



Less Is More: Multiparty Session Types Revisited

ALCESTE SCALAS, Imperial College London, UK

NOBUKO YOSHIDA, Imperial College London, UK

Multiparty Session Types (MPST) are a typing discipline ensuring that a message-passing process implements a given *multiparty session protocol*, without errors. In this paper, we propose a new, generalised MPST theory.

Our contribution is fourfold. (1) We demonstrate that a revision of the theoretical foundations of MPST is *necessary*: classic MPST have a limited *subject reduction* property, with inherent restrictions that are easily overlooked, and in previous work have led to flawed type safety proofs; our new theory removes such restrictions and fixes such flaws. (2) We contribute a new MPST theory that is *less* complicated, and yet *more* general, than the classic one: it does *not* require *global multiparty session types* nor *binary session type duality* – instead, it is grounded on general behavioural type-level properties, and proves type safety of many more protocols and processes. (3) We produce a detailed analysis of type-level properties, showing how, in our new theory, they allow to ensure decidability of type checking, and statically guarantee that processes enjoy, e.g., deadlock-freedom and liveness at run-time. (4) We show how our new theory can integrate type and model checking: type-level properties can be expressed in modal μ -calculus, and verified with well-established tools.

CCS Concepts: • **Theory of computation** → **Process calculi**; **Type structures**; *Verification by model checking*;

Additional Key Words and Phrases: session types, duality, deadlock-freedom, liveness

ACM Reference Format:

Alceste Scalas and Nobuko Yoshida. 2019. Less Is More: Multiparty Session Types Revisited. *Proc. ACM Program. Lang.* 3, POPL, Article 30 (January 2019), 29 pages. <https://doi.org/10.1145/3290343>

1 INTRODUCTION

Session types are a type-based framework for formalising structured communication protocols, and verifying them in concurrent message-passing programs. The original *binary session types* theory [Honda et al. 1998] addresses protocols with two participants (e.g., client and server), and is built on a notion of *duality* in interactions, inspired by linear logic [Girard 1987]; this has led to several studies on the logical foundations for session types, e.g. [Caires et al. 2016; Wadler 2014]. This approach was later generalised to *multiparty* sessions [Bettini et al. 2008; Honda et al. 2008], supporting more sophisticated protocols with *any* number of participants (two or more); correspondingly, binary duality was generalised as *multiparty consistency*, leading to studies on its logical foundations [Caires and Pérez 2016; Carbone et al. 2016, 2015].

Unfortunately, this duality-based framework has intrinsic *limitations*: the consistency requirement is not satisfied by many multiparty protocols – even surprisingly simple ones. Such limitations are subtle: in this paper, we show that they have been overlooked or wrongly bypassed in several previous works, leading to MPST extensions that are *no longer correct*, and have flawed *subject reduction* proofs. Then, we provide a solution: a new, generalised MPST theory that subsumes

Authors' addresses: Alceste Scalas, Imperial College London, UK, alceste.scalas@imperial.ac.uk; Nobuko Yoshida, Imperial College London, UK, n.yoshida@imperial.ac.uk.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2019 Copyright held by the owner/author(s).

2475-1421/2019/1-ART30

<https://doi.org/10.1145/3290343>

classic MPST under a new theoretical foundation, removes its limitations, fixes the aforementioned flaws, and supports a richer set of multiparty protocols and processes.

The Multiparty Session Types (MPST) framework. Bettini et al. [2008]; Honda et al. [2008] introduce the seminal notion of *global types*, which describe multiparty conversations from a global perspective. MPST verification follows a top-down approach based on *endpoint projections*:

- (1) a *multiparty protocol* is formalised as a *global type* G , providing a bird's eye view on the interactions between two or more *roles*;
- (2) G is *projected* onto a set of *endpoint (local) session types* (one per role); and
- (3) session types are assigned to *communication channels*, used by *MPST processes* that can be written and type-checked separately.

E.g., the global type G below models a protocol (based on OAuth 2.0 [OAuth Working Group 2012]) between service \mathbf{s} , client \mathbf{c} , and authorisation server \mathbf{a} :

$$G = \mathbf{s} \rightarrow \mathbf{c}: \left\{ \begin{array}{l} \text{login} . \mathbf{c} \rightarrow \mathbf{a}: \text{passwd}(\text{Str}) . \mathbf{a} \rightarrow \mathbf{s}: \text{auth}(\text{Bool}) . \text{end} , \\ \text{cancel} . \mathbf{c} \rightarrow \mathbf{a}: \text{quit} . \text{end} \end{array} \right\} \quad (1)$$

The protocol of G says that the **service** sends to the **client** *either* a request to **login**, or **cancel**; in the first case, **c** continues by sending **passwd** (carrying a **String**) to the **authorisation server**, who in turn sends **auth** to **s** (with a **Boolean**, telling whether the client is authorised), and the session **ends**; in the second case, **c** sends **quit** to **a**, and the session **ends**. The *projections of G* describe the local I/O actions (i.e., the interfaces) that programs must implement to play the roles in G :

$$S_s = \mathbf{c} \oplus \left\{ \begin{array}{l} \text{login} . \mathbf{a} \& \text{auth}(\text{Bool}) , \\ \text{cancel} \end{array} \right\} \quad S_c = \mathbf{s} \& \left\{ \begin{array}{l} \text{login} . \mathbf{a} \oplus \text{passwd}(\text{Str}) , \\ \text{cancel} . \mathbf{a} \oplus \text{quit} \end{array} \right\} \quad S_a = \mathbf{c} \& \left\{ \begin{array}{l} \text{passwd}(\text{Str}) . \mathbf{s} \oplus \text{auth}(\text{Bool}) , \\ \text{quit} \end{array} \right\} \quad (2)$$

Here, S_s, S_c, S_a are *session types*, obtained by projecting G resp. onto $\mathbf{s}, \mathbf{c}, \mathbf{a}$ (for brevity, we omit final **ends**). S_s represents the interface of **s** in G : it must send (\oplus) to **c** either **login** or **cancel**; in the first case, **s** must then receive ($\&$) message **auth(Bool)** from **a**, and the session ends; otherwise, in the second case, the session just ends. Types S_c and S_a follow the same intuition. The multiparty session type system assigns the types in (2) to *channels*, and checks that *endpoint programs* use them correctly: e.g., the program implementing the **service** is checked against S_s , and the programs implementing **c/a** against S_c/S_a . Endpoint programs, in turn, are formalised as *processes* in a π -calculus extended with multiparty communication primitives. Variations of this framework have been implemented in numerous programming languages (surveyed in Ancona et al. [2017]; Gay and Ravara [2017]), allowing to develop distributed applications with guaranteed protocol conformance.

Limitations and Theoretical Issues of MPST. Theories and implementations based on MPST crucially require “correct by construction” protocols that do not cause deadlocks nor communication errors when endpoint programs interact. This is achieved by imposing *well-formedness* conditions to global types, and *consistency* restrictions when processes are type-checked.

However, such restrictions introduce rather serious problems when proving *subject reduction* – i.e., when proving that typed processes only reduce to typed processes, and thus, no (untypable) error state can be reached (“*typed processes never go wrong*”). Usually, one expects a statement like:

$$\Gamma \vdash P \text{ and } P \rightarrow P' \text{ implies } \exists \Gamma' : \Gamma' \vdash P' \quad (3)$$

where $\Gamma \vdash P$ is a typing judgement stating that process P abides by the typing context Γ , which can map, e.g., the communication channels c_s, c_c, c_a to the types S_s, S_c, S_a in (2).

Unfortunately, (3) *is wrong*. If we take Γ without any constraint as in (3), it might contain types like $\mathbf{c} \oplus \mathbf{m}(\text{Str}) . \text{end}$ and $\mathbf{s} \& \mathbf{m}(\text{Int}) . \text{end}$, and they could type a parallel process $P = P_1 \mid P_2$, where P_1 and P_2 interact according to the types, with P_1 sending a message m (“Hello”) (carrying a **String**), and P_2 receiving m but using its payload as an **Integer**. In this case, P would reduce to a “wrong”

and untypable P' (see also [Coppo et al. 2015a, p. 163], and §3 later on): this means that (3) does *not* hold. For this reason, the MPST theory requires the aforementioned *consistency* restriction, and its actual subject reduction statement reads:

$$\Gamma \vdash P \text{ with } \Gamma \text{ consistent and } P \rightarrow P' \quad \text{implies} \quad \exists \Gamma' \text{ consistent: } \Gamma \rightarrow^* \Gamma' \text{ and } \Gamma' \vdash P' \quad (4)$$

(where $\Gamma \rightarrow^* \Gamma'$ denotes typing context reductions). Consistency is a syntactic constraint ensuring that the potential output messages of each role match the input capabilities of their recipient; as noted above, this requirement was developed by generalising the notion of *binary session duality* [Honda et al. 1998]. However, due to this binary session heritage, multiparty consistency is:

- (1) **overly restrictive.** Consistency does *not* hold for many protocols: even the simple authorisation protocol in (1)/(2) above is *not* consistent. Hence, for such protocols, the MPST framework cannot prove type safety of *any* process, because (4) holds vacuously;
- (2) **inflexible and error-prone.** Some MPST works, e.g. [Deniélou et al. 2012; Deniélou and Yoshida 2012; Yoshida et al. 2010], propose richer global types with flexible well-formedness conditions — but either overlook the consistency requirement, or fail to realise that their extensions do *not* satisfy it. Hence, their subject reduction theorems do not hold (like (3)), or hold vacuously (as above); and worryingly, such results are reused in later works and implementations (more details in §8).

These two claims are based on technical arguments, that we develop in §3. They clearly undermine the expressiveness and applicability of MPST: when the theory cannot ensure type safety for a given protocol, MPST-based implementations should either reject it (thus being overly restrictive), or forfeit the guaranteed absence of run-time errors. To solve these problems, we pose the questions:

Can we remove the duality/consistency requirements of MPST?

Can we use, instead, more flexible properties of session types, thus enlarging the subject reduction property, and the set of provably type-safe processes?

To answer positively, we need a new MPST theory that is *not* rooted in binary session duality — but has more general foundations, that still support duality as a special case.

Contributions. We present a *new theory of multiparty session types*. Its novel theoretical foundations leverage a weak *behavioural safety* invariant that, for the first time, eschews the limitations of duality/consistency, and allows to obtain much more general results than classic MPST.

We summarise MPST definitions and typing rules in §2, highlighting where our new theory diverges from the classic (§2.3): i.e., when establishing the prerequisites for proving type safety.

- (1) We explain how classic MPST establish such prerequisites: i.e., by imposing consistency/duality. We uncover that the resulting severe limitations lead to subtle theoretical issues (§3).
- (2) We present our new MPST theory (§4), with a much weaker prerequisite: a *safety* invariant, *not* depending on global types, *nor* needing projection/duality/consistency from classic MPST.
- (3) By removing consistency, we rebuild the theoretical foundations of MPST on a more general basis. Our rebuilding subsumes classic MPST works, and fixes their theoretical issues, by producing more general typing rules, with just small visible differences (Remark 5.12).
- (4) We design our new type system to be parametric: its safety invariant is abstracted as a parameter φ . We show that φ can be fine-tuned to ensure decidability of type-checking, and statically enforce various run-time properties on processes — e.g., liveness (§5.3, §5.4, §5.5).
- (5) The parameter φ can be a *behavioural* property: this allows for a novel integration of type/model checking techniques for MPST. We show how to express φ as a modal μ -calculus formula, and verify type-level properties via model checking, using the paper's companion artifact (§6). Via point 4 above, the model-checked properties transfer to processes.

- (6) Our theory extends to *asynchronous* communication, to handle richer protocols and programs. Asynchrony makes φ (and type checking) undecidable; still, we present various ways to achieve decidable type checking, with methods based e.g. on communicating automata (§7).

NOTE: the technical report [Scalas and Yoshida 2018a] contains more technical details, proofs, and discussion on related work.

2 MULTIPARTY SESSION TYPES

This section describes the multiparty session π -calculus (§2.1), its types, and typing rules (§2.2). Our streamlined formulation is based on Coppo et al. [2015a] and Scalas et al. [2017a], i.e., the most common in literature; we include subtyping [Dezani-Ciancaglini et al. 2015], to later study its crucial influence on the behavioural properties of types and processes (§5).

Crucially, in this section we leave one typing rule under-specified: the rule for session restriction. The reason is explained in §2.3: the exact form of this rule strictly depends on the theoretical foundations that allow to prove type safety – and the choice of such foundations is the crossroads where our new theory (§4) departs from classic MPST (§3).

2.1 The Multiparty Session π -Calculus

The multiparty session π -calculus models processes that interact via *multiparty channels*. We give a streamlined definition, sufficient for our developments. Extensions with, e.g., ground values (booleans, strings,...), or conditionals, are standard and orthogonal; we use them in examples.

Definition 2.1. The **multiparty session π -calculus** syntax is defined as follows:

$$\begin{array}{ll}
 c, d ::= x \mid s[\mathbf{p}] & \text{(variable, channel with role } \mathbf{p} \text{)} \\
 P, Q ::= \mathbf{0} \mid P \mid Q \mid (vs)P & \text{(inaction, composition, restriction)} \\
 & c[\mathbf{q}] \oplus m\langle d \rangle . P \quad \text{(selection towards role } \mathbf{q} \text{)} \\
 & c[\mathbf{q}] \sum_{i \in I} m_i(x_i) . P_i \quad \text{(branching from role } \mathbf{q} \text{ with } I \neq \emptyset \text{)} \\
 & \mathbf{def} D \mathbf{in} P \mid X(\bar{c}) \mid \mathbf{err} \quad \text{(process definition, process call, error)} \\
 D ::= X(\bar{x}) = P & \text{(declaration of process variable } X \text{)}
 \end{array}$$

Restriction, branching and declarations act as binders, as expected; $\text{fc}(P)$ is the set of *free channels with roles* in P , and $\text{fv}(P)$ is the set of *free variables* in P . We adopt a form of Barendregt convention: bound sessions and process variables are assumed pairwise distinct, and different from free ones.

A **channel** c can be either a variable or a **channel with role** $s[\mathbf{p}]$, i.e., a multiparty communication endpoint whose user plays role \mathbf{p} in the session s . The **inaction** $\mathbf{0}$ represents a terminated process (and is often omitted). The **parallel composition** $P \mid Q$ represents two processes that can execute concurrently, and potentially communicate. The **session restriction** $(vs)P$ declares a new session s with scope limited to process P . Process $c[\mathbf{q}] \oplus m\langle d \rangle . P$ performs a **selection (internal choice)** towards role \mathbf{q} , using the channel c : the *message label* m is sent with the *payload* channel d , and the execution continues as P . Dually, the **branching (external choice)** $c[\mathbf{q}] \sum_{i \in I} m_i(x_i) . P_i$ uses channels c to wait for a message from role \mathbf{q} : if a message label m_k with payload d is received (for some $k \in I$), then the execution continues as P_k , with x_k replaced by d . Note that variable x_i is bound with scope P_i . **Process definition** $\mathbf{def} X(\bar{x}) = P \mathbf{in} Q$ and **process call** $X(\bar{c})$ model recursion: the call invokes X by expanding it into P , and replacing its formal parameters with the actual ones. **err** denotes the **error process**. Note that our simplified syntax does not have “pure” input/output prefixes: they can be easily encoded as singleton branch/selection.

Definition 2.2 (Semantics). A **reduction context** \mathbb{C} is: $\mathbb{C} ::= \mathbb{C} \mid P \mid (vs)\mathbb{C} \mid \mathbf{def} D \mathbf{in} \mathbb{C} \mid []$

$$\begin{aligned}
\text{[R-COMM]} \quad & s[\mathbf{p}][\mathbf{q}] \sum_{i \in I} m_i(x_i).P_i \mid s[\mathbf{q}][\mathbf{p}] \oplus m_k \langle s'[\mathbf{r}] \rangle . Q \rightarrow P_k \{s'[\mathbf{r}]/x_k\} \mid Q \quad \text{if } k \in I \\
\text{[R-X]} \quad & \mathbf{def} X(x_1, \dots, x_n) = P \mathbf{in} (X \langle s_1[\mathbf{p}_1], \dots, s_n[\mathbf{p}_n] \rangle \mid Q) \\
& \rightarrow \mathbf{def} X(x_1, \dots, x_n) = P \mathbf{in} (P \{s_1[\mathbf{p}_1]/x_1\} \cdots \{s_n[\mathbf{p}_n]/x_n\} \mid Q) \\
\text{[R-CTX]} \quad & P \rightarrow P' \text{ implies } \mathbb{C}[P] \rightarrow \mathbb{C}[P'] \\
\text{[R-ERR]} \quad & s[\mathbf{p}][\mathbf{q}] \sum_{i \in I} m_i(x_i).P_i \mid s[\mathbf{q}][\mathbf{p}] \oplus m \langle s'[\mathbf{r}] \rangle . Q \rightarrow \mathbf{err} \quad \text{if } \forall i \in I : m_i \neq m
\end{aligned}$$

Fig. 1. MPST π -calculus semantics, defined up-to standard structural congruence [Scalas and Yoshida 2018a].

Reduction \rightarrow is inductively defined in Fig. 1, up-to a standard **structural congruence** \equiv [Scalas and Yoshida 2018a] including α -conversion. We say that P **has an error** iff, for some \mathbb{C} , $P = \mathbb{C}[\mathbf{err}]$.

In Def. 2.2, the **reduction context** \mathbb{C} defines a process with a single hole $[\]$, occurring in place of some subterm P . The **communication rule** [R-COMM] says that the parallel composition of a branching and a selection process, both operating on the same session s respectively as roles \mathbf{p} and \mathbf{q} , reduces to the corresponding continuations, with the sent channel being substituted on the receiver side. The **process call rule** [R-X] allows to invoke the process P in the definition of X by creating a copy of P , and replacing the formal parameters x_i with actual parameters, i.e., channels with role $s_i[\mathbf{p}_i]$. The standard **context rule** [R-CTX] says that reduction can happen under parallel composition, restriction and process definition (cf. definition of \mathbb{C}). Finally, the **error rule** [R-ERR] says that a parallel composition of mismatching selection and branching processes reduces to **err**: intuitively, it models a scenario where a process implementing role \mathbf{q} is trying to send m to another process implementing \mathbf{p} — who is indeed waiting for an input, but does not expect to receive m .

Example 2.3. The following process interacts on session s using channels with role $s[\mathbf{s}]$, $s[\mathbf{c}]$, $s[\mathbf{a}]$, to play resp. roles \mathbf{s} , \mathbf{c} , \mathbf{a} . For brevity, we omit irrelevant message payloads.

$$(vs) (P_s \mid P_c \mid P_a) \quad \text{where: } \begin{cases} P_s = s[\mathbf{s}][\mathbf{c}] \oplus \text{cancel} \\ P_c = s[\mathbf{c}][\mathbf{s}] \sum \{ \text{login}.s[\mathbf{c}][\mathbf{a}] \oplus \text{passwd} \langle \text{"XYZ"} \rangle, \text{cancel}.s[\mathbf{c}][\mathbf{a}] \oplus \text{quit} \} \\ P_a = s[\mathbf{a}][\mathbf{c}] \sum \{ \text{passwd}(y).s[\mathbf{a}][\mathbf{s}] \oplus \text{auth}(y = \text{"secret"}), \text{quit} \} \end{cases}$$

Here, $(vs) (P_s \mid P_c \mid P_a)$ is the parallel composition of processes P_s, P_c, P_a in the scope of session s . In P_s , “ $s[\mathbf{s}][\mathbf{c}] \oplus \text{cancel}$ ” means: use $s[\mathbf{s}]$ to send cancel to \mathbf{c} . Process P_c uses $s[\mathbf{c}]$ to receive login or cancel from \mathbf{s} ; then, in the first case it uses $s[\mathbf{c}]$ to send passwd to \mathbf{a} ; in the second case, it uses $s[\mathbf{c}]$ to send quit to \mathbf{a} . By Def. 2.2, we have the reductions:

$$(vs) (P_s \mid P_c \mid P_a) \rightarrow (vs) (0 \mid s[\mathbf{c}][\mathbf{a}] \oplus \text{quit} \mid P_a) \rightarrow (vs) (0 \mid 0 \mid 0) \equiv 0$$

2.2 Types, Subtypes, and Typing

Session types (Def. 2.4) describe the intended use of communication channels in the MPST π -calculus (Def. 2.1); channels are mapped to their respective type by session typing contexts (Def. 2.6).

Definition 2.4. The syntax of **multiparty session types** is:

$$S, T ::= \mathbf{p} \&_{i \in I} m_i(S_i).S'_i \mid \mathbf{p} \oplus_{i \in I} m_i(S_i).S'_i \mid \mathbf{end} \mid \mu t.S \mid \mathbf{t} \quad \text{with } I \neq \emptyset, \text{ and } m_i \text{ pairwise distinct}$$

We require types to be closed, and recursion variables to be guarded.

The **branching type** (or **external choice**) $\mathbf{p} \&_{i \in I} m_i(S_i).S'_i$ says that a channel must be used to receive from \mathbf{p} one input of the form $m_i(S_i)$, for any $i \in I$ chosen by \mathbf{p} , where m_i are *message labels* and S_i are *message payload types*; then, the channel must be used following the *continuation type* S'_i . The **selection type** (or **internal choice**) $\mathbf{p} \oplus_{i \in I} m_i(S_i).S'_i$, instead, requires to use a channel to perform one output $m_i(S_i)$ towards \mathbf{p} , for some $i \in I$, and continue using the channel according

to S'_i . Type **end** describes a **terminated** channel allowing no further inputs/outputs. Type $\mu t.S$ models **recursion**: μ binds the **recursion variable** t in S . The guardedness requirement ensures that recursive types are *contractive*: i.e., in $\mu t.S$ we have $S \neq t'$ for all t' . For brevity, we often omit the trailing **end** in types, and **end**-typed message payloads: e.g., $\mathbf{p} \oplus \mathbf{m}$ stands for $\mathbf{p} \oplus \mathbf{m}(\mathbf{end}).\mathbf{end}$.

In Def. 2.5 below, we define the *multiparty session subtyping* relation [Dezani-Ciancaglini et al. 2015].¹ Intuitively, Def. 2.5 says that a type S is smaller than S' when S is “less demanding” than S' – i.e., when S imposes to support less external choices and allows to perform more internal choices. Session subtyping is used in the type system to augment its flexibility.

Definition 2.5. The **session subtyping** \leq is coinductively defined:

$$\frac{\forall i \in I \quad S_i \leq T_i \quad S'_i \leq T'_i}{\mathbf{p} \&_{i \in I} m_i(S_i).S'_i \leq \mathbf{p} \&_{i \in I \cup J} m_i(T_i).T'_i} \text{[SUB-}\&\text{]} \quad \frac{\forall i \in I \quad T_i \leq S_i \quad S'_i \leq T'_i}{\mathbf{p} \oplus_{i \in I \cup J} m_i(S_i).S'_i \leq \mathbf{p} \oplus_{i \in I} m_i(T_i).T'_i} \text{[SUB-}\oplus\text{]}$$

$$\frac{}{\mathbf{end} \leq \mathbf{end}} \text{[SUB-end]} \quad \frac{S\{\mu t.S/t\} \leq T}{\mu t.S \leq T} \text{[SUB-}\mu\text{L]} \quad \frac{S \leq T\{\mu t.T/t\}}{S \leq \mu t.T} \text{[SUB-}\mu\text{R]}$$

In Def. 2.5, rules [SUB- $\&$]/[SUB- \oplus] define **subtyping on branch/select types**: [SUB- $\&$] is covariant in both the carried types and in the number of branches, whereas [SUB- \oplus] is contravariant in both: this formalises the intuition of a smaller type having less external choices, and more internal choices. By rule [SUB-end], **end** is only subtype of itself. The **recursion rules** [SUB- μ L]/[SUB- μ R] relate types up-to their unfoldings, as usual for coinductive subtyping [Pierce 2002, Ch. 21].

Definition 2.6 (Typing Contexts). Θ denotes a partial mapping from process variables to n -tuples of types, and Γ denotes a partial mapping from channels to types, defined as:

$$\Theta ::= \emptyset \mid \Theta, X:S_1, \dots, S_n \quad \Gamma ::= \emptyset \mid \Gamma, x:S \mid \Gamma, s[\mathbf{p}]:S$$

The *composition* Γ_1, Γ_2 is defined iff $\text{dom}(\Gamma_1) \cap \text{dom}(\Gamma_2) = \emptyset$.

We write $s \notin \Gamma$ iff $\forall \mathbf{p} : s[\mathbf{p}] \notin \text{dom}(\Gamma)$ (i.e., session s does not occur in Γ).

We write $\text{dom}(\Gamma) = \{s\}$ iff $\forall c \in \text{dom}(\Gamma)$ there is \mathbf{p} such that $c = s[\mathbf{p}]$ (i.e., Γ only contains session s).

We write $\Gamma \leq \Gamma'$ iff $\text{dom}(\Gamma) = \text{dom}(\Gamma')$ and $\forall c \in \text{dom}(\Gamma): \Gamma(c) \leq \Gamma'(c)$.

The type system uses two kinds of typing contexts: Θ to assign an n -tuple of types to each process variable X (one type per argument), and Γ to map variables and channels with roles to session types. Together, they are used in judgements of the following form:

$$\Theta \cdot \Gamma \vdash P \quad (\text{with } \Theta \text{ omitted when empty}) \quad (5)$$

meaning: “given the process types in Θ , P uses its variables and channels *linearly* according to Γ .”

The **typing judgement** (5) is inductively defined by the rules in Fig. 2. For convenience, we type-annotate channels bound by process definitions and restrictions.

The first three rules in Fig. 2 define auxiliary judgements. By [T-X], $\Theta \vdash X:S_1, \dots, S_n$ holds if Θ maps X to an n -tuple of types S_1, \dots, S_n . By [T-SUB], $\Gamma \vdash c:S'$ holds if Γ only contains *one* entry $c:S$ with $S \leq S'$: i.e., when typing processes, [T-SUB] allows to use a channel of type S whenever a channel with a larger type S' is needed, as per Liskov and Wing [1994]’s substitution principle; note that Def. 2.5 relates types up-to unfolding, hence [T-SUB] makes the type system *equi-recursive* [Pierce 2002, Ch. 21]. Finally, $\text{end}(\Gamma)$ holds if Γ ’s entries are **end**-typed (under [T-SUB]).

The other rules in Fig. 2 define the process typing judgement in (5). The **termination rule** [T-0] says that $\mathbf{0}$ is typed if all channels in Γ are **end**-typed. By the **process definition rule** [T-def], **def** $X(\bar{x}) = P$ in Q is typed if P uses the arguments x_1, \dots, x_n according to S_1, \dots, S_n , and the latter

¹Our \leq is inverted w.r.t. the “process-oriented” subtyping of Dezani-Ciancaglini et al. [2015] because, for convenience, we use the “channel-oriented” order of Gay and Hole [2005]; Scalas et al. [2017a]. For a thorough comparison, see [Gay 2016].

$$\begin{array}{c}
\frac{\Theta(X) = S_1, \dots, S_n}{\Theta \vdash X:S_1, \dots, S_n} \text{ [T-X]} \quad \frac{S \leq S'}{c:S \vdash c:S'} \text{ [T-SUB]} \quad \frac{\forall i \in 1..n \quad c_i:S_i \vdash c_i:\mathbf{end}}{\mathbf{end}(c_1:S_1, \dots, c_n:S_n)} \text{ [T-end]} \\
\frac{\mathbf{end}(\Gamma)}{\Theta \cdot \Gamma \vdash \mathbf{0}} \text{ [T-0]} \quad \frac{\Theta, X:S_1, \dots, S_n \cdot x_1:S_1, \dots, x_n:S_n \vdash P \quad \Theta, X:S_1, \dots, S_n \cdot \Gamma \vdash Q}{\Theta \cdot \Gamma \vdash \mathbf{def} X(x_1:S_1, \dots, x_n:S_n) = P \mathbf{in} Q} \text{ [T-def]} \\
\frac{\Theta \vdash X:S_1, \dots, S_n \quad \mathbf{end}(\Gamma_0) \quad \forall i \in 1..n \quad \Gamma_i \vdash c_i:S_i}{\Theta \cdot \Gamma_0, \Gamma_1, \dots, \Gamma_n \vdash X(c_1, \dots, c_n)} \text{ [T-X]} \\
\frac{\Gamma_1 \vdash c:\mathbf{q}\&_{i \in I} m_i(S_i) \cdot S'_i \quad \forall i \in I \quad \Theta \cdot \Gamma, y_i:S_i, c:S'_i \vdash P_i}{\Theta \cdot \Gamma, \Gamma_1 \vdash c[\mathbf{q}] \sum_{i \in I} m_i(y_i) \cdot P_i} \text{ [T-&]} \\
\frac{\Gamma_1 \vdash c:\mathbf{q}\oplus m(S) \cdot S' \quad \Gamma_2 \vdash d:S \quad \Theta \cdot \Gamma, c:S' \vdash P}{\Theta \cdot \Gamma, \Gamma_1, \Gamma_2 \vdash c[\mathbf{q}] \oplus m(d) \cdot P} \text{ [T-\oplus]} \quad \frac{\Theta \cdot \Gamma_1 \vdash P_1 \quad \Theta \cdot \Gamma_2 \vdash P_2}{\Theta \cdot \Gamma_1, \Gamma_2 \vdash P_1 \mid P_2} \text{ [T-|]} \\
\frac{\Gamma' = \{s[\mathbf{p}]:S_{\mathbf{p}}\}_{\mathbf{p} \in I} \quad \varphi(\Gamma') \quad s \notin \Gamma \quad \Theta \cdot \Gamma, \Gamma' \vdash P}{\Theta \cdot \Gamma \vdash (vs:\Gamma') P} \text{ [T-\nu]} \quad \text{where } \varphi \text{ is a typing context property}
\end{array}$$

Fig. 2. Multiparty session typing rules. Rule [T-ν] for session restriction is discussed in §2.3.

is the type of X when typing both P and Q : this means that P can refer to X , and this allows to type recursive processes. By the **process call rule** [T-X], $X(\bar{c})$ is typed if the types of \bar{c} match those of the formal parameters of X , and any unused channel (in Γ_0) is **end**-typed: this preserves linearity by ensuring that channels requiring more inputs/outputs cannot be forgotten. By the **branching rule** [T-&], $c[\mathbf{q}] \sum_{i \in I} m_i(y_i) \cdot P_i$ is typed if c has type S , where S is an external choice from \mathbf{q} , with the same branching labels m_i . The **selection rule** [T-⊕] says that $c[\mathbf{q}] \oplus m(d) \cdot P$ is typed if c has type S , where S is an internal choice towards \mathbf{q} with message label m . By the **parallel rule** [T-|], two parallel processes are typed by splitting the context in the premises. The **session restriction rule** [T-ν] deserves special attention: we discuss it in §2.3.

Example 2.7. Take the processes from Ex.2.3, and the types S_s, S_c, S_a from §1, eq. (2). With the rules in Fig.2, we have the following typing derivation:

$$\frac{\frac{\vdots}{s[\mathbf{s}]:S_s \vdash P_s} \quad \frac{\vdots}{s[\mathbf{c}]:S_c \vdash P_c}}{s[\mathbf{s}]:S_s, s[\mathbf{c}]:S_c \vdash P_s \mid P_c} \text{ [T-|]} \quad \frac{\vdots}{s[\mathbf{a}]:S_a \vdash P_a}}{\Gamma \vdash P_s \mid P_c \mid P_a} \text{ [T-|]} \quad \text{where } \Gamma = s[\mathbf{s}]:S_s, s[\mathbf{c}]:S_c, s[\mathbf{a}]:S_a$$

The process $P_s \mid P_c \mid P_a$ is typed by rule [T-|], that splits the typing context linearly ensuring that a channel is not used by two parallel sub-processes. In the omitted part of the derivation, processes P_s, P_c, P_a are typed separately, using rules [T-⊕]/[T-&]: each process uses one of the channels with role $s[\mathbf{s}], s[\mathbf{c}], s[\mathbf{a}]$, according to the type S_s, S_c, S_a , respectively.

We conclude with the transitions/reductions of typing contexts (Def. 2.8): intuitively, they abstract the message exchanges that might occur over typed channels. We adopt a standard formulation, with two adaptations: we compare payloads using \leq (to cater for subtyping), and we specify transition labels for inputs, outputs, and communication.

Definition 2.8. Let α have the form $s:\mathbf{p}\&\mathbf{q}:m(S)$, or $s:\mathbf{p}\oplus\mathbf{q}:m(S)$, or $s:\mathbf{p}.\mathbf{q}:m$ (for any roles \mathbf{p}, \mathbf{q} , message label m , and type S). The *typing context transition* $\xrightarrow{\alpha}$ is inductively defined by the rules:

$$\begin{array}{c}
\frac{k \in I}{s[\mathbf{p}]:\mathbf{q}\oplus_{i \in I} m_i(S_i).S'_i \xrightarrow{s:\mathbf{p}\oplus\mathbf{q}:m_k(S_k)} S'_k} \quad [\Gamma\text{-}\oplus] \qquad \frac{k \in I}{s[\mathbf{p}]:\mathbf{q}\&_{i \in I} m_i(S_i).S'_i \xrightarrow{s:\mathbf{p}\&\mathbf{q}:m_k(S_k)} S'_k} \quad [\Gamma\text{-}\&] \\
\frac{\Gamma_1 \xrightarrow{s:\mathbf{p}\oplus\mathbf{q}:m(S)} \Gamma'_1 \quad \Gamma_2 \xrightarrow{s:\mathbf{q}\&\mathbf{p}:m(T)} \Gamma'_2 \quad S \leq T}{\Gamma_1, \Gamma_2 \xrightarrow{s:\mathbf{p}, \mathbf{q}:m} \Gamma'_1, \Gamma'_2} \quad [\Gamma\text{-COMM}] \qquad \frac{\Gamma, c:S\{\mu t.S/t\} \xrightarrow{\alpha} \Gamma'}{\Gamma, c:\mu t.S \xrightarrow{\alpha} \Gamma'} \quad [\Gamma\text{-}\mu] \quad \frac{\Gamma \xrightarrow{\alpha} \Gamma'}{\Gamma, c:S \xrightarrow{\alpha} \Gamma', c:S} \quad [\Gamma\text{-CONG}]}
\end{array}$$

We write $\Gamma \xrightarrow{\alpha}$ iff $\Gamma \xrightarrow{\alpha} \Gamma'$ for some Γ' . The *reduction* $\Gamma \rightarrow \Gamma'$ is defined iff $\Gamma \xrightarrow{s:\mathbf{p}, \mathbf{q}:m} \Gamma'$ for some $s, \mathbf{p}, \mathbf{q}, m$. We write $\Gamma \rightarrow$ iff $\Gamma \rightarrow \Gamma'$ for some Γ' , and $\Gamma \not\rightarrow$ for its negation (i.e., when there is no Γ' such that $\Gamma \rightarrow \Gamma'$). We define \rightarrow^* as the reflexive and transitive closure of \rightarrow .

By $[\Gamma\text{-}\oplus]/[\Gamma\text{-}\&]$ in Def. 2.8, a typing context entry can transition to one of its continuations by firing an output label of the form $s:\mathbf{p}\oplus\mathbf{q}:m(S)$ (in case of selection types), or an input label of the form $s:\mathbf{p}\&\mathbf{q}:m(S)$ (in case of branching types). Rule $[\Gamma\text{-COMM}]$ models type-level communication: e.g., it allows two entries $s[\mathbf{p}]:S_{\mathbf{p}}, s[\mathbf{q}]:S_{\mathbf{q}}$ to interact, provided that: (1) $S_{\mathbf{p}}$ is a selection towards \mathbf{q} (with a corresponding output transition); (2) $S_{\mathbf{q}}$ is a branching from \mathbf{p} (with a corresponding input transition); and (3) they are firing a common message label m , and the carried type S sent by $S_{\mathbf{p}}$ is subtype of the type T expected by $S_{\mathbf{q}}$. When all such conditions hold, $s[\mathbf{p}]:S_{\mathbf{p}}, s[\mathbf{q}]:S_{\mathbf{q}}$ transition to the respective continuations, by firing a communication label $s:\mathbf{p}, \mathbf{q}:m$ that records the session s , and the message sender \mathbf{p} , recipient \mathbf{q} , and label m (the payload types are discarded).

In the rest of the paper, we will mostly use the unlabelled reduction $\Gamma \rightarrow \Gamma'$, which means that Γ transitions to Γ' through some communication. The labelled transitions will be reprised in §5.

2.3 Towards Subject Reduction and Type Safety

In §1, we mentioned that a process naively typed with an arbitrary Γ can “go wrong.” Indeed, by themselves, the typing rules in Fig. 2 do *not* guarantee type safety, as shown by the following (counter-)example:

$$s[\mathbf{p}]:\mathbf{q}\oplus\text{foo}(\text{end}), s[\mathbf{q}]:\mathbf{p}\&\text{bar}(\text{end}), s'[\mathbf{r}]:\text{end} \vdash s[\mathbf{p}][\mathbf{q}]\oplus\text{foo}(s'[\mathbf{r}]) \mid s[\mathbf{q}][\mathbf{p}]\sum\text{bar}(x) \rightarrow \text{err} \quad (6)$$

Intuitively, the problem of this typing judgement can be seen in its typing context: the type of $s[\mathbf{p}]$ outputs `foo` to \mathbf{q} , but the type of $s[\mathbf{q}]$ expects `bar`. This means that we need a criterion to reject (6).

Importantly, the same criterion must be applied for **typing session restriction**. Consider rule $[\Gamma\text{-}\nu]$ in Fig. 2: it types a restricted session s with Γ' , provided that (1) Γ' only contains channels with roles belonging to s ; (2) the restricted s does not occur in the remaining context Γ (to avoid clashes); and (3) Γ' satisfies a (yet unspecified) property φ . How should we define φ ? It cannot be always true, because we would have this counterexample to type-safety, where Γ is the context in (6):

$$\emptyset \vdash (\nu s:\Gamma) (s[\mathbf{p}][\mathbf{q}]\oplus\text{foo}(s'[\mathbf{r}]) \mid s[\mathbf{q}][\mathbf{p}]\sum\text{bar}(x)) \rightarrow (\nu s) \text{err} \quad (\text{by (6) and rule } [\text{R-CTX}] \text{ in Fig. 1}) \quad (7)$$

To achieve type safety, we want the process in (7) to be untypable – which means that, when type-checking $(\nu s:\Gamma) \dots$, we must ensure that φ in rule $[\Gamma\text{-}\nu]$ does *not* hold for Γ , in cases like (6).

Moreover, φ must be technically usable to prove subject reduction; this leads to three *desiderata*:

- (D1) φ must make the typing context “safe:” if the type of $s[\mathbf{p}]$ sends a message to \mathbf{q} , then the type of $s[\mathbf{q}]$ must be able to input such a message;
- (D2) φ must be preserved when the typing rule $[\Gamma\text{-}\parallel]$ splits typing contexts (see derivation in Ex. 2.7);
- (D3) φ must be preserved when processes, and typing contexts, interact and reduce (Def. 2.2/2.8).

Therefore, the choice of the criterion for handling cases like (6) has a deep impact on the theoretical foundations of the type system: it determines how subject reduction and type safety properties are stated and proved, and how general/restrictive they are; it also determines how to define φ in rule $[\Gamma\text{-}\nu]$, to correctly type session restriction $(\nu s) P$, and handle cases like (7).

In §4, we show how our new MPST theory establishes its foundations, and φ in rule [T- ν]. But first, in §3, we show how such choices are made in classic MPST, and what are the consequences.

3 LIMITATIONS AND THEORETICAL ISSUES OF CLASSIC MPST

This section gives a formal basis to our claims in §1: in §3.1 we use our opening example to show the technical issues of classic MPST, caused by *consistency* (also called *coherency*, e.g., by Deniélou et al. [2012]); and in §3.2, we provide further examples that are rejected by classic MPST. Our new MPST system (§4) eschews these problems, by adopting a more general theoretical basis.

REMARK 3.1. *The issues described in this section do not apply to two recent MPST works, by Dezani-Ciancaglini et al. [2015] and Scalas and Yoshida [2018b]: they have different, non-classic MPST theories. However, such works have other limitations, surmounted by this paper: they are detailed in §8.2.*

3.1 Consistency and Subject Reduction

To reject cases like (6) (§2.3), classic MPST require typing contexts to be *consistent*: for each pair of entries $\{s[\mathbf{p}]:S_{\mathbf{p}}, s[\mathbf{q}]:S_{\mathbf{q}}\} \subseteq \Gamma$, the inputs/outputs of $S_{\mathbf{p}}$ from/to \mathbf{q} must be *dual* w.r.t. the outputs/inputs of $S_{\mathbf{q}}$ to/from \mathbf{p} . This guarantees that two roles \mathbf{p}, \mathbf{q} can only send/receive compatible messages in a session s . More precisely, consistency requires to check the duality of the *partial projections* $S_{\mathbf{p}}|\mathbf{q}$ and $S_{\mathbf{q}}|\mathbf{p}$, using Def. 3.5, 3.6, 3.7, and 3.8 (collected in Fig. 3): this clearly shows that MPST were developed by adopting a proof framework based on *binary* session types.

Correspondingly, to reject cases like (7), classic MPST define rule [T- ν] in Fig. 2 by setting $\varphi = \text{consistent}$. This yields the **classic session restriction typing rule**:

$$\frac{\Gamma' = \{s[\mathbf{p}]:S_{\mathbf{p}}\}_{\mathbf{p} \in I} \quad s \notin \Gamma \quad \text{consistent}(\Gamma') \quad \Theta \cdot \Gamma, \Gamma' \vdash P}{\Theta \cdot \Gamma \vdash (\nu s:\Gamma') P} \quad [\text{T-}\nu\text{CLASSIC}]$$

and this is sound (indeed, consistency satisfies the *desiderata* (D1)–(D3) described in §2.3).

E.g., the typing context in (6) is not consistent; correspondingly, no consistent Γ can be assigned to $(\nu s) \dots$ in (7): hence, with rule [T- $\nu\text{CLASSIC}$], the process in (7) is untypable in classic MPST.

Limitations of Consistency. Take the processes from Ex. 2.3, and the typing derivation from Ex. 2.7. Using the rules in Fig. 2 with [T- $\nu\text{CLASSIC}$] above, we might try to type our opening example as:

$$\frac{\Gamma \text{ consistent} \quad \frac{\vdots \text{ (from Ex. 2.7)}}{\Gamma \vdash P_{\mathbf{s}} | P_{\mathbf{c}} | P_{\mathbf{a}}} \quad [\text{T-}]}{\emptyset \vdash (\nu s:\Gamma) (P_{\mathbf{s}} | P_{\mathbf{c}} | P_{\mathbf{a}})} \quad [\text{T-}\nu\text{CLASSIC}] \quad \text{where } \Gamma = s[\mathbf{s}]:S_{\mathbf{s}}, s[\mathbf{c}]:S_{\mathbf{c}}, s[\mathbf{a}]:S_{\mathbf{a}} \quad (8)$$

As shown in §1(2), the types $S_{\mathbf{s}}, S_{\mathbf{c}}, S_{\mathbf{a}}$ assigned to $s[\mathbf{s}], s[\mathbf{c}], s[\mathbf{a}]$ are respectively $G|\mathbf{s}, G|\mathbf{c}, G|\mathbf{a}$, i.e., the projections of G (Def. 3.3). However, *the derivation in (8) is wrong*, because the consistency premise of [T- $\nu\text{CLASSIC}$] does *not* hold. To see why, we need to check all pairs of types for session s :

- $S_{\mathbf{s}}, S_{\mathbf{c}}$ are consistent: the outputs of $S_{\mathbf{s}}$ to \mathbf{c} are *dual* w.r.t. the inputs of $S_{\mathbf{c}}$ from \mathbf{s} ;
- $S_{\mathbf{s}}, S_{\mathbf{a}}$ are *not* consistent, because the partial projections $S_{\mathbf{s}}|\mathbf{a}$ and $S_{\mathbf{a}}|\mathbf{s}$ are *undefined* (Def. 3.6).

Intuitively, $S_{\mathbf{s}}|\mathbf{a}$ and $S_{\mathbf{a}}|\mathbf{s}$ are undefined because the inputs/outputs of $S_{\mathbf{s}}/S_{\mathbf{a}}$ from/to \mathbf{a}/\mathbf{s} depend on previous I/O with \mathbf{c} : i.e., if the service \mathbf{s} sends **login** (resp. **cancel**) to the client \mathbf{c} , then \mathbf{s} will (resp. will *not*) later interact with the authorisation server \mathbf{a} . This is not captured by the syntactic nature of projection/duality checks: i.e., protocols with inter-role dependencies are often *not* consistent — even simple ones, like G in (1). Consequently, the process in Ex. 2.3 is untypable, albeit correct (does not reduce to **err**).

Subject Reduction and Type Safety (or Lack Thereof). As noted in §1, the classic MPST subject reduction statement is (4). Now, consider (8) again: the conclusion is wrong, but the intermediate

Definition 3.2. The syntax of a global type G is:

$$G ::= \mathbf{p} \rightarrow \mathbf{q}; \{m_i(S_i) \cdot G_i\}_{i \in I} \mid \mu t.G \mid t \mid \mathbf{end} \quad \text{with } \mathbf{p} \neq \mathbf{q}, I \neq \emptyset, \text{ and } \forall i \in I: \text{fv}(S_i) = \emptyset$$

We write $\mathbf{p} \in \text{roles}(G)$ (or simply $\mathbf{p} \in G$) iff, for some \mathbf{q} , either $\mathbf{p} \rightarrow \mathbf{q}$ or $\mathbf{q} \rightarrow \mathbf{p}$ occurs in G .

Definition 3.3 (Global Type Projection). The projection of G onto \mathbf{p} , written $G|\mathbf{p}$, is:

$$(\mathbf{q} \rightarrow \mathbf{r}; \{m_i(S_i) \cdot G_i\}_{i \in I})|\mathbf{p} = \begin{cases} \mathbf{r} \oplus_{i \in I} m_i(S_i) \cdot (G_i|\mathbf{p}) & \text{if } \mathbf{p} = \mathbf{q} \\ \mathbf{q} \&_{i \in I} m_i(S_i) \cdot (G_i|\mathbf{p}) & \text{if } \mathbf{p} = \mathbf{r} \\ \prod_{i \in I} G_i|\mathbf{p} & \text{if } \mathbf{q} \neq \mathbf{p} \neq \mathbf{r} \end{cases}$$

$$(\mu t.G)|\mathbf{p} = \begin{cases} \mu t.(G|\mathbf{p}) & \text{if } G|\mathbf{p} \neq t' \ (\forall t') & t|\mathbf{p} = t \\ \mathbf{end} & \text{otherwise} & \mathbf{end}|\mathbf{p} = \mathbf{end} \end{cases}$$

where \prod is the merge operator for session types, that could be either the plain merging defined as $S \sqcap S = S$, or the full merging:

$$\mathbf{p} \&_{i \in I} m_i(S_i) \cdot S'_i \sqcap \mathbf{p} \&_{j \in J} m_j(S_j) \cdot T'_j = \mathbf{p} \&_{k \in I \cup J} m_k(S_k) \cdot (S'_k \sqcap T'_k) \ \& \ \mathbf{p} \&_{i \in I \cup J} m_i(S_i) \cdot S'_i \ \& \ \mathbf{p} \&_{j \in J \cup I} m_j(S_j) \cdot T'_j$$

$$\mathbf{p} \oplus_{i \in I} m_i(S_i) \cdot S'_i \sqcap \mathbf{p} \oplus_{i \in I} m_i(S_i) \cdot S'_i = \mathbf{p} \oplus_{i \in I} m_i(S_i) \cdot S'_i$$

$$\mu t.S \sqcap \mu t.T = \mu t.(S \sqcap T) \quad t \sqcap t = t \quad \mathbf{end} \sqcap \mathbf{end} = \mathbf{end}$$

Definition 3.4 (Partial Session Types). Partial session types, ranged over by H , are:

$$H ::= \&_{i \in I} m_i(S_i) \cdot H_i \mid \oplus_{i \in I} m_i(S_i) \cdot H_i \mid \mathbf{end} \mid \mu t.H \mid t \quad \text{with } I \neq \emptyset \text{ and } \forall i \in I: \text{fv}(S_i) = \emptyset$$

Definition 3.5 (Duality of Partial Session Types). The dual of H , written \bar{H} , is:

$$\overline{\&_{i \in I} m_i(S_i) \cdot H_i} = \oplus_{i \in I} m_i(S_i) \cdot \bar{H}_i \quad \overline{\oplus_{i \in I} m_i(S_i) \cdot H_i} = \&_{i \in I} m_i(S_i) \cdot \bar{H}_i \quad \overline{\mu t.H} = \mu t.\bar{H} \quad \bar{t} = t \quad \overline{\mathbf{end}} = \mathbf{end}$$

Definition 3.6 (Partial Projection). The projection of S onto \mathbf{p} , written $S|\mathbf{p}$, is:

$$(\mathbf{q} \&_{i \in I} m_i(S_i) \cdot S'_i)|\mathbf{p} = \begin{cases} \&_{i \in I} m_i(S_i) \cdot (S'_i|\mathbf{p}) & \text{if } \mathbf{p} = \mathbf{q} \\ \prod_{i \in I} S'_i|\mathbf{p} & \text{if } \mathbf{p} \neq \mathbf{q} \end{cases} \quad (\mathbf{q} \oplus_{i \in I} m_i(S_i) \cdot S'_i)|\mathbf{p} = \begin{cases} \oplus_{i \in I} m_i(S_i) \cdot (S'_i|\mathbf{p}) & \text{if } \mathbf{p} = \mathbf{q} \\ \prod_{i \in I} S'_i|\mathbf{p} & \text{if } \mathbf{p} \neq \mathbf{q} \end{cases}$$

$$(\mu t.S)|\mathbf{p} = \begin{cases} \mu t.(S|\mathbf{p}) & \text{if } S|\mathbf{p} \neq t' \ (\forall t') & t|\mathbf{p} = t \\ \mathbf{end} & \text{otherwise} & \mathbf{end}|\mathbf{p} = \mathbf{end} \end{cases}$$

where \prod is the merge operator for partial session types, defined as:

$$\&_{i \in I} m_i(S_i) \cdot H_i \sqcap \&_{i \in I} m_i(S_i) \cdot H'_i = \&_{i \in I} m_i(S_i) \cdot (H_i \sqcap H'_i)$$

$$\oplus_{i \in I} m_i(S_i) \cdot H_i \sqcap \oplus_{j \in J} m_j(S_j) \cdot H'_j = \oplus_{k \in I \cup J} m_k(S_k) \cdot (H_k \sqcap H'_k) \ \oplus \ \oplus_{i \in I \cup J} m_i(S_i) \cdot H_i \ \oplus \ \oplus_{j \in J \cup I} m_j(S_j) \cdot H'_j$$

$$\mu t.H \sqcap \mu t.H' = \mu t.(H \sqcap H') \quad t \sqcap t = t \quad \mathbf{end} \sqcap \mathbf{end} = \mathbf{end}$$

Definition 3.7. Subtyping for partial types is coinductively defined (we omit unfolding rules, cf. Def. 2.5):

$$\frac{\forall i \in I \quad S_i \leq T_i \quad H'_i \leq H''_i}{\&_{i \in I} m_i(S_i) \cdot H'_i \leq \&_{i \in I \cup J} m_i(T_i) \cdot H''_i} \quad \frac{\forall i \in I \quad T_i \leq S_i \quad H'_i \leq H''_i}{\oplus_{i \in I \cup J} m_i(S_i) \cdot H'_i \leq \oplus_{i \in I} m_i(T_i) \cdot H''_i} \quad \overline{\mathbf{end}} \leq \mathbf{end}$$

Definition 3.8. Γ is consistent iff, $\forall s, \mathbf{p} \neq \mathbf{q}, S, T, \{s[\mathbf{p}]:S, s[\mathbf{q}]:T\} \subseteq \Gamma$ implies $\overline{S|\mathbf{q}} \leq T|\mathbf{p}$.

Fig. 3. Classic MPST: global types, projections, consistency, and duality. Note that all these definitions are **not** necessary in our new theory of multiparty session types (§4).

judgement $\Gamma \vdash P_s \mid P_c \mid P_a$ holds. For this judgement, the subject reduction statement (4) is vacuously true (since Γ is not consistent): hence, we cannot prove that $P_s \mid P_c \mid P_a$ “never goes wrong.”

Interplay Between Consistency and Global Type Projection. The consistency requirement constrains the MPST theory in non-obvious ways, causing subtle issues with *global type projections*. Several MPST papers claim that if Γ is obtained by projecting a global type G , then Γ is consistent (see

e.g.: [Deniérou et al. 2012, p.28], [Coppo et al. 2015a, Prop. 1], [Chen 2015, Prop. 2]). This claim corresponds to introducing the typing rule $[T\text{-}v\text{CLASSIC}G]$ below, that seemingly fixes derivation (8):

$$\frac{\Gamma' = \{s[\mathbf{p}]:G|\mathbf{p}\}_{\mathbf{p} \in \text{roles}(G)} \quad s \notin \Gamma \quad \Theta \cdot \Gamma, \Gamma' \vdash P}{\Theta \cdot \Gamma \vdash (vs:\Gamma')P} \quad [T\text{-}v\text{CLASSIC}G]$$

Unfortunately, our example in §1 shows a global type whose projections are *not* consistent. This is because we use the “full merging” projection (Def. 3.3), introduced in Deniérou et al. [2012]; Yoshida et al. [2010] to type more processes. The intuition is the following. Take the initial choice of the global type G in §1(1) (reported below), that does *not* involve role \mathbf{a} :

$$G = \mathbf{s} \rightarrow \mathbf{c}: \{\text{login}.G_1, \text{cancel}.G_2\} \quad \text{where} \quad \begin{cases} G_1 = \mathbf{c} \rightarrow \mathbf{a}: \text{passwd}(\text{Str}) . \mathbf{a} \rightarrow \mathbf{s}: \text{auth}(\text{Bool}) \\ G_2 = \mathbf{c} \rightarrow \mathbf{a}: \text{quit} \end{cases}$$

To project G onto \mathbf{a} , we must “skip” the first interaction between \mathbf{s} and \mathbf{c} , and *merge* the projections of G_1 and G_2 onto \mathbf{a} , rejecting potentially unsafe local types combinations (thus avoiding cases like (6) above). Consequently, projection works as follows:

$$G \upharpoonright \mathbf{a} = S_1 \sqcap S_2 \quad \text{where} \quad \begin{cases} S_1 = G_1 \upharpoonright \mathbf{a} = \mathbf{c} \& \text{passwd}(\text{Str}) . \mathbf{s} \oplus \text{auth}(\text{Bool}) \\ S_2 = G_2 \upharpoonright \mathbf{a} = \mathbf{c} \& \text{quit} \end{cases}$$

We now have two possibilities, depending on how we choose the *merging operator* \sqcap (Def. 3.3):

- *plain merging*: $S_1 \sqcap S_2 = S_1$ iff $S_1 = S_2$ (undefined otherwise);
- *full merging*: $S_1 \sqcap S_2 = S_{\mathbf{a}}$ (see (2) in §1).

i.e., the restrictive plain merging is undefined for our example G , while full merging yields all desired projections — but they are *not consistent*, as shown above. Consequently, the tentative rule $[T\text{-}v\text{CLASSIC}G]$ with “full merging” projections *breaks subject reduction proofs*. E.g., take P typed by $[T\text{-}v\text{CLASSIC}G]$, and reducing to P' , as follows:

$$\emptyset \cdot \emptyset \vdash P \quad \text{with} \quad P = (vs:\Gamma)P_0 \rightarrow (vs:\Gamma')P_1 = P' \quad (\text{induced by } P_0 \rightarrow P_1 \text{ and rule } [R\text{-}C_{\text{Tx}}] \text{ in Fig.1}) \quad (9)$$

To prove subject reduction as stated in (4), we need to invert P 's typing and apply the induction hypothesis on $\Theta \cdot \Gamma \vdash P_0$ and $P_0 \rightarrow P_1$ (from (9)), to obtain that there is some Γ' such that $\Gamma \rightarrow^* \Gamma'$ and $\Theta \cdot \Gamma' \vdash P_1$; however, to apply (4) in the induction hypothesis we need Γ consistent, and we have shown that this hypothesis might not hold.

We can now revisit our claims in §1, making them precise, and highlighting the resulting impasse:

- (C1) **overly restrictive**: requiring Γ consistent drastically constrains typability: it rejects our simple example in §1, and many other correct protocols (see §3.2 later on). Correspondingly, the restrictive “plain merging” projection of [Honda et al. 2008, Def. 4.1] and [Coppo et al. 2015a, Def. 1], guarantees consistency by rejecting many correct protocols;
- (C2) **inflexible and error-prone**: if we use a “full merging” projection as in, e.g., Yoshida et al. [2010] or Deniérou et al. [2012], then Γ might *not* be consistent. This means that the proofs of subject reduction depending on “full merging” (e.g. [Yoshida et al. 2010, Thm 3.5], [Deniérou et al. 2012, Thm 4.6], and successive papers discussed in §8) do not work; we might fix such proofs by adding a consistency requirement — but then, we would fall back into (C1) above.

In §4, we completely eschew these issues by developing new theoretical foundations for MPST: we cut the ties with binary session types, adopting a more general, *behavioural* safety invariant, that subsumes consistency and binary session duality.

3.2 More Examples of Correct, yet Non-Consistent Protocols

We conclude this section with Fig. 4, that describes various multiparty protocols, formalised as typing contexts. None of such protocols is consistent, because some of their partial projections are

(1) OAuth2 fragment.	
(See global type (1) in §1)	(See types (2) in §1, and Γ in Ex.2.7)
(2) Recursive two-buyers protocol. This is a mild variation of a typical example in MPST literature. Alice (a) queries the store (s) for an item, and the store replies with a price ; then, she asks Bob (b) to split the price: if he says yes , then she buys the item from the store; if he says no , then Alice recursively retries, proposing another split to Bob; at any point, Alice can cancel her bargaining with Bob, and say no to the store.	
N/A	$s[\mathbf{a}] : s \oplus \text{query}(\text{Str}).s \& \text{price}(\text{Int}).\mu t. \mathbf{b} \oplus \left\{ \begin{array}{l} \text{split}(\text{Int}).\mathbf{b} \& \left\{ \begin{array}{l} \text{yes}.s \oplus \text{buy}.\text{end} \\ \text{no}.t \end{array} \right\} \\ \text{cancel}.s \oplus \text{no} \end{array} \right\}$ $s[\mathbf{s}] : \mathbf{a} \& \text{query}(\text{Str}).\mathbf{a} \oplus \text{price}(\text{Int}).\mathbf{a} \& \left\{ \begin{array}{l} \text{buy}.\text{end} \\ \text{no}.\text{end} \end{array} \right\}$ $s[\mathbf{b}] : \mu t. \mathbf{a} \& \left\{ \begin{array}{l} \text{split}(\text{Int}).\mathbf{a} \oplus \left\{ \begin{array}{l} \text{yes}.\text{end} \\ \text{no}.t \end{array} \right\} \\ \text{cancel}.\text{end} \end{array} \right\}$
(3) Recursive map/reduce. The mapper (m) sends a datum to n workers (w_1, \dots, w_n , for some given n), and each one sends a result to the reducer (r); then, the reducer tells the mapper whether to continue with another iteration, or stop : in the first case, the mapper loops, while in the second case, it stops the workers.	
$\mu t. \mathbf{m} \rightarrow w_1 : \text{datum}(\text{Int}) \dots \dots$ $\mathbf{m} \rightarrow w_n : \text{datum}(\text{Int})$ $w_1 \rightarrow \mathbf{r} : \text{result}(\text{Int}) \dots \dots$ $w_n \rightarrow \mathbf{r} : \text{result}(\text{Int})$ $\mathbf{r} \rightarrow \mathbf{m} : \left\{ \begin{array}{l} \text{continue}(\text{Int}).t \\ \text{stop}.\mathbf{m} \rightarrow w_1 : \text{stop} \dots \dots \\ \mathbf{m} \rightarrow w_n : \text{stop} \end{array} \right\}$	$s[\mathbf{m}] : \mu t. w_1 \oplus \text{datum}(\text{Int}) \dots \dots w_n \oplus \text{datum}(\text{Int}).\mathbf{r} \& \left\{ \begin{array}{l} \text{continue}(\text{Int}).t \\ \text{stop}.w_1 \oplus \text{stop} \dots \dots w_n \oplus \text{stop} \end{array} \right\}$ $s[w_i] : \mathbf{m} \& \text{datum}(\text{Int}).\mu t. \mathbf{r} \oplus \text{result}(\text{Int}).\mathbf{m} \& \left\{ \begin{array}{l} \text{datum}(\text{Int}).t \\ \text{stop}.\text{end} \end{array} \right\} \quad (i \in 1..n)$ $s[\mathbf{r}] : \mu t. w_1 \& \text{datum}(\text{Int}) \dots \dots w_n \& \text{datum}(\text{Int}).\mathbf{m} \oplus \left\{ \begin{array}{l} \text{continue}(\text{Int}).t \\ \text{stop}.\text{end} \end{array} \right\}$
(4) Independent multiparty workers. The starter process (s) sends a datum to n worker processes (w_1, \dots, w_n , for some given n), and each one starts exchanging datum/result messages with two other workers (w_{b_i} and w_{c_i} , for $i \in 1..n$). Each triplet of workers $w_{a_i}, w_{b_i}, w_{c_i}$ ($i \in 1..n$) keeps interacting until w_{a_i} sends stop to w_{b_i} , who forwards stop to w_{c_i} .	
$\mathbf{s} \rightarrow w_{a_1} : \text{datum}(\text{Int}) \dots \dots \mathbf{s} \rightarrow w_{a_n} : \text{datum}(\text{Int})$ $\mu t. w_{a_i} \rightarrow w_{b_i} : \left\{ \begin{array}{l} \text{datum}(\text{Int}).w_{b_i} \rightarrow w_{c_i} : \text{datum}(\text{Int}) \\ \mathbf{w}_{c_i} \rightarrow w_{a_i} : \text{result}(\text{Int}).t \\ \text{stop}.w_{b_i} \rightarrow w_{c_i} : \text{stop} \end{array} \right\}$	$s[\mathbf{s}] : w_{a_1} \oplus \text{datum}(\text{Int}) \dots \dots w_{a_n} \oplus \text{datum}(\text{Int}).\text{end}$ $s[w_{a_i}] : \mathbf{s} \& \text{datum}(\text{Int}).\mu t. w_{b_i} \oplus \left\{ \begin{array}{l} \text{datum}(\text{Int}).w_{c_i} \& \text{result}(\text{Int}).t \\ \text{stop}.\text{end} \end{array} \right\}$ $s[w_{b_i}] : \mu t. w_{a_i} \& \left\{ \begin{array}{l} \text{datum}(\text{Int}).w_{c_i} \oplus \text{datum}(\text{Int}).t \\ \text{stop}.w_{c_i} \oplus \text{stop}.\text{end} \end{array} \right\}$ $s[w_{c_i}] : \mu t. w_{b_i} \& \left\{ \begin{array}{l} \text{datum}(\text{Int}).w_{a_i} \oplus \text{result}(\text{Int}).t \\ \text{stop}.\text{end} \end{array} \right\}$

Fig. 4. A selection of multiparty protocols: each one is expressed as a (non-consistent) typing context (on the right); for the sake of clarity, we also outline the shape of a global type with corresponding projections (on the left). The exception is protocol (2), that cannot be projected from *any* global type: see §3.2. Being non-consistent, all these protocols are not supported by classic MPST — but they are all supported by our new general type system (§4); moreover, they have different behavioural properties, analysed in §5.3 (Table 1).

undefined — as a consequence of the issues illustrated in §3.1; moreover, the protocols (2), (3) and (4) trigger further subtle restrictions in the partial projection/merging of recursive types (Def. 3.6).

Notably, Fig. 4 includes an example of multiparty protocol that cannot be projected from *any* global type: the recursive two-buyers protocol (2). The key issue is in the type of $s[\mathbf{a}]$, when **alice** interacts with **bob**: **alice** sends a message to the **store** in one of the branches under recursion $\mu t \dots$ (where **bob** answers **yes**), but not in the other branch (where **bob** says **no**). This is *not* supported by projection and merging (Def. 3.3): they can only generate session types where all branches under recursion syntactically contain a same set of roles. Consequently, no global type can be projected and yield the type of $s[\mathbf{a}]$ in Fig. 4(2). This restriction does not impact our new MPST theory (§4).

4 A NEW, GENERAL MULTIPARTY SESSION TYPE SYSTEM

We now present our new general MPST theory. Its generality comes from the fact that it is based on a weak *typing context safety* invariant, that rejects cases like (6)/(7) (§2.3) without the restrictions

and drawbacks of classic MPST consistency. Moreover, we design the new type system to be *parametric* on the safety invariant itself: by fine-tuning the parameter, the type system can accept or reject MPST processes depending on the properties of the protocols they implement (we will take advantage of this feature in §5). Hence, different instantiations of the parameter yield different type system instances — but we just need to prove type safety *once*, under the *weakest* safety invariant. This design is inspired by Igarashi and Kobayashi [2004]’s Generic Type System for the π -calculus.

We first formalise what a “safety invariant” is, in Def. 4.1 below: it is a *behavioural* property of typing contexts, that depends on how they reduce (cf. Def. 2.8). The fundamental difference with classic MPST (§3) is that our safety is *not* based on binary session types, *nor* duality.

Definition 4.1. φ is a *safety property* of typing contexts iff:

- [S- $\oplus\&$] $\varphi(\Gamma, s[\mathbf{p}]:\mathbf{q}\oplus_{i\in I}m_i(S_i).S'_i, s[\mathbf{q}]:\mathbf{p}\&_{j\in J}m_j(T_j).T'_j)$ implies $I\subseteq J$, and $\forall i\in I : S_i\leq T_i$;
- [S- μ] $\varphi(\Gamma, s[\mathbf{p}]:\mu t.S)$ implies $\varphi(\Gamma, s[\mathbf{p}]:S\{\mu t.S/t\})$;
- [S- \rightarrow] $\varphi(\Gamma)$ and $\Gamma\rightarrow\Gamma'$ implies $\varphi(\Gamma')$.

We say Γ is *safe*, written $\text{safe}(\Gamma)$, if $\varphi(\Gamma)$ for some safety property φ .

The rules of Def. 4.1 directly satisfy the *desiderata* (D1) and (D3) discussed in §2.3 (whereas (D2) is satisfied by Lemma 4.3, as we will see shortly). Rule [S- $\oplus\&$] says that the roles in a safe typing context can only exchange compatible messages (this is *desideratum* (D1)): more precisely, if the typing context contains entries for $s[\mathbf{p}]$ and $s[\mathbf{q}]$, with \mathbf{p} sending to \mathbf{q} and \mathbf{q} receiving from \mathbf{p} , then \mathbf{p} support all \mathbf{q} ’s messages — and thus, they can reduce, by Def. 2.8. Rule [S- μ] says that φ contains all recursive type unfoldings: this allows rule [S- $\oplus\&$] to check unfolded types, where $\oplus/\&$ occur at the the top-level. By rule [S- \rightarrow], safety is preserved whenever Γ reduces (this is *desideratum* (D3)).

Example 4.2. The typing context Γ of (8) in §3 is safe. This can be easily verified by: (1) defining φ as $\varphi = \{\Gamma' \mid \Gamma\rightarrow^*\Gamma'\}$, i.e., containing Γ and all its reductions; (2) checking that φ is a safety property, because all its elements satisfy the clauses of Def. 4.1; and (3) concluding that, since $\varphi(\Gamma)$ holds, Γ is safe. Instead, the typing context in (6) is *not* safe: any property φ containing such typing context is *not* a safety property, as it violates clause [S- $\oplus\&$] of Def. 4.1.

Def. 4.1 also has the properties in Lemma 4.3 below, useful for proving subject reduction: typing context splits preserve safety (item 1, which satisfies the remaining *desideratum* (D2) in §2.3); if Γ is safe, then supertyping/reductions commute (item 2); supertyping preserves safety (item 3).

LEMMA 4.3. For all typing contexts Γ and Γ' :

- (1) if Γ, Γ' is safe, then Γ is safe;
- (2) if Γ safe and $\Gamma\leq\Gamma'\rightarrow\Gamma''$ (for some Γ''), then there is Γ''' such that $\Gamma\rightarrow\Gamma'''\leq\Gamma''$;
- (3) if Γ is safe and $\Gamma\leq\Gamma'$, then Γ' is safe.

We can now define our new multiparty session type system. As explained in §2.3, since we are adopting safety (Def. 4.1) as the criterion for accepting/rejecting typing contexts, we use the same criterion to define a typing rule for session restriction.

Definition 4.4 (General Multiparty Session Type System). The *general MPST typing judgement* is inductively defined by the rules in Fig.2 — with rule [T- ν] restricted as follows:

$$\frac{\Gamma' = \{s[\mathbf{p}]:S_{\mathbf{p}}\}_{\mathbf{p}\in I} \quad \varphi(\Gamma') \quad s\notin\Gamma \quad \Theta\cdot\Gamma, \Gamma' \vdash P}{\Theta\cdot\Gamma \vdash (\nu s:\Gamma')P} \quad \text{[TGEN-}\nu\text{]} \quad \text{where } \varphi \text{ is a safety property}$$

Given a safety property φ , we write “ $\Theta\cdot\Gamma \vdash P$ with φ ” to instantiate φ in [TGEN- ν] above; when “with φ ” is omitted, then the instantiation is $\varphi = \text{safe}$ (i.e., the largest safety property, cf. Def. 4.1).

Example 4.5. Take the (wrong) typing derivation (8) in §3.1, and replace the (wrong) application of rule $[\text{T-VCLASSIC}]$ with $[\text{TGEN-V}]$ from Def. 4.4, instantiating φ with the safety property of Ex. 4.2 (that contains Γ). The resulting typing derivation is correct.

Ex. 4.5 above shows that our new type system is not limited by consistency requirements, and types our opening example. Notably, the only visible difference between our new type system (Def. 4.4) and the classic one (§3.1) is that $[\text{TGEN-V}]$ uses a (parametric) safety property φ , instead of consistency.² As explained in §2.3, this small visible difference between typing rules is a manifestation of a deeper underlying change: by removing the crucial consistency/duality assumption of classic MPST, we are replacing its theoretical underpinnings, and this requires a revision of all MPST soundness proofs. The payoff is that our new MPST theory enjoys a much more general subject reduction property (Thm. 4.6, based on Lemmas 4.3 to 4.3); from this, we get that typed processes “never go wrong” (Cor. 4.7). And again, unlike classic MPST, these results are *not* limited by consistency.

THEOREM 4.6 (SUBJECT REDUCTION). *Assume $\Theta \cdot \Gamma \vdash P$ and Γ safe. Then, $P \rightarrow P'$ implies $\exists \Gamma'$ safe such that $\Gamma \rightarrow^* \Gamma'$ and $\Theta \cdot \Gamma' \vdash P'$.*

COROLLARY 4.7 (TYPE SAFETY). *If $\emptyset \cdot \emptyset \vdash P$ and $P \rightarrow^* P'$, then P' has no error.*

PROOF. We first prove a more general result. Assume $\Theta \cdot \Gamma \vdash P$ with Γ safe, and $P = P_1 \rightarrow \dots \rightarrow P_n = P'$. By induction on n , using Thm. 4.6, we prove $\Theta \cdot \Gamma' \vdash P'$, for some safe Γ' such that $\Gamma \rightarrow^* \Gamma'$. Now, by contradiction, assume that P' has an error (Def. 2.2); then, P' is untypable, since its **err** subterm is untypable: contradiction. Hence, P' has no errors. We obtain Cor. 4.7 as a special case of the result above, with $\Theta = \emptyset$ and $\Gamma = \Gamma' = \emptyset$ (that is vacuously safe). \square

Example 4.8. Take our opening example in §1, and its typed process from Ex. 2.7 and 4.5. Using our new Thm. 4.6 instead of the classic MPST subject reduction (4) in §1, we infer that all process reductions are well-typed. And by Cor. 4.7, we are guaranteed that they do not contain errors.

Finally, note that type checking is decidable, whenever Def. 4.4 is instantiated with a decidable safety property: this mainly follows because typing rules are syntax-directed, and for any P , at most one can be applied. Also note that, since we proved Thm. 4.6 and Cor. 4.7 using the largest (i.e., the weakest) safety property, we do not need to repeat the proof depending on how φ is instantiated in Def. 4.4: subject reduction and type safety hold for any safety property φ .

THEOREM 4.9. *If φ is decidable, then “ $\Theta \cdot \Gamma \vdash P$ with φ ” is decidable.*

5 VERIFYING RUN-TIME PROPERTIES OF PROCESSES, USING TYPES

In this section, we show that by suitably instantiating φ in our type system (Def. 4.4), we can statically enforce desired run-time properties on processes — e.g., deadlock freedom and liveness.

In order to achieve this result, we study several typing context properties, and compare them with safety (Def. 4.1). The main reason for this study is that safety, albeit guaranteeing error-freedom (Thm. 4.6, Cor. 4.7), is otherwise rather weak. E.g., the following typing context is safe but deadlocked (it cannot reduce, because \mathbf{p} is waiting an input from \mathbf{q} , who is waiting for \mathbf{r} , who is waiting for \mathbf{p}):

$$s[\mathbf{p}]:\mathbf{q}\&\mathbf{m}_1.\mathbf{r}\oplus\mathbf{m}_2, s[\mathbf{q}]:\mathbf{r}\&\mathbf{m}_3.\mathbf{p}\oplus\mathbf{m}_1, s[\mathbf{r}]:\mathbf{p}\&\mathbf{m}_2.\mathbf{q}\oplus\mathbf{m}_3$$

and the context above types deadlocked processes that cannot reduce, either. This is undesirable: “real-world” programs should be deadlock-free, or even *live* (i.e., each pending input/output should be fired, eventually). Therefore, stronger typing context properties are needed — and in our new MPST theory, we can use the parameter φ of Def. 4.4 to enforce them, without consistency limitations.

²In §5.4, we show that all typing derivations of classic MPST are valid under Def. 4.4: consistency implies safety, hence in $[\text{TGEN-V}]$ we can let $\varphi = \text{consistent}$; and in §5.5, we show how φ statically determines the run-time properties on processes.

We first discuss several desirable, although undecidable, run-time properties of processes, such as deadlock-freedom and liveness (§5.1); next, we prove *session fidelity*, a crucial result that connects typing context reductions to processes reductions (§5.2). Then, we present various typing context properties (§5.3), and compare them (§5.4); finally, we show that they are decidable, and, with our new type system, they can be used to ensure that processes are, e.g., deadlock-free and live (§5.5).

5.1 Run-Time Properties of Processes

In Def. 5.1 below, we formalise various desirable process properties. All these properties are *undecidable*, because the MPST π -calculus is Turing-powerful [Busi et al. 2009]. To surmount this obstacle, from §5.3 we will reason on analogous properties for types (that are not Turing-powerful).

Definition 5.1 (Process properties). P is **deadlock-free** iff $P \rightarrow^* P' \not\rightarrow$ implies $P' \equiv \mathbf{0}$. P is **terminating** iff it is deadlock-free, and $\exists j$ finite such that, $\forall n \geq j$, $P = P_0 \rightarrow P_1 \rightarrow \dots \rightarrow P_n$ implies $P_n \equiv \mathbf{0}$. P is **never-terminating** iff $P \rightarrow^* P'$ implies $P' \rightarrow$. P is **live** iff $P \rightarrow^* P' \equiv \mathbb{C}[Q]$ implies:

- (1) if $Q = c[\mathbf{q}] \oplus m\langle s'[\mathbf{r}] \rangle . Q'$ (for some m, s', \mathbf{r}, Q'), then $\exists \mathbb{C}' : P' \rightarrow^* \mathbb{C}'[Q']$; and
- (2) if $Q = c[\mathbf{q}] \sum_{i \in I} m_i(x_i) . Q'_i$ (for some m_i, x_i, Q'_i), then $\exists \mathbb{C}', k \in I, s', \mathbf{r} : P' \rightarrow^* \mathbb{C}'[Q'_k\{s'[\mathbf{r}]/x_k\}]$.

P is **strongly live** iff $P \rightarrow^* P' \equiv \mathbb{C}[Q]$ implies:

- (3) item 1 above, and moreover, there is n finite such that, whenever $P' = P'_0 \rightarrow P'_1 \rightarrow \dots \rightarrow P'_n$, then for some $j \leq n$ we have $P'_j \rightarrow \mathbb{C}''[Q']$ (for some \mathbb{C}'');
- (4) item 2 above, and moreover, there is n finite such that, whenever $P' = P'_0 \rightarrow P'_1 \rightarrow \dots \rightarrow P'_n$, then for some $j \leq n$ we have $P'_j \rightarrow \mathbb{C}''[Q'_k\{s'[\mathbf{r}]/x_k\}]$ (for some \mathbb{C}'' , $k \in I, s', \mathbf{r}$).

In Def. 5.1, a process P is deadlock-free when it only stops reducing by becoming $\mathbf{0}$; P is terminating when it always reaches $\mathbf{0}$ after a finite number of reductions; P is never-terminating when it reduces forever; P is live (a.k.a. “lock-free” [Kobayashi and Sangiorgi 2010; Padovani 2014]) when all its pending inputs/outputs *can* always eventually communicate with a corresponding output/input; P is strongly live when all its pending inputs/outputs *will* always find a corresponding output/input, enabling communication after a finite number of reductions.

Example 5.2. We now illustrate the differences among the properties in Def. 5.1. Let:

$$P = P_1 \mid P_2 \quad \text{where} \quad \begin{cases} P_1 = s[\mathbf{p}][\mathbf{q}] \sum \text{resp}.P \\ P_2 = \text{def } X(x) = x[\mathbf{r}] \sum \{m_1.X(x), m_2.x[\mathbf{p}] \oplus \text{resp}.\mathbf{0}\} \text{ in } X\langle s[\mathbf{q}] \rangle \mid Q \end{cases}$$

i.e., P_1 implements \mathbf{p} , and waits a response from \mathbf{q} ; P_2 implements \mathbf{q} , and loops every time role \mathbf{r} (whose omitted implementation is in Q) sends m_1 ; if/when \mathbf{r} chooses to send m_2 , then P_2 sends the response to \mathbf{p} , triggering the input in P_1 . Now, consider the following implementation of Q :

$$Q = \text{def } Y(y) = y[\mathbf{q}] \oplus m_1.Y\langle y \rangle \text{ in } Y\langle s[\mathbf{r}] \rangle$$

i.e., \mathbf{r} sends m_1 to \mathbf{q} forever – hence, P reduces forever, which means that P is never-terminating and deadlock-free. But note that the sub-process P_1 never has a chance to receive the desired response from \mathbf{q} : hence, P is *not* live. To address this, we can instead define Q above as:

$$Q = s[\mathbf{r}][\mathbf{q}] \oplus m_1.s[\mathbf{r}][\mathbf{q}] \oplus m_2.\mathbf{0} \mid Q' \quad \text{where} \quad Q' = \begin{cases} \text{def } Z(z) = z[\mathbf{r}'] \oplus m_3.Z\langle z \rangle \text{ in} \\ \text{def } Z'(z') = z'[\mathbf{r}'] \sum m_3(x).Z'\langle z' \rangle \text{ in} \\ Z\langle s[\mathbf{r}'] \rangle \mid Z'\langle s[\mathbf{r}'] \rangle \end{cases}$$

i.e., \mathbf{r} sends m_1 and then m_2 to \mathbf{q} , and this causes \mathbf{q} to send *resp* to \mathbf{p} (cf. P_2 above); meanwhile, the sub-process Q' loops, with \mathbf{r}' and \mathbf{r}'' exchanging message m_3 . With this definition of Q , we obtain that P is live, because P_1 *can* always eventually receive its input while P_2 reduces.

Still, P is *not* strongly live, because the input of P_1 could be arbitrarily delayed by letting Q' reduce forever, without firing the outputs of Q . We can make P strongly live, e.g., by redefining Q' as $Q' = 0$: this guarantees that P_1 *will* receive its input within 3 reductions.³

5.2 Session Fidelity

We now prove that if a typing context can reduce, then a typed process P simulates the reduction (Thm. 5.4). A related result can be proved for classic MPST – but in our new theory, it is stronger: we do *not* assume consistency of the typing context, *nor* the existence of a global type projecting it. Session fidelity requires P to be (1) not deadlocked, and (2) *productive*, i.e., not trapped in a loop like $\text{def } X(x) = X\langle x \rangle \text{ in } X\langle s[\mathbf{p}] \rangle$, if $s[\mathbf{p}]$ needs to be used for input/output: this is formalised in Def. 5.3.

Definition 5.3. Assume $\emptyset \cdot \Gamma \vdash P$. We say that P :

- (1) **has guarded definitions** iff in each subterm of the form $\text{def } X(x_1:S_1, \dots, x_n:S_n) = Q \text{ in } P'$, for all $i \in 1..n$, $S_i \not\prec \text{end}$ implies that a call $Y\langle \dots, x_i, \dots \rangle$ can only occur in Q as subterm of $x_i[\mathbf{q}] \sum_{j \in J} m_j(y_j).P_j$ or $x_i[\mathbf{q}] \oplus m(c).P''$ (i.e., after using x_i for input/output);
- (2) **only plays role \mathbf{p} in s , by Γ** , iff: (i) P has guarded definitions; (ii) $\text{fv}(P) = \emptyset$; (iii) $\Gamma = \Gamma_0, s[\mathbf{p}] : S$ with $S \not\prec \text{end}$ and $\text{end}(\Gamma_0)$; (iv) in all subterms $(vs':\Gamma') P'$ of P , we have $\text{end}(\Gamma')$.

We say “ P **only plays role \mathbf{p} in s ” iff $\exists \Gamma : \emptyset \cdot \Gamma \vdash P$, and item 2 holds.**

We will explain item 1 of Def. 5.3 shortly (after Thm. 5.4). Item 2 identifies a process that plays exactly *one* role on *one* session: clearly, an ensemble of such processes cannot deadlock by waiting for each other on multiple sessions. All our examples (except a few, duly noted) satisfy Def. 5.3(2).

Now, in Thm. 5.4 we prove that a set of processes involved in a single session simulates the typing context, following its types/protocols. This addresses the typical application scenario of MPST: an ensemble of programs $P_{\mathbf{p}}$ interact on a multiparty session s , each one playing a distinct role \mathbf{p} .

THEOREM 5.4 (SESSION FIDELITY). Assume $\emptyset \cdot \Gamma \vdash P$, where Γ is safe, $P \equiv \prod_{\mathbf{p} \in I} P_{\mathbf{p}}$, and each $P_{\mathbf{p}}$ either is 0 (up-to \equiv), or only plays \mathbf{p} in s . Then, $\Gamma \rightarrow$ implies $\exists \Gamma', P'$ such that $\Gamma \rightarrow \Gamma', P \rightarrow^* P'$ and $\emptyset \cdot \Gamma' \vdash P'$, where $P' \equiv \prod_{\mathbf{p} \in I} P'_{\mathbf{p}}$ and each $P'_{\mathbf{p}}$ either is 0 (up-to \equiv), or only plays \mathbf{p} in s .

Note that in Thm. 5.4, P chooses which reduction of Γ to follow: in fact, a selection type in Γ might allow to choose m_1, \dots, m_n (with different continuations), but P might select only one m_k (by $[\tau\text{-}\oplus]$ in Fig. 2, and subtyping). This observation will be a crucial when reasoning about process liveness (§5.5). Also note that Thm. 5.4 relies on item 1 of Def. 5.3. In fact, by rule $[\tau\text{-}\text{def}]$ (Fig. 2), an unguarded definition $X(x:S) = X\langle x \rangle$ can be typed with *any* S ; therefore, we have e.g.:

$$\emptyset \cdot s[\mathbf{p}]:\mathbf{q}\oplus\mathbf{m}, s[\mathbf{q}]:\mathbf{p}\&\mathbf{m} \vdash \text{def } X(x:\mathbf{q}\oplus\mathbf{m}) = X\langle x \rangle \text{ in } X\langle s[\mathbf{p}] \rangle \mid s[\mathbf{q}][\mathbf{p}] \sum \mathbf{m}$$

and the unguarded process above reduces vacuously by calling X infinitely, without matching any typing context reduction; this explains the need of guarded definitions in Thm. 5.4.

5.3 Typing Context Properties

Fig. 5 lists several behavioural properties of typing contexts. In §5.5, we will show how they can statically enforce the run-time process properties discussed in §5.1.

- Γ is *deadlock-free* iff it stops reducing only when it only contains **ends**;
- Γ is *terminating* iff it always reaches a final configuration, in a finite number of steps;
- Γ is *never-terminating* iff it never stops reducing;
- Γ is *live*, *live⁺* or *live⁺⁺* iff each branching/selection can be eventually fired.

³As a more laborious alternative, we could formalise and assume a notion of *fair scheduling*, that eventually fires any action that is persistently enabled; we adopt a similar intuition for type reductions, in Def. 5.5.

<p>(1) Γ is safe, written $\text{safe}(\Gamma)$, iff:</p> <p>(see Def. 4.1)</p>	$\Gamma \models \nu Z. \left(\begin{array}{l} \forall s, \mathbf{p}, \mathbf{q}, \mathbf{m}, \mathbf{m}', S, S'. \\ \langle s:\mathbf{p}\oplus\mathbf{q}:\mathbf{m}(S)\rangle_{\top} \wedge \langle s:\mathbf{p}\&\mathbf{q}:\mathbf{m}'(S')\rangle_{\top} \Rightarrow \langle s:\mathbf{p},\mathbf{q}:\mathbf{m}\rangle_{\top} \\ \wedge [s:\mathbf{p},\mathbf{q}:\mathbf{m}]Z \end{array} \right)$
<p>(2) Γ is deadlock-free, written $\text{df}(\Gamma)$, iff:</p> <p>$\Gamma \rightarrow^* \Gamma' \not\rightarrow$ implies $\text{end}(\Gamma')$</p>	$\Gamma \models \nu Z. \left(\begin{array}{l} (\forall s, \mathbf{p}, \mathbf{q}, \mathbf{m}. [s:\mathbf{p},\mathbf{q}:\mathbf{m}]_{\perp}) \Rightarrow \\ \forall \mathbf{p}, \mathbf{q}, \mathbf{m}, S. [s:\mathbf{p}\&\mathbf{q}:\mathbf{m}(S)]_{\perp} \wedge [s:\mathbf{p}\oplus\mathbf{q}:\mathbf{m}(S)]_{\perp} \\ \wedge \forall \mathbf{p}, \mathbf{q}, \mathbf{m}. [s:\mathbf{p},\mathbf{q}:\mathbf{m}]Z \end{array} \right)$
<p>(3) Γ is terminating, written $\text{term}(\Gamma)$, iff:</p> <p>Γ is deadlock-free, and there is $j \in \mathbb{N}^0$ such that for all $n \geq j$, $\Gamma = \Gamma_0 \rightarrow \Gamma_1 \rightarrow \dots \rightarrow \Gamma_n$ implies $\text{end}(\Gamma_n)$</p>	$\Gamma \models \mu Z. \left(\begin{array}{l} (\forall s, \mathbf{p}, \mathbf{q}, \mathbf{m}. [s:\mathbf{p},\mathbf{q}:\mathbf{m}]_{\perp}) \Rightarrow \\ \forall s, \mathbf{p}, \mathbf{q}, \mathbf{m}, S. [s:\mathbf{p}\&\mathbf{q}:\mathbf{m}(S)]_{\perp} \wedge [s:\mathbf{p}\oplus\mathbf{q}:\mathbf{m}(S)]_{\perp} \\ \wedge \forall s, \mathbf{p}, \mathbf{q}, \mathbf{m}. [s:\mathbf{p},\mathbf{q}:\mathbf{m}]Z \end{array} \right)$
<p>(4) Γ is never-terminating, written $\text{nterm}(\Gamma)$, iff:</p> <p>$\Gamma \rightarrow^* \Gamma'$ implies $\Gamma' \rightarrow$</p>	$\Gamma \models \nu Z. (\exists s, \mathbf{p}, \mathbf{q}, \mathbf{m}. \langle s:\mathbf{p},\mathbf{q}:\mathbf{m}\rangle_{\top} \wedge \forall s, \mathbf{p}, \mathbf{q}, \mathbf{m}. [s:\mathbf{p},\mathbf{q}:\mathbf{m}]Z)$
<p>(5) Γ is live, written $\text{live}(\Gamma)$, iff:</p> <p>$\varphi(\Gamma)$, for some φ such that</p> <p>[L-&] whenever $\varphi(\Gamma', s[\mathbf{p}]:S)$ with $S = \mathbf{q}\&_{i \in I} \mathbf{m}_i(S_i).S'_i$, $\exists i \in I: \exists \Gamma''$: $\Gamma', s[\mathbf{p}]:S \rightarrow^* \Gamma'', s[\mathbf{p}]:S'_i$</p> <p>[L-$\oplus$] whenever $\varphi(\Gamma', s[\mathbf{p}]:S)$ with $S = \mathbf{q}\oplus_{i \in I} \mathbf{m}_i(S_i).S'_i$, $\forall i \in I: \exists \Gamma''$: $\Gamma', s[\mathbf{p}]:S \rightarrow^* \Gamma'', s[\mathbf{p}]:S'_i$</p> <p>plus clauses [S-$\mu$], [S-$\rightarrow$] (Def. 4.1).</p>	$\Gamma \models \nu Z. \left(\begin{array}{l} \forall s, \mathbf{p}, \mathbf{q}. \\ \left((\exists \mathbf{m}, S. \langle s:\mathbf{p}\&\mathbf{q}:\mathbf{m}(S)\rangle_{\top}) \Rightarrow \right. \\ \left. \mu Z'. \exists \mathbf{m}. \langle s:\mathbf{p},\mathbf{q}:\mathbf{m}\rangle_{\top} \vee \exists \mathbf{p}', \mathbf{q}', \mathbf{m}'. \langle s:\mathbf{p}',\mathbf{q}':\mathbf{m}'\rangle_{Z'} \right) \\ \wedge \\ \forall \mathbf{m}. \left((\exists S. \langle s:\mathbf{p}\oplus\mathbf{q}:\mathbf{m}(S)\rangle_{\top}) \Rightarrow \right. \\ \left. \mu Z'. \langle s:\mathbf{p},\mathbf{q}:\mathbf{m}\rangle_{\top} \vee \exists \mathbf{p}', \mathbf{q}', \mathbf{m}'. \langle s:\mathbf{p}',\mathbf{q}':\mathbf{m}'\rangle_{Z'} \right) \\ \wedge \\ \forall \mathbf{m}. [s:\mathbf{p},\mathbf{q}:\mathbf{m}]Z \end{array} \right)$
<p>(6) Γ is live⁺, written $\text{live}^+(\Gamma)$, iff:</p> <p>$\varphi(\Gamma)$, for φ such that</p> <p>[L-&⁺] clause [L-&] above; <i>moreover</i>, $\Gamma', s[\mathbf{p}]:S$ belongs to some fair traversal set \mathbb{X} with targets \mathbb{Y} (Def. 5.5) such that, $\forall \Gamma_t \in \mathbb{Y}$, we have $\Gamma_t = \Gamma'', s[\mathbf{p}]:S'_i$ (for some $\Gamma'', i \in I$)</p> <p>[L-\oplus⁺] clause [L-\oplus] above, <i>plus</i> the “<i>moreover...</i>” part of [L-&⁺]</p> <p>plus clauses [S-μ], [S-\rightarrow] (Def. 4.1).</p>	$\Gamma \models \nu Z. \left(\begin{array}{l} \forall s, \mathbf{p}, \mathbf{q}. \\ \left((\exists \mathbf{m}, S. \langle s:\mathbf{p}\&\mathbf{q}:\mathbf{m}(S)\rangle_{\top}) \Rightarrow \right. \\ \left. \mu Z'. \exists \mathbf{m}. \langle s:\mathbf{p},\mathbf{q}:\mathbf{m}\rangle_{\top} \vee \exists \mathbf{p}', \mathbf{q}'. \left(\exists \mathbf{m}'. \langle s:\mathbf{p}',\mathbf{q}':\mathbf{m}'\rangle_{\top} \wedge \forall \mathbf{m}''. [s:\mathbf{p}',\mathbf{q}':\mathbf{m}'']_{Z'} \right) \right) \\ \wedge \\ \forall \mathbf{m}. \left((\exists S. \langle s:\mathbf{p}\oplus\mathbf{q}:\mathbf{m}(S)\rangle_{\top}) \Rightarrow \right. \\ \left. \mu Z'. \langle s:\mathbf{p},\mathbf{q}:\mathbf{m}\rangle_{\top} \vee \exists \mathbf{p}', \mathbf{q}'. \left(\exists \mathbf{m}'. \langle s:\mathbf{p}',\mathbf{q}':\mathbf{m}'\rangle_{\top} \wedge \forall \mathbf{m}''. [s:\mathbf{p}',\mathbf{q}':\mathbf{m}'']_{Z'} \right) \right) \\ \wedge \\ \forall \mathbf{m}. [s:\mathbf{p},\mathbf{q}:\mathbf{m}]Z \end{array} \right)$
<p>(7) Γ is live⁺⁺, written $\text{live}^{++}(\Gamma)$, iff:</p> <p>$\varphi(\Gamma)$, for φ such that</p> <p>[L-&⁺⁺] clause [L-&] above; <i>moreover</i>, $\exists n \in \mathbb{N}^0$ such that, whenever $\Gamma', s[\mathbf{p}]:S = \Gamma_0 \rightarrow \Gamma_1 \rightarrow \dots \rightarrow \Gamma_n$, then $\exists j \leq n, \Gamma''$ such that $\Gamma_j \rightarrow \Gamma'', s[\mathbf{p}]:S'_i$ (for some $i \in I$)</p> <p>[L-\oplus⁺⁺] clause [L-\oplus] above, <i>plus</i> the “<i>moreover...</i>” part of [L-&⁺⁺]</p> <p>plus clauses [S-μ], [S-\rightarrow] (Def. 4.1).</p>	$\Gamma \models \nu Z. \left(\begin{array}{l} \forall s, \mathbf{p}, \mathbf{q}. \\ \left((\exists \mathbf{m}, S. \langle s:\mathbf{p}\&\mathbf{q}:\mathbf{m}(S)\rangle_{\top}) \Rightarrow \right. \\ \left. \mu Z'. \exists \mathbf{m}. \langle s:\mathbf{p},\mathbf{q}:\mathbf{m}\rangle_{\top} \vee \left(\exists s', \mathbf{p}', \mathbf{q}', \mathbf{m}'. \langle s':\mathbf{p}',\mathbf{q}':\mathbf{m}'\rangle_{\top} \wedge \forall s'', \mathbf{p}'', \mathbf{q}'', \mathbf{m}''. [s'':\mathbf{p}'',\mathbf{q}'':\mathbf{m}']_{Z'} \right) \right) \\ \wedge \\ \forall \mathbf{m}. \left((\exists S. \langle s:\mathbf{p}\oplus\mathbf{q}:\mathbf{m}(S)\rangle_{\top}) \Rightarrow \right. \\ \left. \mu Z'. \langle s:\mathbf{p},\mathbf{q}:\mathbf{m}\rangle_{\top} \vee \left(\exists s', \mathbf{p}', \mathbf{q}', \mathbf{m}'. \langle s':\mathbf{p}',\mathbf{q}':\mathbf{m}'\rangle_{\top} \wedge \forall s'', \mathbf{p}'', \mathbf{q}'', \mathbf{m}''. [s'':\mathbf{p}'',\mathbf{q}'':\mathbf{m}']_{Z'} \right) \right) \\ \wedge \\ \forall \mathbf{m}. [s:\mathbf{p},\mathbf{q}:\mathbf{m}]Z \end{array} \right)$

Fig. 5. Properties of typing contexts. Each property is presented in two equivalent formalisations: the left-side ones are based on the notation and definitions introduced up to §5.4 (excluded); the right-side ones are μ -calculus formulas (explained in §6), and allow to verify typing contexts via model checking (e.g., with tools like mCRL2 [Groote and Mousavi 2014]).

The intuition behind $\text{live}/\text{live}^+/\text{live}^{++}$ is the following. Take a typing context $\Gamma, s[\mathbf{p}]:S$. If such a context is live , then, by clause $[\text{L-}\&]$ of Fig. 5(5), if S is an external choice, then Γ can reduce until *some* branch of S is triggered; and by clause $[\text{L-}\oplus]$, if S is an internal choice, then Γ can reduce allowing to send *each* message of S . The clauses of liveness^+ are stricter: they ensure that, under “fair scheduling” (details below) the interaction with S will be enabled in a finite number of steps. The clauses of liveness^{++} are even stricter, and ensure that the interaction with S will be enabled within a finite number of steps, no matter how other roles are scheduled. We will give examples and more explanations shortly (Ex. 5.10, Ex. 5.11, Ex. 5.14, Thm. 5.15). But first, we explain what “under fair scheduling” means: roughly, we ensure that there is a set of roles whose interactions *always* cause a desired input/output to meet a corresponding output/input. This requires some sophistication, and the formalisation of the “fair traversal set” mentioned in the definition of liveness^+ (Fig. 5(6)).

Definition 5.5 (Fair traversal set). Let \mathbb{X}, \mathbb{Y} be sets of typing contexts. We say that \mathbb{X} is a *fair traversal set with targets* \mathbb{Y} iff \mathbb{X} is closed under the rules:

$$\frac{\Gamma \in \mathbb{Y}}{\Gamma \in \mathbb{X}} \text{ [TS-TARGET]} \quad \frac{\exists s, \mathbf{p}, \mathbf{q} : \exists m : \Gamma \xrightarrow{s:\mathbf{p}, \mathbf{q}; m} \quad \forall m : \Gamma \xrightarrow{s:\mathbf{p}, \mathbf{q}; m} \Gamma' \text{ implies } \Gamma' \in \mathbb{X}}{\Gamma \in \mathbb{X}} \text{ [TS-COMM]}$$

Def. 5.5 says that if a fair traversal set \mathbb{X} contains a typing context Γ , then \mathbb{X} also contains (part of) Γ 's reductions (inductive rule $[\text{TS-COMM}]$), reaching one of the target contexts in \mathbb{Y} (base rule $[\text{TS-TARGET}]$). Notably, by rule $[\text{TS-COMM}]$, for each reduction of Γ , it is enough to choose just *two* roles \mathbf{p}, \mathbf{q} who can interact (clause “ $\exists m : \Gamma \xrightarrow{s:\mathbf{p}, \mathbf{q}; m}$ ”), as long as, for *all* interactions they can engage in, the corresponding reductum belongs to \mathbb{X} (clause “ $\dots \Gamma' \in \mathbb{X}$ ”). Consequently, if we prove that \mathbb{X} is a fair traversal set with targets \mathbb{Y} , then any $\Gamma \in \mathbb{X}$ is supported by an inductive derivation \mathcal{D} — that, in turn, shows how we can reach some $\Gamma' \in \mathbb{Y}$ in a finite number of steps, by choosing a set of participants and following *any* of their possible interactions (one per instance of $[\text{TS-COMM}]$ in \mathcal{D}).

Example 5.6. By Def. 5.5, fair traversal sets are inductively defined: this excludes cases where target elements are reachable, but can be “infinitely delayed” by choices and recursion. E.g., let:

$$\Gamma = s[\mathbf{p}]:\mu t. \mathbf{q} \oplus m_1. t, m_2, s[\mathbf{q}]:\mu t. \mathbf{p} \& \{m_1. t, m_2. \mathbf{r} \oplus m_3\}, s[\mathbf{r}]:\mathbf{q} \& m_3 \quad \text{and thus, } \Gamma \xrightarrow{s:\mathbf{p}, \mathbf{q}; m_2} \xrightarrow{s:\mathbf{q}, \mathbf{r}; m_3} \Gamma'$$

$$\Gamma' = s[\mathbf{p}]:\text{end}, s[\mathbf{q}]:\text{end}, s[\mathbf{r}]:\text{end}$$

Note that Γ is live , and Γ' is reachable — and yet, we *cannot* define a fair traversal set \mathbb{X} containing Γ , with a target set $\mathbb{Y} = \{\Gamma'\}$. This is because \mathbf{p}, \mathbf{q} can interact infinitely by exchanging m_1 , yielding the infinite run $\Gamma \xrightarrow{s:\mathbf{p}, \mathbf{q}; m_1} \Gamma \xrightarrow{s:\mathbf{p}, \mathbf{q}; m_1} \dots$; consequently, to support $\Gamma \in \mathbb{X}$ we would need an inductive derivation with an *infinite* series of instances of rule $[\text{TS-COMM}]$ — i.e., the derivation would be invalid.

Example 5.7. Fair traversal sets can be defined when elements of the target set are reachable, but can be infinitely delayed by “unfair scheduling.” E.g., consider:

$$\Gamma = s[\mathbf{p}]:\mathbf{q} \oplus m_1. \mathbf{q}' \oplus m_2, s[\mathbf{q}]:\mathbf{p} \& m_1, s[\mathbf{q}']: \mathbf{p} \& m_2, s[\mathbf{r}]:\mu t. \mathbf{r}' \oplus m_2. t, s[\mathbf{r}']: \mu t. \mathbf{r} \& m_2. t$$

$$\Gamma' = s[\mathbf{p}]:\text{end}, s[\mathbf{q}]:\text{end}, s[\mathbf{q}']: \text{end}, s[\mathbf{r}]:\mu t. \mathbf{r}' \oplus m_2. t, s[\mathbf{r}']: \mu t. \mathbf{r} \& m_2. t$$

Note that Γ is live , and Γ' is reachable from Γ , via the reductions $\Gamma \xrightarrow{s:\mathbf{p}, \mathbf{q}; m_1} \xrightarrow{s:\mathbf{p}, \mathbf{q}'; m_2} \Gamma'$; however, Γ' can be infinitely delayed in the unfair run $\Gamma \xrightarrow{s:\mathbf{r}, \mathbf{r}'; m_2} \Gamma \xrightarrow{s:\mathbf{r}, \mathbf{r}'; m_2} \dots$ that never fires the communication between \mathbf{p} and \mathbf{q} , and thus, never enables the interaction between \mathbf{p} and \mathbf{q}' . Yet, unlike Ex. 5.6, we *can* define a fair traversal set $\mathbb{X} = \{\Gamma, \Gamma'\}$, with target $\mathbb{Y} = \{\Gamma'\}$: in fact, we can build a finite derivation that supports $\Gamma \in \mathbb{X}$ by instantiating rule $[\text{TS-COMM}]$ twice — choosing \mathbf{p}, \mathbf{q} for the first reduction, and then \mathbf{p}, \mathbf{q}' to reach the axiom $[\text{TS-TARGET}]$, ignoring the interactions between \mathbf{r}, \mathbf{r}' .

Ex. 5.6 and Ex. 5.7 clarify why live^+ in Fig. 5(6) requires the existence of a certain traversal set: this ensures that, when Γ has some pending input/output, then under “fair scheduling,” Γ can reach a target Γ_t where such input/output has been fired, by interacting with a matching output/input.

Table 1. Verification of the multiparty protocols in Fig.4 against the properties in Fig.5. The results for protocol (3) hold for $n \geq 1$, while the results for protocol (4) hold for $n \geq 2$.

	consistent	safe	deadlock-free	live	live ⁺	live ⁺⁺	never-terminat.	terminat.
(1) OAuth2 fragment	false	true	true	true	true	true	false	true
(2) Rec. two-buyers	false	true	true	true	false	false	false	false
(3) Rec. map/reduce	false	true	true	true	true	true	false	false
(4) MP workers	false	true	true	true	true	false	false	false

5.4 Relationships Between Typing Context Properties


We now study how typing context properties are related: this is formalised in Lemma 5.9 below, that also conveys the expressiveness of our new type system (Remark 5.12).

To cover classic MPST theory, we first define *projected typing contexts*, in Def. 5.8; note that the projections with plain and full merging correspond to claims (C1) and (C2) in §3.1, respectively.

Definition 5.8. We say that Γ is the **full (resp. plain) projection of G for session s** , written $\text{fproj}_{G,s}(\Gamma)$ (resp. $\text{pproj}_{G,s}(\Gamma)$), iff $\Gamma = \{s[\mathbf{p}]:G[\mathbf{p}]_{\mathbf{p} \in \text{roles}(G)}\}$, where $G[\mathbf{p}]$ is the *projection with full merging (resp. plain merging)* in Def. 3.3.

LEMMA 5.9. For all Γ , the following (non-)implications hold:

- (1) $\text{consistent}(\Gamma) \iff \text{safe}(\Gamma)$;
- (2) $\text{live}(\Gamma) \iff \text{safe}(\Gamma)$;
- (3) $\text{live}(\Gamma) \iff \text{df}(\Gamma)$;
- (4) $\text{nterm}(\Gamma) \iff \text{df}(\Gamma)$;
- (5) $\text{consistent}(\Gamma) \iff \text{df}(\Gamma)$;
- (6) $\text{consistent}(\Gamma) \wedge \text{df}(\Gamma) \iff \text{live}(\Gamma)$;
- (7) $\text{live}^{++}(\Gamma) \iff \text{live}^+(\Gamma) \iff \text{live}(\Gamma)$;
- (8) $\text{term}(\Gamma) \iff \text{live}^{++}(\Gamma)$;
- (9) assume $\text{dom}(\Gamma) = \{s\}$ (Def. 2.6). Then:
 $\exists G : \text{fproj}_{G,s}(\Gamma) \iff \text{live}^+(\Gamma)$.



In the diagram, the “safe” set contains all typing contexts supported by our general type system. The red subsets are the classic MPST theory: G contains all contexts projected by some global type; its subset $G-$ only has *consistent* typing contexts, i.e. the only class of global types for which classic MPST proves type safety: this class excludes our example in §1, and also all protocols in Fig.4, and more (see Ex. 5.10 and Ex. 5.11 below). Notably, in item (9), we prove that all projected contexts are live^+ : this is discussed in Remark 5.16 later.

Example 5.10. The protocols described in Fig.4 are verified in Table 1. We observe:

- all protocols are safe and live, but *none of them* is consistent: hence, they are *not* supported by the classic MPST theory;
- all protocols are live^+ , except recursive two-buyers (2): this is because it allows **alice** and **bob** to bargain forever by exchanging **split/no** messages, without ever involving the **store** (that will keep waiting for **alice** to send either **buy** or **no**). This violates clause $[\text{L-}\&^+]$ of Fig. 5(6), because we cannot find any traversal set whose targets trigger the **store**’s pending input (the issue is similar to Ex. 5.6);
- two protocols are not live^{++} : recursive two-buyers (as expected, by the point above and the contrapositive of Lemma 5.9(7)), and MP workers (4). The latter is not live^{++} because each

triplet of workers $\mathbf{wa}_i, \mathbf{wb}_i, \mathbf{wc}_i$ ($i \in 1..n \geq 2$) can loop independently from the others; therefore, the interaction between, e.g., two workers in triplet 1 might be delayed for an unbounded number of transitions, while triplet 2 keeps progressing. Note that this scenario arises if the roles are scheduled unfairly; otherwise, each enabled interaction *will* be eventually fired, and this is reflected by the fact that the MP workers protocol is live^+ ;

- only the OAuth2 fragment (1) is terminating – while the other protocols are *neither* terminating, *nor* never-terminating: i.e., they might loop forever, but depending on the choices of one or more roles, they can reach a terminated state (where all roles have type **end**).

Example 5.11. We now provide some more small examples of multiparty protocols and their properties, complementing those discussed Ex. 5.10.

$\Gamma_A = s[\mathbf{p}]:\mathbf{q}\&\mathbf{m}_1.\mathbf{r}\&\mathbf{m}_3, s[\mathbf{q}]:\mathbf{r}\&\mathbf{m}_2.\mathbf{p}\&\mathbf{m}_1, s[\mathbf{r}]:\mathbf{p}\&\mathbf{m}_3.\mathbf{q}\&\mathbf{m}_2$ is consistent (hence safe), but *not* live *nor* deadlock-free: this is because its inputs/outputs, albeit dual, occur in the wrong order.

$\Gamma_B = s[\mathbf{p}]:\mu\mathbf{t}.\mathbf{q}\&\mathbf{m}_1.\mathbf{t}, s[\mathbf{q}]:\mu\mathbf{t}.\mathbf{p}\&\mathbf{m}_1.\mathbf{t}, s[\mathbf{r}]:\mathbf{p}\&\mathbf{m}_2$ is consistent, deadlock-free and safe, but *not* live: in fact, $s[\mathbf{p}], s[\mathbf{q}]$ reduce infinitely, but $s[\mathbf{r}]$ cannot fire its input (violating $[\mathbf{L}\&]$ in Fig. 5).

$\Gamma_C = s[\mathbf{p}]:S, s[\mathbf{q}]:\mathbf{p}\&\mathbf{m}(S).\mathbf{end}$ with $S = \mu\mathbf{t}.\mathbf{q}\&\mathbf{m}(\mathbf{t}).\mathbf{end}$ (from [Bernardi and Hennessy 2016, Ex. 1.2]) is terminating (hence live^{++} , and safe), but *not* projectable from any global type, *nor* consistent: this is because a recursion variable \mathbf{t} occurs as payload in S , which is disallowed by Def. 3.3 and Def. 3.8. Notably, Γ_C types the process below (from [Bernardi and Hennessy 2016, Ex. 1.2]): it creates infinitely many sessions s' where \mathbf{p} and \mathbf{q} exchange one message \mathbf{m} (note that this process, although deadlock-free, does not satisfy Def. 5.3(2)).

$$\emptyset \cdot \Gamma_C \vdash \mathbf{def} \ X(x:S, y:\mathbf{p}\&\mathbf{m}(S)) = P \ \mathbf{in} \ X\langle s[\mathbf{p}], s[\mathbf{q}] \rangle$$

$$\text{where } P = (\nu s':\Gamma'_C) \left(x[\mathbf{q}]\&\mathbf{m}(s'[\mathbf{p}]).0 \mid y[\mathbf{p}]\Sigma\mathbf{m}(z).X\langle z, s'[\mathbf{q}] \rangle \right) \quad \text{with } \Gamma'_C = s'[\mathbf{p}]:S, s'[\mathbf{q}]:\mathbf{p}\&\mathbf{m}(S).\mathbf{end}$$

REMARK 5.12. By Lemma 5.9(1,9), our general session type system instantiated with $\varphi = \text{fproj}_{G,s}$ subsumes the classic MPST theory, and also proves subject reduction and type safety in presence of “full-merging” global type projections: this is because consistency/projectability are limited syntactic approximations of safety/liveness. Hence, the typing rule $[\mathbf{T}\text{-V}_{\text{CLASSIC}G}]$ in §3 is valid in our theory, and we can type our opening example (Ex. 4.5), and support complex protocols rejected by classic MPST, such as all those listed in Fig. 4. This retroactively fixes some flawed results in literature, described in §3.1 (claim (C2)), and impacting the works listed in §8. Further, we support protocols for which no global type exists: see Ex. 5.10 (case “recursive two-buyers”) and Ex. 5.11 (case Γ_C).

5.5 Static Verification of Run-Time Process Properties

We now show that, by using the type-level properties in Fig. 5, we can predict and constrain the run-time behaviour of processes. Roughly, the intuition is: if we have $\Gamma \vdash P$, and some property in Fig. 5 holds for Γ , then a similar corresponding property from Def. 5.1 holds for P . From this it follows that, to ensure that a closed process $(\nu s)P$ has a desired property from Def. 5.1, we can correspondingly instantiate φ in Def. 4.4, and check if the judgement “ $\emptyset \vdash (\nu s:\Gamma)P$ with φ ” holds.

First, we highlight that all typing context properties mentioned thus far are decidable (Thm. 5.13 below) – unlike the run-time process properties in Def. 5.1. This is clear for consistency and projectability, that are syntactic and inductive; others (safety, liveness, ...) are decidable because, by Def. 2.8, typing contexts have finite-state transition systems. Consequently, by Thm. 4.9, type checking is decidable, if φ is instantiated with any property listed in Thm. 5.13.

THEOREM 5.13 (DECIDABILITY OF φ). $\varphi(\Gamma)$ is decidable, for all Γ , and for all φ such that

$$\varphi \in \left\{ \text{consistent, fproj}_{G,s}, \text{pproj}_{G,s}, \text{safe, term, nterm, df, live, live}^+, \text{live}^{++} \right\} \quad (\text{for any } G)$$

Now, assume $\Gamma \vdash P$. To predict the run-time behaviour of P from Γ , we need to overcome a complication: it might seem that if Γ is live (Fig. 5(5)), then P should be live, too. But this is *not* the case, due to a subtle interaction between the typing rule $[\text{T-SUB}]$ in Fig. 2, and the fact that *supertyping does not preserve liveness*: this issue (that is related to the problem of *fair subtyping*, studied by Padovani [2016]), is illustrated in Ex. 5.14 below. For this reason, in Thm. 5.15 we guarantee process liveness via the stronger type-level property live^+ : this is the payoff of fair traversal sets (Def. 5.5).

Example 5.14. Take Γ with the rec. two-buyer protocol (Fig. 4(2)): it is live (Table 1). Now, let:

$$\Gamma' = \begin{cases} s[\mathbf{a}]: s\oplus\text{query}(\text{Str}).s\&\text{price}(\text{Int}).\mu\mathbf{t}.\mathbf{b}\oplus\text{split}(\text{Int}).\mathbf{b}\&\{\text{yes}(\text{Int}).s\oplus\text{buy}, \text{no}.\mathbf{t}\} \\ s[\mathbf{s}]: \mathbf{a}\&\text{query}(\text{Str}).\mathbf{a}\oplus\text{price}(\text{Int}).\mathbf{a}\&\{\text{buy}.\text{end}, \text{no}.\text{end}\} \\ s[\mathbf{b}]: \mu\mathbf{t}.\mathbf{a}\&\{\text{split}(\text{Int}).\mathbf{a}\oplus\text{no}.\mathbf{t}, \text{cancel}.\text{end}\} \end{cases} \quad (\text{as in Fig. 4(2)})$$

i.e., the types of **alice** and **bob** in Γ' are *supertypes* (Def. 2.5) of those in Γ : **alice** never chooses to send **cancel** to **bob**, who in turn always answers **no** to all **split** proposals. We have $\Gamma \leq \Gamma'$ (Def. 2.5) and Γ' is safe (Lemma 4.3), but *not live*: after sending the **price**, the **store** will wait for either **buy** or **no** from **alice**, but neither message will ever be sent, while **alice** and **bob** loop by exchanging **split/no**. Consequently, a process P typed by Γ' can have two sub-processes implementing **alice** and **bob** that interact forever, while a sub-process implementing the **store** waits for a buy/no message, but will never receive it: hence, P is not live, as it does not satisfy Def. 5.1(2). Now, note that such P is also typed by Γ (via rule $[\text{T-SUB}]$ in Fig. 2): i.e., a live typing context can type a *non-live* process.

THEOREM 5.15. *Assume $\emptyset \cdot \Gamma \vdash P$, with Γ safe, $P \equiv \prod_{\mathbf{p} \in I} P_{\mathbf{p}}$, each $P_{\mathbf{p}}$ having guarded definitions and either being \emptyset (up-to \equiv), or only playing role \mathbf{p} in s . Then, (1) $\text{df}(\Gamma)$ implies that P is deadlock-free; (2) $\text{term}(\Gamma)$ implies that P is terminating; (3) $\text{nterm}(\Gamma)$ implies that P is never-terminating; (4) $\text{live}^+(\Gamma)$ implies that P is live; and (5) $\text{live}^{++}(\Gamma)$ implies that P is strongly live.*

PROOF. The results follow by Thm. 5.4 (session fidelity). For (4) we also use the fact that, if $\text{live}^+(\Gamma)$ and $\Gamma \leq \Gamma'$, then $\text{live}^+(\Gamma')$. \square

REMARK 5.16. *With Lemma 5.9(9) and Thm. 5.15(4), we uncover that global types / projections (Fig. 3) are ways to produce live^+ typing contexts, and ensure that processes are live. Since Thm. 5.15 does not need the technicalities of Fig. 3, our theory and results are more general than classic MPST. And importantly, the premises of all cases of Thm. 5.15 are decidable (by Thm. 5.13 and Thm. 4.9).*

6 VERIFYING TYPE-LEVEL PROPERTIES VIA MODEL CHECKING

Our new MPST theory (§ 4) is parametric on a general property φ , that is not constrained by syntactic duality/consistency. In this section, we leverage this distinguishing feature to integrate type checking and model checking, in two steps: (1) we show how to express φ as a *modal μ -calculus formula*, and (2) we use a model checker (through the paper's companion artifact) to verify whether the transitions of Γ satisfy the μ -calculus version of φ . This provides a practical method to verify whether $\varphi(\Gamma)$ holds — e.g., in rule $[\text{T-GEN-V}]$ (Def. 4.4), and in Thm. 5.15.

We focus on a fragment of the μ -calculus with data, adopting a formulation based on [Groote and Mousavi 2014, §6.5]. Let α range over the labels in Def. 2.8 — i.e., α can have the form $s:\mathbf{p}\&\mathbf{q}:\mathbf{m}(S)$ for input, or $s:\mathbf{p}\oplus\mathbf{q}:\mathbf{m}(S)$ for output, or $s:\mathbf{p},\mathbf{q}:\mathbf{m}$ for communication. Then, *μ -calculus formulas* are defined as follows, where \mathbf{d} (“data”) ranges over sessions, roles, message labels, and session types:

$$\phi ::= \top \mid \perp \mid [\alpha]\phi \mid \langle \alpha \rangle \phi \mid \phi_1 \wedge \phi_2 \mid \phi_1 \vee \phi_2 \mid \phi_1 \Rightarrow \phi_2 \mid \mu Z.\phi \mid \nu Z.\phi \mid Z \mid \forall \mathbf{d}.\phi \mid \exists \mathbf{d}.\phi$$

A formula ϕ accepts or rejects a typing context Γ depending on the sequences of actions that Γ can fire along its transitions. A formula can be either: true/false (\top/\perp), i.e., accept any/no typing context; box modality $[\alpha]\phi$ (“for all transitions with label α , the reached typing context must satisfy

"); diamond modality $\langle \alpha \rangle \phi$ ("for some transition with label α , the reached typing context satisfies ϕ "); implication \Rightarrow ; least/greatest fixed point $\mu Z. \phi / \nu Z. \phi$, allowing to iterate ϕ for a finite/infinite number of times; a variable Z , for iteration; and universal/existential quantification $\forall d. \phi / \exists d. \phi$. When a typing context Γ satisfies a formula ϕ , we write $\Gamma \models \phi$.

Example 6.1. The μ -calculus formula $\phi = \exists s. \exists p. \exists q. \exists m. \exists S. \langle s : p \oplus q : m(S) \rangle \top$ says: "accept a typing context if, for some session s , roles p and q , message label m , and type S , it can perform an output action $s : p \oplus q : m(S)$ " – and after such a transition, the reached typing context is always accepted, by \top . Therefore, if we take the typing context $\Gamma = s[r] : r' \oplus \text{msg}(\text{Str}). \text{end}$, then we have $\Gamma \xrightarrow{s : r \oplus r' : \text{msg}(\text{Str})} \top$ (by Def. 2.8), which means that Γ satisfies ϕ – in symbols, $\Gamma \models \phi$. Moreover, Γ satisfies the formula $\forall s. \forall p. \forall q. \forall m. [s : p, q : m] \perp$, that holds when *no* communication is possible, for any role: in fact, the formula says that any communication would reach a context rejected by \perp .

Instead, if we take the formula $\phi' = \exists s. \exists p. \exists q. \exists m. \langle s : p, q : m \rangle \top$, then Γ above does *not* satisfy ϕ' , because it requires a communication transition to be enabled. However, if we extend Γ as $\Gamma' = \Gamma, s[r'] : r \& \text{msg}(\text{Str}). \text{end}$, then we have both $\Gamma' \models \phi$ and $\Gamma' \models \phi'$ – and thus, $\Gamma' \models \phi \wedge \phi'$.

Example 6.2 (Formulas in Fig. 5). We now describe the μ -calculus formulas in Fig. 5:

- **safety (1)** checks that, if an output m and an input m' are enabled between two roles p and q , then they can communicate via m (i.e., by Def. 2.8, the output message m must be supported by the recipient). This must hold for any context reachable via communication transitions: this is enforced by the greatest fixed point $\nu Z. \dots$ and the clause $\dots \wedge [s : p, q : m] Z$;
- **deadlock-freedom (2)** checks whether communication is possible; if not (" $\forall \dots [s : p, q : m] \perp$ ", that holds only when no roles can interact, cf. Ex. 6.1), then (\Rightarrow) there must be no input nor output transitions enabled – i.e., all typing context entries must be **end**. This must hold for any context reachable via communications: it is enforced by $\nu Z. \dots$ and $\dots \wedge \forall \dots [s : p, q : m] Z$;
- **termination (3)** is similar to deadlock-freedom, but uses a *least* fixed point $\mu Z. \dots$: hence, the clause $\dots \wedge \forall \dots [s : p, q : m] Z$ can only iterate for a *finite* number of times, and then no communications, nor inputs, nor outputs must be enabled – i.e., all context entries are **end**;
- **never-termination (4)** checks that in any context reachable via communication transitions ($\nu Z. \dots$ and $\dots \wedge \forall \dots [s : p, q : m] Z$), some further communication is possible ($\exists \dots \langle s : p, q : m \rangle \top$);
- **liveness (5)** checks that, if an input or output between two roles p and q is enabled, then (\Rightarrow) a corresponding communication can be fired, after a finite sequence of communications among any role. The sequence is built with a least fixed point $\mu Z' \dots$, that can iterate on the clause $\dots \vee \exists \dots \langle s : p', q' : m' \rangle Z'$ for a finite number of times. The top-level greatest fixed point $\nu Z. \dots$ repeats the check for all contexts reachable via communication (clause $\dots \wedge \forall \dots [s : p, q : m] Z$);
- **liveness⁺ (6)** is similar to liveness, but the nested fixed points $\mu Z' \dots$ build finite sequences of communications by picking a pair of roles p', q' at each step, and following *all* their interactions, until a communication between p, q is enabled. This corresponds to building the fair traversal set (Def. 5.5) required by the left-side definition of live^+ in Fig. 5;
- **liveness⁺⁺ (7)** is also similar to liveness, but the nested fixed points $\mu Z' \dots$ build finite sequences by following *any* communication between *any* pair of roles, until a communication between p, q is enabled. This ensures that, along any execution path, after a finite number of steps, p and q will be able to interact, as in the left-side definition of live^{++} in Fig. 5.

Implementation. This paper has a companion artifact: a toolkit, called `mpstk` ("*MultiParty Session Types toolKit*"), that verifies the properties listed Fig. 5 (and described in Ex. 6.2). It is available at:

<https://alcestes.github.io/mpstk>

Internally, `mpstk` uses the mCRL2 model checker [Groote and Mousavi 2014], in combination with the theory in § 2.2 and § 4 (e.g., `mpstk` checks subtyping, as per Def. 2.5). We used `mpstk` to verify

Table 2. Average time (in seconds \pm std. dev.) for the verification of the protocols in Fig. 4. Protocols (3) and (4) are instantiated with $n=3$. The outcome of the verification is shown in Table 1. (Benchmarking specs: Intel Core i7-4790 CPU, 3.60GHz, 16 GB RAM, mCRL2 201808.0 invoked 30 times (by mpstk) with: pbes2bool --strategy=2)

	states	safe	deadlock-free	live	live ⁺	live ⁺⁺	never-terminat.	terminat.
(1) OAuth2 fragment	37	1.00 \pm 0%	1.00 \pm 0%	1.00 \pm 0%	1.00 \pm 0%	1.00 \pm 0%	1.00 \pm 0%	0.98 \pm 9%
(2) Rec. two-buyers	85	1.00 \pm 0%	1.00 \pm 0%	1.00 \pm 0%	1.00 \pm 0%	1.00 \pm 0%	1.00 \pm 0%	0.99 \pm 3%
(3) Rec. map/reduce	2561	1.00 \pm 0%	1.00 \pm 0%	1.00 \pm 0%	1.00 \pm 0%	1.00 \pm 0%	1.00 \pm 0%	0.99 \pm 3%
(4) MP workers	442369	1.01 \pm 4%	0.98 \pm 8%	0.98 \pm 9%	1.03 \pm 14%	1.02 \pm 7%	0.99 \pm 6%	1.00 \pm 1%

the protocols in Fig. 4: the results are in Table 1. We also measured the time needed to verify each case: the results are in Table 2. In all instances, the verification takes around one second. Notably, this also holds for the multiparty workers protocol (4), although it has 12000 \times more states than the OAuth2 fragment (1). This state space explosion is due to the interleaving of multiple parallel components – but still, its impact on verification time is minimal: in fact, the properties in Fig. 5 only follow the communication transitions of a typing context Γ , whereas the input and output transitions of Γ are checked for their presence/absence, but *not* followed to their destination state. Hence, mCRL2 can verify our formulas in Fig. 5 without exploring the whole state space of Γ .

7 ASYNCHRONOUS MULTIPARTY SESSION π -CALCULUS

In its original formulation [Bettini et al. 2008; Honda et al. 2008], the MPST π -calculus has *asynchronous* buffered semantics, to model typical “real-world” distributed message-passing programs. Our new theory extends to asynchrony, overcoming challenges due to queue handling and decidability. Due to space limits, we summarise the main results from [Scalas and Yoshida 2018a].

Asynchronous MPST. We give an intuition of the asynchronous calculus with an example:

$$\begin{aligned} s[\mathbf{p}][\mathbf{q}] \oplus m(s'[\mathbf{r}]).P \mid s[\mathbf{q}][\mathbf{p}] \sum m(x).Q \mid s \blacktriangleright \epsilon \\ \rightarrow P \mid s[\mathbf{q}][\mathbf{p}] \sum m(x).Q \mid s \blacktriangleright (\mathbf{p}, \mathbf{q}, m(s'[\mathbf{r}])) \cdot \epsilon \rightarrow P \mid Q\{s'[\mathbf{r}]/x\} \mid s \blacktriangleright \epsilon \end{aligned} \quad (10)$$

In the topmost process, $s \blacktriangleright \epsilon$ is the (empty) *message queue of session s* (not present in the calculus of §2.1). The first reduction enqueues the *pending message* $(\mathbf{p}, \mathbf{q}, m(s'[\mathbf{r}]))$, meaning that \mathbf{p} has sent to \mathbf{q} a message with label m and payload $s'[\mathbf{r}]$. With the second reduction, the message is received.

The classic async MPST typing judgement has the following form:

$$\Theta \cdot \Gamma \vdash_S P \quad (11)$$

where S is the set of sessions whose queue occurs in P (e.g., to type (10) above, we let $S = \{s\}$). Types are extended to model pending messages; e.g., the processes in (10) are typed by, respectively:

$$\begin{aligned} \Gamma &= s[\mathbf{p}]:\mathbf{q} \oplus m(S').S, s[\mathbf{q}]:\mathbf{p} \& m(S').T, s'[\mathbf{r}]:S' \\ \Gamma' &= s[\mathbf{p}]:(\mathbf{q}!m(S') \cdot \epsilon; S), s[\mathbf{q}]:\mathbf{p} \& m(S').T, s'[\mathbf{r}]:S' \\ \Gamma'' &= s[\mathbf{p}]:S, s[\mathbf{q}]:T, s'[\mathbf{r}]:S' \end{aligned} \quad (12)$$

Note that Γ above is a typing context similar to Def. 2.6. Instead, in Γ' the type of $s[\mathbf{p}]$ is a pair $(M; S)$, where $M = \mathbf{q}!m(S') \cdot \epsilon$ is a *message queue type* (abstracting the pending messages sent through $s[\mathbf{p}]$), followed by the continuation type S . In Γ' , the topmost queued message type matches the branching type of $s[\mathbf{q}]$: their interaction leads to Γ'' , with a reduction similar to Def. 2.8.

The classic async MPST theory has all the issues described in §3 – but the presence of message queues makes its subject reduction statement more complicated [Coppo et al. 2015a, Lemma 1]:

$$\begin{aligned} \text{If } \Theta \cdot \Gamma \vdash_S P \text{ and } \exists \Gamma_0 \text{ such that } \Gamma, \Gamma_0 \text{ consistent and } P \rightarrow P', \\ \text{then } \exists \Gamma', \Gamma'_0 \text{ consistent: } \Gamma, \Gamma_0 \rightarrow^* \Gamma', \Gamma'_0 \text{ and } \Theta \cdot \Gamma' \vdash_S P' \end{aligned}$$

General Asynchronous MPST. We extend our new theory in §4 to asynchronous MPST, and prove a simpler and more general subject reduction statement: Thm. 7.1. To achieve it, we develop async typing rules based on an *async safety property* φ , with a more sophisticated *async typing context reduction* $\rightarrow_{\mathcal{S}}$, where \mathcal{S} is a set of sessions, as in (11); e.g., in (12) we have $\Gamma \rightarrow_{\{s\}} \Gamma' \rightarrow_{\{s\}} \Gamma''$.

THEOREM 7.1 (ASYNC SUBJECT REDUCTION). *Assume $\Theta \cdot \Gamma \vdash_{\mathcal{S}} P$ with Γ asynchronously safe. Then, $P \rightarrow P'$ implies $\exists \Gamma' \text{ asynchronously safe such that } \Gamma \rightarrow_{\mathcal{S}}^* \Gamma' \text{ and } \Theta \cdot \Gamma' \vdash_{\mathcal{S}} P'$.*

We define asynchronous variants of φ , similar to those in Fig. 5; and by suitably instantiating φ , we ensure that typed async processes are deadlock-free/live, similarly to Thm. 5.15.

(Un-)Decidability of Type Checking. A result akin to Thm. 4.9 holds for async MPST.

THEOREM 7.2. *If φ is decidable, then “ $\Theta \cdot \Gamma \vdash_{\mathcal{S}} P$ with φ ” is decidable.*

However, under asynchrony we do *not* have a decidability result for φ as general as Thm. 5.13. On the contrary, async safety and most other properties are *undecidable*: the pairing of a session type with a message queue (cf. Γ' in (12)) corresponds to a Communicating Finite-State Machine (CFSM) [Brand and Zafriopulo 1983], and makes typing contexts Turing-powerful [Bartoletti et al. 2016, Thm. 2.5]. Still, we obtain decidable instances of φ through various sound approximations:

- (M1) consistency is decidable, and implies asynchronous safety;
- (M2) via the session type / CFSM correspondence established in [Deniérou and Yoshida 2013], we show that if Γ is *synchronously live* (Fig. 5(5)), decidable by Thm. 5.13, then Γ is also *asynchronously live*; we extend the result to *live⁺* (Fig. 5(6)); and by Lemma 5.9(9), this means that any Γ projected from a global type is *asynchronously live⁺*;
- (M3) given $n \geq 1$, we can decide if Γ enqueues at most n messages; if so, Γ is finite-state, hence async safety/liveness are decidable. For example, take $\Gamma = s[\mathbf{p}]:\mathbf{q}\oplus\mathbf{m}_1.\mathbf{q}\&\mathbf{m}_2, s[\mathbf{q}]:\mathbf{p}\oplus\mathbf{m}_2.\mathbf{p}\&\mathbf{m}_1$: it is deadlocked under synchronous semantics, and not projectable from any global type – but under asynchrony, the top-level outputs of \mathbf{p} and \mathbf{q} can be both enqueued, and then received; hence, we can decide that Γ enqueues at most 2 messages, and is *asynchronously live*.

REMARK 7.3. *By instantiating φ in Thm. 7.2 with one of the methods above, we obtain an expressive decidable fragment of our new asynchronous MPST theory: (M1) subsumes classic async MPST; (M2) covers all live typing contexts, albeit non-consistent: e.g., it covers all cases in Fig. 4, and all global types (by Lemma 5.9(9)); (M3) covers more typing contexts that are not projectable from global types.*

8 CONCLUSION, RELATED AND FUTURE WORK

We have presented a new theory of multiparty sessions types, with novel foundations that do *not* depend on duality/consistency, *nor* global types, *nor* projections. Our new theory subsumes classic MPST, also fixing subject reduction flaws in previous work (Remark 5.16). Moreover, our new type system is modular and reusable: by fine-tuning its parameter φ , we ensure that type-checking is decidable, and that processes are safe, deadlock-free, and live. A summary of the main results:

- (R1) our type safety results (Thm. 4.6, Cor. 4.7) are much more general than classic MPST;
- (R2) if we instantiate φ with projection/consistency, or any property in Fig. 5, then the type checking judgement “ $\Theta \cdot \Gamma \vdash P$ with φ ” is decidable. This follows from Thm. 4.9 and Thm. 5.13;
- (R3) by suitably choosing φ in (R2) above, we can statically guarantee that P “inherits” φ , and has certain desired run-time properties. This is formalised in Thm. 5.15;
- (R4) we can implement φ in (R2)/(R3) above as a syntactic check (Remark 5.12), or as a μ -calculus formula (Fig. 5). In the latter case, we can verify whether Γ satisfies φ via model checking – e.g., using mCRL2, through the paper’s companion artifact (mpsttk). This is shown in §6;
- (R5) our new theory extends to asynchronous communication, as illustrated in §7.

8.1 Classic Multiparty Session Types (MPST)

The classic MPST framework, and its notions of *global types* and *projections*, were introduced by Honda et al. [2008], with *linearity conditions* to check the well-formedness of global types, and ensure *projectability* of local types. Later, Bettini et al. [2008] proposed a simplified MPST system adopted by most works, including ours.

We now classify some related works w.r.t. their use of projection/consistency:

	papers	projection	consistency	subj. red.	claim
(a)	Bettini et al. [2008]; Carbone et al. [2016, 2015]; Coppo et al. [2015a]; Honda et al. [2008, 2016] Caires and Pérez [2016]; Chen [2015]; Deniérou et al.	\leq plain	yes	correct	(C1)
(b)	[2012]; Deniérou and Yoshida [2012]; Toninho and Yoshida [2016]	\geq full	no	flawed	(C2)
(c)	Scalas et al. [2017a]; Toninho and Yoshida [2017]	full	yes (required)	correct	(C1)

Row (a) lists works using *plain* (or stricter) global type projection (Def. 3.3), guaranteeing consistency. As shown in § 5.4, our theory captures plain projection / consistency by setting its parameter φ as $\varphi = \text{pproj}_{G,s} / \varphi = \text{consistent}$; however, this excludes many valid protocols, as per claim (C1) – e.g., all our examples in Fig. 4.

Row (b) lists works using *full* (or more flexible) global type projection, originally introduced in Yoshida et al. [2010] to support more protocols. Such works overlook the consistency requirement; and in § 3, we reveal that classic MPST subject reduction proofs relying on full projection (without consistency) are flawed, as per claim (C2). To “fix” these works within the classic MPST theory, we must require consistency, as done by works in row (c) – but this restricts typability, thus falling back into claim (C1). Instead, by Remark 5.12, our new MPST theory supports full projections with $\varphi = \text{fproj}_{G,s}$, thus subsuming classic MPST and fixing flaws, without losing expressiveness.

8.2 Non-Classic Multiparty Session Types

To the best of our knowledge, there are two MPST works (mentioned in Remark 3.1) that are *not* based on classic projection+consistency (Fig. 3) – but have other limitations, that we surmount.

The first work is by Dezani-Ciancaglini et al. [2015]: it presents a *single-session* type system, with first-order session types (i.e., without channel-passing); it is rooted on global types and their projections, but does *not* require consistency. Such a type system is subsumed by letting $\varphi = \text{fproj}_{G,s}$ in our Def. 4.4; in addition, our work also supports higher-order types, multiple interleaved sessions, and protocols for which no global type exists (see Table 1(2), and Ex. 5.11, case Γ_C).

The second non-classic MPST work is by Scalas and Yoshida [2018b]: it was our first attempt (and, to the best of our knowledge, the first work in general) to directly address the limitations of consistency (claim (C1)), and propose a behavioural theory of MPST, *not* based on global types and projections. Unfortunately, we could not build upon that work, due to its intrinsic limitations:

- (1) a major goal of this paper is subsuming and fixing classic MPST (cf. claim (C2) in § 1, and § 3). However, the theory of Scalas and Yoshida [2018b] *cannot* achieve this goal: it has different (and more complicated) typing rules that require typing context liveness, and do not support consistency. Our new theory, instead, supports both consistency and liveness, as instances of φ (Lemma 5.9, Remark 5.12);
- (2) from Scalas and Yoshida [2018b], we reuse the definition of typing context liveness (Fig. 5(5)) – but we show that it is insufficient to guarantee *process* liveness (Def. 5.1, Ex. 5.14). Hence, we develop the stronger properties $\text{live}^+ / \text{live}^{++}$ (Fig. 5(6,7)), to obtain the results on run-time process behaviour in Thm. 5.15. Such results are absent in Scalas and Yoshida [2018b];
- (3) the branching/selection typing rules of Scalas and Yoshida [2018b] (Fig. 3) directly inspect typing context reductions. This peculiarity is not problematic under synchronous semantics

(where typing contexts have finite-state transition systems), and in some cases, it enables flexible typing judgements that cannot be obtained in classic MPST [Scalas and Yoshida 2018b, Ex. 5.5]. The drawback is that, when extended to asynchronous semantics, typing contexts become Turing-powerful (§7), and typing rules that inspect their reductions become inherently undecidable; consequently, the theory of Scalas and Yoshida [2018b] does not subsume classic works on asynchronous MPST, and cannot achieve this goal without a major overhaul. Instead, our typing rules do *not* inspect typing context reductions, but only check whether the parametric property φ holds: hence, type checking is decidable whenever φ is decidable (Thm. 7.2), and this allows us to subsume classic asynchronous MPST (Remark 7.3).

By instantiating $\varphi = \text{live}$ in Def. 4.4, this paper largely subsumes Scalas and Yoshida [2018b]’s work – minus some corner cases based on the inspection of typing context reductions (cf. item (3) above).

8.3 Binary Sessions Without Duality

Our work yields a generalised theory of *binary* sessions *not* based on classic duality (Def. 3.5), subsuming classic papers based on [Honda et al. 1998]. If we take a binary session typing context $\Gamma = s[\mathbf{p}]:S, s[\mathbf{q}]:T$, our Lemma 5.9 becomes:

$$\exists G: \text{fproj}_{G,s}(\Gamma) \not\Leftarrow \implies \text{consistent}(\Gamma) \not\Leftarrow \implies \text{live}^{++}(\Gamma) \iff (\text{safe}(\Gamma) \text{ and } \text{df}(\Gamma)) \quad (13)$$

Here, the leftmost “ $\not\Leftarrow$ ” is due to supertyping: e.g., if we take the global type $G = \mathbf{p} \rightarrow \mathbf{q}: \{m, m'\}$, it projects the typing context $\Gamma = s[\mathbf{p}]:\mathbf{q} \oplus \{m, m'\}, s[\mathbf{q}]:\mathbf{p} \& \{m, m'\}$, that is consistent and live^{++} (hence safe); however, if we replace \mathbf{p} ’s entry with the supertype $\mathbf{p} \oplus m$, the resulting context is still live^{++} and consistent, but *not* projectable from any global type. The other “ $\not\Leftarrow$ ” in (13) is due to *non-tail-recursive types* like $\mu t. \mathbf{q} \oplus m(t). \text{end}$: they have no dual in classic binary session types (since t is a forbidden payload): thus, they yield non-consistent typing contexts, and processes like P in Ex. 5.11 (case Γ_C) cannot be typed. This limitation has been addressed by several authors, extending duality with various pitfalls (see e.g. [Bernardi and Hennessy 2016, §5.3]): for a survey, and a logic-based solution, see [Lindley and Morris 2016, §3.2]. By not using duality, our theory eschews these issues.

8.4 Type Systems for the π -Calculus

Many type systems have been proposed for the π -calculus, also influencing MPST: see survey in [Hüttel et al. 2016]. Our new MPST theory is a case of *behavioural* type system: it treats *types as simple processes* that reduce and evolve along a typed computation; and since types are simpler than programs, they can be analysed with simpler methods (e.g., finite model checking via our parameter φ , cf. §6). As stated in §4, the design of our new MPST theory is inspired by Igarashi and Kobayashi [2004]’s Generic Type System (GTS) for the π -calculus: i.e., we define a type system that is parametric on a property φ , and we prove type safety under the weakest φ ; then, we fine-tune φ to statically verify stronger properties of processes, like deadlock-freedom and liveness (§5). Besides this general analogy, our treatment is wholly different: we carefully reuse fundamental MPST definitions (§2.1) and develop new and more general results (§4, §5) to ensure our new theory fully subsumes the classic one; moreover, for async MPST we devise a new treatment of queue types, obtaining a new, more general subject reduction result (Thm. 7.1).

As an alternative, we might have tried to encode MPST in the GTS, and develop our new results from there. However, this appears unfeasible. Gay et al. [2014] tried the approach for *binary* sessions, obtaining drawbacks in terms of complication and loss of abstraction (see “Assessment” in Gay et al. [2014]): such drawbacks would be greatly amplified for multiparty sessions. Moreover, [Igarashi and Kobayashi 2004, §4.2, §5] study process/type correspondence using a temporal logic without fixed points, with limited support for recursion: their logic would not allow, e.g., to model our variants of liveness (Fig. 5) and address the interplay between liveness, subtyping, and recursion

(Ex. 5.14, Thm. 5.15). Further, the encoding approach would not work for async MPST: the types of Igarashi and Kobayashi [2004] lack message queues, and are akin to CCS without restriction, with decidable reachability [He 2011, p. 374]; hence, they cannot encode the Turing-powerful typing contexts of async MPST (§7), whose reachability is undecidable.

8.5 Choreographies and Communicating Finite-State Machines (CFSMs)

Various works model and verify multiparty protocols, a.k.a. *choreographies*, via automata-theoretic methods, by representing each party as a CFM [Brand and Zafiropolo 1983]. The safety their interactions (that is generally undecidable) is verified with two main approaches: (a) assume the decidability of a *synchronisability* property [Basu and Bultan 2011, 2016; Basu et al. 2012], and then check temporal properties of CFM via model checking; (b) check decidable *synchronous* execution conditions on CFMs, and prove that they ensure safe *asynchronous* executions [Bocchi et al. 2015; Deniérou and Yoshida 2013; Lange et al. 2015]. Both methods can help extending our new MPST theory: since we essentially treat async typing contexts as systems of CFMs (§7), new decidable results on CFM safety can yield new decidable instances of our type system (Thm. 7.2). Unfortunately, synchronisability has been recently proven *undecidable* by Finkel and Lozes [2017]: i.e., method (a) above might be unusable — hence, we adopt method (b) (cf. (M2) in §7). Unlike this paper, the above CFM works do *not* study type systems, nor properties of typed processes.

8.6 Future Work

We kept our typing rules close to classic MPST, to easily combine our results with existing works. E.g., we plan to integrate our work with Coppo et al. [2015b], that studies MPST deadlock-freedom in presence of *multiple* interleaved sessions: our generalised typing rules can be a drop-in replacement for the classic rules used by Coppo et al. [2015b], and this integration would combine their global deadlock-freedom checks, with our improved type safety results for individual sessions. We also plan to extend the calculus (e.g., with polymorphism [Caires and Pérez 2016; Goto et al. 2016]), and expand the properties/formulas studied in Fig. 5 and Thm. 5.15. We will investigate the logical foundations of our new MPST theory, aiming at results that generalise those by Carbone et al. [2016, 2015], which are focused on limited global types, projections, and consistency.

Another interesting research topic is the *completeness* of safety (Def. 4.1), i.e., studying whether the inverse implication w.r.t. Thm. 4.6/Cor. 4.7 holds. This corresponds to the following conjecture:

Take any Γ . If $\forall P, P': \Gamma \vdash P$ and $P \rightarrow^ P'$ implies that P' has no error, then $\text{safe}(\Gamma)$.*

We will investigate whether this conjecture holds — and if not, what other completeness results are achievable. Since session subtyping is central for defining safety (via clause $[\text{S} \rightarrow]$ in Def. 4.1, and $[\Gamma\text{-COMM}]$ in Def. 2.8), we will leverage Chen et al. [2017]’s work on the completeness of subtyping.

We will also study how to implement our new MPST theory. A basis is the work by Scalas et al. [2017a,b], that embeds classic MPST in Scala, through a linear π -calculus encoding based on consistency; however, since we do *not* require consistency, the work by Scalas et al. [2017a,b] only covers a fragment of our new theory. Using the μ -calculus formulas illustrated in §6, a new implementation can verify typing context properties by offloading them to a model checker.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their helpful remarks. Thanks to Francisco Ferreira, Sung-Shik Jongmans, and Julien Lange for their comments, and to Simon Castellan for testing the companion artifact. This work was partially supported by EPSRC (projects EP/K034413/1, EP/K011715/1, EP/L00058X/1, EP/N027833/1, EP/N028201/1), and by the EU COST Action CA15123 (“EUTypes”).

REFERENCES

- Davide Ancona, Viviana Bono, Mario Bravetti, Joana Campos, Giuseppe Castagna, Pierre-Malo Deniérou, Simon J. Gay, Nils Gesbert, Elena Giachino, Raymond Hu, Einar Broch Johnsen, Francisco Martins, Viviana Mascardi, Fabrizio Montesi, Rumyana Neykova, Nicholas Ng, Luca Padovani, Vasco T. Vasconcelos, and Nobuko Yoshida. 2017. Behavioral Types in Programming Languages. *Foundations and Trends in Programming Languages* 3(2-3) (2017). <https://doi.org/10.1561/2500000031>
- Massimo Bartoletti, Alceste Scalas, Emilio Tuosto, and Roberto Zunino. 2016. Honesty by Typing. *LMCS* 12(4) (2016). [https://doi.org/10.2168/LMCS-12\(4:7\)2016](https://doi.org/10.2168/LMCS-12(4:7)2016)
- Samik Basu and Tevfik Bultan. 2011. Choreography conformance via synchronizability. In *WWW*.
- Samik Basu and Tevfik Bultan. 2016. On deciding synchronizability for asynchronously communicating systems. *Theor. Comput. Sci.* 656 (2016).
- Samik Basu, Tevfik Bultan, and Meriem Ouederni. 2012. Synchronizability for Verification of Asynchronously Communicating Systems. In *VMCAI*.
- Giovanni Bernardi and Matthew Hennessy. 2016. Using higher-order contracts to model session types. *LMCS* 12(2) (2016). [https://doi.org/10.2168/LMCS-12\(2:10\)2016](https://doi.org/10.2168/LMCS-12(2:10)2016)
- Lorenzo Bettini, Mario Coppo, Loris D'Antoni, Marco De Luca, Mariangiola Dezani-Ciancaglini, and Nobuko Yoshida. 2008. Global Progress in Dynamically Interleaved Multiparty Sessions. In *CONCUR*. https://doi.org/10.1007/978-3-540-85361-9_33
- Laura Bocchi, Julien Lange, and Nobuko Yoshida. 2015. Meeting Deadlines Together. In *CONCUR*. <https://doi.org/10.4230/LIPICs.CONCUR.2015.283>
- Daniel Brand and Pitro Zafiropolo. 1983. On Communicating Finite-State Machines. *JACM* 30, 2 (1983). <https://doi.org/10.1145/322374.322380>
- Nadia Busi, Maurizio Gabbriellini, and Gianluigi Zavattaro. 2009. On the expressive power of recursion, replication and iteration in process calculi. *Mathematical Structures in Computer Science* 19, 6 (2009). <https://doi.org/10.1017/S096012950999017X>
- Luis Caires and Jorge A. Pérez. 2016. Multiparty Session Types Within a Canonical Binary Theory, and Beyond. In *FORTE*. https://doi.org/10.1007/978-3-319-39570-8_6
- Luis Caires, Frank Pfenning, and Bernardo Toninho. 2016. Linear logic propositions as session types. *MSCS* 26, 3 (2016).
- Marco Carbone, Sam Lindley, Fabrizio Montesi, Carsten Schürmann, and Philip Wadler. 2016. Coherence Generalises Duality: A Logical Explanation of Multiparty Session Types. In *CONCUR*. <https://doi.org/10.4230/LIPICs.CONCUR.2016.33>
- Marco Carbone, Fabrizio Montesi, Carsten Schürmann, and Nobuko Yoshida. 2015. Multiparty Session Types as Coherence Proofs. In *CONCUR*. <https://doi.org/10.4230/LIPICs.CONCUR.2015.412>
- Tzu-Chun Chen, Mariangiola Dezani-Ciancaglini, Alceste Scalas, and Nobuko Yoshida. 2017. On the Preciseness of Subtyping in Session Types. *Logical Methods in Computer Science* 13, 2 (2017). [https://doi.org/10.23638/LMCS-13\(2:12\)2017](https://doi.org/10.23638/LMCS-13(2:12)2017)
- Tzu-Chun Chen. 2015. Lightening global types. *JLAMP* 84, 5 (2015). <https://doi.org/10.1016/j.jlamp.2015.06.003>
- Mario Coppo, Mariangiola Dezani-Ciancaglini, Luca Padovani, and Nobuko Yoshida. 2015a. A Gentle Introduction to Multiparty Asynchronous Session Types. In *Formal Methods for Multicore Programming*. https://doi.org/10.1007/978-3-319-18941-3_4
- Mario Coppo, Mariangiola Dezani-Ciancaglini, Nobuko Yoshida, and Luca Padovani. 2015b. Global Progress for Dynamically Interleaved Multiparty Sessions. *MSCS* 760 (2015). <https://doi.org/10.1017/S0960129514000188>
- Pierre-Malo Deniérou, Nobuko Yoshida, Andi Bejleri, and Raymond Hu. 2012. Parameterised Multiparty Session Types. *LMCS* 8, 4 (2012). [https://doi.org/10.2168/LMCS-8\(4:6\)2012](https://doi.org/10.2168/LMCS-8(4:6)2012)
- Pierre-Malo Deniérou and Nobuko Yoshida. 2012. Multiparty Session Types Meet Communicating Automata. In *ESOP*. https://doi.org/10.1007/978-3-642-28869-2_10
- Pierre-Malo Deniérou and Nobuko Yoshida. 2013. Multiparty Compatibility in Communicating Automata: Characterisation and Synthesis of Global Session Types. In *ICALP*. https://doi.org/10.1007/978-3-642-39212-2_18
- Mariangiola Dezani-Ciancaglini, Silvia Ghilezan, Svetlana Jaksic, Jovanka Pantovic, and Nobuko Yoshida. 2015. Precise subtyping for synchronous multiparty sessions. In *PLACES*. <https://doi.org/10.4204/EPTCS.203.3>
- Alain Finkel and Etienne Lozes. 2017. Synchronizability of Communicating Finite State Machines is not Decidable. In *ICALP*. <https://doi.org/10.4230/LIPICs.ICALP.2017.122>
- Simon Gay and António Ravara. 2017. *Behavioural Types: From Theory to Tools*. River Publishers, Series in Automation, Control and Robotics. <https://doi.org/10.13052/rp-9788793519817>
- Simon J. Gay. 2016. Subtyping Supports Safe Session Substitution. In *A List of Successes That Can Change the World: Essays Dedicated to Philip Wadler on the Occasion of His 60th Birthday (LNCS)*, Vol. 9600. https://doi.org/10.1007/978-3-319-30936-1_5
- Simon J. Gay, Nils Gesbert, and António Ravara. 2014. Session Types as Generic Process Types. In *EXPRESS/SOS*. <https://doi.org/10.4204/EPTCS.160.9>

- Simon J. Gay and Malcolm Hole. 2005. Subtyping for session types in the π -calculus. *Acta Inf.* 42, 2-3 (2005). <https://doi.org/10.1007/s00236-005-0177-z>
- Jean-Yves Girard. 1987. Linear Logic. *TCS* 50 (1987), 1–102.
- Matthew Goto, Radha Jagadeesan, Alan Jeffrey, Corin Pitcher, and James Riely. 2016. An extensible approach to session polymorphism. *Mathematical Structures in Computer Science* 26, 3 (2016). <https://doi.org/10.1017/S0960129514000231>
- Jan Friso Groote and Mohammad Reza Mousavi. 2014. *Modeling and Analysis of Communicating Systems*. The MIT Press.
- Chaodong He. 2011. The Decidability of the Reachability Problem for CCS!. In *CONCUR*. https://doi.org/10.1007/978-3-642-23217-6_25
- Kohei Honda, Vasco Thudichum Vasconcelos, and Makoto Kubo. 1998. Language Primitives and Type Discipline for Structured Communication-Based Programming. In *ESOP*. <https://doi.org/10.1007/BFb0053567>
- Kohei Honda, Nobuko Yoshida, and Marco Carbone. 2008. Multiparty asynchronous session types. In *POPL*. <https://doi.org/10.1145/1328438.1328472> Full version in [Honda et al. 2016].
- Kohei Honda, Nobuko Yoshida, and Marco Carbone. 2016. Multiparty Asynchronous Session Types. *J. ACM* 63, 1, Article 9 (2016). <https://doi.org/10.1145/2827695>
- Hans Hüttel, Ivan Lanese, Vasco T. Vasconcelos, Luís Caires, Marco Carbone, Pierre-Malo Deniérou, Dimitris Mostrous, Luca Padovani, António Ravara, Emilio Tuosto, Hugo Torres Vieira, and Gianluigi Zavattaro. 2016. Foundations of Session Types and Behavioural Contracts. *ACM Comput. Surv.* 49, 1, Article 3 (2016). <https://doi.org/10.1145/2873052>
- Atsushi Igarashi and Naoki Kobayashi. 2004. A generic type system for the π -calculus. *TCS* 311, 1 (2004). [https://doi.org/10.1016/S0304-3975\(03\)00325-6](https://doi.org/10.1016/S0304-3975(03)00325-6)
- Naoki Kobayashi and Davide Sangiorgi. 2010. A hybrid type system for lock-freedom of mobile processes. *TOPLAS* 32, 5 (2010). <https://doi.org/10.1145/1745312.1745313>
- Julien Lange, Emilio Tuosto, and Nobuko Yoshida. 2015. From Communicating Machines to Graphical Choreographies. In *POPL*. <https://doi.org/10.1145/2676726.2676964>
- Sam Lindley and J. Garrett Morris. 2016. Talking Bananas: Structural Recursion for Session Types. In *ICFP*. <https://doi.org/10.1145/2951913.2951921>
- Barbara H. Liskov and Jeannette M. Wing. 1994. A Behavioral Notion of Subtyping. *TOPLAS* 16, 6 (1994). <https://doi.org/10.1145/197320.197383>
- OAuth Working Group. 2012. RFC 6749: OAuth 2.0 Framework. <http://tools.ietf.org/html/rfc6749>.
- Luca Padovani. 2014. Deadlock and lock freedom in the linear π -calculus. In *CSL-LICS*. <https://doi.org/10.1145/2603088.2603116>
- Luca Padovani. 2016. Fair Subtyping for Multi-Party Session Types. *Mathematical Structures in Computer Science* 26, 3 (2016). <https://doi.org/10.1017/S096012951400022X>
- Benjamin C. Pierce. 2002. *Types and programming languages*. MIT Press.
- Alceste Scalas, Ornella Dardha, Raymond Hu, and Nobuko Yoshida. 2017a. A Linear Decomposition of Multiparty Sessions for Safe Distributed Programming. In *ECOOP*. <https://doi.org/10.4230/LIPIcs.ECOOP.2017.24>
- Alceste Scalas, Ornella Dardha, Raymond Hu, and Nobuko Yoshida. 2017b. A Linear Decomposition of Multiparty Sessions for Safe Distributed Programming (Artifact). *Dagstuhl Artifacts Series* 3, 1 (2017). <https://doi.org/10.4230/DARTS.3.2.3>
- Alceste Scalas and Nobuko Yoshida. 2018a. *Less is More: Multiparty Session Types Revisited*. Technical Report 6. Imperial College London. <https://www.doc.ic.ac.uk/research/technicalreports/2018/6>
- Alceste Scalas and Nobuko Yoshida. 2018b. Multiparty session types, beyond duality. *Journal of Logical and Algebraic Methods in Programming* 97. <https://doi.org/10.1016/j.jlamp.2018.01.001>
- Bernardo Toninho and Nobuko Yoshida. 2016. Certifying Data in Multiparty Session Types. In *A List of Successes That Can Change the World: Essays Dedicated to Philip Wadler on the Occasion of His 60th Birthday (LNCS)*, Vol. 9600. https://doi.org/10.1007/978-3-319-30936-1_23
- Bernardo Toninho and Nobuko Yoshida. 2017. Certifying data in multiparty session types. *JLAMP* 90 (2017). <https://doi.org/10.1016/j.jlamp.2016.11.005>
- Philip Wadler. 2014. Propositions as sessions. *J. Funct. Program.* 24, 2-3 (2014).
- Nobuko Yoshida, Pierre-Malo Deniérou, Andi Bejleri, and Raymond Hu. 2010. Parameterised Multiparty Session Types. In *FOSSACS*. https://doi.org/10.1007/978-3-642-12032-9_10