

A SESSION TYPE SYSTEM FOR ASYNCHRONOUS UNRELIABLE BROADCAST COMMUNICATION

DIMITRIOS KOUZAPAS ^a, RAMŪNAS FORSBERG GUTKOVAS ^b,
A. LAURA VOINEA ^c, AND SIMON J. GAY ^c

^a KIOS Centre of Excellence, University of Cyprus, Cyprus

^b Ericsson AB, Gothenburg, Sweden

^c School of Computing Science, University of Glasgow, UK

ABSTRACT. Session types are formal specifications of communication protocols, allowing protocol implementations to be verified by typechecking. Up to now, session type disciplines have assumed that the communication medium is reliable, with no loss of messages. However, unreliable broadcast communication is common in a wide class of distributed systems such as ad-hoc and wireless sensor networks. Often such systems have structured communication patterns that should be amenable to analysis by means of session types, but the necessary theory has not previously been developed. We introduce the Unreliable Broadcast Session Calculus, a process calculus with unreliable broadcast communication, and equip it with a session type system that we show is sound. We capture two common operations, broadcast and gather, inhabiting dual session types. Message loss may lead to non-synchronised session endpoints. To further account for unreliability we provide with an autonomous recovery mechanism that does not require acknowledgements from session participants. Our type system ensures soundness, safety, and progress between the synchronised endpoints within a session. We demonstrate the expressiveness of our framework by implementing Paxos, the textbook protocol for reaching consensus in an unreliable, asynchronous network.

1. INTRODUCTION

Networks that use a shared, stateless communication medium, such as wireless sensor networks or ad-hoc networks, are widely used. They are complex in nature, and their design and verification is a challenging topic. In the presence of unreliability, these networks feature structured communication, which lends itself to being formalised. The use of a formalism in a first step can help in understanding and implementing network protocols, and then move on to proving their correctness.

One such formalisation is session types [HVK98]. Session types specify structured communication (protocols) in concurrent, distributed systems. They allow implementations of such protocols to be verified by type-checking, for properties such as session fidelity and deadlock freedom. A prolific topic of research, session types have found their way into many programming languages and paradigms [ABB⁺16], with several session type

Key words and phrases: Session types, type safety, broadcasting, unreliability, Paxos.

technologies developed [GR17]. An assumption of session types so far has been the reliability of communication. That is, that messages are never lost and are always delivered to the receiver. However, it is not realistic to assume reliability in ad-hoc and wireless sensor networks. Such networks use a shared, stateless communication medium and broadcast to deliver messages. Broadcast collisions and the stateless nature of the communication medium will occasionally lead to message loss.

In this paper we introduce *the Unreliable Broadcast Session Calculus* (UBSC) accompanied by a safe and sound session type system. The semantics of the calculus are inspired by the practice of ad-hoc and sensor networks. More concretely, the semantics describe asynchronous broadcast and gather operations between a set of network nodes operating in an unreliable context. Unreliable communication allows for link failures and lost messages, that may lead to situations where a network node is not synchronised with the rest of the computation. To cope with these cases, the semantics of the calculus offer flexible recovery mechanisms that follow the practice of ad-hoc and sensor networks.

We develop a session type system based on binary session types [HVK98] for the UBSC. The syntax and the duality operator for the UBSC session types are identical to the syntax and duality operator for standard binary session types. However, a non-trivial type system describes the interaction between multiple nodes, in contrast to binary session types that describe interaction only between two processes. Moreover, the type system, ensures session safety and soundness of the communication interaction and of the recovery mechanisms. The main idea for the type system is to interpret the case of failure and message loss as a node that is not synchronised with the overall session type protocol. The session type system ensures that the conditions for recovery are adequate to support safe session interaction. The type system is proved to be sound via a type preservation theorem and safe via a type safety theorem stating that a well-typed network will never reduce to an undesirable/error state. The type system also ensures progress within a session, under standard conditions of non-session interleaving. An additional progress result states that every non-synchronised node within a typed process may eventually recover.

Notably, our system does not introduce the description of failure and recovery at the type level as in [APN17, CHY08, CGY16]. Our approach assumes that every interaction described at the type level can fail and can lead to non-synchronised session endpoints. Moreover, all non-synchronised endpoints can eventually recover. In [APN17, CHY08, CGY16], non-synchronised endpoints (or multiparty roles) recover at runtime using safe recovery as permitted by the session type description. Thus, the description of recovery actions at each communication interaction would require long and tedious description of protocols. As shown in Section 2.2, a second advantage of our approach is the lack of global, and often complicated and unnatural, synchronisation between failed session endpoints (or multiparty roles) in order for the session to safely recover.

1.1. Unreliable Broadcast Session Communication. The communication assumptions of the UBSC are based on the practice of networks such as ad-hoc and wireless networks, and Ethernet networks. Specifically, the semantics of the calculus are justified by the following assumptions:

A1. The network nodes operate within a shared, stateless communication medium. The term broadcast in these networks is understood when a node transmits a message within

the shared communication medium, which is then “instantly” received by a set of nodes that share the communication medium.

- A2. The communication medium is considered unreliable, therefore transmitted messages may fail to be delivered to some or to all the receivers. In practice, unreliability arises for various reasons, e.g. weak transmission signal or message collisions.
- A3. A network node transmits a message only once; we assume that a message is received at most once, i.e. there is no duplication of successfully received messages.
- A4. There is no mechanism to acknowledge a successful message reception, i.e. upon reception a receiver network node does not reply with an acknowledgement message. Lack of acknowledgement messages, implies that a transmitting network node cannot possibly know the subset of network nodes that successfully received the transmitted message.
- A5. Messages are either lost, or delivered without being corrupted. This can be achieved by assuming an underlying mechanism that detects, e.g. through parity bit comparison, and rejects all corrupted messages. Rejected messages are considered lost.
- A6. The network nodes operate at an arbitrary speed. Each network node decides to perform an interaction (either transmission, or internal processing) at an arbitrary moment.

The semantics of the UBSC deploy the mechanisms that support safe session interaction and, at the same time, respect assumptions A1-A6. Specifically, the semantics of the calculus deploy mechanisms that respect the following session type principles:

- S1. We assume binary interaction within a session. The interaction within a session name s is defined between a \check{s} -endpoint, uniquely used by a single network node, and a s -endpoint shared by an arbitrary number of network nodes.

The one to many correspondence between endpoints gives rise to a broadcasting operation, as imposed by assumption A1, where the \check{s} -endpoint broadcasts a value towards the s -endpoints, and a gather operation, where the \check{s} -endpoint gathers messages sent from the s -endpoints.

- S2. Safe session interaction requires that messages are delivered in the order they are transmitted. Ordered message delivery is a direct consequence of assumptions A1, A3, and A5.
- S3. Session interaction is subject to unreliability. Assumptions A2 and A6 imply that session endpoints may become non-synchronised with the overall session interaction.
- S4. Safe session interaction requires that two session endpoints can communicate whenever they are synchronised.
- S5. To ensure progress, endpoints that cannot progress, e.g. because the endpoint is not synchronised or because the opposing endpoint is deadlocked, need to be able to autonomously recover. Autonomous recovery is a consequence of assumptions A4 and A6 that imply that a network node does not have global information whether an endpoint can progress or not, a behaviour typical for ad-hoc and wireless sensor networks.

Below, we use a simple example to introduce the basic assumptions of the UBSC. Prior to the example, consider the following informal presentation of the syntax of the calculus. The calculus defines the syntax for *network nodes*. Specifically, a network node,

$$N = [P \mid \prod_{i \in I} s_i[c, \tilde{m}_i]]$$

composes a binary session π -calculus [HVK98] process P together with a finite parallel composition of session buffer terms, $\prod_{i \in I} s_i[c, \tilde{m}_i]$.

A session π -calculus process of the form $s_1\langle v \rangle.P$ denotes a process that is ready to send value v via channel (or session) endpoint s and proceed as process P . Dually, a process of

the form $s?(x).P$ denotes a process that is ready to receive a value on channel endpoint s and substitute it on variable x within process P . A buffer term, $s[c, \tilde{m}]$, represents a first-in first-out message buffer that interacts on session endpoint s . The buffer stores messages \tilde{m} and keeps track of the session endpoint state using integer counter c . Multiple network nodes can be composed in parallel $N_1 \parallel \dots \parallel N_n$ to form a network. We often write term, $\prod_{j \in J} [P_j \mid \prod_{i \in I_j} s_i[c, \tilde{m}_i]]$ to represent a parallel composition of network nodes.

The next example demonstrates the semantics and basic typing ideas for the UBSC. Variations of the example will be used as a running example throughout the paper.

Example 1.1 (A simple Heartbeat Protocol). The heartbeat protocol, is a simple sensor network protocol where one, or more, network nodes periodically broadcast a *heartbeat* message to signal that they are alive. In the UBSC, a simple heartbeat interaction can be specified with the following network.

$$\text{Heartbeat} = [\check{s}_1(\text{hbt}).P_0 \mid \check{s}[0, \varepsilon]] \parallel \prod_{i \in I} [s?(x).P_i \mid s[0, \varepsilon]]$$

The network implements the requirements of S1 on session channel s ; the \check{s} -endpoint is uniquely used by network node $[\check{s}_1(\text{hbt}).P_0 \mid \check{s}[0, \varepsilon]]$, and the s -endpoint is shared among an arbitrary number of network nodes, represented by network $\prod_{i \in I} [s?(x).P_i \mid s[0, \varepsilon]]$.

Under the requirements of S2 for ordered delivery of messages, it is typical for nodes in ad-hoc and wireless networks, to deploy a received message buffer that follows a first-in first-out policy. The use of a message buffer leads to asynchronous communication semantics, where we first observe a message stored in a message buffer and then extracted from the buffer for processing. At each buffer there is also a counter which is increased with every endpoint interaction and keeps track of the session endpoint state. The counter is used by the semantics to maintain interaction between synchronised endpoints, as required by S4. In the example above, each session endpoint is associated with a corresponding empty message buffer, $s[0, \varepsilon]$, (similarly $\check{s}[0, \varepsilon]$ for the \check{s} -endpoint). The state counter designates that all endpoints are in the 0 state.

A heartbeat (*hbt*) message is broadcast on the \check{s} -endpoint and it is received by the s -endpoints, as expected by the requirements of S1. As required by S3, the broadcast interaction is subject to unreliability and is captured by the reduction:

$$\text{Heartbeat} \longrightarrow \text{Heartbeat}' = \begin{aligned} & [P_0 \mid \check{s}[1, \varepsilon]] \\ & \parallel \prod_{j \in J} [s?(x).P_j \mid s[1, \text{hbt}]] \parallel \prod_{k \in K} [s?(x).P_k \mid s[0, \varepsilon]] \end{aligned}$$

where $I = J \cup K$ and J, K are disjoint. The broadcast message is received instantly during the transmission time. Failure of reception is modelled by the fact that only an arbitrary subset of the network nodes, indexed by set J , accept the heartbeat message. Each successful receiver uses the session queue to buffer the heartbeat message and model asynchronous communication.

Following the requirements of S4, the reduction semantics, defined in Section 2.2, allow for interaction only between the \check{s} -endpoint and the s -endpoints that are in the same state. To maintain synchronisation, a successful interaction increases the session state counter on each buffer by 1. The network nodes that failed to receive the heartbeat messages, indexed by set K , do not update their session counter, therefore they are considered non-synchronised and cannot continue to safely interact with the \check{s} -endpoint.

In practice, state counting implies an underlying mechanism where each transmitted message includes a header that tags the message with the corresponding session and state counter. A network node can only accept a message if it implements the corresponding

session endpoint and the session endpoint is synchronised with the state of the message. In a different case the message is dropped and considered lost.

Each receiver will then interact locally with its own queue and extract the message for processing. For example, network node, $[s?(x).P_q \mid s[1, \mathbf{hbt}]]$ will interact with its local buffer using the reduction:

$$\begin{aligned} \text{Heartbeat}' \longrightarrow & \quad [P_0 \mid \check{s}[1, \varepsilon]] \parallel [P_q \{\mathbf{hbt}/x\} \mid s[1, \varepsilon]] \\ & \parallel \prod_{j \in J \setminus \{q\}} [s?(x).P_j \mid s[1, \mathbf{hbt}]] \parallel \prod_{k \in K} [s?(x).P_k \mid s[0, \varepsilon]] \end{aligned}$$

In network $\text{Heartbeat}'$, the nodes described by set K appear to be stuck because the receiving s -endpoints are non-synchronised with the \check{s} -endpoint. The requirements of S5 impose the development of recovery semantics, in order to observe session progress. For example, recovery on network $\text{Heartbeat}'$ is described by the interaction.

$$\begin{aligned} \text{Heartbeat}' \longrightarrow & \quad [P_0 \mid \check{s}[1, \varepsilon]] \parallel [P_q \{\mathbf{1}/x\} \mid s[1, \varepsilon]] \\ & \parallel \prod_{j \in J} [s?(x).P_j \mid s[1, \mathbf{hbt}]] \parallel \prod_{k \in K \setminus \{q\}} [s?(x).P_k \mid s[0, \varepsilon]] \end{aligned}$$

Whenever a session endpoint is possibly unable to progress, the reduction semantics use a default expression, in this case the unit ($\mathbf{1}$) expression, to substitute the receive variable.

However, due to assumptions A4 and A6, a network node does not have knowledge whether an endpoint can progress or not. This situation also exists in the practice of ad-hoc and wireless sensor networks that approximate lack of progress using necessary but not sufficient conditions, combined with internal mechanisms, e.g. timeout signals, to achieve recovery and progress. Here the conditions that are necessary but not sufficient to detect lack of progress are the receiving prefix and the empty s -buffer.

The typing system ensures the conditions for a safe recovery. The full semantics of the calculus offer the programmer multiple and flexible recovery choices to deploy whenever the session safety conditions are met.

1.2. Related Work. The UBSC follows the principles of [KGG14], where the authors provided semantics to a syntax for an unreliable broadcasting session calculus through an encoding into the psi-calculus [BHJ⁺11]. The work also provides with a sound session type system for the aforementioned syntax. In comparison to [KGG14], in this work we define the syntax and semantics of a more sophisticated calculus that includes distinction between processes and nodes, asynchronous semantics through buffers, and more advanced recovery mechanisms. The type system is also non-trivial, despite the fact that the session types syntax remains the same. In contrast to [KGG14], we also prove type safety and progress results.

Asynchronous session semantics and their corresponding session typing systems are studied, both for binary [HKP⁺10, KYHH16] and multiparty [HYC08, CDYP16] session types. However, our calculus is the first calculus that supports asynchronous broadcast semantics, in contrast to point-to-point communication proposed by state-of-the-art session type calculi.

Session types for reliable gather semantics were proposed in [CDP12] but a corresponding session calculus, was never proposed. Structured multiparty session interaction for a parametrised (arbitrary) number of participants is studied in [DYBH12, NY15], where the authors propose a framework for describing the behaviour of an arbitrary number of participants that implement different interacting roles. A parametrised multiparty session type protocol can be instantiated to a specific number of participants. Our framework can

handle an arbitrary number of agents within a session as a result of the shared unreliable communication medium. Agents might arbitrarily lose messages and become non-synchronised and arbitrarily recover and re-synchronise.

Session type systems in an unreliable context are studied in [APN17, CHY08, CGY16]. Adameit et al. [APN17] extend multiparty session types with optional blocks that cover a limited class of link failures. Specifically, the authors extend the typing syntax with constructs that allow the description of a default value each role needs to receive whenever there is a possibility of link of failure. In contrast, our framework does not require to describe failures at the type level, rather than it enforces a number of conditions for a session safe recovery whenever a failure occurs. Furthermore, it assumes that every interaction is subject to failure and proposes several recovery mechanisms, inspired by the practice of wireless sensor-networks. Our calculus, also, supports broadcast semantics in contrast to point-to-point communication supported by [APN17].

Recovery after failure, was introduced in the shape of exception handling in the binary [CHY08] and multiparty [CGY16] session types. Both works present a procedure where upon an exception during communication the session endpoints are informed and safely handle the exception. Similarly to [APN17], the type discipline proposes a complicated syntax for exception description at the type level. Moreover, exception handling within session types assumes strong global synchronisation requirements among the session endpoints. In contrast, our recovery semantics are implemented locally, thus being more natural and general; each network node can autonomously and safely recover from a communication failure, by choosing between several recovery mechanisms.

The report in [PNC23] presents fault-tolerant multiparty session types; an extension to multiparty session types that handles failures such as unreliable communication and process crashes, applied to a calculus that supports failure patterns. Moreover, this work demonstrates fault tolerant multiparty session types with an application on the rotating coordination algorithm. An additional work that describes recovery patterns is found in [BD23], which develops an asynchronous multiparty framework that accommodates non-Byzantine faults in unreliable settings. The work in [PNC23] has similar recovery semantics as our work adapted in the context of multiparty session types. Nevertheless, we are the first to define patterns of broadcast and gather that are found in more dynamic systems with share communication medium, such as wireless sensor networks. Moreover, our work is the first to demonstrate an implementation of the Paxos protocol defined as sessions of structured interaction.

1.3. Overview. In Section 2 we present the Unreliable Broadcast Session Calculus — UBSC. UBSC is accompanied by the first type system for structured communication in an unreliable broadcast setting. The calculus develops all the necessary mechanisms to satisfy the requirements S1-S5 under the assumptions A1-A6.

Section 3 presents the UBSC session type system. Interestingly, we do not introduce the description of unreliability at the type level, keeping the session type syntax identical to standard binary session type syntax. The type system makes use of the synchronisation notion to cope with non-synchronised session endpoints and to ensure session duality.

In Section 4, the type system is proved to be sound via a type preservation theorem and safe via a type safety theorem. A process is safe whenever it can never reduce to an error process. In turn, error processes are a class of processes that do not respect the session

types principles. Section 4 also includes a set of progress results that ensure safe progress within a session and safe session recovery.

We demonstrate the expressiveness of our framework in Section 5 with a session specification of the Paxos consensus algorithm, which is the standard consensus algorithm in distributed systems. We argue that our framework can provide support for the implementation of such protocols and their extensions.

Finally, section 6 discusses the possibility for future work and concludes the article.

2. ASYNCHRONOUS UNRELIABLE BROADCAST SESSION CALCULUS

In this Section, we define the syntax and the semantics for the Unreliable Broadcast Session Calculus. The semantics are extensively demonstrated using several examples.

2.1. Syntax. Assume the following disjoint sets of names/variables: \mathcal{C} is a countable set of *shared channels* ranged over by a, b, \dots ; \mathcal{S} is a countable set of *session channels* ranged over by s, s', \dots , where each session channel has two distinct endpoints s and \check{s} (we write κ to denote either s or \check{s}); \mathcal{V} is a countable set of *variables* ranged over by x, y, z, \dots ; and \mathbf{Lab} is a countable set of labels ranged over by ℓ, ℓ', \dots . We let n range over shared channels or sessions. We write k to denote either κ or x or \check{x} , where \check{x} is used to distinguish a variable used as a \check{s} -endpoint.

Let c to range over natural numbers, \mathbb{N} and let **true** and **false** be the boolean values. Let \mathcal{E} be a non-empty set of *expressions* ranged over by e, e', \dots . Elements of \mathcal{E} contain natural numbers and boolean values, and may contain variables. Function $\text{fv}(e)$ returns the variables in expression e . Expressions that do not contain variables, i.e. $\text{fv}(e) = \emptyset$, are called closed. Assume a binary operation \odot on \mathcal{E} called *aggregation* operator, and an element $\mathbf{1}$ of \mathcal{E} called *unit*. We define an evaluation operator $\downarrow: \mathcal{E} \rightarrow \mathcal{E}$ from closed expressions to single value expressions (natural numbers, boolean values, $\mathbf{1}$ value, etc.). Let $\mathcal{F} \subseteq \mathcal{E}$ be a non-empty set of *conditions* ranged over by φ . Closed conditions are evaluated to boolean values $\varphi \downarrow = \mathbf{true}$ or $\varphi \downarrow = \mathbf{false}$. We use metavariable u to denote either shared names, session names, expressions or variables.

The syntax of processes, $P, Q, R \in \mathcal{P}$, buffers, $B \in \mathcal{B}$, and networks, $N \in \mathcal{N}$, is then defined in Figure 1. Functions returning the set of free names, $\text{fn}(P)$, bound names, $\text{bn}(P)$, and free and bound names, $\text{n}(P)$ are defined in the expected way. Terms **(Req)** and **(Acc)** bind \check{x} and x in P respectively, and term **(Rcv)** binds x in P . In term **(Def)**, $\text{def } \{D_i(\tilde{x}_i) \stackrel{\text{def}}{=} P_i\}_{i \in I}$ in P , $D_i(\tilde{x}_i)$ binds \tilde{x}_i in P_i . Moreover, terms $\{D_i\}_{i \in I}$ are bound in P . Term **(Inact)** is the inactive term. Terms **(Req)** and **(Acc)** express the processes that are ready to initiate a fresh session on a shared channel a via a request/accept interaction, respectively. Term **(Snd)** defines a prefix ready to send an expression on session k . Term **(Rcv)** defines a prefix ready to receive a message on session k and substitute it on x . The prefix also provides with an expression e , called default expression, that will be substituted on x in the case of recovery. We often write $k?(x).P$ for $k?(x)(\mathbf{1}).P$. The select prefix, defined by term **(Sel)**, is ready to send a label ℓ over session k . Dually, the branch prefix, defined by term **(Bra)**, is ready to receive a label from a predefined set of label $\{\ell_i\}_{i \in I}$ on session k . Moreover, the branch prefix defines a default label, **df**, with a process R used for in the case of recovery.

In a highly dynamic unreliable environment, it is convenient to consider non-deterministic choice — term **(Sum)**. We write $\sum_{1 \leq i \leq n} P_i$ for process $P_1 + \dots + P_n$. Term **(Cond)** is a standard

$a, b, \dots \in \mathcal{C}$	(ShCh)	$s, s', \dots \in \mathcal{S}$	(SCh)
$\ell, \ell', \dots \in \mathbf{Lab}$	(Labels)	$x, y, \dots \in \mathcal{V}$	(Var)
$e, e', \dots \in \mathcal{E}$	(Exp)	$c, c', \dots \in \mathbb{N}$	(Cnt)
$\kappa ::= \check{s} \mid s$	(Endp)	$k ::= \kappa \mid x$	(Sess)
$v ::= a \mid \kappa \mid e$	(Identifier)		
$P, Q, R ::= \mathbf{0}$	(Inact)	$k \triangleleft \ell.P$	(Sel)
$a!(\check{x}).P$	(Req)	$k \triangleright \{\ell_i : P_i, \mathbf{df} : R\}_{i \in I}$	(Bra)
$a?(x).P$	(Acc)	$P + P$	(Sum)
$k!(e).P$	(Snd)	$D\langle \check{v} \rangle$	(PVar)
$k?(x)\langle e \rangle.P$	(Rcv)	$\text{if } \varphi \text{ then } P \text{ else } P$	(Cond)
$\mathbf{def} \{D_i(\check{x}_i) \stackrel{\mathbf{def}}{=} P_i\}_{i \in I} \text{ in } P$			(Def)
$N, M ::= [P \mid B]$	(Node)	$m ::= e \mid \ell$	(Msg)
$N \parallel M$	(Par)	$h ::= (c, e)$	(TMsg)
$(\nu n)N$	(Restr)	$B ::= \varepsilon \mid B \mid s[c, \tilde{m}] \mid B \mid \check{s}[c, \tilde{h}]$	(Buffer)

Figure 1: Syntax of Processes, Buffers, and Networks

conditional term. Finally, terms (Def) and (PVar) express a named recursive process definition with parameters, cf. [HVK98]. We assume a standard variable substitution over processes $P\{e/x\}$, inductively defined to include a standard variable substitution over expressions, $e'\{e/x\}$. Moreover, the process variable substitution $P\{P'(\tilde{x})/D\}$ is defined inductively with $D\langle \check{v} \rangle\{P(\tilde{x})/D\} = P\{\tilde{v}/\tilde{x}\}$ as the basic definition case. Concurrency is introduced at the network level, rather than at process level.

Term B , is a parallel composition of session buffers, that are used to store messages and keep track of the session state via natural number c . Buffer terms are used to model a form of asynchrony that preserves the order of received messages, as required by S2. The purpose of counter c is to keep track of the session state in the presence of communication failure, and to synchronise the interaction between session prefixes, as required by S4. Message loss in an unreliable setting leads to session endpoints that are not synchronised with the overall protocol, as expected by S3. To ensure correctness, many frameworks and algorithms that operate in an unreliable setting use message tagging or state counting; for example, the TCP/IP protocol tags packets with unique sequential numbers to maintain consistency in the case of packet loss. In our setting, state counting is necessary to maintain the correct semantics within a session. The type system in Section 3 provides with static guarantees for a session despite the dynamic nature of session reduction.

Buffer terms on s -endpoints store *messages* m that range over expressions e and labels ℓ . Buffer terms on \check{s} -endpoints store messages, h , which are expressions tagged with a session counter, $h = (c, e)$. The session counter in h distinguishes the session state, at which the expression e needs to be received.

Network (Node) consists of a process P , and the necessary buffer terms, B , used for asynchronous session communication. A process may participate in several sessions, and therefore, more than one buffer term may be present in a node. The type system ensures that there is no more than one buffer term on the same session in each network node. We write $[P]$ for node $[P \mid \varepsilon]$.

$$\begin{aligned}
& P_1 + P_2 \equiv P_2 + P_1 \quad (P_1 + P_2) + P_3 \equiv P_1 + (P_2 + P_3) \quad P \equiv_\alpha P' \\
& \text{def } \{D_i(\tilde{x}_i) \stackrel{\text{def}}{=} P_i\}_{i \in I} \text{ in } P \equiv \text{def } \{D_i(\tilde{x}_i) \stackrel{\text{def}}{=} P_i\}_{i \in I} \text{ in } P\{P_k(\tilde{x}_k)/D_k\} \quad k \in I \\
& B | B' \equiv B' | B \quad (B | B') | B'' \equiv B | (B' | B'') \\
& N_1 \parallel N_2 \equiv N_2 \parallel N_1 \quad (N_1 \parallel N_2) \parallel N_3 \equiv N_1 \parallel (N_2 \parallel N_3) \quad N \equiv [\mathbf{0}] \parallel N \\
& (\nu n)(\nu m)N \equiv (\nu m)(\nu n)N \quad (\nu n)N \parallel M \equiv (\nu n)(N \parallel M) \quad \text{if } n \notin \text{fn}(M) \\
& [P | B] \equiv [P' | B'] \quad \text{if } P \equiv P' \text{ and } B \equiv B' \quad N \equiv_\alpha N'
\end{aligned}$$

Figure 2: Structural Congruence for Processes, Buffers, and Networks

A network is a parallel composition of nodes — term (Par). We write $\prod_{i \in I} N_i$ for the parallel composition of $N_1 \parallel \dots \parallel N_n$ for (possibly empty) $I = \{1, \dots, n\}$. Network (Restr) binds both session and shared channels. We write $(\nu \tilde{n})N$ for the network $(\nu n_1) \dots (\nu n_m)N$, where the sequence \tilde{n} may be empty. We also extend the $\text{fn}(\cdot)$ function to networks.

2.2. Operational Semantics. The operational semantics are defined as a reduction relation on networks with the use of a standard structural congruence relation.

Structural Congruence. The structural congruence on processes, resp. buffers and networks, is defined to be the least congruence relation satisfying the rules in Figure 2. Structural congruence on processes considers commutativity and associativity of the $+$ operator, and includes the unfolding of definitions and alpha-conversion. Named definition substitution is defined up-to structural congruence. The parallel composition is commutative and associative for buffer terms and for network terms, with $[\mathbf{0}]$ as the unit for network terms. Name restriction order is irrelevant, and moreover, the scope of restricted channels can be extruded. The clause for network nodes simply bridges the buffer and process congruences with the structural congruence for network. Finally, structural congruence allows alpha renaming for networks.

The operational semantics is defined as the least relation on networks, $N \longrightarrow N'$, satisfying the rules given in Figure 3 (Network Semantics), Figure 4 (Process Communication Semantics), and Figure 5 (Recovery Semantics).

2.2.1. Operational Semantics for Networks. In Figure 3 rules [NDEF], [DEF], [RPAR], [RRRES], and [RCONG] are standard congruence rules for operator $+$, named definition, parallel composition, name restriction, and structural congruence, respectively.

2.2.2. Operational Semantics for Broadcast and Gather Operations. Figure 4 defines the reduction semantics for the broadcast and gather operations. Rule [CONN] establishes a session between a request network node and several accept network nodes. It is a broadcast (one-to-many) communication between the request network node and an arbitrary (possibly empty) set of accepting network nodes described by I . Unreliability is achieved because set I is chosen arbitrarily from a set of parallel nodes and then using the reduction rule [RPAR] to close the reduction of the entire network. The interaction creates a fresh session s

$$\begin{array}{c}
\frac{[P_1 | B] \parallel N \longrightarrow [P' | B'] \parallel N'}{[P_1 + P_2 | B] \parallel N \longrightarrow [P' | B'] \parallel N'} \text{ [NDET]} \\
\\
\text{[DEF]} \\
\frac{[P | B] \parallel N \longrightarrow [P' | B'] \parallel N'}{[\text{def } \{D_i(\tilde{x}_i) \stackrel{\text{def}}{=} P_i\}_{i \in I} \text{ in } P | B] \parallel N \longrightarrow [\text{def } \{D_i(\tilde{x}_i) \stackrel{\text{def}}{=} P_i\}_{i \in I} \text{ in } P' | B'] \parallel N'} \\
\\
\begin{array}{ccc}
\text{[RPAR]} & \text{[RCONG]} & \text{[RRES]} \\
\frac{N \longrightarrow N'}{N \parallel M \longrightarrow N' \parallel M} & \frac{N \equiv N_1 \quad N_1 \longrightarrow N_2 \quad N_2 \equiv N'}{N \longrightarrow N'} & \frac{N \longrightarrow N'}{(\nu n)N \longrightarrow (\nu n)N'}
\end{array}
\end{array}$$

Figure 3: Reduction rules for networks.

$$\begin{array}{c}
\frac{s \text{ fresh}}{[a!(\tilde{x}).P | B] \parallel \prod_{i \in I} [a?(x).P_i | B_i] \longrightarrow (\nu s)([P\{\check{s}/\tilde{x}\} | B | \check{s}[0, \varepsilon]] \parallel \prod_{i \in I} [P_i\{s/x\} | B_i | s[0, \varepsilon]])} \text{ [CONN]} \\
\\
\frac{e \downarrow = e'}{[\check{s}! \langle e \rangle . P | B | \check{s}[c, \tilde{m}]] \parallel \prod_{i \in I} [P_i | B_i | s[c, \tilde{m}_i]] \longrightarrow [P | B | \check{s}[c+1, \tilde{m}]] \parallel \prod_{i \in I} [P_i | B_i | s[c+1, \tilde{m}_i \cdot e']]} \text{ [BCAST]} \\
\\
\frac{c_1 \geq c_2 \quad e \downarrow = e'}{[s! \langle e \rangle . P_1 | B_1 | s[c_1, \tilde{m}]] \parallel [P_2 | B_2 | \check{s}[c_2, \tilde{h}]] \longrightarrow [P_1 | B_1 | s[c_1+1, \tilde{m}]] \parallel [P_2 | \check{s}[c_2, \tilde{h} \cdot (c_1, e')]}] \text{ [UCAST]} \\
\\
[s?(x) \langle e' \rangle . P | B | s[c, e \cdot \tilde{m}]] \longrightarrow [P\{e/x\} | B | s[c, \tilde{m}]] \text{ [RCV]} \\
\\
\frac{\tilde{h}' = \mathbf{B}(\tilde{h}, c) \quad e = \mathbf{V}(\tilde{h}, c)}{[\check{s}?(x).P | B | \check{s}[c, \tilde{h}]] \longrightarrow [P\{e/x\} | B | \check{s}[c+1, \tilde{h}']]} \text{ [GTHR]} \\
\\
[\check{s} \triangleleft \ell . P | B | \check{s}[c, \tilde{m}]] \parallel \prod_{i \in I} [P_i | B_i | s[c, \tilde{m}_i]] \longrightarrow [P | B | \check{s}[c+1, \tilde{m}]] \parallel \prod_{i \in I} [P_i | B_i | s[c+1, \tilde{m}_i \cdot \ell]] \text{ [SEL]} \\
\\
\frac{k \in I}{[s \triangleright \{\ell_i : P_i, \mathbf{df} : R\}_{i \in I} | B | s[c, \ell_k \cdot \tilde{m}]] \longrightarrow [P_k | B | s[c, \tilde{m}]]} \text{ [BRA]}
\end{array}$$

Figure 4: Reduction rules for Broadcast and Gather Operations.

with a unique \check{s} -endpoint in the request side and a shared s -endpoint on the accept sides. A corresponding session buffer term is created in all connected network nodes.

There are two kind of session communication distinguished by the interaction between the \check{s} -endpoint and s -endpoint: broadcast communication; and unicast communication. Rule [BCAST] defines the asynchronous broadcasting semantics; the \check{s} -endpoint broadcasts evaluated message e' , with the message being enqueued to the buffers of the s -endpoint, thus modelling asynchrony. The crucial condition for a broadcast interaction is that all participating nodes are synchronised in the same session state c . After the interaction all participating nodes update their state to $c + 1$. Unreliability is modelled by the fact that the, possibly empty, set I is chosen arbitrarily from a set of parallel nodes, and then using reduction rule [RPAR] to close the reduction for the entire network.

Rule [UCAST] defines the case where an s -endpoint enqueues an evaluated message e' in the buffer of the unique \check{s} -endpoint. The message e' , when enqueued, is tagged with the sender's session state c , as $h = (c, e')$, since all the messages on the same state will be gathered by the \check{s} -endpoint (rule [GTHR]). The session state counter of the s -endpoint will be updated, in contrast to the state counter of the \check{s} -endpoint. The \check{s} -endpoint remains in the same state since it needs to continue the [UCAST] interactions with other network nodes. Increasing the state counter does not disallow the s -endpoint to be involved in a subsequent [UCAST] interaction, i.e the s -endpoint is not considered non-synchronised. This is captured by condition $c_1 \geq c_2$ that ensures that a s -endpoint exists in the same or in a later session state from the \check{s} -endpoint, prior to interacting. The [UCAST] semantics follow the practice of ad-hoc and sensor networks, where a sender node cannot locally know whether the receiving node has performed a gather prior to sending a subsequent message. For further intuition on the [UCAST] rule, see Example 2.1.

Rule [RCV] defines the interaction of a process with its s -endpoint buffer; a (RCV) process on endpoint s receives and substitutes on variable x the next available expression, e , from the s -endpoint buffer. Session state c is not updated, because it was updated by the operation that stored expression e in the buffer. Expression e' is not used in this rule, since it is a default value used for recovery in the case where the session endpoint becomes non-synchronised. The case for recovery is described by recovery rule [REC] defined in Section 2.2.3.

Rule [GTHR] defined the interaction between a process and its \check{s} -endpoint buffer, where it gathers, via the \odot operator, all the expression messages that are tagged with state, c , of the state of the \check{s} -endpoint. After the reduction the state of the \check{s} -endpoint increases by 1. The rule uses auxiliary operations $V(\tilde{h}, c)$ and $B(\tilde{h}, c)$:

$$\begin{aligned} V((c, e) \cdot \tilde{h}, c) &= e \odot V(\tilde{h}, c) & V(\varepsilon, c) &= \mathbf{1} & V((c', e) \cdot \tilde{h}, c) &= V(\tilde{h}, c) & \text{if } c' \neq c \\ B((c, e) \cdot \tilde{h}, c) &= B(\tilde{h}, c) & B(\varepsilon, c) &= \varepsilon & B((c', e) \cdot \tilde{h}, c) &= (c', e) \cdot B(\tilde{h}, c) & \text{if } c' \neq c \end{aligned}$$

Operation $V(\tilde{h}, c)$ goes through the messages \tilde{h} and returns, up to operator \odot , all expressions e tagged with session state c , (c, e) . Operation $B(\tilde{h}, c)$ returns a new \tilde{h}' by removing all messages tagged with session state c , (c, e) . The [GTHR] semantics also describes the case where the \check{s} -endpoint gathers no messages, capturing the case where all message where either lost or not delivered yet; operator $V(\tilde{h}, \check{s})$ will return unit, $\mathbf{1}$, if there are no messages tagged with state c in \tilde{h} . The gather pattern is common in ad-hoc networks; see, for example, the RIME communication stack [DOH07] for wireless sensor networks.

Rules [SEL] and [BRA] are similar to rule [BCAST] and [RCV]; the \check{s} -endpoint selects and broadcasts a label to the corresponding s -endpoint buffer terms, and dually the s -endpoints

$$\begin{array}{c}
\frac{e \downarrow = e'}{[s?(x)\langle e \rangle.P \mid B \mid s[c, \varepsilon]] \longrightarrow [P\{e'/x\} \mid B \mid s[c+1, \varepsilon]]} \quad [\text{REC}] \\
\\
\frac{\text{fs}(B) = \text{fs}(R)}{[s \triangleright \{\ell_i : P_i, \text{df} : R\}_{i \in I} \mid B \mid B' \mid s[c, \varepsilon]] \longrightarrow [R \mid B]} \quad [\text{BREC}] \\
\\
[s!(e).P \mid B \mid s[c, \tilde{m}]] \longrightarrow [P \mid B \mid s[c+1, \tilde{m}]] \quad [\text{LOSS}] \\
\\
\frac{\varphi \downarrow = \text{true} \quad \text{fs}(B) = \text{fs}(P_1)}{[\text{if } \varphi \text{ then } P_1 \text{ else } P_2 \mid B \mid B'] \longrightarrow [P_1 \mid B]} \quad [\text{TRUE}] \\
\\
\frac{\varphi \downarrow = \text{false} \quad \text{fs}(B) = \text{fs}(P_2)}{[\text{if } \varphi \text{ then } P_1 \text{ else } P_2 \mid B \mid B'] \longrightarrow [P_2 \mid B]} \quad [\text{FALSE}]
\end{array}$$

Figure 5: Reduction rules for recovery.

receive a label from its session buffer and proceeds accordingly. The dual case where multiple s -endpoint select a label is not defined, since gathering (i.e. branching on) multiple labels on the \check{s} -endpoint makes no sense in session types semantics. Rule [BRA] also defines a default label, **df**, with a process R , used when there is a need to recover. Recovery on process R is described by recover rule [BREC] defined in Section 2.2.3.

An implementation of the above semantics in real systems, e.g. wireless ad-hoc networks, requires that sent messages are tagged with the the session state of the sender. This allows a prospective receiver to check the session state conditions of the interaction and act accordingly, e.g. a receiver will check its session state against the message tag and only then will store a broadcast message.

2.2.3. Operational Semantics for Recovery. The semantics of the calculus offer flexibility, and recovery in a context of asynchrony and unreliability. Figure 5 defines the semantics for these mechanisms. Crucially, the recovery semantics depend on conditions that are checked locally within a network node, thus following the typical behaviour of ad-hoc and sensor networks.

Rule [REC] describes the recovery conditions for a s -endpoint input prefix. Lack of messages in the corresponding s -buffer may trigger recovery, where the evaluation of the default expression is substituted in the continuation of the process. The interaction results in an increase of the s -endpoint session state counter, aligning the s -endpoint with the overall session interaction.

Rule [BREC] describes the conditions for recovering from a branch prefix. Lack of messages in the corresponding s -buffer may trigger a recovery, by continuing with the default label process R . A branch recover is hard; the reduction drops the corresponding s -endpoint and s -buffer. A process R cannot implement a behaviour for the s -endpoint because it cannot possibly know the corresponding choice from the \check{s} -endpoint, in the case recovery results in a synchronised s -endpoint.

Reduction rules $[\text{REC}]$ and $[\text{BREC}]$ describe recovery from a situation where a s -endpoint might not be able to progress. For example, it may be the case that the s -endpoint is non-synchronised, or the network node using the \check{s} -endpoint is deadlocked due to the interleaving with other behaviour. However, due to assumptions A2, A4 and A6, a network node cannot possibly have the knowledge to decide whether a local session endpoint can progress or not. Assumption A4 implies that a node cannot know whether a send message was received. Assumption A2 implies that a network node cannot know whether an expected message was lost and, moreover, due to assumption A6, if an expected message was not send yet.

Therefore, reduction rules $[\text{REC}]$ and $[\text{BREC}]$ allow nodes to act autonomously, which means that network nodes can recover even if their recovery is not necessary to achieve progress. For example, the following reduction on rule $[\text{REC}]$

$$[\check{s}!\langle e \rangle.\mathbf{0} \mid \check{s}[0, \varepsilon]] \parallel [s?(x).\mathbf{0} \mid \check{s}[0, \varepsilon]] \longrightarrow [\check{s}!\langle e \rangle.\mathbf{0} \mid \check{s}[0, \varepsilon]] \parallel [\mathbf{0} \mid \check{s}[1, \varepsilon]]$$

is not necessary to achieve progress, because the \check{s} -endpoints and s -endpoint are in a state that can eventually interact.

In practice communicating nodes require necessary but not sufficient conditions together with mechanisms, such as time-outs, to approximate lack of progress. This means that, in practice, there is always the possibility to recover even if endpoint interaction is eventually possible. We could define, for example, semantics that use global information to recover but then we would not respect assumptions A4 and A6. For example, rule $[\text{REC}]$ might be defined as:

$$\frac{c < c'}{[s?(x)\langle e \rangle.P \mid B \mid s[c, \varepsilon]] \parallel [P' \mid B' \mid \check{s}[c', \tilde{h}]] \longrightarrow [P\{e/x\}B \mid s[c+1, \varepsilon]] \parallel [P' \mid B' \mid \check{s}[c', \tilde{h}]]}$$

The rule uses global information between network nodes. Condition $c < c'$ checks that the s -endpoint exists in an earlier state than the \check{s} -endpoint. The recovery takes place only after ensuring that the s -endpoint is non-synchronised. Rule $[\text{BREC}]$ can also be accommodated in a similar fashion.

Rule $[\text{LOSS}]$ allows for message loss on the s -endpoint; a process simply proceeds by dropping its sending prefix. Similarly with rule $[\text{REC}]$, the session state is updated/increased, allowing the s -endpoint to re-synchronise with the \check{s} -endpoint. Rules $[\text{LOSS}]$ and $[\text{UCAST}]$ realistically model unreliable unicast communication since messages can always be lost and, moreover, due to assumption A4 a node cannot possibly know if a message was passed to the receiver and should always update its state. The application of rule $[\text{LOSS}]$ implies two cases: i) the s -endpoint exists in a later state than the state of the \check{s} -endpoint, in which case the s -endpoint is considered synchronised, similarly to rule $[\text{UCAST}]$; ii) the s -endpoint prior to $[\text{LOSS}]$ exists in an earlier state than the \check{s} -endpoint, in which case the application of the $[\text{LOSS}]$ rule is the only interaction that can be done towards recovery and synchronisation of the s -endpoint.

Rules $[\text{TRUE}]$ and $[\text{FALSE}]$ offer the programmer additional flexibility when handling shared session endpoints. The conditional process gives the ability to discontinue using some sessions in its branches. Thus, a session can be dropped based on the truth evaluation of condition φ . When the process proceeds to a branch it must drop the buffers not used by the continuation (condition $\text{fs}(B') = \text{fn}(P_i) \setminus \text{fn}(P_j), i \neq j$). The case where no sessions are dropped, corresponds to standard conditional semantics. Explicitly dropping a communication channel is a common pattern in the context of unreliable communication.

2.3. Encoding Node Failure and Recovery Process. The UBSC semantics are powerful enough to express sophisticated patterns of network node failure and recovery. We provide the semantics for a recovery pattern, typically found in ad-hoc and sensor network, through an encoding in the UBSC. The encoding is useful to avoid long and tedious descriptions of recovery interaction.

The recovery pattern takes advantage of the interplay between rules [REC], [TRUE], and [FALSE], to describe a mechanism that recovers by proceeding to a recovery process. We defined the UBSC with recovery terms by extending the terms of UBSC in Figure 1 as:

$$P ::= \dots \mid P \diamond R \text{ (Recov)}$$

Additionally, in the case of a branch prefixed process we write $k \triangleright \{l_i : P_i\}_{i \in I} \diamond R$, instead of $k \triangleright \{l_i : P_i, \mathbf{df} : R'\}_{i \in I} \diamond R$.

The semantics of the processes of the form $P \diamond R$ are then defined through a syntactic encoding into the terms of UBSC (syntax of Figure 1):

$$\begin{aligned} \llbracket \mathbf{0} \diamond R \rrbracket &\stackrel{\text{def}}{=} \mathbf{0} \\ \llbracket D\langle \tilde{v} \rangle \diamond R \rrbracket &\stackrel{\text{def}}{=} D\langle \tilde{v} \rangle \\ \llbracket k?(x).P \diamond R \rrbracket &\stackrel{\text{def}}{=} k?(\mathbf{exc})\langle x \rangle. \text{if } x \neq \mathbf{exc} \text{ then } \llbracket P \diamond R \rrbracket \text{ else } R \\ \llbracket k \triangleright \{l_i : P_i\}_{i \in I} \diamond R \rrbracket &\stackrel{\text{def}}{=} k \triangleright \{l_i : (\llbracket P_i \diamond R \rrbracket), \mathbf{df} : R\}_{i \in I} \\ \llbracket (\mathbf{def} \{D_i(\tilde{x}_i) \stackrel{\text{def}}{=} P_i\}_{i \in I} \text{ in } P) \diamond R \rrbracket &\stackrel{\text{def}}{=} \mathbf{def} \{D_i(\tilde{x}_i) \stackrel{\text{def}}{=} \llbracket P_i \diamond R \rrbracket\}_{i \in I} \text{ in } \llbracket P \diamond R \rrbracket \end{aligned}$$

and homomorphic for the rest of the syntax of the UBSC with recovery processes. The encoding assumes the existence of an *exception* value, \mathbf{exc} , which is used internally for the purpose of detecting recovery. Encoding for terms (Inact) and (PVar) is the identity. The recovery behaviour of the encoding for terms (Rcv) and (Bra) is based on reduction rules [REC], [BREC], [TRUE], and [FALSE]; whenever a process exists in a s -endpoint input prefix (prefixes receive and branch) and the recovery conditions for rules [REC] and [BREC] hold, the network node can recover by continuing to a recovery process R . Note that sessions can be dropped following the semantics of rules [BREC], [TRUE], and [FALSE]. Finally, term (Def) is defined inductively both on the named definitions and on the process body.

2.4. Reduction Semantics Examples. The next two examples demonstrate some rules and basic intuition of the operational semantics. The first example demonstrates the semantics for rule [CONN], and the interplay between the rules [UCAST], [LOSS], and [GTHR].

Example 2.1 (A Heartbeat Protocol). Consider a variant of the Heartbeat protocol, introduced in Example 1.1, where a node periodically gathers heartbeat messages from nodes within the network.

$$\text{Heartbeat} = \llbracket a_l(\check{y}).\check{y}?(x_1).\check{y}?(x_2).P_0 \rrbracket \parallel \llbracket a?(y).y!(\mathbf{hbt}_1).y!(\mathbf{hbt}_1).P_1 \rrbracket \parallel \llbracket a?(y).y!(\mathbf{hbt}_2).y!(\mathbf{hbt}_2).P_2 \rrbracket$$

Node $\llbracket a_l(\check{y}).\check{y}?(x).\check{y}?(x).P_0 \rrbracket$ requests a new session on shared name a . After the establishment of the new session each accepting node will periodically send a heartbeat message to the

requestor node. Consider now the interaction:

$$\begin{aligned}
\text{Heartbeat} &\longrightarrow (\nu s)([\check{s}_?(x_1).\check{s}_?(x_2).P_0 \mid \check{s}[0, \varepsilon]] \\
&\parallel [s_!(\text{hbt}_1).s_!(\text{hbt}_1).P_1 \mid s[0, \varepsilon]] \parallel [s_!(\text{hbt}_2).s_!(\text{hbt}_2).P_2 \mid s[0, \varepsilon]]) \\
&\longrightarrow (\nu s)([\check{s}_?(x_1).\check{s}_?(x_2).P_0 \mid \check{s}[0, (0, \text{hbt}_2)]] \\
&\parallel [s_!(\text{hbt}_1).s_!(\text{hbt}_1).P_1 \mid s[0, \varepsilon]] \parallel [s_!(\text{hbt}_2).P_2 \mid s[1, \varepsilon]]) \\
&\longrightarrow (\nu s)([\check{s}_?(x_1).\check{s}_?(x_2).P_0 \mid \check{s}[0, (0, \text{hbt}_2), (1, \text{hbt}_2)]] \\
&\parallel [s_!(\text{hbt}_1).s_!(\text{hbt}_1).P_1 \mid s[0, \varepsilon]] \parallel [P_2 \mid s[2, \varepsilon]]) \\
&\longrightarrow (\nu s)([\check{s}_?(x_1).\check{s}_?(x_2).P_0 \mid \check{s}[0, (0, \text{hbt}_2), (1, \text{hbt}_2), (0, \text{hbt}_1)]] \\
&\parallel [s_!(\text{hbt}_1).P_1 \mid s[1, \varepsilon]] \parallel [P_2 \mid s[2, \varepsilon]]) \\
&= \text{Heartbeat}_1
\end{aligned}$$

The first interaction is an instance of rule [CONN] and has lead to the establishment of new session involving all the network nodes. The second interaction uses an instance of the rule [UCAST] where network node $[s_!(\text{hbt}_2).s_!(\text{hbt}_2).P_2 \mid s[0, \varepsilon]]$ sends a heartbeat message to the \check{s} -endpoint network node. Observe that the session state of the sender node has increased by one. The third interaction is also an instance of rule [UCAST] and yet another heartbeat message from the same network node is unicast to the \check{s} -endpoint network node. The node can unicast the heartbeat message even if it is in a later session state. This is because messages are gathered at each session state. In the last interaction, network node $[s_!(\text{hbt}_1).s_!(\text{hbt}_1).P_1 \mid s[0, \varepsilon]]$ also unicasts a heartbeat message. Consider then that the \check{s} -endpoint gathers the heartbeat messages:

$$\begin{aligned}
\text{Heartbeat}_1 &\longrightarrow (\nu s)([\check{s}_?(x_2).P_0\{\text{hbt}_1 \odot \text{hbt}_2/x_1\} \mid \check{s}[1, (1, \text{hbt}_2)]] \\
&\parallel [s_!(\text{hbt}_1).P_1 \mid s[1, \varepsilon]] \parallel [P_2 \mid s[2, \varepsilon]]) \\
&\longrightarrow (\nu s)([P_0\{\text{hbt}_1 \odot \text{hbt}_2/x_1\}\{\text{hbt}_2/x_2\} \mid \check{s}[2, \varepsilon]] \\
&\parallel [s_!(\text{hbt}_1).P_1 \mid s[1, \varepsilon]] \parallel [P_2 \mid s[2, \varepsilon]]) \\
&= \text{Heartbeat}_2
\end{aligned}$$

Two instances of the rule [GTHR] allow for the \check{s} -endpoint to consume all the messages from the \check{s} -buffer. The first interaction gathers all heartbeat messages tagged with state 0, whereas the second interaction gathers all the heartbeat messages tagged with state 1. Each interaction updates the state counter of the \check{s} -endpoint. After the last two interactions the s -endpoint $[s_!(\text{hbt}_1).P_1 \mid s[1, \varepsilon]]$ is found in a non-synchronised state. However an instance of rule [LOSS] can be observed that will re-synchronise the non-synchronised s -endpoint:

$$\begin{aligned}
\text{Heartbeat}_2 &\longrightarrow (\nu s)([P_0\{\text{hbt}_1 \odot \text{hbt}_2/x_1\}\{\text{hbt}_2/x_2\} \mid \check{s}[2, \varepsilon]] \\
&\parallel [P_1 \mid s[2, \varepsilon]] \parallel [P_2 \mid s[2, \varepsilon]])
\end{aligned}$$

□

A second example demonstrates a typical pattern of interaction found ad-hoc and sensor networks when running consensus protocols. In this example, rules [TRUE]/[FALSE] are used to drop connections.

Example 2.2 (Dropping Connections). It is typical in consensus algorithms for a node to establish multiple connections with other nodes but maintaining active only the one with the highest id number, e.g. the Paxos consensus algorithm [Lam98, L⁺01] (also see Section 5).

For example, consider network:

$$\begin{aligned} \text{Network} = & \quad [a_1(w).\check{w}_1\langle id_2 \rangle.P_0] \parallel (\nu s)([\check{s}_1\langle id_1 \rangle.P_1 \mid \check{s}[0, \varepsilon]]) \\ & \parallel [s_?(x).(P + a_?(w).w_?(y).\text{if } x > y \text{ then } P \text{ else } P') \mid s[0, \varepsilon]] \end{aligned}$$

with $s \notin \text{fn}(P')$. Node $[s_?(x).(P + a_?(w).w_?(y).\text{if } x > y \text{ then } P \text{ else } P') \mid s[0, \varepsilon]]$ has already established a connection on channel s and awaits for a connection id number; a [BCAST] and a [RCV] interaction will result in:

$$\begin{aligned} \text{Network} \longrightarrow & \quad [a_1(w).\check{w}_1\langle id_2 \rangle.P_0] \parallel (\nu s)([P_1 \mid \check{s}[1, \varepsilon]]) \\ & \parallel [P + a_?(w).w_?(y).\text{if } id_1 > y \text{ then } P \text{ else } P' \mid s[1, \varepsilon]] \\ = & \quad \text{Network}_1 \end{aligned}$$

The resulting network has the option to either continue interaction on the s -endpoint via process P , or establish a new connection on shared channel a . Assume that the latter interaction takes place:

$$\begin{aligned} \text{Network}_1 \longrightarrow & \quad (\nu s, s')([\check{s}'_1\langle id_2 \rangle.P_0 \mid \check{s}[0, \varepsilon]] \parallel [P_1 \mid \check{s}[1, \varepsilon]]) \\ & \parallel [s'_?(y).\text{if } id_1 > y \text{ then } P \text{ else } P' \mid s[1, \varepsilon] \mid s'[0, \varepsilon]] \\ = & \quad \text{Network}_2 \end{aligned}$$

The s' -endpoint also receives a connection id number, via reduction rules [BCAST] and [RCV]:

$$\begin{aligned} \text{Network}_2 \longrightarrow & \quad (\nu s, s')([P_0 \mid \check{s}[1, \varepsilon]] \parallel [P_1 \mid \check{s}[1, \varepsilon]]) \\ & \parallel [\text{if } id_1 > id_2 \text{ then } P \text{ else } P' \mid s[1, \varepsilon] \mid s'[1, \varepsilon]] \\ = & \quad \text{Network}_3 \end{aligned}$$

The receiving node will then compare the two connection id numbers and decide to drop the session with the smallest corresponding id number and continue with the corresponding process; it will continue the interaction on the s -endpoint in process P in the case where the newest connection is dropped, or, otherwise, it will proceed with the interaction on the s' -endpoint via process P' . For example, if $id_2 > id_1$ the network will reduce using an instance of rule [FALSE] and result as in:

$$\text{Network}_3 \longrightarrow (\nu s)([P_1 \mid \check{s}[1, \varepsilon]]) \parallel (\nu s')([P_0 \mid \check{s}[1, \varepsilon]] \parallel [P' \mid s'[1, \varepsilon]]) \quad \square$$

The next example demonstrates a simple recursive behaviour, where a \check{s} -endpoint recurses or terminates an session interaction based on the acknowledgements it receives by the corresponding s -endpoints.

Example 2.3 (Recursive Interaction). The following presents an example where the \check{s} -endpoint sends a message and then requires from the corresponding s -endpoints to reply with an acknowledgement. By inspecting the acknowledgements, the \check{s} -endpoint decides to terminate the interaction or to reiterate. Consider the network:

$$\begin{aligned}
P &= \text{def} \\
&\quad Q(w) \stackrel{\text{def}}{=} \check{w}! \langle v \rangle . \check{w}?(\{x_i\}_{i \in I}). \\
&\quad \quad \text{if cond}(\{x_i\}_{i \in I}) \text{ then } \check{w} \triangleleft \text{accept} . \mathbf{0} \\
&\quad \quad \text{else } \check{w} \triangleleft \text{restart} . P \langle w \rangle \\
&\quad \text{in } Q \langle s \rangle \\
P_k &= \text{def } R_k(w) \stackrel{\text{def}}{=} w?(x) . w! \langle \text{ack}_k \rangle . w \triangleright \{ \text{accept} : \mathbf{0}, \text{restart} : Q \langle w \rangle, \text{df} : \mathbf{0} \} \\
&\quad \text{in } R_k \langle s \rangle \\
\text{Recursive} &= [P \mid \check{s}[0, \varepsilon]] \parallel \prod_{J \in J} [P_J \mid s[0, \varepsilon]]
\end{aligned}$$

where operation \odot aggregates all acknowledgements as a set $\{\text{ack}_i\}_{i \in I}$. A $[\text{BCAST}]$ followed by a series of $[\text{RCV}]$ and $[\text{UCAST}]$ interactions, and finally a $[\text{GTHR}]$ operation, results in (the recursive process definition is omitted):

$$\begin{aligned}
\text{Recursive} &\longrightarrow^* \text{Recursive}_1 \\
&= [\text{if cond}(\{x_i\}_{i \in I}) \text{ then } \check{s} \triangleleft \text{accept} . \mathbf{0} \text{ else } \check{s} \triangleleft \text{restart} . Q \langle s \rangle \mid \check{s}[2, \varepsilon]] \\
&\parallel \prod_{k \in K_1} [s \triangleright \{ \text{accept} : \mathbf{0}, \text{restart} : R_k \langle s \rangle, \text{df} : \mathbf{0} \} \mid s[2, \varepsilon]] \\
&\parallel \prod_{k \in K_2} [P_k \mid s[0, \varepsilon]]
\end{aligned}$$

with $K_1 \cup K_2 = J$ and $K_1 \cap K_2 = \emptyset$. If the condition on $\{\text{ack}\}_{i \in I}$ is true then after a $[\text{SEL}]$ on label accept followed by a series of $[\text{BRA}]$ interactions the network may result in:

$$\text{Recursive}_1 \longrightarrow^* [\mathbf{0} \mid \check{s}[3, \varepsilon]] \parallel \prod_{k \in K_1} [\mathbf{0} \mid s[3, \varepsilon]] \parallel \prod_{k \in K_2} [Q_k \mid s[0, \varepsilon]]$$

Note that each node in network $\prod_{k \in K_2} [Q \mid s[0, \varepsilon]]$ can terminate following a sequence of $[\text{REC}]$, $[\text{LOSS}]$ and $[\text{BREC}]$ interactions. If the condition on $\{\text{ack}\}_{i \in I}$ is false then after a $[\text{SEL}]$ on label restart followed by a series of $[\text{BRA}]$ interactions the network may result in:

$$\text{Recursive}_1 \longrightarrow^* [P \mid \check{s}[3, \varepsilon]] \parallel \prod_{k \in K_1} [Q \mid s[3, \varepsilon]] \parallel \prod_{k \in K_2} [Q_k \mid s[0, \varepsilon]]$$

where the network reiterates. Note that subnetwork $\prod_{k \in K_2} [Q \mid s[0, \varepsilon]]$ will not participate in the next iteration, since it can only terminate its interaction by recovering following a sequence of $[\text{REC}]$, $[\text{LOSS}]$ and $[\text{BREC}]$ interactions. \square

3. SESSION TYPES

We now introduce the session type system for the UBSC. The type system combines ideas from [HKP⁺10, Kou12] to type buffer terms. A novel notion is the notion of endpoint synchronisation used to cope with the presence of non-synchronised session endpoints and with recovery semantics.

3.1. Type Syntax. The session types syntax follows the standard binary session types syntax as introduced by Honda et al. [HVK98]. However, we have no session delegation since channel delegation is not found in systems with unreliable communication such as sensor networks.

Definition 3.1 (Session Type). Let \mathcal{B} be a set of *base types* ranged over by β . Session types are inductively defined by the following grammar:

$$T ::= !\beta.T \mid ?\beta.T \mid \oplus \{\ell_i : T_i\}_{i \in I} \mid \& \{\ell_i : T_i\}_{i \in I} \mid \text{end} \mid t \mid \mu t.T$$

Type $!_1\beta.T$ describes the sending of a value of type β and then proceeding with type T . Type $?_1\beta.T$ describes the reception of a value with type β and then proceeding with type T . Type $\oplus\{\ell_i : T_i\}_{i \in I}$ selects a label from the set of labels $\{\ell_i\}_{i \in I}$ and then proceeds with the corresponding type $\{T_i\}_{i \in I}$. Dually, type $\&\{\ell_i : T_i\}_{i \in I}$ branches on the set of labels $\{\ell_i\}_{i \in I}$ and then proceeds with the corresponding type $\{T_i\}_{i \in I}$. Type **end** is the inactive type, whereas t is the recursive variable. Finally, type $\mu t.T$ is the recursive type, which binds free occurrences of t in T . We define a capture avoiding substitution on types $T\{T'/t\}$ in the usual way. We assume equi-recursive types, $\mu t.T = T\{\mu t.T/t\}$.

Our calculus does not incorporate session delegation. Session delegation is rather unnatural in an unreliable setting supporting broadcasting and gather semantics, and sharing of channel resources. Moreover, session delegation is valid only for \checkmark -endpoints, which are linear, and it would require additional syntax and reduction semantics beyond broadcast and gather operations.

The duality operator also follows the standard binary session type definition [HVK98]. A simple inductive definition is enough to capture session endpoint duality because of the lack of delegation [BDGK14].

Definition 3.2 (Type Duality).

$$\begin{aligned} \overline{\text{end}} &= \text{end} & \overline{!_1\beta.T} &= ?_1\beta.\overline{T} & \overline{?_1\beta.T} &= !_1\beta.\overline{T} & \overline{t} &= t & \overline{\mu t.T} &= \mu t.\overline{T} \\ \overline{\oplus\{\ell_i : T_i\}_{i \in I}} &= \&\{\ell_i : \overline{T}_i\}_{i \in I} & \overline{\&\{\ell_i : T_i\}_{i \in I}} &= \oplus\{\ell_i : \overline{T}_i\}_{i \in I} \end{aligned}$$

Two types T_1 and T_2 are dual if $\overline{T}_1 = T_2$. Note $\overline{\overline{T}} = T$ for any T . Next, we define a buffer type syntax used to type message buffers (cf. [HKP⁺10, Kou12]).

Definition 3.3 (Buffer Types). Buffer types are inductively defined by the following grammar:

$$M ::= \varepsilon \mid !_1\beta.M \mid \oplus\ell.M$$

Buffer types are used to type buffer terms; they describe the types of the values, $!_1\beta.M$, or labels, $\oplus\ell.M$, in a session buffer term.

The \circ operator is used to combine session types and buffer types (cf. [HKP⁺10, Kou12]).

Definition 3.4 (Operator \circ).

$$\begin{aligned} ?_1\beta.T \circ !_1\beta.M &= T \circ M & \&\{\ell_i : T_i\}_{i \in I} \circ \oplus\ell_k.M &= T_k \circ M \\ \mu t.T \circ M &= T\{\mu t.T/t\} \circ M & T \circ \varepsilon &= T \end{aligned}$$

The \circ operator combines a session type with a buffer type and returns a session type (cf. [HKP⁺10, Kou12]). It works inductively by removing the prefix from the input session when the prefix of the buffer type is dual. The intuition for the \circ operator considers that messages within a session buffer have already been received from a node (see typing rule [T_{NODE}] in Figure 8), so the operator consumes a buffer type against a session type.

3.2. Typing System. We define the typing contexts used by the typing system.

Definition 3.5 (Typing Context). We define Γ , Δ , and Θ typing contexts:

$$\begin{aligned} \Gamma &::= \emptyset \mid \Gamma, x : \beta \mid \Gamma, a : T \mid \Gamma, D(\tilde{x}) : (\Gamma; \Delta) \\ \Delta &::= \emptyset \mid \Delta, k : T \mid \kappa : (c, T) \\ \Theta &::= \emptyset \mid \Theta, \kappa : (c, M) \end{aligned}$$

$$\begin{array}{c}
\Gamma; \emptyset \vdash \varepsilon \quad [\text{BEMP}] \quad \Gamma; \Theta, \kappa : (c, \varepsilon) \vdash \kappa[c, \varepsilon] \quad [\text{SEMP}] \quad \frac{\{\Gamma; \Theta_i \vdash B_i\}_{i \in \{1,2\}}}{\Gamma; \Theta_1, \Theta_2 \vdash B_1 \mid B_2} \quad [\text{BPAR}] \\
\\
\frac{\Gamma; \Theta, (c, M) \vdash s[c, \tilde{m}] \quad \Gamma \vdash e : \beta}{\Gamma; \Theta, s : (c, M.\!:\!\beta) \vdash s[c, \tilde{m} \cdot e]} \quad [\text{SEXP}] \quad \frac{\Gamma; \Theta, (c, M) \vdash s[c, \tilde{m}]}{\Gamma; \Theta, s : (c, M.\oplus\ell) \vdash s[c, \tilde{m} \cdot \ell]} \quad [\text{SLAB}] \\
\\
\frac{\tilde{h}' = \mathbf{B}(\tilde{h}, c) \quad \Gamma \vdash \mathbf{V}(\tilde{h}, c) : \beta \quad \Gamma; \Theta, \check{s} : (c, M) \vdash \check{s}[c', \tilde{h}']}{\Gamma; \Theta, \check{s} : (c+1, M.\!:\!\beta) \vdash \check{s}[c', \tilde{h}]} \quad [\text{LEXP}]
\end{array}$$

Figure 6: Typing rules for session buffers

Context Γ is called shared context and maps variables to ground types, shared names to session types, and recursive variables to typing contexts $(\Gamma; \Delta)$. Context Δ is called linear context and maps session names and session variables to session types, and session names to tuples, (c, T) , that combine an integer value, c , that corresponds to session state together with a session type, T . Similarly, context Θ is called buffer context and maps session names to tuples, (c, M) , that combine an integer value, c , that corresponds to session state together with a buffer type, M .

Contexts are treated as sets. We write Γ, Γ' to denote the union of contexts Γ and Γ' . We also write Δ, Δ' , resp. Θ, Θ' , for the disjoint union of contexts Δ and Δ' , resp. Θ and Θ' . We define the domain of $\text{dom}(\cdot)$ of contexts Γ , Δ , and Θ in the expected way.

The \circ operator is lifted to combine contexts Δ and Θ :

$$\Delta \circ \Theta = \{\kappa : (c, T \circ M) \mid \kappa : T \in \Delta \wedge \kappa : (c, M) \in \Theta\}$$

The result of the \circ operator on contexts is a new linear context.

The typing judgement for expressions is defined as the least relation $\Gamma \vdash e : \beta$ that respects: i) the types of value and unit expressions (e.g., $\text{nat}, \text{bool} \in B$); and ii) the types of variables in expressions, e.g. whenever $x \in \text{fv}(e)$, then $\Gamma \vdash e$ is defined if $x : \beta \in \Gamma$ for some β .

The typing judgement for session buffers, $\Gamma; \Theta \vdash B$, is defined as the least relation that satisfies the rules in Fig. 6. Rule $[\text{BEMP}]$ types the empty buffer term with an empty Θ context. Rule $[\text{SEMP}]$ types an empty κ -endpoint buffer. It maps, in context Θ , the κ -endpoint to the state of the κ -buffer together with the empty buffer type. Rule $[\text{BPAR}]$ types a parallel composition of session buffer terms, with the disjoint union of the two respective Θ environments. Buffers for s -endpoints are typed with rules $[\text{SEXP}]$ and $[\text{SLAB}]$; the type of an expression (respectively, label) is appended at the end of the buffer type of s within the Θ context.

Rule $[\text{LEXP}]$ types the \check{s} -endpoints buffers. The rule works by reconstructing the messages within the \check{s} -buffer using operation $\tilde{h}' = \mathbf{B}(\tilde{h}, c)$, with h' appearing in the premise, and h appearing in the conclusion. Rule $[\text{LEXP}]$ appends to the type of \check{s} within context Θ , the type of value $\mathbf{V}(\tilde{h}, c)$, which is \ominus -composition of all expressions associated with session state c , that appears at the type of the premise. The session state is updated at the conclusion.

The typing judgement for processes, $\Gamma; \Delta \vdash P$, is defined as the least relation that satisfies the rules in Fig. 7. Rule $[\text{TINACT}]$ is standard for typing the inactive process. Rule $[\text{SWK}]$ defines standard session type weakening for context Δ . The next three rules are

$$\begin{array}{c}
\Gamma; \emptyset \vdash \mathbf{0} \quad [\text{TINACT}] \quad \frac{\Gamma; \Delta \vdash P}{\Gamma; \Delta, k : \text{end} \vdash P} \quad [\text{SWK}] \quad \frac{\Gamma, a : T; \Delta, \check{x} : \bar{T} \vdash P}{\Gamma, a : T; \Delta \vdash a_1(\check{x}).P} \quad [\text{TREQ}] \\
\\
\frac{\Gamma, a : T; \Delta, x : T \vdash P}{\Gamma, a : T; \Delta \vdash a?(x).P} \quad [\text{TACC}] \quad \frac{\Gamma; \Delta, k : T \vdash P \quad \Gamma \vdash e : \beta}{\Gamma; \Delta, k : !\beta.T \vdash k_1\langle e \rangle.P} \quad [\text{TSND}] \\
\\
\frac{\Gamma \vdash e : \beta \quad \Gamma, x : \beta; \Delta, k : T \vdash P}{\Gamma; \Delta, k : ?\beta.T \vdash k?(x)\langle e \rangle.P} \quad [\text{TRCV}] \quad \frac{\Gamma; \Delta \vdash P_i \quad i \in \{1, 2\}}{\Gamma; \Delta \vdash P_1 + P_2} \quad [\text{TSUM}] \\
\\
\frac{\Gamma; \Delta, k : T_j \vdash P \quad j \in I \quad (\exists s \in \mathcal{S}, k = \check{s}) \vee (\exists x \in \mathcal{V}, k = \check{x})}{\Gamma; \Delta, k : \oplus\{\ell_i : T_i\}_{i \in I} \vdash k \triangleleft \ell_j.P} \quad [\text{TSEL}] \\
\\
\frac{\forall i \in I, \Gamma; \Delta, \Delta', k : T_i \vdash P_i \quad (\exists s \in \mathcal{S}, k = s) \vee (\exists x \in \mathcal{V}, k = x) \quad \Gamma; \Delta' \vdash R \quad \forall s \in \mathcal{S}, \check{s} \notin \text{dom}(\Delta)}{\Gamma; \Delta, \Delta', k : \&\{\ell_i : T_i\}_{i \in I} \vdash k \triangleright \{\ell_i : P_i, \text{df} : R\}_{i \in I}} \quad [\text{TBR}] \\
\\
\frac{\{\Gamma; \Delta, \Delta_i \vdash P_i \quad \forall s \in \mathcal{S}, \check{s} \notin \text{dom}(\Delta_i)\}_{i \in \{1, 2\}}}{\Gamma; \Delta, \Delta_1, \Delta_2 \vdash \text{if } \varphi \text{ then } P_1 \text{ else } P_2} \quad [\text{TCOND}] \\
\\
\frac{\Gamma' \vdash u_i : \beta_i \text{ iff } u_i \in \tilde{v} \quad \tilde{y} = \text{dom}(\Delta') \quad k_i : T_i \in \Delta \text{ iff } y_i : T_i \in \Delta'}{\Gamma, D(\tilde{x}, \tilde{y}) : (\Gamma'; \Delta'); \Delta \vdash D\langle \tilde{v}, \tilde{k} \rangle} \quad [\text{TVAR}] \\
\\
\frac{\forall i \in I, D_i(\tilde{x}_i) : (\Gamma_i, \Delta_i) \in \Gamma \wedge \Gamma, \Gamma_i; \Delta_i \vdash P_i \quad \Gamma; \Delta \vdash P}{\Gamma; \Delta \vdash \text{def } \{D_i(\tilde{x}_i) \stackrel{\text{def}}{=} P_i\}_{i \in I} \text{ in } P} \quad [\text{TREC}]
\end{array}$$

Figure 7: Typing rules for processes

standard session type rules (cf. [HVK98]) for typing session request, rule [TREQ], session accept, rule [TACC], and session send, rule [TSND], respectively.

Rule [TRCV] is used to type the (RCV) process, which requires for the default expression e to have the same type as variable x . Rule [TSUM] types non-deterministic behaviour. The rule expects the same type for each process combined by the [TSUM] operator.

The next two rules are standard rules for typing session select, rule [TSEL], and session branch, rule [TBR], prefixes, respectively. The extra conditions on rules [TSEL] and [TBR] ensure correctness of the selection/branching interaction as an interaction between \check{s} -endpoint and s -endpoint, respectively; only \check{s} -endpoints and \check{x} variables are typed by rule [TSEL], and similarly only s -endpoints and x variables are being typed by rule [TBR]. The recovery process in the branch prefixed process is used to drop session endpoints including the s -endpoint, thus typing rule [TBR] ensures that the recovery process R is typed with a subset of the linear context, excluding the s -endpoint. Moreover, the rule ensures that only s -endpoints are be dropped.

$$\begin{array}{c}
\frac{\Gamma; \Theta \vdash B \quad \Gamma; \Delta \vdash P \quad \text{dom}(\Delta) = \text{dom}(\Theta)}{\Gamma; \Delta \circ \Theta \vdash [P | B]} \text{[TNODE]} \\
\\
\frac{\Gamma, a : T; \Delta \vdash N}{\Gamma; \Delta \vdash (\nu a)N} \text{[TCRES]} \quad \frac{\{\Gamma; \Delta, \Delta_i \vdash N_i\}_{i \in \{1,2\}} \quad \forall s \in \mathcal{S}, \check{s} \notin \text{dom}(\Delta)}{\Gamma; \Delta, \Delta_1, \Delta_2 \vdash N_1 \parallel N_2} \text{[TPAR]} \\
\\
\frac{\Gamma; \Delta, \Delta', \check{s} : (c, \bar{T}) \vdash N \quad \Delta' = s : (c, T) \vee \Delta' = \emptyset \quad s \notin \text{dom}(\Delta)}{\Gamma; \Delta \vdash (\nu s)N} \text{[TSRES]}
\end{array}$$

Figure 8: Typing rules for networks

A novel rule is [TCOND], that types the conditional process; since the conditional process is used to drop session endpoints, the typing rule splits the contexts Δ_i of the two branches into common and non-common, i.e. dropped, session endpoints. The rule also ensures that only s -endpoints can be dropped.

Rules [TVAR], and [TREC] are standard. Rule [TVAR] requires that recursive variables are mapped in the Γ context and checks whether the arguments application has the correct type with respect to context Γ . Similarly rule [TREC] requires to checks context Γ and the recursive definition for type consistency.

The typing judgement for networks, $\Gamma; \Delta \vdash N$, is defined as the least relation that satisfies the rules in Fig. 8. Rule [TNODE] type a network node. The rule combines the contexts of the process, P , and the session buffer terms, B using operator \circ . Rule [TNODE] requires that $\text{dom}(\Delta) = \text{dom}(\Theta)$ implying that: i) all the session names that appear in the network process have their corresponding session buffer present in the network node; and ii) process P does not contain free session variables, i.e. $x : T \notin \Delta$.

Rule [TPAR] requires that parallel components of a network should share the same type for common s -endpoints. The next rule is [TCRES], which restricts shared name a by removing it from Γ . Finally, rule [TSRES] captures the interaction intuition between the \check{s} -endpoint and the s -endpoints. A session name s can be restricted whenever its two endpoints are dual when in the same state c , or if only the \check{s} -endpoint is present in the typing context. The latter condition captures the case where the s -endpoints were not created due to message loss or lost due to recovery.

The following example gives the typing for the Recursive network in Example 2.3.

Example 3.6 (Typing for Example 2.3). Given Γ such that $\Gamma \vdash v : \beta$ and for all $i \in I, \Gamma \vdash \text{ack}_i : \text{ackType}$, and moreover $T = \mu t. \lambda \beta. ?\text{ackType}. \oplus \{\text{accept} : \text{end}, \text{restart} : t\}$, the reader can verify that $\Gamma; \check{s} : T, s : \bar{T} \vdash \text{Recursive}$. \square

Typing System for Runtime Processes. The typing system presented so far can only type networks with s -endpoints that are in the same state with the corresponding \check{s} -endpoint. This is demonstrated with the following example.

Example 3.7 (Heartbeat at Runtime). Consider an instance of Example 1.1 with two s -endpoints:

$$\text{Heartbeat} = [\check{s}_i \langle \text{hbt} \rangle . P_0 \mid \check{s}[0, \varepsilon]] \parallel [s_?(x) . P_1 \mid s[0, \varepsilon]] \parallel [s_?(x) . P_2 \mid s[0, \varepsilon]]$$

All the endpoints in the above network are in the same state. Assuming that value `hbt` has type `b`, the Heartbeat network can be typed as

$$\text{hbt} : \mathbf{b}; \check{s} : (0, \mathbf{!b.end}), s : (0, \mathbf{?b.end}) \vdash \text{Heartbeat}$$

Following reduction rule [BCAST] we can observe an unreliable broadcast operation that results in

$$\text{Heartbeat} \longrightarrow \text{Heartbeat}' = [P_0 \mid \check{s}[1, \varepsilon]] \parallel [s_{?}(x).P_1 \mid s[1, \text{hbt}]] \parallel [s_{?}(x).P_2 \mid s[0, \varepsilon]]$$

The typing system presented so far cannot type the networks that may arise at runtime where some s -endpoints are not synchronised with the corresponding \check{s} -endpoint. In network `Heartbeat'`, the second network node is typed as $\text{hbt} : \mathbf{b}; s : (1, \text{end}) \vdash [s_{?}(x).P_1 \mid s[1, \text{hbt}]]$ and the third network node is typed as $\text{hbt} : \mathbf{b}; s : (0, \mathbf{?b.end}) \vdash [s_{?}(x).P_2 \mid s[0, \varepsilon]]$. Thus, we cannot apply rule [TPAR] that requires the same type for the two s -endpoints.

We can achieve typing by using the session state information to construct the type information that was lost due to unreliable communication. The key is to use a typing rule to synchronise the s -endpoints that are not synchronised with the \check{s} -endpoint.

We define type advancement as a transition relation on types. The type advancement relation is used to define the notion of endpoint synchronisation.

Definition 3.8 (Type Advancement). Relation $T \rightarrow^n T'$ is defined as:

$$\begin{array}{c} \mathbf{?}\beta.T \rightarrow^1 T \qquad \mathbf{!}\beta.T \rightarrow^1 T \qquad \frac{k \in I}{\oplus\{\ell_i : T_i\}_{i \in I} \rightarrow^1 T_k} \qquad \frac{k \in I}{\&\{\ell_i : T_i\}_{i \in I} \rightarrow^1 T_k} \\ \\ T \rightarrow^0 T \qquad \frac{T \rightarrow^n T'' \quad T'' \rightarrow^1 T'}{T \rightarrow^{n+1} T'} \end{array}$$

It is useful to distinguish output advance by writing $\mathbf{!}\beta.T \rightarrow_o T$, extended to $T \rightarrow_o^n T$ in the standard way. Intuitively, $T \rightarrow^n T'$ says that T' is reached in n advancements from T . Type advancement for recursive types is obtained by expansion.

Next, we use type advancement (Definition 3.8) and the session state information within a linear context to define the linear context synchronisation relation over linear contexts.

Definition 3.9 (Linear Context Synchronisation). We define the relation $\Delta \leftrightarrow \Delta'$ inductively as follows.

$$\emptyset \leftrightarrow \emptyset \qquad \frac{\Delta \leftrightarrow \Delta' \quad T \rightarrow^{m-n} T' \quad n \leq m}{\Delta, s : (n, T) \leftrightarrow \Delta', s : (m, T')} \qquad \frac{\Delta \leftrightarrow \Delta' \quad T' \rightarrow_o^{n-m} T \quad n \geq m}{\Delta, s : (n, T) \leftrightarrow \Delta', s : (m, T')}$$

Linear context synchronisation is crucial for the correctness of the type system and it is used for achieving s -endpoint synchronisation. The middle rule states that the type of an s -endpoint can use type advancement to proceed its state within a linear context. The right rule describes the opposite case where the type of an s -endpoint uses output type advancement, to reconstruct a previous session state. The requirement for output type advancement is aligned with the fact that a non-synchronised s -endpoint may use reduction rule [Loss], that drops output session prefixes. Finally the left rule is a basic case rule for the linear context.

Figure 9 defines rule [TSYNCH]. Rule [TSYNCH], is used to provide static guarantees of type duality between all non-synchronised sessions endpoints that may arise during execution. The rule uses linear context synchronisation, $\Delta \leftrightarrow \Delta'$, to synchronise non-synchronised s -endpoints.

$$\frac{\Gamma; \Delta' \vdash N \quad \Delta' \hookrightarrow \Delta}{\Gamma; \Delta \vdash N} \text{ [TSYNCH]}$$

Figure 9: Runtime Typing Rule

Rule [TSYNCH] types only networks that result at runtime, where due to unreliability, session endpoints become unsynchronised, and it is important to prove the properties of progress (Theorem 4.12) and recovery (Theorem 4.15) in Section 4. The main use of the rule is to align non-synchronised s -endpoints in a linear context in order to apply typing rule [TPAR], as well as enforcing the notion of duality between \check{s} -endpoint and s -endpoint as in typing rule [TSRES]. Inspecting rule [TSYNCH] from the point of view of single network node it seems that [TSYNCH] adds a degree of non-determinism during runtime typing. However, at the network level, the state counter of the \check{s} -endpoint can deterministically guide the computation of the linear context synchronisation for the corresponding s -endpoints within a network, i.e., an implementation of runtime checking will seek, during linear context synchronisation to align the counter of the s -endpoints with the counter of the \check{s} -endpoint.

Example 3.10 (Typing Heartbeat at Runtime). For example we can use [TSYNCH] to synchronise the type of network node $[s?(x).P_2 \mid s[0, \varepsilon]]$ in Example (3.7).

$$\frac{\text{hbt} : \mathbf{b}; s : (0, ?\mathbf{b}.\text{end}) \vdash [s?(x).P_2 \mid s[0, \varepsilon]] \quad s : (0, ?\mathbf{b}.\text{end}) \hookrightarrow s : (1, \text{end})}{\text{hbt} : \mathbf{b}; s : (1, \text{end}) \vdash [s?(x).P_2 \mid s[0, \varepsilon]]} \text{ [TSYNCH]}$$

This leads to typing judgement

$$\text{hbt} : \mathbf{b}; \check{s} : (1, \text{end}), s : (1, \text{end}) \vdash \text{Heartbeat}'$$

A network N , resp. process P , is called well-typed, whenever $\Gamma; \Delta \vdash N$, resp. $\Gamma; \Delta \vdash P$, for some contexts Γ and Δ .

3.3. Typing Derivation Example. We present the typing derivation for the Heartbeat_1 network in Example 2.1. The example demonstrates linear context synchronisation which is a main notion of the typing system.

Example 3.11 (Typing derivation for the Heartbeat Protocol). Recall, network Heartbeat_1 in Example 2.1:

$$\begin{aligned} \text{Heartbeat}_1 &= (\nu s)([\check{s}?(x).\check{s}?(x).P_0 \mid \check{s}[0, (0, \text{hbt}_2), (1, \text{hbt}_2), (0, \text{hbt}_1)]) \\ &\parallel [s_1\langle \text{hbt}_1 \rangle.P_1 \mid s[1, \varepsilon]] \parallel [P_2 \mid s[2, \varepsilon]] \end{aligned}$$

Also, consider that the heartbeat message has type \mathbf{b} . We first type network node

$$[\check{s}?(x).\check{s}?(x).P_0 \mid \check{s}[0, (0, \text{hbt}_2), (1, \text{hbt}_2), (0, \text{hbt}_1)]]$$

We give the type derivation for buffer $\check{s}[0, (0, \text{hbt}_2), (1, \text{hbt}_2), (0, \text{hbt}_1)]$.

$$\boxed{
 \begin{array}{c}
 \Gamma; \check{s} : (0, \varepsilon) \vdash \check{s}[0, \varepsilon] \text{ [SEMP]} \quad \varepsilon = \mathbf{B}((0, \text{hbt}_2), (0, \text{hbt}_1), 0) \\
 \Gamma \vdash \mathbf{V}((0, \text{hbt}_2), (0, \text{hbt}_1), 0) : \mathbf{b} \\
 \hline
 \Gamma; \check{s} : (1, \mathbf{!b}) \vdash \check{s}[0, (0, \text{hbt}_2), (0, \text{hbt}_1)] \text{ [LEXP]} \\
 \hline
 (0, \text{hbt}_2), (0, \text{hbt}_1) = \mathbf{B}((0, \text{hbt}_2), (1, \text{hbt}_2), (0, \text{hbt}_1), 1) \\
 \Gamma \vdash \mathbf{V}((0, \text{hbt}_2), (1, \text{hbt}_2), (0, \text{hbt}_1), 1) : \mathbf{b} \\
 \hline
 \Gamma; \check{s} : (2, \mathbf{!b.!\mathbf{b}}) \vdash \check{s}[0, (0, \text{hbt}_2), (1, \text{hbt}_2), (0, \text{hbt}_1)] \text{ [LEXP]}
 \end{array}
 } \quad (3.1)$$

The typing for the buffer constructs a derivation using typing rules: i) [SEMP] to initially type the empty buffer; ii) [LEXP] to buffer $\check{s}[0, (0, \text{hbt}_2), (0, \text{hbt}_1)]$ that contains only the heartbeat messages send at session state 0 of the protocol; and iii) [LEXP] again to type the entire buffer. The instances of the [LEXP] typing rule make use of operations $\mathbf{B}(\cdot, c)$ and $\mathbf{V}(\cdot, c)$ to type the messages at each session state. The network node is then typed using an instance of rule [TNode]:

$$\boxed{
 \begin{array}{c}
 \Gamma; \check{s} : \overline{T} \vdash P_0 \\
 \hline
 \Gamma; \check{s} : \mathbf{?b.?\mathbf{b}.\overline{T}} \vdash \check{s}?(x).\check{s}?(x).P_0 \text{ [TRCV]} \\
 \hline
 \Gamma; \check{s} : (2, \mathbf{!b.!\mathbf{b}}) \vdash \check{s}[0, (0, \text{hbt}_2), (1, \text{hbt}_2), (0, \text{hbt}_1)] \quad 3.1 \\
 \check{s} : \mathbf{?b.?\mathbf{b}.\overline{T}} \circ \check{s} : (2, \mathbf{!b.!\mathbf{b}}) = \check{s} : (2, \overline{T}) \\
 \hline
 \Gamma; \check{s} : (2, \overline{T}) \vdash [\check{s}?(x).\check{s}?(x).P_0 \mid \check{s}[0, (0, \text{hbt}_2), (1, \text{hbt}_2), (0, \text{hbt}_1)]] \text{ [TNode]}
 \end{array}
 } \quad (3.2)$$

The rule uses the \circ operator to combine tuple $(2, \mathbf{!b.!\mathbf{b}})$, derived from typing derivation 3.1, and session type $\mathbf{?b.?\mathbf{b}.\overline{T}}$, derived from typing rule [TRCV]. The result is tuple $(2, \overline{T})$, denoting that the \check{s} -endpoint is in state 2 and has type \overline{T} .

We continue by typing network node $[P_2 \mid s[2, \varepsilon]]$ using an instance of typing rule [TNode]:

$$\boxed{
 \begin{array}{c}
 \Gamma; s : (2, \varepsilon) \vdash s[2, \varepsilon] \text{ [SEMP]} \quad \Gamma; s : T \vdash P_2 \quad s : T \circ s : (2, \varepsilon) = s : (2, T) \\
 \hline
 \Gamma; s : (2, T) \vdash [P_2 \mid s[2, \varepsilon]] \text{ [TNode]}
 \end{array}
 } \quad (3.3)$$

Similarly to the previous derivation, the \circ operator is used to combine tuple $(2, \varepsilon)$, derived from rule [SEMP], and session type T , derived from typing judgement $\Gamma; s : T \vdash P_2$. The result is tuple $(2, T)$, which means that the s -endpoint used by the network node is in state 2 and has type T .

The network node $[s_1(\text{hbt}_1).P_1 \mid s[1, \varepsilon]]$ is also typed using an instance of rule [TNode]:

$$\boxed{
 \begin{array}{c}
 \text{[TSND]} \\
 \Gamma; s : T \vdash P_1 \quad \Gamma \vdash \text{hbt}_1 : \mathbf{b} \\
 \hline
 \Gamma; s : \mathbf{!\mathbf{b}.T} \vdash s_1(\text{hbt}_1).P_1 \\
 \hline
 \Gamma; s : (1, \varepsilon) \vdash s[1, \varepsilon] \text{ [SEMP]} \\
 s : \mathbf{!\mathbf{b}.T} \circ s : (1, \varepsilon) = s : (1, \mathbf{!\mathbf{b}.T}) \\
 \hline
 \Gamma; s : (1, \mathbf{!\mathbf{b}.T}) \vdash [s_1(\text{hbt}_1).P_1 \mid s[1, \varepsilon]] \text{ [TNode]}
 \end{array}
 } \quad (3.4)$$

The s -endpoint used by the network node is in state 1 and has type $\mathfrak{!b}.T$. However, it is not in the same state with the \check{s} -endpoint in derivation 3.2, therefore we apply rule [TSYNCH] to achieve endpoint synchronisation.

$$\frac{\Gamma; s : (1, \mathfrak{!b}.T) \vdash [s_1 \langle \mathfrak{hbt}_1 \rangle . P_1 \mid s[1, \varepsilon]] \quad 3.4 \quad \frac{\mathfrak{!b}.T \rightarrow^{2-1} T}{s : (1, \mathfrak{!b}.T) \hookrightarrow s : (2, T)} \quad \text{[TSYNCH]}}{\Gamma; s : (2, T) \vdash [s_1 \langle \mathfrak{hbt}_1 \rangle . P_1 \mid s[1, \varepsilon]]} \quad (3.5)$$

The [TSYNCH] rule uses linear context synchronisation, Definition 3.9, to synchronise type $\mathfrak{!b}.T$ from state 1 to state 2 and get $(2, T)$.

We then apply rule [TPAR] twice to get:

$$\frac{\Gamma; s : (2, T) \vdash [s_1 \langle \mathfrak{hbt}_1 \rangle . P_1 \mid s[1, \varepsilon]] \quad 3.5 \quad \Gamma; s : (2, T) \vdash [P_2 \mid s[2, \varepsilon]] \quad 3.3}{\Gamma; s : (2, T) \vdash [s_1 \langle \mathfrak{hbt}_1 \rangle . P_1 \mid s[1, \varepsilon]] \parallel [P_2 \mid s[2, \varepsilon]]} \quad \text{[TPAR]} \quad (3.6)$$

and also

$$\frac{\Gamma; s : (2, T) \vdash [s_1 \langle \mathfrak{hbt}_1 \rangle . P_1 \mid s[1, \varepsilon]] \parallel [P_2 \mid s[2, \varepsilon]] \quad 3.6 \quad \Gamma; \check{s} : (2, \bar{T}) \vdash [\check{s}_?(x) . \check{s}_?(x) . P_0 \mid \check{s}[0, (0, \mathfrak{hbt}_2), (1, \mathfrak{hbt}_2), (0, \mathfrak{hbt}_1)]] \quad 3.2}{\Gamma; \check{s} : (2, \bar{T}), s : (2, T) \vdash \begin{array}{l} [\check{s}_?(x) . \check{s}_?(x) . P_0 \mid \check{s}[0, (0, \mathfrak{hbt}_2), (1, \mathfrak{hbt}_2), (0, \mathfrak{hbt}_1)]] \\ \parallel [s_1 \langle \mathfrak{hbt}_1 \rangle . P_1 \mid s[1, \varepsilon]] \parallel [P_2 \mid s[2, \varepsilon]] \end{array}} \quad \text{[TPAR]} \quad (3.7)$$

Finally, we use rule [TSRES] to restrict session s and provide with the typing derivation for network Heartbeat_1 :

$$\frac{\Gamma; \check{s} : (2, \bar{T}), s : (2, T) \vdash \begin{array}{l} [\check{s}_?(x) . \check{s}_?(x) . P_0 \mid \check{s}[0, (0, \mathfrak{hbt}_2), (1, \mathfrak{hbt}_2), (0, \mathfrak{hbt}_1)]] \quad 3.7 \\ \parallel [s_1 \langle \mathfrak{hbt}_1 \rangle . P_1 \mid s[1, \varepsilon]] \parallel [P_2 \mid s[2, \varepsilon]] \end{array}}{\Gamma; \emptyset \vdash \text{Heartbeat}_1} \quad \text{[TSRES]}$$

Both endpoints are in the same state, 2, and have dual types, thus safe interaction endpoint interaction respects session duality and session s can be restricted. \square

4. TYPE SOUNDNESS, TYPE SAFETY, AND PROGRESS

In this Section we prove that the proposed type system is sound via a type preservation theorem and safe via a type safety theorem. Before we proceed with the main results we deploy the necessary technical machinery and auxiliary results.

We define the notion of the well-formed Linear Context.

Definition 4.1 (Well-formed Linear Context). A context Δ is *well-formed* whenever $\check{s} : (c, T) \in \Delta$ implies either

- $s : (c, \bar{T}) \in \Delta$; or
- $s \notin \text{dom}(\Delta)$.

Well-formed linear context Δ requires that whenever the \check{s} -endpoint is present in Δ , then the corresponding s -endpoint, if present in Δ , needs to be synchronised with the \check{s} -endpoint and, additionally, have dual type with respect to the \check{s} -endpoint. The case where

a corresponding s -endpoint is not present in Δ captures the case where no s -endpoints were created or when all s -endpoints were dropped.

The next definition captures the interaction of processes at the type level.

Definition 4.2 (Linear Context Advancement). Advancement relation, \rightarrow , over linear contexts is defined as:

- $\Delta, \check{s} : (c, !\beta.T_1), s : (c, ?\beta.T_2) \rightarrow \Delta, \check{s} : (c+1, T_1), s : (c+1, T_2)$.
- $\Delta, \check{s} : (c, \oplus\{l_i : T_i\}), s : (c, \&\{l_i : T'_i\}) \rightarrow \Delta, \check{s} : (c+1, T_k), s : (c+1, T'_k)$.
- $\Delta, s : (c, !\beta.T_1), \check{s} : (c, ?\beta.T_2) \rightarrow \Delta, s : (c+1, T_1), \check{s} : (c+1, T_2)$
- $\Delta, \{s : (c, T_i)\}_{i \in I} \rightarrow \Delta$

The first two cases of linear context advancement define session interaction at type level and align with broadcasting and selection interactions. The third case aligns with the session gather interaction. The fourth case indicates that a linear context advances by reducing in size. The definition aligns with the fact that processes may drop s -endpoints due to rules [REC], [BREC], [TRUE], and [FALSE].

The next lemma shows that the well-formedness of Δ is preserved by linear context inclusion and linear context advancement.

Lemma 4.3. *Let Δ be well-formed. If $\Delta \rightarrow \Delta'$ then Δ' is well-formed.*

The next Lemma is a consequence of the typing system.

Lemma 4.4. *Consider networks*

- *If $N_1 \equiv [s_1\langle e \rangle.P \mid B \mid \check{s}[c, \tilde{h}]]$ and $\Gamma; \Delta \vdash N_1$ with Δ well-formed, then $\tilde{h} = \varepsilon$.*
- *If $N_2 \equiv [s_1\langle e \rangle.P \mid B \mid s[c, \tilde{m}]]$ and $\Gamma; \Delta \vdash N_2$ with Δ well-formed, then $\tilde{m} = \varepsilon$.*
- *If $N_3 \equiv [s \triangleleft l.P \mid B \mid s[c, \tilde{m}]]$ and $\Gamma; \Delta \vdash N_3$ with Δ well-formed, then $\tilde{m} = \varepsilon$.*

Proof. The proof follows the requirement on rule [TNode] and the fact that operations: $i) !\beta.T \circ !\beta'.M; !\beta.T \circ \oplus \ell.M; \oplus\{\ell : T\} \circ !\beta'.M; \text{ and } \oplus\{\ell : T\} \circ \oplus \ell.M$ are undefined. \square

The above lemma states a standard property for asynchronous session type systems (cf. [GV10]), which requires that in a session typed setting whenever a network node has a send prefix, the corresponding session buffer is necessarily empty. This is because correct send/receive interaction will consume any messages in the session buffer prior to the send prefix.

We can now state and prove the Typing Preservation Theorem.

Theorem 4.5 (Typing Preservation). *If $\Gamma; \Delta \vdash N$ and $N \longrightarrow N'$ and Δ well-formed, then there exist well-formed Δ' such that $\Delta \rightarrow \Delta'$ and $\Gamma; \Delta' \vdash N'$.*

Proof. By induction on the depth of the derivation of $N \longrightarrow N'$. For the full proof see Appendix A.2 in page 38. \square

Type preservation theorem states that a reduction maintains typing and well-formedness.

Towards the statement of a type safety theorem, we proceed by defining the notion of the error network. The class of error networks indicates all networks that should not be typed with a well-formed linear context by the proposed typing system. We define the notion of an error network, cf. [YV07], as the network that contains an invalid s -pair that cannot make a safe communication interaction.

Definition 4.6 (Error Network). Let s -prefix be a network of the form

$$\begin{aligned} \text{BRC}^c &= [\check{s}_! \langle e \rangle . P_1 \mid \check{s}[c, \varepsilon] \mid B] & \text{GTH}^c &= [\check{s}_?(x) \langle e \rangle . P_2 \mid \check{s}[c, \tilde{h}] \mid B] \\ \text{UNI}^c &= [s_! \langle e' \rangle . P_3 \mid s[c, \varepsilon] \mid B] & \text{RCV}^c &= [s_?(x') . P_4 \mid s[c, \tilde{m}] \mid B] \\ \text{SEL}^c &= [\check{s} \triangleleft \ell . P_5 \mid \check{s}[c, \varepsilon] \mid B] & \text{BRA}^c &= [s \triangleright \{\ell_i : P_i, \mathbf{df} : R\}_{i \in I} \mid s[c, \tilde{m}] \mid B] \end{aligned}$$

An *invalid* s -pair is one of the following parallel compositions of s -prefixes:

$$\begin{array}{llll} \text{BRC}^c & \parallel & \text{BRC}^{c'} & \text{GTH}^c & \parallel & \text{GTH}^{c'} & \text{SEL}^c & \parallel & \text{SEL}^{c'} \\ \text{BRC}^c & \parallel & \text{GTH}^{c'} & \text{BRC}^c & \parallel & \text{SEL}^{c'} & \text{BRC}^c & \parallel & \text{UNI}^c & \text{BRC}^c & \parallel & \text{BRA}^c \\ \text{GTH}^c & \parallel & \text{SEL}^{c'} & \text{GTH}^c & \parallel & \text{RCV}^c & \text{GTH}^c & \parallel & \text{BRA}^c \\ \text{SEL}^c & \parallel & \text{RCV}^c & \text{SEL}^c & \parallel & \text{UNI}^c \end{array}$$

A network N is called an *error network* whenever there exists an invalid s -pair M such that for some network N' it holds that $N \equiv (\nu \tilde{n})(N' \parallel M)$.

An invalid s -pair is formed by a parallel composition of two s -prefixes. All other parallel compositions between two s -prefixes are considered valid. An error network is any network that composes in parallel at least one invalid s -pair. Equivalently, a valid network is a parallel composition of network terms of which all s -prefixes that are in the same session state, consists of at most one broadcast prefix (resp. gather, selection) and many receive prefixes (resp. send, branch).

An invalid s -pair is either a composition of a \check{s} -prefix and a s -prefix that cannot safely interact, or a composition of two \check{s} -prefixes that violates linearity conditions. In the latter case the composition of two \check{s} -endpoints is always an error regardless of the state they are in. In the former case we require that the two s -prefixes are in the same state c , because for such s -redexes not in the same state, reduction semantics will never allow an interaction, other than recovery (reduction rules [REC], [BREC], [LOSS]). The next example presents instances of error networks.

Example 4.7 (Error Network). As a first example, consider an instance of error network $\text{BRC}^c \parallel \text{BRA}^c$

$$[\check{s}_! \langle e \rangle . P \mid \check{s}[c, \varepsilon]] \parallel [s \triangleright \{\ell_i : P_i, \mathbf{df} : R\}_{i \in I} \mid s[c, \tilde{m}]]$$

The network cannot perform a safe interaction on channel s , because the \check{s} -endpoint does not select a continuation on the s -endpoint, therefore the above network cannot observe a reduction. The network is not typable, because two session endpoints (\check{s} -endpoint and s -endpoint) do not have dual types.

A more interesting example is given by an instance of the error network $N \parallel \text{BRC}^c \parallel \text{BRC}^c$

$$[\check{s}_! \langle e \rangle . \mathbf{0} \mid \check{s}[c, \varepsilon]] \parallel [\check{s}_! \langle e \rangle . \mathbf{0} \mid \check{s}[c, \varepsilon]] \parallel [s_?(x) . \mathbf{0} \mid s[c, \varepsilon]]$$

The network cannot perform a safe interaction on channel s , because the s -endpoint can interact with either of the \check{s} -endpoints. Network N is not typable, because of the requirement on rule [TPAR] that the \check{s} -endpoint is linear.

We also do not consider as invalid s -pair the case where two s -endpoints are composed, i.e. redexes of the form $\text{RCV}^c \parallel \text{UNI}^c$, $\text{RCV}^c \parallel \text{BRA}^c$ and $\text{UNI}^c \parallel \text{BRA}^c$. Such redexes may be well-typed, due to the requirement for linear context synchronisation on the [TSYNCH] rule. However, composing them with a corresponding \check{s} -endpoint results in a non well-typed network. Disallowing such s -redexes would require more complicated semantics that record interaction sequences instead of session state. The next example clarifies this last intuition.

Example 4.8. Consider for example an instance of network $\text{RCV}^0 \parallel \text{UNI}^0$:

$$[s_?(x).\mathbf{0} \mid s[0, \varepsilon]] \parallel [s_!(e).\mathbf{0} \mid s[0, \varepsilon]]$$

The network may be well-typed, due to linear context synchronisation, with well-formed context $\Delta = s : (1, s : \text{end})$ using derivation:

$$\boxed{\begin{array}{c} \Gamma; s : ?\beta_1.\text{end} \vdash s_?(x).\mathbf{0} \quad \Gamma; s : (0, \varepsilon) \vdash s[0, \varepsilon] \\ \hline s : ?\beta_1.\text{end} \circ s : (0, \varepsilon) = s : (0, ?\beta_1.\text{end}) \quad [\text{TNode}] \\ \Gamma; (0, ?\beta_1.\text{end}) \vdash [s_?(x).\mathbf{0} \mid s[0, \varepsilon]] \\ \hline s : (0, ?\beta_1.\text{end}) \hookrightarrow s : (1, \text{end}) \\ \hline \Gamma; s : (1, s : \text{end}) \vdash [s_?(x).\mathbf{0} \mid s[0, \varepsilon]] \quad [\text{TSynch}] \end{array}} \quad (4.1)$$

$$\boxed{\begin{array}{c} \Gamma; s : !\beta_2.\text{end} \vdash s_!(e).\mathbf{0} \quad \Gamma; s : (0, \varepsilon) \vdash s[0, \varepsilon] \\ \hline s : !\beta_2.\text{end} \circ s : (0, \varepsilon) = s : (0, !\beta_1.\text{end}) \quad [\text{TNode}] \\ \Gamma; s : (0, !\beta_1.\text{end}) \vdash [s_!(e).\mathbf{0} \mid s[0, \varepsilon]] \\ \hline s : (0, !\beta_2.\text{end}) \hookrightarrow s : (1, \text{end}) \\ \hline \Gamma; s : (1, \text{end}) \vdash [s_!(e).\mathbf{0} \mid s[0, \varepsilon]] \quad [\text{TSynch}] \end{array}} \quad (4.2)$$

that allows the usage of rule $[\text{TPAR}]$ to get

$$\boxed{\begin{array}{c} \Gamma; s : (1, \text{end}) \vdash [s_?(x).\mathbf{0} \mid s[0, \varepsilon]] \quad 4.1 \\ \Gamma; s : (1, \text{end}) \vdash [s_!(e).\mathbf{0} \mid s[0, \varepsilon]] \quad 4.2 \\ \hline \Gamma; s : (1, \text{end}) \vdash [s_?(x).\mathbf{0} \mid s[0, \varepsilon]] \parallel [s_!(e).\mathbf{0} \mid s[0, \varepsilon]] \quad [\text{TPAR}] \end{array}}$$

However, attempting to add a network node that implements the \check{s} -endpoint would result in an non well-typed network. For example, networks:

$$\begin{aligned} N_1 &= [\check{s}_!(e).\mathbf{0} \mid \check{s}[0, \varepsilon]] \parallel [s_?(x).\mathbf{0} \mid s[0, \varepsilon]] \parallel [s_!(e).\mathbf{0} \mid s[0, \varepsilon]] \\ N_2 &= [\check{s}_?(x).\mathbf{0} \mid \check{s}[0, \varepsilon]] \parallel [s_?(x).\mathbf{0} \mid s[0, \varepsilon]] \parallel [s_!(e).\mathbf{0} \mid s[0, \varepsilon]] \end{aligned}$$

are both non well-typed. In network N_1 , the \check{s} -endpoint is typed as $\Gamma; \check{s} : (0, !\beta.\text{end}) \vdash [\check{s}_!(e).\mathbf{0} \mid \check{s}[0, \varepsilon]]$ and cannot be synchronised with the type of network node $[s_!(e).\mathbf{0} \mid s[0, \varepsilon]]$. Similarly in network N_2 , the \check{s} -endpoint is typed as $\Gamma; \check{s} : (0, ?\beta.\text{end}) \vdash [\check{s}_?(x).\mathbf{0} \mid \check{s}[0, \varepsilon]]$ and cannot be synchronised with the type of network node $[s_?(x).\mathbf{0} \mid s[0, \varepsilon]]$.

The next theorem shows that our framework enjoys type safety.

Theorem 4.9 (Type Safety). *Let network N such that $\Gamma; \Delta \vdash N$ for some Γ and some well-formed Δ . If there exists network N' such that $N \longrightarrow^* N'$, then network N' is not an error network.*

Proof. From Theorem 4.5 and the fact that an error network is not well-typed with a well-formed linear context. \square

Type safety states that a network typed with a well-formed context can always interact safely, i.e. never reduce to an error.

4.1. Progress. We prove that a network typed with a well-formed session environment can ensure strong progress properties. We formally introduce the notion of a deadlocked network.

Definition 4.10 (Shared name prefixed process). A process is called *shared input prefixed process* if it has the form $a?(x).P$. A process is called *shared output prefixed process* if it has the form $a!(\check{x}).P$.

Definition 4.11 (Deadlocked Network). A network N is called *deadlocked* whenever

$$N \equiv (\nu \tilde{n})(\prod_{i \in I} [\sum_{j \in J_i} P_j \mid B_i])$$

and for all $i \in I$ we have that for all $j \in J_i$, P_j is prefixed with an input shared name.

It is easy to show that whenever a network N is a deadlocked network, it holds that there exists no N' such that $N \longrightarrow N'$. We can now formulate a basic progress result:

Theorem 4.12 (Progress). *Let N be a network such that $\Gamma; \Delta \vdash N$ for some Γ and some well-formed Δ . Then either*

- $N \equiv \prod_{i \in J} [\mathbf{0}]$; or
- N is a deadlocked network; or
- $N \longrightarrow N'$ for some N' .

Proof. The proof is straightforward since all instances of a session prefixed processes in a network node can perform an interaction, e.g. due to recovery semantics, whenever they are well-typed. Processes that are prefixed with an output on shared name in a network can also perform an interaction, whenever they are well-typed. \square

The progress result states that unless a deadlocked occurs due to shared names, a network term is always able to progress. Indeed, due to the recovery semantics sessions can progress even if they are interleaved. If we combine the progress result with the result for 4.9 we can also deduce that a network progresses in a session safe manner.

The progress result, however, does not indicate the nature of progress within a session, nor it provides additional insight on the nature of the recovery semantics in the progress within a session. We state two results on session progress and on session recovery. We first establish the notion of a simple network, following the notion of the simple process in [HYC08].

Definition 4.13. A network N is called *simple* if it is well-typed with a type derivation where the linear environment in the premise and the conclusion for each prefix rule in Figure 7 has at most one element.

Intuitively, a simple network is a network whose composing processes do not implement more than one active session name during their interaction, e.g. session interaction does not interleave. For example, network $N \equiv [a?(x).s!(\text{hbt}).x!(\text{hbt}).\mathbf{0} \mid s[0, \varepsilon]]$ is not simple because session name s and session variable x (which will be substituted with a fresh session at runtime) will be both active at runtime.

The next theorem states a form of session fidelity where a session within a simple network is always able to progress without observing a recovery interaction.

Theorem 4.14 (Session Progress). *Let N be a simple network such that $\Gamma; \Delta \vdash N$ for some Γ and some well-formed Δ . If*

$$N \equiv (\nu \tilde{n})([P \mid \check{s}[c, \tilde{m}] \mid B] \parallel \prod_{i \in I} [P_i \mid s[c, \tilde{m}_i] \mid B_i] \parallel M)$$

with *i*) $\check{s} \in \text{fs}(P)$, *ii*) for all $i \in I$ it holds that $s \in \text{fs}(P_i)$, and *iii*) $s \notin \text{fn}(M)$, then there exists N' such that $N \longrightarrow^* N'$ with

- $N' \equiv (\nu \tilde{n}')([P' | \check{s}[c+1, \tilde{m}'] | B'] \parallel \prod_{j \in J} [P_j | s[c+1, \tilde{m}_j] | B_j] \parallel M')$;
- \longrightarrow^* does not involve reduction rules [REC], [BREC], and [LOSS];
- $J \subseteq I$.

Proof. The proof comes from the fact that $\Gamma; \Delta \vdash N$ and Δ well-formed. The typing judgement implies that sub-network:

$$N' \equiv (\nu \tilde{n})([P | \check{s}[c, \tilde{m}] | B] \parallel \prod_{i \in I} [P_i | s[c, \tilde{m}_i] | B_i])$$

is also typed with a well-typed linear context and from the fact that it is a simple network we can show that it can perform the desired reduction using either: *i*) [BCAST] reduction possibly followed by several [RCV] reductions; *ii*) or a [SEL] reduction possible followed by several [BRA] reductions; *iii*) or several [UCAST] reductions followed by a [GTHR] reduction. \square

The next theorem demonstrates the ability of s -endpoints that are not synchronised with the corresponding \check{s} -endpoint to recover, given that the entire network is a simple network.

Theorem 4.15 (Session Recovery). *Let N be a simple network such that $\Gamma; \Delta \vdash N$ for some Γ and some well-formed Δ . If*

$$N \equiv (\nu \tilde{n})([P | \check{s}[c, \tilde{m}] | B] \parallel \prod_{i \in I} [P_i | s[c_i, \tilde{m}_i] | B_i] \parallel M)$$

and for all $i \in I, c_i < c$, then there exists N' such that $N \longrightarrow^* N'$ with

$$N' \equiv (\nu \tilde{n})([P | \check{s}[c, \tilde{m}] | B] \parallel \prod_{j \in J} [P'_j | s[c, \tilde{m}'_j] | B'_j] \parallel \prod_{k \in K} [P'_k | B'_k] \parallel M)$$

with $J, K \subseteq I$ and $J \cap K = \emptyset$.

Proof. The proof is done by double induction; the first is an induction on the size of I and the second is an induction on size $c - c_i$. The fact that N is simple allows us to observe recovery rules [REC], [BREC] and [LOSS] on the session prefixes of processes P_i that will increase the counter of the s -endpoint until it reaches the value of counter c . However, in the case of a [COND] process term a s -endpoint might be dropped resulting in a partition of the I set into sets J and K . \square

5. THE PAXOS CONSENSUS PROTOCOL

We present a session safe implementation of the Paxos consensus protocol using the UBSC. Paxos [Lam98, L⁺01] is a protocol for reaching consensus in a network that operates under conditions of unreliability. It ensures that network agents can agree on a single value in the presence of failures. Despite being the standard consensus algorithm, Paxos is notoriously difficult to understand, and real-world implementations have brought forth many problems that are not taken into account by the theoretical model of Paxos [CGR07].

The session type representation of Paxos is not used to prove the correctness of the protocol itself, but to check whether the Paxos protocol enjoys session types properties, i.e., type preservation, type safety, and progress, under the particular session specification. Session types can also help to identify subtle interactions such as branching or dropping

sessions. Furthermore, a session type representation allows for the basic algorithm to be easily extended while still providing formal guarantees.

5.1. The Paxos Protocol. We implement the most basic protocol of the Paxos family as described in [L⁺01]. The network agents act autonomously and propose values for consensus to the other agents within the network. If eventually a majority of agents run for long enough without failing, consensus on one of the proposed values is guaranteed. A correct implementation of the protocol ensures that:

- Only a value that has been proposed for consensus is chosen.
- The agents within the network agree on a single value.
- An agent is never informed that a value is chosen for consensus, unless it has been chosen for consensus.

The Paxos setting assumes asynchronous non-Byzantine communication that operates under the following assumptions:

- Messages can take arbitrarily long to be delivered, can be duplicated, but are not delivered corrupted.
- Agents operate at an arbitrary speed, may stop operating and may restart. However, it is assumed that agents maintain persistent storage that survives crashes.

Following the Paxos setting assumptions, an implementation of the Paxos protocol in UBSC is ensured to be correct under the assumptions A1-A6. Specifically: requirement S2 ensures asynchronous communication; requirement A4 ensures non-Byzantine communication; requirement A5 ensures non-corrupted delivered messages; requirement A6 ensures that agents operate at an arbitrary speed; and requirement A3 ensures that messages are never duplicated, which is subsumed by the Paxos requirement that messages might be duplicated, i.e., the properties of the Paxos protocol are also ensured in a setting where messages are never duplicated. Moreover, requirement A6 can express the case where an agent has failed by allowing the agent to take an arbitrary long time to perform an action. If the agent eventually interacts then it is considered to have restarted maintaining, through persistent storage, all the information it had prior to the fail.

Paxos agents implement three roles: i) a *proposer* agent proposes values towards the network for reaching consensus; ii) an *acceptor* accepts a value from those proposed, whereas a majority of acceptors accepting the same value implies consensus and signifies protocol termination; and iii) a *learner* discovers the chosen consensus value. The implementation of the protocol may proceed over several rounds. A successful round has two phases: **Prepare** and **Accept**.

The protocol ensures that in the case where a consensus value v has already been chosen among the majority of the network agents, broadcasting a new proposal request with a higher proposal number will result in choosing the already chosen consensus value v . Following this fact, we assume for simplicity that a learner has the same implementation as a proposer.

5.2. Implementation of the Paxos protocol in UBSC. Figure 10 describes the implementation of the Paxos protocol in our framework. The implementation assumes that expressions contain finite sets of integer tuples. We use notation $\{(r_i, v_i)_{i \in I}\}$ for such a finite set. Moreover, we assume that the aggregation operation is defined as a union of sets of integer tuples.

$$\begin{aligned}
\text{PaxosType} &= \text{!prepare.}\text{?promise.} \oplus \left\{ \begin{array}{l} \text{accept : !(\text{round, value}).end,} \\ \text{restart : end} \end{array} \right\} \\
\text{PaxosNode}_{r,v}^{\text{id}} &= [\text{Paxos}_{r,v}^{\text{id}}] \\
\text{Paxos}_{r,v}^{\text{id}} &= \text{def} \\
&\quad \text{Proposer}(x, y) \stackrel{\text{def}}{=} a_!(s).\check{s}_!(x).\check{s}_?(\{(r_i, v_i)\}_{i \in I}). \\
&\quad \quad \text{if } |I| > \frac{M}{2} \text{ then} \\
&\quad \quad \quad \check{s} \triangleleft \text{accept}.\check{s}_!(x, v = \text{choose}(\{(r_i, v_i)\}_{i \in I}, \text{id})). \\
&\quad \quad \quad \text{Paxos}\langle x, v \rangle \\
&\quad \quad \text{else} \\
&\quad \quad \quad \check{s} \triangleleft \text{restart.Paxos}\langle x, y \rangle \\
&\quad \text{Acceptor}(x, y) \stackrel{\text{def}}{=} a?(s).s?(x'). \\
&\quad \quad \text{if } (x' > x) \text{ then} \\
&\quad \quad \quad s_!(x, y).s \triangleright \left\{ \begin{array}{l} \text{accept : } s?(x', y').\text{Paxos}\langle x', y' \rangle, \\ \text{restart : Paxos}\langle x, y \rangle \end{array} \right\} \\
&\quad \quad \text{else} \\
&\quad \quad \quad \text{Paxos}\langle x, y \rangle \\
&\quad \text{Paxos}(x, y) \stackrel{\text{def}}{=} \text{Proposer}\langle x + 1, y \rangle + (\text{Acceptor}\langle x, y \rangle \diamond \text{Paxos}(x, y)) \\
&\text{in} \\
&\quad \text{Paxos}\langle r, v \rangle
\end{aligned}$$

Figure 10: Implementation of the Paxos consensus protocol

The interaction for establishing a consensus value takes place within a single session that involves the nodes of the network. The communication behaviour of the Acceptor is described by session type PaxosType , whereas the communication behaviour of the proposer is described by the dual type $\overline{\text{PaxosType}}$. The PaxosType session type provides with a type level description of the protocol in [L⁺01]. The fact that the type is enforced to the implementation, through the type system, together with the fact that the underlying assumptions of our calculus are covered by the assumptions of the Paxos setting, provide supporting evidence for proving the correctness of the implementation.

A Paxos agent is described by network node $\text{PaxosNode}_{r,v}^{\text{id}} = [\text{Paxos}_{r,v}^{\text{id}}]$ where r is the number of the current proposal number, v is the consensus value that corresponds to proposal number r , and id is a unique node identity number. A Paxos agent non-deterministically behaves either as a proposer, (definition $\text{Proposer}(x, y)$) or as an acceptor (definition $\text{Acceptor}(x, y) \diamond \text{Paxos}(x, y)$). The definition of the acceptor requires that during computation an acceptor agent recovers when an input endpoint does not progress, by dropping all active sessions and proceeding to process $\text{Paxos}\langle r, v \rangle$.

If a Paxos agent decides to act as a proposer, it does so by increasing its current proposal number and proceeding to process $\text{Proposer}\langle r + 1, v \rangle$. It then requests a new session and enters the Prepare phase. All the Paxos agents that accept a session request act as acceptors. The proposer then broadcasts towards the network a *prepare* message request, type prepare , that contains the proposal number r .

All the acceptors that received the *prepare* message check whether the proposal number is greater than the one they currently have. If not, they drop the session and restart the computation proceeding to process $\text{Paxos}\langle r, v \rangle$. Otherwise, they reply with a *promise* message, type **promise**, not to respond to a prepare message with a lower round number. The promise message contains the current proposal number and the current consensus value of the acceptor. If this is the first time the acceptor is involved in a consensus round and it has no information of consensus value then the promise message will contain empty values (ϵ, ϵ) . Here assume that for all proposal numbers r it holds that $r > \epsilon$.

After all the involved acceptors reply with a promise message, the protocol enters the **Accept** phase. The proposer gathers all the promises as a set of promise messages, $\{(r_i, v_i)\}_{i \in I}$, and then checks whether the majority of acceptors have replied using condition $|I| \leq \frac{M}{2}$, with M being the number of the nodes in the network. Note that for clarity, in the Proposer agent we use the set notation $\{(r_i, v_i)\}_{i \in I}$ in place of variable in an input process.

If the check fails the proposer sends a restart label to all the acceptors, and restarts its own computation by proceeding to process $\text{Paxos}\langle r, v \rangle$. All acceptors that receive label **restart** also restart their computation by proceeding to process $\text{Paxos}\langle r, v \rangle$.

If the majority check is passed, the proposer selects a value to submit to the acceptors by inspecting the promises received and choosing the value corresponding to the highest proposal number received in a promise message. If no value is received then it chooses its id value. This is expressed by computation

$$\begin{aligned} v_k &= \text{choose}(\{(r_i, v_i)\}_{i \in I}, \text{id}) \text{ when } \forall i \in I, r_k \geq r_i & \text{and} \\ \text{id} &= \text{choose}(\{(r_i, v_i)\}_{i \in I}, \text{id}) \text{ when } \forall i \in I, r_i = \epsilon \end{aligned}$$

The proposer then broadcasts an *accept* message, which is an **accept** label followed by a tuple that contains the current proposal number and the chosen highest value. Finally, it updates its own proposal number and consensus value and proceed to process $\text{Paxos}\langle r, v \rangle$. Moreover, all acceptors that receive the *accept* message update their proposal number and their consensus value, and proceed to state $\text{Paxos}\langle r, v \rangle$.

Network node $\text{PaxosNode}_{n,v}$ can be typed using the following typing judgement:

$$a : \text{PaxosType}; \emptyset \vdash \text{PaxosNode}_{n,v}^{\text{id}}$$

Shared channel a uses type PaxosType , thus all establish sessions of the computation follow the behaviour of the Paxos protocol as described by the PaxosType session type. Subsequently, a network that describes a set of nodes that run the Paxos protocol is defined as:

$$N = \prod_{i \in I} \text{PaxosNode}_{\epsilon, \epsilon}^i$$

Network N is typed using typing judgement $a : \text{PaxosType}; \emptyset \vdash N$. Typing is possible due to the typing of network node $\text{PaxosNode}_{n,v}^{\text{id}}$ and multiple applications of rule [TPAR].

The paxos process $\text{Paxos}_{r,v}^{\text{id}}$ is recursive, where each recursion indicates a new propose round that implies the creation of a new session of type PaxosType . Each iteration creates a new session by necessity, since if, due to reliability, a session within a propose round does not include the majority of Paxos acceptors then it is necessary to terminate the session and reiterate. This is reflected by the fact that type PaxosType is not recursive.

Our typing framework guaranties that network N , which is typed with a well-formed (empty) linear context, will never reduce to an error state (Theorem 4.5). Moreover, the fact that a Paxos network is a simple network (Definition 4.13) and the progress results (Theorems 4.14 and 4.14) ensure that the interaction takes place within a session as defined

by the session type. This lifts the burden from the programmer to check for deadlocks and type mismatches, and leaves only the burden for implementing correctly the algorithmic logic. If the algorithmic logic is implemented correctly, then the fact that the assumptions of the UBSC are covered by the execution context of Paxos ensures the properties of the Paxos protocol.

More complicated Paxos networks can be described. For example, we can allow a Paxos agent that acts as an acceptor to establish multiple sessions with different proposers during an execution and explicitly drop the session with the lowest proposal number:

$$\begin{aligned}
\text{Acceptor}(x, y) &\stackrel{\text{def}}{=} a_{?}(s).s_{?}(x').\text{Acc}\langle s, x, x', y \rangle \\
\text{Acc}(w, x, x', y) &\stackrel{\text{def}}{=} \text{if } (x' > x) \text{ then} \\
&\quad w_{!}\langle x, y \rangle. \\
&\quad \left(\begin{array}{l} \text{AcceptPhase}\langle w, x, y \rangle \\ + \\ a_{?}(s').s'_{?}(x''). \\ \text{if } (x'' > x') \\ \quad \text{Acc}\langle s', x, x'', y \rangle \\ \text{else} \\ \quad \text{AcceptPhase}\langle w, x, y \rangle \end{array} \right) \\
&\quad \text{else} \\
&\quad \text{Paxos}\langle x, y \rangle \\
\text{AcceptPhase}(w, x, y) &\stackrel{\text{def}}{=} w \triangleright \left\{ \begin{array}{l} \text{accept} : w_{?}(x', y').\text{Paxos}\langle x', y' \rangle, \\ \text{restart} : \text{Paxos}\langle x, y \rangle \end{array} \right\}
\end{aligned}$$

The definition of the acceptor has the option, expressed as non-deterministic choice, to establish a new session, i.e. enter a prepare phase with a different proposer, while in the accept phase of another proposal. The computation then compares the proposal numbers from the two sessions. Following the result of the comparison, the computation will explicitly drop the session with the lowest proposal number and proceed as described by the type of the session with the highest proposal number. The above process is well-typed; note that the application of rule [TSUM] ensures that the two branches of the non-deterministic operator have the same session type.

The new acceptor definition results in a network that is not simple following Definition 4.13, therefore the conditions for Session Progress (Theorem 4.14) and Session Recovery (Theorem 4.15) do not hold. Nevertheless, the network will never reduce to an error network (Definition 4.6) and it can be proven that it still enjoys the progress property via Theorem 4.12. Moreover we can show that individual sessions can still progress and recover.

6. CONCLUSION AND FUTURE WORK

This paper is motivated by the need to sufficiently define a session type framework for systems such as ad-hoc and wireless sensor networks. The basic characteristic of these networks is a lossy stateless communication medium shared among an arbitrary number of agents.

To this end, we have introduced the asynchronous unreliable broadcast calculus accompanied with a corresponding session type framework. The semantics of the calculus are inspired by the practice of ad-hoc and wireless sensor network and develop the necessary mechanisms to support safe session interaction and recovery. Asynchrony is achieved with the use of

message buffers. The main session communication operations include asynchronous broadcast operation and the asynchronous gather communication pattern in the presence of link failure and message loss. Message loss may lead to session endpoints that are not synchronised with the overall protocol. In such a case the semantics propose the appropriate mechanisms for autonomous session recovery. The recovery semantics find sufficient justification from the practice of networks that operate in an unreliable setting.

Our session type systems ensures safe and sound communication interaction that respects the session types principles. A type preservation theorem ensures the soundness of the system, whereas a type safety theorem ensures that the system will never reduce an undesired/error state. Finally, a set of progress properties ensure that sessions may always progress, and, in the case of non-synchronisation, a session can always safely recover.

The syntax and semantics of the UBSC are expressive enough to describe non-trivial communication protocols that operate in an unreliable setting. We used the UBSC to describe a basic implementation of the Paxos consensus algorithm, the standard protocol for achieving consensus in an asynchronous unreliable setting. Our framework satisfies the underlying operation assumptions of the Paxos protocol. The computation for reaching consensus takes place within a single session, therefore a single session type can describe the interaction of the Paxos agents. The session typing system ensures that a Paxos agent interaction follows the Paxos session protocol.

Moreover, this work follows the syntax of binary session types, where a rather intricate typing system enforces the static properties of soundness, safety, and deadlock freedom to a calculus that supports realistic interactions of broadcast, gather, and recovery. However, it is not in the aim (nor in the static nature) of the UBSC session type system to guarantee more dynamic properties that have to do with the algorithm correctness, e.g., to prove that the Paxos protocol will eventually reach consensus.

Future Work. Our system is open for further extensions in the future. A first direction is to capture more complicated patterns beyond the unreliable broadcast communication and the gather communication pattern. We would also like to investigate more elaborate autonomous recovery mechanisms, e.g. to explicitly define time-out and interrupt routines.

Another interesting direction is the extension of the system with locations and mobility, cf. [HR02], as this would allow the use of session types in a highly dynamic and complex setting. Location and node mobility is an important feature which allows for new nodes to be introduced, disabled, or migrate between locations within a network. Locality and/or mobility might also imply semantics that take into account broadcast/communication range. The feature of communication range further strengthens the need for the existence of autonomous recovery mechanisms.

An important direction for future research is to develop a more robust session type system, based on multiparty session types, where network agents can interact with multiple roles inside a session. The research direction on a multiparty session type framework in the context of unreliable broadcast communication can follow principles from the work on parametrised multiparty session types [NY15], where a session type system is developed for describing interaction among an arbitrary number of agents.

Acknowledgements. Kouzapas and Gay were supported by the UK EPSRC project EP/K034413/1 “From Data Types to Session Types: A Basis for Concurrency and Distribution”. Voinea was supported by an EPSRC PhD studentship. This research was supported by a Short-Term Scientific Mission grant from COST Action IC1201 (Behavioural Types for Reliable Large-Scale Software Systems). Voinea and Gay were supported by the UK EPSRC project “Session Types for Reliable Distributed Systems” (EP/T014628/1 and EP/T014512/1) and by BehAPI (Behavioural Application Program Interfaces), an EU H2020 RISE programme (Marie Skłodowska-Curie grant agreement No 778233).

REFERENCES

- [ABB⁺16] 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. Behavioral types in programming languages. *Foundations and Trends® in Programming Languages*, 3:95–230, 2016.
- [APN17] Manuel Adameit, Kirstin Peters, and Uwe Nestmann. Session types for link failures. In Ahmed Bouajjani and Alexandra Silva, editors, *Formal Techniques for Distributed Objects, Components, and Systems*, pages 1–16, Cham, 2017. Springer International Publishing.
- [BD23] Matthew Alan Le Brun and Ornela Dardha. Mag π : Types for failure-prone communication. In Thomas Wies, editor, *Programming Languages and Systems - 32nd European Symposium on Programming, ESOP 2023, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2023, Paris, France, April 22-27, 2023, Proceedings*, volume 13990 of *Lecture Notes in Computer Science*, pages 363–391. Springer, 2023. doi:10.1007/978-3-031-30044-8_14.
- [BDGK14] Giovanni Bernardi, Ornela Dardha, Simon J. Gay, and Dimitrios Kouzapas. On duality relations for session types. In Matteo Maffei and Emilio Tuosto, editors, *Trustworthy Global Computing*, pages 51–66, Berlin, Heidelberg, 2014. Springer Berlin Heidelberg.
- [BHJ⁺11] Johannes Borgström, Shuqin Huang, Magnus Johansson, Palle Raabjerg, Björn Victor, Johannes Åman Pohjola, and Joachim Parrow. Broadcast psi-calculi with an application to wireless protocols. In *SEFM 2011*, pages 74–89, 2011.
- [CDP12] Giuseppe Castagna, Mariangiola Dezani-Ciancaglini, and Luca Padovani. On global types and multi-party session. *Logical Methods in Computer Science*, 8(1), 2012.
- [CDYP16] Mario Coppo, Mariangiola Dezani-Ciancaglini, Nobuko Yoshida, and Luca Padovani. Global progress for dynamically interleaved multiparty sessions. *Mathematical Structures in Computer Science*, 26(2):238–302, 2016.
- [CGR07] Tushar D Chandra, Robert Griesemer, and Joshua Redstone. Paxos made live: an engineering perspective. In *Proceedings of the twenty-sixth annual ACM symposium on Principles of distributed computing*, pages 398–407. ACM, 2007.
- [CGY16] Sara Capecchi, Elena Giachino, and Nobuko Yoshida. Global escape in multiparty sessions. *Mathematical Structures in Computer Science*, 26, 2 2016.
- [CHY08] Marco Carbone, Kohei Honda, and Nobuko Yoshida. Structured interactional exceptions in session types. In *CONCUR*, volume 5201 of *LNCS*, pages 402–417. Springer, 2008.
- [DOH07] Adam Dunkels, Fredrik Österlind, and Zhitao He. An adaptive communication architecture for wireless sensor networks. In *Proceedings of Embedded Networked Sensor Systems*, SenSys ’07, pages 335–349. ACM, 2007.
- [DYBH12] Pierre-Malo Deniérou, Nobuko Yoshida, Andi Bejleri, and Raymond Hu. Parameterised multiparty session types. *Logical Methods in Computer Science*, 8(4), 2012.
- [GR17] Simon J. Gay and António Ravara, editors. *Behavioural Types: from Theory to Tools*. River Publishers, 2017.
- [GV10] Simon J. Gay and Vasco Thudichum Vasconcelos. Linear type theory for asynchronous session types. *J. Funct. Program.*, 20(1):19–50, 2010. doi:10.1017/S0956796809990268.
- [HKP⁺10] Raymond Hu, Dimitrios Kouzapas, Olivier Pernet, Nobuko Yoshida, and Kohei Honda. Type-safe eventful sessions in java. In *ECOOOP*, pages 329–353, 2010.

- [HR02] Matthew Hennessy and James Riely. Resource access control in systems of mobile agents. *Information and Computation*, 173(1):82 – 120, 2002.
- [HVK98] Kohei Honda, Vasco Thudichum Vasconcelos, and Makoto Kubo. Language primitives and type discipline for structured communication-based programming. In *ESOP*, pages 122–138. Springer-Verlag, 1998.
- [HYC08] Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty asynchronous session types. In *POPL*, pages 273–284, 2008.
- [KGG14] Dimitrios Kouzapas, Ramunas Gutkovas, and Simon J. Gay. Session types for broadcasting. In *PLACES*, pages 25–31, 2014.
- [Kou12] Dimitrios Kouzapas. *A Study of Bisimulation Theory for Session Types*. PhD thesis, Imperial College London, 2012.
- [KYHH16] Dimitrios Kouzapas, Nobuko Yoshida, Raymond Hu, and Kohei Honda. On asynchronous eventful session semantics. *Mathematical Structures in Computer Science*, 26(2):303–364, 2016. doi:10.1017/S096012951400019X.
- [L⁺01] Leslie Lamport et al. Paxos made simple. *ACM Sigact News*, 32(4):18–25, 2001.
- [Lam98] Leslie Lamport. The part-time parliament. *ACM Trans. Comput. Syst.*, 16(2):133–169, 1998.
- [NY15] Nicholas Ng and Nobuko Yoshida. Pabble: parameterised scribble. *Service Oriented Computing and Applications*, 9(3-4):269–284, 2015. doi:10.1007/s11761-014-0172-8.
- [PNC23] Kirstin Peters, Uwe Nestmann, and Wagner Christoph. Ftmpst: Fault-tolerant multiparty session types, 2023. <https://arxiv.org/pdf/2204.07728.pdf>.
- [YV07] Nobuko Yoshida and Vasco T. Vasconcelos. Language primitives and type discipline for structured communication-based programming revisited: Two systems for higher-order session communication. *Electronic Notes in Theoretical Computer Science*, 171(4):73 – 93, 2007.

APPENDIX A. PROOFS

A.1. Congruence invariance proof.

Lemma A.1 (Congruence Invariance). • *If $P \equiv P'$, then $\Gamma; \Delta \vdash P$ if, and only if, $\Gamma; \Delta \vdash P'$.*

• *If $N \equiv N'$, then $\Gamma; \Delta \vdash N$ if, and only if, $\Gamma; \Delta \vdash N'$.*

Proof. By induction on the structural congruence definition. \square

Lemma A.2 (Substitution). • *Whenever $\Gamma \vdash e : \beta$ and $\Gamma, x : \beta; \Delta \vdash P$, then $\Gamma; \Delta \vdash P\{e/x\}$.*

• *Whenever $\Gamma; \Delta, x : T \vdash P$, then $\Gamma; \Delta, s : T \vdash P\{s/x\}$.*

Proof. By induction on the definition of processes. \square

A.2. Typing Preservation proof. Now we are ready to prove the typing preservation property of our system.

Proof of Theorem 4.5. By induction on the depth of derivation of $N \longrightarrow N'$.

• Rule [CONN]
Assume

$$\frac{s \text{ fresh}}{[a_!(\check{x}).P \mid B] \parallel \prod_{i \in I} [a_?(x).Q_i \mid B_i] \longrightarrow (\nu s)([P\{\check{s}/\check{x}\} \mid B \mid \check{s}[0, \varepsilon]] \parallel \prod_{i \in I} [Q_i\{s/x\} \mid B_i \mid s[0, \varepsilon]])}$$

and

$$\Gamma; \Delta \vdash [a_!(\check{x}).P \mid B] \parallel \prod_{i \in I} [a_?(x).Q_i \mid B_i]$$

with Δ well-formed. From the latter typing judgement we get the derivation

$$\frac{\Gamma; \Delta_0, \Delta_1 \vdash [a_!(\check{x}).P \mid B] \quad \Gamma; \Delta_0, \Delta_2 \vdash \prod_{i \in I} [a_?(x).Q_i \mid B_i] \quad \Delta_0 \text{ only } s\text{-endpoints} \quad \Delta = \Delta_0, \Delta_1, \Delta_2}{\Gamma; \Delta \vdash [a_!(\check{x}).P \mid B] \parallel \prod_{i \in I} [a_?(x).Q_i \mid B_i]} \text{ [TPAR]}$$

We continue the above derivation for node $[a_!(\check{x}).P \mid B]$

<table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="padding: 5px;"> $\frac{\Gamma; \Delta_1'', \check{x} : \bar{T} \vdash P \quad \Gamma \vdash a : T}{\Gamma; \Delta_1'' \vdash a_!(\check{x}).P} \text{ [TREQ]}$ </td> <td rowspan="2" style="padding: 5px; vertical-align: middle;">$\Delta_1' \hookrightarrow \Delta_0, \Delta_1$</td> </tr> <tr> <td style="padding: 5px;"> $\frac{\Gamma; \Theta \vdash B \quad \Delta_1' = \Delta_1'' \circ \Theta}{\Gamma; \Delta_1' \vdash [a_!(s).P \mid B]} \text{ [TNODE]}$ </td> </tr> </table>	$\frac{\Gamma; \Delta_1'', \check{x} : \bar{T} \vdash P \quad \Gamma \vdash a : T}{\Gamma; \Delta_1'' \vdash a_!(\check{x}).P} \text{ [TREQ]}$	$\Delta_1' \hookrightarrow \Delta_0, \Delta_1$	$\frac{\Gamma; \Theta \vdash B \quad \Delta_1' = \Delta_1'' \circ \Theta}{\Gamma; \Delta_1' \vdash [a_!(s).P \mid B]} \text{ [TNODE]}$	$\frac{}{\Gamma; \Delta_0, \Delta_1 \vdash [a_!(s).P \mid B]} \text{ [TSYNCH]}$
$\frac{\Gamma; \Delta_1'', \check{x} : \bar{T} \vdash P \quad \Gamma \vdash a : T}{\Gamma; \Delta_1'' \vdash a_!(\check{x}).P} \text{ [TREQ]}$	$\Delta_1' \hookrightarrow \Delta_0, \Delta_1$			
$\frac{\Gamma; \Theta \vdash B \quad \Delta_1' = \Delta_1'' \circ \Theta}{\Gamma; \Delta_1' \vdash [a_!(s).P \mid B]} \text{ [TNODE]}$				

For every node $[a?(x).Q_i | B_i]$ we can derive

$$\boxed{
 \begin{array}{c}
 \frac{\Gamma; \Delta''_i, x : T \vdash Q_i \quad \Gamma \vdash a : T}{\Gamma; \Delta''_i \vdash a?(x).Q_i} \text{ [TACC]} \\
 \frac{\Gamma; \Theta_i \vdash B_i \quad \Delta'_i = \Delta''_i \circ \Theta_i}{\Gamma; \Delta'_i \vdash [a?(x).Q_i | B_i]} \text{ [TNODE]} \\
 \hline
 \Gamma; \Delta_0, \Delta_i \vdash [a?(x).Q_i | B_i] \text{ [TSYNCH]}
 \end{array}
 }$$

and also with multiple applications of the [TPAR] rule we get

$$\frac{\{\Gamma; \Delta_0, \Delta_i \vdash [a?(x).Q_i | B_i]\}_{i \in I} \quad \Delta_2 = \bigcup_{i \in I} \Delta_i}{\Gamma; \Delta_0, \Delta_2 \vdash \prod_{i \in I} [a?(x).Q_i | B_i]} \text{ [TPAR]}$$

We can then combine the information from the latter transition to derive the type for the result of the reduction $(\nu s)([P\{\check{s}/\check{x}\} | B | \check{s}[0, \varepsilon]] \parallel \prod_{i \in I} [Q_i\{s/x\} | B_i | s[0, \varepsilon]])$. We first type network node $[P\{\check{s}/\check{x}\} | B | \check{s}[0, \varepsilon]]$:

$$\boxed{
 \begin{array}{c}
 \frac{\Gamma; \check{s} : (0, \varepsilon) \vdash \check{s}[0, \varepsilon] \text{ [SEMP]} \quad \Gamma; \Theta \vdash B \text{ (from former derivation)}}{\Gamma; \Theta, \check{s} : (0, \varepsilon) \vdash B | \check{s}[\varepsilon, \varepsilon]} \text{ [BPAR]} \\
 \\
 \text{Substitution Lemma A.2 and former derivation} \\
 \Gamma; \Delta''_1, \check{x} : \bar{T} \vdash P \text{ imply } \Gamma; \Delta''_1, \check{s} : \bar{T} \vdash P\{\check{s}/\check{x}\} \\
 \Delta''_1, \check{s} : \bar{T} \circ \Theta, \check{s} : (0, \varepsilon) = \Delta'_1, \check{s} : (0, \bar{T}) \\
 \hline
 \Gamma; \Delta'_1, \check{s} : (0, \bar{T}) \vdash [P\{\check{s}/\check{x}\} | B | \check{s}[0, \varepsilon]] \text{ [TNODE]} \\
 \\
 \text{Linear context synchronisation (Definition 3.9) and former result} \\
 \Delta'_1 \hookrightarrow \Delta_0, \Delta_1 \text{ imply } \Delta'_1, s : 0, \check{s} : \bar{T} \hookrightarrow \Delta_0, \Delta_1, \check{s} : (0, \bar{T}) \\
 \hline
 \Gamma; \Delta_0, \Delta_1, \check{s} : (0, \bar{T}) \vdash [P\{\check{s}/\check{x}\} | B | \check{s}[0, \varepsilon]] \text{ [TSYNCH]}
 \end{array}
 }$$

Similarly we type every network node $[Q_i\{s/x\} | B_i | s[0, \varepsilon]]$

$$\boxed{
\begin{array}{c}
\boxed{
\begin{array}{c}
\Gamma; s : (0, \varepsilon) \vdash s[0, \varepsilon] \quad [\text{SEMP}] \\
\frac{\Gamma; \Theta \vdash B_i \text{ (from former derivation)}}{\Gamma; \Theta, s : (0, \varepsilon) \vdash B_i | s[\varepsilon, \varepsilon]} \quad [\text{BPAR}]
\end{array}
} \\
\text{Substitution Lemma A.2 and former derivation} \\
\Gamma; \Delta_i'', x : T \vdash Q_i \text{ imply } \Gamma; \Delta_i'', s : T \vdash Q_i\{s/x\} \\
\Delta_i'', s : T \circ \Theta, s : (0, \varepsilon) = \Delta_i' s : (0, T) \\
\hline
\Gamma; \Delta_i', s : (0, T) \vdash [Q_i\{s/x\} | B_i | s[0, \varepsilon]] \quad [\text{TNODE}]
\end{array}
} \\
\text{Linear context synchronisation (Definition 3.9) and former result} \\
\Delta_i' \hookrightarrow \Delta_0, \Delta_1 \text{ imply } \Delta_i', s : (0, T) \hookrightarrow \Delta_0, \Delta_i, s : (0, T) \\
\hline
\Gamma; \Delta_0, \Delta_i, s : (0, T) \vdash [Q_i\{s/x\} | B_i | s[0, \varepsilon]] \quad [\text{TSYNCH}]
\end{array}$$

We can now have a multiple application of the $[\text{TPAR}]$ rule to type network $\prod_{i \in I} [Q_i\{s/x\} | B_i | s[0, \varepsilon]]$

$$\frac{\Gamma; \Delta_0, \Delta_i, s : (0, T) \vdash [Q_i\{s/x\} | B_i | s[0, \varepsilon]] \\
\text{from former result } \Delta_2 = \bigcup_{i \in I} \Delta_i}{\Gamma; \Delta_0, \Delta_2, s : (0, T) \vdash \prod_{i \in I} [Q_i\{s/x\} | B_i | s[0, \varepsilon]]} \quad [\text{TPAR}]$$

We then use typing rule $[\text{TPAR}]$ followed by rule $[\text{TSRES}]$ to type the result of the reduction $(\nu s)([P\{\check{s}/\check{x}\} | B | \check{s}[0, \varepsilon]] \parallel \prod_{i \in I} [Q_i\{s/x\} | B_i | s[0, \varepsilon]])$:

$$\boxed{
\begin{array}{c}
[\text{TSRES}] \\
\boxed{
\begin{array}{c}
[\text{TPAR}] \\
\Gamma; \Delta_0, \Delta_1, \check{s} : (0, \bar{T}) \vdash [P\{\check{s}/\check{x}\} | B | \check{s}[0, \varepsilon]] \\
\Gamma; \Delta_0, \Delta_2, s : (0, T) \vdash \prod_{i \in I} [Q_i\{s/x\} | B_i | s[0, \varepsilon]] \\
\hline
\Gamma; \Delta, \check{s} : (0, \bar{T}), s : (0, T) \vdash [P\{\check{s}/\check{x}\} | B | \check{s}[0, \varepsilon]] \parallel \prod_{i \in I} [Q_i\{s/x\} | B_i | s[0, \varepsilon]]
\end{array}
} \\
\hline
\Gamma; \Delta \vdash (\nu s)([P\{\check{s}/\check{x}\} | B | \check{s}[0, \varepsilon]] \parallel \prod_{i \in I} [Q_i\{s/x\} | B_i | s[0, \varepsilon]])
\end{array}$$

The last result together with the trivial result $\Delta \rightarrow \Delta$ concludes the case.

- Rule $[\text{BCAST}]$
Assume

$$[\check{s}_i \langle e \rangle . P | B | \check{s}[c, \varepsilon]] \parallel \prod_{i \in I} [P_i | B_i | s[c, \tilde{m}_i]] \longrightarrow [P | B | s[c+1, \varepsilon]] \parallel \prod_{i \in I} [P_i | B_i | s[c+1, \tilde{m}_i \cdot e]]$$

and

$$\Gamma; \Delta \vdash [\check{s}_i \langle e \rangle . P | B | \check{s}[c, \varepsilon]] \parallel \prod_{i \in I} [P_i | B_i | s[c, \tilde{m}_i]]$$

with Δ well-formed. We proceed with the typing derivation of the latter typing judgement.

well-formed Δ implies $\Delta = \Delta_0, \Delta_1, \check{s} : (c, !\beta.\bar{T}), \Delta_2, s : (c, ?\beta.T)$
with Δ_0 only s -endpoints

$$\frac{\begin{array}{c} \Gamma; \Delta_0, \Delta_1, \check{s} : (c, !\beta.\bar{T}) \vdash [\check{s}!(e).P \mid B \mid \check{s}[c, \varepsilon]] \\ \Gamma; \Delta_0, \Delta_2, s : ?\beta.T \vdash \prod_{i \in I} [P_i \mid B_i \mid s[c, \tilde{m}_i]] \end{array}}{\Gamma; \Delta \vdash [\check{s}!(e).P \mid B \mid \check{s}[c, \varepsilon]] \parallel \prod_{i \in I} [P_i \mid B_i \mid s[c, \tilde{m}_i]]} \text{ [TPAR]}$$

We continue the above derivation for network node $[\check{s}!(e).P \mid B \mid \check{s}[c, \varepsilon]]$:

$$\boxed{\begin{array}{c} \boxed{\frac{\Gamma; \check{s} : (c, \varepsilon) \vdash \check{s}[c, \varepsilon] \quad \text{[SEMP]} \quad \Gamma; \Theta \vdash B}{\Gamma; \Theta, \check{s} : (c, \varepsilon) \vdash B \mid \check{s}[c, \varepsilon]} \text{ [BPAR]}} \\ \frac{\Gamma; \Delta_1'', \check{s} : \bar{T} \vdash P \quad \Gamma \vdash e : \beta}{\Gamma; \Delta_1'', \check{s} : !\beta.\bar{T} \vdash \check{s}!(e).P} \text{ [TSND]} \\ \frac{\Delta_1', \check{s} : (c, !\beta.\bar{T}) = \Delta_1'', \check{s} : !\beta.\bar{T} \circ \Theta, \check{s} : (c, \varepsilon)}{\Gamma; \Delta_1', \check{s} : (c, !\beta.\bar{T}) \vdash [\check{s}!(e).P \mid B \mid \check{s}[c, \varepsilon]]} \text{ [TNODE]} \\ \frac{\Delta_1', \check{s} : (c, !\beta.\bar{T}) \hookrightarrow \Delta_0, \Delta_1, \check{s} : (c, !\beta.\bar{T})}{\Gamma; \Delta_0, \Delta_1, \check{s} : (c, !\beta.\bar{T}) \vdash [\check{s}!(e).P \mid B \mid \check{s}[c, \varepsilon]]} \text{ [TSYNCH]} \end{array}}$$

We then derive the typing for network nodes $[P_i \mid B_i \mid s[c, \tilde{m}_i]]$

$$\boxed{\begin{array}{c} \text{[TNODE]} \\ \boxed{\frac{\Gamma; s : (c, M) \vdash s[c, \tilde{m}_i] \quad \text{[SEXP]} \quad \Gamma; \Theta_i \vdash B_i}{\Gamma; \Theta_i, s : (c, M) \vdash B_i \mid s[c, \tilde{m}_i]} \text{ [BPAR]}} \\ \frac{\Gamma; \Delta_i'', s : T' \vdash P_i \quad \Delta_i', s : (c, ?\beta.T) = \Delta_i'', s : T' \circ \Theta, s : (c, M)}{\Gamma; \Delta_i', s : (c, ?\beta.T) \vdash [P_i \mid B_i \mid s[c, \tilde{m}_i]]} \\ \frac{\Delta_i', s : (c, ?\beta.T) \hookrightarrow \Delta_0, \Delta_i, s : (c, ?\beta.T)}{\Gamma; \Delta_0, \Delta_i, s : (c, ?\beta.T) \vdash [P_i \mid B_i \mid s[c, \tilde{m}_i]]} \text{ [TSYNCH]} \end{array}}$$

We can now have multiple applications of rule [TPAR] to get:

$$\frac{\Gamma; \Delta_0, \Delta_i, s : (c, ?\beta.T) \vdash [P_i \mid B_i \mid s[c, \tilde{m}_i]] \quad \Delta_2 = \bigcup_{i \in I} \Delta_i}{\Gamma; \Delta_0, \Delta_2, s : (c, ?\beta.T) \vdash \prod_{i \in I} [P_i \mid B_i \mid s[c, \tilde{m}_i]]} \text{ [TPAR]}$$

Using the above information we can type the result of the reduction, $[P \mid B \mid s[c + 1, \varepsilon]] \parallel \prod_{i \in I} [P_i \mid B_i \mid s[c + 1, \tilde{m}_i \cdot e]]$. We begin with the typing derivation for network node

$[P | B | s[c + 1, \varepsilon]]$:

$$\begin{array}{c}
\boxed{\text{[TSYNCH]}} \\
\boxed{\frac{\frac{\Gamma; \check{s} : (c + 1, \varepsilon) \vdash \check{s}[c + 1, \varepsilon] \quad \text{[SEMP]} \quad \Gamma; \Theta \vdash B}{\Gamma; \Theta, \check{s} : (c + 1, \varepsilon) \vdash B | \check{s}[c + 1, \varepsilon]} \quad \text{[BPAR]}}{\text{from former derivation } \Gamma; \Delta_1'', \check{s} : \bar{T} \vdash P} \\
\Delta_1', \check{s} : (c + 1, \bar{T}) = \Delta_1'', \check{s} : \bar{T} \circ \Theta, \check{s} : (c + 1, \varepsilon)}{\Gamma; \Delta_1', \check{s} : (c + 1, \bar{T}) \vdash [P | B | \check{s}[c, \varepsilon]]} \quad \text{[TNODE]}} \\
\text{Linear context synchronisation (Definition 3.9) and former result} \\
\Delta_1', \check{s} : (c, \imath\beta.\bar{T}) \hookrightarrow \Delta_0, \Delta_1, \check{s} : (c, \imath\beta.\bar{T}) \text{ implies} \\
\Delta_1', \check{s} : (c + 1, \bar{T}) \hookrightarrow \Delta_0, \Delta_1, \check{s} : (c + 1, \bar{T})} \\
\hline
\Gamma; \Delta_0, \Delta_1, \check{s} : (c + 1, \bar{T}) \vdash [P | B | \check{s}[c + 1, \varepsilon]]
\end{array}$$

We derive the typing judgement for network node $\prod_{i \in I} [P_i | B_i | s[c + 1, \tilde{m}_i \cdot e]]$:

$$\begin{array}{c}
\boxed{\text{[TSYNCH]}} \\
\boxed{\frac{\boxed{\text{[BPAR]}}} \text{from former judgement} \\
\frac{\frac{\Gamma; s : (c, M) \vdash s[c, \tilde{m}_i \cdot e]}{\Gamma; s : (c + 1, M.\imath\beta) \vdash s[c + 1, \tilde{m}_i \cdot e]} \quad \text{[SEXP]} \quad \Gamma; \Theta_i \vdash B_i}{\Gamma; \Theta_i, s : (c + 1, M.\imath\beta) \vdash B_i | s[c + 1, \tilde{m}_i \cdot e]}} \\
\Gamma; \Delta_i'', s : T' \vdash P_i} \\
\text{Operator } \circ \text{ (Definition 3.4) and former result} \\
\Delta_i', s : (c, \imath\beta.T) = \Delta_i'', s : T' \circ \Theta, s : (c, M) \text{ imply} \\
\Delta_i', s : (c + 1, T) = \Delta_i'', s : T' \circ \Theta, s : (c + 1, M.\imath\beta)} \\
\hline
\Gamma; \Delta_i', s : (c + 1, T) \vdash [P_i | B_i | s[c + 1, \tilde{m}_i \cdot e]] \quad \text{[TNODE]} \\
\Delta_i', s : (c + 1, T) \hookrightarrow \Delta_0, \Delta_i, s : (c + 1, T)} \\
\hline
\Gamma; \Delta_0, \Delta_i, s : (c + 1, T) \vdash [P_i | B_i | s[c + 1, \tilde{m}_i \cdot e]]
\end{array}$$

We can now have multiple applications of rule [TPAR] to get:

$$\frac{\Gamma; \Delta_0, \Delta_i, s : (c + 1, T) \vdash [P_i | B_i | s[c + 1, \tilde{m}_i \cdot e]] \quad \Delta_2 = \bigcup_{i \in I} \Delta_i}{\Gamma; \Delta_0, \Delta_2, s : (c + 1, T) \vdash \prod_{i \in I} [P_i | B_i | s[c + 1, \tilde{m}_i \cdot e]]} \quad \text{[TPAR]}$$

The final rule needed to be applied is [TPAR] to get:

$$\frac{\frac{\Gamma; \Delta_0, \Delta_1, \check{s} : (c + 1, \bar{T}) \vdash [P | B | \check{s}[c + 1, \varepsilon]] \quad \Gamma; \Delta_0, \Delta_2, s : (c + 1, T) \vdash \prod_{i \in I} [P_i | B_i | s[c + 1, \tilde{m}_i \cdot e]]}{\Gamma; \Delta_0, \Delta_1, \check{s} : (c + 1, \bar{T}), \Delta_2, s : (c + 1, T) \vdash [P | B | \check{s}[c + 1, \varepsilon]]} \quad \text{[TPAR]}}{\Gamma; \Delta_0, \Delta_1, \check{s} : (c + 1, \bar{T}), \Delta_2, s : (c + 1, T) \vdash [P | B | \check{s}[c + 1, \varepsilon]] \quad \parallel \prod_{i \in I} [P_i | B_i | s[c + 1, \tilde{m}_i \cdot e]]}$$

The case concludes because $\Delta = \Delta_0, \Delta_1, \check{s} : (c, !\beta.\bar{T}), \Delta_2, s : ?\beta.T$ and $\Delta \rightarrow \Delta_0, \Delta_1, \check{s} : (c+1, \bar{T}), \Delta_2, s : (c+1, T)$ as required.

- Rule [UCAST]
Assume

$$\frac{[UCAST] \quad c_1 \geq c_2}{\frac{[s!(e).P_1 \mid B_1 \mid s[c_1, \varepsilon]] \parallel [P_2 \mid B_2 \mid \check{s}[c_2, \tilde{h}]]}{\longrightarrow [P_1 \mid B_1 \mid s[c_1+1, \varepsilon]] \parallel [P_2 \mid \check{s}[c_2, \tilde{h} \cdot (c_1, e)]]}}$$

and

$$\Gamma; \Delta \vdash [s!(e).P_1 \mid B_1 \mid s[c_1, \varepsilon]] \parallel [P_2 \mid B_2 \mid \check{s}[c_2, \tilde{h}]]$$

We provide the derivation for the latter typing judgement.

<p>well-formed Δ implies $\Delta = \Delta_0, \Delta_1, s : (c_3, T), \Delta_2, \check{s} : (c_3, \bar{T})$ with Δ_0 only s-endpoints and</p> $\frac{\Gamma; \Delta_0, \Delta_1, s : (c_3, T) \vdash [s!(e).P_1 \mid B_1 \mid s[c_1, \varepsilon]] \quad \Gamma; \Delta_0, \Delta_0, \check{s} : (c_3, \bar{T}) \vdash [P_2 \mid B_2 \mid \check{s}[c_2, \tilde{h}]]}{\Gamma; \Delta \vdash [s!(e).P_1 \mid B_1 \mid s[c_1, \varepsilon]] \parallel [P_2 \mid B_2 \mid \check{s}[c_2, \tilde{h}]]} \quad [TPAR]$

We continue with the typing derivation for network node: $[s!(e).P_1 \mid B_1 \mid s[c_1, \varepsilon]]$

<table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="padding: 5px;"> $\frac{\Gamma; s : (c_1, \varepsilon) \vdash s[c_1, \varepsilon] \quad [SEMP] \quad \Gamma; \Theta_1 \vdash B_1}{\Gamma; \Theta_1, s : (c_1, \varepsilon) \vdash B_1 \mid s[c_1, \varepsilon]} \quad [BPAR]$ </td> </tr> <tr> <td style="padding: 5px;"> $\frac{\Gamma; \Delta'_1, s : T_1 \vdash P_1 \quad \Gamma \vdash e : \beta}{\Gamma; \Delta'_1, s : !\beta.T_1 \vdash \check{s}!(e).P_1} \quad [TSND]$ </td> </tr> <tr> <td style="padding: 5px;"> $\frac{\Delta'_1, s : (c_1, !\beta.T_1) = \Delta'_1, s : !\beta.T_1 \circ \Theta_1, s : (c_1, \varepsilon)}{\Gamma; \Delta'_1, s : (c_1, !\beta.T_1) \vdash [s!(e).P_1 \mid B_1 \mid s[c_1, \varepsilon]]} \quad [TNODE]$ </td> </tr> </table> $\frac{\Delta'_1 \leftrightarrow \Delta_0, \Delta_1 \quad c_3 \geq c_1 \quad !\beta.T_1 \rightarrow^{c_3-c_1} T}{\Delta'_1, s : (c_1, !\beta.T_1) \leftrightarrow \Delta_0, \Delta_1, s : (c_3, T)} \quad [TSYNCH]$ $\Gamma; \Delta_0, \Delta_1, s : (c_3, T) \vdash [s!(e).P_1 \mid B_1 \mid s[c_1, \varepsilon]]$	$\frac{\Gamma; s : (c_1, \varepsilon) \vdash s[c_1, \varepsilon] \quad [SEMP] \quad \Gamma; \Theta_1 \vdash B_1}{\Gamma; \Theta_1, s : (c_1, \varepsilon) \vdash B_1 \mid s[c_1, \varepsilon]} \quad [BPAR]$	$\frac{\Gamma; \Delta'_1, s : T_1 \vdash P_1 \quad \Gamma \vdash e : \beta}{\Gamma; \Delta'_1, s : !\beta.T_1 \vdash \check{s}!(e).P_1} \quad [TSND]$	$\frac{\Delta'_1, s : (c_1, !\beta.T_1) = \Delta'_1, s : !\beta.T_1 \circ \Theta_1, s : (c_1, \varepsilon)}{\Gamma; \Delta'_1, s : (c_1, !\beta.T_1) \vdash [s!(e).P_1 \mid B_1 \mid s[c_1, \varepsilon]]} \quad [TNODE]$
$\frac{\Gamma; s : (c_1, \varepsilon) \vdash s[c_1, \varepsilon] \quad [SEMP] \quad \Gamma; \Theta_1 \vdash B_1}{\Gamma; \Theta_1, s : (c_1, \varepsilon) \vdash B_1 \mid s[c_1, \varepsilon]} \quad [BPAR]$			
$\frac{\Gamma; \Delta'_1, s : T_1 \vdash P_1 \quad \Gamma \vdash e : \beta}{\Gamma; \Delta'_1, s : !\beta.T_1 \vdash \check{s}!(e).P_1} \quad [TSND]$			
$\frac{\Delta'_1, s : (c_1, !\beta.T_1) = \Delta'_1, s : !\beta.T_1 \circ \Theta_1, s : (c_1, \varepsilon)}{\Gamma; \Delta'_1, s : (c_1, !\beta.T_1) \vdash [s!(e).P_1 \mid B_1 \mid s[c_1, \varepsilon]]} \quad [TNODE]$			

We also continue the derivation of network node $[P_2 | B_2 | \check{s}[c_2, \tilde{h}]]$:

$$\begin{array}{c}
 \boxed{\begin{array}{c}
 \dots \\
 \frac{}{\Gamma; \check{s} : (c_3, M) \vdash \check{s}[c_2, \tilde{h}]} \text{ [LEXP]} \\
 \frac{c_3 \geq c_1 \quad \Gamma; \Theta_2 \vdash B_2}{\Gamma; \Theta_2, \check{s} : (c_3, M) \vdash B_2 | \check{s}[c_2, \tilde{h}]} \text{ [BPAR]}
 \end{array}} \\
 \frac{\Gamma; \Delta_2'', \check{s} : T_2 \vdash P_2 \quad \Delta_2', \check{s} : (c_3, \bar{T}) = \Delta_2'', \check{s} : T_2 \circ \Theta_2, \check{s} : (c_3, M)}{\Delta_2', \check{s} : (c_3, \bar{T}) \vdash [P_2 | B_2 | \check{s}[c_2, \tilde{h}]]} \text{ [TNODE]} \\
 \frac{\Delta_2', \check{s} : (c_3, \bar{T}) \hookrightarrow \Delta_0, \Delta_2, \check{s} : (c_3, \bar{T})}{\Delta_0, \Delta_2, \check{s} : (c_3, \bar{T}) \vdash [P_2 | B_2 | \check{s}[c_2, \tilde{h}]]} \text{ [TSYNCH]}
 \end{array}$$

Using the above information we can type the result of the reduction, $[P_1 | B_1 | s[c_1 + 1, \varepsilon]] \parallel [P_2 | \check{s}[c_2, \tilde{h} \cdot (c_1, e)]]$. We begin with the typing derivation for network node $[P_1 | B_1 | s[c_1 + 1, \varepsilon]]$:

$$\begin{array}{c}
 \text{[TNODE]} \\
 \boxed{\begin{array}{c}
 \frac{\Gamma; s : (c_1 + 1, \varepsilon) \vdash s[c_1 + 1, \varepsilon] \text{ [SEMP]} \quad \Gamma; \Theta_1 \vdash B_1}{\Gamma; \Theta_1, s : (c_1 + 1, \varepsilon) \vdash B_1 | s[c_1 + 1, \varepsilon]} \text{ [BPAR]} \\
 \frac{\Gamma; \Delta_1'', s : T_1 \vdash P_1 \quad \Delta_1', s : (c_1 + 1, T_1) = \Delta_1'', s : T \circ \Theta_1, s : (c_1 + 1, \varepsilon)}{\Gamma; \Delta_1', s : (c_1 + 1, T_1) \vdash [P_1 | B_1 | s[c_1 + 1, \varepsilon]]}
 \end{array}} \\
 \text{results } \Delta_1', s : (c_1 + 1, T_1) \hookrightarrow \Delta_0, \Delta_1, s : (c_3, T) \text{ and } c_3 \geq c_1 \text{ imply} \\
 \frac{\Delta_1' \hookrightarrow \Delta_0, \Delta_1 \quad c_4 = c_3 \vee c_4 = c_3 + 1 \quad T_1 \rightarrow^{c_4 - c_1 + 1} T'}{\Delta_1', s : (c_1 + 1, T_1) \hookrightarrow \Delta_0, \Delta_1, s : (c_4, T')} \\
 \frac{}{\Gamma; \Delta_0, \Delta_1, s : (c_4, T') \vdash [P_1 | B_1 | s[c_1 + 1, \varepsilon]]} \text{ [TSYNCH]}
 \end{array}$$

We continue with the typing derivation for network node $[P_2 \mid \check{s}[c_2, \tilde{h} \cdot (c_1, e)]]$:

$$\boxed{
\begin{array}{c}
\text{[TNode]} \\
\boxed{
\begin{array}{c}
\text{[BPAR]} \\
\frac{\Gamma; \check{s} : (c_3, M) \vdash \check{s}[c_2, \tilde{h}] \quad \Gamma; \Theta \vdash B_2}{\Gamma; \check{s} : (c_4, M, !\beta) \vdash \check{s}[c_2, \tilde{h} \cdot (c_1, e)]} \text{[LEXP]} \\
\frac{\Gamma; \Theta_2, \check{s} : (c_4, M, !\beta) \vdash B_2 \mid \check{s}[c_2, \tilde{h} \cdot (c_1, e)]}{\Gamma; \Theta_2, \check{s} : (c_4, \overline{T}') \vdash [P_2 \mid B_2 \mid \check{s}[c_2, \tilde{h}]]}
\end{array}
} \\
\frac{\Gamma; \Delta'_2, \check{s} : T_2 \vdash P_2 \quad \Delta'_2, \check{s} : (c_4, \overline{T}') = \Delta'_2, \check{s} : T_2 \circ \Theta_2, \check{s} : (c_4, M, !\beta)}{\Delta'_2, \check{s} : (c_4, \overline{T}') \vdash [P_2 \mid B_2 \mid \check{s}[c_2, \tilde{h}]]} \\
\frac{\Delta'_2, \check{s} : (c_4, \overline{T}') \hookrightarrow \Delta_0, \Delta_2, \check{s} : (c_4, \overline{T}')}{\Delta_0, \Delta_2, \check{s} : (c_4, \overline{T}') \vdash [P_2 \mid B_2 \mid \check{s}[c_2, \tilde{h}]]} \text{[TSYNCH]}
\end{array}
}$$

We apply rule [TPAR] to get

$$\frac{\Gamma; \Delta_0, \Delta_1, s : (c_4, T') \vdash [P_1 \mid B_1 \mid s[c_1 + 1, \varepsilon]] \quad \Delta_0, \Delta_2, \check{s} : (c_4, \overline{T}') \vdash [P_2 \mid B_2 \mid \check{s}[c_2, \tilde{h}]]}{\Gamma; \Delta_0, \Delta_1, s : (c_4, T'), \Delta_2, \check{s} : (c_4, \overline{T}') \vdash [P_1 \mid B_1 \mid s[c_1 + 1, \varepsilon]] \parallel [P_2 \mid B_2 \mid \check{s}[c_2, \tilde{h}]]} \text{[TPAR]}$$

From here there are two cases: If $c_4 = c_3$ then $T' = T$ and the result follows from the fact that $\Delta_0, \Delta_1, s : (c_4, T'), \Delta_2, \check{s} : (c_4, \overline{T}') = \Delta_0, \Delta_1, s : (c_3, T), \Delta_2, \check{s} : (c_3, \overline{T})$. If $c_4 = c_3 + 1$ then $T \rightarrow^1 T'$ and the results follows the fact that $\Delta_0, \Delta_1, s : (c_3, T), \Delta_2, \check{s} : (c_3, \overline{T}) \rightarrow^1 \Delta_0, s : (c_4, T'), \Delta_2, \check{s} : (c_4, \overline{T}')$

- Rule [RCV]

Assume

$$[s?(x)\langle e' \rangle.P \mid B \mid s[c, e \cdot \tilde{m}]] \longrightarrow [P\{e/x\} \mid B \mid s[c, \tilde{m}]]$$

with

$$\Gamma; \Delta \vdash [s?(x)\langle e' \rangle.P \mid B \mid s[c, e \cdot \tilde{m}]]$$

and Δ well-formed. We give the derivation for the latter typing judgement

$$\boxed{
\begin{array}{c}
\frac{\Gamma; s : (c, M) \vdash B \mid s[c, \tilde{m}]}{\Gamma; s : (c, !\beta.M) \vdash B \mid s[c, e \cdot \tilde{m}]} \text{[SEXP]} \quad \Gamma; \Theta \vdash B \\
\frac{\Gamma; s : (c, !\beta.M) \vdash B \mid s[c, e \cdot \tilde{m}]}{\Gamma; \Theta, s : (c, !\beta.M) \vdash B \mid s[c, e \cdot \tilde{m}]} \text{[BPAR]} \\
\boxed{
\frac{\Gamma \vdash e' : \beta \quad \Gamma, x : \beta; \Delta', s : T' \vdash P}{\Gamma; \Delta', s : ?\beta.T' \vdash s?(x)\langle e' \rangle.P} \text{[TRCV]}
} \\
\Delta = \Delta', s : (c, T) = \Delta', s : ?\beta.T' \circ \Theta, (c, !\beta.M) \\
= \Delta, s : T' \circ \Theta, s : (c, M) \\
\frac{\Gamma; \Delta', s : ?\beta.T' \vdash s?(x)\langle e' \rangle.P \quad \Delta = \Delta', s : T' \circ \Theta, s : (c, M)}{\Gamma; \Delta \vdash [s?(x)\langle e' \rangle.P \mid B \mid s[c, e \cdot \tilde{m}]]} \text{[TNode]}
\end{array}
}$$

The above result allows us to produce the typing derivation for the result of the reduction, $[P\{e/x\} \mid B \mid s[c, \tilde{m}]]$:

$$\frac{\frac{\Gamma; s : (c, M) \vdash B \mid s[c, \tilde{m}] \quad \Gamma; \Theta \vdash B}{\Gamma; \Theta, s : (c, M) \vdash B \mid s[c, \tilde{m}]} \text{ [BPAR]}}{\Gamma; \Delta', s : T' \vdash P \quad \Delta = \Delta', s : (c, T) = \Delta, s : T' \circ \Theta, s : (c, M)} \text{ [TNode]}}{\Gamma; \Delta \vdash [P\{e/x\} \mid B \mid s[c, \tilde{m}]]}$$

as required.

- Rule [GTHR].

Assume

$$\frac{\tilde{h}' = \mathbf{B}(\tilde{h}, c + 1) \quad e = \mathbf{V}(\tilde{h}, c + 1)}{[\check{s}_?(x).P \mid B \mid \check{s}[c, \tilde{h}]] \longrightarrow [P\{e/x\} \mid B \mid \check{s}[c + 1, \tilde{h}']]} \text{ [GTHR]}$$

with

$$\Gamma; \Delta \vdash [\check{s}_?(x).P \mid B \mid \check{s}[c, \tilde{h}]]$$

and Δ well-formed. We produce the derivation of the latter typing judgement: We first type the session buffer $\check{s}[c, \tilde{h}]$:

$$\left. \begin{array}{c} \frac{\Gamma; \check{s} : (c, \varepsilon) \vdash \check{s}[c, \varepsilon] \text{ [SEMP]}}{\frac{\varepsilon = \mathbf{B}(\tilde{h}_1, c + 1) \quad \Gamma \vdash V(\tilde{h}_1, c + 1) : \beta}{\Gamma; \check{s} : (c + 1, \imath\beta) \vdash \check{s}[c, \tilde{h}_1]} \text{ [LEXP]}} \\ \vdots \\ \frac{\Gamma; \check{s} : (c' - 1, \imath\beta.M) \vdash \check{s}[c, \tilde{h}_{c'-c}] \quad \tilde{h} = \mathbf{B}(\tilde{h}_{c'-c}, c') \quad \Gamma \vdash V(\tilde{h}, c') : \beta'}{\Gamma; \check{s} : (c', \imath\beta.M, \imath\beta') \vdash \check{s}[c, \tilde{h}]} \text{ [LEXP]} \end{array} \right\} \begin{array}{l} c' - c \\ \text{applications} \\ \text{of rule [LEXP]} \end{array}$$

We now produce the typing derivation for term $[\check{s}_?(x).P \mid B \mid \check{s}[c, \tilde{h}]]$:

$$\frac{\frac{\Gamma; \check{s} : (c', \imath\beta.M, \imath\beta') \vdash \check{s}[c, \tilde{h}] \quad \Gamma; \Theta \vdash B}{\Gamma; \Theta, \check{s} : (c', \imath\beta.M, \imath\beta') \vdash B \mid \check{s}[c, \tilde{h}]} \text{ [BPAR]}}{\frac{\Gamma, x : \beta; \Delta', \check{s} : T' \vdash P}{\Gamma; \Delta', \check{s} : \imath\beta.T' \vdash \check{s}_?(x).P} \text{ [TRCV]}}{\Delta = \Delta', \check{s} : (c', T) = \Delta', \check{s} : \imath\beta.T' \circ \Theta, \check{s} : (c', \imath\beta.M, \imath\beta')} = \Delta', \check{s} : T' \circ \Theta, \check{s} : (c', M, \imath\beta') \text{ [TNode]}}{\Gamma; \Delta \vdash [\check{s}_?(x).P \mid B \mid \check{s}[c, \tilde{h}]]}$$

We use the above information to produce the typing derivation for the result of the reduction. We first produce the typing derivation for session buffer $\check{s}[c+1, \tilde{h}']$

$$\left. \begin{array}{c}
 \boxed{\Gamma; \check{s} : (c+1, \varepsilon) \vdash \check{s}[c+1, \varepsilon] \text{ [SEMP]}} \\
 \frac{\varepsilon = \mathbf{B}(\tilde{h}'_2, c+2) \quad \Gamma \vdash V(\tilde{h}'_2, c+2) : \beta''}{\Gamma; \check{s} : (c+2, \imath\beta'') \vdash \check{s}[c+1, \tilde{h}'_2]} \text{ [LEXP]} \\
 \vdots \\
 \frac{\Gamma; \check{s}(c'-1, M) \vdash \check{s}[c+1, \tilde{h}'_{c'-c+1}] \quad \tilde{h}' = \mathbf{B}(\tilde{h}'_{c'-c+1}, c') \quad \Gamma \vdash V(\tilde{h}', c') : \beta'}{\Gamma; \check{s} : (c', M, \imath\beta') \vdash \check{s}[c+1, \tilde{h}']} \text{ [LEXP]}
 \end{array} \right\} \begin{array}{l} c' - c + 1 \\ \text{applications} \\ \text{of rule [LEXP]} \end{array}$$

Finally we give the typing derivation for network node: $[P\{e/x\} \mid B \mid \check{s}[c+1, \tilde{h}']]$

$$\boxed{\begin{array}{c}
 \text{[BPAR]} \\
 \frac{\Gamma; \check{s} : (c', M, \imath\beta') \vdash \check{s}[c+1, \tilde{h}'] \quad \Gamma; \Theta \vdash B}{\Gamma; \Theta, \check{s} : (c', M, \imath\beta') \vdash B \mid \check{s}[c', \tilde{h}']} \\
 \\
 \text{From Substitution Lemma A.2} \\
 \Gamma, x : \beta; \Delta', \check{s} : T' \vdash P \text{ implies } \Gamma; \Delta', \check{s} : T' \vdash P\{e/x\} \\
 \frac{\Gamma; \emptyset \vdash R \quad \Delta = \Delta', \check{s} : (c', T) = \Delta', \check{s} : T' \circ \Theta, \check{s} : (c', M, \imath\beta')}{\Gamma; \Delta \vdash [\check{s}\langle x \rangle.P \mid B \mid \check{s}[c, \tilde{h}]]} \text{ [TNODE]}
 \end{array}}$$

The result follows the fact that $\Delta \rightarrow \Delta$.

- Rule [SEL]. The proof is similar to the proof for rule [BCAST].
- Rule [BRA]. The proof is similar to the proof for rule [RCV].
- Rule [LOSS]

Assume

$$[s_1\langle e \rangle.P \mid B \mid s[c, \varepsilon]] \longrightarrow [P \mid B \mid s[c+1, \varepsilon]]$$

with

$$\Gamma; \Delta \vdash [s_1\langle e \rangle.P \mid B \mid s[c, \varepsilon]]$$

and Δ well-formed. We provide the derivation of the latter typing judgement:

$$\boxed{\begin{array}{c}
 \frac{\Gamma; s : (c, \varepsilon) \vdash s[c, \varepsilon] \text{ [SEMP]} \quad \Gamma; \Theta \vdash B}{\Gamma; \Theta, s : (c, \varepsilon) \vdash B \mid s[c, \varepsilon]} \text{ [BPAR]} \\
 \\
 \frac{\Gamma; \Delta', s : T \vdash P \quad \Gamma \vdash e : \beta}{\Gamma; \Delta', s : \imath\beta. \vdash s_1\langle e \rangle.P} \text{ [TSND]} \\
 \\
 \frac{\Delta = \Delta', s : (c, \imath\beta.T) = \Delta', s : \imath\beta.T \circ \Theta, s : (c, \varepsilon)}{\Gamma; \Delta \vdash [s_1\langle e \rangle.P \mid B \mid s[c, \varepsilon]]} \text{ [TNODE]}
 \end{array}}$$

The above result allows us to produce the typing derivation for the result of the reduction, $[P | B] s[c+1, \varepsilon]$

$$\boxed{
\begin{array}{c}
\text{[TSYNCH]} \\
\boxed{
\frac{\Gamma; s : (c+1, \varepsilon) \vdash s[c+1, \varepsilon] \text{ [SEMP]} \quad \Gamma; \Theta \vdash B \text{ [BPAR]}}{\Gamma; \Theta, s : (c+1, \varepsilon) \vdash B | s[c+1, \varepsilon]} \text{ [BPAR]} \\
\Gamma; \Delta', s : T \vdash P \quad \Delta', s : (c+1, T) = \Delta', s : T \circ \Theta, s : (c+1, \varepsilon) \\
\hline
\Gamma; \Delta', s : (c+1, T) \vdash [P | B] s[c+1, \varepsilon] \text{ [TNODE]} \\
\Delta' \hookrightarrow \Delta' \quad !\beta.T \rightarrow_o^{c+1-c} T \quad \Delta = \Delta', s : (c, !\beta.T) \\
\hline
\Delta', s : (c+1, T) \hookrightarrow \Delta', s : (c, !\beta.T) \\
\hline
\Gamma; \Delta \vdash [P | B] s[c+1, \varepsilon]
\end{array}
}$$

as required.

- Rule [REC]

Assume

$$[s?(x)\langle e \rangle.P | B] s[c, \varepsilon] \longrightarrow [P\{e/x\} | B] s[c+1, \varepsilon] \text{ [REC]}$$

with

$$\Gamma; \Delta \vdash [s?(x)\langle e \rangle.P | B] s[c, \varepsilon]$$

and Δ well-formed. We give the derivation for the latter typing judgement

$$\boxed{
\begin{array}{c}
\boxed{
\frac{\Gamma; s : (c, \varepsilon) \vdash s[c, \varepsilon] \text{ [SEMP]} \quad \Gamma; \Theta \vdash B \text{ [BPAR]}}{\Gamma; \Theta, s : (c, \varepsilon) \vdash B | s[c, \varepsilon]} \text{ [BPAR]} \\
\Gamma \vdash e : \beta \quad \Gamma, x : \beta; \Delta', s : T \vdash P \\
\hline
\Gamma; \Delta', s : ?\beta.T \vdash s?(x)\langle e \rangle.P \text{ [TRCV]} \\
\Delta = \Delta', s : (c, ?\beta.T) = \Delta', s : ?\beta.T \circ \Theta, (c, \varepsilon) \\
\hline
\Gamma; \Delta \vdash [s?(x)\langle e \rangle.P | B] s[c, e \cdot \tilde{m}] \text{ [TNODE]}
\end{array}
}$$

The above result allows us to produce the typing derivation for the result of the reduction, $[P\{e/x\} | B] s[c+1, \varepsilon]$

$$\boxed{
\begin{array}{c}
\text{[TSYNCH]} \\
\boxed{
\frac{\Gamma; s : (c+1, \varepsilon) \vdash s[c+1, \varepsilon] \quad \Gamma; \Theta \vdash B \text{ [BPAR]}}{\Gamma; \Theta, s : (c+1, \varepsilon) \vdash B | s[c, \tilde{m}]} \text{ [BPAR]} \\
\Gamma; \Delta', s : T \vdash P\{e/x\} \quad \Delta', s : (c+1, T) = \Delta, s : T \circ \Theta, s : (c+1, \varepsilon) \\
\hline
\Gamma; \Delta', s : (c+1, T) \vdash [P\{e/x\} | B] s[c, \tilde{m}] \text{ [TNODE]} \\
\Delta', s : (c+1, T) \hookrightarrow \Delta', s : (c, ?\beta.T) \quad \Delta = \Delta', s : (c, ?\beta.T) \\
\hline
\Gamma; \Delta \vdash [P\{e/x\} | B] s[c, \tilde{m}]
\end{array}
}$$

as required.

- Rule [BREC]
Assume

$$\frac{\text{fs}(B) = \text{fs}(R)}{[s \triangleright \{\ell_i : P_i, \mathbf{df} : R\}_{i \in I} \mid B \mid B' \mid s[c, \varepsilon]] \longrightarrow [R \mid B]} \text{ [BREC]}$$

with

$$\Gamma; \Delta \vdash [s \triangleright \{\ell_i : P_i, \mathbf{df} : R\}_{i \in I} \mid B \mid B' \mid s[c, \varepsilon]]$$

and Δ well-formed. The last result follows typing derivation:

$$\boxed{\begin{array}{c} \boxed{\frac{\Gamma; s : (c, \varepsilon) \vdash s[c, \varepsilon] \text{ [SEMP]} \\ \Gamma; \Theta_1 \vdash B' \quad \Gamma; \Theta_2 \vdash B \quad \Theta = \Theta_1, \Theta_2}{\Gamma; \Theta, s : (c, \varepsilon) \vdash B \mid B' \mid s[c, \varepsilon]} \text{ [BPAR]}} \\ \boxed{\frac{\forall i \in I, \Gamma; \Delta_1, \Delta_2, s : T_i \vdash P_i \\ \text{only } s\text{-endpoints in } \Delta_1 \quad \Gamma; \Delta_2 \vdash R \quad \Delta' = \Delta_1, \Delta_2}{\Gamma; \Delta', s : \&\{\ell_i : T_i\}_{i \in I} \vdash s \triangleright \{\ell_i : P_i, \mathbf{df} : R\}_{i \in I}} \text{ [TBR]}} \\ \Delta = \Delta', s : \&\{\ell_i : T_i\}_{i \in I} \circ \Theta, s : (c, \varepsilon) \\ \hline \Gamma; \Delta \vdash [s \triangleright \{\ell_i : P_i, \mathbf{df} : R\}_{i \in I} \mid B \mid B' \mid s[c, \varepsilon]] \text{ [TNODE]} \end{array}}$$

We can now provide the typing derivation for the result of the reduction $[R \mid B]$

$$\boxed{\frac{\Gamma; \Delta_2 \vdash R \quad \Gamma; \Theta_2 \vdash B \quad \text{fs}(R) = \text{fs}(B) \text{ implies } \Delta'' = \Delta_2 \circ \Theta_2}{\Gamma; \Delta'' \vdash [R \mid B]} \text{ [TNODE]}}$$

Because Δ_1 includes only s -endpoints, we can use \rightarrow to drop them:

$$\begin{aligned} \Delta &= \Delta_1, \Delta_2, s : \&\{\ell_i : T_i\}_{i \in I} \circ \Theta_1, \Theta_2, s : (c, \varepsilon) \\ &= (\Delta_1 \circ \Theta_1), (\Delta_2 \circ \Theta_2), s : (c, \&\{\ell_i : T_i\}_{i \in I}) \\ &\rightarrow (\Delta_2 \circ \Theta_2) \\ &= \Delta'' \end{aligned}$$

as required

- Rule [TRUE]
Assume

$$\frac{\Vdash \varphi \quad \text{fs}(P_1) = \text{fs}(B)}{[\text{if } \varphi \text{ then } P_1 \text{ else } P_2 \mid B \mid B'] \longrightarrow [P_1 \mid B]} \text{ [TRUE]}$$

with

$$\Gamma; \Delta \vdash [\text{if } \varphi \text{ then } P_1 \text{ else } P_2 \mid B \mid B']$$

and Δ well-type. The derivation for the latter typing judgement is

$$\boxed{
 \begin{array}{c}
 \boxed{
 \frac{\Gamma; \Theta_1 \vdash B \quad \Gamma; \Theta_2 \vdash B'}{\Gamma; \Theta_1, \Theta_2 \vdash B \mid B'} \text{ [BPAR]}
 } \\
 \boxed{
 \begin{array}{c}
 \text{[TCOND]} \\
 \frac{\Gamma \vdash \varphi : \mathbf{bool} \quad \Gamma; \Delta_0, \Delta_1 \vdash P_1 \quad \Gamma; \Delta_0, \Delta_2 \vdash P_2}{i, j \in \{1, 2\} \wedge i \neq j \implies \text{dom}(\Delta_i) = \text{fs}(P_i) - \text{fs}(P_j)} \\
 \Gamma; \Delta_0, \Delta_1, \Delta_2 \vdash \text{if } \varphi \text{ then } P_1 \text{ else } P_2
 \end{array}
 } \\
 \frac{\Gamma; \emptyset \vdash R \quad \Delta = \Delta_0, \Delta_1, \Delta_2 \circ \Theta_1, \Theta_2}{\Gamma; \Delta \vdash [\text{if } \varphi \text{ then } P_1 \text{ else } P_2 \mid B \mid B']} \text{ [TNODE]}
 \end{array}
 }$$

We use the above information to type the result of the reduction $[P_1 \mid B]$:

$$\frac{\Gamma; \Delta_0, \Delta_1 \vdash P_1 \quad \Gamma; \emptyset \vdash R \quad \Gamma; \Theta_1 \vdash B \quad \Delta' = \Delta_0, \Delta_1 \circ \Theta_1}{\Gamma; \Delta' \vdash [P_1 \mid B]} \text{ [TNODE]}$$

The reduction condition requires that $\text{fs}(B) = \text{fs}(P_1)$. From the last equation and the fact that $\text{fs}(B) = \text{dom}(\Theta_1)$ we have that $\text{fs}(P_1) = \text{dom}(\Theta_1)$. Also $\text{fs}(P_1) = \text{dom}(\Delta_0, \Delta_1)$. From the latter two equalities we can conclude that $\Delta_0, \Delta_1 \circ \Theta_1$ is defined and $\text{fs}(P_1) = \text{dom}(\Delta')$. Moreover, $\text{fs}(P_1) \cup \text{fs}(P_2) = \Delta$. Therefore, $\Delta' \subseteq \Delta$ as required.

- Rule [FALSE]. The proof the same with the proof of case [TRUE].
- Rule (Sum).

Assume

$$\frac{[P_1 \mid B] \parallel N \longrightarrow [P' \mid B'] \parallel N'}{[P_1 + P_2 \mid B] \parallel N \longrightarrow [P' \mid B'] \parallel N'} \text{ [NDET]}$$

with

$$\Gamma; \Delta \vdash [P_1 + P_2 \mid B] \parallel N$$

and Δ well-formed.

We produce the derivation for the latter typing judgement.

$$\boxed{
 \begin{array}{c}
 \boxed{
 \begin{array}{c}
 \text{[TSUM]} \\
 \frac{\Gamma; \Delta_1'' \vdash P_1 \quad \Gamma; \Delta_1'' \vdash P_2}{\Gamma; \Delta_1'' \vdash P_1 + P_2} \\
 \frac{\Gamma; \Theta \vdash B \quad \Delta_1' = \Delta_1'' \circ \Theta}{\Gamma; \Delta_1' \vdash [P_1 + P_2 \mid B]} \text{ [TNODE]}
 \end{array}
 } \quad \Delta_1' \leftrightarrow \Delta_0, \Delta_1 \\
 \frac{\Gamma; \Delta_0, \Delta_1 \vdash [P_1 + P_2 \mid B]}{\Gamma; \Delta_0, \Delta_2 \vdash N \quad \Delta_0 \text{ only } s\text{-endpoints} \quad \Delta = \Delta_0, \Delta_1, \Delta_2} \text{ [TSYNCH]} \\
 \frac{\Gamma; \Delta_0, \Delta_2 \vdash N \quad \Delta_0 \text{ only } s\text{-endpoints} \quad \Delta = \Delta_0, \Delta_1, \Delta_2}{\Gamma; \Delta \vdash [P_1 + P_2 \mid B] \parallel N} \text{ [TPAR]}
 \end{array}
 }$$

We can now produce the typing derivation for the left hand side network of the premise $[P_1 | B] \parallel N$:

$$\boxed{
\begin{array}{c}
\boxed{
\begin{array}{c}
\text{[TNode]} \\
\frac{\Gamma; \Delta'_1 \vdash P_1 \quad \Gamma; \Theta \vdash B \quad \Delta'_1 = \Delta''_1 \circ \Theta}{\Gamma; \Delta'_1 \vdash [P_1 | B]} \quad \Delta'_1 \hookrightarrow \Delta_0, \Delta_1 \\
\text{[TSYNCH]} \\
\Gamma; \Delta_0, \Delta_1 \vdash [P_1 | B]
\end{array}
} \\
\frac{\Gamma; \Delta_0, \Delta_2 \vdash N \quad \Delta_0 \text{ only } s\text{-endpoints} \quad \Delta = \Delta_0, \Delta_1, \Delta_2}{\Gamma; \Delta \vdash [P_1 | B] \parallel N} \text{[TPAR]}
\end{array}
}$$

The latter two derivations have the same typing contexts. The result is then immediate from the induction hypothesis; from the fact that type preservation holds on the premise of the reduction, we have that

$$\Gamma; \Delta' \vdash [P' | B'] \parallel N'$$

and $\Delta \rightarrow \Delta'$. The latter results is what needed to conclude type preservation for the reduction.

- Rule [DEF].

Assume

$$\frac{[P | B] \parallel N \longrightarrow [P' | B'] \parallel N'}{[\text{def } \{D_i(\tilde{x}_i) \stackrel{\text{def}}{=} P_i\}_{i \in I} \text{ in } P | B] \parallel N \longrightarrow [\text{def } \{D_i(\tilde{x}_i) \stackrel{\text{def}}{=} P_i\}_{i \in I} \text{ in } P' | B'] \parallel N'} \text{[DEF]}$$

with

$$\Gamma; \Delta \vdash [\text{def } \{D_i(\tilde{x}_i) \stackrel{\text{def}}{=} P_i\}_{i \in I} \text{ in } P | B] \parallel N$$

and Δ well-formed.

We produce the typing derivation of the latter judgement. We begin with the typing derivation for network node $[\text{def } \{D_i(\tilde{x}_i) \stackrel{\text{def}}{=} P_i\}_{i \in I} \text{ in } P | B]$:

$$\boxed{
\begin{array}{c}
\boxed{
\begin{array}{c}
\text{[TREC]} \\
\frac{\forall i \in I, D_i(\tilde{x}_i) : (\Gamma_i, \Delta_i) \in \Gamma \wedge \Gamma, \Gamma_i; \Delta_i \vdash P_i \quad \Gamma; \Delta''_1 \vdash P}{\Gamma; \Delta'_1 \vdash \text{def } \{D_i(\tilde{x}_i) \stackrel{\text{def}}{=} P_i\}_{i \in I} \text{ in } P} \\
\Gamma; \Theta \vdash B \quad \Delta'_1 = \Delta''_1 \circ \Theta \\
\text{[TNode]} \\
\Gamma; \Delta'_1 \vdash [\text{def } \{D_i(\tilde{x}_i) \stackrel{\text{def}}{=} P_i\}_{i \in I} \text{ in } P | B]
\end{array}
} \\
\frac{\Delta'_1 \hookrightarrow \Delta_0, \Delta_1}{\Gamma; \Delta_0, \Delta_1 \vdash [\text{def } \{D_i(\tilde{x}_i) \stackrel{\text{def}}{=} P_i\}_{i \in I} \text{ in } P | B]} \text{[TSYNCH]}
\end{array}
}$$

We type the entire network node $[\text{def } \{D_i(\tilde{x}_i) \stackrel{\text{def}}{=} P_i\}_{i \in I} \text{ in } P | B] \parallel N$:

$$\frac{\Gamma; \Delta_0, \Delta_1 \vdash [\text{def } \{D_i(\tilde{x}_i) \stackrel{\text{def}}{=} P_i\}_{i \in I} \text{ in } P | B] \quad \Gamma; \Delta_0, \Delta_2 \vdash N \quad \Delta_0 \text{ only } s\text{-endpoint} \quad \Delta = \Delta_0, \Delta_1, \Delta_2}{\Gamma; \Delta \vdash [\text{def } \{D_i(\tilde{x}_i) \stackrel{\text{def}}{=} P_i\}_{i \in I} \text{ in } P | B] \parallel N} \text{[TPAR]}$$

We can use the above information to produce the typing derivation of the network term $[P | B] \parallel N$ in the premise of the reduction:

$$\boxed{
 \begin{array}{c}
 \boxed{
 \begin{array}{c}
 \text{[TNODE]} \\
 \frac{\Gamma; \Delta_1'' \vdash P \quad \Gamma; \Theta \vdash B \quad \Delta_1' = \Delta_1'' \circ \Theta}{\Gamma; \Delta_1' \vdash [P | B]} \quad \Delta_1' \hookrightarrow \Delta_0, \Delta_1
 \end{array}
 } \\
 \frac{\Gamma; \Delta_0, \Delta_1 \vdash [P | B]}{\Gamma; \Delta_0, \Delta_1 \vdash [P | B]} \text{[TSYNCH]} \\
 \frac{\Gamma; \Delta_0, \Delta_2 \vdash N \quad \Delta_0 \text{ only } s\text{-endpoint} \quad \Delta = \Delta_0, \Delta_1, \Delta_2}{\Gamma; \Delta \vdash [\text{def } P | B \text{ in } \parallel N]} \text{[TPAR]}
 \end{array}
 }$$

The results then comes from the induction hypothesis; the latter typing derivation and the reduction on the premise of the case implies typing derivation $\Gamma; \Delta' \vdash [\text{def } P' | B' \text{ in } \parallel N']$ and $\Delta \rightarrow \Delta'$, which concludes the case.

- Rule [TPAR].

Assume

$$\frac{N \longrightarrow N'}{N \parallel M \longrightarrow N' \parallel M} \text{[RPAR]}$$

with

$$\Gamma; \Delta \vdash N \parallel M$$

and Δ well-formed.

We provide with the typing derivation for the latter typing judgement.

$$\frac{\Gamma; \Delta_0, \Delta_1 \vdash N \quad \Gamma; \Delta_0, \Delta_2 \vdash M \quad \text{dom}(\Delta_0) \text{ only } s\text{-endpoints} \quad \Delta = \Delta_0, \Delta_1, \Delta_2}{\Gamma; \Delta \vdash N \parallel M} \text{[TPAR]}$$

From the fact that $\Gamma; \Delta_0, \Delta_1 \vdash N$, the reduction, and the induction hypothesis we can conclude that $\Gamma; \Delta' \vdash N$ and $\Delta_0, \Delta_1 \rightarrow \Delta'$. It remains to type the network term $N \parallel M$ and show the Typing Preservation requirements for Δ' . This can be analysed in several sub-cases.

- $\Delta' = \Delta_0', \Delta_1'$ with $\Delta_0' = \Delta_0, \{s_i : (c_i, T_i)\}_{i \in I}$ and $\Delta_1 = \Delta_1', \{s_j : (c_j, T_j)\}_{j \in J}$. This implies derivation:

$$\boxed{
 \begin{array}{c}
 \Gamma; \Delta_0', \Delta_1' \vdash N' \quad \frac{\Delta_0 = \Delta_0', \Delta_0'' \quad \Gamma; \Delta_0, \Delta_2 \vdash M}{\Gamma; \Delta_0', \Delta_0'', \Delta_2 \vdash M} \\
 \text{dom}(\Delta_0) \text{ only } s\text{-endpoints and } \Delta_0' \subseteq \Delta_0 \text{ implies} \\
 \text{dom}(\Delta_0') \text{ only } s\text{-endpoints} \\
 \frac{\Gamma; \Delta_0', \Delta_0'', \Delta_1', \Delta_2 \vdash N \parallel M}{\Gamma; \Delta_0', \Delta_0'', \Delta_1', \Delta_2 \vdash N \parallel M} \text{[TPAR]}
 \end{array}
 }$$

as required, because $\Delta = \Delta_0, \Delta_1, \Delta_2 \rightarrow \Delta_0', \Delta_0'', \Delta_1', \Delta_2 = \Delta_0, \Delta_1', \Delta_2$

– $\Delta' = \Delta_0, \Delta'_1$ and $\Delta_1 \rightarrow \Delta'_1$ implies:

$$\frac{\Gamma; \Delta_0, \Delta'_1 \vdash N' \quad \Gamma; \Delta_0, \Delta_2 \vdash M \quad \text{dom}(\Delta_0) \text{ only } s\text{-endpoints}}{\Gamma; \Delta_0, \Delta'_1, \Delta_2 \vdash N \parallel M} \text{ [TPAR]}$$

as required, because $\Delta = \Delta_0, \Delta_1, \Delta_2 \rightarrow \Delta_0, \Delta'_1, \Delta_2$.

– $\Delta' = \Delta'_0, \Delta'_1$ with $\Delta_0, \Delta_1 \rightarrow \Delta'_0, \Delta'_1$. The latter assumption, the fact that $\text{dom}(\Delta_0)$ contains only s -endpoints, and Linear Context Advancement (Definition 4.2) imply that $\Delta_0 = \Delta''_0, s : (c, T)$ and $\Delta_1 = \Delta''_1, \check{s} : (c, \overline{T})$ and furthermore $\Delta_0 = \Delta''_0, s : (c+1, T')$ and $\Delta'_1 = \Delta''_1, \check{s} : (c+1, \overline{T'})$. Note that $T \rightarrow T'$. These last equations allows to type the resulting network of the reduction:

$$\boxed{\begin{array}{c} \boxed{\frac{\Gamma; \Delta_0, \Delta_2 \vdash M \quad \Delta_0 = \Delta''_0, s : (c, T)}{\Gamma; \Delta''_0, s : (c, T), \Delta_2 \vdash M}} \\ \boxed{\frac{\Delta''_0, \Delta_2 \hookrightarrow \Delta''_0, \Delta_2 \quad T \rightarrow^{c+1-c} T'}{\Delta''_0, s : (c, T), \Delta_2 \hookrightarrow \Delta''_0, s : (c+1, T'), \Delta_2}} \\ \hline \Gamma; \Delta''_0, s : (c+1, T'), \Delta_2 \vdash M \quad \text{[TSYNCH]} \\ \Gamma; \Delta''_0, s : (c+1, T'), \Delta''_1, \check{s} : (c+1, \overline{T'}) \vdash N' \\ \text{dom}(\Delta''_0, s : (c+1, T')) \text{ only } s\text{-endpoints} \\ \hline \Gamma; \Delta''_0, s : (c+1, T'), \Delta''_1, \check{s} : (c+1, \overline{T'}), \Delta_2 \vdash N \parallel M \quad \text{[TPAR]} \end{array}}$$

as required, because $\Delta = \Delta_0, \Delta_1, \Delta_2 \rightarrow \Delta'_0, \Delta'_1, \Delta_2 = \Delta''_0, s : (c+1, T'), \Delta''_1, \check{s} : (c+1, \overline{T'}), \Delta_2$.

– $\Delta = \Delta'_0, \Delta_1$ and $\Delta_0 \rightarrow \Delta'_0$ cannot happen because $\text{dom}(\Delta_0)$ has only s -endpoints.

- Rule [RCONG]. The result is immediate by applying Congruence Invariant (Lemma A.1) on the induction hypothesis.
- Rule [RRES].

Assume

$$\frac{N \longrightarrow N'}{(\nu n)N \longrightarrow (\nu n)N'} \text{ [RRES]}$$

with

$$\Gamma; \Delta \vdash (\nu n)N$$

and Δ well-formed.

There are two sub-cases for the typing derivation of the latter judgement:

– $n = a$. Restrict shared name. In this sub-case the typing derivation for network term $(\nu a)N$ is as:

$$\frac{\Gamma, a : T; \Delta \vdash N}{\Gamma; \Delta \vdash (\nu a)N} \text{ [TCRES]}$$

From $\Gamma, a : T; \Delta \vdash N$, well-formed Δ , and the induction hypothesis we conclude that $\Gamma, a : T; \Delta' \vdash N'$ and either $\Delta' \subseteq \Delta$ or $\Delta \rightarrow \Delta'$. From the last typing judgement we can

type the result of the reduction $(\nu a)N$:

$$\frac{\Gamma, a : T; \Delta' \vdash N'}{\Gamma; \Delta' \vdash N'} \text{ [TCRES]}$$

to conclude the sub-case.

- $n = s$. Restrict session name. In this sub-case the typing derivation for network term $(\nu s)N$ is as:

$$\frac{\Gamma; \Delta, \Delta''', \check{s} : (c, T) \vdash N \quad \Delta''' = s : (c, \bar{T}) \vee \Delta''' = \emptyset}{\Gamma; \Delta \vdash (\nu s)N} \text{ [TSRES]}$$

Typing context $\Delta, \check{s} : (c, T)$ is well-formed because Δ is well-formed, see well-formedness Definition 4.1.

From $\Gamma; \Delta, \check{s} : (c, T), s : (c, \bar{T}) \vdash N$, well-formed $\Delta, \check{s} : (c, T), s : (c, \bar{T})$, and the induction hypothesis we conclude that $\Gamma; \Delta' \vdash N'$ and $\Delta, \check{s} : (c, T), s : (c, \bar{T}) \rightarrow \Delta'$. All it remains is to conclude the case, is the typing derivation of the result of the reduction $(\nu s)N'$.

If $\Delta, \check{s} : (c, T), s : (c, \bar{T}) \rightarrow \Delta'$ then there are two sub-cases:

- * $\Delta' = \Delta'', \check{s} : (c, T), s : (c, \bar{T})$ with $\Delta \rightarrow \Delta''$, which implies

$$\frac{\Gamma; \Delta'', \check{s} : (c, T), s : (c, \bar{T}) \vdash N'}{\Gamma; \Delta'' \vdash (\nu s)N'} \text{ [TSRES]}$$

as required.

- * $\Delta' = \Delta, \check{s} : (c, T), s : (c, \bar{T})$ with $\check{s} : (c, T), s : (c, \bar{T}) \rightarrow \check{s} : (c + 1, T_1), s : (c + 1, T_2)$, which implies

$$\frac{\begin{array}{l} \Gamma; \Delta, \check{s} : (c + 1, T'), s : (c + 1, \bar{T}') \vdash (\nu n)N' \\ \text{From the definition of } \rightarrow \text{ we get that } T'_1 = \bar{T}'_2 \end{array}}{\Gamma; \Delta \vdash (\nu s)N'} \text{ [TSRES]}$$

as required. □