Typing Liveness in Multiparty Communicating Systems

Luca Padovani¹, Vasco Thudichum Vasconcelos², and Hugo Torres Vieira²

¹ Dipartimento di Informatica, Università di Torino, Italy

² LaSIGE, Faculdade de Ciências, Universidade de Lisboa, Portugal

Abstract. Session type systems are an effective tool to prove that communicating programs do not go wrong, ensuring that the participants of a session follow the protocols described by the types. In a previous work we introduced a typing discipline for the analysis of progress in binary sessions. In this paper we generalize the approach to multiparty sessions following the conversation type approach, while strengthening progress to liveness. We combine the usual session-like fidelity analysis with the liveness analysis and devise an original treatment of recursive types allowing us to address challenging configurations that are out of the reach of existing approaches.

1 Introduction

The importance of error detection in the early cycles of software development, and the consequent savings arising from it, can never be overemphasized. The problem becomes even more acute when concurrency comes into play, for concurrency faults are notoriously hard to track down. This work focuses on early error detection of concurrent message passing systems, and addresses, apart from the usual communication safety, the static identification of states in which liveness is compromised.

The setting in which we operate is that of (multi-party) sessions [2–4, 8, 10–12]. *Sessions* are private conversations occurring between two or more interacting participants. Each participant behaves according to a *session type* that describes the messages that the participant is supposed to send/receive and their relative order. One of the strengths of sessions is that they provide a structuring construct on top of which complex systems can be built in a modular way. The relatively simple typing discipline imposed by session types ensures strong properties such as *liveness*, that is the eventual completion of communication operations. This point in favor of sessions is also, somewhat paradoxically, a weakness: since session types describe only *intra-session* communications, but say nothing on *inter-session* dependencies, it may be the case that a well-typed participant simultaneously involved in two or more sessions finds itself in a deadlocked situation because of mutual dependencies between sessions. We address this problem by identifying potentially dangerous dependencies between sessions, so that liveness is ensured also when communications on several different sessions are interleaved.

To illustrate the basic ingredients of our approach, consider the process

$$(vs)(\text{rec } \mathcal{X}.s?x.\mathcal{X} \mid \text{rec } \mathcal{X}.s?y.\mathcal{X} \mid \text{rec } \mathcal{X}.s!5.s!\text{true}.\mathcal{X}) \quad (1)$$

describing three participants (say A, B, and C, composed in parallel) that interact within the scope of a multiparty session s . The aim of C is to repeatedly send two messages

(here exemplified as the constants 5 and true) respectively to A and B. All participants interact within the same session s . However, the order of the synchronizations cannot be predicted and it may well be the case that a 5 message is received by B and a true message is received by A or, in fact, that one of A or B does not receive any message at all! In order to recover the *linearity* of communications (i.e., at most one possible synchronization per session channel at a given moment) we tag messages with *labels*, following the approach of [4]. In this way, we refine (1) to

$$(vs)(\text{rec } \mathcal{X}.s?1x.\mathcal{X} \mid \text{rec } \mathcal{X}.s?my.\mathcal{X} \mid \text{rec } \mathcal{X}.s!15.s!m\text{true}.\mathcal{X}) \quad (2)$$

so that 1- and m-tagged messages respectively and uniquely identify synchronizations with A and B. We are then able to characterize the overall protocol that takes place on session s with the following type.

$$T_s \triangleq \mu\alpha.\tau 1 \text{int}.\tau m \text{bool}.\alpha$$

The type describes a conversation consisting of an infinite exchange of alternated 1- and m-tagged messages whose payload is described by the int and bool types, respectively. The occurrences of τ in the type denote *synchronizations* that are supposed to occur in a session typed by T_s . To specify the behavior of the participants involved in the conversation, we *split* T_s into “slices” which we distribute among the participants. First of all, we separate the behavior of C from the rest of the system, and obtain

$$T_s = T_C \circ T' \quad \text{where} \quad T_C \triangleq \mu\alpha.!1 \text{int}.!m \text{bool}.\alpha \quad \text{and} \quad T' \triangleq \mu\alpha.?1 \text{int}.?m \text{bool}.\alpha$$

In particular, T_C says that C repeatedly sends alternated 1 and m messages and T' says that the rest of the system should be ready to receive the very same messages, in this order. Then, we further split T' in the behaviors of A and B, thus:

$$T' = T_A \circ T_B \quad \text{where} \quad T_A \triangleq \mu\alpha.?1 \text{int}.\alpha \quad \text{and} \quad T_B \triangleq \mu\alpha.?m \text{bool}.\alpha$$

Note that this splitting is valid *assuming* that the environment in which A and B execute guarantees that the synchronization on each 1 message occurs before the synchronization on each m message. This is indeed guaranteed by the sequential structure of process C. Since T_A , T_B , and T_C match the behaviors of A, B, and C with respect to s we may show that (2) is well typed and consequently that it enjoys *communication safety* (no message with wrong type is ever sent), *session fidelity* (the interactions follow the protocol described by T_s), and *liveness* (each interaction described in T_s eventually occurs).

Of all these properties, liveness is the most delicate one, in the sense that it may easily break up when two or more sessions are interleaved with each other. To illustrate the issue, consider the following refinement of (2)

$$(vs)((vr)(\text{rec } \mathcal{X}.\underline{r?ny.s?1x}.\mathcal{X} \mid \text{rec } \mathcal{X}.\underline{s?my.r!ny}.\mathcal{X}) \mid \text{rec } \mathcal{X}.\underline{s!15.s!m\text{true}}.\mathcal{X}) \quad (3)$$

in which A and B are engaged in another session r , different from s , while C behaves exactly as before. Now B forwards y in a n-tagged message to A, perhaps so that A and B can double-check that they are given consistent information from C. Session s is still well typed according to T_s and session r is well typed according to $T_r \triangleq \mu\alpha.\tau n \text{bool}.\alpha$. Yet, (3) is stuck because A waits for the message from B *before* having received the

message from C, but C sends its message to B only *after* it has successfully delivered the message to A. So, none of the synchronizations in T_s and T_r ever happens, although the structure of the participants in (3) agrees to these types.

One possibility for detecting the problem in (3) stems from the observation that the two sessions s and r are mutually dependent on each other. So, one may devise a static analysis technique that keeps track of inter-session dependencies and flags any system that gives rise to circularities as ill typed. This approach has been pursued, for instance, in [3, 6, 7]. The limit of this approach is that, by considering sessions as atomic units, it is quite coarse grained when it comes to analyzing dependencies. For instance,

$$(vs)((vr)(\text{rec } \mathcal{X} . s?1x.r?ny.\mathcal{X} \mid \text{rec } \mathcal{X} . s?my.r!ny.\mathcal{X})) \mid \text{rec } \mathcal{X} . s!15.s!m.\text{true}.\mathcal{X}) \quad (4)$$

is a simple variation of (3) where A performs the same two inputs, but in the “correct” order. Also in (4) there are actions on session s interleaving with actions on session r and vice versa, so the approach based on session dependencies also flags (4) as ill typed, which is unfortunate because (4), contrarily to (3), enjoys liveness.

The approach we pursue here is based on the idea of tracking the dependencies between *actions* instead of sessions. Towards this aim, we annotate each interaction in a type with an identifier—which we call *event*—and we keep track of the dependencies between events by means of a *strict partial order* $\prec$. To get the flavor of the technique at work, let us apply it to the sessions s and r discussed above. First of all, we annotate the actions in the types of s and r with three events e , f , and g :

$$s : \mu\alpha.e\tau 1.\text{int}.f\tau m.\text{bool}.\alpha \quad r : \mu\alpha.g\tau n.\text{bool}.\alpha$$

Then, we analyze the dependencies between the actions in the participants of (3): it must be $e \prec f$ (read, e precedes f) because C first sends the 1-tagged message, and only then it sends the m-tagged message; it must be $g \prec e$ because A waits for the n-tagged message before waiting for the 1-tagged one; finally, it must be $f \prec g$ by looking at the structure of B. Overall, $\prec$ is not a strict partial order because of the circularity in the relation $g \prec e \prec f \prec g$ between the two sessions s and r , hence (3) is ill typed.

Our approach builds on previous works [16, 22] that use analogous annotations for reasoning on the dependencies between actions. With respect to these works, our contributions are along two major axes. First of all, we show that the techniques can be applied to sessions/conversations with an arbitrary number of participants. Second, we support complex recursive process structures. The latter aspect requires a non-trivial extension of the technique described in [22] because, in order to declare that a system like (4) is well typed, we must be able to distinguish occurrences of the same event that pertain to different iterations of a recursive process.

The next section formally describes our language. Sections 3 and 4 introduce the notion of types, the type system and the main results. Section 5 concludes the paper including a more detailed comparison with related work and hints on future developments. Additional material can be found in the associated technical report [18].

2 Process Model

We consider an infinite set of *names* ranged over by $x, y, \dots$ representing communication channels, an infinite set of *process variables* ranged over by $\mathcal{X}, \dots$, and a set of

$P, Q ::= \mathbf{0}$	(Inaction)	$ \quad x!ly.P$	(Output)
$ \quad P Q$	(Parallel)	$ \quad x?\{l_iy_i.P_i\}_{i \in I}$	(Input Summation)
$ \quad (vx)P$	(Restriction)	$ \quad \text{rec } \mathcal{X}.P$	(Recursion)
$ \quad \mathcal{X}$	(Recursion Variable)		

Fig. 1. Syntax of processes

$$\begin{array}{c}
\frac{k \in I}{x?\{l_iy_i.P_i\}_{i \in I} | x!l_kz.Q \rightarrow P_k\{z/y_k\} | Q} \quad \frac{P \rightarrow Q}{(vx)P \rightarrow (vx)Q} \quad (\text{R-Com, R-New}) \\
\frac{P \rightarrow P'}{P|Q \rightarrow P'|Q} \quad \frac{P \equiv P' \quad P' \rightarrow Q' \quad Q' \equiv Q}{P \rightarrow Q} \quad (\text{R-Par, R-Cong})
\end{array}$$

Fig. 2. Reduction relation

message labels $l, \dots$. Processes, ranged over by $P, Q, \dots$, are the terms defined by the grammar in Fig. 1. The language is that of TyCO [20] that extends the π -calculus [15] by considering labeled communications. The terms $\mathbf{0}$, $P|Q$, and $(vx)P$ respectively denote the inactive process, the parallel composition of P and Q , and the restriction of name x to P . Terms $\text{rec } \mathcal{X}.P$ and $\mathcal{X}$ are used to build recursive processes. The term $x!ly.P$ denotes a process that sends a message on channel x and then continues as P . A *message* is made of a label l and an argument y . The term $x?\{l_iy_i.P_i\}_{i \in I}$ denotes a process that waits for a message from channel x and then continues as P_i according to the label of the received message. The argument of the received message replaces the name y_i in P_i . To keep the setting as simple as possible, we have not included conditional or non-deterministic processes. These constructs can be easily added.

The binders of the language are name restriction $(vx)P$, which binds the name x in P , the input prefix $x!ly.P$, which binds the name y in P , and the recursion $\text{rec } \mathcal{X}.P$, which binds the recursion variable $\mathcal{X}$ in P . The notions of free and bound names (as well as free and bound process variables) are defined in the usual way. We identify processes modulo renaming of bound names and of bound process variables. By convention, we exclude recursive processes where unguarded recursion variables occur.

The semantics of the language is defined via a structural congruence and a reduction relation. Structural congruence is standard, except that it includes the law $\text{rec } \mathcal{X}.P \equiv P\{\text{rec } \mathcal{X}.P/\mathcal{X}\}$ for unfolding recursive processes, where $P\{Q/\mathcal{X}\}$ denotes the capture-avoiding substitution of the free occurrences of $\mathcal{X}$ by process Q in P . Reduction is defined by the rules in Fig. 2. Rule (R-Com) describes the synchronization of two processes exchanging a message: the sender emits a message with a label l_k that is among those accepted by the receiver and the argument z of the message replaces the bound input parameter in the appropriate continuation P_k of the receiver. The remaining rules close the relation under language contexts—name restriction and parallel composition—as well as under structural congruence.

3 Types and Typing Contexts

This section starts by introducing the notion of strict partial orders which allows to identify well-formed communication dependencies in processes. It then introduces types and operations on these, most notably type split which allows to separate a type in two disjoint “slices” of behavior.

$p ::= ! \mid ? \mid \tau$	(Polarity)	$B ::= \text{end}$	(Stop)
$T ::= a^n p l T$	(Shared type, S)	$ B_1 B_2$	(Parallel)
$ B^\prec$	(Linear type, L)	$ \mu \alpha. B$	(Recursion)
$\Gamma ::= \cdot \mid \Gamma, x: T$	(Context)	$ \alpha$	(Variable)
$\Delta ::= \cdot \mid \Delta, X: (\Gamma; \prec)$	(Recursion context)	$ a^n p \{l_i T_i. B_i\}_{i \in I}$	(Prefix summation)

Fig. 3. Syntax of types and typing contexts

Strict Partial Orders. We consider an infinite set of event identifiers $\mathcal{E}$ and the set of natural numbers $\mathbb{N}$, and use $a, b, \dots$ to range over $\mathcal{E}$ and $n, m, \dots$ to range over $\mathbb{N}$. We use a^n to denote an element in set $\mathcal{E} \times \mathbb{N}$. We further introduce a distinguished event, $\top$, use $e, f, \dots$ to range over $(\mathcal{E} \times \mathbb{N}) \cup \{\top\}$, and call this set the *set of events*. A *strict (or irreflexive) partial order* $\prec$ over the set of events is a binary relation that is asymmetric (hence irreflexive) and transitive. We write $e \prec f$ when the pair (e, f) is in $\prec$, and $\text{supp}(\prec)$ for the *support* of $\prec$, namely the set of events that occur in $\prec$.

Next we define two *partial* operations over strict partial orders. We write $e + \prec$ for the strict partial order obtained by *adding a least event* e to $\prec$, provided that e does not occur in $\text{supp}(\prec)$. Formally, $e + \prec \triangleq \prec \cup \{(e, f) \mid f \in \text{supp}(\prec)\} \cup \{(e, \top)\}$, where we explicitly add the pair $(e, \top)$ since $\prec$ may be empty (in which case $e + \emptyset$ is defined as $\{(e, \top)\}$). We write $\prec_1 \sqcup \prec_2$ for the least strict partial order that includes both $\prec_1$ and $\prec_2$, if it exists. We use $\sqcup$ to gather the communication dependency structures of, e.g., two parallel processes.

Types. The syntax of types is given in Fig. 3. Our types are based on conversation types [4] extended with event annotations following the approach introduced in [22]. A polarity p describes a communication capability: $!$ specifies an output; $?$ specifies an input; and τ specifies a synchronization, i.e., a matched communication pair (cf., [4]). At the type level we distinguish two separate categories of channels: *shared* (or unrestricted) channels—ranged over by S —are used for modeling (possibly persistent) services having a publicly known name, with which sessions can be established; *linear* channels—ranged over by L —are used for modeling the private conversations within sessions. Note that such distinction between shared and linear channels appears at the type level only, while they are treated uniformly in the process model. For the sake of simplicity we omit non-channel types (e.g., Int) which could be easily added.

A type $a^n p l T$ describes the behavior of a shared channel via an event a^n , a polarity p , a message label l and a type T describing the message argument. We associate shared types with events to temporally relate shared communications with others, in particular with the communications specified by the message type T .

A type $B^\prec$ captures the linear usage of a channel: B specifies the *behavior* of a process w.r.t. the channel, whereas $\prec$ specifies the ordering of events expected from the external environment. Informally, $\prec$ is used in a type $B^\prec$ to represent the sequentiality information that B admits but does not impose. For example, when typing a process that concurrently sends messages `hello` and `bye` the type *may* specify that the outputs on `hello` and `bye` actually take place one after the other if such order is imposed by the corresponding inputs (present in the external process environment).

Behavioral types B include inaction `end`, parallel composition $B_1 | B_2$ of two independent behaviors B_1 and B_2 , recursive types $\mu \alpha. B$, recursion variables α , and (prefixed)

summation $a^n p\{l_i T_i.B_i\}_{i \in I}$. Sums capture communication capabilities associated with event a^n , polarity p , and a menu of synchronization options. Each entry in the menu is identified by a distinct label l_i , the type of the argument of the message T_i , and the behavior B_i that takes place after the synchronization. We say that a linear type $B^\prec$ is well-formed if $\text{supp}(\prec)$ does not include events associated with communication actions of polarity τ in B (since no further ordering information can be provided for such actions by the external environment). In the remainder, whenever we write $B^\prec$, we assume that $B^\prec$ is well formed. We also identify α -equivalent (recursive) types by convention.

Following the ideas presented in [22], we associate with each linear communication an event a^n so as to temporally relate the communication action described by the summation with respect to others, establishing an overall ordering of communications. In this work, we introduce the notion of *iteration*, by adding to events a natural number n , allowing to describe infinite chains of (related) events. Informally, the index allows to capture the several “stages” of a type by means of an *increment*, so, for example, $\mu\alpha.a^1 \tau 1 T.b^1 \tau m T'.\alpha$ unfolds to $a^1 \tau 1 T.b^1 \tau m T'.\mu\alpha.a^2 \tau 1 T.b^2 \tau m T'.\alpha$ so as to associate the first iteration with index 1 and the second iteration with index 2 and so on and so forth.

Operations on Types. We write $\text{labels}(B)$ for the set of labels occurring in B . We say that B_1 and B_2 are *behaviorally independent*, and denote it by $B_1 \# B_2$, if $\text{labels}(B_1) \cap \text{labels}(B_2) = \emptyset$ so that disjoint message sets ensure behavioral independence. We also need an operation to *remove* part of the partial order in a type, defined as $B^{\prec'} \setminus \prec \triangleq B^{\prec' \setminus \prec}$, and $S \setminus \prec \triangleq S$. Since linear types may contain sequentiality assumptions, we use this operation to clear hypotheses that are proved externally.

In order to capture the several iterations of a communication that may repeat itself in the context of recursion, we introduce an operator that *increments* the index associated with an event by a given factor, defined as $\text{inc}(a^n, m) \triangleq a^{n+m}$ and $\text{inc}(\top, m) \triangleq \top$. We then extend inc to strict partial orders, pointwise, and to behavior types so that $\text{inc}(a^n p\{l_i T_i.B_i\}_{i \in I}, m) \triangleq a^{n+m} p\{l_i \text{inc}(T_i, m). \text{inc}(B_i, m)\}_{i \in I}$. The increment of a behavior is an homomorphism for all other constructs. The increment operation on types affects only linear types, $\text{inc}(B^\prec, m) \triangleq \text{inc}(B, m)^{\text{inc}(\prec, m)}$, since events associated with shared communications are not considered to be repeated in different stages but rather to be repeated always at the same stage, hence $\text{inc}(a^n p l T, m) \triangleq a^n p l T$. Essentially, we model shared communication repetition using replication (via recursion), so shared replicated communication actions have the same temporal ordering, while we model linear recursive repetition using a sequential chain of events. The index is then used to capture repetition (without cycles) in the orderings.

To simplify the typing rules, we define a type equivalence relation $\equiv$ that includes commutativity, associativity and neutral end for $|$, as well as iso-recursive equivalence for recursive types $\mu\alpha.B \equiv B\{\mu\alpha.\text{inc}(B, m)/\alpha\}$ for some $m > 0$, saying that the next iteration of the behavior is captured by the increment of the events (we use any positive m so as to support misalignment between processes and types).

We now introduce operations that capture the temporal ordering prescribed by types. We write $\text{events}(B)$ for the set of elements of $\mathcal{E} \times \mathbb{N}$ occurring in a behavior B , not including the events in message types. Formally:

$$events(B) \triangleq \begin{cases} \emptyset & \text{if } B = \text{end or } B = \alpha \\ events(B_1) \cup events(B_2) & \text{if } B = B_1 \mid B_2 \\ \{e\} \cup \bigcup_{i \in I} events(B_i) & \text{if } B = ep\{l_i T_i . B_i\}_{i \in I} \\ \{inc(e, k) \mid e \in events(B') \text{ and } k \geq 0\} & \text{if } B = \mu\alpha . B' \end{cases}$$

Notice that $events(\mu\alpha.B)$ includes all the events in the body of the recursion, incremented zero or more times so as to capture the first and the following iterations. We extend the operation to types, by defining $events(B^\prec) \triangleq events(B)$, as we are only interested in linear types where $supp(\prec) \subseteq events(B)$, and by defining $events(ep l T) \triangleq \{e\}$.

We write $B\downarrow$ for the strict partial order over $\mathcal{E} \times \mathbb{N}$ induced by a type B . Notice that $B\downarrow$ is a partial operator since it uses $\cup$ and $+$. Formally:

$$B\downarrow \triangleq \begin{cases} \emptyset & \text{if } B = \text{end or } B = \alpha \\ B_1\downarrow \cup B_2\downarrow & \text{if } B = B_1 \mid B_2 \\ e + (\cup_{i \in I} B_i\downarrow) & \text{if } B = ep\{l_i T_i . B_i\}_{i \in I} \\ \{(inc(e, k), inc(f, k)) \mid (e, f) \in B'\downarrow \text{ and } k \geq 0\} \cup \\ \{(inc(e, m), inc(f, n)) \mid e, f \in events(B') \text{ and } 0 \leq m < n\} & \text{if } B = \mu\alpha . B' \end{cases}$$

Notice that the operation adds a least event in the case of the prefix summation, and for recursions it adds all pairs obtained from the body of the recursion (incremented zero or more times) and all pairs that pertain to different iterations. We extend the definition to types by taking $(ep l T)\downarrow \triangleq (e, \top)$ and $B^\prec\downarrow \triangleq B\downarrow \setminus \prec$, where $\setminus$ denotes set difference. The definition for linear types considers the order obtained from the behavioral type removing the ordering expected from the external environment, so $B^\prec\downarrow$ characterizes exclusively the ordering imposed by the type.

To identify types that characterize channels that do not depend on the external environment to evolve, and hence are “self-sustained” communication wise, we introduce a predicate that is true for types containing no unmatched communication actions. We say that a behavioral type B is *matched* if it contains no top-level (i.e., excluding message types) input or output polarities. We extend the definition to linear types by considering $matched(B^\prec) \triangleq matched(B)$ (which, by well-formedness, implies $\prec = \emptyset$), and $matched(ep l T) \triangleq p = ?$. A message type of polarity $?$ says that a shared input is available. Since we are only interested in capturing continuously available shared inputs, $?$ shared types “absorb” (as will be clear from the definition of type splitting) $!$ shared types, so as to capture the fact that (replicated) shared inputs are still available after synchronization. Hence, the definition of $matched()$ for shared types excludes solely (unmatched) shared outputs, and considers the (infinitely available) shared inputs to be matched (regardless whether they are used or not).

Typing Contexts. The syntax of typing contexts is given in Fig. 3. We assume by convention that, in a typing context $\Gamma, x : T$ and in a recursion environment $\Delta, \mathcal{X} : (\Gamma; \prec)$, the name x and the process variable $\mathcal{X}$ do not occur in Γ and in Δ , respectively, as usual. Also, we consider contexts up to permutations of their entries.

We denote by Γ_{un} contexts that contain only outputs on shared channels and linear types with end behavior, that is, if $x : T$ is in Γ_{un} , then T is either $e ! T'$ or $\text{end}^\emptyset$. We use such contexts to describe systems that only use shared resources, namely to describe (the continuation of) processes that input on shared channels. We exclude shared inputs

$$\begin{array}{c}
\frac{}{B = B \circ \text{end}} \quad \frac{B_1 = B'_1 \circ B''_1 \quad B_2 = B'_2 \circ B''_2 \quad B_1 \# B_2}{B_1 | B_2 = B'_1 | B'_2 \circ B''_1 | B''_2} \quad (\text{B-End, B-Par}) \\
\frac{B = B_1 \circ B_2}{\mu \alpha. B = \mu \alpha. B_1 \circ \mu \alpha. B_2} \quad \frac{\forall i \in I \quad B_i = B'_i \circ B \quad e p\{l_i T_i. \text{end}\}_{i \in I} \# B}{e p\{l_i T_i. B_i\}_{i \in I} = e p\{l_i T_i. B'_i\}_{i \in I} \circ B} \quad (\text{B-Rec, B-Break}) \\
\frac{}{\alpha = \alpha \circ \alpha} \quad \frac{\forall i \in I \quad B_i = B'_i \circ B''_i}{e \tau\{l_i T_i. B_i\}_{i \in I} = e ?\{l_i T_i. B'_i\}_{i \in I} \circ e !\{l_i T_i. B''_i\}_{i \in I}} \quad (\text{B-Var, B-Sync})
\end{array}$$

Fig. 4. Behavioural type splitting

$$\frac{B = B_1 \circ B_2 \quad \prec = \prec_1 \setminus B_2 \prec_2 \downarrow \cup \prec_2 \setminus B_1 \prec_1 \downarrow \cup B \downarrow \setminus (B_1 \downarrow \cup B_2 \downarrow)}{B \prec = B_1 \prec_1 \circ B_2 \prec_2} \quad (\text{L-Split})$$

Fig. 5. Linear type splitting

from Γ_{un} in order to avoid “nested” shared inputs, so that inputs on shared channels are continuously active (cf. uniform receptiveness [1, 19]). Similarly, we denote by Γ_{lin} contexts that contain only linear types, that is, types of the form $B^\prec$.

We are interested in systems where all communications are matched, i.e., typed against *matched* contexts, defined as the pointwise extension of the *matched* predicate on types. We also lift the notions of type *increment*, *inc*, type *equivalence*, $\equiv$, and partial order difference, $\setminus$, pointwise to contexts.

Splitting and Conformance. We now introduce two notions crucial to our development, namely *splitting* (inspired by [2] and by the *merge* operation of [4]) that explains how behaviors can be decomposed and safely distributed to distinct parts of a process (e.g., to the branches of a parallel composition), and *conformance* that captures the desired relation between typing contexts and strict partial orders.

We say type T conforms to order $\prec$, noted $\text{conforms}(T, \prec)$, if $T \downarrow \subseteq \prec$. Notice that since $T \downarrow$ excludes the ordering expected from the external environment, conformance focuses on the order imposed by the types (which is the focus of the overall ordering). The *conforms* predicate is defined on typing contexts as the pointwise extension of the predicate on types, so $\text{conforms}(\Gamma, \prec)$ ensures that every communication action specified in Γ is ordered by $\prec$.

Splitting is defined both on types and on typing contexts. We write $T = T_1 \circ T_2$ to mean that type T is split in types T_1 and T_2 , and likewise for $\Gamma = \Gamma_1 \circ \Gamma_2$. Behavioral type splitting, linear type splitting, shared type splitting and context splitting are given by the rules in Figs. 4–7 (where we omit symmetric rules).

We briefly describe the rules in Fig 4. A behavioral type may be split in itself and in end, so as to allow, e.g., to give away the behavior completely to one branch of a parallel composition—rule (B-End). A parallel composition $B_1 | B_2$ (where B_1 and B_2 are apart $\#$) may be split in two parallel compositions, the components of which are obtained by decomposing B_1 and B_2 —rule (B-Par). A recursive type is split in two recursive types, the bodies of which are obtained by splitting the body of the incoming recursive type—rule (B-Rec). Also, a recursion variable may be split in itself—rule (B-Var).

A prefix summation may be split in an independent ($\#$) behavior, obtained by splitting (all) the continuations, and in the prefix whose continuations specify the remaining

$$\frac{p \in \{?, !\}}{e?lT = e?lT \circ eplT} \text{(S-In-L)}$$

$$\frac{}{e!lT = e!lT \circ e!lT} \text{(S-Out)}$$

$$\frac{}{. = . \circ .} \text{(C-Empty)} \quad \frac{\Gamma = \Gamma_1 \circ \Gamma_2}{\Gamma, x: T = \Gamma_1, x: T \circ \Gamma_2} \text{(C-Left)}$$

$$\frac{\Gamma = \Gamma_1 \circ \Gamma_2 \quad T = T_1 \circ T_2}{\Gamma, x: T = \Gamma_1, x: T_1 \circ \Gamma_2, x: T_2} \text{(C-Split)}$$

Fig. 6. Shared type splitting

Fig. 7. Context splitting

behavior—rule (B-Break). A synchronized (τ) prefix summation may be split in prefix with dual polarities ($?$ and $!$) whose continuations are obtained by splitting the synchronized prefix continuations—rule (B-Sync).

Notice that rule (B-Break) may decompose a type in such a way that the overall ordering is not guaranteed by the splitted types. To this end we keep track of the ordering assumptions in the linear type splitting, defined in Fig. 5. A linear type split is defined by the behavioral split and also by a separation of the ordering assumptions $\prec$, such that everything $\prec$ assumes may be assumed by $\prec_1$ or $\prec_2$, but $\prec_1$ and $\prec_2$ may specify other assumptions which are actually ensured by $B_2 \prec_2$ and by $B_1 \prec_1$, respectively. Also, $\prec$ necessarily contains the ordering present in B (i.e., $B \downarrow$) that is not supported by either B_1 or B_2 , hence any sequentiality information that B specifies introduced via (B-Break).

Shared type splitting (Fig. 6) decomposes shared communication capabilities in two distinct ways, depending on whether the polarity of the incoming type is $?$ or $!$. A shared input is split in a shared input and either in an output or another input, via rule (S-In-L). Essentially, the latter allows for typing processes that separately offer the input capability (e.g., a service that is provided by two distinct sites), and the former allows for typing processes that offer the dual communication capabilities (e.g., a service provider and a service client). A shared output is split in two shared outputs—rule (S-Out)—which allows for typing processes that offer the output capability separately (e.g., two clients of some service). Notice type splitting preserves the message types and event association so as to guarantee the dual communication actions agree on the type of what is communicated and on the ordering.

Context splitting (Fig. 7) allows to divide a context in two distinct ways: context entries either go into the left or the right outgoing contexts—(C-Left) as well as the omitted symmetric rule—or they go in both contexts—(C-Split). The latter form lifts the (type) behavior distribution to the context level, while the former allows to delegate the entire behavior to a part of the process, leaving no usage at all to the other part. To lighten notation we use $\Gamma_1 \circ \Gamma_2$ to represent any Γ such that $\Gamma = \Gamma_1 \circ \Gamma_2$ (if such Γ exists). Notice that, given Γ_1 and Γ_2 , there may be more than one Γ such that $\Gamma = \Gamma_1 \circ \Gamma_2$.

4 Typing System

This section introduces our typing system and the main results of the paper, namely soundness of the type system (Theorem 1) and liveness (Theorem 2).

Typing System. The typing system characterizes processes according to typing assumptions for free process variables (Δ) and for names (Γ), as well as an overall ordering of events ($\prec$). We say process P is well-typed if $\Delta; \Gamma; \prec \vdash P$ is derivable using the

$$\begin{array}{c}
\frac{\Delta; \Gamma_1; \prec_1 \vdash P \quad \Delta; \Gamma_2; \prec_2 \vdash Q}{\Delta; \Gamma_1 \circ \Gamma_2; \prec_1 \sqcup \prec_2 \vdash P \mid Q} \quad \frac{\Delta; \Gamma, x: T; \prec \vdash P \quad \text{matched}(T)}{\Delta; \Gamma; \prec \vdash (\nu x)P} \quad \text{(T-Par, T-New)} \\
\\
\frac{}{\Delta; \Gamma_{un}; \emptyset \vdash \mathbf{0}} \quad \frac{\Delta, \mathcal{X} : (\text{inc}(\Gamma_{lin}, n); \text{inc}(\prec, n)); \Gamma_{lin}; \prec \vdash P \quad n \in \mathbb{N}}{\Delta; \Gamma_{lin}; \prec \vdash \text{rec } \mathcal{X}.P} \quad \text{(T-Inact, T-LRec)} \\
\\
\frac{\Delta, \mathcal{X} : (\Gamma; \prec); \Gamma; \prec \vdash P \quad \Gamma = \Gamma_{un}, x: e?lT \quad \prec = (e + \prec'') \sqcup \prec' \quad e \notin \text{supp}(\prec')}{\Delta; \Gamma; \prec \vdash \text{rec } \mathcal{X}.P} \quad \text{(T-URec)} \\
\\
\frac{\Delta(\mathcal{X}) = (\Gamma; \prec) \quad \text{conforms}(\Gamma, \prec)}{\Delta; \Gamma; \prec \vdash \mathcal{X}} \quad \frac{\Delta, \Gamma_2; \prec \vdash P \quad \Gamma_1 \equiv \Gamma_2}{\Delta; \Gamma_1; \prec \vdash P} \quad \text{(T-Var, T-Equiv)} \\
\\
\frac{\forall i \in I \Delta; \Gamma, x: B_i \prec'_i, y_i: T_i; \prec_i \vdash P_i}{\Delta; \Gamma, x: e?\{l_i T_i. B_i\}_{i \in I} \sqcup_{i \in I} \prec'_i; e + (\sqcup_{i \in I} \prec_i) \vdash x?\{l_i y_i. P_i\}_{i \in I}} \quad \text{(T-LinIn)} \\
\\
\frac{\Delta; \Gamma, x: B_k \prec'_k; \prec \vdash P \quad k \in I}{\Delta; (\Gamma, x: e!\{l_i T_i. B_i\}_{i \in I} \prec'^k) \circ y: T_k; e + (\prec \sqcup T_k \downarrow) \vdash x!l_k y.P} \quad \text{(T-LinOut)} \\
\\
\frac{\Delta; \Gamma_{un}, x: e?lT, y: T; e + \prec \vdash P}{\Delta; \Gamma_{un}, x: e?lT; e + \prec \vdash x?l y.P} \quad \frac{\Gamma = (\Gamma_{un}, x: e!lT) \circ y: T \quad \cdot; \cdot; \emptyset \vdash P}{\Delta; \Gamma; e + T \downarrow \vdash x!l y.P} \quad \text{(T-UIn, T-UOut)}
\end{array}$$

Fig. 8. Typing rules

rules in Fig. 8. We briefly comment on the rule. In rule (T-Par) the parallel composition is typed if the branches are typed in splittings ($\circ$) of the context and a decomposition of the order. In rule (T-New) the name restriction is typed if the process in the scope of the restriction is typed in the same contexts together with the typing assumption for the usage of the restricted name which must be a *matched* type (all communication prefixes are matched). Notice that the overall ordering $\prec$ is preserved, hence the ordering prescribed by name x is still present in the conclusion, even if the type T of x is not.

In rule (T-Inact) the inaction process is typed with any usage of recursion variables (Δ), and with only outputs on shared labels and end linear types (Γ_{un}), and an empty overall ordering ($\emptyset$). In rule (T-LRec) a recursive process is well typed if so is the body of the recursion in the same typing context Γ_{lin} (which only includes linear usages) and overall ordering $\prec$, and in the recursion environment augmented with an assumption for the recursion variable: the variable is assumed to have exactly the same usage and overall ordering *up to* an increment (for some n) of the natural exponent of the events. Rule (T-LRec) therefore captures, in a fairly intuitive way, subsequent iterations of a (linear) recursion: the point of the next iteration is characterized by an increment of the typing and ordering.

Rule (T-URec) addresses a recursive process that uses only shared resources where no increment is involved since shared communications do not have iterations (their repetition is considered to happen at the level of a single iteration). So the recursion environment is augmented with the typing and ordering that types the body of the recursion. The typing mentions only shared exponential resources (Γ_{un}) together with a shared input (on x), as we intend to capture replicated shared inputs. In order to ensure that the shared input is an immediate action of the body of the recursion, the ordering makes e a minimal event. Given the above explanation, rule (T-Var) is straightforward:

the assumption for the variable provides the context and ordering for the process. In rule (T-Equiv) we embed the notion of context equivalence in the type system, since we need to unfold recursive types when typing the body of a recursion.

Communication prefixes are also typed in separate rules, depending on the type of the subject of the communication. In rule (T-LinIn) the input on a channel x with linear usage is typed if the continuation processes are typed with the usages for x prescribed in the prefix summation type, together with a separation of the ordering assumptions; also, by adding a typing assumption for the usage of the received name (according to the corresponding message type), and a separation of the events *greater than e* in the overall ordering. The e is the event associated with the prefix summation (notice that we pick *fresh* events since $+$ is undefined otherwise). The fact that events in the continuation are of greater order ensures that the communications in the continuation are in fact prescribed to take place after the prefix itself. Notice that the overall ordering registered in the conclusion is a tree rooted in e . Further notice that the communication dependency structure of the received name is transparently kept in the conclusion (the ordering prescribed by the channel usages is *invariantly* registered in the overall ordering). This allows us to type systems where communications on received channels are interleaved with others, configurations out of reach of related approaches.

The reasoning is similar in rule (T-LinOut). The continuation is typed by considering the continuations of the prefix summation (any prefix summation containing the only label mentioned by the process) which is uniquely associated with event e , together with the same ordering assumptions $\prec'$ (as we are only interested that the environment guarantees the order of one branch). The typing context Γ is actually the result of a split of the context registered in the conclusion, which also mentions the usage delegated in the communication for the sent name. Finally, the overall ordering in the conclusion also registers the ordering (of events greater than e) prescribed by the message type.

Rules (T-UIn) and (T-UOut) explain the typing for communications on shared channels. In rule (T-UIn) the input with shared usage is typed if the continuation process is typed adding the usage for the received name to the context. Notice that since we type the continuation with the shared input usage, the continuation must offer again the shared input behavior (so shared inputs can be typed only in the context of a recursion). Notice also that the overall ordering in the conclusion is that of the premise, as expected in a replicated process, and specifies that the event associated with the shared input is minimal (so as to ensure it is immediately available). Furthermore, we require that the remaining context mentions exclusively shared outputs (Γ_{un}) so that no other shared inputs are defined in the continuation. This would be a problem for liveness since shared inputs on *free* names defined in the continuation might leave a matching output dangling. However, we may freely type processes in the continuation that specify shared inputs in restricted names (or even in the received name).

In rule (T-UOut) we type the output on a shared channel if the continuation is typed in the empty context and empty ordering. This means that our model for shared channel communications is an asynchronous one. There are at least two approaches to guarantee that shared inputs are always active (*uniform receptiveness*): one is to exclude usage of the shared name in the continuations of *both* input and output prefixes [19] (we followed a similar approach in [22] excluding the corresponding *event* in the continuations); the

other relies on an asynchronous model of communication [1]—which we adopt here for shared channels. The advantage of this approach is that it supports processes that specify in the continuation of a shared input a matching output (intuitively, think of a recursive “service” call). Also, looking at (T-UOut) and (T-Par) we argue that every process P that we have used in examples in the continuation of a shared output $x!l.y.P$ can be specified (and typed) using the parallel composition $P|x!l.y.\mathbf{0}$, essentially since the type delegated in the communication is obtained via a split of the context nonetheless. Notice that rule (T-UOut) says that the event associated with the output is minimal w.r.t. the message type in the conclusion. Notice also that the rules for communication prefixes make no distinction whatsoever on the type of the channel communicated.

One can now show that process (4) is well typed. Consider the following types.

$$T_s \triangleq \mu\alpha.e^1 \tau 1 \text{int}.f^1 \tau \text{mbool}.\alpha \quad T_r \triangleq \mu\alpha.g^1 \tau \text{nbool}.\alpha$$

Each unfolding of a type increments the indexes of the events. The splitting of these behaviors produces

$$T_{As} \triangleq \mu\alpha.e^1 ? 1 \text{int}.\alpha \quad T_{Bs} \triangleq \mu\alpha.f^1 ? \text{mbool}.\alpha \quad T_{Cs} \triangleq \mu\alpha.e^1 ! 1 \text{int}.f^1 ! \text{mbool}.\alpha$$

$$T_{Ar} \triangleq \mu\alpha.g^1 ? \text{nbool}.\alpha \quad T_{Br} \triangleq \mu\alpha.f^1 ! \text{nbool}.\alpha$$

regarding sessions s and r . Now, looking at the structure of the participants in (4), we realize that the following relations must hold: the structure of A requires $e^1 \prec g^1 \prec e^2$; the structure of B requires $f^1 \prec g^1 \prec f^2$; finally, the structure of C requires $e^1 \prec f^1 \prec e^2$. Overall, it is possible to find a typing derivation for the whole process by considering the strict partial order

$$\prec \triangleq T_s \downarrow \cup T_r \downarrow \cup \{(f^i, g^i), (g^i, e^{i+1}) \mid i \in \mathbb{N}\}$$

Results. We start by mentioning some auxiliary results, in particular that conformance between the typing context and the overall ordering is ensured for all derivations. This result may be viewed as a sanity check saying that the conditions imposed by our rules are enough to keep conformance invariant in a derivation. We may also show that split is an associative relation, in particular for behavioral types. This result in particular ensures that the derivation (sub-)trees may be moved around, and used in the proof of the following (standard) results.

Lemma 1 (Subject Congruence). *If $\Delta; \Gamma; \prec \vdash P$ and $P \equiv Q$ then $\Delta; \Gamma; \prec \vdash Q$.*

Lemma 2 (Substitution). *If $\Delta; \Gamma_1, x: T; \prec \vdash P$ and $\Gamma_2 = \Gamma_1 \circ y: T$ then $\Delta; \Gamma_2; \prec \vdash P\{x/y\}$.*

The proofs follow by induction on the structure of the process and on the length of the typing derivation (respectively) along unsurprising lines. Notice that substitution uses context splitting to characterize the context that types the resulting process, since name y may already be used by P and the soundness of the substitution is guaranteed by the split. Before presenting our first main result we need to introduce two auxiliary notions that characterize reduction of contexts and of strict partial orders. As expected from a behavioral type system, as processes evolve so must the types that characterize the processes. The reduction relations for behavioral types and contexts are given in

$$\begin{array}{c}
 \frac{k \in I}{e \tau\{l_i T_i.B_i\}_{i \in I} \rightarrow B_k} \quad \frac{B_1 \rightarrow B'_1}{B_1 | B_2 \rightarrow B'_1 | B_2} \quad \frac{B_1 \equiv B'_1 \rightarrow B'_2 \equiv B_2}{B_1 \rightarrow B_2} \\
 \frac{}{\cdot \rightarrow \cdot} \quad \frac{B_1 \rightarrow B_2}{\Gamma, x: B_1 \multimap \rightarrow \Gamma, x: B_2 \multimap} \quad \frac{\Gamma_1 \rightarrow \Gamma_2}{\Gamma_1, x: T \rightarrow \Gamma_2, x: T}
 \end{array}$$

Fig. 9. Type and context reduction

$$\frac{}{\prec \rightarrow \prec} \quad \frac{e \in \text{supp}(\prec)}{\prec \rightarrow \prec \setminus e} \quad \frac{\Gamma_1 \rightarrow \Gamma_2 \quad \prec_1 \rightarrow \prec_2}{\Gamma_1; \prec_1 \rightarrow \Gamma_2; \prec_2}$$

Fig. 10. Order and typing reduction

Fig. 9. Note that τ -prefixed summations (in “active contexts”) may reduce and a context reduces if it has an entry on a linear type prefix that reduces. Also, the empty context reduces so as to mimic synchronizations on restricted and shared channels (embedding reflexivity in context reduction); these synchronizations do not change the types.

Fig. 10 shows the reduction for orders and context/order pairs. Strict partial order reduction is also reflexive to capture both shared synchronizations and communications that *depend* on shared communications (as they take place repeatedly for each of the continuation of the shared input). Reduction is also defined by removing an event of the ordering, so as to capture *one shot* synchronizations (which includes infinite chains of synchronizations). We may now present our first main result.

Theorem 1 (Preservation). *If $\Delta; \Gamma_1; \prec_1 \vdash P_1$ and $P_1 \rightarrow P_2$ then $\Gamma_1; \prec_1 \rightarrow \Gamma_2; \prec_2$ and $\Delta; \Gamma_2; \prec_2 \vdash P_2$.*

The proof follows by induction on the length of the derivation of $P_1 \rightarrow P_2$. The theorem says that typing is preserved under process reduction, up to a reduction in the context and ordering. Fidelity is an immediate consequence of Theorem 1, as usual (cf. [2]), thanks to the precise correspondence between reduction in processes and in typing contexts. We now turn our attention to the liveness result, where we use $\rightarrow^n$ to denote a sequence of n reductions.

Theorem 2 (Liveness). *Let $\Delta; \Gamma_1; \prec_1 \vdash P_1$ with $\text{matched}(\Gamma_1)$, and let $x: L_1$ in Γ_1 with $e \in \text{events}(L_1)$. Then $P_1 \rightarrow^n P_2$ and $(\Gamma_1; \prec_1) \rightarrow^n (\Gamma_2; \prec_2)$ and $\Delta; \Gamma_2; \prec_2 \vdash P_2$ with $x: L_2$ in Γ_2 and $e \notin \text{events}(L_2)$, for some $n > 0$.*

In words, every event e occurring in the type of a linear channel used by a well-typed process can eventually disappear from the type environment. This means that either e is associated with an (inter)action that can eventually be performed by the process, or that e occurs in a branch of a choice which is not selected. This property is akin to *lock freedom* [13] or *progress* [3, 6, 12] except that e in Theorem 2 can be associated with an action that is arbitrarily deep within the process structure, whereas lock freedom and progress are usually formulated for top-level actions only. The proof invariant is that for each linear synchronization prescribed by the types there is either an immediate corresponding synchronization in the process or there are preceding actions which necessarily are of “lesser” order. The fact that behaviors described by linear types

have a correspondence with the communication capabilities of processes is a standard property of linear type theories.

Notice that we are not able to characterize shared usages in the same way, as the events associated with them are persistent. However, we may immediately conclude that since any linear synchronization that depends on a shared synchronization takes place then so does the shared synchronization (in fact, our proof relies on the fact that also shared synchronizations are live, along with communications in restricted channels with matched typings). Notice also that our type-based approach addresses processes with “unmatched” typing, just as long as we consider them up to the composition with any other processes for which the resulting typing is *matched*—in particular via rule (T-Par) of Fig. 8. An immediate consequence of Theorem 1 and Theorem 2 combined is that any configuration reachable from a matched typed one (as the matched predicate is invariant under context reduction) also has the liveness property.

5 Concluding Remarks

We have presented a type system for multiparty session-based communication-centred systems that guarantees *liveness* in addition to session *fidelity* even when multiple sessions are interleaved. Compared to other models for multiparty session communication, our approach strives to achieve minimality of both language and type features. Regarding language features, we rely on message labels for preventing communication races on linear channels, whereas other approaches make use of *channel polarities* [9], of distinct *channel endpoints* [21], or *roles* [3, 6]. Moreover, we do not make use of dedicated session initialization primitives. Regarding type features, our work exploits notions introduced in [2, 4] (e.g., the τ polarity and the split operator), allowing us to use the same type language for specifying both *global* and *local* types. This is in contrast with common multiparty session type theories such as [3, 6, 12], which introduce distinct languages for global and local types connected by a projection operation from the former into the latter.

A number of type-based techniques guaranteeing deadlock freedom, progress, or liveness properties have been proposed. Kobayashi [13, 14] presents type systems for *lock-free* and *deadlock-free* processes written in the pure π -calculus. Roughly speaking, every top-level input/output prefix in a lock-free process is guaranteed to be eventually consumed, whereas a deadlock-free process is one that is always able to reduce, unless it has terminated. The type systems rely on *channel usages*, which are behavioral types resembling session types where actions are annotated with pairs of *obligation/capability* levels, roughly denoting the time at which actions begin/are supposed to end. Top-level actions with a finite capability level are guaranteed to succeed in a finite amount of time (and possibly under some fairness assumption). For session-based languages, the relevant works on binary sessions are [8, 16], while [3, 6] deal with multiparty sessions. The basic idea of [3, 6, 8] is to devise a type system that detects the *dependency graph* between different sessions, where a dependency arises if a (blocking) action in one session guards an action pertaining a different session. Liveness is guaranteed if the dependency graph is acyclic. [16] leverages Kobayashi’s technique (in [13]) from channel usages to session types showing that such technique can achieve a greater accuracy when compared to [3, 6, 8]. The present work differs from these in several minor and major ways.

In particular, our process model is synchronous, while the ones in [3, 6, 16] is asynchronous. Asynchrony has a non-trivial impact in the type system for progress, mainly because output actions are *non-blocking*. The progress property considered in [3, 6] assumes that missing session participants can eventually join the system at any time. In practice, this assumption implies that any action on shared channels is considered non-blocking, because it is always possible to add some (well-typed) processes that provide for the missing messages. Also, [6] defines a syntax-directed type system and automatic inferences are known for the systems described in [13, 14]. In our case, the definition of a syntax-directed type system and of an inference algorithm remain open problems.

One major difference between our work and the aforementioned ones, which constitutes the main technical contribution, regards the treatment of recursive types. In all previous works, annotations such as obligation/capability levels in [13, 14], dependency graphs [3, 6], timestamps [16] are statically associated with types, regardless of their recursive structure. In our case, unfolding a recursive type has the effect to “freshen” the events occurring therein. This significantly increases the range of well-typed processes. In particular, none of the aforementioned works is able to prove progress for non-trivial recursive processes interleaving (blocking) actions on different channels. For example, the (appropriate encoding of the) (4) is ill typed according to all previous type systems. More recently, the first author has studied a type system for deadlock and lock freedom which is capable of addressing non-trivial recursive process configurations, albeit in the context of the linear π -calculus [17]. The type system in [17] can prove that a configuration such as (4) is (dead)lock free, but only encoding the multiparty session s in terms of several binary sessions which, in turn, can be encoded using linear channels. In the present work instead we consider a calculus with a primitive notion of multiparty session, addressing scenarios that cannot be compiled down to binary sessions.

Naturally a type-based approach is only relevant if it can be taken into practice, so decidability is a fundamental property. We may argue that we can extract a decidable type-checking procedure from our type system, if we annotate restricted names with their types (as usual) and process recursion variables with the increment factor (together with confining unfoldings to a “just-in-time” setting). Inference is also an important issue as it allows to save the programmer’s effort to specify the types and increase the probability that such advanced type system can actually be used in practice. Although we believe these are very important questions to address, we decided to leave them to future clarification and focus on the principles of our approach for now, so as to make further efforts worthwhile. Furthermore, observing that types are becoming very rich characterizations of process behavior (in our case how and when channels are used), one may ask if it is possible to deduce processes from types (e.g., [5]) and spare the “programmer” the effort of writing programs and just ask him to write the types.

Acknowledgments. This work was supported by MIUR PRIN CINA 2010LHT4KM, FCT through project Liveness, PTDC/ EIA-CCO/117513/2010, and LaSIGE Strategic Project, PEst-OE/EEI/UI0408/2014. We are grateful to the COORDINATION’14 reviewers whose comments helped us clarifying and improving the paper.

References

1. Amadio, R.M., Boudol, G., Lhoussaine, C.: On message deliverability and non-uniform receptivity. *Fundam. Inform.* 53, 105–129 (2002)
2. Baltazar, P., Caires, L., Vasconcelos, V.T., Vieira, H.T.: A type system for flexible role assignment in multiparty communicating systems. In: Palamidessi, C., Ryan, M.D. (eds.) *TGC 2012. LNCS*, vol. 8191, pp. 82–96. Springer, Heidelberg (2013)
3. Bettini, L., Coppo, M., D’Antoni, L., De Luca, M., Dezani-Ciancaglini, M., Yoshida, N.: Global progress in dynamically interleaved multiparty sessions. In: van Breugel, F., Chechik, M. (eds.) *CONCUR 2008. LNCS*, vol. 5201, pp. 418–433. Springer, Heidelberg (2008)
4. Caires, L., Vieira, H.T.: Conversation types. *Theor. Comput. Sci.* 411, 4399–4440 (2010)
5. Carbone, M., Honda, K., Yoshida, N.: Structured communication-centered programming for web services. *ACM Trans. Program. Lang. Syst.* 34, 8 (2012)
6. Coppo, M., Dezani-Ciancaglini, M., Padovani, L., Yoshida, N.: Inference of Global Progress Properties for Dynamically Interleaved Multiparty Sessions. In: De Nicola, R., Julien, C. (eds.) *COORDINATION 2013. LNCS*, vol. 7890, pp. 45–59. Springer, Heidelberg (2013)
7. Coppo, M., Dezani-Ciancaglini, M., Yoshida, N., Padovani, L.: Global progress for dynamically interleaved multiparty sessions. *MSCS* (to appear)
8. Dezani-Ciancaglini, M., de’Liguoro, U., Yoshida, N.: On progress for structured communications. In: Barthe, G., Fournet, C. (eds.) *TGC 2007. LNCS*, vol. 4912, pp. 257–275. Springer, Heidelberg (2008)
9. Gay, S.J., Hole, M.: Subtyping for session types in the pi calculus. *Acta Inf.* 42, 191–225 (2005)
10. Honda, K.: Types for dyadic interaction. In: Best, E. (ed.) *CONCUR 1993. LNCS*, vol. 715, pp. 509–523. Springer, Heidelberg (1993)
11. Honda, K., Vasconcelos, V.T., Kubo, M.: Language primitives and type discipline for structured communication-based programming. In: Hankin, C. (ed.) *ESOP 1998. LNCS*, vol. 1381, pp. 122–138. Springer, Heidelberg (1998)
12. Honda, K., Yoshida, N., Carbone, M.: Multiparty asynchronous session types. In: *POPL 2008*, pp. 273–284. ACM (2008)
13. Kobayashi, N.: A type system for lock-free processes. *Inf. Comput.* 177, 122–159 (2002)
14. Kobayashi, N.: A new type system for deadlock-free processes. In: Baier, C., Hermanns, H. (eds.) *CONCUR 2006. LNCS*, vol. 4137, pp. 233–247. Springer, Heidelberg (2006)
15. Milner, R., Parrow, J., Walker, D.: A calculus of mobile processes, part I and II. *Inf. Comput.* 100, 1–77 (1992)
16. Padovani, L.: From Lock Freedom to Progress Using Session Types. In: *Proceedings of the 6th Workshop on Programming Language Approaches to Concurrency and Communication-cEntric Software (PLACES 2013). EPTCS*, vol. 137, pp. 3–19 (2013)
17. Padovani, L.: Deadlock and lock freedom in the linear π -calculus. Technical report, HAL (2014)
18. Padovani, L., Vasconcelos, V.T., Vieira, H.T.: Typing liveness in multiparty communicating systems. Technical report (2014), <http://hal.inria.fr/hal-00960879>
19. Sangiorgi, D.: The name discipline of uniform receptiveness. *Theor. Comput. Sci.* 221, 457–493 (1999)
20. Vasconcelos, V.T.: Typed concurrent objects. In: Pareschi, R. (ed.) *ECOOP 1994. LNCS*, vol. 821, pp. 100–117. Springer, Heidelberg (1994)
21. Vasconcelos, V.T.: Fundamentals of session types. *Inf. Comput.* 217, 52–70 (2012)
22. Torres Vieira, H., Thudichum Vasconcelos, V.: Typing progress in communication-centred systems. In: De Nicola, R., Julien, C. (eds.) *COORDINATION 2013. LNCS*, vol. 7890, pp. 236–250. Springer, Heidelberg (2013)