

Session Types with Runtime Adaptation: Overview and Examples

Cinzia Di Giusto

Université d'Evry - Val d'Essonne, Laboratoire IBISC, Evry, France

Jorge A. Pérez

CITI and Departamento de Informática, FCT - Universidade Nova de Lisboa, Portugal

In recent work, we have developed a session types discipline for a calculus that features the usual constructs for session establishment and communication, but also two novel constructs that enable communicating processes to be stopped, duplicated, or discarded at runtime. The aim is to understand whether known techniques for the static analysis of structured communications scale up to the challenging context of context-aware, adaptable distributed systems, in which disciplined interaction and runtime adaptation are intertwined concerns. In this short note, we summarize the main features of our session-typed framework with runtime adaptation, and recall its basic correctness properties. We illustrate our framework by means of examples. In particular, we present a session representation of *supervision trees*, a mechanism for enforcing fault-tolerant applications in the Erlang language.

1 Introduction

The type-theoretic foundations provided by *session types* have proved successful in the analysis of complex communication scenarios. This is witnessed by, e.g., the several extensions and enhancements of early session types proposals [11] with common concepts/idioms in practical distributed programming. Interestingly, while often such extensions have appealed to increasingly sophisticated type structures (which feature, e.g., subtyping [10], dependent/indexed types [5], and kinding [4]), the associated modeling language—the polyadic π -calculus—has remained essentially the same. This “asymmetric” development of session types and process languages should not appear as a surprise: because of its canonicity and expressiveness, the π -calculus represents a rather natural choice for representing concurrent systems which interact by following precise—yet intricate—behavioral patterns.

We are interested in one such patterns, namely *runtime adaptation*: it allows to represent the suspension, restarting, replacement, or abortion of running processes. As such, it can be useful to model and reason about important mechanisms in modern distributed systems, such as code mobility, online software update, failure recovery (as in, e.g., constructs for exceptions/compensations), and scaling (i.e., the ability of acquiring/releasing computing resources based on runtime demand). These mechanisms typically have a global effect over the system and are meant to be executed atomically. Hence, they appear difficult to implement as π -calculus specifications: it is not obvious how name passing/scope extrusion—the central abstraction vehicles of the π -calculus—can be used to model this kind of global reconfiguration primitives while retaining atomicity and an adequate level of abstraction for reasoning.

Based on these limitations, and with the aim of setting a formal framework for reasoning about communicating systems with dynamic reconfiguration, with Bravetti and Zavattaro we have developed a framework of *adaptable processes* [2]. Adaptable processes extend usual process calculi with *located processes* $l[P]$ and *update processes* $\text{upd}(l, U(X))$, where P is a process and U is a context, i.e., a process with zero or more occurrences of a *process variable* X . While located processes explicitly represent

distributed behavior, update processes provide a built-in adaptation mechanism for located processes. In fact, by synchronizing on name l , processes $l[P]$ and $\text{upd}(l, U(X))$ may reduce to $U[P/X]$ —the process which results from substituting all free occurrences of X in U with P . This way, we obtain a mechanism for dynamic process reconfiguration which is performed atomically in a single reduction.

In recent work [7, 8], we have investigated the integration of a session-typed discipline into a π -calculus with located and update processes. The intention is to understand whether session types scale up to the challenging context of context-aware, adaptable distributed systems, in which disciplined communication and dynamic reconfiguration are intertwined concerns. Indeed, processes in our language may evolve either by the usual forms of synchronization but also by performing evolvability actions. Hence, besides showing that well-typed processes enjoy *safety* (i.e. absence of communication errors at runtime), we have proved that typing entails *consistency*, a property that guarantees that an update action for a located process is only enabled if such a process is not already involved in an active session. In this way, consistency ensures a disciplined interplay of communicating and evolvability actions. Our work borrows inspiration from the approach of Garralda et al [9] on (binary) session types for Ambients, and extends it to the case of adaptable processes which run in arbitrary, nested distributed locations.

In this presentation, we summarize the framework in [7, 8] and illustrate it by presenting a process representation of *supervision trees*, a mechanism used in the Erlang language to enforce fault-tolerant programs [1]. Intuitively, supervision trees define a hierarchical structure of program components, in which *workers* perform computation and *supervisors* monitor the behavior of one or more workers. In case of failure or unexpected termination of a worker, its supervisor should restart its behavior. The supervisor module in Erlang defines different *restarting strategies* for supervisors. For instance, in the *one for one* restarting strategy, only the affected worker is restarted, and so supervision realizes a form of local adaptation. In contrast, in the *one for all* strategy, the supervisor restarts both the affected worker and also all of its sibling in the supervision tree, and so supervision enforces a form of global adaptation. We show how both these strategies may be represented in our typed process framework. We believe the resulting models are useful to appreciate the main features of our intended model of runtime adaptation.

Disclaimer and Organization. This paper is intended as a brief overview presentation to the approach and results first presented in [7]. As such, our focus here is in intuitions and examples rather than on technical details; the reader is referred to [7] (and to its extended version [8]) for extended details.

The rest of the paper is structured as follows. In Sect. 2 we present a motivating example for our intended model. Then, Sect. 3 gives a overview of our process framework for runtime adaptation with session types, and describes its main properties. In Sect. 4 we present our model of supervision trees as session-typed processes —both this model and the example in Sect. 2 are original to this presentation. Finally, in Sect. 5 we collect some concluding remarks and discuss directions for future work.

2 A Motivating Example

We begin by discussing a simple model of a *workflow application*, which extends the one in [2, § 4.2]. Our model combines the main features of adaptable processes (nested locations and update processes) with *delegation*, the well-known mechanism for representing reconfiguration in session-typed processes.

Briefly, a workflow is a conceptual unit that describes how a number of *activities* coordinate to achieve a given task. A *workflow-based application* usually consists of a *workflow runtime engine* that contains a number of workflows running concurrently on top of it; a *workflow base library* on which activities may rely on; and of a number of *runtime services*, which are typically application dependent.

Exploiting nested located processes, in [2, § 4.2] we propose the following high-level process representation of a workflow application:

$$App \triangleq \text{wfa} \left[\text{we}[\text{WE} \mid W_1 \mid \dots \mid W_k \mid \text{wbl}[\text{BL}]] \mid S_1 \mid \dots \mid S_j \right]$$

This way, the application is modeled as a located process wfa which contains a workflow engine (located at we) and a number of runtime services S_1, \dots, S_j . In turn, the workflow engine contains a number of workflows W_1, \dots, W_k , a process WE (which represents the engine's behavior and is left unspecified), and a located process wbl representing the base library (also left unspecified). Each workflow W_i (with $1 \leq i \leq k$) contains a process WL_i (representing the workflow's logic) and is composed of n activities. Each activity is formalized as an adaptable process a_j and an *execution environment* env_j . The intention is that while the execution environment hosts the current (session) behavior of the activity (denoted P_j below), location a_j contains an adaptation routine for location env_j (denoted R_j below). More precisely, our process representation for W_i is the following:

$$W_i \triangleq w_i \left[\text{WL}_i \mid \prod_{j=1}^n \text{env}_j[P_j] \mid a_j[R_j] \right] \quad \text{where}$$

$$R_j \triangleq \text{request } a(x).x(b).\text{if } b = \text{true} \text{ then } (\text{upd}(\text{env}_j, \text{env}_j[Q_j]) \mid \text{close}(x)) \text{ else } \text{close}(x)$$

Intuitively, P_j is an unspecified session behavior that may interact with other activities, either inside or outside workflow W_i . Process R_j is meant to establish a session on name a with a dual behavior contained in WL_i . The purpose of such a session is very simple: WL_i will send to R_j a boolean expression which may determine runtime adaptation for location env_j . In fact, if such an expression reduces to true then an update process at env_j is released and the session is closed. The purpose of such an update is to replace the current behavior of env_j with process Q_j . Otherwise, if the boolean reduces to false , then the session between WL_i and R_j is closed and no modification on env_j is performed.

Although the possibility of adapting the behavior of session-typed processes at runtime is quite appealing, arbitrary update actions may jeopardize the communication protocols described by session types. For instance, if the adaptation that R_j may perform takes place while P_j is in the middle of an active session, then a communication error would arise from the resulting unmatched communication steps. Similarly, if P_j contains replicated server definitions which Q_j does not implement, then overall service availability at the workflow level would be affected. It then becomes important to add a certain discipline to update actions, in such a way that (i) prescribed communication behaviors are not disrupted by them, and (ii) service definitions are always uniformly updated. In our session-typed process framework with runtime adaptation (summarized in the following section), we address these two natural requirements for disciplining update actions.

3 A Framework of Disciplined, Adaptable Processes

3.1 Syntax and Semantics

Our process language corresponds to the extension of the usual polyadic π -calculus for binary session types (cf. [11, 12]) with located and update processes. Besides disjoint base sets for *names*, ranged over a, b, x, y, \dots ; *labels*, ranged over n, n', \dots ; and *constants* (integers, booleans, names), ranged over c, c', \dots , we consider *locations*, ranged over l, l', \dots ; *integers*, ranged over j, h, m, \dots ; and *process variables*,

$$\begin{array}{l}
\text{Expressions } e ::= c \mid e_1 + e_2 \mid e_1 - e_2 \mid \dots \\
\text{Processes } P ::= \text{request } a(x).P \mid \text{accept } a(x).P \mid \text{!accept } a(x).P \\
\quad \mid l[P] \mid X \mid \text{upd}(l, P(X)) \\
\quad \mid k\langle\bar{e}\rangle.P \mid k(\bar{e}).P \mid k\langle\langle k' \rangle\rangle.P \mid k\langle\langle d \rangle\rangle.P \\
\quad \mid k \triangleright \{n_1:P_1 \parallel \dots \parallel n_m:P_m\} \mid k \triangleleft n; P \\
\quad \mid P_1 \mid P_1 \mid \text{if } e \text{ then } P \text{ else } Q \mid \text{close}(k).P \mid (\nu k)P \mid \mathbf{0}
\end{array}$$

Figure 1: Process syntax.

ranged over X, X', \dots . Then, *processes*, ranged over P, Q, R, \dots and *expressions*, ranged over e, e', \dots are given by the grammar in Fig. 1. In the figure, we use k, k' to denote names x and (*polarized*) *channels* [10, 12], ranged over $\kappa^p, \kappa_1^p, \dots$ (with $p \in \{+, -\}$). Channels are runtime entities, as formalized by our operational semantics (see below).

Our process syntax contains an almost standard session π -calculus, with session passing (delegation) [12] and replicated session acceptance constructs. Observe that we consider an explicit prefix for session closing, denoted $\text{close}(k)$, which is useful to track the active sessions in a located process. Also, we consider only restriction over channels k . Constructs related to runtime adaptation are given in the second line of the grammar for processes. We have located processes $l[P]$ and update processes $\text{upd}(l, P(X))$, where $P(X)$ denotes a process with one or more occurrences of a process variable X . Locations are transparent: hence, in $l[P]$ process P may evolve autonomously until it is updated by an update action on l . One objective of our type system is then to discipline such update actions, so as to avoid updating processes which contain open sessions.

The operational semantics for our process language is defined as a reduction relation —Fig. 2 reports the required reduction rules. The semantics generates and maintains *runtime annotations* for locations: annotated located processes are written $l^h[P]$, where integer h denotes the number of *active sessions* in P . These annotations are key in avoiding update actions over located processes which are currently engaged in active sessions. Reduction rules rely on (syntactic) contexts (ranged over C, D, E), which represent the nested structure of located processes. Such contexts are defined by the following syntax:

$$C, D, E, \dots ::= \{\bullet\} \mid l^h[C \mid P]$$

This way, given a context $C\{\bullet\}$ and a process P , we write $C\{P\}$ to denote the process obtained by filling in the holes in C with P . The intention is that P may reduce inside C , thus reflecting the transparent containment realized by location nesting. The semantics relies on a structural congruence relation (omitted), which handles scope extrusion for channels in the usual way. Rules (R:OPEN) and (R:CLOSE), which formalize session establishment and termination, resp., use two operations over contexts, denoted C^+ and C^- . Informally, C^+ (resp. C^-) denotes the context obtained from C by increasing (resp. decreasing) the annotation h in all of its located processes. Rule (R:UPD) captures the essence of our notion of consistency: a (located) process Q can be updated only if it contains no active sessions. The other rules are completely standard, generalized to our setting of nested, transparent locations.

3.2 Type System

We have deliberately aimed at retaining a standard session type structure; see Fig. 3. We assume a notion of *duality* over session types α , noted $\bar{\alpha}$, defined as usual. We now comment on our notion of *typing judgment*. We extend usual judgments with an explicit *interface*, denoted \mathcal{I} . Intuitively, interfaces are

$$\begin{array}{l}
\text{(R:OPEN)} \quad E\{C\{\text{accept } a(x).P\} \mid D\{\text{request } a(y).Q\}\} \longrightarrow \\
\quad \quad \quad E^{++}\{(\nu\kappa)(C^+\{P[\kappa^+/x]\} \mid D^+\{Q[\kappa^-/y]\})\} \\
\text{(R:ROPEN)} \quad E\{C\{\text{!accept } a(x).P\} \mid D\{\text{request } a(y).Q\}\} \longrightarrow \\
\quad \quad \quad E^{++}\{(\nu\kappa)(C^+\{P[\kappa^+/x] \mid \text{!accept } a(x).P\} \mid D^+\{Q[\kappa^-/y]\})\} \\
\text{(R:UPD)} \quad E\{C\{l^0[Q]\} \mid D\{\text{upd}(l, P(X))\}\} \longrightarrow E\{C\{P[Q/X]\} \mid D\{0\}\} \\
\text{(R:I/O)} \quad E\{C\{\kappa^p \langle \tilde{e} \rangle . P\} \mid D\{\kappa^{\bar{p}}(\tilde{x}).Q\}\} \longrightarrow E\{C\{P\} \mid D\{Q[\tilde{c}/\tilde{x}]\}\} \quad (\tilde{e} \downarrow \tilde{c}) \\
\text{(R:PASS)} \quad E\{C\{\kappa^p \langle \langle \kappa'^q \rangle \rangle . P\} \mid D\{\kappa^{\bar{p}}((x)).Q\}\} \longrightarrow E\{C^-\{P\} \mid D^+\{Q[\kappa'^q/x]\}\} \\
\text{(R:SEL)} \quad E\{C\{\kappa^p \triangleright \{n_1:P_1 \parallel \dots \parallel n_m:P_m\}\} \mid D\{\kappa^{\bar{p}} \triangleleft n_j; Q\}\} \longrightarrow \\
\quad \quad \quad E\{C\{P_j\} \mid D\{Q\}\} \quad (1 \leq j \leq m) \\
\text{(R:CLOSE)} \quad E\{C\{\text{close}(\kappa^p).P\} \mid D\{\text{close}(\kappa^{\bar{p}}).Q\}\} \longrightarrow E^{--}\{C^-\{P\} \mid D^-\{Q\}\} \\
\text{(R:IFTR)} \quad C\{\text{if } e \text{ then } P \text{ else } Q\} \longrightarrow C\{P\} \quad (e \downarrow \text{true}) \\
\text{(R:IFFA)} \quad C\{\text{if } e \text{ then } P \text{ else } Q\} \longrightarrow C\{Q\} \quad (e \downarrow \text{false}) \\
\text{(R:STR)} \quad \text{if } P \equiv P', P' \longrightarrow Q', \text{ and } Q' \equiv Q \text{ then } P \longrightarrow Q \\
\text{(R:PAR)} \quad \text{if } P \longrightarrow P' \text{ then } P \mid Q \longrightarrow P' \mid Q \\
\text{(R:RES)} \quad \text{if } P \longrightarrow P' \text{ then } (\nu\kappa)P \longrightarrow (\nu\kappa)P'
\end{array}$$

Figure 2: Reduction Rules. We use $\tilde{e} \downarrow \tilde{c}$ to denote an evaluation relation on expressions.

assignments from names to session types which describe the services declared in a given process. We also consider typings Δ and environments Γ and Θ . As typical in session types disciplines, typing Δ collects assignments from channels from session types, therefore describing currently active sessions. In Δ we also include *bracketed assignments*, denoted $[\kappa^p : \alpha]$, which represent active but restricted sessions. As we discuss in [8], bracketed assignments arise in the typing of channel restriction, and are key to keep a precise count of the active sessions in a given location. Γ is a first-order environment which maps expressions to basic types and names to pairs of *qualified* session types. Within the interface, session types can be linear (lin) or unrestricted (un). While lin is used for those session types to be used just once, un annotates those session types intended to feature a persistent behavior. The higher-order environment Θ collects assignments of process variables and locations to interfaces. While the former kind of assignments are relevant to update processes, the latter concern located processes. Given these environments, a *type judgment* is of form

$$\Gamma; \Theta \vdash P \triangleright \Delta \circ \mathcal{I}$$

meaning that, under environments Γ and Θ , process P has active sessions declared in Δ and interface \mathcal{I} .

We only present typing rules for adaptable and update processes; see [8] for a full account.

$$\begin{array}{c}
\text{(T:LOC)} \\
\frac{\Theta \vdash l : \mathcal{I} \quad \Gamma; \Theta \vdash P \triangleright \Delta \circ \mathcal{I}' \quad h = |\Delta| \quad \mathcal{I}' \sqsubseteq \mathcal{I}}{\Gamma; \Theta \vdash l^h[P] \triangleright \Delta \circ \mathcal{I}'} \\
\text{(T:ADAPT)} \\
\frac{\Theta \vdash l : \mathcal{I} \quad \Gamma; \Theta, X : \mathcal{I} \vdash P \triangleright \emptyset \circ \mathcal{I}'}{\Gamma; \Theta \vdash \text{upd}(l, P(X)) \triangleright \emptyset \circ \emptyset}
\end{array}$$

Rule (T:LOC) types located processes and performs two checks. First, the runtime annotation h is computed by counting the assignments (standard and bracketed) declared in Δ . Second, the rule checks that the interface of the located process is less or equal than the declared interface of the given location.

TYPES			
τ, σ	$::=$	int bool ...	basic types
α, β	$::=$	$!(\tilde{\tau}).\alpha$ $?(\tilde{\tau}).\alpha$	send, receive
		$!(\beta).\alpha$ $?(\beta).\alpha$	throw, catch
		$\&\{n_1 : \alpha_1, \dots, n_m : \alpha_m\}$	branch
		$\oplus\{n_1 : \alpha_1, \dots, n_m : \alpha_m\}$	select
		ε	closed session
ENVIRONMENTS			
q	$::=$	lin un	type qualifiers
\mathcal{I}	$::=$	\emptyset $\mathcal{I}, \langle a : \alpha, h \rangle$	interface
Δ	$::=$	\emptyset $\Delta, k : \alpha$ $\Delta, [k : \alpha]$	typing with active sessions
Γ	$::=$	\emptyset $\Gamma, e : \tau$ $\Gamma, a : \langle \alpha_q, \overline{\alpha_q} \rangle$	first-order environment
Θ	$::=$	\emptyset $\Theta, X : \mathcal{I}$ $\Theta, l : \mathcal{I}$	higher-order environment

Figure 3: Type Syntax and Typing Environments

In the rule, \sqsubseteq denotes a preorder over interfaces, defined in [8]; hence, the premise in the rule informally ensures that the process behavior does not “exceed” the expected behavior within the location. It is worth observing how a typed located processes has the exact same typing and interface of its contained process: this is how transparency of locations arises in typing. Rule (T:ADAPT) types update processes. Observe how the interface associated to the process variable of the given runtime adaptation routine should match with the declared interfaces for the given location.

3.3 Safety and Consistency

We now summarize the main properties of well-typed processes. They are thoroughly developed in [8]. We say a typing Δ is balanced iff for all $\kappa^p : \alpha \in \Delta$ (resp. $[\kappa^p : \alpha] \in \Delta$) then also $\kappa^{\overline{p}} : \overline{\alpha} \in \Delta$ (resp. $[\kappa^{\overline{p}} : \overline{\alpha}] \in \Delta$). Then we have the following *subject reduction* theorem:

Theorem 1 (Subject Reduction) *If $\Gamma; \Theta \vdash P \triangleright \Delta \circ \mathcal{I}$ with Δ balanced and $P \longrightarrow Q$ then $\Gamma; \Theta \vdash Q \triangleright \Delta' \circ \mathcal{I}'$, for some \mathcal{I}' and balanced Δ' .*

In order to state runtime safety, we require some auxiliary definitions. Process P is said to be a κ -process if it sends (resp. receives) a value/session on κ^p , if it makes a selection (resp. offers a choice), or if it closes a session on κ^p . Then, two κ -processes constitute a κ -redex if they are complementary to each other, possibly enclosed in suitable contexts: i.e., if one sends, the other receives on κ^p ; if one makes a selection, the other chooses; or if they are closing the same session κ^p . Then, P is an *error* if, up to structural congruence, it contains either exactly two κ -processes that do not form a κ -redex or three or more κ -processes. Our first result is that well-typed process do not lead to communication errors:

Theorem 2 (Typing Ensures Runtime Safety) *If $\Gamma; \Theta \vdash P \triangleright \Delta \circ \mathcal{I}$ with Δ balanced then P never reduces into an error.*

We now state and proof *consistency*, the property that ensures that update actions do not disrupt session behavior. We require an auxiliary notation. Let $\longrightarrow_{\text{upd}}$ stand for an update action, i.e., a reduction inferred using rule (R:UPD), possibly with the interplay of (structural) congruence rules. We have:




Figure 4: A supervision tree: supervisors are depicted as rectangles; workers are represented by circles.

Definition 1 (Consistency) *A process P is update-consistent if and only if, for all P' and κ such that $P \longrightarrow^* P'$ and P' contains a κ -redex, if $P' \longrightarrow_{\text{upd}} P''$ then P'' contains a κ -redex.*

Our type system ensures that both annotations enabling update actions and interface assignments to locations are correctly assigned and maintained along reduction. Indeed, as our second result we have:

Theorem 3 (Typing Ensures Update Consistency) *If $\Gamma; \Theta \vdash P \triangleright \Delta \circ \mathcal{I}$, with Δ balanced, then P is update consistent.*

4 A Session Model of Supervision Trees

In this section, we present encodings of *supervision trees* in our session language with located and update processes. As supervision trees (as provided by the supervisor module of Erlang [1]) define a basic principle for designing and programming fault-tolerant programs, our encodings are useful to understand how these our two constructs for runtime adaptation (a) may integrate into communicating processes and (b) relate to actual programming mechanisms.

As mentioned in the Introduction, supervision trees organize a program’s behavior into two classes: *workers* and *supervisors*. Workers are the processes that actually perform computation—in a communication-centric setting, workers may well correspond to servers which expect and fulfill clients’ requests. Supervisors are meant to monitor the behavior of workers by relying on a hierarchical, tree-like structure. This way, a supervisor may monitor the behavior of a number of children; the supervisor may itself be supervised by another supervisor at a higher level in the hierarchy. As an example, Fig. 4 shows a supervisor S_2 that monitors workers W_3 and W_4 ; in turn, S_2 is supervised by S_1 who is also a supervisor for worker W_2 . In Erlang, a supervisor can implement several *restart strategies* which, given a termination signal (e.g. a failure), offer different possibilities for restarting the supervised worker(s). Such strategies are fully detailed in [1]; in what follows, we concentrate on the following two:

1. *One for one* strategy: if one of the children fails and terminates then only that child is restarted.
2. *One for all* strategy: if one of the children fails and terminates then all the siblings must be terminated and restarted.

It is interesting to observe how the hierarchical model for fault-tolerance realized by supervision trees in Erlang is in line with our intended model for runtime adaptation: in addition to the loose analogy between located processes as workers and update processes as supervisors, we find useful similarities in terms of hierarchical organization of behavior, and in the fact that both supervision trees and our update

processes are in essence *external* recovery mechanisms for precisely delimited behaviors. These intuitions are our starting point for the process representations of supervision trees which we analyze in the rest of this section. For simplicity, we restrict to representing the hierarchy given in Fig. 4, assuming that both S_1 and S_2 realize the same restarting strategy; the representation, however, can be easily generalized to more complex settings. We represent the tree structure by the nesting of located processes, as follows:

$$l_1 \left[S_1 \mid l_2 \left[S_2 \mid w_3[W_3]w_4[W_4] \mid w_2[W_2] \right] \right] \quad (1)$$

In our models, we distinguish between localities for supervisors (denoted l_i) and localities for workers (denoted w_j). In the following, we describe two different implementations of supervisors, each realizing a different restarting strategy.

4.1 Modeling One for one Supervisors

As suggested by the above informal description, the one for one strategy accounts for a form of *local supervision*: whenever a child enters in an error state or fails, the child activity must be terminated and the server restarted. This way, for instance, in the scenario of Fig. 4, if worker W_3 fails then its siblings W_4 and W_2 should be unaffected by (and unaware of) the restarting process that S_2 is expected to initiate.

The one for one strategy can be encoded in a quite natural way within our framework. Indeed, by assuming that workers correspond to servers which are connected to clients, we may represent failures as the reception of an invalid value from the client; the restarting process would then correspond to an update action that simply recreates the located process containing the server. More precisely, we assume worker W_3 is an “arithmetic server” which calculates the division between two natural numbers which are communicated by a client C after establishing a session. In this context, failure corresponds to division by zero. Then, W_3 notifies S_2 whether it has found a failure (i.e., the denominator being equal to zero). Subsequently, S_2 should decide what to do next: (i) if the denominator sent by the client was zero, then there is indeed a failure and session between W_3 and C must terminate immediately; (ii) otherwise, if there is not a failure, then S_2 should allow W_3 to perform the division and communicate the result to C . Observe how in this “local” scenario, supervisor S_2 does not need to interact with its supervisor S_1 .

To implement the supervision model described above, every time a worker (such as W_3) establishes a session with a client (such as C), it also connects to its supervisor. This way, supervisor and worker have a unique communication mechanism. In our case, this mechanism is realized by a session request from W_3 to S_2 on name b . By extending the process in (1) with a client C , we obtain the following system:

$$l_1 \left[S_1 \mid l_2 \left[S_2 \mid w_3[W_3]w_4[W_4] \mid w_2[W_2] \right] \mid C \right] \quad (2)$$

where S_2 , W_3 , and C are defined as follows (the other processes are left unspecified):

$$S_2 \triangleq ! \text{accept } b(x_b).x_b(v).x_b \triangleleft v. \text{if } v = \text{fail} \text{ then close}(x_b).\text{upd}(w_3, w_3[W_3]) \text{ else close}(x_b)$$

$$W_3 \triangleq ! \text{accept } a(x_a).\text{request } b(x_b).x_a(u).x_a(v). \text{if } v = 0 \text{ then } x_b \langle \text{fail} \rangle.P \text{ else } x_b \langle \text{ok} \rangle.P$$

$$P \triangleq x_b \triangleright \{ \text{fail} : x_a \triangleleft \text{fail}.\text{close}(x_a).\text{close}(x_b) \parallel \text{ok} : x_a \triangleleft \text{ok}.x_a \langle u \div v \rangle.\text{close}(x_a).\text{close}(x_b) \}$$

$$C \triangleq \text{request } a(x_a).x_a \langle c_1 \rangle.x_a \langle c_2 \rangle.x_a \triangleright \{ \text{fail} : \text{close}(x_a) \parallel \text{ok} : x_a(v).\text{close}(x_a) \}$$

Observe how process S_2 defines an update action on w_3 if W_3 communicates to him that the invalid value zero was received. Moreover, using our update construct we can implement more interesting

supervisors which go beyond service restarting. For instance, consider the following process:

$$S'_2 \triangleq !\text{accept } b(x_b).x_b(v).x_b \triangleleft v.\text{close}(x_b).\text{if } v = \text{fail then upd}(w_3, w_3^0[W'_3]) \text{ else } 0$$

Note how S'_2 improves over S_2 by placing inside w_3 a *different* server behavior, represented by W'_3 . Suppose that W'_3 represents an upgraded version of W_3 in which the arithmetic server no longer relies on a supervisor in order to deal with failures/invalid values. This way, whenever W_3 communicates a failure to S'_2 , the resulting update action will lead to an improved server W'_3 that will be able to handle the failure locally, without appealing to an external entity.

Finally, we briefly discuss the typing of previous processes. Given session types

$$\alpha \triangleq ?(\text{int}).?(\text{int}).\oplus \{\text{fail} : \varepsilon, \text{ok} : !(\text{int}).\varepsilon\} \quad \text{and} \quad \beta \triangleq ?(\text{str}).\oplus \{\text{fail} : \varepsilon, \text{ok} : \varepsilon\}$$

the processes above are well-typed, given the following environments:

$$\begin{aligned} \Gamma &\supseteq \{a : \langle \alpha_{\text{un}}, \bar{\alpha}_{\text{lin}} \rangle, b : \langle \beta_{\text{un}}, \bar{\beta}_{\text{lin}} \rangle\} \\ \Theta &\supseteq \{l_2 : \langle b : \beta, \infty \rangle, w_3 : \langle a : \alpha, \infty \rangle, \langle b : \bar{\beta}, \infty \rangle\} \end{aligned}$$

In fact, processes S_2 , W_3 , and C can be typed as follows:

$$\begin{aligned} \Gamma; \Theta &\vdash S_2 \triangleright \emptyset \circ \langle b : \beta, \infty \rangle \\ \Gamma; \Theta &\vdash W_3 \triangleright \emptyset \circ \langle a : \alpha, \infty \rangle, \langle b : \bar{\beta}, \infty \rangle \\ \Gamma; \Theta &\vdash C \triangleright \emptyset \circ \langle a : \bar{\alpha}, 1 \rangle \end{aligned}$$

Notice that as a consequence of being well-typed, the supervisor has to wait that all sessions have been closed before proceeding with an update of its worker.

4.2 Modeling One for all Supervisors

While the one for one restarting strategy implements a local relation between a worker and his supervisor, in the one for all strategy supervisors need to globally control their children. Hence, whenever one of the children enters in an error state, the child and its siblings have to be terminated and restarted. For instance, in the scenario of Fig. 4, if W_2 fails then both W_2 and all the workers contained in the subtree rooted in S_2 need to be terminated and restarted. This event should be communicated by S_1 to S_2 , which should handle the termination of its children. On the other hand, if W_3 fails then its supervisor S_2 only has to stop W_4 , and nothing needs to be communicated to S_1 .

The required global control that supervisors should exercise over children in the one for all strategy turns out to be slightly unnatural to express in our framework of binary communication. Nevertheless, it is possible to devise a protocol that achieves the desired synchronization sequences. An alternative to represent global synchronization among services and supervisors is described by the following informal protocol, in which communication between workers and clients (and between supervisors and workers) is managed via session establishment:

1. Each worker receives a data value v from a client. The value is checked for validity (via a boolean expression $e(v)$). The result of this check is sent to the worker's supervisor.
2. Each supervisor waits to receive the result of the validity check from all its children.
3. Each supervisor (apart from the root node S_1) sends a true value to its own supervisor. This is to have a uniform communication pattern, and is related to the fact that supervisors themselves cannot fail.

4. Starting from the root of the tree, a communication cascade takes place: each supervisor communicates to its children if they can continue with their own service or if they have to terminate. Intuitively, a single worker can continue if and only if (i) none of its siblings have failed and (ii) the supervisor has not received from its own supervisor a failure signal.
5. If a worker (server) receives a failure signal, then it stops the communications with its clients; otherwise, it can continue with its intended behavior.
6. Finally, when a worker/server stops it is recreated by its supervisor. In the case of a forced termination this corresponds to the actual restarting strategy; in the case of natural termination, this enables server persistency. In the process Q below, this is realized via an update. Notice that S_1 will need to recreate only W_2 , while S_2 will be responsible for the restart of its own children W_3 and W_4 .

Notice that Fig. 4 already shows in the edges the names of the services established between each pair. As before, we consider a process structure as in (2). Here we just give the encoding of worker W_3 —the other workers have analogous representations. Moreover, we suppose that W_3 can establish a session i with client C . Notice that each worker can establish a session with just one client at the time. We have:

$$\begin{aligned}
W_3 &::= \text{accept } i(x_i).\text{request } c(x_c).x_i(v).\text{if } e(v) \text{ then } x_c\langle\text{ok}\rangle.P \text{ else } x_c\langle\text{fail}\rangle.P \\
P &::= x_c(y).x_c \triangleleft y.x_c \triangleright \{ \text{fail} : x_i \triangleleft \text{fail}.\text{close}(x_i).\text{close}(x_c) \parallel \text{ok} : x_i \triangleleft \text{ok}.\text{Proto}_W.\text{close}(x_i).\text{close}(x_c) \} \\
C &::= \text{request } i(x_i).x_i\langle c \rangle.x_i \triangleright \{ \text{fail} : \text{close}(x_i) \parallel \text{ok} : \text{Proto}_C.\text{close}(x_i) \} \\
\\
S_2 &::= \text{!accept } c(x_c).\text{accept } d(x_d).\text{request } a(x_a).x_c(u).x_d(v).x_a\langle\text{true}\rangle.x_a(z).x_a \triangleleft z. \\
&\quad \text{if } u \wedge v \wedge (z = \text{ok}) \text{ then } x_c\langle\text{ok}\rangle.x_d\langle\text{ok}\rangle.R \text{ else } x_c\langle\text{fail}\rangle.x_d\langle\text{fail}\rangle.R \\
R &::= x_c \triangleright \{ \text{fail} : x_c \triangleleft \text{fail}.x_d \triangleright \{ \text{fail} : x_d \triangleleft \text{fail}.\text{close}(x_c).\text{close}(x_d).\text{close}(x_a).Q \} \parallel \\
&\quad \text{ok} : x_c \triangleleft \text{ok}.x_d \triangleright \{ \text{ok} : x_d \triangleleft \text{ok}.\text{Proto}_{S_2}.\text{close}(x_c).\text{close}(x_d).\text{close}(x_a).Q \} \} \\
Q &::= \text{upd}(w_3, w_3[W_3]) \mid \text{upd}(w_4, w_4[W_4]) \\
\\
S_1 &::= \text{!accept } a(x_a).\text{accept } b(x_b).x_a(u).x_b(v).\text{if } u \wedge v \text{ then } x_a\langle\text{ok}\rangle.x_b\langle\text{ok}\rangle.T \text{ else } x_a\langle\text{fail}\rangle.x_b\langle\text{fail}\rangle.T \\
T &::= x_a \triangleright \{ \text{fail} : x_a \triangleleft \text{fail}.x_b \triangleright \{ \text{fail} : x_b \triangleleft \text{fail}.\text{close}(x_a).\text{close}(x_b).\text{upd}(w_2, w_2[W_2]) \} \parallel \\
&\quad \text{ok} : x_a \triangleleft \text{ok}.x_b \triangleright \{ \text{ok} : x_b \triangleleft \text{ok}.\text{Proto}_{S_1}.\text{close}(x_a).\text{close}(x_b).\text{upd}(w_2, w_2[W_2]) \} \}
\end{aligned}$$

Above, Proto_W , Proto_C , and Proto_{S_i} denote the rest of the behavior of workers, clients and supervisors; as such, we assume that Proto_W and Proto_C are the dual of each other. Observe how as soon as worker W_3 has established a session on i with the client, it contacts its own supervisor through c and implements step (1) of the protocol above, then it waits for the decision of the supervisor. Similarly for its supervisor S_2 : after it has established a session for each worker (W_3 and W_4) through services c and d respectively, it contacts S_1 via a and implements steps (2) and (3) of the protocol. Finally, the root supervisor S_1 will take a decision and communicate to its workers whether they have to terminate or continue. Then it will pass the information to (child) supervisor S_2 , that in turn will rule the behavior of its children (cf. steps (4) and (5)). As a last step (point (6)) each supervisor will take care of recreating of its own children.

The system above can be shown to be well-typed, analogously as we did for the process representation for the one for one strategy. Here again, typing ensures that no session is disrupted while communication with clients is still active.

5 Concluding Remarks and Future Work

In this short note, we have presented the main elements of our work in [7, 8], where we have investigated the integration of constructs for runtime adaptation into a session types discipline. Our starting point was our own work on adaptable processes [2], a process calculi framework for specifying interacting processes which may be suspended, restarted, updated, or discarded at runtime. Our work in [7] appears to be the first attempt in addressing the integration of dynamic evolution issues in models of communicating systems in which interaction is described (and disciplined) by session types.

The original contributions of this presentation are examples which illustrate how the constructs for runtime adaptation fit in a session-typed, communication-centric setting. We presented and discussed a process representation of Erlang’s supervision trees, an existing mechanism for designing and programming fault-tolerant applications. Similarly as update processes, supervisors are external entities which monitor the behavior of one or more workers. Supervisors are parametric in the restarting strategies that are spawned when a failure occurs in a worker. We described representations of supervisors into our session process language, and observed that although a model of local supervision (in which a supervisor monitors only one worker) is better suited to the communication primitives in our model, an alternative model of global supervision (in which a supervisor may monitor more than one worker) can also be encoded. Interestingly, in both cases our session language with runtime adaptation capabilities is able to represent more insightful restarting strategies than those provided by the Erlang language.

The framework in [7, 8], summarized in this note, paves the way for several avenues of future work:

- We would like to investigate *progress (deadlock freedom)* for adaptable, session-typed processes. This includes adapting to our setting known session type systems for ensuring progress [6, 3], but also understanding whether the information added by such systems (e.g., orderings on sessions) can be integrated into update actions so as to prevent/overcome deadlocked sessions at runtime.
- Our nested, transparent locations may be too “open” for some applications. One may need to enforce *secure updates*, so as to protect located processes from unintended evolvability actions —as in, for instance, update actions triggered by *unauthorized* users, or update actions based on *permissions* for replacing, extending, or destroying the behavior of a located process.
- In this presentation, we have refrained from enriching the typing system with information related to updates. We think such an extension would not be difficult, and could offer a way to better control update actions. This way, for instance, one could decree (and statically check) that update actions occur only in selected parts of a communication protocol.

Acknowledgments. We are grateful to the anonymous reviewers: their comments and remarks were most useful in improving the presentation of this short note. Our interest in exploring how supervision trees in Erlang could be modeled in our framework was triggered by a suggestion from Philip Wadler during PLACES’13 to whom we are grateful. This research was supported by the Portuguese Foundation for Science and Technology (FCT) under grants SFRH / BPD / 84067 / 2012 and CITI, and by French project ANR BLAN 0307 01 - SYNBIOTIC.

References

- [1] Ericsson AB (2013): *Erlang/OTP - System Documentation: Supervisor Module*. <http://www.erlang.org/doc/man/supervisor.html>.

- [2] Mario Bravetti, Cinzia Di Giusto, Jorge A. Pérez & Gianluigi Zavattaro (2012): *Adaptable Processes*. *Logical Methods in Computer Science* 8(4). Available at [http://dx.doi.org/10.2168/LMCS-8\(4:13\)2012](http://dx.doi.org/10.2168/LMCS-8(4:13)2012). Extended abstract in Proc. of FMOODS-FORTE'11, Springer, LNCS 6722.
- [3] Marco Carbone & Søren Debois (2010): *A Graphical Approach to Progress for Structured Communication in Web Services*. In Simon Bliudze, Roberto Bruni, Davide Grohmann & Alexandra Silva, editors: *ICE, EPTCS* 38, pp. 13–27. Available at <http://dx.doi.org/10.4204/EPTCS.38.4>.
- [4] Romain Demangeon & Kohei Honda (2012): *Nested Protocols in Session Types*. In Maciej Koutny & Irek Ulidowski, editors: *CONCUR, Lecture Notes in Computer Science* 7454, Springer, pp. 272–286. Available at http://dx.doi.org/10.1007/978-3-642-32940-1_20.
- [5] Pierre-Malo Deniérou, Nobuko Yoshida, Andi Bejleri & Raymond Hu (2012): *Parameterised Multiparty Session Types*. *Logical Methods in Computer Science* 8(4). Available at [http://dx.doi.org/10.2168/LMCS-8\(4:6\)2012](http://dx.doi.org/10.2168/LMCS-8(4:6)2012).
- [6] Mariangiola Dezani-Ciancaglini, Ugo de'Liguoro & Nobuko Yoshida (2007): *On Progress for Structured Communications*. In: *TGC 2007, Lecture Notes in Computer Science* 4912, Springer, pp. 257–275. Available at http://dx.doi.org/10.1007/978-3-540-78663-4_18.
- [7] Cinzia Di Giusto & Jorge A. Pérez (2013): *Disciplined structured communications with consistent runtime adaptation*. In Sung Y. Shin & José Carlos Maldonado, editors: *SAC*, ACM, pp. 1913–1918. Available at <http://doi.acm.org/10.1145/2480362.2480716>.
- [8] Cinzia Di Giusto & Jorge A. Pérez (2013): *Disciplined Structured Communications with Consistent Runtime Adaptation (Extended Version)*. Technical Report, CITI - FCT UNL. Available at <http://www.jorgeaperez.net/publications/sac13ext.pdf>.
- [9] Pablo Garralda, Adriana B. Compagnoni & Mariangiola Dezani-Ciancaglini (2006): *BASS: boxed ambients with safe sessions*. In: *PPDP*, ACM, pp. 61–72. Available at <http://doi.acm.org/10.1145/1140335.1140344>.
- [10] Simon J. Gay & Malcolm Hole (2005): *Subtyping for session types in the pi calculus*. *Acta Inf.* 42(2-3), pp. 191–225. Available at <http://dx.doi.org/10.1007/s00236-005-0177-z>.
- [11] Kohei Honda, Vasco Thudichum Vasconcelos & Makoto Kubo (1998): *Language Primitives and Type Discipline for Structured Communication-Based Programming*. In: *ESOP, LNCS* 1381, Springer, pp. 122–138. Available at <http://dx.doi.org/10.1007/BFb0053567>.
- [12] Nobuko Yoshida & Vasco Thudichum Vasconcelos (2007): *Language Primitives and Type Discipline for Structured Communication-Based Programming Revisited: Two Systems for Higher-Order Session Communication*. *Electr. Notes Theor. Comput. Sci.* 171(4), pp. 73–93. Available at <http://dx.doi.org/10.1016/j.entcs.2007.02.056>.