M}^{\mathcal{G}}$	Local \mathcal{M}_a^r	Local \mathcal{M}_b^r	Local \mathcal{M}_c^r	Composed \mathcal{M}'

and $0 \equiv_{\mathbf{b}} 1$ must hold since there is the transition $0 \xrightarrow{\mathbf{c} \rightarrow \mathbf{b} : \mathbf{m}}_{\mathcal{M}} 1$ where \mathbf{a} does not participate and the transition $0 \xrightarrow{\mathbf{c} \rightarrow \mathbf{a} : \mathbf{m}}_{\mathcal{M}} 1$ where \mathbf{b} does not participate. So we can take 1 as a glue state between the global states $q_1 = 0$ and $q_2 = 0$. Then we consider the transition $0 \xrightarrow{\mathbf{a} \rightarrow \mathbf{b} : \mathbf{m}}_{\mathcal{M}} 1$ one time for q_1 and one time for q_2 . Since we have assumed $RC(\mathcal{M}, \equiv)^r$, there must be a transition $1 \xrightarrow{\mathbf{a} \rightarrow \mathbf{b} : \mathbf{m}}_{\mathcal{M}'}$ leaving the glue state which is, however, not the case. Contradiction! Note that nevertheless the bisimilar global LTS \mathcal{M}' does satisfy $RC(\mathcal{M}')^r$. The example can be checked at \mathcal{M}^{\boxplus} and \mathcal{M}'^{\boxplus} . \triangleright

Discussion. Our realisability condition $RC(\mathcal{M})^r$, based on the notion of an I -equivalence $(\equiv_i)_{i \in I}$, is strongly related to a condition for implementability in [14, Theorem 3.1]. In fact, $RC(\mathcal{M})^r$ can be seen as a generalisation of [14] since we consider multi-interactions with distinguished sets of senders and receivers and also specifications for admissible interactions represented by Γ . Thus we get a generic realisability notion based on Γ -composition rather than full synchronisation. Moreover, our condition ensures realisability modulo bisimulation instead of isomorphism. Technically, implementability with respect to isomorphism is achieved in [14, Theorem 3.1] by requiring that whenever two global states q and q' are i -equivalent, i.e. $(q \equiv_i q')$, for all $i \in I$, then $q = q'$. We do not use this assumption and thus can get realisations modulo bisimilarity which do not realise a global LTS up to isomorphism (cf. [6, Example 8]). Note that [14, Theorem 6.2] also provides a proposal to deal with implementability modulo bisimulation under the assumption of “deterministic product transition systems”. In the next section, we study a realisability condition for the case of poor local actions, which deviates significantly from [14].

5.2 Condition for Realisability Using Poor Local Actions

We return to the question of how to check realisability, now in the case of poor local actions. The notion of I -equivalence is again the key. Note, however, that the computation of an appropriate I -equivalence may differ from the rich case.

The realisability condition below is stronger than the one for rich local actions in Definition 9. Intuitively, the reason is that local LTS with poor local actions have, in general, more choices for synchronisation and therefore a global LTS must support these choices in order to be realisable. For each interaction $(\mathbf{out} \rightarrow \mathbf{in} : \mathbf{m}) \in \Gamma$, we require in more cases the enabledness in a glue state g . More concretely, $\mathbf{out} \rightarrow \mathbf{in} : \mathbf{m}$ must be enabled in g already when in the j -equivalent states, say q_j , component j is able to output/input message m independently of the communication partners named in $\mathbf{out} \cup \mathbf{in}$, since those would anyway not be known from a poor local action. This is formally reflected by considering the interactions $\gamma_1, \dots, \gamma_n$ in the next definition.

Definition 11 (realisability condition (poor case)). *Let \mathcal{M} be a global LTS over (Σ, Γ) . The realisability condition $RC(\mathcal{M})^p$ for \mathcal{M} with respect to poor local actions says that there exists an I -equivalence $(\equiv_i)_{i \in I}$ over \mathcal{M} such that the following property $RC(\mathcal{M}, (\equiv_i)_{i \in I})^p$ holds.*

For all $\gamma = (\text{out} \rightarrow \text{in} : m) \in \Gamma$ with participants $\text{out} \cup \text{in} = \{k_1, \dots, k_n\}$ we get:

$$\forall \left(\begin{array}{l} \gamma_1 = (\text{out}_1 \rightarrow \text{in}_1 : m) \in \Gamma \\ k_1 \in (\text{out}_1 \cap \text{out}) \cup (\text{in}_1 \cap \text{in}) \\ \vdots \\ \gamma_n = (\text{out}_n \rightarrow \text{in}_n : m) \in \Gamma \\ k_n \in (\text{out}_n \cap \text{out}) \cup (\text{in}_n \cap \text{in}) \end{array} \right) \quad \forall \left(\begin{array}{l} q_1 \xrightarrow{\gamma_1}_{\mathcal{M}} q'_1 \\ \vdots \\ q_n \xrightarrow{\gamma_n}_{\mathcal{M}} q'_n \\ g \in \bigcap_{j=1}^n [q_j]_{\equiv_{k_j}} \end{array} \right) \quad \exists g' : \left(\begin{array}{l} g \xrightarrow{\gamma}_{\mathcal{M}} g' \\ g' \in \bigcap_{j=1}^n [q'_j]_{\equiv_{k_j}} \end{array} \right)$$


To prove that the condition $RC(\mathcal{M})^p$ indeed guarantees realisability with poor local actions, the idea is again to consider local quotients. Their construction is, however, different from the rich case.


Definition 12 (local quotients with poor local actions). Let $\mathcal{M} = (Q, q_0, \Gamma, T)$ be a global LTS over (Σ, Γ) and $\equiv = (\equiv_i)_{i \in I}$ an I -equivalence over \mathcal{M} . For each $i \in I$ the local i -quotient of \mathcal{M} with poor local actions is the LTS $(\mathcal{M}/\equiv_i)^p = (Q/\equiv_i, [q_0]_{\equiv_i}, A_i^p(\Gamma), (T/\equiv_i)^p)$ where

- $Q/\equiv_i = \{[q]_{\equiv_i} \mid q \in Q\}$,
- $(T/\equiv_i)^p$ is the least set of transitions generated by the following rules:

$$\frac{q \xrightarrow{\text{out} \rightarrow \text{in} : m}_{\mathcal{M}} q' \quad i \in \text{out}}{[q]_{\equiv_i} \xrightarrow{!m}_{(\mathcal{M}/\equiv_i)^p} [q']_{\equiv_i}} \quad \frac{q \xrightarrow{\text{out} \rightarrow \text{in} : m}_{\mathcal{M}} q' \quad i \in \text{in}}{[q]_{\equiv_i} \xrightarrow{?m}_{(\mathcal{M}/\equiv_i)^p} [q']_{\equiv_i}}$$

Theorem 3. Let \mathcal{M} be a global LTS over (Σ, Γ) and let $\equiv = (\equiv_i)_{i \in I}$ be an I -equivalence over \mathcal{M} . If $RC(\mathcal{M}, \equiv)^p$ holds then $\mathcal{M} \sim \otimes_{\Gamma}^p ((\mathcal{M}/\equiv_i)^p)_{i \in I}$.

Example 6. Consider the global LTS $\mathcal{M}_{\text{Race}}$  (Fig. 2). We show $RC(\mathcal{M}_{\text{Race}})^p$ holds and how, following Theorem 3, a realisation of $\mathcal{M}_{\text{Race}}$ with poor local actions can be constructed. The situation differs from the rich case in Example 4, since the equivalence for **Ctrl** must be chosen differently. We use the family of equivalences $\equiv = (\equiv_i)_{i \in \{\text{Ctrl}, \text{R1}, \text{R2}\}}$ that obeys $RC(\mathcal{M}_{\text{Race}}, \equiv)^p$ (see below) and partitions the state space Q as follows: $Q/\equiv_{\text{Ctrl}} = \{\{0\}, \{1\}, \{2, 3\}\}$, $Q/\equiv_{\text{R1}} = \{\{0, 2\}, \{1, 3\}\}$, and $Q/\equiv_{\text{R2}} = \{\{0, 3\}, \{1, 2\}\}$. Using these equivalences, the local quotients for **Ctrl**, **R1** and **R2** are as follows:



Thus we have obtained a system which is a realisation with poor local actions of $\mathcal{M}_{\text{Race}}$. The local quotients coincide, up to renaming of states, with the local LTS used in the system $\mathcal{S}_{\text{Race}}^p$ considered in Sect. 1.

Now, we illustrate how to verify $RC(\mathcal{M}_{\text{Race}}, \equiv)^r$ using, as an example, the interaction $\text{R1} \rightarrow \text{Ctrl} : \text{finish}$. We have $1 \xrightarrow{\text{R1} \rightarrow \text{Ctrl} : \text{finish}}_{\mathcal{M}_{\text{Race}}} 2$. We also have $1 \xrightarrow{\text{R2} \rightarrow \text{Ctrl} : \text{finish}}_{\mathcal{M}_{\text{Race}}} 3$ (we must consider the interaction $\text{R2} \rightarrow \text{Ctrl} : \text{finish}$ as well

since we are in the poor case). Taking 1 as a (trivial) glue state, we thus have, as required, the existence of $1 \xrightarrow{R1 \rightarrow Ctrl: finish} \mathcal{M}_{Race} 2$ but also it is required that $2 \equiv_{Ctrl} 3$ must hold which is the case. Note that we wouldn't have succeeded here if we would have taken the identity for \equiv_{Ctrl} as done for the rich case. \triangleright

6 Tool Support: Ceta

We developed a supporting prototypical tool *Ceta* (Choreographic Extended Team Automata) to analyse global specifications and produce visualisations of state machines. It is open-source, available at <https://github.com/arcalab/choreo/tree/ceta>, and executable by browsing to <https://lmf.di.uminho.pt/ceta>.

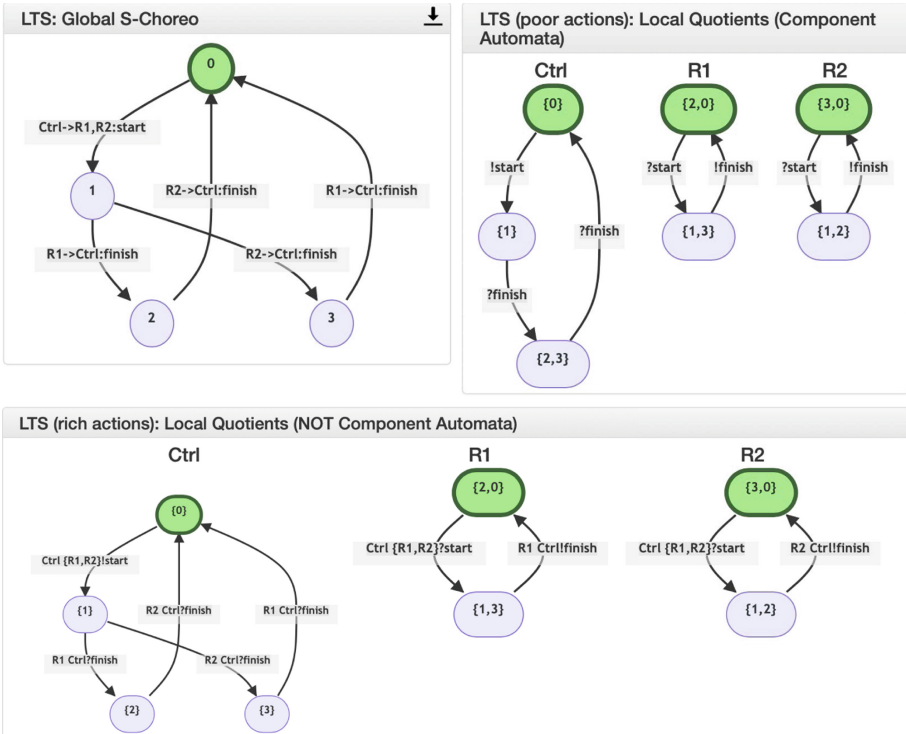


Fig. 4. Screenshots of the Ceta tool (<http://lmf.di.uminho.pt/ceta>)

Ceta starts with a web browser, opening a static webpage that uses our compiled JavaScript built with the Caos framework [32] (cf. screenshot in Fig. 4). The user input is a global protocol described in a choreographic language, resembling regular expressions of interactions. A set of examples with descriptions is also included, covering the examples presented in this paper. The analyses include graphical views of: (i) the global LTS; (ii) local LTS's with rich actions; and

(iii) local LTS's with poor actions. Other widgets provide further insights, such as the composition of the local LTS's, the intermediate equivalence classes for both the rich and poor case, the synchronous composition of local LTS's, and bisimulations between the global protocol and composed systems. Readable error messages are provided when the realisability conditions do not hold.

7 Conclusion

We have proposed a rigorous discipline for developing interaction-based systems. At the heart of our methodology lies the realisation of a global interaction model, i.e. its decomposition into a set of (possibly distributed) components with synchronous communication. We have investigated realisability conditions for two different localisation styles (rich and poor local actions) and techniques to synthesise realisations. Our approach is generic with respect to the choice of admissible interaction sets which may contain arbitrary interactions between multiple senders and receivers but may also be restricted, e.g., to various forms of communication, like multicast or peer-to-peer communication. Due to the generic nature of our notion of an interaction set, our results can be instantiated by different concrete coordination formalisms. For instance, synchronisation type specifications used in the framework of (extended) team automata [5] as well as interactions used in BIP [11] can be represented by interaction sets. Our results should then be directly applicable, to extend the team automata framework as well as BIP by global LTS and to generate distributed component systems for them on the basis of our realisation conditions.

In future research, we plan to (i) integrate the treatment of internal actions using weak bisimulation equivalence for the realisation notions; (ii) consider communication properties (like receptiveness and responsiveness, cf. [5]) when systems are generated from global models; (iii) study open global models (systems) and their composition; and (iv) investigate realisability conditions in the context of asynchronous communication. Moreover, we are still looking for a weaker version of our realisability condition for synchronous systems making it necessary for arbitrary (also non-functional) bisimulations.

Furthermore, we intend to investigate the relation with work in the literature on the decomposition of related formalisms like (Petri net or algebraic) processes into (indecomposable) components [27,31] used to parallelise (verification of) concurrent systems [16,18] or obtain better (optimised) implementations [33,34].

Acknowledgments. *Ter Beek* was supported by MUR PRIN 2020TL3X8X project T-LADIES (Typeful Language Adaptation for Dynamic, Interacting and Evolving Systems) and *Proença* by the CISTER Research Unit (UIDP/UIDB/04234/2020), financed by National Funds through FCT/MCTES (Portuguese Foundation for Science and Technology); by project IBEX (PTDC/CCI-COM/4280/2021) financed by national funds through FCT; and by project Route 25 (ref. TRB/2022/00061 – C645463824-00000063) funded by the EU/Next Generation, within the Recovery and Resilience Plan.

References

1. Barbanera, F., Lanese, I., Tuosto, E.: Choreography automata. In: Bludze, S., Bocchi, L. (eds.) COORDINATION 2020. LNCS, vol. 12134, pp. 86–106. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-50029-0_6
2. Barbanera, F., Lanese, I., Tuosto, E.: Formal choreographic languages. In: ter Beek, M.H., Sirjani, M. (eds.) COORDINATION 2022. LNCS, vol. 13271, pp. 121–139. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-08143-9_8
3. Basile, D., Degano, P., Ferrari, G.-L., Tuosto, E.: Relating two automata-based models of orchestration and choreography. *J. Log. Algebr. Meth. Program.* **85**(3), 425–446 (2016). <https://doi.org/10.1016/J.JLAMP.2015.09.011>
4. ter Beek, M.H., Ellis, C.A., Kleijn, J., Rozenberg, G.: Synchronizations in team automata for groupware systems. *Comput. Sup. Coop. Work* **12**(1), 21–69 (2003). <https://doi.org/10.1023/A:1022407907596>
5. ter Beek, M.H., Hennicker, R., Kleijn, J.: Compositionality of safe communication in systems of team automata. In: Pun, V.K.I., Stolz, V., Simao, A. (eds.) ICTAC 2020. LNCS, vol. 12545, pp. 200–220. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-64276-1_11
6. ter Beek, M.H., Hennicker, R., Proença, J.: Realisability of global models of interaction (extended version). Technical report, Zenodo (2023). <https://doi.org/10.5281/zenodo.8377188>
7. ter Beek, M.H., Cledou, G., Hennicker, R., Proença, J.: Can We communicate? Using dynamic logic to verify team automata. In: Chechik, M., Katoen, J.P., Leucker, M. (eds.) FM 2023. LNCS, vol. 14000, pp. 122–141. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-27481-7_9
8. Bejleri, A., Yoshida, N.: Synchronous multiparty session types. *ENTCS* **241**, 3–33 (2008). <https://doi.org/10.1016/j.entcs.2009.06.002>
9. Ben-David, S., Chechik, M., Gurfinkel, A., Uchitel, S.: CSSL: a logic for specifying conditional scenarios. In: ESEC/FSE, pp. 37–47. ACM (2011). <https://doi.org/10.1145/2025113.2025123>
10. van Benthem, J., van Eijck, J., Stebletsova, V.: Modal logic, transition systems and processes. *J. Log. Comput.* **4**(5), 811–855 (1994). <https://doi.org/10.1093/logcom/4.5.811>
11. Bludze, S., Sifakis, J.: The algebra of connectors: structuring interaction in BIP. *IEEE Trans. Comput.* **57**(10), 1315–1330 (2008). <https://doi.org/10.1109/TC.2008.26>
12. Bunte, O., et al.: The mCRL2 toolset for analysing concurrent systems. In: Vojnar, T., Zhang, L. (eds.) TACAS 2019. LNCS, vol. 11428, pp. 21–39. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-17465-1_2
13. Castagna, G., Dezani-Ciancaglini, M., Padovani, L.: On global types and multiparty sessions. *Log. Methods Comput. Sci.* **8**(1), 24:1–24:45 (2012). [https://doi.org/10.2168/LMCS-8\(1:24\)2012](https://doi.org/10.2168/LMCS-8(1:24)2012)
14. Castellani, I., Mukund, M., Thiagarajan, P.S.: Synthesizing distributed transition systems from global specifications. In: Rangan, C.P., Raman, V., Ramanujam, R. (eds.) FSTTCS 1999. LNCS, vol. 1738, pp. 219–231. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-46691-6_17
15. Cengarle, M.V., Knapp, A., Mühlberger, H.: Interactions. In: Lano, K. (ed.) UML 2 Semantics and Applications, chap. 9, pp. 205–248. Wiley (2009). <https://doi.org/10.1002/9780470522622>
16. Corradini, F., Gorrieri, R., Marchignoli, D.: Towards parallelization of concurrent systems. *RAIRO Theor. Inform. Appl.* **32**(4–6), 99–125 (1998). <https://doi.org/10.1051/ita/1998324-600991>

17. de Alfaro, L., Henzinger, T.A.: Interface automata. In: ESEC/FSE, pp. 109–120. ACM (2001). <https://doi.org/10.1145/503209.503226>
18. Groote, J.F., Moller, F.: Verification of parallel systems via decomposition. In: Cleaveland, W.R. (ed.) CONCUR 1992. LNCS, vol. 630, pp. 62–76. Springer, Heidelberg (1992). <https://doi.org/10.1007/BFb0084783>
19. Harel, D., Kozen, D., Tiuryn, J.: Dynamic Logic. Foundations of Computing. MIT Press, Cambridge (2000). <https://doi.org/10.7551/mitpress/2516.001.0001>
20. Harel, D., Thiagarajan, P.S.: Message sequence charts. In: Lavagno, L., Martin, G., Selic, B. (eds.) UML for Real: Design of Embedded Real-Time Systems, pp. 77–105. Kluwer (2003). https://doi.org/10.1007/0-306-48738-1_4
21. Hennicker, R.: Role-based development of dynamically evolving ensembles. In: Fiadeiro, J.L., Tutu, I. (eds.) WADT 2018. LNCS, vol. 11563, pp. 3–24. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-23220-7_1
22. Hennicker, R., Wirsing, M.: Dynamic logic for ensembles. In: Margaria, T., Steffen, B. (eds.) ISoLA 2018. LNCS, vol. 11246, pp. 32–47. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-03424-5_3
23. Honda, K., Yoshida, N., Carbone, M.: Multiparty asynchronous session types. In: POPL, pp. 273–284. ACM (2008). <https://doi.org/10.1145/1328438.1328472>
24. Hüttel, H., et al.: Foundations of session types and behavioural contracts. ACM Comput. Surv. **49**(1), 3:1–3:36 (2016). <https://doi.org/10.1145/2873052>
25. Jongmans, S.S., Ferreira, F.: Synthetic behavioural typing: sound, regular multiparty sessions via implicit local types. In: ECOOP. LIPIcs, vol. 263, pp. 9:1–9:29. Dagstuhl (2023). <https://doi.org/10.4230/LIPIcs.ECOOP.2023.9>
26. Larsen, K.G., Nyman, U., Wasowski, A.: Modal I/O automata for interface and product line theories. In: De Nicola, R. (ed.) ESOP 2007. LNCS, vol. 4421, pp. 64–79. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-71316-6_6
27. Luttkik, B.: Unique parallel decomposition in branching and weak bisimulation semantics. Theor. Comput. Sci. **612**, 29–44 (2016). <https://doi.org/10.1016/j.tcs.2015.10.013>
28. Lynch, N.A., Tuttle, M.R.: An introduction to Input/Output automata. CWI Q. **2**(3), 219–246 (1989). <https://ir.cwi.nl/pub/18164>
29. Magee, J., Kramer, J.: Concurrency: State Models & Java Programming. Wiley, Hoboken (2006)
30. Micskei, Z., Waeselynck, H.: The many meanings of UML 2 sequence diagrams: a survey. Softw. Syst. Model. **10**(4), 489–514 (2011). <https://doi.org/10.1007/s10270-010-0157-9>
31. Milner, R., Moller, F.: Unique decomposition of processes. Theor. Comput. Sci. **107**(2), 357–363 (1993). [https://doi.org/10.1016/0304-3975\(93\)90176-T](https://doi.org/10.1016/0304-3975(93)90176-T)
32. Proença, J., Edixhoven, L.: Caos: a reusable scala web animator of operational semantics. In: Jongmans, S.S., Lopes, A. (eds.) COORDINATION 2023. LNCS, vol. 13908, pp. 163–171. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-35361-1_9
33. Teren, V., Cortadella, J., Villa, T.: Decomposition of transition systems into sets of synchronizing state machines. In: DSD, pp. 77–81. IEEE (2021). <https://doi.org/10.1109/DSD53832.2021.00021>
34. Teren, V., Cortadella, J., Villa, T.: Decomposition of transition systems into sets of synchronizing free-choice Petri Nets. In: DSD, pp. 165–173. IEEE (2022). <https://doi.org/10.1109/DSD57027.2022.00031>
35. Tuosto, E., Guanciale, R.: Semantics of global view of choreographies. J. Log. Algebr. Meth. Program. **95**, 17–40 (2018). <https://doi.org/10.1016/j.jlamp.2017.11.002>