

Corecursion in Session-Typed Processes

Bernardo Toninho^{*†}, Luís Caires^{*} and Frank Pfenning[†]

^{*}CITI and FCT - Universidade Nova de Lisboa

[†]Carnegie Mellon University

Abstract—Session types are widely accepted as a useful expressive discipline for structuring communications in concurrent and distributed systems. In order to express infinitely unbounded sessions, as required to model the behaviour of realistic distributed systems, session typed languages often introduce general recursion operators at both the program and the type level. Unfortunately, general recursion, in particular when combined with name passing and mobility, may easily introduce undesirable divergence, e.g., infinite unobservable reduction sequences.

In this paper we address, by means of typing, the challenge of ensuring non-divergence in a session-typed π -calculus with general (co)recursion, while still allowing interesting infinite behaviours to be definable, as necessary to model realistic open process networks and distributed systems. Our approach builds on a Curry-Howard correspondence between our type system and a standard system of linear logic extended with co-inductive types, for which our non-divergence property implies consistency.

We prove type safety for our framework, implying protocol compliance and global progress of well-typed processes. Remarkably, we also establish - through a logical relations argument - that well-typed processes are compositionally non-divergent, in the sense that no well-typed composition of processes, including compositions dynamically assembled through name passing, can result in divergent behaviour, a property of practical relevance.

I. INTRODUCTION

This work addresses, by means of a logically motivated typing discipline, the problem of ensuring compositional non-divergence in session-based process systems with infinite behaviour defined using (co)recursive types.

We live in an age of concurrent and distributed software, meant to run not only as local applications but as cooperating parts of larger, distributed and mobile services, meant to run indefinitely with multiple independent clients, which are hard to build and ensure correct. Process models combined with techniques for precisely characterising and analysing system behavior have been exploited to verify properties such as deadlock freedom, protocol compliance, and availability of distributed and service based systems. Among type-based approaches to verification, a rather successful technique has been that of *session types* [1], [2], [3].

Session types structure message-based concurrency around the notion of a session, which is a precise description of the interaction patterns of two (or more) communicating agents with an intrinsic notion of protocol state (e.g. input a string and then output an integer) and duality. Thus, session types are a form of protocol descriptions that can be statically checked for compliance. The recent discovery of a correspondence between (all linear and shared features of) session types and pure Linear Logic in the style of Curry-Howard [4], linking proofs with typed processes and proof reduction as communication,

has sparked a renewed interest in session types and their foundations [5], and in the idea of exploiting logically motivated approaches for providing powerful reasoning techniques about concurrent and distributed systems. This line of work has addressed concepts such as value-dependent types [6], proof-carrying code [7] behavioural polymorphism [8], and higher-order computation [9], approaching a general type theory for session-based concurrency, with strong guarantees by virtue of typing such as deadlock freedom and session fidelity.

Although in the untyped case infinite behavior is potentially encodable using replication in the π -calculus, to express infinitely unbounded session exchanges, as required to model the behaviour of realistic distributed systems, session typed languages often introduce general recursion operators at both the program and the type level. In a typed setting, replication can only capture *finite session behavior* that can be replicated (shared) arbitrarily often. It is not rich enough to model *infinite session behavior*, nor repeating behavioral patterns that depend on evolving state.

Unfortunately, existing session type systems for (channel-passing, higher-order) systems equipped with general recursion operators do not avoid validating systems exhibiting undesirable divergence, e.g., infinite sequences of unobservable internal reduction steps. While this issue already arises at the level of individual systems, it but becomes more serious when one needs to consider realistic dynamically linked distributed systems. For example, plugging together subsystems, e.g. as a result of dynamic channel passing or of linking downloaded (higher-order) code to local services, may undesirably result in a system actually unable to offer its intended services due to divergent behavior, even if the component subsystems are well-typed and divergence-free, and this may be caused by simple programming error or by malicious intentions (e.g. a denial of service attack).

In this work, we tackle the challenge of reconciling general recursion, thus enabling potentially infinite behaviour to be expressed, with local termination and strong normalization within a session typed framework. As a toy example, consider a web service from Twitter offering a replicated service *trends* which is intended to produce a stream of current trends, according to some custom metrics. To that end, the service is parametrized by a filter process that given a stream of tweets produces a stream of trends. The code implementing the service is:

$$\text{Serv} \triangleq !\text{trends}(x).x(f).(\nu y)f\langle y\rangle.(\text{TweetSrc}_y \mid [f \leftrightarrow x])$$

When invoked, the *Serv* replicated service creates a new

session on fresh channel x , inputs on session x the appropriate filter function f , links it to the internal tweet source TweetSrc_y and then repeatedly forward the results back on session x . The session types involved are:

$$\begin{aligned} \text{TrendService} &\triangleq !((\text{Tweets} \multimap \text{Trends}) \multimap \text{Trends}) \\ \text{Tweets} &\triangleq \nu X.(\text{tweet} \wedge X) \quad \text{Trends} \triangleq \nu Y.(\text{trend} \wedge Y) \end{aligned}$$

A possible client for the service Serv is

$$\text{Joe} \triangleq (\nu x)\text{trends}\langle x \rangle.(\nu k)x\langle k \rangle.(\text{F}_k | (\text{corec } Z.x(y).p\langle y \rangle.Z))$$

This process invokes the shared server, resulting in a fresh session in channel x , sends (access to) a custom analytics package F_k , and then sits in a loop printing out (on session p) each trend received from the server in session x . In our type system, we may derive the typing judgment

$$\text{trends}:\text{TrendService} \vdash \text{Joe} :: p:\text{Trends}$$

Notice that we build on the intuitionistic formulation of linear logic based session types [4], where typing judgments have the form $\Gamma; \Delta \vdash P :: x:U$. This states of process P that it provides a session of type U at x , when composed with services / sessions as specified by $\Gamma; \Delta$. Using the cut rule, we may compose the server with a client in a system $\text{Sys} \triangleq (\nu \text{trends})(\text{Serv} | \text{Joe})$. System Sys will (unboundedly) print out trends on channel p , actually, system Sys is typed as $\vdash \text{Sys} :: p:\text{Trends}$. Being typed in our system, and although it generates an infinite stream of trends at p , involves higher-order name passing, dynamic linking, and occurrences of recursive calls, Sys will never get into internal divergence. It is challenging to obtain an expressive typing discipline ensuring the compositional non-divergence property. For instance, consider the very similar looking processes

$$\begin{aligned} \text{Loop} &\triangleq \text{corec } L(c).c(x).(\nu d)(L(d) | d\langle n \rangle.[d \leftrightarrow c]) \\ \text{Good} &\triangleq \text{corec } G(c).c\langle x \rangle.(\nu d)(G(d) | c\langle n \rangle.[d \leftrightarrow c]) \end{aligned}$$

These two processes do not autonomously eventually diverge. It is possible to type Good , thus ensuring that it will never diverge, even when composed with arbitrary (well-typed) processes. However, this is not the case for process Loop , which gets into an infinite internal reduction sequence after the first communication on c , and (of course) is not typeable in our type system.

Our work has important practical and foundational consequences. From a foundational perspective, we proceed by developing a theory of coinductive session types, fully preserving the connections with linear logic in the sense of [4]. In this context, non-divergence implies termination of cut-elimination, and therefore we establish logical consistency of our framework. From a practical perspective, we provide a typing discipline for session types supporting rich infinite interactive behaviour, while rejecting systems that may engage into autistic internal divergence. More precisely, our typing discipline ensures, together with global progress (actually, lock-freedom) and protocol fidelity, the *compositional non-divergence* of infinite behaviors (i.e. that there is no well-typed process context under which a well-typed process will evolve

to a divergent behavior), an important property out of the scope of existing session type systems with recursive types.

We summarise the contributions of our work:

- We introduce a session type system for our process calculus based on linear logic with coinductive types, associating corecursive process definitions with coinductive session types which encode potentially infinite session behavior.
- We show that well-typed processes enjoy very strong safety properties such as type preservation (or session fidelity) and progress, even in the presence of corecursion.
- We prove that well-typed processes are *compositionally non-divergent* by employing a logical relations argument, extended to coinductive session types, which combined with type safety ensures that any well-typed, distributed service implemented in our calculus will not only never “get stuck”, but will also never become unavailable through divergent behavior.

II. PROCESS MODEL

In this section we introduce our process calculus, essentially consisting of a (synchronous) π -calculus with basic data types for convenience, input-guarded replication, labelled choice and selection and corecursion. The syntax of processes is given below:

$$\begin{aligned} M, N & ::= \dots \quad (\text{basic data constructors}) \\ P, Q & ::= x\langle M \rangle.P \mid x(y).P \mid x(y).P \mid (\nu y)P \\ & \quad \mid !x(y).P \mid P \mid Q \mid x.\text{case}(l_j \Rightarrow P_j) \mid x.l_i; P \\ & \quad \mid (\text{corec } X(\bar{y}).P) \bar{c} \mid X(\bar{c}) \mid [x \leftrightarrow y] \mid \mathbf{0} \end{aligned}$$

We range over basic data with M, N and processes with P, Q . We write \bar{y} and \bar{c} for a list of variables and channels, respectively. We write $fn(P)$ for the free names of process P . Basic data type constructors include the typical constructs for manipulating data such as numbers, strings and lists. The process language is a synchronous π -calculus with term input $x(y).P$ and output $x\langle M \rangle.P$, channel output $c\langle y \rangle.P$, input-guarded replication $!x(y).P$, n -ary labelled choice $x.\text{case}(l_j \Rightarrow P_j)$ and selection $x.l_i; P$, channel forwarding $[x \leftrightarrow y]$ and, crucially, a parametrized *corecursion* operator $(\text{corec } X(\bar{y}).P) \bar{c}$, enabling corecursive process definitions (the variables \bar{y} are bound in P). The parameters are used to instantiate channels (and values) in recursive calls accordingly.

The operational behavior of processes is given in terms of reduction and labelled transitions, both defined modulo structural congruence (written \equiv) which captures basic structural identities of processes (Fig. 1). Reduction $P \longrightarrow Q$ is defined by the rules in Fig. 2. Term communication is only done in value form, meaning that all terms are reduced to values before communication takes place (we omit the reduction rules for terms and the compatible closure rules for conciseness). We note the standard unfolding semantics for corecursive definitions. Channel communication is always fresh. In the epilogue, we write \Rightarrow for the reflexive transitive closure of \longrightarrow and $P \Downarrow$ iff P is non-divergent (i.e. there is no infinite reduction sequence starting with P).

$$\begin{array}{l}
x \notin \text{fn}(P) \Rightarrow P \mid (\nu x)Q \equiv (\nu x)(P \mid Q) \quad P \equiv_\alpha Q \Rightarrow P \equiv Q \\
P \mid (Q \mid R) \equiv (P \mid Q) \mid R \quad P \mid Q \equiv Q \mid P \\
(\nu x)(\nu y)P \equiv (\nu y)(\nu x)P \quad [y \leftrightarrow x] \equiv [x \leftrightarrow y] \\
P \mid \mathbf{0} \equiv P \quad (\nu x)\mathbf{0} \equiv \mathbf{0}
\end{array}$$

Fig. 1. Structural Congruence

$$\begin{array}{l}
c\langle V \rangle.P \mid c(x).Q \longrightarrow P \mid Q\{V/x\} \\
c\langle y \rangle.P \mid c(x).Q \longrightarrow P \mid Q\{y/x\} \\
c\langle y \rangle.P \mid !c(x).Q \longrightarrow P \mid Q\{y/x\} \mid !c(x).Q \\
c.\text{case}(\bar{l}_i \Rightarrow \bar{P}_i) \mid c.l_i; Q \longrightarrow P_i \mid Q \\
(\nu x)(P_y \mid [y \leftrightarrow x]) \longrightarrow P\{y/x\} \quad (x \neq y) \\
(\text{corec } X(\bar{y}).P) \bar{c} \longrightarrow P\{\bar{c}/\bar{y}\} \{(\text{corec } X(\bar{y}).P)/X\} \\
\text{If } Q \longrightarrow Q' \text{ then } P \mid Q \longrightarrow P \mid Q' \\
\text{If } P \longrightarrow Q \text{ then } (\nu y)P \longrightarrow (\nu y)Q \\
\text{If } P \equiv P' \text{ and } P' \longrightarrow Q' \text{ and } Q' \equiv Q \text{ then } P \longrightarrow Q
\end{array}$$

Fig. 2. Reduction

The labelled transition system is defined by the rules of Fig. 3, consisting of a typical early transition system with internal actions τ , output $x\langle y \rangle$ and its dual input $x(z)$ action, bound output $(\nu y)x\langle y \rangle$, label selection $x.l$ and choice $x.l$.

III. TYPE SYSTEM

In this section we motivate and present our type system based on intuitionistic linear logic. The syntax of types is given in Fig. 4. We distinguish the types of basic data τ, σ from session types A, B, C .

The language of session types covers the standard session constructs: data input and output ($\tau \supset A$ and $\tau \wedge A$), (fresh) session input and output ($A \multimap B$ and $A \otimes B$), termination ($\mathbf{1}$), labelled choice and selection ($\&\{l_j:A_j\}$ and $\oplus\{l_j:A_j\}$), replication ($!A$) and coinductive session behavior ($\nu X.A$).

A. From General Recursion to Corecursion

While both recursive and coinductive session types denote (potentially) infinite session behavior, the fundamental distinction is that using general recursion a process might generate an infinite sequence of *internal* actions, whereas a valid coinductive definition of a session typed process is guaranteed to always have a finite sequence of internal actions before offering some observable behavior. It is this *external*, observable behavior that may be infinite in a coinductively defined, session typed process. This is of paramount importance since it ensures that a well-typed process offering a coinductive session type will always be able to truly offer its specified service in a divergence-free way, even if the type specifies an infinite sequence of observable actions. Moreover, this non-divergence result is *compositional*: well-typed coinductive sessions may be safely composed, ensuring that the resulting system is itself non-divergent.

To rule out divergence we must impose a discipline on the occurrence of the recursion variable in processes, in line with the work on coinduction in implementations of dependent type theories such as Coq [10] or Agda [11], but here mapped to a

concurrent process setting. Essentially, we must ensure that a corecursive process definition is *productive* – there is always a finite sequence of internal actions between *observable* actions.

In our setting, observable actions are those that take place on the session channel that is being *offered* by a given (well-typed) process, whereas internal actions are those generated through interactions with ambient sessions. Given our definition of productivity, a natural restriction is to require an action on the offered channel before allowing recursive calls to take place (i.e. recursive calls must be guarded by an observable action). For instance, the process $\text{corec } P.c\langle 0 \rangle.P$, is guarded if we assume that c is the session channel that is being offered. However, as known from type theory, guardedness alone is not sufficient to ensure productivity. Process Loop from Section I is guarded but produces divergent behavior when composed with a process that provides it with an output. The problem with the definition of Loop is that after the input along c , we have an occurrence of the recursion variable in parallel with a process that *interacts* with the recursive occurrence locally, essentially destroying the productivity of the original definition.

Thus, ensuring non-divergence in a compositional way requires not only ensuring guardedness but also disallowing interactions with corecursive calls within corecursive definitions, a property we call *co-regular recursion*. In functional type theories, such restrictions typically consist of disallowing pattern matching over the result of recursive calls, or passing such results as arguments to other functions. In our process setting, interactions are generated by communication and so we must impose that processes that are placed in parallel with the corecursive call may not communicate with it, although they may themselves perform other actions. As we discuss at the end of this section, our type system ensures this form of non-interference by not exposing the communication interface of corecursive calls internally, which ensures that processes composed with corecursive calls may perform communication, but *not* with the corecursive call itself.

B. Typing Coinductive Sessions

Having made the informal case for the kinds of restrictions on general recursion that are needed in order to eliminate divergent computation, we now present the type system for our language, essentially made up of the rules of linear logic plus the rules that pertain to corecursive process definitions, coinductive session types and the corecursion variable.

We define the typing judgment: $\Psi; \Gamma; \Delta \vdash_\eta P :: z:A$ denoting that process P offers the session behavior typed with A along channel z , when composed with the (linear) session behaviors specified in Δ , with the (unrestricted, or shared) session behaviors specified in Γ and where η is a mapping from (corecursive) type variables to typing contexts (we detail this further below). We note that the names in Γ, Δ and z are all pairwise distinct. We assume typing to be defined modulo structural congruence by definition. The context Ψ tracks the free variables in P that pertain to basic data values that are to be sent and received. We make use of a second judgment

$$\begin{array}{c}
\frac{P \xrightarrow{\alpha} Q}{P \mid R \xrightarrow{\alpha} Q \mid R} \text{ par} \quad \frac{P \xrightarrow{\bar{\alpha}} P' \quad Q \xrightarrow{\alpha} Q'}{P \mid Q \xrightarrow{\tau} P' \mid Q'} \text{ com} \quad \frac{P \xrightarrow{\alpha} Q}{(\nu y)P \xrightarrow{\alpha} (\nu y)Q} \text{ res} \quad \frac{P \xrightarrow{x(y)} Q}{(\nu y)P \xrightarrow{(\nu y)x(y)} Q} \text{ open} \\
\frac{P \xrightarrow{(\nu y)x(y)} P' \quad Q \xrightarrow{x(y)} Q'}{P \mid Q \xrightarrow{\tau} (\nu y)(P' \mid Q')} \text{ close} \quad !u(x).P \xrightarrow{u(z)} P\{z/x\} \mid !u(x).P \quad x(y).P \xrightarrow{x(y)} P \quad x(y).P \xrightarrow{x(z)} P\{z/y\} \\
(\nu x)([x \leftrightarrow y] \mid P) \xrightarrow{\tau} P\{y/x\} \quad x.l; P \xrightarrow{x.l} P \quad x.\text{case}(\bar{l}_i \Rightarrow \bar{P}_i) \xrightarrow{x.l_i} P_i \quad (\text{corec } X(\bar{y}).P) \bar{c} \xrightarrow{\tau} P\{\bar{c}/\bar{y}\} \{(\text{corec } X(\bar{y}).P)/X\}
\end{array}$$

Fig. 3. Labelled Transition System

τ, σ	::=	nat string ...	(basic data types)
A, B, C	::=	$\tau \supset A$	input value of type τ and continue as A
		$\tau \wedge A$	output value of type τ and continue as A
		$A \multimap B$	input channel of type A and continue as B
		$A \otimes B$	output fresh channel of type A and continue as B
		1	terminate
		$\&\{l_j:A_j\}$	offer choice between l_j and continue as A_j
		$\oplus\{l_j:A_j\}$	provide one of the l_j and continue as A_j
		$!A$	provide replicable service A
		$\nu X.A \mid X$	coinductive session type

Fig. 4. The Syntax of Types

$\Psi \vdash M:\tau$ to denote that term M , denoting a value that is to be communicated, is well typed under the assumptions in Ψ .

The rules that define our type system are given in Fig. 5, consisting essentially of those of [4] with the identity rule, value input and output (present in [9]) and coinductive types, which are associated with corecursion in the process calculus. We note that coinductive types have strictly positive occurrences of the type variable, also excluding coinductive types that have no associated session behavior before the type variable occurrence (such as $\nu X.X$). Moreover, we require coinductive types to mention the type variable.

We refrain from a detailed presentation of every rule for the sake of conciseness, highlighting instead the rules pertaining to corecursive processes and coinductive session types. We begin with the right rule for coinductive sessions, which types (parameterized) corecursive process definitions:

$$\frac{\Psi; \Gamma; \Delta \vdash_{\eta'} P :: c:A \quad \eta' = \eta[X(\bar{y}) \mapsto \Psi; \Gamma; \Delta \vdash c:Y]}{\Psi; \Gamma; \Delta \vdash_{\eta} (\text{corec } X(\bar{y}).P\{\bar{y}/\bar{z}\}) \bar{z} :: c:\nu Y.A} (\nu R)$$

In the rule above, the process P may use the recursion variable X and refer to the parameter list \bar{y} , which is instantiated with the list of (distinct) names \bar{z} which may occur in Ψ , Δ , Γ or c . Moreover, we keep track of the contexts Ψ , Γ and Δ in which the corecursive definition is made, as well as the channel name along which the coinductive behavior is offered and the type variable associated with the corecursive behavior, by extending the mapping η with a binding for X with the appropriate information. This is necessary because, intuitively, each occurrence of the corecursion variable stands for P itself (modulo the parameter instantiations) and therefore we must check that the necessary ambient session behaviors

are available for P to execute in a type correct way, respecting linearity. P itself simply offers along channel c the session behavior A (which is an open type). To type the corecursion variable we use the following rule:

$$\frac{\eta(X(\bar{y})) = \Psi; \Gamma; \Delta \vdash d:Y \quad \rho = \{\bar{z}/\bar{y}\}}{\rho(\Psi); \rho(\Gamma); \rho(\Delta) \vdash_{\eta} X(\bar{z}) :: \rho(d):Y} (\text{VAR})$$

We type a process corecursion variable X by looking up in η the binding for X , which references the typing environments Ψ , Γ and Δ under which the corecursive definition is well defined, the coinductive type variable Y associated with the corecursive behavior and the channel name d along which the behavior is offered. The corecursion variable X is typed with the type variable Y if the parameter instantiation is able to satisfy the typing signature (by renaming available linear and exponential resources or term variables). We also allow for the offered session channel to be a parameter of the corecursion. Finally, the left rule for coinductive session types simply unfolds the type accordingly:

$$\frac{\Psi; \Gamma; \Delta, c:A\{\nu X.A/X\} \vdash_{\eta} Q :: d:D}{\Psi; \Gamma; \Delta, c:\nu X.A \vdash_{\eta} Q :: d:D} (\nu L)$$

While the rules look fairly straightforward, they turn out to introduce quite subtle restrictions on what constitutes a well-formed (i.e. well-typed) corecursive process definition. First, observe that the introduction form for corecursive definitions does not directly unfold the coinductive type in its premise and thus references an *open* type, which means that the (corecursive) definition of P provides the behavior specified in A up to occurrences of the coinductive type variable Y . On the other hand, using a coinductive session type (which is

$$\begin{array}{c}
\begin{array}{ccc}
(\wedge R) & (\wedge L) & (\supset R) \\
\frac{\Psi \vdash M : \tau \quad \Psi; \Gamma; \Delta \vdash_{\eta} P :: c : A}{\Psi; \Gamma; \Delta \vdash_{\eta} c \langle M \rangle . P :: c : \tau \wedge A} & \frac{\Psi, x : \tau; \Gamma; \Delta, c : A \vdash_{\eta} Q :: d : D}{\Psi; \Gamma; \Delta, c : \tau \wedge A \vdash_{\eta} c \langle x \rangle . Q :: d : D} & \frac{\Psi, x : \tau; \Gamma; \Delta \vdash_{\eta} P :: c : A}{\Psi; \Gamma; \Delta \vdash_{\eta} c \langle x \rangle . P :: c : \tau \supset A} \\
(\supset L) & (ID) & (1R) \\
\frac{\Psi \vdash M : \tau \quad \Psi; \Gamma; \Delta, c : A \vdash_{\eta} Q :: d : D}{\Psi; \Gamma; \Delta, c : \tau \supset A \vdash_{\eta} c \langle M \rangle . Q :: d : D} & \frac{}{\Psi; \Gamma; d : A \vdash_{\eta} [d \leftrightarrow c] :: c : A} & \frac{}{\Psi; \Gamma; \cdot \vdash_{\eta} \mathbf{0} :: c : \mathbf{1}} \\
(1L) & & \\
\frac{}{\Psi; \Gamma; \Delta, c : \mathbf{1} \vdash_{\eta} P :: d : D} & & \\
(!R) & (!L) & (COPY) \\
\frac{}{\Psi; \Gamma; \cdot \vdash_{\eta} !c \langle x \rangle . P :: c : !A} & \frac{}{\Psi; \Gamma, u : A; \Delta \vdash_{\eta} P :: d : D} & \frac{}{\Psi; \Gamma, u : A; \Delta, x : A \vdash_{\eta} P :: d : D} \\
\frac{}{\Psi; \Gamma; \Delta, c : !A \vdash_{\eta} P \{c/u\} :: d : D} & & \\
(\otimes R) & & (\otimes L) \\
\frac{\Psi; \Gamma; \Delta_1 \vdash_{\eta} P_1 :: x : A \quad \Psi; \Gamma; \Delta_2 \vdash_{\eta} P_2 :: c : B}{\Psi; \Gamma; \Delta_1, \Delta_2 \vdash_{\eta} (\nu x) c \langle x \rangle . (P_1 \mid P_2) :: c : A \otimes B} & & \frac{\Psi; \Gamma; \Delta, x : A, c : B \vdash_{\eta} Q :: d : D}{\Psi; \Gamma; \Delta, c : A \otimes B \vdash_{\eta} c \langle x \rangle . Q :: d : D} \\
(\multimap R) & (\multimap L) & \\
\frac{\Psi; \Gamma; \Delta, x : A \vdash_{\eta} P :: c : B}{\Psi; \Gamma; \Delta \vdash_{\eta} c \langle x \rangle . P :: c : A \multimap B} & \frac{\Psi; \Gamma; \Delta_1 \vdash_{\eta} Q_1 :: x : A \quad \Psi; \Gamma; \Delta_2, c : B \vdash_{\eta} Q_2 :: d : D}{\Psi; \Gamma; \Delta_1, \Delta_2, c : A \multimap B \vdash_{\eta} (\nu x) c \langle x \rangle . (Q_1 \mid Q_2) :: d : D} & \\
(\& R) & & (\& L) \\
\frac{\Psi; \Gamma; \Delta \vdash_{\eta} P_1 :: c : A_1 \quad \dots \quad \Psi; \Gamma; \Delta \vdash_{\eta} P_n :: c : A_n}{\Psi; \Gamma; \Delta \vdash_{\eta} c . \text{case}(\bar{l}_j \Rightarrow P_j) :: c : \& \{l_j : A_j\}} & & \frac{\Psi; \Gamma; \Delta, c : A_i \vdash_{\eta} Q :: d : D}{\Psi; \Gamma; \Delta, c : \& \{l_j : A_j\} \vdash_{\eta} c . l_i; Q :: d : D} \\
(\oplus R) & & (\oplus L) \\
\frac{}{\Psi; \Gamma; \Delta \vdash_{\eta} P :: c : A_i} & & \frac{\Psi; \Gamma; \Delta, c : A_1 \vdash_{\eta} Q_1 :: d : D \quad \dots \quad \Psi; \Gamma; \Delta, c : A_n \vdash_{\eta} Q_n :: d : D}{\Psi; \Gamma; \Delta, c : \oplus \{l_j : A_j\} \vdash_{\eta} c . \text{case}(\bar{l}_j \Rightarrow Q_j) :: d : D} \\
(\nu L) & (\nu R) & (\text{VAR}) \\
\frac{\Psi; \Gamma; \Delta, c : A \{ \nu X . A / X \} \vdash_{\eta} Q :: d : D}{\Psi; \Gamma; \Delta, c : \nu X . A \vdash_{\eta} Q :: d : D} & \frac{\Psi; \Gamma; \Delta \vdash_{\eta'} P :: c : A \quad \eta' = \eta[X(\bar{y}) \mapsto \Psi; \Gamma; \Delta \vdash c : Y]}{\Psi; \Gamma; \Delta \vdash_{\eta} (\text{corec } X(\bar{y}) . P \{ \bar{y} / \bar{z} \}) \bar{z} :: c : \nu Y . A} & \frac{\eta(X(\bar{y})) = \Psi; \Gamma; \Delta \vdash d : Y \quad \rho = \{ \bar{z} / \bar{y} \}}{\rho(\Psi); \rho(\Gamma); \rho(\Delta) \vdash_{\eta} X(\bar{z}) :: \rho(d) : Y} \\
(\text{CUT}) & & (\text{CUT}^\dagger) \\
\frac{\Psi; \Gamma; \Delta_1 \vdash_{\eta} P :: x : A \quad \Psi; \Gamma; \Delta_2, x : A \vdash_{\eta} Q :: d : D}{\Psi; \Gamma; \Delta_1, \Delta_2 \vdash_{\eta} (\nu x)(P \mid Q) :: d : D} & & \frac{\Psi; \Gamma; \cdot \vdash_{\eta} P :: x : A \quad \Psi; \Gamma, u : A; \Delta \vdash Q :: d : D}{\Psi; \Gamma; \Delta \vdash_{\eta} (\nu u)(!u \langle x \rangle . P \mid Q) :: d : D}
\end{array}
\end{array}$$

Fig. 5. Typing Rules

achieved by the left rule νL) entails unfolding the coinductive type as expected, and so a user of a coinductive behavior may use as many unfoldings of $\nu Y.A$ as required.

Up to this point, we have yet to discuss how our type system enforces the necessary restrictions to ensure non-divergence. It turns out that these restrictions are imposed by the variable rule in quite subtle ways, due to its interaction with the νR rule. First, observe that we enable multiple occurrences of the recursion variable in a well-typed process, for instance in different branches of a case construct, through cuts or the $\otimes R$ rule (e.g. Fig 6). Moreover, we do not syntactically exclude processes without terminal recursion (for instance, process Good from Section I).

However, a well-typed corecursive definition *cannot* interact with its corecursive calls (which could potentially destroy productivity). To see why this is the case, consider how such an interaction might be allowed in our type system: in order for a corecursive definition to interact with its own corecursive call in a well-typed manner, the system would require some form of parallel composition where we obtain a handle to the corecursive call. This would only be possible through a cut, given that the right rule for \otimes ensures that the two parallel

processes cannot interact. However, since the coinductive type is never unfolded when *offering* a coinductive definition, the session interface of the corecursive occurrence is not visible internally. Thus, a cut of the corecursion variable with some other process Q will generate a fresh channel c that offers the coinductive type variable, but not the *coinductive type*, meaning that no left rules that interact with the unfolding of the recursive definition can be applied. In fact, the only rule that can use $c : X$ is the identity rule, which will forward the corecursively defined session. This does not prohibit Q from having additional behavior besides forwarding, it simply excludes (potentially) problematic internal interactions with corecursive calls.

IV. EXAMPLES

In this section we discuss three examples: the Loop process and the hypothetical Twitter web service, combining replication and recursion, from Section I, and a coordinated process network implementing a binary numeral.

A. Excluding Unobservable Divergence

Elaborating on the process Loop given in Section I, we show how it can result in divergent behavior when interacting with a

$$\begin{array}{c}
\frac{}{\vdash X(y) :: y:Y} \text{ (var)} \quad \frac{}{\vdash X(z) :: z:Y} \text{ (var)} \\
\frac{}{\vdash 1 : \text{nat}} \quad \frac{}{\vdash (\nu y)z\langle y \rangle.(X(y) \mid X(z)) :: z:Y \otimes Y} \text{ (\otimes R)} \\
\frac{}{\vdash z\langle 1 \rangle.(\nu y)z\langle y \rangle.(X(y) \mid X(z)) :: z:\text{nat} \wedge (Y \otimes Y)} \text{ (\wedge R)} \\
\frac{}{\vdash (\text{corec } X(z).z\langle 1 \rangle.(\nu y)z\langle y \rangle.(X(y) \mid X(z))) z :: z:\nu Y.\text{nat} \wedge (Y \otimes Y)} \text{ (\nu R)}
\end{array}$$

Fig. 6. Example Typing Derivation

client. We also show why Loop is not typeable in our system. Recall the definition of Loop:

$$\text{Loop} \triangleq \text{corec } L(c).c(x).(\nu d)(L(d) \mid d\langle n \rangle.[d \leftrightarrow c])$$

The intended behavior of Loop is to continuously input along the channel c , which is represented by the session type $\nu X.\text{nat} \supset X$. Consider a process P that provides an output along c , composed with Loop. We have the following reduction(s):

$$(\nu c)(\text{Loop} \mid P) \Rightarrow (\nu c)(\nu d)(\text{Loop}\{c/d\} \mid d\langle n \rangle.[d \leftrightarrow c] \mid P')$$

In the process to the right of the arrow, there is an internal synchronization between the output along d and the input along the unfolding of Loop, resulting in the following process:

$$(\nu c)(\nu d)((\nu d')(\text{Loop}\{d'/c\} \mid d'\langle n \rangle.[d' \leftrightarrow d]) \mid [d \leftrightarrow c] \mid P')$$

The infinite internal reduction is now made clear, where regardless of the behavior of P' we always have available an internal reduction that produces an additional unfolding of Loop and may repeat this process an infinite number of times.

It is easy to see that Loop is not well-typed in our system: if we try to type the process bottom-up, we first apply the νR rule, followed by the $\supset R$ rule. We are then left with trying to type the following cut:

$$\frac{\vdash L(d) :: d:X \quad d:X \vdash d\langle n \rangle.[d \leftrightarrow c] :: c:X}{\vdash (\nu d)(L(d) \mid d\langle n \rangle.[d \leftrightarrow c]) :: c:X}$$

It is then obvious that the right premise of the cut cannot be typed, since all that is known about session d is that it has type X , so no communication along d can be well-typed. A type-correct version of the behavior $c:\nu X.\text{nat} \supset X$ is:

$$\text{Ins} \triangleq \text{corec } L(c).c(x).L(c)$$

Note however that a process that is meant to interact with Ins by just outputting continuously along c is *not* well-typed in our system. To use the behavior of Ins in this way, we must necessarily define a corecursive process, and thus a process that offers a coinductive session type. However, the continuous user of Ins that just outputs along c offers no observable behavior and thus no valid coinductive session type can be assigned to it (recall that valid coinductive types need to be strictly positive and non-trivial). We exclude such processes from our system since the composition of such a continuous user process with Ins would result in a process with no observable actions, but with infinite internal actions:

$$(\nu c)(\text{Ins} \mid \text{corec } S.c\langle 0 \rangle.S) \Rightarrow (\nu c)(\text{Ins} \mid \text{corec } S.c\langle 0 \rangle.S) \Rightarrow \dots$$

Morally, such a user process must offer some observable behavior to be useful (even a divergent process should be receptive to a “kill” signal). For instance, the process $P \equiv \text{corec } S'.c\langle 0 \rangle.d\langle 0 \rangle.S'$ is a well-typed process offering behavior $d:\nu X.\text{nat} \wedge X$ in a context where Ins is available, and so the composition of Ins and P is well-typed:

$$\vdash (\nu c)(\text{Ins} \mid P) :: d:\nu X.\text{nat} \wedge X$$

B. Replication and Corecursion

A web service such as Twitter offers a persistent (and thus replicated) service *trends* which is intended to produce a stream of current trends, according to some custom metrics. To do this, the service is parametrized by a filter process (of type Tweets \multimap Trends) that given a stream of tweets produces a stream of trends. Such a protocol is captured via the following type(s):

$$\begin{array}{lcl}
\text{TrendService} & \triangleq & !((\text{Tweets} \multimap \text{Trends}) \multimap \text{Trends}) \\
\text{Tweets} & \triangleq & \nu X.(\text{tweet} \wedge X) \\
\text{Trends} & \triangleq & \nu Y.(\text{trend} \wedge Y)
\end{array}$$

The use of the exponential $!$ in TrendService ensures that each invocation of the service provides access to a fresh, independent session. Any such session expects to receive a session of type Tweets \multimap Trends able to transduce a stream of tweets to a stream of trends. This is the usual way to represent higher-order communication (e.g., sending an “object” of type Tweets \multimap Trends) in a communication) in the π -calculus. The service will use this session to finally provide the stream of trends, after connecting it to a (private) source of tweets TweetSrc $_y$. The code implementing the service is:

$$\text{Serv} \triangleq !\text{trends}(x).x(f).(\nu y)f\langle y \rangle.(\text{TweetSrc}_y \mid [f \leftrightarrow x])$$

Given that $\vdash \text{TweetSrc}_y :: y:\text{Tweets}$, we can derive the typing $\vdash \text{Serv} :: \text{trends}:\text{TrendService}$.

Possible client code for such a service is:

$$\text{Cl} \triangleq (\nu x)\text{trends}\langle x \rangle.(\nu k)x\langle k \rangle.(\text{Analytics}_k \mid [x \leftrightarrow tr])$$

This process interacts by first invoking the (replicated) trends service, and afterwards interacting with the freshly generated replica (at x). It first provides access to the appropriate analytics process, and afterwards forwards the resulting trend stream (from k) on channel tr . Assuming $\vdash \text{Analytics}_k :: k:\text{Tweets} \multimap \text{Trends}$ we derive

$$\text{trends} : \text{TrendService} \vdash \text{Cl} :: tr:\text{Trends}$$

C. Little Endian Bit Counter

We illustrate how our framework can express fairly general process networks with nodes interacting according to structured session protocols. We consider the implementation of a binary counter, where each bit of a (arbitrary length) binary numeral is implemented by a process node which can only communicate with its neighbouring processes in the network (in [9] we discussed a similar example, but expressed using non co-regular recursion, and not typeable in our system. We present here a version using co-regular recursion, and requiring more sophisticated handshaking). Overall, the network implements a protocol offering three operations: poll the counter for its current integer value; increment the counter value, or terminate the counter. The corresponding session type is expressed coinductively, as we do not want to limit a priori the counter size:

$$\text{Counter} \triangleq \nu X. \&\{\text{val}:\text{int} \wedge X, \text{inc}:X, \text{halt}:1\}$$

Our implementation of Counter is based on a coordinator process that keeps a (linear) network of communicating processes, representing the counter value in little endian form (the tail process in the network holds the least significant bit). Each network node communicates with its two adjacent bit representations, whereas the coordinator communicates with the most significant bit process (and with the counter's external client), spawning new bits as needed. Overall, the coordinator works as follows. To halt the counter, the coordinator halts every node before halting itself. To provide the value of the counter, the coordinator propagates a val message along the network, which will compute the value as an integer. To increment the counter, the coordinator injects an inc message into the network, which will be propagated to the least significant bit process, and, in a second phase, propagated back as a carry message travelling back in the opposite direction, incrementing each bit (modulo 2) as needed. If the carry message, instead of a done message, reaches the coordinator, a new bit process will be dynamically spawned.

The behavior of each node is as follows. When a node receives a val message it receives the integer value computed by all the nodes encoding more significant bits, updates it with its own contribution and propagates the val message to the less significant bit nodes. When it gets a inc message it forwards it to its less significant bit neighbour, and waits for it to send back either a carry or done message. In the latter case, it will just forward the done message network along the most significant bit node up to the coordinator, signalling that nothing more needs to be done. In the former case, it will flip its value and either send a carry or a done message to its most significant bit neighbour. When a carry message reaches the coordinator, it generates a new bit, as mentioned above.

So, the type for each node $\text{Node}(b, x, n)$ will be given by the judgment $x : \text{Clmpl} \vdash \text{Node}(b, x, n) :: n:\text{Clmpl}$ where

$$\text{Clmpl} \triangleq \nu X. \&\{\text{val}:\text{int} \supset \text{int} \wedge X, \text{inc}:\oplus\{\text{carry}:X, \text{done}:X\}, \text{halt}:1\}$$

In $\text{Node}(b, x, n)$, b holds the bit value, x is the session channel connecting to the less significant node (or to the terminal node) and n is the channel connecting to the most significant node (or to the coordinator). Code for $\text{Node}(b, x, n)$ is as follows:

$$\begin{aligned} \text{Node}(b, x, n) &\triangleq \text{corec } X(b, x, n).n.\text{case}(\text{val} \Rightarrow x.\text{val}; n(m).x\langle(2 * m + b)\rangle. \\ &\quad X(b, x, n), \\ \text{inc} &\Rightarrow x.\text{inc}; \\ &\quad x.\text{case}(\text{carry} \Rightarrow \\ &\quad \text{if } (b = 1) \text{ then } n.\text{carry}; X(0, x, n) \\ &\quad \text{else } n.\text{done}; X(1, x, n), \\ \text{done} &\Rightarrow X(1, x, n), \\ \text{halt} &\Rightarrow x.\text{halt}; \mathbf{0})(b, x, n) \end{aligned}$$

The coordinator process code interfaces with clients and wraps the bit process network, generating new bit nodes as needed.

$$\begin{aligned} \text{Coord}(x, z) &\triangleq \text{corec } X(x, z). \\ &\quad z.\text{case}(\text{val} \Rightarrow x.\text{val}; x\langle 0 \rangle. \\ &\quad \quad x(n).z\langle n \rangle.X(x, z), \\ \text{inc} &\Rightarrow x.\text{inc}; \\ &\quad \quad x.\text{case}(\text{carry} \Rightarrow (\nu n')(\text{Node}(1, x, n') \mid \\ &\quad \quad \quad \text{Coord}(n', z)), \\ &\quad \quad \text{done} \Rightarrow X(x, z)) \\ \text{halt} &\Rightarrow x.\text{halt}; \mathbf{0})(x, z) \end{aligned}$$

To complete the system we provide the implementation of the empty bit string epsilon, which will be a closed process of type Clmpl , expressing the “base case” of the protocol. For reading the value, it ping pongs the received value. For incrementing, it emits the (first) carry message. The Counter system is then produced by composing epsilon and Coord.

$$\begin{aligned} \text{epsilon}(x) &\triangleq \text{corec } X(x).x.\text{case}(\text{val} \Rightarrow x(n).x\langle n \rangle.X(x), \\ &\quad \text{inc} \Rightarrow x.\text{carry}; X(x), \\ &\quad \text{halt} \Rightarrow \mathbf{0})x \\ \text{Counter}(c) &\triangleq (\nu e)(\text{epsilon}(e) \mid \text{Coord}(e, c)) \end{aligned}$$

The several relevant typings are

$$\begin{aligned} e : \text{Clmpl} &\vdash \text{Coord}(e, c) :: c:\text{Counter} \\ &\vdash \text{epsilon}(e) :: e:\text{Clmpl} \\ &\vdash \text{Counter}(e) :: c:\text{Counter} \end{aligned}$$

The two examples above showcase the interaction between corecursive session types, which denote potentially infinite interactions (such as streams) on one session channel; and replication, which denotes the potential for an arbitrary number of copies of a given session, but on different channels.

V. MAIN RESULTS

In this section, we establish type safety for our calculus, entailing session fidelity and deadlock freedom. We also develop our main result of compositional non-divergence by extending the linear logical relations of [8], [12] to the coinductive setting, fully restoring the connection of our framework with the logical interpretation.

A. Type Safety

Following [4], our proof of type preservation relies on a simulation between reductions in the session-typed π -calculus and proof reductions from logic.

Theorem 5.1 (Type Preservation): If $\Psi; \Gamma; \Delta \vdash_{\eta} P :: z:A$ and $P \rightarrow Q$ then $\Psi; \Gamma; \Delta \vdash_{\eta} Q :: z:A$.

The proof of progress also follows the lines of [4], but with some additional caveats due to the presence of corecursive definitions. The key technical aspects of the proof are a series of inversion lemmas and a notion of a *live* process, consisting of a process that has not yet fully carried out its ascribed session behavior, and thus is a parallel composition of processes where at least one is a non-replicated process, guarded by some action. We define $live(P)$ if and only if $P \equiv (\nu \tilde{n})(\pi.Q \mid R)$, for some process R , sequence of names \tilde{n} and a non-replicated guarded process $\pi.Q$.

Theorem 5.2 (Progress): If $\Psi; \Gamma; \Delta \vdash P :: z:A$ and $live(P)$ then there is Q with $P \rightarrow Q$ or one of the following holds:

- (a) $\Delta = \Delta', y:B$, for some Δ' and $y:B$. There exists $\Gamma; \Delta'' \vdash R :: y:B$ s.t. $(\nu y)(R \mid P) \rightarrow Q$.
- (b) Exists $\Gamma; z:A, \Delta' \vdash R :: w:C$ s.t. $(\nu z)(P \mid R) \rightarrow Q$.
- (c) $\Gamma = \Gamma', u:B$, for some Γ' and $u:B$. There exists $\Gamma; \cdot \vdash R :: x:B$ s.t. $(\nu u)(!u(x).R \mid P) \rightarrow Q$.

Our notion of progress states that well-typed processes never get stuck even in the presence of infinite session behavior. Either the process progresses outright via internal computation, awaits on an interaction with its environment ((a) or (c)) or is waiting for an interaction along its offered channel ((b)).

B. Compositional Non-Divergence

To prove the key property of compositional non-divergence, we develop a (linear) logical relations argument for coinductive session types. As in prior work for languages without recursion [12], [8], we build on Girard’s technique of reducibility candidates: a set of non-divergent, well-typed terms, closed under reduction and expansion (backward reduction), adapted to the setting of our process calculus.

Definition 5.3 (Reducibility Candidates): Given a type A and a name z , a reducibility candidate at $z:A$, written $\mathcal{R}[z:A]$ is a set of processes satisfying the following properties:

- 1) If $P \in \mathcal{R}[z:A]$ then $\cdot; \cdot \vdash_{\eta} P :: z:A$.
- 2) If $P \in \mathcal{R}[z:A]$ then $P \Downarrow$.
- 3) If $P \in \mathcal{R}[z:A]$ and $P \Rightarrow P'$ then $P' \in \mathcal{R}[z:A]$
- 4) If for all P_i such that $P \Rightarrow P_i$ we have $P_i \in \mathcal{R}[z:A]$ then $P \in \mathcal{R}[z:A]$.

We refer to $\mathbb{R}[-:A]$ as the collection of all sets of reducibility candidates at (closed) type A . Our definition of the logical predicate identifies processes up to the compatible extension of structural congruence with the well-known *sharpened replication axioms* [13], written \equiv_1 . The replication axioms express strong behavioral equivalences in our typed setting [12].

Definition 5.4: We write \equiv_1 for the least congruence relation on process expressions resulting from extending structural congruence \equiv with the following axioms:

- 1) $(\nu u)(!u(z).P \mid (\nu y)(Q \mid R)) \equiv_1$

- 2) $(\nu y)((\nu u)(!u(z).P \mid Q) \mid (\nu u)(!u(z).P \mid R))$
 $\equiv_1 (\nu u)(!u(y).P \mid (\nu v)(!v(z).Q \mid R))$
 $\equiv_1 (\nu v)((!v(z).(\nu u)(!u(y).P \mid Q)) \mid (\nu u)(!u(y).P \mid R))$
- 3) $(\nu u)(!u(y).Q \mid P) \equiv_1 P$ if $u \notin fn(P)$

Intuitively, axioms (1) and (2) represent the distribution of shared servers among “client” processes, and (3) garbage collects shared servers which can no longer be invoked.

1) Logical Predicate: We define a logical predicate on processes by induction on types and the size of typing contexts in the case of processes. The predicate captures the computational behavior of non-divergent processes, as defined by their typing. In the development below, we make use of $\mathcal{L}[\tau]$, which denotes the logical interpretation of the values exchanged in communication (i.e. sets for well-typed values of the appropriate type). We omit this definition due to its simplicity and for the sake of conciseness.

Definition 5.5 (Logical Predicate - Open Process Expressions): Given $\Psi; \Gamma; \Delta \vdash_{\eta} T$ with a non-empty left hand side environment, we define $\mathcal{L}^{\omega}[\Psi; \Gamma; \Delta \vdash_{\eta} T]$, where ω is a mapping from type variables to reducibility candidates, as the set of processes inductively defined by the rules of Fig. 7.

The definition of the logical interpretation for open processes inductively composes the process with the appropriate witnesses in the logical interpretation at the types specified in the three contexts, following [12].

The key part of our development is the definition of the logical predicate for closed processes. Similar to the treatment of type variables in logical relations for polymorphic languages (c.f. [8]), we employ a mapping from type variables to candidates at a given type. More precisely, we map type variables to candidates at the appropriate coinductive type, representing the unfolding of coinductive types.

We interpret a coinductive session type $\nu X.A$ as the union of all reducibility candidates ψ of the appropriate coinductive type that are in the interpretation of the *open* type A , when X is mapped to ψ itself, enabling a principle of proof by coinduction.

Definition 5.6 (Logical Predicate - Closed Process Expressions): For any type $T = z:A$ we define $\mathcal{L}^{\omega}[T]$ as the set of all processes P such that $P \Downarrow$ and $\cdot; \cdot \vdash_{\eta} P :: T$ satisfying the rules of Fig. 8.

We elide several technical aspects and focus mainly on the intuitions behind the development. The key observation is that for *open types*, we may view our logical predicate as a mapping between sets of reducibility candidates, of which the interpretation for coinductive session type turns out to be a greatest fixpoint.

Definition 5.7: Let $\nu X.A$ be a strictly positive type. We define: $\phi_A(s) \triangleq \mathcal{L}^{\omega}[X \mapsto s][z:A]$

Theorem 5.8 (Greatest Fixpoint): $\mathcal{L}^{\omega}[z:\nu X.A]$ is a greatest fixpoint of ϕ_A .

Using Theorem 5.8 we can show an unfolding lemma for coinductive session types, relating type unfolding with the interpretation of open types.

Lemma 5.9 (Unfolding): $P \in \mathcal{L}^{\omega}[z:A\{\nu X.A/X\}]$ iff $P \in \mathcal{L}^{\omega}[X \mapsto \mathcal{L}^{\omega}[-:\nu X.A]][z:A]$

$$\begin{aligned}
P \in \mathcal{L}^\omega[\Psi, x:\tau; \Gamma; \Delta \vdash_\eta T] & \text{ iff } \forall M \in \mathcal{L}[\tau]. P\{M/x\} \in \mathcal{L}^\omega[\Psi; \Gamma; \Delta \vdash_\eta T] \\
P \in \mathcal{L}^\omega[\Gamma; \Delta, y:A \vdash_\eta T] & \text{ iff } \forall R \in \mathcal{L}^\omega[y:A]. (\nu y)(R \mid P) \in \mathcal{L}^\omega[\Gamma; \Delta \vdash_\eta T] \\
P \in \mathcal{L}^\omega[\Gamma, u:A; \Delta \vdash_\eta T] & \text{ iff } \forall R \in \mathcal{L}^\omega[y:A]. (\nu u)(!u(y).R \mid P) \in \mathcal{L}^\omega[\Gamma; \Delta \vdash_\eta T]
\end{aligned}$$

Fig. 7. Logical Predicate - Open Processes

$$\begin{aligned}
\mathcal{L}^\omega[z:\nu X.A] & \triangleq \bigcup \{ \Psi \in \mathbb{R}[-:\nu X.A] \mid \Psi \subseteq \mathcal{L}^\omega[X \mapsto \Psi][z:A] \} \\
\mathcal{L}^\omega[z:X] & \triangleq \omega(z)(X) \\
\mathcal{L}^\omega[z:\mathbf{1}] & \triangleq \{ P \mid \forall P'. (P \Longrightarrow P' \wedge P' \not\rightarrow) \Rightarrow P' \equiv \mathbf{0} \} \\
\mathcal{L}^\omega[z:A \multimap B] & \triangleq \{ P \mid \forall P'. (P \xrightarrow{z(y)} P') \Rightarrow \forall Q \in \mathcal{L}^\omega[y:A]. (\nu y)(P' \mid Q) \in \mathcal{L}^\omega[z:B] \} \\
\mathcal{L}^\omega[z:A \otimes B] & \triangleq \{ P \mid \forall P'. (P \xrightarrow{(\nu y)z(y)} P') \Rightarrow \\
& \quad \exists P_1, P_2. (P' \equiv P_1 \mid P_2 \wedge P_1 \in \mathcal{L}^\omega[y:A] \wedge P_2 \in \mathcal{L}^\omega[z:B]) \} \\
\mathcal{L}^\omega[z:!A] & \triangleq \{ P \mid \forall P'. (P \Longrightarrow P') \Rightarrow \exists P_1. (P' \equiv !z(y).P_1 \wedge P_1 \in \mathcal{L}^\omega[y:A]) \} \\
\mathcal{L}^\omega[z:\& \{ \overline{l_i \Rightarrow A_i} \}] & \triangleq \{ P \mid \bigwedge_i (\forall P'. (P \xrightarrow{z.l_i} P') \Rightarrow P' \in \mathcal{L}^\omega[z:A_i]) \} \\
\mathcal{L}^\omega[z:\oplus \{ \overline{l_i \Rightarrow A_i} \}] & \triangleq \{ P \mid \bigwedge_i (\forall P'. (P \xrightarrow{z.l_i} P') \Rightarrow P' \in \mathcal{L}^\omega[z:A_i]) \} \\
\mathcal{L}^\omega[z:\tau \wedge A] & \triangleq \{ P \mid \forall P'. (P \xrightarrow{z(M)} P') \Rightarrow (M \in \mathcal{L}[\tau] \wedge P' \in \mathcal{L}^\omega[z:A]) \} \\
\mathcal{L}^\omega[z:\tau \supset A] & \triangleq \{ P \mid \forall P', M. (M \in \mathcal{L}[\tau] \wedge P \xrightarrow{z(M)} P') \Rightarrow P' \in \mathcal{L}^\omega[z:A] \}
\end{aligned}$$

Fig. 8. Logical Predicate - Closed Processes

The combination of these results enables us to show the fundamental lemma: all well-typed processes are in the logical predicate, from which follows that all well-typed processes are compositionally non-divergent.

Theorem 5.10 (Fundamental Lemma): If $\Psi; \Gamma; \Delta \vdash_\eta P :: z:A$ then $P \subseteq \mathcal{L}[\Psi; \Gamma; \Delta \vdash_\eta z:A]$.

The proof proceeds by induction on typing. The case for rule νL follows from Lemma 5.9. The most interesting case is the one for the νR rule. Since the interpretation of coinductive types $\nu X.A$ is a greatest fixpoint (Theorem 5.8), we proceed by coinduction: we produce a set of processes $\mathcal{C}_{P'}$, containing the body of the corecursive definition P' closed under composition with representatives for the linear and exponential contexts, where the corecursion variable has been instantiated with an unfolding of the corecursive definition (let us refer to this process a P''). $\mathcal{C}_{P'}$ is also closed under reduction and expansion. We show that $\mathcal{C}_{P'}$ is a reducibility candidate at $\nu X.A$ and that $\mathcal{C}_{P'} \subseteq \mathcal{L}^\omega[X \mapsto \mathcal{C}_{P'}][z:A]$. This is a sufficient condition since $\mathcal{L}[c:\nu X.A]$ is the largest such set. We need essentially focus on P'' .

Showing that $\mathcal{C}_{P'}$ satisfies this condition relies crucially on type preservation, progress and the fact that type variables are ultimately offered by the unfolding of the corecursive definition. The crucial points are the occurrences of the corecursion variable in P' , which in P'' are instantiated with P' itself: if its a terminal occurrence the property is immediate, since the corecursion variable is typed with the type variable X , which is mapped to $\mathcal{C}_{P'}$, containing P' . If the corecursion variable occurs in the left branch of a cut, we know that the only possible use of the fresh channel (typed with the type variable) by the right branch of the cut is to eventually forward it, potentially after a number of internal reductions or observable actions as specified by the session A . When the forwarder is triggered, we must necessarily be at the type variable X and we can conclude since the resulting process is P' , in $\mathcal{C}_{P'}$.

We remark that if the corecursion variable is guarded in the left branch of a cut, these guards must be consumed by the right branch of the cut (this follows by the well-typedness and progress) through internal reductions.

Corollary 5.11 (Non-Divergence): If $\Psi; \Gamma; \Delta \vdash_\eta P :: z:A$ then $P \Downarrow$

Combining type safety (Theorem 5.1 and Theorem 5.2) and non-divergence (Corollary 5.11) we conclude that typing enables strong guarantees on programs written in our calculus. A well-typed program will always be able to fulfil its protocol: no divergence can take place, nor can any deadlocks occur. Moreover, these properties are *compositional*, ensuring that any (well-typed) composition of services produces a deadlock-free, non-divergent service.

VI. RELATED WORK AND CONCLUDING REMARKS

Forms of general recursion have been often introduced in session typed languages, but without much concern on how to conciliate unbounded behaviour with local termination (e.g. [2], [14]). From the perspective of more traditional work on session typed languages, [15] establishes a strong normalization result for action types, which are similar to session types, but without addressing recursive types. Still within the logic based approach to session typed programming, our work is related to [5], which develops a logical interpretation of session types using second-order classical linear logic and a polymorphic session-typed functional language similar to that of [16], but does not consider coinductive sessions. In prior work, we have studied strong normalisation for session types based on linear logic, including extensions to behavioral polymorphism [12], [8], however, this work is the first to address general (co)recursion in the framework.

Expressiveness relationships between replication and recursion in the context of π -calculi have been investigated (e.g. see [17]); in general such constructs are not reducible to each other, in particular in the presence of name scoping constructs.

In our context, they are fundamentally separated by the typing discipline, which reflect different behavioral usages.

The particular issues related to coinductive session types have to the best of our knowledge been overlooked in the literature, which typically deals with general recursive types and definitions [2], [3], [16] and so is inherently non-divergent. In this regard our work is closer to that on termination for λ -calculi with inductive and coinductive types [18]. For instance, [19] develops a strong normalization result for the second-order λ -calculus with explicit inductive and coinductive types using logical relations. The restrictions imposed by our type system to ensure non-divergence are related to those developed in [20], [21] for the Coq proof assistant, but with obvious distinctions given the very different settings. Our interpretation of coinductive types as greatest fixed points is also related to the work of Baelde [22]. The main differences are that his work uses classical linear logic and doesn't consider a proof term assignment. The former makes the system and logical relation techniques substantially different since they rely on orthogonality, whereas ours do not; the latter leads Baelde's work towards proof search related techniques, which is in sharp contrast to our work.

The work on type-based methods for ensuring termination has been generalized to include inductive and coinductive definitions in the style of Coq [23] (although the proof of strong normalization is substantially more complex due to the expressive power of CIC and type annotations). It is not clear how to adapt these type-based methods to our setting since type annotations are a measure of the size of "values", which do not immediately apply to processes. Similar difficulties arise when considering copatterns [24], which seem to be inherently tied to the term structure.

A. Concluding Remarks

We have developed a theory of coinductive definitions for synchronous π -calculus processes with corecursion, establishing a strong logical foundation based on intuitionistic linear logic which provides guarantees of deadlock freedom, session fidelity and, crucially, compositional non-divergence by typing, ensuring that services implemented in our calculus do not "get stuck", and never become unavailable through divergent behavior.

Our type system ensures termination by restricting the occurrences of the corecursion variable in corecursive process definitions. While the restrictions we impose are not particularly oppressive, they naturally still exclude processes that are non-divergent (for instance, the binary counter example of [9]), as often is the case in these settings given that productivity is undecidable in general. Having set forth a first significant benchmark, it is certainly a challenge for future work to find less restrictive or more general conditions for guaranteeing productivity in this setting, potentially adapting type-based termination techniques (viz. [23], [24]).

By establishing a logical foundation for coinductive session-typed processes, the work set forth in this paper is an important stepping stone towards the development of a dependent

type theory rich enough to allow us to express and reason about concurrent session-based concurrency. These concurrent programs are often coinductive in nature, and are typically studied using coinductive proof techniques. Since we establish typed processes as coinductive objects, we may be able to use processes as witnesses to coinductive proofs. A sound (wrt an extensional typed equivalence, c.f. [8], [12]) notion of definitional equality of corecursive processes is also part of future work.

REFERENCES

- [1] K. Honda, "Types for dyadic interaction," in *CONCUR'93*, 1993, pp. 509–523.
- [2] K. Honda, V. T. Vasconcelos, and M. Kubo, "Language primitives and type discipline for structured communication-based programming," in *ESOP*, 1998, pp. 122–138.
- [3] K. Honda, N. Yoshida, and M. Carbone, "Multiparty asynchronous session types," in *POPL*, 2008, pp. 273–284.
- [4] L. Caires and F. Pfenning, "Session types as intuitionistic linear propositions," in *CONCUR'10*, 2010, pp. 222–236.
- [5] P. Wadler, "Propositions as sessions," in *ICFP'12*, 2012, pp. 273–286.
- [6] B. Toninho, L. Caires, and F. Pfenning, "Dependent session types via intuitionistic linear type theory," in *PPDP'11*, 2011, pp. 161–172.
- [7] F. Pfenning, L. Caires, and B. Toninho, "Proof-carrying code in a session-typed process calculus," in *CPP*, 2011, pp. 21–36.
- [8] L. Caires, J. A. Pérez, F. Pfenning, and B. Toninho, "Behavioral polymorphism and parametricity in session-based communication," in *ESOP*, 2013, pp. 330–349.
- [9] B. Toninho, L. Caires, and F. Pfenning, "Higher-order processes, functions, and sessions: A monadic integration," in *ESOP'13*, 2013, pp. 350–369.
- [10] The Coq Development Team, *The Coq Proof Assistant Reference Manual*, 2013. [Online]. Available: <http://coq.inria.fr>
- [11] U. Norell, "Towards a practical programming language based on dependent type theory," Ph.D. dissertation, Chalmers University of Technology, SE-412 96 Göteborg, Sweden, September 2007.
- [12] J. A. Pérez, L. Caires, F. Pfenning, and B. Toninho, "Linear logical relations for session-based concurrency," in *ESOP'12*, 2012, pp. 539–558.
- [13] D. Sangiorgi and D. Walker, *The π -calculus: A Theory of Mobile Processes*. Cambridge University Press, 2001.
- [14] S. Gay and M. Hole, "Subtyping for Session Types in the Pi Calculus," *Acta Informatica*, vol. 42, no. 2-3, pp. 191–225, 2005.
- [15] N. Yoshida, M. Berger, and K. Honda, "Strong normalisation in the pi-calculus," *Inf. Comput.*, vol. 191, no. 2, pp. 145–202, 2004.
- [16] S. Gay and V. T. Vasconcelos, "Linear type theory for asynchronous session types," *J. Funct. Programming*, vol. 20, no. 1, pp. 19–50, 2010.
- [17] J. Aranda, C. D. Giusto, C. Palamidessi, and F. D. Valencia, "On recursion, replication and scope mechanisms in process calculi," in *Formal Methods for Components and Objects, 5th International Symposium, FMCO 2006*, ser. Lecture Notes in Computer Science, vol. 4709. Springer, 2006, pp. 185–206.
- [18] A. Abel, "Strong normalization and equi-(co)inductive types," in *Typed Lambda Calculi and Applications*, 2007, pp. 8–22.
- [19] N. P. Mendler, "Inductive types and type constraints in the second-order lambda calculus," *Annals of Pure and Applied Logic*, vol. 51, no. 12, pp. 159 – 172, 1991.
- [20] E. Giménez, "Codifying guarded definitions with recursive schemes," in *TYPES '94*, pp. 39–59.
- [21] E. Gimenez, "Structural recursive definitions in type theory," in *Automata, Languages and Programming, 25th International Colloquium, ICALP'98*. Springer-Verlag, 1998, pp. 13–17.
- [22] D. Baelde, "Least and greatest fixed points in linear logic," *ACM Transactions on Computational Logic*, vol. 13, no. 1, Jan. 2012. [Online]. Available: <http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/baelde12tocl.pdf>
- [23] B. Grégoire and J. L. Sacchini, "On strong normalization of the calculus of constructions with type-based termination," in *LPAR*, 2010, pp. 333–347.
- [24] A. Abel and B. Pientka, "Wellfounded recursion with copatterns: A unified approach to termination and productivity," in *LICS*, 2013.