

A Typing Discipline for Statically Verified Crash Failure Handling in Distributed Systems [★]

Malte Viering¹(✉), Tzu-Chun Chen¹, Patrick Eugster^{1,2,3},
Raymond Hu⁴, and Lukasz Ziarek⁵

¹ Department of Computer Science, TU Darmstadt, Darmstadt, Germany
`viering@dsp.tu-darmstadt.de`

² Faculty of Informatics, Università della Svizzera italiana, Lugano, Switzerland

³ Department of Computer Science, Purdue University, West Lafayette, USA

⁴ Department of Computing, Imperial College London, London, UK

⁵ Department of Computer Science and Engineering, SUNY Buffalo, Buffalo, USA

Abstract. A key requirement for many distributed systems is to be resilient toward partial failures, allowing a system to progress despite the failure of some components. This makes programming of such systems daunting, particularly in regards to avoiding inconsistencies due to failures and asynchrony. This work introduces a formal model for crash failure handling in asynchronous distributed systems featuring a lightweight coordinator, modeled in the image of widely used systems such as ZooKeeper and Chubby. We develop a typing discipline based on multiparty session types for this model that supports the specification and static verification of multiparty protocols with explicit failure handling. We show that our type system ensures subject reduction and progress in the presence of failures. In other words, in a well-typed system even if some participants crash during execution, the system is guaranteed to progress in a consistent manner with the remaining participants.

1 Introduction

Distributed programs, partial failures, and coordination. Developing programs that execute across a set of physically remote, networked processes is challenging. The correct operation of a *distributed program* requires correctly designed protocols by which concurrent processes interact asynchronously, and correctly implemented processes according to their roles in the protocols. This becomes particularly challenging when distributed programs have to be resilient to *partial failures*, where some processes crashes while others remain operational. Partial failures affect both *safety* and *liveness* of applications. Asynchrony is the key issue, resulting in the inability to distinguish slow processes from failed ones. In general, this makes it impossible for processes to reach agreement, even when only a single process can crash [19].

[★] Financially supported by ERC grant FP7-617805 “LiVeSoft - Lightweight Verification of Software” and NSF grants CNS-1405614 and IIS-1617586.

In practice, such impasses are overcome by making appropriate assumptions for the considered infrastructure and applications. One common approach is to assume the presence of a highly available *coordination service* [26] – realized using a set of replicated processes large enough to survive common rates of process failures (e.g., 1 out of 3, 2 out of 5) – and delegating critical decisions to this service. While this *coordinator model* has been in widespread use for many years (cf. *consensus service* [22]), the advent of cloud computing has recently brought it further into the mainstream, via instances like Chubby [4] and ZooKeeper [26]. Such systems are used not only by end applications but also by a variety of frameworks and middleware systems across the layers of the protocol stack [11,20,31,41].

Typing disciplines for distributed programs. Typing disciplines for distributed programs is a promising and active research area towards addressing the challenges in the correct development of distributed programs. See Hüttel et al. [27] for a broad survey. *Session types* are one of the established typing disciplines for message passing systems. Originally developed in the π -calculus [23], these have been later successfully applied to a range of practical languages, e.g., Java [25,42], Scala [40], Haskell [35,39], and OCaml [28,38]. *Multiparty session types* (MPSTs) [15,24] generalize session types beyond two participants. In a nutshell, a standard MPST framework takes (1) a specification of the whole multiparty message protocol as a *global type*; from which (2) *local types*, describing the protocol from the perspective of each participant, are derived; these are in turn used to (3) statically *type check* the I/O actions of endpoint programs implementing the session participants. A well-typed system of session endpoint programs enjoys important safety and liveness properties, such as *no reception errors* (only expected messages are received) and *session progress*. A basic intuition behind MPSTs is that the design (i.e., restrictions) of the type language constitutes a class of distributed protocols for which these properties can be statically guaranteed by the type system.

Unfortunately, *no* MPST work supports protocols for asynchronous distributed programs dealing with *partial failures due to process crashes*, so the aforementioned properties no longer hold in such an event. Several MPST works have treated communication patterns based on *exception messages* (or *interrupts*) [6,7,16]. In these works, such messages may convey exceptional states in an *application* sense; from a protocol compliance perspective, however, these messages are the same as any other message communicated during a *normal* execution of the session. This is in contrast to *process* failures, which may invalidate already in-transit (*orphan*) messages, and where the task of agreeing on the concerted handling of a crash failure is itself prone to such failures.

Outside of session types and other type-based approaches, there have been a number of advances on verifying fault tolerant distributed protocols and applications (e.g., based on model checking [29], proof assistants [44]); however, little work exists on providing direct compile-time support for *programming* such applications in the spirit of MPSTs.

Contributions and challenges. This paper puts forward a new typing discipline for safe specification and implementation of distributed programs prone to process crash failures based on MPSTs. The following summarizes the key challenges and contributions.

Multiparty session calculus with coordination service. We develop an extended multiparty session calculus as a formal model of processes prone to crash failures in asynchronous message passing systems. Unlike standard session calculi that reflect only “minimal” networking infrastructures, our model introduces a practically-motivated *coordinator* artifact and explicit, asynchronous messages for run-time crash notifications and failure handling.

MPSTs with explicit failure handling. We introduce new global and local type constructs for *explicit failure handling*, designed for specifying protocols tolerating partial failures. Our type system carefully reworks many of the key elements in standard MPSTs to manage the intricacies of handling crash failures. These include the well-formedness of failure-prone global types, and the crucial *coherence* invariant on MPST typing environments to reflect the notion of system consistency in the presence of crash failures and the resulting errors. We show safety and progress for a well-typed MPST session despite potential failures.

To fit our model to practice, we introduce programming constructs similar to well-known and intuitive exception handling mechanisms, for handling concurrent and asynchronous process crash failures in sessions. These constructs serve to integrate user-level session control flow in endpoint processes and the underlying communications with the coordination service, used by the target applications of our work to outsource critical failure management decisions (see Fig. 1). It is important to note that the coordinator does *not* magically solve all problems. Key design challenges are to ensure that communication with it is fully asynchronous as in real-life, and that it is involved only in a “minimal” fashion. Thus we treat the coordinator as a first-class, asynchronous network artifact, as opposed to a convenient but impractical global “oracle” (cf. [6]), and our operational semantics of multiparty sessions remains primarily *choreographic* in the original spirit of distributed MPSTs, unlike works that resort to a centralized *orchestrator* to conduct all actions [5,8]. As depicted in Fig. 1, application-specific communication does not involve the coordinator. Our model lends itself to common practical scenarios where processes monitor each other in a peer-based fashion to detect failures, and rely on a coordinator only to establish agreement on which processes have failed, and when.

We also developed a prototype implementation in Scala, that uses ZooKeeper for coordination. Due to space limitations we present it in the appendix.

Example. As a motivating example, Fig. 2 gives a global formal specification for a big data streaming task between a distributed file system (DFS) dfs , and two workers $w_{1,2}$. The DFS streams data to two workers, which process the data and




Fig. 1. Coordinator model for asynchronous distributed systems. The coordinator is implemented by replicated processes (internals omitted).

$$\begin{aligned}
 [dfs]G = & \mathbf{t}(\mu t. \\
 & dfs \rightarrow w_1 \ l_{d_1}(S).dfs \rightarrow w_2 \ l_{d_2}(S). \\
 & w_1 \rightarrow dfs \ l_{r_1}(S').w_2 \rightarrow dfs \ l_{r_2}(S').t \\
 &)\mathbf{h}(\\
 & \{w_1\} : \mu t'.dfs \rightarrow w_2 \ l'_{d_1}(S). \\
 & \quad w_2 \rightarrow dfs \ l'_{r_1}(S').t', \\
 & \{w_2\} : \dots, \{w_1, w_2\} : \mathbf{end}
 \end{aligned}$$

Fig. 2. Global type for a big data streaming task with failure handling capabilities.

write the result back. Most DFSs have built-in fault tolerance mechanisms [20], so we consider *dfs* to be *robust*, denoted by the annotation $[dfs]$; the workers, however, may individually fail. In the *try-handle* construct $\mathbf{t}(\dots)\mathbf{h}(\dots)$, the *try-block* $\mathbf{t}(\dots)$ gives the *normal* (i.e., failure-free) flow of the protocol, and $\mathbf{h}(\dots)$ contains the explicit *handlers* for potential crashes. In the try-block, the workers receive data from the DFS ($dfs \rightarrow w_i$), perform local computations, and send back the result ($w_i \rightarrow dfs$). If a worker crashes ($\{w_i\} : \dots$), the other worker will also take over the computation of the crashed worker, allowing the system to still produce a valid result. If both workers crash (by any interleaving of their concurrent crash events), the global type specifies that the DFS should safely terminate its role in the session.

We shall refer to this basic example, that focuses on the new failure handling constructs, in explanations in later sections. See App. A.1 for examples of larger protocols featuring multiparty choices and recursion with explicit failure handling. We also give many further examples throughout the following sections to illustrate the potential session errors due to failures exposed by our model, and how our framework resolves them to recover MPST safety and progress.

Roadmap. Sec. 2 describes the adopted system and failure model. Sec. 3 introduces global types for guiding failure handling. Sec. 4 introduces our process calculus with failure handling capabilities and a coordinator. Sec. 5 introduces local types, derived from global types by projection. Sec. 6 describes typing rules, and defines *coherence* of session environments with respect to endpoint crashes. Sec. 7 states properties of our model. Sec. 8 discusses related work. Sec. 9 draws conclusions. App. B contains further formal definitions. App. A contains additional examples. App. C provides details of our prototype implementation in Scala using ZooKeeper. App. D gives full proofs of properties.

2 System and Failure Model

In distributed systems care is required to avoid partial failures affecting liveness (e.g., waiting on messages from crashed processes) or safety (e.g., when processes

manage to communicate with some peers but not others before crashing) properties of applications. Based on the nature of the infrastructure and application, appropriate *system and failure models* are chosen along with judiciously made assumptions to overcome such impasses in practice.

We pinpoint the key characteristics of our model, according to our practical motivations and standard distributed systems literature, that shape the design choices we make later for the process calculus and types. As it is common we augment our system with a *failure detector* (FD) to allow for distinguishing slow and failed processes. The advantage of the FD (1) in terms of reasoning is that it concentrates all assumptions to solve given problems and (2) implementation-wise it yields a single main module where time-outs are set and used.

Concretely we make the following assumptions on failures and the system:

- (1) **Crash-stop failures:** Application processes fail by crashing (halting), and do not recover.
- (2) **Asynchronous system:** Application processes and the network are asynchronous, meaning that there are no upper bounds on processes' relative speeds or message transmission delays.
- (3) **Reliable communication:** Messages transmitted between correct (i.e., non-failed) participants are eventually received.
- (4) **Robust coordinator:** The coordinator (coordination service) is permanently available.
- (5) **Asynchronous reliable failure detection:** Application processes have access to local FDs which eventually detect all failed peers and do not falsely suspect peers.

(1)–(3) are standard in literature on fault-tolerant distributed systems [19].

Note that processes can still recover but will not do so *within* sessions (or will not be re-considered for those). Other failure models, e.g., network partitions [21] or Byzantine failures [33], are subject of future work. The former are not tolerated by ZooKeeper et al., and the latter have often been argued to be a too generic failure model (e.g., [3]).

The assumption on the coordinator (4) implicitly means that the number of concomitant failures among the coordinator replicas is assumed to remain within a minority, and that failed replicas are replaced in time (to tolerate further failures). Without loss of validity, the coordinator internals can be treated as a blackbox (e.g., ZooKeeper uses a variant of Paxos [32]). The final assumption (5) on failure detection is backed in practice by the concept of *program-controlled* crash [10], which consists in communicating decisions to disregard supposedly failed processes also to those processes, prompting them to reset themselves upon false suspicion. In practice systems can be configured to minimize the probability of such events, and by a “two-level” membership consisting in evicting processes from *individual* sessions (cf. recovery above) more quickly than from a system as a whole; several authors have also proposed network support to entirely avoid false suspicions (e.g., [34]).

These assumptions do not make handling of failures trivial, let alone mask them. For instance, the network can arbitrarily delay messages and thus reorder

(Basic type) $S ::= \text{bool} \mid \text{str} \mid \text{int}$
 (Global type) $G ::= p \rightarrow q\{l_i(S_i).G_i\}_{i \in I} \mid \mu t.G \mid t \mid \text{end} \mid \mathbf{t}(G_1)\mathbf{h}(H)^\kappa.G_2$
 (Handling env.) $H ::= F:G \mid H, H$ (Handler sig.) $F ::= \{p_i\}_{i \in I}$

Fig. 3. Syntax of global types with explicit handling of partial failures.

them with respect to their real sending times, and (so) different processes can detect failures at different points in time and in different orders.

3 Global Types for Explicit Handling of Partial Failures

Based on the foundations of MPSTs, we develop *global types* to formalize specifications of distributed protocols with explicit handling of *partial failures due to role crashes*, simply referred to as *failures*. We present global types before introducing the process calculus to provide a high-level intuition of how failure handling works in our model.

The syntax of *global types* is depicted in Fig. 3. We use the following base notations: p, q, \dots for *role* (i.e., participant) names; l_1, l_2, \dots for message *labels*; and t, t', \dots for type variables. *Base types* S may range over, bool, int , etc.

Global types are denoted by G . We first summarize the constructs from standard MPST [15,24]. A *branch* type $p \rightarrow q\{l_i(S_i).G_i\}_{i \in I}$ means that p can send to q *one* of the messages of type S_k with label l_k , where k is a member of the non-empty index set I . The protocol then proceeds according to the continuation G_k . When I is a singleton, we may simply write $p \rightarrow q l(S).G$. We use t for type variables and take an equi-recursive view, i.e., $\mu t.G$ and its unfolding $[\mu t.G/t]$ are equivalent. We assume type variable occurrences are bound and guarded (e.g., $\mu t.t$ is not permitted). end is for termination.

We now introduce our extensions for partial failure handling. A *try-handle* $\mathbf{t}(G_1)\mathbf{h}(H)^\kappa.G_2$ describes a “failure-atomic” protocol unit: all *live* (i.e., non-crashed) roles will eventually reach a consistent protocol state, despite any concurrent and asynchronous role crashes. The try-block G_1 defines the *default* protocol flow, and H is a *handling environment*. Each element of H maps a *handler signature* F , that specifies a set of *failed* roles $\{p_i\}_{i \in I}$, to a *handler body* specified by a G . The handler body G specifies how the live roles should proceed given the failure of roles F . The protocol then proceeds (for live roles) according to the continuation G_2 after the default block G_1 or failure handling defined in H has been completed as appropriate.

To simplify later technical developments, we annotate each try-handle term in a given G by a unique $\kappa \in \mathbb{N}$ that lexically identifies the term within G . These annotations may be assigned mechanically. As a short hand, we refer to the try-block and handling environment of a particular try-handle by its annotation; e.g., we use κ to stand for $\mathbf{t}(G_1)\mathbf{h}(H)^\kappa$. In the running examples (e.g., Fig. 2), if there exists only one try-handle, we omit κ for simplicity.

Top-level global types and robust roles. We use the term *top-level* global type to mean the source protocol specified by a user, following a typical top-down interpretation of MPST frameworks [15,24]. We allow top-level global types to be optionally annotated $[\bar{p}]G$, where $[\bar{p}]$ specifies a set of *robust* roles—i.e., roles that can be assumed to never fail. In practice, a participant may be robust if it is replicated or is made inherently fault tolerant by other means (e.g., the participant that represents the distributed file system in Fig. 2).

Well-formedness. The first stage of validation in standard MPSTs is to check that the top-level global type satisfies the supporting criteria used to ensure the desired properties of the type system. We first list basic syntactic conditions which we assume on any given G : (i) each F is non-empty; (ii) a role in a F cannot occur in the corresponding handler body (a failed role cannot be involved in the handling of its own failure); and (iii) every occurrence of a non-robust role p must be contained within a, possibly outer, try-handle that has a handler signature $\{p\}$ (the protocol must be able to handle its potential failure). Lastly, to simplify the presentation without loss of generality, we impose that separate branch types *not* defined in the same default block or handler body must have disjoint label sets. This can be implicitly achieved by combining label names with try-handle annotations.

Assuming the above, we define *well-formedness* for our extended global types. We write $G' \in G$ to mean that G' syntactically occurs in G (\in is reflexive); similarly for the variations $\kappa \in G$ and $\kappa \in \kappa'$. Recall κ is shorthand for $\mathbf{t}(G_1)\mathbf{h}(H)^\kappa$. We use a lookup function $outer_G(\kappa)$ for the set of all try-handles in G that enclose a given κ , including κ itself, defined by $outer_G(\kappa) = \{\kappa' \mid \kappa \in \kappa' \wedge \kappa' \in G\}$. Full definitions for all the above are presented in App. B.1.

Definition 1 (Well-formedness). Let κ stand for $\mathbf{t}(G_1)\mathbf{h}(H)^\kappa$, and κ' for $\mathbf{t}(G'_1)\mathbf{h}(H')^{\kappa'}$. A global type G is *well-formed* if both of the following conditions hold. For all $\kappa \in G$:

1. $\forall F_1 \in dom(H). \forall F_2 \in dom(H). \exists \kappa' \in outer_G(\kappa)$ s.t. $F_1 \cup F_2 \in dom(H')$
2. $\nexists F \in dom(H). \exists \kappa' \in outer_G(\kappa). \exists F' \in dom(H')$ s.t. $\kappa' \neq \kappa \wedge F' \subseteq F$

The first condition asserts that for any two separate handler signatures of a handling environment of κ , there always exists a handler whose handler signature matches the union of their respective failure sets – this handler is either inside the handling environment of κ itself, or in the handling environment of an outer try-handle. This ensures that if roles are active in different handlers of the same try-handle then there is a handler whose signature corresponds to the union over the signatures of those different handlers. Example 2 together with Example 3 in Sec. 4 illustrate a case where this condition is needed. The second condition asserts that if the handling environment of a try-handle contains a handler for F , then there is no outer try-handle with a handler for F' such that $F' \subseteq F$. The reason for this condition is that in the case of *nested* try-handles, our communication model allows separate try-handles to start failure handling independently (the operational semantics will be detailed in the next section; see




Fig. 4. Challenges under pure asynchronous interactions with a coordinator. Between time (1) and time (2), the task $\phi = (\kappa, \emptyset)$ is interrupted by the crash of P_a . Between time (3) and time (4), due to asynchrony and multiple crashes, P_c starts handling the crash of $\{P_a, P_d\}$ without handling the crash of $\{P_a\}$. Finally after (4) P_b and P_c finish their common task.

(**TryHdl**) in Fig. 6). The aim is to have the relevant roles eventually converge on performing the handling of the outermost try-handle, possibly by interrupting the handling of an inner try-handle. Consider the following example:

Example 1. $G = t(t(G')h(\{p_1, p_2\} : G_1)^2)h(\{p_1\} : G_1)^1$ violates condition 2 because, when p_1 and p_2 both failed, the handler signature $\{p_1\}$ will still be triggered (i.e., the outer try-handle will eventually take over). It is not sensible to run G'_1 instead of G_1 (which is for the crashes of p_1 and p_2).

App. B.1 gives examples of well-formed and ill-formed global types, and further elaborates on issues with the latter.

4 A Process Calculus for Coordinator-based Failure Handling

Fig. 4 depicts a scenario that can occur in practical asynchronous systems with coordinator-based failure handling through frameworks such as ZooKeeper (Sec. 2). Using this scenario, we first illustrate challenges, formally define our model, and then develop a safe type system.

The scenario corresponds to a global type of the form $t(G)h(\{P_a\} : G_a, \{P_a, P_d\} : G_{ad}, \dots)^\kappa$, with processes $P_{a..d}$ and a coordinator Ψ . We define a *task* to mean a unit of interactions, which includes failure handling behaviors. Initially all processes are collaborating on a task ϕ , which we label (κ, \emptyset) (identifying the task context, and the set of failed processes). The shaded boxes signify which task each process is working on. Dotted arrows represent notifications between

| | | | |
|---------------|--|---------------|------------------------------------|
| (Expression) | $e ::= v \mid x \mid e + e \mid -e \mid \dots$ | (Channel) | $c ::= s[p] \mid y$ |
| (Process) | $P ::= a[p](y).P \mid c : \eta$ | (Level) | $\phi ::= (\kappa, F)$ |
| (Statement) | $\eta ::= t(\eta)h(\mathbf{H})^\phi.\eta \mid \underline{0} \mid 0 \mid p! l(e).\eta$ $\mid p?\{l_i(x_i).\eta_i\}_{i \in I} \mid X(e)$ $\mid \text{def } D \text{ in } \eta \mid \text{if } e \eta \text{ else } \eta$ | (Declaration) | $D ::= X(x) = \eta$ |
| (Application) | $N ::= P \mid N \mid N \mid s : h$ | (Queue) | $h ::= \emptyset \mid h \cdot m$ |
| (Message) | $m ::= \langle p, q, l(v) \rangle \mid \langle p, \text{crash } F \rangle \mid dn$ | (Done) | $dn ::= \langle p, q \rangle^\phi$ |
| (System) | $\mathcal{S} ::= \Psi \blacklozenge N \mid (\nu s)\mathcal{S} \mid \mathcal{S} \mid \mathcal{S}$ | (Coordinator) | $\Psi ::= G : (F, d)$ |
| (Context) | $E ::= t(E)h(\mathbf{H})^\phi.\eta \mid \text{def } D \text{ in } E \mid []$ | (Done Queue) | $d ::= \emptyset \mid d \cdot dn$ |

Fig. 5. Grammar for processes, applications, systems, and evaluation contexts.

processes and Ψ related to task completion, and solid arrows for failure notifications from Ψ to processes. During the scenario, P_a first fails, then P_d fails: the execution proceeds through failure handling for $\{P_a\}$ and $\{P_a, P_d\}$.

- (I) When P_b reaches the end of its part in ϕ , the application has P_b notify Ψ . P_b then remains in the context of ϕ (the continuation of the box after notifying) in consideration of other non-robust participants still working on ϕ — P_b may yet need to handle their potential failure(s).
- (II) The processes of synchronizing on the completion of a task or performing failure handling *are themselves subject to failures* that may arise concurrently. In Fig. 4, all processes reach the end of ϕ (i.e., four dotted arrows from ϕ), but P_a fails. Ψ determines this failure and it initiates failure handling at time (1), while *done* notifications for ϕ continue to arrive asynchronously at time (2). The failure handling for crash of P_a is itself interrupted by the second failure at time (3).
- (III) Ψ can receive notifications that are no longer relevant. For example, at time (2), Ψ has received all *done* notifications for ϕ , but the failure of P_a has already triggered failure handling from time (1).
- (IV) Due to multiple concurrent failures, interacting participants may end up in different tasks: around time (2), P_b and P_d are in task $\phi' = (\kappa, \{P_a\})$, whereas P_c is still in ϕ (and asynchronously sending or receiving messages with the others). Moreover, P_c never executes ϕ' because of delayed notifications, so it goes from ϕ directly to $(\kappa, \{P_a, P_d\})$.

Processes. Fig. 5 defines the grammar of processes and (distributed) applications. Expressions e, e_i, \dots can be values v, v_i, \dots , variables x, x_i, \dots , and standard operations. (Application) processes are denoted by P, P_i, \dots . An initialization $a[p](y).P$ agrees to play role p via shared name a and takes actions defined in P ; actions are executed on a session channel $c : \eta$, where c ranges over $s[p]$ (session name and role name) and session variables y ; η represents action statements.

A try-handle $t(\eta)h(\mathbf{H})^\phi$ attempts to execute the local action η , and can handle failures occurring therein as defined in the handling environment \mathbf{H} , analogously to global types. \mathbf{H} thus also maps a handler signature F to a handler body η defining how to handle F . Annotation $\phi = (\kappa, F)$ is composed of two elements:

an identity κ of a *global* try-handle, and an indication of the *current* handler signature which can be empty. $F = \emptyset$ means that the default try-block is executing, whereas $F \neq \emptyset$ means that the handler body for F is executing. Term $\underline{0}$ only occurs in a try-handle during runtime. It denotes a *yielding* for a *notification* from a *coordinator* (introduced shortly).

Other statements are similar to those defined in [15,24]. Term 0 represents an *idle* action. For convention, we omit 0 at the end of a statement. Action $p! l(e).\eta$ represents a sending action that sends p a label l with content e , then it continues as η . Branching $p?\{l_i(x_i).\eta_i\}_{i \in I}$ represents a receiving action from p with several possible branches. When label l_k is selected, the transmitted value v is saved in x_k , and $\eta_k\{v/x_k\}$ continues. For convenience, when there is only one branch, the curly brackets are omitted, e.g., $c : p?l(x).P$ means there is only one branch $l(x)$. $X(e)$ is for a statement variable with one parameter e , and $\text{def } D \text{ in } \eta$ is for recursion, where declaration D defines the recursive body that can be called in η . The conditional statement is standard.

The structure of processes ensures that failure handling is not interleaved between different sessions. However, we note that in standard MPSTs [15,24], session interleaving must anyway be prohibited for the basic progress property. Since our aim will be to show progress, we disallow session interleaving within process bodies. Our model does allow parallel sessions at the top-level, whose actions may be concurrently interleaved during execution.

(Distributed) systems. A (distributed) *system* in our programming framework is a composition of an application, which contains more than one process, and a coordinator (cf. Fig. 1). A system can be running within a private session \mathbf{s} , represented by $(\nu \mathbf{s})\mathcal{S}$, or $\mathcal{S} \mid \mathcal{S}'$ for systems running in different sessions independently and in parallel (i.e., no session interleaving). The job of the coordinator is to ensure that even in the presence of failures there is consensus on whether all participants in a given try-handle completed their local actions, or whether failures need to be handled, and which ones. We use $\Psi = G : (F, d)$ to denote a (robust) coordinator for the global type G , which stores in (F, d) the failures F that occurred in the application, and in d done notifications sent to the coordinator. The coordinator is denoted by ψ when viewed as a role.

A (distributed) *application*⁶ is a process P , a parallel composition $N \mid N'$, or a global queue carrying messages $\mathbf{s} : h$. A global queue $\mathbf{s} : h$ carries a sequence of messages m , sent by participants in session \mathbf{s} . A message is either a regular message $\langle p, q, l(v) \rangle$ with label l and content v sent from p to q or a *notification*. A notification may contain the role of a coordinator. There are *done* and *failure* notifications with two kinds of done notifications dn used for coordination: $\langle p, \psi \rangle^\phi$ notifies ψ that p has finished its local actions of the try-handle ϕ ; $\langle \psi, p \rangle^\phi$ is sent from ψ to notify p that ψ has received all done notifications for the try-handle ϕ so that p shall end its current try-handle and

⁶ Other works use the term *network* which is the reason why we use N instead of, e.g., A . We call it application to avoid confusion with the physical network which interconnects all processes as well as the coordinator.

$$\begin{array}{c}
a[p_I](y_1).P_1 \mid \dots \mid a[p_n](y_n).P_n \rightarrow \\
(\nu s)(G : (\emptyset, \emptyset) \blacklozenge P_1 \{s[p_I]/y_1\} \mid \dots \mid P_n \{s[p_n]/y_n\} \mid s : \emptyset) \quad a : G \quad \mathbf{(Link)} \\
s[p] : E[q! l(e).\eta] \mid s : h \rightarrow s[p] : E[\eta] \mid s : h \cdot \langle p, q, l(v) \rangle \quad e \Downarrow v \quad \mathbf{(Snd)} \\
s[p] : E[q? \{l_i(x_i).\eta_i\}_{i \in I}] \mid s : \langle q, p, l_k(v_k) \rangle \cdot h \rightarrow \\
s[p] : E[\eta_k \{v_k/x_k\}] \mid s : h \quad k \in I \quad \mathbf{(Rcv)} \\
s[p] : E[\mathbf{def} X(x) = \eta \mathbf{in} X\langle e \rangle] \rightarrow s[p] : E[\mathbf{def} X(x) = \eta \mathbf{in} \eta\{v/x\}] \quad e \Downarrow v \quad \mathbf{(Rec)} \\
\frac{N_1 \equiv N_3 \rightarrow N_4 \equiv N_2}{N_1 \rightarrow N_2} \quad \frac{N_1 \rightarrow N_2}{N_1 \mid N \rightarrow N_2 \mid N} \quad \mathbf{(Str, Par)} \\
\frac{N_1 \rightarrow N_2}{\psi \blacklozenge N_1 \rightarrow \psi \blacklozenge N_2} \quad \frac{S \rightarrow S'}{(\nu s)S \rightarrow (\nu s)S'} \quad \mathbf{(Sys, New)} \\
N \mid s : h \rightarrow N \setminus s[p] : \eta \mid s : \mathit{remove}(h, p) \cdot \langle \psi, \mathit{crash} \{p\} \rangle \\
s[p] : \eta \text{ non-robust} \quad \mathbf{(Crash)}
\end{array}$$

Fig. 6. Operational semantics of distributed applications, for local actions.

move to its next task. For example, in Fig. 4 at time (4) the coordinator will inform P_b and P_c via $\langle \psi, P_b \rangle^{(\kappa, \{P_a, P_d\})} \cdot \langle \psi, P_c \rangle^{(\kappa, \{P_a, P_d\})}$ that they can finish the try-handle $(\kappa, \{P_a, P_d\})$. Note that the appearance of $\langle \psi, p \rangle^\phi$ implies that the coordinator has been informed that all participants in ϕ have completed their local actions. We define two kinds of *failure* notifications: $\langle \psi, \mathit{crash} F \rangle$ notifies ψ that F occurred, e.g., $\{q\}$ means q has failed; $\langle p, \mathit{crash} F \rangle$ is sent from ψ to notify p about the failure F for possible handling. We write $\langle \tilde{p}, \mathit{crash} F \rangle$, where $\tilde{p} = p_1, \dots, p_n$ short for $\langle p_1, \mathit{crash} F \rangle \cdot \dots \cdot \langle p_n, \mathit{crash} F \rangle$; similarly for $\langle \psi, \tilde{p} \rangle^\phi$. Following the tradition of other MPST works the global queue provides an abstraction for multiple FIFO queues, each queue being between two endpoints (cf. TCP) with no global ordering. Therefore $m_i \cdot m_j$ can be permuted to $m_j \cdot m_i$ in the global queue if the sender or the receiver differ. For example the following messages are permutable: $\langle p, q, l(v) \rangle \cdot \langle p, q', l(v) \rangle$ if $q \neq q'$ and $\langle p, q, l(v) \rangle \cdot \langle \psi, p \rangle^\phi$ and $\langle p, q, l(v) \rangle \cdot \langle q, \mathit{crash} F \rangle$. But $\langle \psi, p \rangle^\phi \cdot \langle p, \mathit{crash} F \rangle$ is not permutable, both have the same sender and receiver (ψ is the sender of $\langle p, \mathit{crash} F \rangle$). The formal definition of message permutation is in App. B Def. 13.

Basic dynamic semantics for applications. Fig. 6 shows the operational semantics of applications. We use evaluation contexts as defined in Fig. 5. Context E is either a hole $[\]$, a default context $\mathbf{t}(E)\mathbf{h}(\mathbf{H})^\phi.\eta$, or a recursion context $\mathbf{def} D \mathbf{in} E$. We write $E[\eta]$ to denote the action statement obtained by filling the hole in $E[\]$ with η .

Rule **(Link)** says that (local) processes who agree on shared name a , obeying to some protocol (global type), playing certain roles p_i represented by $a[p_i](y_i).P$, together will start a private session s ; this will result in replacing every variable y_i in P_i and, at the same time, creating a new global queue $s : \emptyset$, and appointing a coordinator $G : (\emptyset, \emptyset)$, which is novel in our work.

Rule **(Snd)** in Fig. 6 reduces a sending action $q! l(e)$ by emitting a message $\langle p, q, l(v) \rangle$ to the global queue $s : h$. Rule **(Rcv)** reduces a receiving action if the

message arriving at its end is sent from the expected sender with an expected label. Rule **(Rec)** is for recursion. When the recursive body, defined inside η , is called by $X\langle e \rangle$ where e is evaluated to v , it reduces to the statement $\eta\{v/x\}$ which will again implement the recursive body. Rule **(Str)** says that processes which are structurally congruent have the same reduction. Processes, applications, and systems are considered modulo structural congruence, denoted by \equiv , along with α -renaming. Since the definition of structural congruence is mostly standard, we leave the full definition in App. B.2. Rule **(Par)** and **(Str)** together state that a parallel composition has a reduction if its sub-application can reduce. Rule **(Sys)** states that a system has a reduction if its application has a reduction, and **(New)** says a reduction can proceed under a session. Rule **(Crash)** states that a process on channel $s[p]$ can fail at any point in time. **(Crash)** also adds a notification $\langle \psi, \text{crash } F \rangle$ which is sent to ψ (the coordinator). This is an abstraction for the failure detector described in Sec. 2 (5), the notification $\langle \psi, \text{crash } F \rangle$ is the first such notification issued by a participant based on its local failure detector. Adding the notification into the global queue instead of making the coordinator immediately aware of it models that failures are only detected eventually. Note that a failure is not annotated with a level because failures transcend all levels, and asynchrony makes it impossible to identify “where” exactly they occurred. As a failure is permanent it can affect multiple try-handles. The **(Crash)** rule does not apply to participants which are robust, i.e., that conceptually cannot fail (e.g., dfs in Fig. 2). Rule **(Crash)** removes channel $s[p]$ (the failed process) from application N , and removes messages and notifications delivered from, or heading to, the failed p by function $remove(h, p)$. Function $remove(h, p)$ returns a new queue after removing all regular messages and notifications that contain p , e.g., let $h = \langle p_2, p_1, l(v) \rangle \cdot \langle p_3, p_2, l'(v') \rangle \cdot \langle p_3, p_4, l'(v') \rangle \cdot \langle p_2, \psi \rangle^\phi \cdot \langle p_2, \text{crash } \{p_3\} \rangle \cdot \langle \psi, p_2 \rangle^\phi$ then $remove(h, p_2) = \langle p_3, p_4, l'(v') \rangle$. Messages are removed to model that in a real system send/receive does *not* constitute an atomic action.

Handling at processes. Failure handling, defined in Fig. 7, is based on the observations that (i) a process that fails stays down, and (ii) multiple processes can fail. As a consequence a failure can trigger multiple failure handlers either because these handlers are in different (subsequent) try-handles or because of additional failures. Therefore a process needs to retain the information of *who* failed. For simplicity we do not model state at processes, but instead processes read but do not remove failure notifications from the global queue. We define $Fset(h, p)$ to return the union of failures for which there are notifications heading to p , i.e., $\langle p, \text{crash } F \rangle$, issued by the coordinator in queue h up to the first done notification heading to p :

Definition 2 (Union of Existing Failures $Fset(h, p)$).

$$Fset(\emptyset, p) = \emptyset \quad Fset(h, p) = \begin{cases} F \cup Fset(h', p) & \text{if } h = \langle p, \text{crash } F \rangle \cdot h' \\ \emptyset & \text{if } h = \langle \psi, p \rangle^\phi \cdot h' \\ Fset(h', p) & \text{otherwise} \end{cases}$$

$$\begin{array}{c}
\frac{F' = \cup\{A \mid A \in \text{dom}(\mathbf{H}) \wedge F \subset A \subseteq F\text{set}(h, p)\} \quad F': \eta' \in \mathbf{H}}{\mathfrak{s}[p] : E[\mathfrak{t}(\eta)\mathfrak{h}(\mathbf{H})^{(\kappa, F)}. \eta''] \mid \mathfrak{s} : h \rightarrow \mathfrak{s}[p] : E[\mathfrak{t}(\eta')\mathfrak{h}(\mathbf{H})^{(\kappa, F')}. \eta'']] \mid \mathfrak{s} : h} \quad (\mathbf{TryHdl}) \\
\mathfrak{s}[p] : E[\mathfrak{t}(\mathbf{0})\mathfrak{h}(\mathbf{H})^\phi. \eta] \mid \mathfrak{s} : h \rightarrow \mathfrak{s}[p] : E[\mathfrak{t}(\mathbf{0})\mathfrak{h}(\mathbf{H})^\phi. \eta] \mid \mathfrak{s} : h \cdot \langle p, \psi \rangle^\phi \quad (\mathbf{SndDone}) \\
\frac{\langle \psi, p \rangle^\phi \in h}{\mathfrak{s}[p] : E[\mathfrak{t}(\mathbf{0})\mathfrak{h}(\mathbf{H})^\phi. \eta] \mid \mathfrak{s} : h \rightarrow \mathfrak{s}[p] : E[\eta] \mid \mathfrak{s} : h \setminus \{\langle \psi, p \rangle^\phi\}} \quad (\mathbf{RcvDone}) \\
\mathfrak{s}[p] : E[\eta] \mid \mathfrak{s} : \langle q, p, l(v) \rangle \cdot h \rightarrow \mathfrak{s}[p] : E[\eta] \mid \mathfrak{s} : h \quad l \notin \text{labels}(E[\eta]) \quad (\mathbf{Cln}) \\
\frac{\langle \psi, p \rangle^\phi \in h \quad \phi \notin E[\eta]}{\mathfrak{s}[p] : E[\eta] \mid \mathfrak{s} : h \rightarrow \mathfrak{s}[p] : E[\eta] \mid \mathfrak{s} : h \setminus \langle \psi, p \rangle^\phi} \quad (\mathbf{ClnDone})
\end{array}$$

Fig. 7. Operational semantics of distributed applications, for endpoint handling.

In short, if the global queue is \emptyset , then naturally there are no failure notifications. If the global queue contains a failure notification sent from the coordinator, say $\langle p, \text{crash } F \rangle$, we collect the failure. If the global queue contains done notification $\langle \psi, p \rangle^\phi$ sent from the coordinator then *all* participants in ϕ have finished their local actions, which implies that the try-handle ϕ can be completed. Our failure handling semantics, **(TryHdl)**, allows a try-handle $\phi = (\kappa, F)$ to handle different failures or sets of failures by allowing a try-handle to switch between different handlers. F thus denotes the current set of handled failures. For simplicity we refer to this as the *current(ly handled) failure set*. This is a slight abuse of terminology, done for brevity, as obviously failures are only detected with a certain lag. The handling strategy for a process is to handle the — currently — largest set of failed processes that this process has been informed of and is able to handle. This largest set is calculated by $\cup\{A \mid A \in \text{dom}(\mathbf{H}) \wedge F \subset A \subseteq F\text{set}(h, p)\}$, that selects all failure sets which are larger than the current one ($A \in \text{dom}(\mathbf{H}) \wedge F \subset A$) if they are also triggered by known failures ($A \subseteq F\text{set}(h, p)$). Condition $F': \eta' \in \mathbf{H}$ in **(TryHdl)** ensures that there exists a handler for F' . The following example shows how **(TryHdl)** is applied to switch handlers.

Example 2. Take h such that $F\text{set}(h, p) = \{p_1\}$ and $\mathbf{H} = \{p_1\} : \eta_1, \{p_2\} : \eta_2, \{p_1, p_2\} : \eta_{12}$ in process $P = \mathfrak{s}[p] : \mathfrak{t}(\eta_1)\mathfrak{h}(\mathbf{H})^{(\kappa, \{p_1\})}$, which indicates that P is handling failure $\{p_1\}$. Assume now one more failure occurs and results in a new queue h' such that $F\text{set}(h', p) = \{p_1, p_2\}$. By **(TryHdl)**, the process acting at $\mathfrak{s}[p]$ is handling the failure set $\{p_1, p_2\}$ such that $P = \mathfrak{s}[p] : \mathfrak{t}(\eta_{12})\mathfrak{h}(\mathbf{H})^{(\kappa, \{p_1, p_2\})}$ (also notice the η_{12} inside the try-block). A switch to only handling $\{p_2\}$ does not make sense, since, e.g., η_2 can contain p_1 . Fig. 2 shows a case where the handling strategy differs according to the number of failures.

In Sec. 3 we formally define well-formedness conditions, which guarantee that if there exist two handlers for two different handler signatures in a try-handle, then a handler exists for their union. The following example demonstrates why such a guarantee is needed.

Example 3. Assume a slightly different P compared to the previous examples (no handler for the union of failures): $P = \mathfrak{s}[p] : E[\mathfrak{t}(\eta)\mathfrak{h}(\mathbf{H})^{(\kappa, \emptyset)}]$ with $\mathbf{H} =$

$\{p_1\} : \eta_1, \{p_2\} : \eta_2$. Assume also that $Fset(h, p) = \{p_1, p_2\}$. Here **(TryHdl)** will not apply since there is no failure handling for $\{p_1, p_2\}$ in P . If we would allow a handler for either $\{p_1\}$ or $\{p_2\}$ to be triggered we would have no guarantee that other participants involved in this try-handle will all select the same failure set. Even with a deterministic selection, i.e., all participants in that try-handle selecting the same handling activity, there needs to be a handler with handler signature $= \{p_1, p_2\}$ since it is possible that p_1 is involved in η_2 . Therefore the type system will ensure that there is a handler for $\{p_1, p_2\}$ either at this level or at an outer level.

(I) explains that a process finishing its default action (P_b) cannot leave its current try-handle (κ, \emptyset) immediately because other participants may fail (P_a failed). Below Eq. 1 also shows this issue from the perspective of semantics:

$$\begin{aligned} \mathfrak{s}[p] : \mathfrak{t}(0)\mathfrak{h}(F : q!l(10).q?l'(x))^{(\kappa, \emptyset)}. \eta' \mid \mathfrak{s}[q] : \mathfrak{t}(p?l(x').p!l'(x'+10))\mathfrak{h}(H)^{(\kappa, F)}. \eta'' \\ \mid \mathfrak{s} : \llbracket q, \text{crash } F \rrbracket \cdot \llbracket p, \text{crash } F \rrbracket \cdot h \quad (1) \end{aligned}$$

In Eq. 1 the process acting on $\mathfrak{s}[p]$ ended its try-handle (i.e., the action is 0 in the try-block), and if $\mathfrak{s}[p]$ finishes its try-handle the participant acting on $\mathfrak{s}[q]$ which started handling F would be stuck.

To solve the issue, we use **(SndDone)** and **(RcvDone)** for completing a local try-handle with the help of a coordinator. The rule **(SndDone)** sends out a done notification $\langle p, \psi \rangle^\phi$ if the current action in ϕ is 0 and sets the action to $\underline{0}$, indicating that a done notification from the coordinator is needed for ending the try-handle.

Assume process on channel $\mathfrak{s}[p]$ finished its local actions in the try-block (i.e., as in Eq. 1 above), then by **(SndDone)**, we have

$$\begin{aligned} (1) \rightarrow \mathfrak{s} : \llbracket q, \text{crash } F \rrbracket \cdot \llbracket p, \text{crash } F \rrbracket \cdot \langle p, \psi \rangle^{(\kappa, \emptyset)} \cdot h \mid \\ \mathfrak{s}[p] : \mathfrak{t}(\underline{0})\mathfrak{h}(F : q!l(10).q?l'(x))^{(\kappa, \emptyset)}. \eta' \mid \mathfrak{s}[q] : \mathfrak{t}(p?l(x').p!l'(x'+10))\mathfrak{h}(H)^{(\kappa, F)}. \eta'' \end{aligned}$$

where notification $\langle p, \psi \rangle^{(\kappa, \emptyset)}$ is added to inform the coordinator. Now the process on channel $\mathfrak{s}[p]$ can still handle failures defined in its handling environment. This is similar to the case described in (II).

Rule **(RcvDone)** is the counterpart of **(SndDone)**. Once a process receives a done notification for ϕ from the coordinator it can finish the try-handle ϕ and reduces to the continuation η . Consider Eq. 2 below, which is similar to Eq. 1 but we take a case where the try-handle can be reduced with **(RcvDone)**. In Eq. 2 **(SndDone)** is applied:

$$\begin{aligned} \mathfrak{s}[p] : \mathfrak{t}(\underline{0})\mathfrak{h}(F : q!l(10).q?l'(x))^{(\kappa, \emptyset)}. \eta' \mid \\ \mathfrak{s}[q] : \mathfrak{t}(\underline{0})\mathfrak{h}(F : p?l(x').p!l'(x'+10))^{(\kappa, \emptyset)}. \eta'' \mid \mathfrak{s} : h \quad (2) \end{aligned}$$

With $h = \langle \psi, q \rangle^{(\kappa, \emptyset)} \cdot \langle \psi, p \rangle^{(\kappa, \emptyset)} \cdot \llbracket q, \text{crash } F \rrbracket \cdot \llbracket p, \text{crash } F \rrbracket$ both processes can apply **(RcvDone)** and safely terminate the try-handle (κ, \emptyset) . Note that $Fset(h, p) =$

$$\begin{array}{c}
\frac{\tilde{p} = \text{roles}(G) \setminus F' \quad F' = F \cup \{p\} \quad m = \langle \tilde{p}, \text{crash } \{p\} \rangle}{G : (F, d) \blacklozenge N \mid \mathfrak{s} : \langle \psi, \text{crash } \{p\} \rangle \cdot h \rightarrow G : (F', d) \blacklozenge N \mid \mathfrak{s} : h \cdot m} \text{ (F)} \\
\frac{d' = d \cdot \langle p, \psi \rangle^\phi}{G : (F, d) \blacklozenge \mathfrak{s} : \langle p, \psi \rangle^\phi \cdot h \rightarrow G : (F, d') \blacklozenge \mathfrak{s} : h} \text{ (CollectDone)} \\
\frac{\text{roles}(d, \phi) \supseteq \text{roles}(G, \phi) \setminus F \quad \forall F' \in \text{hdl}(G, \phi). (F' \not\subseteq F)}{G : (F, d) \blacklozenge \mathfrak{s} : h \rightarrow G : (F, \text{remove}(d, \phi)) \blacklozenge \mathfrak{s} : h \cdot \langle \psi, \text{roles}(G, \phi) \setminus F \rangle^\phi} \text{ (IssueDone)}
\end{array}$$

Fig. 8. Operational semantics for the coordinator.

$Fset(h, q) = \emptyset$ (by Def. 2), i.e., rule **(TryHdl)** can not be applied since a done notification suppresses the failure notification. Thus Eq. 2 will reduce to:

$$(2) \rightarrow^* \mathfrak{s}[p] : \eta' \mid \mathfrak{s}[q] : \eta'' \mid \mathfrak{s} : \langle q, \text{crash } F \rangle \cdot \langle p, \text{crash } F \rangle$$

It is possible that η' or η'' have handlers for F . Note that once a queue contains $\langle \psi, p \rangle^{(\kappa, \emptyset)}$, all non-failed process in the try-handle (κ, \emptyset) have sent done notifications to ψ (i.e. applied rule **(SndDone)**). The coordinator which will be introduced shortly ensures this.

Rule **(Cln)** removes a normal message from the queue if the label in the message does not exist in the target process, which can happen when a failure handler was triggered. The function $\text{labels}(\eta)$ returns all labels of receiving actions in η which are able to receive messages now or possible later. (The function is formally defined in App. B. Def. 14). This removal based on the syntactic process is safe because in a global type separate branch types *not* defined in the same default block or handler body must have disjoint sets of labels (c.f., Sec. 3). Let $\phi \in P$ if try-handle ϕ appears inside P . Rule **(ClnDone)** removes a done notification of ϕ from the queue if no try-handle ϕ exists, which can happen in case of nesting when a handler of an outer try-handle is triggered.

Handling at coordinator. Fig. 8 defines the semantics of the coordinator. We firstly give the auxiliary definition of $\text{roles}(G)$ which gives the set of *all* roles appearing in G .

In rule **(F)**, F represents the failures that the coordinator is aware of. This rule states that the coordinator collects and removes a failure notification $\langle \psi, \text{crash } p \rangle$ heading to it, retains this notification by $G : (F', d)$, $F' = F \cup \{p\}$, and issues failure notifications to all non-failed participants.

Rules **(CollectDone, IssueDone)**, in short inform all participants in $\phi = (\kappa, F)$ to finish their try-handle ϕ if the coordinator has received sufficient done notifications of ϕ and did not send out failure notifications that interrupt the task (κ, F) (e.g. see (III)). Rule **(CollectDone)** collects done notifications, i.e., $\langle p, \psi \rangle^\phi$, from the queue and retains these notification; they are used in **(IssueDone)**. For introducing **(IssueDone)**, we first introduce $\text{hdl}(G, (\kappa, F))$ to return a set of handler signatures which can be triggered with respect to the current handler:

Definition 3. $\text{hdl}(G, (\kappa, F)) = \text{dom}(H) \setminus \mathcal{P}(F)$ if $\mathfrak{t}(G_0)\mathfrak{h}(H)^\kappa \in G$ where $\mathcal{P}(F)$ represents a powerset of F .

Also, we abuse the function $roles$ to collect the non-coordinator roles of ϕ in d , written $roles(d, \phi)$; similarly, we write $roles(G, \phi)$ where $\phi = (\kappa, F)$ to collect the roles appearing in the handler body F in the try-handle of κ in G . Remember that d only contains done notifications sent by participants.

Rule **(IssueDone)** is applied for some ϕ when conditions $\forall F' \in hdl(G, \phi).(F' \not\subseteq F)$ and $roles(d, \phi) \supseteq roles(G, \phi) \setminus F$ are both satisfied, where F contains all failures the coordinator is aware of. Intuitively, these two conditions ensure that (1) the coordinator only issues done notifications to the participants in the try-handle ϕ if it did not send failure notifications which will trigger a handler of the try-handle ϕ ; (2) the coordinator has received all done notifications from all non-failed participants of ϕ . We further explain both conditions in the following examples, starting from condition $\forall F' \in hdl(G, \phi).(F' \not\subseteq F)$, which ensures no handler in ϕ can be triggered based on the failure notifications F sent out by the coordinator.

Example 4. Assume a process playing role p_i is $P_i = s[p_i] : t(\eta_i)h(H_i)^{\phi_i}$. where $i \in \{1, 2, 3\}$ and $H_i = \{p_2\} : \eta_{i2}, \{p_3\} : \eta_{i3}, \{p_2, p_3\} : \eta_{i23}$ and the coordinator is $G : (\{p_2, p_3\}, d)$ where $t(\dots)h(H)^\kappa \in G$ and $dom(H) = dom(H_i)$ for any $i \in \{1, 2, 3\}$ and $d = \langle p_1, \psi \rangle^{\kappa, \{p_2\}} \cdot \langle p_1, \psi \rangle^{\kappa, \{p_2, p_3\}} \cdot d'$. For any ϕ in d , the coordinator checks if it has issued any failure notification that can possibly trigger a new handler of ϕ :

1. For $\phi = (\kappa, \{p_2\})$ the coordinator issued failure notifications that can interrupt a handler since

$$hdl(G, (\kappa, \{p_2\})) = dom(H) \setminus \mathcal{P}(\{p_2\}) = \{\{p_3\}, \{p_2, p_3\}\}$$

and $\{p_2, p_3\} \subseteq \{p_2, p_3\}$. That means the failure notifications issued by the coordinator, i.e., $\{p_2, p_3\}$, can trigger the handler with signature $\{p_2, p_3\}$. Thus the coordinator will not issue done notifications for $\phi = (\kappa, \{p_2\})$. A similar case is visualized in Fig. 4 at time (2).

2. For $\phi = (\kappa, \{p_2, p_3\})$ the coordinator did not issue failure notifications that can interrupt a handler since

$$hdl(G, (\kappa, \{p_2, p_3\})) = dom(H) \setminus \mathcal{P}(\{p_2, p_3\}) = \emptyset$$

so that $\forall F' \in hdl(G, (\kappa, \{p_2, p_3\})).(F' \not\subseteq \{p_2, p_3\})$ is true. The coordinator will issue done notifications for $\phi = (\kappa, \{p_2, p_3\})$.

Another condition $roles(d, \phi) \supseteq roles(G, \phi) \setminus F$ states that only when the coordinator sees sufficient done notifications (in d) for ϕ , it issues done notifications to *all* non-failed participants in ϕ , i.e., $\langle \psi, roles(G, \phi) \setminus F \rangle^\phi$. Recall that $roles(d, \phi)$ returns all roles which have sent a done notification for the handling of ϕ and $roles(G, \phi)$ returns all roles involving in the handling of ϕ . Intuitively one might expect the condition to be $roles(d, \phi) = roles(G, \phi)$; the following example shows why this would be wrong.

$$\begin{aligned}
T &::= p! \{l_i(S_i).T_i\}_{i \in I} \mid p? \{l_i(S_i).T_i\}_{i \in I} \mid t \mid \mu t.T \mid \text{end} \mid \underline{\text{end}} \mid \mathbf{t}(T)\mathbf{h}(\mathcal{H})^\phi.T \\
\mathcal{H} &::= F:T \mid \mathcal{H}, \mathcal{H}
\end{aligned}$$

Fig. 9. The grammar of local types.

Example 5. Consider a process P acting on channel $s[p]$ and $\{q\} \notin \text{dom}(\mathbf{H})$:

$$P = s[p] : \mathbf{t}(\dots\mathbf{t}(\dots)\mathbf{h}(\{q\}:\eta, \mathbf{H}')^{\phi'}.\eta')\mathbf{h}(\mathbf{H})^\phi$$

Assume P has already reduced to:

$$P = s[p] : \mathbf{t}(\mathbf{0})\mathbf{h}(\mathbf{H})^\phi$$

We show why $\text{roles}(d, \phi) \supseteq \text{roles}(G, \phi) \setminus F$ is necessary. We start with the simple cases and then move to the more involving ones.

- (a) Assume q did not fail, the coordinator is $G : (\emptyset, d)$, and all roles in ϕ issued a done notification. Then $\text{roles}(d, \phi) = \text{roles}(G, \phi)$ and $F = \emptyset$.
- (b) Assume q failed in the try-handle ϕ' , the coordinator is $G : (\{q\}, d)$, and all roles except q in ϕ issued a done notification. $\text{roles}(d, \phi) \neq \text{roles}(G, \phi)$ however $\text{roles}(d, \phi) = \text{roles}(G, \phi) \setminus \{q\}$. Cases like this are the reason why **(IssueDone)** only requires done notifications from non-failed roles.
- (c) Assume q failed after it has issued a done notification for ϕ (i.e., q finished try-handle ϕ') and the coordinator collected it (by **(CollectDone)**), so we have $G : (\{q\}, d)$ and $q \in \text{roles}(b, \phi)$. Then $\text{roles}(d, \phi) \supset \text{roles}(G, \phi) \setminus \{q\}$. i.e. **(IssueDone)** needs to consider done notifications from failed roles.

Thus rule **(IssueDone)** has the condition $\text{roles}(d, \phi) \supseteq \text{roles}(G, \phi) \setminus F$ because of cases like (b) and (c).

The interplay between issuing of done notification **(IssueDone)** and issuing of failure notifications **(F)** is non-trivial. The following proposition clarifies that the participants in the same try-handle ϕ will never get confused with handling failures or completing the try-handle ϕ .

Proposition 1. *Given $s : h$ with $h = h' \cdot \langle \psi, p \rangle^\phi \cdot h''$ and $F\text{set}(h, p) \neq \emptyset$, the rule **(TryHdl)** is not applicable for the try-handle ϕ at the process playing role p .*

Proof see App. D.1

5 Local Types

Fig. 9 defines local types for typing behaviors of endpoint processes with failure handling. Type $p!$ is the primitive for a sending type, and $p?$ is the primitive for a receiving type, derived from global type $p \rightarrow q\{l_i(S_i).G_i\}_{i \in I}$ by projection. Others correspond straightforwardly to process terms. Note that type **end** only appears in *runtime* type checking. Below we define $G|p$ to project a global type G on p , thus generating p 's local type.

Definition 4 (Projection). Consider a well-formed top-level global type $[\tilde{q}]G$. Then $G \downarrow p$ is defined as follows:

- (1) $G \downarrow p$ where $G = \mathbf{t}(G_0)\mathbf{h}(F_1 : G_1, \dots, F_n : G_n)^\kappa . G' =$

$$\begin{cases} \mathbf{t}(G_0 \downarrow p)\mathbf{h}(F_1 : G_1 \downarrow p, \dots, F_n : G_n \downarrow p)^{(\kappa, \emptyset)} . G' \downarrow p & \text{if } p \in \text{roles}(G) \\ G' \downarrow p & \text{otherwise} \end{cases}$$
- (2) $p_1 \rightarrow p_2 \{l_i(S_i).G_i\}_{i \in I} \downarrow p =$

$$\begin{cases} p_2! \{l_i(S_i).G_i \downarrow p\}_{i \in I} & \text{if } p = p_1 \\ p_1? \{l_i(S_i).G_i \downarrow p\}_{i \in I} & \text{if } p = p_2 \\ G_1 \downarrow p & \text{if } \forall i, j \in I. G_i \downarrow p = G_j \downarrow p \end{cases}$$
- (3) $(\mu t.G) \downarrow p = \mu t.(G \downarrow p)$ if $\nexists \mathbf{t}(G')\mathbf{h}(H) \in G$ and $G \downarrow p \neq t'$ for any t'
- (4) $t \downarrow p = t$ (5) $\text{end} \downarrow p = \text{end}$

Otherwise it is undefined.

The main rule is (1): if p appears somewhere in the target try-handle global type then the endpoint type has a try-handle annotated with κ and the default logic (i.e., $F = \emptyset$). Note that even if $G_0 \downarrow p = \text{end}$ the endpoint still gets such a try-handle because it needs to be ready for (possible) failure handling; if p does not appear anywhere in the target try-handle global type, then the projection skips to the continuation.

Rule (2) produces local types for interaction endpoints. If the endpoint is a sender (i.e., $p = p_1$), then its local type abstracts that it will send something from one of the possible internal choices defined in $\{l_i(S_i)\}_{i \in I}$ to p_2 , then continue as $G_k \downarrow p$, gained from the projection, if $k \in I$ is chosen. If the endpoint is a receiver (i.e., $p = p_2$), then its local type abstracts that it will receive something from one of the possible external choices defined in $\{l_i(S_i)\}_{i \in I}$ sent by p_1 ; the rest is similarly as for the sender. However, if p is not in this interaction, then its local type starts from the next interaction which p is in; moreover, because p does not know what choice that p_1 has made, every path $G_i \downarrow p$ lead by branch l_i shall be the same for p to ensure that interactions are consistent. For example, in $G = p_1 \rightarrow p_2 \{l_1(S_1).p_3 \rightarrow p_1 l_3(S), l_2(S_2).p_3 \rightarrow p_1 l_4(S)\}$, interaction $p_3 \rightarrow p_1$ continues after $p_1 \rightarrow p_2$ takes place. If $l_3 \neq l_4$, then G is not projectable for p_3 because p_3 does not know which branch that p_1 has chosen; if p_1 chose branch l_1 , but p_3 (blindly) sends out label l_4 to p_1 , for p_1 it is a mistake (but it is not a mistake for p_3) because p_1 is expecting to receive label l_3 . To prevent such inconsistencies, we adopt the projection algorithm proposed in [24]. Other session type works [17,40] provide ways to weaken the classical restriction on projection of branching which we use.

Rule (3) forbids a try-handle to appear in a recursive body, e.g., $\mu t.\mathbf{t}(G)\mathbf{h}(F : t)^\kappa . G$ is not allowed, but $\mathbf{t}(\mu t.G)\mathbf{h}(H)^\kappa$ and $\mathbf{t}(G)\mathbf{h}(F : \mu t.G', H)^\kappa$ are allowed. This is because κ is used to avoid confusion of messages from different try-handles. If a recursive body contains a try-handle, we have to dynamically generate different levels to maintain interaction consistency, so static type checking does not suffice.

We are investigating alternative runtime checking mechanisms, but this is beyond the scope of this paper. Other rules are straightforward.

Example 6. Recall the global type G from Fig. 2 in Sec. 1. Applying projection rules defined in Def. 4 to G on every role in G we obtain the following:

$$\begin{aligned}
T_{dfs} &= G \upharpoonright dfs = \mathbf{t}(\mu t. w_1 ! l_{d_1}(S). w_2 ! l_{d_2}(S). w_1 ? l_{r_1}(S'). w_2 ? l_{r_2}(S'). t) \mathbf{h}(\mathcal{H}_{dfs})^{(1, \emptyset)} \\
\mathcal{H}_{dfs} &= \{w_1\} : \mu t'. w_2 ! l'_{d_1}(S). w_2 ? l'_{r_1}(S'). t', \\
&\quad \{w_2\} : \mu t''. w_1 ! l'_{d_2}(S). w_1 ? l'_{r_2}(S'). t'', \{w_1, w_2\} : \mathbf{end} \\
T_{w_1} &= G \upharpoonright w_1 = \mathbf{t}(\mu t. dfs ? l_{d_1}(S). dfs ! l_{r_1}(S'). t) \mathbf{h}(\mathcal{H}_{w_1})^{(1, \emptyset)} \\
\mathcal{H}_{w_1} &= \{w_1\} : \mathbf{end}, \{w_2\} : \mu t'. dfs ? l'_{d_2}(S). dfs ! l'_{r_2}(S'). t', \{w_1, w_2\} : \mathbf{end} \\
T_{w_2} &= G \upharpoonright w_2 = \mathbf{t}(\mu t. dfs ? l_{d_2}(S). dfs ! l_{r_2}(S'). t) \mathbf{h}(\mathcal{H}_{w_2})^{(1, \emptyset)} \\
\mathcal{H}_{w_2} &= \{w_1\} : \mu t''. dfs ? l'_{d_1}(S). dfs ! l'_{r_1}(S'). t'', \{w_2\} : \mathbf{end}, \{w_1, w_2\} : \mathbf{end}
\end{aligned}$$

6 Type System

Next we introduce our type system for typing processes. Fig. 10 and Fig. 11 present typing rules for endpoints processes, and typing judgments for applications and systems respectively.

We define shared environments Γ to keep information on variables and the coordinator, and session environments Δ to keep information on endpoint types:

$$\begin{aligned}
\Gamma &::= \emptyset \mid \Gamma, X : S \mid \Gamma, x : S \mid \Gamma, a : G \mid \Gamma, \Psi & \Delta &::= \emptyset \mid \Delta, c : T \mid \Delta, s : \mathbf{h} \\
\mathbf{m} &::= \langle p, q, l(S) \rangle \mid \langle p, \text{crash } F \rangle \mid \langle p, q \rangle^\phi & \mathbf{h} &::= \emptyset \mid \mathbf{h} \cdot \mathbf{m}
\end{aligned}$$

Γ maps process variables X and content variables x to their types, shared names a to global types G , and a coordinator $\Psi = G : (F, d)$ to failures and done notifications it has observed. Δ maps session channels c to local types and session queues to queue types. We write $\Gamma, \Gamma' = \Gamma \cup \Gamma'$ when $\text{dom}(\Gamma) \cap \text{dom}(\Gamma') = \emptyset$; same for Δ, Δ' . Queue types \mathbf{h} are composed of message types \mathbf{m} . Their permutation is defined analogously to the permutation for messages. The typing judgment for local processes $\Gamma \vdash P \triangleright \Delta$ states that process P is well-typed by Δ under Γ .

Since we do not define sequential composition for processes, our type system implicitly forbids session interleaving by [T-**ini**]. This is different from other session type works [15,24], where session interleaving is prohibited for the progress property; here the restriction is inherent to the type system.

Fig. 10 lists our typing rules for endpoint processes. Rule [T-**ini**] says that if a process's set of actions is well-typed by $G \upharpoonright p$ on some c , this process can play role p in a , which claims to have interactions obeying behaviors defined in G . $\langle G \rangle$ means that G is closed, i.e., devoid of type variables. This rule forbids $a[p].b[q].P$ because a process can only use one session channel. Rule [T-**snd**] states that an action for sending is well-typed to a sending type if the label and the type of the content are expected; [T-**rcv**] states that an action for branching (i.e., for receiving) is well-typed to a branching type if all labels and the types of contents are as expected. Their follow-up actions shall also be well-typed. Rule [T-0] types an idle process. Predicate end-only Δ is defined as stating whether all endpoints in Δ have type **end**:

$$\begin{array}{c}
\frac{\Gamma \vdash a : \langle G \rangle}{\Gamma \vdash P \triangleright \{c : G[p]\}} \quad \frac{k \in I \quad \Gamma \vdash e : S_k}{\Gamma \vdash c : \eta_k \triangleright \{c : T_k\}}}{\Gamma \vdash a[p].P \triangleright \emptyset \quad \Gamma \vdash c : p! l_k(e).\eta_k \triangleright \{c : p! \{l_i(S_i).T_i\}_{i \in I}\}} \quad [\text{T-ini/T-snd}] \\
\frac{\forall i \in I. \Gamma, x_i : S_i \vdash c : \eta_i \triangleright \{c : T_i\}}{\Gamma \vdash c : p? \{l_i(x_i).\eta_i\}_{i \in I} \triangleright \{c : p? \{l_i(S_i).T_i\}_{i \in I}\}} \quad [\text{T-rcv}] \\
\frac{\Delta \text{ end-only}}{\Gamma \vdash c : 0 \triangleright \Delta} \quad \frac{\Gamma \vdash c : \eta \triangleright \{c : \text{end}\}}{\Gamma \vdash c : \underline{0}.\eta \triangleright \{c : \text{end}.\text{end}\}} \quad [\text{T-0/T-yd}] \\
\Gamma \vdash e : \text{bool} \\
\frac{\forall i \in \{1, 2\}. \Gamma \vdash c : \eta_i \triangleright \Delta}{\Gamma \vdash c : \text{if } e \eta_1 \text{ else } \eta_2 \triangleright \Delta} \quad \frac{\Gamma \vdash e : S}{\Gamma, X : S \quad T \vdash c : X(e) \triangleright \{c : T\}} \quad [\text{T-if/T-var}] \\
\frac{\Gamma, X : S \quad t, x : S \vdash c : \eta_1 \triangleright \{c : T'\} \quad \Gamma, X : S \quad \mu t.T' \vdash c : \eta_2 \triangleright \{c : T\}}{\Gamma \vdash c : \text{def } X(x) = \eta_1 \text{ in } \eta_2 \triangleright \{c : T\}} \quad [\text{T-def}] \\
\frac{\Gamma \vdash c : \eta \triangleright \{c : T\} \quad \Gamma \vdash c : \eta' \triangleright \{c : T'\} \quad \text{dom}(\mathbf{H}) = \text{dom}(\mathcal{H}) \quad \forall F \in \text{dom}(\mathbf{H}). \Gamma \vdash c : \mathbf{H}(F) \triangleright \{c : \mathcal{H}(F)\}}{\Gamma \vdash c : \mathbf{t}(\eta)\mathbf{h}(\mathbf{H})^\phi.\eta' \triangleright \{c : \mathbf{t}(T)\mathbf{h}(\mathcal{H})^\phi.T'\}} \quad [\text{T-th}]
\end{array}$$

Fig. 10. Typing rules for processes

Definition 5 (End-only Δ). We say Δ is end-only if and only if $\forall s[p] \in \text{dom}(\Delta), \Delta(s[p]) = \text{end}$.

Rule [T-yd] types yielding actions, which only appear at runtime. Rule [T-if] is standard in the sense that the process is well-typed by Δ if e has boolean type and its sub-processes (i.e., η_1 and η_2) are well-typed by Δ . Rules [T-var, T-def] are based on a recent summary of MPSTs [14]. Note that [T-def] forbids the type $\mu t.t$. Rule [T-th] states that a try-handle is well-typed if it is annotated with the expected level ϕ , its default statement is well-typed, \mathcal{H} and \mathbf{H} have the same handler signatures, and all handling actions are well-typed.

Fig. 11 shows typing rules for applications and systems. Rule [T- \emptyset] types an empty queue. Rules [T-m, T-D, T-F] simply type messages based on their shapes. Rule [T-pa] says two applications composed in parallel are well-typed if they do not share any session channel. Rule [T-s] says a part of a system \mathcal{S} can start a private session, say \mathbf{s} , if \mathcal{S} is well-typed according to a $\Gamma \vdash \Delta_{\mathbf{s}}$ that is *coherent* (defined shortly). The system $(\nu \mathbf{s})\mathcal{S}$ with a part becoming private in \mathbf{s} is well-typed to $\Delta \setminus \Delta_{\mathbf{s}}$, that is, Δ after removing $\Delta_{\mathbf{s}}$.

Definition 6 (A Session Environment Having \mathbf{s} Only: $\Delta_{\mathbf{s}}$).

$$\Delta_{\mathbf{s}} = \{s[p] : T \mid s[p] \in \text{dom}(\Delta)\} \cup \{s : \mathbf{h} \mid s \in \text{dom}(\Delta)\}$$

Rule [T-sys] says that a system $\Psi \blacklozenge N$ is well-typed if application N is well-typed and there exists a coordinator Ψ for handling this application. We say $\Gamma \vdash \Delta$ is coherent under Γ if the local types of all endpoints are dual to each other after their local types are updated because of messages or notifications in $\mathbf{s} : \mathbf{h}$.

$$\begin{array}{c}
\Gamma \vdash \mathbf{s} : \emptyset \triangleright \{\mathbf{s} : \emptyset\} \quad \frac{\Gamma \vdash \mathbf{s} : h \triangleright \{\mathbf{s} : \mathbf{h}\} \quad \Gamma \vdash e : S}{\Gamma \vdash \mathbf{s} : h \cdot \langle p, q, l(e) \rangle \triangleright \{\mathbf{s} : \mathbf{h} \cdot \langle p, q, l(S) \rangle\}} \quad [\text{T-}\emptyset/\text{T-m}] \\
\\
\frac{(p_1, p_2) \in \{(p, \psi), (\psi, p)\} \quad \Gamma \vdash \mathbf{s} : h \triangleright \{\mathbf{s} : \mathbf{h}\}}{\Gamma \vdash \mathbf{s} : h \cdot \langle p_1, p_2 \rangle^\phi \triangleright \{\mathbf{s} : \mathbf{h} \cdot \langle p_1, p_2 \rangle^\phi\}} \quad [\text{T-D}] \\
\\
\frac{p \in \{q, \psi\} \quad \mathbf{m} = \langle p, \text{crash } F \rangle \quad \Gamma \vdash N_1 \triangleright \Delta_1 \quad \Gamma \vdash N_2 \triangleright \Delta_2 \quad \Gamma \vdash \mathbf{s} : h \triangleright \{\mathbf{s} : \mathbf{h}\} \quad \text{dom}(\Delta_1) \cap \text{dom}(\Delta_2) = \emptyset}{\Gamma \vdash \mathbf{s} : h \cdot \langle p, \text{crash } F \rangle \triangleright \{\mathbf{s} : \mathbf{h} \cdot \mathbf{m}\} \quad \Gamma \vdash N_1 \mid N_2 \triangleright \Delta_1, \Delta_2} \quad [\text{T-F/T-pa}] \\
\\
\frac{\Gamma \vdash S \triangleright \Delta \quad \Gamma \vdash \Delta_s \text{ coherent}}{\Gamma \vdash (\nu \mathbf{s})S \triangleright \Delta \setminus \Delta_s} \quad \frac{\Gamma' = \Gamma, \Psi \quad \Gamma \vdash N \triangleright \Delta}{\Gamma' \vdash \Psi \blacklozenge N \triangleright \Delta} \quad [\text{T-s/T-sys}]
\end{array}$$

Fig. 11. Typing rules for applications and systems.

Coherence. We say that a session environment is *coherent* if, at any time, given a session with its latest messages and notifications, every endpoint participating in it is able to find someone to interact with (i.e., its dual party exists) right now or afterwards.

Example 7. Continuing with Example 6 – the session environment $\Gamma \vdash \Delta$ is coherent even if w_2 will not receive any message from dfs at this point. The only possible action to take in Δ is that dfs sends out a message to w_1 . When this action fires, Δ is reduced to Δ' under a coordinator. (The reduction relation $\Gamma \vdash \Delta \rightarrow_T \Gamma' \vdash \Delta'$, where $\Gamma = \Gamma_0, \Psi$ and $\Gamma' = \Gamma_0, \Psi'$, is defined based on the rules of operational semantics of applications in Sec. 4, Fig. 6 and Fig. 7. Due to space limitations we put the complete rules into App. D.2). In Δ' , which abstracts the environment when dfs sends a message to w_1 , w_2 will be able to receive this message.

$$\begin{aligned}
\Delta &= \mathbf{s}[dfs] : T_{dfs}, \mathbf{s}[w_1] : T_{w_1}, \mathbf{s}[w_2] : T_{w_2}, \mathbf{s} : \emptyset \\
\Delta' &= \mathbf{s}[dfs] : \mathbf{t}(w_2!l_{d_2}(S).w_1?l_{r_1}(S').w_2?l_{r_2}(S').T)\mathbf{h}(\mathcal{H})^{(1,\emptyset)}, \\
&\quad \mathbf{s}[w_1] : T_{w_1}, \mathbf{s}[w_2] : T_{w_2}, \mathbf{s} : \langle dfs, w_1, l_{d_1}(S) \rangle \\
&\quad \text{where } T = \mu t.w_1!l_{d_1}(S).w_2!l_{d_2}(S).w_1?l_{r_1}(S').w_2?l_{r_2}(S').t
\end{aligned}$$

We write $\mathbf{s}[p] : T \bowtie \mathbf{s}[q] : T'$ to state that actions of the two types are *dual*:

Definition 7 (Duality). We define $\mathbf{s}[p] : T \bowtie \mathbf{s}[q] : T'$ as follows:

$$\begin{array}{c}
\mathbf{s}[p] : \text{end} \bowtie \mathbf{s}[q] : \text{end} \quad \mathbf{s}[p] : \underline{\text{end}} \bowtie \mathbf{s}[q] : \underline{\text{end}} \quad \mathbf{s}[p] : \text{end} \bowtie \mathbf{s}[q] : \underline{\text{end}} \\
\mathbf{s}[p] : \underline{\text{end}} \bowtie \mathbf{s}[q] : \text{end} \quad \mathbf{s}[p] : t \bowtie \mathbf{s}[q] : t \quad \frac{\mathbf{s}[p] : T \bowtie \mathbf{s}[q] : T'}{\mathbf{s}[p] : \mu t.T \bowtie \mathbf{s}[q] : \mu t.T'} \\
\\
\frac{\forall i \in I. \mathbf{s}[p] : T_i \bowtie \mathbf{s}[q] : T'_i}{\mathbf{s}[p] : q! \{l_i(S_i).T_i\}_{i \in I} \bowtie \mathbf{s}[q] : p? \{l_i(S_i).T'_i\}_{i \in I}} \\
\frac{\mathbf{s}[p] : T_1 \bowtie \mathbf{s}[q] : T_2 \quad \mathbf{s}[p] : T'_1 \bowtie \mathbf{s}[q] : T'_2 \quad \text{dom}(\mathcal{H}_1) = \text{dom}(\mathcal{H}_2) \quad \forall F \in \text{dom}(\mathcal{H}_1). \mathbf{s}[p] : \mathcal{H}_1(F) \bowtie \mathbf{s}[q] : \mathcal{H}_2(F)}{\mathbf{s}[p] : \mathbf{t}(T_1)\mathbf{h}(\mathcal{H}_1)^\phi.T'_1 \bowtie \mathbf{s}[q] : \mathbf{t}(T_2)\mathbf{h}(\mathcal{H}_2)^\phi.T'_2}
\end{array}$$

Due to space limitations, we provide the following formal definitions in App. B.3. Operation $T \downarrow p$ is to filter T to get the partial type which only contains actions of p . For example, $p_1!l'(S').p_2!l(S) \downarrow p_2 = p_2!l(S)$ and $p_1!\{T_1, T_2\} \downarrow p_2 = p_2?l(S)$ where $T_1 = l_1(S_1).p_2?l(S)$ and $T_2 = l_2(S_2).p_2?l(S)$. Next we define $(\mathbf{h})_{p \rightarrow q}$ to filter \mathbf{h} to generate (1) the normal message types sent from p heading to q , and (2) the notifications heading to q . For example $(\langle p, q, l(S) \rangle \cdot \langle q, \text{crash } F \rangle \cdot \langle \psi, q \rangle^\phi \cdot \langle p, \text{crash } F \rangle)_{p \rightarrow q} = p?l(S) \cdot \langle F \rangle \cdot \langle \psi \rangle^\phi$. The message types are abbreviated to contain only necessary information.

We define $T\text{-ht}$ to mean the effect of \mathbf{ht} on T . Its concept is similar to the *session remainder* defined in [36], which returns new local types of participants after participants consume messages from the global queue. Since failure notifications will not be consumed in our system, and we only have to observe the change of a participant's type after receiving or being triggered by some message types in \mathbf{ht} , we say that $T\text{-ht}$ represents the effect of \mathbf{ht} on T . The behaviors follows our operational semantics of applications and systems defined in Fig. 6, Fig. 7, and Fig. 8. For example $\mathbf{t}(q?\{l_i(S_i).T_i\}_{i \in I})\mathbf{h}(\mathcal{H})^\phi.T' - q?l_k(S_k).\mathbf{ht} = \mathbf{t}(T_k)\mathbf{h}(\mathcal{H})^\phi.T'\text{-ht}$ where $k \in I$.

Now we define what it means for Δ to be coherent under Γ :

Definition 8 (Coherence). $\Gamma \vdash \Delta$ coherent if the following conditions hold:

1. If $\mathbf{s} : \mathbf{h} \in \Delta$, then $\exists G : (F, d) \in \Gamma$ and $\{p \mid \mathbf{s}[p] \in \text{dom}(\Delta)\} \subseteq \text{roles}(G)$ and G is well-formed and $\forall p \in \text{roles}(G), G \downarrow p$ is defined.
2. $\forall \mathbf{s}[p] : T, \mathbf{s}[q] : T' \in \Delta$ we have $\mathbf{s}[p] : T \downarrow q - (\mathbf{h})_{q \rightarrow p} \bowtie \mathbf{s}[q] : T' \downarrow p - (\mathbf{h})_{p \rightarrow q}$.

In condition 1, we require a coordinator for every session so that when a failure occurs, the coordinator can announce failure notifications to ask participants to handle the failure. Condition 2 requires that, for any two endpoints, say $\mathbf{s}[p]$ and $\mathbf{s}[q]$, in Δ , equation $\mathbf{s}[p] : T \downarrow q - (\mathbf{h})_{q \rightarrow p} \bowtie \mathbf{s}[q] : T' \downarrow p - (\mathbf{h})_{p \rightarrow q}$, must hold. This condition asserts that interactions of non-failed endpoints are dual to each other after the effect of \mathbf{h} ; while failed endpoints are removed from Δ , thus the condition is satisfied immediately.

7 Properties

We show that our type system ensures properties of subject congruence, subject reduction, and progress. All auxiliary definitions and proof details are in App. D.

The property of subject congruence states that if \mathcal{S} (a system containing an application and a coordinator) is well-typed by some session environment, then a \mathcal{S}' that is structurally congruent to it is also well-typed by the same session environment:

Theorem 1 (Subject Congruence). $\Gamma \vdash \mathcal{S} \triangleright \Delta$ and $\mathcal{S} \equiv \mathcal{S}'$ imply $\Gamma \vdash \mathcal{S}' \triangleright \Delta$.

Subject reduction states that a well-typed \mathcal{S} (coherent session environment respectively) is always well-typed (coherent respectively) after reduction:

Theorem 2 (Subject Reduction).

- $\Gamma \vdash \mathcal{S} \triangleright \Delta$ with $\Gamma \vdash \Delta$ coherent and $\mathcal{S} \rightarrow^* \mathcal{S}'$ imply that $\exists \Delta', \Gamma'$ such that $\Gamma' \vdash \mathcal{S}' \triangleright \Delta'$ and $\Gamma \vdash \Delta \rightarrow_T^* \Gamma' \vdash \Delta'$ or $\Delta \equiv \Delta'$ and $\Gamma \vdash \Delta'$ coherent.
- $\Gamma \vdash \mathcal{S} \triangleright \emptyset$ and $\mathcal{S} \rightarrow^* \mathcal{S}'$ imply that $\Gamma' \vdash \mathcal{S}' \triangleright \emptyset$ for some Γ' .

We allow sessions to run in parallel at the top level, e.g., $\mathcal{S} = (\nu s_1)(\Psi_1 \blacklozenge N_1) \mid \dots \mid (\nu s_n)(\Psi_n \blacklozenge N_n)$. Assume we have \mathcal{S} with $a[p].P \in \mathcal{S}$. If we cannot apply rule **(Link)**, \mathcal{S} cannot reduce. To prevent this kind of situation, we require \mathcal{S} to be *initializable* such that, $\forall a[p].P \in \mathcal{S}$, **(Link)** is applicable.

The following property states that \mathcal{S} never gets stuck (property of progress):

Theorem 3 (Progress). If $\Gamma \vdash \mathcal{S} \triangleright \emptyset$ and \mathcal{S} is initializable, then either $\mathcal{S} \rightarrow^* \mathcal{S}'$ and \mathcal{S}' is initializable or $\mathcal{S}' = \Psi \blacklozenge s : h \mid \dots \mid \Psi' \blacklozenge s' : h'$ and h, \dots, h' only contain failure notifications sent by coordinators and messages heading to failed participants.

After all processes in \mathcal{S} terminate, failure notifications sent by coordinators are left; thus the final system can be of the form $\Psi \blacklozenge s : h \mid \dots \mid \Psi' \blacklozenge s' : h'$, where h, \dots, h' only have failure notifications sent by coordinators and thus reduction rules **(CollectDone)**, **(IssueDone)**, and **(F)** will not be applied.

Minimality. The following proposition points out that, when all roles defined in a global type, say G , are robust, then the application obeying to G will never have interaction with a coordinator (i.e., interactions of the application are equivalent to those without a coordinator). This is an important property, as it states that our model does not incur coordination overhead when all participants are robust, or in failure-agnostic contexts as considered in previous MPST works.

Proposition 2. Assume $\forall p \in \text{roles}(G) = \{p_1, \dots, p_n\}$, p is robust and $P_i = s[p_i] : \eta_i$ for $i \in \{1..n\}$ and $\mathcal{S} = (\nu s)(\Psi \blacklozenge P_1 \mid \dots \mid P_n \mid s : h)$ where $P_i, i \in \{1..n\}$ contains no try-handle. Then we have $\Gamma \vdash \mathcal{S} \triangleright \emptyset$ and whenever $\mathcal{S} \rightarrow^* \mathcal{S}'$ we have $\Psi \in \mathcal{S}', \Psi = G : (\emptyset, \emptyset)$.

Proof. Immediately by typing rules [T-**ini**, T-**s**, T-**sys**], Def. 4 (Projection), and the operational semantics defined in Fig. 6, Fig. 7, and Fig. 8.

8 Related Work

Several session type works study exception handling [7,9,16,30]. However, to the best of our knowledge this is the first theoretical work to develop a formalism and typing discipline for the coordinator-based model of *crash failure* handling in practical asynchronous distributed systems.

Structured interactional exceptions [7] study exception handling for binary sessions. The work extends session types with a *try-catch* construct and a *throw* instruction, allowing participants to raise runtime exceptions. Global escape [6]

extends previous works on exception handling in binary session types to MPSTs. It supports nesting and sequencing of try-catch blocks with restrictions. Reduction rules for exception handling are of the form $\Sigma \vdash P \rightarrow \Sigma' \vdash P'$, where Σ is the *exception environment*. This central environment at the core of the semantics is updated synchronously and atomically. Furthermore, the reduction of a try-catch block to its continuation is done in a synchronous reduction step involving all participants in a block. Lastly this work can only handle exceptions, i.e., explicitly raised application-level failures. These do not affect communication channels [6], unlike participant crashes.

Similarly, our previous work [13] only deals with exceptions. An interaction $p \rightarrow q : S \vee F$ defines that p can send a message of type S to q . If F is not empty then instead of sending a message p can throw F . If a failure is thrown only participants that have casual dependencies to that failure are involved in the failure handling. No concurrent failures are allowed therefore all interaction which can raise failures is executed in a lock step fashion. As a consequence, the model can not be used to deal with crash failures.

Adameit et al. [1] propose session types for link failures, which extend session types with an optional block which surrounds a process and contains default values. The default values are used if a link failure occurs. In contrast to our work, the communication model is overall synchronous whereas our model is asynchronous; the optional block returns default values in case of a failure but it is still the task of the developer to do something useful with it.

Demangeon et al. study interrupts in MPSTs [16]. This work introduces an interruptible block $\{|G|\}^c \langle l \text{ by } \mathbf{r} \rangle; G'$ identified by c ; here the protocol G can be interrupted by a message l from \mathbf{r} and is continued by G' after either a normal or an interrupted completion of G . Interrupts are more a control flow instruction like exceptions than an actual failure handling construct, and the semantics can not model participant crashes.

Neykova and Yoshida [37] show that MPSTs can be used to calculate safe global states for a safe recovery in Erlang's *let it crash* model [2]. That work is well suited for recovery of lightweight processes in an actor setting. However, while it allows for elaborate failure handling by connecting (endpoint) processes with runtime monitors, the model does not address the fault tolerance of runtime monitors themselves. As monitors can be interacting in complex manners replication does not seem straightforwardly applicable, at least not without potentially hampering performance (just as with *straightforward* replication of entire applications).

Failure handling is studied in several process calculi and communication-centered programming languages without typing discipline. The conversation calculus [43] models exception behavior in abstract service-based systems with message-passing based communication. The work does not use channel types but studies the behavioral theory of bisimilarity. Error recovery is also studied in a concurrent object setting [45]; interacting objects are grouped into coordinated atomic actions (CAs) which enable safe error recovery. CAs can however not be nested. PSYNC [18] is a domain specific language based on the *heard-of*

model of distributed computing [12]. Programs written in PSYNC are structured into rounds which are executed in a lock step manner. PSYNC comes with a state-based verification engine which enables checking of safety and liveness properties; for that programmers have to define non-trivial inductive invariants and ranking functions. In contrast to the coordinator model, the heard-of model is not widely deployed in practice. Verdi [44] is a framework for implementing and verifying distributed systems in Coq. It provides the possibility to verify the system against different network models. Verdi enables the verification of properties in an idealized fault model and then transfers the guarantees to more realistic fault models by applying transformation functions. Verdi supports safety properties but no liveness properties.

9 Final Remarks

Implementation. Based on our presented calculus we developed a domain-specific language and corresponding runtime system in Scala, using ZooKeeper as the coordinator. Specifically our implementation provides mechanisms for (1) interacting with ZooKeeper as coordinator, (2) done and failure notification delivery and routing, (3) practical failure detection and dealing with false suspicions and (4) automatically inferring try-handle levels (see more in App. C).

Conclusions. This work introduces a formal model of verified crash failure handling featuring a lightweight coordinator as common in many real-life systems. The model carefully exposes potential problems that may arise in distributed applications due to partial failures, such as inconsistent endpoint behaviors and orphan messages. Our typing discipline addresses these challenges by building on the mechanisms of MPSTs, e.g., global type well-formedness for sound failure handling specifications, modeling asynchronous permutations between regular messages and failure notifications in sessions, and the type-directed mechanisms for determining correct and orphaned messages in the event of failure. We adapt coherence of session typing environments (i.e., endpoint consistency) to consider failed roles and orphan messages, and show that our type system statically ensures subject reduction and progress in the presence of failures.

Future work. We plan to expand our implementation and develop further applications. We believe dynamic role participation and role parameterization would be valuable for failure handling. Also, we are investigating options to enable addressing the coordinator as part of the protocol so that pertinent runtime information can be persisted by the coordinator. We plan to add support to our language and calculus for solving various explicit agreement tasks (e.g., consensus, atomic commit) via the coordinator.

References

1. Adameit, M., Peters, K., Nestmann, U.: Session Types for Link Failures. In: FORTE '17. vol. 10321, pp. 1–16. Springer (2017)

2. Armstrong, J.: Making reliable distributed systems in the presence of software errors. Ph.D. thesis, Royal Institute of Technology, Stockholm, Sweden (2003)
3. Birman, K.P.: Byzantine clients (2017), <https://goo.gl/1Qbc4r>
4. Burrows, M.: The Chubby Lock Service for Loosely-Coupled Distributed Systems. In: OSDI '06. pp. 335–350. USENIX Association (2006)
5. Caires, L., Pérez, J.A.: Multiparty Session Types Within a Canonical Binary Theory, and Beyond. In: FORTE '16. LNCS, vol. 9688, pp. 74–95. Springer (2016)
6. Capecchi, S., Giachino, E., Yoshida, N.: Global escape in multiparty sessions. *MSCS* 26(2), 156–205 (2016)
7. Carbone, M., Honda, K., Yoshida, N.: Structured Interactional Exceptions in Session Types. In: CONCUR '08. LNCS, vol. 5201, pp. 402–417. Springer (2008)
8. Carbone, M., Lindley, S., Montesi, F., Schürmann, C., Wadler, P.: Coherence Generalises Duality: A Logical Explanation of Multiparty Session Types. In: CONCUR '16. *LIPICs*, vol. 59, pp. 33:1–33:15. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2016)
9. Carbone, M., Yoshida, N., Honda, K.: Asynchronous Session Types: Exceptions and Multiparty Interactions. In: SFM 2009. LNCS, vol. 5569, pp. 187–212. Springer (2009)
10. Chandra, T.D., Hadzilacos, V., Toueg, S., Charron-Bost, B.: On the Impossibility of Group Membership. In: PODC '96. pp. 322–330. ACM (1996)
11. Chang, F., Dean, J., Ghemawat, S., Hsieh, W.C., Wallach, D.A., Burrows, M., Chandra, T., Fikes, A., Gruber, R.: Bigtable: A Distributed Storage System for Structured Data. In: OSDI '06. pp. 205–218. USENIX Association (2006)
12. Charron-Bost, B., Schiper, A.: The Heard-Of model: computing in distributed systems with benign faults. *Distributed Computing* 22(1), 49–71 (2009)
13. Chen, T., Viering, M., Bejleri, A., Ziarek, L., Eugster, P.: A Type Theory for Robust Failure Handling in Distributed Systems. In: FORTE '16. vol. 9688, pp. 96–113. Springer (2016)
14. Coppo, M., Dezani-Ciancaglini, M., Padovani, L., Yoshida, N.: A Gentle Introduction to Multiparty Asynchronous Session Types. In: SFM '15. LNCS, vol. 9104, pp. 146–178. Springer (2015)
15. Coppo, M., Dezani-Ciancaglini, M., Yoshida, N., Padovani, L.: Global progress for dynamically interleaved multiparty sessions. *MSCS* 26(2), 238–302 (2016)
16. Demangeon, R., Honda, K., Hu, R., Neykova, R., Yoshida, N.: Practical Interruptible Conversations. *Formal Methods in System Design* 46(3), 197–225 (2015)
17. Deniérou, P.M., Yoshida, N.: Dynamic Multirole Session Types. In: POPL '11. pp. 435–446. ACM (2011)
18. Dragoi, C., Henzinger, T., Zufferey, D.: PSync: A Partially Synchronous Language for Fault-tolerant Distributed Algorithms. In: POPL '16. pp. 400–415. ACM (2016)
19. Fischer, M.J., Lynch, N.A., Paterson, M.S.: Impossibility of Distributed Consensus with One Faulty Process. In: PODS '85. pp. 374–382. ACM (1983)
20. Ghemawat, S., Gobioff, H., Leung, S.T.: The Google File System. In: SOSP '03. pp. 29–43. ACM (2003)
21. Gilbert, S., Lynch, N.: Brewer's Conjecture and the Feasibility of Consistent, Available, Partition-tolerant Web Services. *SIGACT News* 33(2), 51–59 (2002)
22. Guerraoui, R., Schiper, A.: The Generic Consensus Service. *IEEE Trans. Software Eng.* 27(1), 29–41 (2001)
23. Honda, K., Vasconcelos, V.T., Kubo, M.: Language Primitives and Type Discipline for Structured Communication-Based Programming. In: ESOP '98. LNCS, vol. 1381, pp. 122–138. Springer (1998)

24. Honda, K., Yoshida, N., Carbone, M.: Multiparty Asynchronous Session Types. *J. ACM* 63(1), 9:1–9:67 (2016)
25. Hu, R., Yoshida, N.: Hybrid Session Verification Through Endpoint API Generation . In: *FASE '16*. LNCS, vol. 9633, pp. 401–418. Springer (2016)
26. Hunt, P.: ZooKeeper: Wait-free Coordination for Internet-scale Systems. In: *USENIX '10*. USENIX Association (2010)
27. Hüttel, H., et al.: Foundations of Session Types and Behavioural Contracts. *ACM Comput. Surv.* 49(1), 3:1–3:36 (2016)
28. Imai, K., Yoshida, N., Yuen, S.: Session-ocaml: A Session-Based Library with Polarities and Lenses. In: *COORDINATION '17*. LNCS, vol. 10319, pp. 99–118. Springer (2017)
29. Killian, C.E., Anderson, J.W., Braud, R., Jhala, R., Vahdat, A.M.: Mace: Language Support for Building Distributed Systems. In: *PLDI '07*. vol. 42, pp. 179–188. ACM (2007)
30. Kouzapas, D., Yoshida, N.: Globally Governed Session Semantics. *LMCS* 10(4) (2014)
31. Kreps, J., Narkhede, N., Rao, J.: Kafka: A Distributed Messaging System for Log Processing. In: *NetDB '11* (2011)
32. Lamport, L.: The Part-time Parliament. *ACM Trans. Comput. Syst.* 16(2), 133–169 (1998)
33. Lamport, L., Shostak, R., Pease, M.: The Byzantine Generals Problem. *ACM Trans. Program. Lang. Syst.* 4(3), 382–401 (1982)
34. Leners, J.B., Wu, H., Hung, W.L., Aguilera, M.K., Walfish, M.: Detecting failures in distributed systems with the FALCON spy network. In: *SOSP '11*. pp. 279–294. ACM (2011)
35. Lindley, S., Morris, J.G.: Embedding Session Types in Haskell. In: *Haskell '16*. pp. 133–145. ACM (2016)
36. Mostrous, D., Yoshida, N.: Session typing and asynchronous subtyping for the higher-order π -calculus. *Inf. Comput.* 241, 227–263 (2015)
37. Neykova, R., Yoshida, N.: Let It Recover: Multiparty Protocol-Induced Recovery. In: *CC '17*. pp. 98–108. ACM (2017)
38. Padovani, L.: A simple library implementation of binary sessions. *J. Funct. Program.* 27, e4 (2017)
39. Pucella, R., Tov, J.A.: Haskell Session Types with (Almost) No Class. In: *Haskell '08*. pp. 25–36. ACM (2008)
40. Scalas, A., Dardha, O., Hu, R., Yoshida, N.: A Linear Decomposition of Multiparty Sessions for Safe Distributed Programming. In: *ECOOP '17*. LIPIcs, vol. 74, pp. 24:1–24:31. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2017)
41. Shvachko, K., Kuang, H., Radia, S., Chansler, R.: The Hadoop Distributed File System. In: *MSST '10*. pp. 1–10. IEEE Computer Society (2010)
42. Sivaramakrishnan, K.C., Qudeisat, M., Ziarek, L., Nagaraj, K., Eugster, P.: Efficient sessions. *Sci. Comput. Program.* 78(2), 147–167 (2013)
43. Vieira, H.T., Caires, L., Seco, J.C.: The Conversation Calculus: A Model of Service-Oriented Computation. In: *ESOP '08*. LNCS, vol. 4960, pp. 269–283. Springer (2008)
44. Wilcox, J.R., Woos, D., Panchekha, P., Tatlock, Z., Wang, X., Ernst, M.D., Anderson, T.E.: Verdi: A Framework for Implementing and Formally Verifying Distributed Systems. In: *PLDI '15*. pp. 357–368. ACM (2015)
45. Xu, J., Randell, B., Romanovsky, A.B., Rubira, C.M.F., Stroud, R.J., Wu, Z.: Fault Tolerance in Concurrent Object-Oriented Software Through Coordinated Error Recovery. In: *FTCS '95*. pp. 499–508. IEEE Computer Society (1995)

A Examples

This section provides interesting examples which are not included in the main content. We first give examples which are based on MPST and are extended with our explicit failure handling structure. Then we provide more examples of coordinator-based failure handling to show several non-trivial cases.

A.1 Examples of MPST Types and Processes with Coordinator-based Failure Handling

We give versions of the main motivating examples in standard MPST literature [24] extended to support partial failures.

Two-buyers protocol. The Two-Buyers Protocol [24, § 2.3,3.4] derives from a Web services use case. In the original protocol specification, the roles Buyer1 (B1) and Buyer2 (B2) carry out a joint transaction to buy a book from an online Seller (S). The role of B1 is to offer to pay some share of the total price to B2. B2 makes the final *choice* whether to proceed with the purchase (by paying the remaining amount) or not.

$$\begin{aligned}
 G_{\text{TB}} = & \text{t}(\\
 & \text{t}(\\
 & \quad \text{B2} \rightarrow \text{S } l_1(\text{string}). \\
 & \quad \text{S} \rightarrow \text{B1 } l_2(\text{int}). \\
 & \quad \text{S} \rightarrow \text{B2 } l_3(\text{int}). \\
 & \quad \text{B1} \rightarrow \text{B2 } l_4(\text{int}). \\
 & \quad \text{B2} \rightarrow \text{S } \begin{cases} \text{ok}_1(). \text{B2} \rightarrow \text{S } l_5(\text{string}). \text{S} \rightarrow \text{B2 } l_6(\text{date}). \text{end} \\ \text{quit}_1(). \text{end} \end{cases} \\
 & \quad \text{h}(\\
 & \quad \quad \{\text{B1}\} : \text{S} \rightarrow \text{B2 } l_7(\text{int}). \\
 & \quad \quad \text{B2} \rightarrow \text{S } \begin{cases} \text{ok}_2(). \text{B2} \rightarrow \text{S } l_8(\text{string}). \text{S} \rightarrow \text{B2 } l_9(\text{date}). \text{end} \\ \text{quit}_2(). \text{end} \end{cases} \\
 & \quad \quad \text{)}^2. \text{end} \\
 & \quad \text{h}(\\
 & \quad \quad \{\text{B2}\} : \text{end}, \\
 & \quad \quad \{\text{B1}, \text{B2}\} : \text{end} \\
 & \quad \quad \text{)}^1. \\
 & \text{end}
 \end{aligned}$$

Fig. 12. The Two-Buyers Protocol [24, § 3.4] extended with partial failure handling.

Fig. 12 extends the original global type to a fault tolerant version of this application protocol, bearing in mind the asymmetry of the B1 and B2 roles. We shall assume S is robust.

In the initial exchanges, given by the inner try-block, B2 sends S the name of the book (l_1), S sends B1 and B2 the price (l_2), and B1 sends B2 the amount it is willing to contribute (l_3) (as in the original protocol). If B1 crashes during this part, then S acknowledges this event by resending the price to B2 (l_7) in the inner handling environment—asynchrony and inherent concurrency of the interactions between S and the Buyers and this potential failure means that S cannot be certain about the order of the relevant messages arriving at B2.

Whether or not B1 crashes, B2 and S proceed to the choice of B2 to buy (`ok`) or not buy (`quit`) the book (this segment is also as in the original protocol). However, if B2 crashes at any point of the protocol, then the protocol must simply end for S, and also for B1 if it is live, as given by the outer handling environment.

Regarding the unique κ -annotations condition of well-formedness (Def. 1), as noted in Sec. 3, and implemented in App. C, it is straightforward to mechanically infer identifiers for an “unannotated” global type to satisfy this condition. As for many protocols, it is also straightforward to add labels to non-explicitly labelled interactions to satisfy the labelling condition. In this example, it would have been sufficient to use just three labels in total to distinguish the interactions in the try-block, the inner handler, and the choice construct; however, we used unique labels throughout to clarify the explanations given above.

Projection to local types. The local type projection (Def. 4) of G_{TB} onto B1, $G_{\text{TB}}|_{\text{B1}}$, is:

$$\begin{aligned} & \mathbf{t}(\mathbf{t}(\mathbf{S}?l_2(\mathbf{int}).\mathbf{B2}!l_4(\mathbf{int}).\mathbf{end})\mathbf{h}(\{\mathbf{B1}\}:\mathbf{end})^{(2,\emptyset)}.\mathbf{end}) \\ & \mathbf{h}(\{\mathbf{B2}\}:\mathbf{end}, \{\mathbf{B1}, \mathbf{B2}\}:\mathbf{end})^{(1,\emptyset)}.\mathbf{end} \end{aligned}$$

The projection onto B2, $G_{\text{TB}}|_{\text{B2}}$, is:

$$\begin{aligned} & \mathbf{t}(\mathbf{t}(\mathbf{S}!l_1(\mathbf{string}).\mathbf{S}?l_3(\mathbf{int}).\mathbf{B1}?l_4(\mathbf{int}). \\ & \mathbf{S}! \begin{cases} \mathbf{ok}_1().\mathbf{S}!l_5(\mathbf{string}).\mathbf{S}?l_6(\mathbf{date}).\mathbf{end} \\ \mathbf{quit}_1().\mathbf{end} \end{cases} \\ & \mathbf{h}(\{\mathbf{B1}\}:\mathbf{S}?l_7(\mathbf{int}). \\ & \mathbf{S}! \begin{cases} \mathbf{ok}_2().\mathbf{S}!l_8(\mathbf{string}).\mathbf{S}?l_9(\mathbf{date}).\mathbf{end} \\ \mathbf{quit}_2().\mathbf{end} \end{cases} \\ &)^{(2,\emptyset)} \\ & \mathbf{h}(\{\mathbf{B2}\}:\mathbf{end}, \{\mathbf{B1}, \mathbf{B2}\}:\mathbf{end})^{(1,\emptyset)}.\mathbf{end} \end{aligned}$$

The projection onto S, $G_{\text{TB}}|_{\text{S}}$, is:

$$\begin{aligned}
& \mathbf{t}(\\
& \quad \mathbf{t}(\\
& \quad \quad \mathbf{B2?}l_1(\mathbf{string}).\mathbf{B1!}l_2(\mathbf{int}).\mathbf{B2!}l_3(\mathbf{int}). \\
& \quad \quad \mathbf{B2?} \begin{cases} \mathbf{ok}().\mathbf{B2?}l_5(\mathbf{string}).\mathbf{B2!}l_6(\mathbf{date}).\mathbf{end} \\ \mathbf{quit}().\mathbf{end} \end{cases} \\
& \quad)\mathbf{h}(\\
& \quad \quad \{\mathbf{B1}\}:\mathbf{B2!}l_7(\mathbf{int}).\mathbf{end} \\
& \quad \quad \mathbf{B2?} \begin{cases} \mathbf{ok}().\mathbf{B2?}l_8(\mathbf{string}).\mathbf{B2!}l_9(\mathbf{date}).\mathbf{end} \\ \mathbf{quit}().\mathbf{end} \end{cases} \\
& \quad)^{(2,\emptyset)} \\
&)\mathbf{h}(\{\mathbf{B2}\}:\mathbf{end}, \{\mathbf{B1}, \mathbf{B2}\}:\mathbf{end})^{(1,\emptyset)}. \mathbf{end}
\end{aligned}$$

Processes. We give example of well-typed processes implementing each of roles.

For B1:

$$\begin{aligned}
& a[\mathbf{B1}](y).y : \mathbf{t}(\\
& \quad \mathbf{t}(\mathbf{S?}l_2(x_2).\mathbf{B2!}l_4(x_2 \div 2).0)\mathbf{h}(\{\mathbf{B1}\}:0)^{(2,\emptyset)}.0 \\
&)\mathbf{h}(\{\mathbf{B2}\}:0, \{\mathbf{B1}, \mathbf{B2}\}:0)^{(1,\emptyset)}.0
\end{aligned}$$

For B2:

$$\begin{aligned}
& a[\mathbf{B2}](y).y : \mathbf{def} \ X(z) = \mathbf{if} \ z < 100 \ \mathbf{S!ok}().\mathbf{S!}l_5(\text{“addr”}).\mathbf{S?}l_6(x_6).0 \\
& \quad \quad \quad \mathbf{else} \ \mathbf{S!quit}().0 \\
& \mathbf{in} \ \mathbf{t}(\\
& \quad \mathbf{t}(\mathbf{S!}l_1(\text{“title”}).\mathbf{S?}l_2(x_2).\mathbf{B1?}l_4(x_4).X\langle x_2 - x_4 \rangle) \\
& \quad \mathbf{h}(\{\mathbf{B1}\}:\mathbf{S?}l_7(x_7).X\langle x_7 \rangle)^{(2,\emptyset)}.0 \\
&)\mathbf{h}(\{\mathbf{B2}\}:0, \{\mathbf{B1}, \mathbf{B2}\}:0)^{(1,\emptyset)}.0
\end{aligned}$$

For S, assuming some helper functions `getPrice` and `getData` on data:

$$\begin{aligned}
& a[\mathbf{S}](y).y : \mathbf{def} \ X() = \mathbf{B2?}\{\mathbf{ok}().\mathbf{B2?}l_5(x_5).\mathbf{S?}l_6(\mathbf{getDate}(x_5)).0, \mathbf{quit}().0\} \\
& \mathbf{in} \ \mathbf{t}(\\
& \quad \mathbf{t}(\mathbf{B2?}l_1(x_1).\mathbf{B1!}l_2(\mathbf{getPrice}(x_1)).\mathbf{B2!}l_4(\mathbf{getPrice}(x_1)).X\langle \rangle) \\
& \quad \mathbf{h}(\{\mathbf{B1}\}:\mathbf{B1!}l_7(\mathbf{getPrice}(x_1)).X\langle \rangle)^{(2,\emptyset)}.0 \\
&)\mathbf{h}(\{\mathbf{B2}\}:0, \{\mathbf{B1}, \mathbf{B2}\}:0)^{(1,\emptyset)}.0
\end{aligned}$$

Streaming protocol. The original Streaming Protocol [24, §3.4] demonstrated a *recursive* global type for a continuous stream of messages, where two producer roles (DP, KP) independently send to a middleman role (K) in a join pattern, followed by K forwarding a message to a consumer role (C). (The role names are taken from the original protocol.)

Fig. 13 extends the protocol to handle the potential failures of the DP and KP; we assume the other two roles are robust. The idea is if just one of the producers crashes, the now fault tolerant protocol should attempt to continue

$$\begin{array}{l}
\mathbf{t}(\\
 \mu t. \\
 \text{DP} \rightarrow \text{K } l_1(\text{bool}). \\
 \text{KP} \rightarrow \text{K } l_2(\text{bool}). \\
 \text{K} \rightarrow \text{C } l_3(\text{bool}). \\
 \text{DP} \rightarrow \text{K } l_1(\text{bool}). \\
 \text{KP} \rightarrow \text{K } l_2(\text{bool}). \\
 \text{K} \rightarrow \text{C } l_3(\text{bool}). \\
 t \\
)\mathbf{h}(\\
 \{\text{DP}\} : \mu t'. \text{KP} \rightarrow \text{K } l_4(\text{bool}). \text{K} \rightarrow \text{C } l_5(\text{bool}). t' \\
 \{\text{KP}\} : \mu t''. \text{DP} \rightarrow \text{K } l_6(\text{bool}). \text{K} \rightarrow \text{C } l_7(\text{bool}). t'' \\
 \{\text{KP}, \text{KP}\} : \text{end} \\
)^1
\end{array}$$

Fig. 13. The Stream Protocol [24, § 3.4] extended with partial failure handling.

with the other producer. We keep the “once-unfolded” specification from the original protocol definition in the try-block, but use the shorter folded versions in the handling activities.

This example also gives some intuition for the complexity in the design of our failure handling constructs: *without* any explicit choice construct, adding failure handling naturally introduces a *safe* notion of choice into the protocol between distinct paths, and also allows the normally recursive protocol to end (if both producers crash).

A.2 Examples of Coordinator-based Failure Handling

In this part, we give more handling examples for readers who are interested in the coordinator-based failure handling.

Handling failure notifications. The following example shows handling mechanism of **(TryHdl)** and **(F)**:

Example 8. Let P_i play role p_i in session s and h does not contain any notifications. In this example both P_2 and P_3 fail and it shows a reduction for P_1 with a focus on failure handling. Assume we start with:

$$\begin{array}{l}
\dots G : (F, \emptyset) \blacklozenge P_1 \mid P_2 \mid P_3 \mid s : h \cdot \langle p_2, p_1, l(v) \rangle \\
\mathbf{(Crash)} \rightarrow G : (F, \emptyset) \blacklozenge P_1 \mid P_3 \mid s : h_1 \cdot \langle \psi, \text{crash } \{p_2\} \rangle \\
\phantom{\mathbf{(Crash)}} \text{where } h_1 = \text{remove}(h \cdot \langle p_2, p_1, l(v) \rangle, p_2) \\
\equiv G : (F, \emptyset) \blacklozenge P_1 \mid P_3 \mid s : \langle \psi, \text{crash } \{p_2\} \rangle \cdot h_1 \\
\mathbf{(F)} \rightarrow G : (F \cup \{p_2\}, \emptyset) \blacklozenge P_1 \mid P_3 \mid s : h_1 \cdot \langle \{p_1, p_3\}, \text{crash } \{p_2\} \rangle
\end{array}$$

Note that, by Definition 13 (Permutable Messages), because h (therefore also h_1) does not contain any notification message the notification $\langle \psi, \text{crash } \{p_2\} \rangle$

can be permuted to the top of queue. Next the the coordinator receives notification $\langle \psi, \text{crash } \{p_2\} \rangle$ and issues the failure notification $\langle \{p_1, p_3\}, \text{crash } \{p_2\} \rangle$ to P_1 and P_3 for handling the failure of P_2 by applying **(F)**.

Assume (possible concurrently) that P_3 fails, i.e. **(Crash)** is applied to P_3

$$\begin{aligned} \text{(Crash)} \rightarrow G : (F \cup \{p_2\}, \emptyset) \blacklozenge P_1 \mid \mathbf{s} : h_2 \cdot \langle \{p_1\}, \text{crash } \{p_2\} \rangle \cdot \langle \psi, \text{crash } \{p_3\} \rangle \\ \text{where } h_2 \cdot \langle \{p_1\}, \text{crash } \{p_2\} \rangle = \text{remove}(h_1 \cdot \langle \{p_1, p_3\}, \text{crash } \{p_2\} \rangle, p_3) \end{aligned}$$

Assume we have $P_1 = \mathbf{s}[p_1] : \mathbf{t}(\eta)\mathbf{h}(\mathbf{H})^{(\phi, \emptyset)}$ and

$\mathbf{H} = \{p_2\} : \eta_2, \{p_3\} : \eta_3 \{p_1, p_3\} : \eta_{13}$. The reduction shows P_1 starts handling the failure $\{p_2\}$ (P_1 is not aware of the failure of P_3), then the coordinator will pick up the notification that P_3 has failed and forward it. Lastly P_1 will perform the handling for both failures:

$$\begin{aligned} \text{(TryHdl)} \rightarrow G : (F \cup \{p_2\}, \emptyset) \blacklozenge \mathbf{s}[p_1] : \mathbf{t}(\eta_2)\mathbf{h}(\mathbf{H})^{(\phi, \{p_2\})} \mid \\ \mathbf{s} : h_2 \cdot \langle \{p_1\}, \text{crash } \{p_2\} \rangle \cdot \langle \psi, \text{crash } \{p_3\} \rangle \\ \equiv, \text{(F)} \rightarrow G : (F \cup \{p_2, p_3\}, \emptyset) \blacklozenge \mathbf{s}[p_1] : \mathbf{t}(\eta_2)\mathbf{h}(\mathbf{H})^{(\phi, \{p_2\})} \mid \\ \mathbf{s} : h_2 \cdot \langle \{p_1\}, \text{crash } \{p_2\} \rangle \cdot \langle \{p_1\}, \text{crash } \{p_3\} \rangle \\ \text{(TryHdl)} \rightarrow G : (F \cup \{p_2, p_3\}, \emptyset) \blacklozenge \mathbf{s}[p_1] : \mathbf{t}(\eta_{23})\mathbf{h}(\mathbf{H})^{(\phi, \{p_2, p_3\})} \mid \\ \mathbf{s} : h_2 \cdot \langle \{p_1\}, \text{crash } \{p_2\} \rangle \cdot \langle \{p_1\}, \text{crash } \{p_3\} \rangle \\ \dots \end{aligned}$$

The rest of the reduction is straight forward by using **(SndDone)**, **(CollectDone)**, **(IssueDone)** and **(RcvDone)**.

Handling mixed notifications. We illustrate the hardness of consistent handling of mixed notifications (i.e., done and failure notifications), and the need for a coordinator.

Example 9. Assume we have processes P_1, P_2, P_3 acting respectively on channels $\mathbf{s}[p_1], \mathbf{s}[p_2]$ and $\mathbf{s}[p_3]$. Given Eq. (1), in which they have finished their respective actions in the try-handle of $\phi = (\kappa, \emptyset)$, and the queue contains the done notifications sent to the coordinator

$$\begin{aligned} (1) \mathbf{s}[p_1] : \mathbf{t}(\mathbf{0})\mathbf{h}(\mathbf{H}_{p_1})^\phi \cdot \eta_1 \mid \mathbf{s}[p_2] : \mathbf{t}(\mathbf{0})\mathbf{h}(\mathbf{H}_{p_2})^\phi \cdot \eta_2 \mid \mathbf{s}[p_3] : \mathbf{t}(\mathbf{0})\mathbf{h}(\mathbf{H}_{p_3})^\phi \cdot \eta_3 \mid \\ \mathbf{s} : \langle p_2, \psi \rangle^\phi \cdot \langle p_1, \psi \rangle^\phi \cdot \langle p_3, \psi \rangle^\phi \end{aligned}$$

where $\mathbf{H}_{p_1} = \{p_2\} : \eta_1, \mathbf{H}'_{p_1}$ and $\mathbf{H}_{p_3} = \{p_2\} : \eta_3, \mathbf{H}'_{p_3}$. The coordinator is the key in ensuring consistency in the presence of concurrent done and failure notifications. If no participant fails, the coordinator collects $\langle p_2, \psi \rangle^\phi, \dots$ with **(CollectDone)**; and the coordinator issues done notifications $\langle \psi, \{p_1, p_2, p_3\} \rangle^\phi$ via **(IssueDone)**. All participants will then finish the try-handle of ϕ and move to their next actions, as shown by (1) \rightarrow^* (2)

$$\begin{aligned} (2) \mathbf{s}[p_1] : \mathbf{t}(\mathbf{0})\mathbf{h}(\mathbf{H}_{p_1})^\phi \cdot \eta_1 \mid \mathbf{s}[p_2] : \mathbf{t}(\mathbf{0})\mathbf{h}(\mathbf{H}_{p_2})^\phi \cdot \eta_2 \mid \mathbf{s}[p_3] : \mathbf{t}(\mathbf{0})\mathbf{h}(\mathbf{H}_{p_3})^\phi \cdot \eta_3 \mid \\ \mathbf{s} : \langle \psi, p_1 \rangle^\phi \cdot \langle \psi, p_2 \rangle^\phi \cdot \langle \psi, p_3 \rangle^\phi \\ \rightarrow^* \mathbf{s}[p_1] : \eta_1 \mid \mathbf{s}[p_2] : \eta_2 \mid \mathbf{s}[p_3] : \eta_3 \mid \mathbf{s} : \emptyset \end{aligned}$$

If P_2 fails at the moment of Eq. (1), notification $\langle p_2, \psi \rangle^\phi$ is removed from the queue and therefore the try-handle of ϕ cannot finish, as shown by (1) \rightarrow^* (3)

$$(3) \mathbf{s}[p_1] : \mathbf{t}(\mathbf{0})\mathbf{h}(\mathbf{H}_{p_1})^\phi \cdot \eta_1 \mid \mathbf{s}[p_3] : \mathbf{t}(\mathbf{0})\mathbf{h}(\mathbf{H}_{p_3})^\phi \cdot \eta_3 \mid \mathbf{s} : \langle p_1, \psi \rangle^\phi \cdot \langle p_3, \psi \rangle^\phi \cdot \langle \psi, \text{crash } \{p_2\} \rangle$$

and a notification for the failure of p_2 is added to the queue. When the coordinator receives the failure notification $\langle\psi, \text{crash } \{p_2\}\rangle$, it issues $\langle p_1, \text{crash } \{p_2\}\rangle$ and $\langle p_3, \text{crash } \{p_2\}\rangle$ to P_1 and P_3 . Now P_1 and P_3 will perform the handling and change the level of the try-handle to $\phi' = (\kappa, \{p_2\})$, as shown by (3) \rightarrow^* (4)

$$(4) \text{ s}[p_1] : \text{t}(\eta_1)\text{h}(\text{H}_{p_1})^{\phi'}. \eta_1 \mid \text{ s}[p_3] : \text{t}(\eta_3)\text{h}(\text{H}_{p_3})^{\phi'}. \eta_3 \mid \\ \text{ s} : \langle p_1, \text{crash } \{p_2\}\rangle \cdot \langle p_3, \text{crash } \{p_2\}\rangle$$

B Detailed Formal Definitions

This section contains additional detailed formal definitions.

B.1 Definitions for Section 3 (Global Types)

Well-formedness. In the following we provide formal definitions for those used in explaining well-formed global types in Section 3. We formally define the syntactical well-formedness conditions.

For the well-formedness definition we need to address nesting of try-handles. We define the global contexts \mathcal{G} as follows:

$$\mathcal{G} ::= [] \mid \text{t}(\mathcal{G})\text{h}(H)^\kappa.G \mid \text{t}(G)\text{h}(F:\mathcal{G}, H)^\kappa.G' \mid \\ \text{t}(G)\text{h}(H)^\kappa.\mathcal{G} \mid p \rightarrow q\{l_i(S_i).G_i, l(S).\mathcal{G}\}_{i \in I} \mid \mu t.\mathcal{G}$$

Such contexts allow us to reason about parts of global types. Based on these contexts we define a containment relation on global types as follows:

Definition 9 ($G' \in G$). If $\exists \mathcal{G}$ s.t. $G = \mathcal{G}[G']$, then $G' \in G$

$G' \in G$ means G' is a part of global type G . Analogously to $G' \in G$, we write $G \in H$ if the handling environment H contains G ; $\kappa \in G$ if G contains κ (remember κ is shorthand for $\text{t}(G_1)\text{h}(H)^\kappa$); $\kappa \in \kappa'$; $l \in G$ if the label l appears inside G ; and $l \in \mathcal{G}$ if the label l appears inside \mathcal{G} . We use a lookup function $\text{outer}_G(\kappa)$ for the set of all try-handles in G that enclose a given κ (including κ itself):

Definition 10 ($\text{outer}_G(\kappa)$). $\text{outer}_G(\kappa) = \{\kappa' \mid \kappa \in \kappa' \wedge \kappa' \in G\}$.

As we mentioned in Sec. 3, we require that all κ are unique. Based on Def. 9 and Def. 10, we define well-formedness of global types where Conditions 1 and 2 are the same as those stated in Sec. 3 Def. 1, and Conditions 3–5 define other conditions informally described in Sec. 3.

Definition 11 (Well-formedness (Full Version)). Let κ stand for $\text{t}(G_1)\text{h}(H)^\kappa$, and κ' for $\text{t}(G'_1)\text{h}(H')^{\kappa'}$. A top-level $[p]G$ is *well-formed* if it fulfills *all* following conditions. For all $\kappa \in G$:

1. For any two separate handler signatures of a handling environment of κ , there always exists a handler whose handler signature matches the union of the respective failure sets. This handler is either inside the handling environment of κ itself, or in the handling environment of an outer try-handle:

$$\forall F_1 \in \text{dom}(H). \forall F_2 \in \text{dom}(H). \exists \kappa' \in \text{outer}_G(\kappa) \text{ s.t. } F_1 \cup F_2 \in \text{dom}(H')$$

2. If the handling environment of a try-handle κ contains a handler for F , then there is no outer try-handle κ' with a handler for F' such that $F' \subseteq F$:

$$\nexists F \in \text{dom}(H). \exists \kappa' \in \text{outer}_G(\kappa). \exists F' \in \text{dom}(H') \text{ s.t. } \kappa' \neq \kappa \wedge F' \subseteq F$$

3. All κ in G are unique. In addition all labels which appear inside a default try-body or any handler body do not occur outside of the default block/the handler body:

$$\begin{aligned} G = \mathcal{G}[\mathfrak{t}(G_1)\mathfrak{h}(H)^\kappa.G_2] \Rightarrow \\ \forall l \in G_1. l \notin \mathcal{G}, G_2, H \text{ and} \\ \forall F, F' \in H. \text{ s.t. } F \neq F' \forall l' \in H(F). l' \notin \mathcal{G}, G_1, G_2, H(F') \end{aligned}$$

4. A role does not appear in the handling activity of its own failure:

$$\forall F \in \text{dom}(H). \forall p \in F \Rightarrow p \notin \text{roles}(H(F))$$

5. All branching types of non-robust participants ($p \notin \tilde{p}$) must be handled in G , i.e., must be enclosed by try-handles which can handle the potential failures of p :

$$\begin{aligned} \forall q \rightarrow q' \{l_i(S_i).G_i\}_{i \in I} \in G. \forall p \in \{q, q'\}. p \notin \tilde{p} \Rightarrow \\ \exists \kappa' \in G. q \rightarrow q' \{l_i(S_i).G_i\}_{i \in I} \in \kappa' \wedge \{p\} \in \text{dom}(H') \end{aligned}$$

Condition 1 ensures that if roles are active in different handlers of the same try-handle, there is a handler whose signature corresponds to the union over the signatures of those different handlers. Example 2 together with Example 3 in Sec. 4 show why this condition is needed. The reason for Condition 2 is that in case of nested try-handles the operational semantics (see **TryHdl**) in Sec. 4 Fig. 6) allows multiple try-handles to start failure handling; eventually the outer-most try-handle will perform the handling and may interrupt failure handling at an inner try-handle. $G = \mathfrak{t}(\mathfrak{t}(G')\mathfrak{h}(\{p_1, p_2\} : G_1)^2)\mathfrak{h}(\{p_1\} : G_1)^1$ violates condition 2 because, when p_1 and p_2 both failed, the handler signature $\{p_1\}$ will still be triggered (i.e., the outer try-handle will eventually take over). It is not sensible to run G_1' instead of G_1 (which is for the crashes of p_1 and p_2). The reason for Condition 3 is that a try-handle is our handling basis, and we shall not confuse normal messages or done notifications from different try-handles. Condition 4 is straightforward. Condition 5 is also straightforward as we can only handle failures occurring in a try-handle therefore interaction that involves non-robust roles needs to be inside try-handles.

Examples of well-formed and ill-formed global types.

Example 10. $G_{ko} = \mathbf{t}(\mathbf{t}(G')\mathbf{h}(\{p_1\} : G_1)^2)\mathbf{h}(\{p_1\} : G'_1)$ violates Condition 2 because there are two different handling activities for $\{p_1\}$ at different nesting levels. Since the outer try-handle will eventually take over, it is not sensible to have the handling activity $\{p_1\}$ at the inner one.

Example 11. $G_{ko} = \mathbf{t}(p_1 \rightarrow p_2 \ l_1(S_1))\mathbf{h}(\{p_1\} : p_1 \rightarrow p_2 \ l_2(S_2))$ violates Condition 4. It is not well-formed since the handling activity of $\{p_1\}$ has $p_1 \rightarrow p_2 \ l_2(S_2)$ in which p_1 is expected to output a message, yet p_1 would have failed at that point.

Example 12. Consider

$$\begin{aligned} G_{ko} &= \mathbf{t}(p_1 \rightarrow p_2 \ l_1(S_1))\mathbf{h}(\{p_1\} : p_2 \rightarrow p_3 \ l_3(S_3)) \cdot p_1 \rightarrow p_3 \ l_2(S_2) \\ G_{ok} &= \mathbf{t}(p_1 \rightarrow p_2 \ l_1(S_1) \cdot p_1 \rightarrow p_3 \ l_2(S_2))\mathbf{h}(\{p_1\} : p_2 \rightarrow p_3 \ l_3(S_3)) \end{aligned}$$

G_{ko} violates Condition 5 since the non-robust p_1 is expected to perform a message send $p_1 \rightarrow p_3 \ l_2(S_2)$ which is not enclosed by any try-handle. This is not safe because, if p_1 crashes before $p_1 \rightarrow p_3 \ l_2(S_2)$ completes, then p_3 will get stuck. On the other hand, G_{ok} is well-formed because all interactions containing the non-robust participant p_1 are enclosed in appropriate $\mathbf{t}(\cdot)\mathbf{h}(\cdot)$ blocks and all other properties hold.

B.2 Definitions for Section 4 (Process Calculus)

In the following we provide formal definitions for those we have used in Sec. 4 but have not yet formally defined.

Structural congruence. Processes, applications, and systems are considered modulo structural equivalence, denoted by \equiv , and defined by the rules in Def. 12 along with α -renaming.

Definition 12 (Structural Congruence).

$$\begin{aligned} h &\equiv \emptyset \cdot h \equiv h \cdot \emptyset & h_1 \cdot (h_2 \cdot h_3) &\equiv (h_1 \cdot h_2) \cdot h_3 \\ \frac{m \cdot m' \curvearrowright m' \cdot m}{h \cdot m \cdot m' \cdot h' \equiv h \cdot m' \cdot m \cdot h'} & & \frac{h \equiv h'}{s : h \equiv s : h'} \\ \text{def } D \text{ in } 0 &\equiv 0 \\ \text{def } D \text{ in } (\text{def } D' \text{ in } \eta) &\equiv \text{def } D' \text{ in } (\text{def } D \text{ in } \eta) \\ \text{if } dpv(D) \cap (dpv(D') \cup fpv(\eta)) = dpv(D') \cap (dpv(D) \cup fpv(\eta)) = \emptyset & \\ \Psi \blacklozenge 0 &\equiv 0 & \frac{N \equiv N'}{\Psi \blacklozenge N \equiv \Psi \blacklozenge N'} \\ (\nu s)(\nu s')\mathcal{S} &\equiv (\nu s')(\nu s)\mathcal{S} & (\nu s)\mathcal{S}|N \equiv (\nu s)(\mathcal{S}|N) \text{ if } s \notin fn(N) \end{aligned}$$

The rule **(Str)** in Fig. 6 in Sec. 4 uses this definition.

In Def. 12, the rules in the first two lines allow the permutation of messages. $m \cdot m' \curvearrowright m' \cdot m$ means that the order of $m \cdot m'$ can be switched to $m' \cdot m$. Permutation is detailed below. The rules in the next two lines gives structural congruence over recursions. Function $dpv(D)$ gives the set of process variables in declarations, and $fpv(\eta)$ gives the set of process variables which occur free in η . The final two line of rules state structural congruence of systems, each of which combines an application N and a coordinator Ψ . Function $fn(N)$ gives the set of free names in N . The structural congruence rules for restriction and scope extension are standard.

Permutation is possible, in short, as soon as messages and notifications have different sources or destinations they can be permuted ; also the order of done notifications from different levels (e.g., $\phi \neq \phi'$) can be permuted.

Definition 13 (Permutable Messages). We define $m_i \cdot m_j \curvearrowright m_j \cdot m_i$, $i \neq j$, saying $m_i \cdot m_j$ can be permuted to $m_j \cdot m_i$, if *none* of the following conditions holds:

- $m_i = \langle p, q, l(v) \rangle$ and $m_j = \langle p, q, l'(v') \rangle$ for some l, l', v, v' .
- $m_i = \langle \psi, q \rangle^\phi$ and $m_j = \langle q, \text{crash } F \rangle$ for some ϕ, F .

Extracting labels. Rule **(Cln)** in Fig. 7 in Sec. 4 removes messages if the messages contains a label which cannot be reached. (This situation could happen if the process applied **(TryHdl)**).

Def. 14 $labels(\eta)$ extracts all labels of receiving actions in η which potentially can receive messages. The only receiving actions which cannot and never will be able to receive messages in η are those in a handler with signature F in a try-handle (κ, F') where $F \subseteq F'$ (**(TryHdl)** cannot switch the handling to F).

Definition 14 (Extracting Reachable Labels in η).

$$\begin{aligned}
labels(p!\{l_i(e_i).\eta_i\}_{i \in I}) &= \bigcup_{i \in I} labels(\eta_i) \\
labels(p?\{l_i(e_i).\eta_i\}_{i \in I}) &= \bigcup_{i \in I} (\{l_i\} \cup labels(\eta_i)) \\
labels(\mathbf{t}(\eta)\mathbf{h}(\mathbf{H})^{(\kappa, \emptyset)}.\eta') &= labels(\eta) \cup labels(\eta') \bigcup_{(F:\eta'') \in \mathbf{H}} labels(\eta'') \\
labels(\mathbf{t}(\eta)\mathbf{h}(\mathbf{H})^{(\kappa, F)}.\eta') &= labels(\eta) \cup labels(\eta') \bigcup_{(F':\eta'') \in \mathbf{H} \wedge F \subset F'} labels(\eta'') \\
labels(0) &= \emptyset \\
labels(\underline{0}) &= \emptyset \\
labels(X(e)) &= \emptyset \\
labels(\mathbf{def } D \mathbf{ in } \eta) &= labels(\eta) \cup labels(D) \\
labels(\mathbf{if } e \eta \mathbf{ else } \eta') &= labels(\eta) \cup labels(\eta') \\
labels(X(x)) &= labels(\eta)
\end{aligned}$$

B.3 Definitions for Section 6 (Type System)

We define $T \downarrow p$ in Def. 15 as a filter to get the partial type which only contains actions of p in T . For example

$$p_1!l'(S').p_2!l(S) \downarrow p_2 = p_2!l(S)$$

and

$$p_1!\{l_1(S_1).p_2?l(S), l_2(S_2).p_2?l(S)\} \downarrow p_2 = p_2?l(S)$$

Definition 15 (Behaviors for p in T).

$$\begin{aligned} \mathbf{t}(T)\mathbf{h}(F_1:T_1, \dots, F_n:T_n)^\phi.T' \downarrow p &= \\ &\begin{cases} \mathbf{t}(T \downarrow p)\mathbf{h}(F_1:T_1 \downarrow p, \dots, F_n:T_n \downarrow p)^\phi.T' \downarrow p & \text{if } p \in \text{roles}(T) \cup \bigcup_{i=1}^n \text{roles}(T_i) \\ T' \downarrow p & \text{otherwise} \end{cases} \\ q!\{l_i(S_i).T_i\}_{i \in I} \downarrow p &= \begin{cases} p!\{l_i(S_i).T_i \downarrow p\}_{i \in I} & \text{if } p = q \\ T_i \downarrow p & \text{otherwise} \end{cases} \\ q?\{l_i(S_i).T_i\}_{i \in I} \downarrow p &= \begin{cases} p?\{l_i(S_i).T_i \downarrow p\}_{i \in I} & \text{if } p = q \\ T_i \downarrow p & \text{otherwise} \end{cases} \\ \text{end} \downarrow p = \text{end} \quad t \downarrow p = t \quad \mu t.T \downarrow p &= \begin{cases} \mu t.T \downarrow p & \text{if } T \downarrow p \neq t' \text{ for any } t' \\ \text{end} & \text{otherwise} \end{cases} \end{aligned}$$

$\text{roles}(T)$ is used (again by slight abuse of notation) to represent the set of roles appearing in T . By the final case of Def. 4 (Projection).(2) if p does not select or receive a label, then the action of p shall be the same for any branch.

Next we define $(\mathbf{h})_{p \rightarrow q}$ in Def. 16 to generate (1) the normal message types sent from p heading to q , and (2) the notifications heading to q :

Definition 16 (Message Types to q). We define $(\mathbf{h})_{p \rightarrow q}$ by selecting message types and notifications in \mathbf{h} which are (sent from p) heading to q . Let

$$\mathbf{mt} ::= p?l(S) \mid \langle\!\langle F \rangle\!\rangle \mid \langle\psi\rangle^\phi$$

and

$$\mathbf{ht} ::= \emptyset \mid \mathbf{mt} \cdot \mathbf{mt}$$

and \mathbf{ht} range over by $(\mathbf{h})_{p \rightarrow q}$

$$(\emptyset)_{p \rightarrow q} = \emptyset \quad (\mathbf{m} \cdot \mathbf{h})_{p \rightarrow q} = \begin{cases} p?l(S) \cdot (\mathbf{h})_{p \rightarrow q} & \text{if } \mathbf{m} = \langle p, q, l(S) \rangle \cdot \mathbf{h} \\ \langle\!\langle F \rangle\!\rangle \cdot (\mathbf{h})_{p \rightarrow q} & \text{if } \mathbf{m} = \langle\!\langle q, \text{crash } F \rangle\!\rangle \cdot \mathbf{h} \\ \langle\psi\rangle^\phi \cdot (\mathbf{h})_{p \rightarrow q} & \text{if } \mathbf{m} = \langle\psi, q\rangle^\phi \cdot \mathbf{h} \\ (\mathbf{h})_{p \rightarrow q} & \text{otherwise} \end{cases}$$

We define $T\text{-ht}$ in Def. 17, meaning the effect of \mathbf{ht} on T . Its concept is similar to the *session remainder* defined in [36], which returns new local types of participants after participants consume messages in the global queue. Since failure notifications will not be consumed in our system, and we only have to observe the change of a participant's type after receiving or being triggered by some messages types or notifications in \mathbf{ht} , we say that $T\text{-ht}$ represents the effect of \mathbf{ht} on T .

First we define session environment contexts $\mathcal{E} ::= [] \mid \mathbf{t}(\mathcal{E})\mathbf{h}(\mathcal{H})^\phi.T \mid \mu t.\mathcal{E}$ and define $Fset(\mathbf{ht}, p)$ similarly to Def. 2 by replacing queues with \mathbf{ht} . The permutation of \mathbf{ht} is similarly defined as Def. 13 for queues.

Definition 17 (The Effect of ht on T). We define $T\text{-ht}$ as follows:

$$\begin{array}{c}
T\text{-}\emptyset = T \quad \frac{l \notin \text{labels}(T)}{T\text{-}q?l(S_k) \cdot \text{ht} = T\text{-ht}} \quad \frac{\phi \notin T \quad \langle \psi \rangle^\phi \in \text{ht}}{T\text{-ht} = T\text{-ht} \setminus \langle \psi \rangle^\phi} \\
F' = \cup\{A \mid A \in \text{dom}(\mathcal{H}) \wedge F \subset A \subseteq \text{Fset}(\text{ht}, p)\} \quad F':T' \in \mathcal{H} \\
\frac{\mathcal{E}[\text{t}(T)\text{h}(\mathcal{H})^{(\kappa, F)}.T'']\text{-ht} = \mathcal{E}[\text{t}(T')\text{h}(\mathcal{H})^{(\kappa, F')}.T'']\text{-ht}}{\langle \psi \rangle^\phi \in \text{ht}} \\
\frac{\langle \psi \rangle^\phi \in \text{ht}}{\mathcal{E}[\text{t}(\text{end})\text{h}(\mathcal{H})^\phi.T']\text{-ht} = \mathcal{E}[T']\text{-}(\text{ht} \setminus \langle \psi \rangle^\phi)} \\
\frac{k \in I}{\mathcal{E}[q?\{l_i(S_i).T_i\}_{i \in I}\text{-}q?l_k(S_k) \cdot \text{ht} = \mathcal{E}[T_k]\text{-ht}} \quad \frac{\text{ht} \equiv p?l(S) \cdot \text{ht}' \quad \exists k \in I.l = l_k}{\mathcal{E}[q!\{l_i(S_i).T_i\}_{i \in I}\text{-ht} = \mathcal{E}[q!\{l_i(S_i).(T_i\text{-}p?l(S))\}_{i \in I}\text{-ht}']}
\end{array}$$

Other cases are undefined.

Notice that those rules are for *general* session environments (i.e., Δ), not particularly for coherent session environments. Here we use the following example to explain the final rule particularly.

Example 13. Assume in Δ we have $\mathfrak{s}[p] : T \in \Delta$ such that

$$T = p_1!\{l_1(\text{int}).p_2?l(\text{str}).p_1?l_3(\text{int}), l_2(\text{str}).p_2?l(\text{str}).p_1?l_4(\text{str})\}$$

Then, for analyzing the interactions of $\mathfrak{s}[p]$ and $\mathfrak{s}[p_1]$ only, we focus on p_1 's actions by applying Def. 15: $T \downarrow p_1 = p_1!\{l_1(\text{int}).p_1?l_3(\text{int}), l_2(\text{str}).p_1?l_4(\text{str})\}$.

Assume now we have messages $\langle p_2, p, l(\text{str}) \rangle \cdot \langle p_1, p, l_3(\text{int}) \rangle$ heading to $\mathfrak{s}[p]$. By applying Def. 16, we simplify those messages to only $p_1?l_3(\text{int})$ and ignore the other one because it is sent by p_2 . Because an output action $p_1!$ is not able to consume any message, we leave this message to be consumed by the type following $p_1!$, which is either $p_1?l_3(\text{int})$ or $p_1?l_4(\text{str})$. So the result after consuming $\langle p_1, p, l_3(\text{int}) \rangle$, we have $p_1!\{l_1(\text{int}).p_1?l_3(\text{int}), l_2(\text{str}).p_1?l_4(\text{str})\}\text{-}p_1?l_3(\text{int}) = p_1!\{l_1(\text{int}).\text{end}, l_2(\text{str}).\text{end}\}$.

However, very importantly, the existence of $\langle p_1, p, l_3(\text{int}) \rangle$ and T at the same time shall *never* happen in any coherent session environment. Duality, defined in Def. 7 in Sec. 6, makes sure that a coherent session environment can never have this case: If we can have $\langle p_1, p, l_3(\text{int}) \rangle$ before T executes the output $p_1!$, it implies that the local type of $\mathfrak{s}[p_1]$ is $T_1 = p!\{l_3(\text{int}).p?l_1(\text{int}), l_4(\text{str}).p?l_2(\text{str})\}$ so that $\mathfrak{s}[p] : p_1!.T'\text{-}\emptyset \not\triangleleft \mathfrak{s}[p_1] : p!.T_1'\text{-}\emptyset$ from the beginning!

Other rules are trivially defined based on our operational semantics of applications and the system in Fig. 6, Fig. 7, and Fig. 8.

C Implementation Details

Based on the calculus presented previously we developed a domain-specific language (DSL) and corresponding runtime system in Scala (without endpoint type checking), using ZooKeeper as the coordinator. In the following we show that our introduced abstractions match our implementation. In particular we show that Zookeeper together with our runtime system provide the coordinator abstraction and that our failure detector assumption (which in practice can be weakened by program-controlled crash) is reasonable.

ZooKeeper as coordinator. Like the global queue which is an abstraction for pairwise channels (e.g., TCP sockets) and can not model a message that is currently in the network vs a message that is in an endpoint buffer, our coordinator is also an abstraction. Before describing how ZooKeeper and our runtime implement the coordinator abstraction we provide a high level description of important ZooKeeper functionalities we use. In essence ZooKeeper provides a hierarchical name space similar to a classical file system which guarantees sequential consistency and atomicity among other things. Clients that work with ZooKeeper enter into a ZooKeeper session (the runtime system ensures that during **(Link)** every linking process enters a ZooKeeper session). ZooKeeper sessions are important for so-called *ephemeral nodes*, which allow clients to save non-persistent data for as long as their ZooKeeper session is running, in contrast to normal nodes which are persisted. ZooKeeper exchanges heartbeats with all clients and if it does not receive heartbeats from a client for a configuration timespan it disconnects that clients (i.e., ends the ZooKeeper session of that client and removes all its ephemeral nodes).

When participants create a new session via the **(Link)** rule the runtime system creates a unique name space for this session and saves the following information in it (among other): (i) a persistent queue (the notification queue) for all done and failure notification sent to the coordinator, which ensure that all participant see notifications sent to that queue in the same order; (ii) for every participant a node which contains the information whether that participant has terminated (i.e., successfully finished its involvement in the session); (iii) an ephemeral node for every participant that indicates if the participant has an active ZooKeeper session.

Done and failure notification. ZooKeeper by itself does not directly implement the reduction rules **(CollectDone)**, **(IssueDone)** and **(F)**, however ZooKeeper ensures that every participant sees the done and failure notification in the same order. That makes it straightforward for the runtime system to provide the coordinator behavior by applying these rules in order of the notifications in the notification queue. The rules **(CollectDone)**, **(IssueDone)** and **(F)** provide a deterministic outcome ((**(IssueDone)** process done notification based on the order in which they arrive). In cases where more than one of these rules is applicable, the runtime system selects a rule based on the following order: **(CollectDone)** over **(IssueDone)** over **(F)**.

Failure detection and false suspicions. The operational semantics model a perfect asynchronous failure detector, in that upon a process crash the **(Crash)** rule adds a failure notification to the queue which the coordinator eventually processes. In practice false suspicions (i.e., non-perfect failure detection) can happen. The failure detection in our implementation works as follows: if the runtime system of a participant p suspects that participant q has failed it only issues a failure notification if participant q has no active ZooKeeper session (the ephemeral node of q does not exist) and the saved status of q in ZooKeeper is not terminated.

The key point here is that a failure notification for a participant q will only be issued if q lost its ZooKeeper session.

However it is possible that a participant q loses its connection with ZooKeeper without failing, e.g., the network drops too many heartbeat messages. After q loses its ZooKeeper session other participants can issue failure notification for q therefore q realizes that it was suspected and stops, i.e., it performs a controlled crash. As mentioned, this occurs very infrequently in practice but is a feature used in many distributed systems as last resort.

If a participant finishes its actions inside a session it sets its status to terminated before disconnecting from ZooKeeper. This ensures that other participants can not detect it as failed.

Automatically inferring try-handle levels and semi-unique labels. For the simplification of the technique developed we require well-formed global types which require that every try-handle is annotated with a unique level. For the formal development we believe those are reasonable assumptions, however our runtime system provides functionalities to remove this burden from the programmer. Concretely, we implemented a deterministic traversal for global types which assigns unique and deterministic levels to all try-handle in a given global type and we allow writing a global type which only contains labels required for branching, adding others automatically.




Fig. 14. Here two workers implemented in our DSL read data from ZooKeeper, simulate work on it, and write a result back. We inject a failure during the processing of batch number 3, which is clearly visible in the spike. At that point the first worker starts handling the failure of worker two, performing both works.

Evaluation. We evaluated our prototype on a microbenchmark, derived from the big data task shown in Fig. 2. For the evaluation all participants and the different ZooKeeper replicas run on different nodes. In all evaluation ZooKeeper runs in a three-server setting, i.e., we have three different nodes which all have a

ZooKeeper replica running and therefore the ZooKeeper setup can tolerate one process crash while still remaining operational. In this task two workers (w_1, w_2) read data from a distributed file system (*dfs*), simulate working on the data and then writes data back into the file system. The ZooKeeper setup provides reliable data source and sink (*dfs*) used in the benchmarks and the runtime systems are extended so that it can treat ZooKeeper like a regular robust participant. In order to measure the overhead of our system we deployed all nodes in distinct virtual machines on the same server, with the following spec: Intel(R) Xeon(R) CPU E5-2680 v3 2.50GHz and 126 GB of RAM; each node has 4 cores assigned and 8 GB of RAM. In addition ZooKeeper is configured to keep ZooKeeper session for at least 1000 milliseconds alive after the last successfully exchanged heartbeat. Our benchmark differs from the global type in Fig. 2 in that the workers read data, i.e., we have two communication steps instead of one. For evaluation purposes the actions inside the try-block and the handlers are repeated 10 times. Every repetition uses a distinct data batch and each worker reads one partition from this batch. Each batch contains two partitions, the worker w_1 works on partition one and worker w_2 work on partition two. In case of a failure the non-crashed worker processes both partitions.

The microbenchmark measures the cost of handling and recovering from a failure if work is performed. We simulate the processing of a partition of a batch as taking 1000 milliseconds. In this evaluation we manually stop the node of w_2 (the Linux system in which w_2 runs) when w_2 is processing the data from the batch. The graph in Fig. 14 shows a spike at batch three where the crash is injected. This spike shows the combined time of: crash detection, involvement of coordinator and worker on doing the processing of the entire batch three (i.e., also doing the work of w_2). For batches 4 to 10 the processing time doubles since w_1 also does the work of w_2 . Apart from the time where failure handling is triggered one can see that the runtime is entirely dominated by the work, i.e., the overhead in the non-failure case is negligible.

D Proofs of Properties

This part of appendix provides the proofs and auxiliary definitions and lemmas for the proposition and theorems in the paper.

D.1 Proof for Section 4 (Process Calculus)

Proposition 1. Given $s : h$ with $h = h' \cdot \langle \psi, p \rangle^\phi \cdot h''$ and $Fset(h, p) \neq \emptyset$, the rule **(TryHdl)** is not applicable for the try-handle ϕ at the process playing role p .

Proof. W.o.l.g. assume the coordinator is $\Psi = G : (F_q, d)$. Since the failure and done notifications sent by Ψ to the same participant have an order (not permutable by Def. 13), we have the following two cases:

1. If $\exists \langle p, \text{crash } F'' \rangle \in h'$, then we prove the statement by contradiction. Assume we have $F \in hdl(G, \phi)$ and $F = \cup \{A \mid A \in dom(\mathbf{H}) \wedge F_s' \subset A \subseteq Fset(h, p)\}$. Since $Fset(h, p) \subseteq F_q$, immediately rule **(IssueDone)** is violated.
2. If $\exists \langle p, \text{crash } F'' \rangle \in h''$, then we prove the statement by Def. 2 (*Fset*). By Def. 2, $\langle p, \text{crash } F'' \rangle \notin Fset(h, p)$, therefore $F'' \notin \cup \{A \mid A \in dom(\mathbf{H}) \wedge F_s' \subset A \subseteq Fset(h, p)\}$.

D.2 Proofs for Section 7 (Properties)

To prove the properties of subject congruence, subject reduction, and progress, we need to first prove one theorem, Theorem 4 (Preservation of Coherence), which is the basis for proving those properties. This section goes as follows: we prove Theorem 4 (Preservation of Coherence), then prove Theorem 1 (Subject Congruence) and Theorem 2 (Subject Reduction), and finally prove Theorem 3 (Progress). For each proof, several auxiliary formal definitions are introduced, and useful lemmas are stated and proved.

Preservation of Coherence. Before stating and proving preservation of coherence, we first define the reduction relation over session environments

$$\Psi \vdash \Delta \rightarrow_T \Psi' \vdash \Delta'$$

which is defined based on the operational semantics of applications and systems defined in Sec. 4.

Definition 18 (Reduction relation over session environments).

$$\begin{array}{c}
\frac{}{s[p] : \mathcal{E}[q!\{l_i(S_i).T_i\}_{i \in I}], s : \mathbf{h} \rightarrow_T s[p] : \mathcal{E}[T_k], s : \mathbf{h} \cdot \langle p, q, l_k(S_k) \rangle \quad k \in I} \quad \llbracket \text{Snd} \rrbracket \\
\frac{}{s[p] : \mathcal{E}[q?\{l_i(S_i).T_i\}_{i \in I}], s : \langle q, p, l_k(S_k) \rangle \cdot \mathbf{h} \rightarrow_T s[p] : \mathcal{E}[T_k], s : \mathbf{h} \quad k \in I} \quad \llbracket \text{Rcv} \rrbracket \\
\frac{}{\Delta, \{s : \mathbf{h}\} \rightarrow \Delta \setminus \{s[p] : T\}, \{s : \text{remove}(\mathbf{h}, p) \cdot \langle \psi, \text{crash } p \rangle\} \quad p \text{ non-robust}} \quad \llbracket \text{Crash} \rrbracket \\
\frac{F' = \cup\{A \mid A \in \text{dom}(\mathbf{H}) \wedge F \subset A \subseteq \text{Fset}(\mathbf{h}, p)\} \quad F' : T' \in \mathcal{H}}{s[p] : \mathcal{E}[\text{t}(T)\mathbf{h}(\mathcal{H})^{(\kappa, F)}.T''], s : \mathbf{h} \rightarrow s[p] : \mathcal{E}[\text{t}(T')\mathbf{h}(\mathcal{H})^{(\kappa, F')}.T''], s : \mathbf{h}} \quad \llbracket \text{TryHdl} \rrbracket \\
\frac{}{s[p] : \mathcal{E}[\text{t}(\text{end})\mathbf{h}(\mathcal{H})^\phi.T], s : \mathbf{h} \rightarrow s[p] : \mathcal{E}[\text{t}(\text{end})\mathbf{h}(\mathcal{H})^\phi.T], s : \mathbf{h} \cdot \langle p, \psi \rangle^\phi} \quad \llbracket \text{SndDone} \rrbracket \\
\frac{\langle \psi, p \rangle^\phi \in \mathbf{h}}{s[p] : \mathcal{E}[\text{t}(\text{end})\mathbf{h}(\mathcal{H})^\phi.T], s : \mathbf{h} \rightarrow s[p] : \mathcal{E}[T], s : \mathbf{h} \setminus \{\langle \psi, p \rangle^\phi\}} \quad \llbracket \text{RcvDone} \rrbracket \\
\frac{}{s[p] : \mathcal{E}[T], s : \langle q, p, l(S) \rangle \cdot \mathbf{h} \rightarrow s[p] : \mathcal{E}[T], s : \mathbf{h} \quad l \notin \text{labels}(\mathcal{E}[T])} \quad \llbracket \text{Cln} \rrbracket \\
\frac{\langle \psi, p \rangle^\phi \in \mathbf{h} \quad \phi \notin \mathcal{E}[T]}{s[p] : \mathcal{E}[T], s : \mathbf{h} \rightarrow s[p] : \mathcal{E}[T] \mid s : \mathbf{h} \setminus \{\langle \psi, p \rangle^\phi\}} \quad \llbracket \text{ClnDone} \rrbracket \\
\frac{\tilde{p} = \text{roles}(G) \setminus (F \cup \{p\}) \quad \mathbf{h}' = \mathbf{h} \cdot \langle \tilde{p}, \text{crash } \{p\} \rangle}{G : (F, d) \vdash N \mid s : \langle \psi, \text{crash } \{p\} \rangle \cdot \mathbf{h} \rightarrow_T G : (F \cup \{p\}, d) \vdash N \mid s : \mathbf{h}'} \quad \llbracket \text{F} \rrbracket \\
\frac{d' = d \cdot \langle p, \psi \rangle^\phi}{G : (F, d) \vdash s : \langle p, \psi \rangle^\phi \cdot \mathbf{h} \rightarrow_T G : (F, d') \vdash s : \mathbf{h}} \quad \llbracket \text{CollectDone} \rrbracket \\
\frac{\text{roles}(d, \phi) \supseteq \text{roles}(G, \phi) \setminus F \quad \forall F' \in \text{hdl}(G, \phi). (F' \not\subseteq F) \quad d' = \text{remove}(d, \phi)}{G : (F, d) \vdash s : \mathbf{h} \rightarrow_T G : (F, d') \vdash s : \mathbf{h} \cdot \langle \psi, \text{roles}(G, \phi) \setminus F \rangle^\phi} \quad \llbracket \text{IssueDone} \rrbracket \\
\frac{\Delta \rightarrow_T \Delta'}{\Delta, \Delta_0 \rightarrow_T \Delta', \Delta_0} \quad \llbracket \text{Str} \rrbracket
\end{array}$$

When Ψ is not changed during reduction, i.e., $\Psi \vdash \Delta \rightarrow_T \Psi \vdash \Delta'$, we simply write the reduction relation as $\Delta \rightarrow_T \Delta'$. If $\Psi \vdash \Delta \rightarrow_T \Psi' \vdash \Delta'$ and $\Gamma = \Gamma_0, \Psi$ and $\Gamma' = \Gamma_0, \Psi'$, then we write $\Gamma \vdash \Delta \rightarrow_T \Gamma' \vdash \Delta'$. Remember we have defined the context of local types in App. B.3 $\mathcal{E} ::= [] \mid \text{t}(\mathcal{E})\mathbf{h}(\mathcal{H})^\phi.T \mid \mu t.\mathcal{E}$.

Assume G is well-formed and there is a session s such that $\Delta = \{s[p_1] : G \mid p_1, \dots, s[p_n] : G \mid p_n\} \cup \{s : \emptyset\}$ and $\text{roles}(G) = \{p_1, \dots, p_n\}$. The immediate question for Def. 8 (Coherence) is whether this definition ensures the interactions among endpoints do progress with the guidance of G or not.

It is not trivial because, for the endpoints, they can take actions concurrently, failures can occur randomly, and message types/notifications are transmitted in an asynchronous manner. Consider a simple example:

$$G = p_1 \rightarrow p_2 \ l_1(\text{int}).p_3 \rightarrow p_2 \ l_2(\text{str})$$

for a s that obeys to G , we will have

$$\Delta = \{s[p_1] : p_2!l_1(\text{int}).\text{end}, s[p_2] : p_1?l_1(\text{int}).p_3?l_2(\text{str}).\text{end}, s[p_3] : p_2!l_2(\text{str}).\text{end}, s : \emptyset\}$$

and $\Delta \rightarrow_T \Delta'$ such that

$$\Delta' = \{s[p_1] : p_2!l_1(\text{int}).\text{end}, s[p_2] : p_1?l_1(\text{int}).p_3?l_2(\text{str}).\text{end}, s[p_3] : \text{end}, s : \langle p_3, p_2, l_2(\text{str}) \rangle\}$$

At this step, $\langle p_3, p_2, l_2(\text{str}) \rangle$ cannot trigger its receiver $s[p_2]$ because the corresponding receiving action $p_3?l_2(\text{str})$ is guarded by $p_1?l_1(\text{int})$.

Note that, Δ' is still able to reduce because $s[p_1]$ can take action. When $s[p_1]$ takes action and outputs $\langle p_1, p_2, l_2(\text{int}) \rangle$, its receiver $s[p_2]$ can absorb $\langle p_3, p_2, l_2(\text{str}) \rangle$ immediately.

Then in the next reduction, $s[p_2]$ will absorb $\langle p_3, p_2, l_2(\text{str}) \rangle$.

We observe that, not every sending action fired by a well-typed behavior can trigger its expected receiver immediately; but this will happen eventually.

The following simple lemmas and definitions will be used in the proof for Theorem 4 (Preservation of Coherence).

Lemma 1. If $T \downarrow p = \mathcal{E}[p \ w \ \{l_i(S_i).T_i\}_{i \in I}]$ and $w \in \{!, ?\}$ for some \mathcal{E} and $T_i, i \in I$, then for any $p' \neq p$ we have for any $j, k \in I, j \neq k, \mathcal{E}[T_j] \downarrow p' = \mathcal{E}[T_k] \downarrow p'$.

Proof. Immediately by Def. 4 and Def. 7 (Duality).

Lemma 2. For any $G, p \in \text{roles}(G)$ and $q \notin \text{roles}(G)$ imply $\text{end} = G \downarrow q \downarrow p \bowtie G \downarrow p \downarrow q = \text{end}$.

Proof. Immediately by Def. 4 (Projection) and Def. 15.

The next lemma states that, given a session environment, if each of its endpoint's type is projected from a well-formed global type, then this session environment is coherent.

Lemma 3. Given G is well-formed and $\Delta_s = \{s[p_1] : T_1, \dots, s[p_n] : T_n, s : \emptyset\}$ and $\{p_1, \dots, p_n\} = \text{roles}(G)$ and $T_i = G \downarrow p_i$ for $i \in \{1..n\}$. Then $\Gamma, G : (\emptyset, \emptyset) \vdash \Delta_s$ is coherent.

Proof. By the structure of a well-formed G , the proof is immediate by Def. 4 (Projection), Def. 7 (Duality), and Def. 8 (Coherence).

And, after a failure occurs, if a coordinator has not yet issued failure notifications to endpoints in a coherent Δ for handling this failure, the types of all non-failed endpoints are still dual to each other in Δ .

Lemma 4. Assuming $\Delta = \Delta_0, \{s : \mathbf{h}\}$ is coherent and $\Delta \rightarrow_T \Delta_0 \setminus \{s[p] : T\}, \{s : \text{remove}(\mathbf{h}, p) \cdot \langle \psi, \text{crash } p \rangle\} = \Delta'$. Then $\forall s[p] : T \in \Delta$, if $s[q] : T' \in \Delta$ then $s[p] : T \downarrow q^-(\mathbf{h})_{q \rightarrow p} \bowtie s[q] : T' \downarrow p^-(\mathbf{h})_{p \rightarrow q}$.

Proof. Assume the non-failed endpoints are $\{s[p_i] : T_i\}_{i \in I}$. Since Δ is coherent, $\{s[p_i] : T_i\}_{i \in I}$ in Δ are dual to each other after considering the effect of \mathbf{h} on each of them. By the rules defined in Def. 17, we know $\langle \psi, \text{crash } p \rangle$ will not affect any non-failed endpoints. So the endpoints of $\{s[p_i] : T_i\}_{i \in I}$ in Δ' are still dual to each other after considering the effect of $\text{remove}(\mathbf{h}, p) \cdot \langle \psi, \text{crash } p \rangle$ on each of them.

To reason about asynchrony in Δ , we define message types permutation, which is very similar to Def. 13.

Definition 19 (Permutable Message Types). We define $\mathfrak{m}_i \cdot \mathfrak{m}_j \rightsquigarrow \mathfrak{m}_j \cdot \mathfrak{m}_i$, $i \neq j$, saying $\mathfrak{m}_i \cdot \mathfrak{m}_j$ can be permuted to $\mathfrak{m}_j \cdot \mathfrak{m}_i$, if *none* of the following conditions holds:

- $\mathfrak{m}_i = \langle p, q, l(S) \rangle$ and $\mathfrak{m}_j = \langle p, q, l'(S') \rangle$ for some l, l', S, S' .
- $\mathfrak{m}_i = \langle \psi, q \rangle^\phi$ and $\mathfrak{m}_j = \llbracket q, \text{crash } F \rrbracket$ for some ϕ, F .

For convenience we say that participants currently performing actions in try-handle ϕ are *partners*. We also give the following function *HdlLogic* to extract a particular try-handle for a handler. This function is used for proving Case $\llbracket \mathbb{F} \rrbracket$ in Theorem 4.

When $\mathfrak{t}(G_0 \downarrow p) \mathfrak{h}(F_1 G_1 \downarrow p, \dots, F_n G_n \downarrow p)^{(\kappa, \emptyset)}$ is generated by Def. 4 (Projection), for other handler signatures in the try-handle of κ can be automatically generated by

Definition 20. Given $\mathfrak{t}(T) \mathfrak{h}(\mathcal{H})^{(\kappa, \emptyset)}$ and $\mathcal{H} = F_1 T_1, \dots, F_n T_n$. We define

$$\text{HdlLogic}(\mathfrak{t}(T) \mathfrak{h}(\mathcal{H})^{(\kappa, \emptyset)}, F) = \mathfrak{t}(T_i) \mathfrak{h}(\mathcal{H})^{(\kappa, F)} \text{ if } F = F_i, i \in \{1..n\}$$

Also for convenience, we define

$$\text{notify}(p) ::= \llbracket p, \text{crash } F \rrbracket \mid \llbracket \psi, \text{crash } F \rrbracket \mid \langle \psi, p \rangle^\phi \mid \langle p, \psi \rangle^\phi$$

The following proof states that a coherent session environment will reduce to another coherent session environment.

Theorem 4 (Preservation of Coherence). $\Gamma \vdash \Delta$ coherent and $\Gamma = G : (F, d), \Gamma'$ and $G : (F, d) \vdash \Delta \rightarrow_T G : (F', d') \vdash \Delta'$ imply that $\Gamma', G : (F', d') \vdash \Delta'$ is coherent.

Proof. We prove the statement by mechanically proving each cases.

1. Case $\mathfrak{s} : \emptyset$. We have the following subcases.

(I) Case $\llbracket \text{snd} \rrbracket$, there exists $\mathfrak{s}[p] : T \in \Delta, T = \mathcal{E}[q! \{l_i(S_i).T_i\}_{i \in I}]$.

W.o.l.g. assume

$$\Delta = \Delta_0, \mathfrak{s}[p] : T, \mathfrak{s} : \emptyset$$

such that

$$G : (F_q, d) \vdash \Delta_0, \mathfrak{s}[p] : T, \mathfrak{s} : \emptyset \rightarrow_T G : (F_q, d) \vdash \Delta_0, \mathfrak{s}[p] : \mathcal{E}[T_k], \mathfrak{s} : \langle p, q, l_k(S_k) \rangle$$

for some $k \in I$. Let $\Delta' = \Delta_0, \mathfrak{s}[p] : \mathcal{E}[T_k], \mathfrak{s} : \langle p, q, l_k(S_k) \rangle$.

In Δ , by Def. 8.(2), for any $\mathfrak{s}[q'] : T' \in \Delta$, we have

$$\mathfrak{s}[p] : T \downarrow q' \bowtie \mathfrak{s}[q'] : T' \downarrow p$$

In Δ , if $q' = q$, then let $T \downarrow q = \mathcal{E}''[q!\{l_i(S_i).T_i''\}_{i \in I}]$ for some \mathcal{E}'' and $T_i'', i \in I$. By Def. 7 $T' \downarrow p = \mathcal{E}'[p?\{l_i(S_i).T_i'\}_{i \in I}]$ for some \mathcal{E}' and $T_i', i \in I$ and

$$\forall i \in I. \mathfrak{s}[p] : \mathcal{E}''[T_i''] \bowtie \mathfrak{s}[q] : \mathcal{E}'[T_i'].$$

In Δ' , by Def. 17 (The Effect of \mathfrak{ht}) and Def. 7 (Duality), we have $\mathfrak{s}[p] : \mathcal{E}[T_k]$ and $\mathfrak{s}[q] : T'$ such that

$$\begin{aligned} \mathfrak{s}[p] : \mathcal{E}[T_k] \downarrow q - (\langle p, q, l_k(S_k) \rangle)_{q \rightarrow p} &= \mathfrak{s}[p] : \mathcal{E}''[T_k''] \\ &\bowtie \mathfrak{s}[q] : \mathcal{E}'[T_k'] \\ &= \mathfrak{s}[q] : T' \downarrow p - (\langle p, q, l_k(S_k) \rangle)_{p \rightarrow q} \end{aligned}$$

In Δ , for any $q' \neq q$, $q' \in \text{roles}(G)$ and $\mathfrak{s}[q'] : T'' \in \Delta$, by Lemma 1, the fact that $T \downarrow q' = \mathcal{E}[q!\{l_i(S_i).T_i\}_{i \in I}] \downarrow q'$ is defined implies that for any $j, k \in I$, $j \neq k$, we have $\mathcal{E}[T_j] \downarrow q' = \mathcal{E}[T_k] \downarrow q'$ and

$$\mathfrak{s}[p] : T \downarrow q' = \mathfrak{s}[p] : \mathcal{E}[T_j] \downarrow q' \bowtie \mathfrak{s}[q'] : T'' \downarrow p$$

In Δ' , by Def. 17 (The Effect of \mathfrak{ht}) and Def. 7 (Duality), we have $\mathfrak{s}[p] : \mathcal{E}[T_k]$ and $\mathfrak{s}[q'] : T''$ such that

$$\begin{aligned} \mathfrak{s}[p] : \mathcal{E}[T_k] \downarrow q' - (\langle p, q, l_k(S_k) \rangle)_{q' \rightarrow p} &= \mathfrak{s}[p] : \mathcal{E}[T_k] \downarrow q' \\ &\bowtie \mathfrak{s}[q'] : T'' \downarrow p - (\langle p, q, l_k(S_k) \rangle)_{p \rightarrow q'} \\ &= \mathfrak{s}[q'] : T'' \downarrow p \end{aligned}$$

Since the types of other endpoints are not changed, $\Gamma', G : (F_q, d) \vdash \Delta'$ is coherent.

- (II) Case $\llbracket \text{crash} \rrbracket$, some endpoint in Δ crashes. W.o.l.g. assume $\Delta = \Delta_0, \mathfrak{s}[p] : T, \mathfrak{s} : \emptyset$ such that

$$G : (F_q, d) \vdash \Delta_0, \mathfrak{s}[p] : T, \mathfrak{s} : \emptyset \rightarrow_T G : (F_q, d) \vdash \Delta_0, \mathfrak{s} : \langle \psi, \text{crash } F \rangle$$

Let $\Delta' = \Delta_0, \mathfrak{s} : \langle \psi, \text{crash } F \rangle$.

By Lemma 4 we have $\Gamma \vdash \Delta'$ is coherent

- (III) Case $\llbracket \text{sndDone} \rrbracket$, some endpoint in Δ finishes its default action. W.o.l.g. assume $\Delta = \Delta_0, \mathfrak{s}[p] : T, \mathfrak{s} : \emptyset$ and $T = \mathcal{E}[\text{t}(\text{end})\mathfrak{h}(\mathcal{H})^\phi.T']$ such that

$$\begin{aligned} G : (F_q, d) \vdash \Delta_0, \mathfrak{s}[p] : T, \mathfrak{s} : \emptyset \\ \rightarrow_T G : (F_q, d) \vdash \Delta_0, \mathfrak{s}[p] : \mathcal{E}[\text{t}(\text{end})\mathfrak{h}(\mathcal{H})^\phi.T'], \mathfrak{s} : \langle p, \psi \rangle^\phi \end{aligned}$$

Let $\Delta' = \Delta_0, \mathfrak{s}[p] : \mathcal{E}[\text{t}(\text{end})\mathfrak{h}(\mathcal{H})^\phi.T'], \mathfrak{s} : \langle p, \psi \rangle^\phi$.

Since this kind of message type will not affect any endpoints in Δ_0 and by Def. 8.(2) for any $\mathfrak{s}[q] : T' \in \Delta$ we have $\mathfrak{s}[p] : T \downarrow q \bowtie \mathfrak{s}[q] : T' \downarrow p$ so $\Gamma \vdash \Delta'$ is coherent.

2. Case $\mathbf{s} : \mathbf{h} \neq \emptyset$ and $\Gamma \vdash \Delta$ is coherent.⁷

(I) Case $\llbracket \text{snd} \rrbracket$, there exists $\mathbf{s}[p] : T \in \Delta, T = \mathcal{E}[q!\{l_i(S_i).T_i\}_{i \in I}]$.

W.o.l.g. assume

$$\Delta = \Delta_0, \mathbf{s}[p] : T, \mathbf{s} : \mathbf{h}$$

such that

$$\begin{aligned} G : (F_q, d) \vdash \Delta_0, \mathbf{s}[p] : T, \mathbf{s} : \mathbf{h} \\ \rightarrow_T G : (F_q, d) \vdash \Delta_0, \mathbf{s}[p] : \mathcal{E}[T_k], \mathbf{s} : \mathbf{h} \cdot \langle p, q, l_k(S_k) \rangle \end{aligned}$$

for some $k \in I$. Let $\Delta' = \Delta_0, \mathbf{s}[p] : \mathcal{E}[T_k], \mathbf{s} : \mathbf{h} \cdot \langle p, q, l_k(S_k) \rangle$.

Then we have the following cases:

(A) Case $\forall p \in \text{roles}(G)$. $\text{notify}(p) \notin \mathbf{h}$ which implies $\langle \psi, \text{crash } F \rangle \notin \mathbf{h}$ (i.e. \mathbf{h} contains only normal message types).

In Δ , by Def. 17 (The Effect of ht) and Def. 8 (2). (a), for any $\mathbf{s}[q'] : T' \in \Delta$ we have

$$\begin{aligned} \mathbf{s}[p] : T \downarrow q' - (\mathbf{h})_{q' \rightarrow p} = \mathbf{s}[p] : \mathcal{E}[q!\{l_i(S_i).T_i\}_{i \in I}] \downarrow q' - (\mathbf{h})_{q' \rightarrow p} \\ \bowtie \mathbf{s}[q'] : T' \downarrow p - (\mathbf{h})_{p \rightarrow q'} \end{aligned}$$

In Δ , if $q' = q$, we have

$$T \downarrow q - (\mathbf{h})_{q \rightarrow p} = \mathcal{E}[q!\{l_i(S_i).T_i\}_{i \in I}] \downarrow q - (\mathbf{h})_{q \rightarrow p} = \mathcal{E}''[q!\{l_i(S_i).T_i'''\}_{i \in I}]$$

for some \mathcal{E}'' and $T_i''', i \in I$ and by Def. 7 (Duality), we have $T' \downarrow p - (\mathbf{h})_{p \rightarrow q} = \mathcal{E}'[p?\{l_i(S_i).T_i'\}_{i \in I}]$ for some \mathcal{E}' and

$$\forall i \in I. \mathbf{s}[p] : \mathcal{E}''[T_i'''] \bowtie \mathbf{s}[q] : \mathcal{E}'[T_i'].$$

By Def. 17 (The Effect of ht) and Def. 7 (Duality)

$$\begin{aligned} \mathbf{s}[p] : \mathcal{E}[T_k] \downarrow q - (\mathbf{h} \cdot \langle p, q, l_k(S_k) \rangle)_{q \rightarrow p} \\ = \mathbf{s}[p] : \mathcal{E}[T_k] \downarrow q \\ = \mathbf{s}[p] : \mathcal{E}''[T_k'''] \\ \bowtie \mathbf{s}[q] : \mathcal{E}'[T_k'] \\ = \mathbf{s}[q] : T' \downarrow p - (\mathbf{h} \cdot \langle p, q, l_k(S_k) \rangle)_{p \rightarrow q} \end{aligned}$$

In Δ , for any $q' \neq q$, let $\mathbf{s}[q'] : T'' \in \Delta$. By Lemma 1, $\forall j, k \in I, j \neq k$ we have $\mathcal{E}[T_j] \downarrow q' - (\mathbf{h})_{q' \rightarrow p} = \mathcal{E}[T_k] \downarrow q' - (\mathbf{h})_{q' \rightarrow p}$ and

$$\begin{aligned} \mathbf{s}[p] : T \downarrow q' - (\mathbf{h})_{q' \rightarrow p} = \mathbf{s}[p] : \mathcal{E}[q!\{l_i(S_i).T_i\}_{i \in I}] \downarrow q' - (\mathbf{h})_{q' \rightarrow p} \\ = \mathbf{s}[p] : \mathcal{E}[T_j] \downarrow q' - (\mathbf{h})_{q' \rightarrow p} \\ \bowtie \mathbf{s}[q'] : T'' \downarrow p - (\mathbf{h})_{p \rightarrow q'} \end{aligned}$$

⁷ For convenience we sometimes write $\mathbf{s}[p] : T \downarrow p - (\mathbf{h})_{p \rightarrow q} = \mathbf{s}[p] : T \downarrow p$ for $(\mathbf{h})_{p \rightarrow q} \neq \emptyset$ if $(\mathbf{h})_{p \rightarrow q}$ contains only failure notification and they have no effect on $\mathbf{s}[p] : T \downarrow p$

Thus in Δ' we have

$$\begin{aligned}
& \mathfrak{s}[p] : \mathcal{E}[T_k] \downarrow q' - (\mathbf{h} \cdot \langle p, q, l_k(S_k) \rangle)_{q' \rightarrow p} \\
& = \mathfrak{s}[p] : \mathcal{E}[T_k] \downarrow q' - (\mathbf{h})_{q' \rightarrow p} \\
& \bowtie \mathfrak{s}[q'] : T'' \downarrow p - (\mathbf{h})_{p \rightarrow q'} \\
& = \mathfrak{s}[q'] : T'' \downarrow p - (\mathbf{h} \cdot \langle p, q, l_k(S_k) \rangle)_{p \rightarrow q'}
\end{aligned}$$

Since the types of other endpoints are not changed, $\Gamma \vdash \Delta'$ is coherent.

(B) Case $\exists \langle p_0, \text{crash } F' \rangle \in \mathbf{h}$ for some F' and p_0 .

Since $\langle p_0, \text{crash } F' \rangle$ is issued by $\llbracket \mathbb{F} \rrbracket$, we have $\forall \mathfrak{s}[p_i] \in \text{dom}(\Delta_s)$, $\langle p_i, \text{crash } F' \rangle \in \mathbf{h}$ (by $\llbracket \text{crash} \rrbracket$, only the alive endpoints).

W.o.l.g assume $\Delta = \Delta_0, \mathfrak{s}[p] : T, \mathfrak{s}[q] : T', \mathfrak{s} : \mathbf{h}$ and assume $T = \mathcal{E}[q! \{l_i(S_i).T_i''\}_{i \in I}]$ and

$$\begin{aligned}
G & : (F_q, d) \vdash \Delta_0, \mathfrak{s}[p] : T, \mathfrak{s}[q] : T', \mathfrak{s} : \mathbf{h} \rightarrow T \\
G & : (F_q, d) \vdash \Delta_0, \mathfrak{s}[p] : \mathcal{E}[T_k], \mathfrak{s}[q] : T', \mathfrak{s} : \mathbf{h} \cdot \langle p, q, l_k(S_k) \rangle
\end{aligned}$$

In Δ we observe

* If $T = \mathcal{E}'''[\mathfrak{t}(\mathcal{E}''[q! \{l_i(S_i).T_i''\}_{i \in I}])\mathbf{h}(F : T_F, \mathcal{H})^{(\kappa, F_s)}]$, where T_F is the handler body for F , and $F = \cup \{A \mid A \in \text{dom}(\mathcal{H}) \wedge F_s \subset A \subseteq F\text{set}(\mathbf{h}, p)\}$ (notice that $\langle \psi, p \rangle^{(\kappa, F_s)} \notin \mathbf{h}$ because endpoint $\mathfrak{s}[p]$ who is still taking actions in ϕ), then by Def. 8 and Def. 7 (Duality) and Def. 17 (The Effect of ht), we have

$$\begin{aligned}
& \mathfrak{s}[p] : T \downarrow q - (\mathbf{h})_{q \rightarrow p} \\
& = \mathfrak{s}[p] : \mathcal{E}'''[\mathfrak{t}(\mathcal{E}''[q! \{l_i(S_i).T_i''\}_{i \in I}])\mathbf{h}(F : T_F, \mathcal{H})^{(\kappa, F_s)}] \downarrow q - (\mathbf{h})_{q \rightarrow p} \\
& = \mathfrak{s}[p] : \mathcal{E}'''[\mathfrak{t}(T_F)\mathbf{h}(F : T_F, \mathcal{H})^{(\kappa, F)}] \downarrow q - (\mathbf{h})_{q \rightarrow p} \\
& \bowtie \mathfrak{s}[q] : T' \downarrow p - (\mathbf{h})_{p \rightarrow q} \\
& = \mathfrak{s}[q] : \mathcal{E}''''[\mathfrak{t}(T'_F)\mathbf{h}(F : T'_F, \mathcal{H}')^{(\kappa, F)}] \downarrow p - (\mathbf{h})_{p \rightarrow q}
\end{aligned}$$

where $T_F \bowtie T'_F$ by Def. 7.

Now $l_k \notin T'$ due to the occurrence of failure

In Δ' we observe

$$\begin{aligned}
& \mathfrak{s}[p] : \mathcal{E}[T_k] \downarrow q - (\mathbf{h} \cdot \langle p, q, l_k(S_k) \rangle)_{q \rightarrow p} \\
& = \mathfrak{s}[p] : \mathcal{E}''''[\mathfrak{t}(T_F)\mathbf{h}(F : T_F, \mathcal{H})^{(\kappa, F)}] \downarrow q - (\mathbf{h})_{q \rightarrow p} \\
& \bowtie \mathfrak{s}[q] : \mathcal{E}''''[\mathfrak{t}(T'_F)\mathbf{h}(F : T'_F, \mathcal{H}')^{(\kappa, F)}] \downarrow p - (\mathbf{h})_{p \rightarrow q} \\
& = \mathfrak{s}[q] : T' \downarrow p - (\mathbf{h})_{p \rightarrow q} \\
& = \mathfrak{s}[q] : T' \downarrow p - (\mathbf{h} \cdot \langle p, q, l_k(S_k) \rangle)_{p \rightarrow q}
\end{aligned}$$

because $\langle p, q, l_k(S_k) \rangle$ is removed by Def. 11 (Well-formedness).(3) and rule $\llbracket \text{c1n} \rrbracket$ since $l_k \notin \text{labels}(T')$ (i.e., $\langle p, q, l_k(S_k) \rangle$ is a message sent out after failures which can be handled by $\mathfrak{s}[p]$ and $\mathfrak{s}[q]$)

occur).

- * A case like $T = \mathcal{E}'''[\mathfrak{t}(\mathcal{E}''[q!\{l_i(S_i).T_i''\}_{i \in I}])\mathfrak{h}(F : T_F)^{(\kappa, F_s)}]$ and $\langle \psi, p \rangle^{(\kappa, F_s)} \in \mathfrak{h}$ is not possible, immediately by rules $\llbracket \text{SndDone} \rrbracket$ and $\llbracket \text{IssueDone} \rrbracket$.
- * Otherwise (i.e., T cannot handle failures in \mathfrak{h}), it goes to Case (2).(I).(A) or Case (2).(I).(C): Case (2).(I).(A) is for no any done notifications sent from the coordinator in \mathfrak{h} , while Case (2).(I).(C) is for some done notifications sent from the coordinator in \mathfrak{h} .

For other endpoints the proof is similar as above.

Thus $G : (F_q, d) \vdash \Delta'$ is coherent.

(C) Case $\exists \langle \psi, p' \rangle^\phi \in \mathfrak{h}$ for some p' and $\phi = (\kappa, F_s)$.

Let the try-handle of κ in G be $\mathfrak{t}(G_0)\mathfrak{h}(H)^\kappa$. If $\langle \psi, p_i \rangle^\phi \in \mathfrak{h}$, then $p_i \in \text{roles}(\mathfrak{t}(G_0)\mathfrak{h}(H)^\kappa)$.

W.o.l.g assume $\Delta = \Delta_0, \mathfrak{s}[p] : T, \mathfrak{s}[q] : T', \mathfrak{s} : \mathfrak{h}$ and assume $T = \mathcal{E}[q!\{l_i(S_i).T_i''\}_{i \in I}]$

$$\begin{aligned} G : (F_q, d) \vdash \Delta_0, \mathfrak{s}[p] : T, \mathfrak{s}[q] : T', \mathfrak{s} : \mathfrak{h} \rightarrow_T \\ G : (F_q, d) \vdash \Delta_0, \mathfrak{s}[p] : \mathcal{E}[T_k''], \mathfrak{s}[q] : T', \mathfrak{s} : \mathfrak{h} \cdot \langle p, q, l_k(S_k) \rangle \end{aligned}$$

Let $\text{Labels}(T, \phi)$ as a set of labels appearing in the try-handle of $\phi = (\kappa, F_s)$ in T . If $l_k \in \text{Labels}(T, \phi)$, by $\llbracket \text{IssueDone} \rrbracket$ it leads to contradiction because the actions in the try-handle of ϕ is not yet finished.

So we have $l_k \notin \text{Labels}(T, \phi)$, which implies that $l_k \notin \mathfrak{t}(G_0)\mathfrak{h}(H)^\kappa$ by Def. 11(3)(Well-formedness) and $\llbracket \text{IssueDone} \rrbracket$. The rest of the argumentation is similar to case (A) ($\mathfrak{s}[q] : T' \downarrow p - (\mathfrak{h})_{p \rightarrow q} = T''$ and $\# \mathfrak{t}(T''')\mathfrak{h}(\mathcal{H})^{(\kappa, F)} \in T''$).

Thus $G \vdash \Delta'$ is coherent.

(D) Case $\exists \langle \psi, \text{crash } F \rangle \in \mathfrak{h}$ for some F .

W.o.l.g assume $\Delta = \Delta_0, \mathfrak{s}[p] : T, \mathfrak{s} : \mathfrak{h}$.

By assumption, we have

$$\begin{aligned} G : (F_q, d) \vdash \Delta_0, \mathfrak{s}[p] : T, \mathfrak{s} : \mathfrak{h} \rightarrow_T \\ G : (F_q, d) \vdash \Delta_0, \mathfrak{s}[p] : T', \mathfrak{s} : \mathfrak{h} \cdot \langle p, q, l_k(S_k) \rangle \end{aligned}$$

and $\Delta' = \Delta_0, \mathfrak{s}[p] : T, \mathfrak{s} : \mathfrak{h} \cdot \langle p, q, l_k(S_k) \rangle$.

Note that, in Δ , if the reduction comes from $\llbracket \text{snd} \rrbracket$ (by the assumption of this case), then $p \notin F$ (The case that coordinator $G : (F_q, d)$ is affected is Case 2. (V), where $\llbracket \text{F} \rrbracket$ is applied).

- * If $F = \{q\}$: In Δ by Def. 8 we have
 - $\forall s[q'] : T' \in \Delta, s[p] : T \downarrow q' - (\mathbf{h})_{q' \rightarrow p} \bowtie s[q'] : T' \downarrow p - (\mathbf{h})_{p \rightarrow q'}$
 - and $s[p] \notin \Delta$.
- * If $F \neq \{q\}$: Analogous to case (A-C).

Thus $G : (F_q, d) \vdash \Delta'$ is coherent.

(E) Case $\exists \langle p', \psi \rangle^\phi \in \mathbf{h}$ for some p' and ϕ .

By Def. 17, this kind of message type will not trigger any endpoints, the proof goes back to Cases 2.(I).(A,B,C,D).

- (II) Case $\llbracket \text{crash} \rrbracket$, there exists some $s[p] : T \in \Delta$ that crashes, and failure notification $\langle \psi, \text{crash} \{p\} \rangle$ is sent to the coordinator. The proof is similar to Case 1.(II).
- (III) Case $\llbracket \text{sndDone} \rrbracket$, there exists $s[p] : T \in \Delta, T = \mathcal{E}[\mathbf{t}(\text{end})\mathbf{h}(\mathcal{H})^\phi.T']$ and the reduction sends done notification $\langle p, \psi \rangle^\phi$ to the coordinator. The proof is similar to Case 1.(III).

Note that, this notification at this moment, which is abstracted as $\langle p, \psi \rangle^\phi$, is not able to affect any endpoints; moreover, the coordinator $G : (F_q, d)$ will eventually handle this notification.

- (V) Case $\llbracket \text{rcv} \rrbracket$. Every receiving case is trivial because the coherence has been ensured when every message/notification is sent out. We only prove the standard receiving case $\llbracket \text{rcv} \rrbracket$, where a reduction comes from an endpoint receives a message in the queue. We do not need to consider here that the sender failed, since by $\llbracket \text{crash} \rrbracket$ we know that there is no message to be consumed (i.e., $\llbracket \text{rcv} \rrbracket$ is not applicable).

W.o.l.g. assume $s[p] : \mathcal{E}[q? \{l_i(S_i).T_i\}_{i \in I}] \in \Delta$ and $\mathbf{h} = \langle q, p, l_k(S_k) \rangle \cdot \mathbf{h}'$ and $k \in I$ and $\Delta' = \Delta'_0, s : \mathbf{h}'$ such that $s[p] : \mathcal{E}[T_k] \in \Delta'$.

By Def. 8.(2) and Def. 17 (The Effect of \mathbf{ht}), for any $p' \neq p, s[p'] : T' \in \Delta$, we have

$$\begin{aligned}
 & s[p] : \mathcal{E}[q? \{l_i(S_i).T_i\}_{i \in I}] \downarrow p' - (\langle q, p, l_k(S_k) \rangle \cdot \mathbf{h}')_{p' \rightarrow p} \\
 &= s[p] : \mathcal{E}[T_k] \downarrow p' - (\mathbf{h}')_{p' \rightarrow p} \\
 &\bowtie s[p'] : T' \downarrow p - (\langle q, p, l_k(S_k) \rangle \cdot \mathbf{h}')_{p \rightarrow p'} \\
 &= s[p'] : T' \downarrow p - (\mathbf{h}')_{p \rightarrow p'}
 \end{aligned}$$

Thus in Δ' , we immediately have

$$s[p] : \mathcal{E}[T_k] \downarrow p' - (\mathbf{h}')_{p' \rightarrow p} \bowtie s[p'] : T' \downarrow p - (\mathbf{h}')_{p \rightarrow p'}$$

Thus $\Gamma \vdash \Delta'$ is coherent.

- (VI) Case $\llbracket \mathbb{F} \rrbracket$. If the reduction comes from $\llbracket \mathbb{F} \rrbracket$, in which the coordinator adds a notification $\langle p, \text{crash } F \rangle$ into the queue type to notify the endpoint acting as role p that F occurs.

This case, particularly, gives reasons for

- (a) Def. 11 (Well-formedness).(5) ensures either F will be handled or there is no more interactions involving $\forall q \in F$.
- (b) Def. 11 (Well-formedness).(1),(2) work for avoiding any confusion among partners for handling some F .
- (c) With the above reasons, a coordinator needs to keep a well-formed G as a global guidance.

In the following we give our proofs for (a) and (b).

W.o.l.g assume

- $\Delta = \Delta_0, s : \langle \psi, \text{crash } F \rangle \cdot \mathbf{h}''$,
- $\Delta' = \Delta_0, s : \mathbf{h}'' \cdot \langle \tilde{p}, \text{crash } F \rangle$,
- $\mathbf{h}' = \langle \tilde{p}, \text{crash } F \rangle$,
- $\mathbf{h} = \mathbf{h}'' \cdot \mathbf{h}'$
- $\mathbf{h}''' = \langle \psi, \text{crash } F \rangle \cdot \mathbf{h}''$

$$\begin{aligned} G : (F_q, d) \vdash \Delta_0, s : \langle \psi, \text{crash } F \rangle \cdot \mathbf{h}'' &\rightarrow_T \\ G : (F_q \cup F, d) \vdash \Delta_0, s : \mathbf{h}'' \cdot \langle \tilde{p}, \text{crash } F \rangle & \end{aligned}$$

- (a) If G is well-formed $\langle p, \text{crash } F \rangle \in \mathbf{h}$ means $\forall q \in F$ is non-robust, so a well-formed G shall take care of it (somewhere) by Def. 11.5. For convenience we call all participant which occur in a try-handle partners for that try-handle. Moreover, by Def. 4 (Projection), for every *alive* partners (ensured by Def. 11 (Well-formedness).(1)) for handlers we have (for some try-handle of ϕ)
- * The same set of handler signature and
 - * The same try-handles of ϕ and
 - * They are dual to each other in the handler body by given Def. 8.(b) holds.

They imply that, once a participant is triggered for handling F , all partners will soon or later be triggered for handling F too.

- (b) By Def. 11 (Well-formedness).(1) and rules (**TryHdl**) and $\llbracket \text{TryHdl} \rrbracket$, a participant's try-block always gets triggered by the largest set of failures which currently it is able to handle.

Moreover, Def. 11 (Well-formedness).(2) ensures that the structure of G never results in some partners are handling F in an inner try block but some are handling F in the outer try-handle. The following is reasoning about this point.

By Def. 11 (Well-formedness), we know the possible shapes of G for any F are as follows:

- (1) There is only one handler for F : $\exists t(G_0)h(H)^\kappa \in G, F \in \text{dom}(H)$ and $H(F) = G'$ and $\forall t(G'_0)h(H')^{\kappa'} \in G, \kappa' \neq \kappa. F \notin \text{dom}(H')$.
- (2) There are two handlers for F such that

$$G = \mathcal{G}[t(G_0)h(F:G'_0, H_0)^\kappa].\mathcal{G}'[t(G_1)h(F:G'_1, H_1)^{\kappa'}]$$

Note that, by Def. 11 (Well-formedness), the following shape

$$G = \mathcal{G}[t(\mathcal{G}'[t(G_0)h(F:G'_0, H_0)^\kappa])h(F:G'_1, H_1)^{\kappa'}]$$

is forbidden.

- (3) There are more than two handlers for F . This case is by inductively applying Case (2).

In the following we prove (1) and (2) respectively.

For (1)

- (i) If $\forall m \in \mathbf{h}'', m \neq \langle p, \text{crash } F' \rangle, m \neq \langle \psi, p \rangle^\phi$ for some p, ϕ, F' . We can have $m = \langle p', q', l(S) \rangle$ for some p', q', l, S , Consider $G = \mathcal{G}[t(\dots)h(F:G', H)^\kappa]$. Let $G'' = t(\dots)h(F:G', H)^\kappa$.

In Δ , for any $p, q \in \text{roles}(G'')$ and $s[p], s[q] \in \Delta$, let $\Delta(s[p]) = T$ and $\Delta(s[q]) = T'$.

W.o.l.g assume $s[p] : T \downarrow q - (\mathbf{h}''')_{q \rightarrow p} = s[p] : (\mathcal{E}[G'' \downarrow p.T'']) \downarrow q$ for all p and q .

By Def. 8 (Coherence).(2), we have

$$\begin{aligned} s[p] : T \downarrow q - (\mathbf{h}''')_{q \rightarrow p} &= s[p] : (\mathcal{E}[G'' \downarrow p.T'']) \downarrow q \\ &\bowtie s[q] : (\mathcal{E}[G'' \downarrow q.T''']) \downarrow p \\ &= s[q] : T' \downarrow p - (\mathbf{h}''')_{p \rightarrow q} \end{aligned}$$

In Δ' , for those $p, q \in \text{roles}(G'')$ and $s[p], s[q] \in \Delta$, by Def. 20 and Def. 7 (Duality), we have

$$s[p] : (\text{HdlLogic}(G'' \upharpoonright p, F).T'') \downarrow q \bowtie s[q] : (\text{HdlLogic}(G'' \upharpoonright q, F).T''') \downarrow p$$

In Δ , for any $p, q \notin \text{roles}(G'')$, they will not be affected by F , so they remain being dual to each other and so as their endpoint types in Δ' .

In Δ , for any $p \in \text{roles}(G'')$, $q \notin \text{roles}(G'')$, they do not interact with each other when p is affected by F and $k \notin T \downarrow q, \kappa \notin T^{IV}$ and $\kappa \notin T^V$

$$\begin{aligned} s[p] : T \downarrow q - (\mathbf{h}''')_{q \rightarrow p} &= s[p] : T^{IV} \\ &\bowtie s[q] : T^V \\ &= s[q] : T' \downarrow p - (\mathbf{h}''')_{p \rightarrow q} \end{aligned}$$

In Δ' , their endpoint types hold this relation because they still do not interact to each other in κ .

- (ii) If $F\text{set}(\mathbf{h}'', p) = F' \neq \emptyset$ for some p and $\forall \mathbf{m} \in \mathbf{h}'', \mathbf{m} \neq \langle \psi, p' \rangle^\phi$ for some p' and some ϕ .

Then we consider G with the following shapes:

$$\cdot G = \mathcal{G}[\mathbf{t}(\mathcal{G}'[\mathbf{t}(G_0)\mathbf{h}(F : G', H_0)^\kappa])\mathbf{h}(F' : G'', H_1)^{\kappa'}].$$

In Δ , for any $p, q \in \text{roles}(\mathbf{t}(\mathcal{G}'[\mathbf{t}(G_0)\mathbf{h}(F : G', H_0)^\kappa])\mathbf{h}(F' : G'', H_1)^{\kappa'})$ and $s[p], s[q] \in \Delta$, let $s[p] : \mathcal{E}[T_0]$ and $s[q] : \mathcal{E}'[T'_0]$ w.o.l.g such that T_0 and T'_0 both have the try-handle of (κ', F_s) for some F_s at T_0 and T'_0 .

When $\nexists F'' \supset F \cup F'$ such that $F'' \in G$, then by Def. 8 (Coherence).(2) and Def. 20, we have in Δ

$$\begin{aligned} s[p] : \mathcal{E}[T_0] \downarrow q - (\mathbf{h}''')_{q \rightarrow p} &= s[p] : \mathcal{E}[\text{HdlLogic}(T_0, F')] \downarrow q - (\mathbf{h}''')_{q \rightarrow p} \\ &\bowtie s[q] : \mathcal{E}'[T'_0] \downarrow p - (\mathbf{h}''')_{p \rightarrow q} \\ &= s[q] : \mathcal{E}'[\text{HdlLogic}(T'_0, F')] \downarrow p - (\mathbf{h}''')_{p \rightarrow q} \end{aligned}$$

This case shows that, if there exists failure notifications to trigger outer and inner handler, the outer one takes over.

In Δ' their endpoint types hold

$$\begin{aligned} s[p] : \mathcal{E}[\text{HdlLogic}(T_0, F')] \downarrow q - (\mathbf{h}''' \cdot \mathbf{h}')_{q \rightarrow p} \\ \bowtie s[q] : \mathcal{E}'[\text{HdlLogic}(T'_0, F')] \downarrow p - (\mathbf{h}''' \cdot \mathbf{h}')_{p \rightarrow q} \end{aligned}$$

When $\exists F'' \supset F \cup F'$ such that $F'' \in G$, by Def. 11. (2), F'' must be in a try-handle of $outer_G(\kappa')$ (i.e., F'' is in an outer try-handle or the same try-handle of F'). The proof is similar to the previous case.

$$\cdot G = \mathcal{G}[\mathfrak{t}(\mathcal{G}'[\mathfrak{t}(G_0)\mathfrak{h}(F':G'', H_0)^\kappa])\mathfrak{h}(F:G', H_1)^{\kappa'}].$$

In Δ , for any $p, q \in roles(\mathfrak{t}(\mathcal{G}'[\mathfrak{t}(G_0)\mathfrak{h}(F':G'', H_0)^\kappa])\mathfrak{h}(F:G', H_1)^{\kappa'})$ and $s[p], s[q] \in \Delta$, let $s[p] : \mathcal{E}[T_0]$ and $s[q] : \mathcal{E}'[T'_0]$ such that T_0 and T'_0 both have the try-handle of (κ', F_s) for some F_s .

By Def. 17 (The Effect of **ht**) and Def. 7 (Duality),

$$\begin{aligned} & s[p] : T_0 \downarrow q - (\mathfrak{h} \cdot \mathfrak{h}')_{q \rightarrow p} \\ &= s[p] : \mathcal{E}[HdlLogic(T_0, F)] \downarrow q - (\mathfrak{h}'' \cdot \mathfrak{h}')_{q \rightarrow p} \\ &\bowtie s[q] : T'_0 \downarrow p - (\mathfrak{h} \cdot \mathfrak{h}')_{p \rightarrow q} \\ &= s[q] : \mathcal{E}'[HdlLogic(T'_0, F)] \downarrow p - (\mathfrak{h}'' \cdot \mathfrak{h}')_{p \rightarrow q} \end{aligned}$$

Again, in Δ' their endpoint types hold

$$\begin{aligned} & s[p] : \mathcal{E}[HdlLogic(T_0, F)] \downarrow q - (\mathfrak{h}'' \cdot \mathfrak{h}')_{q \rightarrow p} \\ &\bowtie s[q] : \mathcal{E}'[HdlLogic(T'_0, F)] \downarrow p - (\mathfrak{h}'' \cdot \mathfrak{h}')_{p \rightarrow q} \end{aligned}$$

(iii) If $\exists \mathfrak{m} \in \mathfrak{h}'', \mathfrak{m} = \langle \psi, p \rangle^\phi$ for some p and some $\phi = (\kappa, F_s)$.

Then we consider G having the following shapes:

$$\cdot G = \mathcal{G}[\mathfrak{t}(\mathcal{G}'[\mathfrak{t}(G_0)\mathfrak{h}(H_0)^\kappa])\mathfrak{h}(F:G', H_1)^{\kappa'}] \text{ and } F_s \in dom(H_0)$$

W.o.l.g assume

$$s[p] : \mathcal{E}[\mathfrak{t}(\mathcal{E}'[\mathfrak{t}(T)\mathfrak{h}(\mathcal{H})^{(\kappa, F_s)}.T'''])\mathfrak{h}(\mathcal{H})^{(\kappa, F'')}.T''] \text{ for all } s[p] \in \Delta$$

and $p \in \mathfrak{t}(G_0)\mathfrak{h}(H_0)^\kappa$

$$\begin{aligned} & s[p] : \mathcal{E}[\mathfrak{t}(\mathcal{E}'[\mathfrak{t}(T)\mathfrak{h}(\mathcal{H})^{(\kappa, F_s)}.T'''])\mathfrak{h}(\mathcal{H})^{(\kappa, F'')}.T''] \downarrow q - (\mathfrak{h}''')_{q \rightarrow p} \\ &= s[p] : \mathcal{E}[\mathfrak{t}(\mathcal{E}'[T'''])\mathfrak{h}(\mathcal{H})^{(\kappa, F'')}.T''] \downarrow q - (\mathfrak{h}''' \setminus \langle \psi, p \rangle^\phi)_{q \rightarrow p} \end{aligned}$$

The rest of the proofs follows the reasoning in (i) and (ii).

$$\cdot G = \mathcal{G}[\mathfrak{t}(G_0)\mathfrak{h}(F:G', H)^\kappa] \text{ and } F_s \in dom(H).$$

In this case $Fset(\mathfrak{h}, p)$ (Def. 2 second condition) will not collect $\langle p, crash F \rangle$ to trigger p 's handler in the try-handle of κ . Done notification $\langle \psi, p \rangle^\phi$, $\phi = (\kappa, F_s)$, will finish the try-handle of κ .

For (b).(2), there are two handlers for F , say

$$G = \mathcal{G}[\mathfrak{t}(G_0)\mathfrak{h}(F : G', H_0)^\kappa].\mathcal{G}'[\mathfrak{t}(G_1)\mathfrak{h}(F : G'', H_1)^{\kappa'}]$$

If $\text{roles}(\mathfrak{t}(G_0)\mathfrak{h}(F : G', H_0)^\kappa) \oplus \text{roles}(\mathfrak{t}(G_1)\mathfrak{h}(F : G'', H_1)^{\kappa'}) = \emptyset^8$ the proof is the same as the part in (b).(1).

W.o.l.g assume $p \in \text{roles}(\mathfrak{t}(G_0)\mathfrak{h}(F : G', H_0)^\kappa)$, $q \in \text{roles}(\mathfrak{t}(G_1)\mathfrak{h}(F : G'', H_1)^{\kappa'})$, $q \notin \kappa$, and $\mathfrak{s}[p], \mathfrak{s}[q] \in \Delta$, $\Delta(\mathfrak{s}[p]) = T$, $\Delta(\mathfrak{s}[q]) = T'$ and $T = \mathcal{E}[\mathfrak{t}(T_{F_s})\mathfrak{h}(\mathcal{H})^{(\kappa, F_s)}.T'']$.

- (a) $\Delta(\mathfrak{s}[q]) = T' = \mathcal{E}'[\mathfrak{t}(T_{F'_s})\mathfrak{h}(\mathcal{H}')^{(\kappa', F'_s)}.T''']$
 Assume $\exists F''_s \supset \{F'_s, F''_s\} \in \text{dom}(\mathcal{H}') \wedge F''_s \subseteq \text{Fset}(\mathfrak{h}, q) \wedge F \in F''_s$
 otherwise \mathfrak{h}' has no affect on $\Delta(\mathfrak{s}[q])$.
 In Δ by Def. 8

$$\begin{aligned} \mathfrak{s}[p] &: \mathcal{E}[\mathfrak{t}(T_{F_s})\mathfrak{h}(\mathcal{H})^{(\kappa, F_s)}.T''] \downarrow q - (\mathfrak{h}''')_{q \rightarrow p} \\ &= \mathfrak{s}[p] : \mathcal{E}''[T'' \downarrow q] - (\mathfrak{h}^*)_{q \rightarrow p} \\ \bowtie \mathfrak{s}[q] &: \mathcal{E}'[\mathfrak{t}(T_{F'_s})\mathfrak{h}(\mathcal{H}')^{(\kappa', F'_s)}.T'''] \downarrow p - (\mathfrak{h}''')_{p \rightarrow q} \end{aligned}$$

If $p \notin \kappa'$ then $\mathcal{E}'[\mathfrak{t}(T_{F'_s})\mathfrak{h}(\mathcal{H}')^{(\kappa', F'_s)}.T'''] \downarrow p - (\mathfrak{h}''')_{p \rightarrow q} = \mathcal{E}'''[T''' \downarrow p] - (\mathfrak{h}''')_{p \rightarrow q}$ and $\mathcal{E}'''[T''' \downarrow p] - (\mathfrak{h}''')_{p \rightarrow q} = \mathcal{E}'''[T''' \downarrow p] - (\mathfrak{h})_{p \rightarrow q}$

Otherwise ($p \in \kappa'$): Then
 $\mathcal{E}''[T'' \downarrow q] - (\mathfrak{h}^*)_{q \rightarrow p} = \mathcal{E}''[\mathfrak{t}(T_{F''_s})\mathfrak{h}(\mathcal{H})^{(\kappa', F''_s)}.T^{IV}] - (\mathfrak{h}^*)_{q \rightarrow p}$, by Def. 7

$$\begin{aligned} &\mathcal{E}''[\text{HdlLogic}(\mathfrak{t}(T_{F''_s})\mathfrak{h}(\mathcal{H})^{(\kappa', F''_s)}, F''_s).T^{IV}] - (\mathfrak{h}^*)_{q \rightarrow p} \\ &\bowtie \\ &\mathcal{E}'[\text{HdlLogic}(\mathfrak{t}(T_{F'_s})\mathfrak{h}(\mathcal{H}')^{(\kappa', F'_s)}, F''_s).T'''] - (\mathfrak{h}''')_{p \rightarrow q} \end{aligned}$$

- (b) $\Delta(\mathfrak{s}[q]) = T'$ and $\kappa' \in T'$ and $\nexists \mathcal{E}'.T' = \mathcal{E}'[\mathfrak{t}(T_{F'_s})\mathfrak{h}(\mathcal{H}')^{(\kappa', F'_s)}.T''']$
 $T' \downarrow p - (\mathfrak{h}''')_{p \rightarrow q} = T' \downarrow p - (\mathfrak{h})_{p \rightarrow q}$ ($\llbracket \tilde{p}, \text{crash } F \rrbracket$ has no immediately effect on $\Delta(\mathfrak{s}[q])$)
 (c) $\Delta(\mathfrak{s}[q]) = T'$ and $\kappa' \notin T'$
 $T' \downarrow p - (\mathfrak{h}''')_{p \rightarrow q} = T' \downarrow p - (\mathfrak{h})_{p \rightarrow q}$ ($\llbracket \tilde{p}, \text{crash } F \rrbracket$ has no effect on $\Delta(\mathfrak{s}[q])$ since the try-handle has already finished)

Thus overall, $\Gamma', G : (F_q', d) \vdash \Delta'$ is coherent.

(C) Case $\llbracket \text{collectDone} \rrbracket$. This case is trivial because only the coordinator collects done notifications sent from participants; no endpoints will be affected.

(D) Case $\llbracket \text{IssueDone} \rrbracket$. By $\llbracket \text{IssueDone} \rrbracket$,

$$G : (F_q, d) \vdash \Delta = \Delta_0, \mathfrak{s} : \mathfrak{h} \rightarrow_T G : (F_q, d') \vdash \Delta' = \Delta_0, \mathfrak{s} : \mathfrak{h}'' \cdot \mathfrak{h}'$$

⁸ We use \oplus for the symmetric difference

such that \mathbf{h}'' contains no $\langle p'', \psi \rangle^{(\kappa, F_s)}$ for some p'' and $\forall \mathbf{m} \in \mathbf{h}', \mathbf{m} = \langle \psi, p' \rangle^{(\kappa, F_s)}$ and, let $t(\dots)\mathbf{h}(H)^\kappa \in G$, we have $p' \in \text{roles}(H(F_s))$.

- (a) If there exists $\langle p, \text{crash } F \rangle \in \mathbf{h}'' \cdot \mathbf{h}'$ such that $p \in \text{roles}(H(F_s))$, then by Proposition 1, endpoint $s[p] : T$ will not be affected by failures in $\mathbf{h}'' \cdot \mathbf{h}'$. Therefore, $\langle \psi, p' \rangle^{(\kappa, F_s)}$ is the one that will affect $s[p] : T$ and finishes T 's try-handle of (κ, F_s) .
- (b) If there exists $\langle q, p, l(S) \rangle \in \mathbf{h}$, then by the structure of try-handles, this message type have the following relation with the try-handle of $\phi = (\kappa, F_s)$:
 - if it comes from a try-handle of $\phi' = (\kappa', F_s')$ such that $\kappa \in \text{outer}_G(\kappa')$ then it is already orphan; by Def. 11 (Well-formedness).(3), $\llbracket \text{c1n} \rrbracket$ will be applied.
 - if it comes from a try-handle of $\phi' = (\kappa', F_s')$ such that $\kappa' \in \text{outer}_G(\kappa)$, then by rule $\llbracket \text{RcvDone} \rrbracket$, $\langle \psi, p' \rangle^{(\kappa, F_s)}$ will finish T 's try-handle of F_s .
 - if it comes from a try-handle of $\phi' = (\kappa', F_s')$ such that $\kappa' \notin \text{outer}_G(\kappa)$ and $\kappa \notin \text{outer}_G(\kappa')$ then by rule $\llbracket \text{RcvDone} \rrbracket$, $\langle \psi, p' \rangle^{(\kappa, F_s)}$ will finish T 's try-handle of F_s .

Then the rest of proof is similar to Case 2.(V).(B).

- (E) Case $\llbracket \text{TryHdl} \rrbracket$. In Case $\llbracket \text{F} \rrbracket$, we prove that the failure notifications sent out by the coordinator with a well-formed G can trigger endpoints who are able to handle failures; moreover, Def. 8(Coherence).(2) ensures that the handler for any F are coherent in the sense that every participant has a dual interacting party to handle failures.

After triggering handler(s), the prove is as same as the one in Case 1.(I).(B).

- (F) Case $\llbracket \text{RcvDone} \rrbracket$. The proof is similar to Case 1.(I).(C).

Subject congruence and subject reduction. We first state the inversion lemma and several auxiliary lemmas which will be used in our proofs.

Lemma 5 (Inversion Lemma).

1. If $\Gamma \vdash c : 0 \triangleright \Delta$, then Δ is end-only.
2. If $\Gamma \vdash c : \text{if } e \ \eta_1 \ \text{else } \eta_2 \triangleright \Delta$, then $\Gamma \vdash e : \text{bool}$ and $\forall i \in \{1, 2\}$ we have $\Gamma \vdash c : \eta_i \triangleright \Delta$.
3. If $\Gamma \vdash a[p](y).P \triangleright \Delta$, then $\Delta = \emptyset$ and $\Gamma \vdash a : \langle G \rangle$ and $\Gamma \vdash P \triangleright \{c : G[p]\}$.

4. If $\Gamma \vdash c : p! l(e).\eta \triangleright \Delta$, then $\Delta = \{c : T\}$ $l = l_k$ and $k \in I$ and $T = p! \{l_i(S_i).T'_i\}_{i \in I}$ and $\Gamma \vdash e : S_k$ and $\Gamma \vdash c : \eta_k \triangleright \{c : T'_k\}$.
5. If $\Gamma \vdash c : p? \{l_i(e_i).\eta_i\}_{i \in I} \triangleright \Delta$, then $\Delta = \{c : T\}$ and $T = p? \{l_i(S_i).T_i\}_{i \in I}$ and $\forall i \in I. \Gamma, x_i : S_i \vdash c : \eta_i \triangleright \{c : T_i\}$.
6. If $\Gamma \vdash c : \underline{0}.\eta \triangleright \Delta$, then $\Delta = \{c : T\}$ and $T = \underline{\text{end}}.\text{end}$ and $\Gamma \vdash c : \eta \triangleright \{c : \text{end}\}$.
7. If $\Gamma \vdash c : X\langle e \rangle \triangleright \Delta$, then $\Delta = \{c : T\}$ and $\Gamma = \Gamma', X : S T$ and $\Gamma' \vdash e : S$.
8. If $\Gamma \vdash c : \text{def } D \text{ in } \eta_2 \triangleright \Delta$ and $X(x) = \eta_1 \in D$, then $\Delta = \{c : T\}$ and $\Gamma, X : S \mu t.T' \vdash c : \eta_2 \triangleright \{c : T\}$ and $\Gamma, X : S t, x : S \vdash c : \eta_1 \triangleright \{c : T'\}$.
9. If $\Gamma \vdash c : t(\eta)\mathbf{h}(\mathbf{H})^\phi.\eta' \triangleright \Delta$, then $\Delta = \{c : T\}$ and $T = t(T')\mathbf{h}(\mathcal{H})^\phi.T''$ and $\Gamma \vdash c : \eta \triangleright \{c : T'\}$ and $\Gamma \vdash c : \eta' \triangleright \{c : T''\}$ and $\text{dom}(\mathbf{H}) = \text{dom}(\mathcal{H})$ and $\forall F \in \text{dom}(\mathbf{H})$ we have $\Gamma \vdash c : \mathbf{H}(F) \triangleright \{c : \mathcal{H}(F)\}$.
10. If $\Gamma \vdash s : \emptyset \triangleright \Delta$, then $\Delta = \{s : \emptyset\}$.
11. If $\Gamma \vdash s : h \cdot \langle p, q, l(e) \rangle \triangleright \Delta$, then $\Delta = \{s : \mathbf{h} \cdot \langle p, q, l(S) \rangle\}$ and $\Gamma \vdash s : h \triangleright \{s : \mathbf{h}\}$ and $\Gamma \vdash e : S$.
12. If $\Gamma \vdash s : h \cdot \langle p_1, p_2 \rangle^\phi \triangleright \Delta$, then $(p_1, p_2) \in \{(p, \psi), (\psi, p)\}$ for some p and $\Delta = \{s : \mathbf{h} \cdot \langle p_1, p_2 \rangle^\phi\}$ and $\Gamma \vdash s : h \triangleright \{s : \mathbf{h}\}$.
13. If $\Gamma \vdash s : h \cdot \langle q, \text{crash } F \rangle \triangleright \Delta$, then $\Delta = \{s : \mathbf{h} \cdot \langle q, \text{crash } F \rangle\}$ and $q \in \{p, \psi\}$ and $\Gamma \vdash s : h \triangleright \{s : \mathbf{h}\}$.
14. If $\Gamma \vdash N_1 | N_2 \triangleright \Delta$, then $\Gamma \vdash N_1 \triangleright \Delta_1$ and $\Gamma \vdash N_2 \triangleright \Delta_2$ and $\Delta = \Delta_1, \Delta_2$ such that $\text{dom}(\Delta_1) \cap \text{dom}(\Delta_2) = \emptyset$.
15. If $\Gamma \vdash (\nu s)\mathcal{S} \triangleright \Delta$, then $\Gamma \vdash \mathcal{S} \triangleright \Delta'$ and $\Delta = \Delta' \setminus \Delta'_s$ and $\Gamma \vdash \Delta'_s$ is coherent.
16. If $\Gamma \vdash \Psi \blacklozenge N \triangleright \Delta$, then $\Gamma = \Gamma', \Psi$ and $\Gamma' \vdash N \triangleright \Delta$.

Proof. By induction on derivations.

Lemma 6 (Substitution Lemma).

1. If $\Gamma, x : S \vdash \mathcal{S} \triangleright \Delta$ and $\Gamma \vdash v : S$, then $\Gamma \vdash \mathcal{S}\{v/x\} \triangleright \Delta$.
2. If $\Gamma \vdash \mathcal{S} \triangleright \Delta, y : T$, then $\Gamma \vdash \mathcal{S}\{s[p]/y\} \triangleright \Delta, s[p] : T$.

Proof.

- 1 The proof is by induction on derivation of $\Gamma, x : S \vdash \triangleright \Delta$.
- 2 The proof is by induction on derivation of $\Gamma \vdash \triangleright \Delta, y : T$.

Lemma 7 (Types of Queues).

1. If $\Gamma \vdash s : \langle p, q, l(v) \rangle \cdot h \triangleright \Delta$, then $\Delta = \{s : \langle p, q, l(S) \rangle \cdot \mathbf{h}\}$ and $\Gamma \vdash s : h \triangleright \{s : \mathbf{h}\}$.
2. If $\Gamma \vdash s : \langle \text{role}, \text{crash } F \rangle \cdot h \triangleright \Delta$, then $\text{role} \in \{p, \psi\}$ for some p and $\Delta = \{s : \langle \text{role}, \text{crash } F \rangle \cdot \mathbf{h}\}$ and $\Gamma \vdash s : h \triangleright \{s : \mathbf{h}\}$.
3. If $\Gamma \vdash s : \langle \text{role}, \text{role}' \rangle^\phi \cdot h \triangleright \Delta$, then $(\text{role}, \text{role}') \in \{(p, \psi), (\psi, p)\}$ for some p and $\Delta = \{s : \langle \text{role}, \text{role}' \rangle^\phi \cdot \mathbf{h}\}$ and $\Gamma \vdash s : h \triangleright \{s : \mathbf{h}\}$.

Proof. For all cases, the first step follows from Lemma 5.10. The induction step follows from Lemma 5.11.

Lemma 8. If $\Gamma \vdash c : E[\eta] \triangleright \{c : T\}$ and $\Gamma \vdash c : \eta \triangleright \{c : T'\}$, then $T = \mathcal{E}[T']$ for some \mathcal{E} .

Proof. The proof is by structural induction on contexts E .

– $E = []$, then immediately $\mathcal{E} = []$.

– $E = \text{def } D \text{ in } E'$ and $X(x) = \eta' \in D$ and $\Gamma \vdash c : \eta \triangleright \{c : T'\}$. By applying Lemma 5.8 to

$$\Gamma \vdash c : \text{def } D \text{ in } E'[\eta] \triangleright \{c : T\}$$

we have

$$\Gamma, X : S \mu t.T'' \vdash c : E'[\eta] \triangleright \{c : T\} \text{ and } \Gamma, X : S t, x : S \vdash c : \eta' \triangleright \{c : T''\}$$

Since $\Gamma, X : S \mu t.T'' \vdash c : \eta \triangleright \{c : T'\}$, by induction, we have $T = \mathcal{E}[T']$.

– $E = \mathfrak{t}(E')\mathfrak{h}(\mathbf{H})^\phi.\eta'$ and $\Gamma \vdash c : \eta \triangleright \{c : T'\}$. By applying Lemma 5.9 to

$$\Gamma \vdash c : \mathfrak{t}(E'[\eta])\mathfrak{h}(\mathbf{H})^\phi.\eta' \triangleright \{c : T\}$$

we have

$$T = \mathfrak{t}(T'')\mathfrak{h}(\mathcal{H})^\phi.T''' \text{ and } \Gamma \vdash c : E'[\eta] \triangleright \{c : T''\}$$

By induction, we have $T'' = \mathcal{E}'[T']$ for some \mathcal{E}' . So we have

$$T = \mathfrak{t}(\mathcal{E}'[T'])\mathfrak{h}(\mathcal{H})^\phi.T''' = \mathfrak{t}(T'')\mathfrak{h}(\mathcal{H})^\phi.T'''$$

such that $\mathcal{E} = \mathfrak{t}(\mathcal{E}')\mathfrak{h}(\mathcal{H})^\phi.T'''$.

Theorem 1 (Subject Congruence)

- (1) $\Gamma \vdash N \triangleright \Delta$ and $N \equiv N'$ imply $\Gamma \vdash N' \triangleright \Delta$
- (2) $\Gamma \vdash S \triangleright \Delta$ and $S \equiv S'$ imply $\Gamma \vdash S' \triangleright \Delta$.

Proof. Both proofs are by induction on \equiv . We only list the interesting cases.

– $\frac{h \equiv h'}{s : h \equiv s : h'}$.

Given $h \equiv h'$, we firstly prove that $\Gamma \vdash s : h \triangleright \Delta$ implies $\Gamma \vdash s : h' \triangleright \Delta$.

Then by $\frac{N \equiv N'}{\Psi \blacklozenge N \equiv \Psi \blacklozenge N'}$ (which is proved in the next case) and [T-sys], we have $\Gamma, G : (F_q, d) \vdash \Psi \blacklozenge s : h \triangleright \Delta$ then $\Gamma', G : (F_q, d) \vdash \Psi \blacklozenge s : h' \triangleright \Delta$ and $G : (F_q, d) \in \Psi$.

Let $h \equiv h'$ and $\Gamma \vdash s : h \triangleright \Delta$. The equivalence $h \equiv h'$ should come from one of the following cases: (1) $h \equiv h \cdot \emptyset \equiv h'$ or (2) $h \equiv \emptyset \cdot h \equiv h'$ or (3) $h \equiv h_1 \cdot (h_2 \cdot h_3) \equiv (h_1 \cdot h_2) \cdot h_3 \equiv h'$ or (4) $h \equiv h_1 \cdot m \cdot m' \cdot h_2 \equiv h_1 \cdot m' \cdot m \cdot h_2 \equiv h'$ by given $m \cdot m' \curvearrowright m' \cdot m$.

If the structural congruence comes from cases (1), (2), or (3), since the messages in h and h' are the same and they are in the same order, then by Lemma 7, $\Gamma \vdash s : h' \triangleright \Delta$.

If the structural congruence comes from case (4), given $m \cdot m' \curvearrowright m' \cdot m$, then by Lemma 5.11, Lemma 5.12, Lemma 5.13 and Lemma 7,

$$\Gamma \vdash s : h_1 \cdot m \cdot m' \cdot h_2 \triangleright \{s : h_1 \cdot m \cdot m' \cdot h_2\}$$

imply $\Gamma \vdash s : h_1 \triangleright \{s : h_1\}$ and $\Gamma \vdash s : m \triangleright \{s : m\}$ and $\Gamma \vdash s : m' \triangleright \{s : m'\}$ and $\Gamma \vdash s : h_2 \triangleright \{s : h_2\}$.

By [T-m] or [T-D] or [T-F], we derive

$$\Gamma \vdash s : h_1 \cdot m' \cdot m \cdot h_2 \triangleright \{s : h_1 \cdot m' \cdot m \cdot h_2\}.$$

By Definition 19 (Permutable Message Types),

$$\{s : h_1 \cdot m' \cdot m \cdot h_2\} \equiv \{s : h_1 \cdot m \cdot m' \cdot h_2\} = \Delta$$

Thus we conclude this case.

$$- \frac{N \equiv N'}{\Psi \blacklozenge N \equiv \Psi \blacklozenge N'}$$

By Lemma 5.16, then $G : (F_q, d) \in \psi$ and $\Gamma = \Gamma', G : (F_q, d)$ and $\Gamma' \vdash N \triangleright \Delta$.

By induction on \equiv , we have $\Gamma \vdash N' \triangleright \Delta$. By [T-sys], we have

$$\Gamma \vdash \Psi \blacklozenge N' \triangleright \Delta$$

We conclude this case.

Before proving Theorem 2 (Subject Reduction), we give the following lemma for a special case that when $N \rightarrow \mathcal{S}$:

Lemma 9. $\Gamma \vdash N \triangleright \Delta$ with $\Gamma \vdash \Delta$ coherent and $N \rightarrow \mathcal{S}$ imply that $\Gamma, \psi \vdash \mathcal{S} \triangleright \Delta$.

Proof. The only case is **(Link)**. Assume $\Gamma \vdash a[p_1](y_1).P_1 \mid \dots \mid a[p_n](y_n).P_n \triangleright \Delta$ and $a : G$ and $roles(G) = \{p_1, \dots, p_n\}$.

By applying Lemma 5.(14) n times, we have

$$\Delta = \{\Delta_i\}_{i \in \{1..n\}} \text{ and } \forall i \in \{1..n\}. \Gamma \vdash a[p_i](y_i).P_i \triangleright \Delta_i \quad (3)$$

By applying Lemma 5.(3) on Eq. (3), we have

$$\begin{aligned} \Delta_i = \emptyset &\Rightarrow \Delta = \emptyset \\ \forall i \in \{1..n\}. \Gamma \vdash P_i \triangleright \{y_i : G\uparrow p_i\} \text{ and } \Gamma \vdash a : \langle G \rangle \end{aligned} \quad (4)$$

By applying [T-pa] n times on Eq. (4), we derive

$$\Gamma \vdash P_1 \mid \dots \mid P_n \triangleright \{y_1 : G\uparrow p_1, \dots, y_n : G\uparrow p_n\} \quad (5)$$

By applying Lemma 6.(2) to Eq. (5), we derive

$$\Gamma \vdash P_1\{s[p_1]/y_1\} \mid \dots \mid P_n\{s[p_n]/y_n\} \triangleright \{s[p_1] : G\uparrow p_1, \dots, s[p_n] : G\uparrow p_n\} \quad (6)$$

By applying [T-pa] n and [T- \emptyset] to Eq. (6), we derive

$$\begin{aligned} \Gamma \vdash P_1\{s[p_1]/y_1\} \mid \dots \mid P_n\{s[p_n]/y_n\} \mid \mathbf{s} : \emptyset \triangleright \\ \{s[p_1] : G\uparrow p_1, \dots, s[p_n] : G\uparrow p_n, \{\mathbf{s} : \emptyset\}\} \end{aligned} \quad (7)$$

By applying [T-sys] to Eq. (7), we have

$$\begin{aligned} \Gamma, G : (\emptyset, \emptyset) \vdash G : (\emptyset, \emptyset) \blacklozenge P_1\{s[p_1]/y_1\} \mid \dots \mid P_n\{s[p_n]/y_n\} \mid \mathbf{s} : \emptyset \triangleright \\ \{s[p_1] : G\uparrow p_1, \dots, s[p_n] : G\uparrow p_n, \{\mathbf{s} : \emptyset\}\} \end{aligned} \quad (8)$$

By applying Lemma 3 to Eq. (8), we have

$$\Gamma, G : (\emptyset, \emptyset) \vdash \{s[p_1] : G\uparrow p_1, \dots, s[p_n] : G\uparrow p_n, \{\mathbf{s} : \emptyset\}\} \text{ coherent} \quad (9)$$

By applying [T-s] to Eq. (9) we derive

$$\Gamma, G : (\emptyset, \emptyset) \vdash (\nu \mathbf{s})(G : (\emptyset, \emptyset) \blacklozenge P_1\{s[p_1]/y_1\} \mid \dots \mid P_n\{s[p_n]/y_n\} \mid \mathbf{s} : \emptyset) \triangleright \emptyset$$

Thus $\Delta' = \emptyset = \Delta$. We conclude this case.

Now we prove subject reduction theorem.

Theorem 2 (Subject Reduction)

- (a). $\Gamma \vdash \mathcal{S} \triangleright \Delta$ with $\Gamma \vdash \Delta$ coherent and $\mathcal{S} \rightarrow^* \mathcal{S}'$ imply that $\exists \Delta'$ such that $\Gamma' \vdash \mathcal{S}' \triangleright \Delta'$ and $\Gamma \vdash \Delta \rightarrow_T^* \Gamma' \vdash \Delta'$ or $\Delta \equiv \Delta'$ and $\Gamma' \vdash \Delta'$ coherent.
- (b). $\Gamma \vdash \mathcal{S} \triangleright \emptyset$ and $\mathcal{S} \rightarrow \mathcal{S}'$ imply that $\Gamma \vdash \mathcal{S}' \triangleright \emptyset$.

Proof. **For (a).**

The proof is by induction on the derivation of $\mathcal{S} \rightarrow \mathcal{S}'$. We list rules in Fig. 6 (operational semantics of applications and system).

- **(Snd)**. Assume $\Gamma \vdash \mathfrak{s}[p] : E[q! l(e).\eta] \mid \mathfrak{s} : h \triangleright \Delta$ and $e \Downarrow v$ and $\Gamma \vdash \Delta$ is coherent.

By applying Lemma 5.(14), we have

$$\Delta = \Delta_1, \Delta_2 \text{ and } \Gamma \vdash \mathfrak{s}[p] : E[q! l(e).\eta] \triangleright \Delta_1 \text{ and } \Gamma \vdash \mathfrak{s} : h \triangleright \Delta_2 \quad (10)$$

By applying Lemma 7 and Eq. (10) and [T-m], we have

$$\Delta_2 = \{\mathfrak{s} : \mathfrak{h}\} \text{ and } \Gamma \vdash v : S \text{ and } \Gamma \vdash \mathfrak{s} : h \cdot \langle p, q, l(v) \rangle \triangleright \{\mathfrak{s} : \mathfrak{h} \cdot \langle p, q, l(S) \rangle\} \quad (11)$$

By applying Lemma 5.(4) and Lemma 8 to Eq. (10), we have

$$\Delta_1 = \{\mathfrak{s}[p] : \mathcal{E}[q! l(S).T']\} \text{ and } \Gamma \vdash \mathfrak{s}[p] : E[\eta] \triangleright \{\mathfrak{s}[p] : \mathcal{E}[T']\} \quad (12)$$

By [T-pa] and Eq. (11) and Eq. (12), we derive

$$\Gamma \vdash \mathfrak{s}[p] : E[\eta] \mid \mathfrak{s} : h \cdot \langle p, q, l(v) \rangle \triangleright \{\mathfrak{s}[p] : \mathcal{E}[T'], \mathfrak{s} : \mathfrak{h} \cdot \langle p, q, l(S) \rangle\} \quad (13)$$

Let $\Delta' = \{\mathfrak{s}[p] : \mathcal{E}[T'], \mathfrak{s} : \mathfrak{h} \cdot \langle p, q, l(S) \rangle\}$. By $\llbracket \text{snd} \rrbracket$, we have

$$G : (F_q, d) \vdash \Delta \rightarrow_T G : (F_q, d) \vdash \Delta'$$

and by Theorem 4 $\Gamma \vdash \Delta'$ is coherent. We conclude this case.

- **(Rcv)**. Assume $\Gamma \vdash \mathfrak{s}[p] : E[q?\{l_i(x_i) : \eta_i\}_{i \in I}] \mid \mathfrak{s} : \langle q, p, l(v) \rangle \cdot h \triangleright \Delta$ and $\Gamma \vdash \Delta$ is coherent.

By applying Lemma 5.(14), we have $\Delta = \Delta_1, \Delta_2$

$$\Gamma \vdash \mathfrak{s}[p] : E[q?\{l_i(x_i) : \eta_i\}_{i \in I}] \triangleright \Delta_1 \text{ and } \Gamma \vdash \mathfrak{s} : \langle q, p, l(v) \rangle \cdot h \triangleright \Delta_2 \quad (14)$$

By applying Lemma 7 and Eq. (14), we have

$$\Delta_2 = \{\mathfrak{s} : \langle q, p, l(v) \rangle \cdot \mathfrak{h}\} \text{ and } \Gamma \vdash v : S \text{ and } \Gamma \vdash \mathfrak{s} : h \triangleright \{\mathfrak{s} : \mathfrak{h}\} \quad (15)$$

By applying Lemma 5.(5) and Lemma 8 to Eq. (14) we have

$$\begin{aligned} \Gamma \vdash \mathfrak{s}[p] : E[q?\{l_i(x_i) : \eta_i\}_{i \in I}] \triangleright \{\mathfrak{s}[p] : \mathcal{E}[q?\{l_i(S_i).T_i\}_{i \in I}]\} \\ \forall i \in I. \Gamma, x_i : S_i \vdash \mathfrak{s}[p] : E[\eta_i] \triangleright \{\mathfrak{s}[p] : \mathcal{E}[T_i]\} \end{aligned} \quad (16)$$

By applying Lemma 6 (Substitution) and Eq. (16),

$$\Gamma \vdash \mathfrak{s}[p] : E[\eta_k\{v_k/x_k\}] \triangleright \{\mathfrak{s}[p] : \mathcal{E}[T_k]\} \quad (17)$$

By [T-pa] and Eq. (15) and Eq. (17), we derive

$$\Gamma \vdash \mathfrak{s}[p] : E[\eta_k\{v_k/x_k\}] \mid \mathfrak{s} : h \cdot \langle q, p, l(v) \rangle \triangleright \{\mathfrak{s}[p] : \mathcal{E}[T_k], \mathfrak{s} : \mathfrak{h}\}$$

Let $\Delta' = \{\mathfrak{s}[p] : \mathcal{E}[T_k], \mathfrak{s} : \mathfrak{h}\}$. By $\llbracket \text{rcv} \rrbracket$, we have

$$G : (F_q, d) \vdash \Delta \rightarrow_T G : (F_q, d) \vdash \Delta'$$

and by Theorem 4 $\Gamma \vdash \Delta'$ is coherent. We conclude this case.

- **(Rec)**. Assume $\Gamma \vdash \mathfrak{s}[p] : \text{def } X(x) = \eta \text{ in } X\langle e \rangle \triangleright \Delta$.

By applying Lemma 5.(8), we have

$$\begin{aligned} \Delta &= \{\mathfrak{s}[p] : T\} \\ \Gamma, X : S \ \mu t. T' \vdash \mathfrak{s}[p] : X\langle e \rangle \triangleright \Delta \\ \Gamma, X : S \ t, x : S \vdash \mathfrak{s}[p] : \eta \triangleright \{\mathfrak{s}[p] : T'\} \end{aligned} \quad (18)$$

By applying Lemma 5.(8) to Eq. (18), we have

$$\Gamma = \Gamma', X : S \ T \text{ and } \Gamma' \vdash e : S \quad (19)$$

By applying [T-def] to Eq. (18) and Eq. (19), we derive

$$\Gamma, x : S \vdash \mathfrak{s}[p] : \text{def } X(x) = \eta \text{ in } X\langle e \rangle \triangleright \Delta \quad (20)$$

By applying Lemma 6.(1) to Eq. (20), we have

$$\Gamma \vdash \mathfrak{s}[p] : \text{def } X(x) = \eta \text{ in } \eta\{v/x\} \triangleright \Delta \text{ where } e \Downarrow v$$

Thus we conclude this case.

- **(TryHdl)**. Assume $\Gamma \vdash \mathfrak{s}[p] : E[\mathfrak{t}(\eta)\mathfrak{h}(F:\eta', \mathbf{H})^{(\kappa, F_s)}. \eta''] \mid \mathfrak{s} : h \triangleright \Delta$ and $F = \cup\{A \mid A \in \text{dom}(\mathbf{H}) \wedge F_s \subset A \subseteq \text{Fset}(h, p)\}$ and $\Gamma \vdash \Delta$ is coherent.

By applying Lemma 5.(14), we have

$$\begin{aligned} \Delta &= \Delta_1, \Delta_2 \text{ and} \\ \Gamma \vdash \mathfrak{s}[p] : E[\mathfrak{t}(\eta)\mathfrak{h}(F:\eta', \mathbf{H})^{(\kappa, F_s)}. \eta''] \triangleright \Delta_1 \text{ and} \\ \Gamma \vdash \mathfrak{s} : h \triangleright \Delta_2 \end{aligned} \quad (21)$$

By applying Lemma 8 and Lemma 5.(9) to Eq. (21), we have

$$\begin{aligned} \Delta_1 &= \{\mathfrak{s}[p] : T\} \\ \Gamma \vdash \mathfrak{s}[p] : E[\mathfrak{t}(\eta)\mathfrak{h}(F:\eta', \mathbf{H})^{(\kappa, F_s)}. \eta''] \triangleright \{\mathfrak{s}[p] : T\} \\ T &= \mathcal{E}[\mathfrak{t}(T')\mathfrak{h}(F:T'', \mathcal{H})^{(\kappa, F_s)}. T'''] \\ \Gamma \vdash \mathfrak{s}[p] : \eta \triangleright \{\mathfrak{s}[p] : T'\} \text{ and } \Gamma \vdash \mathfrak{s}[p] : \eta'' \triangleright \{\mathfrak{s}[p] : T'''\} \text{ and} \\ \text{dom}(\mathbf{H}) &= \text{dom}(\mathcal{H}) \text{ and } \forall F \in \text{dom}(\mathbf{H}). \Gamma \vdash \mathfrak{s}[p] : \eta' \triangleright \{\mathfrak{s}[p] : T''\} \end{aligned} \quad (22)$$

By applying Lemma 7 to Eq. (21), we have

$$\Delta_2 = \{\mathfrak{s} : h\} \text{ and } \Gamma \vdash \mathfrak{s} : h \triangleright \{\mathfrak{s} : h\} \quad (23)$$

By [T-th] and Eq. (22), we have

$$\begin{aligned} \Gamma \vdash \mathfrak{s}[p] : E[\mathfrak{t}(\eta')\mathfrak{h}(F:\eta', \mathbf{H})^{(\kappa, F)}. \eta''] \triangleright \\ \{\mathfrak{s}[p] : \mathcal{E}[\mathfrak{t}(T'')\mathfrak{h}(F:T'', \mathcal{H})^{(\kappa, F)}. T''']\} \end{aligned} \quad (24)$$

By [T-pa] and Eqs. (22), (23), (24), we derive

$$\begin{aligned} \Gamma \vdash s[p] : E[\mathfrak{t}(\eta')\mathfrak{h}(F:\eta', \mathbf{H})^{(\kappa, F)}. \eta''] \mid s : h \\ \triangleright \{s[p] : \mathcal{E}[\mathfrak{t}(T'')\mathfrak{h}(F:T'', \mathcal{H})^{(\kappa, F)}. T'''] , s : \mathfrak{h}\} \end{aligned} \quad (25)$$

Let $\Delta' = \{s[p] : \mathcal{E}[\mathfrak{t}(T'')\mathfrak{h}(F:T'', \mathcal{H})^{(\kappa, F)}. T'''] , s : \mathfrak{h}\}$. By [[TryHd1]] we have

$$G : (F_q, d) \vdash \Delta \rightarrow_T G : (F_q, d) \vdash \Delta'$$

and by Theorem 4, $\Gamma \vdash \Delta'$ is coherent. We conclude this case.

- **(CIn)**. Assume $\Gamma \vdash s[p] : E[\eta] \mid s : \langle q, p, l(v) \rangle \cdot h \triangleright \Delta$ and $l \notin \text{labels}(E[\eta])$ and Δ is coherent. (Note that, Def. 17 (The Effect of **ht**) defines that this kind of $\langle q, p, l(v) \rangle$ will not affect any endpoints in Δ).

By applying Lemma 5.(14), we have

$$\Delta = \Delta_1, \Delta_2 \text{ and } \Gamma \vdash s[p] : E[\eta] \triangleright \Delta_1 \text{ and } \Gamma \vdash s : \langle q, p, l(v) \rangle \cdot h \triangleright \Delta_2 \quad (26)$$

By applying Lemma 7 and Eq. (26), we have

$$\Delta_2 = \{s : \langle q, p, l(e) \rangle \cdot \mathfrak{h}\} \text{ and } \Gamma \vdash v : S \text{ and } \Gamma \vdash s : h \triangleright \{s : \mathfrak{h}\} \quad (27)$$

By [T-pa] and Eq. (26) and Eq. (27), we have

$$\Gamma \vdash s[p] : E[\eta] \mid s : h \triangleright \Delta_1, \{s : \mathfrak{h}\} \quad (28)$$

Let $\Delta' = \Delta_1, \{s : \mathfrak{h}\}$. By [[cIn]], we have

$$G : (F_q, d) \vdash \Delta \rightarrow_T G : (F_q, d) \vdash \Delta'$$

and by Theorem 4 $\Gamma \vdash \Delta'$ is coherent. We conclude this case.

- **(CInDone)**. The proof is similar to the case **(CIn)**.
- **(Crash)**. Assume $\Gamma \vdash s[p] : \eta \mid N \mid s : h \triangleright \Delta$.

By applying Lemma 5.(14) two times, we have

$$\begin{aligned} \Delta = \Delta_1, \Delta_2, \Delta_3 \text{ and } \Gamma \vdash s[p] : \eta \triangleright \Delta_1 \text{ and} \\ \Gamma \vdash N \triangleright \Delta_2 \text{ and } \Gamma \vdash s : h \triangleright \Delta_3 \end{aligned} \quad (29)$$

By applying Lemma 7 and [T-m] to Eq. (29), we have

$$\Delta_3 = \{s : \mathfrak{h}\} \text{ and } \mathfrak{h} = \mathfrak{h}_0 \cdot \mathfrak{m}_1 \cdot \mathfrak{h}_1 \dots \mathfrak{h}_{n-1} \cdot \mathfrak{m}_n \cdot \mathfrak{h}_n \text{ where } p \in \mathfrak{m}_i, i \in \{1..n\} \quad (30)$$

Let $\text{msg}(h, p)$ collect messages. By applying Lemma 7 and [T-m] to Eq. (30), we have

$$\Gamma \vdash s : \text{remove}(h, p) \triangleright \{s : (\text{remove}(\mathfrak{h}, p))\} \quad (31)$$

By [T-pa] and [T-F] to Eq. (29) and Eq. (31), we derive

$$\Gamma \vdash N \mid \mathbf{s} : \text{remove}(h, p) \cdot \langle \psi, \text{crash } p \rangle \triangleright \quad (32)$$

$$\Delta_2, \{\mathbf{s} : \text{remove}(\mathbf{h}, p) \cdot \langle \psi, \text{crash } p \rangle\} \quad (33)$$

Let $\Delta' = \Delta_2, \{\mathbf{s} : \mathbf{h} \setminus \text{msg}(\mathbf{h}, p) \cdot \langle \psi, \text{crash } p \rangle\}$. By $\llbracket \text{Crash} \rrbracket$ we have

$$G : (F_q, d) \vdash \Delta \rightarrow_T G : (F_q, d) \vdash \Delta'$$

and by Theorem 4, $\Gamma \vdash \Delta'$ is coherent. We conclude this case.

– (SndDone). Assume $\Gamma \vdash \mathbf{s}[p] : E[\mathbf{t}(0)\mathbf{h}(\mathbf{H})^\phi.\eta'] \mid \mathbf{s} : h \triangleright \Delta$

By applying Lemma 5.(14), we have

$$\begin{aligned} \Delta &= \Delta_1, \Delta_2 \text{ and} \\ \Gamma \vdash \mathbf{s}[p] : E[\mathbf{t}(0)\mathbf{h}(\mathbf{H})^\phi.\eta'] &\triangleright \Delta_1 \text{ and} \\ \Gamma \vdash \mathbf{s} : h &\triangleright \Delta_2 \end{aligned} \quad (34)$$

By applying Lemma 8 and Lemma 5.(9) and Lemma 5.(1) to Eq. (34), we have

$$\Delta_1 = \{\mathbf{s}[p] : T\} \quad (35)$$

$$T = \mathcal{E}[\mathbf{t}(\text{end})\mathbf{h}(\mathcal{H})^\phi.T'] \text{ and}$$

$$\Gamma \vdash \mathbf{s}[p] : 0 \triangleright \{\mathbf{s}[p] : \text{end}\} \text{ and } \Gamma \vdash \mathbf{s}[p] : \eta' \triangleright \{\mathbf{s}[p] : T'\} \text{ and}$$

$$\text{dom}(\mathbf{H}) = \text{dom}(\mathcal{H}) \text{ and}$$

$$\forall F \in \text{dom}(\mathbf{H}). \Gamma \vdash \mathbf{s}[p] : \mathbf{H}(F) \triangleright \{\mathbf{s}[p] : \mathcal{H}(F)\} \quad (36)$$

By applying Lemma 7 and Eq. (34), we have

$$\Delta_2 = \{\mathbf{s} : \mathbf{h}\} \quad (37)$$

By [T-yd] and [T-th] and Eq. (35), we have

$$\Gamma \vdash \mathbf{s}[p] : E[\mathbf{t}(0)\mathbf{h}(\mathbf{H})^\phi.\eta'] \triangleright \{\mathbf{s}[p] : \mathcal{E}[\mathbf{t}(0)\mathbf{h}(\mathcal{H})^\phi.T']\} \quad (38)$$

By [T-D] and Eq. (37), we derive

$$\Gamma \vdash \mathbf{s} : h \cdot \langle p, \psi \rangle^\phi \triangleright \{\mathbf{s} : \mathbf{h} \cdot \langle p, \psi \rangle^\phi\} \quad (39)$$

By [T-pa] and Eqs. (35), (37), (38), (39), we derive

$$\begin{aligned} \Gamma \vdash \mathbf{s}[p] : E[\mathbf{t}(0)\mathbf{h}(\mathbf{H})^\phi.\eta'] \mid \mathbf{s} : h \cdot \langle p, \psi \rangle^\phi &\triangleright \\ \{\mathbf{s}[p] : \mathcal{E}[\mathbf{t}(0)\mathbf{h}(\mathcal{H})^\phi.T'], \mathbf{s} : \mathbf{h} \cdot \langle p, \psi \rangle^\phi\} & \end{aligned}$$

Let

$$\Delta' = \{\mathbf{s}[p] : \mathcal{E}[\mathbf{t}(0)\mathbf{h}(\mathcal{H})^\phi.T'], \mathbf{s} : \mathbf{h} \cdot \langle p, \psi \rangle^\phi\}.$$

By $\llbracket \text{SndDone} \rrbracket$ we have

$$G : (F_q, d) \vdash \Delta \rightarrow_T G : (F_q, d) \vdash \Delta'$$

and by Theorem 4, $\Gamma \vdash \Delta'$ is coherent. We conclude this case.

- **(RcvDone)**. Assume $\Gamma \vdash s[p] : E[\mathbf{t}(0)\mathbf{h}(\mathbf{H})^\phi.\eta] \mid s : h \triangleright \Delta$ and $\langle \psi, p \rangle^\phi \in h$.

By applying Lemma 5.(14), we have

$$\begin{aligned} \Delta &= \Delta_1, \Delta_2 \text{ and} \\ \Gamma \vdash s[p] : E[\mathbf{t}(0)\mathbf{h}(\mathbf{H})^\phi.\eta'] \triangleright \Delta_1 \text{ and} \\ \Gamma \vdash s : h \triangleright \Delta_2 \end{aligned} \quad (40)$$

By applying Lemma 8 and Lemma 5.(6) to Eq. (40), we have

$$\Delta_1 = \{s[p] : T\} \quad (41)$$

$$T = \mathcal{E}[\mathbf{t}(\mathbf{end})\mathbf{h}(\mathcal{H})^\phi.T'] \text{ and}$$

$$\Gamma \vdash s[p] : \mathbf{0} \triangleright \{s[p] : \mathbf{end}\} \text{ and } \Gamma \vdash s[p] : \eta' \triangleright \{s[p] : T'\} \text{ and}$$

$$\text{dom}(\mathbf{H}) = \text{dom}(\mathcal{H}) \text{ and } \forall F \in \text{dom}(\mathbf{H}). \Gamma \vdash \mathbf{H}(F) \triangleright \mathcal{H}(F) \quad (42)$$

By Lemma 7 and [T-D] and Eq. (40) and the condition that $\langle \psi, p \rangle^\phi \in h$, we have

$$\Delta_2 = \{s : \mathbf{h}\} \text{ and } \mathbf{h} = \mathbf{h}' \cdot \langle \psi, p \rangle^\phi \cdot \mathbf{h}'' \quad (43)$$

By applying Lemma 8 to Eq. (41), we have

$$\Gamma \vdash s[p] : E[\eta'] \triangleright \{s[p] : \mathcal{E}[T']\} \quad (44)$$

By applying Eq. (43), we derive

$$\Gamma \vdash s : h \setminus \langle \psi, p \rangle^\phi \triangleright \{s : \mathbf{h}' \cdot \mathbf{h}''\} \quad (45)$$

By [T-pa] and Eqs. (44), (45), we derive

$$\Gamma \vdash s[p] : E[\eta'] \mid s : h \setminus \langle \psi, p \rangle^\phi \triangleright \{s[p] : \mathcal{E}[T'], s : \mathbf{h}' \cdot \mathbf{h}''\}$$

Let $\Delta' = \{s[p] : \mathcal{E}[T'], s : \mathbf{h}' \cdot \mathbf{h}''\}$. By [[RcvDone]] we have

$$G : (F_q, d) \vdash \Delta \rightarrow_T G : (F_q, d) \vdash \Delta'$$

and by Theorem 4, $\Gamma \vdash \Delta'$ is coherent.

- **(F)**. Assume $\Gamma \vdash G : (F_q, d) \blacklozenge s : \langle \psi, \text{crash } F \rangle \cdot h \triangleright \Delta$ and $\exists \mathbf{t}(\dots)\mathbf{h}(H)^\kappa \in G$ such that $F \in \text{dom}(H)$.

By applying Lemma 5.(16), we have

$$\begin{aligned} G : (F_q, d) \in \Gamma \text{ and } \Gamma' \vdash s : \langle \psi, \text{crash } F \rangle \cdot h \triangleright \Delta \\ \text{such that } \Gamma = \Gamma', G : (F_q, d) \end{aligned} \quad (46)$$

By applying Lemma 7 and Eq. (46), we have

$$\Delta = \{s : \langle \psi, \text{crash } F \rangle \cdot \mathbf{h}\} \text{ and } \Gamma' \vdash s : h \triangleright \{s : \mathbf{h}\} \quad (47)$$

By [T-sys] and [T-F] and Eq. (47), we derive

$$\begin{aligned} \Gamma', G : (F_q \cup F, d) \vdash G : (F_q \cup F, d) \blacklozenge \mathbf{s} : \\ h \cdot \langle \text{roles}(G) \setminus (F_q \cup F), \text{crash } F \rangle \cdot \triangleright \\ \{ \mathbf{s} : \mathbf{h} \cdot \langle \text{roles}(G) \setminus (F_q \cup F), \text{crash } F \rangle \} \end{aligned} \quad (48)$$

Let

$$\Delta' = \{ \mathbf{s} : \mathbf{h} \cdot \langle \text{roles}(G) \setminus (F_q \cup F), \text{crash } F \rangle \}$$

By [F] we have

$$G : (F_q, d) \vdash \Delta \rightarrow_T G : (F_q \cup F, d) \vdash \Delta'$$

and by Theorem 4, $\Gamma', G : (F_q \cup F, d) \vdash \Delta'$ is coherent. We conclude this case.

– **(CollectDone)**. The proof is trivial.

– **(IssueDone)**. Assume $\Gamma \blacklozenge \mathbf{s} : h \triangleright \Delta$ and $\Gamma = \Gamma', G : (F_q, d)$ and there exists ϕ such that $\text{roles}(d, \phi) \supseteq \text{roles}(G, \phi) \setminus F_q$ and $\forall F \in \text{hdl}(G, \phi). (F \not\subseteq F_q)$

By applying Lemma 5.(16), we have

$$\Gamma' \vdash \mathbf{s} : h \triangleright \Delta \quad (49)$$

By applying Lemma 7 and Eq. (49) and [T-D], we have

$$\begin{aligned} \Delta &= \{ \mathbf{s} : \mathbf{h} \} \quad (50) \\ \mathbf{h} &= \mathbf{h}_0 \cdot \langle p_1, \psi \rangle^\phi \cdot \mathbf{h}_1 \cdot \dots \cdot \langle p_n, \psi \rangle^\phi \cdot \mathbf{h}_n \text{ where } \text{roles}(d, \phi) = \{ p_1, \dots, p_n \} \end{aligned}$$

By [T-sys] and [T-D] and Eq. (50), we derive

$$\begin{aligned} \Gamma', G : (F_q, \text{remove}(d, \phi)) \vdash G : (F_q, \text{remove}(d, \phi)) \blacklozenge \mathbf{s} : \\ h \cdot \langle \psi, \text{roles}(G, \phi) \setminus (F_q) \rangle^\phi \\ \triangleright \{ \mathbf{s} : (\mathbf{h} \cdot \langle \psi, \text{roles}(G, \phi) \setminus (F_q) \rangle^\phi) \} \end{aligned} \quad (51)$$

Let

$$\Delta' = \{ \mathbf{s} : \{ \mathbf{s} : (\mathbf{h} \cdot \langle \psi, \text{roles}(G, \phi) \setminus (F_q) \rangle^\phi) \} \}$$

By [IssueDone] we have

$$G : (F_q, d) \vdash \Delta \rightarrow_T G : (F_q, d') \vdash \Delta'$$

where $d' = \text{remove}(d, \phi)$ and by Theorem 4, $\Gamma, G : (F_q, d') \vdash \Delta'$ is coherent. We conclude this case.

For (b).

The proof is immediately because $\Gamma \vdash \emptyset$ is always coherent.

Progress. This section proves the property of progress. We first state the following auxiliary lemmas which will be used in our proof.

Lemma 10. If $\Gamma \vdash \Delta$ is coherent and $\mathbf{s} : \mathbf{h} \in \Delta$ and $\Gamma = \Gamma', G : (F_q, d)$, then $\Gamma', G : (F_q, d) \vdash \Delta \rightarrow_T \Gamma', G : (F_q', d') \vdash \Delta'$ for some F_q', d' .

Proof.

- (A) Firstly, if Δ is not end-only and it exists $\mathbf{s}[p] \in \text{dom}(\Delta)$ which is non-robust, by $\llbracket \text{crash} \rrbracket$, it is always possible that $\mathbf{s}[p] \in \text{dom}(\Delta)$ crashes. Then $G : (F_q, d) \vdash \Delta, \mathbf{s}[p] : T, \mathbf{s} : \mathbf{h} \rightarrow_T G : (F_q, d) \vdash \Delta, \mathbf{s} : \mathbf{h} \cdot \langle \psi, \text{crash } F \rangle$. If $\mathbf{h} \equiv \langle \psi, \text{crash } F \rangle \cdot \mathbf{h}'$, it implies $\text{crash } F$ has happened; then either
- (a) $\exists \mathbf{t}(G')\mathbf{h}(H)^\kappa \in G.F \in \text{dom}(H)$, then by $\llbracket \mathbf{F} \rrbracket$, $G : (F_q, d) \vdash \Delta \rightarrow_T G : (F_q, d) \vdash \Delta'$ such that Δ' is the result of the trigger of some handler;
 - (b) Otherwise, it violates rule $\llbracket \text{crash} \rrbracket$.
- (B) If no processes crash and Δ is not end-only, we mechanically prove all cases in Fig. 18. By the structure of local types (defined in Section 5), we have the following possible types at endpoints in Δ :
- 1 There exists $\mathbf{s}[p] : \mathcal{E}[q!\{l_i(S_i).T_i\}_{i \in I}]$. By $\llbracket \text{snd} \rrbracket$, $G : (F_q, d) \vdash \Delta \rightarrow_T G : (F_q, d) \vdash \Delta'$.
 - 2 There exists $\mathbf{s}[p] : \mathcal{E}[\mathbf{t}(\text{end})\mathbf{h}(\mathcal{H})^\phi.T']$. By $\llbracket \text{sndDone} \rrbracket$, $G : (F_q, d) \vdash \Delta \rightarrow_T G : (F_q, d) \vdash \Delta'$.
 - 3 $\mathbf{h} \equiv \langle q, p, l_k(S_k) \rangle \cdot \mathbf{h}'$. If $\mathbf{s}[p] : T \in \Delta$ because $\Gamma \vdash \Delta$ is coherent, then either by $\llbracket \text{rcv} \rrbracket$ or $\llbracket \text{cIn} \rrbracket$, $G : (F_q, d) \vdash \Delta \rightarrow_T G : (F_q, d) \vdash \Delta'$.
 - 4 $\langle \psi, p \rangle^\phi \in \mathbf{h}$ by $\llbracket \text{collectDone} \rrbracket$, $G : (F_q, d) \vdash \Delta, \mathbf{s} : \mathbf{h} \rightarrow_T G : (F_q, d) \cup \{\langle \psi, p \rangle^\phi\} \vdash \Delta, \mathbf{s} : \mathbf{h} \setminus \{\langle \psi, p \rangle^\phi\}$.
 - 5 If $\text{roles}(d, \phi) \supseteq \text{roles}(G, \phi) \setminus F_q$ and $\forall F \in \text{hdl}(G, \phi)$ we have $F \not\subseteq F_q$, then $\llbracket \text{issueDone} \rrbracket$ is applied and $G : (F_q, d) \vdash \Delta, \mathbf{s} : \mathbf{h} \rightarrow_T G : (F_q, \text{remove}(d, \phi)) \vdash \Delta, \mathbf{s} : \mathbf{h} \cdot \langle \psi, \text{roles}(G, \phi) \setminus F_q \rangle^\phi$.
 - 6 $\langle \psi, p \rangle^\phi \in \mathbf{h}$. Because $\Gamma \vdash \Delta$ is coherent, so we have $G : (F_q, d) \in \Gamma$ and there exists $G' = \mathbf{t}(G'')\mathbf{h}(H)^\phi, G' \in G$ and $p \in \text{roles}(G') \setminus F_q$ and $\langle \psi, p \rangle^\phi$ is issued only when $\mathbf{s}[p]$ has sent out $\langle p, \psi \rangle^\phi$, which means, by $\llbracket \text{sndDone} \rrbracket$, either (a) there exists $\mathbf{s}[p] : \mathcal{E}[\mathbf{t}(\text{end})\mathbf{h}(\mathcal{H})^\phi.T]$. Moreover, $p \in \text{roles}(G') \setminus F_q$ implies that $\mathbf{s}[p]$ is alive at the moment. Thus we have $\mathbf{s}[p]$ exists to apply $\llbracket \text{rcvDone} \rrbracket$ such that $G : (F_q, d) \vdash \Delta, \mathbf{s}[p] : \mathcal{E}[\mathbf{t}(\text{end})\mathbf{h}(\mathcal{H})^\phi.T] \rightarrow_T G : (F_q, d) \vdash \Delta, \mathbf{s}[p] : T$. Or (b) $\mathbf{s}[p] = T$ and $\phi \notin T$ then by $\llbracket \text{cInDone} \rrbracket$ $G : (F_q, d) \vdash \Delta, \mathbf{s}[p] : T, \mathbf{s} : \mathbf{h} \rightarrow_T G : (F_q, d) \vdash \Delta, \mathbf{s}[p] : T, \mathbf{s} : \mathbf{h} \setminus \langle \psi, p \rangle^\phi$.

7 $\mathbf{h} \equiv \langle \psi, \text{crash } F \rangle \cdot \mathbf{h}'$. Then $\llbracket \mathbf{F} \rrbracket$ is applied and $G : (F_q, d) \vdash \Delta, \mathbf{s} : \mathbf{h} \rightarrow_T G : (F_q \cup F, d) \vdash \Delta, \mathbf{s} : \mathbf{h}', \langle \tilde{q}, \text{crash } F \rangle (q = \text{roles}(G) \setminus (F_q \cup F))$

8 $\mathbf{h} \equiv \langle p, \text{crash } F \rangle \cdot \mathbf{h}'$. Then either

i. there exists $\mathbf{s}[p] : \mathcal{E}[\mathbf{t}(T)\mathbf{h}(F' : T', \mathcal{H})^\phi.T'']$ such that $F' = \cup\{A \mid A \in \text{dom}(\mathbf{H}) \wedge F_s \subset A \subseteq \text{Fset}(h, p)\}$. Then by $\llbracket \text{TRYHANDLE} \rrbracket$, we have $G : (F_q, d) \vdash \Delta \rightarrow_T G : (F_q, d) \vdash \Delta'$; or

ii. $\mathbf{s}[p] : T$ and $\forall \mathbf{t}(T')\mathbf{h}(\mathcal{H})^\phi \in T, F \notin \text{dom}(\mathcal{H})$. Then we shall have either $T = \text{end}$ or T is one of Cases 1–8.

9 There exists $\mathbf{s}[p] : \mathcal{E}[q?\{l_i(S_i).T_i\}_{i \in I}]$ and $\mathbf{s}[q] \notin \Delta$ by rule $\llbracket \text{CRASH} \rrbracket$. Assume we have $q \in F$ and $\mathbf{h} \equiv \langle p, \text{crash } F \rangle \cdot \mathbf{h}$ (If that is not the case we either have $\mathbf{h} \equiv \langle \psi, \text{crash } F \rangle \cdot \mathbf{h}'$ (goes to case 7) or $\mathbf{h} \equiv \mathbf{h}' \cdot \langle \psi, p \rangle^\phi \cdot \mathbf{h}'' \cdot \langle p, \text{crash } F \rangle \cdot \mathbf{h}'''$ (goes to case (6))). By Def. 11 (Well-formedness) (5) $\mathcal{E}[q?\{l_i(S_i).T_i\}_{i \in I}]$ is enclosed by a try-handle which can handle F . The case goes to (8)i.

(C) If Δ is end-only, then we the statement is true.

We conclude the proof after checking all cases.

Lemma 11. If $\Gamma \vdash \mathcal{S} \triangleright \Delta$ and $G : (F_q, d) \in \Gamma$ and $G : (F_q, d) \vdash \Delta_s \rightarrow_T G : (F_q', d') \vdash \Delta'_s$, then there exists \mathcal{S}' such that $\mathcal{S} \rightarrow \mathcal{S}'$.

Proof. W.l.o.g, assume $\mathcal{S} = \Psi \blacklozenge N | N'$ and $\Gamma \vdash \Psi \blacklozenge N \triangleright \Delta_s$. By Lemma 5 and typing rules defined in Fig. 10 and Fig. 11, $G : (F_q, d) \vdash \Delta_s \rightarrow_T G : (F_q', d') \vdash \Delta'_s$ implies that there exists some $\mathbf{s}[p] : \eta \in N$ and $\mathbf{s} : h \in N$ such that $N = \mathbf{s}[p] : \eta \mid \mathbf{s} : h \mid N_0$ and $\mathbf{s}[p] : \eta \mid \mathbf{s} : h \rightarrow \mathbf{s}[p] : \eta' \mid \mathbf{s} : h'$ Thus there exists $\mathcal{S}' = \Psi \blacklozenge \mathbf{s}[p] : \eta' \mid \mathbf{s} : h' \mid N_0 \mid N'$.

Given $\mathcal{S} = (\nu \mathbf{s})(\Psi \blacklozenge N) \mid N'$ and $\mathbf{s} \notin \text{fn}(N')$ is possible, we can have some participants in N' which do not join session \mathbf{s} but may later or immediately create another session concurrently.

To cater this kind of of situation, we define an *initializable* \mathcal{S} such that, $\forall a[p].P \in \mathcal{S}$, **(Link)** is applicable:

Definition 21 (Initializable \mathcal{S}). Given $\Gamma \vdash \mathcal{S} \triangleright \emptyset$. \mathcal{S} is initializable under Γ if $\exists a[p].P \in \mathcal{S}$ implies that $\Gamma \vdash a : G$ for some $G, \forall p_i \in \text{roles}(G), a[p_i].P_i \in \mathcal{S}$ and **(Link)** is applicable to them.

Theorem 3(Progress) If $\Gamma \vdash \mathcal{S} \triangleright \emptyset$ and \mathcal{S} is initializable, then either $\mathcal{S} \rightarrow^* \mathcal{S}'$ and \mathcal{S}' is initializable or $\mathcal{S}' = \Psi \blacklozenge \mathbf{s} : h \mid \dots \mid \Psi' \blacklozenge \mathbf{s}' : h'$ and h, \dots, h' only contain failure notifications sent by coordinators and messages heading to failed participants.

Proof. The proof is by Lemma 10 and Theorem 1 and Theorem 2.

For convenience, we also use Q to range over processes.

1. Assume no sessions in \mathcal{S} have started, for example

$$\Psi \blacklozenge a[p_1].P_1 \mid \dots \mid a[p_n].P_n \mid N$$

Since \mathcal{S} is initializable, \mathcal{S} either terminates (with only global queue left) or $\mathcal{S} \rightarrow \mathcal{S}'$ and \mathcal{S}' is initializable and $\Gamma \vdash \mathcal{S}' \triangleright \emptyset$ by Theorem 2. (2).

2. Assume some session, say \mathfrak{s} , has started its communication.

W.o.l.g. assume $\mathcal{S} = (\nu \mathfrak{s})(\Psi \blacklozenge N) \mid N'$ and $\mathfrak{s} \notin \text{fn}(N')$, i.e., N' has no any session running in it.

For all $P \in N$ and $Q \in N'$, we have $P \neq Q$. E.g.

$$(\nu \mathfrak{s})(\mathfrak{s}[p] : \eta \mid \mathfrak{s}[p'] : \eta' \mid \mathfrak{s} : h) \mid a[q].Q \mid a[q'].Q'$$

where a is a shared name ready to start another session. For the processes in N' , the proof goes to Case (1), which says that due to initializability those processes will start a session without interfering the existing \mathfrak{s} ; for the processes in N , the proof goes to Case (3).

3. If all sessions in \mathcal{S} have started by **(Link)**.

W.o.l.g. assume $\mathcal{S} = (\nu \mathfrak{s})(\Psi \blacklozenge N) \mid \mathcal{S}''$.

Since $(\nu \mathfrak{s})(\Psi \blacklozenge N)$ and \mathcal{S}' do not interfere each other, we can analyze them independently:

- (a) For the part of $(\nu \mathfrak{s})(\Psi \blacklozenge N)$, by applying Lemma 5.(15) to \mathcal{S} , we have

$$\begin{aligned} \Gamma \vdash (\nu \mathfrak{s})(\Psi \blacklozenge N) \triangleright \emptyset \\ \Gamma \vdash \Psi \blacklozenge N \triangleright \Delta = \Delta_{\mathfrak{s}} \\ \Gamma \vdash \Delta_{\mathfrak{s}} \text{ coherent} \end{aligned} \tag{52}$$

By Lemma 10, we have

$$\Gamma \vdash \Delta \rightarrow_T^* \Gamma \vdash \Delta'$$

where Δ' is end-only for session \mathfrak{s} .

By Lemma 11 and Theorem 2, there exists N' such that

$$\Psi \blacklozenge N \rightarrow^* \Psi \blacklozenge N' \text{ and } \Gamma \vdash \Psi \blacklozenge N' \triangleright \Delta' \text{ and } \mathcal{S} \rightarrow^* (\nu \mathfrak{s})(\Psi \blacklozenge N') = \mathcal{S}' \tag{53}$$

By applying [T-s] to Eq. (53), we have $\Gamma \vdash (\nu \mathfrak{s})\mathcal{S}' = \mathcal{S}'' \triangleright \emptyset$.

Since Δ' is end-only for session \mathbf{s} , it implies all interactions at endpoint processes in \mathbf{s} have finished.

Moreover, based on **(CollectDone)**, **(IssueDone)**, and **(F)**, since a coordinator always consume notifications of $\langle p, \psi \rangle^\phi$ and $\langle \psi, \text{crash } F \rangle$ for any p, ϕ, F , so if there is any this kind of message left, the coordinator will consume so $\mathcal{S}' \rightarrow \mathcal{S}''$ until there is no more such messages in \mathcal{S}'' .

- (b) For the part of \mathcal{S}'' , since all participants have joined some sessions (by assumption), the proof is as same as the above.