

Rollback Recovery in Session-based Programming ^{*}

– Technical Report –

Claudio Antares Mezzina¹, Francesco Tiezzi², and Nobuko Yoshida³

¹ Università degli Studi di Urbino Carlo Bo, Italy

claudio.mezzina@uniurb.it

² Università degli Studi di Firenze, Italy

francesco.tiezzi@unifi.it

³ University of Oxford, United Kingdom

nobuko.yoshida@cs.ox.ac.uk

Abstract. To react to unforeseen circumstances or amend abnormal situations in communication-centric systems, programmers are in charge of “undoing” the interactions which led to an undesired state. To assist this task, session-based languages can be endowed with reversibility mechanisms. In this paper we propose a language enriched with programming facilities to *commit* session interactions, to *roll back* the computation to a previous commit point, and to *abort* the session. Rollbacks in our language always bring the system to previous visited states and a rollback cannot bring the system back to a point prior to the last commit. Programmers are relieved from the burden of ensuring that a rollback never restores a checkpoint imposed by a session participant different from the rollback requester. Such undesired situations are prevented at design-time (statically) by relying on a decidable *compliance* check at the type level, implemented in MAUDE. We show that the language satisfies error-freedom and progress of a session.

1 Introduction

Reversible computing [1,27] has gained interest for its application to different fields: from modelling biological/chemical phenomena [19], to simulation [30], debugging [13] and modelling fault-tolerant systems [11,20,34]. Our interest focuses on this latter application and stems from the fact that reversibility can be used to rigorously model, implement and revisit programming abstractions for reliable software systems.

Recent works [4,25,26,6,32] have studied the effect of reversibility in communication-centric scenarios, as a way to correct faulty computations by bringing back the system to

^{*} This research was funded in whole, or in part, by EPSRC EP/T006544/2, EP/K011715/1, EP/K034413/1, EP/L00058X/1, EP/N027833/2, EP/N028201/1, EP/T014709/2, EP/V000462/1, EP/X015955/1, NCSS/EPSRC VeTSS and Horizon EU TaRDIS 101093006. This work was also partially supported by the Italian MUR PRIN 2020 project NiRvAna, the Italian MUR PRIN 2017 project SEDUCE n. 2017TWRCNB, the French ANR project ANR-18-CE25-0007 DCore, the INdAM - GNCS Project “Proprietà qualitative e quantitative di sistemi reversibili” n. CUP E55F2200027001, and the project SERICS (PE00000014) under the NRRP MUR program funded by the EU - NextGenerationEU. For the purpose of Open Access, the authors have applied a CC BY public copyright licence to any Author Accepted Manuscript (AAM) version arising from this submission.

a previous consistent state. In this setting, processes' behaviours are strongly disciplined by their types, prescribing the actions they have to perform within a *session*. A session consists of a structured series of message exchanges, whose flow can be controlled via conditional choices, branching and recursion. Correctness of communication is statically guaranteed by a framework based on a (session) type discipline [17]. None of the aforementioned works addresses systems in which the participants can *explicitly* abort the session, commit a computation and roll it back to a previous checkpoint. In this paper, we aim at filling this gap. We explain below the distinctive aspects of our checkpoint-based rollback recovery approach.

Linguistic primitives to explicitly program reversible sessions. We introduce three primitives to: (i) *commit* a session, preventing undoing the interactions performed so far along the session; (ii) *roll back* a session, restoring the last saved process checkpoints; (iii) *abort* a session, to discard the session, and hence all interactions already performed in it, thus allowing another session of the same protocol to start with possible different participants. Notice that most proposals in the literature (e.g., [2,3,4]) only consider an abstract view, as they focus on reversible contracts (i.e., types). Instead, we focus on programming primitives at process level, and use types for guaranteeing a safe and consistent system evolution.

Asynchronous commits. Our commit primitive does not require a session-wide synchronisation among all participants, as it is a local decision. However, its effect is on the whole session, as it affects the other session participants. This means that each participant can independently decide when to commit. Such flexibility comes at the cost of being error-prone, especially considering that the programmer has not only to deal with the usual forward executions, but also with the backward ones. Our type discipline allows for ruling out programs which may lead to these errors. The key idea of our approach is that *a session participant executing a rollback action is interested in restoring the last checkpoint he/she has committed*. For the success of the rollback recovery it is irrelevant whether the 'passive' participants go back to their own last checkpoints. Instead, if the 'active' participant is unable to restore the last checkpoint he/she has created, because it has been replaced by a checkpoint imposed by another participant, the rollback recovery is considered unsatisfactory.

In our framework, programmers are relieved from the burden of ensuring the satisfaction of rollbacks, since undesired situations are prevented at design time (statically) by relying on a *compliance* check at the type level. To this aim, we introduce *cherry-pi* (*checkpoint-based rollback recovery pi-calculus*), a variant of the session-based π -calculus [38,18] enriched with rollback recovery primitives. We present here a binary version of the calculus, which is more convenient to demonstrate the essence of our rollback recovery approach; the proposed approach can be seamlessly extended to multiparty sessions (see Appendix B). A key difference with respect to the standard binary type discipline is the *relaxation* of the duality requirement. The types of two session participants are not required to be dual, but they will be compared with respect to a compliance relation (as in [5]), which also takes into account the effects of commit and rollback actions. Such relaxation also involves the requirements concerning selection and branching types, and those concerning branches of conditional choices. The *cherry-pi*

type system is used to infer types of session participants, which are then combined together for the compliance check.

Reversibility in *cherry- π* is *controlled* via two specific primitives: a rollback one telling when a reverse computation has to take place, and a commit one limiting the scope of a potential reverse computation. This implies that the calculus is not fully reversible (i.e., backward computations are not always enabled), leading to have properties that are relaxed and different with respect to other reversible calculi [10,9,22,32]. We prove that *cherry- π* satisfies the following properties: (i) a rollback always brings back the system to a previous visited state and (ii) it is not possible to bring the computation back to a point prior to the last checkpoint, which implies that our commits have a persistent effect. Concerning soundness properties, we prove that (a) our compliance check is decidable, (b) compliance-checked *cherry- π* specifications never lead to communication errors (e.g., a blocked communication where there is a receiver without the corresponding sender), and (c) compliance-checked *cherry- π* specifications never activate undesirable rollbacks (according to our notion of rollback recovery mentioned above). Property (b) resembles the type safety property of session-based calculi (see, e.g., [38]), while property (c) is a new property specifically defined for *cherry- π* . The technical development of property proofs turns out to be more intricate than that of standard properties of session-based calculi, due to the combined use of type and compliance checking. To demonstrate feasibility and effectiveness of our rollback recovery approach, we have concretely implemented the compliance check using the MAUDE [8] framework (the code is available at <https://github.com/tiezzi/cherry-pi>).

Outline. Sec. 2 illustrates the key idea of our rollback recovery approach; Sec. 3 introduces the *cherry- π* calculus; Sec. 4 introduces typing and compliance checking; Sec. 5 presents the properties satisfied by *cherry- π* ; Sec. 6 concludes the paper with related and future work. Appendices report on omitted rules, extension to multiparty sessions, proofs of the results, and a further example.

2 A reversible video on demand service example

We discuss the motivations underlying our work by introducing our running example, a Video on Demand (VOD) scenario. The key idea is that a rollback requester is satisfied only if her restored checkpoint was set by herself. In Fig. 1(a), a service (*S*) offers to a user (*U*) videos with two different quality levels, namely high definition (HD) and standard definition (SD). After the *login*, *U* sends her video *request*, and receives the corresponding *price* and *metadata* (actors, directors, description, etc.) from *S*. According to this information, *U* selects the video quality. Then, she receives, first, a short *test* video (to check the audio and video quality in her device) and, finally, the requested *video*. If the vision of the HD test video is not satisfactory, *U* can roll back to her last checkpoint to possibly change the video quality, instead in the SD case *U* can abort the session.

Let us now add commit actions as in the run shown in Fig. 1(b). After receiving the price, *U* commits, while *S* commits after the quality selection. In this scenario, however, if *U* activates the rollback, she is unable to go back to the checkpoint she set with her commit action because the actual effect of rollback is to restore the checkpoint set by the commit action performed by *S*.

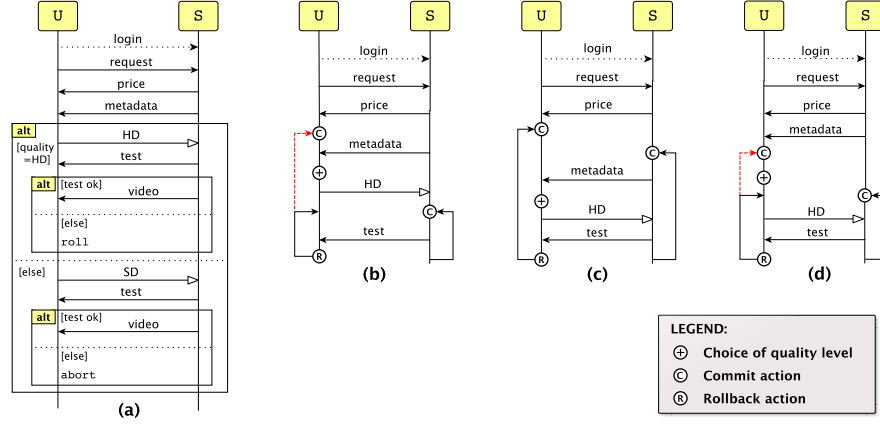


Fig. 1. VOD example: (a) a full description without commit actions; (b,d) runs with undesired rollback; (c) a run with satisfactory rollback.

In the scenario in Fig. 1(c), instead, S commits after sending the price to U . In this case, no matter who first performed the commit action, the rollback results to be satisfactory. Also if S commits later, the checkpoint of U remains unchanged, as U performed no other action between the two commits. This would not be the case if both U and S committed after the communication of the metadata, as in Fig. 1(d). If S commits before U , no rollback issue arises, but if U commits first it may happen that her internal decision is taken before S commits. In this case, U would not be able to go back to the checkpoint set by herself, and she would be unable to change the video quality.

These undesired rollbacks are caused by bad choices of commit points. We propose a compliance check that identifies these situations at design time.

3 The cherry-pi calculus

In this section, we introduce **cherry-pi**, a calculus (extending that in [38]) devised for studying sessions equipped with our checkpoint-based rollback recovery mechanism.

Syntax. The syntax of the **cherry-pi** calculus relies on the following base sets: *shared channels* (ranged over by a), used to initiate sessions; *session channels* (ranged over by s), consisting of pairs of *endpoints* (ranged over, with a slight abuse of notation, by s , \bar{s}) used by the two parties to interact within an established session; *labels* (ranged over by l), used to select and offer branching choices; *values* (ranged over by v), including booleans, integers and strings (whose *sorts*, ranged over by S , are `bool`, `int` and `str`, respectively), which are exchanged within a session; *variables* (ranged over by x, y, z), storing values and session endpoints; *process variables* (ranged over by X), used for recursion.

Collaborations, ranged over by C , are given by the grammar in Fig. 2. The key ingredient of the calculus is the set of actions for controlling the session rollback.

| | | |
|---------|--|--|
| $C ::=$ | $\bar{a}(x).P \mid a(x).P \mid C_1 \mid C_2$ | Collaborations request, accept, parallel |
| $P ::=$ | $x!\langle e \rangle.P \mid x?(y : S).P$ $\mid x \triangleleft l.P \mid x \triangleright \{l_1 : P_1, \dots, l_n : P_n\}$ $\mid \text{if } e \text{ then } P_1 \text{ else } P_2 \mid X \mid \mu X.P \mid \mathbf{0}$ $\mid \text{commit}.P \mid \text{roll} \mid \text{abort}$ | Processes output, input selection, branching choice, recursion, inact commit, roll, abort |
| $e ::=$ | $v \mid +(e_1, e_2) \mid \wedge(e_1, e_2) \mid \dots$ | Expressions |

Fig. 2. cherry- π syntax.

Actions `commit`, `roll` and `abort` are used, respectively, to commit a session (producing a checkpoint for each session participant), to trigger the session rollback (restoring the last committed checkpoints) or to abort the whole session. We discuss below the other constructs of the calculus, which are those typically used for session-based programming [16]. A *cherry- π* collaboration is a collection of *session initiators*, i.e. terms ready to initiate sessions by synchronising on shared channels. A synchronisation of two initiators $\bar{a}(x).P$ and $a(y).Q$ causes the generation of a fresh session channel, whose endpoints replace variables x and y in order to be used by the triggered processes P and Q , respectively, for later communications. No subordinate sessions can be initiated within a running session.

When a session is started, each participant executes a *process*. Processes are built up from the empty process $\mathbf{0}$ and basic actions by means of action prefix $_.$, conditional choice `if e then $_$ else $_$` , and recursion $\mu X. _$. Actions $x!\langle e \rangle$ and $x?(z : S)$ denote output and input via session endpoints replacing x and y , respectively. These communication primitives realise the standard synchronous message passing, where messages result from the evaluation of *expressions*, which are defined by means of standard operators on boolean, integer and string values. Variables that are arguments of input actions are (statically) typed by sorts. There is no need for statically typing the variables occurring as arguments of session initiating actions, as they are always replaced by session endpoints. Notice that in *cherry- π* the exchanged values cannot be endpoints, meaning that session delegation (i.e., channel-passing) is not considered. Actions $x \triangleleft l$ and $x \triangleright \{l_1 : P_1, \dots, l_n : P_n\}$ denote selection and branching (where l_1, \dots, l_n are pairwise distinct).

Example 1. Let us consider the VOD example informally introduced in Sec. 2. The scenario described in Fig. 1(a) with commit actions placed as in Fig. 1(b) is rendered in *cherry- π* as $C_{\text{US}} = \overline{\text{login}}(x). P_{\text{U}} \mid \text{login}(y). P_{\text{S}}$, where:

$$\begin{aligned}
 P_{\text{U}} = & x!\langle v_{\text{req}} \rangle. x?(x_{\text{price}} : \text{int}). \text{commit}. x?(x_{\text{meta}} : \text{str}). \text{if } (f_{\text{eval}}(x_{\text{price}}, x_{\text{meta}})) \\
 & \text{then } x \triangleleft l_{\text{HD}}. x?(x_{\text{testHD}} : \text{str}). \\
 & \quad (\text{if } (f_{\text{HD}}(x_{\text{testHD}})) \text{ then } x?(x_{\text{videoHD}} : \text{str}). \mathbf{0} \text{ else roll}) \\
 & \text{else } x \triangleleft l_{\text{SD}}. x?(x_{\text{testSD}} : \text{str}). \\
 & \quad (\text{if } (f_{\text{SD}}(x_{\text{testSD}})) \text{ then } x?(x_{\text{videoSD}} : \text{str}). \mathbf{0} \text{ else abort})
 \end{aligned}$$

| | |
|--|-----------------------|
| $C ::= \bar{a}(x).P \mid a(x).P \mid C_1 \mid C_2 \mid (\nu s : C_1) C_2 \mid \langle P_1 \rangle \blacktriangleright P_2$ | Collaborations |
| $P ::= r!\langle e \rangle.P \mid r?(y : S).P \mid r \triangleleft l.P \mid r \triangleright \{l_1 : P_1, \dots, l_n : P_n\} \mid \dots$ | Processes |

Fig. 3. cherry-pi runtime syntax (the rest of processes P and expressions e are as in Fig. 2).

$$\begin{aligned}
P_s = & y?(y_{req} : \mathbf{str}). y!\langle f_{price}(y_{req}) \rangle. y!\langle f_{meta}(y_{req}) \rangle. \\
& y \triangleright \{ l_{HD} : \mathbf{commit}. y!\langle f_{testHD}(y_{req}) \rangle. y!\langle f_{videoHD}(y_{req}) \rangle. \mathbf{0} , \\
& \quad l_{SD} : \mathbf{commit}. y!\langle f_{testSD}(y_{req}) \rangle. y!\langle f_{videoSD}(y_{req}) \rangle. \mathbf{0} \}
\end{aligned}$$

Notice that expressions used for decisions and computations are abstracted by relations $f_n(\cdot)$, whose definitions are left unspecified. Considering the placement of commit actions depicted in Fig. 1(c), the cherry-pi specification of the service's process becomes:

$$\begin{aligned}
& y?(y_{req} : \mathbf{str}). y!\langle f_{price}(y_{req}) \rangle. \mathbf{commit}. y!\langle f_{meta}(y_{req}) \rangle. \\
& y \triangleright \{ l_{HD} : y!\langle f_{testHD}(y_{req}) \rangle. y!\langle f_{videoHD}(y_{req}) \rangle. \mathbf{0} , \\
& \quad l_{SD} : y!\langle f_{testSD}(y_{req}) \rangle. y!\langle f_{videoSD}(y_{req}) \rangle. \mathbf{0} \}
\end{aligned}$$

Finally, considering the placement of commit actions depicted in Fig. 1(d), the cherry-pi specification of the user's process becomes:

$$x!\langle v_{req} \rangle. x?(x_{price} : \mathbf{int}). x?(x_{meta} : \mathbf{str}). \mathbf{commit}. \text{if } (f_{eval}(x_{price}, x_{meta})) \text{ then } \dots$$

Semantics. The operational semantics of cherry-pi is defined for *runtime* terms, generated by the extended syntax of the calculus in Fig. 3 (new constructs are highlighted by a grey background). We use r to denote *session identifiers*, i.e. session endpoints and variables. Those runtime terms that can be also generated by the grammar in Fig. 2 are called *initial collaborations*.

At collaboration level, two constructs are introduced: $(\nu s : C_1) C_2$ represents a *session* along the channel s with associated starting checkpoint C_1 (corresponding to the collaboration that has initialised the session) and code C_2 ; $\langle P_1 \rangle \blacktriangleright P_2$ represents a *log* storing the checkpoint P_1 associated to the code P_2 . At process level, the only difference is that session identifiers r are used as first argument of communicating actions. We extend the standard notion of binders to take into account $(\nu s : C_1) C_2$, which binds session endpoints s and \bar{s} in C_2 (in this respect, it acts similarly to the restriction of π -calculus, but its scope cannot be extended in order to avoid involving processes that do not belong to the session in the rollback effect). The derived notions of bound and free names (where *names* stand for variables, process variables and session endpoints), alpha-equivalence, and substitution are standard and we assume that bound names are pairwise distinct. The semantics of the calculus is defined for *closed* terms, i.e. terms without free variables and process variables.

Not all processes allowed by the extended syntax correspond to meaningful collaborations. In a general term the processes stored in logs may not be consistent with the computation that has taken place. We get rid of such malformed terms, as we will only consider those runtime terms, called *reachable* collaborations, obtained by means of reductions from initial collaborations.

$$\begin{array}{l}
k!\langle e \rangle.P \xrightarrow{k!\langle v \rangle} P \quad (e \downarrow v) \quad [\text{P-SND}] \qquad k?(x : S).P \xrightarrow{k?(x)} P \quad [\text{P-RCV}] \\
k \triangleleft l.P \xrightarrow{k \triangleleft l} P \quad [\text{P-SEL}] \qquad k \triangleright \{l_1 : P_1, \dots, l_n : P_n\} \xrightarrow{k \triangleright l_i} P_i \quad (1 \leq i \leq n) \quad [\text{P-BRN}] \\
\text{if } e \text{ then } P_1 \text{ else } P_2 \xrightarrow{\tau} P_1 \quad (e \downarrow \text{true}) \quad [\text{P-IFT}] \\
\text{commit}.P \xrightarrow{\text{cmt}} P \quad [\text{P-CMT}] \qquad \text{roll} \xrightarrow{\text{roll}} \mathbf{0} \quad [\text{P-RLL}] \qquad \text{abort} \xrightarrow{\text{abt}} \mathbf{0} \quad [\text{P-ABT}]
\end{array}$$

Fig. 4. cherry-pi semantics: auxiliary labelled relation.

The operational semantics of cherry-pi is given in terms of a standard *structural congruence* \equiv [18] and a *reduction* relation \rightarrow given as the union of the *forward reduction* relation \rightarrow and *backward reduction* relations \rightsquigarrow . The definition of the relation \rightarrow over closed collaborations relies on an auxiliary labelled relation $\xrightarrow{\ell}$ over processes that specifies the actions that processes can initially perform and the continuation process obtained after each such action. Given a reduction relation \mathcal{R} , we will indicate with \mathcal{R}^+ and \mathcal{R}^* respectively the *transitive* and the *reflexive-transitive* closure of \mathcal{R} .

The operational rules defining the auxiliary labelled relation are in Fig. 4 (omitted rules are in Appendix A). We use k to denote *generic session endpoints* (s or \bar{s}). Action label ℓ stands for either $k!\langle v \rangle$, $k?(x)$, $k \triangleleft l$, $k \triangleright l$, cmt , roll , abt , or τ . The meaning of the rules is straightforward, as they just produce as labels the actions currently enabled in the process. In doing that, expressions of sending actions and conditional choices are evaluated (auxiliary function $e \downarrow v$ says that closed expression e evaluates to value v).

The operational rules defining the reduction relation \rightarrow are reported in Fig. 5 (standard rules for congruence, in the forward and backward case, are omitted). We comment on salient points. Once a session is created, its initiating collaboration is stored in the session construct (rule [F-CON]). Communication, branching selection and internal conditional choices proceed as usual, without affecting logs (rules [F-COM], [F-LAB] and [F-IF]). A commit action updates the checkpoint of a session, by replacing the processes stored in the logs of the two involved parties (rule [F-CMT]). Notably, this form of commit is asynchronous as it does not require the passive participant to explicitly synchronise with the active participant by means of a primitive for accepting the commit. On the other hand, under the hood, a low-level implementation of this mechanism would synchronously update the logs of the involved parties. Conversely, a rollback action restores the processes in the two logs (rule [B-RLL]). The abort action (rule [B-ABT]), instead, kills the session and restores the collaboration stored in the session construct formed by the two initiators that have started the session; this allows the initiators to be involved in new sessions. The other rules simply extend the standard parallel, restriction rules to forward and backward relations.

Example 2. Consider the first cherry-pi specification of the VOD scenario given in Ex. 1. In the initial state C_{US} of the collaboration, U and S can synchronise in order to initialise the session, thus evolving to $C_{\text{US}}^1 = (\nu s : C_{\text{US}}) (\langle P_{\text{U}}[\bar{s}/x] \rangle \blacktriangleright P_{\text{U}}[\bar{s}/x] \mid \langle P_{\text{S}}[s/y] \rangle \blacktriangleright P_{\text{S}}[s/y])$.

$$\begin{array}{c}
\frac{\bar{a}(x_1).P_1 \mid a(x_2).P_2 \rightarrow (\nu s : (\bar{a}(x_1).P_1 \mid a(x_2).P_2))}{(\langle P_1[\bar{s}/x_1] \rangle \blacktriangleright P_1[\bar{s}/x_1] \mid \langle P_2[s/x_2] \rangle \blacktriangleright P_2[s/x_2])} \text{ [F-CON]} \\
\\
\frac{P_1 \xrightarrow{\bar{k}! \langle v \rangle} P'_1 \quad P_2 \xrightarrow{k?(x)} P'_2}{\langle Q_1 \rangle \blacktriangleright P_1 \mid \langle Q_2 \rangle \blacktriangleright P_2 \rightarrow \langle Q_1 \rangle \blacktriangleright P'_1 \mid \langle Q_2 \rangle \blacktriangleright P'_2[v/x]} \text{ [F-COM]} \quad \frac{C_1 \rightarrow C'_1}{C_1 \mid C_2 \rightarrow C'_1 \mid C_2} \text{ [F-PAR]} \\
\\
\frac{P_1 \xrightarrow{\bar{k} \leq l} P'_1 \quad P_2 \xrightarrow{k \triangleright l} P'_2}{\langle Q_1 \rangle \blacktriangleright P_1 \mid \langle Q_2 \rangle \blacktriangleright P_2 \rightarrow \langle Q_1 \rangle \blacktriangleright P'_1 \mid \langle Q_2 \rangle \blacktriangleright P'_2} \text{ [F-LAB]} \quad \frac{C_2 \rightarrow C'_2}{(\nu s : C_1) C_2 \rightarrow (\nu s : C_1) C'_2} \text{ [F-RES]} \\
\\
\frac{P_1 \xrightarrow{cmt} P'_1}{\langle Q_1 \rangle \blacktriangleright P_1 \mid \langle Q_2 \rangle \blacktriangleright P_2 \rightarrow \langle P'_1 \rangle \blacktriangleright P'_1 \mid \langle P_2 \rangle \blacktriangleright P_2} \text{ [F-CMT]} \quad \frac{P \xrightarrow{\tau} P'}{\langle Q \rangle \blacktriangleright P \rightarrow \langle Q \rangle \blacktriangleright P'} \text{ [F-IF]} \\
\\
\frac{P_1 \xrightarrow{roll} P'_1}{\langle Q_1 \rangle \blacktriangleright P_1 \mid \langle Q_2 \rangle \blacktriangleright P_2 \rightsquigarrow \langle Q_1 \rangle \blacktriangleright P'_1 \mid \langle Q_2 \rangle \blacktriangleright P_2} \text{ [B-RLL]} \quad \frac{C_1 \rightsquigarrow C'_1}{C_1 \mid C_2 \rightsquigarrow C'_1 \mid C_2} \text{ [B-PAR]} \\
\\
\frac{P_1 \xrightarrow{abt} P'_1}{(\nu s : C)(\langle Q_1 \rangle \blacktriangleright P_1 \mid \langle Q_2 \rangle \blacktriangleright P_2) \rightsquigarrow C} \text{ [B-ABT]} \quad \frac{C_2 \rightsquigarrow C'_2}{(\nu s : C_1) C_2 \rightsquigarrow (\nu s : C_1) C'_2} \text{ [B-RES]}
\end{array}$$

Fig. 5. cherry-pi semantics: forward and backward reduction relations.

Let us consider now a possible run of the session. After three reduction steps, \mathcal{U} executes the commit action, obtaining the following runtime term:

$$\begin{aligned}
C_{\mathcal{US}}^2 &= (\nu s : C_{\mathcal{US}}) (\langle P'_{\mathcal{U}} \rangle \blacktriangleright P'_{\mathcal{U}} \mid \langle P'_S \rangle \blacktriangleright P'_S) \\
P'_{\mathcal{U}} &= \bar{s}?(x_{meta} : \mathbf{str}). \text{if } (f_{eval}(v_{price}, x_{meta})) \text{ then } \dots \quad P'_S = s!\langle f_{meta}(v_{req}) \rangle. y \triangleright \{ \dots \}
\end{aligned}$$

After four further reduction steps, \mathcal{U} chooses the HD video quality and \mathcal{S} commits as well; the resulting runtime collaboration is as follows:

$$\begin{aligned}
C_{\mathcal{US}}^3 &= (\nu s : C_{\mathcal{US}}) (\langle P''_{\mathcal{U}} \rangle \blacktriangleright P''_{\mathcal{U}} \mid \langle P''_S \rangle \blacktriangleright P''_S) \\
P''_{\mathcal{U}} &= \bar{s}?(x_{testHD} : \mathbf{str}). \text{if } (f_{HD}(x_{testHD})) \text{ then } \bar{s}?(x_{videoHD} : \mathbf{str}). \mathbf{0} \text{ else roll} \\
P''_S &= s!\langle f_{testHD}(v_{req}) \rangle. s!\langle f_{videoHD}(v_{req}) \rangle. \mathbf{0}
\end{aligned}$$

In the next reductions, \mathcal{U} evaluates the test video and decides to revert the session execution, resulting in $C_{\mathcal{US}}^4 = (\nu s : C_{\mathcal{US}}) (\langle P''_{\mathcal{U}} \rangle \blacktriangleright \text{roll} \mid \langle P''_S \rangle \blacktriangleright s!\langle f_{videoHD}(v_{req}) \rangle. \mathbf{0})$. The execution of the roll action restores the checkpoints $P''_{\mathcal{U}}$ and P''_S , that is $C_{\mathcal{US}}^4 \rightsquigarrow C_{\mathcal{US}}^3$. After the rollback, \mathcal{U} is not able to change the video quality as her own commit point would have permitted; in fact, it holds $C_{\mathcal{US}}^4 \not\rightsquigarrow C_{\mathcal{US}}^2$.

4 Rollback safety

The operational semantics of **cherry-pi** provides a description of the functioning of the primitives for programming the checkpoint-based rollback recovery in a session-based language. However, as shown in Ex. 2, it does not guarantee high-level properties about

$$\frac{\emptyset; \emptyset \vdash P \triangleright x : T}{\bar{a}(x).P \triangleright \{\bar{a} : T\}} [\text{T-REQ}] \quad \frac{\emptyset; \emptyset \vdash P \triangleright x : T}{a(x).P \triangleright \{a : T\}} [\text{T-ACC}] \quad \frac{C_1 \triangleright A_1 \quad C_2 \triangleright A_2}{C_1 \mid C_2 \triangleright A_1 \cup A_2} [\text{T-PAR}]$$

Fig. 6. Typing system for cherry-pi collaborations.

the safe execution of the rollback. To prevent such undesired rollbacks, we propose the use of *compliance checking*, to be performed at design time. This check is not done on the full system specification, but only at the level of session types.

Session types and typing. The syntax of the cherry-pi *session types* T is defined as follows. Type $![S].T$ represents the behaviour of first outputting a value of sort S (i.e., `bool`, `int` or `str`), then performing the actions prescribed by type T . Type $?[S].T$ is the dual one, where a value is received instead of sent. Types `end` and `err` represent inaction and faulty termination, respectively. Type $\triangleleft[l].T$ represents the behaviour that selects the label l and then behaves as T . Type $\triangleright[l_1 : T_1, \dots, l_n : T_n]$ describes a branching behaviour: it waits for one of the n options to be selected, and behaves as type T_i if the i -th label is selected (external choice). Type $T_1 \oplus T_2$ behaves as either T_1 or T_2 (internal choice). Type $\mu t.T$ represents a recursive behaviour. Type `cmt`. T represents a commit action followed by the actions prescribed by type T . Finally, types `roll` and `abt` represent rollback and abort actions.

The cherry-pi type system does not perform compliance checks, but only infers the types of collaboration participants, which will be then checked together according to the compliance relation. *Typing judgements* are of the form $C \triangleright A$, where A , called *type associations*, is a set of session type associations of the form $\hat{a} : T$, where \hat{a} stands for either \bar{a} or a . Intuitively, $C \triangleright A$ indicates that from the collaboration C the type associations in A are inferred. The definition of the type system for these judgements relies on auxiliary typing judgements for processes, of the form $\Theta; \Gamma \vdash P \triangleright \Delta$, where Θ , Γ and Δ , called *basis*, *sorting* and *typing* respectively, are finite partial maps from process variables to type variables, from variables to sorts, and from variables to types, respectively. Updates of basis and sorting are denoted, respectively, by $\Theta \cdot X : t$ and $\Gamma \cdot y : S$, where $X \notin \text{dom}(\Theta)$, $t \notin \text{cod}(\Theta)$ and $y \notin \text{dom}(\Gamma)$. The judgement $\Theta; \Gamma \vdash P \triangleright \Delta$ stands for “under the environment $\Theta; \Gamma$, process P has typing Δ ”. In its own turn, the typing of processes relies on auxiliary judgments for expressions, of the form $\Gamma \vdash e \triangleright S$. The axioms and rules defining the typing system for cherry-pi collaborations and processes are given in Fig. 6 and 7; typing rules for expressions are standard (see Appendix A). The type system is defined only for initial collaborations, i.e. for terms generated by the grammar in Fig. 2. Other runtime collaborations are not considered here, as no check will be performed at runtime. We comment on salient points. Typing rules at collaboration level simply collect the type associations of session initiators in the collaboration. Rules at process level instead determine the session type corresponding to each process, by mapping each process operator to the corresponding type operator. Data and expression used in communication actions are abstracted as sorts, and a conditional choice is rendered as an internal non-deterministic choice.

Compliance checking. To check compliance between pairs of session parties, we consider *type configurations* of the form $(T, T') : \langle \tilde{T}_1 \rangle \triangleright T_2 \parallel \langle \tilde{T}_3 \rangle \triangleright T_4$, consisting in a

$$\begin{array}{c}
\frac{\Gamma \vdash e \triangleright S \quad \Theta; \Gamma \vdash P \triangleright x : T}{\Theta; \Gamma \vdash x! \langle e \rangle . P \triangleright x : ! [S]. T} \text{ [T-SND]} \quad \frac{\Theta; \Gamma \cdot y : S \vdash P \triangleright x : T}{\Theta; \Gamma \vdash x?(y : S). P \triangleright x : ? [S]. T} \text{ [T-RCV]} \\
\\
\frac{}{\Theta; \Gamma \vdash \mathbf{0} \triangleright x : \mathbf{end}} \text{ [T-INACT]} \quad \frac{\Gamma \vdash e \triangleright \mathbf{bool} \quad \Theta; \Gamma \vdash P_1 \triangleright x : T_1 \quad \Theta; \Gamma \vdash P_2 \triangleright x : T_2}{\Theta; \Gamma \vdash \text{if } e \text{ then } P_1 \text{ else } P_2 \triangleright x : T_1 \oplus T_2} \text{ [T-IF]} \\
\frac{}{\Gamma \cdot x : S \vdash x \triangleright S} \text{ [T-VAR]} \\
\frac{}{\Theta \cdot X : t; \Gamma \vdash X \triangleright t} \text{ [T-PVAR]} \quad \frac{\Theta \cdot X : t; \Gamma \vdash P \triangleright T}{\Theta; \Gamma \vdash \mu X. P \triangleright \mu t. T} \text{ [T-REC]} \\
\frac{}{\Theta; \Gamma \vdash \mathbf{roll} \triangleright x : \mathbf{roll}} \text{ [T-RLL]} \quad \frac{\Theta; \Gamma \vdash P \triangleright x : T}{\Theta; \Gamma \vdash \mathbf{commit}. P \triangleright x : \mathbf{cmt}. T} \text{ [T-CMT]} \\
\frac{}{\Theta; \Gamma \vdash \mathbf{abort} \triangleright x : \mathbf{abt}} \text{ [T-ABT]}
\end{array}$$

Fig. 7. Typing system for cherry-pi processes.

pair (T, T') of session types, corresponding to the types of the parties at the initiation of the session, and in the parallel composition of two pairs $\langle \tilde{T}_c \rangle \triangleright T$, where T is the session type of a party and \tilde{T}_c is the type of the party's checkpoint. We use \tilde{T} to denote either a type T , representing a checkpoint committed by the party, or \underline{T} , representing a checkpoint imposed by the other party. The semantics of type configurations, necessary for the definition of the compliance relation, is given in Fig. 8, where label λ stands for either $! [S]$, $? [S]$, $\triangleleft l$, $\triangleright l$, τ , \mathbf{cmt} , \mathbf{roll} , or \mathbf{abt} . We comment on the relevant rules. In case of a commit action, the checkpoints of both parties are updated, and the one of the passive party (i.e., the party that has not performed the commit) is marked as 'imposed' (rule [TS-CMT₁]). However, if the passive party did not perform any action from its current checkpoint, this checkpoint is not overwritten by the active party (rule [TS-CMT₂]), as discussed in Sec. 2 (Fig. 1(c)). In case of a roll action (rule [TS-RLL₁]), the reduction step is performed only if the active party (i.e., the party that has performed the rollback action) has a non-imposed checkpoint; in all other situations the configuration cannot proceed with the rollback. Finally, in case of abort (rule [TS-ABT₁]), the configuration goes back to the initial state; this allows the type computation to proceed, in order not to affect the compliance check between the two parties.

On top of the above type semantics, we define the compliance relation, inspired by the relation in [3], and prove its decidability (the proof is reported in Appendix C).

Definition 1 (Compliance). *Relation $\dashv\!\!\dashv$ on configurations is defined as follows: $(T, T') : \langle \tilde{U}_1 \rangle \triangleright T_1 \dashv\!\!\dashv \langle \tilde{U}_2 \rangle \triangleright T_2$ holds if for all U'_1, T'_1, U'_2, T'_2 such that $(T, T') : \langle \tilde{U}_1 \rangle \triangleright T_1 \parallel \langle \tilde{U}_2 \rangle \triangleright T_2 \xrightarrow{*} (T, T') : \langle \tilde{U}'_1 \rangle \triangleright T'_1 \parallel \langle \tilde{U}'_2 \rangle \triangleright T'_2 \not\xrightarrow{*}$ we have that $T'_1 = T'_2 = \mathbf{end}$. Two types T_1 and T_2 are compliant, written $T_1 \dashv\!\!\dashv T_2$, if $(T_1, T_2) : \langle \tilde{T}_1 \rangle \triangleright T_1 \dashv\!\!\dashv \langle \tilde{T}_2 \rangle \triangleright T_2$.*

Theorem 1. *Let T_1 and T_2 be two session types, checking if $T_1 \dashv\!\!\dashv T_2$ holds is decidable.*

This compliance relation is used to define the notion of *rollback safety*.

Definition 2 (Rollback safety). *Let C be an initial collaboration, then C is rollback safe (shortened roll-safe) if $C \triangleright A$ and for all pairs $\bar{a} : T_1$ and $a : T_2$ in A we have $T_1 \dashv\!\!\dashv T_2$.*

$$\begin{array}{c}
\text{cmt}.T \xrightarrow{\text{cmt}} T \text{ [TS-CMT]} \quad \text{roll} \xrightarrow{\text{roll}} \text{end} \text{ [TS-RLL]} \quad \text{abt} \xrightarrow{\text{abt}} \text{end} \text{ [TS-ABT]} \\
\\
\frac{T_1 \xrightarrow{\tau} T'_1}{(T, T') : \langle \tilde{U}_1 \rangle \blacktriangleright T_1 \parallel \langle \tilde{U}_2 \rangle \blacktriangleright T_2 \mapsto (T, T') : \langle \tilde{U}_1 \rangle \blacktriangleright T'_1 \parallel \langle \tilde{U}_2 \rangle \blacktriangleright T_2} \text{ [TS-TAU]} \\
\\
\frac{T_1 \xrightarrow{! [S]} T'_1 \quad T_2 \xrightarrow{? [S]} T'_2}{(T, T') : \langle \tilde{U}_1 \rangle \blacktriangleright T_1 \parallel \langle \tilde{U}_2 \rangle \blacktriangleright T_2 \mapsto (T, T') : \langle \tilde{U}_1 \rangle \blacktriangleright T'_1 \parallel \langle \tilde{U}_2 \rangle \blacktriangleright T'_2} \text{ [TS-COM]} \\
\\
\frac{T_1 \xrightarrow{\text{cmt}} T'_1 \quad \tilde{U}_2 \neq T_2}{(T, T') : \langle \tilde{U}_1 \rangle \blacktriangleright T_1 \parallel \langle \tilde{U}_2 \rangle \blacktriangleright T_2 \mapsto (T, T') : \langle \tilde{U}_1 \rangle \blacktriangleright T'_1 \parallel \langle \tilde{U}_2 \rangle \blacktriangleright T_2} \text{ [TS-CMT}_1\text{]} \\
\\
\frac{T_1 \xrightarrow{\text{cmt}} T'_1 \quad \tilde{U}_2 = T_2}{(T, T') : \langle \tilde{U}_1 \rangle \blacktriangleright T_1 \parallel \langle \tilde{U}_2 \rangle \blacktriangleright T_2 \mapsto (T, T') : \langle \tilde{U}_1 \rangle \blacktriangleright T'_1 \parallel \langle \tilde{U}_2 \rangle \blacktriangleright T_2} \text{ [TS-CMT}_2\text{]} \\
\\
\frac{T_1 \xrightarrow{\text{roll}} T'_1}{(T, T') : \langle U_1 \rangle \blacktriangleright T_1 \parallel \langle \tilde{U}_2 \rangle \blacktriangleright T_2 \mapsto (T, T') : \langle U_1 \rangle \blacktriangleright U_1 \parallel \langle \tilde{U}_2 \rangle \blacktriangleright U_2} \text{ [TS-RLL}_1\text{]} \\
\\
\frac{T_1 \xrightarrow{\text{roll}} T'_1}{(T, T') : \langle U_1 \rangle \blacktriangleright T_1 \parallel \langle \tilde{U}_2 \rangle \blacktriangleright T_2 \mapsto (T, T') : \langle U_1 \rangle \blacktriangleright \text{err} \parallel \langle \tilde{U}_2 \rangle \blacktriangleright \text{err}} \text{ [TS-RLL}_2\text{]} \\
\\
\frac{T_1 \xrightarrow{\text{abt}} T'_1}{(T, T') : \langle \tilde{U}_1 \rangle \blacktriangleright T_1 \parallel \langle \tilde{U}_2 \rangle \blacktriangleright T_2 \mapsto (T, T') : \langle T \rangle \blacktriangleright T \parallel \langle T' \rangle \blacktriangleright T'} \text{ [TS-ABT}_1\text{]}
\end{array}$$

Fig. 8. Semantics of types and type configurations (symmetric rules for configurations are omitted).

Example 3. Let us consider again the VOD example. As expected, the first cherry-pi collaboration defined in Ex. 1, corresponding to the scenario described in Fig. 1(b), is not rollback safe, because the types of the two parties are not compliant. Indeed, the session types T_U and T_S associated by the type system to the user and the service processes, respectively, are as follows:

$$\begin{aligned}
T_U &= ![str].?[int].\text{cmt}.?[str].(\triangleleft[l_{HD}].[str].(?[str].\text{end} \oplus \text{roll}) \\
&\quad \oplus \triangleleft[l_{SD}].[str].(?[str].\text{end} \oplus \text{abt})) \\
T_S &= ?[str].![int].![str].\triangleright[l_{HD}:\text{cmt}.![str].![str].\text{end}, \\
&\quad l_{SD}:\text{cmt}.![str].![str].\text{end}]
\end{aligned}$$

Thus, the resulting initial configuration is $(T_U, T_S) : \langle T_U \rangle \blacktriangleright T_U \parallel \langle T_S \rangle \blacktriangleright T_S$, which can evolve to the configuration $(T_U, T_S) : \langle \underline{T} \rangle \blacktriangleright \text{roll} \parallel \langle U \rangle \blacktriangleright ![str].\text{end}$, with $T = ?[str].(?[str].\text{end} \oplus \text{roll})$ and $U = ![str].![str].\text{end}$. This configuration evolves to $(T_U, T_S) : \langle \underline{T} \rangle \blacktriangleright \text{err} \parallel \langle U \rangle \blacktriangleright \text{err}$, which cannot further evolve and is not in a completed state (in fact, type err is different from end), meaning that T_U and T_S are not compliant.

In the scenario described in Fig. 1(c), instead, the type of the server process is as follows: $T'_S = ?[str].![int].\text{cmt}.![str].\triangleright[l_{HD}:\text{cmt}.![str].![str].\text{end}, l_{SD}:\text{cmt}.![str].![str].\text{end}]$ and we have $T_U \dashv\!\!\dashv T'_S$. Finally, the types of the processes depicted in

Fig. 1(d) are:

$$T'_U = ![str].?[int].?[str].cmt.(\triangleleft[l_{HD}].\dots \oplus \triangleleft[l_{SD}].\dots)$$

$$T''_S = ?[str].![int].![str].cmt.\triangleright[l_{HD}:![str].![str].end, l_{SD}:![str].![str].end]$$

and we have $T'_U \not\vdash T''_S$. Indeed, the corresponding initial configuration can evolve to the configuration $(T'_U, T''_S) : \langle \triangleleft[l_{HD}].\dots \rangle \triangleright roll \parallel \langle \triangleright[l_{HD} : \dots, l_{SD} : \dots] \rangle \triangleright ![str].end$, which again evolves to a configuration that is not in a completed state.

MAUDE implementation. To show the feasibility of our approach, we have implemented the semantics of type configurations in Fig. 8 in the MAUDE framework [8]. MAUDE provides an instantiation of rewriting logic [23] and it has been used to implement the semantics of several formal languages [24].

The syntax of *cherry-pi* types and type configurations is specified by defining algebraic data types, while transitions and reductions are rendered as rewrites and, hence, inference rules are given in terms of (conditional) rewrite rules. Since MAUDE specifications are executable, we have obtained in this way an interpreter for *cherry-pi* type configurations, which permits to explore the reductions arising from the initial configuration of two given session types.

Our implementation consists of two MAUDE modules. The CHERRY-TYPES-SYNTAX module provides the definition of the sorts that characterise the syntax of *cherry-pi* types, such as session types, selection/branching labels, type variables and type configurations. In particular, basic terms of session types are rendered as constant *operations* on the sort *Type*; e.g., the roll type is defined as

```
op roll : -> Type .
```

The other syntactic operators are instead defined as operations with one or more arguments; e.g., the output type takes as input a *Sort* and a continuation type:

```
op ![_]._ : Sort Type -> Type [frozen prec 25] .
```

To prevent undesired rewrites inside operator arguments, following the approach in [35], we have declared these operations as *frozen*. The *prec* attribute has been used to define the precedence among operators. The CHERRY-TYPES-SEMANTICS module provides *rewrite rules*, and additional operators and equations, to define the *cherry-pi* type semantics. For example, the operational rule [TS-SND] is rendered as follows:

```
rl [TS-Snd] : ![S].T => {![S]}T .
```

The correspondence between the operational rule and the rewrite rule is one-to-one; the only peculiarity is the fact that, since rewrites have no labels, we have made the transition label part of the resulting term. Reduction rules for type configurations are instead rendered in terms of conditional rewrite rules with rewrites in their conditions. For example, the [TS-COM] rule is rendered as:

```
cr1 [TS-Com] :
  init(T,T') CT1 > T1 || CT2 > T2 => init(T,T') CT1 > T1' || CT2 > T2'
  if T1 => {![S]}T1' /\ T2 => {?[S]}T2' .
```

Again, there is a close correspondence between the operational rule and the rewrite one.

The compliance check between two session types can be then conveniently realised on top of the implementation described above by resorting to the MAUDE command `search`. This permits indeed to explore the state space of the configurations reachable from an initial configuration. Specifically, the compliance check between types $T1$ and $T2$ is rendered as follows:

```
search
  init(T1,T2) ckp(T1) > T1 || ckp(T2) > T2
  =>!
  init(T:Type,T':Type) CT1:CkpType > T1':Type || CT2:CkpType > T2':Type
  such that T1' /= end or T2' /= end .
```

This command searches for all terminal states ($=>!$), i.e. states that cannot be rewritten any more (see $\not\rightarrow$ in Def. 1), and checks if at least one of the two session types in the corresponding configurations ($T1'$ and $T2'$) is different from the `end` type. Thus, if this search has no solution, $T1$ and $T2$ are compliant; otherwise, they are not compliant and a violating configuration is returned.

Example 4. Let us consider the `cherry-pi` types defined in Ex. 3 for the scenario described in Fig. 1(b). In our MAUDE implementation of the type syntax, the session types T_U and T_S , and the corresponding initial type configuration, are rendered as follows:

```
eq Tuser = ![str]. ?[int]. cmt. ?[str].
           ((sel['hd']. ?[str]. ((?[str]. end) (+) roll))
            (+) (sel['sd']. ?[str]. ((?[str]. end) (+) abt))) .

eq Tservice = ?[str]. ![int]. ![str].
              brn[brnEl('hd, cmt. ![str]. ![str]. end);
              brnEl('sd, cmt. ![str]. ![str]. end)] .

eq InitConfig = init(Tuser,Tservice)
               ckp(Tuser) > Tuser || ckp(Tservice) > Tservice .
```

where `(+)` represents the internal choice operator, `sel` the selection operator, `brn` the branching operator, `brnEl` an option offered in a branching, and `ckp` a non-imposed checkpoint. The compliance between the two session types can be checked by loading the two modules of our MAUDE implementation, and executing the following command:

```
search InitConfig
  =>!
  init(T:Type,T':Type) CT1:CkpType > T1:Type || CT2:CkpType > T2:Type
  such that T1 /= end or T2 /= end .
```

This search command returns the following solution:

```
CT1 --> ickp(?[str]. ((?[str]. end) (+) roll))
T1 --> err
CT2 --> ckp(![str]. ![str]. end)
T2 --> err
```

As explained in Ex. 3, the two types are not compliant. Indeed, the configuration above is a terminal state, and $T1$ and $T2$ are clearly different from `end`.

The scenario in Fig. 1(c) is rendered by the following implementation of the service type:

```

eq Tservice' = ?[str]. ![int]. cmt. ![str].
           brn[brnEl('hd, ![str]. ![str]. end);
           brnEl('sd, ![str]. ![str]. end)] .

```

In this case, as expected, the `search` command returns:

```

No solution.

```

meaning that types `Tuser` and `Tservice'` are compliant. Finally, the `search` command applied to the type configuration related to the scenario depicted in Fig. 1(d) returns a solution, meaning that in that case the user and service types are not compliant.

5 Properties of cherry-pi

This section presents the results regarding the properties of `cherry-pi`. Missing proofs are reported in Appendix C. The statement of some properties exploits labelled transitions that permit to easily distinguish the execution of commit and rollback actions from the other ones. To this end, we can instrument the reduction semantics of collaborations by means of labels of the form *cmt s*, *roll s* and *abt s*, indicating the rule used to derive the reduction and the session on which such operation has been done.

Rollback properties. We show some properties concerning the reversible behaviour of `cherry-pi` related to the interplay between rollback and commit primitives. The first two properties, namely Theorem 2 and Lemma 1, are an adaptation of typical properties of reversible calculi, while Lemma 2 and Lemma 3 are brand new.

The following theorem states that any reachable collaboration is also a *forward only* reachable collaboration. This means that all the states a collaboration reaches via mixed executions (also involving backward reductions) are states that we can reach from the initial configuration with just forward reductions. This assures us that if the system goes back it will reach previous visited states.

Theorem 2. *Let C_0 be an initial collaboration. If $C_0 \rightarrow^* C_1$ then $C_0 \rightarrow^* C_1$.*

We now show a variant of the so-called Loop Lemma [10]. In a fully reversible calculus this lemma states that each computational step, either forward or backward, can be undone. Since reversibility in `cherry-pi` is controlled, we have to state that if a reversible step is possible (e.g., a rollback is *enabled*) then the effects of the rollback can be undone.

Lemma 1 (Safe rollback). *Let C_1 and C_2 be reachable collaborations. If $C_1 \rightsquigarrow C_2$ then $C_2 \rightarrow^* C_1$.*

A rollback always brings the system to the last taken checkpoint. We recall that, since there may be sessions running in parallel, a collaboration may be able to do different rollbacks within different sessions. Thus, determinism only holds relative to a given session, and rollback within one session has no effect on any other parallel session.

Lemma 2 (Determinism). *Let C be a reachable collaboration. If $C \rightsquigarrow^{roll s} C'$ and $C \rightsquigarrow^{roll s} C''$ then $C' \equiv C''$.*

$$\begin{array}{c}
\frac{P_1 \xrightarrow{k! \langle v \rangle} P'_1 \quad \neg P_2 \Downarrow_{k?} \quad \neg P_2 \Downarrow_{roll} \quad [\text{E-COM1}]}{\langle \tilde{Q}_1 \rangle \blacktriangleright P_1 \mid \langle \tilde{Q}_2 \rangle \blacktriangleright P_2 \rightarrow \text{com_error}} \quad \frac{P_1 \xrightarrow{roll} P'_1 \quad [\text{E-RLL2}]}{\langle Q_1 \rangle \blacktriangleright P_1 \mid \langle \tilde{Q}_2 \rangle \blacktriangleright P_2 \rightarrow \text{roll_error}}
\end{array}$$

Fig. 9. cherry-pi semantics: error reductions.

The last rollback property states that a collaboration cannot go back to a state prior to the execution of a commit action, that is commits have a persistent effect. Let us note that recursion does not affect this theorem, since at the beginning of a collaboration computation there is always a new session establishment, leading to a stack of past configurations. Hence it is never the case that from a collaboration C you can reach again C via forward steps.

Theorem 3 (Commit persistency). *Let C be a reachable collaboration. If $C \xrightarrow{cmts} C'$ then there exists no C'' such that $C' \rightarrow^* \rightsquigarrow^{rolls} C''$ and $C'' \rightarrow^+ C$.*

Soundness properties. The second group of properties concerns soundness guarantees. The definition of these properties requires formally characterising the errors that may occur in the execution of an unsound collaboration. We rely on error reduction (as in [7]) rather than on the usual static characterisation of errors (as, e.g., in [38]), since rollback errors cannot be easily detected statically. In particular, we extend the syntax of cherry-pi collaborations with the `roll_error` and `com_error` terms, denoting respectively collaborations in rollback and communication error states:

$$C ::= \dots \mid \langle \tilde{P}_1 \rangle \blacktriangleright P_2 \mid \text{roll_error} \mid \text{com_error}$$

where \tilde{P} denotes either a checkpoint P committed by the party or a checkpoint \underline{P} imposed by the other party of the session. The semantics of cherry-pi is extended as well by the (excerpt) of error reduction rules in Fig. 9. The error semantics does not affect the normal behaviour of cherry-pi specifications, but it is crucial for stating our soundness theorems. Its definition is based on the notion of *barb* predicate: $P \Downarrow_\mu$ holds if there exists P' such that $P \Rightarrow P'$ and P' can perform an action μ , where μ stands for $k?$, $k!$, $k \triangleleft l$, $k \triangleright l$, or *roll* (i.e., input, output, select, branching action along session channel k , or roll action); \Rightarrow is the reflexive and transitive closure of $\xrightarrow{\tau}$. The meaning of the error semantics rules is as follows. A *communication error* takes place in a collaboration when a session participant is willing to perform an output but the other participant is ready to perform neither the corresponding input nor a roll back (rule [E-COM₁]) or vice versa, or one participant is willing to perform a selection but the corresponding branching is not available on the other side or viceversa. Instead, a *rollback error* takes place in a collaboration when a participant is willing to perform a rollback action but her checkpoint has been imposed by the other participant ([E-RLL₂]). To enable this error check, the rules for commit and rollback have been modified to keep track of imposed overwriting of checkpoints. This information is not relevant for the runtime execution of processes, but it is necessary for characterising the rollback errors that our type-based approach prevents.

Besides defining the error semantics, we also need to define erroneous collaborations, based on the following notion of context: $\mathbb{C} ::= [\cdot] \mid \mathbb{C} \mid C \mid (\nu s : C) \mathbb{C}$.

Definition 3 (Erroneous collaborations). A collaboration C is communication (resp. rollback) erroneous if $C = \mathbb{C}[\text{com_error}]$ (resp. $C = \mathbb{C}[\text{roll_error}]$).

The key soundness results follow: a rollback safe collaboration never reduces to either a rollback erroneous collaboration (Theorem 4) or a communication erroneous collaboration (Theorem 5).

Theorem 4 (Rollback soundness). If C is a roll-safe collaboration, then we have that $C \not\rightarrow^* \mathbb{C}[\text{roll_error}]$.

Theorem 5 (Communication soundness). If C is a roll-safe collaboration, then we have that $C \not\rightarrow^* \mathbb{C}[\text{com_error}]$.

We conclude with a progress property of cherry-pi sessions: given a rollback safe collaboration that can initiate a session, each collaboration reachable from it either is able to progress on the session with a forward/backward reduction step or has correctly reached the end of the session. This result follows from Theorems 4 and 5, and from the fact that we consider binary sessions without delegation and subordinate sessions.

Theorem 6 (Session progress). Let $C = (\bar{a}(x_1).P_1 \mid a(x_2).P_2)$ be a roll-safe collaboration. If $C \rightarrow^* C'$ then either $C' \rightarrow C''$ for some C'' or $C' \equiv (\nu s:C)(\langle \bar{Q}_1 \rangle \bullet \mathbf{0} \mid \langle \bar{Q}_2 \rangle \bullet \mathbf{0})$ for some \bar{Q}_1 and \bar{Q}_2 .

6 Conclusion and related work

This paper proposes rollback recovery primitives for session-based programming. These primitives come together with session typing, enabling a design time compliance check which ensures checkpoint persistency properties (Lemma 1 and Theorem 3) and session soundness (Theorems 4 and 5). Our compliance check has been implemented in MAUDE.

In the literature we can distinguish two ways of dealing with rollback: either using explicit rollbacks and implicit commits [21], or by using explicit commits and spontaneous aborts [11,36]. Differently from these works, we have introduced a way to control reversibility by both *triggering* it and *limiting* its scope. Reversibility is triggered via an explicit rollback primitive (as in [21]), while explicit commits limit the scope of potential future reverse executions (as in [11,36]). Differently from [11,36], commit does not require any synchronisation, as it is a local decision. This could lead to run-time misbehaviours where a process willing to roll back to its last checkpoint reaches a point which has been imposed by another participant of the session. Our type discipline rules out such cases.

Reversibility in behavioural types has been studied in different formalisms: *contracts* [2,4], *binary session types* [25], *multiparty session types* [26,6,32,33], and *global graphs* [28,14]. In [2,4] choices can be seen as *implicit* checkpoints and the system can go back to a previous choice and try another branch. In [2] rollback is triggered non-deterministically, while in [4] it is triggered by the system only when the computation is stuck. In both works reversibility (and rollbacks) is used to achieve a relaxed variant of client-server compliance: if there exists an execution in which the client is able to terminate then the client and server are compliant. Hence, reversibility is used as a means

to explore different branches if the current one leads to a deadlock. In [25] reversibility is studied in the context of binary session types. Types information is used at run-time by monitors, for binary [25] and multiparty [26] settings, to keep track of the computational history of the system, allowing to revert any computational step, where global types are enriched with computational history. There, reversibility is uncontrolled, and each computational step can be undone. In [6] global types are enriched with history information, and choices are seen as labelled checkpoints. The information about checkpoints is projected into local types. At any moment, the party who decided which branch to take in a choice may decide to revert it, forcing the entire system to go back to a point prior to the choice. Hence, rollback is confined inside choices and it is spontaneous, meaning that the former can be programmed while the latter cannot. Checkpoints are not seen as commits, and a rollback can bring the system to a state prior to several checkpoints. In [32] an *uncontrolled* reversible variant of session π -calculus is presented, while [33] studies different notions of reversibility for both binary and multiparty single sessions. In [28,14] global graphs are extended with conditions on branches. These conditions at runtime can trigger coordinated rollbacks to revert a distributed choice. Reversibility is confined into branches of a distributed choice and not all the computational steps are reversible; inputs, in fact, are irreversible unless they are inside an ongoing loop. To trigger a rollback several conditions and constraints about loops have to be satisfied. Hence, in order to trigger a rollback a runtime condition should be satisfied.

We detach from these works in several ways. Our checkpoint facility is explicit and checkpointing is not relegated to choices: the programmer can decide at any point when to commit. This is because the programmer may be interested in committing, besides choice points, a series of interactions (e.g., to make a payment irreversible). Once a commit is taken, the system cannot revert to a state prior to it. Our rollback is explicit, meaning that it is the programmer who deliberately triggers a rollback. The extension to the multiparty setting is natural and does not rely on a formalism to describe the global view of the system. Our compliance check, which is decidable, resembles those of [2,3,4], which are defined for different rollback recovery approaches based on implicit checkpoints.

As future work, we plan to extend our approach to deal with sessions where parties can interleave interactions performed along different sessions. This requires to deal with subordinate sessions, which may affect enclosing sessions by performing, e.g., commit actions that make some interaction of the enclosing sessions irreversible, similarly to nested transactions [37]. To tackle this issue it would be necessary to extend the notion of compliance relation to take into account possible partial commits (in case of nested sub-sessions) that could be undone at the top level if a rollback is performed. Also, the way our checkpoints are taken resembles the Communication Induced Checkpoints (CIC) approach [12]; we leave as future work a thoughtful comparison between these two mechanisms.

References

1. Aman, B., Ciobanu, G., Glück, R., Kaarsgaard, R., Kari, J., Kutrib, M., Lanese, I., Mezzina, C.A., Mikulski, L., Nagarajan, R., Phillips, I.C.C., Pinna, G.M., Prigioniero, L., Ulidowski, I.,

- Vidal, G.: Foundations of reversible computation. In: Ulidowski, I., Lanese, I., Schultz, U.P., Ferreira, C. (eds.) *Reversible Computation: Extending Horizons of Computing - Selected Results of the COST Action IC1405, Lecture Notes in Computer Science*, vol. 12070, pp. 1–40. Springer (2020)
2. Barbanera, F., Dezani-Ciancaglini, M., de'Liguoro, U.: Compliance for reversible client/server interactions. In: *BEAT. EPTCS*, vol. 162, pp. 35–42 (2014)
 3. Barbanera, F., Dezani-Ciancaglini, M., Lanese, I., de'Liguoro, U.: Retractable contracts. In: *PLACES 2015. EPTCS*, vol. 203, pp. 61–72 (2016)
 4. Barbanera, F., Lanese, I., de'Liguoro, U.: Retractable and speculative contracts. In: *COORDINATION. LNCS*, vol. 10319, pp. 119–137. Springer (2017)
 5. Barbanera, F., Lanese, I., de'Liguoro, U.: A theory of retractable and speculative contracts. *Sci. Comput. Program.* **167**, 25–50 (2018)
 6. Castellani, I., Dezani-Ciancaglini, M., Giannini, P.: Concurrent reversible sessions. In: *CONCUR. LIPIcs*, vol. 85, pp. 30:1–30:17. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2017)
 7. Chen, T., Dezani-Ciancaglini, M., Scalas, A., Yoshida, N.: On the preciseness of subtyping in session types. *Logical Methods in Computer Science* **13**(2) (2017)
 8. Clavel, M., Durán, F., Eker, S., Lincoln, P., Martí-Oliet, N., Meseguer, J., Talcott, C.L. (eds.): *All About Maude - A High-Performance Logical Framework, How to Specify, Program and Verify Systems in Rewriting Logic*, LNCS, vol. 4350. Springer (2007)
 9. Cristescu, I., Krivine, J., Varacca, D.: A Compositional Semantics for the Reversible π -Calculus. In: *LICS*. pp. 388–397. IEEE (2013)
 10. Danos, V., Krivine, J.: Reversible communicating systems. In: *CONCUR. LNCS*, vol. 3170, pp. 292–307. Springer (2004)
 11. Danos, V., Krivine, J.: Transactions in rccs. In: *CONCUR. LNCS*, vol. 3653, pp. 398–412. Springer (2005)
 12. Elnozahy, E.N., Alvisi, L., Wang, Y., Johnson, D.B.: A survey of rollback-recovery protocols in message-passing systems. *ACM Comput. Surv.* **34**(3), 375–408 (2002)
 13. Engblom, J.: A review of reverse debugging. In: *System, Software, SoC and Silicon Debug Conference (S4D)*. pp. 1–6 (Sept 2012)
 14. Francalanza, A., Mezzina, C.A., Tuosto, E.: Reversible choreographies via monitoring in erlang. In: Bonomi, S., Rivi re, E. (eds.) *Distributed Applications and Interoperable Systems - 18th IFIP WG 6.1 International Conference, DAIS 2018. Lecture Notes in Computer Science*, vol. 10853, pp. 75–92. Springer (2018)
 15. Giachino, E., Lanese, I., Mezzina, C.A., Tiezzi, F.: Causal-consistent rollback in a tuple-based language. *J. Log. Algebr. Meth. Program.* **88**, 99–120 (2017)
 16. Honda, K., Vasconcelos, V.T., Kubo, M.: Language primitives and type discipline for structured communication-based programming. In: *ESOP. LNCS*, vol. 1381, pp. 122–138. Springer (1998)
 17. H ttel, H., Lanese, I., Vasconcelos, V.T., Caires, L., Carbone, M., Deni lou, P., Mostrous, D., Padovani, L., Ravara, A., Tuosto, E., Vieira, H.T., Zavattaro, G.: Foundations of session types and behavioural contracts. *ACM Comput. Surv.* **49**(1), 3:1–3:36 (2016)
 18. Kouzapas, D., Yoshida, N.: Globally governed session semantics. *Logical Methods in Computer Science* **10**(4) (2014)
 19. Kuhn, S., Ulidowski, I.: Local reversibility in a calculus of covalent bonding. *Sci. Comput. Program.* **151**, 18–47 (2018)
 20. Lanese, I., Lienhardt, M., Mezzina, C.A., Schmitt, A., Stefani, J.: Concurrent flexible reversibility. In: *ESOP. LNCS*, vol. 7792, pp. 370–390. Springer (2013)
 21. Lanese, I., Mezzina, C.A., Schmitt, A., Stefani, J.B.: Controlling Reversibility in Higher-Order Π . In: *CONCUR. LNCS*, vol. 6901, pp. 297–311. Springer (2011)

22. Lanese, I., Mezzina, C.A., Stefani, J.B.: Reversing higher-order pi. In: CONCUR. LNCS, vol. 6269, pp. 478–493. Springer (2010)
23. Meseguer, J.: Conditional rewriting logic as a unified model of concurrency. *Theoretical Computer Science* **96**(1), 73–155 (1992)
24. Meseguer, J.: Twenty years of rewriting logic. *J. Log. Algebr. Program.* **81**(7-8), 721–781 (2012)
25. Mezzina, C.A., Pérez, J.A.: Reversibility in session-based concurrency: A fresh look. *J. Log. Algebr. Meth. Program.* **90**, 2–30 (2017)
26. Mezzina, C.A., Pérez, J.A.: Causal consistency for reversible multiparty protocols. *Log. Methods Comput. Sci.* **17**(4) (2021)
27. Mezzina, C.A., Schlatte, R., Glück, R., Haulund, T., Hoey, J., Cservenka, M.H., Lanese, I., Mogensen, T.Æ., Siljak, H., Schultz, U.P., Ulidowski, I.: Software and reversible systems: A survey of recent activities. In: *Reversible Computation: Extending Horizons of Computing - Selected Results of the COST Action IC1405*, Lecture Notes in Computer Science, vol. 12070, pp. 41–59. Springer (2020)
28. Mezzina, C.A., Tuosto, E.: Choreographies for automatic recovery. CoRR **abs/1705.09525** (2017), <http://arxiv.org/abs/1705.09525>
29. Milner, R.: *Communication and concurrency*. Prentice-Hall (1989)
30. Perumalla, K.S., Protopopescu, V.A.: Reversible simulations of elastic collisions. *ACM Trans. Model. Comput. Simul.* **23**(2), 12:1–12:25 (2013)
31. Prabhu, P., Ramalingam, G., Vaswani, K.: Safe programmable speculative parallelism. In: Zorn, B.G., Aiken, A. (eds.) PLDI. pp. 50–61. ACM (2010)
32. Tiezzi, F., Yoshida, N.: Reversible session-based pi-calculus. *J. Log. Algebr. Meth. Program.* **84**(5), 684–707 (2015)
33. Tiezzi, F., Yoshida, N.: Reversing single sessions. In: RC. LNCS, vol. 9720, pp. 52–69. Springer (2016)
34. Vassor, M., Stefani, J.: Checkpoint/rollback vs causally-consistent reversibility. In: *Reversible Computation - 10th International Conference, RC 2018*, Lecture Notes in Computer Science, vol. 11106, pp. 286–303. Springer (2018)
35. Verdejo, A., Martí-Oliet, N.: Implementing CCS in Maude 2. In: WRLA. ENTCS, vol. 71, pp. 239–257. Elsevier (2002)
36. de Vries, E., Koutavas, V., Hennessy, M.: Communicating transactions - (extended abstract). In: Gastin, P., Laroussinie, F. (eds.) CONCUR. LNCS, vol. 6269. Springer (2010)
37. Weikum, G., Schek, H.J.: Concepts and applications of multilevel transactions and open nested transactions. In: *Database Transaction Models for Advanced Applications*. pp. 515–553. Morgan Kaufmann (1992)
38. Yoshida, N., Vasconcelos, V.T.: Language Primitives and Type Discipline for Structured Communication-Based Programming Revisited: Two Systems for Higher-Order Session Communication. *Electr. Notes Theor. Comp. Sci.* **171**(4), 73–93 (2007)

A Omitted Rules

For reviewers' convenience, we report in this appendix those rules that have been omitted in the paper.

The omitted operational rules of the auxiliary labelled relation $\xrightarrow{\ell}$ are in Fig. 10.

$$\frac{P_1 \equiv P'_1 \xrightarrow{\ell} P'_2 \equiv P_2 \quad [\text{P-STR}] \quad \text{if } e \text{ then } P_1 \text{ else } P_2 \xrightarrow{\tau} P_2 \quad (e \downarrow \text{false}) \quad [\text{P-IF}]}{P_1 \xrightarrow{\ell} P_2}$$

Fig. 10. cherry-pi semantics: auxiliary labelled relation (omitted rules).

The omitted rules defining the typing system for cherry-pi processes and expressions are given in Fig. 11 and 12.

$$\frac{\Theta; \Gamma \vdash P \blacktriangleright x : T}{\Theta; \Gamma \vdash x \triangleleft l.P \blacktriangleright x : \triangleleft[l].T} \quad [\text{T-SEL}]$$

$$\frac{\Theta; \Gamma \vdash P_1 \blacktriangleright x : T_1 \quad \dots \quad \Theta; \Gamma \vdash P_n \blacktriangleright x : T_n}{\Theta; \Gamma \vdash x \triangleright \{l_1 : P_1, \dots, l_n : P_n\} \blacktriangleright x : \triangleright[l_1 : T_1, \dots, l_n : T_n]} \quad [\text{T-BR}]$$

Fig. 11. Typing system for cherry-pi processes (omitted rules).

$$\Gamma \vdash \text{true} \blacktriangleright \text{bool} \quad [\text{T-BOOL}_{tt}] \quad \Gamma \vdash \text{false} \blacktriangleright \text{bool} \quad [\text{T-BOOL}_{ff}]$$

$$\Gamma \vdash x : S \vdash x \blacktriangleright S \quad [\text{T-VAR}] \quad \Gamma \vdash 1 \blacktriangleright \text{int} \quad [\text{T-INT}] \quad \Gamma \vdash "a" \blacktriangleright \text{str} \quad [\text{T-STR}]$$

$$\frac{\Gamma \vdash e_1 \blacktriangleright \text{int} \quad \Gamma \vdash e_2 \blacktriangleright \text{int}}{\Gamma \vdash +(e_1, e_2) \blacktriangleright \text{int}} \quad [\text{T-SUM}] \quad \frac{\Gamma \vdash e_1 \blacktriangleright \text{bool} \quad \Gamma \vdash e_2 \blacktriangleright \text{bool}}{\Gamma \vdash \wedge(e_1, e_2) \blacktriangleright \text{bool}} \quad [\text{T-AND}]$$

Fig. 12. Typing system for cherry-pi expressions (excerpt of omitted rules).

The syntax of the cherry-pi *session types* is defined in Fig. 13.

| | |
|--|----------------------|
| $S ::= \text{bool} \mid \text{int} \mid \text{str}$ | Sorts |
| $T ::= ![S].T \mid ?[S].T \mid \triangleleft[l].T \mid \triangleright[l_1 : T_1, \dots, l_n : T_n]$ | Session types |
| $\mid T_1 \oplus T_2 \mid t \mid \mu t.T \mid \text{end} \mid \text{err} \mid \text{cmt}.T \mid \text{roll} \mid \text{abt}$ | |

Fig. 13. cherry-pi type syntax.

The omitted rules defining the semantics of types and type configurations are given in Fig. 14 (symmetric rules for configurations are omitted).

The omitted rules defining the error semantics are given in Fig. 15. The semantics of cherry-pi is extended by the additional rules, where [E-CMT₁] and [E-CMT₂] replace

$$\begin{array}{c}
! [S].T \xrightarrow{! [S]} T \text{ [TS-SND]} \quad ? [S].T \xrightarrow{? [S]} T \text{ [TS-Rcv]} \quad \triangleleft [l].T \xrightarrow{\triangleleft l} T \text{ [TS-SEL]} \\
\triangleright [l_1 : T_1, \dots, l_n : T_n] \xrightarrow{\triangleright l_i} T_i \quad (1 \leq i \leq n) \text{ [TS-BR]} \quad \frac{T[\mu t.T/t] \xrightarrow{\lambda} T'}{\mu t.T \xrightarrow{\lambda} T'} \text{ [TS-REC]} \\
T_1 \oplus T_2 \xrightarrow{\tau} T_1 \text{ [TS-IF}_1\text{]} \quad T_1 \oplus T_2 \xrightarrow{\tau} T_2 \text{ [TS-IF}_2\text{]} \\
\frac{T_1 \xrightarrow{\triangleleft l} T'_1 \quad T_2 \xrightarrow{\triangleright l} T'_2}{(T, T') : \langle \tilde{U}_1 \rangle \blacktriangleright T_1 \parallel \langle \tilde{U}_2 \rangle \blacktriangleright T_2 \mapsto (T, T') : \langle \tilde{U}_1 \rangle \blacktriangleright T'_1 \parallel \langle \tilde{U}_2 \rangle \blacktriangleright T'_2} \text{ [TS-LAB]}
\end{array}$$

Fig. 14. Semantics of types and type configurations (omitted rules).

$$\begin{array}{c}
\frac{P_1 \xrightarrow{k?(x)} P'_1 \quad \neg P_2 \Downarrow_{k!} \quad \neg P_2 \Downarrow_{roll} \text{ [E-COM2]}}{\langle \tilde{Q}_1 \rangle \blacktriangleright P_1 \mid \langle \tilde{Q}_2 \rangle \blacktriangleright P_2 \twoheadrightarrow \text{com_error}} \\
\frac{P_1 \xrightarrow{k \triangleleft l} P'_1 \quad \neg P_2 \Downarrow_{k \triangleright l} \quad \neg P_2 \Downarrow_{roll} \text{ [E-LAB}_1\text{]}}{\langle \tilde{Q}_1 \rangle \blacktriangleright P_1 \mid \langle \tilde{Q}_2 \rangle \blacktriangleright P_2 \twoheadrightarrow \text{com_error}} \quad \frac{P_1 \xrightarrow{k \triangleright l} P'_1 \quad \neg P_2 \Downarrow_{k \triangleleft l} \quad \neg P_2 \Downarrow_{roll} \text{ [E-LAB}_2\text{]}}{\langle \tilde{Q}_1 \rangle \blacktriangleright P_1 \mid \langle \tilde{Q}_2 \rangle \blacktriangleright P_2 \twoheadrightarrow \text{com_error}} \\
\frac{P_1 \xrightarrow{cmt} P'_1 \quad \tilde{Q}_2 \neq P_2}{\langle \tilde{Q}_1 \rangle \blacktriangleright P_1 \mid \langle \tilde{Q}_2 \rangle \blacktriangleright P_2 \twoheadrightarrow \langle P'_1 \rangle \blacktriangleright P'_1 \mid \langle \tilde{Q}_2 \rangle \blacktriangleright P_2} \text{ [E-CMT}_1\text{]} \\
\frac{P_1 \xrightarrow{cmt} P'_1 \quad \tilde{Q}_2 = P_2}{\langle \tilde{Q}_1 \rangle \blacktriangleright P_1 \mid \langle \tilde{Q}_2 \rangle \blacktriangleright P_2 \twoheadrightarrow \langle P'_1 \rangle \blacktriangleright P'_1 \mid \langle \tilde{Q}_2 \rangle \blacktriangleright P_2} \text{ [E-CMT}_2\text{]} \\
\frac{P_1 \xrightarrow{roll} P'_1}{\langle \tilde{Q}_1 \rangle \blacktriangleright P_1 \mid \langle \tilde{Q}_2 \rangle \blacktriangleright P_2 \rightsquigarrow \langle Q_1 \rangle \blacktriangleright Q_1 \mid \langle \tilde{Q}_2 \rangle \blacktriangleright Q_2} \text{ [E-RLL}_1\text{]}
\end{array}$$

Fig. 15. cherry-pi semantics: error reductions (omitted rules).

[F-CMT], and [E-RLL₁] replaces [B-RLL], and \sim is used in the checkpoints of the other rules.

| | |
|--|-----------------------|
| $C ::= \dots \mid \bar{a}[p](x).P \mid a[p](x).P$ | Collaborations |
| $P ::= \dots \mid r[p]!\langle e \rangle.P \mid r[p]?(y : S).P \mid r[p] \triangleleft l.P \mid r[p] \triangleright \{l_1 : P_1, \dots, l_n : P_n\}$ | Processes |

Fig. 16. Multiparty cherry- π runtime syntax (the omitted parts are as in Fig. 3).

B Extension to multiparty sessions

We show in this Appendix how to extend cherry- π , its type discipline, the compliance checking and the related results, to multiparty sessions [18].

The base sets for the multiparty syntax of cherry- π are the same of the binary case, except for *session endpoints*, which now are denoted by $s[p]$, with p, q ranging over *roles* (represented as natural numbers). Thus, *session identifiers* r now range over session endpoints $s[p]$ or variables x . The runtime syntax of multiparty cherry- π is defined by the grammar in Fig. 16, where expressions e are defined as in the binary case (with values that extends to multiparty session endpoints). Primitive $\bar{a}[p](x).P$ initiates a new session through identifier a on the other multiple participants, each one of the form $a[q](x).P_q$ where $1 \leq q \leq p - 1$. Variable x will be substituted with the session endpoint used for the interactions inside the established session. Primitive $r[p]!\langle e \rangle.P$ denotes the intention of sending a value to role p ; similarly, process $r[p]?(y : S).P$ denotes the intention of receiving a value from role p . Selection and branching are extended in a similar way.

As usual the operational semantics is given in terms of a structural congruence and of a reduction relation. The rules defining the structural congruence are standard, while the forward and backward reduction relations are given by the rules in Fig. 17. We comment on salient points. Rule [M-F-CON] synchronously initiates a session by requiring all session endpoints be present for a forward reduction, where each role p creates a session endpoint $s[p]$ on a fresh session channel s . The participant with the maximum role is responsible for requesting a session initiation. Rule [M-F-COM] defines how a party with role p synchronously sends a value to the receiving party with role q . Rules [M-F-CMT], [M-B-RLL] and [M-B-ABT] are similar to those of the binary case, and affect all participants within the considered session.

The syntax of session types extends to multiparty as shown in Fig. 18. The session types for output ($[p][q]!\langle S \rangle.T$) and input ($[p][q]?(S).T$) are extended with information about the interacting roles; selection and branching types are similarly extended. In the type inference, when one of such role is unknown, it is used a placeholder $_$ to be filled with a given role; $T \cdot p$ denotes the type obtained from T by filling all its placeholders with the role p . The cherry- π type system extends accordingly, as shown in Fig. 19.

The semantics of type configurations is defined only for filled types. Semantic rules in Fig. 20 are the natural extension of those for the binary case. Rule [M-TS-COM] shows that communication affects only the two interacting parties, without modifying any checkpoint. Rule [M-TS-CMT] sets the checkpoint of the committing party and sets an imposed checkpoint for each other party that has performed at least an action from its current checkpoint. Rule [M-TS-RLL] rolls all parties back to their checkpoints, provided that the checkpoint of the party requesting the rollback is not imposed. Rule [M-TS-ABT] brings all parties back to the initial configuration.

$$\begin{array}{c}
s[p][q]!\langle e \rangle.P \xrightarrow{s[p][q]!\langle v \rangle} P \quad (e \downarrow v) \text{ [M-P-SND]} \quad s[p][q]?(y:S).P \xrightarrow{s[p][q]?(y)} P \text{ [M-P-RCV]} \\
\\
\bar{a}[n](x).P_n \mid \prod_{i \in I} a[i](x).P_i \twoheadrightarrow \\
(\nu s : (\bar{a}[n](x).P_n \mid \prod_{i \in I} a[i](x).P_i)) \quad I = \{1, \dots, n-1\} \text{ [M-F-CON]} \\
(\langle P_n[s[n]/x] \rangle \blacktriangleright P_n[s[n]/x] \mid \prod_{i \in I} \langle P_i[s[i]/x] \rangle \blacktriangleright P_i[s[i]/x]) \\
\\
\frac{P_1 \xrightarrow{s[p][q]!\langle v \rangle} P'_1 \quad P_2 \xrightarrow{s[q][p]?(x)} P'_2}{\langle Q_1 \rangle \blacktriangleright P_1 \mid \langle Q_2 \rangle \blacktriangleright P_2 \twoheadrightarrow \langle Q_1 \rangle \blacktriangleright P'_1 \mid \langle Q_2 \rangle \blacktriangleright P'_2[v/x]} \text{ [M-F-COM]} \\
\\
\frac{P \xrightarrow{cmt} P'}{(\nu s : C)(\langle Q \rangle \blacktriangleright P \mid \prod_{i \in I} \langle Q_i \rangle \blacktriangleright P_i) \twoheadrightarrow (\nu s : C)(\langle P' \rangle \blacktriangleright P' \mid \prod_{i \in I} \langle P_i \rangle \blacktriangleright P_i)} \text{ [M-F-CMT]} \\
\\
\frac{P \xrightarrow{roll} P'}{(\nu s : C)(\langle Q \rangle \blacktriangleright P \mid \prod_{i \in I} \langle Q_i \rangle \blacktriangleright P_i) \rightsquigarrow (\nu s : C)(\langle Q \rangle \blacktriangleright Q \mid \prod_{i \in I} \langle Q_i \rangle \blacktriangleright Q_i)} \text{ [M-B-RLL]} \\
\\
\frac{P \xrightarrow{abt} P'}{(\nu s : C)(\langle Q \rangle \blacktriangleright P \mid \prod_{i \in I} \langle Q_i \rangle \blacktriangleright P_i) \rightsquigarrow C} \text{ [M-B-ABT]}
\end{array}$$

Fig. 17. Multiparty cherry-pi semantics: forward and backward reductions (excerpt of rules).

$T ::= \dots \mid [p][q]!\langle S \rangle.T \mid [p][q]?(S).T \mid [p][q] \triangleleft [l].T \mid [p][q] \triangleright [l_1:T_1, \dots, l_n:T_n]$ **Types**

Fig. 18. Multiparty cherry-pi type syntax (the omitted parts are as in Fig. 13).

$$\begin{array}{c}
\frac{\emptyset; \emptyset \vdash P \blacktriangleright x : T}{\bar{a}[p](x).P \blacktriangleright \{a[p] : T\}} \text{ [M-T-REQ]} \quad \frac{\emptyset; \emptyset \vdash P \blacktriangleright x : T}{a[p](x).P \blacktriangleright \{a[p] : T\}} \text{ [M-T-ACC]} \\
\\
\frac{\Gamma \vdash e \blacktriangleright S \quad \Theta; \Gamma \vdash P \blacktriangleright x : T}{\Theta; \Gamma \vdash x[p]!\langle e \rangle.P \blacktriangleright x : [-][p]!\langle S \rangle.T} \text{ [M-T-SND]} \\
\\
\frac{\Theta; \Gamma \cdot y : S \vdash P \blacktriangleright x : T}{\Theta; \Gamma \vdash x[p]?(y:S).P \blacktriangleright x : [-][p]?(S).T} \text{ [M-T-RCV]} \\
\\
\frac{\Theta; \Gamma \vdash P \blacktriangleright x : T}{\Theta; \Gamma \vdash x[p] \triangleleft l.P \blacktriangleright x : [-][p] \triangleleft [l].T} \text{ [M-T-SEL]} \\
\\
\frac{\Theta; \Gamma \vdash P_1 \blacktriangleright x : T_1 \quad \dots \quad \Theta; \Gamma \vdash P_n \blacktriangleright x : T_n}{\Theta; \Gamma \vdash x[p] \triangleright \{l_1:P_1, \dots, l_n:P_n\} \blacktriangleright x : [-][p] \triangleright [l_1:T_1, \dots, l_n:T_n]} \text{ [M-T-BR]}
\end{array}$$

Fig. 19. Multiparty cherry-pi typing system (the omitted rules are as in Fig. 6, 7 and 12).

Our notion of *rollback safety*, and the related compliance relation, extend to multiparty sessions as follows. Notice that types T_i in Def. 4 contain placeholders, while in Def. 5 all types are filled.

$$\begin{array}{c}
[p][q]!\langle S \rangle.T \xrightarrow{[p][q]!\langle S \rangle} T \quad [\text{M-TS-SND}] \qquad [p][q]?(S).T \xrightarrow{[p][q]?(S)} T \quad [\text{M-TS-Rcv}] \\
\\
\frac{T_i \xrightarrow{[p][q]!\langle S \rangle} T'_i \quad T_j \xrightarrow{[q][p]?(S)} T'_j}{(T^I) : \langle \tilde{U}_i \rangle \blacktriangleright T_i \parallel \langle \tilde{U}_j \rangle \blacktriangleright T_j \parallel \prod_{h \in I - \{i,j\}} \langle \tilde{U}_h \rangle \blacktriangleright T_h \mapsto (T^I) : \langle \tilde{U}_i \rangle \blacktriangleright T'_i \parallel \langle \tilde{U}_j \rangle \blacktriangleright T'_j \parallel \prod_{h \in I - \{i,j\}} \langle \tilde{U}_h \rangle \blacktriangleright T_h} \quad [\text{M-TS-COM}] \\
\\
\frac{T_i \xrightarrow{cmt} T'_i}{(T^I) : \langle \tilde{U}_i \rangle \blacktriangleright T_i \parallel \prod_{h \in I - \{i\}} \langle \tilde{U}_h \rangle \blacktriangleright T_h \mapsto (T^I) : \langle \tilde{U}_i \rangle \blacktriangleright T'_i \parallel \prod_{h \in I - \{i\}} \langle \tilde{U}_h \rangle \blacktriangleright T_h} \quad \tilde{U}'_h = \begin{cases} \tilde{U}_h & \text{if } \tilde{U}_h = T_h \\ \underline{T}_h & \text{otherwise} \end{cases} \quad [\text{M-TS-CMT}] \\
\\
\frac{T_i \xrightarrow{roll} T'_i}{(T^I) : \langle U_i \rangle \blacktriangleright T_i \parallel \prod_{h \in I - \{i\}} \langle \tilde{U}_h \rangle \blacktriangleright T_h \mapsto (T^I) : \langle U_i \rangle \blacktriangleright U_i \parallel \prod_{h \in I - \{i\}} \langle \tilde{U}_h \rangle \blacktriangleright T_h} \quad [\text{M-TS-RLL}] \\
\\
\frac{T_i \xrightarrow{abt} T'_i}{(T^I) : \langle \tilde{U}_i \rangle \blacktriangleright T_i \parallel \prod_{h \in I - \{i\}} \langle \tilde{U}_h \rangle \blacktriangleright T_h \mapsto (T^I) : \prod_{k \in I} \langle T^k \rangle \blacktriangleright T^k} \quad [\text{M-TS-ABT}]
\end{array}$$

Fig. 20. Multiparty cherry- π type semantics (excerpt of rules, where $I = \{1, \dots, n\}$, $i, j \in I$, and T^I denotes T^1, \dots, T^n).

Definition 4 (Multiparty Rollback safety). Let C be an initial collaboration, then C is rollback safe (shortened roll-safe) if $C \blacktriangleright A$ and for each n -tuple $\bar{a}[n] : T_n, \dots, a[1] : T_1$ in A we have $\dashv\!\!\dashv (T_n \cdot n, \dots, T_1 \cdot 1)$.

Definition 5 (Compliance for Multiparty Sessions). Types T_1, \dots, T_n are compliant, written $\dashv\!\!\dashv (T_1, \dots, T_n)$, if $\dashv\!\!\dashv ((T_1, \dots, T_n) : \langle T_1 \rangle \blacktriangleright T_1, \dots, \langle T_n \rangle \blacktriangleright T_n)$. Relation $\dashv\!\!\dashv$ on type configurations is defined as follows: $\dashv\!\!\dashv ((T^1, \dots, T^n) : \langle \tilde{U}_1 \rangle \blacktriangleright T_1, \dots, \langle \tilde{U}_n \rangle \blacktriangleright T_n)$ holds if for all $\tilde{U}'_1, T'_1, \dots, \tilde{U}'_n, T'_n$ such that $(T^1, \dots, T^n) : \prod_{h \in \{1, \dots, n\}} \langle \tilde{U}_h \rangle \blacktriangleright T_h \mapsto^* (T^1, \dots, T^n) : \prod_{h \in \{1, \dots, n\}} \langle \tilde{U}'_h \rangle \blacktriangleright T'_h \not\vdash$ we have that $T'_1 = \dots = T'_n = \text{end}$.

All notions and concepts of our rollback recovery approach smoothly extend to the multiparty case. As consequence, all properties in Sec. 5 still hold in the extended setting; their proofs are in Appendix C.4. We report below just the key theorem concerning session progress.

Theorem 7 (Multiparty session progress). Let C be a roll-safe collaboration of the form $(\bar{a}[n](x).P_n \mid \prod_{i \in \{1, \dots, n-1\}} a[i](x).P_i)$. If $C \mapsto^* C'$ then either $C' \mapsto C''$ for some C'' or $C' \equiv (\nu s : C) \prod_{i \in \{1, \dots, n\}} \langle \tilde{Q}_i \rangle \blacktriangleright \mathbf{0}$ for some $\tilde{Q}_1, \dots, \tilde{Q}_n$.

C Proofs

C.1 Decidability result

Theorem 1. Let T_1 and T_2 be two session types, checking if $T_1 \dashv\!\!\dashv T_2$ holds is decidable.

Proof. By Def. 1, checking $T_1 \dashv\!\!\dashv T_2$ consists in checking that types T'_1 and T'_2 of each configuration $(T_1, T_2) : \langle \tilde{U}'_1 \rangle \blacktriangleright T'_1 \parallel \langle \tilde{U}'_2 \rangle \blacktriangleright T'_2$ such that $(T_1, T_2) : \langle T_1 \rangle \blacktriangleright T_1 \parallel \langle T_2 \rangle \blacktriangleright T_2 \xrightarrow{*} (T_1, T_2) : \langle \tilde{U}'_1 \rangle \blacktriangleright T'_1 \parallel \langle \tilde{U}'_2 \rangle \blacktriangleright T'_2 \not\vdash$ (i.e., type configurations that are reachable from the initial one and that cannot further evolve) are end types. Thus, to prove that the compliance check is decidable we have to show that the number of these reachable configurations is finite. Let us consider the transition system $TS = \langle S, \mathcal{R} \rangle$ associated to the type configuration $t = (T_1, T_2) : \langle T_1 \rangle \blacktriangleright T_1 \parallel \langle T_2 \rangle \blacktriangleright T_2$ by the reduction semantics of types (Fig. 8): the set S of states corresponds to the set of type configurations reachable from t , i.e. $S = \{ t' \mid t \xrightarrow{*} t' \}$, while the set \mathcal{R} of system transitions corresponds to set of the type reductions involving configurations in S , i.e. $\mathcal{R} = \{ (t', t'') \in S \times S \mid t' \xrightarrow{} t'' \}$. Hence, checking $T_1 \dashv\!\!\dashv T_2$ boils down to check the type configurations corresponding to the leaves (i.e., states without outgoing transitions) of TS . Specifically, given a leaf of TS corresponding to $(T_1, T_2) : \langle \tilde{V}_1 \rangle \blacktriangleright V_2 \parallel \langle \tilde{V}_3 \rangle \blacktriangleright V_4$, we have to check if $V_2 = V_4 = \text{end}$. The decidability of this check therefore depends on the finiteness of TS . This result is ensured by the fact that: (i) backward reductions connect states of TS only to previously visited states of TS (Theorem 2), and (ii) our language of types (Fig. 13) corresponds to a CCS-like process algebra without static operators (i.e., parallel and restriction operators) within recursion (see [29, Sec. 7.5]).

C.2 Reversibility results

We can instrument our semantics in order to carry the information on which session the reduction is taking place. Hence we will indicate with $C \xrightarrow{s} C'$ the fact that the reduction is taking place on session s . We can show that sessions are independent:

Lemma 3 (Swap Lemma). *Let C be a collaboration and s and r two sessions. If $C \xrightarrow{s} C_1 \xrightarrow{r} C_2$ then there exists a collaboration C_3 such that $C \xrightarrow{r} C_3 \xrightarrow{s} C_2$.*

Proof. By case analysis on the reductions \xrightarrow{s} and \xrightarrow{r} .

Lemma 4. *Let C be a collaboration. If $C \xrightarrow{*} C_1$, then for any session s in C_1 there exists a collaboration C_0 such that $C \xrightarrow{*} C_0 \xrightarrow{s} C_1$ and s is never used in the trace $C \xrightarrow{*} C_0$.*

Proof. By induction on the number n of reduction on s . If there are no reductions then the thesis banally holds. Otherwise we can take the very last reduction on s , that is the closest one to C_1 and iteratively apply Lemma 3 in order to bring it to the very end. Then we can conclude by induction on a trace with less occurrences of reductions on s .

Thanks to Lemma 4 we can rearrange any trace as a sequence of *independent* sessions. Moreover given an initial collaboration C for any reachable collaboration C_1 and session s such that $C \xrightarrow{*} C_1 \xrightarrow{s} *$ and $s \notin \text{used}(C)$, we indicate C_1 as the *initial* collaboration for s . This will allows us to focus just on a single session, say s , and to consider collaboration initial for s without loosing of generality.

Lemma 5. *Let C be an initial collaboration such that $C \xrightarrow{*} C_1$. If $C_1 \xrightarrow{abt} C_2$ then $C_2 \equiv C$.*

Proof. Since C is *initial*, without losing of generality we can assume

$$C \equiv \bar{a}(x_1).P_1 \mid a(x_2).P_2$$

The first reduction of $C \rightarrow^* C_1$ has to be an application of rule [F-CON], that is

$$\begin{aligned} C &\rightarrow (\nu s : (\bar{a}(x_1).P_1 \mid a(x_2).P_2)) \\ &\quad (\langle P_1[\bar{s}/x_1] \rangle \blacktriangleright P_1[\bar{s}/x_1] \mid \langle P_2[s/x_2] \rangle \blacktriangleright P_2[s/x_2]) = C' \end{aligned}$$

and, by hypothesis, $C' \rightarrow^* C_1$.

Now, no matter the shape of processes in C_1 by applying rule [B-ABT], and possibly [B-STR], we will go back to C , that is $C_1 \xrightarrow{abt} C$, as desired.

The following lemma states that a rollback leads back to the last committed check-point.

Lemma 6. *Let C be a reachable collaboration, such that $C \xrightarrow{cmt} C_1$. If $C_1 \rightarrow^* C_2 \xrightarrow{roll}$ and there is no commit in $C_1 \rightarrow^* C_2$, then $C_3 \equiv C_1$.*

Proof. Since C is a reachable collaboration, this implies it has been generated from an initial collaboration C_0 . Without losing of generality, similarly to the Lemma 5's proof, we can assume $C \equiv (\nu s : C_0)(\langle P_1 \rangle \blacktriangleright P_2 \mid \langle Q_1 \rangle \blacktriangleright Q_2) \mid C^\circ$. Therefore, we have that $C_1 = (\nu s : C_0)(\langle P_2 \rangle \blacktriangleright P_2 \mid \langle Q_2 \rangle \blacktriangleright Q_2) \mid C^\circ$. By hypothesis, there is no commits in $C_1 \rightarrow^* C_2$, and this implies that the log part of the C_1 will never change. Hence, we have that $C_2 \equiv (\nu s : C_0)(\langle P_2 \rangle \blacktriangleright P \mid \langle Q_2 \rangle \blacktriangleright Q) \mid C^\circ$ for some processes P and Q . By applying [B-RLL] and [B-PAR] we have that $C_2 \xrightarrow{roll} (\nu s : C_0)(\langle P_2 \rangle \blacktriangleright P_2 \mid \langle Q_2 \rangle \blacktriangleright Q_2) \mid C^\circ \equiv C_1$, as desired.

Any reachable collaboration is also a *forward only* reachable collaboration. Formally:

Theorem 2. Let C_0 be an initial collaboration. If $C_0 \rightarrow^* C_1$ then $C_0 \rightarrow^* C_1$.

Proof. By induction on the number n of backward reductions contained into $C_0 \rightarrow^* C_1$. The base case ($n = 0$) trivially holds. In the inductive case, let us take the backward reduction which is the nearest to C_0 . That is:

$$C_0 \rightarrow^* C' \rightsquigarrow C'' \rightarrow^* C_1$$

Depending whether it is an \xrightarrow{abt} or a \xrightarrow{roll} we can apply respectively Lemma 5 or Lemma 6 to obtain a forward trace of the form

$$C_0 \rightarrow^* C'' \rightarrow^* C_1$$

and we can conclude by applying the inductive hypothesis on the obtained trace which contains less backward moves with respect to the original one.

Lemma 1. Let C_1 and C_2 be two reachable collaborations. $C_1 \rightsquigarrow C_2$ then $C_2 \rightarrow^* C_1$.

Proof. Since C_1 is a reachable collaboration, we have that there exists an initial collaboration C_0 such that $C_0 \rightarrow^* C_1$. By applying Theorem 2 we can rearrange the trace such that it contains just forward transitions as follows

$$C_0 \rightarrow^* C_1$$

If the backward reduction is obtained by applying [B-ABT], by Lemma 5 we have $C_2 \equiv C_0$, from which the thesis trivially follows. Instead, if the backward reduction is obtained by applying [B-RLL], we proceed by case analysis depending on the presence of commit reductions in the trace. If they are present, we select the last of such commit, that is we can decompose the trace in the following way:

$$C_0 \rightarrow^* \xrightarrow{cmt} C_{cmt} \rightarrow^* C_1 \xrightarrow{roll} C_2$$

and by applying Lemma 6 we have that $C_2 \rightarrow^* \equiv C_1$ as desired.

In the case there is no commit in the trace, without losing of generality we can assume

$$\begin{aligned} C_0 &\equiv \bar{a}(x_1).P_1 \mid a(x_2).P_2 \mid C^\circ \\ C_0 &\rightarrow (\nu s : C_0)(\langle P_2 \rangle \blacktriangleright P_2 \mid \langle Q_2 \rangle \blacktriangleright Q_2) \mid C^\circ \end{aligned}$$

By rule [B-RLL], $C_2 \equiv (\nu s : C_0)(\langle P_2 \rangle \blacktriangleright P_2 \mid \langle Q_2 \rangle \blacktriangleright Q_2) \mid C^\circ$. Thus, we can conclude by noticing that $C_0 \rightarrow C_2$ must be the first reduction in $C_0 \rightarrow^* C_1$.

Lemma 2. Let C be a reachable collaboration. If $C \xrightarrow{rolls} C'$ and $C \xrightarrow{rolls} C''$ then $C' \equiv C''$.

Proof. Since C is a reachable collaboration, it is has been generated by an initial collaboration C_0 of the form $C_0 = \bar{a}(x).P_1 \mid a(x).P_2$, and by Theorem 2 we have that $C_0 \rightarrow^* C$. We distinguish two cases, whether in the trace there has been at least one commit or not. In the first case, we can decompose the trace in such a way to single out the last commit as follows:

$$C_0 \rightarrow^* C_{cmt} \rightarrow^* C$$

so that in the reduction $C_{cmt} \rightarrow^* C$ there is no commit. If from C the rollbacks $C \xrightarrow{roll} C'$ and $C \xrightarrow{roll} C''$ are triggered by the same process, the thesis trivially follows. In the other case, we have that:

$$C_{cmt} \equiv (\nu s : C_0)(\langle P_1 \rangle \blacktriangleright P_2 \mid \langle Q_1 \rangle \blacktriangleright Q_2) \mid C^\circ$$

with both P_2 and Q_2 are able to trigger a rollback. If the roll action is executed by P_2 we have that

$$(\nu s : C_0)(\langle P_1 \rangle \blacktriangleright P_2 \mid \langle Q_1 \rangle \blacktriangleright Q_2) \mid C^\circ \xrightarrow{roll} (\nu s : C_0)(\langle P_1 \rangle \blacktriangleright P_1 \mid \langle Q_1 \rangle \blacktriangleright Q_1) \mid C^\circ = C'$$

If the roll is triggered by Q_2 we have that

$$(\nu s : C_0)(\langle P_1 \rangle \blacktriangleright P_2 \mid \langle Q_1 \rangle \blacktriangleright Q_2) \mid C^\circ \xrightarrow{roll} (\nu s : C_0)(\langle P_1 \rangle \blacktriangleright P_1 \mid \langle Q_1 \rangle \blacktriangleright Q_1) \mid C^\circ = C''$$

And we can conclude by noticing that $C' \equiv C''$, as desired.

Theorem 3. Let C be a reachable collaboration. If $C \xrightarrow{cmt} C'$ then there exists no C'' such that $C' \xrightarrow{*} \xrightarrow{roll} C''$ and $C'' \xrightarrow{+} C$.

Proof. We proceed by contradiction. Suppose that there exists C'' such that $C' \xrightarrow{*} \xrightarrow{roll} C''$ and $C'' \xrightarrow{+} C$. Since C is a reachable collaboration, thanks to Theorem 2 we have that there exists an initial collaboration C_0 such that $C_0 \xrightarrow{*} C \xrightarrow{cmt} C'$. Since a rollback brings back the collaboration to a point before a commit, this means it has been restored a commit previous to the last one (by [B-RLL], indeed, only processes stored in logs can be committed). This implies that there exist at least two different commits in the trace such that

$$C_0 \xrightarrow{*} C_{cmt} \xrightarrow{*} C \xrightarrow{cmt} C' \xrightarrow{*} \xrightarrow{roll} C''$$

with $C'' = C_{cmt}$. We have that $C'' \xrightarrow{+} C \xrightarrow{cmt} C' \xrightarrow{*} C_{rl}$, where C_{rl} is the collaboration that perform the roll reduction. Now since the last commit has been done by C , supposing that the commit is triggered by P_c , which evolves to P'_c in doing that, we have that:

$$C \equiv (\nu s : C_0)(\langle P \rangle \blacktriangleright P_c \mid \langle Q \rangle \blacktriangleright Q_c) \text{ and } C' = (\nu s : C_0)(\langle P'_c \rangle \blacktriangleright P'_c \mid \langle Q_c \rangle \blacktriangleright Q_c)$$

Now, by hypothesis we have that $C' \xrightarrow{*} C_{rl}$ without any commit being present in the trace, hence:

$$C_{rl} \equiv (\nu s : C_0)(\langle P'_c \rangle \blacktriangleright P_{rl} \mid \langle Q_c \rangle \blacktriangleright Q_{rl})$$

By hypothesis, from C_{rl} a rollback is possible. Regardless the rollback is triggered by P_{rl} or Q_{rl} , we have that $C_{rl} \xrightarrow{roll} C'$. Now, from C' we cannot reach C ($C' \not\xrightarrow{+} C$), as C' is derived from C and the rollback can only bring the collaboration back to C' . This violates the hypothesis, and hence we conclude.

C.3 Soundness results

To prove our soundness results, we need to introduce some auxiliary lemmas, which rely on the following definitions:

- a process \tilde{P} and a type \tilde{T} are in *checkpoint accordance* if $\tilde{P} = P$ implies $\tilde{T} = T$, and $\tilde{P} = \underline{P}$ implies $\tilde{T} = \underline{T}$;
- let ℓ a process label, its *dual label* $\bar{\ell}$ is defined as follows: $\overline{k!\langle v \rangle} = k?(x)$ for some x , $\overline{k?(x)} = k!\langle v \rangle$ for some v , $\overline{k \triangleleft l} = k \triangleright l$, $\overline{k \triangleright l} = k \triangleleft l$; this notion of duality straightforwardly extends to type labels;
- the function $tl_\Gamma(\cdot)$, mapping process labels to type labels under sorting Γ , is defined as follows: $tl_\Gamma(k!\langle v \rangle) = ![S]$ with $\Gamma \vdash v \blacktriangleright S$, $tl_\Gamma(k?(x)) = ?[S]$ with $\Gamma \vdash x \blacktriangleright S$, $tl_\Gamma(k \triangleleft l) = \triangleleft l$, $tl_\Gamma(k \triangleright l) = \triangleright l$, $tl_\Gamma(cmt) = cmt$, $tl_\Gamma(roll) = roll$, $tl_\Gamma(abt) = abt$, and $tl_\Gamma(\tau) = \tau$.

The following lemma states that each reduction of a reachable collaboration corresponds to a reduction of its configuration types.

Lemma 7. *Let $C = (\nu s : C')(\langle \tilde{P}_1 \rangle \blacktriangleright P_2 \mid \langle \tilde{Q}_1 \rangle \blacktriangleright Q_2)$ be a reachable collaboration, with $C' = (\bar{a}(x).P \mid a(y).Q)$, $(\emptyset; \emptyset \vdash P \blacktriangleright x : T)$, $(\emptyset; \emptyset \vdash Q \blacktriangleright y : T')$, $T \dashv\dashv T'$, $(\Theta_1; \Gamma_1 \vdash P_1[x/\bar{s}] \blacktriangleright x : T_1)$, $(\Theta_2; \Gamma_2 \vdash P_2[x/\bar{s}] \blacktriangleright x : T_2)$, $(\Theta'_1; \Gamma'_1 \vdash Q_1[y/\bar{s}] \blacktriangleright y : U_1)$, and $(\Theta'_2; \Gamma'_2 \vdash Q_2[y/\bar{s}] \blacktriangleright y : U_2)$. If $C \mapsto (\nu s : C')(\langle \tilde{P}'_1 \rangle \blacktriangleright P'_2 \mid \langle \tilde{Q}'_1 \rangle \blacktriangleright Q'_2)$ then there exist T'_1, T'_2, U'_1, U'_2 such that $(T, T') : \langle \tilde{T}_1 \rangle \blacktriangleright T_2 \parallel \langle \tilde{U}_1 \rangle \blacktriangleright U_2 \mapsto (T, T') : \langle \tilde{T}'_1 \rangle \blacktriangleright T'_2 \parallel \langle \tilde{U}'_1 \rangle \blacktriangleright U'_2$ with \tilde{P}'_1 (resp. \tilde{Q}'_1) in checkpoint accordance with \tilde{T}'_1 (resp. \tilde{U}'_1), $(\hat{\Theta}_1; \hat{\Gamma}_1 \vdash P'_1[x/\bar{s}] \blacktriangleright x : T'_1)$, $(\hat{\Theta}_2; \hat{\Gamma}_2 \vdash P'_2[x/\bar{s}] \blacktriangleright x : T'_2)$, $(\hat{\Theta}'_1; \hat{\Gamma}'_1 \vdash Q'_1[y/\bar{s}] \blacktriangleright y : U'_1)$, and $(\hat{\Theta}'_2; \hat{\Gamma}'_2 \vdash Q'_2[y/\bar{s}] \blacktriangleright y : U'_2)$.*

Proof. We have two cases depending whether the reduction \mapsto has forward or backward direction.

($\mapsto = \Rightarrow$). From rule [F-RES], we have $\langle \tilde{P}_1 \rangle \blacktriangleright P_2 \mid \langle \tilde{Q}_1 \rangle \blacktriangleright Q_2 \Rightarrow \langle \tilde{P}'_1 \rangle \blacktriangleright P'_2 \mid \langle \tilde{Q}'_1 \rangle \blacktriangleright Q'_2$. We prove the result by case analysis on the last rule applied in the inference of the above reduction.

- [F-COM]. In this case we have $P_2 \xrightarrow{\bar{s}!(v)} P'_2$, $Q_2 \xrightarrow{s?(z)} Q'_2$, with $Q'_2 = Q''_2[v/z]$, $\tilde{P}'_1 = \tilde{P}_1$ and $\tilde{Q}'_1 = \tilde{Q}_1$. Thus, $P_2[x/\bar{s}] = x!(e).P'_2[x/\bar{s}]$ for some e such that $e \downarrow v$, and $Q_2[y/\bar{s}] = y?(z : S').Q'_2[y/\bar{s}]$. By rule [T-SND], we have that $T_2 = ![S].T'_2$, with $\Gamma_2 \vdash e \blacktriangleright S$, (hence $\Gamma_2 \vdash v \blacktriangleright S$), and $\Theta_2; \Gamma_2 \vdash P'_2[x/\bar{s}] \blacktriangleright x : T'_2$ (hence $\hat{\Theta}_2 = \Theta_2$ and $\hat{\Gamma}_2 = \Gamma_2$). Similarly, by rule [T-RCV], we have that $U_2 = ?[S'].U'_2$ and $\Theta'_2; \Gamma'_2 \vdash z : S' \vdash Q''_2[y/\bar{s}] \blacktriangleright y : U'_2$ (hence $\hat{\Theta}'_2 = \Theta'_2$ and $\hat{\Gamma}'_2 = \Gamma'_2 \cdot z : S'$). By rules [TS-SND] and [TS-RCV], we get $T_2 \xrightarrow{![S]} T'_2$ and $U_2 \xrightarrow{?[S']} U'_2$. Now, reasoning by contradiction, let us suppose that $S \neq S'$. Thus, the term $(T, T') : \langle \tilde{T}_1 \rangle \blacktriangleright T_2 \parallel \langle \tilde{U}_1 \rangle \blacktriangleright U_2 \nrightarrow$, since no rule in Fig. 8 can be applied. However, since C is a reachable collaboration, this type configuration is originated from $(T, T') : \langle T \rangle \blacktriangleright T \parallel \langle T' \rangle \blacktriangleright T'$. By Def. 1, $T \dashv\dashv T'$ implies $T_2 = U_2 = \text{end}$, which is a contradiction. Therefore, it holds that $S = S'$. Hence, by applying rule [TS-COM] we can conclude.
- [F-LAB], [E-CMT₁] and [E-CMT₂]. Similar to the previous case.
- [F-STR]. The use of \equiv leads us back to one of the other cases.
- [F-PAR]. In this case we have that $\langle \tilde{P}_1 \rangle \blacktriangleright P_2 \Rightarrow \langle \tilde{P}'_1 \rangle \blacktriangleright P'_2$. Since this transition involves only one log term, it can be inferred only by applying rule [F-IF], from which we have $P_2 \xrightarrow{\tau} P'_2$ and $\tilde{P}'_1 = \tilde{P}_1$. By rule [P-IFT] (the case of rule [P-IFF] is similar), we have $P_2[x/\bar{s}] = \text{if } e \text{ then } P'_2[x/\bar{s}] \text{ else } R$ with $e \downarrow \text{true}$. By rule [T-IF] we get $T_2 = T'_2 \oplus V$ and $\Theta_2; \Gamma_2 \vdash P'_2[x/\bar{s}] \blacktriangleright x : T'_2$. By rule [TS-IF₁], $T_2 \xrightarrow{\tau} T'_2$. By applying rule [TS-TAU] we can conclude.

($\mapsto = \rightsquigarrow$). From rule [B-RES], we have $\langle \tilde{P}_1 \rangle \blacktriangleright P_2 \mid \langle \tilde{Q}_1 \rangle \blacktriangleright Q_2 \rightsquigarrow \langle \tilde{P}'_1 \rangle \blacktriangleright P'_2 \mid \langle \tilde{Q}'_1 \rangle \blacktriangleright Q'_2$. We prove the result by case analysis on the last rule applied in the inference of the above reduction.

- [B-RLL]. In this case $P_2 \xrightarrow{\text{roll}} P'_2$, $\tilde{P}'_1 = \tilde{P}_1$, $\tilde{Q}'_1 = Q_1$ and $Q'_2 = Q_1$. By rule [P-RLL], we have $P_2[x/\bar{s}] = \text{roll}$ and $P'_2[x/\bar{s}] = \mathbf{0}$. By rule [T-RLL] we get $T_2 = \text{roll}$. By rule [TS-RLL], $T_2 \xrightarrow{\text{roll}} T'_2$, with $T'_2 = \text{end}$. By rule [T-INACT], we have $\Theta_2; \Gamma_2 \vdash P'_2[x/\bar{s}] \blacktriangleright x : T'_2$. Finally, by applying rule [TS-RLL₁] we can conclude.
- [B-STR] and [B-PAR]. Similarly to the forward cases.

The following lemma relates collaboration reductions to type reductions when a `roll_error` is produced.

Lemma 8. *Let $C = (\bar{a}(x).P \mid a(y).Q)$ such that $C \triangleright \{\bar{a} : T_1, a : T_2\}$ and $T_1 \dashv\vdash T_2$. If $C \mapsto^* (\nu s : C)(\langle \tilde{P}_1 \rangle \triangleright P_2 \mid \langle \tilde{Q}_1 \rangle \triangleright Q_2)$ and $P_2 \xrightarrow{\text{roll}} P'_2$, then there exist $U_1, U_2, U'_1, U'_2, U''_1$ such that $(T_1, T_2) : \langle T_1 \rangle \triangleright T_1 \parallel \langle T_2 \rangle \triangleright T_2 \mapsto^* (T_1, T_2) : \langle U_1 \rangle \triangleright U'_1 \parallel \langle \tilde{U}_2 \rangle \triangleright U'_2$ and $U'_1 \xrightarrow{\text{roll}} U''_1$.*

Proof. From $C \triangleright \{\bar{a} : T_1, a : T_2\}$, by applying [T-PAR], [T-ACC] and [T-REQ], we have that $\emptyset; \emptyset \vdash P \triangleright x : T_1$ and $\emptyset; \emptyset \vdash Q \triangleright x : T_2$. By applying rule [F-CON] to the collaboration C , we obtain $C \mapsto (\nu s : C)(\langle P[\bar{s}/x] \rangle \triangleright P[\bar{s}/x] \mid \langle Q[s/y] \rangle \triangleright Q[s/y]) = C'$. Now, by repeatedly applying Lemma 7, from $C' \mapsto^* (\nu s : C)(\langle \tilde{P}_1 \rangle \triangleright P_2 \mid \langle \tilde{Q}_1 \rangle \triangleright Q_2)$ we get $(T_1, T_2) : \langle T_1 \rangle \triangleright T_1 \parallel \langle T_2 \rangle \triangleright T_2 \mapsto^* (T_1, T_2) : \langle U_1 \rangle \triangleright U'_1 \parallel \langle \tilde{U}_2 \rangle \triangleright U'_2$ for some U_1, U_2, U'_1, U'_2 , with $\Theta_2; \Gamma_2 \vdash P_2[x/\bar{s}] \triangleright x : U'_1$. Now, let us consider the transition $P_2 \xrightarrow{\text{roll}} P'_2$. This can be derived only by the application of rules [P-RLL]. Thus, $P_2 = \text{roll}$, from which we have $P_2[x/\bar{s}] = \text{roll}$. From $\Theta_2; \Gamma_2 \vdash \text{roll} \triangleright x : U'_1$, by rule [T-RLL], we get $U'_1 = \text{roll}$. Therefore, by rule [TS-RLL], we can conclude $U'_1 \xrightarrow{\text{roll}} U''_1$ with $U''_1 = \text{end}$.

The following lemma relates collaboration reductions to type reductions when a `com_error` is produced.

Lemma 9. *Let $C = (\bar{a}(x).P \mid a(y).Q)$ such that $C \triangleright \{\bar{a} : T_1, a : T_2\}$ and $T_1 \dashv\vdash T_2$. If $C \mapsto^* (\nu s : C)(\langle \tilde{P}_1 \rangle \triangleright P_2 \mid \langle \tilde{Q}_1 \rangle \triangleright Q_2)$, $P_2 \xrightarrow{\ell} P'_2$ and $\neg Q_2 \Downarrow_{\bar{\ell}}$ with ℓ of the form $k!\langle v \rangle$, $k?(x)$, $k \triangleleft l$ or $k \triangleright l$, then there exist $U_1, U_2, U'_1, U'_2, U''_1$ such that $(T_1, T_2) : \langle T_1 \rangle \triangleright T_1 \parallel \langle T_2 \rangle \triangleright T_2 \mapsto^* (T_1, T_2) : \langle \tilde{U}_1 \rangle \triangleright U'_1 \parallel \langle \tilde{U}_2 \rangle \triangleright U'_2$ with \tilde{P}_1 (resp. \tilde{Q}_1) in checkpoint accordance with \tilde{U}_1 (resp. \tilde{U}_2), $U'_1 \xrightarrow{\text{tlr}(\ell)} U''_1$, with Γ sorting for typing P'_2 , and for all U''_2 such that $U'_2 \xrightarrow{\tau} U''_2$ we have $U''_2 \xrightarrow{\text{tlr}(\ell)}$.*

Proof. From $C \triangleright \{\bar{a} : T_1, a : T_2\}$, by applying [T-PAR], [T-ACC] and [T-REQ], we have that $\emptyset; \emptyset \vdash P \triangleright x : T_1$ and $\emptyset; \emptyset \vdash Q \triangleright x : T_2$. By applying rule [F-CON] to the collaboration C , we obtain $C \mapsto (\nu s : C)(\langle P[\bar{s}/x] \rangle \triangleright P[\bar{s}/x] \mid \langle Q[s/y] \rangle \triangleright Q[s/y]) = C'$. Now, by repeatedly applying Lemma 7, from $C' \mapsto^* (\nu s : C)(\langle \tilde{P}_1 \rangle \triangleright P_2 \mid \langle \tilde{Q}_1 \rangle \triangleright Q_2)$ we get $(T_1, T_2) : \langle T_1 \rangle \triangleright T_1 \parallel \langle T_2 \rangle \triangleright T_2 \mapsto^* (T_1, T_2) : \langle \tilde{U}_1 \rangle \triangleright U'_1 \parallel \langle \tilde{U}_2 \rangle \triangleright U'_2$ for some U_1, U_2, U'_1, U'_2 , with $\Theta_2; \Gamma_2 \vdash P_2[x/\bar{s}] \triangleright x : U'_1$ and $\Theta'_2; \Gamma'_2 \vdash Q_2[y/s] \triangleright y : U'_2$. Now, let us reason by case analysis on the rule for deriving the transition $P_2 \xrightarrow{\ell} P'_2$.

Rule [P-SND]. Thus, $P_2 = \bar{s}!\langle e \rangle.P'_2$ and $\ell = \bar{s}!\langle v \rangle$ with $e \downarrow v$. From $\Theta_2; \Gamma_2 \vdash P_2[x/\bar{s}] \triangleright x : U'_1$, by rule [T-SND], we get $U'_1 = ![S].U''_1$ with $\Gamma \vdash e \triangleright S$.

Therefore, by rule [TS-SND], we get $U'_1 \xrightarrow{![S]} U''_1$.

Rule [P-RCV]. Thus, $P_2 = \bar{s}?(y : S).P'_2$ and $\ell = \bar{s}?(y)$. From $\Theta_2; \Gamma_2 \vdash P_2[x/\bar{s}] \triangleright x : U'_1$, by rule [T-RCV], we get $U'_1 = ?[S].U''_1$. Therefore, by rule [TS-RCV], we get

$U'_1 \xrightarrow{?[S]} U''_1$.

Rules [P-SEL] and [P-BRN]. Similar to the previous cases.

Rule [P-STR]. According to the definition of \equiv , the only relevant case is when $P_2 = \mu X.R$. Thus, by rule [P-STR] we have $P_2 \equiv R[\mu X.R/X] = P_2''$, and $P_2'' \xrightarrow{\ell} P_2'''$ with $P_2''' \equiv P_2'$. If P_2'' has not the form $\mu Y.R'$, we can proceed as in one of the cases above, otherwise we repeat this reasoning.

Finally, from $\neg Q_2 \Downarrow_{\bar{\ell}}$, following a similarly reasoning, we can conclude $U_2' \xrightarrow{\tau} \overline{tl_{\Gamma}(\ell)}$.

We can now prove our soundness results.

Theorem 4. If C is a rollback safe collaboration, then $C \not\rightarrow^* \mathbb{C}[\text{roll_error}]$.

Proof. The proof proceeds by contradiction. Suppose that there exists an initial collaboration C that is rollback safe and such that $C \rightarrow^* \mathbb{C}[\text{roll_error}]$. The erroneous collaboration `roll_error` can be only produced by applying rule [E-RLL₂]. Thus, to infer at least one reduction of the sequence $C \rightarrow^* \mathbb{C}[\text{roll_error}]$, rule [E-RLL₂] must be used. From this, we have that there exists a runtime collaboration $C' \equiv \mathbb{C}'[C'']$, with $C'' = (\langle \underline{Q}_1 \rangle \blacktriangleright P_1 \mid \langle \tilde{Q}_2 \rangle \blacktriangleright P_2)$, such that $C \rightarrow^* C'$, $P_1 \xrightarrow{\text{roll}} P_1'$, and $C' \rightarrow \mathbb{C}'[\text{roll_error}] \rightarrow^* \mathbb{C}[\text{roll_error}]$. By rules [F-CON], [F-RES] and [B-RES], and the fact that the scope of operator $(\nu s : _)$ is statically defined (i.e., neither the operational rules nor \equiv allow scope extension), the term C'' can only be the argument of the operator $(\nu s : C_1)$, i.e. $\mathbb{C}' = C_2 \mid (\nu s : C_1)[_]$, with $C_1 = \bar{a}(x).P \mid a(y).Q$ for some a, x, y, P and Q . In its own turn, the term $(\nu s : C_1)C''$ can only be generated by applying rule [F-CON] from C_1 , which must be a subterm of C , i.e. $C \equiv C_1 \mid C_2'$ for some C_2' . Since the scope of $(\nu s : _)$ operator cannot be extended, all reductions performed by terms in parallel with it by applying rules [F-PAR] and [B-PAR] do not affect the argument of such operator. Therefore, focussing on the subterm C_1 of C , by exploiting rules [F-PAR] and [B-PAR] we can set apart the reductions in $C \rightarrow^* \mathbb{C}[\text{roll_error}]$ involving C_1 and its derivatives, thus obtaining $C_1 \rightarrow^* (\nu s : C_1)C'' \rightarrow (\nu s : C_1)\text{roll_error}$.

Now, since C is rollback safe, by Def. 2 we have that $C \blacktriangleright A$ and for all pairs $\bar{b} : V_1$ and $b : V_2$ in A we have $V_1 \dashv V_2$. Since $C \equiv C_1 \mid C_2'$, by rule [T-PAR] we obtain $C_1 \blacktriangleright A_1$ with $A_1 \subseteq A$. By rules [T-REQ] and [T-ACC], we have $A_1 = \{\bar{a} : T_1, a : T_2\}$. Since A_1 is a subset of A_2 , that is $A_1 \subseteq A_2$, we have that $T_1 \dashv T_2$.

By Lemma 8, we have that there exist $U_1, U_2, U_1', U_2', U_1''$ such that $(T_1, T_2) : \langle T_1 \rangle \blacktriangleright T_1 \parallel \langle T_2 \rangle \blacktriangleright T_2 \rightarrow^* (T_1, T_2) : \langle \underline{U}_1 \rangle \blacktriangleright U_1' \parallel \langle \tilde{U}_2 \rangle \blacktriangleright U_2' = t$ and $U_1' \xrightarrow{\text{roll}} U_1''$. Since U_1' can only perform `roll`, rule [TS-RLL₁] cannot be applied due to the imposed checkpoint \underline{U}_1 . Therefore, the only rule that can be applied is [TS-RLL₂], leading to the configuration $t' = (T_1, T_2) : \langle \underline{U}_1 \rangle \blacktriangleright U_1'' \parallel \langle \tilde{U}_2 \rangle \blacktriangleright U_2''$ with $U_1'' = U_2'' = \text{err}$. Now, no rule in Fig. 8 allows the term t' to evolve, i.e. $t' \not\rightarrow$. Since $T_1 \dashv T_2$, by Def. 1 it must hold that $U_1'' = U_2'' = \text{end}$. However, $U_1'' = \text{err} \neq \text{end}$ and $U_2'' = \text{err} \neq \text{end}$, which is a contradiction.

Theorem 5. If C is a rollback safe collaboration, then $C \not\rightarrow^* \mathbb{C}[\text{com_error}]$.

Proof. The proof proceeds by contradiction. Suppose that there exists an initial collaboration C that is rollback safe and such that $C \rightarrow^* \mathbb{C}[\text{com_error}]$. The erroneous collaboration com_error can be produced by applying one of the rules [E-COM1], [E-COM2], [E-LAB1] and [E-LAB2]. Let us consider the case [E-COM1], the other cases are similar. Proceeding as in the proof of Theorem 4, without loss of generality we can focus on the subterm $C_1 = (\bar{a}(x).P \mid a(y).Q)$ of C , such that $C_1 \rightarrow^* (\nu s : C_1)C'' \rightarrow (\nu s : C_1)\text{com_error}$, with $C'' = (\langle \tilde{Q}_1 \rangle \blacktriangleright P_1 \mid \langle \tilde{Q}_2 \rangle \blacktriangleright P_2)$, and $C_1 \blacktriangleright \{\bar{a} : T_1, a : T_2\}$, with $T_1 \dashv\!\!\dashv T_2$.

By Lemma 9, we have that there exist $U_1, U_2, U'_1, U'_2, U''_1$ such that $U_1, U_2, U'_1, U'_2, U''_1$ such that $(T_1, T_2) : \langle T_1 \rangle \blacktriangleright T_1 \parallel \langle T_2 \rangle \blacktriangleright T_2 \mapsto^* (T_1, T_2) : \langle \tilde{U}_1 \rangle \blacktriangleright U'_1 \parallel \langle \tilde{U}_2 \rangle \blacktriangleright U'_2 = t$ with \tilde{Q}_1 (resp. \tilde{Q}_2) in checkpoint accordance with \tilde{U}_1 (resp. \tilde{U}_2), $U'_1 \xrightarrow{! [S]} U''_1$, and for all U''_2 such that $U'_2 \xrightarrow{\tau}^* U''_2$ we have $U''_2 \xrightarrow{? [S]} t$. Thus, for all U''_2 as above, we have $t \mapsto^* (T_1, T_2) : \langle \tilde{U}_1 \rangle \blacktriangleright U'_1 \parallel \langle \tilde{U}_2 \rangle \blacktriangleright U''_2 = t'$. No rule in Fig. 8 allows the term t' to evolve, i.e. $t' \not\mapsto$, because U'_1 can only perform $! [S]$ and rule [TS-COM] cannot be applied since $U''_2 \xrightarrow{? [S]} t$. Since $T_1 \dashv\!\!\dashv T_2$, by Def. 1 it must hold that $U'_1 = \text{end}$. However, since U'_1 is able to perform an action (as it holds that $U'_1 \xrightarrow{! [S]} U''_1$), we get that it cannot be an end type, i.e. $U'_1 \neq \text{end}$, which is a contradiction.

We conclude with the session progress theorem.

Theorem 6. Let $C = (\bar{a}(x_1).P_1 \mid a(x_2).P_2)$ be a rollback safe collaboration. If $C \rightarrow^* C'$ then either $C' \rightarrow C''$ for some C'' or $C' \equiv (\nu s : C)(\langle \tilde{Q}_1 \rangle \blacktriangleright \mathbf{0} \mid \langle \tilde{Q}_2 \rangle \blacktriangleright \mathbf{0})$ for some \tilde{Q}_1 and \tilde{Q}_2 .

Proof. The proof proceeds by contradiction. Suppose that C is rollback safe and $C \rightarrow^* C'$ with $C' \not\mapsto$ and $C' \equiv (\nu s : C)(\langle \tilde{Q}_1 \rangle \blacktriangleright \mathbf{0} \mid \langle \tilde{Q}_2 \rangle \blacktriangleright \mathbf{0})$ for any Q_1 and Q_2 . The only situations that prevents C' from progressing are $C' = \mathbb{C}[\text{roll_error}]$ and $C' = \mathbb{C}[\text{com_error}]$. However, from Theorems 4 and 5, respectively, we have $C' \neq \mathbb{C}[\text{roll_error}]$ and $C' \neq \mathbb{C}[\text{com_error}]$, which is a contradiction.

C.4 Multiparty session results

All notions and concepts of our rollback recovery approach smoothly extend to the multiparty case. As consequence, all properties in Sec. 5 still hold in the extended setting. Their proofs indeed follow the same structure of the binary case and only differ for the technicalities concerning the extended definitions.

Decidability result Theorem 1. Let T_1, \dots, T_n be filled multiparty session types, checking if $\dashv\!\!\dashv (T_1, \dots, T_n)$ holds is decidable.

Proof. By definition of Compliance in Sec. B, checking $\dashv\!\!\dashv (T_1, \dots, T_n)$ consists in checking that types T'_1, \dots, T'_n of each configuration $(T^1, \dots, T^n) : \prod_{h \in \{1, \dots, n\}} \langle \tilde{U}'_h \rangle \blacktriangleright T'_h$ such that $(T^1, \dots, T^n) : \prod_{h \in \{1, \dots, n\}} \langle T_h \rangle \blacktriangleright T_h \mapsto^* (T^1, \dots, T^n) : \prod_{h \in \{1, \dots, n\}} \langle \tilde{U}'_h \rangle \blacktriangleright$

$T'_h \not\mapsto$ (i.e., type configurations that are reachable from the initial one and that cannot further evolve) are end types. Thus, to prove that the compliance check is decidable we have to show that the number of these reachable configurations is finite. Let us consider the transition system $TS = \langle \mathcal{S}, \mathcal{R} \rangle$ associated to the type configuration $t = (T^1, \dots, T^n) : \prod_{h \in \{1, \dots, n\}} \langle T_h \rangle \blacktriangleright T_h$ by the reduction semantics of types (Fig. 20): the set \mathcal{S} of states corresponds to the set of type configurations reachable from t , i.e. $\mathcal{S} = \{t' \mid t \mapsto^* t'\}$, while the set \mathcal{R} of system transitions corresponds to set of the type reductions involving configurations in \mathcal{S} , i.e. $\mathcal{R} = \{(t', t'') \in \mathcal{S} \times \mathcal{S} \mid t' \mapsto t''\}$. Hence, checking $\dashv\!\!\dashv (T_1, \dots, T_n)$ boils down to check the type configurations corresponding to the leaves (i.e., states without outgoing transitions) of TS . Specifically, given a leaf of TS corresponding to $t = (T^1, \dots, T^n) : \prod_{h \in \{1, \dots, n\}} \langle \tilde{V}_h' \rangle \blacktriangleright V_h$, we have to check if $V_1 = \dots = V_n = \text{end}$. The decidability of this check therefore depends on the finiteness of TS . This result is ensured by the fact that: (i) backward reductions connect states of TS only to previously visited states of TS (Theorem 8), and (ii) our language of types (Fig. 18) corresponds to a CCS-like process algebra without static operators (i.e., parallel and restriction operators) within recursion (see [29, Sec. 7.5]).

Reversibility results As in Appendix C.2, we will indicate with $C \xrightarrow{s} C'$ the fact that the reduction is taking place on session s .

Lemma 10 (Swap Lemma). *Let C be a collaboration and s and r two sessions. If $C \xrightarrow{s} C_1 \xrightarrow{r} C_2$ then there exists a collaboration C_3 such that $C \xrightarrow{r} C_3 \xrightarrow{s} C_2$.*

Proof. By case analysis on the reductions \xrightarrow{s} and \xrightarrow{r} .

Lemma 11. *Let C be a collaboration. If $C \mapsto^* C_1$, then for any session s in C_1 there exists a collaboration C_0 such that $C \mapsto^* C_0 \xrightarrow{s}^* C_1$ and s is never used in the trace $C \mapsto^* C_0$.*

Proof. By induction on the number n of reduction on s . If there are no reductions then the thesis banally holds. Otherwise we can take the very last reduction on s , that is the closest one to C_1 and iteratively apply Lemma 10 in order to bring it to the very end. Then we can conclude by induction on a trace with less occurrences of reductions on s .

As in Appendix C.2, thanks to Lemma 11 without loosing of generality we can focus just on a single session, say s , and to consider collaboration initial for s .

Lemma 12. *Let C be an initial collaboration such that $C \mapsto^* C_1$. If $C_1 \xrightarrow{abt} C_2$ then $C_2 \equiv C$.*

Proof. Since C is initial, without losing of generality we can assume

$$C \equiv \bar{a}[n](x).P_n \mid \prod_{i \in I} a[i](x).P_i$$

with $I = \{1, \dots, n-1\}$.

The first reduction of $C \rightarrow^* C_1$ has to be an application of rule [M-F-CON], that is

$$C \rightarrow (\nu s : (\bar{a}[n](x).P_n \mid \prod_{i \in I} a[i](x).P_i))$$

$$(\langle P_n[s[n]/x] \rangle \blacktriangleright P_n[s[n]/x] \mid \prod_{i \in I} \langle P_i[s[i]/x] \rangle \blacktriangleright P_i[s[i]/x]) = C'$$

and, by hypothesis, $C' \rightarrow^* C_1$.

Now, no matter the shape of processes in C_1 by applying rule [M-B-ABT], and possibly [M-B-STR], we will go back to C , that is $C_1 \xrightarrow{abt} C$, as desired.

Lemma 13. *Let C be a reachable collaboration, such that $C \xrightarrow{cmt} C_1$. If $C_1 \rightarrow^* C_2 \xrightarrow{roll} C_3$ and there is no commit in $C_1 \rightarrow^* C_2$, then $C_3 \equiv C_1$.*

Proof. As in Lemma 6's proof, by hypothesis, there is no commits in $C_1 \rightarrow^* C_2$, and this implies that the log part of the C_1 will never change. Hence, by applying [M-B-RLL] and [M-B-PAR] to C_2 we can conclude.

Theorem 8. *Let C_0 be an initial collaboration. If $C_0 \rightarrow^* C_1$ then $C_0 \twoheadrightarrow^* C_1$.*

Proof. By induction on the number n of backward reductions contained into $C_0 \rightarrow^* C_1$. The base case ($n = 0$) trivially holds. In the inductive case, let us take the backward reduction which is the nearest to C_0 . That is:

$$C_0 \twoheadrightarrow^* C' \rightsquigarrow C'' \rightarrow^* C_1$$

Depending whether it is an \xrightarrow{abt} or a \xrightarrow{roll} we can apply respectively Lemma 12 or Lemma 13 to obtain a forward trace of the form

$$C_0 \twoheadrightarrow^* C'' \rightarrow^* C_1$$

and we can conclude by applying the inductive hypothesis on the obtained trace which contains less backward moves with respect to the original one.

Lemma 14. *Let C_1 and C_2 be two reachable collaborations. $C_1 \rightsquigarrow C_2$ then $C_2 \twoheadrightarrow^* C_1$.*

Proof. (Sketch) Since C_1 is a reachable collaboration, we have that there exists an initial collaboration C_0 such that $C_0 \rightarrow^* C_1$. By applying Theorem 8 we can rearrange the trace such that it contains just forward transitions as follows

$$C_0 \twoheadrightarrow^* C_1$$

If the backward reduction is obtained by applying [M-B-ABT], by Lemma 12 we have $C_2 \equiv C_0$, from which the thesis trivially follows. Instead, if the backward reduction is obtained by applying [M-B-RLL], we proceed by case analysis depending on the presence of commit reductions in the trace. If they are present, we select the last of such commit, that is we can decompose the trace in the following way:

$$C_0 \twoheadrightarrow^* \xrightarrow{cmt} C_{cmt} \twoheadrightarrow^* C_1 \rightsquigarrow C_2$$

and by applying Lemma 13 we have that $C_2 \rightarrow^* \equiv C_1$ as desired. In the case there is no commit in the trace, we can conclude by noticing that $C_0 \rightarrow C_2$ must be the first reduction in $C_0 \rightarrow^* C_1$.

Lemma 15. *Let C be a reachable collaboration. If $C \xrightarrow{\text{roll}} C'$ and $C \xrightarrow{\text{roll}} C''$ then $C' \equiv C''$.*

Proof. (Sketch) Since C is a reachable collaboration, it has been generated by an initial collaboration C_0 , and by Theorem 8 we have that $C_0 \rightarrow^* C$. We distinguish two cases, whether in the trace there has been at least one commit or not. In the first case, we can decompose the trace in such a way to single out the last commit as follows:

$$C_0 \rightarrow^* C_{\text{cmt}} \rightarrow^* C$$

so that in the reduction $C_{\text{cmt}} \rightarrow^* C$ there is no commit. If from C the rollbacks $C \xrightarrow{\text{roll}} C'$ and $C \xrightarrow{\text{roll}} C''$ are triggered by the same process, the thesis trivially follows. In the other case, we have that at least two processes, say P and Q , are able to trigger a rollback. If the roll action is executed by P we have that $C_{\text{cmt}} \xrightarrow{\text{roll}} C'$. If the roll is triggered by Q we have that $C_{\text{cmt}} \xrightarrow{\text{roll}} C''$. And we can conclude by noticing that $C' \equiv C''$, as desired.

Theorem 3. *Let C be a reachable collaboration. If $C \xrightarrow{\text{cmt}} C'$ then there exists no C'' such that $C' \xrightarrow{\text{roll}} C''$ and $C'' \rightarrow^+ C$.*

Proof. We proceed by contradiction. Suppose that there exists C'' such that $C' \xrightarrow{\text{roll}} C''$ and $C'' \rightarrow^+ C$. Since C is a reachable collaboration, thanks to Theorem 8 we have that there exists an initial collaboration C_0 such that $C_0 \rightarrow^* C \xrightarrow{\text{cmt}} C'$. Since a rollback brings back the collaboration to a point before a commit, this means it has been restored a commit previous to the last one (by [M-B-RLL], indeed, only processes stored in logs can be committed). This implies that there exist at least two different commits in the trace such that

$$C_0 \rightarrow^* C_{\text{cmt}} \rightarrow^* C \xrightarrow{\text{cmt}} C' \xrightarrow{\text{roll}} C''$$

with $C'' = C_{\text{cmt}}$. We have that $C'' \rightarrow^+ C \xrightarrow{\text{cmt}} C' \rightarrow^* C_{rl}$, where C_{rl} is the collaboration that perform the roll reduction. Now since the last commit has been done by C , supposing that the commit is triggered by P^c , which evolves to P^{cl} in doing that, we have that:

$$C \equiv (\nu s : C_0)(\langle P \rangle \triangleright P^c \mid \prod_{i \in I} \langle Q_i \rangle \triangleright Q_i^c)$$

and

$$C' = (\nu s : C_0)(\langle P^{cl} \rangle \triangleright P^{cl} \mid \prod_{i \in I} \langle Q_i^c \rangle \triangleright Q_i^c)$$

Now, by hypothesis we have that $C' \rightarrow^* C_{rl}$ without any commit being present in the trace, hence:

$$C_{rl} \equiv (\nu s : C_0)(\langle P^{cl} \rangle \triangleright P^{rl} \mid \prod_{i \in I} \langle Q_i^c \rangle \triangleright Q_i^{rl})$$

By hypothesis, from C_{rl} a rollback is possible. Regardless the rollback is triggered by P_{rl} or Q_{rl} , we have that $C_{rl} \xrightarrow{\text{roll}} C'$. Now, from C' we cannot reach C ($C' \not\rightarrow^+ C$), as C' is derived from C and the rollback can only bring the collaboration back to C' . This violates the hypothesis, and hence we conclude.

Soundness results

Lemma 16. *Let $C = (\nu s : C') \prod_{i \in I} \langle \tilde{Q}_i \rangle \blacktriangleright P_i$ be a reachable collaboration, with $I = \{1, \dots, n\}$, $C' = (\bar{a}[n](x).R_n \mid \prod_{i \in I - \{n\}} a[i](x).R_i)$, and for all $i \in I$ ($\emptyset; \emptyset \vdash R_i \blacktriangleright x : T_i$), $\Vdash (T_1, \dots, T_n)$, $(\Theta_i; \Gamma_i \vdash Q_i[x/s[i]] \blacktriangleright x : V_i)$, $(\Theta'_i; \Gamma'_i \vdash P_i[x/s[i]] \blacktriangleright x : U_i)$. If $C \rightarrow (\nu s : C') \prod_{i \in I} \langle \tilde{Q}'_i \rangle \blacktriangleright P'_i$ then there exist V'_i and U'_i such that $(T^1, \dots, T^n) : \prod_{h \in I} \langle \tilde{V}_h \rangle \blacktriangleright U_h \mapsto (T^1, \dots, T^n) : \prod_{h \in I} \langle \tilde{V}'_h \rangle \blacktriangleright U'_h$ with \tilde{Q}'_i in checkpoint accordance with \tilde{V}'_h , and for all $i \in I$ ($\hat{\Theta}_i; \hat{\Gamma}_i \vdash Q'_i[x/s[i]] \blacktriangleright x : V'_i$), $(\hat{\Theta}'_i; \hat{\Gamma}'_i \vdash P'_i[x/s[i]] \blacktriangleright x : U'_i)$.*

Proof. (Sketch) We have two cases depending whether the reduction \rightarrow has forward or backward direction.

($\rightarrow = \Rightarrow$). From rule [M-F-RES], we have $\prod_{i \in I} \langle \tilde{Q}_i \rangle \blacktriangleright P_i \rightarrow \prod_{i \in I} \langle \tilde{Q}'_i \rangle \blacktriangleright P'_i$. We prove the result by case analysis on the last rule applied in the inference of the above reduction.

($\rightarrow = \rightsquigarrow$). From rule [B-RES], we have $\prod_{i \in I} \langle \tilde{Q}_i \rangle \blacktriangleright P_i \rightsquigarrow \prod_{i \in I} \langle \tilde{Q}'_i \rangle \blacktriangleright P'_i$. We prove the result by case analysis on the last rule applied in the inference of the above reduction.

Lemma 17. *Let $C = \bar{a}[n](x).P_n \mid \prod_{i \in I - \{n\}} a[i](x).P_i$, with $I = \{1, \dots, n\}$, such that $C \blacktriangleright \{\bar{a}[n] : V_n, \dots, a[1] : V_1\}$, $T_i = V_i \cdot i$ for all $i \in I$, and $\Vdash (T_n, \dots, T_1)$. If $C \rightarrow^* (\nu s : C) (\langle \tilde{R}_1 \rangle \blacktriangleright Q_1 \mid \prod_{i \in \{2, \dots, n\}} \langle \tilde{R}_i \rangle \blacktriangleright Q_i)$ and $Q_1 \xrightarrow{\text{roll}} Q'_1$, then there exist $\{U_i\}_{i \in I}$, $\{U'_i\}_{i \in I}$ and U''_1 such that $(T_1, \dots, T_n) : \prod_{i \in I} \langle \tilde{T}_i \rangle \blacktriangleright T_i \mapsto^* (T_1, \dots, T_n) : \langle \tilde{U}_1 \rangle \blacktriangleright U'_1 \parallel \prod_{i \in \{2, \dots, n\}} \langle \tilde{U}_i \rangle \blacktriangleright U'_i$ and $U'_1 \xrightarrow{\text{roll}} U''_1$.*

Proof. (Sketch) As in Lemma 8's proof, this result can be obtained by repeatedly applying Lemma 16.

Lemma 18. *Let $C = \bar{a}[n](x).P_n \mid \prod_{i \in I - \{n\}} a[i](x).P_i$, with $I = \{1, \dots, n\}$, such that $C \blacktriangleright \{\bar{a}[n] : V_n, \dots, a[1] : V_1\}$, $T_i = V_i \cdot i$ for all $i \in I$, and $\Vdash (T_n, \dots, T_1)$. If $C \rightarrow^* (\nu s : C) \prod_{i \in I} \langle \tilde{Q}_i \rangle \blacktriangleright P_i$, $P_1 \xrightarrow{\ell} P'_1$ and there is no P_j with $j \in \{2, \dots, n\}$ such that $P_j \Downarrow_{\bar{\ell}}$ with ℓ of the form $s[p][q]! \langle v \rangle$, $s[p][q]?(x)$, $s[p][q] \triangleleft l$ or $s[p][q] \triangleright l$, then there exist $\{U_i\}_{i \in I}$, $\{U'_i\}_{i \in I}$ and U''_1 such that $(T_1, \dots, T_n) : \prod_{i \in I} \langle \tilde{T}_i \rangle \blacktriangleright T_i \mapsto^* (T_1, \dots, T_n) : \prod_{i \in I} \langle \tilde{U}_i \rangle \blacktriangleright U'_i$ with \tilde{Q}_i in checkpoint accordance with \tilde{U}_i , $U'_1 \xrightarrow{\text{tl}_\Gamma(\ell)} U''_1$, with Γ sorting for typing P'_1 , and for all U''_i such that $U'_i \xrightarrow{\tau}^* U''_i$ we have $U''_i \xrightarrow{\text{tl}_\Gamma(\ell)}$.*

Proof. (Sketch) As in Lemma 9's proof, this result can be obtained by repeatedly applying Lemma 16 and reasoning by case analysis on the rule for deriving the transition $P_1 \xrightarrow{\ell} P'_1$.




Fig. 21. Producer-consumer scenario with non-speculative (a) and speculative (b) consumers.

Theorem 9. *If C is a rollback safe collaboration, then $C \not\vdash^* \mathbb{C}[\text{roll_error}]$.*

Proof. (Sketch) The proof proceeds by contradiction and relies on Lemma 17.

Theorem 10. *If C is a rollback safe collaboration, then $C \not\vdash^* \mathbb{C}[\text{com_error}]$.*

Proof. (Sketch) The proof proceeds by contradiction and relies on Lemma 18.

Theorem 7. Let $C = (\bar{a}[n](x).P_n \mid \prod_{i \in \{1, \dots, n-1\}} a[i](x).P_i)$ be a roll-safe collaboration. If $C \vdash^* C'$ then either $C' \vdash C''$ for some C'' or $C' \equiv (\nu s : C) \prod_{i \in \{1, \dots, n\}} \langle \tilde{Q}_i \rangle \blacktriangleright \mathbf{0}$ for some $\tilde{Q}_1, \dots, \tilde{Q}_n$.

Proof. The proof proceeds by contradiction. Suppose that C is rollback safe and $C \vdash^* C'$ with $C' \not\vdash$ and $C' \not\equiv (\nu s : C) \prod_{i \in \{1, \dots, n\}} \langle \tilde{Q}_i \rangle \blacktriangleright \mathbf{0}$ for any Q_1, \dots, Q_n . The only situations that prevents C' from progressing are $C' = \mathbb{C}[\text{roll_error}]$ and $C' = \mathbb{C}[\text{com_error}]$. However, from Theorems 9 and 10, respectively, we have $C' \neq \mathbb{C}[\text{roll_error}]$ and $C' \neq \mathbb{C}[\text{com_error}]$, which is a contradiction.

D cherry-pi at work on a speculative parallelism scenario

To shed light on the practical effectiveness of cherry-pi and the related notion of rollback safety, we consider in this section a simple, yet realistic, scenario concerning a form of speculative execution borrowed from [31]. In this scenario, *value speculation* is used as a mechanism for increasing parallelism, hence system performance, by predicting values of data dependencies between tasks. Whenever a value prediction is incorrect, corrective actions must be taken in order to re-execute the data consumer code with the correct data value. In this regard, as shown in [15] for a shared-memory setting, reversible execution can permit to relieve programmers from the burden of properly undoing the actions subsequent to an incorrect prediction. Here, we tailor the scenario to the channel-based communication model of session-based programming, and show how our rollback safety checking supports programmers in identifying erroneous rollback recovery settings.

In the producer-consumer scenario depicted in Fig. 21(a) the session participant P produces a value and the participant C consumes it. The data dependence between

P and C serialises their executions, thus forcing C to wait for the completion of the value production that requires a fairly long time. In the scenario in Fig. 21(b), instead, C enacts a speculative behaviour, as it predicts ahead of time the value computed by P from a partial information. By using the predicted value, C can execute speculatively and concurrently with P . When P completes the production, C validates the prediction by comparing the actual value computed by P and the predicted one; if the prediction is precise, we gain performance because the execution of C and P overlapped in time, otherwise rollback is used to move C and P back to a state that precedes the speculative behaviour, in order to re-execute C using the correct value. The behaviours of C and P can be recursively defined in order to repeat the overall execution once a value is correctly consumed.

The scenario informally described above is rendered in `cherry-pi` as

$$\overline{start}(x).P_C \mid start(y).P_P$$

where the consumer and producer processes are:

$$P_C = \mu X. x! \langle f_{req}() \rangle. x \triangleright \{ l_{spec} : x?(x_{partial} : \mathbf{str}). x?(x_{final} : \mathbf{str}). \\ \text{if } (f_{compare}(x_{partial}, x_{final})) \text{ then roll else commit. } X, \\ l_{nonSpec} : x?(x_{computed} : \mathbf{str}). \text{commit. } X \}$$

$$P_P = \mu Y. y?(y_{req} : \mathbf{str}). \\ \text{if } (f_{eval}(y_{req})) \text{ then } y < l_{spec}. y! \langle f_{partial}(y_{req}) \rangle. y! \langle f_{final}(y_{req}) \rangle. Y \\ \text{else } y < l_{nonSpec}. y! \langle f_{compute}(y_{req}) \rangle. Y$$

The producer evaluates each consumer's request in order to establish whether to provide directly the produced value or the partial information for the prediction. In the former case the consumer commits the session and both participants restart, while in the latter one the consumer commits or rolls back depending on the result of the comparison between the predicted value and the produced one.

According to our compliance check, the above collaboration is rollback safe. In case of incorrect prediction the session execution is moved back to the last checkpoint, corresponding to the successfully consumption of the previous requested value. When the producer receives again the same request, it can immediately send the already produced value. In case the first prediction is wrong, and hence no commit action is performed by the consumer yet, according to the `cherry-pi` semantics the checkpoint corresponds to the beginning of the session.

Let us consider instead the case of a producer that commits each time a value production is completed, which could apparently seem a reasonable behaviour from the producer side. The resulting collaboration, however, is not rollback safe: while the commit action in the non-speculative case does not affect the compliance between the two session participants, the other commit action overwrites the checkpoint set by the consumer, making it impossible to re-execute the consumer with the correct value. This situation, undesirable for the consumer, is detected by our compliance check.